



# Red Hat Enterprise Virtualization 3.3 Evaluation Guide

---

Evaluating Red Hat Enterprise Virtualization

Zac Dover  
Dan Macpherson

Steve Gordon  
Dayle Parker

Tim Hildred



# Red Hat Enterprise Virtualization 3.3 Evaluation Guide

---

## Evaluating Red Hat Enterprise Virtualization

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

## Legal Notice

Copyright © 2014 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.

## Abstract

This book explains how to deploy a Red Hat Enterprise Virtualization Environment for evaluation purposes.

# Table of Contents

<b>Preface</b> .....	<b>5</b>
1. Document Conventions	5
1.1. Typographic Conventions	5
1.2. Pull-quote Conventions	6
1.3. Notes and Warnings	7
2. Getting Help and Giving Feedback	7
2.1. Do You Need Help?	7
2.2. We Need Feedback	8
<b>Introduction</b> .....	<b>8</b>
1. Introduction to Red Hat Enterprise Virtualization Evaluation Guide	8
2. Evaluation Tracks	8
3. Track A	9
3.1. Track A: Standard Setup	9
3.2. Track A Requirements	12
4. Track B	13
4.1. Track B: Minimal Setup	13
4.2. Track B Requirements	15
<b>Chapter 1. Lab 1 - Basic Setup for Installation and Configuration</b> .....	<b>18</b>
1.1. Lab 1 - Objectives	18
1.2. Installing Red Hat Enterprise Virtualization Manager	19
1.3. Installing Red Hat Enterprise Virtualization Hypervisor	24
1.3.1. Registering the Host on RHN and Acquiring ISO Hypervisor Images	25
1.3.2. Preparing Optical Hypervisor Installation Media	26
1.3.3. Install Red Hat Enterprise Virtualization Hosts from Optical Installation Media	27
1.3.4. Install Red Hat Enterprise Virtualization Hypervisors	27
1.4. Connecting to Red Hat Enterprise Virtualization Web Administration Portal	30
1.5. Graphical User Interface Elements	32
1.6. Approve the Red Hat Enterprise Virtualization Hypervisor	33
1.7. Configure Logical Networks	34
1.8. Configuring Storage	36
1.9. Configuring Storage	37
1.9.1. Creating an NFS Data Domain	37
1.9.2. Creating an iSCSI Data Domain	39
1.9.3. Creating an FCP Data Domain	42
1.10. Attach and Populate ISO Domains	44
1.11. Lab 1 - summary	45
<b>Chapter 2. Lab 2 - Creating virtual machines</b> .....	<b>47</b>
2.1. Lab 2 - Objectives	47
2.2. Creating a Red Hat Enterprise Linux virtual machine	47
2.3. Creating Red Hat Enterprise Linux templates	50
2.4. Creating a Template from a Red Hat Enterprise Linux Virtual Machine	51
2.5. Sealing a Linux Virtual Machine Manually for Deployment as a Template	52
2.6. Sealing a Linux Virtual Machine for Deployment as a Template using sys-unconfig	52
2.7. Cloning a Red Hat Enterprise Linux virtual machine from a template	53
2.8. Lab 2 - Summary	54
<b>Chapter 3. Lab 3 - Live migration scenarios</b> .....	<b>56</b>
3.1. Lab 3 - Objectives	56
3.2. Live migration scenarios	56
3.3. Activating live migration	56

3.3. Activating live migration	55
3.4. Moving hosts into maintenance mode	57
3.5. Defining cluster policies	58
3.6. Lab 3 - Summary	60
<b>Chapter 4. Lab 4 - Power User Portal</b>	<b>61</b>
4.1. Lab 4 - Objectives	61
4.2. Adding IdM domains	61
4.3. Adding new users in the IdM directory	62
4.4. PowerUserRole permission assignment	63
4.5. Assigning PowerUserRole permissions on existing virtual machines	64
4.6. Logging in to the Power User Portal	65
4.7. Logging in to the User Portal	67
4.8. Creating Linux desktop virtual machines	67
4.9. Opening virtual machine consoles	68
4.10. Making templates from virtual machines	68
4.11. Verifying virtual machine template permissions	70
4.12. Lab 4 - Summary	70
<b>Chapter 5. Lab 5 - Managing Multi-Level Administrators</b>	<b>72</b>
5.1. Lab 5 - Objectives	72
5.2. Lab 5 - requirements	72
5.3. Storage administrator definition	72
5.4. Assigning system administrator roles to storage domains	73
5.5. Virtual machine administrator definition	73
5.6. User permission verification	74
5.7. Custom role creation	75
5.8. Lab 5 - Summary	76
<b>Chapter 6. Lab 6 - High Availability scenarios</b>	<b>78</b>
6.1. Lab 6 - Objectives	78
6.2. Lab 6 - Requirements	78
6.3. Lab 6 - Prerequisites	78
6.4. Power Management Configuration	79
6.4.1. Power management configuration	79
6.4.2. Disabling cluster policy	79
6.4.3. Setting up power management on a host	80
6.5. Virtual Machine High-Availability Configuration	82
6.5.1. Virtual Machine High-availability Configuration	83
6.6. High Availability: Host-Initiated Reboot	85
6.6.1. High availability -- host initiated reboot	85
6.7. High Availability: Virtual Machine Interruption	87
6.7.1. High availability -- Virtual machine interruption	87
6.7.2. Demonstrating Virtual Machine High Availability when its Processes are Killed	88
6.8. High-Availability: Non-Operational Host	90
6.8.1. High Availability - Non-Operational Hosts	90
6.8.2. Demonstrating Virtual Machine High Availability when Storage Network is Down (Host: SPM)	91
6.8.3. Demonstrating virtual-machine high availability when storage networks are down (Host: non-SPM)	92
6.9. High Availability: Non-responsive host	93
6.9.1. High availability -- non-responsive host	93
6.9.2. Demonstrating high-availability when host connections are disrupted	93
6.9.3. Demonstrating virtual-machine high availability of incorrectly-fenced hosts	94
6.10. Lab 6 - Summary	96

<b>Chapter 7. Lab 7 - Adding Additional Data Centers</b>	<b>97</b>
7.1. Lab 7 - Objectives	97
7.2. Add Additional Data Center	97
7.3. Installing Red Hat Enterprise Linux 6 hosts	100
7.4. Create a New Data Center	101
7.5. Create a New Cluster	103
7.6. Attach New Red Hat Enterprise Linux Hosts	104
7.7. Configure Logical Networks	106
7.8. Configuring Storage	107
7.8.1. Creating an NFS Data Domain	107
7.8.2. Creating an iSCSI Data Domain	109
7.8.3. Creating an FCP Data Domain	111
7.9. Configure ISO Domain	113
7.10. Lab 7 - Summary	114
<b>Chapter 8. Lab 8 - Virtual desktops</b>	<b>115</b>
8.1. Lab 8 - Objectives	115
8.2. Create a Windows Virtual Machine	115
8.3. Creating Windows virtual machines from templates	119
8.4. Create a Windows Template	121
8.5. Assigning UserRole permissions on virtual machines	122
8.6. Connecting to virtual machines	122
8.7. Use Virtual Machine Pools	123
8.7.1. Creating Virtual Machine Pools	123
8.7.2. Assigning UserRole Permissions	124
8.7.3. Allocating Virtual Machines	125
8.7.4. Deallocating Virtual Machines	126
8.8. Lab 8 - Summary	127
<b>Chapter 9. Lab 9 - Installing and configuring minimal setup</b>	<b>128</b>
9.1. Lab 9 - Objectives	128
9.2. Installing Red Hat Enterprise Virtualization Manager	128
9.3. Installing Red Hat Enterprise Virtualization Hypervisor	133
9.3.1. Registering the Host on RHN and Acquiring ISO Hypervisor Images	134
9.3.2. Preparing Optical Hypervisor Installation Media	135
9.3.3. Install Red Hat Enterprise Virtualization Hosts from Optical Installation Media	136
9.3.4. Install Red Hat Enterprise Virtualization Hypervisors	136
9.4. Connecting to Red Hat Enterprise Virtualization Web Administration Portal	139
9.5. Graphical User Interface Elements	141
9.6. Approve the Red Hat Enterprise Virtualization Hypervisor	142
9.7. Creating Local Storage	143
9.8. Creating local storage domains	144
9.9. Attach and Populate ISO Domains	144
9.10. Lab 9 - Summary	146
<b>Chapter 10. Lab 10 - Advanced storage features</b>	<b>147</b>
10.1. Lab 10 - Objectives	147
10.2. Advanced Storage Features	147
10.3. Requirements	147
10.4. Creating Floating Disks	147
10.5. Associating a Floating Disk with a Virtual Machine	148
10.6. Marking a Virtual Disk Shared	148
10.7. Associating a Shared Virtual Disk with a Second Virtual Machine	149

10.8. Creating a Snapshot of a Running Virtual Machine	149
10.9. Creating a Virtual Machine from a Snapshot	150
10.10. Associating LUNs with Virtual Machines	151
10.11. Associating Fibre Channel (FC) LUNs with Virtual Machine	156
10.12. Lab 10 - Summary	156
<b>Revision History</b> .....	<b>158</b>



---

# Preface

## 1. Document Conventions

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

### 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 submenu 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:

Publican 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>. From the Customer Portal, you can:

- ✧ Search or browse through a knowledge base 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 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.

When submitting a bug report, be sure to mention the manual's identifier: [Guides-Evaluation](#)

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.

## Introduction

### 1. Introduction to Red Hat Enterprise Virtualization Evaluation Guide

This guide gets you started on a full featured Red Hat Enterprise Virtualization deployment using your existing resources. From installing a hypervisor, setting up shared storage and running a fully functional virtual machine, see how you can implement virtualization in your own organization. The tutorials are organized to reflect typical deployments from a small office with a few hosts, to a large enterprise with multiple data centers, and you can even test out a single host version at home. The tracks and tutorials are color coded to provide clear paths through this guide.

Once you have completed the tutorials (labs), arrange your own tracks for further evaluation or training. The tracks are suggested examples only. Use them to understand how to deploy Red Hat Enterprise Virtualization, and then work with our Red Hat Enterprise Virtualization solution architects to build your own unique environments.

[Report a bug](#)

## 2. Evaluation Tracks

Depending on your needs, use one of these tracks to evaluate Red Hat Enterprise Virtualization for your organization:

- ✳ If you have shared storage and two physical servers to install hosts, use Track A: Standard Setup.
- ✳ If you have one physical server on which to install a host, use Track B: Minimal Setup.

For both tracks, you need an evaluation license and a valid Red Hat Network subscription to:

- ✳ the Red Hat Enterprise Virtualization channel
- ✳ the Red Hat Enterprise Linux channel

Refer to the *Red Hat Enterprise Virtualization Manager Release Notes* for specific channel names current to your system. Contact your sales representative if you do not have both of the above.

**See Also:**

- [Section 3.1, “Track A: Standard Setup”](#)
- [Section 4.1, “Track B: Minimal Setup”](#)

[Report a bug](#)

## 3. Track A

### 3.1. Track A: Standard Setup

Track A describes the installation and configuration of a basic Red Hat Enterprise Virtualization environment. Track A requires you to have more than one host and shared storage.

Track A consists of:

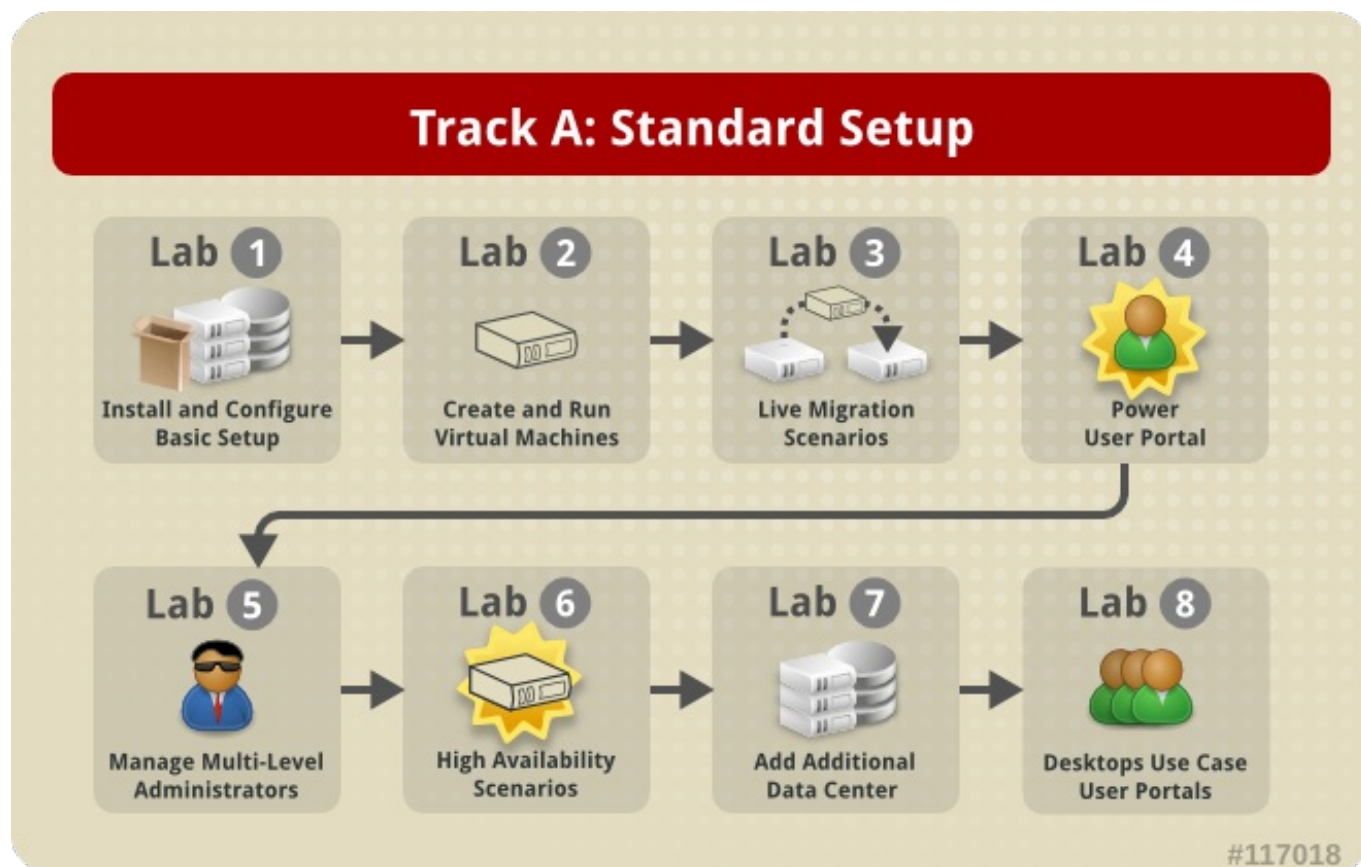
- four basic labs
- four advanced labs (optional)

The basic labs describe how to create virtual machines and assign them to users.

The advanced labs describe how to apply Red Hat Enterprise Virtualization to real-life enterprise operations. This includes:

- protecting against hardware failure by using high-availability
- assigning different levels of user permissions to virtual machines, storage, and servers to reflect the needs of your organization
- reconfiguring red hat enterprise linux servers so that they function as virtual machine hosts
- provisioning virtual desktops for users

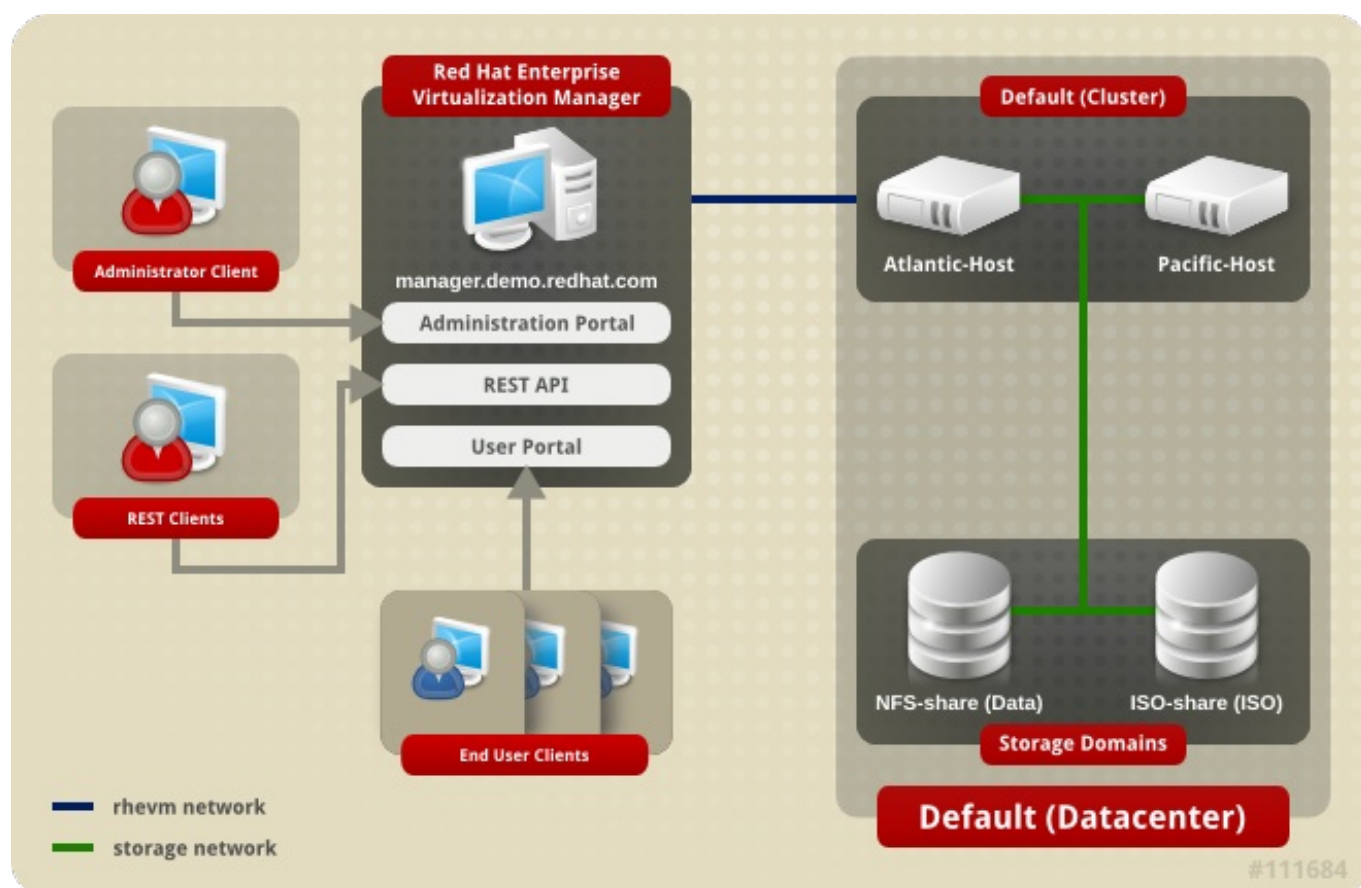
If you have three to four physical servers and one shared storage resource, set up a Red Hat Enterprise Virtualization environment by following the labs in Track A - Standard Setup. The labs in this track are color coded red, so to determine if a lab is part of Track A, look for a red square in the diagram at the start of each lab.



**Figure 1. Evaluation Track A**

After you complete Track A, you will have an environment that includes the following:

- ✧ two virtual machine hosts
- ✧ shared storage
- ✧ a network
- ✧ two portals (User Portal and Web Administration Portal)
- ✧ User clients
- ✧ a host running the Red Hat Enterprise Virtualization Manager



**Figure 2. Red Hat Enterprise Virtualization Standard Setup**

### Track A - Standard Setup Labs

- **Install and Configure Basic Setup:** Install the Red Hat Enterprise Virtualization Manager and Red Hat Enterprise Virtualization Hypervisor, configure storage and define networks (75 minutes).
- **Create Virtual Machines:** Create virtual machines and templates from the administration portal (25 minutes).
- **Live Migration Scenarios:** Configure automatic virtual machine live migration during hardware downtime (10 minutes).
- **Power User Portal:** Create and manage virtual machines from the power user portal (35 minutes).

The basic labs allow you to evaluate how Red Hat Enterprise Virtualization can be deployed in your environment.

The advanced labs show you how to optimize your Red Hat Enterprise Virtualization setup. The advanced labs have additional hardware requirements, which are listed at the start of each advanced lab.

### Track A - Advanced Labs (optional)

- **Manage Multi-Level Administrators:** Manage administrators for each component of Red Hat Enterprise Virtualization (10 minutes).
- **High Availability Scenarios:** Configure power management and high availability (30 minutes).
- **Add Additional Data Center:** Create an additional data center with Red Hat Enterprise Linux hosts. (35 minutes)



- Virtual Desktops: Access desktop pools using the SPICE connection protocol (50 minutes).

[Report a bug](#)

## 3.2. Track A Requirements

To work through the labs in Track A, you must have:

- an evaluation license
- a valid Red Hat Network subscription to
  - the Red Hat Enterprise Virtualization channel
  - the Red Hat Enterprise Linux channel



### Note

Refer to the *Red Hat Enterprise Virtualization Manager Release Notes* for specific channel names current to your system.

Contact your sales representative if you do not have both of the above.

Before you begin Track A, ensure that you have the following:

### Red Hat Enterprise Virtualization Manager Requirements

- Minimum - Dual core server with 4 GB RAM, 25 GB free disk space and 1 Gbps network interface.
- Recommended - Dual Sockets/Quad core server with 16 GB RAM, 50 GB free disk space on multiple disk spindles and 1 Gbps network interface.

The breakdown of the server requirements is:

- Red Hat Enterprise Linux 6 operating system: 1 GB RAM and 5 GB local disk space
- Manager: 3 GB RAM, 3 GB local disk space and 1 Gbps network controller bandwidth
- Local ISO domain: 15 GB disk space
- One client for connecting to Red Hat Enterprise Virtualization Manager.
  - A machine with Firefox 17 or higher installed on Red Hat Enterprise Linux.
  - Internet Explorer 9 or higher on Microsoft Windows.

### Red Hat Enterprise Virtualization Hosts Requirements

- Minimum - Dual Core server, 2 GB RAM and 10 GB Storage, 1 Gbps network interface
- Recommended - Dual socket server, 16 GB RAM and 50 GB storage, two 1 Gbps network interfaces

Server requirements:

- For each host: AMD-V or Intel VT enabled, AMD64 or Intel 64 extensions, minimum 1 GB RAM, 3 GB free storage and 1 Gbps network interface.



- For virtual machines running on each host: 8 GB RAM to run four virtual machines.

## Storage and Networking Requirements

- ✦ At least one of the supported storage types (NFS, iSCSI or FCP).
- ✦ At least three static IP addresses: One for the Red Hat Enterprise Virtualization Manager server and one for each server running Red Hat Enterprise Virtualization Hypervisor.
- ✦ DNS service which can resolve (forward and reverse) all the IP addresses.
- ✦ An existing DHCP server which can allocate network addresses for the virtual machines.
- ✦ Display subnet (extra Network Interface Card on both servers) to create a new display network in addition to the default existing management network.

## Virtual Machine Requirements

- ✦ Installation images for creating virtual machines. These images will be installed on the virtual machines you create.
  - Red Hat Enterprise Linux 3, 4, 5 or 6.
  - Microsoft Windows XP, 7, 8, 2003 or 2008.
- ✦ Valid licenses or subscription entitlements for each operating system.
- ✦ At least one valid user account in any of the supported directory services (IdM, AD, RHDS, or OpenLDAP).

[Report a bug](#)

## 4. Track B

### 4.1. Track B: Minimal Setup

Track B describes the installation and configuration of a minimal Red Hat Enterprise Virtualization environment on a single host, using local storage.

Track B consists of:

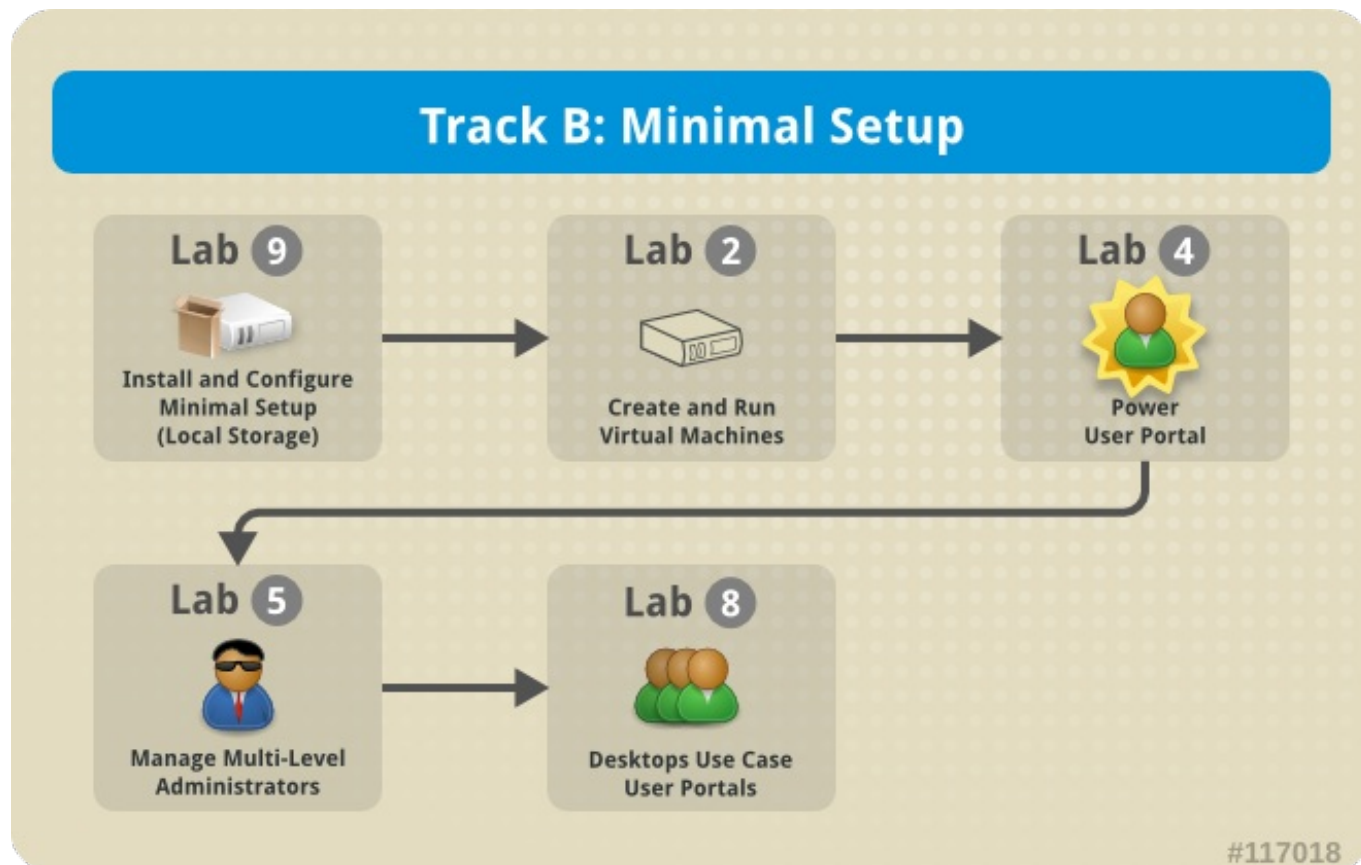
- ✦ three basic labs
- ✦ two advanced labs (optional)

The basic labs describe how to create virtual machines and assign them to users.

The advanced labs include:

- ✦ using the multi-level administration system to assign different levels of user permissions (this is ideal for companies with diverse employee roles)
- ✦ provisioning virtual desktops for users

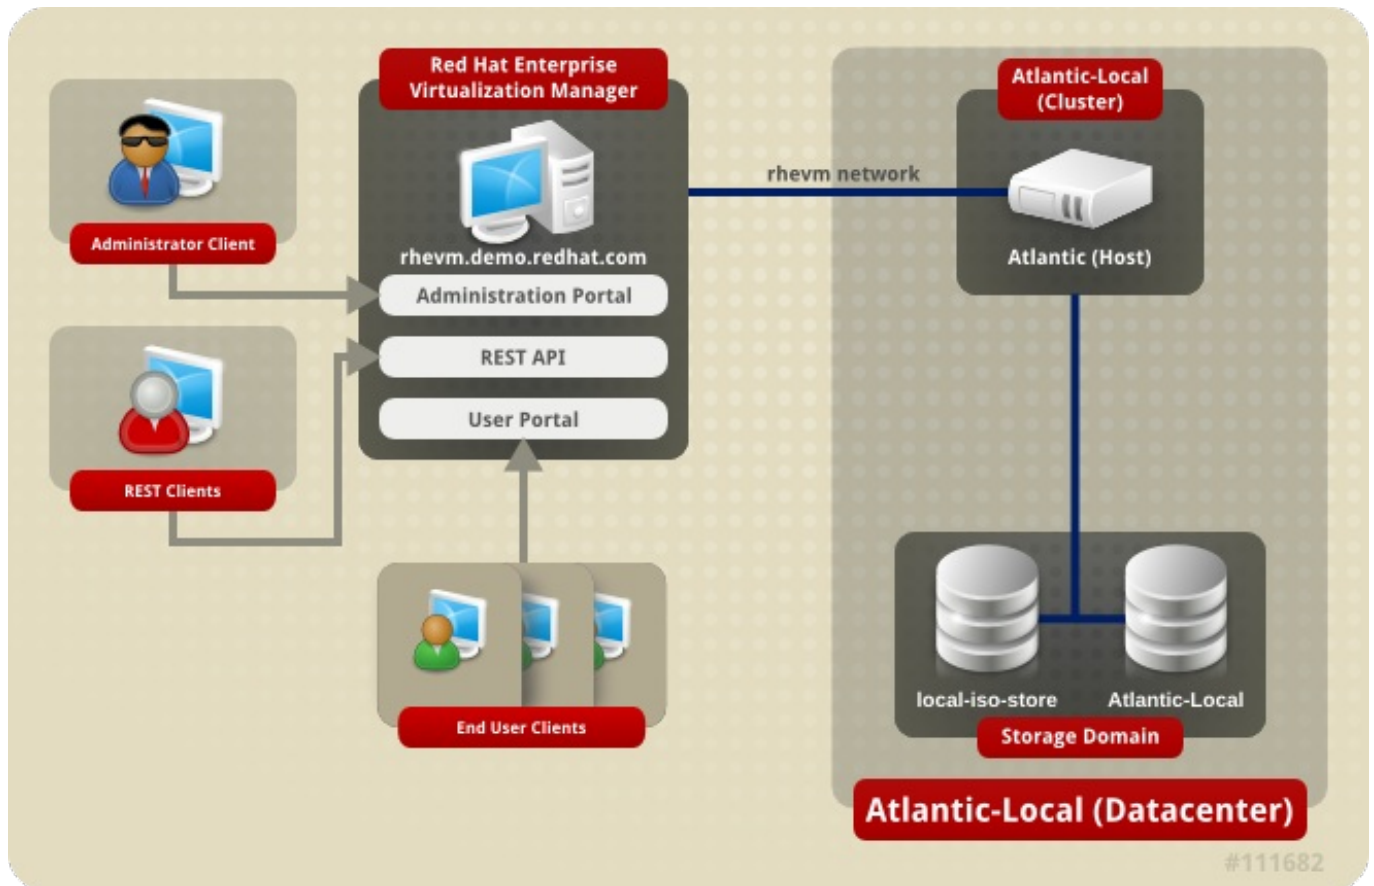
If you have two physical servers and no shared storage, set up a Red Hat Enterprise Virtualization environment by following the labs in Track B - Minimal Setup. The labs in this track are color coded blue. To determine if a lab is part of Track B, look for a blue square in the diagram at the start of each lab.



**Figure 3. Workflow for Evaluation Track B**

After you complete Track B, you will have the following environment that includes the following:

- ✧ one virtual machine host
- ✧ local storage
- ✧ a network
- ✧ two portals
- ✧ user clients
- ✧ a host running Red Hat Enterprise Virtualization Manager



**Figure 4. Red Hat Enterprise Virtualization Minimal Setup**

### Track B - Minimal Setup Labs

- ✦ Install and Configure Minimal Setup: Install the Red Hat Enterprise Virtualization Manager and Red Hat Enterprise Virtualization Hypervisor using minimal hardware (60 minutes).
- ✦ Create Virtual Machines: Create virtual machines and templates from the administration portal (25 minutes).
- ✦ Power User Portal: Create and manage virtual machines from the power user portal (35 minutes).

### Track B - Advanced Lab (optional)

- ✦ Manage Multi-Level Administrators: Manage administrators for each component of Red Hat Enterprise Virtualization (10 minutes).
- ✦ Virtual Desktops: Access desktop pools using the SPICE connection protocol (50 minutes).

[Report a bug](#)

## 4.2. Track B Requirements

To work through the labs in Track B, you must have:

- ✦ an evaluation license
- ✦ a valid Red Hat Network subscription to
  - the Red Hat Enterprise Virtualization channel
  - the Red Hat Enterprise Linux channel



## Note

Refer to the *Red Hat Enterprise Virtualization Manager Release Notes* for specific channel names current to your system.

Contact your sales representative if you do not have both of the above.

Before you begin Track B, ensure that you have the following:

### Red Hat Enterprise Virtualization Manager Requirements

- ✧ Minimum - Dual core server with 4 GB RAM, with 25 GB free disk space and 1 Gbps network interface.
- ✧ Recommended - Dual Sockets/Quad core server with 16 GB RAM, 50 GB free disk space on multiple disk spindles and 1 Gbps network interface.

The breakdown of the server requirements are as below:

- For the Manager: 3 GB memory, 3 GB local disk space, 1 Gbps network controller bandwidth
- For the Red Hat Enterprise Linux 6 operating system: 1 GB memory, 5 GB local disk space
- For the local ISO domain: 15 GB disk space
- ✧ One client for connecting to Red Hat Enterprise Virtualization Manager.
  - A machine with Firefox 17 or higher installed on Red Hat Enterprise Linux.
  - Internet Explorer 9 or higher on Microsoft Windows.

### Red Hat Enterprise Virtualization Host Requirements

- ✧ Minimum - Dual Core server, 2 GB RAM and 10 GB Storage, 1 Gbps network interface.
- ✧ Recommended - Dual socket server, 16 GB RAM and 50 GB storage, two 1 Gbps network interfaces.

The breakdown of the server requirements is:

- For each host: AMD-V or Intel VT enabled, AMD64 or Intel 64 extensions, minimum 1 GB RAM, 3 GB free storage and 1 Gbps network interface.
- For virtual machines running on the host: 8 GB RAM to run four virtual machines.

### Storage and Networking Requirements

- ✧ Two static IP addresses: One for the Red Hat Enterprise Virtualization Manager server and one for the Red Hat Enterprise Virtualization Hypervisor.
- ✧ One name (DNS) Server which can resolve (forward and reverse) all the IP addresses.
- ✧ A DHCP server for the virtual machines.

### Virtual Machine Requirements

- ✧ Installation images for creating virtual machines. These images will be installed on the virtual

machines you create.

- Red Hat Enterprise Linux 3, 4, 5 or 6.
- Microsoft Windows XP, 7, 8, 2003 or 2008.
- Valid licenses or subscription entitlements for each operating system.
- At least one valid user account in any of the supported directory services (IdM, AD, RHDS, or OpenLDAP).

[Report a bug](#)  
.....

## Chapter 1. Lab 1 - Basic Setup for Installation and Configuration

### 1.1. Lab 1 - Objectives



Lab 1 on Track A sets up the basic infrastructure to support virtualization, and shows you how to install and configure the hosts, storage and networks in readiness for the virtual machines. The goal of this lab is an environment that is ready for you to create and provision virtual machines.

This lab is intended for Track A, and requires three servers and shared storage.

#### Lab 1 - Objectives

To achieve the goal of this lab, you will install and set up Red Hat Enterprise Virtualization with multiple hosts and shared storage, you will learn to configure networks and add ISOs. This lab should take you about 75 minutes.

**Install Red Hat Enterprise Virtualization Manager** shows you how to install the Red Hat Enterprise Virtualization Manager on a server running Red Hat Enterprise Linux. (12 minutes\*)

**Install Red Hat Enterprise Virtualization Hypervisor** shows you how to install and configure Red Hat Enterprise Virtualization Hypervisors for use with Red Hat Enterprise Virtualization Manager. (20 minutes\*)

**Connect to Red Hat Enterprise Virtualization Manager** shows you how to configure a client machine to connect to the Red Hat Enterprise Virtualization Manager administration portal. (8 minutes)

**Approve the Red Hat Enterprise Virtualization Hypervisor** shows you how to approve the hosts for use from the Red Hat Enterprise Virtualization Manager. (10 minutes)

**Configure Logical Networks** shows you how to define networks for the storage devices and add them to the hosts. (5 minutes)

**Configure Storage** shows you how to define NFS, iSCSI or FCP storage and attach the domains to the data center. (10 minutes)

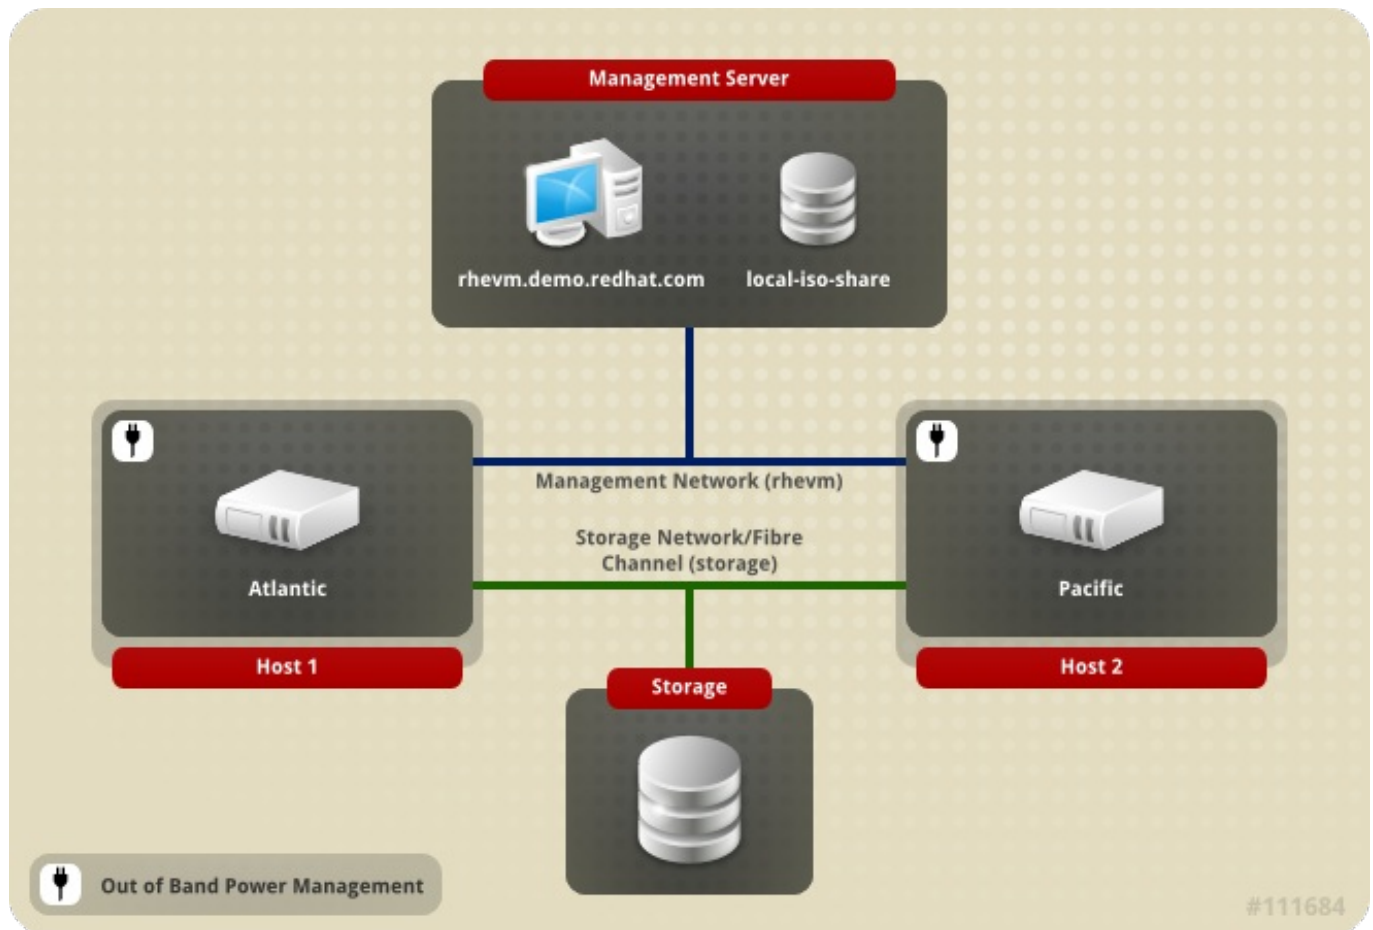
**Attach and Populate ISO Domain** shows you how to attach the predefined ISO domain to the data center and upload ISO images to the repository. (10 minutes)

\* The time required to download packages from the Red Hat Network depends on the bandwidth of your connection to RHN, therefore it has not been included in the estimated time.

#### Lab 1 - Configuration

The following figure shows the environment parameters and object names which will be used consistently throughout this lab. It is strongly recommended that you use these entities in your evaluation environment to ensure that the names are resolvable. You may alter them if necessary, but

make sure you have an equivalent name for each component.



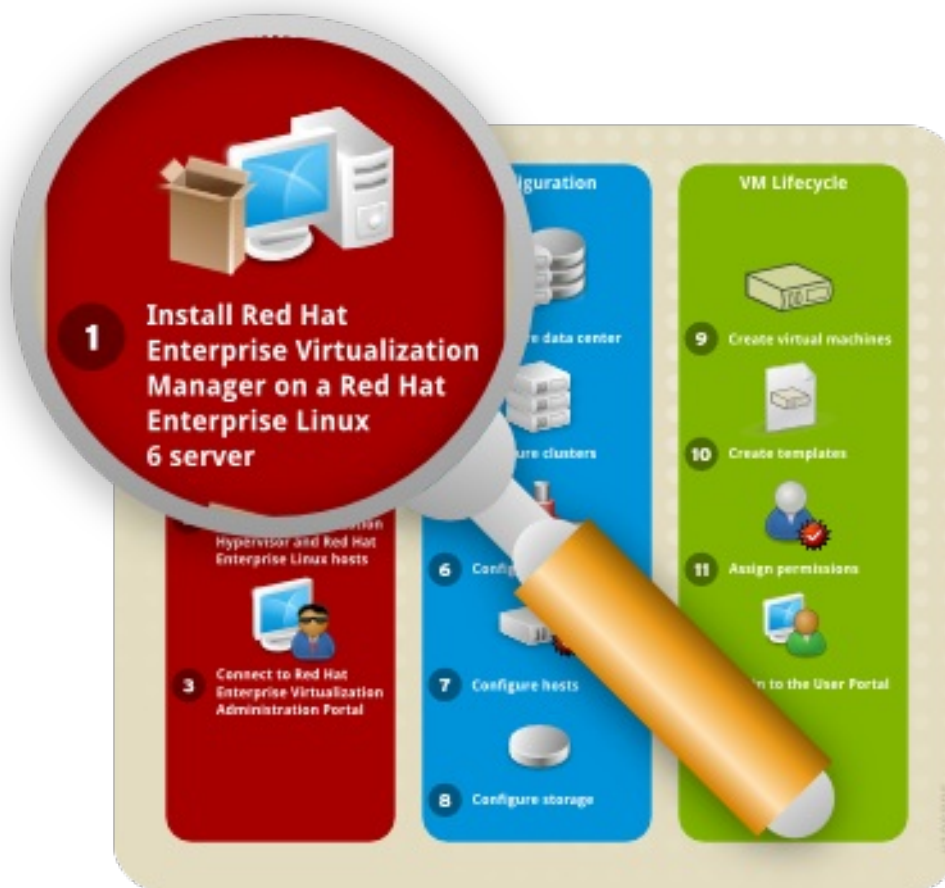
#### See Also:

- [Section 1.2, “Installing Red Hat Enterprise Virtualization Manager”](#)
- [Section 1.3, “Installing Red Hat Enterprise Virtualization Hypervisor”](#)
- [Section 1.4, “Connecting to Red Hat Enterprise Virtualization Web Administration Portal”](#)
- [Section 1.6, “Approve the Red Hat Enterprise Virtualization Hypervisor”](#)
- [Section 1.7, “Configure Logical Networks”](#)
- [Section 1.8, “Configuring Storage”](#)
- [Section 1.10, “Attach and Populate ISO Domains”](#)

[Report a bug](#)

## 1.2. Installing Red Hat Enterprise Virtualization Manager





**Figure 1.1. Install Red Hat Enterprise Virtualization Manager**

The Red Hat Enterprise Virtualization Manager is the control center of the Red Hat Enterprise Virtualization environment. It allows you to define hosts, configure data centers, add storage, define networks, create virtual machines, manage user permissions, and use templates from one central location.

The Red Hat Enterprise Virtualization Manager must be installed on a server running Red Hat Enterprise Linux 6, with minimum 4 GB RAM, 25 GB free disk space and 1 Gbps network interface.

#### **Procedure 1.1. To install Red Hat Enterprise Virtualization Manager**

1. Install Red Hat Enterprise Linux 6 on a server. When prompted for the software packages to install, select the default **Basic Server** option. See the *Red Hat Enterprise Linux Installation Guide* for more details.



#### **Important**

During installation, remember to set the fully qualified domain name (FQDN) and IP for the server.

2. If your server has not been registered with the Red Hat Network, run:

```
# rhn_register
```

To complete registration successfully you need to supply your Red Hat Network user name and password. Follow the onscreen prompts to complete registration of the system.



After you have registered your server, update all the packages on it. Run:

```
# yum -y update
```

Reboot your server for the updates to be applied.

3. Subscribe the server to the required Red Hat Network channels. See the *Red Hat Enterprise Virtualization Manager Release Notes* for a list of required channels.
4. You are now ready to install the Red Hat Enterprise Virtualization Manager. Run the following command:

```
# yum -y install rhevm
```

This command will download the Red Hat Enterprise Virtualization Manager installation software and resolve all dependencies.

5. When the packages have finished downloading, run the installer:

```
# engine-setup
```

6. The installer will take you through a series of interactive questions as listed in the following example. If you do not enter a value when prompted, the installer uses the default settings which are stated in [ ] brackets.

### Example 1.1. Red Hat Enterprise Virtualization Manager installation

```
[ INFO ] Stage: Initializing
[ INFO ] Stage: Environment setup
          Configuration files: ['/etc/ovirt-engine-
setup.conf.d/10-packaging.conf']
          Log file: /var/log/ovirt-engine/setup/ovirt-engine-
setup-20131219122349.log
          Version: otopi-1.1.2 (otopi-1.1.2-1.el6ev)
[ INFO ] Stage: Environment packages setup
[ INFO ] Stage: Programs detection
[ INFO ] Stage: Environment setup
[ INFO ] Stage: Environment customization

--== PACKAGES ==--

[ INFO ] Checking for product updates...
[ INFO ] No product updates found

--== NETWORK CONFIGURATION ==--
          Host fully qualified DNS name of this server
[localhost.localdomain]:
          Setup can automatically configure the firewall on this
system.
          Note: automatic configuration of the firewall may
overwrite current settings.
          Do you want Setup to configure the firewall? (Yes, No)
[Yes]:
```

```

--== DATABASE CONFIGURATION ==--
Where is the database located? (Local, Remote) [Local]:
Setup can configure the local postgresql server
automatically for the engine to run. This may conflict with
existing applications.
Would you like Setup to automatically configure
postgresql, or prefer to perform that manually? (Automatic,
Manual) [Automatic]:

--== OVIRT ENGINE CONFIGURATION ==--
Engine admin password:
Confirm engine admin password:
Application mode (Both, Virt, Gluster) [Both]:
Default storage type: (NFS, FC, ISCSI, POSIXFS) [NFS]:

--== PKI CONFIGURATION ==--
Organization name for certificate [localdomain]:

--== APACHE CONFIGURATION ==--
Setup can configure the default page of the web server
to present the application home page. This may conflict with
existing applications.
Do you wish to set the application as the default page
of the web server? (Yes, No) [Yes]:
Setup can configure apache to use SSL using a
certificate issued from the internal CA.
Do you wish Setup to configure that, or prefer to
perform that manually? (Automatic, Manual) [Automatic]:

--== SYSTEM CONFIGURATION ==--
Configure an NFS share on this server to be used as an
ISO Domain? (Yes, No) [Yes]:
Local ISO domain path [/var/lib/exports/iso-
20131219172449]:
Local ISO domain name [ISO_DOMAIN]:
Configure WebSocket Proxy on this machine? (Yes, No)
[Yes]:

--== END OF CONFIGURATION ==--
Would you like transactions from the Red Hat Access
Plugin sent from the RHEV Manager to be brokered through a proxy
server? (Yes, No) [No]:

```

Important points to note:

- The default ports **80** and **443** must be available to access the Manager on HTTP and HTTPS respectively.
- If you elect to configure an NFS share it will be exported from the machine on which the Manager is being installed.
- The storage type that you select will be used to create a data center and cluster. You will then be able to attach storage to these from the Administration Portal.

7. You are then presented with a summary of the configurations you have selected.

**Example 1.2. Confirm Manager installation settings**

```

--== CONFIGURATION PREVIEW ==--

Database name                : engine
Database secured connection  : False
Database host                 : localhost
Database user name           : engine
Database host name validation : False
Database port                 : 5432
NFS setup                     : True
PKI organization              : localdomain
Application mode              : both
Firewall manager              : iptables
Update Firewall               : True
Configure WebSocket Proxy     : True
Host FQDN                     :
localhost.localdomain
NFS mount point               :
/var/lib/exports/iso-20131219172449
Datacenter storage type       : nfs
Configure local database      : True
Set application as default page : True
Configure Apache SSL          : True

Please confirm installation settings (OK, Cancel) [OK]:

```

8. The installation commences. The following message displays, indicating that the installation was successful.

**Example 1.3. Successful installation**

```

[ INFO ] Stage: Transaction setup
[ INFO ] Stopping engine service
[ INFO ] Stopping websocket-proxy service
[ INFO ] Stage: Misc configuration
[ INFO ] Stage: Package installation
[ INFO ] Stage: Misc configuration
[ INFO ] Creating PostgreSQL database
[ INFO ] Configuring PostgreSQL
[ INFO ] Creating database schema
[ INFO ] Creating CA
[ INFO ] Configuring WebSocket Proxy
[ INFO ] Generating post install configuration file '/etc/ovirt-
engine-setup.conf.d/20-setup-ovirt-post.conf'
[ INFO ] Stage: Transaction commit
[ INFO ] Stage: Closing up

--== SUMMARY ==--
An ISO NFS share has been created on this host.
    If IP based access restrictions are required, edit:
    entry /var/lib/exports/iso-20131219172449 in
/etc/exports
SSH fingerprint: fingerprint

```

```

Internal CA A9: certificate
Web access is enabled at:
    http://localhost.localdomain:80/ovirt-engine
    https://localhost.localdomain:443/ovirt-engine
Please use the user "admin" and password specified in
order to login into oVirt Engine

```

```
--== END OF SUMMARY ==--
```

```

[ INFO ] Starting engine service
[ INFO ] Restarting httpd
[ INFO ] Restarting nfs services
[ INFO ] Generating answer file '/var/lib/ovirt-
engine/setup/answers/20131219122536-setup.conf'
[ INFO ] Stage: Clean up
        Log file is located at /var/log/ovirt-
engine/setup/ovirt-engine-setup-20131219122349.log
[ INFO ] Stage: Pre-termination
[ INFO ] Stage: Termination
[ INFO ] Execution of setup completed successfully

```

Your Red Hat Enterprise Virtualization Manager is now up and running. You can log in to the Red Hat Enterprise Virtualization Manager's web administration portal with the user name **admin** (the administrative user configured during installation) in the **internal** domain. Instructions to do so are provided at the end of this chapter.



### Important

The **internal** domain is automatically created upon installation, however no new users can be added to this domain. To authenticate new users, you need an external directory service. Red Hat Enterprise Virtualization supports Identity Management (IdM), Active Directory, RHDS, and OpenLDAP, and provides a utility called **engine-manage-domains** for attaching new directories to the system. Use of this tool is covered in the *Red Hat Enterprise Virtualization Installation Guide*.



### Note

The Red Hat Enterprise Virtualization engine (that is, the Manager) can be hosted in a virtual machine controlled by the engine that hosts it. This setup is called a "self-hosted engine". In a self-hosted engine system, the virtual machine on which the Manager resides is defined as highly available. Two packages, *ovirt-hosted-engine-setup* and *ovirt-hosted-engine-ha* provide the setup and services necessary to deploy a self-hosted-engine Red Hat Enterprise Virtualization environment. For more on Red Hat Enterprise Virtualization Self-Hosted Engine, see the *Red Hat Enterprise Virtualization Installation Guide*.

[Report a bug](#)

## 1.3. Installing Red Hat Enterprise Virtualization Hypervisor

### 1.3.1. Registering the Host on RHN and Acquiring ISO Hypervisor Images

#### Summary

The **Red Hat Enterprise Virtualization Manager** Red Hat Network channel contains the Hypervisor packages. The Hypervisor itself is contained in the **rhev-hypervisor** package. Additional tools supporting USB and PXE installations are installed as dependencies. Install the Hypervisor packages on the system you plan to use to create Hypervisor boot media.

Select one of the two options below:

#### Procedure 1.2. Subscribing to RHN Entitlement Pools and Installing the Red Hat Enterprise Virtualization Hypervisor Packages

##### 1. Subscribing to download the Hypervisor using certificate-based RHN

###### a. Identify Available Entitlement Pools

To subscribe the system to the Red Hat Enterprise Virtualization channels you need you must locate the identifier for the relevant entitlement pool. Use the **list** action in the **subscription-manager** to find these:

```
# subscription-manager list --available | grep -A8 "Red Hat Enterprise Virtualization"
```

###### b. Subscribe System to Entitlement Pools

Using the pool identifiers located in the previous step, subscribe the system to **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** entitlements. Use the **subscribe** parameter of the **subscription-manager** command, and replace **POOLID** with one of the pool identifiers.

```
# subscription-manager subscribe --pool=POOLID
```

##### 2. Subscribing to download the Hypervisor using RHN Classic

- 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 page, and then click **Registered Systems** in the menu that appears.
  - c. Select the system to which you are adding channels from the list 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 **Red Hat Enterprise Virtualization Manager** channel from the list on the screen, then click the **Change Subscription** button to finalize the change.
3. Log in to the system on which the Red Hat Enterprise Virtualization Manager is installed. Log in as **root**.
  4. Use **yum** to install the **rhev-hypervisor**.

```
# yum install rhev-hypervisor
```

## Result

The Hypervisor ISO image is installed into the `/usr/share/rhev-hypervisor/` directory. The `lived -iso -to -disk` and `lived -iso -to -pxeboot` scripts are installed to the `/usr/bin/` directory.



### Note

All version of Red Hat Enterprise Linux 6.2 and higher allow more than one version of the ISO image to be installed at one time. Because of this, `/usr/share/rhev-hypervisor/rhev-hypervisor.iso` is now a symbolic link to a uniquely-named version of the Hypervisor ISO image, for instance `/usr/share/rhev-hypervisor/rhev-6.2-20111006.0.el6.iso`. Different versions of the image can now be installed alongside each other, allowing administrators to run and maintain a cluster on a previous version of the Hypervisor while upgrading another cluster for testing.

The symbolic links `/usr/share/rhev-hypervisor/rhev-latest6.iso` and `/usr/share/rhev-hypervisor/rhev-hypervisor6.iso` are created. These links target the most-recently installed version of the Red Hat Enterprise Virtualization ISO image.

[Report a bug](#)

## 1.3.2. Preparing Optical Hypervisor Installation Media

### Summary

Burn the Hypervisor image to a CD with the `wodim` command. The `wodim` command is part of the `wodim` package.

### Procedure 1.3. Preparing Optical Hypervisor Installation Media

1. Verify that the `wodim` package is installed on the system.

#### Example 1.4. Verify Installation of `wodim` Package

```
# rpm -q wodim
wodim-1.1.9-11.el6.x86_64
```

If the package version is in the output the package is available.

If nothing is listed, install `wodim`:

```
# yum install wodim
```

2. Insert a blank CD or DVD into your CD or DVD writer.
3. Record the ISO file to the disc. The `wodim` command uses the following:

```
wodim dev=device image.iso
```

This example uses the first CD-RW (`/dev/cdrw`) device available and the default hypervisor image location, `/usr/share/rhev-hypervisor/rhev-hypervisor.iso`.

#### Example 1.5. Use of `wodim` Command

```
# wodim dev=/dev/cdrw /usr/share/rhev-hypervisor/rhev-hypervisor.iso
```

### Result

If no errors occurred, the Hypervisor is ready to boot. Errors sometimes occur during the recording process due to errors on the media itself. If this occurs insert another writable disk and repeat the command above.

The Hypervisor uses a program (`isomd5sum`) to verify the integrity of the installation media every time the Hypervisor is booted. If media errors are reported in the boot sequence you have a bad CD. Follow the procedure above to create a new CD or DVD.

[Report a bug](#)

### 1.3.3. Install Red Hat Enterprise Virtualization Hosts from Optical Installation Media

Now that you have registered the host with RHN, acquired the Hypervisor images, and used the Hypervisor images to create optical installation media, you will boot the system using the optical installation media.

#### Summary

Booting the Hypervisor from optical installation media requires the system to have a correctly defined BIOS boot configuration.

1. Ensure that the system's BIOS is configured to boot from the CD or DVD drive before proceeding.



#### Note

See your manufacturer's manuals for further information on modifying the system's BIOS boot configuration.

2. Insert the Hypervisor CD in the CD or DVD drive.
3. Reboot the system.

### Result

The host's screen will display the Hypervisor boot screen.

[Report a bug](#)

### 1.3.4. Install Red Hat Enterprise Virtualization Hypervisors

## Red Hat Enterprise Virtualization Hypervisor Menu navigation keys

- ✦ Use the Up and Down arrow keys to navigate between selections. Your selections are highlighted in white.
- ✦ The Tab key allows you to move between fields.
- ✦ Use the Spacebar to tick check boxes, which are represented by [ ] brackets. A marked check box displays with an asterisk (\*).
- ✦ To proceed with the selected configurations, press the Enter key.



### Important

Red Hat Enterprise Virtualization Hypervisor is a closed appliance, and does not allow the installation of custom RPMs.

If you require custom RPMs, you must use Red Hat Enterprise Linux hosts. Red Hat Enterprise Linux hosts are not closed appliances, and they allow the installation of custom RPMs.

## To install Red Hat Enterprise Virtualization Hypervisors

1. Insert the Red Hat Enterprise Virtualization Hypervisor installation CD into your CD-ROM drive of the machine designated as a host. Reboot the machine. When the boot splash screen displays, press the **Tab** key and select **Boot** to boot from the hypervisor installation media. Press **Enter**.
2. On the installation confirmation screen, select **Install RHEV Hypervisor** and press **Enter**.
3. The installer automatically detects the drives attached to the system. The disk selected for booting the hypervisor is highlighted in white. Ensure that the local disk is highlighted, otherwise use the arrow keys to select the correct disk. Select **Continue** and press **Enter**.
4. You are prompted to select the drive on which the hypervisor is to be installed. Ensure that the local disk is highlighted, otherwise use the arrow keys to select the correct disk. While multiple installation drives can be used, select only one for this evaluation. Select **Continue** and press **Enter**.
5. Enter a password for local console access and confirm it. Select **Install** and press **Enter**. The Red Hat Enterprise Virtualization Hypervisor partitions the local drive, then commences installation.
6. Once installation is complete, a dialog prompts you to **Reboot** the hypervisor. Press **Enter** to confirm. Remove the installation disc.
7. After the hypervisor has rebooted, you will be taken to a login shell. Log in as the **admin** user with the password you provided during installation to enter the Red Hat Enterprise Virtualization Hypervisor management console.
8. On the hypervisor management console, there are eight tabs on the left. Press the **Up** and **Down** keys to navigate between them and **Enter** to access them.



- a. Select the **Network** tab. Fill in the required fields as shown in the following example. Substitute the **DNS Server** address according to your environment.

```

RHEV Hypervisor 6.3-20120920.0.rhev31.auto92
localhost.localdomain

Status
Network
Security
Keyboard
SNMP
Logging
Kernel Dump
Remote Storage
CIM
RHEV-M
Red Hat Network

System Identification
Hostname: atlantic.demo.redhat.com
DNS Server 1: 192.168.0.2
DNS Server 2:
NTP Server 1: 0.rhel.pool.ntp.org
NTP Server 2: 1.rhel.pool.ntp.org

Device Status Model MAC Address
eth0 Unconfigured Red Hat I 52:54:00:7b:0a:d6

<Flash Lights to Identify> <Apply> <Reset>

```

**Figure 1.2. Configure Hypervisor network settings**

After you have filled in the fields, select **Apply** and press **Enter**. This saves your network settings.

- b. For this document, the **eth0** device will be used to set up the management network. Select it and press **Enter** to access the interface configuration menu. Fill in the required fields as shown in the following example.

```

RHEV Hypervisor 6.3-20120920.0.rhev31.auto92
atlantic.demo.redhat.com

Interface: eth0 Driver: virtio_net
Protocol: Disabled Vendor: Red Hat Inc Virtio network
Link Status: Active MAC Address: 52:54:00:7b:0a:d6

IPv4 Settings
[ ] Disabled [ ] DHCP [*] Static
IP Address: 192.168.0.5 Netmask: 255.255.255.0
Gateway: 192.168.0.254

VLAN ID:
<Apply> <Back> <Reset>

```

**Figure 1.3. Configure management network interface**

Under **IPv4 Settings**, select **DHCP** or **Static** IP addressing and press **Spacebar** to mark the option as enabled. If using static IP addressing you must also provide the **IP Address**, **Netmask**, and **Gateway**. Select **Apply** and press **Enter**.

A dialog prompts you to confirm your network settings, select **OK** and press **Enter**.

c. Select the **oVirt Engine** tab. Configure the following options:

- ✳ In the **Management Server** field, enter **rhev.m.demo.redhat.com**.
- ✳ In the **Management Server Port** field, enter **443**.
- ✳ Tick the **Connect to the RHEV Manager and Validate Certificate** check box.
- ✳ The **Set RHEV-M Admin Password** field allows you to specify the root password for the hypervisor, and enable SSH password authentication from the Red Hat Enterprise Virtualization Manager. You do not have to fill in this field for this document.

Select **Apply** and press **Enter**. A dialog displays, asking you to connect the hypervisor to the Red Hat Enterprise Virtualization Manager and validate its certificate. Select **Approve** and press **Enter**. A message will display notifying you that the Manager configuration has been successfully updated.

- d. Under the **Red Hat Network** tab, you can register the host with the Red Hat Network. This enables the host to run Red Hat Enterprise Linux virtual machines with proper RHN entitlements. However, for the purposes of this document, the evaluation subscriptions will be used for the guests.
- e. Accept all other default settings.
- f. Finally, select the **Status** tab. Select **Restart** and press **Enter** to reboot the host and apply all changes.

You have now successfully installed a Red Hat Enterprise Virtualization Hypervisor. Repeat the above steps for each hypervisor you wish to use. The following sections will provide instructions on how to approve the hypervisors for use with the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

## 1.4. Connecting to Red Hat Enterprise Virtualization Web Administration Portal



**Figure 1.4. Connect to the Manager administration portal**

Now that you have installed the Red Hat Enterprise Virtualization Manager and hosts, you can log in to the Manager administration portal to start configuring your virtualization environment. Use a client running Firefox to access the web-based administration portal.

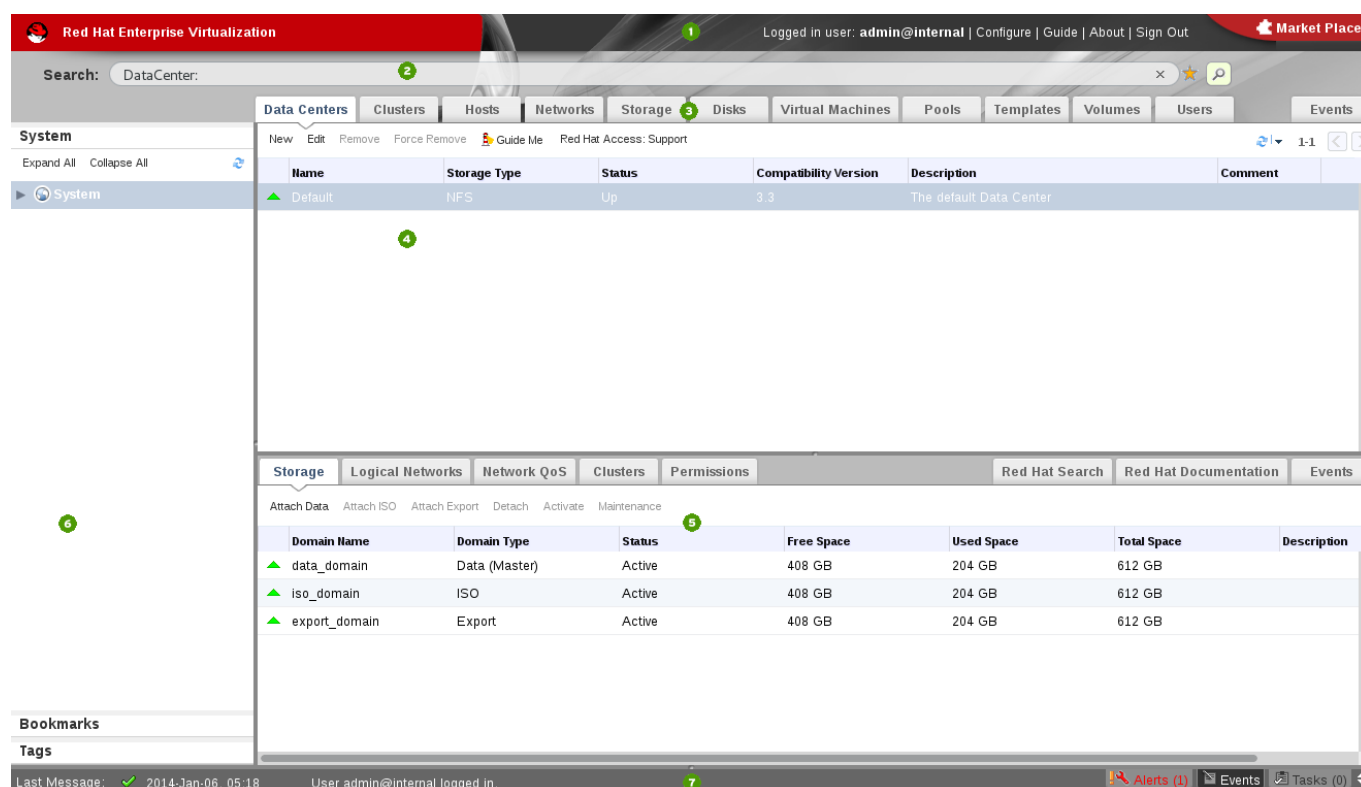
1. Open a browser and navigate to **https://domain.example.com**. Substitute **domain.example.com** with the URL provided during installation.
2. Under the **Portals** heading, click **Web Admin Portal**
3. If this is your first time connecting to the administration portal, Red Hat Enterprise Virtualization Manager will issue security certificates for your browser. Click the link labeled **this certificate** to trust the **ca.cer** certificate. A pop-up displays, click **Open** to launch the **Certificate** dialog. Click **Install Certificate** and select to place the certificate in **Trusted Root Certification Authorities** store.
4. The portal login screen displays. Enter **admin** as your **User Name**, and enter the **Password** that you provided during installation. Ensure that your domain is set to **Internal**. Click **Login**.

You have now successfully logged in to the Red Hat Enterprise Virtualization web administration portal. Here, you can configure and manage all your virtual resources.

[Report a bug](#)

## 1.5. 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.



**Figure 1.5. User Interface Elements of the Administration Portal**

### 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](#)

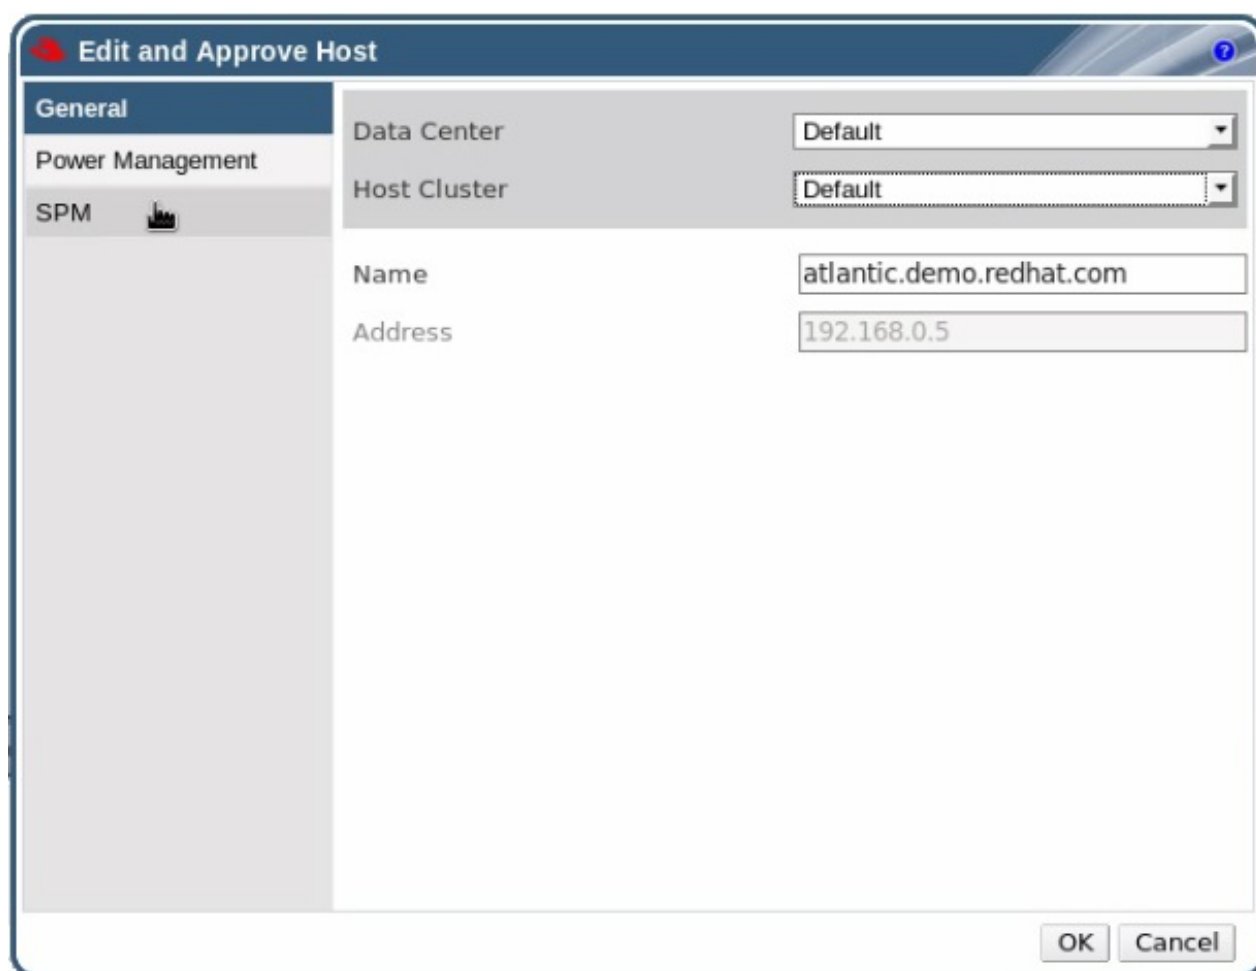
## 1.6. Approve the Red Hat Enterprise Virtualization Hypervisor

At this point you should already have a Default data center and a Default cluster, which have been automatically created during the Manager installation. In addition, the Red Hat Enterprise Virtualization Hypervisors you installed earlier should have been automatically detected by the Red Hat Enterprise Virtualization Manager and attached to the Default cluster of the Default data center.

However, before they can be used, they require a click of approval from the administration portal. Perform the following procedure for each hypervisor.

### To approve the Red Hat Enterprise Virtualization Hypervisor hosts

1. Navigate to the **Tree** pane and click the **Expand All** button. Under the **Default** cluster, click the **Hosts** icon. The **Hosts** tab displays a list of available hypervisors.
2. Select your hypervisor and click the **Approve** button. The **Edit and Approve Host** dialog displays. Accept the defaults or make changes as necessary, then click **OK**.



**Figure 1.6. Approve Red Hat Enterprise Virtualization Hypervisor**

3. A dialog appears, indicating that you have not configured Power Management for this host. For the purpose of this lab, click **OK** to continue. The host goes through a brief installation cycle. When complete, the host status changes from **Non Operational** to **Up**.

Note that both the Red Hat Enterprise Virtualization Hypervisors that you have approved were attached to the same host cluster, which means that they share the same network infrastructure, the same storage and the same type of CPU, therefore they can migrate virtual machines from one to the other. You can learn how to create new host clusters in Advanced Lab 7 - Add Additional Data Center.

Now that you have finished configuring your physical servers for use as the Manager, Hypervisors and administration portal client respectively, you are ready to customize and deploy virtual resources including logical networks, storage domains and virtual machines.

[Report a bug](#)

## 1.7. Configure Logical Networks

Now that you have a data center with hosts grouped in a cluster, you need to define and apply the

networking layer. When you installed the system, a management network was already defined. However new networks, for example data, storage or display can be added to enhance network speed and performance. In addition, other networks can be used to segregate virtual machine traffic from the management networks, or isolate traffic between groups of virtual machines in the same cluster. In Red Hat Enterprise Virtualization Manager, network definition, type and function are encapsulated in a logical entity called a **Logical Network**.

A logical network is assigned as a required resource of a cluster in a data center, and by extension all hosts in a cluster must have the same set of logical networks implemented. The implementation itself may vary from host to host (IP and bonding properties). Therefore, to configure a network, you need to first define the network and then apply this network to each host. By default the management network (rhevms) is defined for a data center.

In this lab, you will create an additional storage network and add it to your hosts. In the following example you will define additional network for the storage, which will be useful when using NAS storage like NFS or iSCSI.

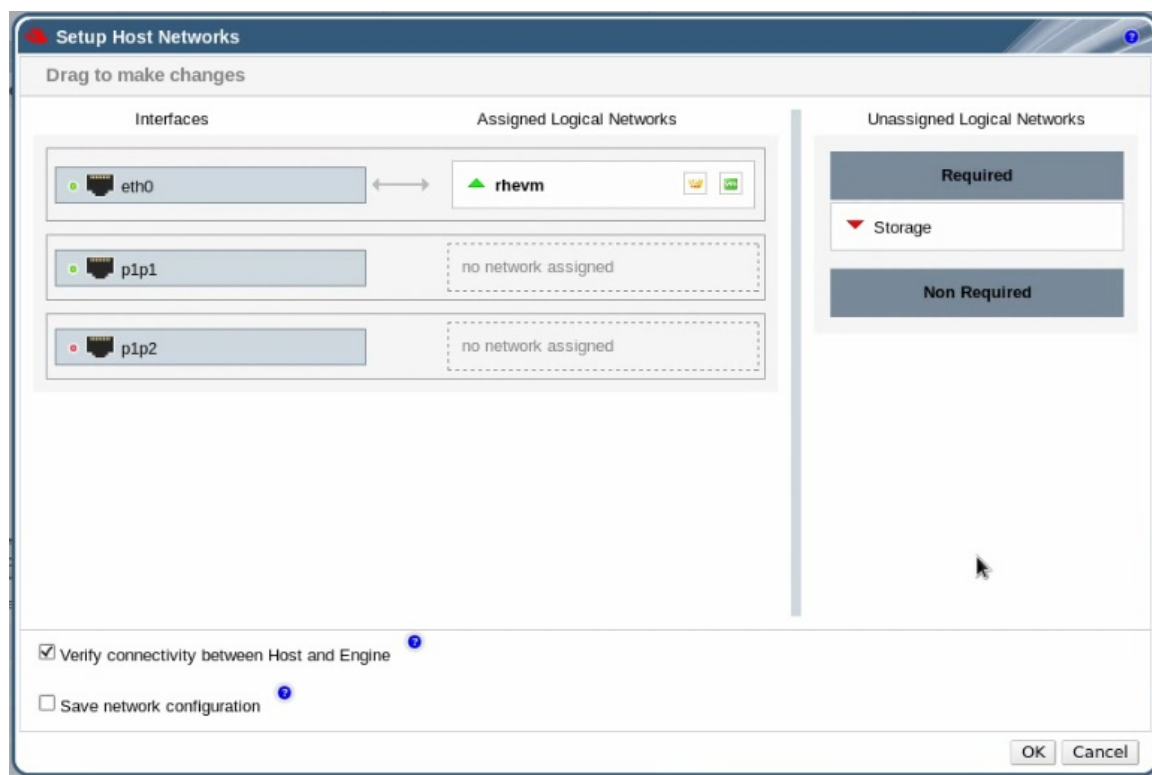
### Defining Logical Networks in a Cluster

1. Navigate to the **Tree** pane and click the **Expand All** button. Under System, click **Default**. On the results list, the **Default** data center displays.
2. On the details pane, select the **Logical Networks** subtab. This displays the existing logical networks. At this stage only the default **rhevms** network is listed.
3. Click **New**. The **New Logical Network** dialog displays. Fill in the **Name** and **Description** fields, clear the **VM network** check box, and select the **Attach** check box under **Attach/Detach Network to/From Cluster(s)** to add the Storage network to the Default data center.
4. Click **OK** to create the new logical network.

Now that you have defined this network as a resource required by the default cluster in the data center, it is time to add this resource to the hosts in the cluster.

### Adding a Network to a Host

1. Back on the **Tree** pane, click **Default** → **Clusters** → **Default** → **Hosts**. The **Hosts** tab displays a list of available hosts.
2. For each of your installed hosts, perform the following tasks:
  - a. Click on the host. On the details pane, select the **Network Interfaces** tab.
  - b. A list of network interfaces available for this host displays. One of them will already have the management network (rhevms) configured.
  - c. Select the interface on which to configure the Storage network and click the **Setup Host Networks** button. The **Setup Host Networks** dialog displays.



**Figure 1.7. Setup Host Networks interface**

Configure the following options:

- ✧ Drag the network "Storage" from the **Unassigned Logical Networks** column to the **Assigned Logical Networks** column and drop it. The network called "Storage" is now associated with the virtual network interface listed in the row to the left of where you dropped the logical network.
- ✧ Click the pencil-shaped icon on the "Storage" network in the **Assigned Logical Networks** column.
- ✧ Select the **Static** radio button. Enter the **IP** and **Subnet Mask** you have prepared as part of the prerequisites to this lab.
- ✧ Select the **Save network configuration** check box.

d. Click **OK**.

You have now added a new storage network to your data center, and attached the network to your hosts. On the **Logical Networks** tab of the **Default** data center, you should have at least two networks - rhevm and storage. Now, you can add storage to the system.

[Report a bug](#)

## 1.8. Configuring Storage

After configuring your logical networks, you need to add storage to your data center.

Red Hat Enterprise Virtualization uses a centralized shared storage system for virtual machine disk images and snapshots. Storage can be implemented using Network File System (NFS), Internet Small Computer System Interface (iSCSI) or Fibre Channel Protocol (FCP). Storage definition, type and function, are encapsulated in a logical entity called a **Storage Domain**. Multiple storage domains can be used in a Red Hat Enterprise Virtualization environment.





## Note

Red Hat Enterprise Virtualization 3.1 supports passing mount options for all POSIX-compliant filesystems. This includes GPFS, GFS, pNFS, and Gluster. For more information, see *Red Hat Enterprise Virtualization 3.1 Administration Guide*.

For this lab you will use two types of storage domains. The first is an NFS share for ISO images of installation media. You have already created this ISO domain during the Red Hat Enterprise Virtualization Manager installation.

The second storage domain will be used to hold virtual machine disk images. For this domain, you need at least one of the supported storage types. You have already set a default storage type during installation. Ensure that you use the same type when creating your data domain.

### Select your next step by checking the storage type you should use:

1. Navigate to the **Tree** pane and click the **Expand All** button. Under System, click **Default**. On the results list, the **Default** data center displays.
2. On the data center entry, the **Storage Type** column displays the type you should add.
3. Now that you have verified the storage type, create the storage domain:
  - ✦ For NFS storage, see **Creating an NFS Data Domain**.
  - ✦ For iSCSI storage, see **Creating an iSCSI Data Domain**.
  - ✦ For FCP storage, see **Creating an FCP Data Domain**.

### See Also:

- ✦ [Section 1.9.1, “Creating an NFS Data Domain”](#)
- ✦ [Section 1.9.2, “Creating an iSCSI Data Domain”](#)
- ✦ [Section 1.9.3, “Creating an FCP Data Domain”](#)

[Report a bug](#)

## 1.9. Configuring Storage

### 1.9.1. Creating an NFS Data Domain

Because you have selected NFS as your default storage type during the Manager installation, you will now create an NFS storage domain. An NFS type storage domain is a mounted NFS share that is attached to a data center and used to provide storage for virtual machine disk images.



## Important

If you are using NFS storage, you must first create and export the directories to be used as storage domains from the NFS server. These directories must have their numerical user and group ownership set to 36:36 on the NFS server, to correspond to the vdsmd user and kvm group respectively on the Red Hat Enterprise Virtualization Manager server. In addition, these directories must be exported with the read write options (rw). For more information see the *Red Hat Enterprise Virtualization Installation Guide*.

### To add NFS storage

1. Navigate to the **Tree** pane and click the **Expand All** button. Under System, select the **Default** data center and click on **Storage**. The available storage domains display on the results list. Click **New Domain**.
2. The **New Domain** window displays.

**New Domain**

Name:  Description:

Data Center: **Default (NFS)** Comment:

Domain Function / Storage Type: **Data / NFS** Format: **V3**

Use Host:

Export Path:   
Remote path to NFS export, takes either the form: FQDN:/path or IP:/path e.g. server.example.com:/export/VMs

☒ Advanced Parameters

**\* It is recommended to keep the default values in the fields below unchanged.**

☐ Override Default Options

NFS Version: **V3 (default)**

Retransmissions (#):

Timeout (deciseconds):

OK Cancel

**Figure 1.8. Add NFS Storage - New Domain Window**

Configure the following options:

- a. **Name:** Enter **NFS-share**.
  - b. **Data Center:** The **Default** data center is already pre-selected.
  - c. **Domain Function / Storage Type:** The **Data/ NFS** option is already pre-selected because during installation you set NFS as your data center's default storage type. The storage domain types which are not compatible with the Data Center will not be available.
  - d. **Use Host:** Select any of the hosts from the drop down menu. Only hosts which belong in this data center will display in this list.
  - e. **Export Path:** Enter the IP address or a resolvable hostname of the chosen host. The export path should be in the format of **192.168.0.10 : /Images/NFS-Share** or **domain.example.com : /Images/NFS-Share**.
3. Click **OK**. The new **NFS-share** 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.

You have created an NFS storage domain. Now, you need to attach an ISO domain to the data center and upload installation images so you can use them to create virtual machines.

[Report a bug](#)

### 1.9.2. Creating an iSCSI Data Domain

Because you have selected iSCSI as your default storage type during the Manager installation, you will now create an iSCSI storage domain. Red Hat Enterprise Virtualization platform supports iSCSI storage domains spanning multiple pre-defined Logical Unit Numbers (LUNs).

#### To add iSCSI storage

1. Navigate to the **Tree** pane and click the **Expand All** button. Under System, select the **Default** data center and click on **Storage**. The available storage domains display on the results list. Click **New Domain**.
2. The **New Domain** dialog box displays.

**New Domain**

Name:

Data Center:

Domain Function / Storage Type:  Format:

Use Host:

**Discover Targets**

Address:

Port:

☐ User Authentication:

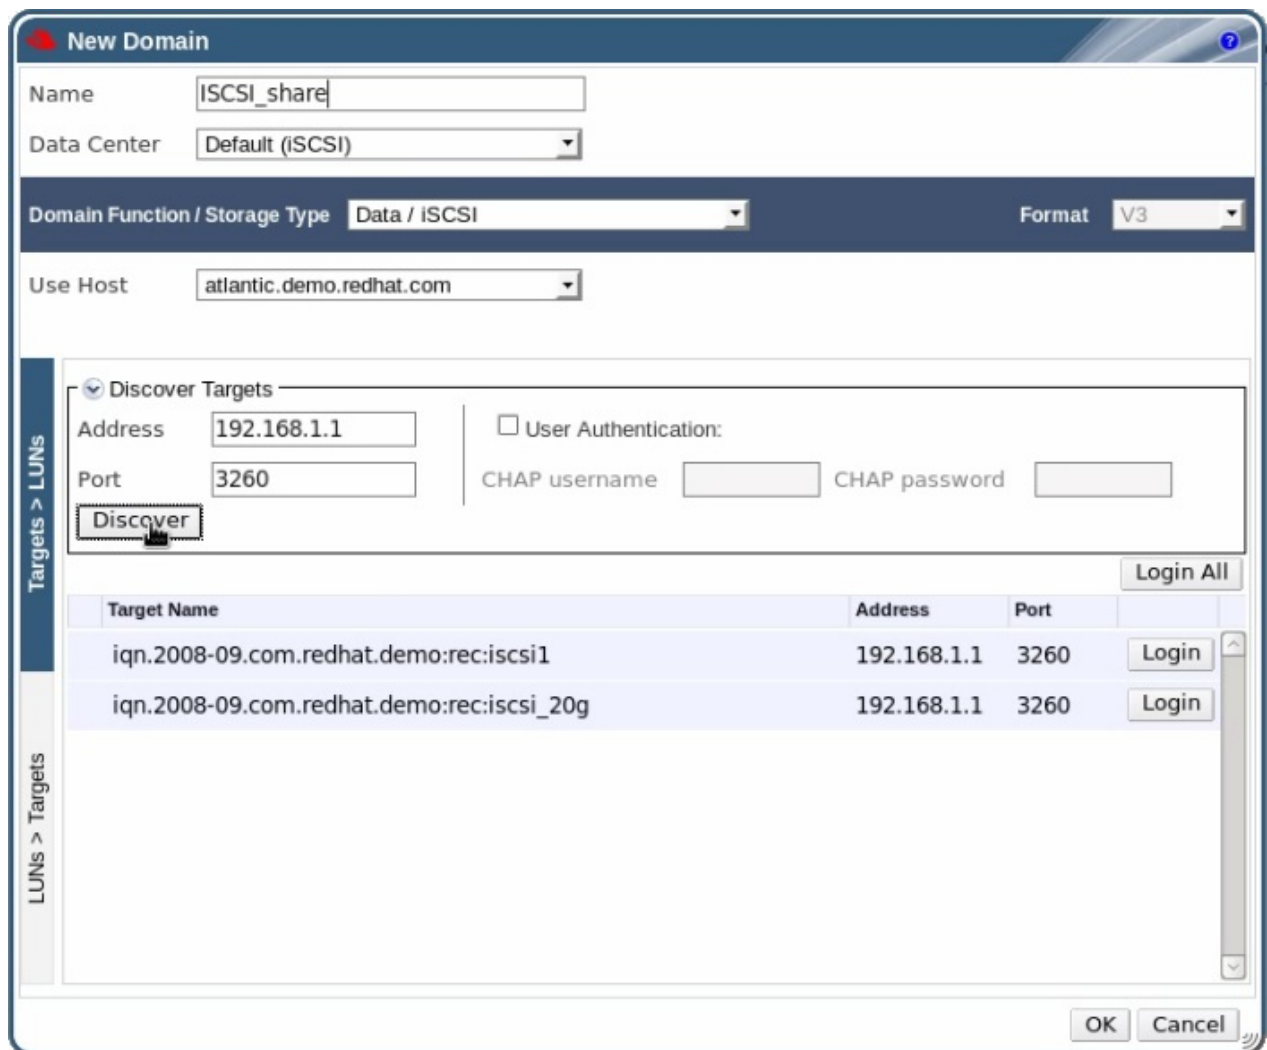
CHAP username:  CHAP password:

Target Name	Address	Port
-------------	---------	------

**Figure 1.9. Add iSCSI Storage**

Configure the following options:

- a. **Name:** Enter **ISCSI-share**.
  - b. **Data Center:** The **Default** data center is already pre-selected.
  - c. **Domain Function / Storage Type:** The **Data/ iSCSI** option is already pre-selected because during installation you set iSCSI as your data center's default storage type. The storage domain types which are not compatible with the Data Center will not be available.
  - d. **Use Host:** Select any of the available hosts from the drop down menu. Only hosts which belong in this data center will display in this list.
3. To connect to the iSCSI target, enter the required information under the **Discover Targets** bar.

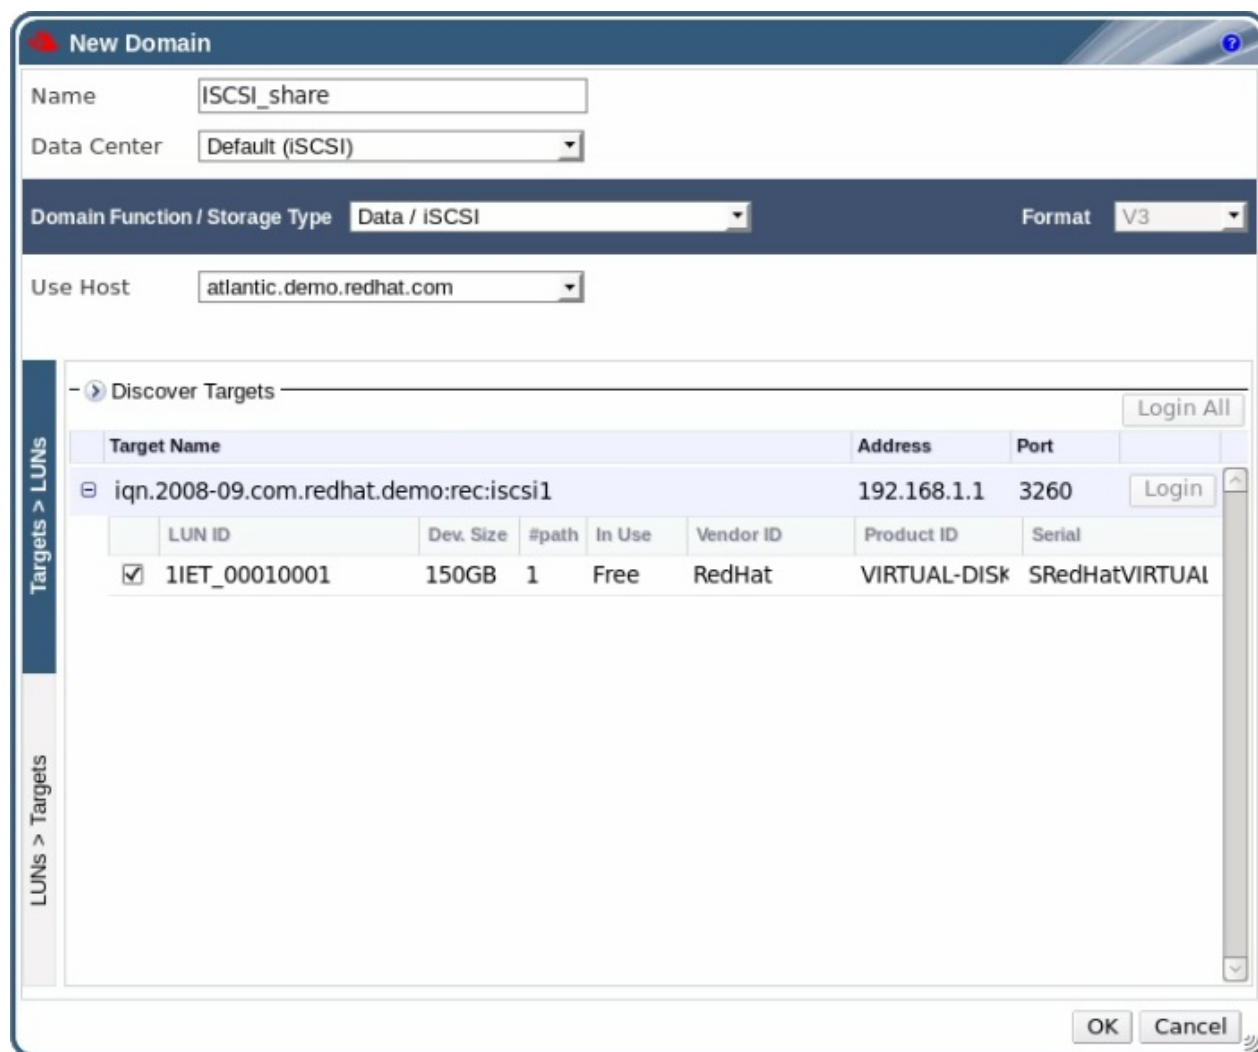


**Figure 1.10. Discover iSCSI target**

- a. **Address:** Enter the address of the iSCSI target.
- b. **Port:** Select the port to connect to. The default is 3260.
- c. **User Authentication:** If required, enter the user name and password.

Click the **Discover** button to find the targets. The iSCSI targets display in the results list with a **Login** button for each target.

4. Click **Login** on the first target to display the list of existing LUNs. Click the + icon under the Target Name to expand the LUN list. Tick the **Add LUN** check box to use the selected LUN as the iSCSI data domain. LUNs that are part of a storage domain in the current setup are disabled. LUNs used by the host (that is, LUNs that are either part of a volume group or that are used as partitions by other devices) display as being in use. You can choose LUNs that are in use, but you will have to forcefully override their contents.



**Figure 1.11. Attach LUNs to iSCSI domain**

- Click **OK**. The new **iSCSI-share** 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.

You have created an iSCSI storage domain. Now, you need to attach an ISO domain to the data center and upload installation images so you can use them to create virtual machines.

[Report a bug](#)

### 1.9.3. Creating an FCP Data Domain

Because you have selected FCP as your default storage type during the Manager installation, you will now create an FCP storage domain. Red Hat Enterprise Virtualization platform supports FCP storage domains spanning multiple pre-defined Logical Unit Numbers (LUNs).

#### To add FCP storage

- Navigate to the **Tree** pane and click the **Expand All** button. Under System, select the **Default** data center and click on **Storage**. The available storage domains display on the results list. Click **New Domain**.
- The **New Domain** dialog box displays.

**New Domain**

Name

Data Center

Domain Function / Storage Type  Format

Use Host

Export Path

Remote path to NFS export, takes either the form: FQDN:/path or IP:/path e.g. server.example.com:/export/VMs

[Advanced Parameters](#)

OK Cancel

**Figure 1.12. Add FCP Storage**

Configure the following options:

- a. **Name:** Enter **FCP - share**.
- b. **Data Center:** The **Default** data center is already pre-selected.
- c. **Domain Function / Storage Type:** The **Data/ Fibre Channel** option is already pre-selected because during installation you set FC as your data center's default storage type. The storage domain types which are not compatible with the Data Center will not be available.
- d. **Use Host:** Select any of the hosts from the drop down menu. Only hosts which belong in this data center will display in this list. Ensure that the LUN you intend to use is available on the host you select. LUNs that are part of a storage domain in the current setup are disabled. LUNs used by the host (that is, LUNs that are either part of a volume group or that are used as partitions by other devices) display as being unavailable. You can choose LUNs that are in use, but you will have to forcefully override their contents.
- e. The list of existing LUNs display. On the selected LUN, select the **Add LUN** check box to use it as the FCP data domain.

3. Click **OK**. The new **FCP - share** 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.

You have created an FCP storage domain. Now, you need to attach an ISO domain to the data center and upload installation images so you can use them to create virtual machines.

[Report a bug](#)

## 1.10. Attach and Populate ISO Domains

You have defined your first storage domain to store virtual guest data, now it is time to configure your second storage domain, which will be used to store installation images for creating virtual machines. You have already created an ISO domain during the installation of the Red Hat Enterprise Virtualization Manager. To use this ISO domain, attach it to the same data center which contains the local storage domain.

### To attach the ISO domain

1. Navigate to the **Tree** pane and click the **Expand All** button. Click the **Atlantic-Local** data center. On the results list, the **Atlantic-Local** data center displays.
2. On the details pane, select the **Storage** tab and click the **Attach ISO** button.
3. The **Attach ISO Library** dialog appears with the available ISO domain. Select the **local-iso-share** domain and click **OK**.



Figure 1.13. Attach ISO Library

4. The ISO domain appears in the results list of the **Storage** tab. It displays with the *Locked* status as the domain is being validated before becoming *Active*.



Media images (CD-ROM or DVD-ROM in the form of ISO images) must be available in the ISO repository for the virtual machines to use. To do so, Red Hat Enterprise Virtualization provides a utility that copies the images and sets the appropriate permissions on the file. For this lab, both the file provided to the utility and the ISO share have to be accessible from the Red Hat Enterprise Virtualization Manager.

Log in to the Red Hat Enterprise Virtualization Manager server console to upload images to the ISO domain.

### To upload ISO images

1. Create or acquire the appropriate ISO images from boot media. Ensure the path to these images is accessible from the Red Hat Enterprise Virtualization Manager server.
2. The next step is to upload these files. First, determine the available ISO domains by running:

```
# engine-iso-uploader list
```

You will be prompted to provide the admin user password which you are using to connect to the administration portal. The tool lists the name of the ISO domain that you have already attached in the previous lab:

```
ISO Storage Domain List:
local-iso-share
```

Now you have all the information required to upload the required files. To copy your installation images to the ISO domain, run:

```
# engine-iso-uploader upload -i local-iso-share [file1] [file2] ....
[fileN]
```

3. After the images have been loaded, check that they are available for use in the Manager administration portal.
  - a. Navigate to the **Tree** and click the **Expand All** button. Click **Storage**
  - b. On the **Storage** tab, click **local-iso-share** to display its details pane.
  - c. Select the **Images** subtab. The list of available images should be populated with the files which you have uploaded. In addition, the **RHEV-toolsSetup.iso** and **virtio-win.vfd** images should have been automatically uploaded during installation.


Now that you have successfully prepared the ISO domain for use, you have completed Lab 1 and are ready to start creating virtual machines.

[Report a bug](#)

## 1.11. Lab 1 - summary

You have reached your first Track A goal of setting up and configuring the infrastructure required to create and run virtual machines. You have successfully installed the Red Hat Enterprise Virtualization Manager and Red Hat Enterprise Virtualization Hypervisors, defined logical networks for storage devices, attached storage domains to the data center, and prepared ISO images.

The next lab on Track A teaches you how to create Red Hat Enterprise Linux virtual machines and templates.



[Report a bug](#)

## Chapter 2. Lab 2 - Creating virtual machines

### 2.1. Lab 2 - Objectives

This lab takes you through the tasks necessary to create Red Hat Enterprise Linux virtual machines. In addition, you will learn how to create and use virtual machine templates. This lab should take you about 25 minutes.

The first section shows you how to create a new Red Hat Enterprise Linux virtual machine, configure storage and networking, and install the operating system. (10 minutes)

The second section shows you how to use the Red Hat Enterprise Linux virtual machine as a basis to create a template. (10 minutes)

Finally, you will learn how to clone a virtual machine from the Red Hat Enterprise Linux template. (5 minutes)

#### See Also:

- ✳ [Section 2.3, “Creating Red Hat Enterprise Linux templates”](#)
- ✳ [Section 2.4, “Creating a Template from a Red Hat Enterprise Linux Virtual Machine”](#)
- ✳ [Section 2.5, “Sealing a Linux Virtual Machine Manually for Deployment as a Template”](#)
- ✳ [Section 2.7, “Cloning a Red Hat Enterprise Linux virtual machine from a template”](#)

[Report a bug](#)

### 2.2. Creating a Red Hat Enterprise Linux virtual machine

In your current configuration, you should have at least one host available for running virtual machines, and the required installation images in your ISO domain. This section guides you through the creation of a Red Hat Enterprise Linux 6 virtual server. You will perform a normal attended installation using a virtual DVD.

You have now created your first Red Hat Enterprise Linux virtual machine. Before you can use your virtual machine, install an operating system on it.

#### To create a Red Hat Enterprise Linux server

1. Navigate to the **Tree** pane and click **Expand All**. Click the **VMs** icon in the **Default** cluster under the **Default** data center. On the **Virtual Machines** tab, click **New VM**.

**New Virtual Machine** ?

**General**

Cluster: Default/Default

Based on Template: Blank

Operating System: Other OS

Optimized for: Server

Name:

Description:

Comment:

☐ Stateless ☐ Start in Pause Mode ☐ Delete Protection

VM has no network interfaces. To add one, assign a profile.

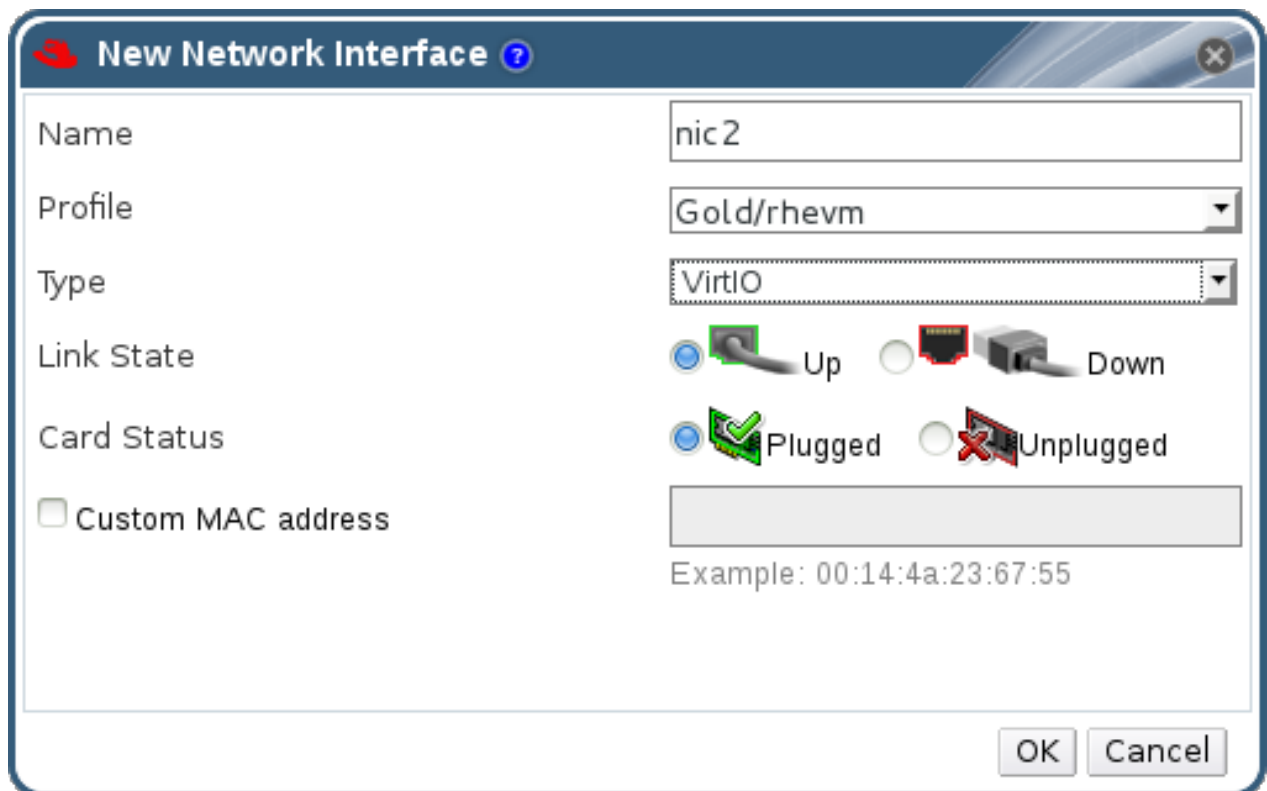
nic1:

Show Advanced Options OK Cancel

**Figure 2.1. Create new Red Hat Enterprise Linux server**

Fill in the **Name** field, set the **Memory Size** to **1 GB**, and set the **Operating System** to **Red Hat Enterprise Linux 6.x**. You may alter other settings but this example retains the defaults. Click **OK** to create the virtual machine.

2. A **New Virtual Machine - Guide Me** window opens. This allows you to define networks for the virtual machine. Click **Configure Network Interfaces**.



**Figure 2.2. New Network Interface configurations**

Retain all the default settings and click **OK**.



### Note

Red Hat VirtIO is the default selection for Red Hat Enterprise Linux virtual machines. This para-virtualizes network traffic to the guest and improves performance.

3. You are returned to the **Guide Me** window. This time, click **Configure Virtual Disks** to add storage to the virtual machine.

**Add Virtual Disk**

☒ Internal    ☐ External (Direct Lun)

Size(GB)

Alias

Description

Interface

Format

Data Center

Storage Domain

☐ Wipe after delete

☐ Is bootable

☐ Is shareable

OK Cancel

**Figure 2.3. New Virtual Disk configurations**

In the **Size (GB)** field, enter *15*. Retain all other default settings and click **OK**.

- Click **Configure Later** to close the Guide Me window. Your new Red Hat Enterprise Linux 6 virtual machine displays in the **Virtual Machines** tab.

You can now connect to your Red Hat Enterprise Linux 6 virtual machine and start using it.

[Report a bug](#)

## 2.3. Creating Red Hat Enterprise Linux templates

Now that you have created a Red Hat Enterprise Linux virtual machine, you can save its settings into a template. This template will retain the original virtual machine's configurations, including virtual disk and network interface settings, operating systems and applications. You can use this template to rapidly create replicas of the original virtual machine.

Before your virtual machine can be used to create a template, it has to be sealed . This ensures that machine-specific settings are not propagated through the template.

You have prepared a virtual machine template for use. You can now clone Red Hat Enterprise Linux virtual machines using this template.

[Report a bug](#)

## 2.4. Creating a Template from a Red Hat Enterprise Linux Virtual Machine

### Creating a Template from a Red Hat Enterprise Linux Virtual Machine

1. Back in the administration portal, click the **VMs** icon on the **Virtual Machines** tab, select the sealed Red Hat Enterprise Linux 6 virtual machine. Ensure that its status is **Down**.
2. Click **Make Template**. The **New Virtual Machine Template** displays.

**New Template** ?

Name

Description

Comment

Cluster

Disks Allocation:

Alias	Virtual Size	Target
<input type="text" value="Disk1"/>	50 GB	<input type="text" value="storage_data_z"/>

☒ Allow all users to access this Template

☐ Copy VM permissions

OK Cancel

**Figure 2.4. Make new virtual machine template**

Enter a suitable **Name** and **Description** for the template. Retain all the other default settings and click **OK**.

3. On the **Tree** pane, click **Templates**. On the **Templates** tab, the template displays the "Image Locked" status icon while it is being created. During this time, the action buttons for the template remain disabled. Once created, the action buttons are enabled and the template is ready for use.

You have prepared a virtual machine template for use. You can now clone Red Hat Enterprise Linux virtual machines using this template.

[Report a bug](#)

## 2.5. Sealing a Linux Virtual Machine Manually 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 2.1. Sealing a Linux Virtual Machine

1. Log in to the virtual machine. Flag the system for re-configuration by running the following command as root:

```
# touch /.unconfigured
```

2. Remove ssh host keys. Run:

```
# rm -rf /etc/ssh/ssh_host_*
```

3. Set **HOSTNAME=localhost.localdomain** in **/etc/sysconfig/network**

4. Remove **/etc/udev/rules.d/70-\***. Run:

```
# rm -rf /etc/udev/rules.d/70-*
```

5. Remove the **HWADDR=** line from **/etc/sysconfig/network-scripts/ifcfg-eth\***.
6. Optionally delete all the logs from **/var/log** and build logs from **/root**.
7. 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](#)

## 2.6. Sealing a Linux Virtual Machine for Deployment as a Template using sys-unconfig



## Summary

Generalize (seal) a Linux virtual machine using the **sys-unconfig** command before making it into a template. This prevents conflicts between virtual machines deployed from the template.

### Procedure 2.2. Sealing a Linux Virtual Machine using sys-unconfig

1. Log in to the virtual machine.
2. Remove ssh host keys. Run:

```
# rm -rf /etc/ssh/ssh_host_*
```

3. Set **HOSTNAME=localhost.localdomain** in **/etc/sysconfig/network**
4. Remove the **HWADDR=** line from **/etc/sysconfig/network-scripts/ifcfg-eth\***.
5. Optionally delete all the logs from **/var/log** and build logs from **/root**.
6. Run the following command:

```
# sys-unconfig
```

## Result

The virtual machine shuts down; it is now 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](#)

## 2.7. Cloning a Red Hat Enterprise Linux virtual machine from a template

1. Navigate to the **Tree** pane and click **Expand All**. Click the **VMs** icon in the **Default** cluster under the **Default** data center. On the **Virtual Machines** tab, click **New VM**.

**New Virtual Machine** ?

**General**

Cluster: Default/Default

Based on Template: Blank

Operating System: Other OS

Optimized for: Server

Name:

Description:

Comment:


☐ Stateless ☐ Start in Pause Mode ☐ Delete Protection

VM has no network interfaces. To add one, assign a profile.

nic1:

Show Advanced Options OK Cancel

**Figure 2.5. Clone a Linux server**


- On the **General** tab, select your newly created Red Hat Enterprise Linux template from the **Based on Template** list.
  - Enter a suitable **Name** and **Description**, and accept the default values inherited from the template in the rest of the fields. You can change them if needed.
  - Click the **Resource Allocation** tab. On the **Provisioning** field, click the drop down menu and select the *Clone* option. Set the clone method under **Disk** to *Preallocated*.
- Retain all other default settings and click **OK** to create the virtual machine. The virtual machine displays in the Virtual Machines list with a status of "Image Locked" until the virtual disk is created. The virtual disk and networking settings are inherited from the template, and do not have to be reconfigured.
  - Click the **Run**  icon to turn it on. This time, the **Run Once** steps are not required as the operating system has already been installed onto the virtual machine hard drive. Click the green **Console** button to connect to the virtual machine.

[Report a bug](#)

## 2.8. Lab 2 - Summary

You have reached your second Track A and second Track B goals respectively to create and prepare virtual machines for use. In this lab, you have successfully created Red Hat Enterprise Linux virtual machines, installed operating systems, defined storage and networks for the virtual machines, and utilized templates to deploy virtual machines.

- The next lab on Track A teaches you how to activate virtual machine live migration for occasions when your hardware is experiencing downtime.
- The next lab on Track B teaches you how to create and manage virtual machines in the Power User Portal.



[Report a bug](#)

## Chapter 3. Lab 3 - Live migration scenarios

### 3.1. Lab 3 - Objectives

This lab takes you through the tasks of optimizing your host clusters so you can activate virtual machine live migration when necessary. This lab should take you about 10 minutes.


It shows you how to live migrate virtual machines from host to host. (2 minutes)

It demonstrates virtual machine automatic live migration when a host is placed in maintenance mode. (3 minutes)

It shows you how to define a cluster power saving policy, then demonstrates virtual machine live migration when a host is loaded beyond its threshold. (5 minutes)

[Report a bug](#)

### 3.2. Live migration scenarios



Live migration enables you to move a virtual machine to another host while the machine and its applications are still running. You can activate automatic live migration for occasions when you perform hardware maintenance or if your hosts become non-operational, and your virtual machines will be migrated with no interruptions or delay. At the end of this lab you will have configured cluster policies and be able to balance your hosts' workloads.

This lab is intended for Track A. It assumes that you have successfully completed Red Hat Enterprise Virtual Machine setup and that you have at least two virtual machines running on your hosts.

[Report a bug](#)

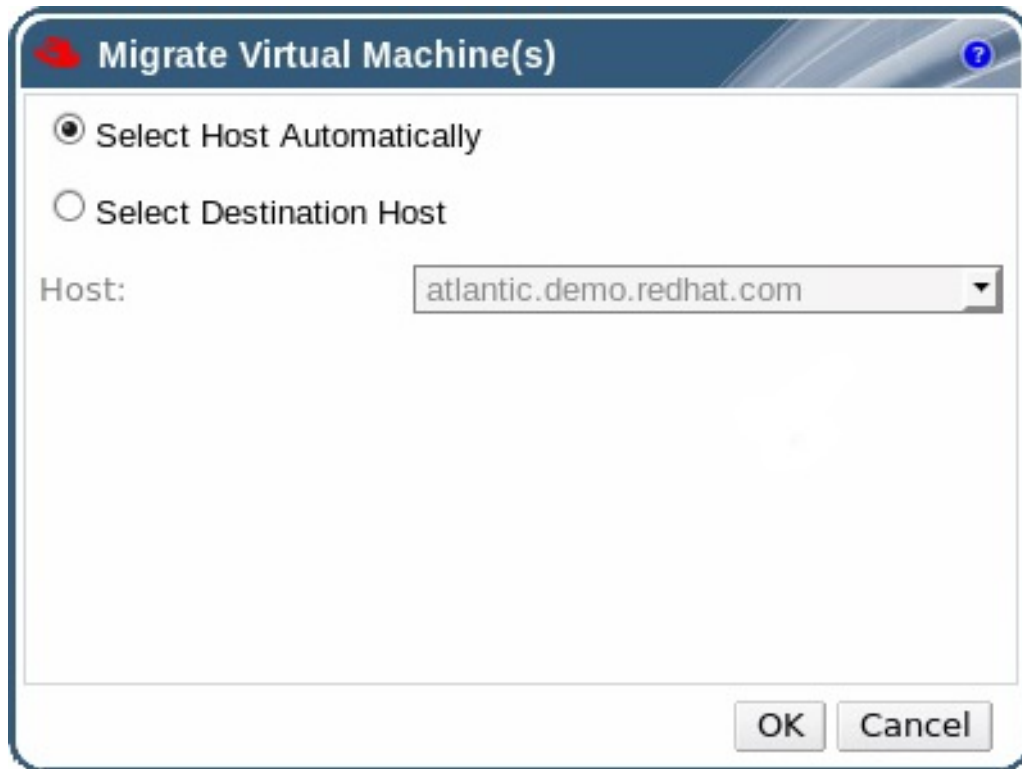
### 3.3. Activating live migration

Now that you have several virtual machines running on each host, you can configure live migration. This provides a backup for the virtual machines - if they are running mission critical workloads, and cannot be powered off, you need to ensure that they must always have a host available to run on. Live migration allows you to move virtual machines to different hosts should the host they are running on experience scheduled or unscheduled downtime, without causing the virtual machine's operations to be suspended.

In this example, you will see how virtual machines are live migrated from one host to any of the available hosts in the same cluster.

### To live migrate virtual machines

1. Navigate to the **Tree** pane and click **Expand All**. Under **Clusters**, click the **VMs** icon. The list of available virtual machines is displayed in the **Virtual Machines** tab.
2. Select one or more virtual machines and click **Migrate**. The **Migrate Virtual Machines** dialog displays.



**Figure 3.1. Live migrate a virtual machine to another host**

You can select a target host, or allow the system to automatically select one of the available hosts in the cluster. In this case, use the default option.

3. Click **OK**. The virtual machines will live migrate to another host in the cluster.

[Report a bug](#)

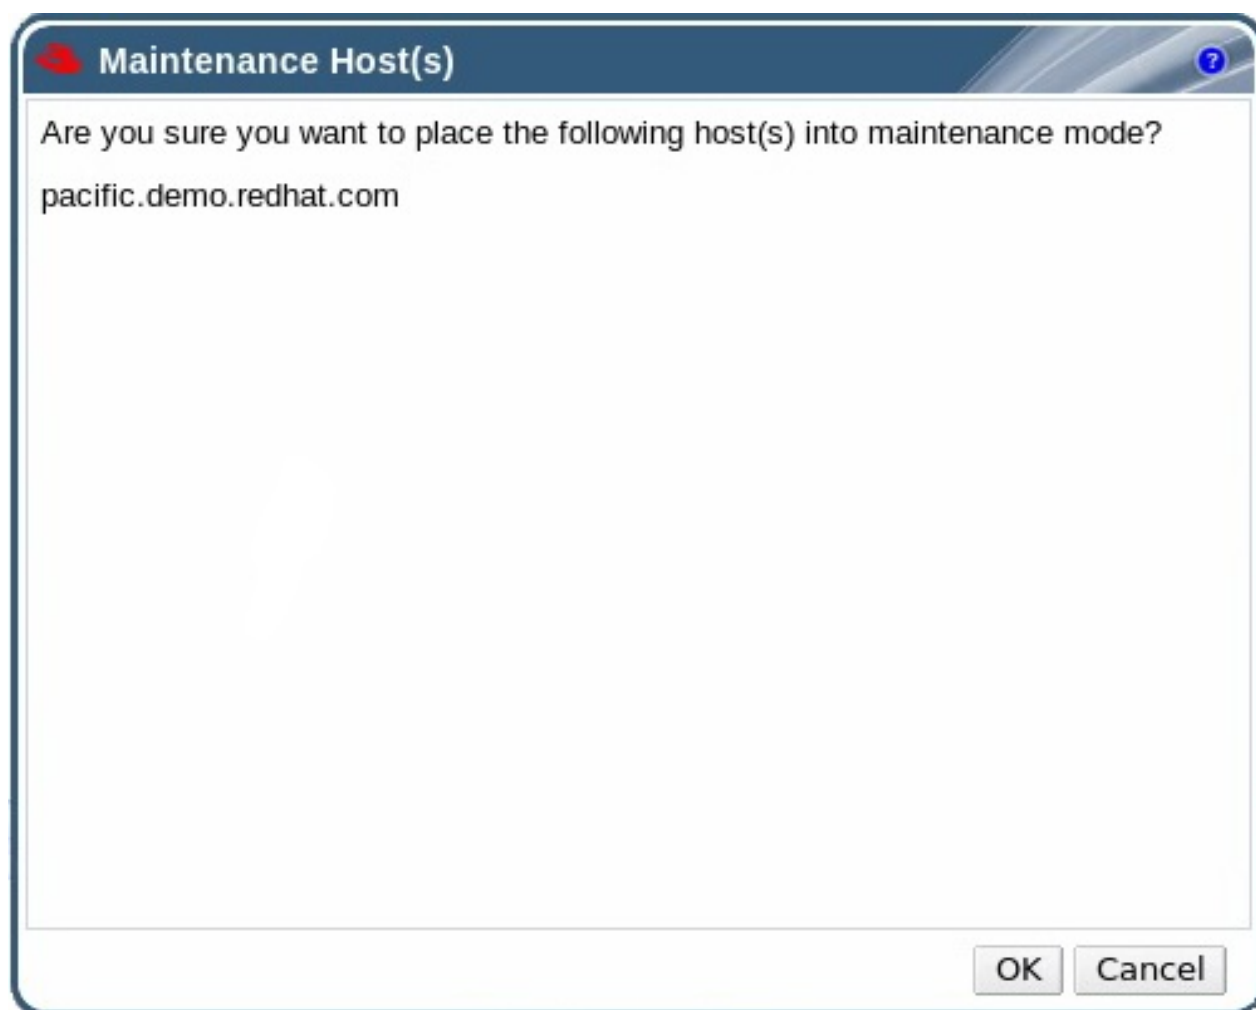
## 3.4. Moving hosts into maintenance mode

In the previous section you watched a user-triggered live migration of virtual machines. Manually activating live migration is useful for occasions when hosts are brought down for maintenance. The virtual machines on the host which is being placed into maintenance can be automatically live migrated to another host in the cluster, provided that power management has been properly configured for the hosts. More information on configuring power management can be found in [Advanced Lab 6 - High Availability Scenarios](#).

### To move a host into maintenance mode

1. Navigate to the **Tree** pane and click **Expand All**. Under **Clusters**, click the **Hosts** icon. The list of available hosts is displayed in the **Hosts** tab.

2. Select the host and click the **Maintenance** button. In this example, select the **Pacific** host.



**Figure 3.2. Move a host into maintenance**

3. Click **OK**. This migrates all virtual machines to alternative hosts, and places the **Pacific** host into maintenance. The **Status** field of the host changes to *Preparing for Maintenance*, followed by *Maintenance*. The icon changes to indicate that the host is in maintenance.
4. Now, it is safe to perform an upgrade, or any other schedule maintenance, on the host. After you have completed your maintenance, reactivate the host by selecting it and clicking the **Activate** button.

[Report a bug](#)

## 3.5. Defining cluster policies

Until now, you have worked without a specific cluster policy. The virtual machines were started on a host selected on round robin logic; as long as there are available resources on the host, virtual machines can be migrated onto it. Furthermore, there was no load balancing of running virtual machines. For example, a virtual machine will remain on the host it was run initially unless it is specifically moved, either by user triggered migration (maintenance mode or manual virtual machine migration) or by a system high availability event (this will be explained in the advanced High Availability lab). Cluster policy allows you to better control the scheduling of virtual machines and to perform load balancing.

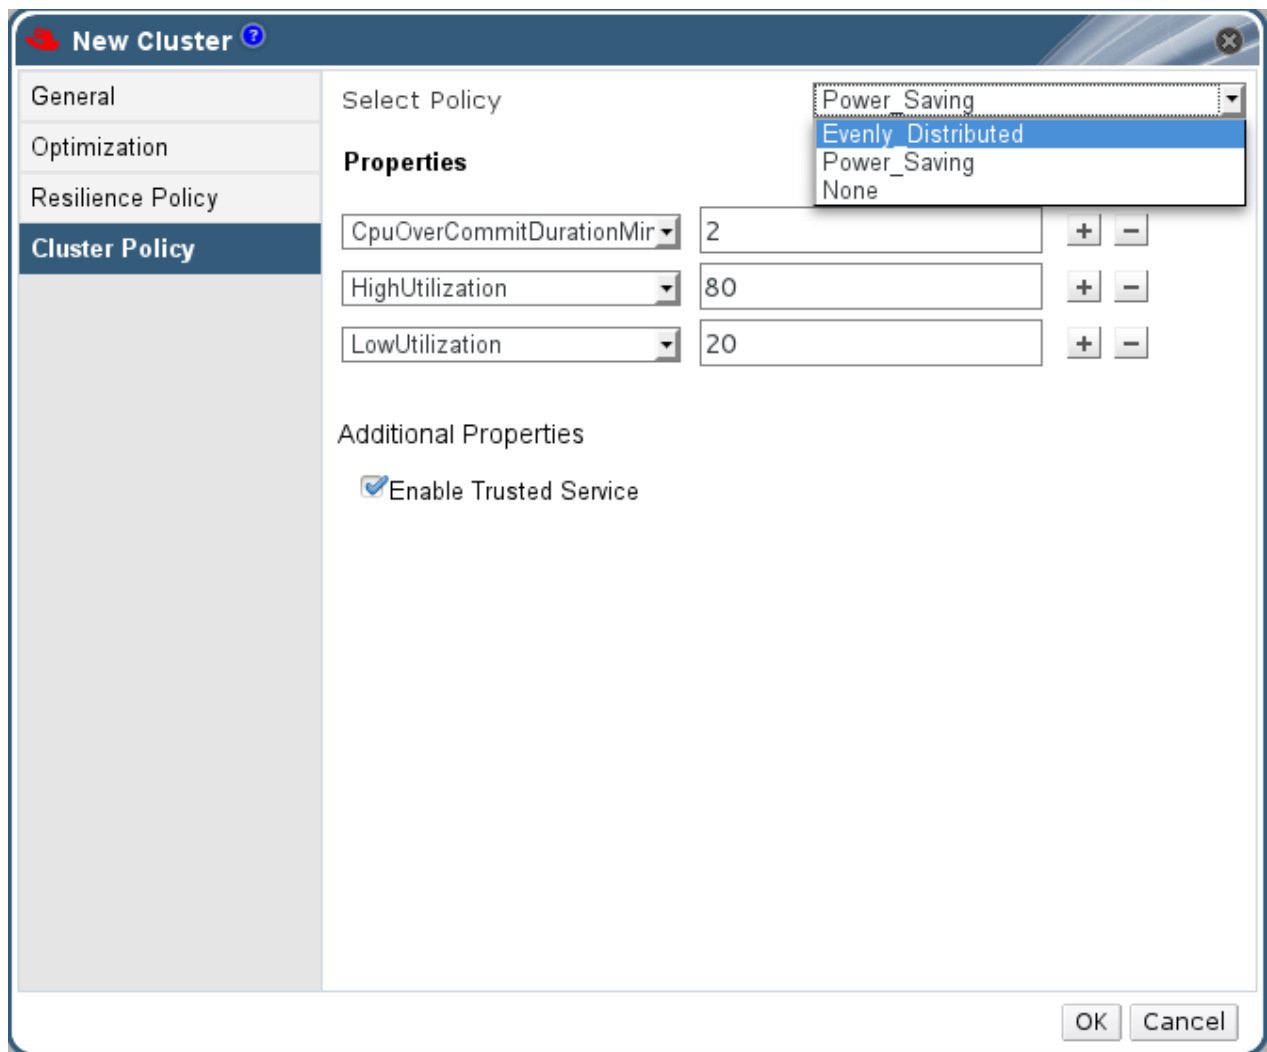
Two types of policies are available:

- **Even Distribution:** You can set a maximum service level threshold for the hosts in the cluster. If a host is above the threshold, then the Red Hat Enterprise Virtualization Manager will live migrate virtual machines off this host to other hosts in the cluster, as long as the other hosts' loads are below this threshold. This affects the scheduling of running new machines in a similar manner.
- **Power Saving:** This is a superset of the even distribution policy. A low threshold is set for consolidating virtual machines. If a server in the cluster drops below the threshold, then the virtual machines on this server are live migrated to other servers in the cluster, as long as the other servers are not above the maximum service level threshold.

In the following example you will set and test both thresholds.

### Procedure 3.1. 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** button to open the **Edit Cluster** window.



**Figure 3.3. Edit Cluster Policy**

3. Select one of the following policies:

- **None**

- ✱ **Evenly\_Distributed** - Enter CPU utilization percentage at which virtual machines start migrating to other hosts in the **HighUtilization** text field.
  - ✱ **Power Saving** - Enter the CPU utilization percentage below which the host will be considered under-utilized in the **LowUtilization** text field. Enter the CPU utilization percentage at which virtual machines start migrating to other hosts in the **HighUtilization** text field
4. Specify the time interval in minutes at which the selected policy will be triggered in the **CpuOverCommitDurationMinutes** text field.
  5. If you are using an OpenAttestation server to verify your hosts, and have set up the server's details using the **engine-config** tool, select the **Enable Trusted Service** check box.
  6. Click **OK**.

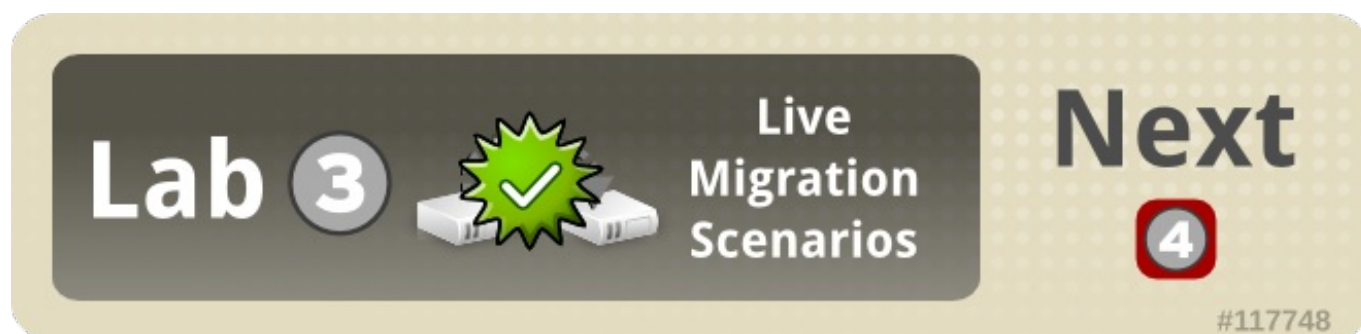
Now you can test these new settings. After the maintenance lab you now have a host (Atlantic) that does not have any virtual machines running. Select one of the Red Hat Enterprise Linux virtual machines and load it, then watch as the hosts in the cluster balance the virtual machines' workloads.

[Report a bug](#)

### 3.6. Lab 3 - Summary

You have reached your third Track A goal to configure cluster policies for virtual machine live migration. In this lab, you have live migrated virtual machines when a host is brought down for maintenance, and configured a cluster policy to activate virtual machine live migration when the host's workload exceeds its defined threshold.

The next lab on Track A teaches you how to create and manage virtual machines in the Power User Portal.



[Report a bug](#)



## Chapter 4. Lab 4 - Power User Portal

### 4.1. Lab 4 - Objectives

The power user portal is a trimmed-down version of the administration portal, tailored for end user self provisioning of virtual machines. It is simultaneously a gateway for logging in to virtual machines, and also a platform to create virtual machines and manage resources specific to those virtual machines. This lab illustrates the functions of the power user portal, which is available to users with `PowerUserRole` permissions. At the end of this lab you will be able to create and manage virtual machines from the power user portal.

This lab assumes that you have correctly installed and configured Red Hat Enterprise Virtualization so you can log in to the power user portal. You should have at least two running virtual machines in your environment.

To log in to the power user portal, you need a Red Hat Enterprise Linux client running Mozilla Firefox 3.5 and higher (you can use your Manager server).

#### Lab 4 - Objectives

This lab takes you through the tasks necessary to create virtual machines from the power user portal and assign user permissions. This lab should take you about 35 minutes.

#### See Also:

- [Section 4.2, “Adding IdM domains”](#)
- [Section 4.3, “Adding new users in the IdM directory”](#)
- [Section 4.4, “PowerUserRole permission assignment”](#)
- [Section 4.5, “Assigning PowerUserRole permissions on existing virtual machines”](#)
- [Section 4.6, “Logging in to the Power User Portal”](#)
- [Section 4.7, “Logging in to the User Portal”](#)
- [Section 4.8, “Creating Linux desktop virtual machines”](#)
- [Section 4.9, “Opening virtual machine consoles”](#)
- [Section 4.10, “Making templates from virtual machines”](#)
- [Section 4.11, “Verifying virtual machine template permissions”](#)

[Report a bug](#)

### 4.2. Adding IdM domains

Previously, you have logged in to the administration portal as the **admin** user on the **internal** domain, which was automatically set up during the installation of the Red Hat Enterprise Virtualization Manager. However, to authenticate new users, you need an external directory service. The term directory service refers to the collection of software, hardware, and processes that store information about an enterprise, subscribers, or both, and make that information available to users. Red Hat Enterprise Virtualization supports Identity Management (IdM), Red Hat Directory Service (RHDS), and Active Directory.

This lab assumes that you already have an existing IdM directory service. However if you need further assistance to install and configure IdM, see the *Red Hat Enterprise Linux Enterprise Identity Management Guide*. In this lab, you will attach an IdM domain to the Red Hat Enterprise Virtualization Manager using the **engine-manage-domains** tool, and create users in the IdM directory. Alternatively, if you have an Active Directory setup, you can attach it to the manager and use it for this lab.

Perform the following procedure on the Red Hat Enterprise Virtualization Manager server.

1. Log in to the Red Hat Enterprise Virtualization Manager server console.
2. Run the following command, and provide the domain administrator password when prompted:

```
# engine-manage-domains -action=add -
domain=ipadomain.demo.redhat.com -user=admin -interactive
```

3. Restart the service for the changes to be applied across the system.

```
# service ovirt-engine restart
```

Restarting the **ovirt-engine** service disconnects you from the administration portal. After a few minutes, the restart completes, and the IdM domain is added to the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

## 4.3. Adding new users in the IdM directory

Before you can add users in the Red Hat Enterprise Virtualization Manager, you must first add them to the IdM directory. For this track, you need at least two users. The names used in this guide are **rhevpower** and **rhevuser**. Perform the following procedure on the Manager server.

1. To add users you must first authenticate as the directory server administrator. Use the **kinit admin** command to do this, entering the administrator password when prompted.

```
$ kinit admin
Password for admin@DIRECTORY.DEMO.REDHAT.COM:
```

2. To add a user interactively use the **ipa user-add** command. The command will prompt you for all values required to create the user.

```
$ ipa user-add
First name: RHEV
Last name: User
User login [rhevuser]:
-----
Added user "rhevuser"
-----
User login: rhevuser
First name: RHEV
Last name: User
Full name: RHEV User
Display name: RHEV User
Initials: RU
```

```
Home directory: /home/rhevuser
GECOS field: rhevuser
Login shell: /bin/sh
Kerberos principal: rhevuser@DIRECTORY.DEMO.REDHAT.COM
UID: 1316000004
```

3. To allow the new user to log in you must set their initial password. Use the **ipa passwd** command, followed by the user name for which you are setting the password, to do this.

```
$ ipa passwd rhevuser
Password:
Enter Password again to verify:
-----
Changed password for "rhevuser@DIRECTORY.DEMO.REDHAT.COM"
-----
```

4. A new user has been added to the directory server and their password has been set. You are now able to add them to the Red Hat Enterprise Virtualization Manager. For this track, repeat Step 2 and Step 3 for another user called **rhevpower** and as necessary for other users.

Now you know how to create users for Red Hat Enterprise Virtualization. Next you will learn how to assign roles and privileges to these users.

[Report a bug](#)

## 4.4. PowerUserRole permission assignment

A Power User can perform some administrative functions in the User Portal, including creating and editing virtual machines, creating templates, and working with snapshots. In addition, a Power User who creates a virtual machine from the User Portal will be automatically assigned to the virtual machine. A Power User has permissions for the assigned virtual machine only, not for all virtual machines in the enterprise.

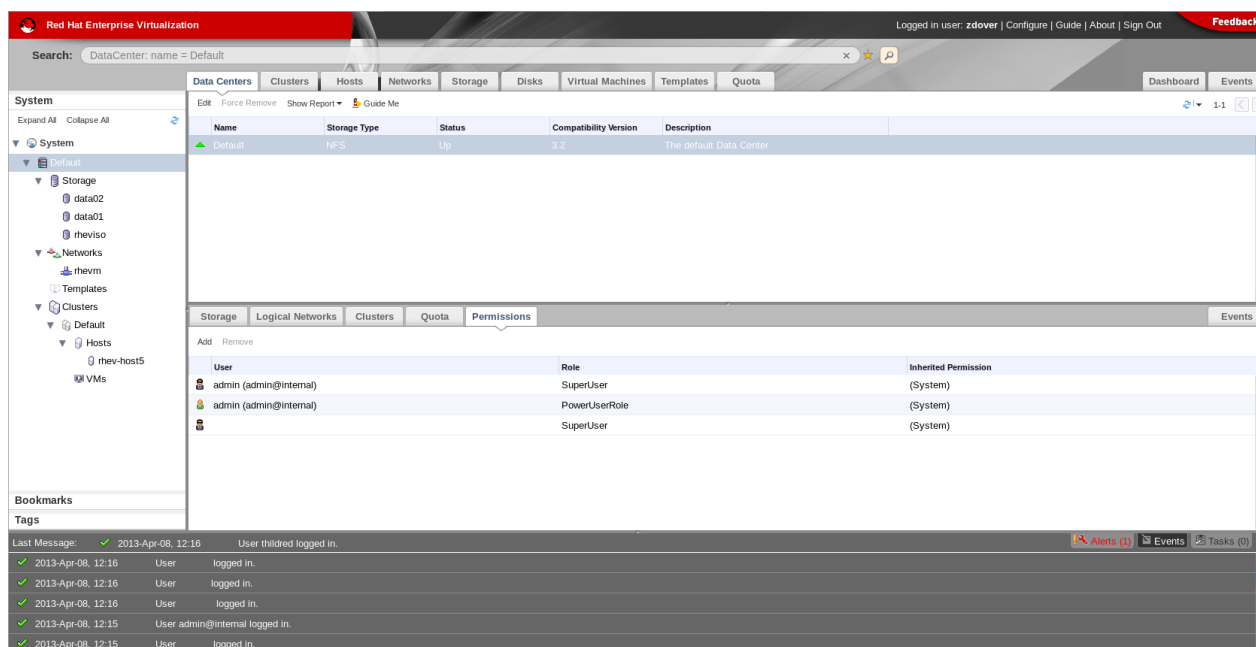
In Red Hat Enterprise Virtualization Manager, permissions are set on system objects, including data centers, clusters and virtual machines. This permission is user specific, while the scope of permission is defined by the user role.

For a user to log in to the power user portal, the user must have power user permissions on any system object. However, in order to allow a user to create both virtual machines and templates, as you are going to do in this lab, the permission has to be assigned for the data center level.

Perform this procedure in the administration portal as the **admin** user in the **internal** domain.

### To assign PowerUserRole permissions on a data center

1. On the **Tree** pane, click **Expand All** and select the **Default** data center. On the **Data Center** tab, select the **Default** data center to display the details pane, and click the **Permissions** subtab.



**Figure 4.1. Add permission to user**

2. Click **Add** to add an existing user. The **Add Permission to User** dialog displays. Under the **Search** fields, select **ipadomain.demo.redhat.com** on the drop-down menu and enter **rhevpower** in the text box. Click **Go**.
3. Select the check box of **rhevpower** to be assigned the permissions. Select the **Assign role to user** drop-down list and select **PowerUserRole**.
4. Click **OK**. The name of the user displays in the **Permissions** tab, with an icon and the assigned role.

While you have assigned permissions for the user **rhevpower** to make administrative changes to the data center, this user does not automatically inherit power user permissions for the existing virtual machines in the data center. To do so, repeat the previous procedure, substituting data center for the existing virtual machines.

[Report a bug](#)

## 4.5. Assigning PowerUserRole permissions on existing virtual machines

While you have assigned permissions for the user **rhevpower** to make administrative changes to the data center, this user does not automatically inherit power user permissions for the existing virtual machines in the data center. To do so, repeat the previous procedure, substituting data center for the existing virtual machines.

1. On the **Tree** pane, select the **VMs** icon under the **Default** data center. A list of virtual machine displays in the **Virtual Machines** tab. Select the first virtual machine you created, **RHEL6Thames**.
2. Click the **Permissions** subtab on the details pane and click **Add**. The **Add Permission to User** dialog displays. Under the **Search** fields, select **ipadomain.demo.redhat.com** on the drop-down menu and enter **rhevpower** in the text box. Click **Go**.
3. Select the check box of **rhevpower** to be assigned the permissions and select **PowerUserRole** from the list of permissions. Click **OK**.

You have now enabled **rhevpower** to create virtual machines in the Default data center from the power user portal.

[Report a bug](#)

## 4.6. Logging in to the Power User Portal

Now that you have created a Power User who has full permissions to an existing virtual machine, log in as **rhevpower** to the User Portal. You can use either the same client that you are using for the administration portal, or any Red Hat Enterprise Linux client running Mozilla Firefox.

If you are using a Red Hat Enterprise Linux client, install the SPICE plug-in before logging in to the User Portal. Run:

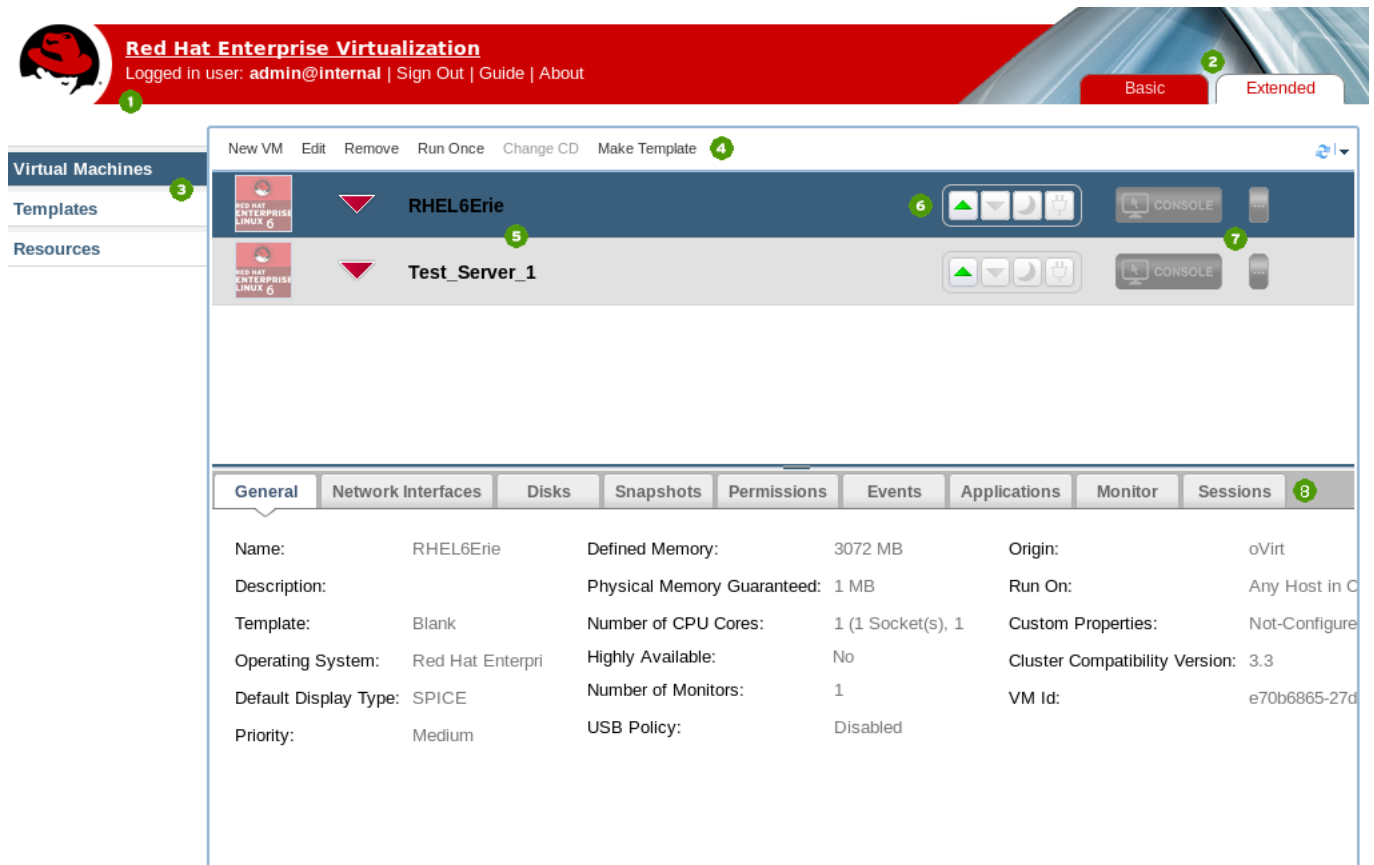
```
# yum install spice-xpi
```

You have now successfully logged into the User Portal. Here, you can access and monitor all the virtual machines that are available to you. The functions are briefly described in the following figure and list:

### To log in to the User Portal

1. Open your browser and navigate to **https://rhevm.demo.redhat.com/UserPortal**.
2. The login screen displays. Enter your **User Name** and **Password**, and select the **ipadomain.demo.redhat.com** domain in the drop-down menu. Click **Login**.

You have now successfully logged into the User Portal. Here, you can access and monitor all the virtual machines that are available to you. The functions are briefly described in the following figure and list:



**Figure 4.2. The Power User Portal**

1. **Title bar:** This bar displays the name of the User logged in to the portal and the Sign out button.
2. **Portal view:** The Extended view of the User Portal displays by default for power users. You can also switch to the Basic view, which is the default for users with basic permissions. You will be using the Basic portal in Advanced Lab 8 - Virtual Desktops.
3. **Navigation pane:** This pane toggles between displaying Virtual Machines, Templates and Resources. When a tab is selected, the available virtual resources display. This example uses the Virtual Machines tab.
4. **Management bar:** The buttons on this bar enable you to create and make changes to virtual machines.
5. **Virtual machine name:** This list displays the virtual machine's name, operating system logo and status of the virtual machine (running, paused or powered off).
6. **Power buttons:** These buttons allow you to play, stop, pause or power off a virtual machine.
7. **Connect button:** The Console button allows you to connect to virtual machines.
8. **Details pane:** Clicking on a virtual machine displays its statistics in this pane. You can view a virtual machine's details in the General, Events, Applications and Monitor subtabs, and make configuration changes in the Network Interfaces, Virtual Disks, Snapshots and Permissions subtabs.

Once you are familiar with the layout of the power user portal, you can start creating virtual machines, as instructed in the next section.

[Report a bug](#)

## 4.7. Logging in to the User Portal

1. Open your browser and navigate to **`https://rhev.m.demo.redhat.com:8443/UserPortal`**.
2. The login screen displays. Enter your **User Name** and **Password**, and select the **ipadomain.demo.redhat.com** domain in the drop-down menu. Click **Login**.

[Report a bug](#)

## 4.8. Creating Linux desktop virtual machines

1. Ensure that you are on the **Extended** User Portal. Click **New VM**. The **New Virtual Machine** window opens.

**New Virtual Machine**

**General**

Cluster: Default/Default

Based on Template: Blank

Operating System: Other OS

Optimized for: Server

Name:

Description:

Comment:

☐ Stateless ☐ Start in Pause Mode ☐ Delete Protection

VM has no network interfaces. To add one, assign a profile.

nic1:

Hide Advanced Options OK Cancel


**Figure 4.3. New Server**

- a. On the **General** tab, select the *EngRHEL6Base* template from the **Based On Template** field.

- b. Enter a suitable **Name** and **Description**, and accept the default values inherited from the template in the rest of the fields. You can change them if needed.
  - c. Click the **Resource Allocation** tab. On the **Provisioning** field, click the drop down menu and select the *Clone* option. Set the clone method under **Disks** to *Preallocated*.
2. Retain all other default settings and click **OK** to create the virtual machine.

[Report a bug](#)

## 4.9. Opening virtual machine consoles

1. Select **RHEL6Erie** and click the Play  button. The virtual machine powers up.
2. When the virtual machine is turned on, the green Play symbol appears next to its name. The virtual machine is now ready for connection.



**Figure 4.4. Virtual machine turned on**

3. Click the **Console** button to connect to the virtual machine.
  4. A SPICE console window of the virtual machine displays. You can now use the virtual machine in the same way you would use a physical desktop.

[Report a bug](#)

## 4.10. Making templates from virtual machines

1. On the **Virtual Machines** tab, select the virtual machine you wish to use as a basis to create a template. Ensure that it is powered down, then click **Make Template**.
2. The **New Template** dialog displays.



**New Template** ?

Name

Description

Comment

Cluster

Disks Allocation:

Alias	Virtual Size	Target
<input type="text" value="Disk1"/>	50 GB	<input type="text" value="storage_data_z"/>

☒ Allow all users to access this Template

☐ Copy VM permissions

OK Cancel

**Figure 4.5. Create New Template**

Enter information in the following fields:

- ✧ **Name:** Name of the new template.
  - ✧ **Description:** Description of the new template.
  - ✧ **Host Cluster:** The host cluster for the virtual machines using this template.
  - ✧ **Storage Domain:** The storage domain for the virtual machines using the template.
  - ✧ **Allow all users to access this template:** Ticking this checkbox allows the template to be accessed by all users. For this lab, leave this checkbox unmarked.
3. Click **OK**. The virtual machine will be locked while the template is being created. Once created, select the **Templates** tab in the navigation pane. The newly created template displays in the results list.

## 4.11. Verifying virtual machine template permissions

Now that you have created a virtual machine and a template in the power user portal, make sure that the changes you made have been applied in the Red Hat Enterprise Virtualization Manager administration portal.

Perform this procedure in the administration portal as the **admin** user.

### Virtual Machine

1. Select the **Virtual Machines** tab. Verify that the virtual machine you created from the User Portal, **RHEL6Erie**, appears in the results list.
2. Click on the **RHEL6Erie** virtual machine and look at the details pane. Click on the **Permissions** tab and check if the user named **rhevpower** has inherited **UserVmManager** permissions for this virtual machine.

### Template

1. Select the **Templates** tab. Verify that the template you created from the User Portal, appears in the results list.
2. Click on the **PrivateRHEL6** template and look at the details pane. Click on the **Permissions** tab. As the creator of the template, **rhevpower** should have inherited **PowerUserRole** permissions. Meanwhile **admin**, as the administrative user for all components in the system, should have inherited **SuperUser** permissions for the template. Nobody else can use this template unless they are given permissions by the existing permitted users.
3. In contrast, click on the **EngRHEL6Base** template, which was marked public. On the **Permissions** tab, **Everyone** has been assigned **UserTemplateBasedVm** permissions, meaning that all users in the system can access this template.

[Report a bug](#)

## 4.12. Lab 4 - Summary

You have now completed the User Portal Lab. In this lab, you successfully created virtual machines from the power user portal and assigned permissions for users to access virtual machines and templates.

This is the last basic lab for Tracks A and B. You can proceed with the Advanced Labs, which demonstrate further features of Red Hat Enterprise Virtualization, including a flexible system to manage administrative users, power saving policies, and methods of deploying Red Hat Enterprise Virtualization in a large organization.

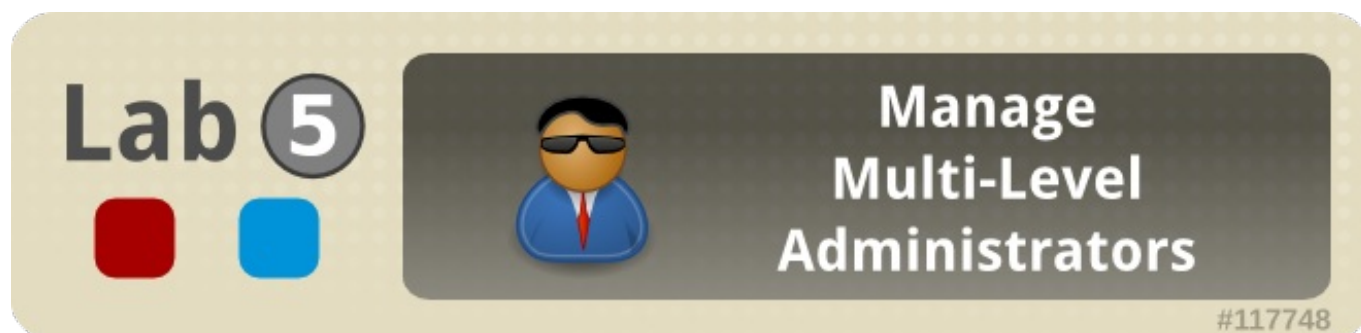


---

[Report a bug](#)

## Chapter 5. Lab 5 - Managing Multi-Level Administrators

### 5.1. Lab 5 - Objectives



This lab introduces you to Red Hat Enterprise Virtualization's multi-level administration system. Multi-level administration presents a hierarchy of permissions that can be configured to provide finely grained levels of permissions as required by your enterprise. You have already been partially introduced to this system when you granted permissions to users on virtual machines and data centers in the Power User Portal chapter.

Permissions enable users to perform actions on objects, where objects are either individual objects or container objects. Any permissions that apply to a container object also apply to all members of that container. For example, when a host administrator role is applied to a user on a specific host, the user will have permissions to perform any of the available host operations, but on the assigned host only. However, if a host administrator role is applied on a data center to a user, the user will gain permissions to perform host operations on all hosts within the cluster of the data center. If there are additional host clusters in the data center, the user will not be able to make changes to the hosts.

This lab assumes that you have successfully completed the basic labs. You should have correctly installed and configured Red Hat Enterprise Virtualization, and have several user accounts in the IdM, AD, or RHDS domain.

[Report a bug](#)

### 5.2. Lab 5 - requirements

In addition to the requirements stipulated in (for Track A) or (for Track B) ensure that you have at least two users in an external directory service.

[Report a bug](#)

### 5.3. Storage administrator definition

A Storage Administrator can manage, create and remove storage domains. This is useful in enterprises comprising multiple storage domains, each of which requires its own system administrator. Storage Administrators have permissions for the assigned storage domain only, not for all storage domains in the enterprise.

To assign user permissions, log in to the Red Hat Enterprise Virtualization Manager administration portal as the SuperUser. In this example, SuperUser permissions have been assigned to **admin**.

You have now assigned administrative privileges for the **local-iso-domain** storage domain to the user named **rhevuser**. Next, you will assign PowerUserRole permissions for the same user.

[Report a bug](#)

## 5.4. Assigning system administrator roles to storage domains

1. Navigate to the **Tree** pane and click the **Expand All** button. Under **Default**, click **Storage**. The available storage domains displays in the **Storage** tab.
2. Select the storage domain to which to assign users and click the **Permissions** sub tab on the details pane. This example uses the **local-iso-share** domain.
3. Click **Add** to add an existing user. The **Add Permission to User** dialog displays. Enter **rhevuser** in the **Search** text box, and click **Go**.
4. Tick the check box of **rhevuser**. Select the **Assign role to user** drop-down list and select **StorageAdmin**.

**Add Permission to User**

☒ Specific User/Group ☐ Everyone

Search: ecs-cloud.lab.eng.bne.redhat.com

First Name	Last Name	User Name
------------	-----------	-----------

Role to Assign: StorageAdmin

**Figure 5.1. Add StorageAdmin permission**

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

[Report a bug](#)

## 5.5. Virtual machine administrator definition

A power user can create, manage and delete virtual machines, configure storage and networking, as well as migrate the virtual machine. These administrative privileges apply only to the virtual machines on which the user has `PowerUserRole` permissions, not all virtual machines in the system.

You should still be logged in to the Red Hat Enterprise Virtualization Manager administration portal as **admin**. Repeat the same procedure as before to assign permissions.

You have now assigned administrative privileges for the **RHEL6Thames** virtual machine and also for the **local-iso-domain** storage to the user named **rhevuser**. The next step is to test if the permissions have been assigned as expected.

### Assigning system administrator roles to virtual domains

1. On the **Tree** pane, click **Expand All**. Under the **Default** data center and **Default** host cluster, click the **Virtual Machines** icon. The available virtual machines display in the **Virtual Machines** tab.
2. Select the virtual machine that you want to assign users to, and click the **Permissions** subtab on the details pane. This example uses the **RHEL6Thames** virtual machine.
3. Click **Add** to add an existing user. The **Add Permission to User** dialog displays. Enter **rhevuser** in the **Search** text box, and click **Go**.
4. Tick the check box of **rhevuser**. Select the **Assign role to user** drop-down list and select **PowerUserRole**.
5. Click **OK**. The name of the user displays in the **Permissions** tab, with an icon and the assigned role.

[Report a bug](#)

## 5.6. User permission verification

To verify that the Storage Administrator role has been correctly assigned, sign out and log in to the Red Hat Enterprise Virtualization Manager administration portal as **rhevuser**, then perform a few storage configuration options.

Next, verify that the `PowerUserRole` has been correctly assigned. You should still be logged in as **rhevuser**.

### To verify StorageAdmin permissions

1. Navigate to the **Tree** pane and click the **Expand All** button. Under **Default**, click **Storage**. The available storage domains displays in the **Storage** tab.
2. Select **local-iso-domain**. On the details pane, select the **Data Center** tab and click the **Maintenance** button. The ISO storage domain is deactivated, and appears as *Inactive* in the Storage pane.
3. Select the Data storage domain. On the details pane, select the **Data Center** tab and click the **Maintenance** button. This time you will get an error message stating "User is not authorized to perform this action". This happened because you gave administrative permissions only for the ISO domain, not the Data domain.
4. Select the ISO domain again and click the **Data Center** tab on the details pane. Click **Activate**. The domain is activated, and displays as Active in the Storage pane.

Next, verify that the **PowerUserRole** has been correctly assigned. You should still be logged in as **rhevuser**.

### To verify **PowerUserRole** permissions

1. On the **Tree** pane, click **Expand All**. Under the **Default** data center and **Default** host cluster, click the **Virtual Machines** icon. The available virtual machines display in the **Virtual Machines** tab.
2. Select the **RHEL6Thames** virtual machine and try to stop it if it is running, or play it if it is not running. You should succeed.
3. Select a different virtual machine, and try to do the same. This time you should get the "User is not authorized to perform this action" message. This is because you only assigned power user permissions for **RHEL6Thames** and not the second virtual machine.

As you have seen in this lab, you can configure different levels of permissions for different users on multiple objects. This multi-level administration system is ideal for organizations with a diverse range of users who have different needs, and allows for enhanced security in that only specifically assigned users will be able to make system-wide changes.

[Report a bug](#)

## 5.7. Custom role creation

In addition to pre-defined roles, you can also create your own custom role with permission levels for various objects to suit your organization's needs. You can create administrator and user roles to permit access and management to different levels of resources in the data center. To perform the following procedure, log in to the administration portal as the **admin** user.

### To create a new role

1. On the header bar, click **Configure**. The **Configure** dialog displays. Under the **Roles** tab, the dialog includes a list of default User and Administrator roles, and any custom roles.
2. Click **New**. The **New Role** dialog displays.

**Figure 5.2. Create a new user role**

- ✧ Enter a **Name** and **Description** for your new role.
- ✧ Select either **User** or **Admin** as the **Account Type**.
- ✧ Click the **Expand All** button to view more 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.
- ✧ For each of the objects, select the actions you wish to permit for the role you are setting up.

3. Click **OK** to apply the changes you have made. The new role displays on the list of roles.

You have now created a new role, and can assign it to users as required. In addition to creating new roles from scratch, you can also clone default roles and modify the cloned roles as you wish. As a result, you have a wide range of possibilities to create a versatile, fine-grained model of permissions according to your organization's requirements.

[Report a bug](#)

## 5.8. Lab 5 - Summary

In this lab, you have successfully assigned multiple levels of administrative permissions to users.

- ✧ If you are on Track A, go to the High Availability lab. This lab teaches you how to configure power saving and resource management policies.



- » If you are on Track B, go to the Virtual Desktops lab. This lab shows you how to create and allocate desktop pools, and how to connect to virtual machines from the basic user portal.



[Report a bug](#)

## Chapter 6. Lab 6 - High Availability scenarios

### 6.1. Lab 6 - Objectives

#### Lab 6 - Objectives

This lab takes you through the best practices of configuring and testing a reliable Red Hat Enterprise Virtualization environment where virtual machines running critical workloads are not easily interrupted. This lab should take you about 30 minutes.

[Report a bug](#)

### 6.2. Lab 6 - Requirements

Before you attempt this lab, you must have a working Red Hat Enterprise Virtualization environment. To successfully complete this lab, you must have:

- An out of band power management device for each host. This lab uses the Intelligent Power Management Interface (IPMI) device.
- Access to the administration portal and access to your hosts (for simulating failure scenarios). You will be instructed on how to enable remote login for your hypervisors later in this lab.

[Report a bug](#)

### 6.3. Lab 6 - Prerequisites



The Red Hat Enterprise Virtualization Manager offers various high availability features which can be applied in a granular manner, from the level of a single virtual machine up to protection against multiple host failure scenarios. In addition, you can protect your virtual machines against various failures by combining virtual machine high availability with out of band power management, Red Hat Enterprise Virtualization Manager's failure detection and failure recovery solutions.

This lab enables you to configure virtual machine high availability, and demonstrates its use in several common enterprise scenarios. It is assumed that you have successfully completed the basic labs, meaning that you have correctly installed and configured Red Hat Enterprise Virtualization, and have several running virtual machines.

In addition, to successfully complete this lab you must have a power management card for each of your hosts. This lab uses the Intelligent Power Management Interface (IPMI) device as an example. If you have a different device see the *Red Hat Enterprise Virtualization Administration Guide*.

[Report a bug](#)

## 6.4. Power Management Configuration

### 6.4.1. Power management configuration

#### Configure Power Management

At this point, you should have two hosts and at least two virtual machines, however this lab uses six virtual machines. You can use as many virtual machines as you want, but for optimal demonstration of high availability features, it is recommended that you add four new virtual machines to your environment.

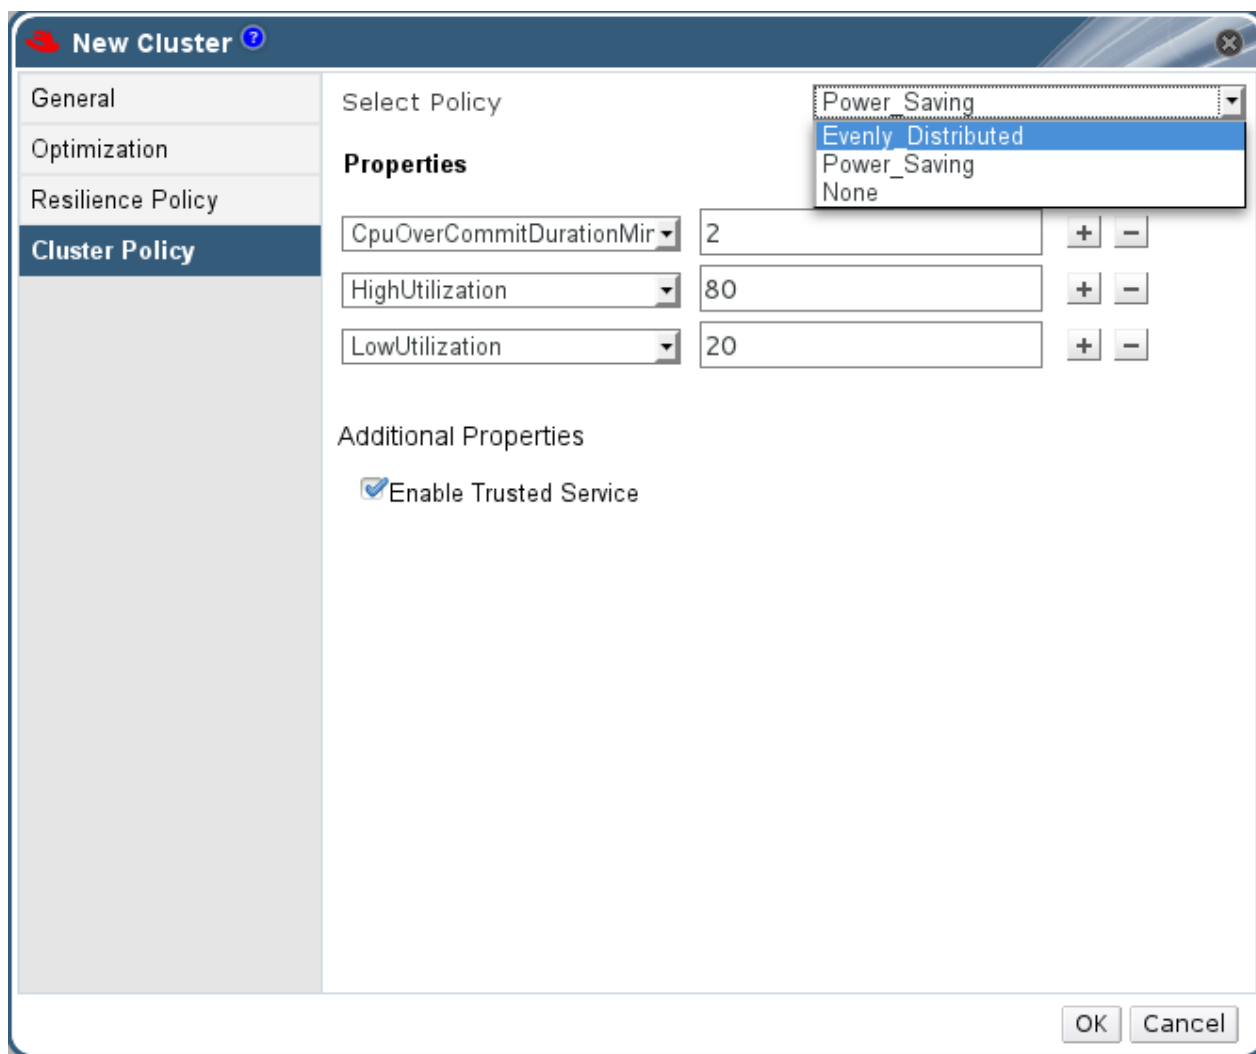
Highly available virtual machine settings will only be effective if power management is enabled on the hosts. However before configuring power management, recall that you have previously defined a power saving cluster policy. Cluster policies and high availability can be used concurrently; however to best demonstrate Red Hat Enterprise Virtualization's high availability features for this lab, reset the cluster policy. This ensures that the virtual machines remain where they are staged before each demonstration, so the migration is triggered by high availability rather than the cluster's load balancing policies.

[Report a bug](#)

### 6.4.2. Disabling cluster policy

#### To disable cluster policy

1. On the administration portal, navigate to the **Tree** pane. Click the **Expand All** button. Click the **Default** cluster. The **Clusters** tab displays. Select the **Default** cluster to display its details pane.
2. On the **Policy** subtab, the policy is set to *Power Saving*. Click **Edit Policy**.



**Figure 6.1. Edit cluster policy**

3. The **Edit Policy** dialog displays. Select the **None** button to remove the previously configured policy. Click **OK**.

Now, you can configure power management on your hosts. Power management enables the system to fence a troublesome host using an additional interface such as an Intelligent Platform Management Interface (IPMI) device. Perform this procedure for each host.

[Report a bug](#)

### 6.4.3. Setting up power management on a host

#### Summary

In this section, you will set up power management on a host.

#### To set up power management on a host

1. On the **Tree** pane, click the **Hosts** icon under the **Default** cluster. The **Hosts** tab displays a list of available hosts.

2. Select a host, in this example the **Atlantic** host is used. Notice that there is an exclamation mark next to the hostname which you were asked to ignore in previous labs. Click the **Edit** button to display the **Edit Host** dialog.
3. Select the **Power Management** tab. Tick the **Enable Power Management** check box and provide the required information in the following fields:
  - ✧ **Address**: The address of the power management card. This card does not have to be on any of the logical networks defined on the Red Hat Enterprise Virtualization Manager. However, all hosts in the cluster must have a valid route to the card's IP. This ensures that fencing is correctly configured, as fencing is done by any available host in the cluster.
  - ✧ **User Name**: The user allowed to log in to the power management device.
  - ✧ **Password**: The password of the user allowed to log in.
  - ✧ **Type**: The type of management device. For this lab, select **ipmilan**.
  - ✧ **Options**: These additional parameters depend on the specific implementation of each device. Detailed documentation of the options available is provided in the man page for each fence agent. For this lab, enter **power\_wait=4,lanplus=yes** in the provided text box.

Click the **Test** button to test the settings. (Note: in the figure below, the **Test** button is not shown. You can find the **Test** button at the bottom of the **Power Management** tab.) If the power management options can be verified, the text *Test Succeeded, Host Status is: on* displays.

**Edit Host**

General

**Power Management**

SPM

Console

☒ Enable Power Management

Primary

Concurrent ☐

Address

User Name

Password

Type **apc**

SSH Port

Slot

Options

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

Secure ☐

Source cluster  
dc

Up

Down

Test

OK Cancel

**Figure 6.2. Enabling Power Management on a Host**

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

## Result

You have now configured power management for your hosts, meaning that the status of power on your host can be verified and controlled by the Red Hat Enterprise Virtualization Manager. Power management checks that a host is properly powered down, then restarts its virtual machines on another host in the same cluster. However, if the host's status cannot be verified, the virtual machines that were originally running on it will not be restarted.

[Report a bug](#)

## 6.5. Virtual Machine High-Availability Configuration

### 6.5.1. Virtual Machine High-availability Configuration

Now that you have configured power management on your hosts, they are ready to run highly available virtual machines. High availability means that a virtual machine will automatically be restarted if its process is interrupted. This interruption occurs when a virtual machine is terminated by methods other than shutting down within the guest or from the administration portal. When these events occur, the highly available virtual machine will be automatically restarted, either on its original host or another host in the cluster.



#### Note

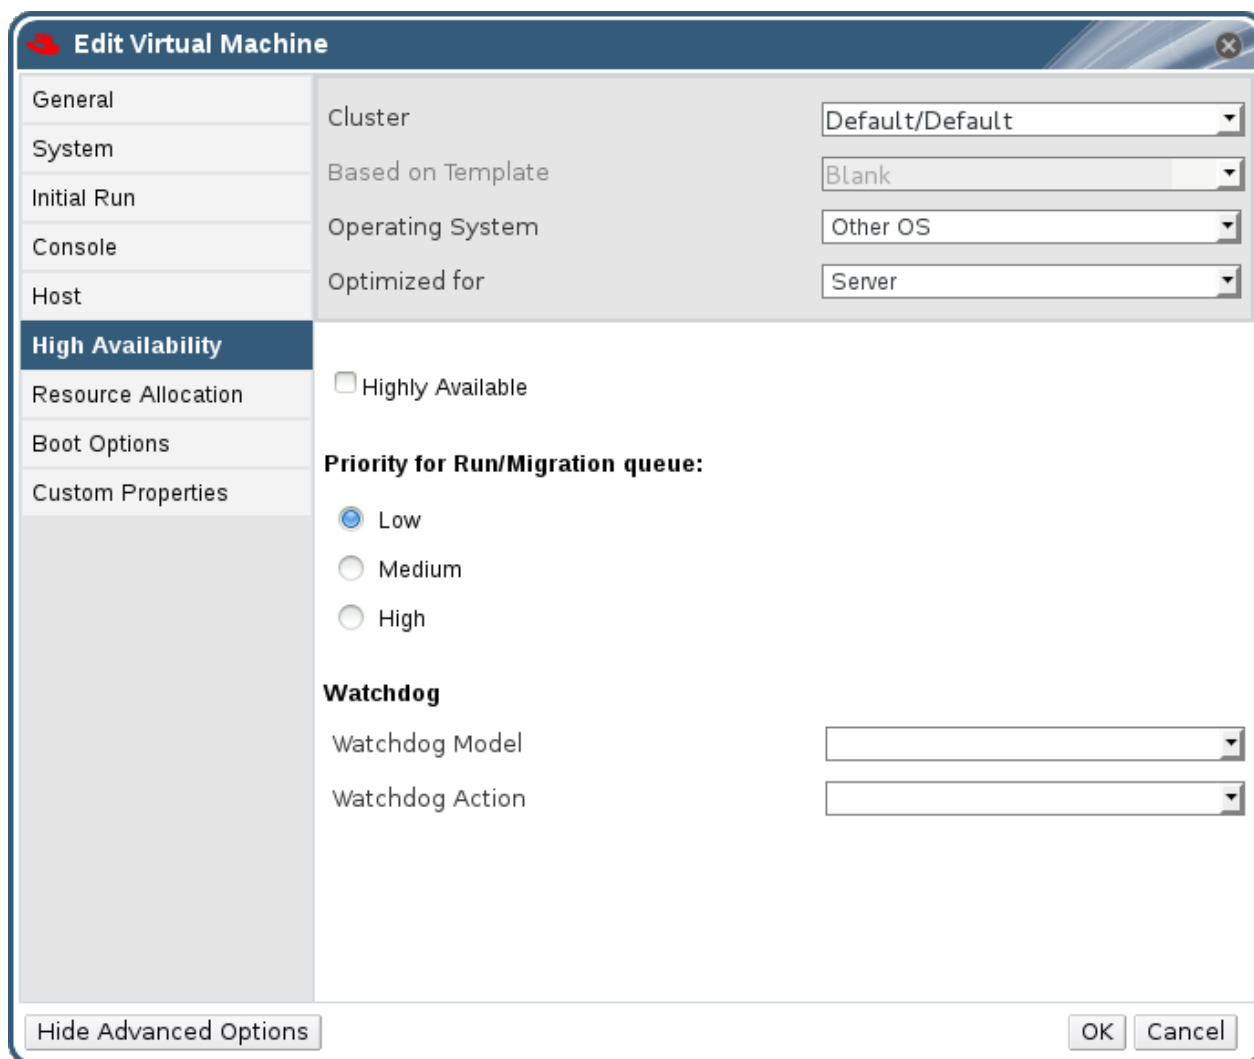
High availability can only be configured for virtual servers, not virtual desktops.

This lab uses six virtual machines, and configures three of them to be highly available. You can use different names, or a different number of virtual machines - just be aware that this lab uses examples which consistently refer to these allocated names. The virtual machines used in this lab are:

- Highly Available Virtual Machines: RHEL6Nile, RHEL6RioGrande, RHEL6Thames
- Non-Highly Available Virtual Machines: RHEL6Congo, RHEL6Erie, RHEL6Ganges

#### To configure High Availability for a virtual machine

1. In the **Navigation** pane, click the **VMs** tab. Select the virtual machine you wish to mark as highly available. Click **Edit**.
2. The **Edit Virtual Machine** dialog displays. Select the **High Availability** tab and tick the **Highly Available** check box. Change the **Priority for Run/ Migrate Queue** to **High**. This means that this virtual machine will take precedence in the queue of virtual machines to be migrated.

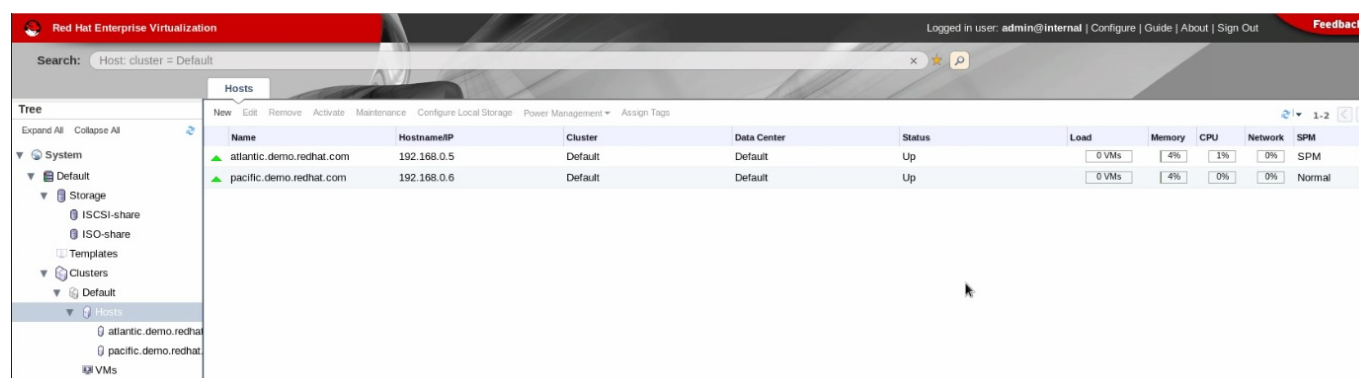


**Figure 6.3. Set a server to be highly available**

3. Click **OK** to save your settings. Repeat this procedure for other virtual machines you wish to mark as highly available, in this example they are **RHEL6Nile** and **RHEL6RioGrande**.

Now, you have three virtual machines which are highly available, and three which are not. Hold down the **Shift** key and click the virtual machines to select all of them, then click the **Run** icon to start them all. The virtual machines may not be evenly distributed across the hosts, because one of them is the allocated Storage Pool Manager (SPM) which manages access between hosts and storage. The SPM's CPU resources are utilized more heavily than that of other hosts. Therefore, when virtual machines are started, the host with a lighter workload is selected to run the machines.

On the **Tree** pane, click **Hosts** to display the available hosts and their workloads on the **Hosts** tab.





**Figure 6.4. Virtual machines running on different hosts**

You have now successfully configured power management for your hosts, and set several of your virtual machines to be highly available. You can now experiment with virtual machine high availability in four scenarios that may happen in an enterprise data center host manual reboot (user error), virtual machine crash (unexpected fault), partial failure (non-operational) and host disconnection full failure (non-responsive).

[Report a bug](#)

## 6.6. High Availability: Host-Initiated Reboot

### 6.6.1. High availability -- host initiated reboot

At this stage you now have six running virtual machines, three of which are highly available. This section simulates an event in which a user error has caused a temporary host failure, after which the host recovers, but the virtual machines which were running on the host are terminated. In this case, the host is manually fenced before being placed into maintenance mode.

**To demonstrate virtual machine high availability when host is incorrectly fenced**

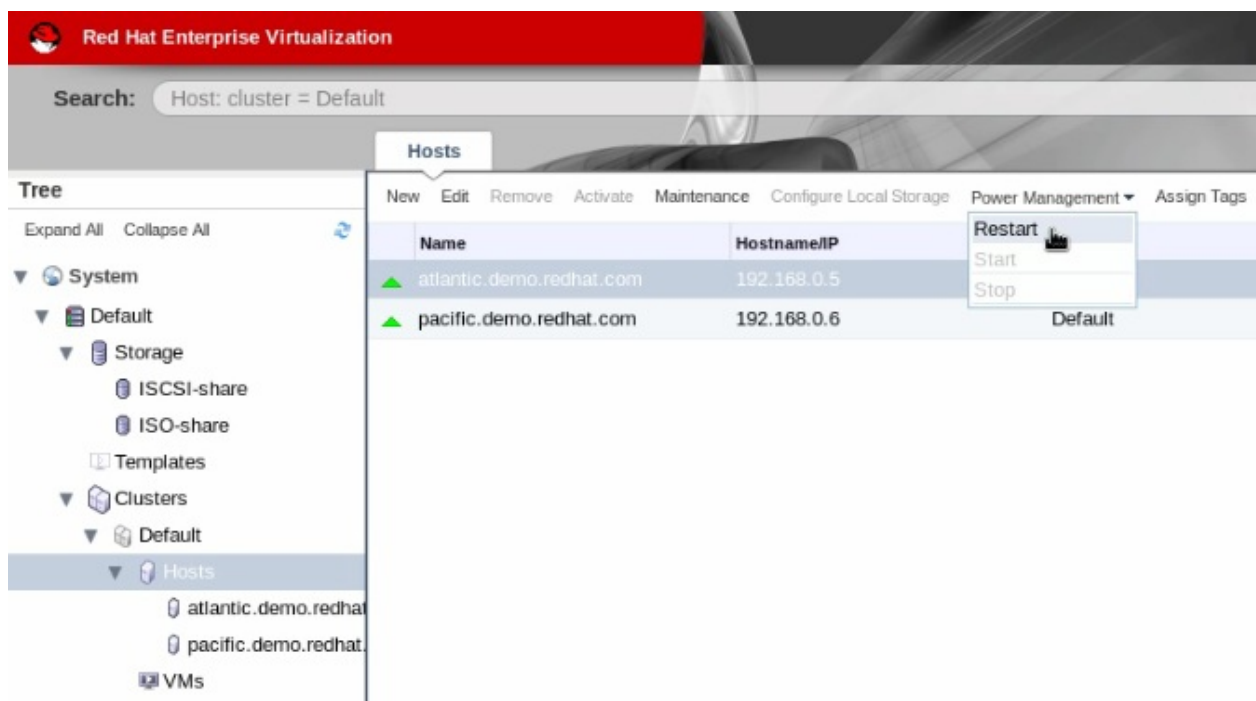
1. On the **Tree** pane, select **VMs**. The available virtual machines display on the **Virtual Machines** tab.

Name	Cluster	Data Center	Host	IP Address	Memory	CPU	Network	Display	Status	Uptime	Logged-in User
RHEL6-1	Default	Default	atlantic.demo.redhat.c	192.168.0.101	13%	0%	0%	Spice	Up	4 min	
RHEL6-3	Default	Default	pacific.demo.redhat.c	192.168.0.148	13%	0%	0%	Spice	Up	5 min	
RHEL6-newVM	Default	Default	pacific.demo.redhat.c	192.168.0.141	13%	3%	0%	Spice	Up	10 min	
RHEL6-poweruser	Default	Default	atlantic.demo.redhat.c	192.168.0.102	13%	3%	0%	Spice	Up	5 min	
RHEL6-VM	Default	Default	pacific.demo.redhat.c	192.168.0.140	13%	0%	0%	Spice	Up	10 min	
RHEV6-2	Default	Default	pacific.demo.redhat.c	192.168.0.104	13%	0%	0%	Spice	Up	3 min	

**Figure 6.5. Virtual machines running on hosts**

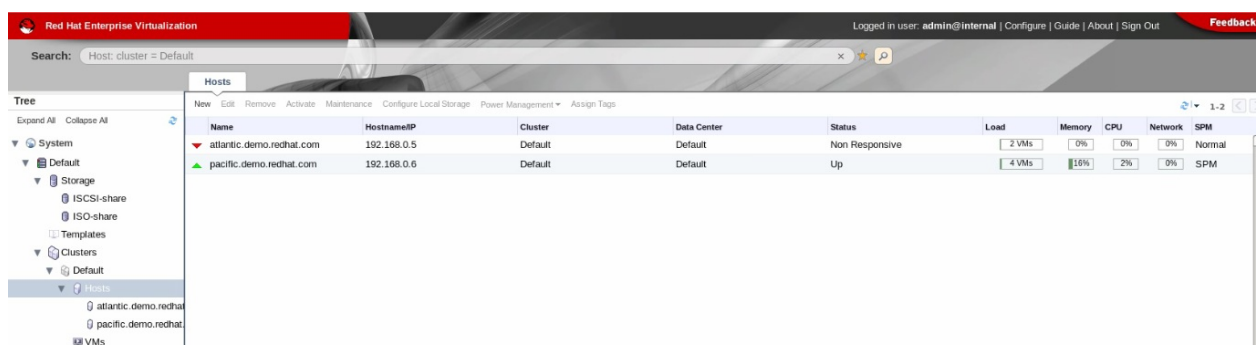
In this example, there are currently two virtual machines running on the **Atlantic** host; and another four on **Pacific**. You can check whether a virtual machine is highly available by selecting it looking at the **General** tab of the details pane. Here, on the **Atlantic** host **RHEL6Thames** is set to be highly available, while **RHEL6Congo** is not.

2. On the **Tree** pane, select **Hosts**. The available hosts display on the **Hosts** tab. Select the **Atlantic** host, click the **Power Management** drop-down menu and select **Restart**.



**Figure 6.6. Restart host**

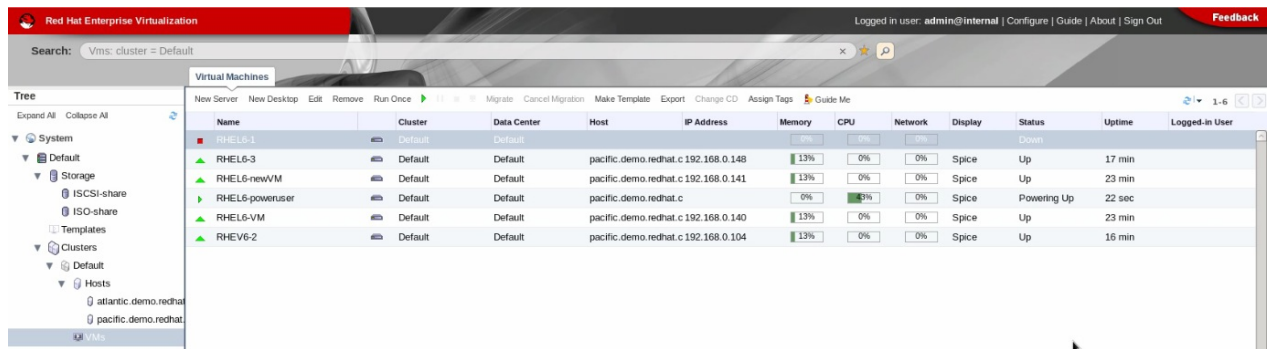
The **Restart Host(s)** dialog displays, click **OK** to confirm and proceed. The host's status changes to *Reboot*, then *Non-Responsive*.



**Figure 6.7. Non-responsive host**

You have now simulated an environment where a host is manually fenced before it was placed into maintenance. Since power management has been configured for this host, it will automatically reboot after a short period.

- While the host is being restarted, observe what has happened to the virtual machines which were running on it. On the **Tree** pane, click **VMs** to display the **Virtual Machines** tab. Notice that both the virtual machines running on **Atlantic** were shut down as soon as the host was restarted.



**Figure 6.8. Virtual machines starting on another host**

The highly available virtual machine, **RHEL6Thames**, is automatically restarted. Its status changes from *Down* to *Wait for Launch*, and then to *Powering Up*. It runs on the **Pacific** host in the interim period while the **Atlantic** host is still rebooting. In contrast, **RHEL6Congo** remains turned off, its status displays as *Down*.

You have just run a demonstration where a host was manually fenced before the virtual machines on it were properly shut down, causing the virtual machines to crash. When virtual machines are not properly stopped, only the highly available ones are restarted on another host in the cluster. In contrast, the non-highly available machines remain powered down until they are manually restarted.

[Report a bug](#)

## 6.7. High Availability: Virtual Machine Interruption

### 6.7.1. High availability -- Virtual machine interruption

In the previous section, you demonstrated virtual machine high availability when a host becomes non-responsive and then rebooted. This section simulates a virtual machine crash. The expected outcome is for the highly available virtual machines to restart automatically, while the non-highly available ones will remain shut down until they are manually restarted.

In your current environment, the host that you restarted in the previous section (**Atlantic**) should be running again. The highly available virtual machine (**RHEL6Thames**) has been restarted on another host (**Pacific**), while the non-highly available virtual machine (**RHEL6Congo**) is still powered down.

For this demonstration, restart **RHEL6Congo**. The **Atlantic** host is automatically selected to run this virtual machine as its workload is currently lighter than **Pacific**'s. Next, live migrate two more virtual machines to the **Atlantic** host so the workload is equally balanced. Select **RHEL6RioGrande** and **RHEL6Thames**, and click **Migrate**. On the **Migrate Virtual Machine(s)** dialog, leave the option as **Select Host Automatically** and click **OK**.

To recap, you now have three virtual machines running on each host. Click on the **Host** label to alphabetically sort the virtual machines according to host.

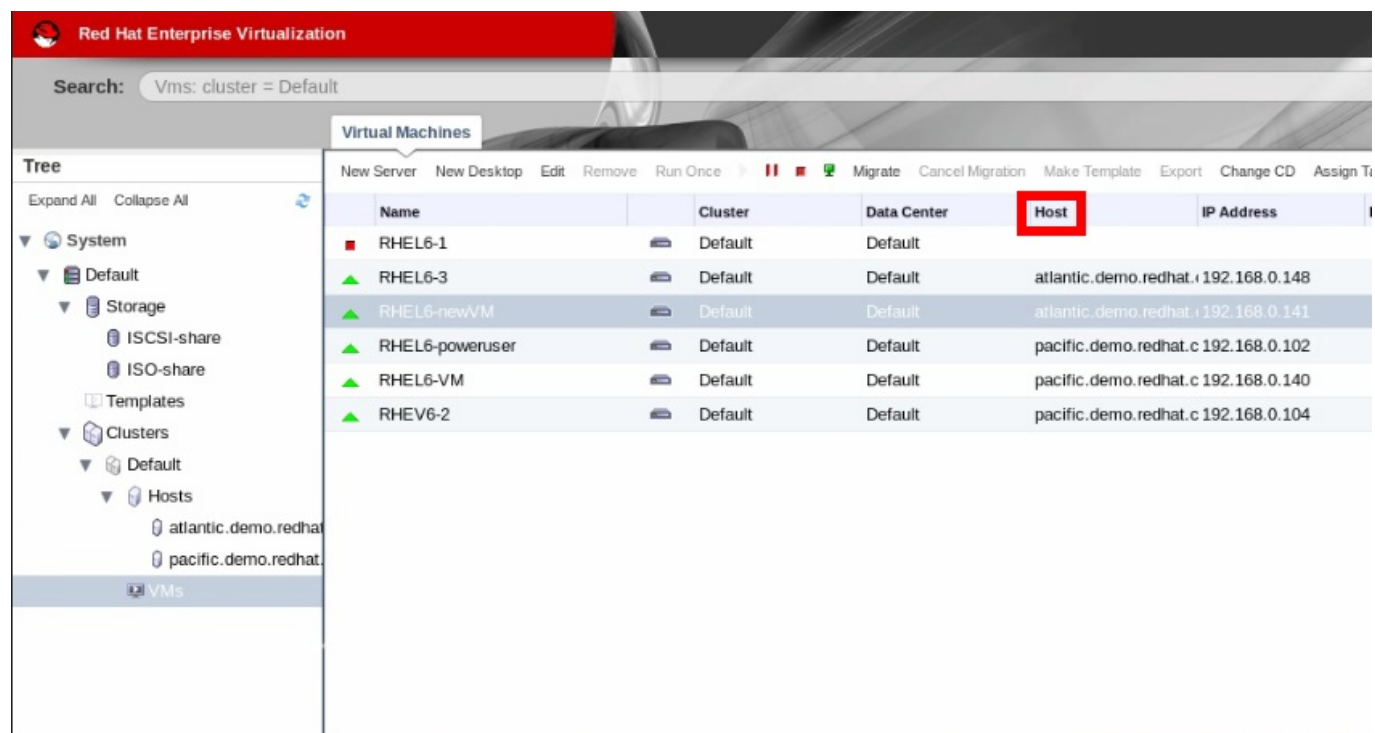


Figure 6.9. Sort virtual machines by host

[Report a bug](#)

## 6.7.2. Demonstrating Virtual Machine High Availability when its Processes are Killed

### Procedure 6.1. Demonstrating Virtual Machine High Availability By Killing the Processes Associated with the Virtual Machine

1. To simulate a virtual machine crash, we will terminate the virtual machine's process from within its host. In this example, we use the **Atlantic** host, and we provide instructions for remotely accessing both hypervisor hosts and Linux hosts.
  - A.
    - a. If you are using a Red Hat Enterprise Virtualization Hypervisor host, you must first enable remote login. On the hypervisor management console, navigate to the **Security** tab and select the **Enable ssh password authentication** check box. Enter your password in the listed fields, and then click **Apply**.



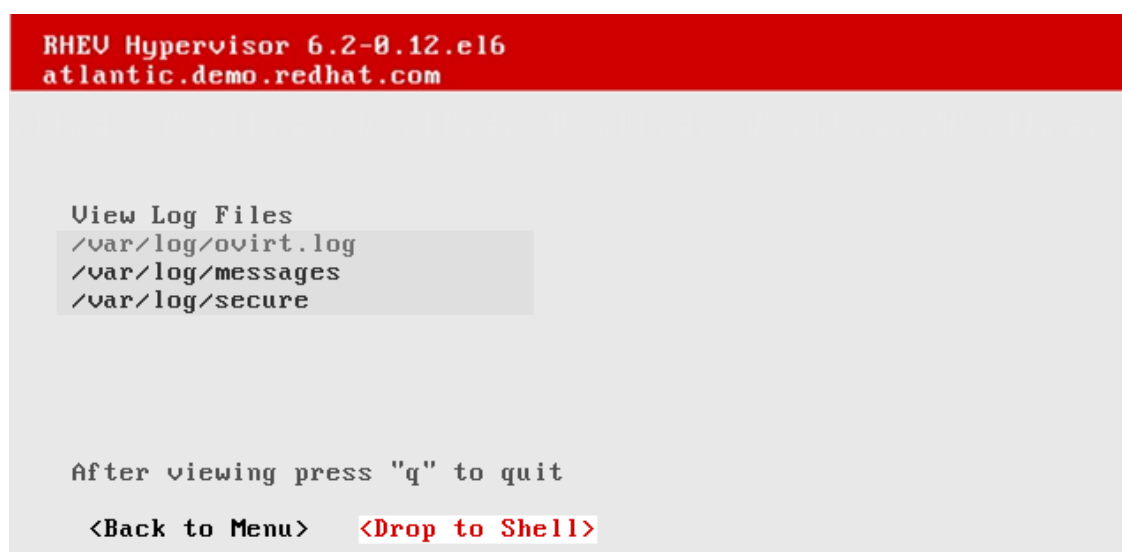
**Figure 6.10. Enable SSH on hypervisor**

SSH login to the hypervisor is not recommended, for security reasons, but we allow it during this lab in order to demonstrate high availability. Disable SSH login to the hypervisor when you complete this lab.

- b. Open a terminal and run:

```
# ssh admin@atlantic.demo.redhat.com
```

- c. Accept the authentication key and enter the password for the hypervisor. When you are logged in to the management console, press the **F2** key and select **Drop to Shell**. You can now execute commands as the root user.



**Figure 6.11. Access hypervisor shell**

- B. a. If you are using a Red Hat Enterprise Linux host, you are able by default to log in to your host remotely. Open a terminal and run:

- b. Accept the authentication key and enter the password for the host.
2. When you have established a remote connection to your host, use the **vdsClient** utility to access the Red Hat Enterprise Virtualization Manager daemon running in the host:

```
[root@atlantic]# vdsClient -s 0 list table
```

This command lists the GUID, process ID, name, and status of each virtual machine on the host. For example:

6aba36aa-d0a9-4d51-bcf3-73544bc95964	2783	RHEL6Congo	Up
6f8bd3cc-3c54-413a-9a28-0bd2f65eeede	3341	RHEL6RioGrande	Up
4e1e0e4f-63e1-4a24-b1a6-76d7ae46e8c4	3015	RHEL6Thames	Up

3. Send the termination signal to two virtual machines - one which is highly available, and one which is not highly available. This example uses **RHEL6RioGrande** and **RHEL6Congo**.

```
[root@atlantic ~]# kill -9 3341
[root@atlantic ~]# kill -9 2783
```

You have now simulated an environment in which a virtual machine has unexpectedly crashed and is inaccessible to users.

4. On the **Virtual Machines** tab in the administration portal, you can see that **RHEL6RioGrande**'s status is **Powering Up**. Because it was marked as highly available, it instantly restarts when it experiences a crash.

In contrast, the status of **RHEL6Congo** remains **Down** because it was not marked as highly available. It remains turned off until it is manually restarted.

This demonstrates that if the virtual machine process is interrupted - whether due to user error, insufficient memory on hosts or other issues - highly available virtual machines will be restarted, while non-highly available ones require manual input. Therefore, if your environment has machines which must be accessible at all times, it is good practice to configure them as highly available, so they will automatically restart if they experience unexpected interruptions.

[Report a bug](#)

## 6.8. High-Availability: Non-Operational Host

### 6.8.1. High Availability - Non-Operational Hosts

Now that you have tested high availability in the event of a failure occurring in a virtual machine, you can examine a scenario in which the failure occurs in a host.

When the status of a host is *Non-Operational*, the host remains accessible from the Red Hat Enterprise Virtualization Manager but cannot serve as a member of a cluster. A *Non-Operational* status signifies that a logical network is down, storage is inaccessible to the host, the type of the host CPU is incompatible with the cluster, or that the host has no logical network.

A non-operational host effects the virtualization environment in several ways. First, if a virtual machines attempts to access a non-operational host, it will be placed into a *Pause* state. Second, all highly available virtual machines will live migrate to another, available host in a cluster. However, only virtual machines that are in an *Up* state are migrated in this manner; virtual machines in a *Pause*

state are not live migrated to prevent data corruption. Last of all, if the non-operational host is an SPM host, the SPM role will be transferred to another host in a data center.

This section demonstrates high availability when a host is non-operational due to being disconnected from the system's external storage resource. You can examine the outcomes of two different cases - when the host is acting as the Storage Pool Manager (SPM), and when it is not.

Before you begin these demonstrations, ensure that all your virtual machines are running. In addition, check which host has been configured as the SPM by clicking **Hosts** on the **Tree** pane. In this example, the **Pacific** host is the SPM.

[Report a bug](#)

### 6.8.2. Demonstrating Virtual Machine High Availability when Storage Network is Down (Host: SPM)

**Procedure 6.2. To demonstrate virtual machine high availability when storage network is down (Host: SPM)**

1. On the **Hosts** tab, select the **Pacific** host to display its details pane. Click the **Network Interfaces** subtab. You should have at least one rhevm network and one storage network, as configured earlier. In this example, the storage network is allocated on **p1p1**.

Next, on the **Tree** pane click **ISCSI-share**. On the **Storage** tab, select the **ISCSI-share** domain and click **Edit**. On the **Edit Domain** dialog, click the **+** symbol to display the iSCSI target. Note that the storage target's subnet mask address is the same as the storage network's, as seen under **Address**.

2. Now that you have determined the name and physical interface of your storage network, connect to the **Pacific** host via SSH. Check your available networks by running the following command:

```
[root@pacific ~]# ifconfig
```

Once you have determined the name of your storage network, run:

```
[root@pacific ~]# ifdown storage
[root@pacific ~]# ifdown p1p1
```

You have now shut down the network between the **Pacific** host and **ISCSI-share** storage. On the **Hosts** tab, the **Pacific** host changes to *Non Operational*, then to *Reboot*.

3. Because the **Pacific** host was configured as the SPM, it is automatically rebooted. The highly available virtual machines are restarted on the other available host in the cluster - in this case it is **Atlantic** - while the non highly available ones are suspended. However, once the **Pacific** host is up and running again, the virtual machines which were originally running on it are migrated back to it in order to balance the workload between all hosts in the cluster.

You have now demonstrated high availability when the connection is disrupted between the storage and the host which is the SPM. However, storage disconnection can also occur with a host which is not acting as the SPM. In this case, the host moves into non-operational status, and its virtual machines migrate to other hosts in the cluster.

[Report a bug](#)



### 6.8.3. Demonstrating virtual-machine high availability when storage networks are down (Host: non-SPM)

Previously, the Pacific host was running as the SPM. However, the role of SPM can be transferred, because it has to be filled by a running host. After Pacific had been rebooted, the Atlantic host gained the status of SPM. Therefore, in this example the **Pacific** host is used again, as it can now play the role of the host which is not the SPM. Before running the next procedure, migrate several virtual machines from Atlantic to Pacific.

#### To demonstrate virtual machine high availability when storage network is down (Host: non-SPM)

1. As before, connect to the **Pacific** host via SSH. Check your available network by running the following command:

```
[root@pacific ~]# ifconfig
```

Once you have determined the name of your storage network, run:

```
[root@pacific ~]# ifdown storage
[root@pacific ~]# ifdown p1p1
```

You have now shut down the network between the **Pacific** host and the **ISCSI-share** storage. If new data is being written onto the virtual machine's disk, the virtual machine detects that the storage connection has been lost, and pauses itself to prevent loss of data. When this happens, the **Hosts** tab shows that the status of the **Pacific** host has changed to *Non Operational*.

2. Click the **pacific.demo.redhat.com** icon on the **Tree** pane to examine the virtual machines. All the virtual machines which were originally running on the **Pacific** host are automatically migrated to the **Atlantic** host. The highly available machines, which were set as high priority, are migrated before the non-highly available ones.

When virtual machines are live migrated, they do not experience any downtime. In rare cases, they will be paused, and then continued on the host they have been migrated to.

The screenshot shows the Red Hat Enterprise Virtualization management console. The top navigation bar includes 'Hosts', 'Storage', 'Virtual Machines', and 'Templates'. The 'Hosts' tab is selected, displaying a table of hosts. The host 'pacific.demo.redhat.com' is highlighted, showing its status as 'Non Operational'.

Name	Hostname/IP	Cluster	Data Center	Status	Load	Memory	CPU	Network	SPM
pacific.demo.redhat.com	192.168.0.5	Default	Default	Non Operational	0 VMs	0%	0%	0%	SPM

The left sidebar shows a tree view with 'System', 'Storage', 'Templates', 'Clusters', and 'Hosts'. The 'Hosts' section is expanded, showing 'atlantic.demo.redhat.com' and 'pacific.demo.redhat.com'.

The bottom pane shows the 'Events' tab for the selected host. It displays a list of events with timestamps and messages:

Time	Message	Correlation Id
2012-Sep-27, 06:11:21	VM RHEL6-powertuser is down. Exit message: Migration succeeded	
2012-Sep-27, 06:11:21	VM RHEL6-VM is down. Exit message: Migration succeeded	
2012-Sep-27, 06:11:21	VM RHEV6-2 is down. Exit message: Migration succeeded	
2012-Sep-27, 06:11:20	Migration complete (VM: RHEL6-powertuser, Source Host: pacific.demo.redhat.com).	64b0d40d
2012-Sep-27, 06:11:20	Migration complete (VM: RHEV6-2, Source Host: pacific.demo.redhat.com).	19471066
2012-Sep-27, 06:11:20	Migration complete (VM: RHEV6-2, Source Host: pacific.demo.redhat.com).	2c1f3401
2012-Sep-27, 06:11:04	Host pacific.demo.redhat.com moved to Non-Operational state because interfaces 'p1p1' are down which needed by networks 'Storage' in the current cluster	
2012-Sep-27, 06:05:54	VM RHEL6-VM started on Host pacific.demo.redhat.com	6f2a3da9



## Figure 6.12. Virtual machines automatically migrated

When the host that holds the SPM role loses connectivity to storage, that host is fenced. VDSM is restarted on the host (restarting VDSM on the host is called "self-fencing").

When the host that runs virtual machines loses network connectivity, that host moves to non-operational status and all of the virtual machines running on it migrate to a new host.

[Report a bug](#)

## 6.9. High Availability: Non-responsive host

### 6.9.1. High availability -- non-responsive host

A host is deemed non-responsive when the Red Hat Enterprise Virtualization Manager cannot communicate with the Red Hat Enterprise Virtualization agent on the host. This can be either due to a networking issue, or failure on the host side (kernel panic, power failure and such) which stops all communication with the host.

When a host is non-responsive, it will be fenced to ensure that virtual machines are allowed to restart on other hosts in the cluster while avoiding "split brain" -- a situation in which communication with the host is lost while the virtual machines are still partially running. This scenario is simulated in the following section, where you will disconnect the host's management network while the storage connection remains functional.

At this stage, the **Pacific** host is non-operational as its storage connection was cut in the previous section.

Restart it for the next demonstration. On the **Tree** pane, select the **Pacific** host. Click the **Power Management** button and select **Restart**. Because you have fenced the host, it automatically brings the **storage** and **p1p1** networks up again, and allows the host to run as normal.

When the host's status changes to *Up*, migrate several machines onto it. This example uses **RHEL6RioGrande**, **RHEL6Thames** (both highly available machines) and **RHEL6Erie**. As you have disabled cluster policy at the beginning of this lab, these virtual machines will not auto-migrate as soon as the host is back up. Therefore, they need to be manually migrated to the **Pacific** host.

[Report a bug](#)

### 6.9.2. Demonstrating high-availability when host connections are disrupted

1. On the **Tree** pane, click **Hosts**. On the **Hosts** tab, select the **Pacific** host, and click the **Network Interfaces** subtab on the details pane. Check the physical interface name of the **rhev** network; in this example it is the **eth0** network.
2. As before, connect to the **Pacific** host via SSH. Disable the management network by running:

```
# ifdown rhvm
```

You have now shut down the network connecting the **Pacific** host to the Red Hat Enterprise Virtualization Manager. The next time that the Manager attempts to transmit signals to the host, it triggers the automatic fencing operation.

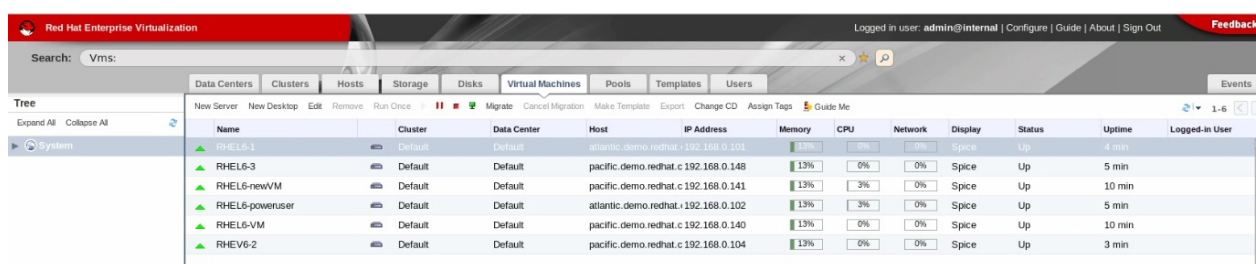
- From the **Tree** pane, click **VMs** to display the **Virtual Machines** tab. The highly available virtual machines, **RHEL6RioGrande** and **RHEL6Thames**, have restarted on the **Atlantic** host. Conversely, **RHEL6Erie** did not restart because it was not configured to be highly available.
- Finally, go to the **Tree** pane and click **Hosts** to examine the status of the hosts. After a short period, the **Pacific** host will be rebooted, assuming that power management was successfully configured on this host.

You have just run a demonstration where a non-responsive host was automatically fenced and rebooted. As you had simulated a non-persistent network failure, the host will recover from the fault following its reboot. In the interim period while it is being restarted, the highly available virtual machines originally running on it are restarted on another available host in the cluster. Conversely, non-highly available virtual machines need to be manually restarted.

[Report a bug](#)

### 6.9.3. Demonstrating virtual-machine high availability of incorrectly-fenced hosts

- On the **Tree** pane, select **VMs**. The available virtual machines display on the **Virtual Machines** tab.

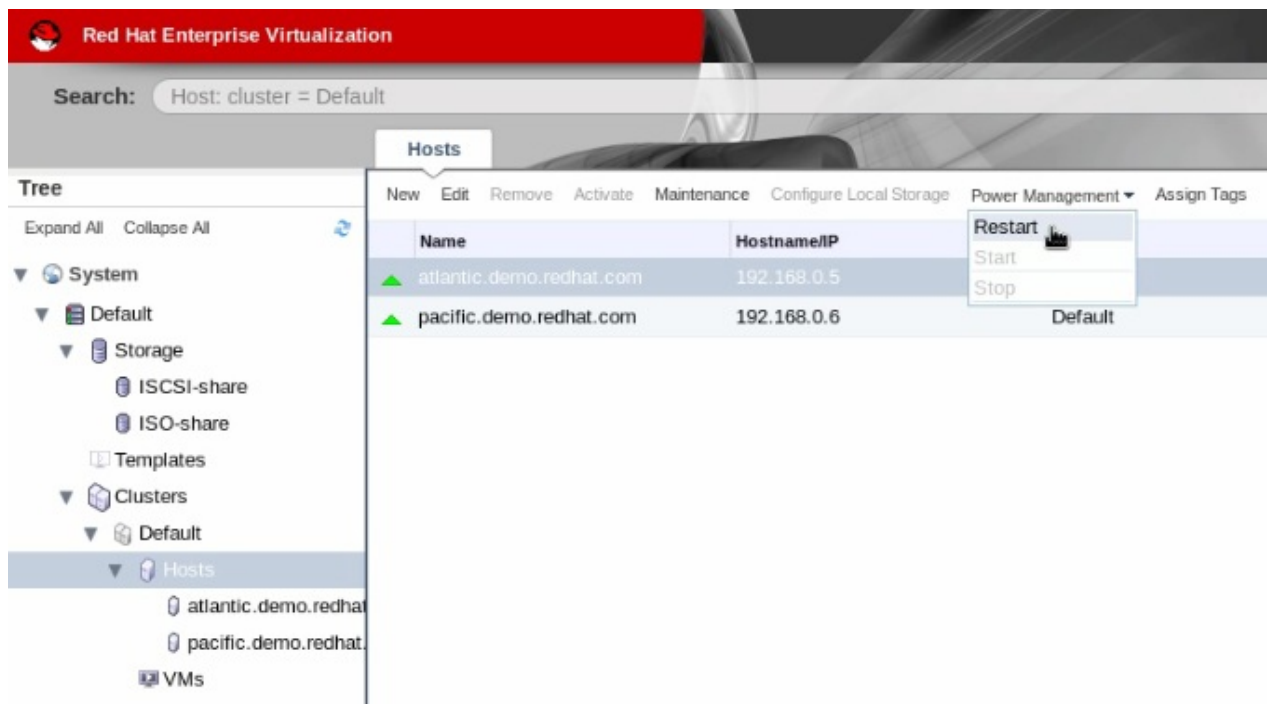


Name	Cluster	Data Center	Host	IP Address	Memory	CPU	Network	Display	Status	Uptime	Logged-In User
RHEL6-1	Default	Default	atlantic.demo.redhat.c	192.168.0.101	13%	0%	0%	Spice	Up	4 min	
RHEL6-3	Default	Default	pacific.demo.redhat.c	192.168.0.148	13%	0%	0%	Spice	Up	5 min	
RHEL6-newVM	Default	Default	pacific.demo.redhat.c	192.168.0.141	13%	3%	0%	Spice	Up	10 min	
RHEL6-poweruser	Default	Default	atlantic.demo.redhat.c	192.168.0.102	13%	3%	0%	Spice	Up	5 min	
RHEL6-VM	Default	Default	pacific.demo.redhat.c	192.168.0.140	13%	0%	0%	Spice	Up	10 min	
RHEV6-2	Default	Default	pacific.demo.redhat.c	192.168.0.104	13%	0%	0%	Spice	Up	3 min	

**Figure 6.13. Virtual machines running on hosts**

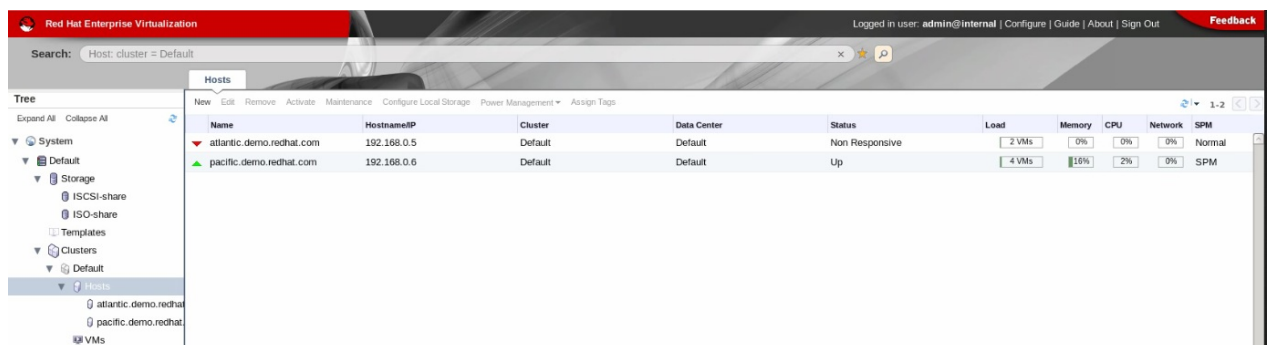
In this example, there are currently two virtual machines running on the **Atlantic** host; and another four on **Pacific**. You can check whether a virtual machine is highly available by selecting it looking at the **General** tab of the details pane. Here, on the **Atlantic** host **RHEL6Thames** is set to be highly available, while **RHEL6Congo** is not.

- On the **Tree** pane, select **Hosts**. The available hosts display on the **Hosts** tab. Select the **Atlantic** host, click the **Power Management** drop-down menu and select **Restart**.



**Figure 6.14. Restart host**

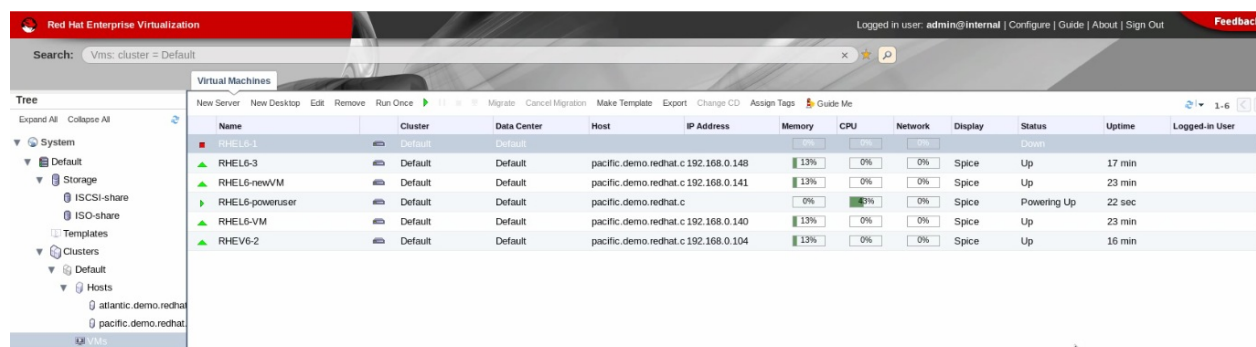
The **Restart Host(s)** dialog displays, click **OK** to confirm and proceed. The host's status changes to *Reboot*, then *Non-Responsive*.



**Figure 6.15. Non-responsive host**

You have now simulated an environment where a host is manually fenced before it was placed into maintenance. Since power management has been configured for this host, it will automatically reboot after a short period.

- While the host is being restarted, observe what has happened to the virtual machines which were running on it. On the **Tree** pane, click **VMs** to display the **Virtual Machines** tab. Notice that both the virtual machines running on **Atlantic** were shut down as soon as the host was restarted.



**Figure 6.16. Virtual machines starting on another host**

The highly available virtual machine, **RHEL6Thames**, is automatically restarted. Its status changes from *Down* to *Wait for Launch*, and then to *Powering Up*. It runs on the **Pacific** host in the interim period while the **Atlantic** host is still rebooting. In contrast, **RHEL6Congo** remains turned off, its status displays as *Down*.


[Report a bug](#)

## 6.10. Lab 6 - Summary

You have now completed the High Availability lab. In this advanced lab, you have successfully enabled power management for your hosts and configured high availability for your virtual machines, and tested their effects in five typical enterprise use cases.

In conclusion, when a host malfunctions, but can still communicate with the Manager, it will try to migrate high priority virtual machines to other hosts in the cluster. If the host cannot connect to the Manager, it will trigger a fencing operation and reboot. In this case, highly available virtual machines will restart in a different host in the same cluster; while non-highly available virtual machines will remain shut down unless manually restarted even when host has finished reboot and is back up.

The next lab in the Advanced Track for Multiple Hosts concerns data centers. It will teach you how to create a new data center with Red Hat Enterprise Linux hosts.



[Report a bug](#)

## Chapter 7. Lab 7 - Adding Additional Data Centers

### 7.1. Lab 7 - Objectives

This lab takes you through the tasks necessary to install and set up Red Hat Enterprise Virtualization with multiple hosts and shared storage. In addition, you will learn how to configure networks and add ISOs in order to create virtual machines. This lab should take you about 35 minutes.

This chapter shows you how to install and configure Red Hat Enterprise Linux hosts for use with Red Hat Enterprise Virtualization Manager. (10 minutes)

This chapter shows you how to create a new Data Center for the Red Hat Enterprise Linux hosts. (1 minute)

This chapter shows you how to create a new host cluster for your Red Hat Enterprise Linux hosts to it. (1 minute)

This chapter shows you how to manually attach the hosts to the Red Hat Enterprise Virtualization Manager. (4 minutes)


This chapter shows you how to define new networks for the storage devices and add them to the hosts. (5 minutes)

This chapter shows you how to define NFS, iSCSI or FCP storage and attach the domains to the data center. (8 minutes)

This chapter shows you how to create a new ISO domain or attach the existing ISO domain to your new data center. (6 minutes)

[Report a bug](#)

### 7.2. Add Additional Data Center



This lab shows you how to customize your standard Red Hat Enterprise Virtualization environment by adding an additional data center. This lab uses Red Hat Enterprise Linux hosts, but you can use Red Hat Enterprise Virtualization Hypervisor hosts instead, if you prefer. In this lab, you will learn how to create virtual machines in your new data center.

This lab assumes that you have successfully completed all the sections in the basic setup chapter. You should have correctly installed and configured Red Hat Enterprise Virtualization, and have configured shared storage and logical networks.

#### Lab 7 - Objectives

This lab takes you through the tasks necessary to install and set up Red Hat Enterprise Virtualization with multiple hosts and shared storage. You will also learn how to configure networks and how to

add ISOs in order to create virtual machines. This lab should take you about 35 minutes.

## Lab 7 - Requirements

Ensure that you have the following:

### Red Hat Enterprise Linux Hosts

- ✧ Minimum - Dual Core server, 2 GB RAM and 10 GB Storage, 1 Gbps network interface
- ✧ Recommended - Dual socket server, 16 GB RAM and 50 GB storage, two 1 Gbps network interfaces.

The breakdown of the server requirements is:

- For each host: AMD-V or Intel VT enabled, AMD64 or Intel 64 extensions, minimum 1 GB RAM, 3 GB free storage and 1 Gbps network interface.
- For virtual machines running on each host: minimum 1 GB RAM per virtual machine.
- ✧ For each host, a valid Red Hat Network subscription to the **rhel-x86\_64-rhev-mgmt-agent-6** channel.

### Storage and Networking

- ✧ One or more of the supported storage types (NFS, iSCSI and FCP).
- ✧ One static IP per host which is resolvable by the DNS server.
- ✧ An NFS mount point for an ISO directory if you choose to configure an ISO share external to the Red Hat Enterprise Virtualization Manager Server.

## Lab 7 - Configuration

The following figure and table list the environment parameters and object names which will be used consistently throughout this lab. It is strongly recommended that you use these entities in your evaluation environment to ensure the names are resolvable. You may alter them if necessary, but make sure you have an equivalent name for each component.



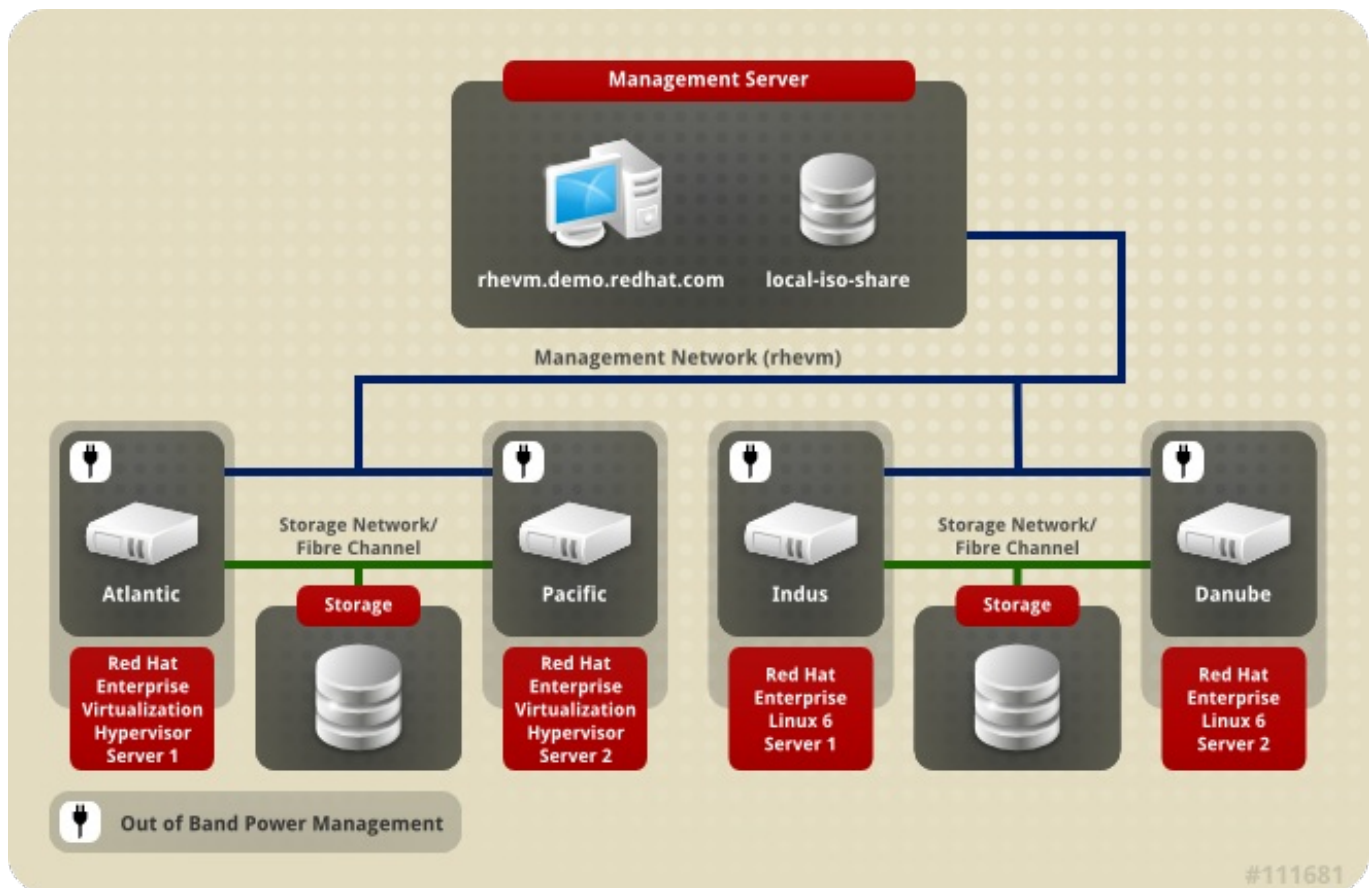


Table 7.1. Lab component names

Component	Name	IP (if applicable)	Fully Qualified Domain Name
Domain Services	-	-	demo.redhat.com
Red Hat Enterprise Virtualization Manager	-	-	rhev.demo.redhat.com
Data Center	-	-	FinanceDataCenter
Cluster	-	-	FinanceCluster
Storage Network	storage	10.23.1.0/24	-
Management Network	rhev	10.35.3.0/24	-
Red Hat Enterprise Virtualization Linux Host 1	Danube	-	danube.demo.redhat.com
Red Hat Enterprise Virtualization Linux Host 2	Indus	-	indus.demo.redhat.com
Red Hat Enterprise Virtualization Hypervisor 1	Atlantic	-	atlantic.demo.redhat.com
Red Hat Enterprise Virtualization Hypervisor 2	Pacific	-	pacific.demo.redhat.com
Administrator User Name	admin	-	-
NFS Storage Domain	NFS-share	-	-
iSCSI Storage Domain	iSCSI-share	-	-
FCP Storage Domain	FCP-share	-	-

Component	Name	IP (if applicable)	Fully Qualified Domain Name
ISO Storage Domain	local-iso-share	-	-

[Report a bug](#)

## 7.3. Installing Red Hat Enterprise Linux 6 hosts

In your current environment, you should have Red Hat Enterprise Virtualization Manager and at least two hosts configured in the default data center. In this section, you will learn how to configure Red Hat Enterprise Linux servers as hosts to run virtual machines. You have to perform the installation on each physical server you wish to use as a host.

However, if you do not wish to use Red Hat Enterprise Linux hosts, you can still create an additional data center and add Red Hat Enterprise Virtualization Hypervisor hosts to it.



### Important

Red Hat Enterprise Virtualization Hypervisor is a closed appliance, and does not allow the installation of custom RPMs.

If you require custom RPMs, you must use Red Hat Enterprise Linux hosts. Red Hat Enterprise Linux hosts are not closed appliances, and they allow the installation of custom RPMs.

### To install a Red Hat Enterprise Linux 6 host

1. On the machine designated as your Red Hat Enterprise Linux host, install the latest version of Red Hat Enterprise Linux. Select only the Base package group during installation. For more comprehensive instructions, refer to the *Red Hat Enterprise Linux Installation Guide*.
2. If your server has not been registered with the Red Hat Network, run the **rhn\_register** command as **root** to register it. To complete registration successfully you will need to supply your Red Hat Network user name and password. Follow the onscreen prompts to complete registration of the system.

```
# rhn_register
```

3. Subscribe the server to the required channels using the Red Hat Network web interface.
  - a. Log on to Red Hat Network (<http://rhn.redhat.com/>).
  - b. Click **Systems** at the top of the page.
  - 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 **Red Hat Enterprise Virt Management Agent (v 6 x86\_64)** channel from the list presented on the screen, then click the **Change Subscription** button to finalize the change.
4. Ensure the kernel and all the packages are up to date. This may take some time. Run:

```
# yum -y update
```

5. Red Hat Enterprise Virtualization platform uses a number of network ports for management and other virtualization features. Adjust your Red Hat Enterprise Linux host's firewall settings to allow access to the required ports by configuring **iptables** rules. Modify the **/etc/sysconfig/iptables** file so it resembles the following example:

```
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [10765:598664]
-A INPUT -m state --state RELATED,ESTABLISHED -j ACCEPT
-A INPUT -p icmp -j ACCEPT
-A INPUT -i lo -j ACCEPT
-A INPUT -p tcp --dport 22 -j ACCEPT
-A INPUT -p tcp --dport 16514 -j ACCEPT
-A INPUT -p tcp --dport 54321 -j ACCEPT
-A INPUT -p tcp -m multiport --dports 5634:6166 -j ACCEPT
-A INPUT -p tcp -m multiport --dports 49152:49216 -j ACCEPT
-A INPUT -p tcp -m state --state NEW
-A INPUT -j REJECT --reject-with icmp-host-prohibited
-A FORWARD -m physdev ! --physdev-is-bridged -j REJECT --reject-
with icmp-host-prohibited
COMMIT
```

6. Ensure that the **iptables** service is configured to start on boot and has been restarted, or started for the first time if it was not already running. Run the following commands:

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

You have now successfully installed your Red Hat Enterprise Virtualization Linux hosts. The following sections will provide instructions on how to attach the hosts to the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

## 7.4. Create a New Data Center

In the basic labs, you used the Default data center which was automatically created upon installation. However, in this lab you will create a new data center. A data center is a logical entity that defines the set of physical and logical resources used in a managed virtual environment. It is a container for clusters of virtual machines, storage and networks.



## Note

While this lab provides the example of creating a new data center consisting only of Red Hat Enterprise Linux hosts, in reality you can actually add both types of hosts (Red Hat Enterprise Virtualization Hypervisor and Red Hat Enterprise Linux) to the same cluster.

From this point onwards, you should be logged in to the administration portal as the **admin** user.

### To add a new Data Center

1. Navigate to the **Tree** pane and click **Expand All**. Click **System** and select the **Data Centers** tab.
2. Click the **New** button. The **New Data Center** dialog displays.

The screenshot shows a 'New Data Center' dialog box. It has a title bar with a red 'x' icon and a question mark. The dialog contains several input fields: 'Name', 'Description', 'Type' (a dropdown menu showing 'NFS'), 'Compatibility Version' (a dropdown menu showing '3.3'), 'Quota Mode' (a dropdown menu showing 'Disabled'), and 'Comment'. At the bottom right are 'OK' and 'Cancel' buttons.

**Figure 7.1. New Data Center**

Fill in the **Name** and **Description** fields. Select the data center **Type** from the drop-down menu. Data center type must be compatible with the storage type of the new data center.

3. Click **OK**. The **New Data Center - Guide Me** dialog displays. This dialog provides a context-sensitive reference for the components which have not been configured for your data center. Leave it open for now.



Figure 7.2. New Data Center - Guide Me

[Report a bug](#)

## 7.5. Create a New Cluster

Now that you have created a new data center, you can populate it with a new cluster. A cluster is a set of physical hosts that are treated as a resource pool for a set of virtual machines. A cluster constitutes a domain for virtual machines to migrate from one host to another, therefore all hosts in a cluster must share the same network infrastructure and the same storage. However, the hosts do not need to share the exact same CPU model, as long as the cluster's CPU name is set to the family that is the most common denominator.

### To add a new host cluster

1. You should still be on the **New Data Center - Guide Me** dialog. If you are not, select your newly created data center and click **Guide Me**. On the dialog, click **Configure Cluster**.
2. The **New Cluster** dialog displays. Note that there are multiple tabs on the left side of this dialog box. Click on each tab to access them in turn.

**Figure 7.3. New Cluster**

Configure the following options:

- ✳ On the **General** tab, enter a suitable **Name** and **Description**.

On the **CPU Name** field, select the family that best describes your host CPUs. Note that for the Default cluster, you did not need to perform this step because it had been automatically set according to the first host you added to the cluster.

- ✳ On the **Optimization** tab, click the **For Desktop Load - Enable memory page sharing to 200%** radio button.
- ✳ On the **Resilience Policy** tab, select **Do Not Migrate Virtual Machines**.

3. Click **OK**. The **New Data Center - Guide Me** dialog displays again. Leave it open for now.

[Report a bug](#)

## 7.6. Attach New Red Hat Enterprise Linux Hosts

You have created a new data center and a new cluster within it, now you need to populate the cluster with the Red Hat Enterprise Linux hosts you installed earlier. Before they can be used, they have to be manually attached to the Red Hat Enterprise Virtualization platform. Perform the following procedure for each host you have installed.

Note that this step is required for Red Hat Enterprise Linux hosts only, if you are using Red Hat Enterprise Virtualization Hypervisors for this lab you only have to approve the hosts with a click, as you did in the Multiple Host chapter. This time, attach the hosts to your newly created cluster.

### To attach Red Hat Enterprise Linux hosts

1. You should still be on the **New Data Center - Guide Me** dialog. If you are not, select your newly created data center and click **Guide Me**. On the dialog, click **Configure Host**.
2. The **New Host** dialog displays.

**New Host** ?

**General**

Power Management

SPM

Console

Network Provider

Data Center: 32\_dc

Host Cluster: 32\_dc

☐ Use External Providers

Name:

Comment:

Address:

SSH Port: 22

**Authentication**

User Name: root

☒ Password

☐ SSH PublicKey

☐ Advanced Parameters

OK Cancel

**Figure 7.4. Attach a Red Hat Enterprise Linux Host**

Configure the following options:

- ✧ The **Data Center** and **Host Cluster** are already pre-selected.

- ✱ Enter the **Name** and the IP **Address** or resolvable hostname of the host. Enter the **Root Password** used to log in to the host.
  - ✱ If you wish to configure Out of Band (OOB) power management for this host, see the Clustering Power Management section.
3. Click **OK**. On the **Tree** pane, click **Expand All** and select the **Danube** host. It displays on the **Hosts** tab with a status of "Installing". The installation may take a few minutes as the required packages are being downloaded. Once installation is complete, the status will update to *Reboot* and then *Awaiting*.
  4. When the host is ready for use, its status changes to **Up**.

[Report a bug](#)

## 7.7. Configure Logical Networks

Now that you have a data center with hosts grouped in a cluster, you need to define and apply the networking layer. By default, the Management network is already defined for the new data center you created. You can define new logical networks, for example data, storage or display networks.

You have already defined a storage network for the Default data center and added it to the hosts in the Default cluster. Now, you will repeat the same procedure for the FinanceCluster in the FinanceDataCenter.

### To define logical networks in a cluster

1. Navigate to the **Tree** pane and click the **Expand All** button. Click **FinanceDataCenter**. On the results list, the **FinanceDataCenter** data center displays. Click on it to display the details pane.
2. On the details pane, select the **Logical Networks** tab and click **New**. The **New Logical Network** dialog displays. Fill in the **Name** and **Description** fields, and tick the **Assign Networks to Cluster(s) - FinanceCluster** check box to automatically add the Storage network to the FinanceCluster.
3. Click **OK** to create the new logical network.

Now that you have defined this network as a resource required by the FinanceCluster in the data center, it is time to add this resource to the hosts in the cluster.

### To add network to hosts

1. Back on the **Tree** pane, click **FinanceDataCenter** → **Clusters** → **FinanceCluster** → **Hosts**. The **Hosts** tab displays a list of available hosts.
2. For each of your installed hosts, perform the following tasks:
  - a. Click on the host. On the details pane, select the **Network Interfaces** tab.
  - b. A list of network interfaces available for this host displays. One of them will already have the management network (rhevmm) configured.

- c. Select the interface on which to configure the Storage network and click the **Setup Host Networks** button. The **Setup Host Networks** window displays. Configure the following options:
  - ✦ Under **Network**, select *storage*.
  - ✦ Select the **Static** radio button. Enter the **IP** and **Subnet Mask** you have prepared as part of the prerequisites to this lab.
  - ✦ Tick the **Save network configuration** check box.
- d. Click **OK**.

You have now added a new storage network to your data center, and attached the network to your hosts. On the **Logical Networks** tab of **FinanceDataCenter**, you should have at least two networks - *rhev* and *storage*. Now, you can add storage to the system.

[Report a bug](#)

## 7.8. Configuring Storage

### 7.8.1. Creating an NFS Data Domain

Because you have selected NFS as your default storage type during the Manager installation, you will now create an NFS storage domain. An NFS type storage domain is a mounted NFS share that is attached to a data center and used to provide storage for virtual machine disk images.



#### Important

If you are using NFS storage, you must first create and export the directories to be used as storage domains from the NFS server. These directories must have their numerical user and group ownership set to 36:36 on the NFS server, to correspond to the *vdsm* user and *kvm* group respectively on the Red Hat Enterprise Virtualization Manager server. In addition, these directories must be exported with the read write options (*rw*). For more information see the *Red Hat Enterprise Virtualization Installation Guide*.

#### To add NFS storage

1. Navigate to the **Tree** pane and click the **Expand All** button. Under System, select the **Default** data center and click on **Storage**. The available storage domains display on the results list. Click **New Domain**.
2. The **New Domain** window displays.

**Figure 7.5. Add NFS Storage - New Domain Window**

Configure the following options:

- a. **Name:** Enter **NFS-share**.
  - b. **Data Center:** The **Default** data center is already pre-selected.
  - c. **Domain Function / Storage Type:** The **Data/ NFS** option is already pre-selected because during installation you set NFS as your data center's default storage type. The storage domain types which are not compatible with the Data Center will not be available.
  - d. **Use Host:** Select any of the hosts from the drop down menu. Only hosts which belong in this data center will display in this list.
  - e. **Export Path:** Enter the IP address or a resolvable hostname of the chosen host. The export path should be in the format of **192.168.0.10 : /Images/NFS-Share** or **domain.example.com : /Images/NFS-Share**.
3. Click **OK**. The new **NFS-share** 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.

You have created an NFS storage domain. Now, you need to attach an ISO domain to the data center and upload installation images so you can use them to create virtual machines.



[Report a bug](#)

## 7.8.2. Creating an iSCSI Data Domain

Because you have selected iSCSI as your default storage type during the Manager installation, you will now create an iSCSI storage domain. Red Hat Enterprise Virtualization platform supports iSCSI storage domains spanning multiple pre-defined Logical Unit Numbers (LUNs).

### To add iSCSI storage

1. Navigate to the **Tree** pane and click the **Expand All** button. Under System, select the **Default** data center and click on **Storage**. The available storage domains display on the results list. Click **New Domain**.
2. The **New Domain** dialog box displays.

**Figure 7.6. Add iSCSI Storage**

Configure the following options:

- a. **Name:** Enter **ISCSI-share**.
- b. **Data Center:** The **Default** data center is already pre-selected.

- c. **Domain Function / Storage Type:** The **Data/ iSCSI** option is already pre-selected because during installation you set iSCSI as your data center's default storage type. The storage domain types which are not compatible with the Data Center will not be available.
  - d. **Use Host:** Select any of the available hosts from the drop down menu. Only hosts which belong in this data center will display in this list.
3. To connect to the iSCSI target, enter the required information under the **Discover Targets** bar.

**New Domain**

Name:

Data Center:

Domain Function / Storage Type:  Format:

Use Host:

**Discover Targets**

Address:  ☐ User Authentication:

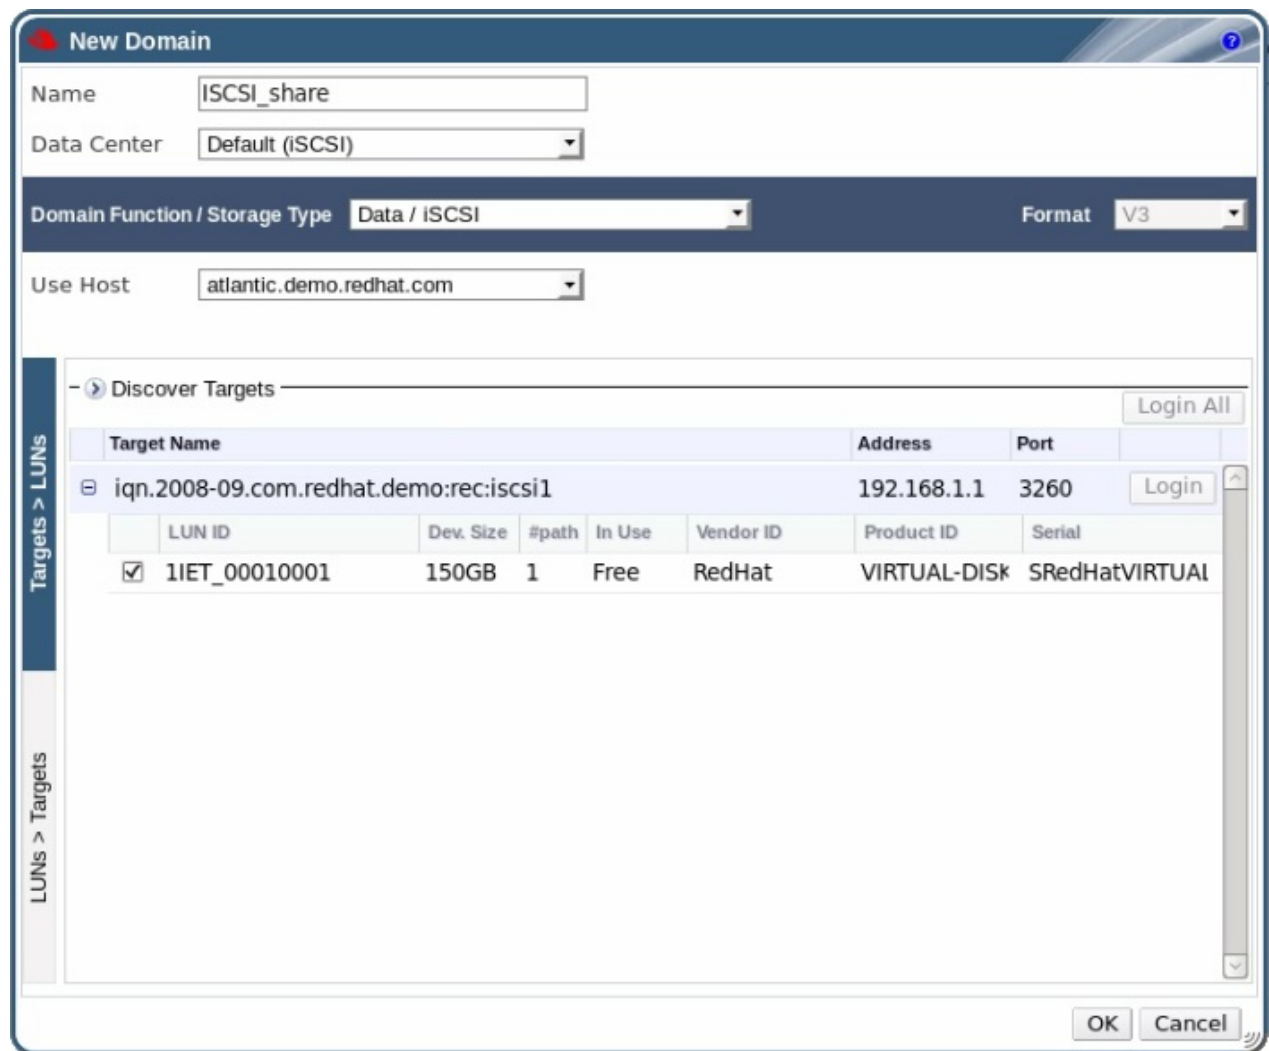
Port:  CHAP username:  CHAP password:

Target Name	Address	Port	
iqn.2008-09.com.redhat.demo:rec:iscsi1	192.168.1.1	3260	<input type="button" value="Login"/>
iqn.2008-09.com.redhat.demo:rec:iscsi_20g	192.168.1.1	3260	<input type="button" value="Login"/>

**Figure 7.7. Discover iSCSI target**

- a. **Address:** Enter the address of the iSCSI target.
  - b. **Port:** Select the port to connect to. The default is 3260.
  - c. **User Authentication:** If required, enter the user name and password.
- Click the **Discover** button to find the targets. The iSCSI targets display in the results list with a **Login** button for each target.
4. Click **Login** on the first target to display the list of existing LUNs. Click the **+** icon under the Target Name to expand the LUN list. Tick the **Add LUN** check box to use the selected LUN as the iSCSI data domain. LUNs that are part of a storage domain in the current setup are disabled. LUNs used by the host (that is, LUNs that are either part of a volume group or that

are used as partitions by other devices) display as being in use. You can choose LUNs that are in use, but you will have to forcefully override their contents.



**Figure 7.8. Attach LUNs to iSCSI domain**

5. Click **OK**. The new **iSCSI-share** 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.

You have created an iSCSI storage domain. Now, you need to attach an ISO domain to the data center and upload installation images so you can use them to create virtual machines.

[Report a bug](#)

### 7.8.3. Creating an FCP Data Domain

Because you have selected FCP as your default storage type during the Manager installation, you will now create an FCP storage domain. Red Hat Enterprise Virtualization platform supports FCP storage domains spanning multiple pre-defined Logical Unit Numbers (LUNs).

#### To add FCP storage

1. Navigate to the **Tree** pane and click the **Expand All** button. Under System, select the **Default** data center and click on **Storage**. The available storage domains display on the results list. Click **New Domain**.
2. The **New Domain** dialog box displays.

**New Domain**

Name

Data Center

Domain Function / Storage Type  Format

Use Host

Export Path   
Remote path to NFS export, takes either the form: FQDN:/path or IP:/path e.g. server.example.com:/export/VMs

[Advanced Parameters](#)

OK Cancel

**Figure 7.9. Add FCP Storage**

Configure the following options:

- a. **Name:** Enter **FCP - share**.
- b. **Data Center:** The **Default** data center is already pre-selected.
- c. **Domain Function / Storage Type:** The **Data/ Fibre Channel** option is already pre-selected because during installation you set FC as your data center's default storage type. The storage domain types which are not compatible with the Data Center will not be available.
- d. **Use Host:** Select any of the hosts from the drop down menu. Only hosts which belong in this data center will display in this list. Ensure that the LUN you intend to use is available on the host you select. LUNs that are part of a storage domain in the current setup are disabled. LUNs used by the host (that is, LUNs that are either part of a volume group or that are used as partitions by other devices) display as being unavailable. You can choose LUNs that are in use, but you will have to forcefully override their contents.

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

You have created an FCP storage domain. Now, you need to attach an ISO domain to the data center and upload installation images so you can use them to create virtual machines.

[Report a bug](#)

## 7.9. Configure ISO Domain

Now that you have configured your data storage domains, the next step is to define an ISO domain for your new data center. If you have an additional NFS mount point, you can choose to create a new ISO domain and attach it to your new data center. Alternatively, you can also use the local ISO domain created during the Manager installation. An ISO domain can be shared across data centers, even if the data center storage types are different.

### Option 1: To create a new ISO domain

1. On the side pane, select the **Tree** tab and click **Expand All**. Under the **FinanceDataCenter** data center, select **Storage**. On the results list, click **New Domain**.
2. The **New Storage** dialog box displays. Configure the following options:
  - a. **Name**: Enter **iso-domain**.
  - b. **Data Center**: The **FinanceDataCenter** data center is already pre-selected.
  - c. **Domain Function / Storage Type**: In the drop down menu, select **Data** → **NFS**. The storage domain types which are not compatible with the Default data center are grayed out. After you select your domain type, the **Export Path** field appears.
  - d. **Use Host**: Select any of the hosts from the drop down menu. Only hosts which belong in this data center will display in this list.
  - e. **Export Path**: Enter the IP address or a resolvable hostname of the chosen host. The export path should be in the format of **192.168.0.10 : /Images/NFS-Share** or **domain.example.com : /Images/NFS-Share**.
3. Click **OK**. The new **iso-domain** displays on the Storage tab. It displays with a *Locked* status as the domain is being validated before becoming *Active*.

You have now successfully defined an ISO domain for your new data center, and can now upload ISO images to it. Once you have populated your ISO domain, you can create virtual machines.

Alternatively, if you do not have an NFS mount point for a new ISO domain, use the **local-iso-share** domain which was created on the Manager server.

### Option 2: To attach the existing ISO domain to a new data center

1. On the side pane, select the **Tree** tab and click **Expand All**. Click **FinanceDataCenter**. On the results list, the **FinanceDataCenter** data center displays. Click on it to display the details pane.

2. On the details pane, select the **Storage** tab and click the **Attach ISO** button. Select the **local-iso-share** domain and click **OK**.
3. The ISO domain appears in the results list of the **Storage** tab. It displays with the *Locked* status as the domain is being validated, then transits to *Inactive*.
4. Select the ISO domain and click the **Activate** button. The status changes to *Locked* and then to *Active*.

You have now successfully attached the ISO domain from the Manager server to your new data center. All the installation images which were available in the Default data center can now be used for the **FinanceDataCenter** data center. You can create more virtual machines on your new hosts. Just remember to substitute the default cluster for **FinanceCluster** and the default data center for **FinanceDataCenter**.

[Report a bug](#)

## 7.10. Lab 7 - Summary

You have now completed the Adding Additional Data Centers lab.. In this lab, you have successfully installed Red Hat Enterprise Linux hosts, attached them to an additional data center, and configured additional resources for your new data center.

[Report a bug](#)

## Chapter 8. Lab 8 - Virtual desktops

### 8.1. Lab 8 - Objectives

This lab takes you through the tasks necessary to create and then connect to a Windows virtual machine or desktop pool with UserRole privileges. This lab should take you about 50 minutes.

This chapter shows you how to add an Active Directory domain to the Red Hat Enterprise Virtualization Manager in order to use Windows virtual desktop pools. (5 minutes)

This chapter shows you how to create a new Windows virtual machine, configure storage and networking, and install the operating system. (15 minutes)

This chapter shows you how to create a thin provisioned virtual machine from the Windows template and apply the sysprep configuration settings. (5 minutes)

This chapter shows you how to assign UserRole permissions for a virtual machine to a user. (2 minutes)

This chapter shows you how to log in to the User Portal and connect to a virtual machine. (3 minutes)

This chapter shows you how to create a Windows desktop pool, assign UserRole permission for the desktop pool and how to connect to a virtual machine in a pool. (10 minutes)

[Report a bug](#)

### 8.2. Create a Windows Virtual Machine

You already know how to create a Red Hat Enterprise Linux virtual machine from scratch. The procedure for creating a Windows virtual machine is similar, except that it requires additional virtio drivers. This example uses Windows 7, but you can also use other Windows operating systems. You will perform a normal attended installation using a virtual DVD.

#### To create a Windows server

1. Navigate to the **Tree** pane and click **Expand All**. Click the **VMs** icon in the **Default** cluster under the **Default** data center. On the **Virtual Machines** tab, click **New VM**.

**New Virtual Machine**

**General**

Cluster: Default/Default

Based on Template: Blank

Operating System: Windows 7

Optimized for: Desktop

Name: Win7-desktop1

Description:

Comment:

☐ Stateless ☐ Start in Pause Mode ☐ Delete Protection

VM has no network interfaces. To add one, assign a profile.

nic1: [dropdown] +

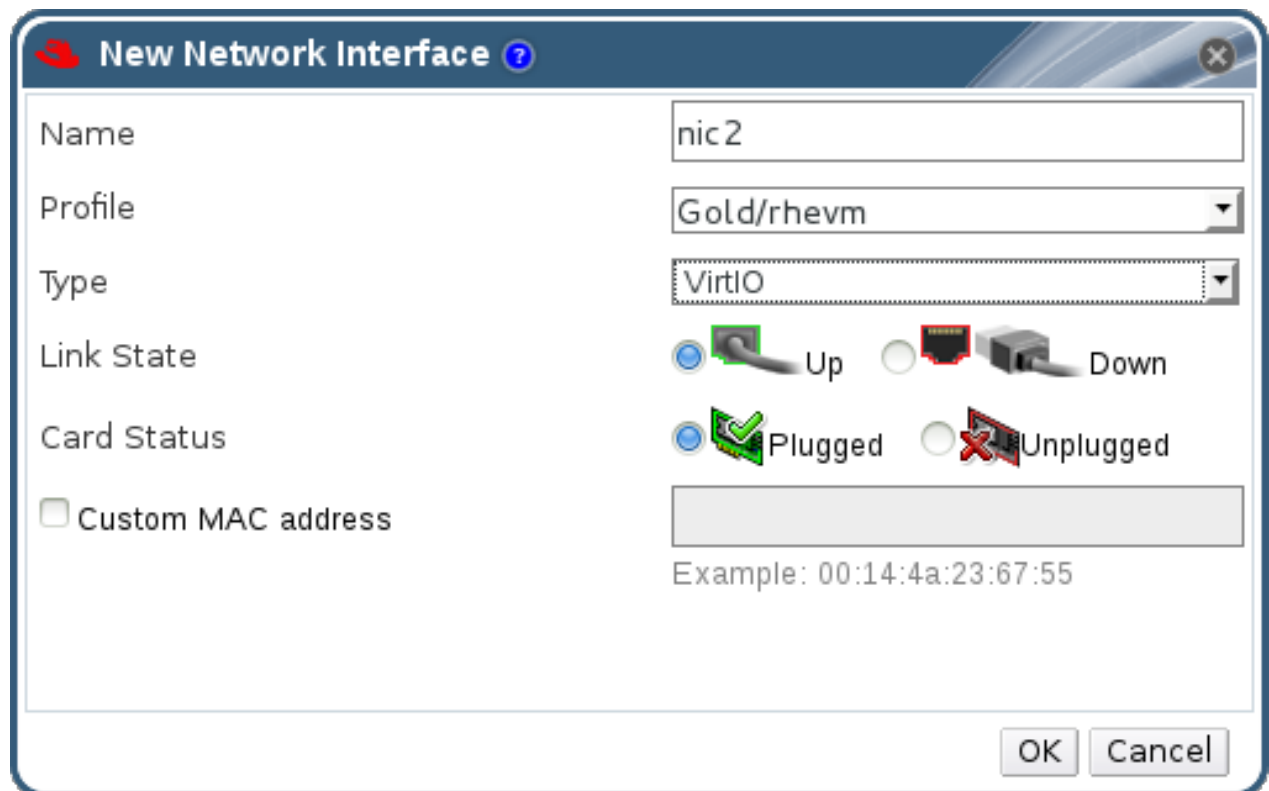
Show Advanced Options OK Cancel

**Figure 8.1. Create new Windows desktop**

You only need to fill in the **Name** field and select **Windows 7** as your **Operating System**. You may alter other settings, but this example retains the defaults. Click **OK** to create the virtual machine.

2. A **New Virtual Machine - Guide Me** window opens. This allows you to define networks for the virtual machine. Click **Configure Network Interfaces**.





**Figure 8.2. New Network Interface configurations**

Retain all the default settings and click **OK**.

3. You are returned to the **Guide Me** window. This time, click **Configure Virtual Disks** to add storage to the virtual machine.

**Add Virtual Disk**

☒ Internal ☐ External (Direct Lun)

Size(GB)

Alias

Description

Interface

Format

Data Center

Storage Domain

☐ Wipe after delete

☐ Is bootable

☐ Is shareable

OK Cancel

**Figure 8.3. New Virtual Disk configurations**

In the **Size (GB)** field, enter *15*. Retain all other default settings and click **OK**.

4. Close the Guide Me window. Your new Windows 7 virtual machine displays in the **Virtual Machines** tab.

You have now created your first Windows desktop virtual machine. Before you can use your virtual machine, install an operating system on it.

### To install the Windows guest operating system

1. Right click the virtual machine and select **Run Once**. Configure the following options:
  - ✧ **Attach Floppy:** virtio-win
  - ✧ **Attach CD:** Your Windows installation media
  - ✧ **Boot sequence:** CD-ROM
  - ✧ **Display protocol:** SPICE

Retain the default settings for the other options and click **OK** to start the virtual machine.

2. Select the virtual machine and click the **Console** icon. If this is your first time connecting to the virtual machine, allow the installation of SPICE ActiveX and the SPICE client.
3. After the SPICE plug-ins have been installed, select the virtual machine and click **Console** again. This displays a window to the virtual machine, where you will be prompted to begin the installation process.
4. Accept the default settings and enter the required information as necessary. The only change you must make is to manually install the VirtIO drivers from the virtual floppy disk (vfd) image. To do so, select the **Custom (advanced)** installation option and click **Load Driver**. Press Ctrl and select:

➤ **Red Hat VirtIO Ethernet Adapter**

➤ **Red Hat VirtIO SCSI Controller**

5. The installation process commences, and the system will reboot itself several times.

When the virtual machine has finished installing, it is time to install **RHEV-tools** which provide some required features including optimizing SPICE display resolution. Press **F12** to release the mouse and click on the SPICE symbol on the upper left hand corner of the SPICE window. In the menu, select **Change CD → RHEV-toolsSetup** to attach the Guest Tools ISO to the virtual machine.

6. On the virtual machine, locate the CD drive (usually D drive) and install the Red Hat Enterprise Virtualization Application provisioning tool. This tool checks for RHEV-tools updates on every reboot.
7. After the installation of the tool, wait for a few minutes. During this time, you will not see any apparent activity because the tool runs as a service. The guest will automatically reboot to apply changes.

You have completed the installation of your Windows virtual machine. You can now log in and start using it. In the next lab you will convert this machine into a template from which you will be able to provision multiple similar machines.

[Report a bug](#)

## 8.3. Creating Windows virtual machines from templates

In the previous section, you created a Windows template complete with pre-configured storage, networking and operating system settings. Now, you will use this template to deploy a pre-installed virtual machine.

### To create a Windows virtual machine from a template

1. Navigate to the **Tree** pane and click **Expand All**. Click the **VMs** icon in the **Default** cluster under the **Default** data center. On the **Virtual Machines** tab, click **New VM**.

**New Virtual Machine**

**General**

Cluster: Default/Default

Based on Template: Win7-template

Operating System: Windows 7

Optimized for: Desktop

Name: Win7-desktop2

Description:

Comment:

☐ Stateless ☐ Start in Pause Mode ☐ Delete Protection

☐ Copy Template Permissions

VM has 1 network interface. Assign a profile to it.

nic1: rhev/rhev

nic2:

Hide Advanced Options OK Cancel

**Figure 8.4. Clone a Windows desktop**

- a. On the **General** tab, select your newly created Windows template called **FinanWinDesk** from the **Based on Template** list.
- b. Enter a suitable **Name** and **Description**, and accept the default values inherited from the template in the rest of the fields. You can change them if needed.
2. Retain all other default setting and click **OK** to create the virtual machine. The virtual machine displays in the Virtual Machines list with a status of "Image Locked" until the virtual disk is created. The virtual disk and networking settings are inherited from the template, and do not have to be reconfigured.
3. Click the **Run** icon to turn it on. This time, the **Run Once** steps are not required as the operating system has already been installed onto the virtual machine hard drive. Click the green **Console** button to connect to the virtual machine.
4. The Windows virtual machine goes through the unattended setup process to apply the customized settings set by sysprep using the answer file provided. When this process ends, log in to the machine. You can see that the virtual machine has joined the Active Directory domain.

You have now learned how to create Windows virtual machines with and without templates. Next, you will learn how to access these virtual machines from a user portal.

[Report a bug](#)

## 8.4. Create a Windows Template

Now that you have created a Windows virtual machine, you can save its settings into a template. This template will retain the original virtual machine's configurations, including virtual disk and network interface settings, operating systems and applications. You can use this template to rapidly create replicas of the original virtual machine. For this lab, name your template **FinanWinDesk**.

Before your virtual machine can be used to create a template, it has to be sealed with **sysprep**. This ensures that machine-specific settings are not propagated through the template.

Note that the procedure below is applicable for creating Windows 7 and Windows 2008 R2 templates. To seal a Windows XP template, see the *Red Hat Enterprise Virtualization Administration Guide*.

### To seal a Windows virtual machine with sysprep

1. In the Windows virtual machine to be used as a template, open a command line terminal and type **regedit**.
2. The **Registry Editor** window displays. On the left pane, expand **HKEY\_LOCAL\_MACHINE** → **SYSTEM** → **SETUP**.
3. On the main window, right-click to add a new string value using **New** → **String Value**. Right-click on the string value file and select **Modify**. When the **Edit String** dialog box displays, fill in the provided text boxes:

✎ Value name: **UnattendFile**

✎ Value data: **a:\sysprep.inf**

4. Launch sysprep from **C:\Windows\System32\sysprep\sysprep.exe**

✎ Under **System Cleanup Action**, select **Enter System Out-of-Box-Experience (OOBE)**.

✎ Tick the **Generalize** check box if you need to change the computer's system identification number (SID).

✎ Under **Shutdown Options**, select **Shutdown**.

Click **OK**. The virtual machine will now go through the sealing process and shut down automatically.

### To create a template from an existing Windows machine

1. In the administration portal, click the **Virtual Machines** tab. Select the sealed Windows 7 virtual machine. Ensure that it has a status of **Down** and click **Make Template**.
2. The **New Virtual Machine Template** displays. Enter information into the following fields:

✎ **Name:** Name of the new template.

Name your new template **FinanWinDesk**.

✎ **Description:** Description of the new template.

- » **Host Cluster:** The Host Cluster for the virtual machines using this template.
  - » **Allow all users to access this template:** Check this box to allow all users to access this template.
3. Click **OK**. In the **Templates** tab, the template displays the "Image Locked" status icon while it is being created. During this time, the action buttons for the template remain disabled. Once created, the action buttons are enabled and the template is ready for use.
  4. Now that you have the template, you can further modify its properties. In this example you will set its time zone and enable it to join the Active Directory. These parameters will be passed to the sysprep answers file while creating a virtual machine using this template.

On the **Tree** pane navigate to the data center where you have created the template. Double-click the template and click the **sysprep** tab. Fill in the Active Directory domain that you have added previously, select your time zone, and click **OK**.

5. You have to activate the Windows operating system for each virtual machine created using this template. To automate the activation process, configure sysprep with the volume license key you have prepared for this track.

Log in to the Red Hat Enterprise Virtualization Manager's server console and run:

```
# engine-config -s ProductKeyWindow7x64=<enter your key here>
# service ovirt-engine restart
```

You can now create new Windows machines using this template.

[Report a bug](#)

## 8.5. Assigning UserRole permissions on virtual machines

To access virtual machines from the Red Hat Enterprise Virtualization User Portal, users must have the relevant permissions. Permissions are added from the Manager administration portal or the Power User Portal. A basic user has permissions for the assigned virtual machine only, not for all virtual machines in the system.

### To assign UserRole permissions on a virtual machine

1. From the administration portal, click the **Virtual Machines** tab. Select the virtual machine you created, **Win7Huron**.
2. On the **Details** pane, navigate to the **Permissions** tab. Click the **Add** button.
3. The **Add Permission to User** dialog displays. Enter **desktopuser** in the **Search** text box, and click **Go**. Select the check box of the user to be assigned the permissions. Scroll through the **Assign role to user** list and select **UserRole**.
4. Click **OK**. The user now has permissions to access the virtual desktop.

[Report a bug](#)


## 8.6. Connecting to virtual machines

Now that you have created a user who has basic permissions, log in as **desktopuser** to the User

Portal. Ensure that you have the **addomain.demo.redhat.com** domain selected. For instructions refer to the section regarding logging in to the User Portal.

The Basic User Portal displays for users who have basic permissions. Note the difference between this portal and the one you accessed in a previous lab as a Power User. The Basic User Portal is designed for end users with little need for performing system administration changes to the virtual machines.

### To connect to a virtual machine

1. Select the virtual machine you wish to connect to and click the Play  button.
2. The virtual machine powers up. When it is turned on, it will no longer be grayed out. The text **Machine is Ready** displays under the virtual machine logo. Double click the virtual machine to connect to console.



**Figure 8.5. Add UserRole permission**

3. A SPICE console window of the virtual machine displays.

A virtual desktop can be used the same way as a physical computer. Red Hat Enterprise Virtualization supports the use of USB devices and CD-ROMs on virtual machines. For more information on using virtual machines refer to the *Red Hat Enterprise Virtualization User Portal Guide*.

4. Once you have finished using the virtual machine, log out according to the instructions specific to the operating system. This returns you to the user portal.

[Report a bug](#)

## 8.7. Use Virtual Machine Pools

### 8.7.1. Creating Virtual Machine Pools

#### Procedure 8.1. To Create a Virtual Machine Pool

1. Log in to the administration portal. 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.

**New Pool**

**General**

Cluster: test

Based on Template: FinanWinDesk

Operating System: Red Hat Enterprise Linux 6.x x64

Optimized for: Desktop

Name: FinWinPool

Description: Finance Windows 7 Shared

Comment:

Number of VMs: 1

Prestared: 0

Maximum number of VMs per user: 1

☐ Delete Protection

Hide Advanced Options

OK Cancel

**Figure 8.6. Create a new virtual machine pool**

Enter *FinWinPool* in **Name** and 5 in **Number of VMs**. In the **Based On Template** field, select **FinanWinDesk**. You may alter other settings, but this example retains the defaults. Click **OK** to create the virtual machine pool.

3. In the **Virtual Machines** tab, you can see that five new virtual machines have been created. You must boot each virtual machine to allow **sysprep** to set up the virtual machines and add them to the domain as specified when you created the template. This process will take a few minutes; you can either open consoles to all the virtual machines to monitor the progress of setup or wait until you see the virtual machine information (like IP addresses) display on the results list, which indicates that setup is complete.
4. Power off all virtual machines in the pool. To do so, select all virtual machines in the pool and click the **Stop** button. The virtual machines will move to a *Powering Down* state, which later changes to *Down*.

[Report a bug](#)

### 8.7.2. Assigning UserRole Permissions

To access virtual machines from the Red Hat Enterprise Virtualization User Portal, users must have the relevant permissions. Assigning UserRole permissions for a pool allows a user to access all of



the virtual machines in the pool, but not other virtual machines in the system.

### Procedure 8.2. To Assign UserRole Permissions

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. On the **Details** pane, navigate to the **Permissions** tab. Click the **Add** button.
3. The **Add Permission to User** window opens.

Select **addomain.demo.redhat.com** as the domain to search and enter **desktopuser** in the **Search** text box, then click **Go**. Select the check box of the **desktopuser**. Scroll through the **Assign role to user** list and select **UserRole**.

4. Click **OK**. The user now has permissions to access the virtual machine pool.


[Report a bug](#)

### 8.7.3. Allocating Virtual Machines

The virtual machines in a pool have the same operating system and installed applications. The pool will allocate identical virtual machines on demand to each user that can access the pool.

Log in to the user portal as **desktopuser**. Ensure you have the **addomain.demo.redhat.com** domain selected.

### Procedure 8.3. To Allocate a Virtual Machine

1. Select the pool you created. Turn it on by clicking on the play  button, which in this case allocates a virtual machine from the pool.




**Figure 8.7. Take a virtual machine from pool**

2. The virtual machine assigned to you powers up. Its name reflects the pool it was taken from.



**Figure 8.8. Take a virtual machine from pool**

3. When the virtual machine displays the text **Machine is Ready**, double-click the virtual machine logo to access its console window.



**Figure 8.9. Connect to virtual machine**



4. A virtual machine in a pool is stateless; any changes made to the virtual machine are discarded after it has been shut down. In the next section, you will test whether this virtual machine is stateless by seeing changes that you make to it get removed when the virtual machine is returned to the pool. First, make a change to the default virtual machine. For example, create a file on the desktop called **test.txt** or add a new user account.

[Report a bug](#)

### 8.7.4. Deallocating Virtual Machines

Now, you will shut down the virtual machine and return it to the pool. Because the virtual machine you used is part of a shared resource, the changes you made will not be persistent.

#### Procedure 8.4. To Deallocate a Virtual Machine

1. Make a note of the name of the virtual machine, and then shut it down by clicking the  button. It is now deallocated and returned to the pool.
2. Now, select the virtual machine pool and click  to allocate a virtual machine again.
3. You might get a different virtual machine from the pool, or the same virtual machine; check the name of the virtual machine that has been allocated.

If it is the same virtual machine, you can check that the desktop has been rebuilt and thus retains only the settings specified in the template from which it was created. To do so, log in to the virtual desktop and verify that the changes you made previously have been discarded.

[Report a bug](#)

## 8.8. Lab 8 - Summary

In this lab, you have successfully created and used Windows virtual machines and virtual machine pools.

The next lab in the Advanced Track is Advanced Storage Features. This lab shows you how to create floating disks, associate them with virtual machines, mark virtual disks shared, and associate shared virtual disks with other machines. You will also take snapshots of running virtual machines and use them to create clones of the virtual machine.



[Report a bug](#)

## Chapter 9. Lab 9 - Installing and configuring minimal setup

### 9.1. Lab 9 - Objectives

#### Lab 9 - Objectives

This lab takes you through the tasks necessary to install and set up Red Hat Enterprise Virtualization on a single host and using local storage. In addition, you will learn how to add ISOs in order to create virtual machines. This lab should take you about 60 minutes.

This chapter shows you how to install the Red Hat Enterprise Virtualization Manager on a server running Red Hat Enterprise Linux. (12 minutes\*)

This chapter shows you how to install and configure a Red Hat Enterprise Virtualization Hypervisor for use with Red Hat Enterprise Virtualization Manager. (13 minutes\*)

This chapter shows you how to configure a client machine to connect to the Red Hat Enterprise Virtualization Manager administration portal. (8 minutes)

This chapter shows you how to approve the host for use from the Red Hat Enterprise Virtualization Manager. (7 minutes)

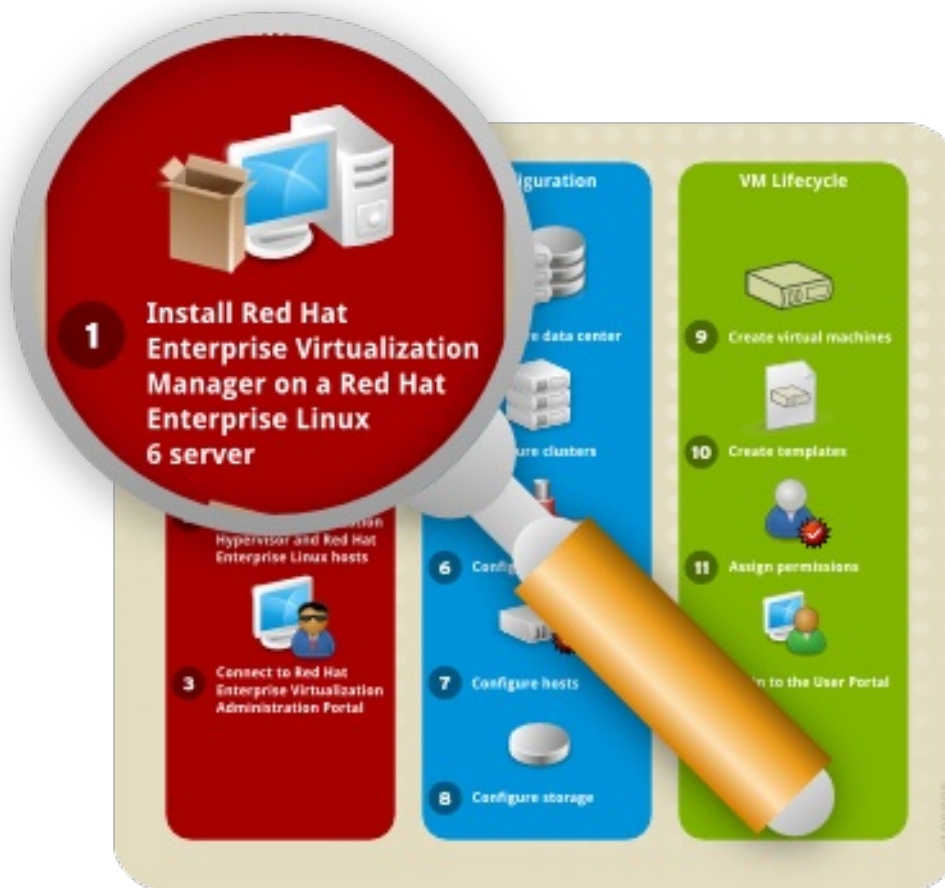
This chapter shows you how to define NFS, iSCSI or FCP storage and attach the domains to the data center. (10 minutes)

This chapter shows you how to attach the predefined ISO domain to the data center and upload ISO images to the repository. (10 minutes)

\* The time required to download packages from the Red Hat Network depends on the bandwidth of your connection to RHN, therefore it has not been included in the estimated time.

[Report a bug](#)

### 9.2. Installing Red Hat Enterprise Virtualization Manager



**Figure 9.1. Install Red Hat Enterprise Virtualization Manager**

The Red Hat Enterprise Virtualization Manager is the control center of the Red Hat Enterprise Virtualization environment. It allows you to define hosts, configure data centers, add storage, define networks, create virtual machines, manage user permissions, and use templates from one central location.

The Red Hat Enterprise Virtualization Manager must be installed on a server running Red Hat Enterprise Linux 6, with minimum 4 GB RAM, 25 GB free disk space and 1 Gbps network interface.

#### Procedure 9.1. To install Red Hat Enterprise Virtualization Manager

1. Install Red Hat Enterprise Linux 6 on a server. When prompted for the software packages to install, select the default **Basic Server** option. See the *Red Hat Enterprise Linux Installation Guide* for more details.



#### Important

During installation, remember to set the fully qualified domain name (FQDN) and IP for the server.

2. If your server has not been registered with the Red Hat Network, run:

```
# rhn_register
```

To complete registration successfully you need to supply your Red Hat Network user name and password. Follow the onscreen prompts to complete registration of the system.

After you have registered your server, update all the packages on it. Run:

```
# yum -y update
```

Reboot your server for the updates to be applied.

3. Subscribe the server to the required Red Hat Network channels. See the *Red Hat Enterprise Virtualization Manager Release Notes* for a list of required channels.
4. You are now ready to install the Red Hat Enterprise Virtualization Manager. Run the following command:

```
# yum -y install rhevm
```

This command will download the Red Hat Enterprise Virtualization Manager installation software and resolve all dependencies.

5. When the packages have finished downloading, run the installer:

```
# engine-setup
```

6. The installer will take you through a series of interactive questions as listed in the following example. If you do not enter a value when prompted, the installer uses the default settings which are stated in [ ] brackets.

#### Example 9.1. Red Hat Enterprise Virtualization Manager installation

```
[ INFO ] Stage: Initializing
[ INFO ] Stage: Environment setup
          Configuration files: ['/etc/ovirt-engine-
setup.conf.d/10-packaging.conf']
          Log file: /var/log/ovirt-engine/setup/ovirt-engine-
setup-20131219122349.log
          Version: otopi-1.1.2 (otopi-1.1.2-1.el6ev)
[ INFO ] Stage: Environment packages setup
[ INFO ] Stage: Programs detection
[ INFO ] Stage: Environment setup
[ INFO ] Stage: Environment customization

--== PACKAGES ==--

[ INFO ] Checking for product updates...
[ INFO ] No product updates found

--== NETWORK CONFIGURATION ==--
          Host fully qualified DNS name of this server
[localhost.localdomain]:
          Setup can automatically configure the firewall on this
system.
          Note: automatic configuration of the firewall may
overwrite current settings.
          Do you want Setup to configure the firewall? (Yes, No)
[Yes]:
```

```

--== DATABASE CONFIGURATION ==--
Where is the database located? (Local, Remote) [Local]:
Setup can configure the local postgresql server
automatically for the engine to run. This may conflict with
existing applications.
Would you like Setup to automatically configure
postgresql, or prefer to perform that manually? (Automatic,
Manual) [Automatic]:

--== OVIRT ENGINE CONFIGURATION ==--
Engine admin password:
Confirm engine admin password:
Application mode (Both, Virt, Gluster) [Both]:
Default storage type: (NFS, FC, ISCSI, POSIXFS) [NFS]:

--== PKI CONFIGURATION ==--
Organization name for certificate [localdomain]:

--== APACHE CONFIGURATION ==--
Setup can configure the default page of the web server
to present the application home page. This may conflict with
existing applications.
Do you wish to set the application as the default page
of the web server? (Yes, No) [Yes]:
Setup can configure apache to use SSL using a
certificate issued from the internal CA.
Do you wish Setup to configure that, or prefer to
perform that manually? (Automatic, Manual) [Automatic]:

--== SYSTEM CONFIGURATION ==--
Configure an NFS share on this server to be used as an
ISO Domain? (Yes, No) [Yes]:
Local ISO domain path [/var/lib/exports/iso-
20131219172449]:
Local ISO domain name [ISO_DOMAIN]:
Configure WebSocket Proxy on this machine? (Yes, No)
[Yes]:

--== END OF CONFIGURATION ==--
Would you like transactions from the Red Hat Access
Plugin sent from the RHEV Manager to be brokered through a proxy
server? (Yes, No) [No]:

```

Important points to note:

- The default ports **80** and **443** must be available to access the Manager on HTTP and HTTPS respectively.
- If you elect to configure an NFS share it will be exported from the machine on which the Manager is being installed.
- The storage type that you select will be used to create a data center and cluster. You will then be able to attach storage to these from the Administration Portal.

7. You are then presented with a summary of the configurations you have selected.

**Example 9.2. Confirm Manager installation settings**

```

--== CONFIGURATION PREVIEW ==--

Database name                : engine
Database secured connection  : False
Database host                 : localhost
Database user name           : engine
Database host name validation : False
Database port                 : 5432
NFS setup                     : True
PKI organization              : localdomain
Application mode              : both
Firewall manager              : iptables
Update Firewall               : True
Configure WebSocket Proxy     : True
Host FQDN                     :
localhost.localdomain
NFS mount point               :
/var/lib/exports/iso-20131219172449
Datacenter storage type       : nfs
Configure local database      : True
Set application as default page : True
Configure Apache SSL          : True

Please confirm installation settings (OK, Cancel) [OK]:

```

8. The installation commences. The following message displays, indicating that the installation was successful.

**Example 9.3. Successful installation**

```

[ INFO ] Stage: Transaction setup
[ INFO ] Stopping engine service
[ INFO ] Stopping websocket-proxy service
[ INFO ] Stage: Misc configuration
[ INFO ] Stage: Package installation
[ INFO ] Stage: Misc configuration
[ INFO ] Creating PostgreSQL database
[ INFO ] Configuring PostgreSQL
[ INFO ] Creating database schema
[ INFO ] Creating CA
[ INFO ] Configuring WebSocket Proxy
[ INFO ] Generating post install configuration file '/etc/ovirt-
engine-setup.conf.d/20-setup-ovirt-post.conf'
[ INFO ] Stage: Transaction commit
[ INFO ] Stage: Closing up

--== SUMMARY ==--
An ISO NFS share has been created on this host.
    If IP based access restrictions are required, edit:
    entry /var/lib/exports/iso-20131219172449 in
/etc/exports
SSH fingerprint: fingerprint

```



```

Internal CA A9: certificate
Web access is enabled at:
    http://localhost.localdomain:80/ovirt-engine
    https://localhost.localdomain:443/ovirt-engine
Please use the user "admin" and password specified in
order to login into oVirt Engine

```

```
--== END OF SUMMARY ==--
```

```

[ INFO ] Starting engine service
[ INFO ] Restarting httpd
[ INFO ] Restarting nfs services
[ INFO ] Generating answer file '/var/lib/ovirt-
engine/setup/answers/20131219122536-setup.conf'
[ INFO ] Stage: Clean up
        Log file is located at /var/log/ovirt-
engine/setup/ovirt-engine-setup-20131219122349.log
[ INFO ] Stage: Pre-termination
[ INFO ] Stage: Termination
[ INFO ] Execution of setup completed successfully

```

Your Red Hat Enterprise Virtualization Manager is now up and running. You can log in to the Red Hat Enterprise Virtualization Manager's web administration portal with the user name **admin** (the administrative user configured during installation) in the **internal** domain. Instructions to do so are provided at the end of this chapter.



### Important

The **internal** domain is automatically created upon installation, however no new users can be added to this domain. To authenticate new users, you need an external directory service. Red Hat Enterprise Virtualization supports Identity Management (IdM), Active Directory, RHDS, and OpenLDAP, and provides a utility called **engine-manage-domains** for attaching new directories to the system. Use of this tool is covered in the *Red Hat Enterprise Virtualization Installation Guide*.



### Note

The Red Hat Enterprise Virtualization engine (that is, the Manager) can be hosted in a virtual machine controlled by the engine that hosts it. This setup is called a "self-hosted engine". In a self-hosted engine system, the virtual machine on which the Manager resides is defined as highly available. Two packages, *ovirt-hosted-engine-setup* and *ovirt-hosted-engine-ha* provide the setup and services necessary to deploy a self-hosted-engine Red Hat Enterprise Virtualization environment. For more on Red Hat Enterprise Virtualization Self-Hosted Engine, see the *Red Hat Enterprise Virtualization Installation Guide*.

[Report a bug](#)

## 9.3. Installing Red Hat Enterprise Virtualization Hypervisor

### 9.3.1. Registering the Host on RHN and Acquiring ISO Hypervisor Images

#### Summary

The **Red Hat Enterprise Virtualization Manager** Red Hat Network channel contains the Hypervisor packages. The Hypervisor itself is contained in the **rhev-hypervisor** package. Additional tools supporting USB and PXE installations are installed as dependencies. Install the Hypervisor packages on the system you plan to use to create Hypervisor boot media.

Select one of the two options below:

#### Procedure 9.2. Subscribing to RHN Entitlement Pools and Installing the Red Hat Enterprise Virtualization Hypervisor Packages

##### 1. Subscribing to download the Hypervisor using certificate-based RHN

###### a. Identify Available Entitlement Pools

To subscribe the system to the Red Hat Enterprise Virtualization channels you need you must locate the identifier for the relevant entitlement pool. Use the **list** action in the **subscription-manager** to find these:

```
# subscription-manager list --available | grep -A8 "Red Hat Enterprise Virtualization"
```

###### b. Subscribe System to Entitlement Pools

Using the pool identifiers located in the previous step, subscribe the system to **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** entitlements. Use the **subscribe** parameter of the **subscription-manager** command, and replace **POOLID** with one of the pool identifiers.

```
# subscription-manager subscribe --pool=POOLID
```

##### 2. Subscribing to download the Hypervisor using RHN Classic

- 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 page, and then click **Registered Systems** in the menu that appears.
  - c. Select the system to which you are adding channels from the list 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 **Red Hat Enterprise Virtualization Manager** channel from the list on the screen, then click the **Change Subscription** button to finalize the change.
3. Log in to the system on which the Red Hat Enterprise Virtualization Manager is installed. Log in as **root**.
  4. Use **yum** to install the **rhev-hypervisor**.

```
# yum install rhev-hypervisor
```

## Result

The Hypervisor ISO image is installed into the `/usr/share/rhev-hypervisor/` directory. The `lived -iso -to -disk` and `lived -iso -to -pxeboot` scripts are installed to the `/usr/bin/` directory.



### Note

All version of Red Hat Enterprise Linux 6.2 and higher allow more than one version of the ISO image to be installed at one time. Because of this, `/usr/share/rhev-hypervisor/rhev-hypervisor.iso` is now a symbolic link to a uniquely-named version of the Hypervisor ISO image, for instance `/usr/share/rhev-hypervisor/rhev-6.2-20111006.0.el6.iso`. Different versions of the image can now be installed alongside each other, allowing administrators to run and maintain a cluster on a previous version of the Hypervisor while upgrading another cluster for testing.

The symbolic links `/usr/share/rhev-hypervisor/rhev-latest6.iso` and `/usr/share/rhev-hypervisor/rhev-hypervisor6.iso` are created. These links target the most-recently installed version of the Red Hat Enterprise Virtualization ISO image.

[Report a bug](#)

## 9.3.2. Preparing Optical Hypervisor Installation Media

### Summary

Burn the Hypervisor image to a CD with the `wodim` command. The `wodim` command is part of the `wodim` package.

### Procedure 9.3. Preparing Optical Hypervisor Installation Media

1. Verify that the `wodim` package is installed on the system.

#### Example 9.4. Verify Installation of `wodim` Package

```
# rpm -q wodim
wodim-1.1.9-11.el6.x86_64
```

If the package version is in the output the package is available.

If nothing is listed, install `wodim`:

```
# yum install wodim
```

2. Insert a blank CD or DVD into your CD or DVD writer.
3. Record the ISO file to the disc. The `wodim` command uses the following:

```
wodim dev=device image.iso
```

This example uses the first CD-RW (`/dev/cdrw`) device available and the default hypervisor image location, `/usr/share/rhev-hypervisor/rhev-hypervisor.iso`.

#### Example 9.5. Use of `wodim` Command

```
# wodim dev=/dev/cdrw /usr/share/rhev-hypervisor/rhev-hypervisor.iso
```

### Result

If no errors occurred, the Hypervisor is ready to boot. Errors sometimes occur during the recording process due to errors on the media itself. If this occurs insert another writable disk and repeat the command above.

The Hypervisor uses a program (`isomd5sum`) to verify the integrity of the installation media every time the Hypervisor is booted. If media errors are reported in the boot sequence you have a bad CD. Follow the procedure above to create a new CD or DVD.

[Report a bug](#)

### 9.3.3. Install Red Hat Enterprise Virtualization Hosts from Optical Installation Media

Now that you have registered the host with RHN, acquired the Hypervisor images, and used the Hypervisor images to create optical installation media, you will boot the system using the optical installation media.

#### Summary

Booting the Hypervisor from optical installation media requires the system to have a correctly defined BIOS boot configuration.

1. Ensure that the system's BIOS is configured to boot from the CD or DVD drive before proceeding.



#### Note

See your manufacturer's manuals for further information on modifying the system's BIOS boot configuration.

2. Insert the Hypervisor CD in the CD or DVD drive.
3. Reboot the system.

### Result

The host's screen will display the Hypervisor boot screen.

[Report a bug](#)

### 9.3.4. Install Red Hat Enterprise Virtualization Hypervisors

## Red Hat Enterprise Virtualization Hypervisor Menu navigation keys

- ✦ Use the Up and Down arrow keys to navigate between selections. Your selections are highlighted in white.
- ✦ The Tab key allows you to move between fields.
- ✦ Use the Spacebar to tick check boxes, which are represented by [ ] brackets. A marked check box displays with an asterisk (\*).
- ✦ To proceed with the selected configurations, press the Enter key.



### Important

Red Hat Enterprise Virtualization Hypervisor is a closed appliance, and does not allow the installation of custom RPMs.

If you require custom RPMs, you must use Red Hat Enterprise Linux hosts. Red Hat Enterprise Linux hosts are not closed appliances, and they allow the installation of custom RPMs.

## To install Red Hat Enterprise Virtualization Hypervisors

1. Insert the Red Hat Enterprise Virtualization Hypervisor installation CD into your CD-ROM drive of the machine designated as a host. Reboot the machine. When the boot splash screen displays, press the **Tab** key and select **Boot** to boot from the hypervisor installation media. Press **Enter**.
2. On the installation confirmation screen, select **Install RHEV Hypervisor** and press **Enter**.
3. The installer automatically detects the drives attached to the system. The disk selected for booting the hypervisor is highlighted in white. Ensure that the local disk is highlighted, otherwise use the arrow keys to select the correct disk. Select **Continue** and press **Enter**.
4. You are prompted to select the drive on which the hypervisor is to be installed. Ensure that the local disk is highlighted, otherwise use the arrow keys to select the correct disk. While multiple installation drives can be used, select only one for this evaluation. Select **Continue** and press **Enter**.
5. Enter a password for local console access and confirm it. Select **Install** and press **Enter**. The Red Hat Enterprise Virtualization Hypervisor partitions the local drive, then commences installation.
6. Once installation is complete, a dialog prompts you to **Reboot** the hypervisor. Press **Enter** to confirm. Remove the installation disc.
7. After the hypervisor has rebooted, you will be taken to a login shell. Log in as the **admin** user with the password you provided during installation to enter the Red Hat Enterprise Virtualization Hypervisor management console.
8. On the hypervisor management console, there are eight tabs on the left. Press the **Up** and **Down** keys to navigate between them and **Enter** to access them.

- a. Select the **Network** tab. Fill in the required fields as shown in the following example. Substitute the **DNS Server** address according to your environment.

```

RHEV Hypervisor 6.3-20120920.0.rhev31.auto92
localhost.localdomain

Status
Network
Security
Keyboard
SNMP
Logging
Kernel Dump
Remote Storage
CIM
RHEV-M
Red Hat Network

System Identification
Hostname: atlantic.demo.redhat.com
DNS Server 1: 192.168.0.2
DNS Server 2:
NTP Server 1: 0.rhel.pool.ntp.org
NTP Server 2: 1.rhel.pool.ntp.org

Device Status Model MAC Address
eth0 Unconfigured Red Hat I 52:54:00:7b:0a:d6

<Flash Lights to Identify> <Apply> <Reset>

```

**Figure 9.2. Configure Hypervisor network settings**

After you have filled in the fields, select **Apply** and press **Enter**. This saves your network settings.

- b. For this document, the **eth0** device will be used to set up the management network. Select it and press **Enter** to access the interface configuration menu. Fill in the required fields as shown in the following example.

```

RHEV Hypervisor 6.3-20120920.0.rhev31.auto92
atlantic.demo.redhat.com

Interface: eth0 Driver: virtio_net
Protocol: Disabled Vendor: Red Hat Inc Virtio network
Link Status: Active MAC Address: 52:54:00:7b:0a:d6

IPv4 Settings
[ ] Disabled [ ] DHCP [*] Static
IP Address: 192.168.0.5 Netmask: 255.255.255.0
Gateway: 192.168.0.254

VLAN ID:
<Apply> <Back> <Reset>

```

**Figure 9.3. Configure management network interface**

Under **IPv4 Settings**, select **DHCP** or **Static** IP addressing and press **Spacebar** to mark the option as enabled. If using static IP addressing you must also provide the **IP Address**, **Netmask**, and **Gateway**. Select **Apply** and press **Enter**.

A dialog prompts you to confirm your network settings, select **OK** and press **Enter**.

c. Select the **oVirt Engine** tab. Configure the following options:

- ✳ In the **Management Server** field, enter **rhev.m.demo.redhat.com**.
- ✳ In the **Management Server Port** field, enter **443**.
- ✳ Tick the **Connect to the RHEV Manager and Validate Certificate** check box.
- ✳ The **Set RHEV-M Admin Password** field allows you to specify the root password for the hypervisor, and enable SSH password authentication from the Red Hat Enterprise Virtualization Manager. You do not have to fill in this field for this document.

Select **Apply** and press **Enter**. A dialog displays, asking you to connect the hypervisor to the Red Hat Enterprise Virtualization Manager and validate its certificate. Select **Approve** and press **Enter**. A message will display notifying you that the Manager configuration has been successfully updated.

- d. Under the **Red Hat Network** tab, you can register the host with the Red Hat Network. This enables the host to run Red Hat Enterprise Linux virtual machines with proper RHN entitlements. However, for the purposes of this document, the evaluation subscriptions will be used for the guests.
- e. Accept all other default settings.
- f. Finally, select the **Status** tab. Select **Restart** and press **Enter** to reboot the host and apply all changes.

You have now successfully installed a Red Hat Enterprise Virtualization Hypervisor. Repeat the above steps for each hypervisor you wish to use. The following sections will provide instructions on how to approve the hypervisors for use with the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

## 9.4. Connecting to Red Hat Enterprise Virtualization Web Administration Portal





**Figure 9.4. Connect to the Manager administration portal**

Now that you have installed the Red Hat Enterprise Virtualization Manager and hosts, you can log in to the Manager administration portal to start configuring your virtualization environment. Use a client running Firefox to access the web-based administration portal.

1. Open a browser and navigate to **`https://domain.example.com`**. Substitute **`domain.example.com`** with the URL provided during installation.
2. Under the **Portals** heading, click **Web Admin Portal**
3. If this is your first time connecting to the administration portal, Red Hat Enterprise Virtualization Manager will issue security certificates for your browser. Click the link labeled **this certificate** to trust the **ca.cer** certificate. A pop-up displays, click **Open** to launch the **Certificate** dialog. Click **Install Certificate** and select to place the certificate in **Trusted Root Certification Authorities** store.
4. The portal login screen displays. Enter **admin** as your **User Name**, and enter the **Password** that you provided during installation. Ensure that your domain is set to **Internal**. Click **Login**.

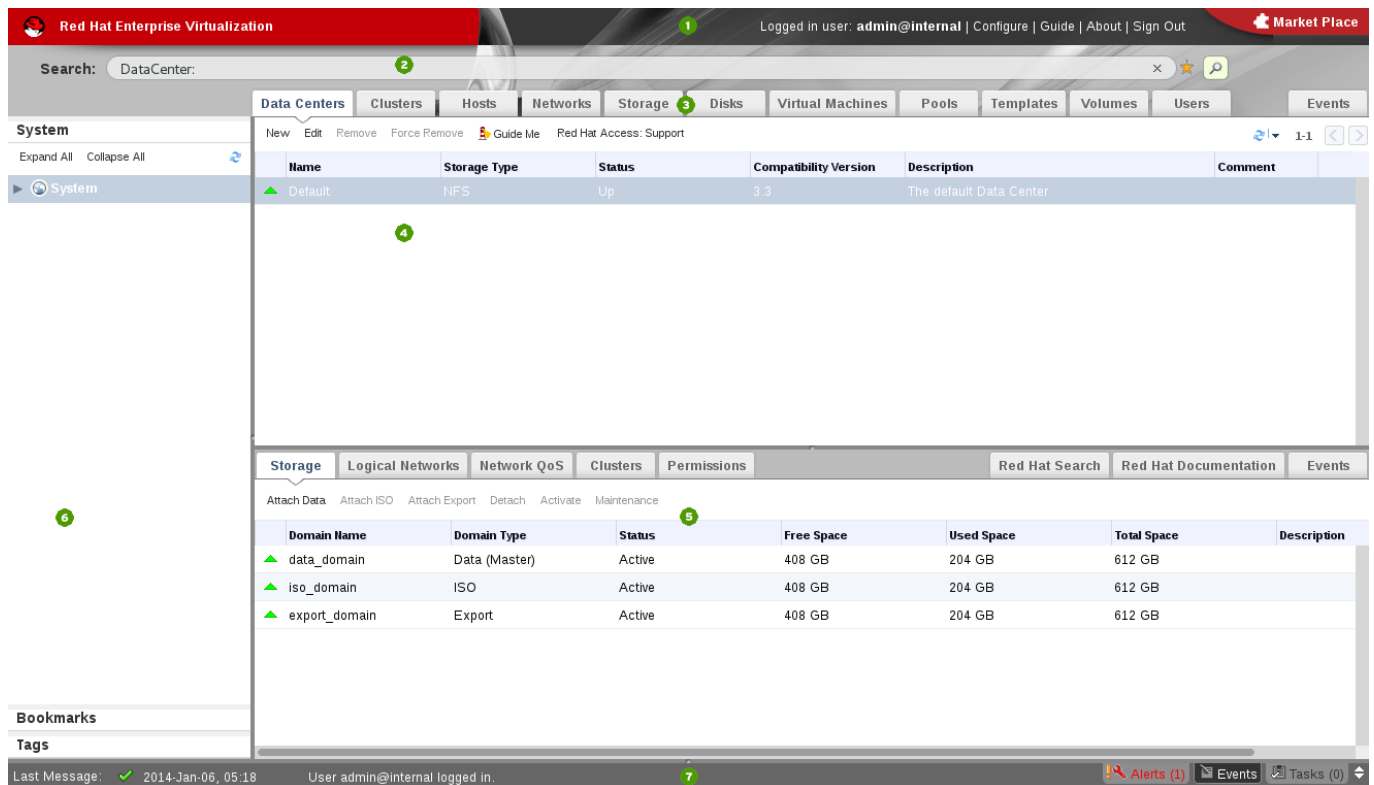
You have now successfully logged in to the Red Hat Enterprise Virtualization web administration portal. Here, you can configure and manage all your virtual resources.



[Report a bug](#)

## 9.5. 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.



**Figure 9.5. User Interface Elements of the Administration Portal**

### 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](#)

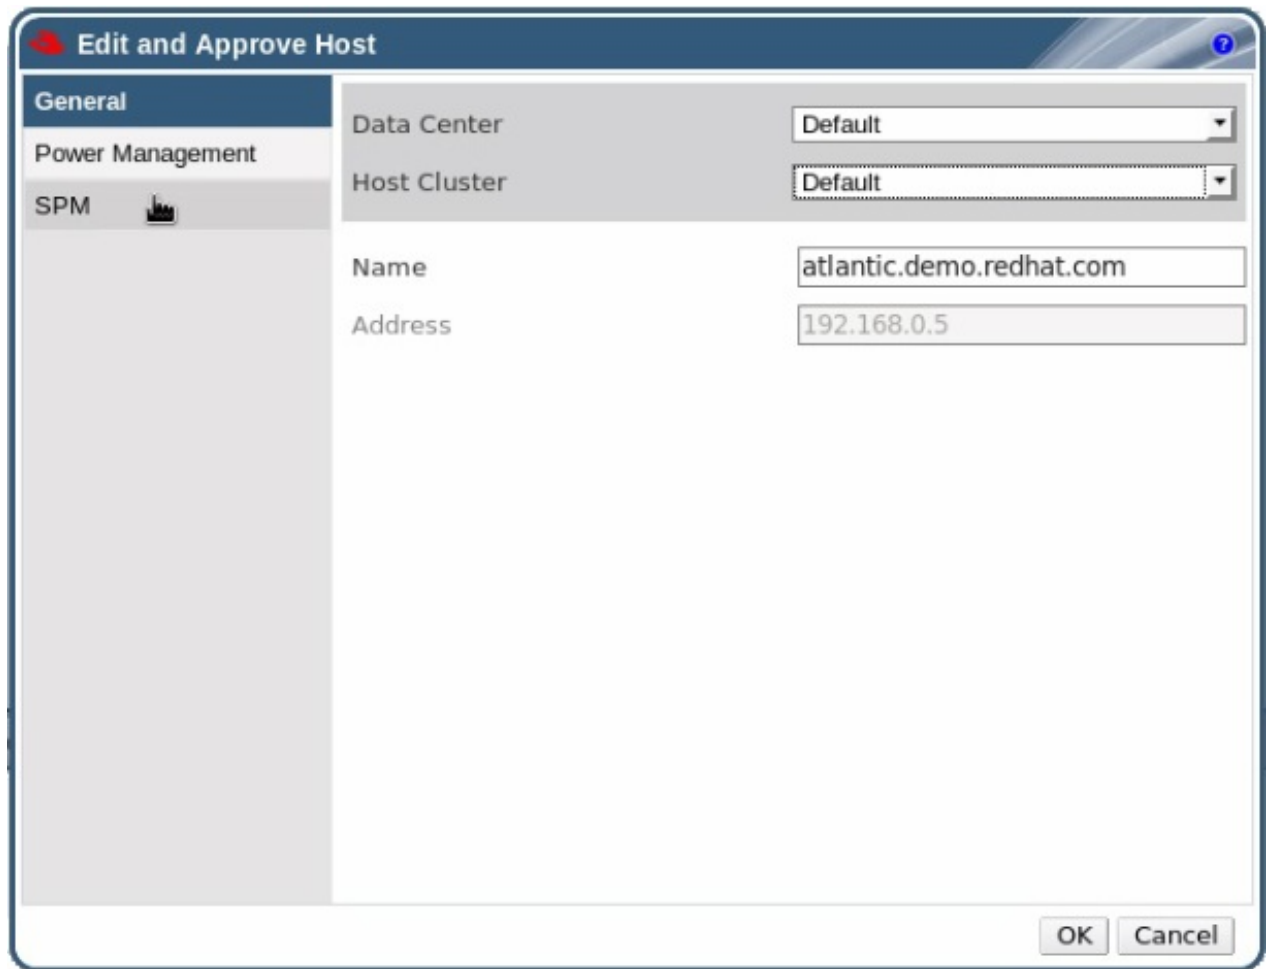
## 9.6. Approve the Red Hat Enterprise Virtualization Hypervisor

At this point you should already have a Default data center and a Default cluster, which have been automatically created during the Manager installation. In addition, the Red Hat Enterprise Virtualization Hypervisors you installed earlier should have been automatically detected by the Red Hat Enterprise Virtualization Manager and attached to the Default cluster of the Default data center.

However, before they can be used, they require a click of approval from the administration portal. Perform the following procedure for each hypervisor.

### To approve the Red Hat Enterprise Virtualization Hypervisor hosts

1. Navigate to the **Tree** pane and click the **Expand All** button. Under the **Default** cluster, click the **Hosts** icon. The **Hosts** tab displays a list of available hypervisors.
2. Select your hypervisor and click the **Approve** button. The **Edit and Approve Host** dialog displays. Accept the defaults or make changes as necessary, then click **OK**.



**Figure 9.6. Approve Red Hat Enterprise Virtualization Hypervisor**

3. A dialog appears, indicating that you have not configured Power Management for this host. For the purpose of this lab, click **OK** to continue. The host goes through a brief installation cycle. When complete, the host status changes from **Non Operational** to **Up**.

Note that both the Red Hat Enterprise Virtualization Hypervisors that you have approved were attached to the same host cluster, which means that they share the same network infrastructure, the same storage and the same type of CPU, therefore they can migrate virtual machines from one to the other. You can learn how to create new host clusters in Advanced Lab 7 - Add Additional Data Center.

Now that you have finished configuring your physical servers for use as the Manager, Hypervisors and administration portal client respectively, you are ready to customize and deploy virtual resources including logical networks, storage domains and virtual machines.

[Report a bug](#)

## 9.7. Creating Local Storage

At this point, you have installed your Red Hat Enterprise Virtualization Hypervisor and approved it for

use with the Red Hat Enterprise Virtualization Manager. Now, you have to set up the host to provide local storage for the virtual machine disk images. Note that setting up local storage is suitable for a small environment, such as this evaluation setup. If you were deploying Red Hat Enterprise Virtualization across a large organization, it is recommended that you use shared storage as described in the section regarding configuring storage. However, for this lab, use the following procedure to add storage to your system.

### To create a local storage domain

1. Navigate to the **Tree** pane and click the **Expand All** button. Under the **Default** cluster, click the **Hosts** icon. The **Hosts** tab displays a list of available hypervisors.
2. Place the host into maintenance mode by clicking the **Maintenance** button. The **Status** field of the host changes to *Preparing for Maintenance*, followed by *Maintenance*. The icon changes to indicate that the host is in maintenance.
3. Click **Configure Local Storage**. The **Configure Local Storage** window opens. On the **General** tab, specify the path to your local storage domain. For this lab, enter **/data/images/rhev**. Retain all the other default settings, and click **OK**.
4. Your local storage domain is created, and attached to the automatically created **Atlantic-Local** data center and **Atlantic-Local** cluster.

You have created a data storage domain. Now, you need to attach an ISO domain to the data center and upload installation images so you can use them to create virtual machines.

[Report a bug](#)

## 9.8. Creating local storage domains

1. Navigate to the **Tree** pane and click the **Expand All** button. Under the **Default** cluster, click the **Hosts** icon. The **Hosts** tab displays a list of available hypervisors.
2. Place the host into maintenance mode by clicking the **Maintenance** button. The **Status** field of the host changes to *Preparing for Maintenance*, followed by *Maintenance*. The icon changes to indicate that the host is in maintenance.
3. Click **Configure Local Storage**. The **Configure Local Storage** window opens. On the **General** tab, specify the path to your local storage domain. For this lab, enter **/data/images/rhev**. Retain all the other default settings, and click **OK**.
4. Your local storage domain is created, and attached to the automatically created **Atlantic-Local** data center and **Atlantic-Local** cluster.

[Report a bug](#)

## 9.9. Attach and Populate ISO Domains

You have defined your first storage domain to store virtual guest data, now it is time to configure your second storage domain, which will be used to store installation images for creating virtual machines. You have already created an ISO domain during the installation of the Red Hat Enterprise Virtualization Manager. To use this ISO domain, attach it to the same data center which contains the local storage domain.

## To attach the ISO domain

1. Navigate to the **Tree** pane and click the **Expand All** button. Click the **Atlantic-Local** data center. On the results list, the **Atlantic-Local** data center displays.
2. On the details pane, select the **Storage** tab and click the **Attach ISO** button.
3. The **Attach ISO Library** dialog appears with the available ISO domain. Select the **local-iso-share** domain and click **OK**.



**Figure 9.7. Attach ISO Library**

4. The ISO domain appears in the results list of the **Storage** tab. It displays with the *Locked* status as the domain is being validated before becoming *Active*.

Media images (CD-ROM or DVD-ROM in the form of ISO images) must be available in the ISO repository for the virtual machines to use. To do so, Red Hat Enterprise Virtualization provides a utility that copies the images and sets the appropriate permissions on the file. For this lab, both the file provided to the utility and the ISO share have to be accessible from the Red Hat Enterprise Virtualization Manager.

Log in to the Red Hat Enterprise Virtualization Manager server console to upload images to the ISO domain.

## To upload ISO images

1. Create or acquire the appropriate ISO images from boot media. Ensure the path to these images is accessible from the Red Hat Enterprise Virtualization Manager server.
2. The next step is to upload these files. First, determine the available ISO domains by running:

```
# engine-iso-uploader list
```

You will be prompted to provide the admin user password which you are using to connect to the administration portal. The tool lists the name of the ISO domain that you have already attached in the previous lab:

```
ISO Storage Domain List:  
local-iso-share
```

Now you have all the information required to upload the required files. To copy your installation images to the ISO domain, run:

```
# engine-iso-uploader upload -i local-iso-share [file1] [file2] ....  
[fileN]
```

3. After the images have been loaded, check that they are available for use in the Manager administration portal.
  - a. Navigate to the **Tree** and click the **Expand All** button. Click **Storage**
  - b. On the **Storage** tab, click **local-iso-share** to display its details pane.
  - c. Select the **Images** subtab. The list of available images should be populated with the files which you have uploaded. In addition, the **RHEV-toolsSetup.iso** and **virtio-win.vfd** images should have been automatically uploaded during installation.

Now that you have successfully prepared the ISO domain for use, you have completed Lab 1 and are ready to start creating virtual machines.

[Report a bug](#)

## 9.10. Lab 9 - Summary

Congratulations, you have achieved the first goal of Track B! You now have an infrastructure ready to create and run virtual machines. To recap, you have successfully installed the Red Hat Enterprise Virtualization Manager and Red Hat Enterprise Virtualization Hypervisor, attached storage domains to the data center, and prepared ISO images.

The next lab on Track B teaches you how to create Red Hat Enterprise Linux virtual machines and templates.

The subsequent labs for this track are also used for **Track A: Standard Setup**. Therefore, most examples, instructions and screen captures are specific to Track A. When you use the instructions in the following labs, substitute the instances of default data center and default cluster with Atlantic-Local data center and Atlantic-Local cluster respectively.



[Report a bug](#)

## Chapter 10. Lab 10 - Advanced storage features

### 10.1. Lab 10 - Objectives

In this lab, you will learn about the advanced storage features supported by Red Hat Enterprise Virtualization.

These include:

- » Floating Disks
- » Shared Disks
- » Creating snapshots of running virtual machines
- » Creating virtual machines from snapshots
- » Associating LUNs with virtual machines
- » Associating Fibre Channel (FC) LUNs with virtual machines

[Report a bug](#)

### 10.2. Advanced Storage Features

Red Hat Enterprise Virtualization has a number of advanced storage features. In this lab, you will explore some of these advanced storage features.

You will create a floating disk, associate it with a virtual machine, mark the virtual disk shared, and associate the shared virtual disk with a second machine.

You will also take a snapshot of a running virtual machine and use it to create a clone of the virtual machine.

[Report a bug](#)

### 10.3. Requirements

This lab assumes that you have successfully completed the basic labs of Section 1, Track A: Standard Setup or Section 2: Track B, Minimal Setup. This lab assumes that you have correctly installed and configured Red Hat Enterprise Virtualization and that you are able to use it to create virtual machines.

[Report a bug](#)

### 10.4. Creating Floating Disks

#### Summary

The first task you will undertake in the Advanced Storage Lab is the creation of a floating disk.

#### Procedure 10.1. Creating Floating Disks

1. Click the **D**isks tab in the navigation pane.

2. Click the **Add** button in the top left of the navigation pane. The **Add Virtual Disk** window opens.
3. Enter **10** in the **Size(GB)** field.
4. Enter **Disk1** in the **Alias** field.
5. Enter **Description placeholder** in the **Description** field.
6. Leave all other settings at their defaults.
7. Click **OK** at the bottom right of the **Add Virtual Disk** window.

## Result

You have created the floating disk **Disk1**. Now that you have created the floating disk **Disk1**, you will associate it with a machine.

## Disks on Multiple Storage Domains

Red Hat Enterprise Virtualization allows you to have disks on multiple storage domains. In the **Add Virtual Disk** window, select the data domain you want your disk to be on. Using this method, you can attach disks from different storage domains to your guest machine.

[Report a bug](#)

## 10.5. Associating a Floating Disk with a Virtual Machine

### Summary

You have just created a floating disk. Now you will associate it with a virtual machine.

### Procedure 10.2. Associating a Floating Disk with a Virtual Machine

1. Click the **Virtual Machines** tab in the navigation pane.
2. Select **RHEL6Thames** from the list. The details pane displays the virtual disks associated with the virtual machine **RHEL6Thames**.
3. Click **Add** in the top left of the details pane. The **Add Virtual Disk** window opens.
4. Select the **Attach Disk** check box. A list of disks appears.
5. Select the check box associated with **Disk1**.
6. Click **OK** at the bottom right of the **Add Virtual Disk** window.

## Result

You have now associated a floating disk with the virtual machine named **RHEL6Thames**. Next you will mark the disk shared, so that it can be associated with a second virtual machine.

[Report a bug](#)

## 10.6. Marking a Virtual Disk Shared

### Summary



You have just associated a floating disk with a virtual machine. Now you will mark the disk shared, so that it can be associated with a second virtual machine.

### Procedure 10.3. Marking a disk shared

1. Click the **Virtual Machines** tab in the navigation pane.
2. Click **RHEL6Thames** in the navigation pane.
3. Click the **Virtual Disks** tab in the details pane.
4. Click **Disk1** in the list in the details pane.
5. Click **Edit** in the upper left of the details pane. The **Edit Virtual Disk** window opens.
6. Check the **is shareable** check box on the right of the **Edit Virtual Disk** window.
7. Click **OK** in the bottom right of the **Edit Virtual Disk** window.

### Result

You have now marked a virtual disk shared. In the next task in this lab, you will associate the disk you just marked shared with a second virtual machine.

[Report a bug](#)

## 10.7. Associating a Shared Virtual Disk with a Second Virtual Machine

### Summary

You have just marked a virtual disk shared. You will now associate the marked disk with a second virtual machine.

### Procedure 10.4. Associating a Shared Virtual Disk with a Second Virtual Machine

1. Click the **Virtual Machines** tab.
2. Select **RHEL6Ganges**. The details pane displays the disks associated with the virtual machine.
3. Click **Add** in the top left of the details pane. The **Add Virtual Disk** window opens.
4. Select the **Attach Disk** check box. A list of disks appears in the **Add Virtual Disk** window.
5. Select the check box associated with **Disk1**.
6. Click **OK** at the bottom right of the **Add Virtual Disk** window.

### Result

You have now associated a shared virtual disk with a second virtual machine.

[Report a bug](#)

## 10.8. Creating a Snapshot of a Running Virtual Machine



## Note

Before you perform the Snapshotting tasks in this lab, you must create a virtual machine named **RHEL6N1le**.

### Summary

In Lab 10, you have created a floating disk, associated it with a virtual machine, flagged it shared, and associated it with a second virtual machine. Now you will create a snapshot of a running virtual machine.

#### Procedure 10.5. Creating a Snapshot of a Running Virtual Machine

1. Click the **Virtual Machines** tab in the navigation pane.
2. Select **RHEL6N1le** in the navigation pane.
3. Ensure that **RHEL6N1le** is running. If it has a status of **Down**, click the green Play button at the top of the navigation pane.
4. Click the **Snapshots** tab in the details pane.
5. Click **Create** in the upper left of the details pane.
6. After a short time, the snapshot you just created will appear in the details pane.

### Result

You have now created a snapshot of a running virtual machine. Next, you will use the snapshot you have just created to create a virtual machine based on this snapshot.

[Report a bug](#)

## 10.9. Creating a Virtual Machine from a Snapshot

### Summary

You have created a snapshot from a virtual machine. Now you can use that snapshot to create another virtual machine.

#### Procedure 10.6. 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 10.1. 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](#)

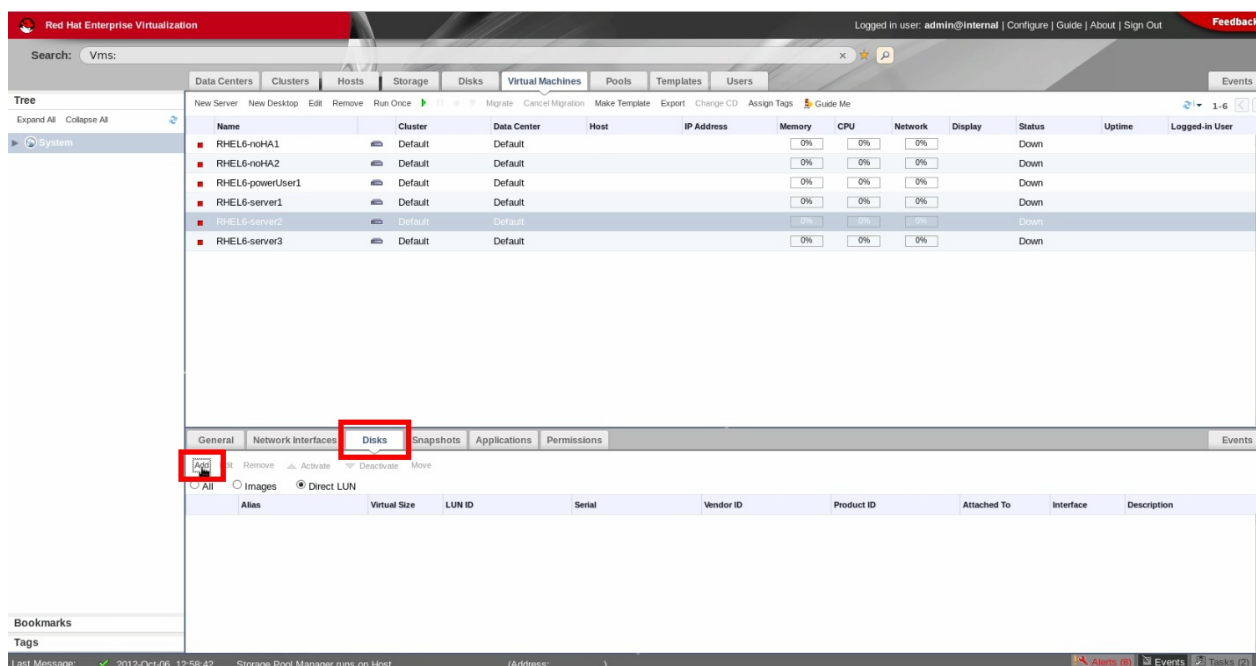
## 10.10. Associating LUNs with Virtual Machines

### Summary

In this procedure, you associate a iSCSI LUN (logical unit number) with a virtual machine.

### Procedure 10.7. Associating iSCSI LUNs with Virtual Machines

1. Select a virtual machine from the navigation pane.
2. Select **Disks** in the details pane.
3. Click **Add**.



**Figure 10.2. Selecting a virtual machine from the navigation pane and opening the Add Virtual Disk window"**

4. Click the **External (Direct LUN)** radio button.
5. Unclick the **Is bootable** check box.



Figure 10.3. De-selecting the "is bootable" check box in the Add Virtual Disk window

6. Give it a description.
7. In the **Targets > LUNs** window, under **Discover Targets** add the address of the iSCSI server.
8. Click **Discover**.

**Add Virtual Disk**

☐ Attach Disk

☐ Internal ☒ External (Direct Lun)

Alias: RHEL6Thames\_Disk2

Description:

Interface: VirtIO

Use Host: rhev-host5

Storage Type: iSCSI

☐ Is bootable

☐ Is shareable

**Discover Targets**

Address: 192.168.1.1

Port: 3260

☐ User Authentication:

CHAP username:

CHAP password:

**Discover**

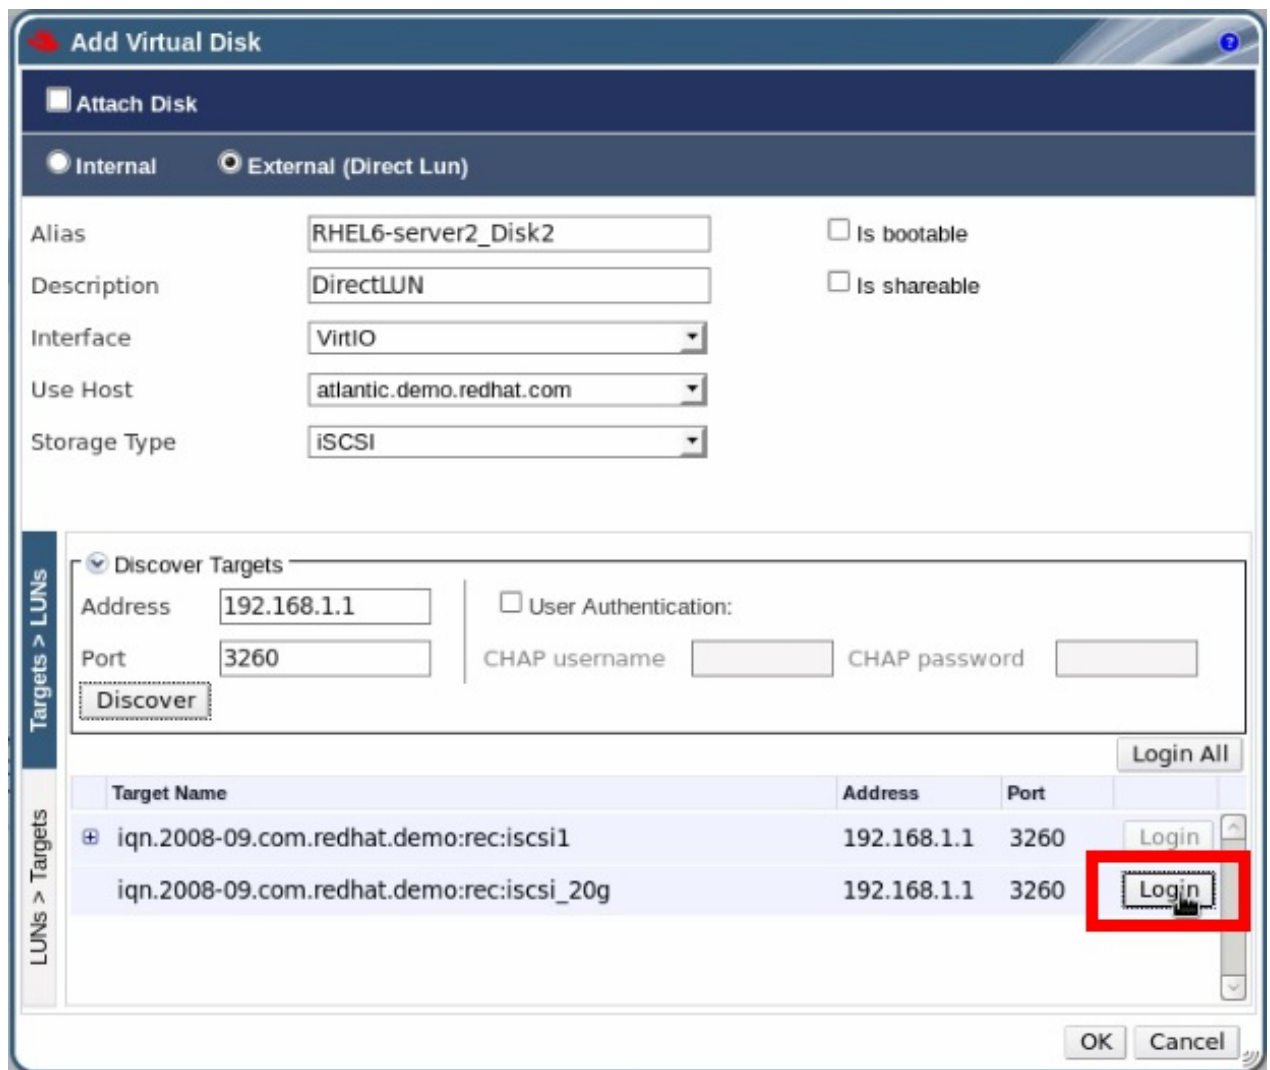
Login All

Target Name	Address	Port
-------------	---------	------

OK Cancel

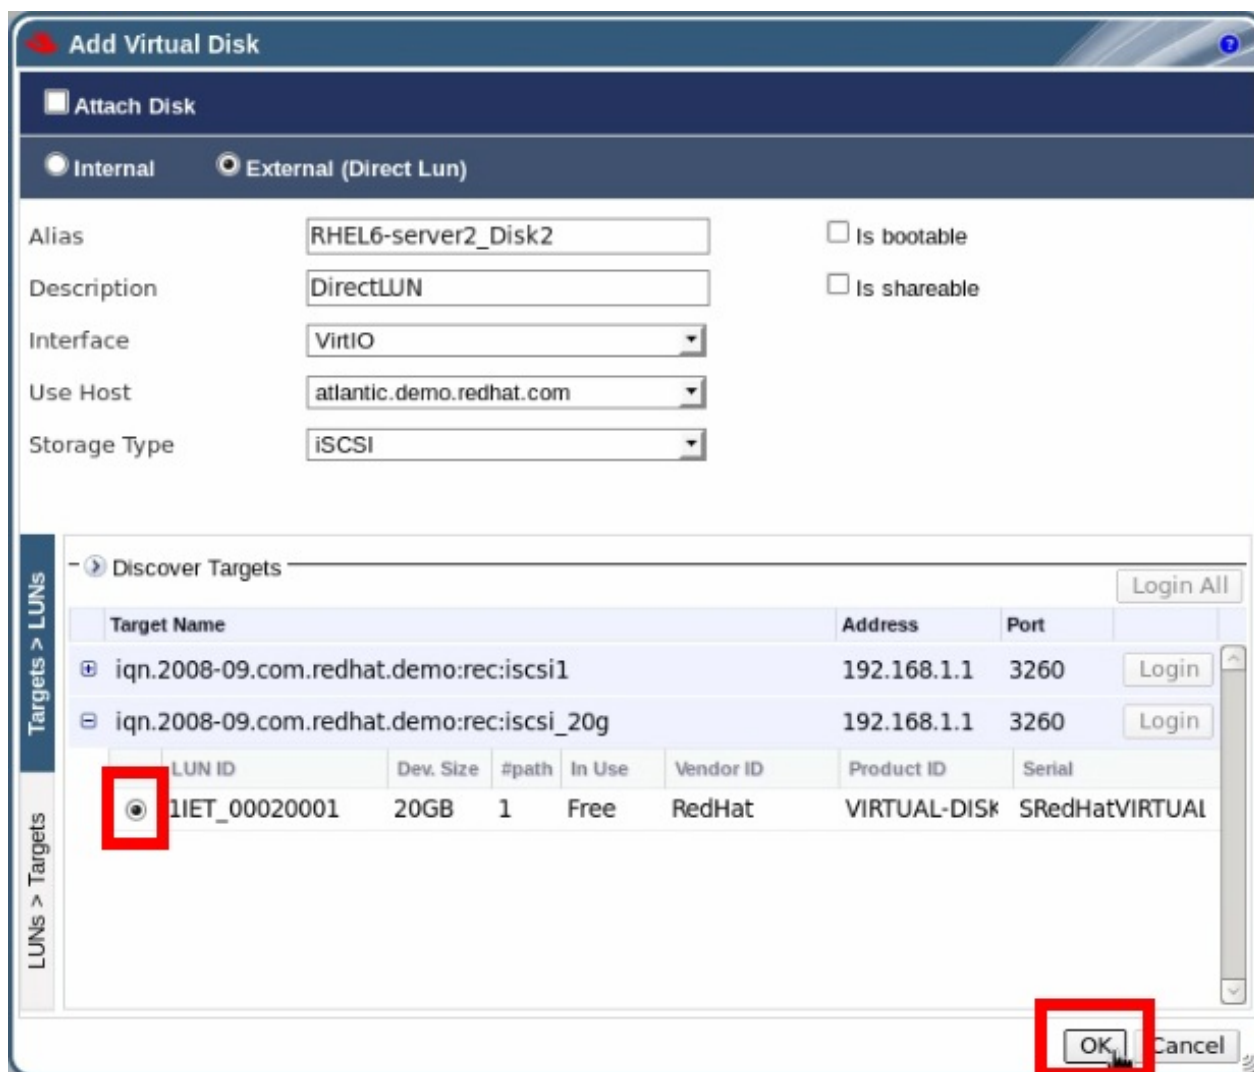
**Figure 10.4. Discovering a LUN**

9. In the **LUNs > Targets** window, click **Login** on the target that you are not logged in to.



**Figure 10.5. Logging in to a Discovered LUN**

10. Click the plus to the left of the target you just logged in to. A list of LUN IDs displays.
11. Click the radio button to the left of the LUN ID of the target.



**Figure 10.6. Selecting a LUN ID and Associating that LUN with a Virtual Machine**

12. Click **OK** in the bottom right hand corner of the **Add Virtual Disk Window**.
13. After a short time, the LUN appears in the **Disks** tab of the details pane.

## Result

You have associated a LUN with a virtual machine.

[Report a bug](#)

## 10.11. Associating Fibre Channel (FC) LUNs with Virtual Machine

Associating Fibre Channel (FC) LUNs with virtual machines is the same as associating iSCSI LUNs with virtual machines, except that it is unnecessary to discover the Fibre Channel LUNs.

[Report a bug](#)

## 10.12. Lab 10 - Summary

You have completed the Advanced Storage lab.





[Report a bug](#)

## Revision History

<b>Revision 3.3-13</b>	<b>Mon 13 Apr 2015</b>	<b>Tahlia Richardson</b>
Fixed broken image links.		
<b>Revision 3.3-12</b>	<b>Fri 20 Mar 2015</b>	<b>Tahlia Richardson</b>
<a href="#">BZ#1203488</a> - Removed references to now-defunct Hypervisor Deployment Guide.		
<b>Revision 3.3-11.1</b>	<b>Tue 29 Apr 2014</b>	<b>Andrew Burden</b>
Publishing for 3.3.z release.		
<b>Revision 3.3-11</b>	<b>Wed 5 Mar 2014</b>	<b>Andrew Dahms</b>
Minor grammatical edits and standardization work.		
<b>Revision 3.3-10</b>	<b>Mon 10 Feb 2014</b>	<b>Zac Dover</b>
<a href="#">BZ#978176</a> - Standardizing prerequisites text		
<b>Revision 3.3-9</b>	<b>Fri 31 Jan 2014</b>	<b>Andrew Dahms</b>
Updated the screen capture for the New Virtual Machine Pool window.		
<b>Revision 3.3-8</b>	<b>Wed 15 Jan 2014</b>	<b>Andrew Dahms</b>
<a href="#">BZ#1053298</a> - Changed all references to command-line tools from 'rhev-' to 'engine-'. <a href="#">BZ#1029567</a> - Updated the package names for the self-hosted engine from 'rhev-' to 'engine-'.		
<b>Revision 3.3-7</b>	<b>Tue 07 Jan 2014</b>	<b>Andrew Dahms</b>
<a href="#">BZ#1024143</a> - Revised the description of the effects of a non-operational host on a virtualization environment.		
<b>Revision 3.3-6</b>	<b>Tue 07 Jan 2014</b>	<b>Zac Dover</b>
<a href="#">BZ#973915</a> - High-availability during storage and network disconnections.		
<b>Revision 3.3-5</b>	<b>Tue 17 Dec 2013</b>	<b>Andrew Dahms</b>
<a href="#">BZ#1012723</a> - Updated all topics in the virtual machine section to reflect current settings.		
<b>Revision 3.3-4</b>	<b>Tue 05 Nov 2013</b>	<b>Andrew Dahms</b>
<a href="#">BZ#1026639</a> - Added procedure for using sys-unconfig to seal Linux Virtual machines for deployment as templates.		
<b>Revision 3.3-3</b>	<b>Sat 24 Aug 2013</b>	<b>Zac Dover</b>
<a href="#">BZ#978176</a> - OpenLDAP now supported as a Directory Service		
<b>Revision 3.3-2</b>	<b>Tue 20 Aug 2013</b>	<b>Andrew Burden</b>
<a href="#">BZ#998807</a> - Updated topics to reflect automatic activation of storage domains when attached.		
<b>Revision 3.3-1</b>	<b>Thu 18 Jul 2013</b>	<b>Tim Hildred</b>
Initial creation for 3.3 release.		