Getting Started Guide

ConVirt

Comprehensive Management for Open Source Virtualization

Version 1.1

Convirture Corporation

Legal notice

This document and the software described in this document are furnished under and are subject to the terms of GNU General Public License version 2 (GPLv2). For details about this license, see: http://www.gnu.org/licenses/old-licenses/gpl-2.0.txt.

This document could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein. These changes may be incorporated in new editions of this document. Convirture Corporation may make improvements in or changes to the software described in this document at any time.

Acknowledgments: Xen is registered trademark of XenSource Inc, rPath is trademark of rPath.

The names of any other companies and products mentioned in this document may be the trademarks or registered trademarks of their respective owners. No association with any real company, organization, product, domain name, e-mail address, logo, person, or event is intended or should be inferred.

Contents

Chapter 3   **Creating virtual machines**

About this guide

ConVirt provides a centralized management console for managing the life cycle of virtual machines and the virtualization infrastructure.

ConVirt is designed to be both easy-to-use and sophisticated, with features that appeal to experienced administrators as well as those who are completely new to virtualization. Its comprehensive feature set includes a performance and configuration dashboard, default server pools and virtual machine images that can be provisioned as-is or customized, a fully integrated virtual appliance catalogue, and an extensible Image Store that you can customize with the virtual machine images you need.

Intended audience

This *Getting Started Guide* provides information to ensure a successful installation and initial deployment of ConVirt for new users. This guide is intended for Workstation or Data Center administrators who are responsible for configuring, provisioning, and managing virtual machine resources on computers running open source operating environments.

For more detailed information about working with ConVirt and using advanced features, see the ConVirt Wiki. For information about creating your own plug-in for ConVirt, see the *ConVirt Developer's Guide* on the ConVirt Wiki.

Using this guide

Depending on your environment and role, you may want to read portions of this guide selectively.

- Chapter 1, "Introduction to ConVirt," provides an overview of ConVirt, the virtualization infrastructure, and the key concepts and terms you should know to get started.
- Chapter 2, "Installing ConVirt," describes the steps for installing ConVirt and starting the ConVirt Workstation.
- Chapter 3, "Creating virtual machines," creating and provisioning your first virtual machine.

Conventions used in this guide

The following conventions are used in this guide:

- `Fixed-width font` is used for sample code, program names or output, file names, and commands that you type at the command line. When *`italicized`*, the fixed-width font is used to indicate variables.
- **Bold** text is used to emphasize commands, tabs, buttons, or user interface options you need to select, and to introduce new terms.
- *Italics* are used for book titles and to emphasize specific words.

Introduction to ConVirt

ConVirt provides a centralized management console for managing the life cycle of virtual machines and the virtualization infrastructure.

ConVirt is designed to be both easy-to-use and sophisticated, with features that appeal to experienced administrators as well as those who are completely new to virtualization. Its comprehensive feature set includes a performance and configuration dashboard, default server pools and virtual machine images that can be provisioned as-is or customized, a fully integrated virtual appliance catalogue, and an extensible Image Store that you can customize with the virtual machine images and appliances you need.

With ConVirt, you can manage any number of virtual machines on local or remote computers with either Xen hypervisor (Xen) or the Kernel-based Virtual Machine (KVM) virtualization platform.

Understanding key terms and concepts

The following table lists the key terms and concepts you should be familiar with to begin using ConVirt.

When you see this term	It indicates
Virtual machine (VM)	A software instance of a computer that operates like a physical computer but runs on a virtualization platform.
Server	The physical computer. The server may be a local or remote host computer, but always refers to a physical machine.
ConVirt Workstation	The computer on which the ConVirt is invoked.

When you see this term	It indicates
Managed server	A server with the Xen hypervisor or Kernel-based Virtual Machine (KVM) virtualization platform where you create and manage virtual machines using the ConVirt Workstation.
Server pool	A collection of managed servers that are grouped together for a particular reason, for example because they are owned or managed by a specific department or used to provide related services or applications.
Image	A configuration template from which Virtual Machines of a specific type can be created.
Image groups	A group of related Images.
Image Store	The repository that contains Images organized in various Image groups.
Storage pool	The list of Storage Definitions describing shared storage accessible to managed servers.

Understanding the ConVirt deployment scenarios

The typical ConVirt deployment consists of at least one **ConVirt Workstation** where ConVirt software package is installed and running. The ConVirt Workstation provides the management console for managing the virtual machine life cycle, including managing images, provisioning new virtual machines, monitoring machine operations and resources, deploying applications and appliances, migrating virtual machines from one server to another, and managing shared storage for server pools. From the ConVirt Workstation, you can then discover one or more **managed servers** where you want to begin deploying virtual machines and start managing the virtualized environment.

Although they can vary greatly in scope and size, there are essentially two deployment scenarios:

- A basic configuration in which the Xen or KVM virtualization platform is on the **local** computer where ConVirt is installed.
- An advanced configuration in which the Xen or KVM virtualization platform is on one or more **remote** servers.
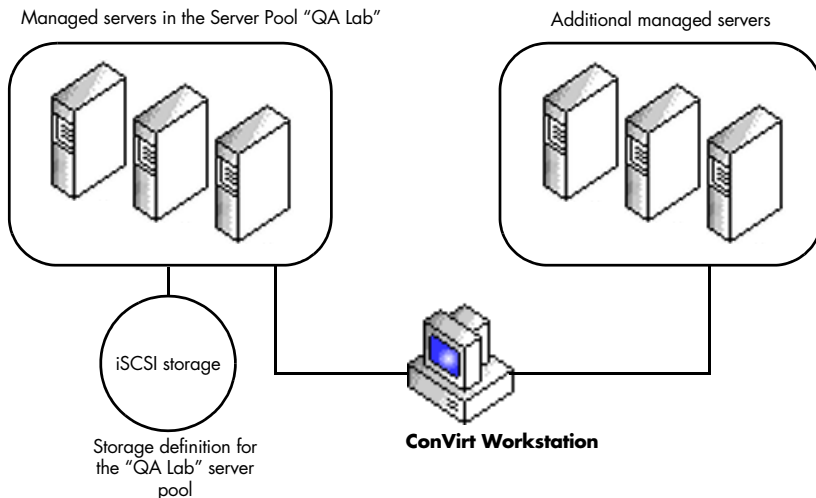
Understanding a local deployment

If ConVirt is deployed on a physical computer that has Xen or KVM virtualization capabilities, the server automatically shows up as the **localhost** managed server in the ConVirt Workstation and you can immediately start provisioning and managing Virtual Machines on the localhost node.

Understanding an advanced deployment

In larger organizations, it is more common for the ConVirt Workstation to be installed on a dedicated server that may or may not have virtualization capabilities and use it primarily to manage virtual machines on multiple remote servers. In this scenario, you add remote servers to a **server pool** in the ConVirt Workstation, then manage the virtual machines on each of the remote servers in the server pool individually or collectively.

In addition, larger scale deployments usually set up one or more **shared storage** definitions on which the disks for the virtual machines are stored. In ConVirt, each storage definition is typically associated with a specific server pool. The configuration of shared storage is necessary for using advanced ConVirt features such as Live Migration. The following figure provides a simplified view of an advanced deployment.

Managed servers in the Server Pool "QA Lab"

Additional managed servers

iSCSI storage

Storage definition for the "QA Lab" server pool

ConVirt Workstation

Installing ConVirt

Setting up the ConVirt environment involves:

- Installing **ConVirt** software on at least one computer.
- Preparing each **Managed Server** to be managed by ConVirt. If ConVirt is only used to manage virtual machines on the local computer, you can skip this step.
- Starting **ConVirt** and discovering the managed servers you have prepared.

To make the installation process as easy as possible, Convirture hosts a repository of ConVirt packages that include distribution-specific dependencies.

Before you begin

You should take the following steps to prepare for installing and using ConVirt:

1 Identify the type of configuration scenario you want to support:

- Basic configuration in which the Xen or KVM virtualization platform is on the **local** computer where ConVirt is installed.
- Advanced configuration in which one ConVirt Workstation manages virtual machines running on one or more **remote** servers.

2 Check whether the host computer where you plan to install ConVirt is a certified Linux platform. ConVirt can be installed and run on most Linux platforms, but is only tested and certified on a subset of the Linux platforms available.

To check whether your distribution is a certified or community-tested Linux platform, see:

- Certified Linux distributions
- Community-tested distributions

3 If you plan to managing virtual machines on the local ConVirt Workstation, you should also verify that computers includes a supported version of Xen hypervisor (Xen) or a Kernel-based Virtual Machine (KVM).

4 Verify the ConVirt Workstation has X server installed to enable it to display the ConVirt console.

5 Verify the ConVirt Workstation can connect to its remote managed servers using SSH.

6 Check that any remote computer where you plan to manage virtual machines includes a supported version of Xen hypervisor (Xen) or the Kernel-based Virtual Machine (KVM).

See "Preparing remote computers to be managed" on page 19 and the ConVirt Wiki for additional operating environment and distribution-specific details.

Installing the ConVirt software package

To set up ConVirt, you need get the appropriate repository for your operating environment distribution, then use a native package manager like `yum` or `apt-get` to check for dependencies and install the package. For specific instructions for your distribution, see one of the following sections:

- CentOS/Red Hat Enterprise Linux 5.x
- SUSE Enterprise Linux Server (SLES 10.x)
- Debian 5.0
- Ubuntu 8.10 / 9.04
- Other Linux distributions or the tarball

CentOS/Red Hat Enterprise Linux 5.x

To install on CentOS or Red Hat Linux 5.x:

1 Log on or switch to the `root` user.

2 Change to the `yum.repos.d` directory. For example:

```
cd /etc/yum.repos.d
```

3 Get the ConVirt repository definition from Convirture. For example:

```
wget --no-cache http://www.convirture.com/repos/definitions/rhel/5.x/
convirt.repo
```

4 Run the install command. For example:

```
yum install convirt
```

5 Start ConVirt. For example:

```
convirt
```

SUSE Enterprise Linux Server (SLES 10.x)

To install on SUSE Enterprise Linux Server 10.x:

1 Log on or switch to the `root` user.

2 Get the ConVirt repository definition from Convirture. For example:

```
zypper sa http://www.convirture.com/repos/convirt/SLES/10.x convirt
zypper sa http://www.convirture.com/repos/deps/SLES/10.x convirt-dep
```

3 Run the install command. For example:

```
zypper install convirt
```

4 Start ConVirt. For example:

```
convirt
```

Debian 5.0

To install on Debian Linux 5.0:

1 Log on or switch to the `root` user.

2 Change to the `/tmp` directory. For example:

```
cd /tmp
```

3 Get the packaging key for the ConVirt repository from Convirture. For example:

```
wget --no-cache http://www.convirture.com/repos/
convirture_packaging_pub_key
```

4 Install the key. For example:

```
apt-key add convirture_packaging_pub_key
```

5 Change to the `/etc/apt/source.list.d` directory. For example:

```
cd /etc/apt/sources.list.d
```

6 Get the ConVirt repository definition from Convirture. For example:

```
wget --no-cache http://www.convirture.com/repos/definitions/debian/5.x/
convirt.list
```

7 Run the install commands. For example:

```
apt-get update
apt-get install convirt
```

8 Start ConVirt. For example:

```
convirt
```

Ubuntu 8.10 / 9.04

To install on Ubuntu 8.10 / 9.04 :

1 Log on or switch to the `root` user.

2 Change to the `/tmp` directory. For example:

```
cd /tmp
```

3 Get the packaging key for the ConVirt repository from Convirture. For example:

```
wget --no-cache http://www.convirture.com/repos/
convirture_packaging_pub_key
```

4 Install the key. For example:

```
apt-key add convirture_packaging_pub_key
```

5 Change to the `/etc/apt/source.list.d` directory. For example:

```
cd /etc/apt/sources.list.d
```

6 Get the ConVirt repository definition from Convirture. For example: (For 9.04 use 9.x in the url)

```
wget --no-cache http://www.convirture.com/repos/definitions/ubuntu/8.x/
convirt.list
```

7 Run the install commands. For example:

```
apt-get update
apt-get install convirt
```

8 Start ConVirt. For example:

```
convirt
```

Other Linux distributions or the tarball

For other Linux distributions and versions, you can do the following:

- Download the ConVirt package (`.rpm`, `.deb`, or tarball)
- Download and install required dependencies
- Install and perform post-install steps for ConVirt
- Start ConVirt

Download the ConVirt package

To download the ConVirt software package from the Convirture Web site, go to:
`http://www.convirture.com/downloads/convirt`

Download and install the required dependencies

ConVirt requires some additional packages to be installed. This section outlines detail about installing the required packages.

You will need to be a logged in as root to install these packages.

ConVirt requires the `python-paramiko` and `socat` packages to run. In addition, the `python-paramiko` package requires the `python-crypto` package, so you may need to install that package first.

- python-paramiko and socat

  To install them on Debian or Ubuntu, use the following command:

  ```
  apt-get python-paramiko socat
  ```

  For other platforms try the native package manager, such as `yum` or `yast`, first. If the package is not available, download and install them from one of the following locations:

  - `python-crypto` (Fedora, Cent OS, and Red Hat, SUSE/SLES): http://rpm.pbone.net; Virtualization Repository

- `python-paramiko` rpms (Fedora, Cent OS, and Red Hat, SUSE/SLES): http://rpm.pbone.net; Virtualization Repository; Paramiko site
- `socat` rpms: http://rpm.pbone.net

Note Be sure to pick the right package for your distribution and architecture (32 bit or 64 bit).

■ Check and install other required packages.

In addition to `python-paramiko` and `socat` packages, the table below outlines the other packages required for various distributions. In most cases, these packages should be available in the distribution or can be installed using the native package manager, such as `yum`, `apt-get`, `zypper`, or `yast2`.

Distribution	Packages
SUSE/SLES	xen xen-tools xen-libs python xml python-gtk vte
Fedora/RHEL/CentOS	xen pygtk2 vte vnc
Ubuntu	ssh python-xen-3.3 libxen3 vncviewer
Debian	ssh xen-utils vncviewer

Note If you are only managing the KVM virtualization platform, you can ignore all `*xen*` dependencies.

Install and perform post the install steps for ConVirt

1 Run the appropriate command to install the package.

For the `.rpm` file, log on or switch to the `root` user then run:
```
rpm -Uvh convirt-1.0-1.fedora.noarch.rpm
```

For the `.deb` file, log on or switch to the `root` user then run:
```
dpkg -i convirt-1.0_1_all.deb
```

For the compressed tarball file, run:
```
tar -xzf convirt-1.0.tar.gz
```

2 Check and patch `paramiko`, if required. For example, run the following command to check the version of `paramiko` installed:
```
python -c 'import paramiko; print paramiko.__version__'
```

If the version returned is < 1.7.2, then you need to patch `paramiko` using following commands.

For the `.rpm` or `.deb` package, log on or switch to the `root` user then run:

```
cd /usr/share/convirt/install/common/patches
./patch_paramiko
```

For the tarball, log on or switch to the `root` user then run:

```
cd <convirt_install_dir>/install/common/patches
./patch_paramiko
```

where *<convirt_install_dir>* is the directory where you extracted the ConVirt files from the tarball.

Start ConVirt

If you installed using the .rpm and .deb file, you can start ConVirt with superuser privileges:

```
sudo convirt
```

Alternatively, you can log in as a `root` user and run:

```
convirt
```

If you installed using the tarball, change to the directory where you extracted the ConVirt files from the tarball, then run the `ConVirt` command. For example:

```
cd <convirt_install_dir>
./ConVirt
```

where *<convirt_install_dir>* is the directory where you extracted the ConVirt files from the tarball.

When you invoke ConVirt from the command line, you should see a ConVirt console that looks similar to this:
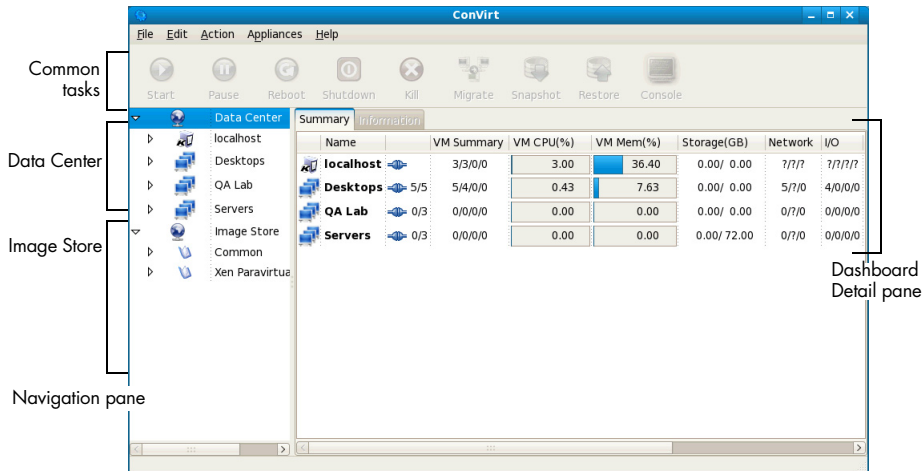
Running ConVirt as a non-root user

You can run convirt as `root` or as a non-root user. If you want to run `convirt` as a non-root user, you first need to create a private image store for the non-root user to use. However, if run `convirt` as a non-root user, you cannot manage local virtual machines.

Setting up a private image store

If you want to run convirt as a non-root user, you must create a private image store for that user account.

To create a private image store:

1 Log on as a non-root user.

2 Change to the ConVirt installation directory. For example, if you installed from the `.rpm` or `.deb` file:

```
cd /usr/share/convirt
```

3 Run the `mk_image_store` utility located in the `/install/common` directory. For example:

```
cd install/common/
./mk_image_store
```

Preparing remote computers to be managed

Before you can manage a remote server, you need to perform the following steps:

1 Install and configure Xen (3.0.4 or higher) or KVM (kvm-70 or higher). Also install `dnsmasq` and `tunctl` for your platform.

2 Verify that you can remotely log on to the managed server from the ConVirt Workstation using `ssh`.

3 Download the `convirture-tools` tar ball from:

`http://www.convirture.com/downloads/convirture-tools/`

4 Extract the contents of the archive into a directory. For example:

`tar -xzf convirture-tools-1.0.tar.gz`

5 Log on or switch to the `root` user.

6 Change to the `<install_dir>/convirture-tools/install/managed_server/scripts` directory where *`<install_dir>`* is the directory where you extracted the Convirture tools from the tarball. For example:

`cd ./convirture-tools/install/managed_server/scripts`

7 Use the `convirt-tool` command to check or configure the managed server.

To view the usage message, run:

`./convirt-tool -h`

To validate the platform without making any changes, run:

`./convirt-tool --detect_only setup`

To configure the server, run:

`./convirt-tool setup`

- For the Xen platform, running this command configures the Xend Server to listen on port 8006 and opens port 8002 for migration. The command also detects the default bridge and writes a summary of its operations to the `/var/cache/convirt/server_info` file.

- For the KVM platform, the command detects the bridge and writes a summary of its operations to the `/var/cache/convirt/server_info` file. If you are using a firewall with

KVM, you should verify that TCP port 8002 is open to allow migration.

To configure the `xend` daemon on a Xen-enabled server to listen on a Secure Socket Layer (SSL) socket, run:

```
./convirt-tool --xen_ssl setup
```

8 If the managed server is a KVM platform, install the `socat` utility using native package manager for your distribution, for example, with `yum` or `apt-get`.

Setting up a public bridge on KVM managed servers

Note `convirt-tool` version 1.1 sets this bridge up. You can skip the rest of the section.

Although setting up a bridge is not required your virtual machines are intended to be isolated from the network or each other, setting up a bridge is required if you want Virtual Machines to communicate with each other or over the local area network (LAN) to addresses outside of the virtual environment.

The steps for setting up a public bridge are distribution-specific. You should consult the documentation for the distribution you are using for information about how to set up the bridge in that environment. For example, see `https://help.ubuntu.com/community/KVM/Networking` for information about setting up the bridge on Ubuntu.

For more general instructions about setting up a bridge, see the steps described in `http://kvm.qumranet.com/kvmwiki/Networking`. If using these general KVM Wiki instructions to set up the bridge, keep the following in mind:

- You may need to install `tunctl`, if it is not included in your distribution.
- You should set the name of the bridge as `bro` instead of `switch`.
- You should remove the `/sbin/sudo` as the ConVirt-generated `qemu` command line is not compatible with the use of `sudo`.

Once the bridge setup is done, run the `convirt-tool` again. This will set the bridge name in the `/var/cache/convirt/server_info` file.

Adding managed servers to ConVirt

Once you have prepared managed servers, you need to add them to ConVirt.
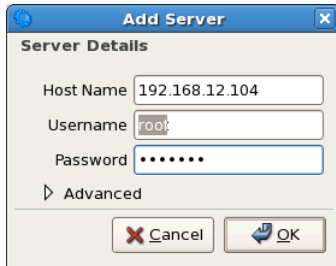
To add a managed server:

1 Open the ConVirt console. For example:

```
convirt
```

2 Select a Server Pool in the left pane, then right-click. For example, select the default server pool **QA Lab**, then right-click to display the context menu:
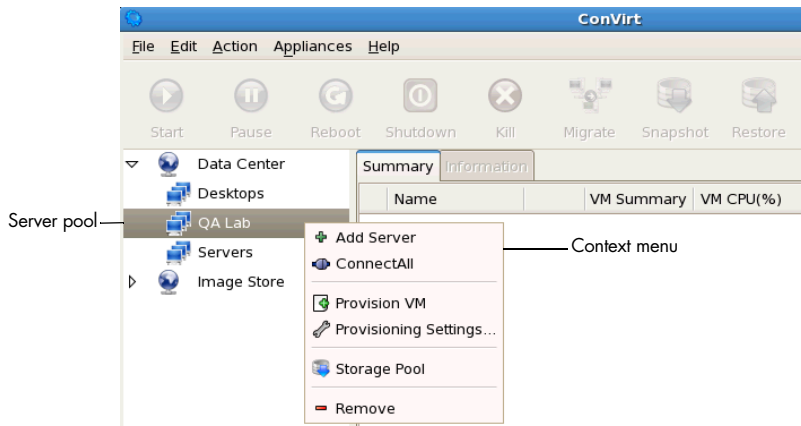
3 Click **Add Server**.

4 Select the Virtualization Platform, then click **OK**. For example, select **KVM**, then click **OK**:
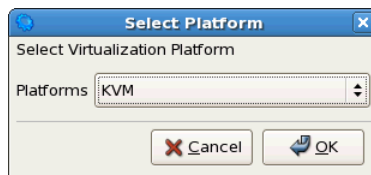
5 Type the host name or IP address of the server, and a user name and password for account with permission to access the managed server, then click **OK**. For example:

If the connection is successful, the managed server is added to the Convirt console.

You can now create and manage the complete life cycle of virtual machines on this managed server.

Creating virtual machines

You are now ready to start creating and managing virtual machines on the local or remote managed server. This chapter provides sample tasks to illustrate using the basic features of ConVirt in your first virtual machine.

Viewing servers in the ConVirt console

When you start ConVirt after installation, you should see a ConVirt console that looks similar to this:
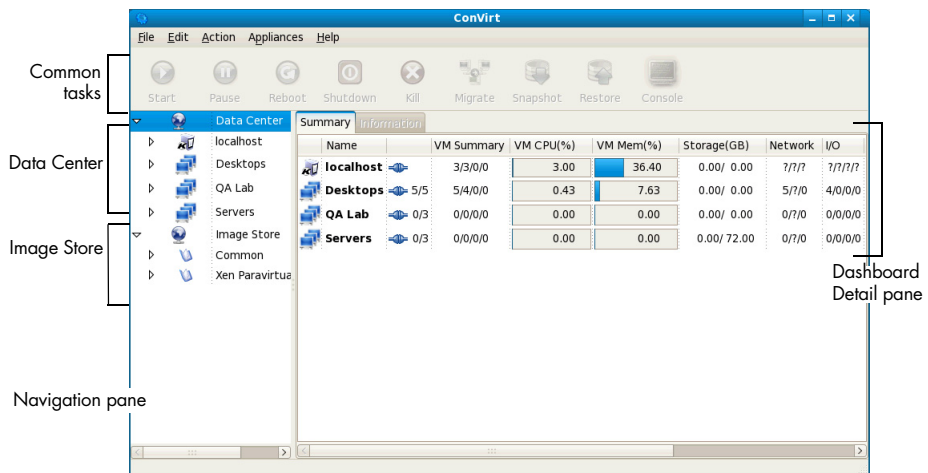
In the left navigation pane, there are three server pools created by default in the **Data Center**:

- **Desktops** server pool
- **QA Lab** server pool
- **Servers** server pool.

You can create your own or modify these default server pools to suit your needs.

If you expand the **Image Store**, you will also see two default image groups:

- **Xen Paravirtual** image group.
- **Common** image group.

The **Xen Paravirtual** group contains default Paravirtual images that allow you to do virtual machine installation over the network for the Xen Platform. The images in the **Common** group are Hardware-assisted Virtual Machine (HVM) images. The Common group images can be deployed on either Xen or KVM. Installing Common images requires the underlying processor to support virtualization.
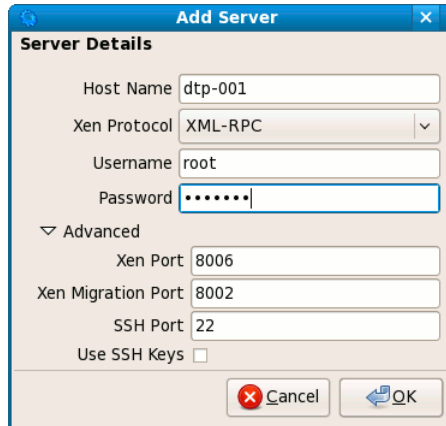
Adding a managed server

If ConVirt is installed on a computer with virtualization capabilities, the **localhost** managed server is automatically added to the Data Center. You can also manually add remote managed servers that have been configured with Xen or KVM.

To add a server to ConVirt:

1 Select the Data Center or a Server Pool in the left pane to specify where the new managed server should be placed.

2 Click **File > Add Server**.

3 Select the Virtualization Platform, then click **OK**. For example, select **Xen**, then click **OK**.

4 Type the host name or IP address of the server, a user name and password for account with permission to access the managed server, and other details, then click **OK**. For example:

If the connection is successful, the managed server is added to the Convirt console.

You can now create and manage the complete life cycle of virtual machines on this managed server.

Viewing and editing default images

An **image** is a template from which you can create one or more **virtual machines**. For example, one of the default images included with ConVirt is the **Linux _CD_Install** in the Common image group.

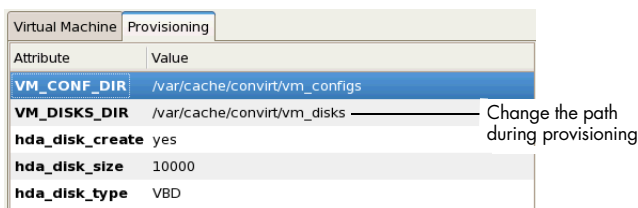To view or edit the existing Linux_CD_Install image in the ConVirt console:

1 Expand **Image Store > Common**, then select **Linux_CD_Install**.

2 Right-click, then select **Edit Settings** on the context menu. By default, this image is configured to create a virtual machine with 512 MB of memory, one virtual CPU, the disk `$VM_DISKS_DIR/$VM_NAME.disk.xm`, and the `cdrom` device defined.

For example:
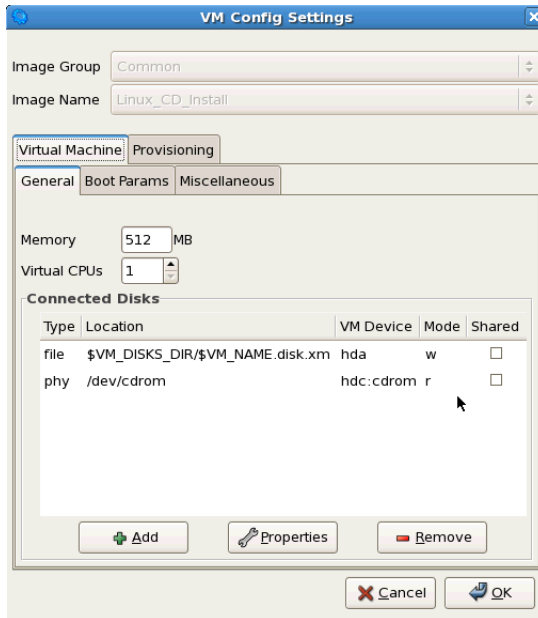
The default VM_DISKS_DIR and VM_DISK_NAME variables are substituted with appropriate values when you provision the virtual machine. For example, by default the VM_DISKS_DIR variable is replaced with /var/cache/convirt/vm_disks if you provision a virtual machine using this image without making any changes to its settings.

You can set another location for the disk for the virtual machine by clicking the **Provisioning** tab and changing the path. For example:

Creating a new image based on an existing image

The default images included with ConVirt give you a good starting point for creating new virtual machines without making changes. In most cases, however, you may find it useful to create custom images by modifying an existing image to suit your needs. For example, you can create a new custom image for Fedora virtual machines by modifying the existing Linux_CD_Install image in the ConVirt console.

To create a new Fedora image by modifying the Linux_CD_Install image:

1 Expand **Image Store > Common**, then select **Linux_CD_Install**.

2 Right-click, then select **Create Like** on the context menu.

3 Type a descriptive name the new image, then click **OK**. For example, type Fedora_8 if you are creating a new image for deploying Fedora Core 8 virtual machines.

4 Select the newly created image, right-click, then select **Edit Description** on the context menu to display an editor window containing HTML-like syntax.

   This editor enables you to document the image-specific details or management processes that need to be followed in your organization.

5 Change the title from `Linux CD Install` to `Fedora 8` and make any other changes to the description of this image, then click **Save**.

6 With the Fedore_8 image selected, right-click, then select **Edit Settings** on the context menu. You can modify specific properties to customize the image to your needs. For example:

   • Change the **Virtual CPUs** to 2 to specify that virtual machines using this image should be configured with two CPUs.

   • Select the `/dev/cdrom` in the Connected Disks list, then click **Properties** and select **Use ISO File** so that virtual machines created using this image use an ISO file instead of the

physical `cdrom` device. In the Disk Location, enter the full path to the ISO file. For example:
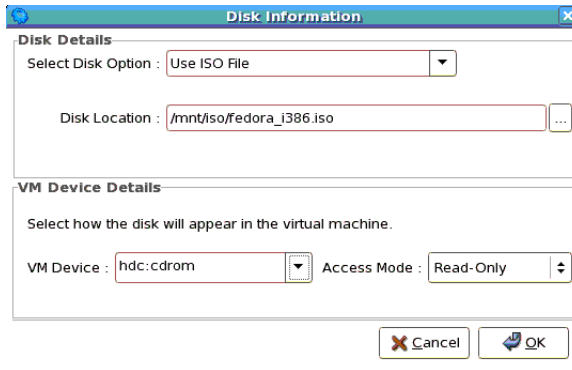
Make sure that there are no spaces in the path to the ISO file location.

7 Click the **Provisioning** tab and change `VM_DISKS_DIR` to `/mnt/storage/vm_disks`. Make sure that there are no spaces in the directory location.
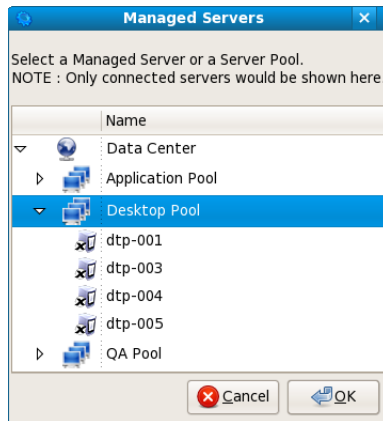
Note During provisioning, the `VM_DISKS_DIR` directory and the `.iso` file location must be accessible on the server where the virtual machine is being deployed. In most cases, the `VM_DISKS_DIR` directory and the `.iso` file location are set up on shared storage to which the managed server has access. This ensures that when you create a virtual machine using the image on a particular managed server, the server can locate its disks and ISO file. If you are not using shared storage for the `VM_DISKS_DIR` directory and the `.iso` file location, you should verify the locations specified in the image exist on the managed server or change the settings during the provisioning process.

Provisioning a virtual machine from the custom image

To provision a virtual machine from the Fedora_8 image you created:

1 Select the Fedora_8 image.

2 Right-click, then select **Provision** from the context menu.

3 Select a managed server or a server pool. For example:

Note If you select a Server Pool, ConVirt will find a suitable managed server within that server pool for you. This simplifies the provisioning process when there are many servers and virtual machines already deployed.

4 Type the Virtual Machine name, then click **OK**. For example, type Fedora8-VM, then click **OK**.

Make sure there are no spaces in Virtual Machine name.

Note Keep in mind that the VM_DISKS_DIR directory and the .iso file location must be accessible on the server where the virtual machine is being deployed during provisioning. If the VM_DISKS_DIR directory and the .iso file location are set up on shared storage, you should be able to create a virtual machine without modifying these properties. If you are not using shared storage for the VM_DISKS_DIR directory and the .iso file location, you should verify the locations exist on the managed server or change the properties, as needed.

5 Expand the managed server selected in Step 3 to see the newly created Virtual Machine.

6 Select the Virtual Machine, right-click, then select **Start** on the context sensitive menu to start the virtual machine.

Alternatively, you can select the Virtual Machine and click **Start** on the common tasks tool bar.

7 Click **Console** on the common tasks tool bar to open a new tab with the virtual machine's name.

8 Type the credential to log on to the new virtual machine, if prompted.

If the `.iso` file you specified for the image is readable, the installation process for the virtual machine will begin. Follow the instructions displayed on the virtual machine tab to complete the installation.

9 When installation is complete, click **Shutdown** on the common tasks tool bar, then click **Yes** to confirm shutting down.

10 With the new virtual machine selected, right-click, then select **Set Boot Device** on the context menu.

11 Change the boot device to **Disk**. For example:

12 Click **Start** to restart the new virtual machine, then click **Console** to display the terminal console on the virtual machine.

You should see the log in prompt for the newly created Fedora installation. Congratulations, you are now ready to begin deploying applications and managing your new virtual machine.

For information about using more advanced ConVirt features, see the ConVirt Wiki.

Getting help and more information

Additional static and dynamic resources are available on the Web at www.convirture.com. These additional online resources include the ConVirt Wiki and Community resources. For example, click **Community** on the Convirture Web site to access:

- Community Forums
- Community Wiki
- Certified Configurations
- Product Documentation
- Developer Resources
- Downloads
- Mailing List
- How To topics
- Bug Tracking and Feature Request systems
- Frequently Asked Questions (FAQ)

To access online resources from the ConVirt console, click the **Help** menu.

If you are interested in extending the virtualization infrastructure with an additional virtualization plug-in to work with ConVirt, you may want to download the *Developer's Guide* from the Developer Resources or join a ConVirt forum.

If you are new to virtualization, some additional resources you may find useful include:

- http://www.xen.org
- http://kvm.qumranet.com/kvmwiki

Questions or suggestions? Contact us...

If you have questions or comments, we look forward to hearing from you. For information about contacting Convirture, visit our Web site at www.convirture.com. From the Web site, you can get the latest news and information about product support and services, community activity, upcoming events, and sales.