

Table of Contents

- [Introduction](#)
- [Definitions](#)
- [Whats New in v0.9 ?](#)
- [Deployment Instructions](#)
 - [I. Deployment from tar ball](#)
 - [II. Deployment from an RPM package](#)
- [Prerequisites](#)
- [Environment setup](#)
 - [SSH setup](#)
 - [Remote file browsing](#)
- [Features](#)
 - [The Dashboard](#)
 - [Server Pool operations](#)
 - [Server operations](#)
 - [Virtual Machine operations](#)
 - [Summary Information](#)
 - [ConVirt and user accounts](#)
 - [VM Migration](#)
 - [Smart VM Placement](#)
 - [The Image Store](#)
 - [Appliance Management](#)
- [The ConVirt config file \(convirt.conf\)](#)
 - [PATH section](#)
 - [APPLICATION section](#)
 - [ENVIRONMENT section](#)
 - [Managed Server specific Sections](#)
 - [CLIENT CONFIGURATION section](#)
- [Tips](#)
- [Known bugs/caveats](#)
- [Platforms](#)
- [Licensing](#)
- [Feedback](#)
- [Copyright/Acknowledgments](#)
- [Appendix](#)
 - [CentOS/RHEL 5.1](#)
 - [Ubuntu Gutsy](#)
 - [Debian Etch](#)
 - [OpenSUSE 10.3](#)
 - [Xen 3.2](#)

Introduction

ConVirt is an intuitive, graphical management tool providing comprehensive life cycle management for Virtual Machines and virtualization infrastructures. ConVirt is built on the firm design philosophy that ease-of-use and sophistication can, and should, co-exist in a single management tool. With its central console paradigm, performance and configuration dashboard, soup-to-nuts Virtual Machine lifecycle management, integrated Virtual Appliance Catalogue, and many more great features, ConVirt should prove an invaluable

tool for seasoned administrators as well as those new to virtualization.

Definitions

VM	Virtual machine (Xen Virtual Machine in this context)
Server	Refers to a physical host or machine
Client Machine	Machine on which the ConVirt is invoked is referred to as client machine.
Managed Server	Server to be managed using ConVirt.
Server Pool	A collection of managed Servers typically owned/managed by a department or used to provide related services/applications.

Whats New in v0.9?

- **Plugin Development Guide** : With ConVirt 0.9, one can develop a plugin for their favorite virtualization platform and integrate it with ConVirt. This enables ConVirt users to be a single user interface for managing multiple platforms. The details on how to develop such a plugin are in the Plugin Development Guide in the doc set. The Xen management now implemented as a plugin provides a good reference.
- **Server Pool Settings** : Now users can define their own settings at server pool, which get used while provisioning any image in that server pool. For example a storage destination can be specified at server pool level. All the vms provisioned in that server pool will get space allocated from that storage.
- **Image Groups** : Organize your images in to groups for easy management.
- **Enhanced Server information** : The Server Information tab now shows rich information about the network interfaces and storage.
- **Summary Information** : The summary information gives you at a glance about your managed environment.
- **Image Cache consistency** : The convirt creates a cache on each managed node for simplifying repeated provisioning. This version addresses some consistency problems between the image store and the caches on remote servers.

Deployment Instructions

ConVirt is distributed both as a source/binary tarball and as a binary rpm for a few specific distributions. Please checkout the distribution specific notes in the [Appendix](#)

I. Deployment from tar ball

1. cd to the location where you've extracted the tarball.

```
e.g.> cd ~/convirt.0.9/
```

2 . If you haven't already, deploy python-paramiko (an SSH client library ConVirt uses for remote management) in your environment. Python-Paramiko is available for download at:

<http://www.lag.net/paramiko>.

(follow the installation instructions at the site carefully).

Here are useful links for the rpms.

<http://rpm.pbone.net/index.php3/stat/4/idpl/3714841/com/python-crypto-2.0.1-15.i586.rpm.html>

<http://rpm.pbone.net/index.php3/stat/4/idpl/5215518/com/python-paramiko-1.7.1-0.pm.1.i586.rpm.html>

There is a bug in paramiko library used by ConVirt. The ConVirt distribution contains necessary files to patch this bug. In order to apply this patch, run patch_paramiko script as root. This patch is needed only on the ConVirt/client node.

```
# cd <convirt install>/install/common/patches  
# ./patch_paramiko
```

3 . Make sure ConVirt has execute permissions:

```
e.g. > chmod 0755 ./ConVirt
```

4 . Run ConVirt and enjoy!

```
e.g. > ./ConVirt
```

By default ConVirt would use the image store that is in the location where you untared the tarball, so this step is not required. But if you want to use a different location with more disk space, use the mk_image_store utility to create a copy of the Image Store (ConVirt's provisionable Virtual Machine Image repository)

```
e.g. > sh ./install/common/mk_image_store <path for image store>
```

This will create the Image Store at the location specified.

If you run the script without any argument, it would create the image store under ~/.convirt location. If you are a root user, it would create the image store at /var/cache/convirt.

Note : You need to modify the convirt.conf to point to the new location of the image store. Change the image_store and appliance_store values in the convirt.conf. Refer to [Initial Setup](#) section under the Image Store.

II. Deployment from an RPM package

The following steps need to be done as a root user.

1. Remove any earlier versions of `convirt` from your environment.

```
e.g.> # rpm -e xenman convirt
```

2. `cd` to the directory you've downloaded the rpm package to and install ConVirt:

```
e.g.> # rpm -Uvh convirt-0.9-1.fedora.noarch.rpm
```

3. If you intend to manage (Xen enabled) local machine, you are good to go. Run ConVirt as a root user.

```
e.g.> # convirt
```

If you intend to use `convirt` from a non-root user, then you can not manage the local machine. Also, in this case, you need to set up your own private image store. The following command would create an image store under `~/convirt` directory. Note: This would be executed as a non-root user.

```
e.g.> $ /usr/share/convirt/install/common/mk_image_store
e.g.> $ convirt
```

NOTE: ConVirt requires the `python-paramiko` package for remote operations over SSH. A `python-paramiko` rpm package is available for most distro's; however, if you cannot locate a suitable package for your distro, you can manually deploy it (see section I.2 above), and then run `rpm` with dependency checks disabled.

```
e.g.> rpm -Uvh convirt-0.9-1.fedora.noarch.rpm --nodeps
```

Refer to [Initial Setup](#) section under the Image Store.

Prerequisites

Prerequisites for running ConVirt : (Don't forget to consult Appendix for Linux distribution specific notes)

Client Machine :

- Xen 3.0.4 or later installed.
- X server (to display ConVirt UI)
- Ability to connect to Managed Server through SSH
- Paramiko library (if you have problems installing along with ConVirt)
- ConVirt installed
- Patching paramiko

Managed Server :

- Xen 3.0.4 or later (booted in to Xen kernel and Xend running). Xen 3.1 and above is recommended.
- SSH daemon running
- Run the script as 'root' user to configure the managed node. These scripts can be found under config-

scripts directory. This directory is a top level directory for the tarball users and resides under /usr/share/convirt for rpm users.

- For Xen 3.0 there is a script configure-xend.sh under xen-3.0 directory

```
# cd xen-3.0
# ./configure-xend.sh <xen_version>
xen_version can be 3.0.3 or 3.0.4
```

- For Xen 3.1 there is a script configure-xend.sh under xen-3.1 directory

```
# cd xen-3.1
# ./configure-xend.sh 3.1
OR
As Xen 3.1 supports SSL, one can set it up with a test certificate as follows.
# ./configure-xend.sh 3.1 SSL
```

NOTE : For modifying parameters used to generate the test certificate, modify the script before executing it.

- For Xen 3.2 there is a script configure-xend.sh under xen-3.2 directory

```
# cd xen-3.2
# ./configure-xend.sh 3.2
OR
As Xen 3.2 supports SSL, one can set it up with a test certificate as follows.
# ./configure-xend.sh 3.2 SSL
```

NOTE : For modifying parameters used to generate the test certificate, modify the script before executing it.

- Vncviewer for hvm domains.
- X server is required to be installed, but need not be started. This is required only for VMs using vncviewer, usually Vms using hvm method.
- Firewall is either shutdown or allows for specific ports 8005 or 8006 (for xml-rpc) and 8002 (used for migration)

Environment setup

This section describes some additional setup for managing multiple Servers. This assumes that you have configured each of your managed servers as per previous section.

This is NOT required if you want to manage virtual machines on local host/client machine only.

SSH setup

ConVirt uses ssh to read a bunch of configuration files, create VBDs, LVMs on the remote node. Basically all managed servers need to trust all client machines. (Yes, u can have more than one client machines)

From client machine, ssh to managed server from the client machine using the account

from which convirt would be started.

```
# ssh <managed server name>
```

This will prompt you to add the key to the known_hosts. Say yes.
This will add the /etc/ssh_host_key.pub from the managed server
to users \$HOME/.ssh/known_hosts on client machine

(Alternatively you can manually add it)

```
## Repeat above steps for each managed server.
```

If you want to use password based authentication, then you are done.
You can skip the rest of the section.

For a small environment it may be OK to use password based authentication, but in a large setup we recommend using key based authentication for convenience and tractability.

Refer to SSH manuals and on line material for setting up key based trust and using ssh-agent.

Here are couple of useful URL's

<http://www.suso.org/docs/shell/ssh.sdf>

<http://www.linux.ie/articles/tutorials/ssh.php>

Remote file browsing

While managing remote servers, there are operations like, "Open VM File" etc that would require selecting a file on a remote server. For GNOME users, ConVirt uses gnome-vfs to browse files on the remote server. As this is done on a separate channel it would require user to enter the password again for the managed server. When prompted for saving the password, it is recommended not to save password in key-ring for security reasons.

The user is NOT prompted for password if the key based authentication is used. The user experience is quite seamless between localhost vs remote managed server management.

Patching paramiko

There is a bug in paramiko library used by ConVirt. The ConVirt distribution contains necessary files to patch this bug. In order to apply this patch, run patch_paramiko script as root. This patch is needed only on the ConVirt/client node.

```
# cd <convirt install>/patches
```

```
# ./patch_paramiko
```

ConVirt Features

The Dashboard

The Dashboard is a consolidated listing of all known managed servers along with critical performance, availability and configuration metrics for each. It provides the user the ability to ascertain the state of his/her entire deployment at a glance. In addition, most common administrative tasks can be launched by right clicking on a server.

- **Launch.** The Dashboard may be launched by selecting 'Server Pool' in the navigator on the left and then clicking on the Summary Tab on the right hand side. (Note: Upon startup, ConVirt launches the Dashboard by default).
- **Operations:** Left-Clicking a row in the Dashboard selects the associated managed server. The following actions may be then be performed:
 - Double-Click: Connect to the managed server (if necessary) and drill down into a more detailed view. This selects the server's node in the navigator on the left hand side and brings up the Summary tab for the server on the right.
 - Right-Click: Context sensitive menu. Most server operations can be executed directly here.
 - Sorting: Clicking on the column header will re-sort the listing based on the clicked column. (not available for all columns)
- **Data.** Each row in the Dashboard corresponds to a managed server. The fields are:
 - Server. The name of the managed server.
 - Connection. Connectivity status to the managed server (i.e. whether ConVirt has an active connection to the server).
 - VM Summary. A compact listing of VM status on the server. Total(known)/Running/Paused/Crashed respectively.
 - VM CPU(%). Aggregate processor usage by VM's running on the server. (Does not include the host OS/Domain-0's processor usage).
 - VM Mem(%). Aggregate memory usage by VM's running on the server. (Does not include the host OS/Domain-0's memory usage).
 - Server CPU(s). Number and clock speed of the *physical* processors on the managed server. (if available)
 - Server Mem. Total, usable *physical* memory installed on the managed server. (if available)
 - Version. The Virtualization platform version being reported by the managed server.

* Most of the data fields in the Dashboard listing are available for managed servers running Xen v3.0.4 or above. For servers running earlier versions of Xen, most fields will display N/A.

Server Pool Operations

A Server Pool node on left navigator represents a group of managed servers.

- Desktops, QA Lab and Servers server pools are created by default. They serve as a good starting point. You can create your own Server Pools as well.
- If ConVirt is installed on a machine running Virtualization enabled kernel, it adds Local machine automatically and would manage Virtual machines on them.
- Settings : This allows users to set server pool level settings. These are simple definitions that would override or fulfill variables in the image at the time of provisioning. For example, the Desktop group is associated with class B storage and can be specified with `VM_DISKS_DIR=/mnt/class_b_storage/`. Now at the time of provisioning, the

VM_DISKS_DIR would get substituted by the class b storage location. This helps to manage all servers in a uniform fashion.

- Add Server : Adds additional remote managed servers to a pool

You are prompted for basic information:

- Host name : Name of the server/host to be added.
- Username : (typically root)
- Password : (NOT SAVED ANYWHERE ON DISK/can be left blank if authentication is done using keys)

In the advance section one can configure the following

- Xen Port : Port on which Xen Daemon (xend) is listening. Typically 8006.
- SSH port : Typically 22
- Use Keys : If you wish to use key based authentication instead of the password.

Notes:

- Make sure that Xen daemon (xend) listening on port 8005/8006 on managed server. If not, Run the configuration script mentioned in the prerequisites section
- Make sure that you can use ssh to login to the managed server
- Remove Server : A server can be removed from the list, by selecting it and choosing "Remove Server"
- Change Group membership : Simply select the Server, drag and drop it to its new group.
- Connect All : This is a convenient feature to make ConVirt connect to all servers in the pool and start collecting measurements and show them in the dashboard.

Server operations

The following operations are allowed on a managed server.

- Start All VMs : Start all VMs on a selected server.
- Shutdown All VMs : Shutdown all VMs on a selected server.
- Kill All VMs : Kill all VMs on a selected server.
- Provision VM: Allows you to create a new VM with very few parameters. (see section on Image Store for more details)
- Import VM Config File(s) : Allows you to import one or more VM files to be managed by ConVirt.
- Migrate All : Migrate all VM configurations and running Vms to a selected server.

Virtual Machine/VM operations

- Edit VM Settings : Change the configuration settings for the selected VM
- Edit VM Config File : Edit the VM's configuration file directly.
- Start : Start the selected VM
- Pause : Toggle button to Pause/Resume running VM
- Reboot : Reboot a selected VM
- Shutdown : Shutdown a running VM
- Kill : Hard kill on a VM
- Snapshot : Save state of a running VM to a file
- Restore : Restore a VM from a stored snapshot.
- Delete : Delete VM file and associated VBD/LVMs.
- Remove VM Config File : Removes VM file name from the list of VM known to convert.

- Migrate VM : Migrate a running VM or VM configuration to a new Server Node with simple drag and drop operation or using a menu. operation.

Summary Information :

A new menu option under the Help menu allows you to see at a glance summary information about number of servers and vms being managed in each group.

ConVirt and user accounts

ConVirt can be invoked from either root user or non-root user account

- Root user : When ConVirt is invoked from root user, it can manage local machine as well as other managed servers for which credentials need to be provided.
- Non root user : When ConVirt is invoked from non-root account, it can not manage VMs on on the local machine. However, this user can manage other managed servers. Wherever necessary, ConVirt prompts for credentials..

VM Migration

ConVirt now allows VM configurations or running Virtual machines to be migrated to another server. This allows the administrator to re-organize the virtual machine to physical machine relationship to balance the workload. For example, a VM needing more CPU can be moved to a machine having more cpu cycles. For VM Migration to succeed the following points must be considered.

- Shared storage for all GuestOS disks
- Identical mountpoints on all servers (hosts)
- The kernel and ramdisk when using para-virtualized virtual machines should also be shared.(This is not required if pygrub is used.)
- Centrally accessible installation media (iso)
- Preferably use identical machines with same version of Virtualization platform running (preferably Xen 3.1 and above)
- Migration needs to be done within the same subnet.

Migrate All : There is also a convenient feature to allow migrating all Virtual machines. This is useful particularly when a Server is to be upgraded or brought down for maintenance.

Smart VM Placement

Administrators are constantly faced with selecting an appropriate Server to create a new VM. They have to look at the utilization of resources across various servers and pick one so as to optimize the resource utilization. This becomes very difficult as the number of servers grow. ConVirt helps with this problem by

providing a "Provision VM" operation on a Server Pool. This advanced feature, selects an appropriate server for the provisioning operations. It uses CPU, Memory and number of configured Vms on the server while making the selection. This smart placement of VM help achieve even distribution of workload among available servers.

The Image Store

ConVirt allows administrators to define their images and create Virtual Machine configurations from them. For example, in a particular data center, four types of machines are to be provisioned frequently. They are Red hat, CentOS, Suse and Ubuntu. You can configure ConVirt to point to kernel and ramdisk of each of these distributions and install/deploy many virtual machines using predefined images. The collection of images are referred to as an Image Store.

ConVirt ships with a default image store containing a few useful provisionable images. One can either edit them in place or use the new "Create Like" feature to create an image that is similar to one of the existing images.

Starting with the 0.9 version, the image store allows you to organize your images in to various image groups. One can drag and drop the image to move it from one group to another.

Sophisticated users may also construct their own, arbitrarily complex image descriptions and provisioning schemes and add them to the Image Store. (For instructions on how to build custom provisionable images, please consult the 'Image Builder's Guide,' a part of this documentation set).

- **Location.** The Image Store is listed in the navigator. Clicking or expanding the Image Store node results in a listing of the available, provisionable images. Right click on the Image Store presents a menu with "Import Appliance" see the next section for details on appliance management.
- **Image Operations.**
 - Click. selects the image and displays useful information about it in on the right hand side. This information typically included a brief description of the provisionable image, pre-requisites and deployment instructions. So, read this information carefully before attempting to provision the image.
 - Right Click. brings up a context menu with various image specific tasks. These are:
 - *Provision.* Start the provisioning process. Launches a dialog prompting for the managed server to target for provisioning.
 - *Edit Settings.* Alter the definition of the image. Any changes made here are permanent and apply to all subsequent VM's provisioned using this image. (more detail below)
 - *Edit Script.* (Advanced Feature). Alter the mechanism used to provision VM's using this image. This is an advanced feature and should be left alone unless you are designing your own provisionable image.
 - *Edit Description.* (Advanced Feature). Edit the description associated with this image. This is the description that appears on the right hand side once the user left-clicks on an image name in the navigator. Unless you are designing your own provisionable image, leave this alone.
 - *Create Like .* (Advanced Feature). Creates a copy of the selected image. This is particularly useful in creating customized images to suit a particular data center needs
 - *Delete..* Deletes the selected image from the library. CAUTION : All the content related to the image would be lost permanently.
- **Edit Settings Dialog.** This is a multi-purpose, context sensitive dialog that can be launched to manipulate settings for provisionable images or existing VM configurations. When launched in context of an image, this dialog presents two top-level tabs, each affecting a different aspect of the provisioning process:

- **Virtual Machine.** This tab controls the various parameters and settings controlling the *configuration** of the VM's provisioned with this image. In typical usage, user will likely not make any changes here, save for perhaps changing disk settings to better match the deployment environment.
- **Provisioning.** This tab controls the various parameters and settings controlling the *mechanism** by which VM's are provisioned using the selected image. A typical alteration here would be to change the disk size and type parameters governing the default sizes and types (VBD/LVM) of system disks created during provisioning process.

* a detailed description of the configuration/mechanism distinction and how parameters may be shared between the two is an advanced topic, more fully addressed in the 'Image Builder's Guide.'

Initial setup : Out of the box the images supplied would work with default values. You might want to do some initial setup to suite to your environment. There are few options in the provisioning area that are worth mentioning.

- VM_DISKS_DIR : Location where disks for newly created VMs would be placed. Default /tmp.
- VM_CONF_DIR : Directory where the VM config files would be placed. Default /etc/xen
- http_proxy : Proxy to be used while provisioning. Default to no proxy (blank). The format for the proxy is http://myproxyserver:myproxyport

Another thing you might want to customize is to make sure that the device for the cdrom is correct for images that install from a CDROM. It is usually /dev/cdrom, but it might be different on your machine. You can make this change in the disk area under the general tab.

Template : The image store also contains a `_template_` image. This is a skeleton for an image and can not be provisioned directly. One should populate the template with suitable values and then use create like from the template to create one or multiple images. This particular template is encouraging shared area for the vm disks, one of the primary requirements for migration.

Appliance Templates : See the section under Appliance Management.

Appliance Management

Appliances are fully configured application images that can be deployed on either physical or virtual machines. Their out of the box experience makes them very attractive solution. They are becoming exceedingly popular as they reduce the costs for the software vendors as well as IT cells using them. This version of ConVirt aims to simplify Appliance Lifecycle Management.

Appliance Catalog : This version has an online appliance catalog containing few popular rPath appliances. Each entry in the catalog contains url from where the appliance can be downloaded, description and a bunch of technical metadata. The catalog also has a search/filter capability to find the appliance that you are looking for. Once you find what you are looking for, simply pick import it into the ConVirt Image Store. Once in the store, multiple instances of appliances can be provisioned as any other image. It shares all the flexibility and configurability offered by the provisioning scheme described in the previous section. We would be adding more appliance vendors soon.

Appliance Management shortcuts : Once the rPath Appliance is provisioned, the context sensitive menu for Appliance Lifecycle Management shows up. It contains short cuts to the application as well as following common appliance management operations.

- Backup Configuration
- Email setup for notifications

- Network configurations
- Setting Root Password
- Setting Time Zone
- Check for Appliance updates
- User Management
- Viewing Logs

This menu is available only when the appliance is running. Out of the box the appliances are configured to obtain a dynamic address using dhcp. When any of the above mentioned menu is invoked for the first time, you are prompted for entering the ip address/hostname of the VM. You can find this information mostly from the console messages when the appliance boots up. You can use the Specify Details menu item to enter this information apriori or change the information after the fact.

Manual Specification : For importing appliances that not in the catalog, one can use the "Specify Manually" option to specify the URL manually and specifying some additional information. For example :

- Provider : Appliance provider, rPath.
- Package : Describes the package as being file system image or appliance in XVA/OVA format.
- Architecture : 32 Bit or 64 Bit
- PAE : Check if you know that the kernel uses physical address extension.
- HVM : If the image is contains xen paravirtualized kernels or is to be instantiated in hvm mode. (For rPath Appliances always leave hvm box blank.)

Importing a reference Disk / Cloning a VM:

If you have a reference disk or want to clone an already provisioned VM, you can use the Manual Specification method mentioned above. Simply specify the path of the reference disk as a url and choose 'Other' as a provider. It is preferred to gzip the disk image.

For example, I can specify the following parameters for cloning the gold image for our employee desktop.

- URL : /path/to/my/disk/gold.winxp.img.gz
- Image Name : Employee_Desktop
- Provider : Other
- Package : File System
- Architecture : X86_32
- PAE : Leave checkbox unchecked.
- HVM : Select the checkbox as we will be using HVM method to create employee desktop

Appliance Templates : When an Appliance is imported, appliance_vm_conf.template, appliance_hvm_conf.template, appliance_image_conf.template and appliance_desc.template are used. You may want to customize these to suite your environment.

ConVirt config file (convirt.conf)

ConVirt's default configuration is kept in the convirt.conf file. This file is in simple text format, so it can be easily changed to suit your environment.

Here is a brief description of important configuration parameters:

PATH section

This section contains default directory/file locations. Some important one's include:

- `image_store`: This is the path variable setting for image store.
- `log_dir`: This is the location where ConVirt writes its logs. By default, this is set to `/var/log/convirt`
- `exec_path`: This is the path variable setting for remote shell operations. Default value is `$PATH:/usr/sbin`.
- `xenconf_dir` : Default location where the new vm config files would be placed.

APPLICATION section

This section keeps application specific data.

- `groups` : Mapping between Server Pools and Servers.
- `vms` : List of virtual machine configuration files being managed by ConVirt.

ENVIRONMENT section

This section contains server environment information.

- `lvm_enabled` : if the server has lvm and want ConVirt to create LVM for new VMs.

Managed Server specific Sections

The config file has section for each managed server. This contains information required by ConVirt to connected managed servers. This sections are relevant only on client machines.

example :

```
[192.168.0.102]
is_remote = True
login = root
xen_port = 8005
ssh_port = 22
use_keys=False
```

CLIENT CONFIGURATION section

This section contains items specific to client/user preferences.

- `gnome_vfs_enabled` : True/False. Set this to True if you are using ConVirt on Gnome desktop and have patched paramiko. See section on Remote File Browsing for more details.
- `log_file` : `convirt.log` file. This can be set to a value of “console” for getting all the verbose messages

on the console instead of a file.

- `enable_log` : Enable logging to `convirt.log`.
- `http_proxy` : Proxy url. example : `http://proxyserver:proxyport`
- `enable_log` : if the log file should be created or not.
- `log_file` : name of the log file. If it is just a file name, it is contained under the `log_dir` (PATH section). `~` representations for the home directory are not allowed. A special value of `console` can be used to put the log in the console window.

Additional Notes:

Upon startup, ConVirt looks for `convirt.conf`, in order, at the following locations:

- `./convirt.conf` - (current directory)
- `~/convirt/convirt.conf` - (user's home directory)
- `/etc/convirt.conf` - (global location)

If it doesn't find a valid, writable configuration file, ConVirt creates a default file under the current directory.

Tips

- Always specify full path of the files, as if seen from the managed server. Avoid using `~` to refer to home directory.
- No spaces in path specification, VM name or Image name. Commonly used path are location of kernels, `VM_DISKS`, path to isos, location of Image Store, snapshot directory.
- If you change the location of the kernel in an image, you might want to consider changing the `KERNEL_TYPE` parameter to appropriate value. The `KERNEL_TYPE` is used as a name space for storing cached kernels and ramdisks.
- ConVirt does not allow editing VM config files NOT generated via ConVirt. It looks at special header (first line) to make this determination. If the config files contains simple name value pairs and are not parameterized, they you may selectively choose to add the ConVirt file header.
- Select a locations with lots of disk space for the image store as well as where you store disks for the VMs. Playing around with appliances would definitely require disk space.

Known bugs/caveats

- When using ConVirt in Xen 3.1+/3.0.4 mixes environments, please ensure that ConVirt is run on a 3.1+ node. ConVirt running on a 3.0.4 node isn't forward compatible for managing Xen 3.1+ servers.
- You might have to enter credentials for the same server multiple times when password based authentication scheme is used for managed server as opposed to key based authentication.
For example :
 - When the server is expanded
 - When terminal window is opened for selected VM
 - When an operation requires remote file access and `gnome-vfs` is enabled in ConVirt configuration.
- Sometimes, the guest VM tab showing the virtual machine terminal fails to redraw properly. Clicking on the terminal and sending a `Ctrl+l` and/or some keystrokes often helps.
- When the Virtual machine configuration file is changed, the directives may get saved in different order.
- On some platforms the password field generates `gtk/pango` warnings and displays "?" characters or

blank/nothing as you type. Ignore this, the password that you type in, is read correctly by ConVirt.

Platforms

Here is a list of platforms on which ConVirt is tested.

Client Platforms

Fedora 7, Fedora 8, Ubuntu Gutsy7.10, OpenSUSE 10.3

Server Platforms

Fedora 7, Fedora 8, CentOS/RHEL 5.1, Debian Etch 4, OpenSUSE 10.3

Xen

Xen 3.0.4 / 3.1/3.2

(Xen 3.2 tested on FC8 only)

Licensing

GNU General Public License (GPL)

For details, see: <http://www.gnu.org/licenses/gpl.html>

Feedback

Do drop us a line, if you download/evaluate/use ConVirt. We would appreciate feedback on the current release as well as suggestions for future releases.

Also, we are hoping for active community assistance in the following areas:

- packaging for more platforms and distributions. (e.g. debian, windows, etc.)
- ImageStore images (kernel/ramdisk pairs) for various pre-packaged Guest Os/VM's.

The best way to reach us is to pop in and say hi at our (low frequency) mailing list. To sign up, visit <http://convirt.sf.net/>. We look forward to hearing from you!

Copyright/Acknowledgments

Xen is registered trademark of XenSource Inc
rPath is trademark of rPath.

paramiko library : By Robey Pointer
htmltextview.py : By Gustavo J. A. M. Carneiro
TreeViewToolTip.py : By Daniel J. Popowich
Migration Icon : By Luis Vinay
Appliance Feed Setup : By Hozefa Shiyaji

Appendix

This Appendix contains some notes and useful suggestions to get Xen and ConVirt environment running on different distributions.

I. Cent OS / RHEL 5.1 :

This distribution seems to have Xen 3.1 but some of the API are not available.

To work around this,

- a. use the xen 3.0.3 scripts for preparing server for remote management.
- b. Once prepared as above, you get the following error while discovering the server using Add Node.

```
<Fault 1: 'method "xend.domains_with_state" is not supported'>
```

To fix this :

- Shutdown ConVirt.
- ssh into the Managed Server, open /etc/convirt.conf and add use_3_0_api=True under the default section. If the default section is not there, add it. After the changes it should look like

```
[DEFAULT]
default_computed_options = ['arch', 'arch_libdir', 'device_model']
use_3_0_api=True
```

- Restart ConVirt
- tds

II. Ubuntu 10.2 Gutsy :

The Ubuntu Gutsy repository has Xen available in the repository. To get it,

```
# apt-get install ubuntu-xen-server
```

```
# apt-get install python-paramiko
```

And follow the tarball installation instruction for ConVirt with the following caveats. (till ConVirt becomes available in the repository)

The patch_paramiko script assumes paramiko to be installed under /usr/lib/python directory. This is not true for Ubuntu. These are found unde /var/lib/python-support/python2.5/paramiko. The paramiko version currently in is version 1.6.4. So to apply the patch manually do the following

```
# cd /var/lib/python-support/python2.5/paramiko
# cp packet.py packet.py.orig
# cp /path/to/convirt/patches/paramiko.packet.py.1.6.4
```

To use it as a managed server or do management of the local host on which ConVirt is installed,

1. run the config script script under config-script directory.

```
# ./configure-xend.sh 3.1
```

2. The hvmloder and the qemu-dm path are not as expected. Create the following symlinks to fix the problem.

```
# cd /usr/lib/xen
# ln -s ../xen-ioemu-3.1/boot boot
# cd bin
# ln -s ../../xen-ioemu-3.1/bin/qemu-dm qemu-dm
```

3. Take a look at the [Validate the Bridge Setup](#) under Debian section.

III. Debian 4 Etch with 3.1 Xen as a Managed Node :

Xen 3.1 Setup: The default xen distrobution for the Debian 4 seems to be Xen 3.0.3-1. This being too old, follow the procedure given below to setup Xen 3.1 from binary distributions. This based on an article from howtoforge.

Preparation

```
# apt-get remove exim4 exim4-base lpr nfs-common portmap pidentd pcmcia-
cs pppoe pppoeconf ppp pppconfig
```

```
# apt-get install screen ssh debootstrap python python-twisted iproute
bridge-utils libcurl3-dev libssl0.9.7
```

Download the binary bits and install it. This example shows downloading a PAE xen bits.

```
# cd /usr/src
# wget http://bits.xensource.com/oss-xen/release/3.1.0/bin.tgz/xen-3.1.0-
install-x86_32p.tgz
# tar xvzf xen-3.1.0-install-x86_32.tgz
# cd dist/
# ./install.sh
# mv /lib/tls /lib/tls.disabled
```

Lets create the ramdisk

```
# depmod 2.6.18-xen
# apt-get install yaird
# mkinitramfs -o /boot/initrd.img-2.6.18-xen 2.6.18-xen
```

Add Xen in the start up scripts

```
# update-rc.d defaults xend 20 21
# update-rc.d defaults xendomains 20 21
```

Update the grub file

```
# update-grub
```

Loop devices setup : To use file based disks for virtual machines there should be enough number of /dev/loop devices

Check if the kernel has loop module

```
# lsmod | grep loop
```

If your kernel has loop compiled in, the above command would not show anything. In this case,

Edit the /boot/grub/menu.lst file to add max_loop=64 to the kernel (in a module line). example

```
module /boot/vmlinuz-2.6.18-xen root=/dev/sda7 ro console=tty0
max_loop=64
```

Otherwise you will see a line containing loop with couple of numbers. In this case,

Edit /etc/modules file and add the following line

```
loop max_loop=64
```

Restart the machine. This should show Xen 3.1.0 as one of the menu item.

Configure the managed node : Get the ConVirt tarball (if you are reading this online) and run the following from the config-scripts directory.

```
# ./configure-xend.sh 3.1
```

Validate the Bridge setup :

The VMs use the bridge to communicate to other vm or other servers. This section help validate/setup the bridge.

```
# brctl show
```

This command should show vif0.0 and peth0 in the last column (interfaces). If it is the case, you are done. But, If not, then your VMs may not be able to access the network. Here are couple of suggestions to fix this Add the following to the /etc/network/interfaces file

```
## To use dhcp:
```

```
##
```

```
auto eth0
```

```
iface eth0 inet dhcp
```

Move a couple of udev rules to prevent different eth being created on every boot.

(Some discussion on this [Ethernet numbering](#), [Eth numbering](#))

```
# cd /etc/udev/rules.d
```

```
# mkdir backup
```

```
# mv *_persistent-net.rules backup
# mv *_persistent-net-generator.rules backup
```

Reboot the machine and check `brctl show` output.

IV. OpenSUSE 10.3

The pygrub bootloader is not available on the SUSE platform. Instead, there is a `domUloader.py` located at `/usr/lib/xen/boot` directory. Once the installation of PV images is complete, you would have to change the bootloader manually to `/usr/lib/xen/boot/domUloader.py`. Also, the bootloader requires `bootargs` param containing the location of path to kernel and ramdisk within the VM image. The `bootargs` can be specified in the Miscellaneous tab.

Here is an exmple entry :

```
bootargs = - -entry=xvda1:/vmlinuz,/initrd.img
```

Note: The `domUloader` is not able to mount the rPath Appliance images. This prevents it from being able to fetch the kernel and ramdisk from the appliance and start it. We would publish a set of steps once the `domUloader` issue is resolved.

For SLES environment, you would have to add `use_3_1_api=True` in the DEFAULT section of `convirt.conf` on the managed server machine.

V. Xen 3.2

The Xen 3.2 has been tested on the Fedora 8 platform. One of the thing notably changed is the the name of the bridge. So if you are using the Xen 3.2, change the name of the bridge in the images to `eth0`.

If the managed server is Xen 3.2, use the config scripts under `config-scripts/xen-3.2` directory.