

Xen Cloud Platform Installation Guide

Release 0.1

0.1

Published October 2009

0.1 Edition

Xen Cloud Platform Installation Guide: Release 0.1

Published October 2009

Copyright © 2009 Xen.org

Xen®, Xen.org®, Xen Cloud Platform™, and logos are either registered trademarks or trademarks of Xen.org in the United States and/or other countries. Other company or product names are for informational purposes only and may be trademarks of their respective owners.

This product contains an embodiment of the following patent pending intellectual property of Xen.org:

1. United States Non-Provisional Utility Patent Application Serial Number 11/487,945, filed on July 17, 2006, and entitled "Using Writeable Page Tables for Memory Address Translation in a Hypervisor Environment".
2. United States Non-Provisional Utility Patent Application Serial Number 11/879,338, filed on July 17, 2007, and entitled "Tracking Current Time on Multiprocessor Hosts and Virtual Machines".

Contents

1. Introducing Xen Cloud Platform	5
About this document	5
How this Guide relates to other documentation	5
New features in Xen Cloud Platform 0.1	5
2. System Requirements	7
Xen Cloud Platform host system requirements	7
VM support	8
3. Installing Xen Cloud Platform	9
Installing the Xen Cloud Platform host	9
Installation and deployment scenarios	14
Xen Cloud Platform hosts with local storage	14
Xen Cloud Platform hosts with shared NFS storage	14
Xen Cloud Platform hosts with shared iSCSI storage	16
4. Updating and upgrading Xen Cloud Platform	19
Rolling pool upgrade between major versions	19
Backup	19
Performing a rolling pool upgrade	19
Upgrading a single server	21
Reinstalling the same version	21
Preparing Xen Cloud Platform hosts for maintenance	21
Applying updates	23
Applying updates using the CLI	23
Backing up and restoring Xen Cloud Platform hosts and VMs	24
Backing up Virtual Machine metadata	25
Backing up Xen Cloud Platform hosts	26
Backing up VMs	27
A. Troubleshooting	29
B. PXE installation of Xen Cloud Platform host	31
Setting up the PXE boot environment	31
Creating an answerfile for unattended PXE installation	34
C. Xen Memory Usage	39
Setting Control Domain Memory	39
Index	41

Chapter 1. Introducing Xen Cloud Platform

Thank you for choosing Xen Cloud Platform™ from Xen.org

About this document

This document is an installation guide for Xen Cloud Platform™, the platform virtualization solution from Xen.org®. The Xen Cloud Platform package contains all you need for creating a network of virtual x86 computers running on Xen®, the open-source paravirtualizing hypervisor with near-native performance.

This document contains procedures to guide you through the installation, configuration, and initial operation of Xen Cloud Platform. This document also contains information about troubleshooting problems that might occur during installation, and where to get further information.

How this Guide relates to other documentation

This document is primarily aimed at system administrators who need to set up Xen Cloud Platform hosts on physical servers. Other documentation shipped with this release includes:

- *Xen Cloud Platform Virtual Machine Installation Guide* describes how to install Linux and Windows VMs on top of a Xen Cloud Platform deployment. As well as installing new VMs from install media (or using the VM *templates* provided with the Xen Cloud Platform release), this guide also explains how to create VMs from existing physical machines, using a process called *P2V*.
- *Xen Cloud Platform Administrator's Guide* describes the tasks involved in configuring a Xen Cloud Platform deployment -- how to set up storage, networking and resource pools, and how to administer Xen Cloud Platform hosts using the **xe** command line interface (CLI).
- *Xen Cloud Platform Software Development Kit Guide* presents an overview of the Xen Cloud Platform SDK -- a selection of code samples that demonstrate how to write applications that interface with Xen Cloud Platform hosts.
- *XenAPI Specification* provides a programmer's reference guide to the Xen Cloud Platform API.
- *Release notes* provide a list of known issues that affect this release.

New features in Xen Cloud Platform 0.1

Xen Cloud Platform 0.1 includes a number of new features and ongoing improvements, including:

- *Improved backup and snapshot support* allows you to perform live snapshot and clone operations on all storage types. This enables better support for backup utilities and prac-

tises, and helps to reduce disk space requirements for storing clones of virtual disks. You can now take snapshots from the `xe` CLI. For more information on this feature, please refer to the *Storage* chapter and the *VM Snapshots* section of the *Backup and recovery* chapter in the *Xen Cloud Platform Administrator's Guide*.

- *Active Directory integration* to allow credentials to be verified against an AD server. This allows granting and revocation of access to Xen Cloud Platform pools easily and securely using existing IT infrastructure. For more information on this feature, please refer to the *Xen Cloud Platform hosts and resource pools* chapter in the *Xen Cloud Platform Administrator's Guide*.
- *Workload Balancing* to optimize VM placement and assist with the balancing of workloads within a pool. This is done using the Workload Balancing server, available as a download with Xen Cloud Platform. For customers interested in creating custom queries on Workload Balancing data, Citrix provides customer accessible SQL views. For more information, see [CTX121348](#). For more information on this feature, please refer to the *Workload Balancing* chapter in the *Xen Cloud Platform Administrator's Guide*.
- *Integration with StorageLink services through the xe CLI* enabling advanced storage management capabilities in the Citrix StorageLink service, available as a download with Xen Cloud Platform. For more information on this feature, please refer to the *Citrix StorageLink Gateway (CSLG) SRs* section of the *Storage* chapter in the *Xen Cloud Platform Administrator's Guide*.
- *Improved operating system support* including new support for Red Hat Enterprise Linux 5.3, Novell SLES 11, and Debian Lenny.

Chapter 2. System Requirements

Xen Cloud Platform requires one physical x86 computers to be the Xen Cloud Platform host. The Xen Cloud Platform host machine is dedicated entirely to the task of hosting VMs and is not used for other applications.

Xen Cloud Platform host system requirements

The Xen Cloud Platform host is a 64-bit x86 server-class machine devoted to hosting multiple VMs. This machine runs a stripped-down Linux operating system with a Xen-enabled kernel which controls the interaction between the virtualized devices seen by VMs and the physical hardware.

Xen Cloud Platform can make use of:

- up to 128 GB of RAM
- up to 6 NICs
- up to 32 CPU cores

The following are the system requirements for the Xen Cloud Platform host:

CPUs	<p>One or more 64-bit x86 CPU(s), 1.5 GHz minimum, 2 GHz or faster multicore CPU recommended</p> <p>To support VMs running Windows, an Intel VT or AMD-V 64-bit x86-based system with one or more CPUs is required.</p> <hr/> <p>Note</p> <p>To run Windows VMs, hardware support for virtualization must be enabled on the Xen Cloud Platform host. This is an option in the BIOS. It is possible your BIOS might have virtualization support disabled. Consult your BIOS documentation for more details.</p> <hr/> <p>To support VMs running supported paravirtualized Linux, a standard 64-bit x86-based system with one or more CPUs is required.</p>
RAM	1 GB minimum, 2 GB or more recommended
Disk space	<p>Locally attached storage (PATA, SATA, SCSI) with 16 GB of disk space minimum, 60 GB of disk space recommended</p> <p>General disk space requirements for VMs:</p> <ul style="list-style-type: none">• Product installation creates two 4GB partitions for the Xen Cloud Platform host control domain; remaining space is available for VMs

	<ul style="list-style-type: none"> • VMs based on the Debian templates are allocated a 4GB root device, and a 512MB swap device • Linux VMs are allocated a root device of 8 GB • Windows Vista and Server 2008 VMs are allocated a root device of 24 GB; other Windows VMs default to 8 GB.
Network	100 Mbit/s or faster <i>network interface card</i> (NIC). A gigabit NIC is recommended for faster P2V and export/import data transfers and for live relocation of VMs.

VM support

Windows VMs can be created only on Xen Cloud Platform hosts equipped with Intel VT-enabled or AMD-V CPUs. All Windows VMs are created by installing the operating system from either the Microsoft installation media in the Xen Cloud Platform host physical CD/DVD-ROM drive or a network-accessible ISO image of the appropriate template.

Linux VMs do not require Xen Cloud Platform hosts that are equipped with Intel VT-enabled or AMD-V CPUs.

For a list of supported Windows and Linux distributions, see the *Xen Cloud Platform Virtual Machine Installation Guide*.

Chapter 3. Installing Xen Cloud Platform

This chapter describes how to install Xen Cloud Platform host software on physical servers.

The first sections describe the installation of a Xen Cloud Platform. The following sections describe several common installation and deployment scenarios and provide information that is specific to each scenario.

Installers for the Xen Cloud Platform host are located on the installation media. The installation media also includes:

- a set of Xen Cloud Platform product documents in Adobe Acrobat PDF format
- a P2V tool for creating VM templates from an existing instances of supported Linux distributions running on physical servers. See the *Xen Cloud Platform Virtual Machine Installation Guide* for details.
- a tool for restoring a backed-up Xen Cloud Platform host control domain filesystem. See the section called “Backing up Xen Cloud Platform hosts” for details.

Installing the Xen Cloud Platform host

The Xen Cloud Platform host consists of a Xen-enabled Linux operating system, a management agent, VM templates, and a local storage repository reserved for VMs. The Xen Cloud Platform host must be installed on a *dedicated* 64-bit x86 server.

You can install the Xen Cloud Platform host from the installation CDs or set up a network-accessible *TFTP* server to boot using *PXE*. For details about setting up a TFTP server for PXE-booting the installer, see [Appendix B, PXE installation of Xen Cloud Platform host](#).

Note

Do not install any other operating system in a dual-boot configuration with the Xen Cloud Platform host; this is an unsupported configuration.

The main installation CD contains the basic packages needed to set up the Xen Cloud Platform host on a physical host, and to create Windows VMs using the Windows installation CDs. The Xen Cloud Platform package also contains a separate CD containing support for creating Linux VMs and six CDs containing source code for the included open-source software.

The installer presents an upgrade choice if it is run on a server that already has a previous version of Xen Cloud Platform on it. The upgrade process follows the first-time installation process but several setup steps are bypassed, and the existing settings for networking configuration, system time setting and so on are retained.

If you only want to create Windows VMs, you can install Xen Cloud Platform using just the first CD. If you want to install Linux VMs, be sure to

1. Download the Linux Pack ISO
2. Burn it to a physical CD if installing from a DVD/CD drive, or set it up for PXE installation as described in [Appendix B, PXE installation of Xen Cloud Platform host](#)

Note

If, after installing Xen Cloud Platform without Linux support, you decide later to add it, mount the Linux Pack installation CD or ISO image on the Xen Cloud Platform host and run the script `install.sh`, located in the root directory of the CD.

To install or upgrade the Xen Cloud Platform host

Warning

If you are performing an upgrade, please ensure you do not have any suspended virtual machines as these may not be resumable after the upgrade. Please ensure that all CD drives of virtual machines have been ejected and are empty, and that HA is disabled before proceeding with an upgrade.

1. Boot the computer from the main installation CD, or PXE-boot from your TFTP server if applicable. See [Appendix B, PXE installation of Xen Cloud Platform host](#) for details on how to set up the Xen Cloud Platform media for PXE installation.
2. After the initial boot messages, the installer does some hardware detection and initialization, then presents a screen asking you to select which keyboard keymap you want to use for the installation. In this and the screens that follow, use **Tab** or **Alt+Tab** to move between elements, **Space** to select, and **F12** to move to the next screen.

Select the desired keymap and choose **OK** to proceed.

3. Next, the **Welcome to Xen Cloud Platform** screen is displayed. If you have been provided with any driver disks, select **Driver Disks**. You will then be lead through the process of installing the necessary drivers.

Once you have installed all necessary drivers, select **Install or upgrade Xen Cloud Platform host** and choose **OK** to proceed.

4. The next screen displays a message telling you that the setup program will install Xen Cloud Platform on the computer, and warns that it will overwrite data on any hard drives that you select to use for the installation. Choose **OK** to proceed.
 5. At this point, if the computer on which you are installing the Xen Cloud Platform host does not have a CPU which supports hardware virtualization, or if the support is disabled in the BIOS, a message appears to warn you that you will not be able to run Windows VMs. Choose **OK** to proceed.
-

Warning

Some systems have bugs in their BIOS software which can result in the setting being incorrect. If you get a spurious warning about a lack of hardware virtualization (or do not see a warning when you expected one), hard reboot the host and restart the

installation. You should also check the support site of the hardware manufacturer for BIOS upgrades.

6. If the installer detects a previously-installed version of Xen Cloud Platform host, you are offered the choice to perform a clean installation, or to upgrade the existing version, which preserves any of the VMs present. Select an installation type as appropriate and choose **OK** to proceed.

If you select to upgrade an existing version, you will get a message that the installer is going to create a backup of the existing installation. Choose **Continue** to proceed.

7. If you have multiple local hard disks, you are asked to choose the Primary Disk for the installation. Select the desired disk and choose **OK** to proceed. After selecting the primary one, you are also prompted to choose if you want any of the other drives to be formatted for use by Xen Cloud Platform for VM storage. Format all required drives and choose **OK** to proceed.

If the computer has a single hard disk, these two screens do not appear.

8. The next screen asks you to specify the source of the installation packages. If you are installing off the CD, you will most likely want to select **Local media**. If you are installing using PXE you will most likely want to select **HTTP or FTP** or **NFS**, as appropriate.

If you select **HTTP or FTP** or **NFS**, you are next prompted to set up Networking so that the installation script can connect to the product repository.

If the computer has multiple network interfaces, you are prompted to select one of them to be used to access the Xen Cloud Platform product repository. Select one and choose **OK** to proceed.

If the computer has a single network interface, that interface is used to access the Xen Cloud Platform product repository and no prompt is displayed.

You can select **Automatic configuration (DHCP)** to configure the NIC using DHCP, or **Static configuration**, which prompts you to configure the properties of the NIC manually. Following that, you are prompted to provide the URL or NFS server and path to where the installation media are, as appropriate.

Note

Xen Cloud Platform hosts must have static IP addresses to be part of a resource pool. When using DHCP ensure that a static DHCP reservation policy is in place.

If you selected **Local media**, this networking setup appears later in the installation process.

If you selected **Local media**, the next screen asks if you want to install the Linux Pack from a second CD. If you are planning to install VMs that run Linux operating systems, choose **Yes**. If you are planning to install only Windows VMs, you can choose **No**.

Important

In a pooled setup, the Linux Pack must be installed either on *all* of the pool Xen Cloud Platform hosts, or on *none* of them, so that they are homogeneous.

9. The next screen asks if you want to verify the integrity of the installation media. If you select **Verify installation source**, the MD5 checksum of the packages is calculated and checked against the known value. This may take some time. If you select **Skip verification**, this check is bypassed. Make your selection and choose **OK** to proceed.
10. You are next prompted to set a root password. Enter the desired password and enter it again to verify it.
11. If you elected to perform a clean installation, you are next prompted to set up Networking for the management NIC, which is the interface that will be used to connect to a management console.

If you selected to upgrade an existing installation, the existing management NIC configuration is used and these screens are bypassed.

If the computer has multiple network interfaces, you are prompted to select one of them to be used as the management NIC for the Xen Cloud Platform host software. Select one and choose **OK** to proceed.

If the computer has a single network interface, that interface is used as the management NIC and no prompt is displayed.

Next you can select **Automatic configuration (DHCP)** to configure the NIC using DHCP, or **Static configuration**, which prompts you to configure the NIC's properties manually.

Note

Xen Cloud Platform hosts need to have static IP addresses to be part of a resource pool.

12. If you elected to perform a clean installation, you are next prompted to specify the hostname and the configuration for the name service.

If you elected to upgrade an existing installation, the existing hostname and name service configuration is used and these screens are bypassed.

In the **Hostname Configuration** section, if you select **Automatically set via DHCP**, the DHCP server will provide the hostname along with the IP address. If you select **Manually specify**, enter the desired hostname for the server in the field provided.

In the **DNS Configuration** section, if you select **Manually specify**, enter the IP addresses of your primary (required), secondary (optional), and tertiary (optional) name-servers in the fields provided. Otherwise, choose **Automatically set up via DHCP** to get name service configuration using DHCP.

Select **OK** to proceed.

13. If you elected to perform a clean installation, you are prompted to select the general geographical area for the time zone. Choose from the displayed list of geographical areas, then choose **OK** to proceed. You are then prompted to select the specific locale for the time zone. You can type the first letter of the desired locale to jump to the first entry that begins with this letter. Choose from the displayed list of locales, then choose **OK** to proceed.

If you elected to upgrade an existing installation, the existing time zone and locale is used and these screens are bypassed.

14. If you elected to perform a clean installation, you are prompted to choose a method of setting the system time. You can select **Using NTP** or **Manual time entry**. Make your selection and choose **OK** to proceed.

If you selected to upgrade an existing installation, the existing method of setting system time is used and this screen is not displayed.

15. If you selected **Using NTP** in the preceding step, you are prompted to identify the time server or servers you want to use. You can check **NTP is configured by my DHCP server** and the time server will be set by DHCP. Otherwise, enter at least one NTP server name or IP address in the fields below. Choose **OK** to proceed.

Otherwise, the installation script moves to the next step; you are prompted for the manually-entered time later, near the end of the installation.

Warning

Xen Cloud Platform assumes that the time setting in the BIOS of the server is the current time in UTC, and that the time for the VMs reflects the local time based on the specified time zone offset.

16. A message is displayed that the installation is ready to proceed and that this will format the primary disk and any other disks selected for VM storage, destroying any data that is currently on them. Choose **Install Xen Cloud Platform** to proceed.

A progress bar is displayed as the installation commences. If you chose to set the system date and time manually, a dialog box appears when the progress bar has reached about 90%. Enter the correct numbers in the fields and choose **OK** to proceed.

17. If you are installing from CD and elected to include support for Linux VMs, you will be prompted to put in the Linux Pack disk. Eject the main disk, and insert the Linux Pack disk. Choose **OK**. A screen appears, identifying that this disk contains the Linux Pack. Select **Use media** to proceed with installing it. Another progress bar is displayed, and when it reaches 100%, a completion message is displayed.

If you selected not to install support for Linux VMs, a completion message is displayed.

Note

If you decide later to add Linux support, mount the Linux Pack installation CD or ISO image on the Xen Cloud Platform host and run the script **install.sh**, located in the root of the CD.

18. From the **Installation Complete** screen, eject the installation CD from the drive and select **OK** to reboot the server.

After the server boots, Xen Cloud Platform displays a splash screen, followed by **xs-console**, a system configuration console.

- To access a local shell from **xsconsole**, press **Alt+F3**.
- To return to **xsconsole**, press **Alt+F1**.

Installation and deployment scenarios

This section describes several common installation and deployment scenarios, and details the steps that differ between them:

- one or more Xen Cloud Platform hosts with local storage
- two or more Xen Cloud Platform hosts with shared NFS storage
- two or more Xen Cloud Platform hosts with shared iSCSI storage

Xen Cloud Platform hosts with local storage

The simplest use of Xen Cloud Platform is to set up a simple network of VMs running on one or more Xen Cloud Platform hosts without shared storage. This means that live relocation of VMs from one Xen Cloud Platform host to another is not possible, as this requires shared storage.

Requirements

- one or more 64-bit x86 servers with local storage
- one or more Windows workstations, on the same network as the Xen Cloud Platform hosts

Basic procedure

- Install Xen Cloud Platform host software on the server(s)

Xen Cloud Platform hosts with shared NFS storage

Adding shared storage to the Xen Cloud Platform network enables grouping of Xen Cloud Platform hosts into resource pools, enabling live relocation of VMs and sharing of server resources.

Requirements

- two or more 64-bit x86 servers with local storage
- one or more Windows workstations, on the same network as the Xen Cloud Platform hosts
- a server exporting a shared directory over NFS

Note

To be part of a resource pool, the Xen Cloud Platform hosts and the server or servers providing the shared NFS storage must have static IP addresses.

Basic procedure

1. Install Xen Cloud Platform host software on the server(s)
2. Set up the NFS server
3. Create an SR on the NFS share at the pool level
4. Choose one Xen Cloud Platform host as a pool master and join other Xen Cloud Platform hosts to the pool.

For this procedure, the NFS server is assumed to be running a typical Linux distribution. Consult your Linux distribution documentation for further information.

Setting up an NFS share on an NFS server

1. Check to see if the portmap daemon is installed and running:

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

In the preceding example, runlevels 3, 4, and 5 are set to `on`. That means that at boot, for runlevels 3, 4 and 5, the portmap daemon is started automatically. If either 3, 4 or 5 are set to `off`, set them to `on` by running the following command:

```
chkconfig portmap on
```

2. Check to see if the NFS daemon is installed and running:

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

If either 3, 4 or 5 are set to `off`, set them to `on` by running the following command:

```
chkconfig nfs on
```

3. Create a directory for the shared storage:

```
mkdir /<vm_share_dir>
```

4. Edit the `/etc/exports` file to add the line

```
/<vm_share_dir> *(rw,no_root_squash, sync)
```

Save and close the file.

5. Restart the portmap and nfs daemons as follows:

```
service portmap restart
service nfs restart
```

Creating an SR on the NFS share at the pool level

1. Open a console on any Xen Cloud Platform host in the pool.
2. Create the storage repository on `server:/path`:

```
xe sr-create content-type=user type=nfs name-label=<sr_name=> \
  shared=true device-config:server=<server> \
  device-config:serverpath=<path>
```

The `device-config-server` refers to the hostname of the NFS server and `device-config-serverpath` refers to the path on the server. Since `shared` is set to true, the shared storage is automatically connected to every host in the pool and any hosts that subsequently join are also connected to the storage. The UUID of the created storage repository is printed to the console.

3. Find the UUID of the pool

```
xe pool-list
```

4. Set the shared storage as the pool-wide default

```
xe pool-param-set uuid=<pool_uuid> \
  default-SR=<storage_repository_uuid>
```

Since the shared storage has been set as the pool-wide default, all future VMs will have their disks created on this shared storage by default.

Xen Cloud Platform hosts with shared iSCSI storage

Adding shared storage to the Xen Cloud Platform network enables grouping of Xen Cloud Platform hosts into resource pools, enabling live relocation of VMs and sharing of server resources.

Requirements

- two or more 64-bit x86 servers with local storage
- one or more Windows workstations, on the same network as the Xen Cloud Platform hosts
- a server providing a shared directory over iSCSI

Note

To be part of a resource pool, the Xen Cloud Platform hosts and the server or servers providing the shared iSCSI storage must have static IP addresses.

Basic procedure

1. Install Xen Cloud Platform host software on the server(s)
2. Prepare the iSCSI storage
3. If necessary, enable multiple initiators on your iSCSI device
4. Choose one Xen Cloud Platform host as a pool master and join other Xen Cloud Platform hosts to the pool
5. Configure the iSCSI IQN for each Xen Cloud Platform host
6. Create an SR on the iSCSI share at the pool level

The details of how to set up iSCSI storage differ between the various iSCSI solutions on the market. In general, though, you need to provide an iSCSI target on the SAN for the VM storage, and then configure Xen Cloud Platform hosts to be able to see and connect to it. This is done by providing a valid iSCSI Qualified Name (IQN) to the iSCSI target and to the iSCSI initiator on each Xen Cloud Platform host.

Prepare the iSCSI storage

1. Assign a virtual storage volume on the iSCSI device for VM storage
2. Create IQNs on the SAN for each Xen Cloud Platform host that will use the storage

You can use the CLI to configure the IQN for each Xen Cloud Platform host and to create the SR.

To configure the iSCSI IQN for each Xen Cloud Platform host using the CLI

1. On the host console, run the command:

```
xe-set-iscsi-iqn <iscsi_iqn>
```

Alternatively, used the CLI directly:

```
xe host-param-set uuid=<host_uuid> other-config-iscsi_iqn=<iscsi_iqn>
```

2. Repeat for each Xen Cloud Platform host in the pool.

To create an SR on the iSCSI share at the pool level using the CLI

1. On the console of any server in the pool, run the command:

```
xe sr-create name-label=<name_for_sr> \  
content-type=user device-config-target=<iscsi_server_ip_address> \  
device-config-targetIQN=<iscsi_target_iqn> \  
device-config-localIQN=<iscsi_local_iqn> \  
type=lvmoiscsi shared=true device-config-LUNid=<lun_id>
```

The *device-config-target* argument refers to the hostname or IP address of the iSCSI server. The *device-config-LUNid* argument can be a comma-separated list of LUN IDs. Since the *shared* argument is set to *true*, the shared storage is automatically connected to every host in the pool and any hosts that subsequently join are also connected to the storage.

The command returns the UUID of the created storage repository.

2. Find the UUID of the pool by running the command:

```
xe pool-list
```

3. Set the shared storage as the pool-wide default as follows:

```
xe pool-param-set uuid=<pool_uuid> default-SR=<iscsi_shared_sr_uuid>
```

Now that the shared storage has been set as the pool-wide default, all future VMs will have their disks created on shared storage by default.

Chapter 4. Updating and upgrading Xen Cloud Platform

This chapter documents how to update (apply minor update patches) or upgrade from an earlier version.

Rolling pool upgrade between major versions

Xen Cloud Platform allows you to upgrade a pool of Xen Cloud Platform hosts to the next major version, while keeping VMs on that pool running and thus avoids service downtime. This is achieved by upgrading on a host-by-host basis, with only one Xen Cloud Platform host offline at a time.

You can use the command line interface to migrate VMs running on a Xen Cloud Platform host running an older version of the product to one running either the same version or higher. It is not possible to migrate VMs located on a Xen Cloud Platform host with a newer Xen Cloud Platform version to one running an older version.

You should plan your upgrade path carefully. Xen.org strongly advise against running a mixed-mode pool (one with multiple versions of Xen Cloud Platform co-existing) for longer than necessary. This is because the pool will be operating in a degraded state during the upgrade: all VMs will continue to function as normal, but control operations other than migration might not be available. Operations such as `vm-copy`, `vm-start`, and `vm-export` are not available. In particular, it is not safe to perform storage-related operations such as adding, removing or resizing virtual disks in this mode.

Backup

Xen.org strongly recommends that you take a backup of the state of your existing pool using the **pool-dump-database** CLI command (see the *Xen Cloud Platform Administrator's Guide*). This allows you revert a partially complete rolling upgrade back to its original state without losing any VM data.

Note

Because it is not possible to migrate a VM from an upgraded Xen Cloud Platform host to a Xen Cloud Platform host running an older version of Xen Cloud Platform, it may be necessary to shut down VMs if you need to revert the rolling upgrade for any reason.

Performing a rolling pool upgrade

Warning

Disable HA functionality before performing a rolling upgrade.

The correct sequence for upgrading a pool of Xen Cloud Platform installations to a newer version is as follows:

1. Eject any CDs from your Virtual Machines before starting the rolling upgrade. Having CDs inserted during rolling upgrade can prevent migrations from working correctly, and due to the mode of operation of the pool while the rolling upgrade is taking place, it is required that this be done before the rolling upgrade is started.
2. Verify that there are no VMs in the *Suspended* state. Any suspended VM with a CD drive attached (with the Tools ISO or a physical CD in the local physical drive, for example) cannot be resumed after performing an upgrade. To get a suspended VM back into a usable state, you have to perform a **Force Shutdown** on it and then restart it.
3. Migrate all VMs running on the pool master to other Xen Cloud Platform hosts using XenMotion.
4. Shut down the pool master using the CLI. This will cause your pool to enter *emergency mode*. VMs will continue to run, but you will be unable to perform control operations. This is expected behavior.
5. Boot the pool master using your Xen Cloud Platform installation media or network and follow the instructions for doing an upgrade (see [Chapter 3, Installing Xen Cloud Platform](#)).
6. On restarting your pool master, your pool will leave emergency mode and normal service will be restored after a few minutes.
7. You are now ready to upgrade a second Xen Cloud Platform host. You should select a Xen Cloud Platform host still running an old version of Xen Cloud Platform and migrate the VMs running on this Xen Cloud Platform host to the one you have just upgraded. Do not attempt to migrate a VM from an upgraded Xen Cloud Platform host to one that has not yet been upgraded. You will see an error message if you attempt to do this, and your VM will continue running without being migrated.
8. Upgrade the member Xen Cloud Platform host you have just freed up (following a similar procedure as for the master); shut down the host using the CLI (your pool will not enter emergency mode this time); then upgrade the server software using your product media or remote installation repository.
9. Repeat the previous two steps for each member Xen Cloud Platform host in the pool.
10. Now that you have upgraded the Xen Cloud Platform host software on your pool, it is important to upgrade the Xen.org Tools for Virtual Machines on each VM. This will enable new functionality and ensure the stability of your VMs. Running old versions of Xen.org Tools for Virtual Machines on newer Xen Cloud Platform installations is not a supported configuration except during the upgrade process. Please refer to the Xen-Center Help, or the *Xen Cloud Platform Virtual Machine Installation Guide* for details on how to perform the upgrade of Xen.org Tools for Virtual Machines for Windows and Linux VMs.

Note

In the unlikely event that a host fails (hardware failure, for example) during the rolling upgrade process, it is necessary to use the **xe host-forget** command to forget the host.

Failure to do so will result in Xen Cloud Platform remaining stuck in rolling upgrade mode indefinitely.

Upgrading a single server

Upgrading a single server is a simple process. Simply shut down any running VMs, shut down the server, and follow the normal installation procedures. The installer will offer an option to upgrade your existing installation – be sure to select this option otherwise you will lose any existing data. Check that the final confirmation screen displays has 'upgrade' mentioned. You will not be required to re-enter any settings during the setup procedure.

Reinstalling the same version

The following procedure describes how to "freshen," or reinstall the current version of the Xen Cloud Platform host over an existing installation of Xen Cloud Platform host 0.1, and preserve settings on VMs.

Warning

When reinstalling your host, be aware that

- any custom RPMs which you might have installed on the Xen Cloud Platform host control domain will *not* be preserved
 - if you have installed any Xen Cloud Platform hotfixes on the server, do *not* reinstall the original version; this is not supported
-

Reinstalling Xen Cloud Platform host from version 0.1

1. Perform an orderly shutdown on the VMs hosted on the Xen Cloud Platform host. If any of your VMs are in the suspended state, resume them first, and then perform an orderly shutdown on them too.

To shut down all the VMs automatically, you can also run the **service xapi-domains stop** command on the control domain terminal.

2. Reboot the Xen Cloud Platform host, and boot from the Installation CD.
3. The installation script identifies the version and prompts you whether you want to re-install over the existing installation and preserve VMs. Select **OK** to proceed with the installation.

Preparing Xen Cloud Platform hosts for maintenance

Before performing maintenance operations on a Xen Cloud Platform host that is part of a resource pool, you should disable it (which prevents any VMs from being started on it), then migrate its VMs to another Xen Cloud Platform host in the pool. This can most readily be accomplished by placing the Xen Cloud Platform host into Maintenance mode using CLI.

Warning

Xen.org highly recommends rebooting all Xen Cloud Platforms prior to installing an update and then verifying their configuration. This is because some configuration changes only take effect when a Xen Cloud Platform is rebooted, so the reboot may uncover configuration problems that would cause the update to fail.

To prepare a Xen Cloud Platform host in a pool for maintenance operations using the CLI

1. Run the command

```
xe host-disable uuid=<xenserver_host_uuid>
xe host-evacuate uuid=<xenserver_host_uuid>
```

This will disable the Xen Cloud Platform host and then migrate any running VMs to other Xen Cloud Platform hosts in the pool.

2. Perform the desired maintenance operation.
3. Once the maintenance operation is completed, enable the Xen Cloud Platform host:

```
xe host-enable
```

Restart any halted VMs and/or resume any suspended VMs.

Before performing maintenance operations on a Xen Cloud Platform host that is not part of a resource pool, you should disable it (which prevents any VMs from being started on it), and then either shutdown or suspend its VMs.

Warning

Any suspended VM with a CD drive attached (with the Tools ISO or a physical CD in the local physical drive, for example) cannot be resumed after performing an upgrade. To get a suspended VM back into a usable state, perform a **Force Shutdown** on it and then restart it.

To prepare an unpooled Xen Cloud Platform host for upgrade using the CLI

1. Disable the Xen Cloud Platform host:

```
xe host-disable
```

2. Shut down or suspend any running VMs using the **xe vm-shutdown** or **xe vm-suspend** commands. If you choose to suspend any VMs, check that no CDs are attached to them.
3. Perform the desired maintenance operation.
4. Once the maintenance operation is completed, enable the Xen Cloud Platform host:

```
xe host-enable
```

Restart any halted VMs and/or resume any suspended VMs.

Applying updates

Between releases of Xen Cloud Platform software, Xen.org occasionally releases updates to the software. These updates typically contain accumulated bug fixes and feature improvements. When an update is released, it is made accessible on the internet and an email announcement is sent to all Xen Cloud Platform customers.

Once downloaded apply the updates using the CLI.

Updates sometimes have special steps that have to be performed after the update is applied, such as requiring that the control domain is restarted. Whenever possible, updates will be such that they can be applied without interruption, but in some cases they might require Xen Cloud Platform host or VM restarts to be performed. In cases where a Xen Cloud Platform host restart is required, you can avoid downtime of virtual machines in a pooled environment by applying the update to each server in turn, migrating VMs away from each server in turn as the update is applied. If you are using the CLI, you will have to do this manually using the **host-evacuate** command.

If you are using the CLI to perform the update, Xen Cloud Platform hosts to be updated should be prepared for this operation by performing the procedures in [the section called "Preparing Xen Cloud Platform hosts for maintenance"](#).

Applying updates using the CLI

First, the update must be uploaded to the pool or server to which it will be applied. This will cause a UUID (identifier) to be assigned to the update, and information about which servers it has been applied to will be tracked. Once an update has been uploaded to a pool or server, you can use the **patch-list** and **patch-param-list** commands to view information about the update. The second stage is to apply the update. Xen.org recommends that the **patch-pool-apply** command is used to do this; this will result in the update being applied on all servers in the pool. Alternatively, the **patch-apply** command may be used to apply the update to one server in a pool - this may be useful when applying the update and then restarting individual servers in the pool. Pools should not be left in an inconsistent update state (one where updates have been installed on some servers and not others).

Discussion of procedures using the CLI below assume a basic knowledge of the usage of the **xe** tool. For information about this, please see the *Xen Cloud Platform Administrator's Guide*.

The update procedure is basically the same for both a single server and pool scenario, except that in a pooled scenario you must ensure that the update is applied to all servers in the pool. This will be achieved either by using the **patch-pool-apply** command, or by executing the **patch-apply** once for each host. These methods are described below.

Applying an update to a single Xen Cloud Platform host or a pool using the CLI:

1. Download the update to a local directory. Note the path to the update file you have downloaded. (It is also possible to download the update directly to an appropriate location on the server, e.g. `/root`, using standard Linux commands, but it is usually best to download it to a remote client first.)
2. Upload the update to your server or pool:

```
xe -s <my_server> -u <user_name> -pw <password> patch-upload \  
file-name=<update_file>
```

Here, the **-s** **-u**, and **-pw** options refer to the server, the username (which would usually be root), and the password, as usual - these would be omitted if running the command directly from a command shell on the Xen Cloud Platform host local console.

Once you have executed the above command, you will be given the UUID of the uploaded update. This UUID will be used to specify the update that is to be applied.

3. Be sure to follow any guidance regarding the update before continuing, in particular any information provided about whether VMs should be moved away from the server or that the server should be restarted after applying the update. As always, we recommend taking appropriate backup measures before making modifications to system software. To automatically move VMs to other servers, you can use the **host-evacuate** CLI command.
4. Apply the update to the pool. A command like the following may be used to do this:

```
xe patch-pool-apply uuid=<b89249c7-feba-41c5-8838-911ded969add>
```

This applies the update to all servers in the pool. Alternatively, if you need to restart servers and perform the update in a rolling manner, you can apply the update to an individual server by running a command like the following:

```
xe patch-apply host-uuid=<ebf17583-d8c5-4835-999a-e0a29716207d> \  
uuid=<b89249c7-feba-41c5-8838-911ded969add>
```

5. Verify that the update was applied by using the **patch-list** command again. Now the *hosts* field should contain the host UUID.

After an update is applied to a Xen Cloud Platform host, a small file containing the same information stored on the master from the **xe patch-upload** command is written to a sub-directory of the machine's patch directory. This enables Xen Cloud Platform hosts later ejected from the pool to repopulate their databases with information about updates already applied.

To reclaim space on the master, large updates can be deleted from disk using the **xe patch-clean** command. (The update information stored in the database of the master, though, is always retained.) These updates can be uploaded again using **xe patch-upload** if required.

Backing up and restoring Xen Cloud Platform hosts and VMs

Xen.org recommends that, whenever possible, you leave the installed state of Xen Cloud Platform hosts unaltered. That is, do not install any additional packages or start additional services on Xen Cloud Platform hosts, and treat them as if they are appliances. The best way to restore, then, is to re-install Xen Cloud Platform host software from the installation media. If you have multiple Xen Cloud Platform hosts, the best approach is to configure a PXE boot server and appropriate answerfiles for this purpose (see [Appendix B, PXE installation of Xen Cloud Platform host](#)).

For VMs, the best approach is to install backup agents on them, just as if they were standard physical servers. For Windows VMs, as of this release we have tested CA [BrightStor ARCserve Backup](#), and Symantec [NetBackup](#) and [Backup Exec](#).

For more information about backup tools tested, best practices, and backups in general, see the [Xen.org Knowledge Base](#).

Backing up Virtual Machine metadata

Xen Cloud Platform hosts use a database on each host to store metadata about VMs and associated resources such as storage and networking. When combined with storage repositories, this database forms the complete view of all VMs available across the pool. Therefore it is important to understand how to backup this database in order to recover from physical hardware failure and other disaster scenarios.

This section first describes how to backup metadata for single-host installations, and then for more complex pool setups.

Backing up single host installations

Use the CLI to backup the pool database. To obtain a consistent pool metadata backup file, run **pool-dump-database** on the Xen Cloud Platform host and archive the resulting file. The backup file will contain sensitive authentication information about the pool, so ensure it is securely stored.

To restore the pool database, use the **xe pool-restore-database** command from a previous dump file. If your Xen Cloud Platform host has died completely, then you must first do a fresh install, and then run the **pool-restore-database** command against the freshly installed Xen Cloud Platform host.

After a restoration of the pool database, some VMs may still be registered as being `Suspended`, but if the storage repository with their suspended memory state (defined in the `suspend-VDI-uuid` field) was a local SR, it will no longer be available since the host has been reinstalled. To reset these VMs back to the `Halted` state so that they can be started up again, use the **xe vm-shutdown vm=vm_name -force** command, or use the **xe vm-reset-powerstate vm=<vm_name> -force** command.

Warning

Xen Cloud Platform hosts restored using this method will have their UUIDs preserved. If you restore to a different physical machine while the original Xen Cloud Platform host is still running, there will be duplicate UUIDs. Pool database backup is not the recommended mechanism for cloning physical hosts; use the automated installation support for that (see [Appendix B, PXE installation of Xen Cloud Platform host](#)).

Backing up pooled installations

In a pool scenario, the master host provides an authoritative database that is synchronously mirrored to all the member hosts in the pool. This provides a degree of built-in redundancy to a pool; the master can be replaced by any member since each of them have an accurate

version of the pool database. Please refer to the *Xen Cloud Platform Administrator's Guide* for more information on how to transition a member into becoming a master host.

This level of protection may not be sufficient; for example, if your shared storage containing the VM data is backed up in multiple sites, but your local server storage (containing the pool metadata) is not. To fully recreate a pool given just a set of shared storage, you must first backup the **pool-dump-database** file on the master host, and archive this file.

Subsequently restoring this backup on a brand new set of hosts

1. Install a fresh set of Xen Cloud Platform hosts from the installation media, or over PXE.
2. Use the **xe pool-restore-database** on the host designated to be the new master.
3. Run the **xe host-forget** command on the new master to remove the old member machines.
4. Use the **xe pool-join** command on the member hosts to connect them to the new pool.

See the section called “Pool failures” in *Xen Cloud Platform Administrator's Guide* for specific restoration scenarios.

Backing up Xen Cloud Platform hosts

This section describes the Xen Cloud Platform host control domain backup and restore procedures. These procedures do *not* back up the storage repositories that house the VMs, but only the privileged control domain that runs Xen and the Xen Cloud Platform agent.

Note

Because the privileged control domain is best left as installed, without customizing it with other packages, Xen.org recommends that you set up a PXE boot environment to cleanly perform a fresh installation from the Xen Cloud Platform media as a recovery strategy. In many cases you will not need to backup the control domain at all, but just save the pool metadata (see [the section called “Backing up Virtual Machine metadata”](#)). This backup method should always be considered complementary to backing up the pool metadata.

Another approach is to run the Xen Cloud Platform installation twice, selecting to back up the existing installation when prompted. This will create a pristine copy of the freshly-installed control domain that can later be restored if necessary by using the installation CD and choosing the Restore option.

Using the **xe** commands **host-backup** and **host-restore** is another approach that you can take. The **xe host-backup** command archives the active partition to a file you specify, and the **xe host-restore** command extracts an archive created by **xe host-backup** over the currently inactive disk partition of the host. This partition can then be made active by booting off the installation CD and choosing the **Restore** option.

After completing the above steps and rebooting the host, you must ensure that the VM meta-data is restored to a consistent state. This can be achieved by running **xe pool-restore-database** on `/var/backup/pool-database- $\{DATE\}$` . This file is created by **xe host-backup** using **xe pool-dump-database** command before archiving the running filesystem, to snapshot a consistent state of the VM metadata.

To back up a Xen Cloud Platform host

- On a remote host with enough disk space, run the command:

```
xe host-backup file-name=<filename> -h <hostname> -u root -pw <password>
```

This creates a compressed image of the control domain file system in the location specified by the *file-name* argument.

To restore a running Xen Cloud Platform host

1. If you want to restore a Xen Cloud Platform host from a specific backup, run the following command while the Xen Cloud Platform host is up and reachable:

```
xe host-restore file-name=<filename> -h <hostname> -u root -pw <password>;
```

This command restores the compressed image back to the hard disk of the Xen Cloud Platform host on which the command is run (not the host on which *filename* resides). In this context “restore” is something of a misnomer, as the word usually suggests that the backed-up state has been put fully in place. The restore command here only unpacks the compressed backup file and restores it to its normal form, but it is written to another partition (*/dev/sda2*) and does *not* overwrite the current version of the filesystem.

2. To use the restored version of the root filesystem, reboot the Xen Cloud Platform host using the Xen Cloud Platform installation CD and select the **Restore from backup** option.

After the restore from backup is completed, reboot the Xen Cloud Platform host and it will start up from the restored image.

Finally, restore the VM meta-data using

```
xe pool-restore-database file-name=/var/backup/pool-database-*
```

Note

Restoring from a backup as described here does *not* destroy the backup partition.

Restarting a crashed Xen Cloud Platform host

1. If your Xen Cloud Platform host has crashed and is not reachable anymore, you need to use the Xen Cloud Platform installation CD to do an upgrade install. When that is completed, reboot the machine and make sure your host is reachable remote CLI.
2. Then proceed with the section called “Backing up Xen Cloud Platform hosts” above.

Backing up VMs

VMs are best backed up using standard backup tools running on them individually. For Windows VMs, we have tested CA BrightStor ARCserve Backup.

Appendix A. Troubleshooting

If you experience odd behavior, crashes, or have other issues during installation, this chapter is meant to help you solve the problem if possible and, failing that, describes where logs are located and other information that can help the Xen.org community track and resolve the issue.

Xen.org provides several forms of support: you can receive free self-help support information on the [Support site](#)

The Xen Cloud Platform host installation CD runs Linux, so most standard Linux commands can be used to diagnose installation problems. There are three virtual terminals available during installation, which display the installation menu, an interactive console and an event log, respectively. Use the **ALT + F1-F3** keys to toggle between the virtual terminals.

You can check some basic things in the interactive terminal:

- **fdisk** lists all disks that can be seen as a result of the loaded storage device drivers. If a particular device driver did not load, for example, the driver for a RAID card, then the disks attached to that card will not appear in the output from the **fdisk** command.
- **ifconfig** shows the network configuration of physical NICs, including their IP addresses, netmasks, and gateway.
- **ping** can be used to verify network connectivity from the Xen Cloud Platform host to a remote IP address and vice-versa.

Installation logs are written to `/install/tmp/`

Appendix B. PXE installation of Xen Cloud Platform host

This appendix describes how to set up a TFTP server to enable PXE booting of Xen Cloud Platform host installations. It also describes how to create an XML answerfile, which allows you to perform unattended installations.

Setting up the PXE boot environment

To create a PXE environment, you need:

- a TFTP server to enable PXE booting
- a DHCP server to provide IP addresses to the systems that are going to PXE-boot
- an NFS, FTP, or HTTP server to house the installation files

These can all co-exist on the same server, or be distributed on different servers on the network.

Additionally, each system that you want to PXE boot and install Xen Cloud Platform on needs a PXE boot-enabled Ethernet card.

The following steps assume that the Linux server or servers you will use have RPM support.

To set up a TFTP server for PXE booting

1. TFTP requires SYSLINUX 3.11 or above. SYSLINUX is a collection of boot loaders for the Linux operating system which operates on Linux EXT2/EXT3 file systems, MS-DOS FAT file systems, network servers using PXE firmware, and CD-ROMs. Make sure you have SYSLINUX version 3.11 or above installed on your system by running the command

```
rpm -q syslinux
```

If you have an earlier version, you can download an appropriate later version from <ftp://ftp.kernel.org/pub/linux/utils/boot/syslinux/RPMS/i386/>, then install it by running the command

```
rpm -Uvh syslinux.-.rpm
```

2. Check if the tftp server package is installed:

```
rpm -q tftp-server
```

If not, use system-config-packages and install.

3. Edit the file `/etc/xinetd.d/tftp` to change the line

```
disable = yes
```

to

```
disable = no
```

- Restart the xinetd service, which manages tftp:

```
service xinetd restart
```

- Make a directory inside `/tftpboot` called `xenserver`.
- Copy the files `mboot.c32` and `pxelinux.0` files from `/usr/lib/syslinux` to the `/tftpboot` directory.
- Copy the files `install.img`, `vmlinuz`, and `xen.gz` from the Base Pack CD (found in the root of the Base Pack CD, and in its `/boot` directory respectively), and place them in the `/tftpboot/xenserver` directory.
- Create a directory called `pxelinux.cfg` inside `/tftpboot` and create a file named `default`. The file contents depend on how you want to configure your PXE boot environment. For example, you might have a configuration file like the following:

Note

The backslashes at the ends of lines in the example PXE configuration files shown below denote continuation of lines; do *not* actually include them in your PXE configuration file.

The three hyphens in the examples are necessary parts of the `mboot.c32` loader syntax; not including them will cause PXE boot attempts to fail.

```
default xenserver
label xenserver
kernel mboot.c32
append /tftpboot/xenserver/xen.gz dom0_mem=752M com1=115200,8n1 \
console=com1,tty --- /tftpboot/xenserver/vmlinuz \
xencons=hvc console=hvc0 console=tty0 \
--- /tftpboot/xenserver/install.img
```

This will start an installation on any machine that boots from this server. Someone would then need to manually respond to the prompts to complete the installation. Alternatively, you might have a configuration file like the following:

```
default xenserver-auto
label xenserver-auto
kernel mboot.c32
append /tftpboot/xenserver/xen.gz dom0_mem=752M com1=115200,8n1 \
console=com1,tty --- /tftpboot/xenserver/vmlinuz \
xencons=hvc console=hvc0 console=tty0 \
answerfile=http://pxehost.example.com/5.0.0-answerfile \
install --- /tftpboot/xenserver/install.img
```

This will perform an unattended installation using the answerfile at the URL specified.

Note

The above examples show how to configure the installer to run on the physical console, `tty0`. To use a different default ensure that the console you want to use is the rightmost.

For details on creating an answerfile for unattended installation, see [the section called “Creating an answerfile for unattended PXE installation”](#). For more information on PXE configuration file syntax, see [the SYSLINUX website](#).

Please refer to your server operating system manual for details for your specific operating system. The information here is a guide that can be used for Red Hat, Fedora, and some other RPM-based distributions.

To set up a DHCP server

1. On the server that you will be using for DHCP, check if you have DHCP installed by running the command:

```
rpm -qa dhcp
```

If not, install it using **system-config-packages**.

2. Configure the DHCP server. Refer to [document 1673 in the Red Hat Knowledge base](#) for details.
3. Add these lines to the end of the existing `dhcpd.conf` file:

```
allow booting;
allow bootp;
class "pxeclients" {
    match if substring(option vendor-class-identifier, 0, 9) = "PXEClient";
    next-server ;
    filename "pxelinux.0";
}
```

4. Restart the `dhcpd` service:

```
service dhcpd restart
```

Setting up the installation media host

1. On the server where you are going to house the installation files, copy the contents of the packages directories from the base pack CD to a location where they are exported by HTTP, FTP, or NFS. For example, you might make a directory in the document root of a webserver named `Xen Cloud Platform_0.1` and then copy the directory `packages.main` from the Base Pack disk to `Xen Cloud Platform_0.1/packages.main`.
2. If Linux support is also desired, copy `packages.linux` from the Linux pack disk to `Xen Cloud Platform_0.1/packages.linux`. This structure allows you to install either both packages by having the `source` element of the answerfile contain the enclosing directory `Xen Cloud Platform_0.1`, or you can install the base

pack only (no support for Linux VMs) by putting in the path to Xen Cloud Platform_0.1/packages.main.

For example, to install both packages from the webserver `http://pxehost.example.com` where the packages are in the directories mentioned above relative to the server's document root, the answerfile would contain this source element:

```
<source type="url">http://pxehost.example.com/XenServer_5.0.0</source>
```

or, to install just the base pack:

```
<source type="url">  
http://pxehost.example.com/XenServer_&PRODUCT_VERSION;/packages.main  
</source>
```

In either case you can specify a username and password if required. For example:

```
<source type="url">  
http://<username>:<password>@pxehost.example.com/XenServer_0.1/packages.main  
</source>
```

To make driver disks available during installation copy the contents of each driver disk ISO into a separate directory in the main installation repository. For example, copy the files in `helloworld.iso` to `Xenserver-<version>/driver.helloworld`. For unattended install of driver disks create an answerfile as described in the section called “Creating an answerfile for unattended PXE installation”.

Preparing the destination system

1. Start the system and enter the Boot Menu (**F12** in most BIOSes) and select to boot from your Ethernet card.
2. The system should then PXE boot from the installation source you set up, and the installation script will commence. If you have set up an answerfile, the installation can proceed unattended.

Creating an answerfile for unattended PXE installation

To perform installations in an unattended fashion, you need to create an XML answerfile. Here is an example answerfile:

```

<?xml version="1.0"?>
  <installation>
    <primary-disk>sda</primary-disk>
    <guest-disk>sdb</guest-disk>
    <guest-disk>sdc</guest-disk>
    <keymap>us</keymap>
    <root-password>mypassword</root-password>
    <source type="url">http://pxehost.example.com</source>
    <post-install-script type="url">
      http://pxehost.example.com/myscripts/post-install-script
    </post-install-script>
    <admin-interface name="eth0" proto="dhcp" />
    <timezone>Europe/London</timezone>
  </installation>

```

All nodes should be within a root node named *installation*.

The following is a summary of the elements. All values should be text in the nodes, unless otherwise stated. Required elements are indicated.

Element	Description	Re-quired?
<primary-disk>	<p>The name of the storage device where the control domain should be installed, equivalent to the choice made on the <i>Select Primary Disk</i> step of the interactive installation process.</p> <p>Attributes:</p> <p>You can specify a <i>gueststorage</i> attribute with possible values <i>yes</i> and <i>no</i>. For example:</p> <pre><primary-disk gueststorage="no">sda</primary-disk></pre> <p>If this attribute is not specified, the default is <i>yes</i>. If you specify <i>no</i>, it is possible to automate an installation scenario where no storage repository is created, if, in addition, no <i>guest-disk</i> keys are specified.</p>	Y
<guest-disk>	<p>The name of a storage device to be used for storing guests. You should use one of these elements for each extra disk.</p>	N
<keymap>	<p>The name of the keymap to use during installation.</p> <pre><keymap>us</keymap></pre>	Y
<root-password>	<p>The desired root password for the Xen Cloud Platform host.</p>	Y

Element	Description	Re-quired?
<source>	<p>Where the packages should be installed from.</p> <p>Attributes:</p> <p><i>type: url, nfs, or local</i></p> <p>If local, leave the element empty. For example,</p> <pre data-bbox="592 625 1180 844"> <source type="url"> http://server/packages </source> <source type="local" /> <source type="nfs"> server:packages </source> </pre>	Y
<driver-source>	<p>Where the device driver packages should be installed from. Optional. Element may occur multiple times.</p> <p>Attributes:</p> <p><i>type: url, nfs, or local</i></p> <p>If local, leave the element empty. For example,</p> <pre data-bbox="592 1163 1180 1381"> <driver-source type="url"> http://server/drivers </driver-source> <driver-source type="local" /> <driver-source type="nfs"> server:drivers </driver-source> </pre>	N
<post-install-script>	<p>Where the post-install-script is.</p> <p>Attributes:</p> <p><i>type: url, nfs, or local</i></p> <p>If <i>url</i> or <i>nfs</i>, put the url or NFS path in the PCDATA; if <i>local</i>, leave the PCDATA empty. For example,</p>	Y

Element	Description	Re-quired?
	<pre> <source type="url"> http://server/scripts </source> <source type="local" /> <source type="nfs"> server:scripts </source> </pre>	
<admin-interface>	<p>The single network interface to be used as the host administration interface.</p> <p>Attributes:</p> <p>proto: <i>dhcp</i> or <i>static</i></p> <p>name: <i>eth0</i> for example.</p> <p>Children:</p> <ul style="list-style-type: none"> • <ip>: The IP address, if proto="static" • <subnet-mask>: The subnet mask, if proto="static" • <gateway>: The gateway, if proto="static" <p>All three child elements are <i>required</i> if proto="static"</p>	N
<timezone>	In the format used by the TZ variable, e.g. Europe/London, or America/Los_Angeles.	Y
<nameserver>	The IP address of a nameserver. You should use one of these elements for each nameserver you want to use.	N
<hostname>	Specify if you want to manually set a hostname.	N
<bootloader>	Specify which bootloader to install for startup time. Only change this if you have problems booting. Currently either <i>extlinux</i> (the default) or <i>grub</i> .	N

You can also perform automated upgrades by changing the answerfile appropriately. Set the mode attribute of the installation element to *reinstall*, specify the disk on which the existing installation lives with the *existing-installation* element, and leave the *primary-disk* and *guest-disk* elements unspecified. For example:

```
<?xml version="1.0"?>
  <installation mode="reinstall">
    <existing-installation>sda</existing-installation>
    <keymap>us</keymap>
    <root-password>mypassword</root-password>
    <source type="url">http://pxehost.example.com</source>
    <post-install-script type="url">
      http://pxehost.example.com/myscripts/post-install-script
    </post-install-script>
    <admin-interface name="eth0" proto="dhcp" />
    <timezone>Europe/London</timezone>
  </installation>
```

Appendix C. Xen Memory Usage

When calculating the memory footprint of a Xen host there are two components that must be taken into consideration. First there is the memory consumed by the Xen hypervisor itself; then there is the memory consumed by the host's *control domain*. The control domain is a privileged VM that provides low-level services to other VMs, such as providing access to physical devices. It also runs the management tool stack.

Setting Control Domain Memory

If your control domain requires more allocated memory, this can be set using the Xen CLI.

Use the **xe vm-memory-target-set** command to set the amount of memory available to the control domain.

The **xe vm-memory-target-wait** command can be used to check if the control domain has managed to attain the requested memory target specified at the last use of the **xe vm-memory-target-set** command. **xe vm-memory-target-wait** will not return until the memory target has been reached, or will time out if the target cannot be reached, for example when the target is lower than the actual memory requirements of the VM.

The following fields on a VM define how much memory will be allocated. The default values shown are indicative of a machine with 8 GiB of RAM:

name	default	description
memory-actual	411041792	The actual amount of memory current available for use by the VM
memory-target	411041792	The target amount of memory as set by using xe vm-memory-target-set
memory-static-max	790102016	The maximum possible physical memory
memory-dynamic-max	790102016	The desired maximum memory to be made available
memory-dynamic-min	306184192	The desired minimum memory to be made available
memory-static-min	306184192	The minimum possible physical memory

Dynamic memory values must be within the boundaries set by the static memory values. Additionally the memory target must fall in the range between the dynamic memory values.

Note

The amount of memory reported in may exceed the values set using this mechanism. This is because the amount reported includes the memory used by the control domain, the hypervisor itself, and the crash kernel. The amount of memory used by the hypervisor will be larger for hosts with more memory.

To find out how much host memory is actually available to be assigned to VMs, get the value of the *memory-free* field of the host, and then use the **vm-compute-maximum-memory** command to get the actual amount of free memory that can be allocated to the VM:

```
xe host-list uuid=<host_uuid> params=memory-free
xe vm-compute-maximum-memory vm=<vm_name> total=<host_memory_free_value>
```

Index

A

AMD-V (AMD hardware virtualization), 7, 8

B

Backup/restore

 pooled installations, 25

 VMs, 27

 Xen Cloud Platform host control domain, 9

 Xen Cloud Platform host metadata, 24, 25

BIOS, enabling VT support in, 7

I

Installation and deployment scenarios

 two or more Xen Cloud Platform hosts with shared iSCSI storage, 16

 two or more Xen Cloud Platform hosts with shared NFS storage, 14, 14

Installation media contents,

M

Maintenance procedures

 Applying updates, 23

 Backing up a Xen Cloud Platform host, 27

 Backing up and restoring Xen Cloud Platform host control domain, 26

 Reinstalling on top of the same version, 21

 Restart a crashed Xen Cloud Platform host, 27

 Restore a running Xen Cloud Platform host, 27

P

P2V (physical to virtual conversion), 9

PXE installation

 answerfile, for unattended installation,

 installation media repository, 33

R

Rolling upgrade between major versions, 19

S

Setting up a DHCP server, 33

Setting up a TFTP installation server, 31

Setting up installation media host for PXE, 33

T

Troubleshooting, 29

U

Updating, 23

Upgrading
rolling upgrade between major versions, 19

X

Xen
memory usage, 39
Xen Cloud Platform host
installation, 9
Management NIC, 12
requirements, for Windows VMs, 7
system requirements, 7