

Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Hypervisor Configuration Guide

Revised on July 31, 2015 11:00 am IST



Citrix CloudPlatform

Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Hypervisor Configuration Guide

Revised on July 31, 2015 11:00 am IST

Author

Citrix CloudPlatform

© 2014 Citrix Systems, Inc. All rights reserved. Specifications are subject to change without notice. Citrix Systems, Inc., the Citrix logo, Citrix XenServer, Citrix XenCenter, and CloudPlatform are trademarks or registered trademarks of Citrix Systems, Inc. All other brands or products are trademarks or registered trademarks of their respective holders.

If you have already installed CloudPlatform and you want to learn about configuring hypervisors that CloudPlatform supports, read this document.

1. About this Guide	1
1.1. About the Audience for this Guide	1
1.2. Using the Product Documentation	1
1.3. Experimental Features	1
1.4. Additional Information and Help	1
1.5. Contacting Support	1
2. General Requirements for Hypervisor Hosts	3
3. Installing XenServer for CloudPlatform	5
3.1. System Requirements for XenServer Hosts	5
3.2. XenServer Installation Steps	6
3.3. Configure XenServer dom0 Memory	6
3.4. Username and Password	6
3.5. Time Synchronization	6
3.6. Licensing	7
3.6.1. Getting and Deploying a License	7
3.7. Install CloudPlatform XenServer Support Package (CSP)	7
3.8. Primary Storage Setup for XenServer	8
3.9. iSCSI Multipath Setup for XenServer (Optional)	10
3.10. Physical Networking Setup for XenServer	10
3.10.1. Configuring Public Network with a Dedicated NIC (Optional)	10
3.10.2. Using NIC Bonding in CloudPlatform (Optional)	11
3.10.3. Configuring Multiple Guest Networks for XenServer (Optional)	13
3.10.4. Separate Storage Network for XenServer (Optional)	13
4. Installing Hyper-V for CloudPlatform	15
4.1. System Requirements for Hyper-V Hypervisor Hosts	15
4.1.1. Supported Operating Systems for Hyper-V Hosts	15
4.1.2. Minimum System Requirements for Hyper-V Hosts	15
4.1.3. Supported Storage	15
4.2. Preparation Checklist for Hyper-V	15
4.3. Hyper-V Installation Steps	17
4.4. Installing the CloudPlatform Role on a Hyper-V Host	18
4.5. Physical Network Configuration for Hyper-V	20
4.6. Storage Preparation for Hyper-V (Optional)	20
5. Installing KVM for CloudPlatform	21
5.1. System Requirements for KVM Hypervisor Hosts	21
5.1.1. Supported Operating Systems for KVM Hosts	21
5.1.2. System Requirements for KVM Hosts	21
5.2. Install and configure the Agent	22
5.3. Installing the CloudPlatform Agent on a KVM Host	22
5.4. Physical Network Configuration for KVM	23
5.5. Time Synchronization for KVM Hosts	24
5.6. Primary Storage Setup for KVM (Optional)	24
6. Installing VMware vSphere for CloudPlatform	27
6.1. System Requirements for vSphere Hosts	27
6.1.1. Software requirements	27
6.1.2. Hardware requirements	27
6.1.3. vCenter Server requirements:	28
6.1.4. Other requirements:	28
6.2. Preparation Checklist for VMware	29
6.2.1. vCenter Checklist	29
6.2.2. Networking Checklist for VMware	29

6.3. vSphere Installation Steps	30
6.4. ESXi Host setup	30
6.5. Physical Host Networking	30
6.5.1. Configure Virtual Switch	30
6.5.2. Configure vCenter Management Network	32
6.5.3. Configure NIC Bonding for vSphere	32
6.6. Configuring a vSphere Cluster with Nexus 1000v Virtual Switch	32
6.6.1. About Cisco Nexus 1000v Distributed Virtual Switch	32
6.6.2. Prerequisites and Guidelines	33
6.6.3. Nexus 1000v Virtual Switch Preconfiguration	33
6.6.4. Enabling Nexus Virtual Switch in CloudPlatform	37
6.6.5. Configuring Nexus 1000v Virtual Switch in CloudPlatform	37
6.6.6. Removing Nexus Virtual Switch	38
6.6.7. Configuring a VMware Datacenter with VMware Distributed Virtual Switch	38
6.7. Storage Preparation for vSphere (iSCSI only)	49
6.7.1. Enable iSCSI initiator for ESXi hosts	49
6.7.2. Add iSCSI target	49
6.7.3. Create an iSCSI datastore	49
6.7.4. Multipathing for vSphere (Optional)	50
6.8. Add Hosts or Configure Clusters (vSphere)	50
6.9. Creating Custom Roles in vCenter for CloudPlatform	50
6.9.1. System Requirements	50
6.9.2. Minimum Permissions	50
6.9.3. Creating Roles	50
7. (Experimental Feature) Installing LXC for CloudPlatform	53
7.1. System Requirements for LXC Hosts	53
7.1.1. Software Requirements	53
7.1.2. Hardware Requirements	53
7.2. LXC Installation Overview	53
7.2.1. LXC Installation Considerations	54
7.3. Preparing the Operating System	54
7.4. Installing and Configuring Libvirt	54
7.5. Installing and Configuring the LXC Agent	55
7.6. Configuring Network Bridges	56
7.6.1. Network example	56
7.6.2. Configuring Network Bridges on RHEL	57
7.7. (Optional)Primary Storage Setup for LXC	59
8. Installing Baremetal for CloudPlatform	61
8.1. Baremetal Host System Requirements	61
8.2. About Baremetal Kickstart Installation	61
8.2.1. Limitations of Baremetal Installation	62
8.2.2. Prerequisites for Baremetal Host with Kickstart	62
8.2.3. Example CentOS 6.x Kickstart File	68
8.2.4. Example Fedora 17 Kickstart File	69
8.2.5. Example Ubuntu 12.04 Kickstart File	70
8.3. Using Cisco UCS as a Bare Metal Host	72
8.3.1. Limitation on Using UCS Manager Profile Templates	72
8.3.2. Registering a UCS Manager	73
8.3.3. Associating a Profile with a UCS Blade	73
8.3.4. Disassociating a Profile from a UCS Blade	74
8.3.5. Synchronizing UCS Manager Changes with CloudPlatform	75
Index	77

About this Guide

1.1. About the Audience for this Guide

This guide is meant for anyone responsible for configuring and administering hypervisors on CloudPlatform such as cloud administrators and Information Technology (IT) administrators.

1.2. Using the Product Documentation

The following guides provide information about CloudPlatform:

- *Citrix CloudPlatform (powered by Apache CloudStack) Installation Guide*
- *Citrix CloudPlatform (powered by Apache CloudStack) Concepts Guide*
- *Citrix CloudPlatform (powered by Apache CloudStack) Getting Started Guide*
- *Citrix CloudPlatform (powered by Apache CloudStack) Administration Guide*
- *Citrix CloudPlatform (powered by Apache CloudStack) Hypervisor Configuration Guide*
- *Citrix CloudPlatform (powered by Apache CloudStack) Developer's Guide*

For complete information on any known limitations or issues in this release, see the *Citrix CloudPlatform (powered by Apache CloudStack) Release Notes*.

For information about the Application Programming Interfaces (APIs) that is used in this product, see the API documents that are available with CloudPlatform.

1.3. Experimental Features

CloudPlatform product releases include some experimental features for customers to test and experiment with in non-production environments, and share any feedback with Citrix. For any issues with these experimental features, customers can open a support ticket but Citrix cannot commit to debugging or providing fixes for them.

The following experimental features are included in this release:

- Linux Containers
- Supported Management Server OS and Supported Hypervisors: RHEL7/CentOS 7 for experimental use with Linux Containers

1.4. Additional Information and Help

Information on accessing Citrix Knowledge Center and about contacting technical support.

1.5. Contacting Support

The support team is available to help customers plan and execute their installations. To contact the support team, log in to the support portal at support.citrix.com/cloudsupport¹ by using the account credentials you received when you purchased your support contract.

¹ <http://support.citrix.com/cloudsupport>

General Requirements for Hypervisor Hosts

The host is a guest virtual machines where the cloud services run. Each host must support the following requirements:

- Support for HVM (Intel-VT or AMD-V enabled).
- CPU - 64-bit x86.
- Hardware virtualization.
- Memory - 4 GB.
- Hard disk - 36 GB of local disk.
- One NIC with static IP
- Hypervisor software with the latest hotfixes applied.
- No VMs running on the hypervisor at the time of deployment.
- Homogenous hosts in a cluster. The CPUs must contain the same type, count, and the feature flags.

Hosts have additional requirements depending on the hypervisor. For more information, refer to the requirements documented at the top of each hypervisor configuration chapter in this guide.

Installing XenServer for CloudPlatform

If you want to use the Citrix XenServer hypervisor to run guest virtual machines, install XenServer on the host(s) in your cloud. For an initial installation, follow the steps below.

3.1. System Requirements for XenServer Hosts

- The following versions of XenServer are supported:
 - XenServer 6.5 SP1
 - XenServer 6.5
 - XenServer version 6.2 SPI with Hotfix XS62ESP1004 and beyond
- The host must be certified as compatible with the XenServer version you are using. See the Citrix Hardware Compatibility Guide: <http://hcl.xensource.com>
- You must re-install XenServer if you are going to re-use a host from a previous install.
- Must support HVM (Intel-VT or AMD-V enabled)
- Be sure all the hotfixes provided by the hypervisor vendor are applied. Apply patches as soon as possible after they are released. It is essential that your hosts are completely up to date with the provided hypervisor patches.
- All hosts within a cluster must be homogenous. The CPUs must be of the same type, count, and feature flags.
- Must support HVM (Intel-VT or AMD-V enabled in BIOS)
- 64-bit x86 CPU (more cores results in better performance)
- Hardware virtualization support required
- 4 GB of memory
- 36 GB of local disk
- At least 1 NIC
- Statically allocated IP Address
- When you deploy CloudPlatform, the hypervisor host must not have any VMs already running



Warning

The lack of up-to-date hotfixes can lead to data corruption and lost VMs.

3.2. XenServer Installation Steps

1. From <https://www.citrix.com/English/ss/downloads/>, download the appropriate version of XenServer for your CloudPlatform version (see [Section 3.1, "System Requirements for XenServer Hosts"](#)). Install it using the Citrix XenServer Installation Guide.
2. After installation, perform the following configuration steps, which are described in the next few sections:

Required	Optional
Section 3.3, "Configure XenServer dom0 Memory"	Section 3.7, "Install CloudPlatform XenServer Support Package (CSP)"
Section 3.4, "Username and Password"	Set up SR if not using NFS, iSCSI, or local disk; see Section 3.8, "Primary Storage Setup for XenServer"
Section 3.5, "Time Synchronization"	Section 3.9, "iSCSI Multipath Setup for XenServer (Optional)"
Section 3.6, "Licensing"	Section 3.10, "Physical Networking Setup for XenServer"

3.3. Configure XenServer dom0 Memory



Note

The following configuration applies to the versions lower than XenServer 6.2.0

Configure the XenServer dom0 settings to allocate more memory to dom0. This can enable XenServer to handle larger numbers of virtual machines. We recommend 2940 MB of RAM for XenServer dom0. For instructions on how to do this, see <http://support.citrix.com/article/CTX126531>. The article refers to XenServer 5.6, but the same information applies to XenServer 6.0.

3.4. Username and Password

All XenServers in a cluster must have the same username and password as configured in CloudPlatform.

3.5. Time Synchronization

The host must be set to use NTP. All hosts in a pod must have the same time.

1. Install NTP.

```
# yum install ntp
```

2. Edit the NTP configuration file to point to your NTP server.

```
# vi /etc/ntp.conf
```

Add one or more server lines in this file with the names of the NTP servers you want to use. For example:

```
server 0.xenserver.pool.ntp.org
server 1.xenserver.pool.ntp.org
server 2.xenserver.pool.ntp.org
server 3.xenserver.pool.ntp.org
```

3. Restart the NTP client.

```
# service ntpd restart
```

4. Make sure NTP will start again upon reboot.

```
# chkconfig ntpd on
```

3.6. Licensing

Citrix XenServer Free version provides 30 days usage without a license. Following the 30 day trial, XenServer requires a free activation and license. You can choose to install a license now or skip this step. If you skip this step, you will need to install a license when you activate and license the XenServer.

3.6.1. Getting and Deploying a License

If you choose to install a license now you will need to use the XenCenter to activate and get a license.

1. In XenCenter, click Tools > License manager.
2. Select your XenServer and select Activate Free XenServer.
3. Request a license.

You can install the license with XenCenter or using the xe command line tool.

3.7. Install CloudPlatform XenServer Support Package (CSP)

To enable security groups, elastic load balancing, and elastic IP on XenServer, download and install the CloudPlatform XenServer Support Package (CSP). After installing XenServer, perform the following additional steps on each XenServer host.



Note

CSP must be installed on XenServer host before the host can be added to a basic zone.

1. If you are using a version prior to XenServer 6.1, perform the following to get the CSP packages. Beginning with XenServer 6.1, the CSP packages are available by default, so you can skip to the next step if you are using one of these more recent versions.

- a. Download the CSP software onto the XenServer host from one of the following link:

For XenServer 6.0.2:

<http://download.cloud.com/releases/3.0.1/XS-6.0.2/xenserver-cloud-supp.tgz>

- b. Extract the file:

```
# tar xf xenserver-cloud-supp.tgz
```

- c. Run the following script:

```
# xe-install-supplemental-pack xenserver-cloud-supp.iso
```

2. If the XenServer host is part of a zone that uses basic networking, disable Open vSwitch (OVS):

```
# xe-switch-network-backend bridge
```

Restart the host machine when prompted.

3. If you are using XenServer 6.1 or greater, perform the following:

- a. Run the following commands:

```
# echo 1 > /proc/sys/net/bridge/bridge-nf-call-iptables
# echo 1 > /proc/sys/net/bridge/bridge-nf-call-arptables
```

- b. To persist the above changes across reboots, set the following values in the `/etc/sysctl.conf` file to 1:

```
net.bridge.bridge-nf-call-iptables = 1
net.bridge.bridge-nf-call-arptables = 1
```

- c. Run the following command:

```
# sysctl -p /etc/sysctl.conf
```

The XenServer host is now ready to be added to CloudPlatform.

3.8. Primary Storage Setup for XenServer

CloudPlatform natively supports NFS, iSCSI and local storage. If you are using one of these storage types, there is no need to create the XenServer Storage Repository ("SR").

If, however, you would like to use storage connected via some other technology, such as FiberChannel, you must set up the SR yourself. To do so, perform the following steps. If you have your hosts in a XenServer pool, perform the steps on the master node. If you are working with a single XenServer which is not part of a cluster, perform the steps on that XenServer.

1. Connect FiberChannel cable to all hosts in the cluster and to the FiberChannel storage host.
2. Rescan the SCSI bus. Either use the following command or use XenCenter to perform an HBA rescan.

```
# scsi-rescan
```

3. Repeat step 2 on every host.
4. Check to be sure you see the new SCSI disk.

```
# ls /dev/disk/by-id/scsi-360a98000503365344e6f6177615a516b -l
```

The output should look like this, although the specific file name will be different (scsi-<scsiID>):

```
lrwxrwxrwx 1 root root 9 Mar 16 13:47
/dev/disk/by-id/scsi-360a98000503365344e6f6177615a516b -> ../../sdc
```

5. Repeat step 4 on every host.
6. On the XenServer host, run this command to get a unique ID for the new SR.

```
# uuidgen
```

The output should look like this, although the specific ID will be different:

```
e6849e96-86c3-4f2c-8fcc-350cc711be3d
```

7. Create the FiberChannel SR. In name-label, use the unique ID you just generated.

```
# xe sr-create type=lvmohba shared=true
device-config:SCSIid=360a98000503365344e6f6177615a516b
name-label="e6849e96-86c3-4f2c-8fcc-350cc711be3d"
```

This command returns a unique ID for the SR, like the following example (your ID will be different):

```
7a143820-e893-6c6a-236e-472da6ee66bf
```

8. To create a human-readable description for the SR, use the following command. In uuid, use the SR ID returned by the previous command. In name-description, set whatever friendly text you prefer.

```
# xe sr-param-set uuid=7a143820-e893-6c6a-236e-472da6ee66bf name-description="Fiber
Channel storage repository"
```

Make note of the values you will need when you add this storage to CloudPlatform later . In the Add Primary Storage dialog, in Protocol, you will choose PreSetup. In SR Name-Label, you will enter the name-label you set earlier (in this example, e6849e96-86c3-4f2c-8fcc-350cc711be3d).

9. (Optional) If you want to enable multipath I/O on a FiberChannel SAN, refer to the documentation provided by the SAN vendor.

3.9. iSCSI Multipath Setup for XenServer (Optional)

When setting up the storage repository on a Citrix XenServer, you can enable multipath I/O, which uses redundant physical components to provide greater reliability in the connection between the server and the SAN. To enable multipathing, use a SAN solution that is supported for Citrix servers and follow the procedures in Citrix documentation. The following links provide a starting point:

- <http://support.citrix.com/article/CTX118791>
- <http://support.citrix.com/article/CTX125403>

You can also ask your SAN vendor for advice about setting up your Citrix repository for multipathing.

Make note of the values you will need when you add this storage to the CloudPlatform later. In the Add Primary Storage dialog, in Protocol, you will choose PreSetup. In SR Name-Label, you will enter the same name used to create the SR.

If you encounter difficulty, address the support team for the SAN provided by your vendor. If they are not able to solve your issue, see Contacting Support.

3.10. Physical Networking Setup for XenServer

Once XenServer has been installed, you may need to do some additional network configuration. At this point in the installation, you should have a plan for what NICs the host will have and what traffic each NIC will carry. The NICs should be cabled as necessary to implement your plan.

If you plan on using NIC bonding, the NICs on all hosts in the cluster must be cabled exactly the same. For example, if eth0 is in the private bond on one host in a cluster, then eth0 must be in the private bond on all hosts in the cluster.

The IP address assigned for the management network interface must be static. It can be set on the host itself or obtained via static DHCP.

CloudPlatform configures network traffic of various types to use different NICs or bonds on the XenServer host. You can control this process and provide input to the Management Server through the use of XenServer network name labels. The name labels are placed on physical interfaces or bonds and configured in CloudPlatform. In some simple cases the name labels are not required.

For information on configuring network and bonds on XenServer, see the following:

- [Network Interface Card Bonds in XenServer¹](#)
- [XenServer Active/Active Bonding -- Switch Configuration²](#)
- [How to Configure Network and Bonds for XenServer³](#)

3.10.1. Configuring Public Network with a Dedicated NIC (Optional)

CloudPlatform supports the use of a second NIC or bonded pair of NICs, as described in [Section 3.10.2, "Using NIC Bonding in CloudPlatform \(Optional\)"](#), for the public network. If bonding is not used, the public network can be on any NIC and can be on different NICs on the hosts in a cluster. For example, the public network can be on eth0 on node A and eth1 on node B. However, the XenServer name-label for the public network must be identical across all hosts. The following

¹ <http://support.citrix.com/article/CTX137599>

² <http://support.citrix.com/article/CTX132559>

³ <http://support.citrix.com/article/CTX132002>

examples set the network label to "cloud-public". After the management server is installed and running you must configure it with the name of the chosen network label.

If you are using two NICs bonded together to create a public network, see [Section 3.10.2, "Using NIC Bonding in CloudPlatform \(Optional\)"](#).

If you are using a single dedicated NIC to provide public network access, follow this procedure on each new host that is added to CloudPlatform before adding the host.

1. Run `xe network-list` and find the public network.

This is usually attached to the NIC that is public. Once you find the network make note of its UUID. Call this <UUID-Public>.

2. Run the following command.

```
# xe network-param-set name-label=cloud-public uuid=<UUID-Public>
```

3.10.2. Using NIC Bonding in CloudPlatform (Optional)

XenServer supports Source Level Balancing (SLB) NIC bonding. Two NICs can be bonded together to carry public, private, and guest traffic, or some combination of these. Separate storage networks are also possible. Here are some example supported configurations:

- 2 NICs on private, 2 NICs on public, 2 NICs on storage
- 2 NICs on private, 1 NIC on public, storage uses management network
- 2 NICs on private, 2 NICs on public, storage uses management network
- 1 NIC for private, public, and storage

All NIC bonding is optional.

XenServer expects all nodes in a cluster will have the same network cabling and same bonds implemented. In an installation the master will be the first host that was added to the cluster and the slave hosts will be all subsequent hosts added to the cluster. The bonds present on the master set the expectation for hosts added to the cluster later. The procedure to set up bonds on the master and slaves are different, and are described below. There are several important implications of this:

- You must set bonds on the first host added to a cluster. Then you must use `xe` commands as below to establish the same bonds in the second and subsequent hosts added to a cluster.
- Slave hosts in a cluster must be cabled exactly the same as the master. For example, if `eth0` is in the private bond on the master, it must be in the management network for added slave hosts.

3.10.2.1. Management Network Bonding

The administrator must bond the management network NICs prior to adding the host to CloudPlatform.

3.10.2.2. Creating a Private Bond on the First Host in the Cluster

Use the following steps to create a bond in XenServer. These steps should be run on only the first host in a cluster. This example creates the cloud-private network with two physical NICs (`eth0` and `eth1`) bonded into it.

1. Find the physical NICs that you want to bond together.

```
# xe pif-list host-name-label='hostname' device=eth0
# xe pif-list host-name-label='hostname' device=eth1
```

These command shows the eth0 and eth1 NICs and their UUIDs. Substitute the ethX devices of your choice. Call the UUID's returned by the above command slave1-UUID and slave2-UUID.

2. Create a new network for the bond. For example, a new network with name "cloud-private".

This label is important. CloudPlatform looks for a network by a name you configure. You must use the same name-label for all hosts in the cloud for the management network.

```
# xe network-create name-label=cloud-private
# xe bond-create network-uuid=[uuid of cloud-private created above]
pif-uuids=[slave1-uuid],[slave2-uuid]
```

Now you have a bonded pair that can be recognized by CloudPlatform as the management network.

3.10.2.3. Public Network Bonding

Bonding can be implemented on a separate, public network. The administrator is responsible for creating a bond for the public network if that network will be bonded and will be separate from the management network.

3.10.2.4. Creating a Public Bond on the First Host in the Cluster

These steps should be run on only the first host in a cluster. This example creates the cloud-public network with two physical NICs (eth2 and eth3) bonded into it.

1. Find the physical NICs that you want to bond together.

```
# xe pif-list host-name-label='hostname' device=eth2
# xe pif-list host-name-label='hostname' device=eth3
```

These command shows the eth2 and eth3 NICs and their UUIDs. Substitute the ethX devices of your choice. Call the UUID's returned by the above command slave1-UUID and slave2-UUID.

2. Create a new network for the bond. For example, a new network with name "cloud-public".

This label is important. CloudPlatform looks for a network by a name you configure. You must use the same name-label for all hosts in the cloud for the public network.

```
# xe network-create name-label=cloud-public
# xe bond-create network-uuid=[uuid of cloud-public created above]
pif-uuids=[slave1-uuid],[slave2-uuid]
```

Now you have a bonded pair that can be recognized by CloudPlatform as the public network.

3.10.2.5. Adding More Hosts to the Cluster

With the bonds (if any) established on the master, you should add additional, slave hosts. Run the following command for all additional hosts to be added to the cluster. This will cause the host to join the master in a single XenServer pool.

```
# xe pool-join master-address=[master IP] master-username=root
```

```
master-password=[your password]
```

3.10.2.6. Complete the Bonding Setup Across the Cluster

With all hosts added to the pool, run the `cloud-setup-bonding.sh` script. This script will complete the configuration and set up of the bonds across all hosts in the cluster.

1. Copy the script from the Management Server in `/usr/share/cloudstack-common/scripts/vm/hypervisor/xenserver/` to the master host and ensure it is executable.
2. Run the script:

```
# ./cloud-setup-bonding.sh
```

Now the bonds are set up and configured properly across the cluster.

3.10.3. Configuring Multiple Guest Networks for XenServer (Optional)

CloudPlatform supports the use of multiple guest networks with the XenServer hypervisor. Each network is assigned a name-label in XenServer. For example, you might have two networks with the labels "cloud-guest" and "cloud-guest2". After the management server is installed and running, you must add the networks and use these labels so that CloudPlatform is aware of the networks.

Follow this procedure on each new host before adding the host to CloudPlatform:

1. Run `xe network-list` and find one of the guest networks. Once you find the network make note of its UUID. Call this <UUID-Guest>.
2. Run the following command, substituting your own name-label and uuid values.

```
# xe network-param-set name-label=<cloud-guestN> uuid=<UUID-Guest>
```

3. Repeat these steps for each additional guest network, using a different name-label and uuid each time.

3.10.4. Separate Storage Network for XenServer (Optional)

You can optionally set up a separate storage network. This should be done first on the host, before implementing the bonding steps below. This can be done using one or two available NICs. With two NICs bonding may be done as above. It is the administrator's responsibility to set up a separate storage network.

Give the storage network a different name-label than what will be given for other networks.

For the separate storage network to work correctly, it must be the only interface that can ping the primary storage device's IP address. For example, if `eth0` is the management network NIC, `ping -I eth0 <primary storage device IP>` must fail. In all deployments, secondary storage devices must be pingable from the management network NIC or bond. If a secondary storage device has been placed on the storage network, it must also be pingable via the storage network NIC or bond on the hosts as well.

You can set up two separate storage networks as well. For example, if you intend to implement iSCSI multipath, dedicate two non-bonded NICs to multipath. Each of the two networks needs a unique name-label.

Chapter 3. Installing XenServer for CloudPlatform

If no bonding is done, the administrator must set up and name-label the separate storage network on all hosts (masters and slaves).

Here is an example to set up eth5 to access a storage network on 172.16.0.0/24.

```
# xe pif-list host-name-label='hostname' device=eth5
uuid(RO): ab0d3dd4-5744-8fae-9693-a022c7a3471d
device ( RO): eth5
#xe pif-reconfigure-ip DNS=172.16.3.3 gateway=172.16.0.1 IP=172.16.0.55 mode=static
netmask=255.255.255.0 uuid=ab0d3dd4-5744-8fae-9693-a022c7a3471d
```

Installing Hyper-V for CloudPlatform

If you want to use Hyper-V hypervisor to run guest virtual machines, install Hyper-V on the hosts in your cloud. The instructions in this section doesn't duplicate Hyper-V Installation documentation. It provides the CloudPlatform-specific steps that are needed to prepare a Hyper-V host to work with CloudPlatform.

4.1. System Requirements for Hyper-V Hypervisor Hosts

4.1.1. Supported Operating Systems for Hyper-V Hosts

- Windows Server 2012 R2 Standard
- Windows Server 2012 R2 Datacenter
- Hyper-V 2012 R2

4.1.2. Minimum System Requirements for Hyper-V Hosts

- 1.4 GHz 64-bit processor with hardware-assisted virtualization.
- 800 MB of RAM
- 32 GB of disk space
- Gigabit (10/100/1000baseT) Ethernet adapter

4.1.3. Supported Storage

- Primary Storage: Server Message Block (SMB) Version 3, Local

SMB Version 3 is available on Windows 2012 and Windows 2012 R2.

- Secondary Storage: SMB

4.2. Preparation Checklist for Hyper-V

For a smoother installation, gather the following information before you start:

Hyper-V Requirements	Value	Description
Server Roles	Hyper-V	After the Windows Server 2012 R2 installation, ensure that Hyper-V is selected from Server Roles. For more information, see Installing Hyper-V¹ .
Share Location	New folders in the /Share directory	Ensure that folders are created for Primary and Secondary storage. The SMB share and

¹ http://technet.microsoft.com/en-us/library/jj134187.aspx#BKMK_Step2

Hyper-V Requirements	Value	Description
		<p>the hosts should be part of the same domain.</p> <p>If you are using Windows SMB share, the location of the file share for the Hyper-V deployment will be the new folder created in the \Shares on the selected volume. You can create sub-folders for both CloudPlatform Primary and Secondary storage within the share location. When you select the profile for the file shares, ensure that you select SMB Share -Applications. This creates the file shares with settings appropriate for Hyper-V.</p>
Domain and Hosts		<p>Hosts should be part of the same Active Directory domain.</p> <p>You must install Hyper-V role on the designated host that you want to use as Hyper-V hypervisor host.</p>
Hyper-V Users	Full control	Full control on the SMB file share.
Virtual Switch		<p>If you are using Hyper-V 2012 R2, manually create an external virtual switch before adding the host to CloudPlatform. If the Hyper-V host is added to the Hyper-V manager, select the host, then click Virtual Switch Manager, then New Virtual Switch. In the External Network, select the desired NIC adapter and click Apply.</p> <p>If you are using Windows 2012 R2, virtual switch is created automatically.</p>
Virtual Switch Name		Take a note of the name of the virtual switch. You need to specify that when configuring CloudPlatform physical network labels.

Hyper-V Requirements	Value	Description
Hyper-V Domain Users		<ul style="list-style-type: none"> • Add the Hyper-V domain users to the Hyper-V Administrators group. • A domain user should have full control on the SMB share that is exported for primary and secondary storage. • This domain user should be part of the Hyper-V Administrators and Local Administrators group on the Hyper-V hosts that are to be managed by CloudPlatform. • The Hyper-V Agent service runs with the credentials of this domain user account. • Specify the credential of the domain user while adding a host to CloudPlatform so that it can manage it. • Specify the credential of the domain user while adding a shared SMB primary or secondary storage.
Migration	Migration	Enable Migration.
Migration	Delegation	If you want to use Live Migration, enable Delegation. Enable the following services of other hosts participating in Live Migration: CIFS and Microsoft Virtual System Migration Service.
Migration	Kerberos	Enable Kerberos for Live Migration.
Network Access Permission for Dial-in	Allow access	Allow access for Dial-in connections.

4.3. Hyper-V Installation Steps

1. Download the operating system from [Windows Server 2012 R2²](http://technet.microsoft.com/en-us/windowsserver/hh534429).
2. Install it on the host as given in [Install and Deploy Windows Server 2012 R2³](http://technet.microsoft.com/library/hh831620).

² <http://technet.microsoft.com/en-us/windowsserver/hh534429>

³ <http://technet.microsoft.com/library/hh831620>

3. Post installation, ensure that you enable Hyper-V role in the server.
4. If no Active Directory domain exists in your deployment, create one and add users to the domain.
5. In the Active Directory domain, ensure that all the Hyper-v hosts are added so that all the hosts are part of the domain.
6. Add the domain user to the following groups on the Hyper-V host: Hyper-V Administrators and Local Administrators.
7. After installation, perform the following configuration tasks, which are described in the next few sections.

Required	Optional
Section 4.4, “Installing the CloudPlatform Role on a Hyper-V Host ”	Section 4.6, “Storage Preparation for Hyper-V (Optional) ”
Section 4.5, “Physical Network Configuration for Hyper-V ”	

4.4. Installing the CloudPlatform Role on a Hyper-V Host

The CloudStack Hyper-V Agent helps CloudPlatform perform operations on the Hyper-V hosts. The CloudStack Hyper-V Agent communicates with the Management Server and controls all the instances on the host. Each Hyper-V host must have the Hyper-V Agent installed on it for successful interaction between the host and CloudPlatform. The Hyper-V Agent runs as a Windows service. For event logs, see Applications in Windows Logs on the host machine. Install the Agent on each host using the following steps.

CloudPlatform Management Server communicates with Hyper-V Agent by using HTTPS. For secure communication between the Management Server and the host, install a self-signed certificate on port 8250.

Prerequisite:

The domain user should be provided with the Log on as a Service permissions before installing the agent. To do that, Open Local Security Policy, select Local Policies, then select User Rights Assignment, and in Logon As a Service add the domain users.



Note

The Agent installer automatically perform this operation. You have not selected this option during the Agent installation, it can also be done manually as given in step [a](#).

You must first create and add a self signed certificate to the port before you install CloudPlatform agent on a Hyper-V host. The Hyper-V Agent installer provides you an option to do this task. This option on the Hyper-V Agent installer to create and add a self signed certificate is selected by default. You can proceed with the installer to install CloudPlatform agent on a Hyper-V host.

**Note**

If you want to create and add a self signed certificate to the port before you run the installer, perform Step 1 in the following procedure. If you have performed step 1, you must clear the option on the installer to create and add a self signed certificate to the port. Otherwise, this will result in the failure of installation.

1. (Optional) Create and add a self-signed SSL certificate on port 8250:

- a. Create A self-signed SSL certificate. Run the following Power shell command:

```
# New-SelfSignedCertificate -DnsName apachecloudstack -CertStoreLocation Cert:
\LocalMachine\My
```

This command creates the self-signed certificate and add that to the certificate store **LocalMachine\My**.

- b. Add the created certificate to port 8250 for https communication:

```
netsh http add sslcert iport=0.0.0.0:8250 certhash=<thumbprint>
appid="{727beb1c-6e7c-49b2-8fbd-f03dbe481b08}"
```

Thumbprint is the thumbprint of the certificate you created.

2. At the Citrix CloudPlatform installation directory, identify the hyper-v agent MSI file (**CloudPlatform-<version>-N-hypervagent.msi**).
3. Copy the CloudPlatform Agent for Hyper-V from this directory to all the Hyper-V host machines.
4. Run the installer as an Administrator.

Press the Ctrl + Shift key while right-clicking the Agent Installer MSI to run as another user. You can select Administrator from the given options.

5. Provide the Domain user credentials when prompted.

The Domain user is part of the Hyper-V Administrators and local Administrators group on the host.

When the agent installation is finished, the agent runs as a service on the host machine.

To install Hyper-V msi through command line:

```
# msiexec /i CloudStackAgentSetup.msi /quiet /qn /norestart /log install.log
SERVICE_USERNAME=>username< SERVICE_PASSWORD=>password<
```

If you do not want to install certificate with the installer:

```
msiexec /i CloudStackAgentSetup.msi /quiet /qn /norestart /log install.log
SERVICE_USERNAME=>username< SERVICE_PASSWORD=>password<INSTALL_CERTIFICATE="False"
```

4.5. Physical Network Configuration for Hyper-V

You should have a plan for how the hosts will be cabled and which physical NICs will carry what types of traffic. By default, CloudPlatform will use the device that is used for the default route.

If you are using Hyper-V 2012 R2, manually create an external virtual switch before adding the host to CloudPlatform. If the Hyper-V host is added to the Hyper-V manager, select the host, then click Virtual Switch Manager, then New Virtual Switch. In the External Network, select the desired NIC adapter and click Apply.

If you are using Windows 2012 R2, virtual switch is created automatically.

4.6. Storage Preparation for Hyper-V (Optional)

CloudPlatform allows administrators to set up shared Primary Storage and Secondary Storage that uses SMB.

1. Create a SMB storage and expose it over SMB Version 3.

For more information, see [Deploying Hyper-V over SMB⁴](#).

You can also create and export SMB share using Windows. After the Windows Server 2012 R2 installation, select File and Storage Services from Server Roles to create an SMB file share. For more information, see [Creating an SMB File Share Using Server Manager⁵](#).

2. Add the SMB share to the Active Directory domain.

The SMB share and the hosts managed by CloudPlatform need to be in the same domain. However, the storage should be accessible from the Management Server with the domain user privileges.

3. While adding storage to CloudPlatform, ensure that the correct domain, and credentials are supplied. This user should be able to access the storage from the Management Server.

⁴ <http://technet.microsoft.com/en-us/library/jj134187.aspx>

⁵ http://technet.microsoft.com/en-us/library/jj134187.aspx#BKMK_Step3

Installing KVM for CloudPlatform

If you want to use the Linux Kernel Virtual Machine (KVM) hypervisor to run guest virtual machines, install KVM on the host(s) in your cloud. The material in this section doesn't duplicate KVM installation documentation. It provides the CloudPlatform-specific steps that are needed to prepare a KVM host to work with CloudPlatform.

5.1. System Requirements for KVM Hypervisor Hosts

5.1.1. Supported Operating Systems for KVM Hosts

KVM is included with a variety of Linux-based operating systems. The OS supported for use with CloudPlatform can be downloaded from the following website and installed by following the Installation Guide provided with the operating system.

- RHEL 6.2, 6.3, or 6.5: <https://access.redhat.com/downloads>

RHEL 6.4 is not supported. Upgrade to 6.5 to use KVM hosts.

- It is highly recommended that you purchase a RHEL support license. Citrix support can not be responsible for helping fix issues with the underlying OS.



Warning

Within a cluster, all KVM hosts must be running the same operating system.

5.1.2. System Requirements for KVM Hosts

- Must be certified as compatible with the selected operating system. See the RHEL Hardware Compatibility Guide at <https://hardware.redhat.com/>.
- Must support HVM (Intel-VT or AMD-V enabled)
- Use libvirt-0.10.2-41 version and above for KVM on RHEL 6.x versions
- All hosts within a cluster must be homogenous. The CPUs must be of the same type, count, and feature flags.
- Within a single cluster, the hosts must be of the same kernel version. For example, if one host is RHEL6.2 64-bit, they must all be RHEL6.2 64-bit.
- 64-bit x86 CPU (more cores results in better performance)
- 4 GB of memory
- 36 GB of local disk
- At least 1 NIC
- Statically allocated IP address

- When you deploy CloudPlatform, the hypervisor host must not have any VMs already running.
- Be sure all the hotfixes provided by the hypervisor vendor are applied. Track the release of hypervisor patches through your hypervisor vendor's support channel, and apply patches as soon as possible after they are released. CloudPlatform will not track or notify you of required hypervisor patches. It is essential that your hosts are completely up to date with the provided hypervisor patches. The hypervisor vendor is likely to refuse to support any system that is not up to date with patches.



Warning

The lack of up-to-date hotfixes can lead to data corruption and lost VMs.

5.2. Install and configure the Agent

1. Download the operating system that includes KVM (see [Section 5.1, "System Requirements for KVM Hypervisor Hosts"](#)) and install it on the host by following the Installation Guide provided with your chosen operating system.
2. After installation, perform the following configuration tasks, which are described in the next few sections.

Required	Optional
Section 5.3, "Installing the CloudPlatform Agent on a KVM Host"	Section 5.6, "Primary Storage Setup for KVM (Optional)"
Section 5.4, "Physical Network Configuration for KVM"	
Section 5.5, "Time Synchronization for KVM Hosts"	

5.3. Installing the CloudPlatform Agent on a KVM Host

Each KVM host must have the CloudPlatform Agent installed on it. This Agent communicates with the Management Server and controls all the instances on the host. Install the CloudPlatform Agent on each host using the following steps.

1. Check for a fully qualified hostname.

```
# hostname --fqdn
```

This should return a fully qualified hostname such as "kvm1.lab.example.org". If it does not, edit /etc/hosts so that it does.

2. Install qemu-kvm. CloudPlatform provides a patched version.

```
# yum install qemu-kvm
```

3. If you do not have a Red Hat Network account, you need to prepare a local Yum repository.
 - a. If you are working with a physical host, insert the RHEL installation CD. If you are using a VM, attach the RHEL ISO.
 - b. Mount the CDROM to /media.
 - c. Create a repo file at /etc/yum.repos.d/rhel6.repo. In the file, insert the following lines:

```
[rhel]
name=rhel6
baseurl=file:///media
enabled=1
gpgcheck=0
```

4. Install the CloudPlatform packages. You should have a file in the form of “CloudPlatform-VERSION-N-OSVERSION.tar.gz”.

Untar the file and then run the install.sh script inside it. Replace the file and directory names below with those you are using:

```
# tar xzf CloudPlatform-VERSION-N-OSVERSION.tar.gz
# cd CloudPlatform-VERSION-N-OSVERSION
# ./install.sh
```

You should see a few messages as the installer prepares, followed by a list of choices.

5. Choose “A” to install the Agent software.

```
> A
```

6. When the agent installation is finished, log in to the host as root and run the following commands to start essential services:

```
# service rpcbind start
# service nfs start
# chkconfig nfs on
# chkconfig rpcbind on
```

7. For ICMP rules to work, run the following command after KVM agent is installed.

```
# iptables -D FORWARD -p icmp -j ACCEPT
```

The CloudPlatform Agent is now installed.

If you find you need to stop or start the Agent, use these commands:

```
# service cloudstack-agent start
# service cloudstack-agent stop
```

5.4. Physical Network Configuration for KVM

You should have a plan for how the hosts will be cabled and which physical NICs will carry what types of traffic. By default, CloudPlatform will use the device that is used for the default route. This device will be placed in a CloudPlatform-created bridge. A bridge is necessary for all configurations.

The following network configuration should be done after installing the CloudPlatform Agent on the host.

If a system has multiple NICs or bonding is desired, you may configure the networking on the host. Manually create multiple bridges, then add them to CloudPlatform. You must place the desired device in the bridge. This may be done for each of the public network and the management network. Then edit `/etc/cloudstack/agent/agent.properties` and add values for the following:

- `public.network.device`
- `private.network.device`

These should be set to the name of the bridge that the user created for the respective traffic type. For example:

- `public.network.device=publicbondbr0`

5.5. Time Synchronization for KVM Hosts

The host must be set to use NTP. All hosts in a pod must have the same time.

1. Log in to the KVM host as root.
2. Install NTP.

```
# yum install ntp
```

3. Edit the NTP configuration file to point to your NTP server.

```
# vi /etc/ntp.conf
```

Add one or more server lines in this file with the names of the NTP servers you want to use. For example:

```
server 0.xenserver.pool.ntp.org
server 1.xenserver.pool.ntp.org
server 2.xenserver.pool.ntp.org
server 3.xenserver.pool.ntp.org
```

4. Restart the NTP client.

```
# service ntpd restart
```

5. Make sure NTP will start again upon reboot.

```
# chkconfig ntpd on
```

5.6. Primary Storage Setup for KVM (Optional)

CloudPlatform allows administrators to set up shared Primary Storage that uses iSCSI or fiber channel. With KVM, the storage is mounted on each host. This is called "SharedMountPoint" storage and is an alternative to NFS. The storage is based on some clustered file system technology, such as OCFS2.



Note

The use of the Cluster Logical Volume Manager (CLVM) is not officially supported with CloudPlatform.

With SharedMountPoint storage:

- Each node in the KVM cluster mounts the storage in the same local location (e.g., /mnt/primary)
- A shared clustered file system is used
- The administrator manages the mounting and unmounting of the storage
- If you want to use SharedMountPoint storage you should set it up on the KVM hosts now. Note the mountpoint that you have used on each host; you will use that later to configure CloudPlatform.

Installing VMware vSphere for CloudPlatform

If you want to use the VMware vSphere hypervisor to run guest virtual machines, install vSphere on the host(s) in your cloud.

6.1. System Requirements for vSphere Hosts

6.1.1. Software requirements

- VMware vCenter versions 5.0 upto Update 3a
- VMware vCenter versions 5.1 upto Update 2a
- VMware vCenter 5.5 version upto update 2 and 1c.

vSphere Standard is recommended. Note however that customers need to consider the CPU constraints in place with vSphere licensing. See http://www.vmware.com/files/pdf/vsphere_pricing.pdf and discuss with your VMware sales representative.

vCenter Server Standard is recommended.

- VMware vCenter versions 6.0
- Be sure all the hotfixes provided by the hypervisor vendor are applied. Track the release of hypervisor patches through your hypervisor vendor's support channel, and apply patches as soon as possible after they are released. CloudPlatform will not track or notify you of required hypervisor patches. It is essential that your hosts are completely up to date with the provided hypervisor patches. The hypervisor vendor is likely to refuse to support any system that is not up to date with patches.



Apply All Necessary Hotfixes

The lack of up-do-date hotfixes can lead to data corruption and lost VMs.

6.1.2. Hardware requirements

- The host must be certified as compatible with the vSphere version you are using. See the VMware Hardware Compatibility Guide at <http://www.vmware.com/resources/compatibility/search.php>.
- All hosts must be 64-bit and must support HVM (Intel-VT or AMD-V enabled).
- All hosts within a cluster must be homogenous. That means the CPUs must be of the same type, count, and feature flags.
- 64-bit x86 CPU (more cores results in better performance)
- Hardware virtualization support required

- 4 GB of memory
- 36 GB of local disk
- At least 1 NIC
- Statically allocated IP Address

6.1.3. vCenter Server requirements:

- Processor - 2 CPUs 2.0GHz or higher Intel or AMD x86 processors. Processor requirements may be higher if the database runs on the same machine.
- Memory - 3GB RAM. RAM requirements may be higher if your database runs on the same machine.
- Disk storage - 2GB. Disk requirements may be higher if your database runs on the same machine.
- Microsoft SQL Server 2005 Express disk requirements. The bundled database requires up to 2GB free disk space to decompress the installation archive.
- Networking - 1Gbit or 10Gbit.

For more information, see "vCenter Server and the vSphere Client Hardware Requirements" at http://pubs.vmware.com/vsp40/wwhelp/wwhimpl/js/html/wwhelp.htm#href=install/c_vc_hw.html.

6.1.4. Other requirements:

- VMware vCenter Standard Edition must be installed and available to manage the vSphere hosts.
- vCenter must be configured to use the standard port 443 so that it can communicate with the CloudPlatform Management Server.
- You must re-install VMware ESXi if you are going to re-use a host from a previous install.
- CloudPlatform requires VMware vSphere 5.0, 5.1 or 5.5. VMware vSphere 4.0 and 4.1 are not supported.
- All hosts must be 64-bit and must support HVM (Intel-VT or AMD-V enabled). All hosts within a cluster must be homogenous. That means the CPUs must be of the same type, count, and feature flags.
- The CloudPlatform management network must not be configured as a separate virtual network. The CloudPlatform management network is the same as the vCenter management network, and will inherit its configuration. See [Section 6.5.2, "Configure vCenter Management Network"](#).
- CloudPlatform requires ESXi. ESX is not supported.
- All resources used for CloudPlatform must be used for CloudPlatform only. CloudPlatform cannot share instance of ESXi or storage with other management consoles. Do not share the same storage volumes that will be used by CloudPlatform with a different set of ESXi servers that are not managed by CloudPlatform.
- Put all target ESXi hypervisors in a cluster in a separate Datacenter in vCenter.
- The cluster that will be managed by CloudPlatform should not contain any VMs. Do not run the management server, vCenter or any other VMs on the cluster that is designated for CloudPlatform use. Create a separate cluster for use of CloudPlatform and make sure that they are no VMs in this cluster.

- All the required VLANS must be trunked into all network switches that are connected to the ESXi hypervisor hosts. These would include the VLANS for Management, Storage, vMotion, and guest VLANs. The guest VLAN (used in Advanced Networking; see Network Setup) is a contiguous range of VLANs that will be managed by CloudPlatform.

6.2. Preparation Checklist for VMware

For a smoother installation, gather the following information before you start:

- Information listed in [Section 6.2.1, “vCenter Checklist ”](#)
- Information listed in [Section 6.2.2, “Networking Checklist for VMware ”](#)

6.2.1. vCenter Checklist

You will need the following information about vCenter.

vCenter Requirement	Value	Notes
vCenter User		This user must have admin privileges.
vCenter User Password		Password for the above user.
vCenter Datacenter Name		Name of the datacenter.
vCenter Cluster Name		Name of the cluster.

6.2.2. Networking Checklist for VMware

You will need the following information about the VLAN.

VLAN Information	Value	Notes
ESXi VLAN		VLAN on which all your ESXi hypervisors reside.
ESXi VLAN IP Address		IP Address Range in the ESXi VLAN. One address per Virtual Router is used from this range.
ESXi VLAN IP Gateway		
ESXi VLAN Netmask		
Management Server VLAN		VLAN on which the CloudPlatform Management server is installed.
Public VLAN		VLAN for the Public Network.
Public VLAN Gateway		
Public VLAN Netmask		
Public VLAN IP Address Range		Range of Public IP Addresses available for CloudPlatform use. These addresses will be used for virtual router on CloudPlatform to route private traffic to external networks.

VLAN Information	Value	Notes
VLAN Range for Customer use		A contiguous range of non-routable VLANs. One VLAN will be assigned for each customer.

6.3. vSphere Installation Steps

1. If you haven't already, you'll need to download and purchase vSphere from the VMware Website (<https://www.vmware.com/tryvmware/index.php?p=vmware-vsphere&lp=1>) and install it by following the VMware vSphere Installation Guide.
2. Following installation, perform the following configuration steps, which are described in the next few sections:

Required	Optional
ESXi host setup	NIC bonding
Configure host physical networking, virtual switch, vCenter Management Network, and extended port range	Multipath storage
Prepare storage for iSCSI	
Configure clusters in vCenter and add hosts to them, or add hosts without clusters to vCenter	

6.4. ESXi Host setup

All ESXi hosts should enable CPU hardware virtualization support in BIOS. Please note hardware virtualization support is not enabled by default on most servers.

6.5. Physical Host Networking

You should have a plan for cabling the vSphere hosts. Proper network configuration is required before adding a vSphere host to CloudPlatform. To configure an ESXi host, you can use vClient to add it as standalone host to vCenter first. Once you see the host appearing in the vCenter inventory tree, click the host node in the inventory tree, and navigate to the Configuration tab.

In the host configuration tab, click the "Hardware/Networking" link to bring up the networking configuration page as above.

6.5.1. Configure Virtual Switch

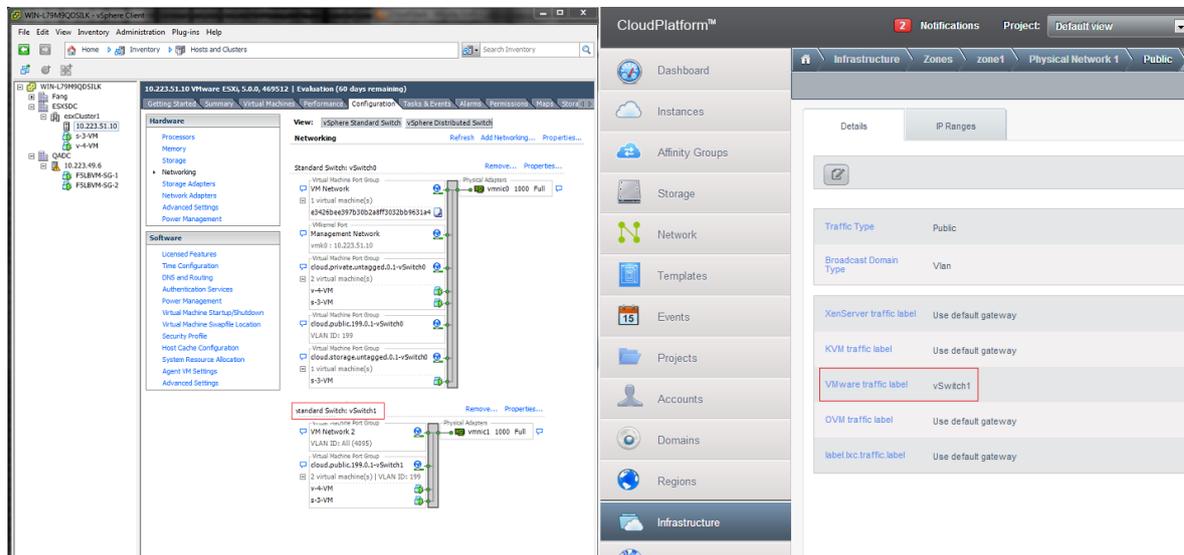
A default virtual switch vSwitch0 is created. CloudPlatform requires all ESXi hosts in the cloud to use the same set of virtual switch names. If you change the default virtual switch name, you will need to configure one or more CloudPlatform configuration variables as well.

6.5.1.1. Separating Traffic

CloudPlatform allows you to use vCenter to configure three separate networks per ESXi host. These networks are identified by the name of the vSwitch they are connected to. The allowed networks for configuration are public (for traffic to/from the public internet), guest (for guest-guest traffic), and private (for management and usually storage traffic). You can use the default virtual switch for all three, or create one or two other vSwitches for those traffic types.

If you want to separate traffic in this way you should first create and configure vSwitches in vCenter according to the vCenter instructions. Take note of the vSwitch names you have used for each traffic type. You will configure CloudPlatform to use these vSwitches.

For example, in the following figure, you can see that the Standard vSwitch name is used in CloudPlatform as the VMware traffic label.



6.5.1.2. Increasing Ports

!
Note

On a virtual switch on an ESXi 5.5 host, the number of ports are dynamically scaled up or down depending on the usage. You cannot manually modify it.

By default a virtual switch on ESXi hosts is created with 56 ports. We recommend setting it to 4088, the maximum number of ports allowed. To do that, click the "Properties..." link for virtual switch (note this is not the Properties link for Networking).

In vSwitch properties dialog, select the vSwitch and click Edit.

In the dialog, you can change the number of switch ports. After you have done that, ESXi hosts are required to reboot in order for the setting to take effect.

Adding Additional Ports to a VNC

You might need to extend the range of firewall ports that the console proxy works with on those hosts. This is to enable the console proxy to work with VMware-based VMs by way of gdbserver. The default additional port range is 59000-60000.

To extend the port range:

1. Log in to the CloudStack.
2. From the left navigational bar, click Global Settings.

3. Set the following:

- `vmware.additional.vnc.portrange.size = 1000`
- `vmware.additional.vnc.portrange.start = 59000`

4. On each ESX host, log in to the VMware ESX service console and run the following commands:

```
# esxcfg-firewall -o 59000-60000,tcp,in,vncextras
# esxcfg-firewall -o 59000-60000,tcp,out,vncextras
```

6.5.2. Configure vCenter Management Network

In the vSwitch properties dialog box, you may see a vCenter management network. This same network will also be used as the CloudPlatform management network. CloudPlatform requires the vCenter management network to be configured properly. Select the management network item in the dialog, then click Edit.

Make sure the following values are set:

- VLAN ID set to the desired ID
- vMotion enabled.
- Management traffic enabled.

If the ESXi hosts have multiple VMKernel ports, and ESXi is not using the default value "Management Network" as the management network name, you must follow these guidelines to configure the management network port group so that CloudPlatform can find it:

- Use one label for the management network port across all ESXi hosts.
- In the CloudPlatform UI, go to Global Settings and set `vmware.management.portgroup` to the management network label from the ESXi hosts.

6.5.3. Configure NIC Bonding for vSphere

NIC bonding on vSphere hosts may be done according to the vSphere installation guide.

6.6. Configuring a vSphere Cluster with Nexus 1000v Virtual Switch

CloudPlatform supports Cisco Nexus 1000v dvSwitch (Distributed Virtual Switch) for virtual network configuration in a VMware vSphere environment. This section helps you configure a vSphere cluster with Nexus 1000v virtual switch in a VMware vCenter environment. For information on creating a vSphere cluster, see [Chapter 6, Installing VMware vSphere for CloudPlatform](#)

6.6.1. About Cisco Nexus 1000v Distributed Virtual Switch

The Cisco Nexus 1000V virtual switch is a software-based virtual machine access switch for VMware vSphere environments. It can span multiple hosts running VMware ESXi 4.0 and later. A Nexus virtual switch consists of two components: the Virtual Supervisor Module (VSM) and the Virtual Ethernet Module (VEM). The VSM is a virtual appliance that acts as the switch's supervisor. It controls multiple VEMs as a single network device. The VSM is installed independent of the VEM and is deployed in redundancy mode as pairs or as a standalone appliance. The VEM is installed on each VMware ESXi

server to provide packet-forwarding capability. It provides each virtual machine with dedicated switch ports. This VSM-VEM architecture is analogous to a physical Cisco switch's supervisor (standalone or configured in high-availability mode) and multiple linecards architecture.

Nexus 1000v switch uses vEthernet port profiles to simplify network provisioning for virtual machines. There are two types of port profiles: Ethernet port profile and vEthernet port profile. The Ethernet port profile is applied to the physical uplink ports-the NIC ports of the physical NIC adapter on an ESXi server. The vEthernet port profile is associated with the virtual NIC (vNIC) that is plumbed on a guest VM on the ESXi server. The port profiles help the network administrators define network policies which can be reused for new virtual machines. The Ethernet port profiles are created on the VSM and are represented as port groups on the vCenter server.

6.6.2. Prerequisites and Guidelines

Nexus virtual switch in CloudPlatform supports management and storage traffic. This section discusses prerequisites and guidelines for using Nexus virtual switch in CloudPlatform. Before configuring Nexus virtual switch, ensure that your system meets the following requirements:

- A cluster of servers (ESXi 4.1 or later) is configured in the vCenter.
- Each cluster managed by CloudPlatform is the only cluster in its vCenter datacenter.
- A Cisco Nexus 1000v virtual switch is installed to serve the datacenter that contains the vCenter cluster. This ensures that CloudPlatform doesn't have to deal with dynamic migration of virtual adapters or networks across other existing virtual switches. See [Cisco Nexus 1000V Installation and Upgrade Guide](#)¹ for guidelines on how to install the Nexus 1000v VSM and VEM modules.
- The Nexus 1000v VSM is not deployed on a vSphere host that is managed by CloudPlatform.
- When the maximum number of VEM modules per VSM instance is reached, an additional VSM instance is created before introducing any more ESXi hosts. The limit is 64 VEM modules for each VSM instance.
- CloudPlatform expects that the Management Network of the ESXi host is configured on the standard vSwitch and searches for it in the standard vSwitch. Therefore, ensure that you do not migrate the management network to Nexus 1000v virtual switch during configuration.
- All information given in [Section 6.6.3, "Nexus 1000v Virtual Switch Preconfiguration"](#)

6.6.3. Nexus 1000v Virtual Switch Preconfiguration

6.6.3.1. Preparation Checklist

For a smoother configuration of Nexus 1000v switch, gather the following information before you start:

- vCenter Credentials
- Nexus 1000v VSM IP address
- Nexus 1000v VSM Credentials
- Ethernet port profile names

¹ http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_1_5_1/install_upgrade/vsm_vem/guide/n1000v_installupgrade.html

6.6.3.1.1. vCenter Credentials Checklist

You will need the following information about vCenter:

Nexus vSwitch Requirements	Value	Notes
vCenter IP		The IP address of the vCenter.
Secure HTTP Port Number	443	Port 443 is configured by default; however, you can change the port if needed.
vCenter User ID		The vCenter user with administrator-level privileges. The vCenter User ID is required when you configure the virtual switch in CloudPlatform.
vCenter Password		The password for the vCenter user specified above. The password for this vCenter user is required when you configure the switch in CloudPlatform.

6.6.3.1.2. Network Configuration Checklist

The following information specified in the Nexus Configure Networking screen is displayed in the Details tab of the Nexus dvSwitch in the CloudPlatform UI:

Network Requirements	Value	Notes
Control Port Group VLAN ID		The VLAN ID of the Control Port Group. The control VLAN is used for communication between the VSM and the VEMs.
Management Port Group VLAN ID		The VLAN ID of the Management Port Group. The management VLAN corresponds to the mgmt0 interface that is used to establish and maintain the connection between the VSM and VMware vCenter Server.
Packet Port Group VLAN ID		The VLAN ID of the Packet Port Group. The packet VLAN forwards relevant data packets from the VEMs to the VSM.



Note

The VLANs used for control, packet, and management port groups can be the same.

For more information, see [Cisco Nexus 1000V Getting Started Guide²](#).

6.6.3.1.3. VSM Configuration Checklist

You will need the following information about network configuration:

VSM Configuration Parameters Value Notes	Value	Notes
Admin Name and Password		The admin name and password to connect to the VSM appliance. You must specify these credentials while configuring Nexus virtual switch.
Management IP Address		This is the IP address of the VSM appliance. This is the IP address you specify in the virtual switch IP Address field while configuring Nexus virtual switch.
SSL	Enable	Always enable SSL. SSH is usually enabled by default during the VSM installation. However, check whether the SSH connection to the VSM is working, without which CloudPlatform fails to connect to the VSM.

6.6.3.2. Creating a Port Profile

- Whether you create a Basic or Advanced zone configuration, ensure that you always create an Ethernet port profile on the VSM after you install it and before you create the zone.
 - The Ethernet port profile created to represent the physical network or networks used by an Advanced zone configuration trunk all the VLANs including guest VLANs, the VLANs that serve the native VLAN, and the packet/control/data/management VLANs of the VSM.
 - The Ethernet port profile created for a Basic zone configuration does not trunk the guest VLANs because the guest VMs do not get their own VLANs provisioned on their network interfaces in a Basic zone.
- An Ethernet port profile configured on the Nexus 1000v virtual switch should not use in its set of system VLANs, or any of the VLANs configured or intended to be configured for use towards VMs or VM resources in the CloudPlatform environment.
- You do not have to create any vEthernet port profiles – CloudPlatform does that during VM deployment.

² http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_1_4_b/getting_started/configuration/guide/n1000v_gsg.pdf

- Ensure that you create required port profiles to be used by CloudPlatform for different traffic types of CloudPlatform, such as Management traffic, Guest traffic, Storage traffic, and Public traffic. The physical networks configured during zone creation should have a one-to-one relation with the Ethernet port profiles.

For information on creating a port profile, see [Cisco Nexus 1000V Port Profile Configuration Guide](#)³.

6.6.3.3. Assigning Physical NIC Adapters

Assign ESXi host's physical NIC adapters, which correspond to each physical network, to the port profiles. In each ESXi host that is part of the vCenter cluster, observe the physical networks assigned to each port profile and note down the names of the port profile for future use. This mapping information helps you when configuring physical networks during the zone configuration on CloudPlatform. These Ethernet port profile names are later specified as VMware Traffic Labels for different traffic types when configuring physical networks during the zone configuration.

Traffic label format in the last case is `[["Name of vSwitch/dvSwitch/EthernetPortProfile"],["VLANID"],["vSwitch Type"]]`. The possible value for Nexus traffic label is `myEthernetPortProfile,,nexusdvs`.

For more information on configuring physical networks, see [Section 6.6, "Configuring a vSphere Cluster with Nexus 1000v Virtual Switch"](#).

6.6.3.4. Adding VLAN Ranges

Determine the public VLAN, System VLAN, and Guest VLANs to be used by the CloudPlatform. Ensure that you add them to the port profile database. Corresponding to each physical network, add the VLAN range to port profiles. In the VSM command prompt, run the `switchport trunk allowed vlan<range>` command to add the VLAN ranges to the port profile.

For example:

```
switchport trunk allowed vlan 1,140-147,196-203
```

In this example, the allowed VLANs added are 1, 140-147, and 196-203

You must also add all the public and private VLANs or VLAN ranges to the switch. This range is the VLAN range you specify in your zone.



Note

Before you run the `vlan` command, ensure that the configuration mode is enabled in Nexus 1000v virtual switch.

For example:

If you want the VLAN 200 to be used on the switch, run the following command:

³ http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_1_4_a/port_profile/configuration/guide/n1000v_port_profile.html

```
vlan 200
```

If you want the VLAN range 1350-1750 to be used on the switch, run the following command:

```
vlan 1350-1750
```

Refer to Cisco Nexus 1000V Command Reference of specific product version.

6.6.4. Enabling Nexus Virtual Switch in CloudPlatform

To make a CloudPlatform deployment Nexus enabled, you must set the `vmware.use.nexus.vswitch` parameter true by using the Global Settings page in the CloudPlatform UI. Unless this parameter is set to "true" and restart the management server, you cannot see any UI options specific to Nexus virtual switch, and CloudPlatform ignores the Nexus virtual switch specific parameters specified in the `AddTrafficTypeCmd`, `UpdateTrafficTypeCmd`, and `AddClusterCmd` API calls.

Unless the CloudPlatform global parameter "vmware.use.nexus.vswitch" is set to "true", CloudPlatform by default uses VMware standard vSwitch for virtual network infrastructure. In this release, CloudPlatform doesn't support configuring virtual networks in a deployment with a mix of standard vSwitch and Nexus 1000v virtual switch. The deployment can have either standard vSwitch or Nexus 1000v virtual switch.

6.6.5. Configuring Nexus 1000v Virtual Switch in CloudPlatform

You can configure Nexus dvSwitch by adding the necessary resources while the zone is being created.

After the zone is created, if you want to create an additional cluster along with Nexus 1000v virtual switch in the existing zone, use the Add Cluster option.

In both these cases, you must specify the following parameters to configure Nexus virtual switch:

Parameters	Description
Cluster Name	Enter the name of the cluster you created in vCenter. For example, "cloud.cluster".
vCenter Host	Enter the host name or the IP address of the vCenter host where you have deployed the Nexus virtual switch.
vCenter User name	Enter the username that CloudPlatform should use to connect to vCenter. This user must have all administrative privileges.
vCenter Password	Enter the password for the user named above.
vCenter Datacenter	Enter the vCenter datacenter that the cluster is in. For example, "cloud.dc.VM".
Nexus dvSwitch IP Address	The IP address of the VSM component of the Nexus 1000v virtual switch.
Nexus dvSwitch Username	The admin name to connect to the VSM appliance.
Nexus dvSwitch Password	The corresponding password for the admin user specified above.

6.6.6. Removing Nexus Virtual Switch

1. In the vCenter datacenter that is served by the Nexus virtual switch, ensure that you delete all the hosts in the corresponding cluster.
2. Log in with Admin permissions to the CloudPlatform administrator UI.
3. In the left navigation bar, select Infrastructure.
4. In the Infrastructure page, click View all under Clusters.
5. Select the cluster where you want to remove the virtual switch.
6. In the dvSwitch tab, click **Nexus**.
7. In the Details page, click Delete Nexus dvSwitch icon. 

Click Yes in the confirmation dialog box.

For more information, see the following procedures in [Cisco's documentation web site](#)⁴.

- Removing the Hosts from the Cisco Nexus 1000V DVS
- Removing the Cisco Nexus 1000V From the vCenter Server
- Unregistering the Extension Key in the vCenter Server

6.6.7. Configuring a VMware Datacenter with VMware Distributed Virtual Switch

CloudPlatform supports VMware vNetwork Distributed Switch (VDS) for virtual network configuration in a VMware vSphere environment. This section helps you configure VMware VDS in a CloudPlatform deployment. Each vCenter server instance can support up to 128 VDS instances and each VDS instance can manage up to 500 VMware hosts.

6.6.7.1. About VMware Distributed Virtual Switch

VMware VDS is an aggregation of host-level virtual switches on a VMware vCenter server. VDS abstracts the configuration of individual virtual switches that span across a large number of hosts, and enables centralized provisioning, administration, and monitoring for your entire datacenter from a centralized interface. In effect, a VDS acts as a single virtual switch at the datacenter level and manages networking for a number of hosts in a datacenter from a centralized VMware vCenter server. Each VDS maintains network runtime state for VMs as they move across multiple hosts, enabling inline monitoring and centralized firewall services. A VDS can be deployed with or without Virtual Standard Switch and a Nexus 1000V virtual switch.

6.6.7.2. Prerequisites and Guidelines

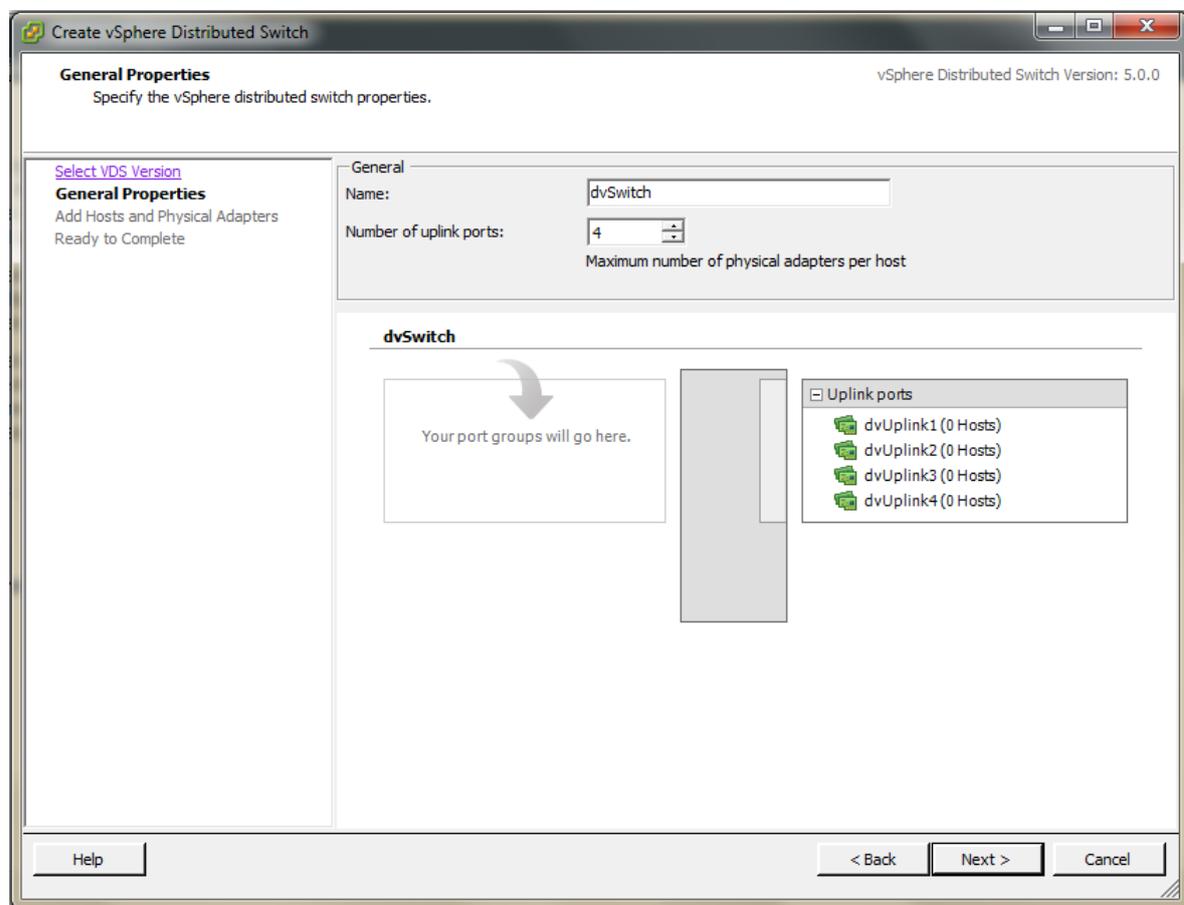
- Availability of one or more ESXi hosts with VMware vSphere 5.0 or later.
- Availability of one or more vCenter 5.0 or later installed with the license required to enable VDS.

⁴ http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus1000/sw/4_0/troubleshooting/configuration/guide/n1000v_troubleshooting/trouble_3install.html#wp1225757

- Create VDS in the vCenter data center associated with the CloudPlatform zone.
- Add all the ESXi hosts in a vSphere cluster to VDS before you add the vSphere cluster to a pod in the CloudPlatform zone.
- VMware VDS is supported for all the four types of traffic - Public, Guest, Management, and Storage - in CloudPlatform.
- When you remove a guest network, the corresponding dvportgroup will not be removed on the vCenter. You must manually delete them on the vCenter.

6.6.7.3. Preparation Checklist

For a smoother configuration of VMware VDS, note down the VDS name you have added in the datacenter before you start:



Use this VDS name in the following:

- The switch name in the Edit traffic label dialog while configuring public, guest, management, and storage traffic during zone creation.

During a zone creation, ensure that you select VMware vNetwork Distributed Virtual Switch when you configure public, guest, management, and storage traffic type.

- The Public Traffic vSwitch Type field when you add a VMware VDS-enabled cluster.
- The switch name in the traffic label while updating the switch type in a zone.

Traffic label format in the last case is `[["Name of vSwitch/dvSwitch/EthernetPortProfile"],["VLAN ID"],["vSwitch Type"]]`

The possible values for traffic labels are:

- empty string
- dvSwitch0
- dvSwitch0,200
- dvSwitch1,300,vmwaredvs
- dvSwitch0,,vmwaredvs

Fields	Name	Description
1	Represents the name of the virtual / distributed virtual switch at vCenter.	<p>The default value depends on the type of virtual switch:</p> <p>vSwitch0: If type of virtual switch is VMware vNetwork Standard virtual switch</p> <p>dvSwitch0: If type of virtual switch is VMware vNetwork Distributed virtual switch</p> <p>epp0: If type of virtual switch is Cisco Nexus 1000v Distributed virtual switch</p>
2	VLAN ID to be used for this traffic wherever applicable.	<p>This field would be used for only public traffic as of now. In case of guest traffic this field would be ignored and could be left empty for guest traffic. By default empty string would be assumed which translates to untagged VLAN for that specific traffic type.</p>
3	Type of virtual switch. Specified as string.	<p>Possible valid values are vmwaredvs, vmwaresvs, nexUSDvs.</p> <p>vmwaresvs: Represents VMware vNetwork Standard virtual switch</p> <p>vmwaredvs: Represents VMware vNetwork distributed virtual switch</p> <p>nexUSDvs: Represents Cisco Nexus 1000v distributed virtual switch.</p> <p>If nothing specified (left empty), zone-level default virtual switch would be defaulted, based on</p>

Fields	Name	Description
		<p>the value of global parameter you specify.</p> <p>Following are the global configuration parameters:</p> <p>vmware.use.dvswitch: Set to true to enable any kind (VMware DVS and Cisco Nexus 1000v) of distributed virtual switch in a CloudPlatform deployment. If set to false, the virtual switch that can be used in that CloudPlatform deployment is Standard virtual switch.</p> <p>vmware.use.nexus.vswitch: This parameter is ignored if vmware.use.dvswitch is set to false. Set to true to enable Cisco Nexus 1000v distributed virtual switch in a CloudPlatform deployment.</p>

6.6.7.4. Enabling Virtual Distributed Switch in CloudPlatform

- To make a CloudPlatform deployment VDS enabled, set the **vmware.use.dvswitch** parameter to true by using the Global Settings page in the CloudPlatform UI and restart the Management Server. Unless you enable the **vmware.use.dvswitch** parameter, you cannot see any UI options specific to VDS, and CloudPlatform ignores the VDS-specific parameters that you specify. Additionally, CloudPlatform uses VDS for virtual network infrastructure if the value of **vmware.use.dvswitch** parameter is true and the value of **vmware.use.nexus.dvswitch** parameter is false.
- For VDS version 4.1 or earlier, you need to use the global configuration parameter **vmware.ports.per.dvportgroup**. This is the default number of ports per VMware dvPortGroup in a VMware environment. By default, the value of this parameter is set to 256. That is, each dvPortGroup created by CloudPlatform will contain 256 dvPorts. This number is associated with number of virtual NICs in a network.
- With vSphere version 5.0 or later, CloudPlatform leverages the autoExpand feature that is available in vSphere. This feature enables scaling up or scaling down of dvPorts in dvPortGroup. Default value is 256.

CloudPlatform supports orchestration of virtual networks in a deployment with a mix of Virtual Distributed Switch, Standard Virtual Switch and Nexus 1000v Virtual Switch. However, Citrix has verified only the following combinations:

- Standard Virtual Switch with Nexus dvSwitch
- Standard Virtual Switch with VMWare Distributed Switch

The term dvPortGroup that is associated with VDS is analogous to the term virtual network that is associated with CloudPlatform. Similarly, dvPort is analogous to vNIC and dvSwitch is analogous to

the physical network. After configuring VDS, CloudPlatform does the following operations that are limited to dvPortGroups:

- Create dvPortGroup over designated VDS.
- Modify dvPortGroup over designated VDS.



Note

CloudPlatform does not delete dvPortGroups.

6.6.7.5. Configuring Distributed Virtual Switch in CloudPlatform

You can configure VDS by adding the necessary resources while a zone is created.

Alternatively, at the cluster level, you can create an additional cluster with VDS enabled in the existing zone. Use the Add Cluster option.

In both these cases, you must specify the following parameters to configure VDS:



Note

In CloudPlatform, the Override option is supported for the public and the guest traffic types. The option is not supported for the management and the storage traffic types.

Parameters	Description
Cluster Name	Enter the name of the cluster you created in vCenter. For example, "cloudcluster".
vCenter Host	Enter the name or the IP address of the vCenter host where you have deployed the VMware VDS.
vCenter User name	Enter the username that CloudPlatform should use to connect to vCenter. This user must have all administrative privileges.
vCenter Password	Enter the password for the user named above.
vCenter Datacenter	Enter the vCenter datacenter that the cluster is in. For example, "cloudcVM".
Override Public Traffic	Enable this option to override the zone-wide public traffic for the cluster you are creating.
Public Traffic vSwitch Type	This option is displayed only if you enable the Override Public Traffic option. Select VMware vNetwork Distributed Virtual Switch. If the <code>vmware.use.dvswitch</code> global parameter is true, the default option will be VMware vNetwork Distributed Virtual Switch.

Parameters	Description
Public Traffic vSwitch Name	Name of virtual switch to be used for the public traffic.
Override Guest Traffic	Enable the option to override the zone-wide guest traffic for the cluster you are creating.
Guest Traffic vSwitch Type	This option is displayed only if you enable the Override Guest Traffic option. Select VMware vNetwork Distributed Virtual Switch. If the <code>vmware.use.dvswitch</code> global parameter is true, the default option will be VMware vNetwork Distributed Virtual Switch.
Guest Traffic vSwitch Name	Name of virtual switch to be used for guest traffic.

6.6.7.6. Use Cases for Migrating Virtual Networks from standard vSwitch to VDS

To use VDS for your virtual networks, you need to migrate the existing virtual networks from a standard vSwitch to VDS. The following use cases describes how you can migrate:

- Public traffic from standard vSwitch to VDS
- Guest traffic from standard vSwitch to VDS
- Management traffic from standard vSwitch to VDS
- Storage traffic from standard vSwitch to VDS

Prerequisites

You must preserve the following:

- Virtual NIC ordering
- VLAN membership
- MAC Adress of each virtual NIC
- Shaping policies

Pre-migration Tasks

- Disable zone.
- Stop all user VMs.
- Stop all VRs.
- Stop all system VMs.
- Set the value of the global configuration parameter `vmware.use.dvswitch` to `true`.

Migrating Public Traffic from Standard vSwitch to VDS

To migrate public traffic from Standard vSwitch to VDS, do the following:

1. In CloudPlatform UI, on the left navigation panel, click **Infrastructure**.
2. On the right-side panel, in the Zones box, click **View all**.
3. In the Zones page that lists all the zones that you have provisioned, click the name of the zone in which you want to migrate the virtual networks to VDS.
4. Click the **Physical Network** tab.
5. In the page that shows the details of the physical network, click the name of the physical network.
6. In the page that displays the network configuration, click **Configure** in the **Public** box. This represents the public traffic type.
7. In the **Details** page that displays the backend configuration, click the **Edit** icon and do the following:
 - a. Update the VMware traffic label to reflect the name of the VDS that is assigned to carry the public traffic in the zone.
 - b. Update vSwitch type as **vmwaredvs** in last field of traffic label.
 - c. Click Apply.

The traffic label appears in the following format:

```
<vswitch_name>[,,vswitch_type]
```

You must ensure that this VDS has already been available in the VMware datacenter that is associated with this zone.

For example, before you perform migration, the details of physical network for public traffic appears as follows:

Physical Network	Traffic Type	VMWare Traffic Label
Physical Network 2	Public	vSwitch1,,vmwaresvs

After you perform migration, the physical network for public traffic appears as follows:

Physical Network	Traffic Type	VMWare Traffic Label
Physical Network 2	Public	dvS-public,,vmwaredvs

8. On each ESXi host of the zone, do the following using vSphere client or web client:
 - a. Add host to the VDS - dvS-public - that is designated by the administrator to carry public traffic without adding or selecting any physical NICs.
 - b. Identify the standard vSwitch that carries the public traffic and check for the VMKernel port on this vSwitch.
 - c. If the standard vSwitch does not contain any VMKernel port, remove the physical NICs from that standard vSwitch and then, add those physical NICs to the VDS designated for public traffic. (To achieve this, go to dvSwitch and click **Manage hosts** and choose this specific host. Then, click **Next** to see the list of all physical NIC(s) associated with that host. Then, select the **physical** NIC(s) removed from the Standard vSwitch and click **Next**.)

⁵ http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1010614

- d. If VMKernel port configured on the physical NIC(s), you must migrate the VMKernel port to VDS while adding physical NICs to this VDS. For more information on migrating VMKernel port on standard vSwitch to VDS, refer to [VMware Knowledge Base article⁵](#).
- 9. In the database, update the cluster_details table to change type and name of vSwitches that are used for the public and the guest traffic.

```
update cluster_details set value='vmwaredvs' where name = 'publicvswitchtype' and cluster_id in (select id from cluster where data_center_id=<ZONE_ID>);
```

```
update cluster_details set value='<NAME_OF_DVSWITCH_FOR_PUBLIC_TRAFFIC>' where name = 'publicvswitchname' and cluster_id in (select id from cluster where data_center_id=<ZONE_ID>);
```

Migrating Guest Traffic from Standard vSwitch to VDS

To migrate guest traffic from Standard vSwitch to VDS, do the following:

1. In CloudPlatform UI, on the left navigation panel, click **Infrastructure**.
2. On the right-side panel, in the Zones box, click **View all**.
3. In the Zones page that lists all the zones that you have provisioned, click the name of the zone in which you want to migrate the virtual networks to VDS.
4. Click the **Physical Network** tab.
5. In the page that shows the details of the physical network, click the name of the physical network.
6. In the page that displays the network configuration, click **Configure** in the **Guest** box. This represents the guest traffic type.
7. In the **Details** page that displays the backend configuration, click the **Edit** icon and do the following:
 - a. Update the VMware traffic label to reflect the name of the VDS that is assigned to carry the guest traffic in the zone.
 - b. Update vSwitch type as **vmwaredvs** in last field of traffic label.
 - c. Click **Apply**.

The traffic label appears in the following format:

<vswitch_name>[, ,vswitch_type]

You must ensure that this VDS has already been available in the VMware datacenter that is associated with this zone.

For example, before you perform migration, the details of physical network for guest traffic appears as follows:

Physical Network	Traffic Type	VMWare Traffic Label
Physical Network 2	Guest	vSwitch2,,vmwaresvs

After you perform migration, the physical network for guest traffic appears as follows:

Physical Network	Traffic Type	VMWare Traffic Label
Physical Network 2	Guest	dvS-guest,,vmwaredvs

8. On each ESXi host of the zone, do the following using vSphere client or web client:
 - a. Add host to the VDS - dvS-guest - that is designated by the administrator to carry guest traffic without adding or selecting any physical NICs.
 - b. Identify the standard vSwitch that carries the guest traffic and check for the VMKernel port on this vSwitch.
 - c. If the standard vSwitch does not contain any VMKernel port, remove the physical NICs from that standard vSwitch and then, add those physical NICs to the VDS designated for guest traffic. (To achieve this, go to dvSwitch and click **Manage hosts** and choose this specific host. Then, click **Next** to see the list of all physical NIC(s) associated with that host. Then, select the physical NIC(s) removed from the Standard vSwitch and click **Next**.)
 - d. If VMKernel port configured on the physical NIC(s), you must migrate the VMKernel port to VDS while adding physical NICs to this VDS. For more information on migrating VMKernel port on standard vSwitch to VDS, refer to [VMware Knowledge Base article⁶](#).
9. In the database, update the cluster_details table to change type and name of vSwitches that are used for the public and the guest traffic.

```
update cluster_details set value='vmwaredvs' where name = 'guestvswitchtype' and cluster_id in (select id from cluster where data_center_id=<ZONE_ID>);
```

```
update cluster_details set value='<NAME_OF_DVSWITCH_FOR_PUBLIC_TRAFFIC>' where name = 'publicvswitchname' and cluster_id in (select id from cluster where data_center_id=<ZONE_ID>);
```

Migrating Management Traffic from Standard vSwitch to VDS

To migrate management traffic from Standard vSwitch to VDS, do the following:

1. In CloudPlatform UI, on the left navigation panel, click **Infrastructure**.
2. On the right-side panel, in the Zones box, click **View all**.
3. In the Zones page that lists all the zones that you have provisioned, click the name of the zone in which you want to migrate the virtual networks to VDS.
4. Click the **Physical Network** tab.
5. In the page that shows the details of the physical network, click the name of the physical network.
6. In the page that displays the network configuration, click **Configure** in the **Management** box. This represents the management traffic type.
7. In the **Details** page that displays the backend configuration, click the **Edit** icon and do the following:

⁶ http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1010614

- a. Update the VMware traffic label to reflect the name of the VDS that is assigned to carry the management traffic in the zone.
- b. Update vSwitch type as `vmwaredvs` in last field of traffic label.
- c. Click **Apply**.

The traffic label appears in the following format:

`<VSWITCH_NAME> [, [<MGMT_VLAN>] , <VSWITCH_TYPE>]`

You must ensure that this VDS has already been available in the VMware datacenter that is associated with this zone.

For example, before you perform migration, the details of physical network for management traffic appears as follows:

Physical Network	Traffic Type	VMWare Traffic Label
Physical Network 2	Management	vSwitch0,120,vmwaresvs

After you perform migration, the physical network for guest traffic appears as follows:

Physical Network	Traffic Type	VMWare Traffic Label
Physical Network 2	Management	dvS-mgmt,120,vmwaredvs

8. On each ESXi host of the zone, do the following using vSphere client or web client:
 - a. Add host to the VDS - dvS-mgmt - that is designated by the administrator to carry management traffic without adding or selecting any physical NICs.
 - b. Identify the standard vSwitch that carries the management traffic and check for the VMKernel port on this vSwitch.
 - c. If the standard vSwitch does not contain any VMKernel port, remove the physical NICs from that standard vSwitch and then, add those physical NICs to the VDS designated for management traffic. (To achieve this, go to dvSwitch and click **Manage hosts** and choose this specific host. Then, click **Next** to see the list of all physical NIC(s) associated with that host. Then, select the physical NIC(s) removed from the Standard vSwitch and click **Next**.)
 - d. If VMKernel port configured on the physical NIC(s), you must migrate the VMKernel port to VDS while adding physical NICs to this VDS. For more information on migrating VMKernel port on standard vSwitch to VDS, refer to [VMware Knowledge Base article⁷](#).

Migrating Storage Traffic from Standard vSwitch to VDS

To migrate storage traffic from Standard vSwitch to VDS, do the following:

1. In CloudPlatform UI, on the left navigation panel, click **Infrastructure**.
2. On the right-side panel, in the Zones box, click **View all**.
3. In the Zones page that lists all the zones that you have provisioned, click the name of the zone in which you want to migrate the virtual networks to VDS.

⁷ http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1010614

4. Click the **Physical Network** tab.
5. In the page that shows the details of the physical network, click the name of the physical network.
6. In the page that displays the network configuration, click **Configure** in the **Storage** box. This represents the storage traffic type.
7. In the **Details** page that displays the backend configuration, click the **Edit** icon and do the following:
 - a. Update the VMware traffic label to reflect the name of the VDS that is assigned to carry the storage traffic in the zone.
 - b. Update vSwitch type as `vmwaredvs` in last field of traffic label.
 - c. Click **Apply**.

The traffic label appears in the following format:

`<VSWITCH_NAME> [, , <VSWITCH_TYPE>]`

You must ensure that this VDS has already been available in the VMware datacenter that is associated with this zone.

For example, before you perform migration, the details of physical network for storage traffic appears as follows:

Physical Network	Traffic Type	VMWare Traffic Label
Physical Network 2	Storage	vSwitch0,120,vmwaresvs

After you perform migration, the physical network for guest traffic appears as follows:

Physical Network	Traffic Type	VMWare Traffic Label
Physical Network 2	Storage	dvS-storage,120,vmwaredvs

8. On each ESXi host of the zone, do the following using vSphere client or web client:
 - a. Add host to the VDS - dvS-storage - that is designated by the administrator to carry management traffic without adding or selecting any physical NICs.
 - b. Identify the standard vSwitch that carries the storage traffic and check for the VMKernel port on this vSwitch.
 - c. If the standard vSwitch does not contain any VMKernel port, remove the physical NICs from that standard vSwitch and then, add those physical NICs to the VDS designated for storage traffic. (To achieve this, go to dvSwitch and click **Manage hosts** and choose this specific host. Then, click **Next** to see the list of all physical NIC(s) associated with that host. Then, select the physical NIC(s) removed from the Standard vSwitch and click **Next**.)
 - d. If VMKernel port configured on the physical NIC(s), you must migrate the VMKernel port to VDS while adding physical NICs to this VDS. For more information on migrating VMKernel port on standard vSwitch to VDS, refer to [VMware Knowledge Base article](#)⁸.

⁸ http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1010614

Post-Migration Tasks

1. Start the Management Server.
2. Enable zone and wait for system VMs to come to the Running state.
3. Start each user VM and wait until the associated Router VMs and User VM to move to the Running state.
4. Verify each System VM, Router VM, and User VM to verify that all virtual NICs are migrated to VMWare DVS.

After successfully migrating to VDS, users can optionally remove the standard vSwitches. The physical NICs have already been removed from the standard vSwitches during migration.

For more information on removing standard vSwitches from ESXi host, refer to VMware documentation.

6.7. Storage Preparation for vSphere (iSCSI only)

Use of iSCSI requires preparatory work in vCenter. You must add an iSCSI target and create an iSCSI datastore.

If you are using NFS, skip this section.

6.7.1. Enable iSCSI initiator for ESXi hosts

1. In vCenter, go to hosts and Clusters/Configuration, and click Storage Adapters link.
2. Select iSCSI software adapter and click Properties.
3. Click the Configure... button.
4. Check Enabled to enable the initiator.
5. Click OK to save.

6.7.2. Add iSCSI target

Under the properties dialog, add the iSCSI target info.

Repeat these steps for all ESXi hosts in the cluster.

6.7.3. Create an iSCSI datastore

You should now create a VMFS datastore. Follow these steps to do so:

1. Select Home/Inventory/Datastores.
2. Right click on the datacenter node.
3. Choose Add Datastore... command.
4. Follow the wizard to create a iSCSI datastore.

This procedure should be done on one host in the cluster. It is not necessary to do this on all hosts.

6.7.4. Multipathing for vSphere (Optional)

Storage multipathing on vSphere nodes may be done according to the vSphere installation guide.

6.8. Add Hosts or Configure Clusters (vSphere)

Use vCenter to create a vCenter cluster and add your desired hosts to the cluster. You will later add the entire cluster to CloudPlatform. .

6.9. Creating Custom Roles in vCenter for CloudPlatform

If you are planning to use CloudPlatform to manage virtual machines on VMware vCenter, you must create a user account on vCenter with certain minimum permissions. This user account is used by CloudPlatform to manage the infrastructure resources on the vCenter datacenter.

6.9.1. System Requirements

Before you create your VMs, check your environment meets the minimum requirements as given in the *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Installation Guide*.

6.9.2. Minimum Permissions

The VMware user account you create should have the following minimum permissions at the DataCenter level:

- Manage clusters and hosts
- Manage datastores, disks, and files
- Manage port groups
- Manage dvPort groups
- Manage templates
- Import appliances
- Export templates
- Manage VMs
- Manage snapshot of VM
- Manage custom fields

6.9.3. Creating Roles

1. Create a VMware user account to be used by CloudPlatform.
2. Create the following roles:
 - Global role: This role manages the custom attributes.
 - Datacenter role: This role manages the datacenter.
3. Add the following list of granular permissions to the Global role:

SDK	User Interface
<i>Global.Manage custom attributes</i>	Global > Manage custom attributes

4. Add the following list of granular permissions to the datacenter role:

SDK	User Interface
<i>Global.set custom attributes</i>	Global > Set custom attributes
<i>Datastore.AllocateSpace</i>	Datastore > Allocate space
<i>Datastore.Browse</i>	Datastore > Browse datastore
<i>Datastore.Configure</i>	Datastore > Configure
<i>Datastore.Remove file</i>	Datastore > Remove File
<i>Datastore.FileManagement</i>	Datastore > Low level file operations Datastore > Update virtual machine files
<i>DVPortgroup.Create</i>	dvPort group > Create
<i>DVPortgroup.Modify</i>	dvPort group > Modify
<i>DVPortgroup.Policy</i>	dvPort group > Policy operation
<i>DVPortgroup.Delete</i>	dvPort group > Delete
<i>Folder.Create</i>	Folder > Create folder
<i>Folder.Delete</i>	Folder > Delete folder
<i>Network.Assign</i>	Network > Assign Network
<i>Network.Configure</i>	Network > Configure
<i>Network.Remove</i>	Network > Remove
<i>Resource.HotMigrate</i>	Resource > Migrate powered on virtual machines
<i>Resource.ColdMigrate</i>	Resource > Migrate powered off virtual machines
<i>Resource.AssignVM</i>	Resource > Assign virtual machines to resource pool
<i>Resource.AssignVApp</i>	Resource > Assign vApps to resource pool
<i>Sessions.ValidateSession</i>	Session > Validate session
<i>Host.Configuration</i>	All permissions under Host > Configuration
<i>Host.LocalOperations.Create</i>	Host > Local operations > Create
<i>Host.LocalOperations.Delete</i>	Host > Local operations > Delete
<i>Host.LocalOperations.Reconfigure</i>	Host > Local operations > Reconfigure
<i>ScheduledTask</i>	All permissions under Scheduled task
<i>vApp.Export</i>	vApp > Export
<i>vApp.Import</i>	vApp > Import
<i>vApp.Clone</i>	vApp > Clone
<i>VirtualMachine</i>	All permissions under virtual machine
<i>DVSwitch.PolicyOp</i>	Distributed switch > Policy operations

SDK	User Interface
<i>DVSwitch.PortConfig</i>	Distributed switch > Port configuration
<i>DVSwitch.HostOp</i>	Distributed switch > Host operation
<i>DVSwitch.PortSetting</i>	Distributed switch > Port setting

5. Add the permission to the vCenter object and map the Global role to the user account you created. Do not propagate the rights to the child objects.
6. Add the permission to the datacenter and map the datacenter role to the user account you created. Propagate the rights to the child objects.

(Experimental Feature) Installing LXC for CloudPlatform

Linux Containers (LXC) is an OS-level virtualization technology that employs resource isolation method to run multiple Linux containers on a single host. This technique is different from the hardware virtualization used by KVM and XenServer hypervisors. LXC requires the Linux kernel cgroups functionality which is available starting version 2.6.24. If you want to use the LXC to run guest virtual machines, install LXC on all the hosts in your cloud.

The material in this section doesn't duplicate LXC installation documentation. It provides the CloudPlatform-specific steps that are needed to prepare a LXC host to work with CloudPlatform.

7.1. System Requirements for LXC Hosts

7.1.1. Software Requirements

- RHEL 7

Use the following libvirt and Qemu versions irrespective of the Linux distribution you use:

- libvirt 1.0.0 or higher

For information on libvirt drivers, see [LXC container driver](#)¹.

- Qemu/KVM 1.0 or higher

7.1.2. Hardware Requirements

- Within a single cluster, the hosts must be of the same Linux distribution version.
- All hosts within a cluster must be homogenous. The CPUs must be of the same type, count, and feature flags.
- HVM (Intel-VT or AMD-V enabled)
- 64-bit x86 CPU (more cores results in better performance)
- 4 GB of memory
- A minimum of 1 NIC

7.2. LXC Installation Overview

Because LXC does not have any native systemVMs, KVM is used to run systemVMs. This implies that your host supports both LXC and KVM, thus installation and configuration procedure is identical to the KVM installation.

The procedure for installing an LXC host is:

1. Prepare OS.

¹ <http://libvirt.org/drvlxc.html>

See [Section 7.3, “Preparing the Operating System”](#).

2. Install and configure libvirt.

See [Section 7.4, “Installing and Configuring Libvirt”](#).

3. Install and configure the Agent.

See [Section 7.5, “Installing and Configuring the LXC Agent”](#).

7.2.1. LXC Installation Considerations

- Ensure that you have applied the latest updates to the host.
- On the host, do not run services that are not controlled by CloudPlatform.

7.3. Preparing the Operating System

The OS of the host must be prepared to host the CloudPlatform Agent and run LXC containers.

1. Log in to the OS as root.
2. Check the fully qualified hostname.

```
$ hostname --fqdn
```

This should return a fully-qualified hostname, such as `kvm.test.example.org`. If it does not, edit `/etc/hosts` so that it does.

3. Make sure that the machine can reach the Internet.

```
$ ping www.CloudPlatform.org
```

4. Turn on NTP for time synchronization.

NTP is required to synchronize the clocks of the servers in your cloud. Unsynchronized clocks can cause unexpected problems.

Install NTP:

```
$ yum install ntp
```

5. Repeat all of these steps on every hypervisor host.

7.4. Installing and Configuring Libvirt

CloudPlatform uses libvirt for managing VMs (containers). Therefore it is vital that libvirt is configured correctly. Libvirt is a dependency of CloudPlatform agent and should already be installed.

1. For libvirt to listen on unsecured TCP connections, turn off libvirtd's attempts to use Multicast DNS advertising. Both of these settings are in `/etc/libvirt/libvirtd.conf`:

Set the parameters as follows:

```
listen_tls = 0
```

```
listen_tcp = 1
tcp_port = "16509"
auth_tcp = "none"
mdns_adv = 0
```

2. Turn on `listen_tcp` in `libvirtd.conf` and change the parameters:

Modify `/etc/sysconfig/libvirtd` as follows:

Uncomment the following line:

```
#LIBVIRT_ARGS="--listen"
```

Now, the line reads as follows:

```
libvirtd_opts="-d -l"
```

3. Edit `/etc/libvirt/qemu.conf` to set `vnc_listen` to `0.0.0.0`:

In order to have the VNC console work bind on `0.0.0.0`.

```
vnc_listen = "0.0.0.0"
```

4. Restart `libvirt`:

```
$ service libvirtd restart
```

7.5. Installing and Configuring the LXC Agent

To manage VMs (containers) on an LXC host CloudPlatform uses an Agent. This Agent communicates with the Management Server and controls all the VMs on the host.

1. Check for a fully qualified hostname.

```
# hostname --fqdn
```

This should return a fully qualified hostname, such as `lxc.lab.example.org`. If it does not, edit `/etc/hosts` so that it does.

2. Remove `qemu-kvm`. CloudPlatform provides a patched version.

```
# yum erase qemu-kvm
```

3. If you do not have a Red Hat Network account, register it online from the Red Hat site, then run the following:

```
# subscription-manager register --username <username> --password <password> --auto-attach
# subscription-manager repos --enable=rhel-7-server-rpms
# subscription-manager repos --enable=rhel-7-server-optional-rpms
```

4. Install the CloudPlatform packages. You should have a file in the form of "CloudPlatform-VERSION-N-OSVERSION.tar.gz".

Chapter 7. (Experimental Feature) Installing LXC for CloudPlatform

Untar the file and then run the install.sh script inside it. Replace the file and directory names below with those you are using:

```
# tar xzf CloudPlatform-VERSION-N-OSVERSION.tar.gz
# cd CloudPlatform-VERSION-N-OSVERSION
# ./install.sh
```

You should see a few messages as the installer prepares, followed by a list of choices.

5. Choose “A” to install the Agent software.

```
> A
```

6. When the agent installation is finished, log in to the host as root and run the following commands to start essential services:

```
# service rpcbind start
# service nfs start
# chkconfig nfs on
# chkconfig rpcbind on
```

7. For ICMP rules to work, run the following command after KVM agent is installed.

```
# iptables -D FORWARD -p icmp -j ACCEPT
```

The CloudPlatform Agent is now installed.

If you find you need to stop or start the Agent, use these commands:

```
# service cloudstack-agent start
# service cloudstack-agent stop
```

7.6. Configuring Network Bridges

You can configure bridges by using the native implementation in Linux. To forward traffic to the VMs (containers) a minimum of two bridges are required: public and private. By default these bridges are called cloudbr0 and cloudbr1, but ensure that they are available on each host. Additionally, ensure that the configuration is consistent on all your hypervisors.

7.6.1. Network example

You can configure a network in different ways. In the Basic networking mode you should have two VLANs, one for your private network and one for the public network. Typically, a hypervisor has one NIC (eth0) with three tagged VLANs:

1. VLAN 100 for management of the hypervisor
2. VLAN 200 for public network of the instances (cloudbr0)
3. VLAN 300 for private network of the instances (cloudbr1)

On VLAN 100 provide the host with the IP 100.100.10.11/24 with the gateway 100.100.10.1.

**Note**

The host and the Management Server can be part of different subnets.

7.6.2. Configuring Network Bridges on RHEL

Given below are the examples to create two bridges, cloudbr0 and cloudbr1. All the required packages are automatically installed when you install libvirt. You can continue with configuring the network.

1. Configure eth0:

```
$ vi /etc/sysconfig/network-scripts/ifcfg-eth0
```

The ifcfg-eth0 file should read as follows:

```
DEVICE=eth0
HWADDR=00:04:xx:xx:xx:xx
ONBOOT=yes
HOTPLUG=no
BOOTPROTO=none
TYPE=Ethernet
```

2. Configure the VLAN interfaces:

```
$ vi /etc/sysconfig/network-scripts/ifcfg-eth0.100
```

The **ifcfg-eth0.100** file should read as follows:

```
DEVICE=eth0.100
HWADDR=00:04:xx:xx:xx:xx
ONBOOT=yes
HOTPLUG=no
BOOTPROTO=none
TYPE=Ethernet
VLAN=yes
IPADDR=192.168.42.11
GATEWAY=192.168.42.1
NETMASK=255.255.255.0
```

```
$ vi /etc/sysconfig/network-scripts/ifcfg-eth0.200
```

The **ifcfg-eth0.200** file should read as follows:

```
DEVICE=eth0.200
HWADDR=00:04:xx:xx:xx:xx
ONBOOT=yes
HOTPLUG=no
BOOTPROTO=none
TYPE=Ethernet
VLAN=yes
BRIDGE=cloudbr0
```

```
$ vi /etc/sysconfig/network-scripts/ifcfg-eth0.300
```

The `ifcfg-eth0.300` file should read as follows:

```
DEVICE=eth0.300
HWADDR=00:04:xx:xx:xx:xx
ONBOOT=yes
HOTPLUG=no
BOOTPROTO=none
TYPE=Ethernet
VLAN=yes
BRIDGE=cloudbr1
```

3. Add the bridges on top of the VLAN interfaces:

```
$ vi /etc/sysconfig/network-scripts/ifcfg-cloudbr0
```

4. Configure a plain bridge without an IP for cloudbr0:

```
$ vi /etc/sysconfig/network-scripts/ifcfg-cloudbr0
```

The file should read as follows:

```
DEVICE=cloudbr0
TYPE=Bridge
ONBOOT=yes
BOOTPROTO=none
IPV6INIT=no
IPV6_AUTOCONF=no
DELAY=5
STP=yes
```

5. Configure a plain bridge without an IP for cloudbr1:

```
$ vi /etc/sysconfig/network-scripts/ifcfg-cloudbr1
```

```
DEVICE=cloudbr1
TYPE=Bridge
ONBOOT=yes
BOOTPROTO=none
IPV6INIT=no
IPV6_AUTOCONF=no
DELAY=5
STP=yes
```

6. Restart the network.



Warning

Ensure that you have an alternative option, such as Intelligent Platform Management Interface (IPMI) or Integrated Lights Out (ILO), to reach the machine. If a configuration error occurs and the network stops functioning, you can revert or fix the configuration settings by using any of these methods.

7.7. (Optional)Primary Storage Setup for LXC

CloudPlatform allows administrators to set up shared Primary Storage that uses RADOS Block Devices (RBD). If you want to use RBD-based data disks on LXC, perform the following on each VM (containers):



Note

For Data disk, only RBD storage is supported on LXC. For root volume, use NFS or local storage.

For more information, see [Get Ceph Packages²](#):

1. Configure the Ceph repository.
2. Install the RBD-enabled libvirt 1.1.1 packages.

For more information, see <http://docs.ceph.com/docs/master/rbd/rbd-cloudstack/>

3. Install the Ceph and RBD kernel module.

```
# yum install ceph kmod-rbd
```

4. Verify that the RBD modules are installed.

```
# modprobe rbd
```

5. Copy `ceph.conf` and `ceph admin key` to the `/etc/ceph` directory.

² <http://ceph.com/docs/v0.80.5/install/get-packages/>

Installing Baremetal for CloudPlatform

You can configure Baremetal hosts in CloudPlatform. Baremetal hosts do not run hypervisor software. As part of this installation, you will create a Baremetal template. You can create an instance using this template. The guest operating system is installed using PXE when you create an instance. Baremetal hosts use both basic and advanced networking. A cloud can contain a mix of Baremetal instances and virtual machine instances.

CloudPlatform supports the kick start installation method for RPM-based Linux operating systems on Baremetal hosts in basic/advanced zones. Users can provision a Baremetal host managed by CloudPlatform as long as they have the kickstart file and corresponding OS.

8.1. Baremetal Host System Requirements

Baremetal hosts can run any of the following operating systems. The hardware must meet the requirements published by the OS vendor. Refer to the OS documentation for details. Baremetal kickstart installation is tested on CentOS 5.5, CentOS 6.2, and CentOS 6.3.

Aside from the requirements of the selected OS, baremetal hosts additionally must meet the following requirements:

- All hosts within a cluster must be homogenous. The CPUs must be of the same type, count, and feature flags.
- 32-bit or 64-bit x86 CPU (more cores results in better performance)
- 4 GB of memory
- 36 GB of local disk
- At least 1 NIC
- (Optional) If using security groups, get http://download.cloud.com/releases/4.5.1.0/security_group_agent-1.0-1.noarch.rpm and the packages it depends on: python-cherrypy, ipset, and libmnl. For more information, see *CloudPlatform Getting Started Guide*.
- PXE bootable kernel and initrd for each OS you want to make available for use on Baremetal hosts

8.2. About Baremetal Kickstart Installation

Kickstart installation eliminates manual intervention during OS installation. It uses a text file as a script to automate installation. The kickstart file contains responses to all the user input prompts that are displayed when you install an operating system. With kickstart installation, you can automate the installation of operating system software on large numbers of hosts.

Support for kickstart is provided by the anaconda installer. You can find out more at <http://fedoraproject.org/wiki/Anaconda/Kickstart>. Anaconda is used by the various Linux distributions supported for CloudPlatform Baremetal hosts (see [Section 8.1, “Baremetal Host System Requirements”](#)). A complete description of kickstart files is outside the scope of this documentation. Luckily, there is plentiful documentation available. We have also provided some example kickstart files later in this document.

- Red Hat / CentOS

Docs: http://www.centos.org/docs/6/html/Installation_Guide-en-US/ch-kickstart2.html

Example: [Section 8.2.3, “Example CentOS 6.x Kickstart File”](#)

- Fedora

Docs: http://docs.fedoraproject.org/en-US/Fedora/17/html/Installation_Guide/ch-kickstart2.html

Example: [Section 8.2.3, “Example CentOS 6.x Kickstart File”](#)

- Ubuntu

Docs: <https://help.ubuntu.com/its/installation-guide/i386/automatic-install.html>

Example: [Section 8.2.5, “Example Ubuntu 12.04 Kickstart File”](#)

8.2.1. Limitations of Baremetal Installation

When this feature is used, the following are not supported:

- CloudPlatform storage concepts: primary storage, secondary storage, volume, snapshot
- System VMs: SSVM, CPVM, VR
- Template copy or template download
- VM migration (if host is placed under maintenance mode)
- Live migration
- High availability
- Values from CPU and Memory are not honoured. Use host tag.
- Multiple NICs
- A stopped VM (the OS running on host) can only start on the host it was most recently on
- Console proxy view
- Dedicated resources: cluster and host
- Affinity and anti-affinity group
- Compute offering for Baremetal honours only BaremetalPlanner.
- (Advanced zone) Do not configure description on the Force10 VLAN. If description is set on the VLAN on Force10 switch, the XML document generated by the switch will have a <description> element before <vlan-id> element, which is not recognized by the Force10 switch.

This issue is caused by a defect in the Force10 REST API parser. The parser requires <vlan-id> to be the first item in the <vlan> body, otherwise it throws 'missing element: vlan-id in /ftos:ftos/ftos:interface/ftos:vlan' error.

8.2.2. Prerequisites for Baremetal Host with Kickstart

Follow the steps in all the following sections in order.

1. [Section 8.2.2.1, “Prerequisites for Baremetal Host”](#)
2. [Section 8.2.2.2, “Prerequisites for Kickstart Template Creation ”](#)
3. [Section 8.2.2.3, “Prerequisites for Setting Up Basic Networking ”](#)

4. [Section 8.2.2.4, “Prerequisites for Setting Up Advanced Networking ”](#)

8.2.2.1. Prerequisites for Baremetal Host

- [Section 8.2.2.1.1, “Setting Up IPMI”](#)
- [Section 8.2.2.1.2, “Enabling PXE on the Baremetal Host ”](#)

8.2.2.1.1. Setting Up IPMI

The procedure to access IPMI settings varies depending on the type of hardware. Consult your manufacturer's documentation if you do not already know how to display the IPMI settings screen.

Once you are there, set the following:

- IP address of IPMI NIC
- Netmask
- Gateway
- Username and password for IPMI NIC

CloudPlatform uses `ipmitool` to control the lifecycle of baremetal hosts. By default, `ipmitool` uses the interface `'lan'` to issue `ipmi` commands. Depending on your motherboard, the interface may need to be `'lanplus'`. Consult your hardware documentation to find out if this is the case. If so, modify the script `/usr/lib64/cloud/agent/scripts/util/ipmi.py`.

```
# vi /usr/lib64/cloud/agent/scripts/util/ipmi.py
```

Modify all lines calling `ipmitool`. For example:

```
// Change this:
o = ipmitool("-H", hostname, "-U", username, "-P", password, "chassis", "power", "status")

// To this:
o = ipmitool("-H", hostname, "-I", "lanplus", "-U", username, "-P", password, "chassis",
"power", "status")
```

You do not have to restart the CloudPlatform Management Server for this to take effect.

8.2.2.1.2. Enabling PXE on the Baremetal Host

The Baremetal host needs to use PXE to boot over the network. Access the BIOS setup screen (or equivalent for your hardware) and do the following:

1. Set hard disk as the first priority device in the boot order.
2. Make sure the connected NIC on the Baremetal machine is PXE-enabled.
3. Make a note of the MAC address of the PXE-enabled NIC. You will need it later.

8.2.2.2. Prerequisites for Kickstart Template Creation

- [Section 8.2.2.2.1, “Setting Up a File Server ”](#)
- [Section 8.2.2.2.2, “Create a Baremetal Image ”](#)

8.2.2.2.1. Setting Up a File Server

The kickstart Baremetal image and kickstart file will be stored on an NFS file server. The following steps tell how to set up the NFS server for use with CloudPlatform Baremetal hosts.



Note

This short step-by-step section doesn't attempt to cover all the intricacies of setting up an NFS server. As you go through these steps, keep in mind that this is just a quick checklist. If at any point you find yourself thinking "there ought to be more options available" or "I wonder if wildcards are allowed," please check the Internet and the documentation for the particular type of NFS server you are using.

1. Set up the NFS configuration file `/etc/exports`. This file contains list of entries that describe the shared directories. Each entry specifies which hosts can access the directory, and under what conditions.

The entry for a shared directory follows one of these formats:

```
# Simple listing of hosts, with no options. Default settings are used.
directory host1 host2

# Options are specified to override access permissions and other settings.
directory host1(option1, option2) host2(option3, option4)
```

- `directory` - the directory to be shared; for example, `Share\Baremetal_Backup`
- `host1`, `host2` - clients that have access to the directory, listed by fully qualified domain name, hostname, or IP address
- `option1`, `option2`, etc. - the conditions that restrict the hosts's access to the directory (all are optional)
 - `ro`: read only access to directory
 - `rw`: read and write access to directory
 - `no_root_squash`: root on client have same level of access to files as root on server
 - `no_subtree_check`: only part of volume is exported
 - `sync`: exportfs notify client when file write is complete instead of async notify



Warning

Be careful with space characters in these NFS configuration files. They must be used exactly as shown in the syntax.

2. In `/etc/hosts.deny`, list the clients that are not permitted access to the NFS server by default. For example, you might want to start by denying access to everyone:

```
portmap:ALL
```

In the next step, you'll override this to allow specific hosts.

3. In `/etc/hosts.allow`, list the clients that are allowed to access the NFS server. This list takes precedence over the list in `/etc/hosts.deny`. For example (note the placement of space characters):

```
portmap: host1 , host2
```



Note

Clients that are not listed in either file are allowed access to the NFS server.

4. Verify that NFS is running on the NFS server:

```
# rpcinfo -p
```

The output should show the following services running:

```
portmapper
rquotad
mountd
nfs
nlockmgr
status
```

If so, then you are finished setting up the NFS server.

5. If the services are not already running, you need to start the following NFS daemons on the NFS server:
 - `rpc.portmap`
 - `rpc.mountd`
 - `rpc.nfsd`
 - `rpc.statd`
 - `rpc.lockd`
 - `rpc.rquotad`

8.2.2.2.2. Create a Baremetal Image

Create an image which can be installed on Baremetal hosts later, when Baremetal instances are provisioned in your cloud. On the NFS file server, create a folder and put a PXE bootable kernel and `initrd` in it. For example:

```
# mkdir -p /home/centos63
# cp iso_mount_path_to_centos63/images/pxeboot/{ initrd.img, vmlinuz } /home/centos63
```

For Ubuntu:

```
iso_mount_path_to_ubuntu/install/netboot/ubuntu-installer/amd64/
```

8.2.2.3. Prerequisites for Setting Up Basic Networking

To set up basic networking in Baremetal cloud, you must first install and set up PXE and DHCP servers. You can continue with **Configuring Basic Zone for Baremetal** section in the *CloudPlatform Getting Started Guide*.

Installing the PXE and DHCP Servers

Each Baremetal host must be able to reach a PXE server and a DHCP server. The PXE and DHCP servers must be installed on a separate machine, or a virtual machine, residing in the same L2 network with the baremetal hosts.

1. Log in as root to a host or virtual machine running RHEL or CentOS v6.2 or 6.3.
2. You should have access to a file in the form of "CloudPlatform-VERSION-N-OSVERSION.tar.gz." Copy that file to the machine. The same file is used for either RHEL or CentOS installation.
3. Untar the file and then run the install.sh script inside it. Replace the file and directory names below with those you are using:

```
# tar xzf CloudPlatform-VERSION-N-OSVERSION.tar.gz
# cd CloudPlatform-VERSION-N-OSVERSION
# ./install.sh
```

You should see a few messages as the installer prepares, followed by a list of choices.

4. Choose "B" to install the software that is needed for Baremetal.

```
> B
```

5. Run the Baremetal setup script.

```
# cloudstack-setup-baremetal
```

6. Make note of the TFTP root directory that is displayed by this script. You will need it later.

8.2.2.4. Prerequisites for Setting Up Advanced Networking

- [Section 8.2.2.4.1, "HTTP Rack Configuration Repo"](#)
- [Section 8.2.2.4.2, "Provisioning Completion Notification"](#)
- [Section 8.2.2.4.3, "Advanced Zone with VMware Cluster"](#)
- Force 10 Switch

8.2.2.4.1. HTTP Rack Configuration Repo

To program VLAN for each Baremetal instance, CloudPlatform must understand the network topology at least on rack level. To gather the topology, CloudPlatform needs two sets of information: switch identity and credential and host-switch port mapping. The switch identity and credential allows CloudPlatform to access the switch, usually, using an IP address with username/password combination, which might vary from one switch vendor to another. Host to switch port mapping indicates which host, identified by the host MAC address, connects to which switch port, identified by port number. Both switch identity/credential and switch-host port mapping can be provided by HTTP Rack Configuration Repo. The repo is an ordinary http link whose target is a structured text file. This is the rack configuration text, known as RCT, in JSON format.

Use the `addBaremetalRCT` API to register RCT into CloudPlatform. RCT is globally available to all the Baremetal zones. Except the host MAC, no cloud infrastructure details are required in the RCT. When RCT is changed on the HTTP server, update it instantly by using the `addBaremetalRCT`.

```
{
  "racks": [
    {
      "l2Switch": {
        "ip": "100.100.0.1",
        "username": "root",
        "password": "password",
        "type": "Force10"
      },
      "hosts": [
        {
          "mac": "b8:ac:6f:9a:fa:6b",
          "port": "tengigabytesinterface:1/3/5"
        }
      ]
    }
  ]
}
```

The RCT file is a map with a single entry, `racks`. It is an array that contains a rack definition. A rack definition is comprised of two parts: an `l2Switch` map and a `hosts` array.

- IP: The IP address of the switch management port.
- Username: The username to log in to the switch management port.
- Password: The password associated with the username specified above.
- Type: The switch vendor type. The only vendor type supported is Force10.

The `hosts` array contains the following set of host-mac-port-pair:

- Mac: The mac address of host nic that connects to the switch port.
- Port: The switch port identity.

For Force10, the port identity is in `port type: port id` format. Force10 S4810 has three port types: gigabitethernet, tengigabitethernet, fortyGigE. The port ID is defined as `stackUnit/slot/port`. See the [S4810 manuals](#)¹ for detailed information.

¹ <http://www.force10networks.com/CSPortal20/KnowledgeBase/Documentation.aspx>

8.2.2.4.2. Provisioning Completion Notification

Ensure that the script to generate notification is configured in the kickstart post install section. This is the the script which notifies CloudPlatform of provisioning completion status. When VR receives the notification, it identifies the instance by its IP address. The VR communicate with the Management Server about the provisioning status. Management Server then changes the state from Starting to Running. If no notification is sent to the Management Server or is failed to reach for some reason, for example, a network outage, Management Server shuts down the instance and transmit the Error state. For security reason, keep the credential to send http notification request safely, and never let any customer script run before running the provisioning completion notification script.

To enable provisioning completion notification, ensure that the following script is added in the beginning of the post-install section of the kickstart file:

```
cmdline=`cat /proc/cmdline`
for pair in $cmdline
do
    set -- `echo $pair | tr '=' ' '`
    if [[ $1 == 'ksdevice' ]]; then
        mac=$2
    fi
done
gw=`ip route | grep 'default' | cut -d " " -f 3`
curl "http://$gw:10086/baremetal/provisiondone/$mac"
```

8.2.2.4.3. Advanced Zone with VMware Cluster

Advanced zone with a VMware cluster is up with SSVM and CPVM are up and running. For more information on setting up a VMware infrastructure, see *CloudPlatform Getting Started Guide*.

In Baremetal, only VMware is supported as the VR provider, because it provides management NIC, which is needed for Baremetal to access the internal HTTP server which stores the kickstart file and installation ISO. The VR of Xenserver/KVM uses link local address for inter-communication; therefore, cannot be used for Baremetal.

Continue with **Configuring Advanced Networking in Baremetal** section in the *CloudPlatform Getting Started Guide*.

8.2.3. Example CentOS 6.x Kickstart File

```
# centos 6.x based kickstart file. Disk layout assumes a 4GB sda
install
url --url=http://10.223.110.231/baremetal/centos62/
lang en_US.UTF-8
keyboard us

network --bootproto=dhcp --onboot=yes --hostname=baremetal-test --noipv6

#network --bootproto=dhcp --device=eth0 --onboot=no --noipv6
#network --bootproto=dhcp --device=eth1 --onboot=no --noipv6
#network --bootproto=dhcp --device=eth2 --onboot=yes --hostname=baremetal-test --noipv6
#network --bootproto=dhcp --device=eth3 --onboot=no --noipv6
#network --bootproto=dhcp --device=eth4 --onboot=no --noipv6
#network --bootproto=dhcp --device=eth5 --onboot=no --noipv6

firewall --enabled --port=22:tcp
services --disabled ip6tables
rootpw password
authconfig --enablesshadow --enablemd5
autopart
```

```

selinux --permissive
timezone --utc Europe/London
bootloader --location=mbr --driveorder=sda
clearpart --initlabel --linux --drives=sda
part /boot --fstype ext3 --size=500 --ondisk=sda
part pv.2 --size=1 --grow --ondisk=sda
volgroup vg00 --pesize=32768 pv.2
logvol swap --fstype swap --name=swap00 --vgname=vg00 --size=1024
logvol / --fstype ext3 --name=lv00 --vgname=vg00 --size=2560
#repo --name=epel --baseurl=http://download.fedoraproject.org/pub/epel/6/x86_64/
repo --name=cs-scurity --baseurl=http://nfs1.lab.vmops.com/baremetal/securitygroupagentrepo/
reboot

%packages --ignoremissing
@base
@core
libmnl
wget
cloud-baremetal-securitygroup-agent
%post

#really disable ipv6

echo "install ipv6 /bin/true" > /etc/modprobe.d/blacklist-ipv6.conf

echo "blacklist ipv6" >> /etc/modprobe.d/blacklist-ipv6.conf

yum -y install libmnl

```

8.2.4. Example Fedora 17 Kickstart File

```

# install, not upgrade
install

# Install from a friendly mirror and add updates
url --url=http://10.223.110.231/baremetal/fedora17/

repo --name=updates

# Language and keyboard setup
lang en_US.UTF-8

keyboard us

# Configure DHCP networking w/optional IPv6, firewall on
# network --onboot yes --device eth0 --bootproto dhcp --ipv6 auto --hostname fedora.local
network --bootproto=dhcp --onboot=yes --hostname=baremetal-test --noipv6

firewall --service=ssh

# Set timezone
timezone --utc Etc/UTC

# Authentication
rootpw password

authconfig --enableshadow --passalgo=sha512

autopart

```

```
# SELinux
#selinux --enforcing

selinux --permissive

# Services running at boot
services --enabled network,sshd
services --disabled sendmail

# Disable anything graphical

skipx

text

# Set up the disk

zerombr

clearpart --all

part / --fstype=ext4 --grow --size=1024 --asprimary

#part swap --size=512 # lets do no swap partition for now

bootloader --location=mbr --timeout=5

# Shut down when the kickstart is done

reboot

# Minimal package set

%packages --excludedocs --nobase

@Core

%end

# Nothing for now.

#%post

#%end
```

8.2.5. Example Ubuntu 12.04 Kickstart File

```
#!/var/lib/cobbler/kickstarts/lucid.ks

#System language

lang en_US

#Language modules to install

langsupport en_US

#System keyboard

keyboard us

#System mouse
```

```
mouse

#System timezone

timezone America/New_York

#Root password

rootpw --iscrypted password

#Initial user

user --disabled

#Reboot after installation

reboot

#Use text mode install

text

#Install OS instead of upgrade

install
# Use network installation

url --url=http://10.223.110.231/baremetal/ubuntu1204

#System bootloader configuration

bootloader --location=mbr

#Clear the Master Boot Record

zerombr yes

#Partition clearing information

clearpart --all --initlabel

autopart

#Disk partitioning information

part swap --size 512

part / --fstype ext3 --size 1 --grow

#System authorization information

auth --useshadow --enablemd5

#Network information

network --bootproto=dhcp --device=eth0 --hostname=baremetal-test --noipv6

#Firewall configuration

firewall --enabled --trust=eth0 --ssh

#Do not configure the X Window System

skipx
```

```
%pre

#services

services --enabled=ntpd,nscd,puppet

#Package install information

%packages

ubuntu-standard

man-db

wget

postfix

openssh-server

sysstat

nfs-common

nscd

postfix

quota

ntp

%post
```

8.3. Using Cisco UCS as a Bare Metal Host

(Supported only for use in CloudPlatform zones with basic networking.)

You can provision Cisco UCS server blades into CloudPlatform for use as bare metal hosts. The goal is to enable easy expansion of the cloud by leveraging the programmability of the UCS converged infrastructure and CloudPlatform's knowledge of the cloud architecture and ability to orchestrate. CloudPlatform can automatically understand the UCS environment so CloudPlatform administrators can deploy a bare metal OS on a Cisco UCS.

An overview of the steps involved in using UCS with CloudPlatform:

1. Set up your UCS blades, profile templates, and UCS Manager according to Cisco documentation.
2. Register the UCS Manager with CloudPlatform.
3. Associate a profile with a UCS blade by choosing one of the profile templates created in step 1.
4. Provision the blade as a bare metal host as described in [Section 8.2.2, "Prerequisites for Baremetal Host with Kickstart"](#). Provisioning a Bare Metal Host with Kickstart in the CloudPlatform Installation Guide.

8.3.1. Limitation on Using UCS Manager Profile Templates

You can use profile templates only when first provisioning a blade into CloudPlatform. Updating the template later and modifying the blade's profile is not supported.

8.3.2. Registering a UCS Manager

Register the UCS Manager with CloudPlatform by following these steps:

1. Install the UCS hardware (blades) and UCS Manager according to the vendor's instructions. Make a note of the following information:
 - UCS manager IP address
 - UCS manager username
 - UCS manager password
2. Log in to the CloudPlatform UI as administrator.
3. In the left navigation bar, click Infrastructure, then click Zones.
4. Click the name of a zone where Network Type is Basic.
5. Click the Compute and Storage tab.
6. Scroll down in the diagram and click UCS.
7. Click the Add UCS Manager button.
8. In the dialog box, provide a display name, then the IP address, username, and password that you made a note of in step 1.
9. Click OK.

CloudPlatform will register the UCS Manager, then automatically discover the blades on this UCS Manager and add them to the resource pool.

8.3.3. Associating a Profile with a UCS Blade

Before associating a profile with a UCS blade, you must first do the steps in [Section 8.3.2, "Registering a UCS Manager"](#).

To associate a profile with a UCS blade, start the process by selecting a profile template from a dropdown list in the CloudPlatform UI. The list shows the profile templates that were previously defined on the UCS Manager side. CloudPlatform then creates a profile based on that template and associates the profile with the blade. In the CloudPlatform UI, this is referred to as instantiating and associating a profile. The profile itself is stored on the UCS Manager.

1. Log in to the CloudPlatform UI as administrator.
2. In the left navigation bar, click Infrastructure, then click Zones.
3. Click the name of a zone where you have registered a UCS Manager.
4. Click the Compute and Storage tab.
5. Scroll down in the diagram and click UCS.
6. Click the name of the UCS Manager.
7. Click the Blades tab.

A list is displayed that shows the names of the blades that are installed under the selected manager.

8. In the Actions column, click the Instantiate and Associate icon. 

9. In the dialog, make the following selections:

- Select the name of the template for the profile you want to associate with this blade.

The dropdown list in the dialog box lists the profile templates that are currently defined in the UCS Manager where this blade resides. The list is refreshed any time you add or remove profile templates on the UCS Manager.

- (Optional) In the Profile field, you can provide a user-friendly display name for the profile. This can make it easier to work with the profile later. If you don't provide a value here, CloudPlatform will auto-generate an alphanumeric ID for the profile.
- You might need to wait a few minutes for this operation to finish. The operation might take a long time, depending on the complexity of the setup. The timeout is 60 minutes.

8.3.4. Disassociating a Profile from a UCS Blade

You can remove the association between a profile and a UCS blade. When you do, only the association and, optionally, the profile instance are removed. The profile template remains in place on the UCS Manager.

1. Log in to the CloudPlatform UI as administrator.
2. In the left navigation bar, click Infrastructure, then click Zones.
3. Click the name of a zone where you have registered a UCS Manager.
4. Click the Compute and Storage tab.
5. Scroll down in the diagram and click UCS.
6. Click the name of the UCS Manager.
7. Click the Blades tab.

A list is displayed that shows the names of the blades that are installed under the selected manager.

8. Select the name of a blade that has been associated with a profile.

9. In the Actions column, click the Disassociate Profile icon. 

10. If you want to disassociate the profile and also remove the profile instance from its storage place on UCS Manager, click the Delete Profile checkbox. If you want to disassociate the profile but still keep it in storage, such as to use it in another context, uncheck this box.

In either case, the profile template itself will not be deleted from UCS Manager.

11. Click OK.

You might need to wait a few minutes for this operation to finish. The operation might take a long time, depending on the complexity of the setup. The timeout is 60 minutes.

8.3.5. Synchronizing UCS Manager Changes with CloudPlatform

At any time, CloudPlatform users might directly make changes on the Cisco UCS Manager, and CloudPlatform would not be aware of these changes. For example, users can add or remove blades, and they can associate or dissociate profiles with blades. Periodically, or whenever you become aware that such changes have been made, you can force CloudPlatform to synchronize itself with UCS Manager in order to become aware of any changes that are made manually.

1. Log in to the CloudPlatform UI as administrator.
2. In the left navigation bar, click Infrastructure, then click Zones.
3. Click the name of a zone where you have registered a UCS Manager.
4. Click the Compute and Storage tab.
5. Scroll down in the diagram and click UCS.
6. Click the name of the UCS Manager.
7. Click the Blades tab.
8. Click the Refresh Blades button.
9. Click OK.

You might need to wait a few minutes for this operation to finish. The operation might take a long time, depending on the complexity of the setup. The timeout is 60 minutes.



Note

You might like to click the Refresh button anytime when you are about to make changes to the UCS blade configuration, such as associating or dissociating profiles. This way, you can make sure you are seeing the most up-to-date status of the UCS Manager in the CloudPlatform UI.

Index

B

Baremetal

- advanced zone with VMWare cluster, 68
- create image, 65
- create kickstart template, 63
- enable PXE, 63
- kickstart installation, 61
- prerequisites for advanced networking, 66
- prerequisites for basic networking, 66
- setup file server, 64
- system requirement, 61

H

Hyper-V

- configure physical network, 20
- install role on host, 18
- installation, 17
- preparation checklist, 15
- requirements, 15
- storage, 20

K

KVM

- configure agent, 22
- configure physical network, 23
- configure primary storage, 24
- install agent, 22
- install CloudPlatform agent on host, 22
- requirements, 21
- synchronize time for hosts, 24

L

LXC

- configure agent, 55
- configure libvirt, 54
- configure network bridge, 56
- hardware requirements, 53
- install agent, 55
- install libvirt, 54
- installation, 53
- prepare operating system, 54
- primary storage, 59
- system requirements, 53

V

vSphere

- add hosts, 50
- configure cluster with Nexus 1000v virtual switch, 32
- configure clusters, 50

- configure datacenter with distributed virtual switch, 38
- configure vCenter management network, 32
- configure virtual switch, 30
- create custom roles in vCenter, 50
- installation, 30
- multipathing, 50
- Networking checklist, 29
- physical host networking, 30
- storage preparation, 49
- system requirements, 27
- traffic separation, 30
- vCenter checklist, 29

X

XenServer

- deploy license, 7
- install support package, 7
- installation, 6
- iSCSI multipath setup, 10
- physical networking, 10
- primary storage, 8
- synchronize time, 6
- system requirements, 5
- user credentials, 6

