Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5 Getting Started Guide

Revised January 30, 2015 06:00 pm IST

Citrix CloudPlatform

Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5 Getting Started Guide
Revised January 30, 2015 06:00 pm IST

Author Citrix CloudPlatform

If you have already installed CloudPlatform and you want to learn more about getting started with CloudPlatform, read this document. This document will help you learn about the CloudPlatform user interface, provisioning your cloud infrastructure using the hypervisors that CloudPlatform supports, using accounts in CloudPlatform, and configuring projects in CloudPlatform.

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 2

2. User Interface 3
2.1. Using SSH Keys for Authentication 3
2.1.1. Creating an Instance from a Template that Supports SSH Keys 3
2.1.2. Creating the SSH Keypair 3
2.1.3. Creating an Instance 4
2.1.4. Logging On Using the SSH Keypair 5
2.1.5. Resetting SSH Keys 5

3. Provisioning Your Cloud infrastructure on KVM 7
3.1. Overview of Provisioning Steps 7
3.2. Adding a Zone 7
3.2.1. Creating a Secondary Storage Mount Point for the Zone 8
3.2.2. Adding a New Zone 8
3.3. Adding a Pod 17
3.4. Adding a KVM Cluster 18
3.4.1. Adding a KVM Host 19
3.5. Adding Primary Storage 20
3.6. Adding Secondary Storage 21
3.6.1. Adding an NFS Secondary Storage for Each Zone 22
3.6.2. Configuring S3 Object Store for Secondary Storage 23
3.6.3. Upgrading from NFS to Object Storage 26
3.7. Initialize and Test 27

4. Provisioning Your Cloud infrastructure on XenServer 29
4.1. Overview of Provisioning Steps 29
4.2. Adding a Zone 29
4.2.1. Creating a Secondary Storage Mount Point for the Zone 30
4.2.2. Adding a New Zone 30
4.3. Adding a Pod 39
4.4. Adding a XenServer Cluster 40
4.4.1. Running a Cluster with a Single Host 40
4.4.2. Adding a XenServer Host 40
4.5. Adding Primary Storage 46
4.6. Adding Secondary Storage 47
4.6.1. Adding an NFS Secondary Storage for Each Zone 48
4.6.2. Configuring S3 Object Store for Secondary Storage 49
4.6.3. Upgrading from NFS to Object Storage 52
4.7. Initialize and Test 53

5. Provisioning Your Cloud infrastructure on VMware vSphere 55
5.1. Overview of Provisioning Steps 55
5.2. Adding a Zone 55
5.2.1. Creating a Secondary Storage Mount Point for the Zone 56
5.2.2. Adding a New Zone 56
5.3. Adding a Pod 63
5.4. Add Cluster: vSphere 63
5.4.1. VMware Cluster Size Limit 64
5.4.2. Adding a vSphere Cluster 64
5.4.3. Adding a Host (vSphere) 65

5.5. Adding Primary Storage 66
5.6. Adding Secondary Storage 67
5.6.1. Adding an NFS Secondary Storage for Each Zone 68
5.6.2. Configuring S3 Object Store for Secondary Storage 69
5.6.3. Upgrading from NFS to Object Storage 72
5.7. Initialize and Test 73

6. Provisioning Your Cloud Infrastructure on Hyper-V 75
6.1. Overview of Provisioning Steps 75
6.2. Adding a Zone 75
6.2.1. Create a Secondary Storage Mount Point for the New Zone 76
6.2.2. Adding a New Zone 76
6.3. Adding a Pod 83
6.4. Adding a Hyper-V Cluster 84
6.4.1. Adding a Hyper-V Host 84
6.5. Adding Primary Storage 86
6.6. Adding Secondary Storage 87
6.7. Initialize and Test 88

7. (Experimental Feature) Provisioning Your Cloud Infrastructure on LXC 91
7.1. Overview of Provisioning Steps 91
7.2. Adding a Zone 91
7.2.1. Creating a Secondary Storage Mount Point for the Zone 92
7.2.2. Adding a New Zone 92
7.3. Adding a Pod 100
7.4. Adding an LXC Cluster 101
7.4.1. (Experimental Feature) Adding an LXC Host 102
7.5. Adding Primary Storage 103
7.6. Adding Secondary Storage 105
7.6.1. Adding an NFS Secondary Storage for Each Zone 106
7.6.2. Upgrading from NFS to Object Storage 106
7.7. Initialize and Test 107

8. Provisioning Cloud Infrastructure on Baremetal 109
8.1. Adding a Zone 109
8.1.1. Limitations of Baremetal Deployment 109
8.1.2. Adding a New Zone 109
8.2. Add a Baremetal Cluster 120
8.3. Add a Baremetal Host 120
8.4. Create a Baremetal Template 121
8.5. Create a Baremetal Compute Offering 122
8.6. (Optional) Setting Baremetal Configuration Parameters 123
8.7. Provisioning a Baremetal Instance 124
8.8. Accessing Your Baremetal Instance 124

A. Important Global Configuration Parameters 125
A.1. Important Healthcheck Parameters 127

Index 137

About this Guide

1.1. About the Audience for this Guide

This guide is meant for anyone responsible for configuring and administering the public cloud infrastructure and the private cloud infrastructure of enterprises using 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 featues are inluded in this release:

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

1.4. Additional Information and Help

Troubleshooting articles by the Citrix support team are available in the Citrix Knowledge Center at *support.citrix.com/product/cs/*.

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/cloudsupport*[1] by using the account credentials you received when you purchased your support contract.

[1] http://support.citrix.com/cloudsupport

Chapter 2.

User Interface

2.1. Using SSH Keys for Authentication

In addition to the username and password authentication, CloudPlatform supports using SSH keys to log in to the cloud infrastructure for additional security for your cloud infrastructure. You can use the createSSHKeyPair API to generate the SSH keys.

Because each cloud user has their own ssh key, one cloud user cannot log in to another cloud user's instances unless they share their ssh key files. Using a single SSH key pair, you can manage multiple instances.

2.1.1. Creating an Instance from a Template that Supports SSH Keys

To create an instance from a template that supports SSH keys, do the following:

1. Create a new instance by using the template provided by CloudPlatform.

   For more information on creating a new instance, see the **Creating VMs** section in the *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5 Administration Guide*.

2. Download the **`cloud-set-guest-sshkey`** script file from the following link:

   http://download.cloud.com/templates/4.2/bindir/cloud-set-guest-sshkey.in

3. Copy the file to **`/etc/init.d`**.

4. Give the necessary permissions on the script:

   ```
   chmod +x /etc/init.d/cloud-set-guest-sshkey
   ```

5. Run the script while starting up the operating system:

   ```
   chkconfig --add cloud-set-guest-sshkey
   ```

6. Stop the instance.

2.1.2. Creating the SSH Keypair

You must make a call to the **`createSSHKeyPair`** API method. You can either use the CloudPlatform python API library or the curl commands to make the call to the CloudPlatform API.

For example, make a call from the CloudPlatform server to create a SSH keypair called "keypair-doc" for the administrator account in the root domain:

> **Note**
>
> Ensure that you adjust these values to meet your needs. If you are making the API call from a different server, your URL or port number will be different, and you will need to use the API keys.

1. Run the following curl command:

   ```
   curl --globoff "http://localhost:8080/?command=createSSHKeyPair&name=keypair-
   doc&account=admin&domainid=1"
   ```

   The output appears similar to the following:

   ```
   <?xml version="1.0" encoding="ISO-8859-1"?><createsshkeypairresponse
    cloud-stack-version="3.0.0.20120228045507"><keypair><name>keypair-
   doc</name><fingerprint>f6:77:39:d5:5e:77:02:22:6a:d8:7f:ce:ab:cd:b3:56</
   fingerprint><privatekey>-----BEGIN RSA PRIVATE KEY-----
   MIICXQIBAAKBgQCSydmnQ67jP6lNoXdX3noZjQdrMAWNQZ7y5SrEu4wDxplvhYci
   dXYBeZVwakDVsU2MLGl/K+wefwefwefwefwefJyKJaogMKn7BperPD6n1wIDAQAB
   AoGAdXaJ7uyZKeRDoy6wA0UmF0kSPbMZCR+UTIHNkS/E0/4U+6lhMokmFSHtu
   mfDZ1kGGDYhMsdytjDBztljawfawfeawefawfawfawQQDCjEsoRdgkduTy
   QpbSGDIa11Jsc+XNDx2fgRinDsxXI/zJYXTKRhSl/LIPHBw/brW8vzxhOlSOrwm7
   VvemkkgpAkEAwSeEw394LYZiEVv395ar9MLRVTVLwpo54jC4tsOxQCBlloocK
   lYaocpk0yBqqOUSBawfIiDCuLXSdvBo1Xz5ICTM19vgvEp/+kMuECQBzm
   nVo8b2Gvyagqt/KEQo8wzH2THghZ1qQ1QRhIeJG2aissEacF6bGB2oZ7Igim5L14
   4KR7OeEToyCLC2k+02UCQQCrniSnWKtDVoVqeK/zbB32JhW3Wullv5p5zUEcd
   KfEEuzcCUIxtJYTahJ1pvlFkQ8anpuxjSEDp8x/18bq3
   -----END RSA PRIVATE KEY-----
   </privatekey></keypair></createsshkeypairresponse>
   ```

2. Copy the key data into a file. The file looks like this:

   ```
   -----BEGIN RSA PRIVATE KEY-----
   MIICXQIBAAKBgQCSydmnQ67jP6lNoXdX3noZjQdrMAWNQZ7y5SrEu4wDxplvhYci
   dXYBeZVwakDVsU2MLGl/K+wefwefwefwefwefJyKJaogMKn7BperPD6n1wIDAQAB
   AoGAdXaJ7uyZKeRDoy6wA0UmF0kSPbMZCR+UTIHNkS/E0/4U+6lhMokmFSHtu
   mfDZ1kGGDYhMsdytjDBztljawfawfeawefawfawfawQQDCjEsoRdgkduTy
   QpbSGDIa11Jsc+XNDx2fgRinDsxXI/zJYXTKRhSl/LIPHBw/brW8vzxhOlSOrwm7
   VvemkkgpAkEAwSeEw394LYZiEVv395ar9MLRVTVLwpo54jC4tsOxQCBlloocK
   lYaocpk0yBqqOUSBawfIiDCuLXSdvBo1Xz5ICTM19vgvEp/+kMuECQBzm
   nVo8b2Gvyagqt/KEQo8wzH2THghZ1qQ1QRhIeJG2aissEacF6bGB2oZ7Igim5L14
   4KR7OeEToyCLC2k+02UCQQCrniSnWKtDVoVqeK/zbB32JhW3Wullv5p5zUEcd
   KfEEuzcCUIxtJYTahJ1pvlFkQ8anpuxjSEDp8x/18bq3
   -----END RSA PRIVATE KEY-----
   ```

3. Save the file.

2.1.3. Creating an Instance

Ensure that you use the same SSH key name that you created.

Note

You cannot create the instance by using the GUI at this time and associate the instance with the newly created SSH keypair.

A sample curl command to create a new instance is:

```
curl --globoff http://localhost:<port number>/?
command=deployVirtualMachine&zoneId=1&serviceOfferingId=18727021-7556-4110-9322-
d625b52e0813&templateId=e899c18a-
ce13-4bbf-98a9-625c5026e0b5&securitygroupids=ff03f02f-9e3b-48f8-834d-91b822da40c5&account
 =admin\&domainid=1&keypair=keypair-doc
```

Substitute the template, service offering and security group IDs (if you are using the security group feature) that are in your cloud environment.

2.1.4. Logging On Using the SSH Keypair

To test the successful generation of the your SSH key, verify whether you can log on to the CloudPlatform setup.

For example, on a Linux OS, run the following command:

```
ssh -i ~/.ssh/keypair-doc <ip address>
```

The -i parameter directs the SSH client to use an SSH key found at ~/.ssh/keypair-doc.

2.1.5. Resetting SSH Keys

With the **`resetSSHKeyForVirtualMachine`** API command, a user can set or reset the SSH keypair assigned to a virtual machine. A lost or compromised SSH keypair can be changed, and the user can access the VM by using the new keypair. Just create or register a new keypair, then call **`resetSSHKeyForVirtualMachine`**.

Provisioning Your Cloud infrastructure on KVM

This section describes how to add zones, pods, clusters, hosts, storage, and networks to your cloud using KVM hypervisor.

For conceptual information about zones, pods, clusters, hosts, storage, and networks in CloudPlatform, refer to *CloudPlatform (powered by CloudStack) Version 4.5 Concepts Guide*.

3.1. Overview of Provisioning Steps

After installing Management Server, you can access the CloudPlatform UI and add the compute resources for CloudPlatform to manage.

Then, you can provision the cloud infrastructure, or scale the cloud infrastructure up at any time.

After you complete provisioning the cloud infrastructure, you will have a deployment with the following basic structure:

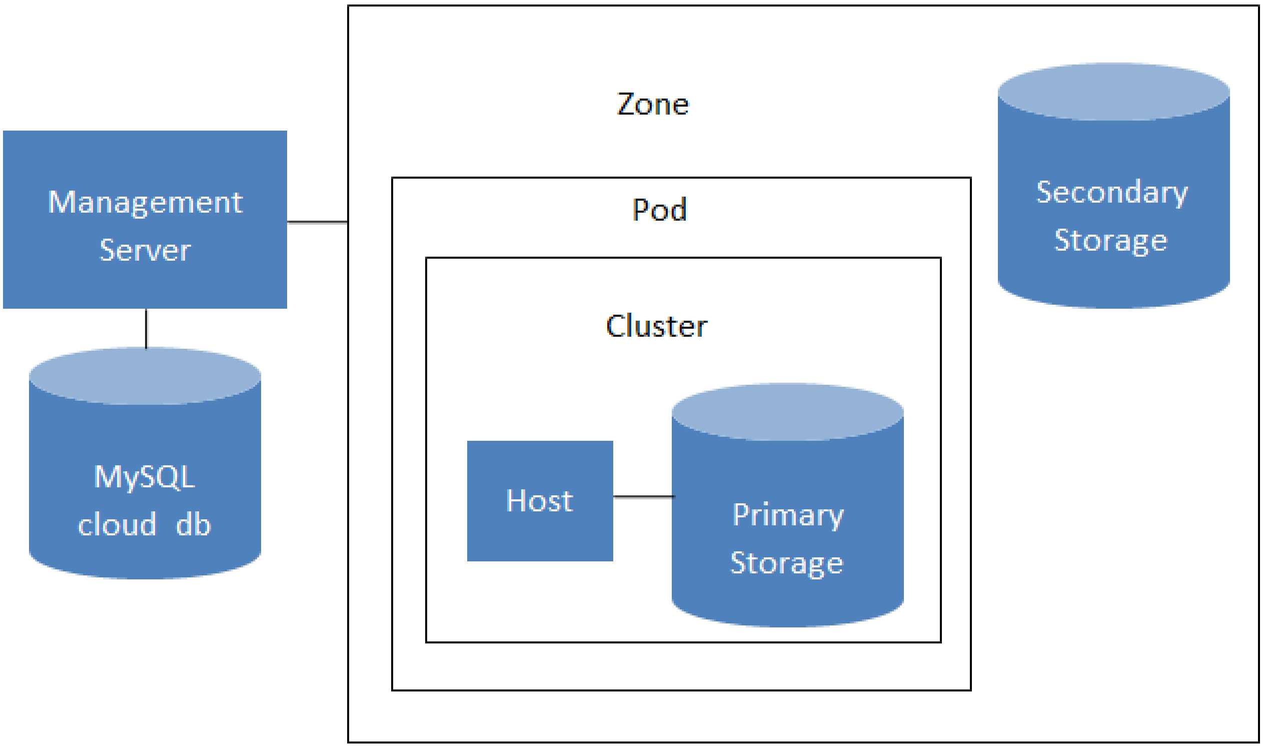

Conceptual view of a basic deployment

For information on adding a region to your cloud infrastructure, refer to **Chapter 3 Adding Regions to Your Cloud Infrastructure (optional)** of *CloudPlatform (powered by CloudStack) Version 4.5 Administration Guide*.

3.2. Adding a Zone

Adding a zone consists of three phases:

- Create a secondary storage mount point for the zone
- Seed the system VM template on the secondary storage.
- Add the zone.

3.2.1. Creating a Secondary Storage Mount Point for the Zone

To ensure that you deploy the latest system VMs in a new zone, you must seed the latest system VM template to the secondary storage of the zone. For this, you must first create a mount point for the secondary storage. Then, you can seed the latest system VM template to the secondary storage.

1. On Management Server, create a mount point for secondary storage. For example:

   ```
   # mkdir -p /mnt/secondary
   ```

2. Mount the secondary storage on your Management Server. Replace NFS server name and NFS share paths in the following example with the server name and path that you use.

   ```
   # mount -t nfs nfsservername:/nfs/share/secondary /mnt/secondary
   ```

3. Seed the secondary storage with the latest template that is used for CloudPlatform system VMs. For more information about seeding the secondary storage, refer to the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Installation Guide*.

   After you seed the secondary storage with the latest system VM template, continue with adding the new zone.

3.2.2. Adding a New Zone

To add a zone, you must first configure the zone’s physical network. Then, you need to add the first pod, cluster, host, primary storage, and secondary storage.

> **Note**
>
> Before you proceed with adding a new zone, you must ensure that you have performed the steps to seed the system VM template.

> **Note**
>
> Citrix strongly recommends using the same type of hypervisors in a zone to avoid operational issues.

1. Log-in to the CloudPlatform UI using the root administrator account.
2. In the left navigation bar, click **Infrastructure**.
3. In the right side panel, under **Zones**, click **View all**.

4. In the next page, click **Add Zone**.

   The **Add zone** wizard panel appears.

5. Select one of the following network types:

   - **Basic**: (For AWS-style networking). Provides a single network where each VM instance is assigned with an IP directly from the network. You can provide guest isolation through layer-3 means such as security groups (IP address source filtering).

   - **Advanced**: Used for more sophisticated network topologies. This network model provides the most flexibility in defining guest networks and providing custom network offerings such as firewall, VPN, or load balancer support.

   For more information, refer to **Chapter 4. Cloud Infrastructure Concepts** of the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Concepts Guide*.

6. Based on the option (Basic or Advanced) that you selected, do one of the following:

   - For **Basic**: *Section 3.2.2.1, “Configuring Basic Zone ”*

   - For **Advanced**: *Section 3.2.2.2, “Advanced Zone Configuration ”*

3.2.2.1. Configuring Basic Zone

1. After you select **Basic** in the **Add Zone** wizard and click **Next**, you need to enter the following details. After you enter the details, click **Next**.

   - **Name**: A name that you can use to identify the zone.

   - **IPv4 DNS 1** and **IPv4 DNS 2**: These are DNS servers for use by guest VMs in the zone. These DNS servers will be accessed via the public network that you will add later. The public IP addresses for the zone must have a route to the DNS server named here.

   - **Internal DNS 1** and **Internal DNS 2**: These are DNS servers for use by system VMs in the zone (these are VMs used by CloudPlatform itself, such as virtual routers, console proxies, and Secondary Storage VMs.) These DNS servers will be accessed via the management traffic network interface of the System VMs. The private IP address you provide for the pods must have a route to the internal DNS server named here.

   - **Hypervisor**: Choose KVM as the hypervisor for the first cluster in the zone.

   - **Network Offering**: Select the network services that will be available on the network for guest VMs.

Network Offering	Description
DefaultSharedNetworkOfferingWithSGService	If you want to enable security groups for guest traffic isolation, choose this. (See Using Security Groups to Control Traffic to VMs.)
DefaultSharedNetworkOffering	If you do not need security groups, choose this.
DefaultSharedNetscalerEIPandELBNetworkOffering	If you have installed a Citrix NetScaler appliance as part of your zone network, and you will be using its Elastic IP and Elastic Load Balancing features, choose this. With the EIP and ELB features, a basic zone with

Network Offering	Description
	security groups enabled can offer 1:1 static NAT and load balancing.
QuickCloudNoServices	TBD

- **Network Domain**: (Optional) If you want to assign a special domain name to the guest VM network, specify the DNS suffix.

- **Public**: A public zone is available to all users. A zone that is not public will be assigned to a particular domain. Only users in that domain will be allowed to create guest VMs in this zone.

2. Select the traffic type that the physical network will carry and click **Next**.

   The traffic types are management, public, guest, and storage traffic. For more information about the types, roll over the icons to display their tool tips. This screen displays some traffic types that are already assigned. To add more, drag and drop traffic types onto the physical network. You can also change the network name, if required.

3. Click the **Edit** button under each traffic type icon that you added to the physical netwok to assign a network traffic label to it. These labels must match the labels you have already defined on the hypervisor host. The **Edit traffic type** dialog appears where you can enter the label and click **OK**.

   These traffic labels will be defined only for the hypervisor selected for the first cluster. For all other hypervisors, the labels can be configured after the zone is created.

4. Click **Next**.

5. (applies to NetScaler only) If you chose the network offering for NetScaler, you get an additional screen where you need to enter information. Provide the requested details to set up the NetScaler, then click **Next**.

   - **IP address:** The NSIP (NetScaler IP) address of the NetScaler device.

   - **Username/Password:** The authentication credentials to access the device. CloudPlatform uses these credentials to access the device.

   - **Type:** NetScaler device type that is being added. It could be NetScaler VPX, NetScaler MPX, or NetScaler SDX. For a comparison of the types, see About Using a NetScaler Load Balancer.

   - **Public interface:** Interface of NetScaler that is configured to be part of the public network.

   - **Private interface:** Interface of NetScaler that is configured to be part of the private network.

   - **Number of retries:** Number of times a command run on the device before considering the operation failed. Default is 2.

   - **Capacity:** Number of guest networks/accounts that will share the NetScaler device.

   - **Dedicated:** Select to indicate that this device is to be dedicated to a single account. When you select **Dedicated**, the value in the **Capacity** field has no significance – implicitly, its value is 1.

6. (applies to NetScaler only) Configure the IP range for public traffic. The IPs in this range will be used for the static NAT capability, which you enabled by selecting the network offering for NetScaler with EIP and ELB. Enter the following details and click **Add**. If required, you can repeat this step to add more IP ranges. Click **Next**.

   - **Gateway:** The gateway in use for these IP addresses.

- **Netmask:** The netmask associated with this IP range.
- **VLAN:** The VLAN that will be used for public traffic.
- **Start IP/End IP:** A range of IP addresses that must be accessible from the Internet and will be allocated for access to guest VMs.

7. In a new zone, CloudPlatform adds the first pod. You can always add more pods later.

   To configure the first pod, enter the following and click **Next**:

   - **Pod Name:** A name for the pod.
   - **Reserved system gateway:** The gateway for the hosts in that pod.
   - **Reserved system netmask:** The network prefix that defines the pod's subnet. Use CIDR notation.
   - **Start/End Reserved System IP:** The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP.

8. Configure the network for guest traffic. Enter the following and click **Next**:

   - **Guest gateway:** The gateway that the guests should use.
   - **Guest netmask:** The netmask in use on the subnet that the guests will use.
   - **Guest start IP/End IP:** Enter the first and the last IP addresses that define a range, which CloudPlatform can assign to guests.
     - Citrix strongly recommends the use of multiple NICs. If multiple NICs are used, they may be in different subnets.
     - If one NIC is used, these IPs must be in the same CIDR as the pod CIDR.

9. In a new pod, CloudPlatform adds the first cluster. You can always add more clusters later. For conceptual information about clusters, in CloudPlatform, refer to *CloudPlatform (powered by CloudStack) Version 4.5 Concepts Guide*

   To configure the first cluster, enter the following and click **Next**:

   - **Hypervisor:** The type of hypervisor software that all hosts in this cluster will run.
   - **Cluster name:** Enter a name for the cluster.

10. In a new cluster, CloudPlatform adds the first host. You can always add more hosts later.
    For conceptual information about hosts in CloudPlatform, refer to *CloudPlatform (powered by CloudStack) Version 4.5 Concepts Guide*

> **Note**
>
> When you add a hypervisor host to CloudPlatform, the host must not have any VMs running on it.

Before you configure the host, you need to install the hypervisor software on the host. You will need to know which version of the hypervisor software version is supported by CloudPlatform and what additional configuration is required to ensure the host will work with CloudPlatform. For more information, refer to the **Chapter 4 Configuring KVM for CloudPlatform** section in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Hypervisor Configuration Guide*.

To configure the first host, enter the following and click **Next**:

- **Host Name:** The DNS name or IP address of the host.
- **Username:** The username is root.
- **Password:** This is the password associated with the user name (from your KVM hypervisor configuration).
- **Host Tags:** (Optional) Any labels that you use to categorize hosts for ease of maintenance. For example, you can set this to the cloud's HA tag (set in the **`ha.tag`** global configuration parameter) if you want this host to be used only for VMs with the "high availability" feature enabled. For more information, refer to the **HA-Enabled Virtual Machines** and the **HA for Hosts** sections in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Administration Guide*.

11. In a new cluster, CloudPlatform adds the first primary storage server. You can always add more servers later. For conceptual information about primary storage in CloudPlatform, refer to *CloudPlatform (powered by CloudStack) Version 4.5 Concepts Guide*.

    To configure the first primary storage server, enter the following and click **Next**:

    - **Name:** The name of the storage device.
    - **Protocol:** For KVM, choose NFS or SharedMountPoint.

3.2.2.2. Advanced Zone Configuration

1. After you select **Advanced** in the **Add Zone** wizard and click **Next**, you can enter the following details. After you enter these details, click **Next**.

   - **Name:** A name for the zone.
   - **IPv4 DNS 1** and **IPv4 DNS 2**: These are DNS servers for use by guest VMs in the zone. These DNS servers will be accessed via the public network that you will add later. The public IP addresses for the zone must have a route to the DNS server named here.
   - **Internal DNS 1** and **Internal DNS 2**: These are DNS servers for use by system VMs in the zone (these are VMs used by CloudPlatform itself, such as virtual routers, console proxies, and Secondary Storage VMs.) These DNS servers will be accessed via the management traffic

network interface of the System VMs. The private IP address you provide for the pods must have a route to the internal DNS server named here.

- **Hypervisor**: Choose **KVM** as the hypervisor for the first cluster in the zone.
- **Network Domain:** (Optional) If you want to assign a special domain name to the guest VM network, specify the DNS suffix.
- **Guest CIDR:** This is the CIDR that describes the IP addresses in use in the guest virtual networks in this zone. For example, 10.1.1.0/24. As a matter of good practice you should set different CIDRs for different zones. This will make it easier to set up VPNs between networks in different zones.
- **Public:** A public zone is available to all users. A zone that is not public will be assigned to a particular domain. Only users in that domain will be allowed to create guest VMs in this zone.

2. Select the traffic type that the physical network will carry and click **Next**.

   The traffic types are management, public, guest, and storage traffic. For more information about the types, roll over the icons to display their tool tips. This screen displays that one network has been already configured. If you have multiple physical networks, you need to add traffic types to them. Drag and drop traffic types onto a greyed-out physical network to make it active. You can move the traffic type icons from one network to another; for example, if the default traffic types shown for Network 1 do not match your actual setup, you can move them down. You can also change the network names if desired.

3. Click the **Edit** button under each traffic type icon that you added to the physical netwok to assign a network traffic label to it. These labels must match the labels you have already defined on the hypervisor host. The **Edit traffic type** dialog appears where you can enter the label and click **OK**.

   These traffic labels will be defined only for the hypervisor selected for the first cluster.

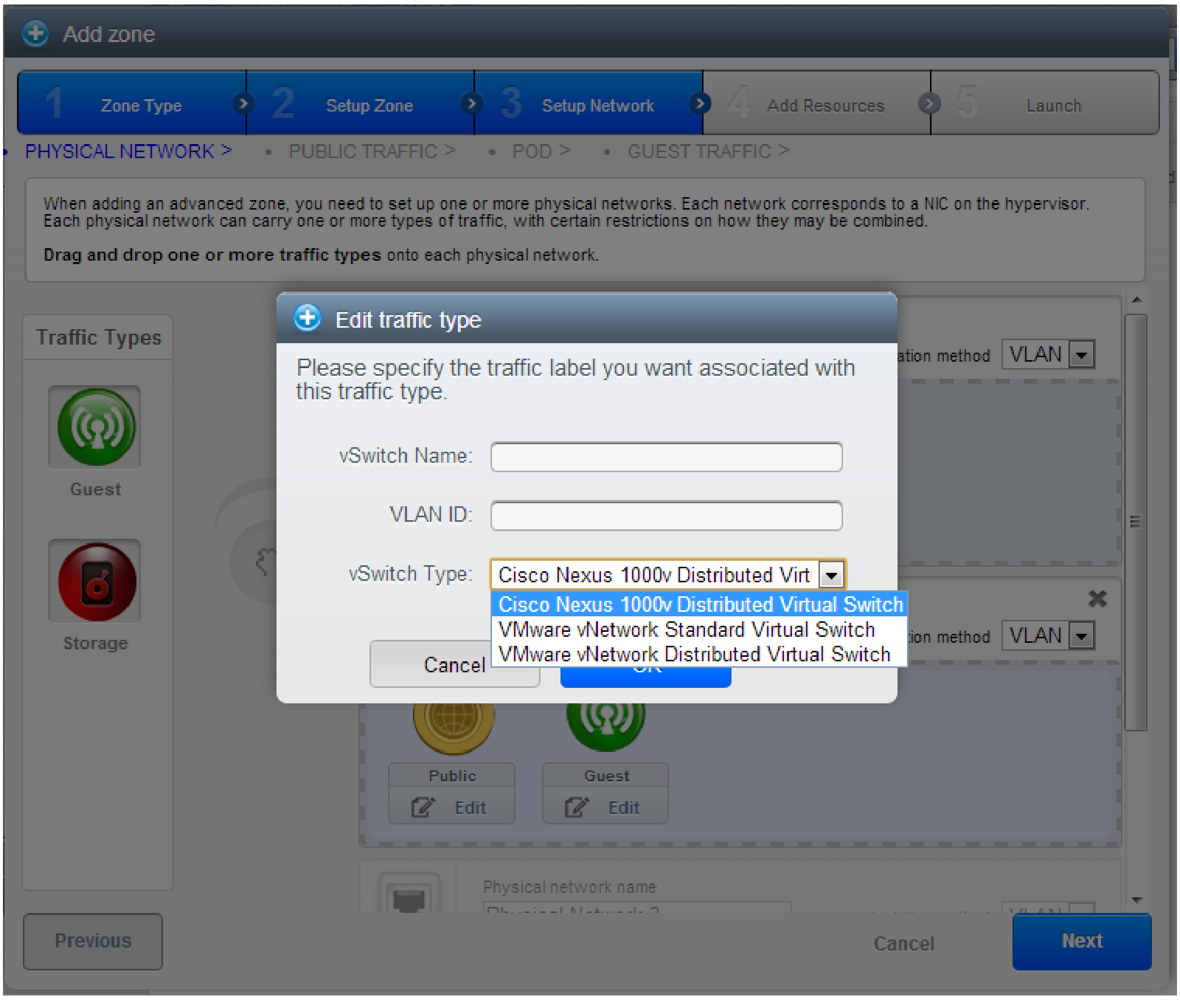

When adding an advanced zone, you need to set up one or more physical networks. Each network corresponds to a NIC on the hypervisor. Each physical network can carry one or more types of traffic, with certain restrictions on how they may be combined.

Drag and drop one or more traffic types onto each physical network.

Edit traffic type

Please specify the traffic label you want associated with this traffic type.

vSwitch Name:

VLAN ID:

vSwitch Type: Cisco Nexus 1000v Distributed Virt

- Cisco Nexus 1000v Distributed Virtual Switch
- VMware vNetwork Standard Virtual Switch
- VMware vNetwork Distributed Virtual Switch

4. Click **Next**.

5. Configure the IP range for the public (Internet) traffic. Enter the following details and click **Add**. If desired, you can repeat this step to add more public Internet IP ranges. Click **Next**.

   - **Gateway:** The gateway in use for these IP addresses.

   - **Netmask:** The netmask associated with this IP range.

   - **VLAN:** The VLAN that will be used for public traffic.

   - **Start IP/End IP:** A range of IP addresses that are assumed to be accessible from the Internet and will be allocated for access to guest networks.

6. In a new zone, CloudPlatform adds the first pod. You can always add more pods later.

   To configure the first pod, enter the following and click **Next**:

   - **Pod Name:** A name that you can use to identify the pod.

   - **Reserved system gateway:** The gateway for the hosts in that pod.

   - **Reserved system netmask:** The network prefix that defines the pod's subnet. Use CIDR notation.

- **Start/End Reserved System IP:** The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP.

- Specify a range of VLAN IDs to carry guest traffic for each physical network and click **Next**.

- In the new pod, CloudPlatform adds the first cluster. You can always add more clusters later.

  To configure the first cluster, enter the following and click **Next**:

  - **Hypervisor:** Select **KVM**.

  - **Cluster name:** Enter a name for the cluster.

- In a new cluster, CloudPlatform adds the first host. You can always add more hosts later.

  > **Note**
  >
  > When you deploy CloudPlatform, the hypervisor host must not have any VMs running on it.

  Before you configure the host, you need to install the hypervisor software on the host. You will need to know which version of the hypervisor software version is supported by CloudPlatform and what additional configuration is required to ensure the host will work with CloudPlatform. For more information, refer to the **Chapter 4 Configuring KVM for CloudPlatform** section in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Hypervisor Configuration Guide*.

  To configure the first host, enter the following, then click **Next**:

  - **Host Name:** The DNS name or IP address of the host.

  - **Username:** Usually root.

  - **Password.** This is the password associated with the user name (from your KVM install).

  - **Host Tags:** (Optional) Any labels that you use to categorize hosts for ease of maintenance. For example, you can set to the cloud's HA tag (set in the **`ha.tag`** global configuration parameter) if you want this host to be used only for VMs with the "high availability" feature enabled. For more information, refer to the **HA-Enabled Virtual Machines** and the **HA for Hosts** sections in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Administration Guide*.

- In a new cluster, CloudPlatform adds the first primary storage server. You can always add more servers later.

  To configure the first primary storage server, enter the following and click **Next**:

  - **Name:** The name of the storage device.

  - **Protocol:** For KVM, choose NFS or SharedMountPoint.

CIFS	• **Server:** The IP address or DNS name of the storage device.

	• **Path:** The exported path from the server. • **Tags (optional):** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
NFS	• **Server:** The IP address or DNS name of the storage device. • **Path:** The exported path from the server. • **Tags (optional):** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
iSCSI	• **Server:** The IP address or DNS name of the storage device. • **Target IQN:** The IQN of the target. For example, iqn.1986-03.com.sun:02:01ec9bb549-1271378984. • **Lun:** The LUN number. For example, 3. • **Tags (optional):** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
preSetup	• **Server:** The IP address or DNS name of the storage device. • **SR Name-Label:** Enter the name-label of the SR that has been set up outside CloudPlatform. • **Tags (optional):** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the

	Zone must also provide primary storage that has tags T1 and T2.
SharedMountPoint	• **Path:** The path on each host that is where this primary storage is mounted. For example, "/mnt/primary". • **Tags (optional):** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
VMFS	• **Server:** The IP address or DNS name of the vCenter server. • **Path:** A combination of the datacenter name and the datastore name. The format is "/" datacenter name "/" datastore name. For example, "/cloud.dc.VM/cluster1datastore". • **Tags (optional):** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.

- In a new zone, CloudPlatform adds the first secondary storage server.

  Before you enter information in the fields on this screen, you need to prepare the secondary storage by setting up NFS shares and installing the latest CloudPlatform System VM template.

  To configure the first secondary storage server on hosts, enter the following and click **Next**:

  - **NFS Server:** The IP address of the server.
  - **Path:** The exported path from the server.

- Click **Launch**.

3.3. Adding a Pod

After you create a new zone, CloudPlatform adds the first pod to it. You can perform the following procedure to add more pods to the zone at anytime.

1. Log-in to the CloudPlatform UI.
2. In the left navigation bar, select **Infrastructure**.
3. In the right-side panel, under **Zones**, click **View all**.

4. In the page that lists the zones that you configured, click the zone where you want to add a pod.

5. Click the **Compute and Storage** tab. At the **Pods** node in the diagram, click **View all**.

6. In the page that lists the pods configured in the zone that you selected, click **Add Pod**.

7. In the **Add Pod** dialog bozx, enter the following details:

    - **Zone:** Select the name of the zone where you want to add the new pod.

    - **Pod name:** The name that you can use to identify the pod.

    - **Reserved system gateway:** The gateway for the hosts in that pod.

    - **Reserved system netmask:** The network prefix that defines the pod's subnet. Use CIDR notation.

    - **Start/End Reserved System IP:** The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP. For more information, see System Reserved IP Addresses.

    - **Dedicate:** Select if you want to dedicate the pod to a specific domain.

    - **Domain:** Select the domain to which you want to dedicate the pod.

    - **Account:** Enter an account name that belongs to the above selected domain so that you can dedicate the pod to this account.

8. Click **OK**.

3.4. Adding a KVM Cluster

You need to tell CloudPlatform about the hosts that it will manage. Hosts exist inside clusters, so before you begin adding hosts to the cloud, you must add at least one cluster.

Before you perform the following steps, you must have installed the hypervisor on the hosts and logged-in to the CloudPlatform UI.

1. In the CloudPlatform UI, in the left navigation bar, click **Infrastructure**.

2. In the right-side panel, under **Zones**, click **View all**.

3. In the page that lists the zones that you have configured, click the zone where you want to add the cluster.

4. In the details page of the zone, click the **Compute and Storage** tab.

5. At the **Clusters** node of the diagram, click **View all**.

6. In the page that list the clusters, click **Add Cluster**.

7. In the **Add Cluster** dialog box, do the following and click **OK**:

    - **Zone Name:** Select the zone where you want to create the cluster.

    - **Hypervisor:** Select KVM as the hypervisor for this cluster.

    - **Pod Name:** Select the pod where you want to create the cluster.

- **Cluster Name:** Enter a name that you can use to identify the cluster.
- **Dedicate**: Select if you want to dedicate the cluster to a specific domain.
- **Domain**: Select the domain to which you want to dedicate the cluster.
- **Account**: Enter an account name that belongs to the domain so that you can dedicate the cluster to this account.

3.4.1. Adding a KVM Host

You can add KVM hosts to a cluster at any time.

3.4.1.1. Requirements for KVM Hosts

Warning

Ensure that the KVM host does not have any VMs running on it before you add it to CloudPlatform.

Configuration requirements:

- Each cluster must contain only hosts with KVM hypervisor.
- Do not keep more than 16 hosts in a cluster.
- If shared mountpoint storage is in use, the administrator must ensure that the new host has all the same mountpoints (with storage mounted) as the other hosts in the cluster.
- Make sure the new host has the same network configuration (guest, private, and public network) as the other hosts in the cluster.

For hardware requirements, refer to *CloudPlatform (powered by Apache CloudStack) Version 4.5 Hypervisor Configuration Guide*.

3.4.1.2. Steps for Adding a KVM Host

1. You must ensure that you have installed the KVM Hypervisor software on the host. You will need to know the version of the KVM software that CloudPlatform supports and the additional configuration that is required to ensure that the host will work with CloudPlatform.
2. Log-in to the CloudPlatform UI using an administrator account.
3. In the left navigation bar, click **Infrastructure**.
4. In the right-side panel, under **Zones**, click **View all**.
5. In the page that lists the zones that are configured with CloudPlatform, click the zone where you want to add the hosts.
6. In the details page of the zone, click the **Compute and Storage** tab.
7. An the **Clusters** node, click **View all**.

8. In the page that lists the clusters available with the zone, click the cluster where you want to add the host.

9. Under the **Details** tab, click the **View Hosts** link.

10. In the page that lists the hosts available with the cluster, click **Add Host**.

11. In the **Add Host** panel, provide the following information:

    - **Zone**: Select the zone where you want to add the host.

    - **Pod**: Select the pod in the zone where you want to add the host.

    - **Cluster**: Select the cluster in the pod where you want to add the host.

    - **Host Name**: The DNS name or IP address of the host.

    - **Username**: Usually root.

    - **Password**: This is the password associated with the user name from your KVM install.

    - **Dedicate**: Select to indicate that this hiost is to be dedicated to a specific domain and account

      Domain: Select the domain to which you want to dedicate the host.

      Account: Select the account that is associated with the domain so that you can dedicate the host to this account.

    - **Host Tags** (Optional). Any labels that you use to categorize hosts for ease of maintenance. For example, you can set to the cloud's HA tag (set in the **`ha.tag`** global configuration parameter) if you want this host to be used only for VMs with the "high availability" feature enabled. For more information, refer to the **HA-Enabled Virtual Machines** and the **HA for Hosts** sections in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Administration Guide*.

    There may be a slight delay while the host is provisioned. It will display automatically in the UI.

12. Repeat the procedure for adding additional hosts.

3.5. Adding Primary Storage

Warning

Ensure that the preallocated storage is empty and contains no data before you use it for primary storage (for example, you must possess an empty SAN volume or an empty NFS share). When you add the storage to CloudPlatform any existing data will be destroyed.

When you create a new zone, the first primary storage is added as part of that procedure. You can add primary storage servers at any time, such as when adding a new cluster or adding more servers to an existing cluster.

1. Log-in to the CloudPlatform UI.

2. In the left navigation bar, click **Infrastructure**.

3. In the right-side panel, under **Zones**, click **View all**.

4. In the page that lists the zones, click the zone where you want to add the primary storage.

5. Click the **Compute and Storage** tab.

6. In the **Primary Storage** node of the diagram, click **View all**.

7. In the page that lists the primary stiorages, click **Add Primary Storage**.

8. In the **Add Primary Storage** dialog box, provide the following information. The information can vary depending on your selection of the protocol in the **Protocol** field.

   - **Scope**: Indicate whether the storage is available to all hosts in the zone or only to hosts in a single cluster.

   - **Pod**: (Visible only if you choose **Cluster** in the **Scope** field.) The pod for the storage device.

   - **Cluster**: (Visible only if you choose **Cluster** in the **Scope** field.) The cluster for the storage device.

   - **Name**: The name of the storage device.

   - **Protocol**: For KVM, select NFS or SharedMountPoint.

   - **Server** (for NFS, iSCSI, SMB/CIFS or PreSetup): The IP address or DNS name of the storage device.

   - **Server** (NFS/iSCSI, PreSetup) IP address or DNS name of the storage device.

   - **Path** (for NFS): In NFS this is the exported path from the server.

   - **Path** (for SharedMountPoint): With KVM this is the path on each host that is where this primary storage is mounted. For example, "/mnt/primary".

   - **Tags** (optional): The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings

   The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.

9. Click **OK**.

3.6. Adding Secondary Storage

> **Note**
>
> Ensure that the storage is empty and contains no data before you use it for secondary storage. When you add the storage to CloudPlatform any existing data will be destroyed.

When you create a new zone, the first secondary storage is added as part of that procedure. You can add secondary storage servers at any time to add more servers to an existing zone.

1. To prepare for the zone-based Secondary Storage, you should have created and mounted an NFS share during Management Server installation.

2. Ensure that you have prepared the system VM template during Management Server installation.

3. Log-in to the CloudPlatform UI using root administrator account.

4. In the left navigation bar, click **Infrastructure**.

5. In the right-side panel, under **Zones**, click **View all**.

6. In the page that lists the zones, click the zone where you want to add the secondary storage.

7. Click the **Compute and Storage** tab.

8. In **Secondary Storage** node of the diagram, click **View all**.

9. In the page that lists the secondary stiorages, click **Add Secondary Storage**.

10. In the **Add Secondary Storage** dialog box, enter the following details:

    - **Name**: Give the storage a descriptive name.

    - **Provider**: Choose the type of storage provider (such as S3, or NFS). NFS can be used for zone-based storage, and the others for region-wide object storage. Depending on which provider you choose, additional fields will appear. Complete all the required fields for your selected provider. For more information, consult the provider's documentation, such as the S3 website.

      Warning

      You can use only a single region-wide object storage account per region. For example, you can not use S3 accounts from different users.

    - **Zone**: The zone where the NFS Secondary Storage is to be located.

    - **Server**: The IP address or DNS name of the storage device.

    - **Path**: The exported path from the server.

    - **NFS server**: The name of the zone's Secondary Storage.

    - **Path**: The path to the zone's Secondary Storage.

3.6.1. Adding an NFS Secondary Storage for Each Zone

You can skip this section if you are upgrading an existing zone from NFS to object storage. You only need to perform the steps below when setting up a new zone that does not yet have its NFS server.

Every zone must have at least one NFS store provisioned; multiple NFS servers are allowed per zone. To provision an NFS Staging Store for a zone:

1. To prepare for the zone-based Secondary Storage, you should have created and mounted an NFS share during Management Server installation.

2. Make sure you prepared the system VM template during Management Server installation.
3. Log-in to the CloudPlatform UI using root administrator account.
4. In the left navigation bar, click **Infrastructure**.
5. In the right-side panel, under **Zones**, click **View all**.
6. In the page that lists the zones, click the zone where you want to add the secondary storage.
7. Click the **Compute and Storage** tab.
8. In **Secondary Storage** node of the diagram, click **View all**.
9. In the **Select View** list box, select **Secondary Storage**.
10. Click the Add NFS Secondary Storage button.
11. Fill out the dialog box fields, then click OK:
    - Zone. The zone where the NFS Secondary Storage is to be located.
    - NFS server. The name of the zone's Secondary Storage.
    - Path. The path to the zone's Secondary Storage.

3.6.2. Configuring S3 Object Store for Secondary Storage

You can configure CloudPlatform to use Amazon S3 Object Store as a secondary storage. S3 Object Store can be used with Amazon Simple Storage Service or any other provider that supports the S3 interface.

1. Make sure you prepared the system VM template during Management Server installation.
2. Log in to the CloudPlatform UI as root administrator.
3. In the left navigation bar, click Infrastructure.
4. In Secondary Storage, click View All.
5. Click Add Secondary Storage.
6. Specify the following:

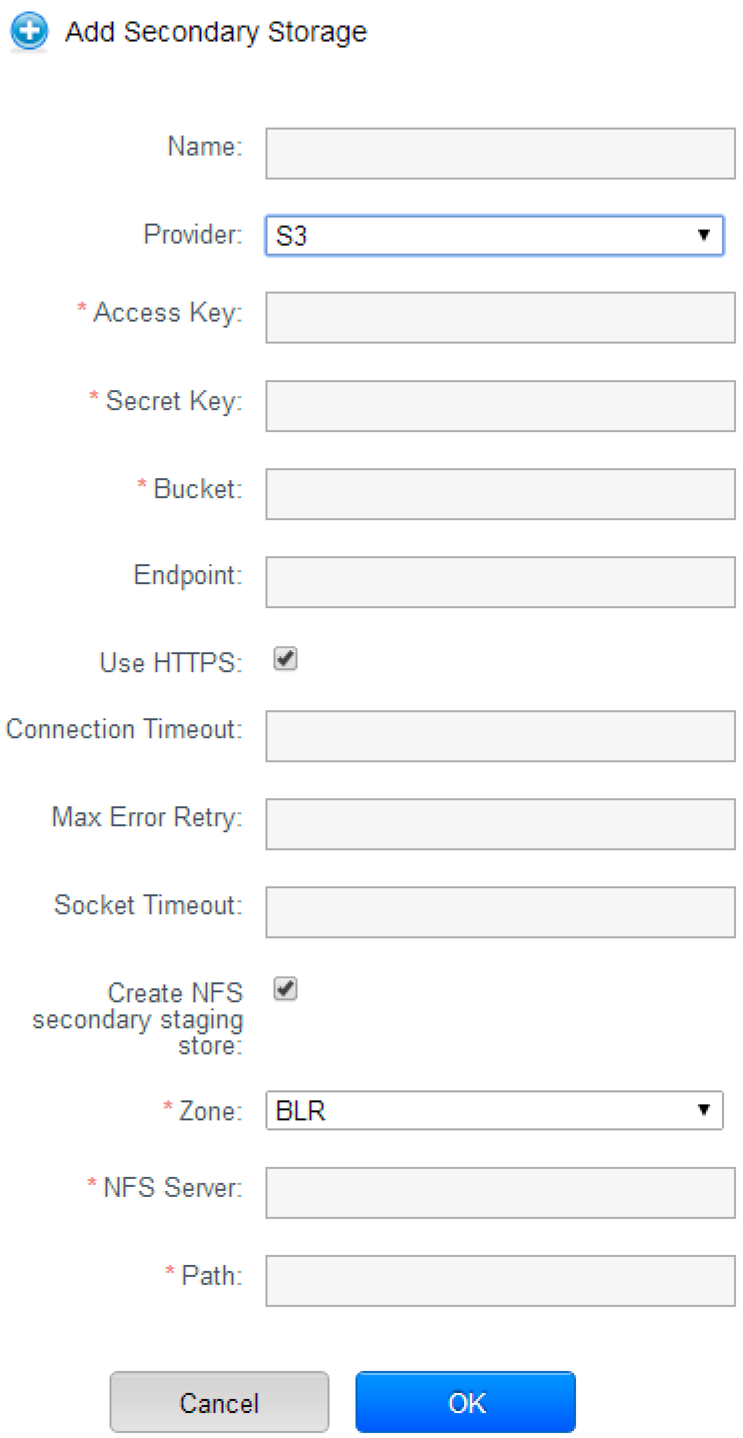

Add Secondary Storage

Name:	
Provider:	S3 ▾
* Access Key:	
* Secret Key:	
* Bucket:	
Endpoint:	
Use HTTPS:	☑
Connection Timeout:	
Max Error Retry:	
Socket Timeout:	
Create NFS secondary staging store:	☑
* Zone:	BLR ▾
* NFS Server:	
* Path:	

- **Name**: Give the storage a descriptive name.
- **Provider**: Select S3 for region-wide object storage. S3 can be used with Amazon Simple Storage Service or any other provider that supports the S3 interface.

Warning

You can use only a single region-wide object storage account per region. For example, you can not use S3 accounts from different users.

- **Access Key**: The Access Key ID of the administrator. These credentials are used to securely sign the requests through a REST of Query API to the CloudPlatform services. You can get this from the admin user Details tab in the Accounts page. Because you include it in each request, the ID is a secret. Each Access Key ID has a Secret Access Key associated with it.

- **Secret Key**: The secret key ID of the administrator. You can get this from the admin user Details tab in the Accounts page.

  This key is just a long string of characters (and not a file) that you use to calculate the digital signature that you include in the request. Your Secret Access Key is a secret, and only you and AWS should have it. Don't e-mail it to anyone, include it any AWS requests, or post it on the AWS Discussion Forums. No authorized person from AWS will ever ask for your Secret Access Key.

- **Bucket** : The container of the objects stored in Amazon S3. Enter the name of the bucket where you store your files.

  Your files are stored as objects in a location called a bucket. When you configure your Amazon S3 bucket as a website, the service delivers the files in your bucket to web browsers as if they were hosted on a web server.

- **End Point**: The IP address or DNS name of the S3 storage server.

  For example: 10.10.29.1:8080, where 8080 is the listening port of the S3 storage server.

- **Use HTTPS**: Specify if you want a secure connection with the S3 storage.

- **Connection Timeout**: The default timeout for creating new connections.

- **Max Error Retry**: The number of retry after service exceptions due to internal errors.

- **Socket Timeout**: The default timeout for reading from a connected socket.

- **Create NFS Secondary Staging Store**: If the zone already contains a secondary staging store, do not select this option. Select if you are upgrading an existing NFS secondary storage into an object storage, as described in *Section 3.6.3, “Upgrading from NFS to Object Storage ”*. Upgrading from NFS to Object Storage in the Installation Guide. In this case, you can skip the rest of the fields described below (Zone, NFS Server, and Path).

- **Zone**: The zone where S3 the Object Store is to be located.

- **Path**: The path to the zone's Secondary Staging Store.

3.6.3. Upgrading from NFS to Object Storage

In an existing zone that is using NFS for secondary storage, you can upgrade the zone to use a region-wide object storage without causing downtime. The existing NFS storage in the zone will be converted to an NFS Staging Store.

After upgrade, all newly created templates, ISOs, volumes, snapshots are moved to the object store. All previously created templates, ISOs, volumes, snapshots are migrated on an on-demand basis based on when they are accessed, rather than as a batch job. Unused objects in the NFS staging store are garbage collected over time.

1. Log-in to the CloudPlatform UI using an administrator account.

2. Fire an admin API to update CloudPlatform to use object storage:

   ```
   http://<MGMTIP>:8096/client/api?command=updateCloudToUseObjectStore&name=<S3
    storage name>&provider=S3&details[0].key=accesskey&details[0].value=<access
    key from .s3cfg file>&details[1].key=secretkey&details[1].value=<secretKey
    from .s3cfg file>&details[2].key=bucket&details[2].value=<bucketname>&details[3].
    key=usehttps&details[3].value=<trueorfalse>&details[4].key=endpoint&details[4].
    value=<S3 server IP:8080>
   ```

   All existing NFS secondary storages has been converted to NFS staging stores for each zone, and your S3 object store specified in the command has been added as a new region-wide secondary storage.

3. Locate the secondary storage that you want to upgrade to object storage.

   Perform either of the following in the Infrastructure page:

   - In Zones, click View All, then locate the desired zone, and select Secondary Storage in the Compute and Storage tab.

   - In Secondary Storage, click View All, then select the desired secondary storage.

Post migration, consider the following:

- For each new snapshot taken after migration, ensure that you take a full snapshot to newly added S3.

  This would help coalesce delta snapshots across NFS and S3 stores. The snapshots taken before migration are pushed to S3 store when you try to use that snapshot by executing createVolumeCmd by passing snapshot id.

- You can deploy VM from templates because they are already in the NFS staging store. A copy of the template is generated in the new S3 object store when ExtractTemplate or CopyTemplate is executed.

- For volume, a copy of the volume is generated in the new S3 object store when ExtractVolume command is executed.

- All the items in the NFS storage is not migrated to S3. Therefore, if you want to completely shut down the NFS storage you have previously used, write your own script to migrate those remaining items to S3.

3.7. Initialize and Test

After you complete all the configuration, CloudPlatform starts the initialization. This can take 30 minutes or more, depending on the speed of your network. When the initialization has completed successfully, the administrator's dashboard should be displayed in the CloudPlatform UI.

1. Verify that the system is ready. In the left navigation bar, click **Templates**. Click on the CentOS 5.5 (64bit) no Gui (KVM) template. Check to be sure that the status is "Download Complete." Do not proceed to the next step until this status is displayed.

2. Go to the Instances tab, and filter by My Instances.

3. Click Add Instance and follow the steps in the wizard.

   a. Choose the zone you just added.

   b. In the template selection, choose the template to use in the VM. If this is a fresh installation, likely only the provided CentOS template is available.

   c. Select a service offering. Be sure that the hardware you have allows starting the selected service offering.

   d. In data disk offering, if desired, add another data disk. This is a second volume that will be available to but not mounted in the guest. A reboot is not required if you have a PV-enabled OS kernel in use.

   e. In default network, choose the primary network for the guest. In a trial installation, you would have only one option here.

   f. Optionally give your VM a name and a group. Use any descriptive text you would like.

   g. Click Launch VM. Your VM will be created and started. It might take some time to download the template and complete the VM startup. You can watch the VMâ€™s progress in the Instances screen.

4. To use the VM, click the View Console button.

   For more information about using VMs, including instructions for how to allow incoming network traffic to the VM, start, stop, and delete VMs, and move a VM from one host to another, see Working With Virtual Machines in the Administratorâ€™s Guide.

Congratulations! You have successfully completed a CloudPlatform Installation.

If you decide to grow your deployment, you can add more hosts, primary storage, zones, pods, and clusters.

Provisioning Your Cloud infrastructure on XenServer

This section describes how to add zones, pods, clusters, hosts, storage, and networks to your cloud using XenServer hypervisor.

For conceptual information about zones, pods, clusters, hosts, storage, and networks in CloudPlatform, refer to *CloudPlatform (powered by CloudStack) Version 4.5 Concepts Guide*.

4.1. Overview of Provisioning Steps

After installing Management Server, you can access the CloudPlatform UI and add the compute resources for CloudPlatform to manage.

Then, you can provision the cloud infrastructure, or scale the cloud infrastructure up at any time.

After you complete provisioning the cloud infrastructure, you will have a deployment with the following basic structure:

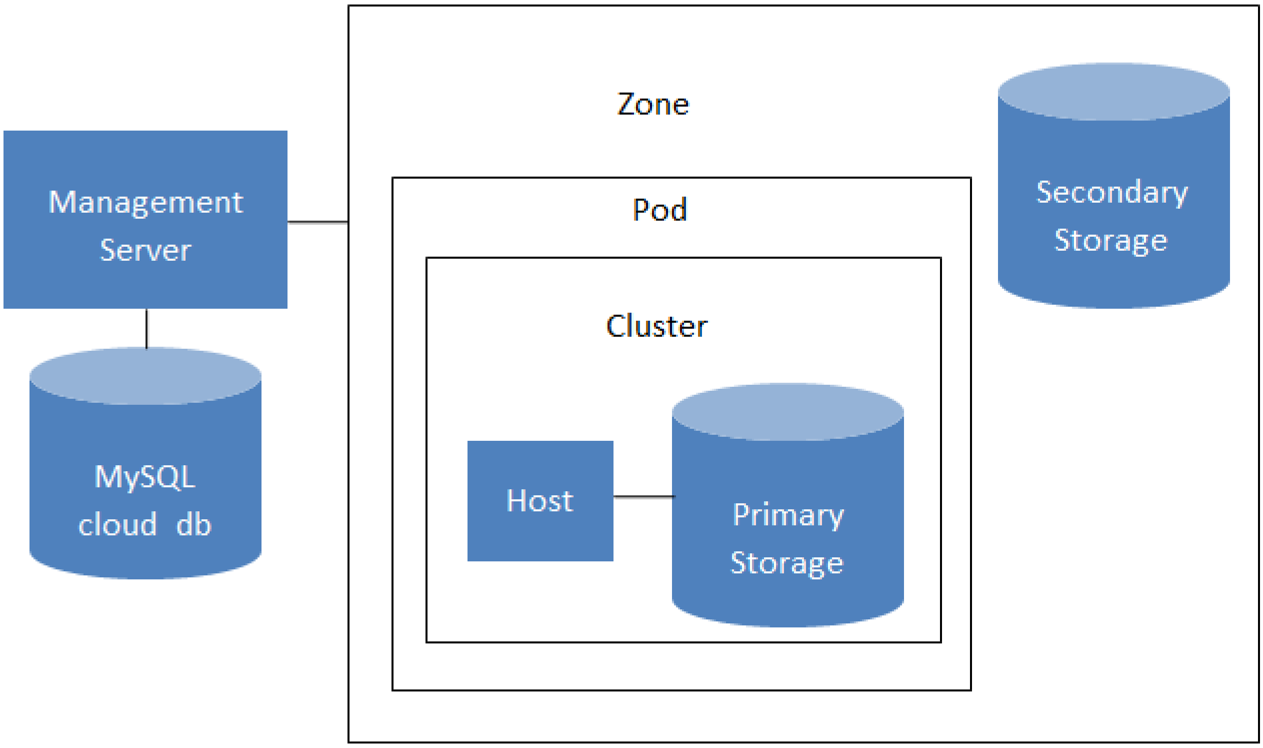

Conceptual view of a basic deployment

For information on adding a region to your cloud infrastructure, refer to **Chapter 3 Adding Regions to Your Cloud Infrastructure (optional)** of *CloudPlatform (powered by CloudStack) Version 4.5 Administration Guide*.

4.2. Adding a Zone

Adding a zone consists of three phases:

- Create a secondary storage mount point for the zone
- Seed the system VM template on the secondary storage.
- Add the zone.

4.2.1. Creating a Secondary Storage Mount Point for the Zone

To ensure that you deploy the latest system VMs in a new zone, you must seed the latest system VM template to the secondary storage of the zone. For this, you must first create a mount point for the secondary storage. Then, you can seed the latest system VM template to the secondary storage.

1. On Management Server, create a mount point for secondary storage. For example:

   ```
   # mkdir -p /mnt/secondary
   ```

2. Mount the secondary storage on your Management Server. Replace NFS server name and NFS share paths in the following example with the server name and path that you use.

   ```
   # mount -t nfs nfsservername:/nfs/share/secondary /mnt/secondary
   ```

3. Seed the secondary storage with the latest template that is used for CloudPlatform system VMs. For more information about seeding the secondary storage, refer to the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Installation Guide*.

   After you seed the secondary storage with the latest system VM template, continue with adding the new zone.

4.2.2. Adding a New Zone

To add a zone, you must first configure the zone's physical network. Then, you need to add the first pod, cluster, host, primary storage, and secondary storage.

> **Note**
>
> Before you proceed with adding a new zone, you must ensure that you have performed the steps to seed the system VM template.

> **Note**
>
> Citrix strongly recommends using the same type of hypervisors in a zone to avoid operational issues.

1. Log-in to the CloudPlatform UI using the root administrator account.
2. In the left navigation bar, click **Infrastructure**.
3. In the right side panel, under **Zones**, click **View all**.

4. In the next page, click **Add Zone**.

   The **Add zone** wizard panel appears.

5. Select one of the following network types:

   - **Basic**: (For AWS-style networking). Provides a single network where each VM instance is assigned with an IP directly from the network. You can provide guest isolation through layer-3 means such as security groups (IP address source filtering).

   - **Advanced**: Used for more sophisticated network topologies. This network model provides the most flexibility in defining guest networks and providing custom network offerings such as firewall, VPN, or load balancer support.

   For more information, refer to **Chapter 4. Cloud Infrastructure Concepts** of the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Concepts Guide*.

6. Based on the option (Basic or Advanced) that you selected, do one of the following:

   - *Section 4.2.2.1, “Basic Zone Configuration ”*

   - *Section 4.2.2.2, “Advanced Zone Configuration ”*

4.2.2.1. Basic Zone Configuration

1. After you select Basic in the Add Zone wizard and click Next, you will be asked to enter the following details. Then click Next.

   - **Name.** A name for the zone.

   - **DNS 1 and 2.** These are DNS servers for use by guest VMs in the zone. These DNS servers will be accessed via the public network you will add later. The public IP addresses for the zone must have a route to the DNS server named here.

   - **Internal DNS 1 and Internal DNS 2.** These are DNS servers for use by system VMs in the zone (these are VMs used by CloudPlatform itself, such as virtual routers, console proxies, and Secondary Storage VMs.) These DNS servers will be accessed via the management traffic network interface of the System VMs. The private IP address you provide for the pods must have a route to the internal DNS server named here.

   - **Hypervisor.** Choose XenServer as the hypervisor for the first cluster in the zone.

   - **Network Offering.** Your choice here determines what network services will be available on the network for guest VMs.

Network Offering	Description
DefaultSharedNetworkOfferingWithSGService	If you want to enable security groups for guest traffic isolation, choose this. (See Using Security Groups to Control Traffic to VMs.)
DefaultSharedNetworkOffering	If you do not need security groups, choose this.
DefaultSharedNetscalerEIPandELBNetworkOffering	If you have installed a Citrix NetScaler appliance as part of your zone network, and you will be using its Elastic IP and Elastic Load Balancing features, choose this. With the EIP and ELB features, a basic zone with

Network Offering	Description
	security groups enabled can offer 1:1 static NAT and load balancing.

   - **Network Domain.** (Optional) If you want to assign a special domain name to the guest VM network, specify the DNS suffix.

   - **Public.** A public zone is available to all users. A zone that is not public will be assigned to a particular domain. Only users in that domain will be allowed to create guest VMs in this zone.

2. Choose which traffic types will be carried by the physical network.

   The traffic types are management, public, guest, and storage traffic. For more information about the types, roll over the icons to display their tool tips, or see Basic Zone Network Traffic Types. This screen starts out with some traffic types already assigned. To add more, drag and drop traffic types onto the network. You can also change the network name if desired.

3. Assign a network traffic label to each traffic type on the physical network. These labels must match the labels you have already defined on the hypervisor host. To assign each label, click the Edit button under the traffic type icon. A popup dialog appears where you can type the label, then click OK.

   These traffic labels will be defined only for the hypervisor selected for the first cluster.

4. Click Next.

5. (NetScaler only) If you chose the network offering for NetScaler, you have an additional screen to fill out. Provide the requested details to set up the NetScaler, then click Next.

   - **IP address.** The NSIP (NetScaler IP) address of the NetScaler device.

   - **Username/Password.** The authentication credentials to access the device. CloudPlatform uses these credentials to access the device.

   - **Type.** NetScaler device type that is being added. It could be NetScaler VPX, NetScaler MPX, or NetScaler SDX. For a comparison of the types, see About Using a NetScaler Load Balancer.

   - **Public interface.** Interface of NetScaler that is configured to be part of the public network.

   - **Private interface.** Interface of NetScaler that is configured to be part of the private network.

   - **Number of retries.** Number of times to attempt a command on the device before considering the operation failed. Default is 2.

   - **Capacity.** Number of guest networks/accounts that will share this NetScaler device.

   - **Dedicated.** When marked as dedicated, this device will be dedicated to a single account. When Dedicated is checked, the value in the Capacity field has no significance – implicitly, its value is 1.

6. (NetScaler only) Configure the IP range for public traffic. The IPs in this range will be used for the static NAT capability which you enabled by selecting the network offering for NetScaler with EIP and ELB. Enter the following details, then click Add. If desired, you can repeat this step to add more IP ranges. When done, click Next.

   - **Gateway.** The gateway in use for these IP addresses.

- **Netmask.** The netmask associated with this IP range.
- **VLAN.** The VLAN that will be used for public traffic.
- **Start IP/End IP.** A range of IP addresses that are assumed to be accessible from the Internet and will be allocated for access to guest VMs.

7. In a new zone, CloudPlatform adds the first pod for you. You can always add more pods later.

   To configure the first pod, enter the following, then click Next:

   - **Pod Name.** A name for the pod.
   - **Reserved system gateway.** The gateway for the hosts in that pod.
   - **Reserved system netmask.** The network prefix that defines the pod's subnet. Use CIDR notation.
   - **Start/End Reserved System IP.** The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP. For more information, see System Reserved IP Addresses.

8. Configure the network for guest traffic. Provide the following, then click Next:

   - **Guest gateway.** The gateway that the guests should use.
   - **Guest netmask.** The netmask in use on the subnet the guests will use.
   - **Guest start IP/End IP.** Enter the first and last IP addresses that define a range that CloudPlatform can assign to guests.
     - We strongly recommend the use of multiple NICs. If multiple NICs are used, they may be in a different subnet.
     - If one NIC is used, these IPs should be in the same CIDR as the pod CIDR.

9. In a new pod, CloudPlatform adds the first cluster for you. You can always add more clusters later.

   To configure the first cluster, enter the following, then click **Next**:

   - **Hypervisor.** Select **XenServer**.
   - **Cluster name.** Enter a name for the cluster. This can be text of your choosing and is not used by CloudPlatform.

10. In a new cluster, CloudPlatform adds the first host for you. You can always add more hosts later. For an overview of what a host is, see About Hosts.

    Note

    When you add a hypervisor host to CloudPlatform, the host must not have any VMs already running.

Before you can configure the host, you need to install the hypervisor software on the host. You will need to know which version of the hypervisor software version is supported by CloudPlatform and what additional configuration is required to ensure the host will work with CloudPlatform. For more information, refer to **Chapter 2 Installing XenServer for CloudPlatform** in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Hypervisor Configuration Guide*

To configure the first host, enter the following, then click Next:

- **Host Name.** The DNS name or IP address of the host.
- **Username.** The username is root.
- **Password.** This is the password for the user named above (from your XenServer install).
- **Host Tags.** (Optional) Any labels that you use to categorize hosts for ease of maintenance. For example, you can set this to the cloud's HA tag (set in the ha.tag global configuration parameter) if you want this host to be used only for VMs with the "high availability" feature enabled. For more information, see HA-Enabled Virtual Machines as well as HA for Hosts.

11. In a new cluster, CloudPlatform adds the first primary storage server for you. You can always add more servers later. For an overview of what primary storage is, see About Primary Storage.

    To configure the first primary storage server, enter the following, then click Next:

    - **Name.** The name of the storage device.
    - **Protocol.** For XenServer, choose either NFS, iSCSI, or PreSetup. The remaining fields in the screen vary depending on what you choose here.

4.2.2.2. Advanced Zone Configuration

1. After you select Advanced in the Add Zone wizard and click Next, you will be asked to enter the following details. Then click Next.

    - **Name.** A name for the zone.
    - **DNS 1 and 2.** These are DNS servers for use by guest VMs in the zone. These DNS servers will be accessed via the public network you will add later. The public IP addresses for the zone must have a route to the DNS server named here.
    - **Internal DNS 1 and Internal DNS 2.** These are DNS servers for use by system VMs in the zone(these are VMs used by CloudPlatform itself, such as virtual routers, console proxies,and Secondary Storage VMs.) These DNS servers will be accessed via the management traffic network interface of the System VMs. The private IP address you provide for the pods must have a route to the internal DNS server named here.
    - **Network Domain.** (Optional) If you want to assign a special domain name to the guest VM network, specify the DNS suffix.
    - **Guest CIDR.** This is the CIDR that describes the IP addresses in use in the guest virtual networks in this zone. For example, 10.1.1.0/24. As a matter of good practice you should set different CIDRs for different zones. This will make it easier to set up VPNs between networks in different zones.
    - **Hypervisor.** Choose **XenServer** as the hypervisor for the first cluster in the zone.

- **Public.** A public zone is available to all users. A zone that is not public will be assigned to a particular domain. Only users in that domain will be allowed to create guest VMs in this zone.

2. Choose which traffic types will be carried by the physical network.

   The traffic types are management, public, guest, and storage traffic. For more information about the types, roll over the icons to display their tool tips. This screen starts out with one network already configured. If you have multiple physical networks, you need to add more. Drag and drop traffic types onto a greyed-out network and it will become active. You can move the traffic icons from one network to another; for example, if the default traffic types shown for Network 1 do not match your actual setup, you can move them down. You can also change the network names if desired.

3. Assign a network traffic label to each traffic type on each physical network. These labels must match the labels you have already defined on the hypervisor host. To assign each label, click the Edit button under the traffic type icon within each physical network. A popup dialog appears where you can type the label, then click OK.

   These traffic labels will be defined only for the hypervisor selected for the first cluster.

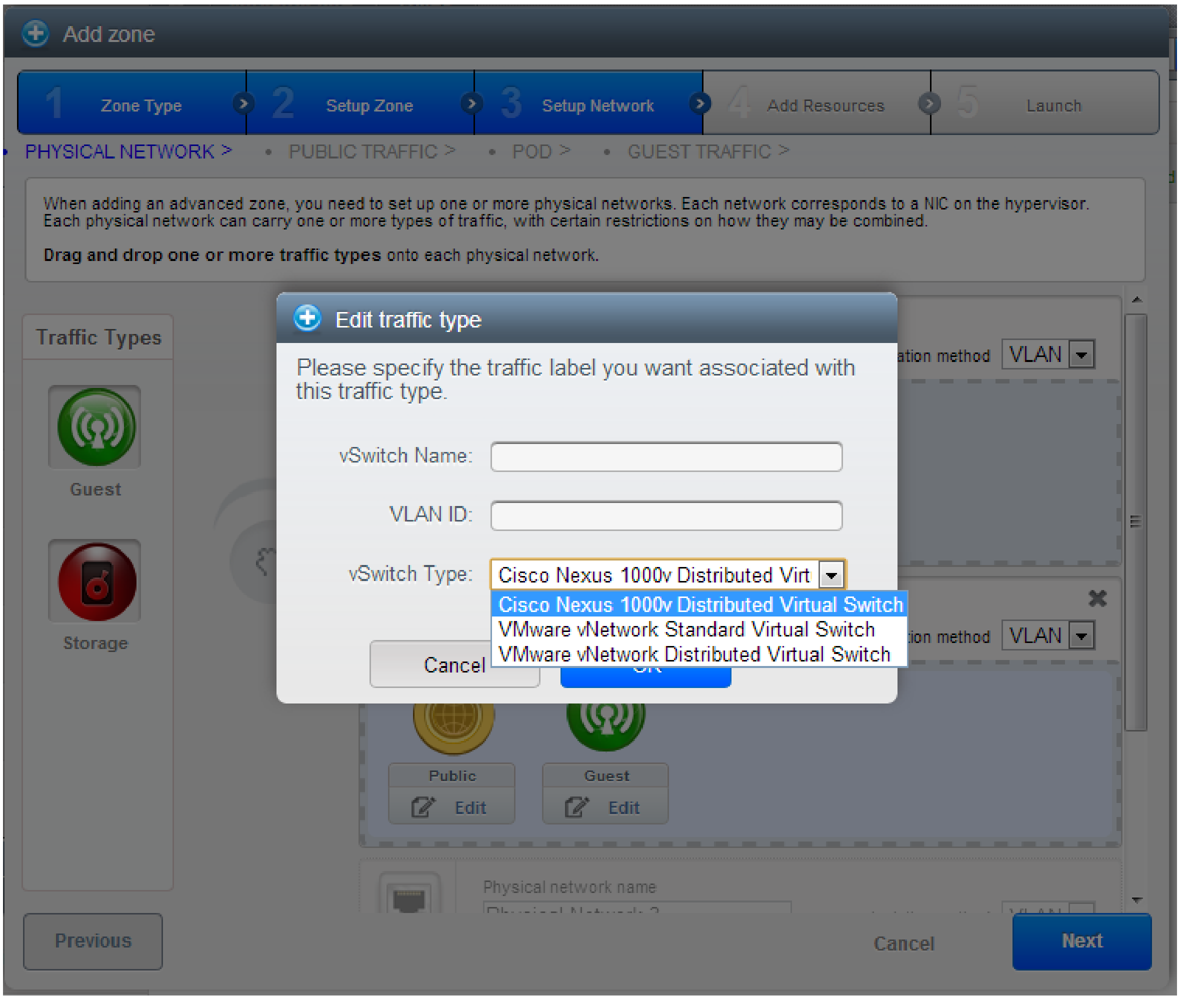

When adding an advanced zone, you need to set up one or more physical networks. Each network corresponds to a NIC on the hypervisor. Each physical network can carry one or more types of traffic, with certain restrictions on how they may be combined.

Drag and drop one or more traffic types onto each physical network.

Edit traffic type

Please specify the traffic label you want associated with this traffic type.

vSwitch Name:

VLAN ID:

vSwitch Type: Cisco Nexus 1000v Distributed Virt

- Cisco Nexus 1000v Distributed Virtual Switch
- VMware vNetwork Standard Virtual Switch
- VMware vNetwork Distributed Virtual Switch

Cancel

4. Click Next.

5. Configure the IP range for public Internet traffic. Enter the following details, then click Add. If desired, you can repeat this step to add more public Internet IP ranges. When done, click Next.

- **Gateway.** The gateway in use for these IP addresses.
- **Netmask.** The netmask associated with this IP range.
- **VLAN.** The VLAN that will be used for public traffic.
- **Start IP/End IP.** A range of IP addresses that are assumed to be accessible from the Internet and will be allocated for access to guest networks.

6. In a new zone, CloudPlatform adds the first pod for you. You can always add more pods later.

   To configure the first pod, enter the following, then click Next:

   - **Pod Name.** A name for the pod.
   - **Reserved system gateway.** The gateway for the hosts in that pod.
   - **Reserved system netmask.** The network prefix that defines the pod's subnet. Use CIDR notation.
   - **Start/End Reserved System IP.** The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP.
   - Specify a range of VLAN IDs to carry guest traffic for each physical network (for more information, see **section: 6.11.1. VLAN Allocation Example** in the *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5 Administration Guide*), then click Next.
   - In a new pod, CloudPlatform adds the first cluster for you. You can always add more clusters later.

     To configure the first cluster, enter the following, then click **Next**:

     - **Hypervisor.** Select **XenServer**.
     - **Cluster name.** Enter a name for the cluster. This can be text of your choosing and is not used by CloudPlatform.
   - In a new cluster, CloudPlatform adds the first host for you. You can always add more hosts later.

     > **Note**
     >
     > When you deploy CloudPlatform, the hypervisor host must not have any VMs already running.

     Before you can configure the host, you need to install the hypervisor software on the host. You will need to know which version of the hypervisor software version is supported by CloudPlatform and what additional configuration is required to ensure the host will work with CloudPlatform. For more information, refer to **Chapter 2 Installing XenServer for CloudPlatform** in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Hypervisor Configuration Guide*.

  To configure the first host, enter the following, then click Next:

  - **Host Name.** The DNS name or IP address of the host.
  - **Username.** Usually root.
  - **Password.** This is the password for the user named above (from your XenServer install).
  - **Host Tags.** (Optional) Any labels that you use to categorize hosts for ease of maintenance. For example, you can set to the cloud's HA tag (set in the ha.tag global configuration parameter) if you want this host to be used only for VMs with the "high availability" feature enabled. For more information, see HA-Enabled Virtual Machines as well as HA for Hosts, both in the Administration Guide.

- In a new cluster, CloudPlatform adds the first primary storage server for you. You can always add more servers later.

  To configure the first primary storage server, enter the following, then click Next:

  - **Name.** The name of the storage device.
  - **Protocol.** For XenServer, choose either NFS, iSCSI, or PreSetup.

CIFS	• **Server.** The IP address or DNS name of the storage device. • **Path.** The exported path from the server. • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
NFS	• **Server.** The IP address or DNS name of the storage device. • **Path.** The exported path from the server. • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
iSCSI	• **Server.** The IP address or DNS name of the storage device. • **Target IQN.** The IQN of the target. For example, iqn.1986-03.com.sun:02:01ec9bb549-1271378984.

	• **Lun.** The LUN number. For example, 3. • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
preSetup	• **Server.** The IP address or DNS name of the storage device. • **SR Name-Label.** Enter the name-label of the SR that has been set up outside CloudPlatform. • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
SharedMountPoint	• **Path.** The path on each host that is where this primary storage is mounted. For example, "/mnt/primary". • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
VMFS	• **Server.** The IP address or DNS name of the vCenter server. • **Path.** A combination of the datacenter name and the datastore name. The format is "/" datacenter name "/" datastore name. For example, "/cloud.dc.VM/cluster1datastore". • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the

8.

	Zone must also provide primary storage that has tags T1 and T2.

- In a new zone, CloudPlatform adds the first secondary storage server for you.

  Before you can fill out this screen, you need to prepare the secondary storage by setting up NFS shares and installing the latest CloudPlatform System VM template.

  To configure the first secondary storage server on XenServer hosts, enter the following, then click Next:

  - **NFS Server.** The IP address of the server.
  - **Path.** The exported path from the server.
- Click Launch.

4.3. Adding a Pod

After you create a new zone, CloudPlatform adds the first pod to it. You can perform the following procedure to add more pods to the zone at anytime.

1. Log-in to the CloudPlatform UI.
2. In the left navigation bar, select **Infrastructure**.
3. In the right-side panel, under **Zones**, click **View all**.
4. In the page that lists the zones that you configured, click the zone where you want to add a pod.
5. Click the **Compute and Storage** tab. At the **Pods** node in the diagram, click **View all**.
6. In the page that lists the pods configured in the zone that you selected, click **Add Pod**.
7. In the **Add Pod** dialog bozx, enter the following details:
   - **Zone:** Select the name of the zone where you want to add the new pod.
   - **Pod name:** The name that you can use to identify the pod.
   - **Reserved system gateway:** The gateway for the hosts in that pod.
   - **Reserved system netmask:** The network prefix that defines the pod's subnet. Use CIDR notation.
   - **Start/End Reserved System IP:** The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP. For more information, see System Reserved IP Addresses.
   - **Dedicate:** Select if you want to dedicate the pod to a specific domain.
   - **Domain:** Select the domain to which you want to dedicate the pod.
   - **Account:** Enter an account name that belongs to the above selected domain so that you can dedicate the pod to this account.
8. Click **OK**.

4.4. Adding a XenServer Cluster

You need to tell CloudPlatform about the hosts that it will manage. Hosts exist inside clusters, so before you begin adding hosts to the cloud, you must add at least one cluster.

Before you perform the following steps, you must have installed the hypervisor on the hosts and logged-in to the CloudPlatform UI.

1. In the CloudPlatform UI, in the left navigation bar, click **Infrastructure**.

2. In the right-side panel, under **Zones**, click **View all**.

3. In the page that lists the zones that you have configured, click the zone where you want to add the cluster.

4. In the details page of the zone, click the **Compute and Storage** tab.

5. At the **Clusters** node of the diagram, click **View all**.

6. In the page that list the clusters, click **Add Cluster**.

7. In the **Add Cluster** dialog box, do the following and click **OK**:

   - **Zone Name:** Select the zone where you want to create the cluster.

   - **Hypervisor:** Select XenServer as the hypervisor for this cluster.

   - **Pod Name:** Select the pod where you want to create the cluster.

   - **Cluster Name:** Enter a name that you can use to identify the cluster.

   - **Dedicate**: Select if you want to dedicate the cluster to a specific domain.

   - **Domain**: Select the domain to which you want to dedicate the cluster.

   - **Account**: Enter an account name that belongs to the domain so that you can dedicate the cluster to this account.

4.4.1. Running a Cluster with a Single Host

With a single host in a cluster, CloudPlatform cannot investigate whether the host is truly down. To detect if XenServer host is down, CloudPlatform performs this: all the hosts are configured to write a timestamp to a file on the shared storage pool. Therefore, if there are two hosts and one shared storage pool in a cluster in CloudPlatform, both the hosts are writing a timestamp to a file by the name of uuid of the host on the shared storage pool. When CloudPlatform finds a host in not responding to PingCommands, it picks another host from the cluster and directs it to investigate if the first host is truly down. As part of the investigation, the second host checks when was the last time the first host wrote a timestamp to its file on the shared pool. If the timestamp is greater than 60 seconds (default value) it concludes that the host is down. In such cases, hosts are placed in Alert state in CloudPlatform.

4.4.2. Adding a XenServer Host

XenServer hosts can be added to a cluster at any time.

4.4.2.1. General Requirements for XenServer Hosts

Consider the following requirements before you add a XenServer host:

- Make sure the hypervisor host does not have any VMs already running before you add it to CloudPlatform.
- Each cluster must contain only hosts with the identical hypervisor.
- Do not add more than 8 hosts in a cluster.
- If network bonding is in use, connect the new host identically to other hosts in the cluster.
- On fresh installation of CloudPlatform, you are recommended to use XenServer 6.2 SP1 Hotfix XS62ESP1004. For host HA support, manually enable Pool HA for XenServer 6.2 SP1 Hotfix XS62ESP1004.
- If you are upgrading to CloudPlatform 4.5, you are recommended to upgrade all existing XenServer clusters to XenServer 6.2 SP1 Hotfix XS62ESP1004. For HA support, manually enable Pool HA for XenServer 6.2 SP1 Hotfix XS62ESP1004.
- CloudPlatform does not support Pool HA for versions prior to XenServer 6.2 SP1 Hotfix XS62ESP1004 release. When master host goes down, CloudPlatform cannot talk to the entire pool, and therefore is not operational from CloudPlatform perspective. In this case, CloudPlatform cannot reflect the right state of the VMs.
- Host HA is manually enabled for XenServer 6.2 SP1 Hotfix XS62ESP1004 release.
- CloudPlatform no longer performs pool join and pool eject. Therefore, procedure for adding and removing hosts in CloudPlatform has been changed. Perform pool join and pool eject by using Citrix XenCenter before you add or delete hosts from CloudPlatform.

For hardware requirements, see the installation section for your hypervisor in the CloudPlatform Installation Guide.

4.4.2.2. Additional Requirements for XenServer Hosts Before Adding to CloudPlatform

4.4.2.2.1. XenServer Version 6.2 SPI Hotfix XS62ESP1004

Before hosts are added to CloudPlatform via UI, perform the following as per your requirement.

4.4.2.2.1.1. Adding Hosts to a New Cluster

1. If you want to add only one host to the cluster, continue to step 5 that provide information on how to enabled pool HA.

2. If you want to add multiple hosts simultaneously, choose one of the hosts to be the master host.

3. Join all the other slave hosts to the master host:

   ```
   # xe pool-join master-address=<masterhost ipaddress> master-username=<username> master-
   password=<password>
   ```

4. Ensure that the hosts are successfully joined the master pool.

   The following command list all the hosts in the pool:

   ```
   # xe host-list
   ```

The following command checks whether pool's master is set to the new master:

```
# xe pool-list params=master
```

5. Enable pool HA by providing the heartbeat Storage Repository.

   For more information, see *Section 4.4.2.2.1.2, "Enabling Pool HA"*.

6. Add the master host to the new cluster in CloudPlatform as explained in *Section 4.4.2.3, "Adding a XenServer Host to CloudPlatform "*

   CloudPlatform automatically adds all the hosts in the pool to the CloudPlatform cluster.

4.4.2.2.1.2. Enabling Pool HA

If you are using XenServer 6.2 SP1 Hotfix XS62ESP1004 clusters, pool HA has to be enabled outside of CloudPlatform.

1. Create a Storage Repository for the XenServer pool.

   Configure a dedicated shared Storage Repository as HA Storage Repository. You can use any shared Storage Repository that XenServer supports. This Storage Repository is not managed by CloudPlatform.

2. To enable XenServer HA, run the following:

```
# xe pool-ha-enable heartbeat-sr-uuids=<sr_uuid>
```

Note

Storage Repository used for heart beat should be a dedicated Storage Repository for HA . Primary storage Storage Repository used by CloudPlatform should not be used for HA purpose.

Do not enable XenServer HA in hosts on versions prior to XenServer 6.2 SP1 Hotfix XS62ESP1004.

4.4.2.2.1.3. Adding a XenServer Host to an Existing CloudPlatform Cluster

When you add a host to a HA-enabled pool, perform the following:

1. Disable Pool HA.

```
# xe pool-ha-disable
```

2. Find the master host:

```
# xe pool-list params=master
```

3. Find the IP of the master host:

```
# xe host-list uuid="host uuid return in #b" params=address
```

4. Join the host to an existing pool:

   ```
   # xe pool-join master-address="master host ip address from #c" master-username=root
    master-password="password for root"
   ```

   Wait 10 minute for the operation to successfully be completed.

5. Enable pool HA by providing the heartbeat Storage Repository:

   ```
   # xe pool-ha-enable heartbeat-sr-uuids="uuid of the HA SR"
   ```

   Note

   When you re-enable pool HA, ensure that you use **`xe pool-ha-enable`** with the *`heartbeat-sr-uuids`* parameter pointing to the correct HA Storage Repository. If the *`heartbeat-sr-uuids`* parameter is skipped, any Storage Repository is randomly picked up for HA, which should be avoided.

6. Continue with *Section 4.4.2.3, “Adding a XenServer Host to CloudPlatform ”*.

   Note

   Adding host to a cluster will fail if the host is not added to XenServer pool.

4.4.2.2.2. XenServer Versions Prior to 6.2 SP1 Hotfix XS62ESP1004

Addition of the first host in a XenServer cluster will succeed. There is no manual steps required in this case. For adding additional hosts to an existing cluster, perform the following before hosts are added to CloudPlatform via UI.

1. Manually join the current host to the existing pool of the first host that have been added to the cluster:

   ```
   # xe pool-join master-address=<masterhost ipaddress> master-username=<username> master-
   password=<password>
   ```

2. Ensure that the hosts have joined the master pool successfully.

   The following command list all the hosts in the pool:

   ```
   # xe host-list
   ```

The following command checks whether pool's master is set to the new master:

```
# xe pool-list params=master
```

3. Continue with *Section 4.4.2.3, "Adding a XenServer Host to CloudPlatform "*.

4.4.2.3. Adding a XenServer Host to CloudPlatform

1. If you have not already done so, install the XenServer software on the host.

   You will need to know which version of the XenServer software version is supported by CloudPlatform and what additional configuration is required to ensure the host will work with CloudPlatform. To find these installation details, see the appropriate section for XenServer in the CloudPlatform Hypervisor Configuration Guide.

2. Review the XenServer requirements listed in this chapter.

3. Depending on host version used, complete the configuration requirements listed in *Section 4.4.2.2, "Additional Requirements for XenServer Hosts Before Adding to CloudPlatform"* and *Section 4.4.2.1, "General Requirements for XenServer Hosts "*.

4. Log-in to the CloudPlatform UI using an administrator account.

5. In the left navigation bar, click **Infrastructure**.

6. In the right-side panel, under **Zones**, click **View all**.

7. In the page that lists the zones that are configured with CloudPlatform, click the zone where you want to add the hosts.

8. In the details page of the zone, click the **Compute and Storage** tab.

9. An the **Clusters** node, click **View all**.

10. In the page that lists the clusters available with the zone, click the cluster where you want to add the host.

11. Under the **Details** tab, click the **View Hosts** link.

12. In the page that lists the hosts available with the cluster, click **Add Host**.

13. In the **Add Host** panel, provide the following information:

    - **Zone**: Select the zone where you want to add the host.

    - **Pod**: Select the pod in the zone where you want to add the host.

    - **Cluster**: Select the cluster in the pod where you want to add the host.

    - **Host Name**: The DNS name or IP address of the host.

    - **Username**: Usually root.

    - **Password**: This is the password associated with the user name from your KVM install.

    - **Dedicate**: Select to indicate that this hiost is to be dedicated to a specific domain and account

      - **Domain**: Select the domain to which you want to dedicate the host.

   - **Account**: Select the account that is associated with the domain so that you can dedicate the host to this account.

 - **Host Tags** (Optional). Any labels that you use to categorize hosts for ease of maintenance. For example, you can set to the cloud's HA tag (set in the **`ha.tag`** global configuration parameter) if you want this host to be used only for VMs with the "high availability" feature enabled. For more information, refer to the **HA-Enabled Virtual Machines** and the **HA for Hosts** sections in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Administration Guide*.

 There may be a slight delay while the host is provisioned. It will display automatically in the UI.

14. Repeat for additional hosts.

4.4.2.4. Recovering When Master Goes Down for XenServer Cluster Versions Prior to 6.2 SP1 Hotfix XS62ESP1004

CloudPlatform does not provide support for pool HA for any versions prior to XenServer 6.2 SP1 Hotfix XS62ESP1004 release. If master hosts on versions prior to XenServer 6.2 SP1 Hotfix XS62ESP1004 go down, CloudPlatform cannot connect to the pool and therefore is not operational from CloudPlatform perspective.

To recover, attempt to bring up the master host. If for some reason, master host cannot be brought up, manually perform the following to designate an existing slave host as master:

1. Make a slave as the master by running the following command on a slave:

   ```
   xe pool-emergency-transition-to-master
   ```

2. Ensure that the new master is effective:

   The following command checks whether pool’s master is set to the new master:

   ```
   # xe pool-list params=master
   ```

3. Point other slaves to the new master by running the following command on the master:

   ```
   # xe pool-recover-slaves
   ```

4. Ensure that all the slaves are pointed to the new master by running the command on all the slaves:

   ```
   # xe pool-list params=master
   ```

4.5. Adding Primary Storage

> **Warning**
>
> When using preallocated storage for primary storage, be sure there is nothing on the storage (ex. you have an empty SAN volume or an empty NFS share). Adding the storage to CloudPlatform will destroy any existing data.

When you create a new zone, the first primary storage is added as part of that procedure. You can add primary storage servers at any time, such as when adding a new cluster or adding more servers to an existing cluster.

1. Log in to the CloudPlatform UI.

2. In the left navigation, choose Infrastructure. In Zones, click View All, then click the zone in which you want to add the primary storage.

3. Click the Compute and Storage tab.

4. In the Primary Storage node of the diagram, click View All.

5. Click Add Primary Storage.

6. Provide the following information in the dialog. The information required varies depending on your choice in Protocol.

   - **Scope**: Indicate whether the storage is available to all hosts in the zone or only to hosts in a single cluster.

   - **Pod**: (Visible only if you choose Cluster in the Scope field.) The pod for the storage device.

   - **Cluster**: (Visible only if you choose Cluster in the Scope field.) The cluster for the storage device.

   - **Name**: The name of the storage device.

   - **Protocol**: Select NFS, iSCSI, or PreSetup.

   - **Server** (for NFS, iSCSI, or PreSetup): The IP address or DNS name of the storage device.

   - **Path** (for NFS): In NFS this is the exported path from the server.

   - **SR Name-Label** (for PreSetup): Enter the name-label of the SR that has been set up outside CloudPlatform.

   - **Target IQN** (for iSCSI): In iSCSI this is the IQN of the target. For example, iqn.1986-03.com.sun:02:01ec9bb549-1271378984

   - **Lun #** (for iSCSI): In iSCSI this is the LUN number. For example, 3.

   - **Tags** (optional): The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings

The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.

7. Click OK.

4.6. Adding Secondary Storage

> **Note**
>
> Be sure there is nothing stored on the server. Adding the server to CloudPlatform will destroy any existing data.

When you create a new zone, the first secondary storage is added as part of that procedure. You can add secondary storage servers at any time to add more servers to an existing zone.

1. To prepare for the zone-based Secondary Storage, you should have created and mounted an NFS share during Management Server installation.

2. Make sure you prepared the system VM template during Management Server installation.

3. Log in to the CloudPlatform UI as root administrator.

4. In the left navigation bar, click Infrastructure.

5. In Secondary Storage, click View All.

6. Click Add Secondary Storage.

7. Fill in the following fields:

   - **Name**: Give the storage a descriptive name.

   - **Provider**: Choose the type of storage provider (such as S3, or NFS). NFS can be used for zone-based storage, and the others for region-wide object storage. Depending on which provider you choose, additional fields will appear. Fill in all the required fields for your selected provider. For more information, consult the provider's documentation, such as the S3 website.

     > **Warning**
     >
     > You can use only a single region-wide object storage account per region. For example, you can not use S3 accounts from different users.

   - **Create NFS Secondary Storage**: Be sure this box is checked, unless the zone already contains a secondary staging store. This option is not required if you are upgrading an existing NFS secondary storage into an object storage, as described in *Section 4.6.3, “Upgrading from*

NFS to Object Storage ". In this case, you can skip the rest of the fields described below (Zone, NFS Server, and Path).

> **Warning**
>
> If you are setting up a new zone, be sure the box is checked. This checkbox and the three fields below it must be filled in. Even when object storage (such as S3) is used as the secondary storage provider, an NFS staging storage in each zone is still required.

- **Zone**: The zone where the NFS Secondary Storage is to be located.
- **Server.** The IP address or DNS name of the storage device.
- **Path.** The exported path from the server.
- **NFS server**: The name of the zone's Secondary Storage.
- **Path**: The path to the zone's Secondary Storage.

4.6.1. Adding an NFS Secondary Storage for Each Zone

You can skip this section if you are upgrading an existing zone from NFS to object storage. You only need to perform the steps below when setting up a new zone that does not yet have its NFS server.

Every zone must have at least one NFS store provisioned; multiple NFS servers are allowed per zone. To provision an NFS Staging Store for a zone:

1. To prepare for the zone-based Secondary Storage, you should have created and mounted an NFS share during Management Server installation.
2. Make sure you prepared the system VM template during Management Server installation.
3. Log in to the CloudPlatform UI as root administrator.
4. In the left navigation bar, click Infrastructure.
5. In Secondary Storage, click View All.
6. In Select View, choose Secondary Storage.
7. Click the Add NFS Secondary Storage button.
8. Fill out the dialog box fields, then click OK:
    - Zone. The zone where the NFS Secondary Storage is to be located.
    - NFS server. The name of the zone's Secondary Storage.
    - Path. The path to the zone's Secondary Storage.

4.6.2. Configuring S3 Object Store for Secondary Storage

You can configure CloudPlatform to use Amazon S3 Object Store as a secondary storage. S3 Object Store can be used with Amazon Simple Storage Service or any other provider that supports the S3 interface.

1. Make sure you prepared the system VM template during Management Server installation.

2. Log in to the CloudPlatform UI as root administrator.

3. In the left navigation bar, click Infrastructure.

4. In Secondary Storage, click View All.

5. Click Add Secondary Storage.

6. Specify the following:

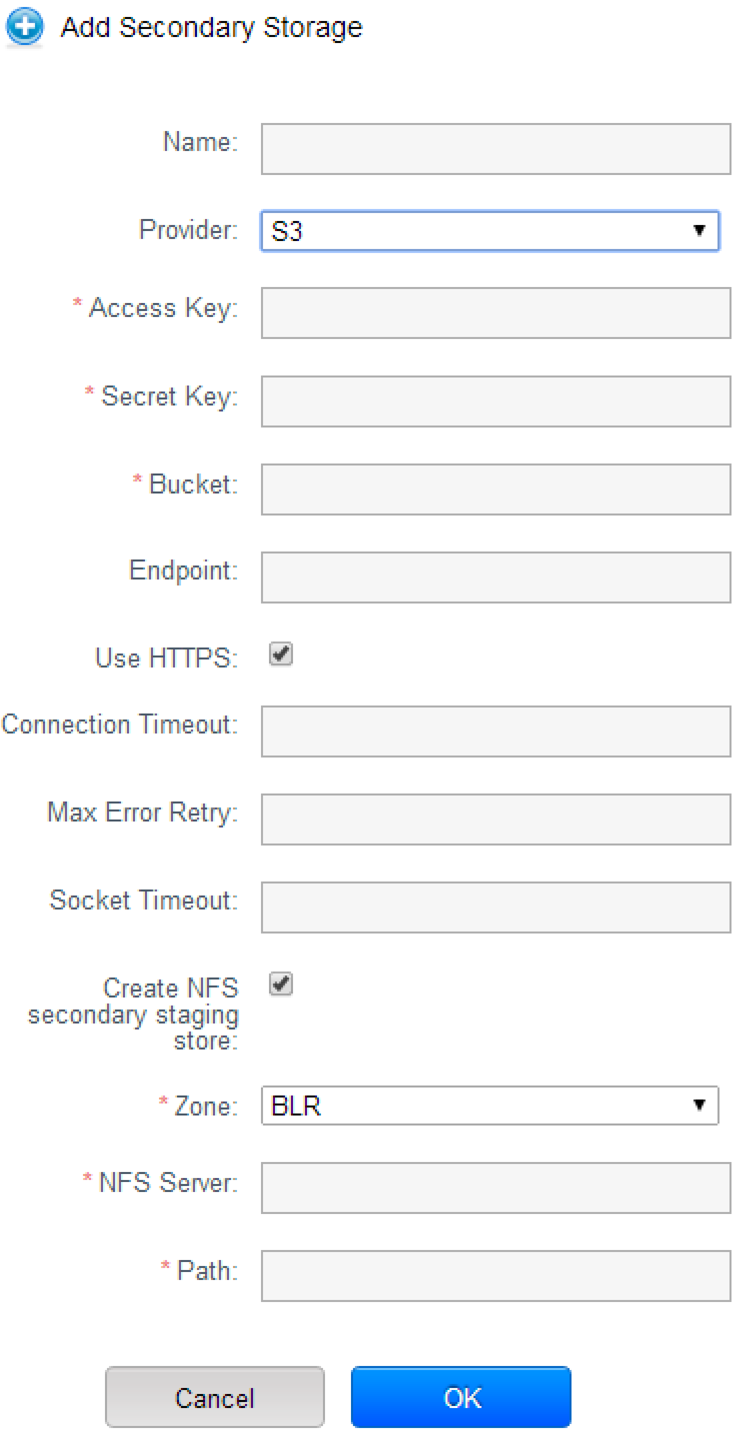

Add Secondary Storage

Name:

Provider: S3

* Access Key:

* Secret Key:

* Bucket:

Endpoint:

Use HTTPS: ☑

Connection Timeout:

Max Error Retry:

Socket Timeout:

Create NFS secondary staging store: ☑

* Zone: BLR

* NFS Server:

* Path:

- **Name**: Give the storage a descriptive name.
- **Provider**: Select S3 for region-wide object storage. S3 can be used with Amazon Simple Storage Service or any other provider that supports the S3 interface.

Warning

You can use only a single region-wide object storage account per region. For example, you can not use S3 accounts from different users.

- **Access Key**: The Access Key ID of the administrator. These credentials are used to securely sign the requests through a REST of Query API to the CloudPlatform services. You can get this from the admin user Details tab in the Accounts page. Because you include it in each request, the ID is a secret. Each Access Key ID has a Secret Access Key associated with it.

- **Secret Key**: The secret key ID of the administrator. You can get this from the admin user Details tab in the Accounts page.

  This key is just a long string of characters (and not a file) that you use to calculate the digital signature that you include in the request. Your Secret Access Key is a secret, and only you and AWS should have it. Don't e-mail it to anyone, include it any AWS requests, or post it on the AWS Discussion Forums. No authorized person from AWS will ever ask for your Secret Access Key.

- **Bucket** : The container of the objects stored in Amazon S3. Enter the name of the bucket where you store your files.

  Your files are stored as objects in a location called a bucket. When you configure your Amazon S3 bucket as a website, the service delivers the files in your bucket to web browsers as if they were hosted on a web server.

- **End Point**: The IP address or DNS name of the S3 storage server.

  For example: 10.10.29.1:8080, where 8080 is the listening port of the S3 storage server.

- **Use HTTPS**: Specify if you want a secure connection with the S3 storage.

- **Connection Timeout**: The default timeout for creating new connections.

- **Max Error Retry**: The number of retry after service exceptions due to internal errors.

- **Socket Timeout**: The default timeout for reading from a connected socket.

- **Create NFS Secondary Staging Store**: If the zone already contains a secondary staging store, do not select this option. Select if you are upgrading an existing NFS secondary storage into an object storage, as described in *Section 4.6.3, “Upgrading from NFS to Object Storage ”*. Upgrading from NFS to Object Storage in the Installation Guide. In this case, you can skip the rest of the fields described below (Zone, NFS Server, and Path).

- **Zone**: The zone where S3 the Object Store is to be located.

- **Path**: The path to the zone's Secondary Staging Store.

4.6.3. Upgrading from NFS to Object Storage

In an existing zone that is using NFS for secondary storage, you can upgrade the zone to use a region-wide object storage without causing downtime. The existing NFS storage in the zone will be converted to an NFS Staging Store.

After upgrade, all newly created templates, ISOs, volumes, snapshots are moved to the object store. All previously created templates, ISOs, volumes, snapshots are migrated on an on-demand basis based on when they are accessed, rather than as a batch job. Unused objects in the NFS staging store are garbage collected over time.

1. Log in as admin to the CloudPlatform UI.

2. Fire an admin API to update CloudPlatform to use object storage:

   ```
   http://<MGMTIP>:8096/client/api?command=updateCloudToUseObjectStore&name=<S3
    storage name>&provider=S3&details[0].key=accesskey&details[0].value=<access
    key from .s3cfg file>&details[1].key=secretkey&details[1].value=<secretKey
    from .s3cfg file>&details[2].key=bucket&details[2].value=<bucketname>&details[3].
    key=usehttps&details[3].value=<trueorfalse>&details[4].key=endpoint&details[4].
    value=<S3 server IP:8080>
   ```

   All existing NFS secondary storages has been converted to NFS staging stores for each zone, and your S3 object store specified in the command has been added as a new region-wide secondary storage.

3. Locate the secondary storage that you want to upgrade to object storage.

   Perform either of the following in the Infrastructure page:

   - In Zones, click View All, then locate the desired zone, and select Secondary Storage in the Compute and Storage tab.

   - In Secondary Storage, click View All, then select the desired secondary storage.

Post migration, consider the following:

- For each new snapshot taken after migration, ensure that you take a full snapshot to newly added S3.

  This would help coalesce delta snapshots across NFS and S3 stores. The snapshots taken before migration are pushed to S3 store when you try to use that snapshot by executing createVolumeCmd by passing snapshot id.

- You can deploy VM from templates because they are already in the NFS staging store. A copy of the template is generated in the new S3 object store when ExtractTemplate or CopyTemplate is executed.

- For volume, a copy of the volume is generated in the new S3 object store when ExtractVolume command is executed.

- All the items in the NFS storage is not migrated to S3. Therefore, if you want to completely shut down the NFS storage you have previously used, write your own script to migrate those remaining items to S3.

4.7. Initialize and Test

After everything is configured, CloudPlatform will perform its initialization. This can take 30 minutes or more, depending on the speed of your network. When the initialization has completed successfully, the administrator's Dashboard should be displayed in the CloudPlatform UI.

1. Verify that the system is ready. In the left navigation bar, select Templates. Click the template. Check to be sure that the status is "Download Complete." Do not proceed to the next step until this status is displayed.

2. Go to the Instances tab, and filter by My Instances.

3. Click Add Instance and follow the steps in the wizard.

    a. Choose the zone you just added.

    b. In the template selection, choose the template to use in the VM. If this is a fresh installation, likely only the provided CentOS template is available.

    c. Select a service offering. Be sure that the hardware you have allows starting the selected service offering.

    d. In data disk offering, if desired, add another data disk. This is a second volume that will be available to but not mounted in the guest. For example, in Linux on XenServer you will see /dev/xvdb in the guest after rebooting the VM. A reboot is not required if you have a PV-enabled OS kernel in use.

    e. In default network, choose the primary network for the guest. In a trial installation, you would have only one option here.

    f. Optionally give your VM a name and a group. Use any descriptive text you would like.

    g. Click Launch VM. Your VM will be created and started. It might take some time to download the template and complete the VM startup. You can watch the VMâ€™s progress in the Instances screen.

4. To use the VM, click the View Console button. >_

    For more information about using VMs, including instructions for how to allow incoming network traffic to the VM, start, stop, and delete VMs, and move a VM from one host to another, see Working With Virtual Machines in the Administratorâ€™s Guide.

Congratulations! You have successfully completed a CloudPlatform Installation.

If you decide to grow your deployment, you can add more hosts, primary storage, zones, pods, and clusters.

Provisioning Your Cloud infrastructure on VMware vSphere

This section describes how to add zones, pods, clusters, hosts, storage, and networks to your cloud using VMware vSphere hypervisor.

For conceptual information about zones, pods, clusters, hosts, storage, and networks in CloudPlatform, refer to *CloudPlatform (powered by CloudStack) Version 4.5 Concepts Guide*.

5.1. Overview of Provisioning Steps

After installing Management Server, you can access the CloudPlatform UI and add the compute resources for CloudPlatform to manage.

Then, you can provision the cloud infrastructure, or scale the cloud infrastructure up at any time.

After you complete provisioning the cloud infrastructure, you will have a deployment with the following basic structure:

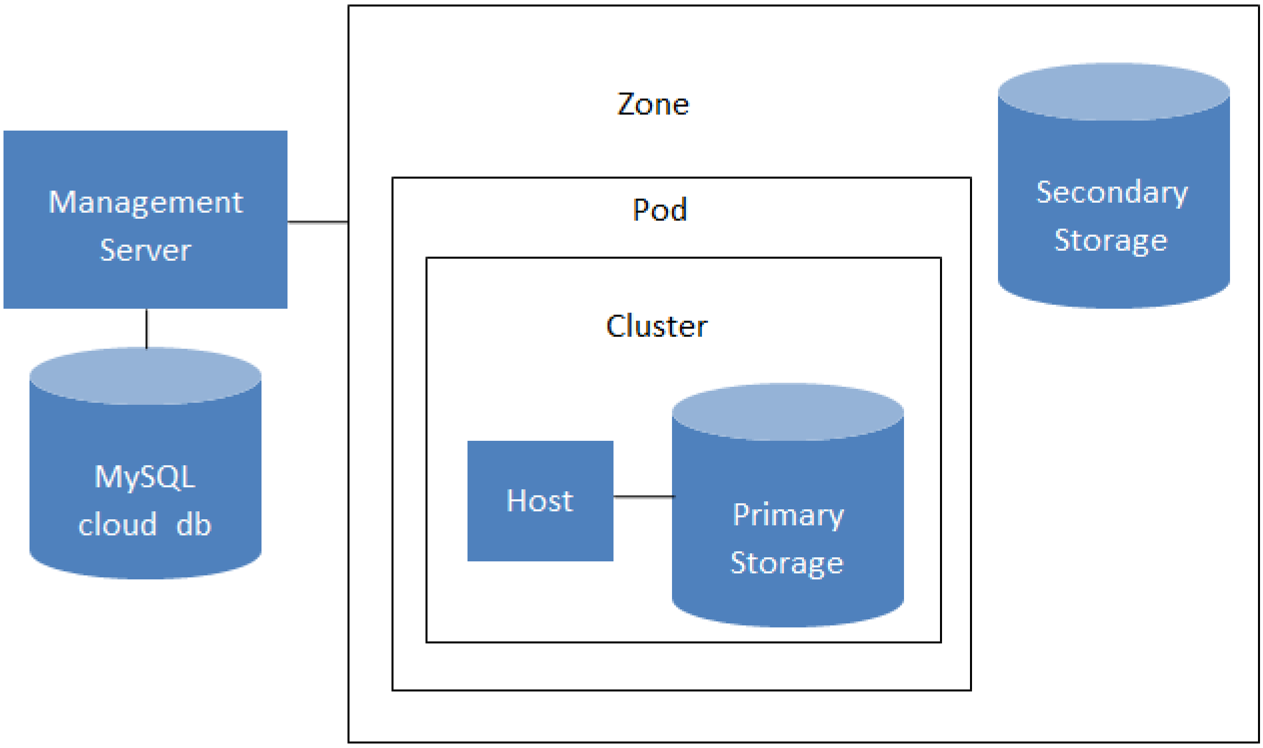

Conceptual view of a basic deployment

For information on adding a region to your cloud infrastructure, refer to **Chapter 3 Adding Regions to Your Cloud Infrastructure (optional)** of *CloudPlatform (powered by CloudStack) Version 4.5 Administration Guide*.

5.2. Adding a Zone

Adding a zone consists of three phases:

- Create a secondary storage mount point for the zone
- Seed the system VM template on the secondary storage.
- Add the zone.

5.2.1. Creating a Secondary Storage Mount Point for the Zone

To ensure that you deploy the latest system VMs in a new zone, you must seed the latest system VM template to the secondary storage of the zone. For this, you must first create a mount point for the secondary storage. Then, you can seed the latest system VM template to the secondary storage.

1. On Management Server, create a mount point for secondary storage. For example:

   ```
   # mkdir -p /mnt/secondary
   ```

2. Mount the secondary storage on your Management Server. Replace NFS server name and NFS share paths in the following example with the server name and path that you use.

   ```
   # mount -t nfs nfsservername:/nfs/share/secondary /mnt/secondary
   ```

3. Seed the secondary storage with the latest template that is used for CloudPlatform system VMs. For more information about seeding the secondary storage, refer to the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Installation Guide*.

   After you seed the secondary storage with the latest system VM template, continue with adding the new zone.

5.2.2. Adding a New Zone

To add a zone, you must first configure the zone's physical network. Then, you need to add the first pod, cluster, host, primary storage, and secondary storage.

Note

Before you proceed with adding a new zone, you must ensure that you have performed the steps to seed the system VM template.

Note

Citrix strongly recommends using the same type of hypervisors in a zone to avoid operational issues.

1. Log-in to the CloudPlatform UI using the root administrator account.
2. In the left navigation bar, click **Infrastructure**.
3. In the right side panel, under **Zones**, click **View all**.

4. In the next page, click **Add Zone**.

   The **Add zone** wizard panel appears.

5. Select **Advanced** as the network type.

   The Advanced network type is used for more sophisticated network topologies. This network model provides the most flexibility in defining guest networks and providing custom network offerings such as firewall, VPN, or load balancer support.

   > **Note**
   >
   > The **Basic** zone type is not supported on VMWare ESXi hosts.

   For more information, refer to **Chapter 4. Cloud Infrastructure Concepts** of the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Concepts Guide*.

6. After you select the Adsvanced network type, proceed as described in *Section 5.2.2.1, “Advanced Zone Configuration ”*

5.2.2.1. Advanced Zone Configuration

1. After you select Advanced in the Add Zone wizard and click Next, you will be asked to enter the following details. Then click Next.

   - **Name.** A name for the zone.

   - **DNS 1 and 2.** These are DNS servers for use by guest VMs in the zone. These DNS servers will be accessed via the public network you will add later. The public IP addresses for the zone must have a route to the DNS server named here.

   - **Internal DNS 1 and Internal DNS 2.** These are DNS servers for use by system VMs in the zone(these are VMs used by CloudPlatform itself, such as virtual routers, console proxies,and Secondary Storage VMs.) These DNS servers will be accessed via the management traffic network interface of the System VMs. The private IP address you provide for the pods must have a route to the internal DNS server named here.

   - **Network Domain.** (Optional) If you want to assign a special domain name to the guest VM network, specify the DNS suffix.

   - **Guest CIDR.** This is the CIDR that describes the IP addresses in use in the guest virtual networks in this zone. For example, 10.1.1.0/24. As a matter of good practice you should set different CIDRs for different zones. This will make it easier to set up VPNs between networks in different zones.

   - **Hypervisor.** Choose **vSphere** as the hypervisor for the first cluster in the zone.

   - **Public.** A public zone is available to all users. A zone that is not public will be assigned to a particular domain. Only users in that domain will be allowed to create guest VMs in this zone.

2. Choose which traffic types will be carried by the physical network.

   The traffic types are management, public, guest, and storage traffic. For more information about the types, roll over the icons to display their tool tips. This screen starts out with one network already configured. If you have multiple physical networks, you need to add more. Drag and drop traffic types onto a greyed-out network and it will become active. You can move the traffic icons from one network to another; for example, if the default traffic types shown for Network 1 do not match your actual setup, you can move them down. You can also change the network names if desired.

3. Assign a network traffic label to each traffic type on each physical network. These labels must match the labels you have already defined on the hypervisor host. To assign each label, click the Edit button under the traffic type icon within each physical network. A popup dialog appears where you can type the label, then click **OK**.

   These traffic labels will be defined only for the hypervisor selected for the first cluster.

   If you have enabled Nexus dvSwitch in the environment, you must specify the corresponding Ethernet port profile names as network traffic label for each traffic type on the physical network. For more information on Nexus dvSwitch, see Configuring a vSphere Cluster with Nexus 1000v Virtual Switch. If you have enabled VMware dvSwitch in the environment, you must specify the corresponding Switch name as network traffic label for each traffic type on the physical network.

   Note

   VMware dvSwitch is supported only for public and guest networks. It's not yet supported for management and storage networks.

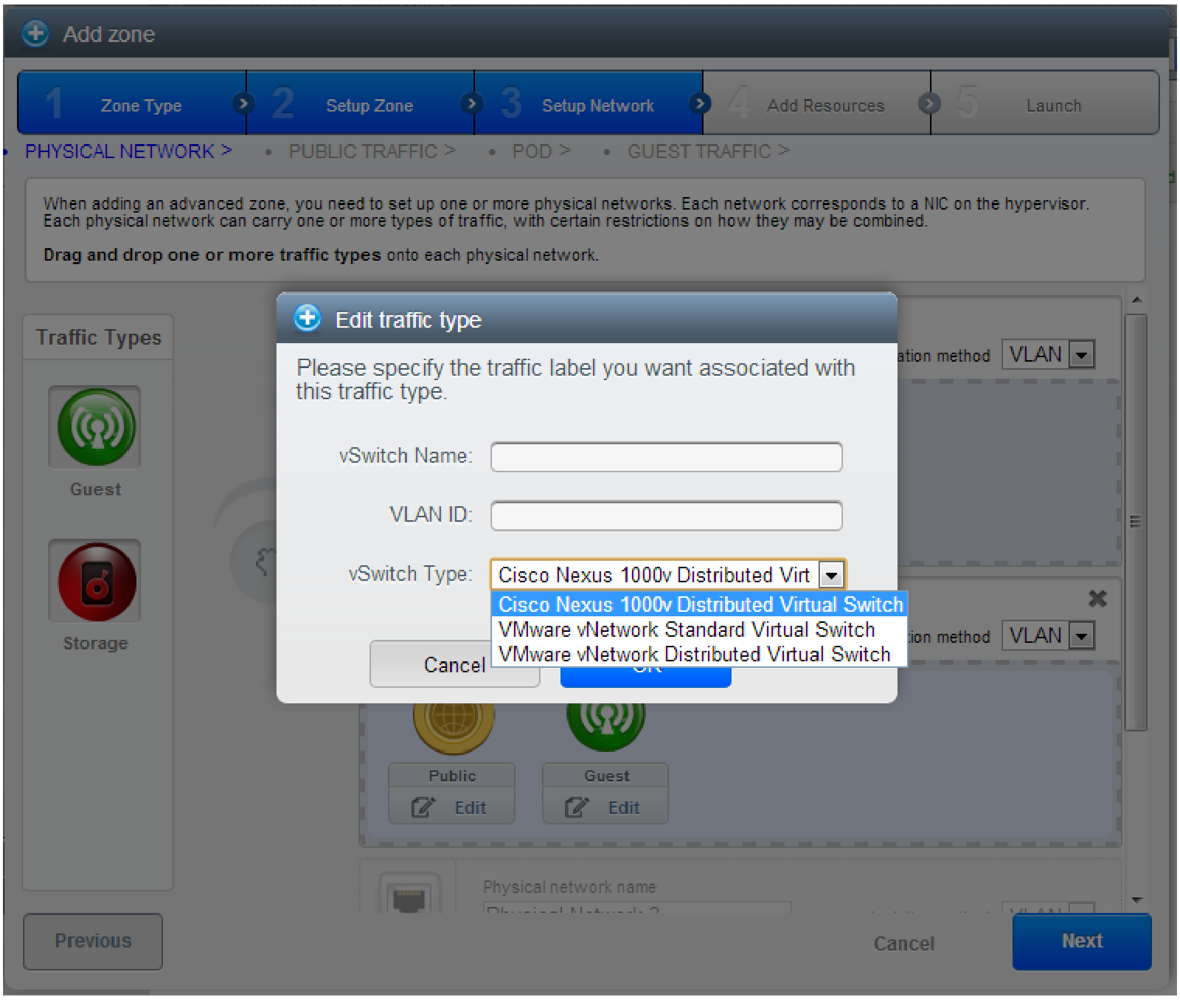

When adding an advanced zone, you need to set up one or more physical networks. Each network corresponds to a NIC on the hypervisor. Each physical network can carry one or more types of traffic, with certain restrictions on how they may be combined.

Drag and drop one or more traffic types onto each physical network.

Edit traffic type

Please specify the traffic label you want associated with this traffic type.

vSwitch Name:

VLAN ID:

vSwitch Type: Cisco Nexus 1000v Distributed Virt

Cisco Nexus 1000v Distributed Virtual Switch
VMware vNetwork Standard Virtual Switch
VMware vNetwork Distributed Virtual Switch
Cancel
OK

4. Click **Next**.

5. Configure the IP range for public Internet traffic. Enter the following details, then click Add. If desired, you can repeat this step to add more public Internet IP ranges. When done, click Next.

   - **Gateway.** The gateway in use for these IP addresses.

   - **Netmask.** The netmask associated with this IP range.

   - **VLAN.** The VLAN that will be used for public traffic.

   - **Start IP/End IP.** A range of IP addresses that are assumed to be accessible from the Internet and will be allocated for access to guest networks.

6. In a new zone, CloudPlatform adds the first pod for you. You can always add more pods later.

   To configure the first pod, enter the following, then click Next:

   - **Pod Name.** A name for the pod.

   - **Reserved system gateway.** The gateway for the hosts in that pod.

   - **Reserved system netmask.** The network prefix that defines the pod's subnet. Use CIDR notation.

- **Start/End Reserved System IP.** The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP.
- Specify a range of VLAN IDs to carry guest traffic for each physical network (for more information, see **section: 6.11.1. VLAN Allocation Example** in the *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5 Administration Guide*), then click Next.
- In a new pod, CloudPlatform adds the first cluster for you. You can always add more clusters later.

  To configure the first cluster, enter the following, then click Next:

  - **Hypervisor.** Select VMware. Then, enter the required information about a vSphere cluster in the additional fields that appear. For vSphere servers, Citrix recommends creating the cluster of hosts in vCenter and then adding the entire cluster to CloudPlatform.
  - **Cluster name.** Enter a name for the cluster. This can be text of your choosing and is not used by CloudPlatform.
- In a new cluster, CloudPlatform adds the first host for you. You can always add more hosts later.

  > **Note**
  >
  > When you deploy CloudPlatform, the hypervisor host must not have any VMs already running.

  Before you can configure the host, you need to install the hypervisor software on the host. You will need to know which version of the hypervisor software version is supported by CloudPlatform and what additional configuration is required to ensure the host will work with CloudPlatform. For more information, refer to **Chapter 5 Installing VMWare for CloudPlatform** in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Hypervisor Configuration Guide*.

  To configure the first host, enter the following, then click **Next**:

  - **Host Name.** The DNS name or IP address of the host.
  - **Username.** Usually root.
  - **Password.** This is the password for the user named above (from your VMWare install).
  - **Host Tags.** (Optional) Any labels that you use to categorize hosts for ease of maintenance. For example, you can set to the cloud's HA tag (set in the ha.tag global configuration parameter) if you want this host to be used only for VMs with the "high availability" feature enabled. For more information, see HA-Enabled Virtual Machines as well as HA for Hosts, both in the Administration Guide.
- In a new cluster, CloudPlatform adds the first primary storage server for you. You can always add more servers later.

  To configure the first primary storage server, enter the following, then click Next:

- **Name.** The name of the storage device.
- **Protocol.**Select VMFS (iSCSI or FiberChannel) or NFS.

CIFS	• **Server.** The IP address or DNS name of the storage device. • **Path.** The exported path from the server. • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
NFS	• **Server.** The IP address or DNS name of the storage device. • **Path.** The exported path from the server. • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
iSCSI	• **Server.** The IP address or DNS name of the storage device. • **Target IQN.** The IQN of the target. For example, iqn.1986-03.com.sun:02:01ec9bb549-1271378984. • **Lun.** The LUN number. For example, 3. • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
preSetup	• **Server.** The IP address or DNS name of the storage device. • **SR Name-Label.** Enter the name-label of the SR that has been set up outside CloudPlatform.

	• **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
SharedMountPoint	• **Path.** The path on each host that is where this primary storage is mounted. For example, "/mnt/primary". • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
VMFS	• **Server.** The IP address or DNS name of the vCenter server. • **Path.** A combination of the datacenter name and the datastore name. The format is "/" datacenter name "/" datastore name. For example, "/cloud.dc.VM/cluster1datastore". • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.

- In a new zone, CloudPlatform adds the first secondary storage server for you.

  Before you can fill out this screen, you need to prepare the secondary storage by setting up NFS shares and installing the latest CloudPlatform System VM template.

  To configure the first secondary storage server on hosts, enter the following, then click **Next**:

  - **NFS Server.** The IP address of the server.

  - **Path.** The exported path from the server.

- Click **Launch**.

5.3. Adding a Pod

After you create a new zone, CloudPlatform adds the first pod to it. You can perform the following procedure to add more pods to the zone at anytime.

1. Log-in to the CloudPlatform UI.
2. In the left navigation bar, select **Infrastructure**.
3. In the right-side panel, under **Zones**, click **View all**.
4. In the page that lists the zones that you configured, click the zone where you want to add a pod.
5. Click the **Compute and Storage** tab. At the **Pods** node in the diagram, click **View all**.
6. In the page that lists the pods configured in the zone that you selected, click **Add Pod**.
7. In the **Add Pod** dialog bozx, enter the following details:
   - **Zone:** Select the name of the zone where you want to add the new pod.
   - **Pod name:** The name that you can use to identify the pod.
   - **Reserved system gateway:** The gateway for the hosts in that pod.
   - **Reserved system netmask:** The network prefix that defines the pod's subnet. Use CIDR notation.
   - **Start/End Reserved System IP:** The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP. For more information, see System Reserved IP Addresses.
   - **Dedicate:** Select if you want to dedicate the pod to a specific domain.
   - **Domain:** Select the domain to which you want to dedicate the pod.
   - **Account:** Enter an account name that belongs to the above selected domain so that you can dedicate the pod to this account.
8. Click **OK**.

5.4. Add Cluster: vSphere

Host management for vSphere is done through a combination of vCenter and the CloudPlatform UI. CloudPlatform requires that all hosts be in a CloudPlatform cluster, but the cluster may consist of a single host. As an administrator you must decide if you would like to use clusters of one host or of multiple hosts. Clusters of multiple hosts allow for features like live migration. Clusters also require shared storage.

> **Note**
>
> Do not use Nexus dvSwitch for management and storage networks. It is supported only for public and guest networks.

For vSphere servers, we recommend creating the cluster of hosts in vCenter and then adding the entire cluster to CloudPlatform.

5.4.1. VMware Cluster Size Limit

The maximum number of hosts in a vSphere cluster is determined by the VMware hypervisor software. For VMware versions 4.2, 4.1, 5.0, and 5.1, the limit is 32 hosts. CloudPlatform adheres to this maximum.

> **Note**
>
> Best Practice: It is advisable for VMware clusters in CloudPlatform to be smaller than the VMware hypervisor's maximum size. A cluster size of up to 8 hosts has been found optimal for most real-world situations.

5.4.2. Adding a vSphere Cluster

To add a vSphere cluster to CloudPlatform:

1. Create the cluster of hosts in vCenter. Follow the vCenter instructions to do this. You will create a cluster that looks something like this in vCenter.

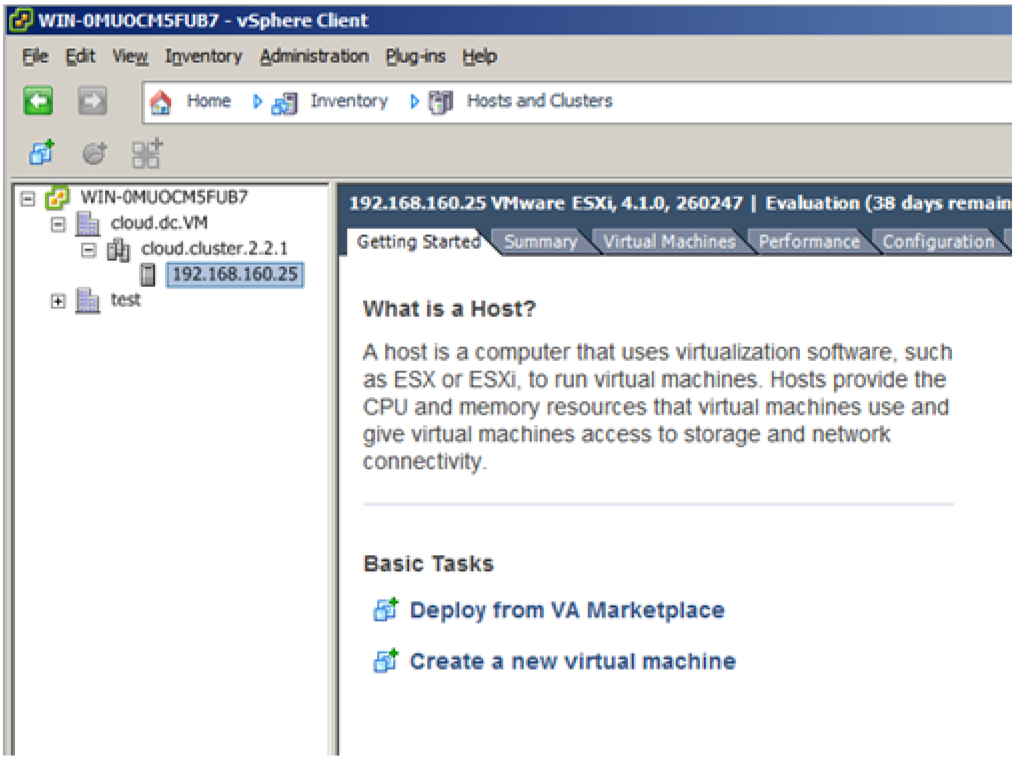
WIN-0MUOCM5FUB7 - vSphere Client
File Edit View Inventory Administration Plug-ins Help
Home
Inventory
Hosts and Clusters
WIN-0MUOCM5FUB7
cloud.dc.VM
cloud.cluster.2.2.1
192.168.160.25
test
192.168.160.25 VMware ESXi, 4.1.0, 260247 | Evaluation (38 days remain
Getting Started
Summary
Virtual Machines
Performance
Configuration
Deploy from VA Marketplace
Create a new virtual machine

What is a Host?

A host is a computer that uses virtualization software, such as ESX or ESXi, to run virtual machines. Hosts provide the CPU and memory resources that virtual machines use and give virtual machines access to storage and network connectivity.

Basic Tasks

2. Log-in to the CloudPlatform UI.

3. In the CloudPlatform UI, in the left navigation bar, click **Infrastructure**.

4. In the right-side panel, under **Zones**, click **View all**.

5. In the page that lists the zones that you have configured, click the zone where you want to add the cluster.

6. In the details page of the zone, click the **Compute and Storage** tab.

7. At the **Clusters** node of the diagram, click **View all**.

8. In the page that list the clusters, click **Add Cluster**.

9. In the **Add Cluster** dialog box, do the following and click **OK**:

10. Provide the following information in the dialog. The fields below make reference to values from vCenter.

    - **Hypervisor**: Select VmWare.

    - **Cluster Name**: Enter the name of the cluster you created in vCenter. For example, "cloud.cluster.2.2.1"

    - **vCenter Host**: Enter the hostname or IP address of the vCenter server.

    - **vCenter Username**: 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".

      If you have enabled Nexus dvSwitch in the environment, the following parameters for dvSwitch configuration are displayed:

    - **Nexus dvSwitch IP Address**: The IP address of the Nexus VSM appliance.

    - **Nexus dvSwitch Username**: The username required to access the Nexus VSM applicance.

    - **Nexus dvSwitch Password**: The password associated with the username specified above.

      There might be a slight delay while the cluster is provisioned. It will automatically display in the UI

5.4.3. Adding a Host (vSphere)

For vSphere servers, we recommend creating the cluster of hosts in vCenter and then adding the entire cluster to CloudPlatform. See *Section 5.4, “Add Cluster: vSphere”*.

5.5. Adding Primary Storage

> **Warning**
>
> When using preallocated storage for primary storage, be sure there is nothing on the storage (ex. you have an empty SAN volume or an empty NFS share). Adding the storage to CloudPlatform will destroy any existing data.

When you create a new zone, the first primary storage is added as part of that procedure. You can add primary storage servers at any time, such as when adding a new cluster or adding more servers to an existing cluster.

1. Log in to the CloudPlatform UI.

2. In the left navigation, choose Infrastructure. In Zones, click View All, then click the zone in which you want to add the primary storage.

3. Click the Compute and Storage tab.

4. In the Primary Storage node of the diagram, click View All.

5. Click Add Primary Storage.

6. Provide the following information in the dialog. The information required varies depending on your choice in Protocol.

   - **Scope**: Indicate whether the storage is available to all hosts in the zone or only to hosts in a single cluster.

   - **Pod**: (Visible only if you choose Cluster in the Scope field.) The pod for the storage device.

   - **Cluster**: (Visible only if you choose Cluster in the Scope field.) The cluster for the storage device.

   - **Name**: The name of the storage device.

   - **Protocol**: For vSphere, choose VMFS (iSCSI or FiberChannel) or NFS.

   - **Server** (for NFS, or iSCSI): The IP address or DNS name of the storage device.

   - **Server** (for VMFS). The IP address or DNS name of the vCenter server.

   - **Path** (for NFS): In NFS this is the exported path from the server.

   - **Path** (for VMFS): In vSphere this is a combination of the datacenter name and the datastore name. The format is "/" datacenter name "/" datastore name. For example, "/cloud.dc.VM/cluster1datastore".

   - **SR Name-Label** (for PreSetup): Enter the name-label of the SR that has been set up outside CloudPlatform.

   - **Target IQN** (for iSCSI): In iSCSI this is the IQN of the target. For example, iqn.1986-03.com.sun:02:01ec9bb549-1271378984

- **Lun #** (for iSCSI): In iSCSI this is the LUN number. For example, 3.
- **Tags** (optional): The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings

The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.

7. Click OK.

5.6. Adding Secondary Storage

Note

Be sure there is nothing stored on the server. Adding the server to CloudPlatform will destroy any existing data.

When you create a new zone, the first secondary storage is added as part of that procedure. You can add secondary storage servers at any time to add more servers to an existing zone.

1. To prepare for the zone-based Secondary Storage, you should have created and mounted an NFS share during Management Server installation.
2. Make sure you prepared the system VM template during Management Server installation.
3. Log in to the CloudPlatform UI as root administrator.
4. In the left navigation bar, click Infrastructure.
5. In Secondary Storage, click View All.
6. Click Add Secondary Storage.
7. Fill in the following fields:
   - **Name**: Give the storage a descriptive name.
   - **Provider**: Choose the type of storage provider (such as S3, or NFS). NFS can be used for zone-based storage, and the others for region-wide object storage. Depending on which provider you choose, additional fields will appear. Fill in all the required fields for your selected provider. For more information, consult the provider's documentation, such as the S3 website.

     Warning

     You can use only a single region-wide object storage account per region. For example, you can not use S3 accounts from different users.

- **Create NFS Secondary Storage**: Be sure this box is checked, unless the zone already contains a secondary staging store. This option is not required if you are upgrading an existing NFS secondary storage into an object storage, as described in *Section 5.6.3, “Upgrading from NFS to Object Storage ”*. Upgrading from NFS to Object Storage in the Installation Guide. In this case, you can skip the rest of the fields described below (Zone, NFS Server, and Path).

  Warning

  If you are setting up a new zone, be sure the box is checked. This checkbox and the three fields below it must be filled in. Even when object storage (such as S3) is used as the secondary storage provider, an NFS staging storage in each zone is still required.

- **Zone**: The zone where the NFS Secondary Storage is to be located.

- **Server.** The IP address or DNS name of the storage device.

- **Path.** The exported path from the server.

- **NFS server**: The name of the zone's Secondary Storage.

- **Path**: The path to the zone's Secondary Storage.

5.6.1. Adding an NFS Secondary Storage for Each Zone

You can skip this section if you are upgrading an existing zone from NFS to object storage. You only need to perform the steps below when setting up a new zone that does not yet have its NFS server.

Every zone must have at least one NFS store provisioned; multiple NFS servers are allowed per zone. To provision an NFS Staging Store for a zone:

1. To prepare for the zone-based Secondary Storage, you should have created and mounted an NFS share during Management Server installation.

2. Make sure you prepared the system VM template during Management Server installation.

3. Log in to the CloudPlatform UI as root administrator.

4. In the left navigation bar, click Infrastructure.

5. In Secondary Storage, click View All.

6. In Select View, choose Secondary Storage.

7. Click the Add NFS Secondary Storage button.

8. Fill out the dialog box fields, then click OK:

   - Zone. The zone where the NFS Secondary Storage is to be located.

   - NFS server. The name of the zone's Secondary Storage.

   - Path. The path to the zone's Secondary Storage.

5.6.2. Configuring S3 Object Store for Secondary Storage

You can configure CloudPlatform to use Amazon S3 Object Store as a secondary storage. S3 Object Store can be used with Amazon Simple Storage Service or any other provider that supports the S3 interface.

1. Make sure you prepared the system VM template during Management Server installation.

2. Log in to the CloudPlatform UI as root administrator.

3. In the left navigation bar, click Infrastructure.

4. In Secondary Storage, click View All.

5. Click Add Secondary Storage.

6. Specify the following:

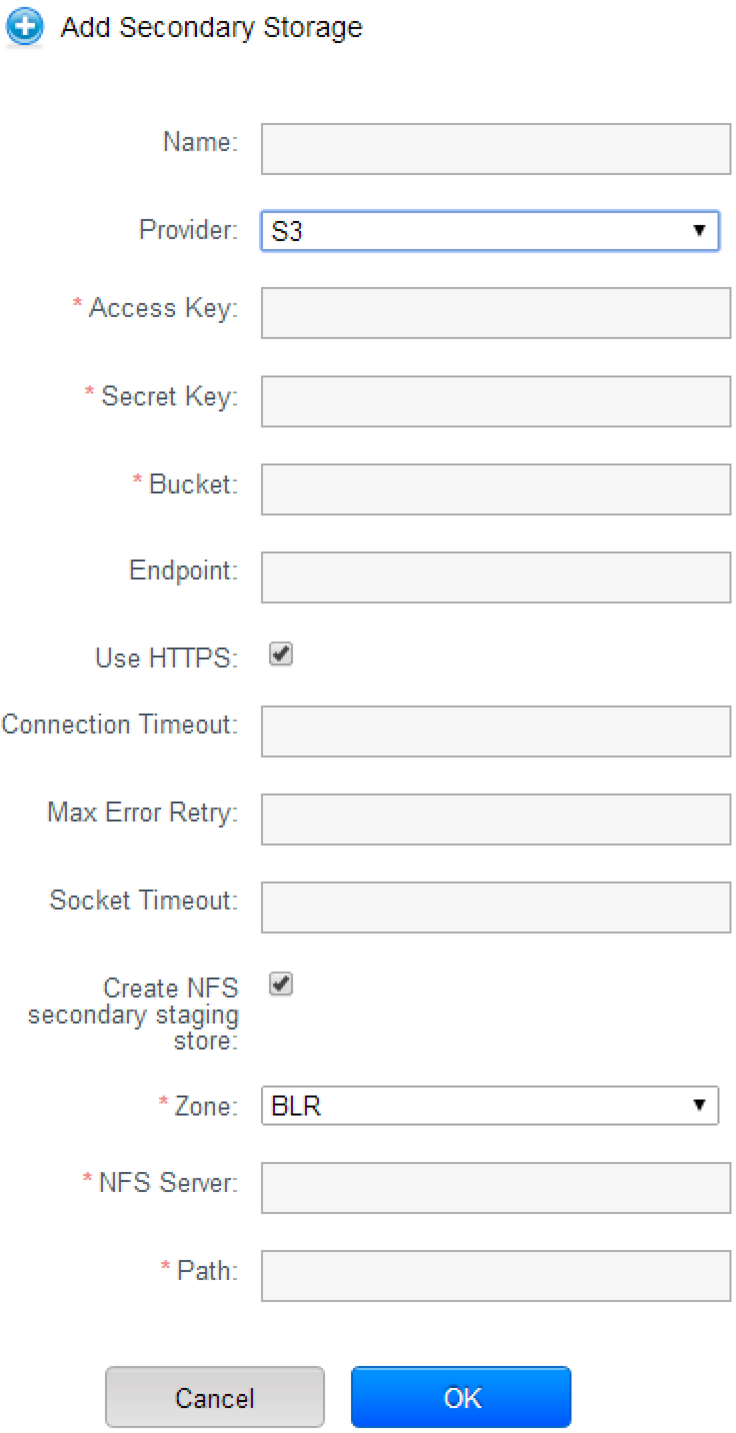

Add Secondary Storage

Name:

Provider: S3

* Access Key:

* Secret Key:

* Bucket:

Endpoint:

Use HTTPS: ☑

Connection Timeout:

Max Error Retry:

Socket Timeout:

Create NFS secondary staging store: ☑

* Zone: BLR

* NFS Server:

* Path:

- **Name**: Give the storage a descriptive name.

- **Provider**: Select S3 for region-wide object storage. S3 can be used with Amazon Simple Storage Service or any other provider that supports the S3 interface.

Warning

You can use only a single region-wide object storage account per region. For example, you can not use S3 accounts from different users.

- **Access Key**: The Access Key ID of the administrator. These credentials are used to securely sign the requests through a REST of Query API to the CloudPlatform services. You can get this from the admin user Details tab in the Accounts page. Because you include it in each request, the ID is a secret. Each Access Key ID has a Secret Access Key associated with it.

- **Secret Key**: The secret key ID of the administrator. You can get this from the admin user Details tab in the Accounts page.

  This key is just a long string of characters (and not a file) that you use to calculate the digital signature that you include in the request. Your Secret Access Key is a secret, and only you and AWS should have it. Don't e-mail it to anyone, include it any AWS requests, or post it on the AWS Discussion Forums. No authorized person from AWS will ever ask for your Secret Access Key.

- **Bucket** : The container of the objects stored in Amazon S3. Enter the name of the bucket where you store your files.

  Your files are stored as objects in a location called a bucket. When you configure your Amazon S3 bucket as a website, the service delivers the files in your bucket to web browsers as if they were hosted on a web server.

- **End Point**: The IP address or DNS name of the S3 storage server.

  For example: 10.10.29.1:8080, where 8080 is the listening port of the S3 storage server.

- **Use HTTPS**: Specify if you want a secure connection with the S3 storage.

- **Connection Timeout**: The default timeout for creating new connections.

- **Max Error Retry**: The number of retry after service exceptions due to internal errors.

- **Socket Timeout**: The default timeout for reading from a connected socket.

- **Create NFS Secondary Staging Store**: If the zone already contains a secondary staging store, do not select this option. Select if you are upgrading an existing NFS secondary storage into an object storage, as described in *Section 5.6.3, “Upgrading from NFS to Object Storage ”*. Upgrading from NFS to Object Storage in the Installation Guide. In this case, you can skip the rest of the fields described below (Zone, NFS Server, and Path).

- **Zone**: The zone where S3 the Object Store is to be located.

- **Path**: The path to the zone's Secondary Staging Store.

5.6.3. Upgrading from NFS to Object Storage

In an existing zone that is using NFS for secondary storage, you can upgrade the zone to use a region-wide object storage without causing downtime. The existing NFS storage in the zone will be converted to an NFS Staging Store.

After upgrade, all newly created templates, ISOs, volumes, snapshots are moved to the object store. All previously created templates, ISOs, volumes, snapshots are migrated on an on-demand basis based on when they are accessed, rather than as a batch job. Unused objects in the NFS staging store are garbage collected over time.

1. Log in as admin to the CloudPlatform UI.

2. Fire an admin API to update CloudPlatform to use object storage:

   ```
   http://<MGMTIP>:8096/client/api?command=updateCloudToUseObjectStore&name=<S3
    storage name>&provider=S3&details[0].key=accesskey&details[0].value=<access
    key from .s3cfg file>&details[1].key=secretkey&details[1].value=<secretKey
    from .s3cfg file>&details[2].key=bucket&details[2].value=<bucketname>&details[3]
    key=usehttps&details[3].value=<trueorfalse>&details[4].key=endpoint&details[4]. value=
    <S3 server IP:8080>
   ```

   All existing NFS secondary storages has been converted to NFS staging stores for each zone, and your S3 object store specified in the command has been added as a new region-wide secondary storage.

3. Locate the secondary storage that you want to upgrade to object storage.

   Perform either of the following in the Infrastructure page:

   - In Zones, click View All, then locate the desired zone, and select Secondary Storage in the Compute and Storage tab.

   - In Secondary Storage, click View All, then select the desired secondary storage.

Post migration, consider the following:

- For each new snapshot taken after migration, ensure that you take a full snapshot to newly added S3.

  This would help coalesce delta snapshots across NFS and S3 stores. The snapshots taken before migration are pushed to S3 store when you try to use that snapshot by executing createVolumeCmd by passing snapshot id.

- You can deploy VM from templates because they are already in the NFS staging store. A copy of the template is generated in the new S3 object store when ExtractTemplate or CopyTemplate is executed.

- For volume, a copy of the volume is generated in the new S3 object store when ExtractVolume command is executed.

- All the items in the NFS storage is not migrated to S3. Therefore, if you want to completely shut down the NFS storage you have previously used, write your own script to migrate those remaining items to S3.

5.7. Initialize and Test

After everything is configured, CloudPlatform will perform its initialization. This can take 30 minutes or more, depending on the speed of your network. When the initialization has completed successfully, the administrator's Dashboard should be displayed in the CloudPlatform UI.

1. Verify that the system is ready. In the left navigation bar, select Templates. Click the template. Check to be sure that the status is "Download Complete." Do not proceed to the next step until this status is displayed.

2. Go to the Instances tab, and filter by My Instances.

3. Click Add Instance and follow the steps in the wizard.

   a. Choose the zone you just added.

   b. In the template selection, choose the template to use in the VM. If this is a fresh installation, likely only the provided CentOS template is available.

   c. Select a service offering. Be sure that the hardware you have allows starting the selected service offering.

   d. In data disk offering, if desired, add another data disk. This is a second volume that will be available to but not mounted in the guest. A reboot is not required if you have a PV-enabled OS kernel in use.

   e. In default network, choose the primary network for the guest. In a trial installation, you would have only one option here.

   f. Optionally give your VM a name and a group. Use any descriptive text you would like.

   g. Click Launch VM. Your VM will be created and started. It might take some time to download the template and complete the VM startup. You can watch the VMâ€™s progress in the Instances screen.

4. To use the VM, click the View Console button.

   For more information about using VMs, including instructions for how to allow incoming network traffic to the VM, start, stop, and delete VMs, and move a VM from one host to another, see Working With Virtual Machines in the Administratorâ€™s Guide.

Congratulations! You have successfully completed a CloudPlatform Installation.

If you decide to grow your deployment, you can add more hosts, primary storage, zones, pods, and clusters.

Provisioning Your Cloud Infrastructure on Hyper-V

This section describes how to add zones, pods, clusters, hosts, storage, and networks to your cloud using Hyper-V hypervisor.

For conceptual information about zones, pods, clusters, hosts, storage, and networks in CloudPlatform, refer to *CloudPlatform (powered by CloudStack) Version 4.5 Concepts Guide*.

6.1. Overview of Provisioning Steps

After installing Management Server, you can access the CloudPlatform UI and add the compute resources for CloudPlatform to manage.

Then, you can provision the cloud infrastructure, or scale the cloud infrastructure up at any time.

After you complete provisioning the cloud infrastructure, you will have a deployment with the following basic structure:

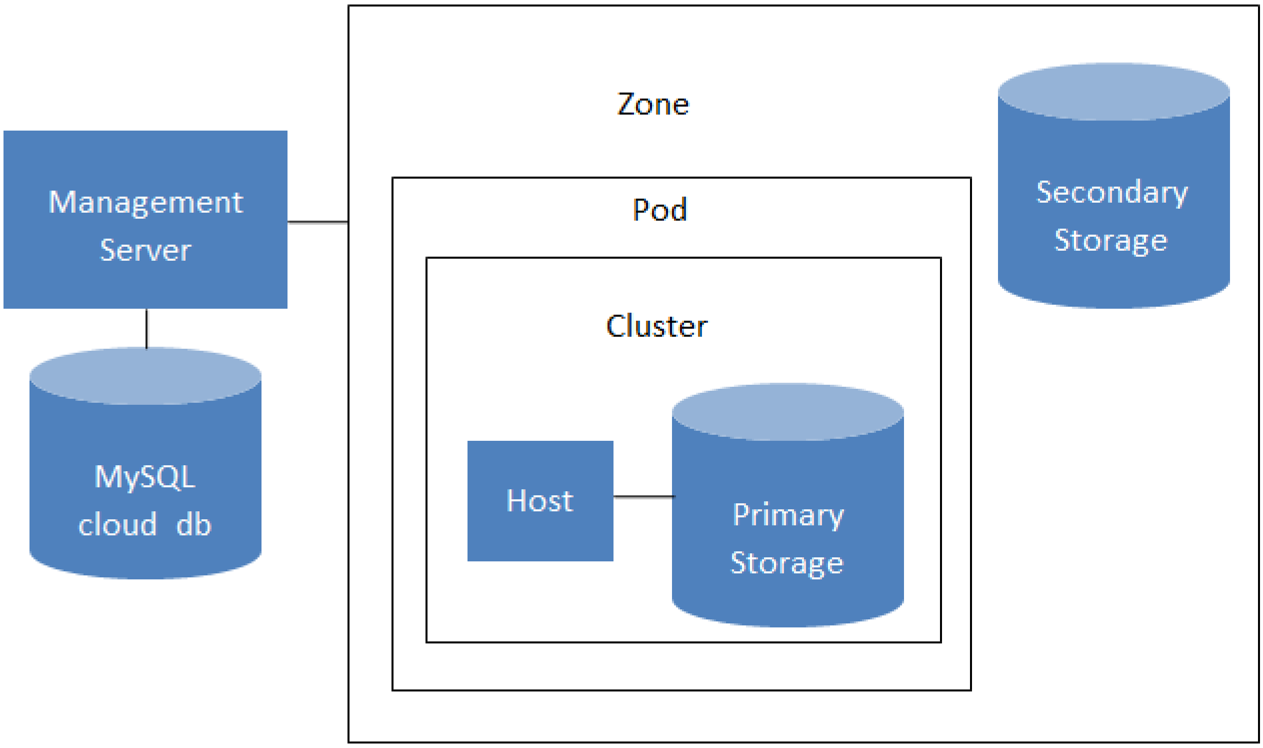

Conceptual view of a basic deployment

For information on adding a region to your cloud infrastructure, refer to **Chapter 3 Adding Regions to Your Cloud Infrastructure (optional)** of *CloudPlatform (powered by CloudStack) Version 4.5 Administration Guide*.

6.2. Adding a Zone

Adding a zone consists of three phases:

- Create a mount point for secondary storage on the Management Server.
- Seed the system VM template on the secondary storage.
- Add the zone.

6.2.1. Create a Secondary Storage Mount Point for the New Zone

To ensure that you have deployed the latest system VMs in new zones, you need to seed the latest system VM template to the zone's secondary storage. The first step is to create a mount point for the secondary storage. Then seed the system VM template. As a prerequisite, you must install CIFS packege.

1. On the management server, create a mount point for secondary storage. For example:

   ```
   # mkdir -p /mnt/secondary
   ```

2. Mount the secondary storage on your Management Server. Use the following command and replace the example CIFS server name and CIFS share paths below with your own:

   You must have Samba client to mount CIFS on Linux. You can run the command **`yum install samba-client samba-common cifs-utils`**.

   ```
   mount -t cifs <cifsServer://cifsserver/cifsShare/Secondary_storage -o
    username=<domainuser>,password=<password>,domain=<domain name> /mnt/secondary
   ```

3. Secondary storage must be seeded with a template that is used for CloudPlatform system VMs. For more information, refer to the CloudPlatform Installation Guide. After you seed the secondary storage with a system VM template, continue with adding the zone.

6.2.2. Adding a New Zone

To add a zone, you must first configure the zone’s physical network. Then, you need to add the first pod, cluster, host, primary storage, and secondary storage.

Note

Before you proceed with adding a new zone, you must ensure that you have performed the steps to seed the system VM template.

Note

Citrix strongly recommends using the same type of hypervisors in a zone to avoid operational issues.

Note

Hyper-V clusters are supported in dedicated zones.

1. Log-in to the CloudPlatform UI using the root administrator account.

2. In the left navigation bar, click **Infrastructure**.

3. In the right side panel, under **Zones**, click **View all**.

4. In the next page, click **Add Zone**. The **Add zone** wizard panel appears.

5. Select **Advanced** as the network type.

   The Advanced network type is used for more sophisticated network topologies. This network model provides the most flexibility in defining guest networks and providing custom network offerings such as firewall, VPN, or load balancer support.

   Note

   The **Basic** zone type is not supported on Hyper-V hosts.

   For more information, refer to **Chapter 4. Cloud Infrastructure Concepts** of the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Concepts Guide*.

6. After you select the Advanced network type, proceed as described in *Section 6.2.2.1, "Advanced Zone Configuration "*

6.2.2.1. Advanced Zone Configuration

1. After you select Advanced in the Add Zone wizard and click Next, you will be asked to enter the following details. Then click Next.

   - **Name.** A name for the zone.

   - **DNS 1 and 2.** These are DNS servers for use by guest VMs in the zone. These DNS servers will be accessed via the public network you will add later. The public IP addresses for the zone must have a route to the DNS server named here.

   - **Internal DNS 1 and Internal DNS 2.** These are DNS servers for use by system VMs in the zone(these are VMs used by CloudPlatform itself, such as virtual routers, console proxies,and Secondary Storage VMs.) These DNS servers will be accessed via the management traffic network interface of the System VMs. The private IP address you provide for the pods must have a route to the internal DNS server named here.

   - **Network Domain.** (Optional) If you want to assign a special domain name to the guest VM network, specify the DNS suffix.

- **Guest CIDR.** This is the CIDR that describes the IP addresses in use in the guest virtual networks in this zone. For example, 10.1.1.0/24. As a matter of good practice you should set different CIDRs for different zones. This will make it easier to set up VPNs between networks in different zones.

- **Hypervisor.** Choose **Hyper-V** as the hypervisor for the first cluster in the zone.

- **Public.** A public zone is available to all users. A zone that is not public will be assigned to a particular domain. Only users in that domain will be allowed to create guest VMs in this zone.

2. Select the traffic types that the physical network will carry.

   The traffic types are management, public, guest, and storage traffic. For more information about the types, roll over the icons to display their tool tips. This screen starts out with one network already configured. If you have multiple physical networks, you need to add more. Drag and drop traffic types onto a greyed-out network and it will become active. You can move the traffic icons from one network to another; for example, if the default traffic types shown for Network 1 do not match your actual setup, you can move them down. You can also change the network names if desired.

3. Assign a network traffic label to each traffic type on each physical network. The traffic label must match the Hyper-V vswitch name that is configured on the Hyper-V host. These labels must match the labels you have already defined on the hypervisor host. To assign each label, click the Edit button under the traffic type icon within each physical network. A popup dialog appears where you can type the label, then click OK.

   These traffic labels will be defined only for the hypervisor selected for the first cluster.

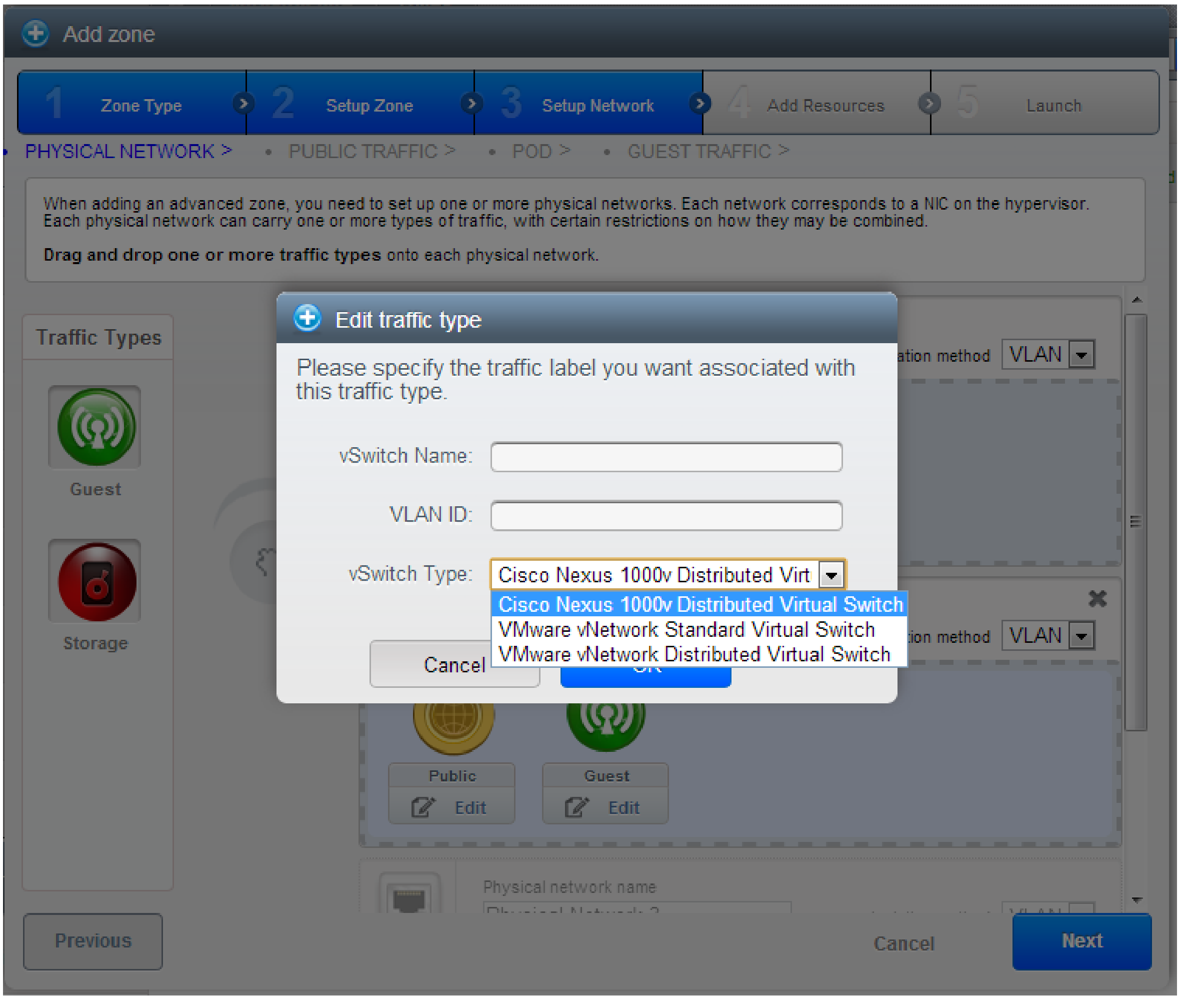

When adding an advanced zone, you need to set up one or more physical networks. Each network corresponds to a NIC on the hypervisor. Each physical network can carry one or more types of traffic, with certain restrictions on how they may be combined.

Drag and drop one or more traffic types onto each physical network.

Please specify the traffic label you want associated with this traffic type.

vSwitch Name:

VLAN ID:

vSwitch Type: Cisco Nexus 1000v Distributed Virt

4. Click Next.

5. Configure the IP range for public Internet traffic. Enter the following details, then click Add. If desired, you can repeat this step to add more public Internet IP ranges. When done, click Next.

   - **Gateway.** The gateway in use for these IP addresses.
   - **Netmask.** The netmask associated with this IP range.
   - **VLAN.** The VLAN that will be used for public traffic.
   - **Start IP/End IP.** A range of IP addresses that are assumed to be accessible from the Internet and will be allocated for access to guest networks.

6. In a new zone, CloudPlatform adds the first pod for you. You can always add more pods later.

   To configure the first pod, enter the following, then click Next:

   - **Pod Name.** A name for the pod.
   - **Reserved system gateway.** The gateway for the hosts in that pod.
   - **Reserved system netmask.** The network prefix that defines the pod's subnet. Use CIDR notation.

- **Start/End Reserved System IP.** The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP.

- Specify a range of VLAN IDs to carry guest traffic for each physical network (for more information, see **section: 6.11.1. VLAN Allocation Example** in the *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5 Administration Guide*), then click Next.

- In a new pod, CloudPlatform adds the first cluster for you. You can always add more clusters later.

  To configure the first cluster, enter the following, then click Next:

  - **Hypervisor.** Select Hyper-V.

  - **Cluster name.** Enter a name for the cluster. This can be text of your choosing and is not used by CloudPlatform.

- In a new cluster, CloudPlatform adds the first host for you. You can always add more hosts later.

  > **Note**
  >
  > When you deploy CloudPlatform, the hypervisor host must not have any VMs already running.

  Before you can configure the host, you need to install the hypervisor software on the host. You will need to know which version of the hypervisor software version is supported by CloudPlatform and what additional configuration is required to ensure the host will work with CloudPlatform. For more information, refer to **Chapter 3 Installing Hyper-V for CloudPlatform** in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Hypervisor Configuration Guide*.

  To configure the first host, enter the following, then click Next:

  - **Host Name.** The DNS name or IP address of the host.

  - **Username.** Usually root.

  - **Password.** This is the password for the user named above (from your Hyper-V install).

  - **Host Tags.** (Optional) Any labels that you use to categorize hosts for ease of maintenance. For example, you can set to the cloud's HA tag (set in the ha.tag global configuration parameter) if you want this host to be used only for VMs with the "high availability" feature enabled. For more information, see HA-Enabled Virtual Machines as well as HA for Hosts, both in the Administration Guide.

- In a new cluster, CloudPlatform adds the first primary storage server for you. You can always add more servers later.

  To configure the first primary storage server, enter the following, then click Next:

  - **Name.** The name of the storage device.

- **Protocol.**Select CIFS. The remaining fields in the screen vary depending on the protocol that you selected.

CIFS	• **Server.** The IP address or DNS name of the storage device. • **Path.** The exported path from the server. • **SMB Username**: The username of the account which has the necessary permissions to the SMB shares. The user must be part of the Hyper-V administrator group. • **SMB Password**: The password associated with the account. • **SMB Domain**: The Active Directory domain that the SMB share is a part of. • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
NFS	• **Server.** The IP address or DNS name of the storage device. • **Path.** The exported path from the server. • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
iSCSI	• **Server.** The IP address or DNS name of the storage device. • **Target IQN.** The IQN of the target. For example, iqn.1986-03.com.sun:02:01ec9bb549-1271378984. • **Lun.** The LUN number. For example, 3. • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary

	storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
preSetup	• **Server.** The IP address or DNS name of the storage device. • **SR Name-Label.** Enter the name-label of the SR that has been set up outside CloudPlatform. • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
SharedMountPoint	• **Path.** The path on each host that is where this primary storage is mounted. For example, "/mnt/primary". • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
VMFS	• **Server.** The IP address or DNS name of the vCenter server. • **Path.** A combination of the datacenter name and the datastore name. The format is "/" datacenter name "/" datastore name. For example, "/cloud.dc.VM/cluster1datastore". • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.

- In a new zone, CloudPlatform adds the first secondary storage server for you.

  Before you can fill out this screen, you need to prepare the secondary storage by setting up NFS shares and installing the latest CloudPlatform System VM template.

  (Hyper-V) To configure the first secondary storage server, enter the following, then click Next:

  - **Server.** The IP address or DNS name of the storage device.
  - **Path.** The exported path from the server.
  - **SMB Username**: The username of the account which has the necessary permissions to the SMB shares. The user must be part of the Hyper-V administrator group.
  - **SMB Password**: The password associated with the account.
  - **SMB Domain**: The Active Directory domain that the SMB share is a part of.

  To configure the first secondary storage server on hosts other than Hyper-V, enter the following, then click Next:

  - **NFS Server.** The IP address of the server.
  - **Path.** The exported path from the server.

- Click Launch.

6.3. Adding a Pod

After you create a new zone, CloudPlatform adds the first pod to it. You can perform the following procedure to add more pods to the zone at anytime.

1. Log-in to the CloudPlatform UI.
2. In the left navigation bar, select **Infrastructure**.
3. In the right-side panel, under **Zones**, click **View all**.
4. In the page that lists the zones that you configured, click the zone where you want to add a pod.
5. Click the **Compute and Storage** tab. At the **Pods** node in the diagram, click **View all**.
6. In the page that lists the pods configured in the zone that you selected, click **Add Pod**.
7. In the **Add Pod** dialog bozx, enter the following details:
   - **Zone:** Select the name of the zone where you want to add the new pod.
   - **Pod name:** The name that you can use to identify the pod.
   - **Reserved system gateway:** The gateway for the hosts in that pod.
   - **Reserved system netmask:** The network prefix that defines the pod's subnet. Use CIDR notation.
   - **Start/End Reserved System IP:** The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP. For more information, see System Reserved IP Addresses.
   - **Dedicate:** Select if you want to dedicate the pod to a specific domain.
   - **Domain:** Select the domain to which you want to dedicate the pod.
   - **Account:** Enter an account name that belongs to the above selected domain so that you can dedicate the pod to this account.

8. Click **OK**.

6.4. Adding a Hyper-V Cluster

You need to tell CloudPlatform about the hosts that it will manage. Hosts exist inside clusters, so before you begin adding hosts to the cloud, you must add at least one cluster.

Before you perform the following steps, you must have installed the Hyper-V on the hosts and logged-in to the CloudPlatform UI.

1. In the CloudPlatform UI, in the left navigation bar, click **Infrastructure**.

2. In the right-side panel, under **Zones**, click **View all**.

3. In the page that lists the zones that you have configured, click the zone where you want to add the cluster.

4. In the details page of the zone, click the **Compute and Storage** tab.

5. At the **Clusters** node of the diagram, click **View all**.

6. In the page that list the clusters, click **Add Cluster**.

7. In the **Add Cluster** dialog box, do the following and click **OK**:

   - **Zone Name:** Select the zone where you want to create the cluster.

   - **Hypervisor:** Select Hyper-V as the hypervisor for this cluster.

   - **Pod Name:** Select the pod where you want to create the cluster.

   - **Cluster Name:** Enter a name that you can use to identify the cluster.

   - **Dedicate**: Select if you want to dedicate the cluster to a specific domain.

   - **Domain**: Select the domain to which you want to dedicate the cluster.

   - **Account**: Enter an account name that belongs to the domain so that you can dedicate the cluster to this account.

6.4.1. Adding a Hyper-V Host

1. You must install Hyper-V on the host. You must know the version of the Hyper-V software that CloudPlatform supports and the additional configurations that are required to ensure that the host will work with CloudPlatform.

   > **Note**
   >
   > Ensure that you have performed the additional CloudPlatform-specific configuration steps described in **Chapter 3. Installing Hyper-V for CloudPlatform** of the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Hypervisor Configuration Guide*.

2. Log-in to the CloudPlatform UI using an administrator account.

3. In the left navigation bar, click **Infrastructure**.
4. In the right-side panel, under **Zones**, click **View all**.
5. In the page that lists the zones that are configured with CloudPlatform, click the zone where you want to add the hosts.
6. In the details page of the zone, click the **Compute and Storage** tab.
7. An the **Clusters** node, click **View all**.
8. In the page that lists the clusters available with the zone, click the cluster where you want to add the host.
9. Under the **Details** tab, click the **View Hosts** link.
10. In the page that lists the hosts available with the cluster, click **Add Host**.
11. In the **Add Host** panel, provide the following information:

    - **Zone**: Select the zone where you want to add the host.
    - **Pod**: Select the pod in the zone where you want to add the host.
    - **Cluster**: Select the cluster in the pod where you want to add the host.
    - **Host Name**: The DNS name or IP address of the host.
    - **Username**: Usually root.
    - **Password**: This is the password associated with the user name from your Hyper-V install.
    - **Dedicate**: Select to indicate that this hiost is to be dedicated to a specific domain and account
        - **Domain**: Select the domain to which you want to dedicate the host.
        - **Account**: Select the account that is associated with the domain so that you can dedicate the host to this account.
    - **Host Tags** (Optional). Any labels that you use to categorize hosts for ease of maintenance. For example, you can set to the cloud's HA tag (set in the **`ha.tag`** global configuration parameter) if you want this host to be used only for VMs with the "high availability" feature enabled. For more information, refer to the **HA-Enabled Virtual Machines** and the **HA for Hosts** sections in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Administration Guide*.

    There may be a slight delay while the host is provisioned. It will display automatically in the UI.

12. Repeat the procedure for adding additional hosts.

6.5. Adding Primary Storage

> **Warning**
>
> When using preallocated storage for primary storage, be sure there is nothing on the storage (ex. you have an empty SAN volume or an empty NFS share). Adding the storage to CloudPlatform will destroy any existing data.

When you create a new zone, the first primary storage is added as part of that procedure. You can add primary storage servers at any time, such as when adding a new cluster or adding more servers to an existing cluster.

1. Log in to the CloudPlatform UI.
2. In the left navigation, choose Infrastructure. In Zones, click View All, then click the zone in which you want to add the primary storage.
3. Click the Compute and Storage tab.
4. In the Primary Storage node of the diagram, click View All.
5. Click Add Primary Storage.
6. Provide the following information in the dialog. The information required varies depending on your choice in Protocol.
   - **Scope**: Indicate whether the storage is available to all hosts in the zone or only to hosts in a single cluster.
   - **Pod**: (Visible only if you choose Cluster in the Scope field.) The pod for the storage device.
   - **Cluster**: (Visible only if you choose Cluster in the Scope field.) The cluster for the storage device.
   - **Name**: The name of the storage device.
   - **Protocol**:Select SMB/CIFS.
   - **Path** (for SMB/CIFS): The exported path from the server.
   - **SMB Username**: Applicable only if you select SMB/CIFS provider. The username of the account which has the necessary permissions to the SMB shares. The user must be part of the Hyper-V administrator group.
   - **SMB Password**: Applicable only if you select SMB/CIFS provider. The password associated with the account.
   - **SMB Domain**: Applicable only if you select SMB/CIFS provider. The Active Directory domain that the SMB share is a part of.
   - **SR Name-Label** (for PreSetup): Enter the name-label of the SR that has been set up outside CloudPlatform.

- **Target IQN** (for iSCSI): In iSCSI this is the IQN of the target. For example, iqn.1986-03.com.sun:02:01ec9bb549-1271378984
- **Lun #** (for iSCSI): In iSCSI this is the LUN number. For example, 3.
- **Tags** (optional): The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings

   The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.

7. Click **OK**.

6.6. Adding Secondary Storage

> **Note**
>
> Be sure there is nothing stored on the server. Adding the server to CloudPlatform will destroy any existing data.

When you create a new zone, the first secondary storage is added as part of that procedure. You can add secondary storage servers at any time to add more servers to an existing zone.

1. To prepare for the zone-based Secondary Storage, you should have created and mounted an NFS share during Management Server installation.
2. Make sure you prepared the system VM template during Management Server installation.
3. Log in to the CloudPlatform UI as root administrator.
4. In the left navigation bar, click Infrastructure.
5. In Secondary Storage, click View All.
6. Click Add Secondary Storage.
7. Fill in the following fields:
   - **Name**: Give the storage a descriptive name.
   - **Provider**: Select SMB as the storage provider. Depending on the provider that you selected, additional fields will appear.

     > **Warning**
     >
     > You can use only a single region-wide object storage account per region.

- **Zone**: Select the zone where you want to create secondary storage
- **Server.** The IP address or DNS name of the storage server.
- **Path.** The exported path from the server.
- **SMB Username**: Applicable only if you select SMB/CIFS provider. The username of the account which has the necessary permissions to the SMB shares. The user must be part of the Hyper-V administrator group.
- **SMB Password**: Applicable only if you select SMB/CIFS provider. The password associated with the account.
- **SMB Domain**: Applicable only if you select SMB/CIFS provider. The Active Directory domain that the SMB share is a part of.

6.7. Initialize and Test

After everything is configured, CloudPlatform will perform its initialization. This can take 30 minutes or more, depending on the speed of your network. When the initialization has completed successfully, the administrator's Dashboard should be displayed in the CloudPlatform UI.

1. Verify that the system is ready. In the left navigation bar, select Templates. Click the template. Check to be sure that the status is "Download Complete." Do not proceed to the next step until this status is displayed.

2. Go to the Instances tab, and filter by My Instances.

3. Click Add Instance and follow the steps in the wizard.

   a. Choose the zone you just added.

   b. In the template selection, choose the template to use in the VM. If this is a fresh installation, likely only the provided CentOS template is available.

   c. Select a service offering. Be sure that the hardware you have allows starting the selected service offering.

   d. In data disk offering, if desired, add another data disk. This is a second volume that will be available to but not mounted in the guest. A reboot is not required if you have a PV-enabled OS kernel in use.

   e. In default network, choose the primary network for the guest. In a trial installation, you would have only one option here.

   f. Optionally give your VM a name and a group. Use any descriptive text you would like.

   g. Click Launch VM. Your VM will be created and started. It might take some time to download the template and complete the VM startup. You can watch the VMâ€™s progress in the Instances screen.

4. To use the VM, click the View Console button. >_

   For more information about using VMs, including instructions for how to allow incoming network traffic to the VM, start, stop, and delete VMs, and move a VM from one host to another, see Working With Virtual Machines in the Administratorâ€™s Guide.

Congratulations! You have successfully completed a CloudPlatform Installation.

If you decide to grow your deployment, you can add more hosts, primary storage, zones, pods, and clusters.

(Experimental Feature) Provisioning Your Cloud Infrastructure on LXC

This section describes how to add zones, pods, clusters, hosts, storage, and networks to your cloud using an LXC hypervisor.

For conceptual information about zones, pods, clusters, hosts, storage, and networks in CloudPlatform, refer to *CloudPlatform (powered by CloudStack) Version 4.5 Concepts Guide*.

7.1. Overview of Provisioning Steps

After installing Management Server, you can access the CloudPlatform UI and add the compute resources for CloudPlatform to manage.

Then, you can provision the cloud infrastructure, or scale the cloud infrastructure up at any time.

After you complete provisioning the cloud infrastructure, you will have a deployment with the following basic structure:

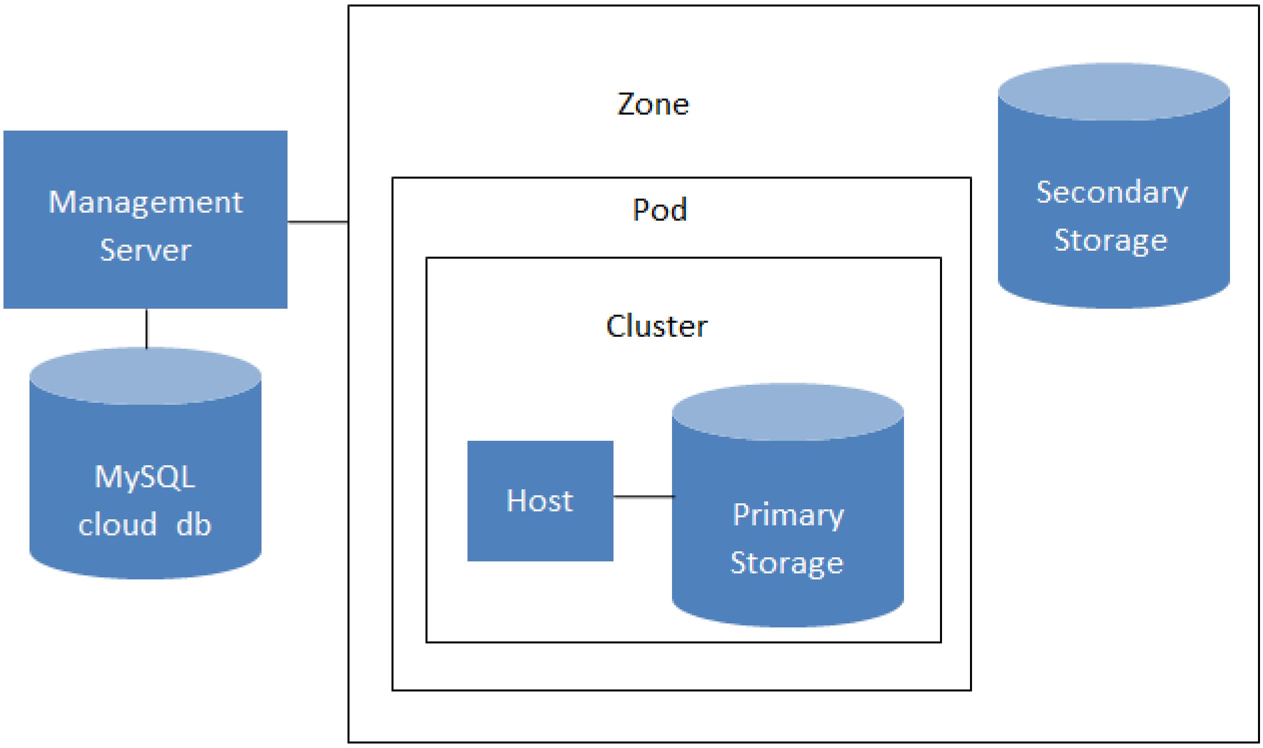

Conceptual view of a basic deployment

For information on adding a region to your cloud infrastructure, refer to **Chapter 3 Adding Regions to Your Cloud Infrastructure (optional)** of *CloudPlatform (powered by CloudStack) Version 4.5 Administration Guide*.

7.2. Adding a Zone

Adding a zone consists of three phases:

- Create a secondary storage mount point for the zone
- Seed the system VM template on the secondary storage.
- Add the zone.

7.2.1. Creating a Secondary Storage Mount Point for the Zone

To ensure that you deploy the latest system VMs in a new zone, you must seed the latest system VM template to the secondary storage of the zone. For this, you must first create a mount point for the secondary storage. Then, you can seed the latest system VM template to the secondary storage.

1. On Management Server, create a mount point for secondary storage. For example:

   ```
   # mkdir -p /mnt/secondary
   ```

2. Mount the secondary storage on your Management Server. Replace NFS server name and NFS share paths in the following example with the server name and path that you use.

   ```
   # mount -t nfs nfsservername:/nfs/share/secondary /mnt/secondary
   ```

3. Seed the secondary storage with the latest template that is used for CloudPlatform system VMs. For more information about seeding the secondary storage, refer to the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Installation Guide*.

   After you seed the secondary storage with the latest system VM template, continue with adding the new zone.

7.2.2. Adding a New Zone

To add a zone, you must first configure the zone's physical network. Then, you need to add the first pod, cluster, host, primary storage, and secondary storage.

> **Note**
>
> Before you proceed with adding a new zone, you must ensure that you have performed the steps to seed the system VM template.

> **Note**
>
> Citrix strongly recommends using the same type of hypervisors in a zone to avoid operational issues.

1. Log in to the CloudPlatform UI using the root administrator account.
2. In the left navigation bar, click **Infrastructure**.
3. In the right side panel, under **Zones**, click **View all**.

4. In the next page, click **Add Zone**.

   The **Add zone** wizard panel appears.

5. Select one of the following network types:

   - **Basic**: Provides a single network where each VM instance is assigned with an IP directly from the network. You can provide guest isolation through layer-3 means, such as security groups (IP address source filtering).
   - **Advanced**: Used for more sophisticated network topologies. This network model provides the most flexibility in defining guest networks and providing custom network offerings such as firewall, VPN, or load balancer support.

   For more information, refer to **Chapter 4. Cloud Infrastructure Concepts** of the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Concepts Guide*.

6. Based on the option (Basic or Advanced) that you selected, do one of the following:

   - For **Basic**: *Section 7.2.2.1, "Basic Zone Configuration "*
   - For **Advanced**: *Section 7.2.2.2, "Advanced Zone Configuration "*

7.2.2.1. Basic Zone Configuration

1. After you select Basic in the Add Zone wizard and click Next, you will be asked to enter the following details. Specify the details, then click Next.

   - **Name**: A name for the zone.
   - **DNS 1 and 2**: These are DNS servers for use by guest VMs in the zone. These DNS servers will be accessed via the public network you will add later. The public IP addresses for the zone must have a route to the DNS server named here.
   - **Internal DNS 1 and Internal DNS 2**: These are DNS servers for use by system VMs in the zone (these are VMs used by CloudPlatform itself, such as virtual routers, console proxies, and Secondary Storage VMs.) These DNS servers will be accessed via the management traffic network interface of the System VMs. The private IP address you provide for the pods must have a route to the internal DNS server named here.
   - **Hypervisor**: Choose LXC for the first cluster in the zone.
   - **Network Offering**: Your choice here determines what network services will be available on the network for guest VMs.

Network Offering	Description
DefaultSharedNetworkOfferingWithSGService	If you want to enable security groups for guest traffic isolation, choose this. See Using Security Groups to Control Traffic to VMs.
DefaultSharedNetworkOffering	If you do not need security groups, choose this.
DefaultSharedNetscalerEIPandELBNetworkOffering	If you have installed a Citrix NetScaler appliance as part of your zone network, and you will be using its Elastic IP and Elastic Load Balancing features, choose this. With the EIP and ELB features, a basic zone with

Network Offering	Description
	security groups enabled can offer 1:1 static NAT and load balancing.

- **Network Domain.** (Optional) If you want to assign a special domain name to the guest VM network, specify the DNS suffix.

- **Public.** A public zone is available to all users. A zone that is not public will be assigned to a particular domain. Only users in that domain will be allowed to create guest VMs in this zone.

2. Choose which traffic types will be carried by the physical network.

   The traffic types are management, public, guest, and storage traffic. For more information about the types, roll over the icons to display their tool tips, or see Basic Zone Network Traffic Types. This screen starts out with some traffic types already assigned. To add more, drag and drop traffic types onto the network. You can also change the network name if desired.

3. Assign a network traffic label to each traffic type on the physical network.

   These labels must match the labels you have already defined on the hypervisor host. To assign each label, click the Edit button under the traffic type icon. A popup dialog appears where you can type the label, then click OK.

   These traffic labels will be defined only for the hypervisor selected for the first cluster.

4. Click Next.

5. (NetScaler only) If you chose the network offering for NetScaler, you have an additional screen to fill out. Provide the requested details to set up the NetScaler, then click Next.

   - **IP address.** The NSIP (NetScaler IP) address of the NetScaler device.

   - **Username/Password.** The authentication credentials to access the device. CloudPlatform uses these credentials to access the device.

   - **Type.** NetScaler device type that is being added. It could be NetScaler VPX, NetScaler MPX, or NetScaler SDX. For a comparison of the types, see About Using a NetScaler Load Balancer.

   - **Public interface.** Interface of NetScaler that is configured to be part of the public network.

   - **Private interface.** Interface of NetScaler that is configured to be part of the private network.

   - **Number of retries.** Number of times to attempt a command on the device before considering the operation failed. Default is 2.

   - **Capacity.** Number of guest networks/accounts that will share this NetScaler device.

   - **Dedicated.** When marked as dedicated, this device will be dedicated to a single account. When Dedicated is checked, the value in the Capacity field has no significance – implicitly, its value is 1.

6. (NetScaler only) Configure the IP range for public traffic. The IPs in this range will be used for the static NAT capability which you enabled by selecting the network offering for NetScaler with EIP and ELB. Enter the following details, then click Add. If desired, you can repeat this step to add more IP ranges. When done, click Next.

   - **Gateway.** The gateway in use for these IP addresses.

- **Netmask.** The netmask associated with this IP range.
- **VLAN.** The VLAN that will be used for public traffic.
- **Start IP/End IP.** A range of IP addresses that are assumed to be accessible from the Internet and will be allocated for access to guest VMs.

7. In a new zone, CloudPlatform adds the first pod for you. You can always add more pods later.

   To configure the first pod, enter the following, then click Next:

   - **Pod Name.** A name for the pod.
   - **Reserved system gateway.** The gateway for the hosts in that pod.
   - **Reserved system netmask.** The network prefix that defines the pod's subnet. Use CIDR notation.
   - **Start/End Reserved System IP.** The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP. For more information, see System Reserved IP Addresses.

8. Configure the network for guest traffic. Provide the following, then click Next:

   - **Guest gateway**: The gateway that the guests should use.
   - **Guest netmask**: The netmask in use on the subnet the guests will use.
   - **Guest start IP/End IP**: Enter the first and last IP addresses that define a range that CloudPlatform can assign to guests.
     - We strongly recommend the use of multiple NICs. If multiple NICs are used, they may be in a different subnet.
     - If one NIC is used, these IPs should be in the same CIDR as the pod CIDR.

9. In a new pod, CloudPlatform adds the first cluster for you. You can always add more clusters later. For an overview of what a cluster is, see About Clusters.

   To configure the first cluster, enter the following, then click Next:

   - **Hypervisor**: The type of hypervisor software that all hosts in this cluster will run.
   - **Cluster name**: Enter a name for the cluster. This can be text of your choosing and is not used by CloudPlatform.

10. In a new cluster, CloudPlatform adds the first host for you. You can always add more hosts later. For an overview of what a host is, see About Hosts.

    > **Note**
    >
    > When you add a hypervisor host to CloudPlatform, the host must not have any VMs that are already running.

Before you can configure the host, you need to install the hypervisor software on the host. You will need to know which version of the hypervisor software version is supported by CloudPlatform and what additional configuration is required to ensure the host will work with CloudPlatform. For more information, refer to the **Chapter 6 Configuring LXC for CloudPlatform** section in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Hypervisor Configuration Guide*.

To configure the first host, enter the following, then click Next:

- **Host Name**: The DNS name or IP address of the host.
- **Username**: The username is root.
- **Password**: This is the password for the user named above, from your LXC install.
- **Host Tags**: (Optional) Any labels that you use to categorize hosts for ease of maintenance. For example, you can set this to the cloud's HA tag, set in the ha.tag global configuration parameter, if you want this host to be used only for VMs with the high availability feature enabled. For more information, refer to the **HA-Enabled Virtual Machines** and the **HA for Hosts** sections in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Administration Guide*.

11. In a new cluster, CloudPlatform adds the first primary storage server for you.

You can always add more servers later. For an overview of what primary storage is, see About Primary Storage.

To configure the first primary storage server, enter the following, then click Next:

- **Name**: The name of the storage device.
- **Protocol**: For LXC, choose NFS or SharedMountPoint. The remaining fields in the screen vary depending on what you choose here.

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

7.2.2.2. Advanced Zone Configuration

1. After you select **Advanced** in the **Add Zone** wizard and click **Next**, you can enter the following details. After you enter these details, click **Next**.

- **Name:** A name for the zone.
- **IPv4 DNS 1** and **IPv4 DNS 2**: These are DNS servers for use by guest VMs in the zone. These DNS servers will be accessed via the public network that you will add later. The public IP addresses for the zone must have a route to the DNS server named here.
- **Internal DNS 1** and **Internal DNS 2**: These are DNS servers for use by system VMs in the zone. These are the VMs used by CloudPlatform itself, such as virtual routers, console proxies, and Secondary Storage VMs. These DNS servers will be accessed via the management traffic

network interface of the System VMs. The private IP address you provide for the pods must have a route to the internal DNS server named here.

- **Hypervisor**: Choose **LXC** as the hypervisor for the first cluster in the zone.
- **Network Domain:** (Optional) If you want to assign a special domain name to the guest VM network, specify the DNS suffix.
- **Guest CIDR:** This is the CIDR that describes the IP addresses in use in the guest virtual networks in this zone. For example, 10.1.1.0/24. As a matter of good practice you should set different CIDRs for different zones. This will make it easier to set up VPNs between networks in different zones.
- **Public:** A public zone is available to all users. A zone that is not public will be assigned to a particular domain. Only users in that domain will be allowed to create guest VMs in this zone.

2. Select the traffic type that the physical network will carry and click **Next**.

   The traffic types are management, public, guest, and storage traffic. For more information about the types, roll over the icons to display their tool tips. This screen displays that one network has been already configured. If you have multiple physical networks, you need to add traffic types to them. Drag and drop traffic types onto a greyed-out physical network to make it active. You can move the traffic type icons from one network to another; for example, if the default traffic types shown for Network 1 do not match your actual setup, you can move them down. You can also change the network names if desired.

3. Click the **Edit** button under each traffic type icon that you added to the physical netwok to assign a network traffic label to it. These labels must match the labels you have already defined on the hypervisor host. The **Edit traffic type** dialog appears where you can enter the label and click **OK**.

   These traffic labels will be defined only for the hypervisor selected for the first cluster.

4. Click **Next**.

5. Configure the IP range for the public (Internet) traffic. Enter the following details and click **Add**. If desired, you can repeat this step to add more public Internet IP ranges. Click **Next**.

   - **Gateway:** The gateway in use for these IP addresses.
   - **Netmask:** The netmask associated with this IP range.
   - **VLAN:** The VLAN that will be used for public traffic.
   - **Start IP/End IP:** A range of IP addresses that are assumed to be accessible from the Internet and will be allocated for access to guest networks.

6. In a new zone, CloudPlatform adds the first pod. You can always add more pods later.

   To configure the first pod, enter the following and click **Next**:

   - **Pod Name:** A name that you can use to identify the pod.
   - **Reserved system gateway:** The gateway for the hosts in that pod.
   - **Reserved system netmask:** The network prefix that defines the pod's subnet. Use CIDR notation.

- **Start/End Reserved System IP:** The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP.
- Specify a range of VLAN IDs to carry guest traffic for each physical network and click **Next**.
- In the new pod, CloudPlatform adds the first cluster. You can always add more clusters later.

  To configure the first cluster, enter the following and click **Next**:

  - **Hypervisor:** The type of hypervisor software that all hosts in this cluster will run.
  - **Cluster name:** Enter a name for the cluster.
- In a new cluster, CloudPlatform adds the first host. You can always add more hosts later.

  > **Note**
  >
  > When you deploy CloudPlatform, the hypervisor host must not have any VMs running on it.

  Before you configure the host, you need to install the hypervisor software on the host. You will need to know which version of the hypervisor software version is supported by CloudPlatform and what additional configuration is required to ensure the host will work with CloudPlatform. For more information, refer to the **Chapter 4 Configuring LXC for CloudPlatform** section in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Hypervisor Configuration Guide*.

  To configure the first host, enter the following, then click **Next**:

  - **Host Name:** The DNS name or IP address of the host.
  - **Username:** Usually root.
  - **Password.** This is the password associated with the user name (from your KVM install).
  - **Host Tags:** (Optional) Any labels that you use to categorize hosts for ease of maintenance. For example, you can set to the cloud's HA tag (set in the **`ha.tag`** global configuration parameter) if you want this host to be used only for VMs with the "high availability" feature enabled. For more information, refer to the **HA-Enabled Virtual Machines** and the **HA for Hosts** sections in the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Administration Guide*.
- In a new cluster, CloudPlatform adds the first primary storage server. You can always add more servers later.

  To configure the first primary storage server, enter the following and click **Next**:

  - **Name:** The name of the storage device.
  - **Protocol:** For LXC, choose NFS, SharedMountPoint, or RBD Ceph.

RBD	**Note** For Data disk, only RBD storage is supported on LXC. For root volume, use NFS or local storage. • **Provider** : Select a provider of the RBD storage. • **RADOS Monitor**: Specify the IP address of the RADOS Monitoring Daemon server. RADOS Monior is responsible for handling communication with all external applications, and handles the decision making and health check of the RBD cluster. Ideally, three Monitor daemons should be run on three separate physical machines, isolated from each other; for example, in different racks. • **RADOS Pool**: Specify the RADOS pool: RADOS pool represents individual fragments of object store. It is a logical layer of binary data tagged as a pool. Pools allow configurations in which individual users can only access specific pools, for example, metadata,data, and rbd are available only in the default configuration. • **RADOS User**: Specify the RADOS user you created. The user with the necessary permissions to access the RADOS pool for CloudPlatform. Although you could use client.admin credentials, it’s recommended to create a user with access only to the CloudPlatform pool. • **RADOS Secret**: The secret pass code associated with the RADOS user you specified above. • **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
NFS	• **Server:** The IP address or DNS name of the storage device. • **Path:** The exported path from the server.

	• **Tags (optional):** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.
SharedMountPoint	• **Path:** The path on each host that is where this primary storage is mounted. For example, "/mnt/primary". • **Tags (optional):** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.

- In a new zone, CloudPlatform adds the first secondary storage server.

  Before you enter information in the fields on this screen, you need to prepare the secondary storage by setting up NFS shares and installing the latest CloudPlatform System VM template.

  To configure the first secondary storage server on hosts, enter the following and click **Next**:

  - **NFS Server:** The IP address of the server.
  - **Path:** The exported path from the server.

- Click **Launch**.

7.3. Adding a Pod

After you create a new zone, CloudPlatform adds the first pod to it. You can perform the following procedure to add more pods to the zone at anytime.

1. Log-in to the CloudPlatform UI.
2. In the left navigation bar, select **Infrastructure**.
3. In the right-side panel, under **Zones**, click **View all**.
4. In the page that lists the zones that you configured, click the zone where you want to add a pod.
5. Click the **Compute and Storage** tab. At the **Pods** node in the diagram, click **View all**.
6. In the page that lists the pods configured in the zone that you selected, click **Add Pod**.
7. In the **Add Pod** dialog box, enter the following details:
   - **Zone:** Select the name of the zone where you want to add the new pod.
   - **Pod name:** The name that you can use to identify the pod.

- **Reserved system gateway:** The gateway for the hosts in that pod.
- **Reserved system netmask:** The network prefix that defines the pod's subnet. Use CIDR notation.
- **Start/End Reserved System IP:** The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP. For more information, see System Reserved IP Addresses.
- **Dedicate:** Select if you want to dedicate the pod to a specific domain.
- **Domain:** Select the domain to which you want to dedicate the pod.
- **Account:** Enter an account name that belongs to the above selected domain so that you can dedicate the pod to this account.

8. Click **OK**.

7.4. Adding an LXC Cluster

You need to tell CloudPlatform about the hosts that it will manage. Hosts exist inside clusters, so before you begin adding hosts to the cloud, you must add at least one cluster.

Before you perform the following steps, you must have installed the hypervisor on the hosts and logged-in to the CloudPlatform UI.

1. In the CloudPlatform UI, in the left navigation bar, click **Infrastructure**.
2. In the right-side panel, under **Zones**, click **View all**.
3. In the page that lists the zones that you have configured, click the zone where you want to add the cluster.
4. In the details page of the zone, click the **Compute and Storage** tab.
5. At the **Clusters** node of the diagram, click **View all**.
6. In the page that list the clusters, click **Add Cluster**.
7. In the **Add Cluster** dialog box, do the following and click **OK**:
   - **Zone Name:** Select the zone where you want to create the cluster.
   - **Hypervisor:** Select LXC as the hypervisor for this cluster.
   - **Pod Name:** Select the pod where you want to create the cluster.
   - **Cluster Name:** Enter a name that you can use to identify the cluster.
   - **Dedicate**: Select if you want to dedicate the cluster to a specific domain.
   - **Domain**: Select the domain to which you want to dedicate the cluster.
   - **Account**: Enter an account name that belongs to the domain so that you can dedicate the cluster to this account.

7.4.1. (Experimental Feature) Adding an LXC Host

LXC hosts can be added to a cluster at any time.

7.4.1.1. Requirements for LXC Hosts

Configuration requirements:

- Make sure that the host does not have any running VMs before you add it to CloudPlatform.
- Each cluster must contain only hosts with the identical hypervisor.
- Do not put more than 16 hosts in a cluster.
- If shared mountpoint storage is in use, the administrator should ensure that the new host has all the same mountpoints (with storage mounted) as the other hosts in the cluster.
- Make sure the new host has the same network configuration (guest, private, and public network) as other hosts in the cluster.

For hardware requirements, see the installation section for your hypervisor in the CloudPlatform Installation Guide.

7.4.1.2. (Optional)Enabling RBD-Based Primary Storage on LXC:

If you want to use RBD-based data disks on LXC, perform the following on each VM (containers):

For more information, see *Get Ceph Packages*[1]:

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. From ceph monitor node, copy ceph.conf and ceph admin key to the **`/etc/ceph`** directory.

7.4.1.3. Adding an LXC Host

1. Log in to the CloudPlatform UI as administrator.
2. In the left navigation, choose Infrastructure.
3. In Zones, click View More, then click the zone in which you want to add the host.
4. Click the Compute tab. In the Clusters node, click View All.

[1] http://ceph.com/docs/v0.80.5/install/get-packages/

5. Click the cluster where you want to add the host.

6. Click View Hosts.

7. Click Add Host.

8. Provide the following information.

   - **Host Name**: The DNS name or IP address of the host.

   - **Username**: Usually root.

   - **Password**: This is the password for the user named above from your LXC install.

   - **Host Tags** (Optional): Any labels that you use to categorize hosts for ease of maintenance. For example, you can set to the cloud's HA tag, which is set in the ha.tag global configuration parameter, if you want this host to be used only for VMs (containers) with the high availability feature enabled. For more information, see HA-Enabled Virtual Machines as well as HA for Hosts sections.

   - **Dedicate**: When marked as dedicated, this host is dedicated to a single account that you select.

   There may be a slight delay while the host is provisioned. It should automatically displayed in the UI.

9. Repeat these steps for adding additional hosts.

7.5. Adding Primary Storage

> **Warning**
>
> When using preallocated storage for primary storage, be sure there is nothing on the storage, for example, you have an empty SAN volume or an empty NFS share. Adding the storage to CloudPlatform will destroy any existing data.

When you create a new zone, the first primary storage is added as part of that procedure. You can add primary storage servers at any time, such as when adding a new cluster or adding more servers to an existing cluster.

1. Log in to the CloudPlatform UI.

2. In the left navigation, choose Infrastructure. In Zones, click View All, then click the zone in which you want to add the primary storage.

3. Click the Compute and Storage tab.

4. In the Primary Storage node of the diagram, click View All.

5. Click Add Primary Storage.

6. Provide the following information in the dialog. The information required varies depending on your choice in Protocol.

- **Scope**: Indicate whether the storage is available to all hosts in the zone or only to hosts in a single cluster.
- **Pod**: Visible only if you choose Cluster in the Scope field. The pod for the storage device.
- **Cluster**: Visible only if you choose Cluster in the Scope field. The cluster for the storage device.
- **Name**: The name of the storage device.
- **Protocol**: For LXC, choose NFS, RADOS Block Devices (RBD) or SharedMountPoint.

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

- **Server**: (NFS)The IP address or DNS name of the storage device.
- **Path**: (NFS) On NFS this is the exported path from the server.
- **Path**: (SharedMountPoint) On LXC, this is the path on each host where this primary storage is mounted. For example, "/mnt/primary".
- **Provider**: (RBD) Select a provider of the RBD storage.
- **RADOS Monitor**: (RBD) Specify the IP address of the RADOS Monitoring Daemon server.

  RADOS Monior is responsible for handling communication with all external applications, and handles the decision making and health check of the RBD cluster. Ideally, three Monitor daemons should be run on three separate physical machines, isolated from each other; for example, in different racks.

- **RADOS Pool**: (RBD) Specify the RADOS pool:

  RADOS pool represents individual fragments of object store. It is a logical layer of binary data tagged as a pool. Pools allow configurations in which individual users can only access specific pools, for example, metadata,data, and rbd are available only in the default configuration.

- **RADOS User**: (RBD) Specify the RADOS user you created.

  The user with the necessary permissions to access the RADOS pool for CloudPlatform. Although you could use client.admin credentials, it’s recommended to create a user with access only to the CloudPlatform pool.

- **RADOS Secret**: (RBD) The secret pass code associated with the RADOS user you specified above.
- **Tags**: (optional) The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings

The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.

7. Click OK.

7.6. Adding Secondary Storage

Note

Be sure there is nothing stored on the server. Adding the server to CloudPlatform will destroy any existing data.

When you create a new zone, the first secondary storage is added as part of that procedure. You can add secondary storage servers at any time to add more servers to an existing zone.

1. To prepare for the zone-based Secondary Storage, you should have created and mounted an NFS share during Management Server installation.

2. Make sure you prepared the system VM template during Management Server installation.

3. Log in to the CloudPlatform UI as root administrator.

4. In the left navigation bar, click Infrastructure.

5. In Secondary Storage, click View All.

6. Click Add Secondary Storage.

7. Fill in the following fields:

   - **Name**: Give the storage a descriptive name.
   - **Provider**: Choose the type of storage provider.
   - **Create NFS Secondary Storage**: Be sure this box is checked, unless the zone already contains a secondary staging store.

     Warning

     If you are setting up a new zone, be sure the box is checked. This checkbox and the three fields below it must be filled in.

   - **Zone**: The zone where the NFS Secondary Storage is to be located.
   - **Server.** The IP address or DNS name of the storage device.

- **Path.** The exported path from the server.
- **NFS server**: The name of the zone's Secondary Storage.
- **Path**: The path to the zone's Secondary Storage.

7.6.1. Adding an NFS Secondary Storage for Each Zone

You can skip this section if you are upgrading an existing zone from NFS to object storage. You only need to perform the steps below when setting up a new zone that does not yet have its NFS server.

Every zone must have at least one NFS store provisioned; multiple NFS servers are allowed per zone. To provision an NFS Staging Store for a zone:

1. To prepare for the zone-based Secondary Storage, you should have created and mounted an NFS share during Management Server installation.
2. Make sure you prepared the system VM template during Management Server installation.
3. Log in to the CloudPlatform UI as root administrator.
4. In the left navigation bar, click Infrastructure.
5. In Secondary Storage, click View All.
6. In Select View, choose Secondary Storage.
7. Click the Add NFS Secondary Storage button.
8. Fill out the dialog box fields, then click OK:
    - Zone. The zone where the NFS Secondary Storage is to be located.
    - NFS server. The name of the zone's Secondary Storage.
    - Path. The path to the zone's Secondary Storage.

7.6.2. Upgrading from NFS to Object Storage

In an existing zone that is using NFS for secondary storage, you can upgrade the zone to use a region-wide object storage without causing downtime. The existing NFS storage in the zone will be converted to an NFS Staging Store.

After upgrade, all newly created templates, ISOs, volumes, snapshots are moved to the object store. All previously created templates, ISOs, volumes, snapshots are migrated on an on-demand basis based on when they are accessed, rather than as a batch job. Unused objects in the NFS staging store are garbage collected over time.

1. Log in as admin to the CloudPlatform UI.
2. Fire an admin API to update CloudPlatform to use object storage:

```
http://<MGMTIP>:8096/client/api?command=updateCloudToUseObjectStore&name=<S3
 storage name>&provider=S3&details[0].key=accesskey&details[0].value=<access
 key from .s3cfg file>&details[1].key=secretkey&details[1].value=<secretKey
 from .s3cfg file>&details[2].key=bucket&details[2].value=<bucketname>&details[3].
 key=usehttps&details[3].value=<trueorfalse>&details[4].key=endpoint&details[4].
 value=<S3 server IP:8080>
```

All existing NFS secondary storages has been converted to NFS staging stores for each zone, and your S3 object store specified in the command has been added as a new region-wide secondary storage.

3. Locate the secondary storage that you want to upgrade to object storage.

   Perform either of the following in the Infrastructure page:

   - In Zones, click View All, then locate the desired zone, and select Secondary Storage in the Compute and Storage tab.
   - In Secondary Storage, click View All, then select the desired secondary storage.

Post migration, consider the following:

- For each new snapshot taken after migration, ensure that you take a full snapshot to newly added S3.

  This would help coalesce delta snapshots across NFS and S3 stores. The snapshots taken before migration are pushed to S3 store when you try to use that snapshot by executing createVolumeCmd by passing snapshot id.
- You can deploy VM from templates because they are already in the NFS staging store. A copy of the template is generated in the new S3 object store when ExtractTemplate or CopyTemplate is executed.
- For volume, a copy of the volume is generated in the new S3 object store when ExtractVolume command is executed.
- All the items in the NFS storage is not migrated to S3. Therefore, if you want to completely shut down the NFS storage you have previously used, write your own script to migrate those remaining items to S3.

7.7. Initialize and Test

After you complete all the configuration, CloudPlatform starts the initialization. This can take 30 minutes or more, depending on the speed of your network. When the initialization has completed successfully, the administrator's dashboard should be displayed in the CloudPlatform UI.

1. Verify that the system is ready. In the left navigation bar, click **Templates**. Click the template. Check to be sure that the status is "Download Complete." Do not proceed to the next step until this status is displayed.
2. Go to the Instances tab, and filter by My Instances.
3. Click Add Instance and follow the steps in the wizard.

   a. Choose the zone you just added.

   b. In the template selection, choose the template to use in the VM. If this is a fresh installation, likely only the provided CentOS template is available.

   c. Select a service offering. Be sure that the hardware you have allows starting the selected service offering.

   d. In data disk offering, if desired, add another data disk. This is a second volume that will be available to but not mounted in the guest.

  For example, in Linux on XenServer you will see /dev/xvdb in the guest after rebooting the VM. A reboot is not required if you have a PV-enabled OS kernel in use.

  e. In default network, choose the primary network for the guest. In a trial installation, you would have only one option here.

  f. Optionally give your VM a name and a group. Use any descriptive text you would like.

  g. Click Launch VM. Your VM will be created and started. It might take some time to download the template and complete the VM startup. You can watch the progress in the Instances screen.

4. To use the VM, click the View Console button. >_

   For more information about using VMs, including instructions for how to allow incoming network traffic to the VM, start, stop, and delete VMs, and move a VM from one host to another, see

Congratulations! You have successfully completed a CloudPlatform Installation.

If you decide to grow your deployment, you can add more hosts, primary storage, zones, pods, and clusters.

Provisioning Cloud Infrastructure on Baremetal

This section describes how to add a host and networks to your cloud using Baremetal hypervisor. Note that Baremetal host requires a VMware advanced zone up and running.

For conceptual information about zones, pods, clusters, hosts, and networks in CloudPlatform, refer to *CloudPlatform (powered by CloudStack) Version 4.5 Concepts Guide*.

8.1. Adding a Zone

Adding a zone consists of three phases:

- Consider the Baremetal limitations.
- Ensure that you meet the prerequisites.
- Add the zone.

8.1.1. Limitations of Baremetal Deployment

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.

8.1.2. Adding a New Zone

To add a zone, you must first configure the zone’s physical network. Then, you need to add the first pod, cluster, host, primary storage, and secondary storage.

Note

Before you proceed with adding a new zone, you must ensure that you have performed the steps to seed the system VM template.

Note

Citrix strongly recommends using the same type of hypervisors in a zone to avoid operational issues.

1. Log-in to the CloudPlatform UI using the root administrator account.

2. In the left navigation bar, click **Infrastructure**.

3. In the right side panel, under **Zones**, click **View all**.

4. In the next page, click **Add Zone**.

   The **Add zone** wizard panel appears.

5. Select one of the following network types:

   - **Basic**: (For AWS-style networking). Provides a single network where each VM instance is assigned with an IP directly from the network. You can provide guest isolation through layer-3 means such as security groups (IP address source filtering).

   - **Advanced**: Used for more sophisticated network topologies. This network model provides the most flexibility in defining guest networks and providing custom network offerings such as firewall, VPN, or load balancer support.

   For more information, refer to **Chapter 4. Cloud Infrastructure Concepts** of the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Concepts Guide*.

6. Based on the option (Basic or Advanced) that you selected, do one of the following:

   - For **Basic**: *Section 8.1.2.1, “Configuring Basic Zone for Baremetal”*

   - For **Advanced**: *Section 8.1.2.2, “(Experimental Feature) Configuring Advanced Networking in Baremetal ”*

8.1.2.1. Configuring Basic Zone for Baremetal

Follow the steps in all the following sections in order.

1. *Section 8.1.2.1.1, “Prerequisites for Configuring a Baremetal Basic Zone”*

2. *Section 8.1.2.1.2, “Creating a Basic Zone ”*

8.1.2.1.1. Prerequisites for Configuring a Baremetal Basic Zone

To set up basic networking zone in Baremetal cloud, you must first create a network offering and using which create a shared network. Not all network services are supported in basic zone. Security groups are supported, for which install the Security group agent.

- *Section 8.1.2.1.1.1, “Adding the PXE Server and DHCP Server to Your Deployment”*
- *Section 8.1.2.1.1.2, “Creating a Baremetal Network Offering ”*
- *Section 8.1.2.1.1.3, “Setting Up the Security Group Agent (Optional)”*

8.1.2.1.1.1. Adding the PXE Server and DHCP Server to Your Deployment

As part of describing your deployment to CloudPlatform, you will need to add the PXE server and DHCP servers.

1. Ensure that you have installed the PXE and DHCP Servers.

   For more information, see *CloudPlatform Hypervisor Configuration Guide*.

2. Log in as admin to the CloudPlatform UI.

3. In the left navigation, choose Infrastructure. In Zones, click View All, then click the zone in which you want to add the Baremetal PXE / DHCP server.

4. Click the Physical Network tab. Click the Physical Network entry.

5. In the Network node, click Network Service Providers Configure.

6. In the list of Network service providers, clik Baremetal PXE. In the Details node, click Add Baremetal PXE Device.

   The Add Baremetal PXE Device dialog will appear.

7. In the Add Baremetal PXE Device dialog, make the following choices:

   - URL: http://<PXE DHCP server IP address>
   - Username: login username
   - Password: password
   - Tftp root directory: /var/lib/tftpboot

8. In the list of Network service providers, click Baremetal DHCP. In the Details node, click Add Baremetal DHCP Device button. The Add Baremetal DHCP Device dialog will appear.

9. In the Add Baremetal DHCP Device dialog:

   - URL: http://<PXE DHCP server IP address>
   - Username: login username
   - Password: password

8.1.2.1.1.2. Creating a Baremetal Network Offering

1. Log in as admin to the CloudPlatform UI.

2. In the left navigation bar, click Service Offerings.

3. In Select Offering, choose Network Offering.

4. Click Add Network Offering.

5. In the dialog, make the following choices:

    - Name: You can give the offering any desired name. For example, Baremetal.

    - Guest Type: Shared

    - Supported Services:

        - DHCP checkbox: checked

        - DHCP Provider: Baremetal

        - User Data checkbox: checked

        - User Data Provider: Baremetal

        - Security Groups: checked

        - BaremetalPxeServer: checked

    - Additional choices in this dialog are described in "Creating a New Network Offering" in the Administrator's Guide.

6. Click OK.

7. Verify:

    a. In the left navigation bar, click Service Offerings.

    b. In the Select Offering dropdown, choose Network Offerings.

    c. Click the name of the offering you just created, and check the details. In State, be sure the offering is Enabled. If not, click the Enable button.

8.1.2.1.1.3. Setting Up the Security Group Agent (Optional)

If you are not using security groups, you can skip this section.

If you plan to use security groups to control traffic to Baremetal instances, you need to install security group agent software on each Baremetal host. This involves downloading the software, making it available in an accessible repository, and modifying the kickstart file to go get this software during installation.

1. Download the agent software from the following link:

    http://download.cloud.com/releases/4.2.0/cloudstack-baremetal-agent-4.2.0-1.el6.x86_64.rpm

    The agent software depends on several other RPMs:

    - python-cherrypy: A Python HTTP server which is distributed by default with most Linux distributions. For example, both CentOS and Ubuntu have this package.

    - ipset: An iptables tool which provides ipset match. In Ubuntu, ipset is provided by default. In Cent OS, it is not provided by default; you need to download it from a third party. For example: *http://www.wandin.net/dotclear/index.php?post/2012/05/26/Installing-ipset-on-CentOS-6*

- libmnl: ipset dependent library. it's usually available with ipset rpm for downloading

2. Place the RPMs in a directory that is accessable by browser through HTTP.

3. Create a repo where the kickstart installer can find the security group agent when it's time to install. Run the following command:

   ```
   # createrepo <path_to_rpms>
   ```

   For example, if the RPMs are in the following directory:

   ```
   /var/www/html/securitygroupagent/
   ```

   The command would be:

   ```
   createrepo /var/www/html/securitygroupagent/
   ```

   The repo file will be created in /var/www/html/securitygroupagent/.

4. Add the security group agent package to the kickstart file that you are using for your Baremetal template. Make the following modifications in the kickstart file:

   a. Add the repo that you created in the previous step. Insert the following command above the %package section, before reboot. Substitute the desired repo name, IP address, and the directory in the base URL (if you didn't use the name securitygroupagent in the previous step).

   ```
   repo --name=<repo_name> --baseurl=http://<ip_address>/securitygroupagent/
   ```

   b. In the %package section of the kickstart file, add all the RPMs. For example:

   ```
   %package
   libmnl
   ipset
   python-cherrypy
   security_group_agent
   ```

   c. In the %post section, add the following:

   ```
   %post
   chkconfig iptables off
   chkconfig cs-sgagent on
   service cs-sgagent start
   ```

   This will close iptables to flush the default iptables rules set by the OS (CloudPlatform does not need them), then set the security group agent to "on" and immediately start the agent.

8.1.2.1.2. Creating a Basic Zone

Your cluster(s) of Baremetal hosts must be organized into a zone. This zone can contain only Baremetal host. You can have one or more Baremetal zones in your cloud.

1. After you select **Basic** in the **Add Zone** wizard and click **Next**, you need to enter the following details. After you enter the details, click **Next**.

   - Name. A name for the zone.
   - DNS 1 and 2. These are DNS servers for use by guests in the zone. These DNS servers will be accessed via the public network you will add later. The public IP addresses for the zone must have a route to the DNS server named here.
   - Hypervisor. Choose Baremetal.
   - Network Offering. Choose the network offering you created in *Section 8.1.2.1.1.2, “Creating a Baremetal Network Offering ”*.
   - Network Domain: (Optional) If you want to assign a special domain name to the guest VM network, specify the DNS suffix.
   - Public. A public zone is available to all users. A zone that is not public will be assigned to a particular domain. Only users in that domain will be allowed to use this zone.

2. Click Next.

3. In a new zone, CloudPlatform adds the first pod for you. You can always add more pods later.

   To configure the first pod, enter the following:

   - Pod Name. A name for the pod.
   - Reserved system gateway. The gateway for the hosts in that pod.
   - Reserved system netmask. The network prefix that defines the pod's subnet. Use CIDR notation.
   - Start/End Reserved System IP. The IP range in the management network that CloudPlatform uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP.

4. Click Next, or OK.

   The UI will show a progress indicator.

   Troubleshooting: After a few moments, if this indicator does not finish, click Refresh in the browser.

5. Be sure the zone is enabled:

   a. In the left navigation bar, click Infrastructure.

   b. In Zones, click View All.

   c. Click the name of the zone you just created, and check the details. In Allocation State, be sure the zone is Enabled. If not, click the Enable button.

8.1.2.2. (Experimental Feature) Configuring Advanced Networking in Baremetal

CloudPlatform supports Advanced networking capabilities for Baremetal guest VMs. A new plugin has been introduced in CloudPlatform to enable automatic VLAN programming on a physical switch to which Baremetal VMs are connected. In an Advanced zone, Barmetal VMs gain VLAN isolation provided by CloudPlatform which is particularly used to provide Baremetal As a Service with advanced networking capabilities for public clouds. Baremetal As a Service cannot function standalone; it works in confunction with a physical switch either from vendor's SDK or from an in-switch agent for whitebox switch. A new global parameter, `baremetal.vlan.sync.interval`, has been introduced to control the interval of periodical tasks of syncing VLAN configuration on switch ports and VLAN allocation in CloudPlatform database.

The following networking capabilities are provided:

- SourceNAT
- VPN
- Load Balancing
- Firewall

8.1.2.2.1. Prerequisites and Guidelines

- Advanced zone with a VMware cluster is up and running. SSVM and CPVM are up.
- 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. To use XenServer/KVM with Baremetal, deploy a Baremetal instance before deploying any instance to ensure that VR is created on the VMware host. For VR HA, continue using the same means of virtualization. The VR template size is increased from 2G to 4G in order to accomodate required number of kernel/initrd of Baremetal templates.
- Baremetal uses VR to provide all network services, such as PXE/DHCP, StaticNAT, and Port Forwarding.
- If Baremetal host has multiple NICs, only one NIC which is part of the same VLAN gets the IP address; however, you can configure kickstart file to set a static IP or use a DHCP server outside CloudPlatform for extra NICs.
- Baremetal cluster is a simple aggregation of hosts which have similar hardware. We recommend you to place the hosts of the same cluster in the same IPMI subnet. Connect all the guest NICs, the NIC that plays role as guest NIC after provisioning the guest OS, to the same TOR switch.

8.1.2.2.2. Known Limitations

- Concepts of CloudPlatform storage, including primary storage, secondary storage, volume, and snapshot are not supported.
- Link Layer Discovery Protocol (LLDP) is not supported.
- IPv6 is not supported.
- Deploying more than 10 Baremetal instances concurrently is not supported.
- (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.

8.1.2.2.3. Deploying Baremetal Instances in an Advanced Zone

1. Set IP of the internal HTTP server in the network offering and in the *`baremetal.internal.storage.server.ip`* global parameter.

   This is location where kickstart file, kernel, initrd, ISO for advanced networking baremetal provisioning are stored.

   The IP specified in the network offering overrides the one specified in the global parameter.

2. Create the HTTP Rack Configuration Repo.

   This is the text file, referred as RCT, in JSON format that specifes the L2 switch identity and switch-host-port mapping.

   For more information, see **HTTP Rack Configuration Repo** in the *CloudPlatform Hypervisor Configuration Guide*.

3. Specify the URL of the RCT file in the Global Settings.

   a. Log in to the CloudPlatform UI.

   b. From the left navigational bar, click Global Settings.

   c. From the Select view drop down, select Baremetal Rack Configuration.

   d. Click Add Baremetal Rack Configuration.

      The Add Baremetal Rack Configuration dialog is displayed.

   e. Specify the URL of the location where the RCT configuration file is stored. For example: http://10.20.20.10/baremetal/rct.json.

   f. Click OK.

4. Create a compute offering.

   Fore more information, see *Section 8.1.2.2.4, “Create a Baremetal Compute Offering”*.

5. Create a network offering with PXE and DHCP services, and VR as the service provider.

   For more information, see *Section 8.1.2.2.5, “Creating a Baremetal Network Offering”*.

6. Create an isolated network with the above network offering.

   For more information, see *Section 8.1.2.2.6, “Creating an Isolated Network”*.

7. Add a Baremetal cluster in a VMware Advanced zone.

8. Add a Baremetal host to the cluster.

9. Deploy Baremetal instances in the isolated network you have created.

8.1.2.2.4. Create a Baremetal Compute Offering

1. Log in as admin to the CloudPlatform UI at the URL below. Substitute the IP address of your own Management Server:

   ```
   http://<management-server-ip-address>:8080/client
   ```

2. In the left navigation bar, click Service Offerings.

3. In Select Offering, choose Compute Offerings.

4. Click Add compute offering.

5. In the dialog box, fill in these values:

   - **Name**: Any desired name for the service offering.

   - **Description**: A short description of the offering that can be displayed to users.

   - **Storage Type**: Shared.

   - **Custom**: Custom compute offerings can be used in following cases: deploying a VM, changing the compute offering of a stopped VM and running VMs, which is nothing but scaling up.

     If the Custom field is checked, during the VM deployment following parameters need to be specified:

     - **# of CPU cores**: The number of cores which should be allocated to a system VM with this offering. Use the same value as when you added the host.

     - **CPU (in MHz)**: The CPU speed of the cores that the system VM is allocated. For example, “2000” would provide for a 2 GHz clock. Use the same value as when you added the host.

     - **Memory (in MB)**: The amount of memory in megabytes that the system VM should be allocated. For example, “2048” would provide for a 2 GB RAM allocation. Use the same value as when you added the host.

   - **Network Rate**: Allowed data transfer rate in MB per second.

   - **Disk Read Rate**: Allowed disk read rate in bits per second.

   - **Disk Write Rate**: Allowed disk write rate in bits per second.

   - **Disk Read Rate**: Allowed disk read rate in IOPS (input/output operations per second).

   - **Disk Write Rate**: Allowed disk write rate in IOPS (input/output operations per second).

   - **Offer HA**: Unchecked. High availability services are not supported for Baremetal hosts.

   - **Storage Tags**: The tags that should be associated with the primary storage used by the system VM.

   - **Host Tags**: Any tags that you use to organize your hosts.

   - **CPU cap**: Whether to limit the level of CPU usage even if spare capacity is available.

   - **Public**: Indicate whether the service offering should be available all domains or only some domains. Choose Yes to make it available to all domains.

   - **isVolatile**: If checked, VMs created from this service offering will have their root disks reset upon reboot. This is useful for secure environments that need a fresh start on every boot and for desktops that should not retain state.

   - **Deployment Planner**: Choose the technique that you would like CloudPlatform to use when deploying VMs based on this service offering. Select BareMetal Planner.

6. Click OK.

8.1.2.2.5. Creating a Baremetal Network Offering

1. Log in as admin to the CloudPlatform UI.
2. In the left navigation bar, click Service Offerings.
3. In Select Offering, choose Network Offering.
4. Click Add Network Offering.
5. In the dialog, make the following choices based on the type of network offering you wanted to create:

   - Name: You can give the offering any desired name. For example, Baremetal.
   - Guest Type: Select Shared to create an offering for a basic network.

     Select Isolated to create an offering for an advanced network.
   - Supported Services:
     - DHCP checkbox: checked

       DHCP Provider: Select Virtual Router.
     - DNS: checked

       DNS Provider: Select Virtual Router.
     - Firewall: checked

       Firewall Provider: Select Virtual Router.
     - Load Balancer: checked

       Load Balancer Provider: Select Virtual Router.
     - Source NAT: checked

       Source NAT Provider: Select Virtual Router.
     - Static NAT: checked

       Static NAT Provider: Select Virtual Router.
     - Port Forwarding: checked

       Port Forwarding Provider: Select Virtual Router.
     - User Data checkbox: checked

       User Data Provider: Select Virtual Router.
     - Security Groups: Not supported.
     - BaremetalPxeServer: checked

- Additional choices in this dialog are described in "Creating a New Network Offering" in the Administrator's Guide.

6. Click OK.

7. Verify:

    a. In the left navigation bar, click Service Offerings.

    b. In the Select Offering dropdown, choose Network Offerings.

    c. Click the name of the offering you just created, and check the details. In State, be sure the offering is Enabled. If not, click the Enable button.

8.1.2.2.6. Creating an Isolated Network

This section is applicable only if you are using an advanced network for the Baremetal instances. These steps assume you have already logged in to the CloudPlatform UI. To configure the base guest network:

1. In the left navigation, select Network, then click Add Isolated Network.

   The Add Isolated Guest Network window is displayed:

2. Provide the following information:

   - **Name**. The name of the network. This will be user-visible.
   - **Display Text**: The description of the network. This will be displayed to the user.
   - **Zone**: The Baremetal zone in which you are configuring the guest network.
   - **Network offering**: Select the network offering you created in *Section 8.1.2.2.5, “Creating a Baremetal Network Offering”*.
   - **Guest Gateway**: The gateway that the guests should use.
   - **Guest Netmask**: The netmask in use on the subnet the guests will use.
   - **Network Domain**: A custom DNS suffix at the level of a network. If you want to assign a special domain name to the guest VM network, specify a DNS suffix.

3. Click OK.

8.1.2.2.7. Troubleshooting Tips

- To troubleshoot connectivity issues, access the physical switch by using a direct connected cable or remote ssh which is decided by the physical switch vendor. Log in to the physical switch, compare the VLAN on the port where Baremetal VMs connects to and the VLAN allocated to the VMs by CloudPlatform.

- To reduce possible faults caused by misconfigurating the physical switch, CloudPlatform performs a periodical operation to full sync VLAN configuration on the switch ports and VLAN allocation in the CloudPlatform database. You can configure the interval of this operation through the global setting, *`baremetal.vlan.sync.interval`*. The value 0 indicates shutdown Baremetal support in Advanced zone.

8.2. Add a Baremetal Cluster

1. Log in as admin to the CloudPlatform UI.

2. In the left navigation, choose Infrastructure.

3. Click View Clusters, then click Add Cluster.

   The Add Cluster dialog will appear.

4. Specify the following:

   - Zone: Select the zone where VMware cluster is set up.

   - Hypervisor: Select Baremetal.

   - Pod name: Select the desired pod.

   - Cluster name: Enter a name for the cluster. This can be any text you like.

5. Click OK.

8.3. Add a Baremetal Host

1. Log in as admin to the CloudPlatform UI.

2. In the left navigation, click Infrastructure.

3. In Zoes, click View All, then click the name of the Baremetal zone you added earlier.

4. Click the Compute and Storage tab.

5. In Clusters, click View All, then click the name of the Baremetal cluster you added earlier.

6. Click View Hosts.

7. Click the Add Host button.

   The Add Host dialog will appear.

8. In the Add Host dialog, make the following choices:

   - **Host name**: The IPMI IP address of the machine.

   - **Username**: User name you set for IPMI.

   - **Password**: Password you set for IPMI.

   - **# CPU Cores**: Number of CPUs on the machine.

   - **CPU (in MHz)**: Frequency of CPU.

   - **Memory (in MB)**: Memory capacity of the new host.

   - **Host MAC**: MAC address of the PXE NIC.

   - **Host Tags**: Set to large. You will use this tag later when you create the service offering.

It may take a minute for the host to be provisioned. It should automatically display in the UI.

Repeat for additional Baremetal hosts.

8.4. Create a Baremetal Template

In these steps, it is assumed you already have a directory on your NFS server containing the image for the Baremetal instance, as well as the kickstart file.

1. Ensure that you have met the prerequisites for Kickstart Template creation.

   For more information, see *CloudPlatform Hypervisor Configuration Guide*.

2. Log into the UI as either an end user or administrator.

3. In the left navigation bar, click Templates.

4. Click Create Template.

5. In the dialog box, enter the following values.

   - Name. Short name for the template.

   - Display Text. Description of the template.

   - URL. The location of the image file on your NFS server in the format:

     ```
     ks=<http_link_to_kickstart_file>;kernel=<nfs_path_to_pxe_bootable_kernel>
     ;initrd=<nfs_path_to_pxe_initrd>
     ```

     For example:

     ```
     ks=http://nfs1.lab.vmops.com/baremetal/ubuntu.ks;
     kernel=10.10.10.10:/var/www/html/baremetal/linux;initrd=
     10.10.10.10:/var/www/html/baremetal/initrd.gz
     ```

     Note

     The kickstart file is located on an HTTP server. We use the link to it here.

   - Zone: All Zones.

   - OS Type: Select the OS type of the ISO image. Choose other if the OS Type of the ISO is not listed or if the ISO is not bootable.

   - Hypervisor: BareMetal.

   - Format: BareMetal.

   - Password Enabled: No.

   - Public: No.

   - Featured: Choose Yes if you would like this template to be more prominent for users to select. Only administrators may make templates featured.

6. Click OK.

8.5. Create a Baremetal Compute Offering

1. Log in as admin to the CloudPlatform UI at the URL below. Substitute the IP address of your own Management Server:

   ```
   http://<management-server-ip-address>:8080/client
   ```

2. In the left navigation bar, click Service Offerings.

3. In Select Offering, choose Compute Offerings.

4. Click Add compute offering.

5. In the dialog box, fill in these values:

   - **Name**: Any desired name for the service offering.

   - **Description**: A short description of the offering that can be displayed to users.

   - **Storage Type**: Shared.

   - **Custom**: Custom compute offerings can be used in following cases: deploying a VM, changing the compute offering of a stopped VM and running VMs, which is nothing but scaling up.

     If the Custom field is checked, during the VM deployment following parameters need to be specified:

     - **# of CPU cores**: The number of cores which should be allocated to a system VM with this offering. Use the same value as when you added the host.

     - **CPU (in MHz)**: The CPU speed of the cores that the system VM is allocated. For example, “2000” would provide for a 2 GHz clock. Use the same value as when you added the host.

     - **Memory (in MB)**: The amount of memory in megabytes that the system VM should be allocated. For example, “2048” would provide for a 2 GB RAM allocation. Use the same value as when you added the host.

   - **Network Rate**: Allowed data transfer rate in MB per second.

   - **Disk Read Rate**: Allowed disk read rate in bits per second.

   - **Disk Write Rate**: Allowed disk write rate in bits per second.

   - **Disk Read Rate**: Allowed disk read rate in IOPS (input/output operations per second).

   - **Disk Write Rate**: Allowed disk write rate in IOPS (input/output operations per second).

   - **Offer HA**: Unchecked. High availability services are not supported for Baremetal hosts.

   - **Storage Tags**: The tags that should be associated with the primary storage used by the system VM.

   - **Host Tags**: Any tags that you use to organize your hosts.

   - **CPU cap**: Whether to limit the level of CPU usage even if spare capacity is available.

- **Public**: Indicate whether the service offering should be available all domains or only some domains. Choose Yes to make it available to all domains.

- **isVolatile**: If checked, VMs created from this service offering will have their root disks reset upon reboot. This is useful for secure environments that need a fresh start on every boot and for desktops that should not retain state.

- **Deployment Planner**: Choose the technique that you would like CloudPlatform to use when deploying VMs based on this service offering. Select BareMetal Planner.

6. Click OK.

8.6. (Optional) Setting Baremetal Configuration Parameters

1. Log in as admin to the CloudPlatform UI.

2. Click Global Settings.

3. Make any desired modifications to the Baremetal configuration parameters.

   - Basic Zone

     - enable.baremetal.securitygroup.agent.echo (default: false)

     - external.baremetal.resource.classname

     - external.baremetal.system.url

     - interval.baremetal.securitygroup.agent.echo (default: 10)

     - timeout.baremetal.securitygroup.agent.echo (default: 3600)

     - baremetal.internal.http.server.ip

       This is specified with the gateway IP of an instance, which is SourceNATed by the CloudPlatform Management Server to the mangement NIC. The traffic towards the internal HTTP server goes through VR's management NIC instead of public NIC. The source NAT will only bind to guest IP of the provisioning instance to prevent network sniffer from other VMs in the same network.

   - Advanced Zone

     - external.baremetal.resource.classname

     - external.baremetal.system.url

     - baremetal.internal.server.ip

     - baremetal.internal.http.server.ip

       This is specified with the gateway IP of an instance, which is SourceNATed by the CloudPlatform Management Server to the mangement NIC. The traffic towards the internal HTTP server goes through VR's management NIC instead of public NIC. The source NAT will only bind to guest IP of the provisioning instance to prevent network sniffer from other VMs in the same network.

4. Restart the CloudPlatform Management Server to put the new settings into effect.

8.7. Provisioning a Baremetal Instance

Deploy one Baremetal instance per host using these steps.

1. Ensure that you meet the prerequisites.

   > **Note**
   >
   > Before creating any BM instance on Basic zone, change `baremetal.provision.done.notification.enabled` to false; otherwise your instance will stay in 'Starting' state and finally get destroyed after 30 minutes.

   For Advanced zone-specific instructions, see *Section 8.1.2.2.3, "Deploying Baremetal Instances in an Advanced Zone "*.

2. Log in to the CloudPlatform UI as an administrator or user.

3. In the left navigation bar, click Instances.

4. Click Add Instance.

5. Select a zone.

6. Click Template.

7. Click Next.

8. Select the template that you created earlier, in *Section 8.4, "Create a Baremetal Template"*, and click Next.

9. Select the compute offering you created earlier, in *Section 8.5, "Create a Baremetal Compute Offering "*, and click Next.

10. (Advanced network) Select the Isolated network you created earlier, in *Section 8.1.2.2.6, "Creating an Isolated Network"*

11. Click Launch, and the instance will be created.

12. (Basic network) Set up security groups with ingress and egress rules to control inbound and outbound network traffic.

    Follow the steps in Using Security Groups in the Administrator's Guide. If you want to allow inbound network traffic to the Baremetal instances through public IPs, set up public IPs and port forwarding rules. Follow the steps in How to Set Up Port Forwarding in the Administrator's Guide.

8.8. Accessing Your Baremetal Instance

In the navigation bar of your browser, specify the IPMI address of the Baremetal host, and launch the virtual console. The Baremetal host should be PXE booted to the specified installation.

Appendix A. Important Global Configuration Parameters

The following tables display some important global configuration parameters in CloudPlatform.

Parameter	Description	Default Value
check.pod.cidrs	By default, different pods must belong to different CIDR subnets. You must set this parameter based on the available network configuration.	**`TRUE`**
direct.agent.load.size	Determines the number of direct agents to be loaded each time. Determines how many orphaned hosts must be picked up by a management server at a time. Primarily, this applies to the number of VMWare hosts that can reconnect at a time.	**`16`**
guest.domain.suffix	Default domain name for the VMs inside the virtualized networks fronted by a router.	**`cloud.internal`**
direct.agent.pool.size	Determines the default size for DirectAgentPool. Determines the number of threads that are used to process the commands sent to VMWare. This parameter requires the number of ESXi hosts and API requests sent to VMWare.	**`500`**
ha.workers	Number of the HA-Worker threads in CloudPlatform. This parameter determines the number of VM HA operations that can occur simultaneously.	**`5`**
hypervisor.list	The list of hypervisors that this deployment uses. You can change this to **`KVM`** and **`VMWare`** only. This parameter is used when the listHypervisors	**`XenServer,KVM, VMware, Hyper-V, LXC, and BareMetal`**

Parameter	Description	Default Value
	API is used (with optional parameter **`ZoneId`**) to list all the hypervisor types that are available in the current deployment.	
secstorage.ssl.cert.domain	The domain name of the certificate used for SSL communication. If you enable SSL communication for SSVM-related functionalities, such as copy template/ISO and download volume/template/ISO by setting the secstorage.encrypt.copy parameter to true then secstorage.ssl.cert.domain would signify the domain name of the certificate used for SSL communication. The value you enter here has to be exactly same as the value entered for the global parameter, consoleproxy.url.domain. After you change this value, ensure that you upload the certificates belonging to this domain.	
vmware.percluster.host.max	Maxmium hosts on a vCenter cluster. Do not let it grow over 8.	**8**
vmware.recycle.hung.wokervm	Specifies whether to recycle hung worker VMs. Worker VMs are dummy wrapper VMs that are created to perform certain volume related tasks in VMware, For example, worker VMs back-up a volume into secondary storage. After a worker VM completes the task that it has been assigned to, CS destroys it. However, when it faces some unexpected problems (such as Management Server stops while the worker VM performs a task), CloudPlatform ends up with a hung worker VM that needs to	**FALSE**

Parameter	Description	Default Value
	be cleaned up. This parameter enables administrators to recycle the hung worker VMs. You can distinguish Worker VMs from regular VMs by their names. Worker VMs carry UUID names.	

A.1. Important Healthcheck Parameters

The healthcheck parameters evaluate whether your CloudPlatform configuration conforms to a predefined standard. You can view the important healthcheck parameters that CloudPlatform uses in the following table:

Parameter	Description	Default Value
cluster.cpu.allocated. capacity.notificationthreshold	The CPU utilization threshold that is represented as a percentage value between 0 and 1. If the CPU utilization exceeds this value, CloudPlatform will send alerts indicating the low CPU availability.	**`0.75`**
cluster.localStorage. capacity.notificationthreshold	The local storage utilization threshold that is represented as a percentage value between 0 and 1. If the local storage utilization exceeds this value, CloudPlatform will send alerts indicating the low availability of local storage.	0.75
cluster.memory.allocated. capacity.notificationthreshold	The memory utilization threshold that is represented as a percentage value between 0 and 1. If the memory utilization exceeds this value, CloudPlatform will send alerts indicating the low availability of memory.	**`0.75`**
cluster.message.timeout.seconds	Time (in seconds) to wait before an inter-management server message post times out. The messages between management servers in a cluster must not experience time out.	300
cluster.storage.allocated. capacity.notificationthreshold	The allocated storage utilization threshold that is represented as a percentage value between 0	**`0.75`**

Parameter	Description	Default Value
	and 1. If the allocated storage utilization exceeds this value, CloudPlatform will send alerts indicating the low availability of allocated storage.	
cluster.storage.capacity.notificationthreshold	The storage utilization threshold that is represented as a percentage value between 0 and 1. If the storage utilization exceeds this value, CloudPlatform will send alerts indicating the low availability of storage.	**0.75**
consoleproxy.disable.rpfilter	By default, the **rp_filter** is disabled on the public interface of console proxy VM.	TRUE
consoleproxy.loadscan.interval	The time interval(in milliseconds) to scan console proxy working-load information.	10000
consoleproxy.session.max	The maximum number of viewer sessions that console proxy is configured to serve.	50
custom.diskoffering.size.max	Maximum size represented in Gigabyte (GB) for the custom disk offering.	**1024**
event.purge.delay	The number of days for which the entries are retained in the Events table. CloudPlatform deletes the entries that are older than the number of days specified for this parameter from the Events table. If you want to retain all the entries in the Events table, set this value to 0.	**15**
extract.url.cleanup.interval	The interval (in seconds) for which CloudPlatform waits before it cleans up the extracted URLs. Do not modify the default value for this parameter unless you want to disable the clean up or you want to trigger the clean up more frequently.	**120**
host.retry	Number of times the host tries to create the volume to attach to the VM.	2

Parameter	Description	Default Value
host.stats.interval	The time interval (in milliseconds) when the host statistics are retrieved from agents. Host statistics are collected for all hypervisors (KVM/Xen/ VmWare/HyperV).	60000
job.cancel.threshold.minutes	Time threshold (in minutes) for which CloudPlatform waits before it forces the cancellation of asynchronized jobs that have been in process for a long time. An asynchronized job that has been in process for a long time can block other jobs that must run. This parameter does not control the timeout of any orchestration commands.	**`60`**
linkLocalIp.nums	The number of link local IP addresses that the Domain Routers (domR) needed (in power of 2). If the number of Domain Routers (XenServer, KVM)exceeds the value set by this parameter, system VM deployment fails due to insufficient link local IP addresses.	**`10`**
migratewait	Time (in seconds) for which CloudPlatform waits to complete the VM migration. This parameter applies to VM migration only.	**`3600`**
network.disable.rpfilter	By default, the rp_filter is disabled on the [public interface of the Domain Router VM. If you enable rp_filter, the kernel does source validation by confirming the reverse path.	TRUE
network.gc.interval	The time interval (in seconds) between identifying an idle VR (that does not contain any User	600

Parameter	Description	Default Value
	VMs in the 'Running' state) and deciding to shut it down.	
network.gc.wait	Time (in seconds) to wait before actually shut down a idle VR (which does not contain any User VMs in the 'Running' state) that is not in use.	600
network.guest.cidr.limit	Size limit for guest CIDR. This cannot be less than this value..	22
pod.privateip.capacity.notificationthreshold	The private IP address space utilization threshold that is represented as a percentage value between 0 and 1. If the private IP address space utilization exceeds this value, CloudPlatform will send alerts indicating the low availability of private IP address space.	0.75
pool.storage.allocated.capacity.disablethreshold	The allocated storage pool utilization threshold that is represented as a percentage value between 0 and 1. If the allocated storage pool utilization exceeds this value, the allocators will disable using the pool indicating the low availability of allocated storage.	0.85
pool.storage.capacity.disablethreshold	The storage utilization capacity threshold that is represented as a percentage value between 0 and 1. If the storage utilization exceeds this value, the allocators will disable the storage indicating the low availability of storage.	0.85
restart.retry.interval	Time (in seconds) between retries to restart a VM. This flags controls the interval between HA restart attempts.	600
router.ram.size	Default RAM size of router VM (in MB). This is a hidden parameter. CloudPlatform persists the encrypted values of hidden parameters in the database.	8QhDMmJeVzoP eyAdDkQRQw==
secondary.storage.vm	If rue, deploys a VM per zone to manage secondary storage. Otherwise, secondary storage	hbPAm6DyrYJtYkM ed0I0ng==

Parameter	Description	Default Value
	is mounted on management server. This flag should have the value "TRUE/FALSE". If it is TRUE, CloudPlatform will start SSVM per zone to manage secondary storage. Otherwise, all secondary storage-related operations will be done on management server.	
secstorage.cmd.execution.time.max	The maximum time (in minutes) for executing a command. This flag is only used by VMware to clean up some commands from the **`cmd_exec_log`** table that are not executed properly. VMware secondary storage manager will use entries in this table to determine whether the SSVMs be expanded to handle current load.	30
secstorage.vm.mtu.size	MTU size (in Bytes) of the storage network in the secondary storage VMs. This value will be passed as boot arguments when we start and boot up SSVM.	1500
snapshot.delta.max	The maximum number of delta snapshots that are allowed in the time interval between two full snapshots. Delta snapshot is the difference between the parent snapshot and the current volume.	16
start.retry	Number of retries in deploying a VM or starting a VM when the placement of VM to a physical host fails due to runtime error. No time interval is required between the retries.	10
stop.retry.interval	Time in seconds between two consecutive retries to stop or destroy a VM.	600

Parameter	Description	Default Value
	This flag does not control the time that is waiting for graceful guest OS shutdown.	
storage.cleanup.enabled	Enables/disables the storage cleanup thread. if this flag is TRUE, a background thread will start on intervals specified by **`storage.cleanup.interval`**. This background thread will cleanup unused templates and destroyed volumes from primary storage, remove all templates, volumes, and snapshots that are marked as destroryed in **`template_store_ref`**, **`volume_store_ref`**, and **`snapshot_store_ref`** from secondary storage, and remove snapshots marked as Error from DB tables.	TRUE
storage.pool.max.waitseconds	Time in seconds to synchronize storage pool operations. This is the timeout value used in copying template from secondary to primary storage.	3600
sync.interval	Cluster Delta sync interval in seconds This is to keep the VM states in sync between the host and management server. Only applicable in case of Xenserver.	60
`system.vm.default.hypervisor`	Hypervisor type used to create system vm. If this parameter is null, CloudPlatform considers one of the supported hypervisors in that zone.	**`NULL`**
task.cleanup.retry.interval	Time (in seconds) to wait before retrying the cleaning up of tasks if the clean up operaton had failed previously. '0' means no retry. This is a GC (Garbage Collecting) process. The	600

Parameter	Description	Default Value
	interval is used to control the frequecy of running the clean up process.	
traffic.sentinel.include.zones	Traffic that goes into the specified list of zones is metered. For metering all traffic, leave this parameter empty. Used for metering data in the shared network. Zones listed here are traffic sentinel zones. All traffic going from guest network to listed zones in Traffic Sentinel will be metered.	EXTERNAL
update.wait	Time to wait (in seconds) before alerting an updating agent	600
vm.op.cancel.interval	Time (in seconds) to wait before cancelling an operation. This parameter is generic to all VM operation commands. When a VM operation command has been submitted and for whatever reason it could not finish, it won't be picked up again after the canceling interval has passed. This parameter is not related to VMware.(--)	3600
vm.op.cleanup.wait	Time (in seconds) to wait before cleanuping up any vm work items. It is the wait time to cleanup tasks that are already completed (success or failure).	3600
vm.op.cleanup.interval	Interval (in seconds) to run the thread that cleans up the VM operations.	86400
vm.op.wait.interval	Time (in seconds) to wait before checking if a previous operation has successfully performed.	120
router.check.poolsize	Number of threads to check the status of redundant router.	10
vm.op.lock.state.retry	Times to retry locking the state of a VM for operations	5

Parameter	Description	Default Value
	This parameter specifies the number of retries when trying to lock-down the VM for state transition to happen.	
vm.tranisition.wait.interval.	Time (in seconds) to wait before taking over a VM in transition state	3600
wait	Time in seconds to wait for control commands to return This flag gives the default timeout value for commands sent from management server to host agent s. If it is changed, it affects commands that are using the default timeout setting, however, there are commands that are using command specific timeout settings, for those commands, it does not have effect on them.	1800
`xapiwait`	Time (in seconds) to wait for XAPI to return	**`600`**
`zone.directnetwork.publicip.capacity.notificationthreshold`	The Direct network public IP address utilization threshold that is represented as a percentage value between 0 and 1. If the Direct network public IP utilization exceeds this value, CloudPlatform will send alerts indicating the low availability Direct network public IP addresses.	**`0.75`**
`zone.secstorage.capacity.notificationthreshold`	The secondary storage utilization threshold that is represented as a percentage value between 0 and 1. If the secondary storage utilization exceeds this value, CloudPlatform will send alerts indicating the low availability of secondary storage.	**`0.75`**
`zone.virtualnetwork.publicip.capacity.notificationthreshold`	The public IP address space utilization threshold that is represented as a percentage value between 0 and 1. If the public IP address space utilization exceeds this value,	**`0.75`**

Parameter	Description	Default Value
	CloudPlatform will send alerts indicating the low availability of public IP address space.	
`zone.vlan.capacity.notificationthreshold`	The Zone VLAN utilization threshold that is represented as a percentage value between 0 and 1. If the Zone VLAN utilization exceeds this value, CloudPlatform will send alerts indicating the low number of Zone VLANs.	**`0.75`**
`network.loadbalancer.basiczone.elb.gc.interval.minutes`	Garbage collection interval to destroy unused ELB VMs in minutes. Minimum is 5 minutes.	**`30`**
`secstorage.allowed.internal.sites`	Comma separated list of CIDRs internal to the datacenter that can host template download servers, please note 0.0.0.0 is not a valid site.	**`Null`**

XenSever and KVM hypervisors do not support the following parameters:

Parameter	Description	Default Value
secstorage.session.max	The maximum number of command execution sessions that a secondary storage VM (SSVM) can handle.	**`50`**
secstorage.capacity.standby	The minimum number of command execution sessions that the system is able to serve immediately (standby capacity).	**`10`**

Index

B

Baremetal
- add cluster, 120
- add host, 120
- add zone, 109
- configure advanced networking, 114
- configure basic zone, 113
- create compute offering, 122
- create network offering, 111
- deploy in advanced zone, 116
- limitations, 109
- provision instance, 124
- set parameters, 123

C

create SSH keypair, 3

H

Hyper-V
- add cluster, 84
- add host, 84
- add pod, 83
- add primary storage, 86
- add secondary storage, 87
- add zone, 76
- configure advanced zone, 77
- create secondary storage mount point, 76
- provisioning overview, 75

K

KVM
- add cluster, 18
- add host, 19
- add NFS storage, 22
- add pod, 17
- add primary storage, 20
- add secondary storage, 21
- add zone, 8
- configure advanced zone, 12
- configure basic zone, 9
- configure S3 Object Store, 23
- create secondary storage mount point, 8
- host requirements, 19
- provisioning overview, 7
- upgrade to object storage, 26

L

LXC
- add cluster, 101
- add host, 102
- add NFS storage, 106
- add pod, 100
- add primary storage, 103
- add secondary storage, 105
- add zone, 92
- configure advanced zone, 96
- configure basic zone, 93
- create secondary storage mount point, 92
- host requirements, 102
- provisioning overview, 91
- upgrade to object storage, 106

R

reset SSH key, 5

S

SSH key, 3

V

vSphere
- add cluster, 64
- add NFS storage, 68
- add pod, 63
- add primary storage, 66
- add secondary storage, 67
- add zone, 56
- cluster size limit, 64
- configure advanced zone, 57
- configure S3 Object Store, 69
- create secondary storage mount point, 56
- provisioning overview, 55
- upgrade to object storage, 72

X

XenServer
- add cluster, 40
- add host, 44
- add NFS storage, 48
- add pod, 39
- add primary storage, 46
- add secondary storage, 47
- add zone, 30
- configure advanced zone, 34
- configure basic zone, 31
- configure S3 Object Store, 49
- create secondary storage mount point, 30
- host requirements, 40
- provisioning overview, 29
- upgrade to object storage, 52