Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5 Administration Guide

Revised January 30, 2015 06:00 pm IST

Citrix CloudPlatform

Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5 Administration 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 the ongoing operation and maintenance of a CloudPlatform-powered cloud, read this document. This document will help you start using, configuring, and managing the ongoing operation of your cloud.

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

Managing Service Offerings

A service offering is a set of virtual hardware features such as CPU core count and speed, memory, and disk size. As a CloudPlatform administrator, you can set up various service offerings. The end users can choose from the available service offerings when they create a new VM. Based on the offering that the end users select, CloudPlatform generates usage records that they can integrated with billing systems.

In this chapter we discuss compute, disk, and system service offerings. Network offerings are discussed in **Chapter 4. Managing Network Offerings**.

For conceptual information about compute, disk, and system service offerings, refer to **Chapter 6: Service Offerings** in the *Citrix CloudPlatform 4.5 (powered by Apache CloudStack) Concepts Guide*.

2.1. Creating a New Compute Offering

To create a new compute offering, you must do the following:

1. Log-in to the CloudPlatform UI using administrator privileges .
2. In the left navigation bar, click **Service Offerings**.
3. On the right-side panel, in the **Select offering** list box, select **Compute Offering**.
4. Click **Add compute offering**.
5. In the **Add compute offering** dialog, enter the following:
   - **Name**: A name to identify the service offering.
   - **Description**: A description of the offering that you want to display to users.
   - **Storage Type**: Select the type of disk that you must allocate.
     - **Shared**: allocates from the storage that is accessible via NFS.
     - **Local**: allocates from the storage that is attached directly to the host where the VM is running.
   - **Custom**: You can use custom compute offerings in the following cases: deploying a VM, changing the compute offering of a stopped VM and running VMs, which is nothing but scaling up.

     If you select **Custom**, you need to specify the following parameters during the VM deployment:

     - **# of CPU cores**: The number of cores that is to be allocated to a VM.

       CloudPlatform supports hyperthreading. It's admin's discretion whether to use it or not. If hyperthreading is enabled on the hypervisor and you want to use it, specify the corresponding value. You can view the effective number of CPUs in the Statistics tab corresponding to the host under Infrastructure page in the CloudPlatform UI. For example, 4 x 2.26 GHz, where 4 represents the number of cores.
     - **CPU (in MHz)**: The CPU speed of the cores that are allocated to VM. For example, “2000” would provide for a 2 GHz clock.
     - **Memory (in MB)**: The amount of memory in megabytes that the VM should be allocated. For example, “2048” would provide for a 2 GB RAM allocation.

- **Network Rate (Mb/s)**: Allowed rate of data transfer in MB per second (megabits per second).
- **Disk Read Rate**: Allowed rate of disk read in bits per second (byte per second).
- **Disk Write Rate**: Allowed rate of disk write in bits per second (byte per second).
- **Disk Read Rate**: Allowed rate of disk read in IOPS (Input/Output Operations Per Second).
- **Disk Write Rate**: Allowed rate of disk write in IOPS.
- **Offer HA**: Select to enable the administrator to monitor the VM as highly available as possible.

Note

Enable native HA for VMware deployments. Because CloudPlatform relies on native HA for VMware, this option is ignored.

- **Storage Tags**: The tags to be associated with the primary storage that the VM uses.
- **Host Tags**: (Optional) Any tags that you use to organize your hosts.
- **CPU cap**: Select if you want to limit the level of CPU usage even if spare capacity is available. The impact of CPU Cap, which you configure as part of creating a computer offering, on a VM differs based on the hypervisor that hosts the VM. The following explanation will help you assess the impact of CPU Cap value on the VMs running in your environment:

Hypervisor	Impact of CPU Cap on VMs
KVM	No impact.
XenServer	The CPU cap is not visible in XenCenter. It is set in the "cap" option in the "VCPUs-params" parameter for a VM. The value of "cap" is a percentage of one physical CPU core and is calculated using values from the Compute Offering and the physical CPU speed of the host using the following formula: **`( ( ( CPU (in MHz) ) * 0.99 * ( # of CPU cores ) ) / host CPU speed ) * 100`** For example: A Compute Offering with 2200 MHz and one core on a host with a quad-core 2200 MHz CPU: CPU cap = (2200*0.99*1/2200)*100 = 99
vSphere	The CPU cap is set as a limit in the VM properties and can be seen with vSphere Client.

- **Public**: Select to make the service offering available across all domains. If you have not selected this check box, the availability of the service offering will be limited to a sub domain. Then, CloudPlatform prompts for the sub domain's name.
- **Volatile**: Select to enable the VMs that are created from this service offering to reset the root disks 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**: Select the technique that you want CloudPlatform to use when deploying VMs based on this service offering.

  FirstFitPlanner places new VMs on the first host that has sufficient capacity to support the VM's requirements.

  UserDispersingPlanner makes the best effort to evenly distribute VMs that belong to the same account on different clusters or pods.

  UserConcentratedPodPlanner prefers to deploy VMs that belong to the same account within a single pod.

  ImplicitDedicationPlanner deploys VMs on private infrastructure that is dedicated to a specific domain or account. After you select this planner, you must select a value in the Planner mode field..

  BareMetalPlanner is used with Bare Metal hosts. See the section on configuring Bare Metal in the *CloudPlatform Hypervisor Configuration Guide*
- **Planner mode**: Select a value after you select ImplicitDedicationPlanner in the previous field. The planner mode determines how VMs are deployed on private infrastructure that is dedicated to a single domain or account.

  Strict: Select this option if you do not want the host to be shared across multiple accounts. For example, strict implicit dedication is useful for deployment of certain types of applications such as desktops, where no host can be shared between different accounts without violating the desktop software's terms of license. Used when ImplicitDedicationPlanner is selected in the previous field. The planner mode determines how VMs will be deployed on private infrastructure that is dedicated to a single domain or account.

  Preferred: The VM will be deployed in dedicated infrastructure if possible. Otherwise, the VM can be deployed in shared infrastructure.

6. Click **Add**.

2.2. Creating a New Disk Offering

To create a new disk offering, you must do the following:

1. Log-in to the CloudPlatform UI using administrator privileges .
2. In the left navigation bar, click **Service Offerings**.
3. On the right-side panel, in the **Select offering** list box, select **Disk Offerings**.
4. Click **Add Disk Offering**.
5. In the **Add Disk Offering** dialog, enter the following:

- **Name**: A name to identify the disk offering.
- **Description**: A description of the offering that you want to display to users.
- **Storage Type**: Select the type of disk that you want to allocate to the VM. **Local** is attached to the hypervisor host where the VM is running. **Shared** is the storage that is accessible via the secondary storage provider.
- **Custom Disk Size**: Select this check box to set your own disk size when creating a VM based on this disk offering. If not selected, the root administrator must define a value in the **Disk Size** field.
- **Disk Size**: This field appears only if you have not selected the **Custom Disk Size** check box. You can define the volume size in GB.
- **QoS Type**: You can select one of the following:
  - Keep this field empty if you do not want to specify any Quality of Service.
  - Select **hypervisor** if you want to go with the rate limiting enforced on the hypervisor side.
  - Select **storage** if you want to go with the guaranteed minimum and maximum IOPS enforced on the storage side.

  If using QoS, make sure that the hypervisor or storage system supports this feature.
- (Optional) **Storage Tags**: The tags that you can associate with the primary storage for this disk. You can use a comma separated list of attributes of the storage as tags. For example, "ssd,blue".

  Tags are also added on Primary Storage. CloudPlatform matches tags on a disk offering to tags on the storage. A tag that is available on a disk offering must be available on Primary Storage as well for the volume to be provisioned. If no such primary storage exists, allocation from the disk offering will fail.
- **Public**: Select to make the service offering available across all domains. If you have not selected this check box, the availability of the service offering will be limited to a sub domain. Then, CloudPlatform prompts for the sub domain's name.

6. Click **OK**.

2.3. Modifying or Deleting a Service Offering

You cannot modify a Service Offering. This applies to both the compute and the disk offerings.

However, you can remove a service offering. If it is no longer in use, you can delete it immediately and permanently. If the service offering is still in use, it will remain in the database until all the virtual machines referencing to it have been deleted. After you delete a service offering, it will not be available to end users who create new instances.

2.4. Creating a New System Service Offering

To create a system service offering, do the following:

1. Log-in to the CloudPlatform UI using administrator privileges.
2. In the left navigation bar, click **Service Offerings**.

3. On the right-side panel, in the **Select offering** list box, select **System Offerings**.
4. Click **Add System Service Offering**.
5. In the **Add System Service Offering** dialog, enter the following:
   - **Name**: A name to identify the system offering.
   - **Description**: A description of the offering that you want to display to users.
   - **System VM Type**: Select the type of system virtual machine that this offering is intended to support.
   - **Storage Type**: Select the type of disk that you must allocate.

     Shared: allocates from the storage that is accessible via NFS.

     Local: allocates from the storage that is attached directly to the host where the system VM is running.
   - **# of CPU cores**: The number of cores that is to be allocated to a system VM with the offering.
   - **CPU (in MHz)**: The CPU speed of the cores that are allocated to system VM. For example, “2000” would provide for a 2 GHz clock.
   - **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.
   - **Network Rate (Mb/s)**: Allowed rate of data transfer in MB per second.
   - **Disk Read Rate (BPS)**: Allowed rate of disk read in bits per second.
   - **Disk Write Rate (BPS)**: Allowed rate of disk write in bits per second.
   - **Disk Read Rate (IOPS)**: Allowed rate of disk read in IOPS (Input/Output Operations Per Second).
   - **Disk Write Rate (IOPS)**: Allowed rate of disk write in IOPS.
   - **Offer HA**: Select to enable the administrator to monitor the system VM as highly available as possible.
   - **Storage Tags**: The tags to be associated with the primary storage that the system VM uses.
   - **Host Tags**: (Optional) Any tags that you can use to organize your hosts.
   - **CPU cap**: Select if you want to limit the level of CPU usage even if spare capacity is available. The impact of CPU Cap, which you configure as part of creating a computer offering, on a VM differs based on the hypervisor that hosts the VM. For more information, refer to *Section 2.1, “Creating a New Compute Offering ”*
   - **Public**: Select to make the service offering available across all domains. If you have not selected this check box, the availability of the service offering will be limited to a sub domain. Then, CloudPlatform prompts for the sub domain's name.
6. Click **OK**.

2.5. Changing the Secondary Storage VM Service Offering on a Guest Network

To change the SSVM service offering that is associated with an existing guest network, do the following:

1. Log-in to the CloudPlatform UI as an administrator or as a user.
2. Before you change the SSVM service offering, you must first stop all the SSVMs on the network. For more information, see *Section 6.5, "Stopping and Starting Virtual Machines"*
3. In the left navigation bar, click **Instances**.
4. In the right-side panel, select the VM that you want to work with.
5. In the **Name** column, click the name of the VM.
6. Under the **Details** tab, do the following:
    a. Click the **Stop Instance** button to stop the VM. In the dialog box to confirm the action, click **OK**.
    b. Click the **Change service offering** button
    c. In the **Change service offerings** dialog box, in the **Compute offering** list box, select the service
    d. Click **OK**.
    e. Click **Start Instance** to start the VM.

Deployment Planners in CloudPlatform

In CloudPlatform, each VM must be deployed on a suitable deployment destination. A deployment destination is a set of recommended resources that you can choose for deploying a VM. The deployment destination provides the datacenter, pod, cluster, host, and the storage pool that you can use for deploying the VM.

A deployment planner, which is part of the CloudPlatform adapters, provides the suitable deployment destination that is required for a VM. The deployment planner uses two internal allocators - HostAllocators and StoragePoolAllocators - to figure out suitable hosts and storage pools for deploying the VM. The HostAllocator lets you create custom rules for determining the physical host for the guest virtual machine. The StoragePoolAllocator lets you create custom rules for determining the storage pool for the guest virtual machines. Deployment Planner applies custom heuristics to order the list of clusters. After the cluster list is formed, a planner uses host and storage allocators for each cluster until suitable resources are returned by both these allocators. Each deployment planner uses its own algorithm to order the list of clusters.

Note

Deployment planners are not specific to any hypervisors. However, CloudPlatform uses a separate deployment planner for Baremetal.

The following table displays the deployment planners that are available with CloudPlatform:

Deployment Planner	Description
`FirstFit`	Ensures that clusters are ordered by available capacity and first host/pool that has enough capacity is selected within a cluster.
`Random`	Shuffles the list of clusters/hosts/pools that the DB lookup returns. **`Random`** does not follow the round-robin method.
`UserDispersing`	Ensures that VM's are dispersed for a given account. It first selects the clusters/hosts with minimum number of running VMs for that account. Similarly, storage Pool with minimum number of Ready volumes for the account is selected first.
`Userconcentratedpod`	Always selects the pod/cluster with maximum number of VMs for the account.

As an administrator, you can use the global configuration parameter *`vm.deployment.planner`* to specify the default deployment planner that CloudPlatform can use. The default value for the *`vm.deployment.planner`* parameter is **`FirstFitPlanner`** deployment planner.

A service offering in CloudPlatform can specify a deployment planner. This deployment planner is used while deploying VMs using that service offering.

3.1. Allocators in CloudPlatform

Allocators are part of the CloudPlatform adapters. An Allocator typically posses the following two base interfaces:

HostAllocator

A HostAllocator returns suitable hosts in the cluster where you can deploy the given VM. To decide the hosts, a HostAllocator considers factors such as the host tags on the service offering, CPU and RAM capacity of the host, current state of the host, and any specific capability of the host such as HVM, GPU, and so on.

An allocator uses the following algorithms to select a host in the cluster that contains a set of hosts that meet all the conditions:

- Random - Selects a host in the cluster randomly.
- FirstFit- Selects the first available host.
- UserDispersing - Selects the host with least number of VMs for the given account.

You can use the global configuration parameter: **`vm.allocation.algorithm`** to specify the algorith that the Allocator must use.

StoragePoolAllocator

A StoragePoolAllocator returns suitable storage pools available in the cluster where the volumes of the given VM can be created. To decide the storage pools, a StoragePoolAllocator considers factors such as storage tags on the disk offering, storage capacity of the pool, availability scope of the pool: local/ cluster/zone. (There is a separate allocator for each scope.), and any specific capability such as IOPS.

Adding Regions to Your Cloud Infrastructure (optional)

Grouping your cloud resources into geographic regions is an optional step when provisioning the cloud.

For conceptual information about cloud infrastructure, refer to **Chapter 4: Cloud Infrastructure Concepts** in the *Citrix CloudPlatform 4.5 (powered by Apache CloudStack) Concepts Guide*.

4.1. The First Region: The Default Region

If you do not specify a region, all the zones in your cloud will be automatically grouped into a single, default region. The region ID assigned to this default region is 1. If required, you can modify the name or URL of the default region from the CloudPlatform UI.

4.2. Adding a Region

Each region has its own CloudPlatform instance. Therefore, the first step of creating a new region is to install the Management Server software in the geographic area where you want to set up the new region. Then, you must assign a region ID to the second region that you want to add in addition to the default region. To assign the region ID, you can use the additional command-line flag **`-r <region_id>`** when you set up the database. The default region is automatically assigned a region ID of 1, so your first additional region might be region 2.

```
cloudstack-setup-databases cloud:<dbpassword>@localhost --deploy-as=root:<password> -e
 <encryption_type> -m <management_server_key> -k <database_key> -r <region_id>
```

Ensure that installation has been completed and the Management Server has been started.

For more information on installing CloudPlatform, refer to *CloudPlatform (powered by Apache CloudStack) Version 4.5 Installation Guide*.

1. After you add the region ID, you must add the new region to region 1 in CloudPlatform. Do the following:

    a. Log-in to CloudPlatform in the first region using the root administrator privileges (that is, log-in to <region.1.IP.address>:8080/client).

    b. In the left navigation bar, click **Regions**.

    c. On the right-side panel, click **Add Region**. In the **Add Region** dialog, enter the following:

    - **ID**: A number that uniquely identifies the region. Use the same number you set in the database during Management Server installation in the new region; for example, 2.
    - **Name**: A descriptive name that you can use to identify the new region.
    - **Endpoint**: The URL where you can log-in to Management Server in the new region. This URL will follow the format: <region.2.IP.address>:8080/client.

    d. Click **OK**.

2. Now, perform similar steps to log-in to region 2, and add region 1.

3. Copy the account, user, and domain tables from the region 1 database to the region 2 database.

In the following commands, it is assumed that you have set the root password on the database, which is a CloudPlatform recommended best practice. Substitute your own MySQL root password.

a. First, run the following command to copy the contents of the database:

```
# mysqldump -u root -p<mysql_password> -h <region1_db_host> cloud account user domain
 > region1.sql
```

b. Then, run the following command to put the data onto the region 2 database:

```
# mysql -u root -p<mysql_password> -h <region2_db_host> cloud < region1.sql
```

4. Remove project accounts. Run the following command on the region 2 database:

```
mysql> delete from account where type = 5;
```

5. Set the default zone as null:

```
mysql> update account set default_zone_id = null;
```

6. Restart Management Servers in region 2.

4.3. Adding Third and Subsequent Regions

To add the third region, and subsequent regions, the steps are similar to those for adding the second region. However, you must repeat certain steps multiple times for each additional region:

1. Install CloudPlatform in each additional region. Set the region ID for each region during the database setup step.

```
cloudstack-setup-databases cloud:<dbpassword>@localhost --deploy-as=root:<password> -e
 <encryption_type> -m <management_server_key> -k <database_key> -r <region_id>
```

2. After the Management Server is running, add your new region to all existing regions by using the **Add Region** button in the UI. For example, if you were adding region 3:

   a. Log-in to CloudPlatform in the first region (that is, log-in to <region.1.IP.address>:8080/client) using the root administrator privileges. Then, add a region with ID 3, which represents region 3, and the endpoint <region.3.IP.address>:8080/client.

   b. Log-in to CloudPlatform in the second region (that is, log in to <region.2.IP.address>:8080/client) using the root administrator privileges, and add a region with ID 3, which represents region 3, and the endpoint <region.3.IP.address>:8080/client.

3. Repeat the procedure in reverse to add all existing regions to the new region. For example, for the third region, add the other two existing regions:

   a. Log-in to CloudPlatform in the third region (that is, log in to <region.3.IP.address>:8080/client) using root administrator privileges.

   b. Add a region with ID 1, which represents region 1, and the endpoint <region.1.IP.address>:8080/client.

c. Add a region with ID 2, which represents region 2, and the endpoint <region.2.IP.address>:8080/client.

4. Copy the account, user, and domain tables from any existing region's database to the new region's database.

   In the following commands, it is assumed that you have set the root password on the database, which is a CloudPlatform recommended best practice. Substitute your own MySQL root password.

   a. First, run the following command to copy the contents of the database:

   ```
   # mysqldump -u root -p<mysql_password> -h <region1_db_host> cloud account user domain
    > region1.sql
   ```

   b. Then run the following command to put the data onto the new region's database. For example, for region 3:

   ```
   # mysql -u root -p<mysql_password> -h <region3_db_host> cloud < region1.sql
   ```

5. Remove project accounts. Run the following command on the region 3 database:

   ```
   mysql> delete from account where type = 5;
   ```

6. Set the default zone as null:

   ```
   mysql> update account set default_zone_id = null;
   ```

7. Restart Management Servers in the new region.

4.4. Deleting a Region

You must remove the region from all the other regions in your environment where it is added. You can log-in to each of the regions, navigate to the one you want to delete, and click **Remove Region**. For example, to remove the third region in a 3-region cloud, do the following:

1. Log-in to <region.1.IP.address>:8080/client.
2. In the left navigation bar, click **Regions**.
3. On the right-side panel, click the name of the region that you want to delete.
4. Click **Remove Region** .
5. Repeat these steps for <region.2.IP.address>:8080/client.

Managing Network Offerings

This chapter describes how to manage network offerings in CloudPlatform.

A network offering is a named set of network services, such as DHCP, DNS, Firewall, VPN, etc. As the CloudPlatform administrator, you can create any number of custom network offerings in addition to the default network offerings that CloudPlatform provides. When the end users create a new VM, they can choose one of the available network offerings that determines the network services the VM can use.

For more conceptual information about network offerings, refer to **Chapter 8: Networking for Users** in the *Citrix CloudPlatform 4.5 (powered by Apache CloudStack) Concepts Guide*.

5.1. Creating a New Network Offering

To create a network offering, you must do the following:

1. Log-in to the CloudPlatform UI using administrator privileges.
2. In the left navigation bar, click **Service Offerings**.
3. On the right-side panel, in the **Select offering** list box, select **Network Offerings**.
4. Click **Add network offering**.
5. In the **Add network offering** dialog, enter the following:
   - **Name**: A name to identify the network offering.
   - **Description**: A description of the offering that you want to display to users.
   - **Network Rate (Mb/s)**: Allowed rate of data transfer in MB per second (megabits per second).
   - **Disk Read Rate**: Allowed rate of disk read in bits per second (byte per second).
   - **Disk Write Rate**: Allowed rate of disk write in bits per second (byte per second).
   - **Guest Type**: Select the Guest Network type. You can select from the options: **Isolated** and **Shared**.

     For a description of this term, see the Concepts Guide.
   - **Persistent**: Select to indicate that the guest network is persistent. The network that you can provision without having to deploy a VM on it is termed persistent network.

     For more information, see *Section 9.25, “Persistent Networks”*.
   - **Specify VLAN**: (applies to Isolated guest networks only) Select to indicate that a VLAN can be specified when this offering is used. If you select this option and later use this network offering while creating a VPC tier or an isolated network, you will be able to specify a VLAN ID for the network you create.
   - **VPC**: Select to indicate that the guest network is Virtual Private Cloud-enabled. A Virtual Private Cloud (VPC) is a private, isolated part of CloudPlatform. A VPC can have its own virtual network topology that resembles a traditional physical network.

     For more information on VPCs, see *Section 9.24.1, “About Virtual Private Clouds”*.

- **Supported Services**: Select one or more possible network services. For some services, you must also choose the service provider. For example, if you select **Load Balancer**, you can choose the CloudPlatform virtual router or any other load balancers that have been configured in the cloud. Depending on the services that you choose, additional fields may appear in the rest of the dialog box.

  Based on the guest network type selected, you can see the following supported services:

Supported Services	Description	Isolated	Shared
VPN	Supported	Supported	
DHCP	For more information, see *Section 9.20, "DNS and DHCP"*.	Supported	
DNS	Supported	Supported	
Firewall	Supported	Supported	
Load Balancer	If you select Load Balancer, you can choose the CloudPlatform virtual router or any other load balancers that have been configured in the cloud.  **Note** VR supports load balancing only on TCP protocol. HAProxy used by VR does not balance the load on UDP protocol.	Supported	Supported
User Data	Not Supported	Supported	
Source NAT	If you select Source NAT, you can choose the CloudPlatform virtual router or any other Source NAT providers that have been configured in the cloud.	Supported	Supported
Static NAT	If you select Static NAT, you can choose	Supported	Supported

Supported Services	Description	Isolated	Shared
	the CloudPlatform virtual router or any other Static NAT providers that have been configured in the cloud.		
Port Forwarding	If you select Port Forwarding, you can choose the CloudPlatform virtual router or any other Port Forwarding providers that have been configured in the cloud.	Supported	Supported
Security Groups	Not Supported	Supported	
Network ACL	Supported	Not Supported	
Virtual Networking	Supported	Not Supported	
BaremetalPxeService	Not Supported	Supported	

- **System Offering**: This field appears when the service provider for any of the services selected in the **Supported Services** field is a virtual router. Select the system service offering that you want virtual routers to use in this network. For example, if you select **Load Balancer** in the **Supported Services** field and then select a virtual router to provide load balancing, the **System Offering** field appears. This enables you to choose between the CloudPlatform default system service offering and any custom system service offerings that have been defined by the CloudPlatform root administrator.
- **LB Isolation**: Specify the type of load balancer isolation you want for the network. The options are **Shared** or **Dedicated**.

  Dedicated: If you select dedicated LB isolation, a dedicated load balancer device is assigned for the network from the pool of dedicated load balancer devices that are provisioned in the zone. If no sufficient dedicated load balancer devices are available in the zone, network creation fails. Dedicated device is a good choice for the high-traffic networks that make full use of the device's resources.

  Shared: If you select shared LB isolation, a shared load balancer device is assigned for the network from the pool of shared load balancer devices that are provisioned in the zone. While provisioning, CloudPlatform picks the shared load balancer device that is used by the least number of accounts. Once the device reaches its maximum capacity, the device will not be allocated to a new account.
- **Mode**: You can select either Inline mode or Side by Side mode:

  Inline mode: Supported only for Juniper SRX firewall and BigF5 load balancer devices. In inline mode, a firewall device is placed in front of a load balancing device. The firewall acts as the gateway for all the incoming traffic, and then redirects the load balancing traffic to the load balancer behind it. The load balancer in this case will not have direct access to the public network.

Side by Side: In the Side by Side mode, a firewall device is deployed in parallel with the load balancer device so that the traffic to the load balancer public IP is not routed through the firewall, and therefore, is exposed to the public network.

- **Associate Public IP**: Select this option if you want to assign a public IP address to the VMs deployed in the guest network. To avail this option, you must do the following:
  - Share Guest network.
  - Enable StaticNAT.
  - Enable Elastic IP.
- **Redundant router capability**: Available only when you select Virtual Router as the Source NAT provider. Select this option if you want to use two virtual routers in the network for uninterrupted connection: one operating as the master virtual router and the other as the backup. The master virtual router receives requests from the user's VM and sends responses to the user's VM. The backup virtual router is activated only when the master is down. After the failover, the backup becomes the master virtual router. CloudPlatform deploys the routers on different hosts to ensure reliability if one host is down.
- **Conserve mode**: When the conserve mode is on, you define more than one service on the same public IP. A network offering with conserve mode enabled implies single public IP can be used for multiple services (PF/LB/StaticNat) at the same time. When conserve mode is off, you can use the public IP only for a single service. For example, a public IP used for a port forwarding rule cannot be used for defining other services, such as StaticNAT or load balancing.

  

  Note

  If StaticNAT is enabled, you cannot create port forwarding or load balancing rules for the IP irrespective of the status of the conserve mode. However, you can add the firewall rules by using the **`createFirewallRule`** command.

- **Tags**: Enter the network tag specifying the physical network to use.
- **Default egress policy**: Configure the default policy for firewall egress rule. Options are **Allow** and **Deny**. Default is **Allow** if no egress policy is specified. This policy indicates that the guest network that you create from this offering accepts all the egress traffic.

  To block the egress traffic for a guest network, select **Deny**. In this case, when you configure an egress rule for an isolated guest network, rules are added to allow the specified traffic.

6. Click **Add**.

5.2. Changing the Network Offering on a Guest Network

To change the network offering that is associated with an existing guest network, do the following:

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

Note

If you are changing from a network offering that uses the CloudPlatform virtual router to one that uses external devices as network service providers, you must first stop all the VMs on the network.

See *Section 6.5, "Stopping and Starting Virtual Machines"*.

2. In the left navigation bar, click **Network**.
3. On the right-side panel, in the **Name** column, click the name of the network that you want to modify.
4. Under **Details**, click the **Edit** icon.
5. In Network Offering, choose the new network offering, then click **Apply**.

   A message box that asks whether you want to keep the existing CIDR appears. The CIDR will be affected when you change the network offering.

   If you upgrade between virtual router as a provider and an external network device as provider, acknowledge the change of CIDR to continue. Click **Yes**.
6. Wait for the update to complete. Don't restart the VMs until the network change is completed.
7. After the network change is completed, restart the VMs that you have stopped.

5.3. Creating and Changing a Virtual Router Network Offering

To create the network offering in association with a virtual router system service offering, do the following:

1. Log-in to the CloudPlatform UI using an administrator account or a user account.
2. Create a system service offering, for example: VRsystemofferingHA.

   For more information on creating a system service offering, see *Section 2.4, "Creating a New System Service Offering"*.
3. On the right-side panel, in the **Select offering** list box, select **Network Offerings**.
4. Click **Add network offering**.
5. In the Add network offering dialog box, enter the details and click **OK** to create the network offering. For more information, see *Section 5.1, "Creating a New Network Offering"*

To change the network offering of a guest network to the virtual router service offering, do the following:

1. In the CloudPlatform UI, on the left navigation bar, select **Network**.

2. On the right-side panel, in the **Select view** list box, select **Guest networks**.
3. In the table that lists the guest networks, in the **Name** column, click the name of the guest network that you want to offer.
4. Click the Edit button.
5. From the Network Offering drop-down, select the virtual router network offering you have just created.
6. Click OK.

Working With Virtual Machines

As a CloudPlatform administrator, you can manage the life cycle of all the guest VMs that are executing in the cloud. This chapter describes how you can work with virtual machines in CloudPlatform.

For the conceptual information about working with virtual machines in CloudPlatform and virtual machine life cycle, refer to **Chapter 9: About Virtual Machines in CloudPlatform** in the *Citrix CloudPlatform 4.5 (powered by Apache CloudStack) Concepts Guide*.

6.1. Best Practices for Virtual Machines

For VMs to work as expected and provide excellent service, follow these guidelines.

6.1.1. Monitoring VMs for Max Capacity

The CloudPlatform administrator should monitor the total number of VM instances in each cluster, and disable allocation to the cluster if the total is approaching the maximum that the hypervisor can handle. Be sure to leave a safety margin to allow for the possibility of one or more hosts failing, which would increase the VM load on the other hosts as the VMs are automatically redeployed. Consult the documentation for your chosen hypervisor to find the maximum permitted number of VMs per host, then use CloudPlatform global configuration settings to set this as the default limit. Monitor the VM activity in each cluster at all times. Keep the total number of VMs below a safe level that allows for the occasional host failure. For example, if there are N hosts in the cluster, and you want to allow for one host in the cluster to be down at any given time, the total number of VM instances you can permit in the cluster is at most (N-1) * (per-host-limit). Once a cluster reaches this number of VMs, use the CloudPlatform UI to disable allocation of more VMs to the cluster.

6.1.2. Installing Required Tools and Drivers

Ensure that the following are installed on each VM:

- For XenServer, install PV drivers and Xen tools on each VM. This will enable live migration and clean guest shutdown. Xen tools are required for dynamic CPU and RAM scaling to work.
- For vSphere, install VMware Tools on each VM. This will enable console view to work properly. VMware Tools are required for dynamic CPU and RAM scaling to work.
- For Hyper-V, see *About Virtual Machines and Guest Operating Systems*[1].

  For Linux VMs on Hyper-V, see *Linux Virtual Machines on Hyper-V*[2].

To be sure that XenServer tools or VMware Tools is installed, use one of the following techniques:

- Create each VM from a template that already has the tools installed; or,
- When registering a new template, the administrator or user can indicate whether tools are installed on the template. This can be done through the UI or using the updateTemplate API; or,
- If a user deploys a virtual machine with a template that does not have Xen tools or VMware Tools, and later installs the tools on the VM, then the user can inform CloudPlatform using the

[1] http://technet.microsoft.com/en-us/library/cc794868(WS.10).aspx
[2] http://technet.microsoft.com/library/dn531030.aspx

updateVirtualMachine API. After installing the tools and updating the virtual machine, stop and start the VM.

6.1.3. VM Sync Consideration

CloudPlatform synchronizes out of the band changes for VM state and VM to host mapping. This is consistent across hypervisors. However, limitations exist in synchronizing states on hypervisors, such as on KVM. For example, paused state of a VM on KVM is not synchronized to CloudPlatform.

CloudPlatform also supports out of the band operations, such as DRS and native HA. However, dependency operations, such as starting corresponding router VM or implementing a network, for a VM are not automatically synchronized if performed outside CloudPlatform. Perform these operations by using CloudPlatform.

6.2. Creating Virtual Machines

Virtual machines (VMs) are usually created from a template. Users can also create blank virtual machines. A blank virtual machine is a virtual machine without an OS template. Users can attach an ISO file and install the OS from the CD/DVD-ROM.

For more information on the operating systems that the hypervisors support, refer to the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Hypervisor Configuration Guide*.

Note

You can create a VM without starting it. You can determine whether the VM needs to be started as part of the VM deployment. A new request parameter, **`startVM`** that is introduced in the **`deployVm`** API supports this feature. For more information, see the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Developer's Guide*e

6.2.1. Creating a Virtual Machine from a Template

1. Log-in to the CloudPlatform UI using an administrator account or a user account.
2. In the left navigation bar, click **Instances**.
3. On the right-side panel, click **Add Instance**.
4. In the Add Instance wizard panel, Select a zone.
5. Select a template, then follow the steps in the wizard.

   For more information about the templates, see *Chapter 10, Working with Templates*.
6. Be sure that the hardware you have allows starting the selected service offering.
7. Click Submit and your VM will be created and started.

Note

For security reasons, the internal name of the VM is visible only to the root admin.

6.2.2. Creating a Virtual Machine from an ISO

1. Log in to the CloudPlatform UI as an administrator or user.
2. In the left navigation bar, click Instances.
3. Click Add Instance.
4. Select a zone.
5. Select ISO Boot, and follow the steps in the wizard.
6. Click Submit and your VM will be created and started.

6.2.3. Configuring Usage of Linked Clones on VMware

(For ESX hypervisor in conjunction with vCenter)

VMs can be created as either linked clones or full clones on VMware.

A full clone is a copy of an existing virtual machine that does not depend on the original virtual machine after you create it. A linked clone is also a copy of an existing virtual machine, but it has ongoing dependency on the original. A linked clone shares the virtual disk of the original VM, and retains access to all files that were present at the time when you created the clone.

For more information on clone types, refer to VMware documentation.

The use of these clone types involves some side effects and tradeoffs. So, as an administrator, you must choose the clone type that you want to use in a CloudPlatform deployment.

A new global configuration setting, **`vmware.create.full.clone`**, has been added. If the administrator sets this to true, the end users can create guest VMs only as full clones. If the value is set to false, linked clones are used. For fresh installations of CloudPlatform, the default value of this setting is true. For customers upgrading from CloudPlatform 2.x or 3.x, the default value of the **`vmware.create.full.clone`** setting depends on whether the existing setup has any deployed datacenters. If there are no deployed datacenters in the previous-version of the CloudPlatform setup, the default value is true. If there are pre-existing data centers, the default is false.

Citrix recommends you not to change the value of the **`vmware.create.full.clone`** setting in a cloud with running VMs. However, if you change the value, it will not affect the existing VMs. Only the VMs that you create after you change the setting will be subjected to the restriction.

6.3. Accessing Virtual Machines

Any user can access the virtual machines that they created. The administrator can access all the VMs that run in the cloud.

To access a VM through the CloudPlatform UI, do the following:

1. Log-in to the CloudPlatform UI using an administrator account or a user account.
2. On the left navigation bar, click **Instances**.
3. On the right-side panel, in the table that lists the VMs, click the name of a running VM.
4. Under the **Details** tab, click the **View Console** icon.

You can access a virtual machine directly over the network. Before you do this, you must ensure the following:

1. The VM has a port to listen to incoming traffic. For example, in a basic zone, a new VM might be assigned to a security group, which allows incoming traffic. This depends on the security group that you selected for creating the VM. In other cases, you can open a port by setting up a port forwarding policy. For more information, refer to *Section 9.18, "IP Forwarding and Firewalling"*.
2. SSH is enabled on the VM. If the ssh is not enabled on the VM, you cannot access the VM even after you open a port on the VM. This depends on whether ssh is enabled in the template you chose when creating the VM. Access the VM through the CloudPlatform UI and enable ssh on the machine using the commands for the VM's operating system.
3. A firewall rule has been created to allow access if the network has an external firewall device. For more information, refer to *Section 9.18, "IP Forwarding and Firewalling"*.

6.4. Providing a Display Name to the Guest Virtual Machines

Every guest VM has an internal name. The host uses the internal name to identify the guest VMs. CloudPlatform gives you an option to provide a display name to guest VMs. You can add this display name to the internal name so that it is displayed in vCenter. This feature is intended to make the correlation between instance names and internal names easier in large datacenter deployments.

To provide display names to a VM, you need to set the global configuration parameter **`vm.instancename.flag`** to **`true`**. The default value of this parameter is **`false`**.

The default format of the internal name is **`i-<user_id>-<vm_id>-<instance.name>`**, where **`instance.name`** is a global parameter. After you set the *`vm.instancename.flag`* parameter to **`true`** and provide a display name during the creation of a guest VM, the display name will be displayed in vCenter for the guest.

Note

The VMs that are deployed for CloudPlatform version 3.0.7 will continue displaying the VM names in vCenter in the **`InternalName-DisplayName`** format. The VMs that are deployed for the higher versions of CloudPlatform will display their display name in vCenter.

The following table explains how a VM name is displayed in different scenarios.

In the following table, Display Name represents the user-supplied display name.

User-Provided Display Name	vm.instancename.flag	Host name on the VM	Name on vCenter
Yes	True	Display name	Display Name
No	True	<instance.name>-<UUID>	<instance.name>-<UUID>
Yes	False	Display name	i-<user_id>-<vm_id>-<instance.name>
No	False	<instance.name>-<UUID>	i-<user_id>-<vm_id>-<instance.name>

6.5. Stopping and Starting Virtual Machines

After a VM instance is created, you can do the following operations from the **Details** tab of the VM:

- Stop Instance
- Reboot Instance
- Take VM Snapshot
- Destroy Instance
- Attach ISO
- Reset Password
- Migrate instance to another host
- Change service offering
- View console
- Reset VM
- View Volume
- View Snapshot
- View Affinity Groups
- View Hosts

To stop or start VMs, do the following:

1. Log-in to the CloudPlatform UI using an administrator account or a user account.
2. On the left navigation bar, click **Instances**.
3. On the right-side panel, in the table that lists the VMs, click the name of the VM that you want to stop or start.

   You can stop the VMs that are in the Running state and you can start the VMs that are in the Stopped state.
4. Under the **Details** tab, click the **Stop Instance** or the **Start Instance** icon as required.

To reboot the VM, do the following:

1. Navigate to the **Details** view of the VM.
2. Under the **Details** tab, click **Reboot Instance** to reboot the VM.

6.6. Assigning VMs to Hosts

Each virtual machine instance runs on a single host.

CloudPlatform allows you to define affinity groups and assign VMs to them. This helps you influence (but not dictate) the VMs that must run on separate hosts. This feature lets you to specify that certain VMs will not be placed on the same host.

6.6.1. Affinity Groups

By defining affinity groups and assigning VMs to them, the user or administrator can influence (but not dictate) the VMs that should run on separate hosts. This feature lets users specify that VMs with the same “host anti-affinity” type will not be on the same host. This serves to increase fault tolerance. If a host fails, another VM offering the same service (for example, hosting the user's website) is still up and running on another host.

Currently, the affinity groups functionality only works at the user account level. It does not work at the project level.

6.6.2. Creating a New Affinity Group

To add an affinity group, do the following:

1. Log-in to the CloudPlatform UI using an administrator or a user account.
2. In the left navigation bar, click **Affinity Groups**.
3. On the right-side panel, click **Add new affinity group**.
4. In the **Add new affinity group** dialog box, enter the following details:
   - **Name**: A name to identify the affinity group.
   - **Description**: A description of the affinity group that you want to display to users.
   - **Type**: The only supported type shipped with CloudPlatform is Host Anti-Affinity. This indicates that the VMs in this group should avoid being placed on the same VM with each other. If you see other types in this list, it means that your installation of CloudPlatform has been extended with customized affinity group plug-ins.

6.6.3. Assigning a New Virtual Machine to an Affinity Group

To assign a new VM to an affinity group, do the following:

- Create a VM as described in *Section 6.2, “Creating Virtual Machines ”*. In the **Add Instance** wizard, select the affinity group from the **Affinity** tab.

6.6.4. Changing Affinity Group for an Existing Virtual Machine

To assign an existing VM to an affinity group, do the following:

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

2. In the left navigation bar, click **Instances**.
3. On the right side panel, click the name of the VM you want to work with.
4. Under the **Details** tab, click the **Stop Instance** icon to stop the VM.
5. Click the **Change affinity** icon.

6.6.5. Viewing Members of an Affinity Group

To view the VMs that are currently assigned to an affinity group, do the following:

1. Log-in to the CloudPlatform UI using an administrator or a user account.
2. In the left navigation bar, click **Affinity Groups**.
3. On the right side panel, click the name of the affinity group to view its member VMs.
4. Click **View Instances**.

   The members of the group are listed.

   You can click the name of any VM in the list to access its details and controls.

6.6.6. Deleting an Affinity Group

To delete an affinity group, do the following:

1. Log-in to the CloudPlatform UI using an administrator or a user account.
2. In the left navigation bar, click **Affinity Groups**.
3. On the right side panel, click the name of the affinity group you want to delete.
4. Click **Delete**.

   Any VM that is a member of the affinity group will be disassociated from the affinity group. The former group members will continue running normally on the current hosts until the VM is restarted. After you restart the VM, it will not follow the host allocation rules from its former affinity group.

6.7. Configuring VM Snapshots

As a cloud administrator, you can use global configuration variables to control the behavior of VM snapshots. You can click Global Settings on the left navigation bar of the CloudPlatform UI and set the following variables:

Configuration Setting Name	Description
vmsnapshots.max	The maximum number of VM snapshots that can be saved for any given virtual machine in the cloud. The total possible number of VM snapshots in the cloud is (number of VMs) * vmsnapshots.max. If the number of snapshots for any VM ever hits the maximum, the older ones are removed by the snapshot expunge job.

Configuration Setting Name	Description
vmsnapshot.create.wait	Number of seconds to wait for a snapshot job to succeed before declaring failure and issuing an error.

6.7.1. Using VM Snapshots

To create a VM snapshot using the CloudPlatform UI, do the following:

1. Log-in to the CloudPlatform UI using an administrator or a user account.
2. In the left navigation bar, click **Instances**.
3. On the right-side panel, click the name of the VM that you want to snapshot.
4. Under the **Details** tab, click the **Take VM Snapshot** icon.

Note

If a snapshot is already in progress, then clicking this icon will have no effect.

5. In the **Take VM Snapshot** dialog box, do the following:
   - Enter a name and a description for the snapshot. These information will be displayed in the VM Snapshots list.
   - (For running VMs only) If you want to include the VM's memory in the snapshot, click the **Snapshot memory** check box. This saves the CPU and memory state of the virtual machine. If you don't check this box, then only the current state of the VM disk is saved. Checking this box makes the snapshot take longer.
   - Click **OK**.
6. Make selection for the Quiesce VM option.

   Check this box if you want to quiesce the file system on the VM before taking the snapshot.

   When this option is used with CloudPlatform-provided primary storage, the quiesce operation is performed by the underlying hypervisor. When used with another primary storage vendor's plug-in, the quiesce operation is provided according to the vendor's implementation.

   This option is supported only on VMware hypervisor.
7. Click OK.

To delete a snapshot or restore a VM to the state saved in a particular snapshot:

1. Navigate to the VM as described in the earlier steps.
2. Under the **Details** tab, click **View Snapshots**.
3. In the list of snapshots, click the name of the snapshot you want to work with.

4. Do one of the following:

   To delete the snapshot, click the **Delete** icon.

   To revert to the snapshot, click the **Revert** icon.

Note

When a VM is destroyed, the snapshots associated with that VM are deleted automatically. You do not need to manually delete the snapshots in this case.

6.7.2. Limitations on VM Snapshots

You cannot perform the following tasks on a VM that contains VM snapshot:

- Attach a volume to the VM or detach a volume from the VM.
- Attach the VM to a network or detach the VM from the network.
- Resize a volume associated with the VM.
- Change an offering on the VM.
- Create volume snapshots for the volumes that are associated with the VM.
- Perform volume migration on the VM.
- Perform storage migration (both live and cold storage migrations)on the VM.
- Scale the CPU/memory on the VM.
- Reset the VM.

Also, you cannot perform the following actions:

- Creating VM snapshots for the VMs that are associated with a volume, which already contains volume snapshots.
- Reverting to a VM snapshot that will modify the VM status. You are not allowed to do the following actions:
  - Reverting a VM in the 'running' state to a VM snapshot without memory. This changes the VM status to 'stopped'.
  - Reverting a VM in the 'stopped' state to a VM snapshot with memory. This changes the VM status to 'running'.

6.8. Changing the Display Name, OS, or Group of the Virtual Machine

After a VM is created, you can modify its display name, operating system, and the group it belongs to.

To access a VM through the CloudPlatform UI:

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

2. In the left navigation bar, click **Instances**.
3. On the right-side panel, click the name of the VM that you want to modify.
4. Under the **Details** tab, click the **Stop Instance** icon to stop the VM.
5. Click the **Edit** icon.
6. Make the desired changes in the following fields as required:
   - **Display name**: Enter a new display name if you want to change the name of the VM.
   - **OS Type**: Select the desired operating system.
   - **Group**: Enter the group name for the VM.
7. Click **Apply**.

6.9. Changing the Service Offering for a Stopped Virtual Machine

To upgrade or downgrade the level of compute resources available to a virtual machine, you can change the VM's compute offering.

1. Log-in to the CloudPlatform UI using an administrator or a user account.
2. In the left navigation bar, click **Instances**.
3. On the right-side panel, click the name of the VM that you want to work with.
4. Click the **Stop Instance** button to stop the VM.
5. Click the **Change service offering** icon.

   The **Change service offering** dialog box displays.
6. In the **Compute offering** list box, select the offering you want to apply to the VM.
7. Click **OK**.

6.9.1. Scaling CPU and Memory for Running Virtual Machines

(Supported on VMware and XenServer)

It is not always possible to accurately predict the CPU and RAM requirements when you first deploy a VM. You might need to increase these resources at any time during the life of a VM. You can dynamically modify CPU and RAM levels to scale up these resources for a running VM without incurring any downtime.

You can use dynamic CPU and RAM scaling in the following cases:

- User VMs on hosts running VMware and XenServer (XenServer license must support Dynamic Memory Control (DMC)).
- System VMs on VMware.
- VMware Tools or XenServer Tools installed on the virtual machine.

- The new requested CPU and RAM values are greater than or equal to current values within the constraints allowed by the hypervisor and the VM operating system.
- New VMs that are created after the installation of CloudPlatform 4.2 can use the dynamic scaling feature. If you are upgrading from a previous version of CloudPlatform, your existing VMs created with previous versions will not have the dynamic scaling capability unless you update them using the following procedure.

6.9.2. Updating Existing Virtual Machines to Enable Dynamic Scaling Capability

If you are upgrading from a previous version of CloudPlatform, and you want your existing VMs created with previous versions to have the dynamic scaling capability, update the VMs using the following steps:

1. Make sure the zone-level setting **`enable.dynamic.scale.vm`** is set to true. In the left navigation bar of the CloudPlatform UI, click **Infrastructure**.
2. On the right-side panel, under **Zones**, click **View all**.
3. In the page that lists the zones, in the Zone column, click the name of the zone.
4. In the details page, click the **Settings** tab.
5. Install Xen tools (for XenServer hosts) or VMware Tools (for VMware hosts) on each VM if they are not already installed.
6. Stop the VM.
7. Click the Edit button.
8. Click the Dynamically Scalable check box.
9. Click Apply.
10. Restart the VM.

6.9.3. Configuring Dynamic CPU and RAM Scaling

To configure dynamic CPU and RAM scaling, use the following global configuration variables:

- **`enable.dynamic.scale.vm`**: Set to True to enable the feature. By default, the feature is turned off.
- **`scale.retry`**: Number of attempts to perform scaling operation. Default = 2.

6.9.4. How to Dynamically Scale CPU and RAM

To modify the CPU and/or RAM capacity of a virtual machine, you need to change the compute offering of the VM to a new compute offering that has the desired CPU and RAM values. You can use the same steps described in *Section 6.9, "Changing the Service Offering for a Stopped Virtual Machine"*, but skip the step where you stop the virtual machine. You might have to create a new compute offering first.

When you submit a dynamic scaling request, the resources will be scaled up on the current host if possible. If the host does not have enough resources, the VM will be migrated to another host in the

same cluster in real-time. If there is no host in the cluster that can fulfill the requested level of CPU and RAM, the scaling operation will fail. The VM will continue to run as it was before.

6.9.5. Adding Dynamically Scalable Templates

Before you add a template, mark the template as "dynamicallyscalable" and mention that the template contains XS/VMWare tools to support dynamic scaling of VM CPU/memory. When a VM is deployed using this templates, the VM can be dynamically scaled. If such a template is not used, you need to perform the steps in the *Section 6.9.2, "Updating Existing Virtual Machines to Enable Dynamic Scaling Capability"* section to enable dynamic scaling on the existing VMs.

6.9.6. Limitations

- You cannot perform dynamic scaling for system VMs on XenServer.
- Refer to hypervisor and guest operating system documentation before you scale a running VM to the desired offering.
- CloudPlatform will not check that the new CPU and RAM levels are compatible with the OS running on the VM.
- When scaling memory or CPU for a Linux VM on VMware, you might need to run scripts in addition to the other steps mentioned in the procedure. For more information, see VMWare documentation.
- (VMware) If resources are not available on the current host, scaling up will fail on VMware because of a known issue where CloudPlatform and vCenter calculate the available capacity differently. For more information, see *https://issues.apache.org/jira/browse/CLOUDSTACK-1809*.
- On VMs running Linux 64-bit and Windows 7 32-bit operating systems, if the VM is initially assigned a RAM of less than 3 GB, it can be dynamically scaled up to 3 GB, but not more. This is due to a known issue with these operating systems, which will freeze if an attempt is made to dynamically scale from less than 3 GB to more than 3 GB.

6.10. Resetting the Virtual Machine Root Volume on Reboot

For secure environments, and to ensure that VM state is not persisted across reboots, you can reset the root disk.

6.11. Moving VMs Between Hosts (Manual Live Migration)

As CloudPlatform administrator, you can move a running VM from one host to another without interrupting service to users or going into maintenance mode. This is called manual live migration. You must ensure the following before you perform the manual live migration:

- You must have logged-in as root administrator. Domain administrators and users cannot perform manual live migration of VMs.
- The VM must be in the Running state. You cannot perform live migration on VMs in the Stopped state.
- The destination host must have enough available capacity. If not, the VM will remain in the "migrating" state until memory becomes available.
- A VM can be migrated from a tagged host to an untagged host. This is because migration of a VM is an admin operation; therefore, the admin is free to choose where to place a VM.

- (KVM) The VM must not be using local disk storage. (On XenServer and VMware, VM live migration with local disk is enabled by CloudPlatform support for XenMotion and vMotion.)
- (KVM) The destination host must be in the same cluster as the original host. (On XenServer and VMware, VM live migration from one cluster to another is enabled by CloudPlatform support for XenMotion and vMotion.)

To Perform manual live migration on a virtual machine, do the following:

1. Log-in to the CloudPlatform UI using the root administrator account.
2. In the left navigation, click **Instances**.
3. On the right-side panel, click the name of the VM that you want to migrate in real-time.
4. Under the **Details** tab, click the **Migrate instance to another host** icon.
5. From the list of suitable hosts, choose the one where you want to migrate the VM.

Note

If the VM's storage has to be migrated along with the VM, this will be noted in the host list. CloudPlatform will take care of the storage migration for you.

6. Click **OK**.

6.12. Deleting Virtual Machines

Users can delete their own virtual machines. A running virtual machine will be stopped before it is deleted. Administrators can delete any virtual machines.

To delete a virtual machine, do the following:

1. Log-in to the CloudPlatform UI using an administrator or a user account.
2. In the left navigation bar, click **Instances**.
3. On the right-side panel, select the VM that you want to delete.
4. Under the **Details** tab, click **Stop Instance** if the VM has not already been stopped.
5. Click the **Destroy Instance** button.

6.13. Recovering a Destroyed VM

Users can recover their virtual machines that are destroyed. Administrators can recover any destroyed virtual machines.

To recover a virtual machine, do the following:

1. Log-in to the CloudPlatform UI using an administrator or a user account.
2. In the left navigation bar, click **Instances**.

3. On the right-side panel, click the name of the VM that you want to recover.
4. Under the **Details** tab, click the **Restore Instance** icon.

Working With Hosts

7.1. Adding Hosts

A host is a single computer that provides the computing resources that run guest virtual machines. Each host will have a hypervisor software installed on it to manage the guest VMs. The host is the smallest organizational unit within a CloudPlatform deployment. As an administrator, you can add additional hosts at any time to provide more capacity for guest VMs.

For conceptual information about hosts and cloud infrastructure, refer to **Chapter 4: Cloud Infrastructure Concepts** in the *Citrix CloudPlatform 4.5 (powered by Apache CloudStack) Concepts Guide*.

7.2. Scheduled Maintenance and Maintenance Mode for Hosts

You can place a host into maintenance mode. When maintenance mode is activated, the host becomes unavailable to receive new guest VMs, and the guest VMs already running on the host are seamlessly migrated to another host not in maintenance mode. This migration uses live migration technology and does not interrupt the execution of the guest.

7.2.1. vCenter and Maintenance Mode

To enter maintenance mode on a vCenter host, both vCenter and CloudPlatform must be used in concert. CloudPlatform and vCenter have separate maintenance modes that work closely together.

1. Place the host into CloudPlatform's "scheduled maintenance" mode. This does not invoke the vCenter maintenance mode, but only causes VMs to be migrated off the host

   When the CloudPlatform maintenance mode is requested, the host first moves into the Prepare for Maintenance state. In this state it cannot be the target of new guest VM starts. Then all VMs will be migrated off the server. Live migration will be used to move VMs off the host. This allows the guests to be migrated to other hosts with no disruption to the guests. After this migration is completed, the host will enter the Ready for Maintenance mode.

2. Wait for the "Ready for Maintenance" indicator to appear in the UI.

3. Now use vCenter to perform whatever actions are necessary to maintain the host. During this time, the host cannot be the target of new VM allocations.

4. When the maintenance tasks are complete, take the host out of maintenance mode as follows:

   a. First use vCenter to exit the vCenter maintenance mode.

      This makes the host ready for CloudPlatform to reactivate it.

   b. Then use CloudPlatform's administrator UI to cancel the CloudPlatform maintenance mode.

      When the host comes back online, the VMs that were migrated off of it are migrated back to it and new VMs can be added.

7.2.2. XenServer Maintenance Mode

For XenServer, you can take a server offline temporarily from the CloudPlatform UI. When you place a server in maintenance mode, all VMs running on the server automatically migrate to another host in the same pool. You cannot create or start any VMs on a server that is in maintenance mode.

To place a server in maintenance mode:

1. In the CloudPlatform UI, on the left panel, click **Infrastructure**.
2. On the right-panel, under **Hosts**, click **View All**.
3. In the table that lists the hosts, identify the XenServer host that you want to place in maintenance mode.
4. In the **Name** column, click the name of the XenServer host.
5. Under **Details** tab, click the **Enable Maintenance Mode** icon to enable maintenance mode on the host.
6. In the **Confirmation** message box, click **Yes**.

To disable maintenance mode on a server:

1. In the CloudPlatform UI, on the left panel, click **Infrastructure**.
2. On the right-panel, under **Hosts**, click **View All**.
3. In the table that lists the hosts, identify the XenServer host where you want to disable maintenance mode.
4. In the **Name** column, click the name of the XenServer host.
5. Under **Details** tab, click the **Cancel Maintenance Mode** icon to disable maintenance mode on the host.
6. In the **Confirmation** message box, click **Yes**.

7.2.3. Hyper-V Maintenance Mode

For CloudPlatform 4.5, the latest Hyper-V version - Windows Server 2012 R2 (with Hyper-V Role enabled) and Hyper-V Server 2012 R2 - is certified. If you apply patches from Microsoft in future and you want to perform the maintenance, you can follow a maintenance procedure similar to the following for Hyper-V hosts:

1. Place a host to maintenance mode.
2. Perform necessary upgrades.
3. Add the host back.

7.2.4. Baremetal Maintenance Mode

For BM, you can take a host offline temporarily by using the Enable Maintenance option in CloudPlatform. When you place a host into Maintenance mode, the VM is automatically stopped. While a host is in Maintenance mode, you cannot create or start a VM on it. Use the Cancel Maintenance option to take the host out of Maintenance mode.

To place a server in maintenance mode:

1. In the CloudPlatform UI, on the left panel, click **Infrastructure**.
2. On the right-panel, under **Hosts**, click **View All**.
3. In the table that lists the hosts, identify the Baremetal host that you want to place in maintenance mode.
4. In the **Name** column, click the name of the Baremetal host.
5. Under **Details** tab, click the **Enter Maintenance Mode** icon to enable maintenance mode on the host.
6. In the **Confirmation** message box, click **Yes**.

To disable maintenance mode on a server:

1. In the CloudPlatform UI, on the left panel, click **Infrastructure**.
2. On the right-panel, under **Hosts**, click **View All**.
3. In the table that lists the hosts, identify the Baremetal host where you want to disable maintenance mode.
4. In the **Name** column, click the name of the Baremetal host.
5. Under **Details** tab, click the **Cancel Maintenance Mode** icon to disable maintenance mode on the host.
6. In the **Confirmation** message box, click **Yes**.

7.3. Disabling and Enabling Zones, Pods, and Clusters

You can enable or disable a zone, pod, or cluster without permanently removing it from the cloud. This is useful for maintenance or when there are problems that make a portion of the cloud infrastructure unreliable. No new allocations will be made to a disabled zone, pod, or cluster until its state is returned to Enabled. When a zone, pod, or cluster is first added to the cloud, it is Disabled by default.

Note

Disabling a specific infrastructure component on a VM that is in the running state will not affect the functioning of the VM. The VM will continue functioning normally.

To disable and enable a zone, pod, or cluster:

1. Log in to the CloudPlatform UI as administrator
2. In the left navigation bar, click Infrastructure.
3. In Zones, click View More.
4. If you are disabling or enabling a zone, find the name of the zone in the list, and click the Enable/Disable button. 

5. If you are disabling or enabling a pod or cluster, click the name of the zone that contains the pod or cluster.
6. Click the Compute tab.
7. In the Pods or Clusters node of the diagram, click View All.
8. Click the pod or cluster name in the list.
9. Click the Enable/Disable button. 

7.4. Disabling Hosts

From the CloudPlatform UI, you can disable hosts without permanently removing them from the cloud.

To disable hosts from the CloudPlatform UI, do the following:

1. Log-in to the CloudPlatform UI using administrator privileges.
2. In the left navigation panel, click **Infrastructure**.
3. On the right-panel, under **Hosts**, click **View All**.
4. In the table that lists the hosts, identify the host that you want to disable.
5. In the **Name** column, click the name of the host.
6. Under **Details** tab, click the **Disable Host** icon.
7. In the Confirmation message box, click **Yes**.

7.5. Removing Hosts

Hosts can be removed from the cloud as needed. The procedure to remove a host depends on the hypervisor type.

7.5.1. Removing a XenServer Host

1. Remove the host from the CloudPlatform UI.
    a. Log in to the CloudPlatform UI.
    b. In the left navigation bar, select Infrastructure.
    c. Select View All under Hosts.
    d. Locate the host for deletion.
    e. Click Enable Maintenance mode.
    f. Click Remove.
2. Remove this host from the XenServer pool.
    a. If the host is XenServer 6.2 SP1 Hotfix XS62ESP1004, disable Pool HA, else continue with the next step.

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

b. If you are removing a master host, designate one of the slave hosts as the new master.

```
#  xe pool-designate-new-master host-uuid="new master uuid"
```

Wait for 10 minute for slave to master transition.

c. Eject the host from the pool.

```
# xe pool-eject host-uuid=" uuid of the host"
```

d. If the host is XenServer 6.2 SP1 Hotfix XS62ESP1004, enable Pool HA.

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

Note

When enabling the pool HA again, ensure that you run the **`xe pool-ha-enable`** command with the *`heartbeat-sr-uuids`* parameter pointing to the correct HA Storage Repository. If the *`heartbeat-sr-uuids`* parameter is skipped, any Storage Repository might randomly be considered for HA, which needs to be avoided.

7.5.2. Removing KVM Hosts

A node cannot be removed from a cluster until it has been placed in maintenance mode. This will ensure that all of the VMs on it have been migrated to other Hosts. To remove a Host from the cloud:

1. Place the node in maintenance mode.

   See *Section 7.2, “Scheduled Maintenance and Maintenance Mode for Hosts ”*.

2. Stop the cloud-agent service.

3. Use the UI option to remove the node.

   Then you may power down the Host, re-use its IP address, re-install it, etc

7.5.3. Removing vSphere Hosts

To remove this type of host, first place it in maintenance mode, as described in *Section 7.2, “Scheduled Maintenance and Maintenance Mode for Hosts ”*. Then use CloudPlatform to remove the host. CloudPlatform will not direct commands to a host that has been removed using CloudPlatform. However, the host may still exist in the vCenter cluster.

7.5.4. Removing Hyper-V Hosts

To remove this type of host, first place it in maintenance mode, as described in *Section 7.2, “Scheduled Maintenance and Maintenance Mode for Hosts ”*. Then use CloudPlatform UI to

remove the host. CloudPlatform will not direct commands to a host that has been removed using CloudPlatform.

7.6. Re-Installing Hosts

You can re-install a host after placing it in maintenance mode and then removing it. If a host is down and cannot be placed in maintenance mode, it should still be removed before the re-install.

7.7. Maintaining Hypervisors on Hosts

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

Note

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

7.8. Changing Host Password

The password for a XenServer Node, KVM Node, or vSphere Node may be changed in the database. Note that all Nodes in a Cluster must have the same password.

To change a Node's password:

1. Identify all hosts in the cluster.
2. Change the password on all hosts in the cluster. Now the password for the host and the password known to CloudPlatform will not match. Operations on the cluster will fail until the two passwords match.

   

   Note

   As a best practice, you must unmanage the cluster before you change the password on the hosts. If you do not unmanage the cluster, you will view many error messages.

3. Get the list of host IDs for the host in the cluster where you are changing the password. You will need to access the database to determine these host IDs. For each host name "h" (or vSphere cluster) that you are changing the password for, execute:

   ```
   mysql> select id from cloud.host where name like '%h%';
   ```

4. This should return a single ID. Record the set of such IDs for these hosts.

5. Update the passwords for the host in the database. In this example, we change the passwords for hosts with IDs 5, 10, and 12 to "password".

```
mysql> update cloud.host_details set value='password' where host_id is (5, 10, 12) and
 name=password;
```

6. Manage the cluster through the CloudPlatform UI.

7.9. Over-Provisioning and Service Offering Limits

(Supported for XenServer, KVM, and VMware)

CPU and memory (RAM) over-provisioning factors can be set for each cluster to change the number of VMs that can run on each host in the cluster. This helps optimize the use of resources. By increasing the over-provisioning ratio, more resource capacity will be used. If the ratio is set to 1, no over-provisioning is done.

The administrator can also set global default over-provisioning ratios in the cpu.overprovisioning.factor and mem.overprovisioning.factor global configuration variables. The default value of these variables is 1: over-provisioning is turned off by default.

Over-provisioning ratios are dynamically substituted in CloudPlatform's capacity calculations. For example:

Capacity = 2 GB

Over-provisioning factor = 2

Capacity after over-provisioning = 4 GB

With this configuration, suppose you deploy 3 VMs of 1 GB each:

Used = 3 GB

Free = 1 GB

The administrator can specify both CPU and memory over-provisioning ratios on a per-cluster basis.

In any given cloud, the optimum number of VMs for each host is affected by such things as the hypervisor, storage, and hardware configuration. These may be different for each cluster in the same cloud. A single global over-provisioning setting can not provide the best utilization for all the different clusters in the cloud. It has to be set for the lowest common denominator. The per-cluster setting provides a finer granularity for better utilization of resources, no matter where the CloudPlatform placement algorithm decides to place a VM.

The overprovisioning settings can be used along with dedicated resources (assigning a specific cluster to an account) to effectively offer different levels of service to different accounts. For example, an account paying for a more expensive level of service could be assigned to a dedicated cluster with an over-provisioning ratio of 1, and a lower-paying account to a cluster with a ratio of 2.

When a new host is added to a cluster, CloudPlatform will assume the host has the capability to perform the CPU and RAM over-provisioning which is configured for that cluster. It is up to the administrator to be sure the host is actually suitable for the level of over-provisioning which has been set.

7.9.1. Limitations on Over-Provisioning in XenServer and KVM

- In XenServer, due to a constraint of this hypervisor, you can not use an over-provisioning factor greater than 4.
- The KVM hypervisor can not manage memory allocation to VMs dynamically. CloudPlatform sets the minimum and maximum amount of memory that a VM can use. The hypervisor adjusts the memory within the set limits based on the memory contention.

7.9.2. Requirements for Over-Provisioning

Several prerequisites are required in order for over-provisioning to function properly. The feature is dependent on the OS type, hypervisor capabilities, and certain scripts. It is the administrator's responsibility to ensure that these requirements are met.

7.9.2.1. Balloon Driver

All VMs should have a balloon driver installed in them. The hypervisor communicates with the balloon driver to free up and make the memory available to a VM.

XenServer

The balloon driver can be found as a part of xen pv or PVHVM drivers. The xen pvhvm drivers are included in upstream linux kernels 2.6.36+.

VMware

The balloon driver can be found as a part of the VMware tools. All the VMs that are deployed in a over-provisioned cluster should have the VMware tools installed.

KVM

All VMs are required to support the virtio drivers. These drivers are installed in all Linux kernel versions 2.6.25 and greater. The administrator must set CONFIG_VIRTIO_BALLOON=y in the virtio configuration.

7.9.2.2. Hypervisor capabilities

The hypervisor must be capable of using the memory ballooning.

XenServer

The DMC (Dynamic Memory Control) capability of the hypervisor should be enabled. Only XenServer Advanced and above versions have this feature.

VMware, KVM

Memory ballooning is supported by default.

7.9.3. Setting Over-Provisioning Ratios

There are two ways the root admin can set CPU and RAM over-provisioning ratios. First, the global configuration settings cpu.overprovisioning.factor and mem.overprovisioning.factor will be applied when a new cluster is created. Later, the ratios can be modified for an existing cluster.

Only VMs deployed after the change are affected by the new setting. If you want VMs deployed before the change to adopt the new over-provisioning ratio, you must stop and restart the VMs. When this is

done, CloudPlatform recalculates or scales the used and reserved capacities based on the new over-provisioning ratios, to ensure that CloudPlatform is correctly tracking the amount of free capacity.

Note

It is safer not to deploy additional new VMs while the capacity recalculation is underway, in case the new values for available capacity are not high enough to accommodate the new VMs. Just wait for the new used/available values to become available, to be sure there is room for all the new VMs you want.

To change the over-provisioning ratios for an existing cluster:

1. Log in as administrator to the CloudPlatform UI.
2. In the left navigation bar, click Infrastructure.
3. Under Clusters, click View All.
4. Select the cluster you want to work with, and click the Edit button.
5. Fill in your desired over-provisioning multipliers in the fields CPU overcommit ratio and RAM overcommit ratio. The value which is initially shown in these fields is the default value inherited from the global configuration settings.

   

   Note

   In XenServer, due to a constraint of this hypervisor, you can not use an over-provisioning factor greater than 4.

7.9.4. Changing Over-Provisioning Ratios with Running VMs

Ideally, you shouldn't change the over-provisioning factor in a cluster with VMs running, because the existing VMs were deployed with the previously specified factor. However, it is possible to change the over-provisioning factor, if you keep the following considerations in mind.

When you change the over-provisioning factor for a cluster where VMs are running, both used and total capacity are multiplied by this factor to keep track of available capacity. For example:

```
Cluster – c,
cpu over provisioning = 1,
Total cpu = 2GHZ
```

Suppose you deploy 2 VMs which each have a 512Mhz service offering. Then:

```
totalCapacity = 2GHz
AvailableCapacity = 1GHz
UsedCapacity  = 1GHZ
```

Now suppose you change the CPU over-provisioning ratio of cluster c to 2:

```
totalCapacity = 4GHz
AvailableCapacity = 2GHz
UsedCapacity  = 2GHZ
```

Notice the difference. Both used and total capacity are multiplied by the over-provisioning factor. Used Capacity after changing the factor is:

```
(service offering of VM / overcommit it was originally deployed with) * new overcommit
```

Or, to plug in our specific example values:

```
(1GHZ/1)*2
```

CloudPlatform uses this calculation in order to guarantee minimum CPU in case of contention. When a VM is deployed with an overprovisioing factor of "x", CloudPlatform aims to guarantee (service offering of VM / x) during the VM's lifecycle, even though the overprovisioning factor might have changed. CloudPlatform scales the total used CPU value in order to keep track of the actual amount of CPU left on the host.

Now suppose you launch 2 VMs which each have a 1Ghz CPU service offering:

```
totalCapacity = 4GHz
AvailableCapacity = 0GHz
UsedCapacity  = 4GHZ
```

The calculation for used capacity for all 4 VMs is:

```
((service offering of vm / overcommit it got deployed with) * new overcommit)
```

Or to plug in our specific example values:

```
(512Mhz/1)*2 + (512Mhz/1)*2 + (1Ghz/2)*2 + (1Ghz/2)*2 = 4Ghz
```

Now suppose you change the over-provisioning to 3:

```
totalCapacity = 6 GHz
AvailableCapacity = 0 GHz
UsedCapacity  = 6 GHZ
```

The calculation for used capacity for 4 VMs is:

```
((service offering of vm / overcommit it got deployed with) * new overcommit)
```

Or to plug in our specific example values:

```
(512Mhz/1)*3 +(512Mhz/1)*3 +(1Ghz/2)*3 + (1Ghz/2)*3 = 6Ghz
```

So far, this example is assuming you haven't stopped and started the VMs all this while. Suppose now you stop and start 1 VM at 512Mhz and another VM at 1Ghz. The over-provisioning factor changes for these VMs to 3 each. Note the denominator in the calculation:

```
totalCapacity = 6 GHz
AvailableCapacity = 1.5 GHz
```

```
UsedCapacity  = 4.5 GHZ
```

The calculation for used capacity for 4 VMs is:

```
((service offering of vm / overcommit it got deployed with) * new overcommit)
```

Or to plug in our specific example values:

```
(512Mhz/3)*3 +(512Mhz/1)*3 +(1Ghz/3)*3 + (1Ghz/2)*3 = 4.5 Ghz
```

This over-provisioning model guarantees QOS as the following during the lifecycle of any VM:

```
(service offering of vm / x )
```

7.9.5. Enforcement of Service Offering Limits

Service offering limits (e.g. 1 GHz, 1 core) are strictly enforced for core count. For example, a guest with a service offering of one core will have only one core available to it regardless of other activity on the Host.

Service offering limits for gigahertz are enforced only in the presence of contention for CPU resources. For example, suppose that a guest was created with a service offering of 1 GHz on a Host that has 2 GHz cores, and that guest is the only guest running on the Host. The guest will have the full 2 GHz available to it. When multiple guests are attempting to use the CPU a weighting factor is used to schedule CPU resources. The weight is based on the clock speed in the service offering. Guests receive a CPU allocation that is proportionate to the GHz in the service offering. For example, a guest created from a 2 GHz service offering will receive twice the CPU allocation as a guest created from a 1 GHz service offering.

7.10. VLAN Provisioning

CloudPlatform automatically creates and destroys interfaces bridged to VLANs on the hosts. In general the administrator does not need to manage this process.

CloudPlatform manages VLANs differently based on hypervisor type. For XenServer or KVM, the VLANs are created on only the hosts where they will be used and then they are destroyed when all guests that require them have been terminated or moved to another host.

For vSphere the VLANs are provisioned on all hosts in the cluster even if there is no guest running on a particular Host that requires the VLAN. This allows the administrator to perform live migration and other functions in vCenter without having to create the VLAN on the destination Host. Additionally, the VLANs are not removed from the Hosts when they are no longer needed.

You can use the same VLANs on different physical networks provided that each physical network has its own underlying layer-2 infrastructure, such as switches. For example, you can specify VLAN range 500 to 1000 while deploying physical networks A and B in an Advanced zone setup. This capability allows you to set up an additional layer-2 physical infrastructure on a different physical NIC and use the same set of VLANs if you run out of VLANs. Another advantage is that you can use the same set of IPs for different customers, each one with their own routers and the guest networks on different physical NICs.

7.10.1. VLAN Allocation Example

VLANs are required for public and guest traffic. The following is an example of a VLAN allocation scheme:

VLAN IDs	Traffic type	Scope
less than 500	Management traffic. Reserved for administrative purposes.	CloudPlatform software, hypervisors, and system VMs can access these VLANs.
500-599	VLAN carrying public traffic.	CloudPlatform accounts.
600-799	VLANs carrying guest traffic.	CloudPlatform accounts. Account-specific VLAN is chosen from this pool.
800-899	VLANs carrying guest traffic.	CloudPlatform accounts. Account-specific VLAN chosen by CloudPlatform admin to assign to that account.
900-999	VLAN carrying guest traffic	CloudPlatform accounts. Can be scoped by project, domain, or all accounts.
greater than 1000	Reserved for future use	

7.10.2. Adding Non Contiguous VLAN Ranges

CloudPlatform provides you with the flexibility to add non contiguous VLAN ranges to your network. The administrator can either update an existing VLAN range or add multiple non contiguous VLAN ranges while creating a zone. You can also use the UpdatephysicalNetwork API to extend the VLAN range.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. Ensure that the VLAN range does not already exist.
3. In the left navigation, choose Infrastructure.
4. On Zones, click View More, then click the zone with which you want to work.
5. Click Physical Network.
6. In the Guest node of the diagram, click Configure.
7. Click **Edit**.

   The **VLAN Ranges** field becomes editable.
8. Enter the start and end of the VLAN range. If you have multiple ranges, separate them by a comma.

   For example: 200-210,300-350,500-600, 100-110

   Specify all the VLANs you want to use, VLANs not specified will be removed if you are adding new ranges to the existing list.
9. Click Apply.

7.10.3. Assigning VLANs to Isolated Networks

CloudPlatform provides you the ability to control VLAN assignment to isolated networks. You can assign a VLAN ID when a network is created, just the way it's done for Shared networks.

The former behaviour also is supported — VLAN is randomly allocated to a network from the VNET range of the physical network when the network turns to Implemented state. The VLAN is released back to the VNET pool when the network shuts down as a part of the Network Garbage Collection. The VLAN can be re-used either by the same network when it is implemented again, or by any other network. On each subsequent implementation of a network, a new VLAN can be assigned.

To assign VLANs to Isolated networks,

1. In the CloudPlatform UI, on the left navigation bar, click **Service Offerings**.
2. In the right-side panel, in the **Select offering** list box, select **Network Offerings**.
3. Click **Add network offering**.
4. To create a network offering, enter the details in the **Add network offering** dialog box and click **OK**
   - **Guest Type**: Select Isolated.
   - **Specify VLAN**: Select the option.
5. Using this network offering, create a network.

   You can create a VPC tier or an Isolated network.
6. Specify the VLAN when you create the network.

   When VLAN is specified, a CIDR and gateway are assigned to this network and the state is changed to Setup. In this state, the network will not be garbage collected.

Note

You cannot change a VLAN once it's assigned to the network. The VLAN remains with the network for its entire life cycle.

Working With Storage

CloudPlatform defines two types of storage: primary and secondary.

Primary storage can be accessed by either iSCSI or NFS. Additionally, direct attached storage may be used for primary storage. Secondary storage is always accessed using NFS or a combination of NFS and object storage. Secondary storage stores the templates, ISO images, and disk volume snapshots. The items in secondary storage are available to all hosts in the scope of the secondary storage, which may be defined as per zone or per region.

Best Practices for Primary Storage

- The speed of primary storage will impact guest performance. If possible, choose smaller, higher RPM drives for primary storage.
- Ensure that nothing is stored on the server. Adding the server to CloudPlatform will destroy any existing data

This chapter decsribes how to work with secondary storage and storage volumes in CloudPlatform.

For conceptual information about storage in CloudPlatform, refer to **Chapter 7: Storage Concepts Used in CloudPlatform** in the *Citrix CloudPlatform 4.5 (powered by Apache CloudStack) Concepts Guide*.

8.1. Secondary Storage

This section discusses the tasks related to CloudPlatform secondary storage.

8.1.1. Best Practices for Secondary Storage

- Each Zone can have one or more secondary storage servers. Multiple secondary storage servers provide increased scalability to the system.
- Secondary storage has a high read:write ratio and is expected to consist of larger drives with lower IOPS than primary storage.
- Ensure that nothing is stored on the server. Adding the server to CloudPlatform will destroy any existing data.

8.1.2. Changing the Secondary Storage IP Address

You can change the secondary storage IP address after it has been provisioned. After changing the IP address on the host, log in to your management server and execute the following commands. Replace HOSTID below with your own value, and change the URL to use the appropriate IP address and path for your server:

```
# mysql -p
mysql> use cloud;
mysql> select id from image_store where url like '<old ip address<';
mysql> update image_store set name = 'nfs://192.168.160.20/export/mike-ss1' where id = #;
mysql> update image_store set url = 'nfs://192.168.160.20/export/mike-ss1' where id = #;
```

Note

When copying and pasting a command, be sure the command has pasted as a single line before executing. Some document viewers may introduce unwanted line breaks in copied text.

Then log in to the CloudPlatform UI and stop and start (not reboot) the Secondary Storage VM for that Zone.

8.1.3. Changing Secondary Storage Servers

You can change the secondary storage NFS mount. Perform the following steps to do so:

1. Stop all running Management Servers.
2. Wait 30 minutes. This allows any writes to secondary storage to complete.
3. Copy all files from the old secondary storage mount to the new.
4. Change the IP address for secondary storage if required. See *Section 8.1.2, "Changing the Secondary Storage IP Address"*.
5. Start the Management Server.

8.2. Working With Volumes

A volume provides storage to a guest VM. The volume can provide for a root disk or an additional data disk. CloudPlatform supports additional volumes for guest VMs.

Volumes are created for a specific hypervisor type. A volume that has been attached to guest using one hypervisor type (for example, XenServer) may not be attached to a guest that is using another hypervisor type (for example, vSphere, Oracle VM, KVM). This is because the different hypervisors use different disk image formats.

CloudPlatform defines a volume as a unit of storage available to a guest VM. Volumes are either root disks or data disks. The root disk has "/" in the file system and is usually the boot device. Data disks provide for additional storage (for example, as "/opt" or "D:"). Every guest VM has a root disk, and VMs can also optionally have a data disk. End users can mount multiple data disks to guest VMs. Users choose data disks from the disk offerings created by administrators. The user can create a template from a volume as well; this is the standard procedure for private template creation. Volumes are hypervisor-specific: a volume from one hypervisor type may not be used on a guest of another hypervisor type.

Note

CloudPlatform supports attaching up to 13 data disks to a VM on XenServer hypervisor versions 6.0 and above. For the VMs on other hypervisor types, the data disk limit is 6.

8.2.1. Creating a New Volume

You can add more data disk volumes to a guest VM at any time, up to the limits of your storage capacity. Both CloudPlatform administrators and users can add volumes to VM instances. When you create a new volume, it is stored as an entity in CloudPlatform, but the actual storage resources are not allocated on the physical storage device until you attach the volume. This optimization allows the CloudPlatform to provision the volume nearest to the guest that will use it when the first attachment is made.

8.2.1.1. Using Local Storage for Data Volumes

You can create data volumes on local storage (supported with XenServer, KVM, and VMware). The data volume is placed on the same host as the VM instance that is attached to the data volume. These local data volumes can be attached to virtual machines, detached, re-attached, and deleted just as with the other types of data volume.

Local storage is ideal for scenarios where persistence of data volumes and HA is not required. Some of the benefits include reduced disk I/O latency and cost reduction from using inexpensive local disks.

In order for local volumes to be used, the feature must be enabled for the zone.

You can create a data disk offering for local storage. When a user creates a new VM, they can select this disk offering in order to cause the data disk volume to be placed in local storage.

You can not migrate a VM that has a volume in local storage to a different host, nor migrate the volume itself away to a different host. If you want to put a host into maintenance mode, you must first stop any VMs with local data volumes on that host.

8.2.1.2. To Create a New Volume

1. Log in to the CloudPlatform UI as a user or admin.
2. In the left navigation bar, click Storage.
3. In Select View, choose Volumes.
4. To create a new volume, click Add Volume, provide the following details, and click OK.
   - Name. Give the volume a unique name so you can find it later.
   - Availability Zone. Where do you want the storage to reside? This should be close to the VM that will use the volume.
   - Disk Offering. Choose the characteristics of the storage.

   The new volume appears in the list of volumes with the state “Allocated.” The volume data is stored in CloudPlatform, but the volume is not yet ready for use
5. To start using the volume, continue to Attaching a Volume

8.2.2. Uploading an Existing Volume to a Virtual Machine

Existing data can be made accessible to a virtual machine. This is called uploading a volume to the VM. For example, this is useful to upload data from a local file system and attach it to a VM. Root administrators, domain administrators, and end users can all upload existing volumes to VMs.

The upload is performed using HTTP. The uploaded volume is placed in the zone's secondary storage

You cannot upload a volume if the preconfigured volume limit has already been reached. The default limit for the cloud is set in the global configuration parameter max.account.volumes, but administrators can also set per-domain limits that are different from the global default. See Setting Usage Limits

To upload a volume:

1. (Optional) Create an MD5 hash (checksum) of the disk image file that you are going to upload. After uploading the data disk, CloudPlatform will use this value to verify that no data corruption has occurred.
2. Log in to the CloudPlatform UI as an administrator or user
3. In the left navigation bar, click Storage.
4. Click Upload Volume.
5. Provide the following:
   - Name and Description. Any desired name and a brief description that can be shown in the UI.
   - Availability Zone. Choose the zone where you want to store the volume. VMs running on hosts in this zone can attach the volume.
   - Format. Choose one of the following to indicate the disk image format of the volume.

Hypervisor	Disk Image Format
XenServer	VHD
Hyper-V	VHD
VMware	OVA
KVM	QCOW2

   - URL. The secure HTTP or HTTPS URL that CloudPlatform can use to access your disk. The type of file at the URL must match the value chosen in Format. For example, if Format is VHD, the URL might look like the following:

     http://yourFileServerIP/userdata/myDataDisk.vhd

   - MD5 checksum. (Optional) Use the hash that you created in step 1.
6. Wait until the status of the volume shows that the upload is complete.
7. Click Instances - Volumes, find the name you specified in step 5, and make sure the status is Uploaded.

8.2.3. Attaching a Volume

You can attach a volume to a guest VM to provide extra disk storage. Attach a volume when you first create a new volume, when you are moving an existing volume from one VM to another, or after you have migrated a volume from one storage pool to another.

1. Log in to the CloudPlatform UI as a user or admin.
2. In the left navigation, click Storage.
3. In Select View, choose Volumes.

4. Click the volume name in the Volumes list, then click the Attach Disk button 

5. In the Instance window, choose the VM to which you want to attach the volume. You will only see instances to which you are allowed to attach volumes; for example, a user will see only instances created by that user, but the administrator will have more choices.

6. When the volume has been attached, you should be able to see it by clicking Instances, the instance name, and View Volumes.

8.2.4. Detaching and Moving Volumes

Note

This procedure is different from moving volumes from one storage pool to another as described in *Section 8.2.5, "VM Storage Migration"*.

A volume can be detached from a guest VM and attached to another guest. Both CloudPlatform administrators and users can detach volumes from VMs and move them to other VMs.

If the two VMs are in different clusters, and the volume is large, it may take several minutes for the volume to be moved to the new VM.

1. Log in to the CloudPlatform UI as a user or admin.

2. In the left navigation bar, click Storage, and choose Volumes in Select View. Alternatively, if you know which VM the volume is attached to, you can click Instances, click the VM name, and click View Volumes.

3. Click the name of the volume you want to detach, then click the Detach Disk button. 

4. To move the volume to another VM, follow the steps in *Section 8.2.3, "Attaching a Volume"*.

8.2.5. VM Storage Migration

Supported in XenServer, KVM, and VMware.

Note

This procedure is different from moving disk volumes from one VM to another as described in *Section 8.2.4, "Detaching and Moving Volumes"*.

You can migrate a virtual machine's root disk volume or any additional data disk volume from one storage pool to another in the same zone.

You can use the storage migration feature to achieve some commonly desired administration goals, such as balancing the load on storage pools and increasing the reliability of virtual machines by moving them away from any storage pool that is experiencing issues.

On XenServer and VMware, live migration of VM storage is enabled through CloudPlatform support for XenMotion and vMotion. Live storage migration allows VMs to be moved from one host to another, where the VMs are not located on storage shared between the two hosts. It provides the option to live migrate a VM's disks along with the VM itself. It is possible to migrate a VM from one XenServer resource pool / VMware cluster to another, or to migrate a VM whose disks are on local storage, or even to migrate a VM's disks from one storage repository to another, all while the VM is running.

Note

Because of a limitation in VMware, live migration of storage for a VM is allowed only if the source and target storage pool are accessible to the source host; that is, the host where the VM is running when the live migration operation is requested.

8.2.5.1. Migrating a Data Volume to a New Storage Pool

There are two situations when you might want to migrate a disk:

- Move the disk to new storage, but leave it attached to the same running VM.
- Detach the disk from its current VM, move it to new storage, and attach it to a new VM.

8.2.5.1.1. Migrating Storage For a Running VM

(Supported on XenServer and VMware)

1. Log in to the CloudPlatform UI as a user or admin.
2. In the left navigation bar, click Instances, click the VM name, and click View Volumes.
3. Click the volume you want to migrate.
4. Detach the disk from the VM. See *Section 8.2.4, "Detaching and Moving Volumes"* but skip the "reattach" step at the end. You will do that after migrating to new storage.
5. Click the Migrate Volume button  and choose the destination from the dropdown list.
6. Watch for the volume status to change to Migrating, then back to Ready.

8.2.5.1.2. Migrating Storage and Attaching to a Different VM

1. Log-in to the CloudPlatform UI as a user or administrator.
2. Detach the disk from the VM. See *Section 8.2.4, "Detaching and Moving Volumes"*, but skip the "reattach" step at the end. You will do that after migrating to new storage.
3. In the **Details** page of the volume, click **Migrate Volume**, select the destination storage pool from the dropdown list, and click **OK**.

Note

Normal users do not have the permission to execute the volume migration operation. If you want to do this operation, contact the owner of this VM.

4. Watch for the volume status to change to Migrating, then back to Ready. You can find the volume by clicking Storage in the left navigation bar. Make sure that Volumes is displayed at the top of the window, in the Select View dropdown.

5. Attach the volume to any desired VM running in the same cluster as the new storage server. See *Section 8.2.3, "Attaching a Volume"*

8.2.5.2. Migrating a VM Root Volume to a New Storage Pool

(XenServer, VMware) You can live migrate a VM's root disk from one storage pool to another, without stopping the VM first.

(KVM) When migrating the root disk volume, the VM must first be stopped, and users can not access the VM. After migration is complete, the VM can be restarted.

1. Log in to the CloudPlatform UI as a user or administrator.

Note

Normal users do not have the permission to execute the live volume migration operation.

2. In the left navigation bar, click Instances, and click the VM name.

3. (KVM only) Stop the VM.

4. Click the Migrate button  and choose the destination from the dropdown list.

Note

If the VM's storage has to be migrated along with the VM, this will be noted in the host list. CloudPlatform will take care of the storage migration for you.

5. Watch for the volume status to change to Migrating, then back to Running (or Stopped, in the case of KVM). This can take some time.

6. (KVM only) Restart the VM.

8.2.6. Resizing Volumes

CloudPlatform provides the ability to resize data disks; CloudPlatform controls volume size by using disk offerings. This provides CloudPlatform administrators with the flexibility to choose how much space they want to make available to the end users. Volumes within the disk offerings with the same storage tag can be resized. For example, if you only want to offer 10, 50, and 100 GB offerings, the allowed resize should stay within those limits. That implies if you define a 10 GB, a 50 GB and a 100 GB disk offerings, a user can upgrade from 10 GB to 50 GB, or 50 GB to 100 GB. If you create a custom-sized disk offering, then you have the option to resize the volume by specifying a new, larger size.

Additionally, using the resizeVolume API, a data volume can be moved from a static disk offering to a custom disk offering with the size specified. This functionality allows those who might be billing by certain volume sizes or disk offerings to stick to that model, while providing the flexibility to migrate to whatever custom size necessary.

This feature is supported on XenServer, and VMware hosts. However, shrinking volumes is not supported on Hyper-V and KVM hosts. Consider the following table before you perform the operation.

Hypervisor	Storage Format	Support for Volume Shrink
XenServer	Offline Grow and Shrink  **Note** The supported storage repositories that can resize virtual disks are Logical Volume Management on Local Disk (lvm), Logical Volume Management over Fibre Channel or iSCSI Host Bus Adapter (lvmohba), and Logical Volume Management over iSCSI using software initiator (lvmoiscsi). The following storage repositories that use VHD format for storing virtual disks are not supported: VHD on Local Disk (ext), VHD on Network File System (nfs).	No
VMware	Online Grow	Yes
KVM	Not supported	No
Hyper-V	Not supported	No

Before you try to resize a volume, consider the following:

- The VMs associated with the volume are stopped.
- The data disks associated with the volume are removed.
- When a volume is shrunk, the disk associated with it is simply truncated, and doing so would put its content at risk of data loss. Therefore, resize any partitions or file systems before you shrink a data disk so that all the data is moved off from that disk.

To resize a volume:

1. Log in to the CloudPlatform UI as a user or admin.
2. In the left navigation bar, click Storage.
3. In Select View, choose Volumes.
4. Select the volume name in the Volumes list, then click the Resize Volume button 
5. In the Resize Volume pop-up, choose desired characteristics for the storage.

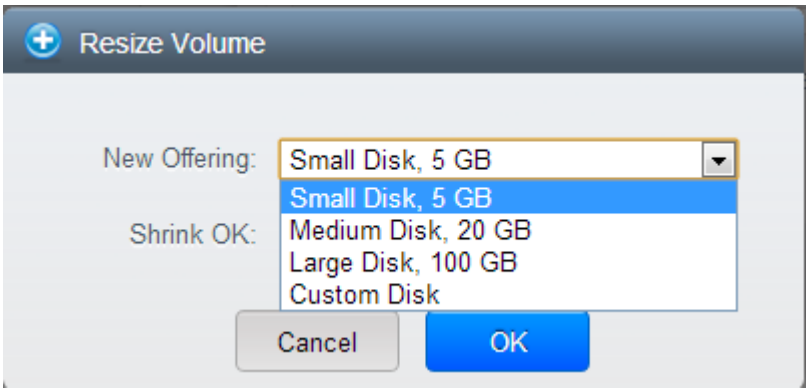

   a. If you select Custom Disk, specify a custom size.

   b. Click Shrink OK to confirm that you are reducing the size of a volume.

      This parameter protects against inadvertent shrinking of a disk, which might lead to the risk of data loss. You must sign off that you know what you are doing.

6. Click OK.

8.2.7. Reset VM to New Root Disk on Reboot

You can specify that you want to discard the root disk and create a new one whenever a given VM is rebooted. This is useful for secure environments that need a fresh start on every boot and for desktops that should not retain state. The IP address of the VM will not change due to this operation.

To enable root disk reset on VM reboot:

When creating a new service offering, set the parameter Volatile VM to True. VMs created from this service offering will have their disks reset upon reboot.

8.2.8. Volume Deletion and Garbage Collection

The deletion of a volume does not delete the snapshots that have been created from the volume

When a VM is destroyed, data disk volumes that are attached to the VM are not deleted.

Volumes are permanently destroyed using a garbage collection process. The global configuration variables expunge.delay and expunge.interval determine when the physical deletion of volumes will occur.

- expunge.delay: determines how old the volume must be before it is destroyed, in seconds
- expunge.interval: determines how often to run the garbage collection check

Administrators should adjust these values depending on site policies around data retention.

8.3. Creating Snapshot a Volume

1. Log in to the CloudPlatform UI as a user or administrator.
2. In the left navigation bar, click Storage.
3. In Select View, be sure Volumes is selected.
4. Click the name of the volume you want to snapshot.
5. Click the Snapshot button. 

Managing Networks and Traffic

For more conceptual information about networks and traffic, refer to **Chapter 8: Networking for Users** in the *Citrix CloudPlatform 4.5 (powered by Apache CloudStack) Concepts Guide*.

In a CloudPlatform, guest VMs can communicate with each other using shared infrastructure with the security and user perception that the guests have a private LAN. The CloudPlatform virtual router is the main component providing networking features for guest traffic.

9.1. Network Throttling in CloudPlatform

Network throttling is the process of controlling network access and bandwidth usage based on rules. CloudPlatform performs network throttling by setting network throttling parameter on NICs of virtual machines. This parameter is defined as the default data transfer rate in Mbps (Megabits Per Second).

9.1.1. Global Configuration for Network Throttling

You can configure network throttling rate parameter as part of the following:

- Service offering
- Network offering
- Global Configuration

If the value is set to '0', it indicates unlimited network throttling. If the value is set to 'NULL' in Service Offerings or Network Offerings, it indicates that the network throttling rate is dictated by the values set in the global configuration parameters. For Service Offerings, the value for the **`vm.network.throttling.rate`** parameter is considered. For Network Offerings, the value for the **`network.throttling.rate`** parameter is considered.

By default, network throttling rate for the Public, Storage, Control, and Management network offerings is set to 0 (unlimited). The default network throttling rate for Guest network offerings is set to NULL.

9.1.2. Network Throttling on Different Types of Virtual Machines

Network throttling is configured on the NICs associated with the virtual machines. The type of vitual machine decides the network throttling value that is to be configured on its NIC.

- A user VM will have one default network and many additional networks. The default network uses the network throttling rate from the Service Offering that you used. The additional networks use the network throttling rate from the Network Offerings that you used.
- The guest and the public networks of a VR use the network throttling rate from the guest virtual Network Offering. The Control and the Management networks of the VR use the network throttling rate from the corresponding System Network Offering.
- All NICs on SSVM or CPVM are configured with unlimited network throttling rates.

The following figure explains how the network throttling rate, which is configured on the VR's guest NIC and VM's NIC, affects the traffic flows from User VM to User VM and from User VM to the Internet:

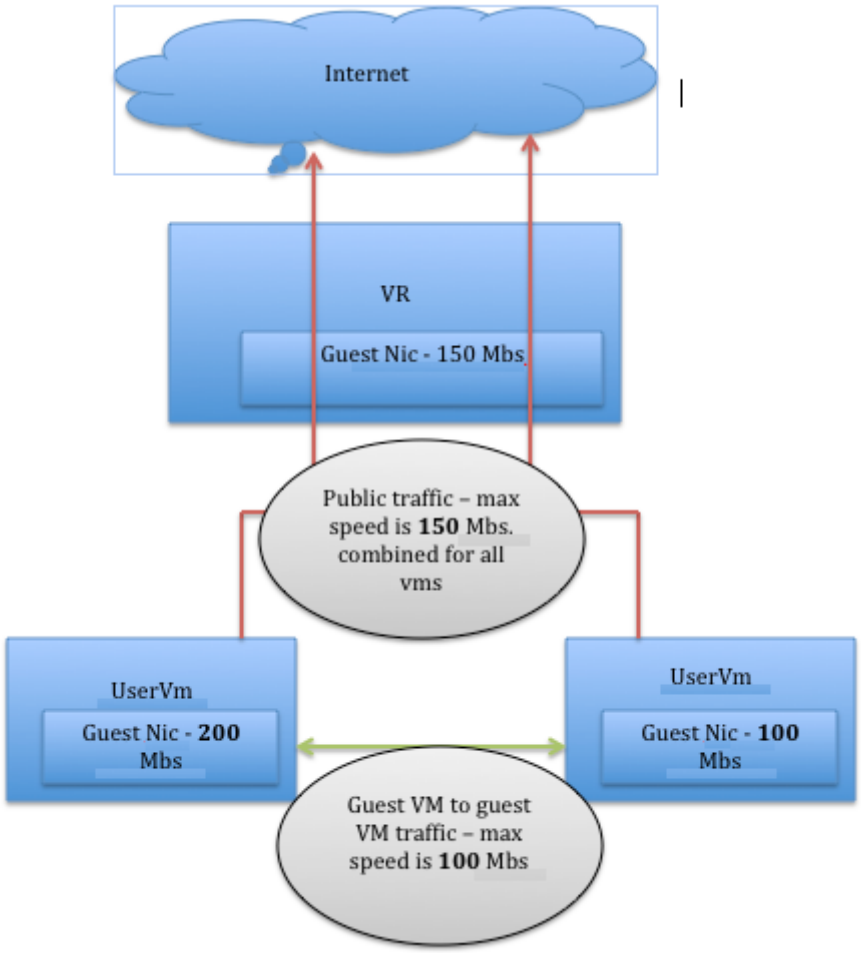

Note

The network throttling affects both the ingress (from the guest VMs to the Virtual Router) and the egress (from the Virtual Router to the Internet) communication.

For the hypervisors, network throttling is configured on a VM's NIC. The implementation of network throttling on a hypervisor depends on the type of hypervisor. In the following table, you can see the hypervisors that support network throttling in CloudPlatform and their implementation details:

Hypervisor	Implementation Details

XenServer	Network throttling is configured in a VIF object:
	`vifr.qosAlgorithmType = "ratelimit";`
	`vifr.qosAlgorithmParams = new HashMap<String, String>();`
	`// convert mbs to kilobyte per second`
	`vifr.qosAlgorithmParams.put("kbps", Integer.toString(nic.getNetworkRateMbps() * 128));`

Hypervisor	Implementation Details
VMWare ESXi	Network throttling is implemented through traffic shaping control at the portgroup level. All VM's NICs are connected to the same portgroup that share a traffic shaping policy. Following gives an example that demonstartes how the traffic shaping policy is created and used when creating the portgroup: `HostNetworkTrafficShapingPolicy shapingPolicy = null; if(networkRateMbps != null && networkRateMbps.intValue() > 0) { shapingPolicy = new HostNetworkTrafficShapingPolicy(); shapingPolicy.setEnabled(true); shapingPolicy.setAverageBandwidth((long) networkRateMbps.intValue()*1024L*1024L); // give 50% premium to peek shapingPolicy.setPeakBandwidth((long) (shapingPolicy.getAv erageBandwidth()*1.5)); // allow 5 seconds of burst transfer shapingPolicy.setBurstSize(5*shapingPolicy. getAverageBandwidth()/8); } // creating a port group with specified shaping policy hostMo.createPortGroup(vSwitch, networkName, vid, secPolicy, shapingPolicy);`
Microsoft Hyper-V	Network throttling is implemented by applying the network rate value on the VM's NIC Bandwidth Management property. The Bandwith Management Property allows you to configure minimum and maximum bandwidth values. When network rate is infinite, the Bandwith Management Property is disabled. That is, no limit for the bandwidth is configured. When the network rate value is configured, the minimum bandwidth is configured as "0" and the maximum bandwidth is configured to the value provided in Mbps metric. The following class is used to manage the bandwidth settings of VM's NIC: ROOT.virtualization.v2.Msvm_ EthernetSwitchPortBandwidthSettingData.cs The following method in Agent code configures the bandwidth or the throttling rate on the NIC by reading the **`network.rate`** value received from the agent command: SetBandWidthLimit(ulong limit, EthernetPortAllocationSettingData portPath)

9.2. Basic Zone Physical Network Configuration

In a basic network, configuring the physical network is fairly straightforward. You only need to configure one guest network to carry traffic that is generated by guest VMs. When you first add a zone to CloudPlatform, you set up the guest network through the Add Zone screens.

9.3. Advanced Zone Physical Network Configuration

Within a zone that uses advanced networking, you need to tell the Management Server how the physical network is set up to carry different kinds of traffic in isolation.

9.3.1. Configuring Isolated Guest Network

These steps assume you have already logged in to the CloudPlatform UI. To configure the base guest network:

1. In the left navigation, choose Infrastructure. On Zones, click View More, then click the zone to which you want to add a network.
2. Click the Network tab.
3. Click Add Isolated Guest Network.

   The Add Isolated Guest Network window is displayed:

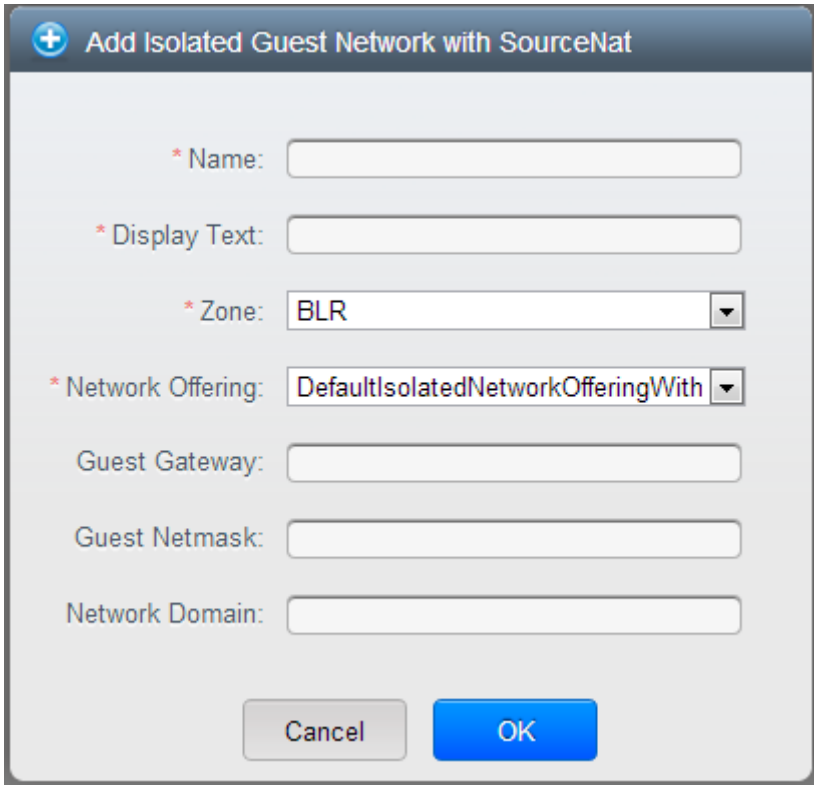

4. Provide the following information:
   - **Name**. The name of the network. This will be user-visible.
   - **Display Text**: The description of the network. This will be displayed to the user.
   - **Zone**: The zone in which you are configuring the guest network.

- **Network offering**: If the administrator has configured multiple network offerings, select the one you want to use for this network.
- **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.

5. Click OK.

9.3.2. Configure Public Traffic in an Advanced Zone

In a zone that uses advanced networking, you need to configure at least one range of IP addresses for Internet traffic.

9.3.3. Configuring a Shared Guest Network

1. Log in to the CloudPlatform UI as administrator.
2. In the left navigation, choose Infrastructure.
3. On Zones, click View More.
4. Click the zone to which you want to add a guest network.
5. Click the Physical Network tab.
6. Click the physical network you want to work with.
7. On the Guest node of the diagram, click Configure.
8. Click the Network tab.
9. Click Add guest network.

   The Add guest network window is displayed.

10. Specify the following:

    - **Name**: The name of the network. This will be visible to the user.
    - **Description**: The short description of the network that can be displayed to users.
    - **VLAN ID**: The unique ID of the VLAN.
    - **Isolated VLAN ID**: The unique ID of the Secondary Isolated VLAN.

      Applies only to a Private VLAN setup.

    - **Scope**: The available scopes are Domain, Account, Project, and All.
      - **Domain**: Selecting Domain limits the scope of this guest network to the domain you specify. The network will not be available for other domains. If you select Subdomain Access, the guest network is available to all the sub domains within the selected domain.

- **Account**: The account for which the guest network is being created for. You must specify the domain the account belongs to.
- **Project**: The project for which the guest network is being created for. You must specify the domain the project belongs to.
- **All**: The guest network is available for all the domains, account, projects within the selected zone.

- **Network Offering**: If the administrator has configured multiple network offerings, select the one you want to use for this network.
- **Gateway**: The gateway that the guests should use.
- **Netmask**: The netmask in use on the subnet the guests will use.
- **IP Range**: A range of IP addresses that are accessible from the Internet and are assigned to the guest VMs.
- **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.

11. Click OK to confirm.

9.4. Using Security Groups to Control Traffic to VMs

9.4.1. About Security Groups

Security groups provide a way to isolate traffic to VMs. A security group is a group of VMs that filter their incoming and outgoing traffic according to a set of rules, called ingress and egress rules. These rules filter network traffic according to the IP address that is attempting to communicate with the VM. Security groups are particularly useful in zones that use basic networking, because there is a single guest network for all guest VMs. In advanced zones, security groups are supported only on the KVM hypervisor.

Note

In a zone that uses advanced networking, you can instead define multiple guest networks to isolate traffic to VMs.

Each CloudPlatform account comes with a default security group that denies all inbound traffic and allows all outbound traffic. The default security group can be modified so that all new VMs inherit some other desired set of rules.

Any CloudPlatform user can set up any number of additional security groups. When a new VM is launched, it is assigned to the default security group unless another user-defined security group is specified. A VM can be a member of any number of security groups. Once a VM is assigned to a security group, it remains in that group for its entire lifetime; you can not move a running VM from one security group to another.

You can modify a security group by deleting or adding any number of ingress and egress rules. When you do, the new rules apply to all VMs in the group, whether running or stopped.

If no ingress rules are specified, then no traffic will be allowed in, except for responses to any traffic that has been allowed out through an egress rule.

9.4.2. Security Groups in Advanced Zones (KVM Only)

CloudPlatform provides the ability to use security groups to provide isolation between guests on a single shared, zone-wide network in an advanced zone where KVM is the hypervisor. Using security groups in advanced zones rather than multiple VLANs allows a greater range of options for setting up guest isolation in a cloud.

Limitation

Multiple VLAN ranges in a security group-enabled shared network are not supported.

Security groups must be enabled in the zone in order for this feature to be used.

9.4.3. Enabling Security Groups

In order for security groups to function in a zone, the security groups feature must first be enabled for the zone. The administrator can do this when creating a new zone, by selecting a network offering that includes security groups. The procedure is described in Zone Configuration in the Installation Guide. The administrator can not enable security groups for an existing zone, only when creating a new zone.

9.4.4. Adding a Security Group

A user or administrator can define a new security group.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network
3. In Select view, choose Security Groups.
4. Click Add Security Group.
5. Provide a name and description.
6. Click OK.

   The new security group appears in the Security Groups Details tab.
7. To make the security group useful, continue to Adding Ingress and Egress Rules to a Security Group.

9.4.5. Adding Ingress and Egress Rules to a Security Group

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network
3. In Select view, choose Security Groups, then click the security group you want .
4. To add an ingress rule, click the Ingress Rules tab and fill out the following fields to specify what network traffic is allowed into VM instances in this security group. If no ingress rules are specified,

then no traffic will be allowed in, except for responses to any traffic that has been allowed out through an egress rule.

- **Add by CIDR/Account**. Indicate whether the source of the traffic will be defined by IP address (CIDR) or an existing security group in a CloudPlatform account (Account). Choose Account if you want to allow incoming traffic from all VMs in another security group
- **Protocol**. The networking protocol that sources will use to send traffic to the security group. TCP and UDP are typically used for data exchange and end-user communications. ICMP is typically used to send error messages or network monitoring data.
- **Start Port, End Port**. (TCP, UDP only) A range of listening ports that are the destination for the incoming traffic. If you are opening a single port, use the same number in both fields.
- **ICMP Type, ICMP Code**. (ICMP only) The type of message and error code that will be accepted.
- **CIDR**. (Add by CIDR only) To accept only traffic from IP addresses within a particular address block, enter a CIDR or a comma-separated list of CIDRs. The CIDR is the base IP address of the incoming traffic. For example, 192.168.0.0/22. To allow all CIDRs, set to 0.0.0.0/0.
- **Account, Security Group**. (Add by Account only) To accept only traffic from another security group, enter the CloudPlatform account and name of a security group that has already been defined in that account. To allow traffic between VMs within the security group you are editing now, enter its name (that is, the same name you chose in step *3*).

The following example allows inbound HTTP access from anywhere:

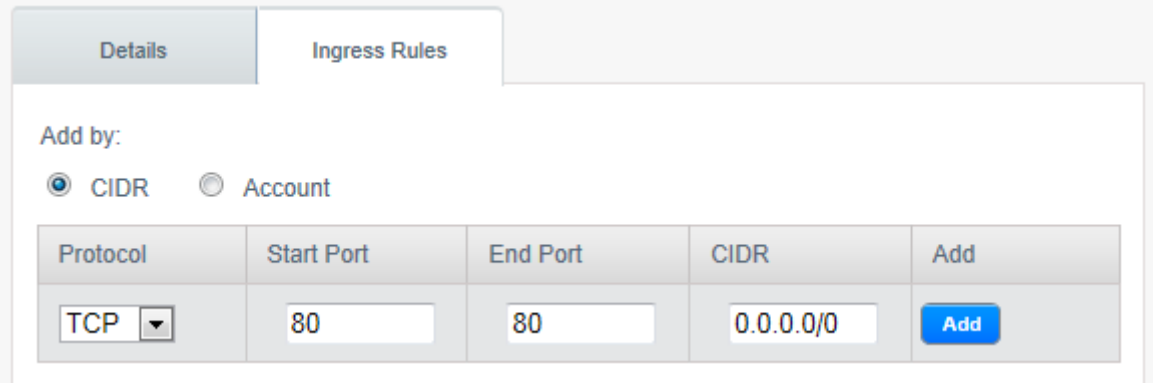

5. To add an egress rule, click the Egress Rules tab and fill out the following fields to specify what type of traffic is allowed to be sent out of VM instances in this security group. If no egress rules are specified, then all traffic will be allowed out. Once egress rules are specified, the following types of traffic are allowed out: traffic specified in egress rules; queries to DNS and DHCP servers; and responses to any traffic that has been allowed in through an ingress rule

   - **Add by CIDR/Account**. Indicate whether the destination of the traffic will be defined by IP address (CIDR) or an existing security group in a CloudPlatform account (Account). Choose Account if you want to allow outgoing traffic to all VMs in another security group.
   - **Protocol**. The networking protocol that VMs will use to send outgoing traffic. TCP and UDP are typically used for data exchange and end-user communications. ICMP is typically used to send error messages or network monitoring data.
   - **Start Port, End Port**. (TCP, UDP only) A range of listening ports that are the destination for the outgoing traffic. If you are opening a single port, use the same number in both fields.

- **ICMP Type, ICMP Code**. (ICMP only) The type of message and error code that will be sent
- **CIDR**. (Add by CIDR only) To send traffic only to IP addresses within a particular address block, enter a CIDR or a comma-separated list of CIDRs. The CIDR is the base IP address of the destination. For example, 192.168.0.0/22. To allow all CIDRs, set to 0.0.0.0/0.
- **Account, Security Group**. (Add by Account only) To allow traffic to be sent to another security group, enter the CloudPlatform account and name of a security group that has already been defined in that account. To allow traffic between VMs within the security group you are editing now, enter its name (that is, the same name you chose in step *3*).

6. Click Add.

9.5. External Firewalls and Load Balancers

CloudPlatform is capable of replacing its Virtual Router with an external Juniper SRX or Cisco ASA 1000v Cloud Firewall device and an optional external Citrix NetScaler or BigIP F5 load balancer for gateway and load balancing services. In this case, the VMs use the SRX or ASA as their gateway.

An external Juniper SRX or Cisco ASA can be used for:

- Source NAT
- Static NAT
- Firewall
- Port forwarding

A NetScaler or F5 can be used for:

- Load balancing

For details about installing and setting up these external network service providers, see the CloudPlatform Installation Guide.

9.5.1. Configuring SNMPCommunity String on a RHEL Server

The SNMP Community string is similar to a user id or password that provides access to a network device, such as router. This string is sent along with all SNMP requests. If the community string is correct, the device responds with the requested information. If the community string is incorrect, the device discards the request and does not respond.

The NetScaler device uses SNMP to communicate with the VMs. You must install SNMP and configure SNMP Community string for a secure communication between the NetScaler device and the RHEL machine.

1. Ensure that you installed SNMP on RedHat. If not, run the following command:

```
yum install net-snmp-utils
```

2. Edit the /etc/snmp/snmpd.conf file to allow the SNMP polling from the NetScaler device.

   a. Map the community name into a security name (local and mynetwork, depending on where the request is coming from):

Note

Use a strong password instead of public when you edit the following table.

```
#          sec.name    source          community
com2sec     local       localhost       public
com2sec    mynetwork    0.0.0.0         public
```

Note

Setting to 0.0.0.0 allows all IPs to poll the NetScaler server.

b. Map the security names into group names:

```
#      group.name    sec.model  sec.name
group    MyRWGroup      v1          local
group    MyRWGroup      v2c         local
group    MyROGroup      v1        mynetwork
group    MyROGroup      v2c       mynetwork
```

c. Create a view to allow the groups to have the permission to:

```
incl/excl subtree mask view all included .1
```

d. Grant access with different write permissions to the two groups to the view you created.

```
# context     sec.model     sec.level     prefix     read     write     notif
  access      MyROGroup ""  any noauth     exact      all      none      none
  access      MyRWGroup ""  any noauth     exact      all      all       all
```

3. Unblock SNMP in iptables. To do this, open the **iptables** file using vi editor and add the line **-A INPUT -p udp --dport 161 -j ACCEPT** in it. Save and restart the **iptables** file.

4. Start the SNMP service:

```
service snmpd start
```

5. Ensure that the SNMP service is started automatically during the system startup:

```
chkconfig snmpd on
```

9.5.2. Initial Setup of External Firewalls and Load Balancers

When the first VM is created for a new account, CloudPlatform programs the external firewall and load balancer to work with the VM. The following objects are created on the firewall:

- A new logical interface to connect to the account's private VLAN. The interface IP is always the first IP of the account's private subnet (e.g. 10.1.1.1).
- A source NAT rule that forwards all outgoing traffic from the account's private VLAN to the public Internet, using the account's public IP address as the source address
- A firewall filter counter that measures the number of bytes of outgoing traffic for the account

The following objects are created on the load balancer:

- A new VLAN that matches the account's provisioned Zone VLAN
- A self IP for the VLAN. This is always the second IP of the account's private subnet (e.g. 10.1.1.2).

9.5.3. Ongoing Configuration of External Firewalls and Load Balancers

Additional user actions (e.g. setting a port forward) will cause further programming of the firewall and load balancer. A user may request additional public IP addresses and forward traffic received at these IPs to specific VMs. This is accomplished by enabling static NAT for a public IP address, assigning the IP to a VM, and specifying a set of protocols and port ranges to open. When a static NAT rule is created, CloudPlatform programs the zone's external firewall with the following objects:

- A static NAT rule that maps the public IP address to the private IP address of a VM.
- A security policy that allows traffic within the set of protocols and port ranges that are specified.
- A firewall filter counter that measures the number of bytes of incoming traffic to the public IP.

The number of incoming and outgoing bytes through source NAT, static NAT, and load balancing rules is measured and saved on each external element. This data is collected on a regular basis and stored in the CloudPlatform database.

9.6. Load Balancer Rules

A CloudPlatform user or administrator may create rules that balance traffic received at a public IP address to one or more VMs. A load balancer rule is useful for distributing requests evenly among a pool of services. A service in this context means an application running on a virtual machine. The pool of services consists of multiple VMs running the same application. A user or cloud administrator creates a load balancer rule, specifies an algorithm, and assigns the rule to a set of VMs. Once the rule is in effect, each incoming request might be forwarded to any one of these redundant application instances, depending on the load balancing algorithm that has been specified in the rule.

Note

If you create load balancing rules while using a network service offering that includes an external load balancer device such as NetScaler, and later change the network service offering to one that uses the CloudPlatform virtual router, you must create a firewall rule on the virtual router for each of your existing load balancing rules so that they continue to function.

9.6.1. Adding a Load Balancer Rule

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. Click the name of the network where you want to load balance the traffic.
4. Click View IP Addresses.
5. Click the IP address for which you want to create the rule, then click the Configuration tab.
6. In the Load Balancing node of the diagram, click View All.

   In a Basic zone, you can also create a load balancing rule without acquiring or selecting an IP address. CloudPlatform internally assigns an IP when you create the load balancing rule, which is listed in the IP Addresses page when the rule is created. To do that, select the name of the network, then click the Add Load Balancer tab. Continue with 7.

7. Fill in the following:

   - **Name**: A name for the load balancer rule.
   - **Public Port**: The port receiving incoming traffic to be balanced.
   - **Private Port**: The port that the VMs will use to receive the traffic.
   - **Algorithm**: Choose the load balancing algorithm you want CloudPlatform to use. CloudPlatform supports a variety of well-known algorithms. If you are not familiar with these choices, you will find plenty of information about them on the Internet.
   - **Stickiness**: (Optional) Click Configure and choose the algorithm for the stickiness policy. See *Section 9.6.3, “Sticky Session Policies for Load Balancer Rules”*.
   - **AutoScale**: Click Configure and complete the AutoScale configuration as explained in *Section 9.6.2, “Configuring AutoScale”*.
   - **Health Check**: (Optional; NetScaler load balancers only) Click Configure and fill in the characteristics of the health check policy. See *Section 9.6.4, “Health Checks for Load Balancer Rules”*.
     - **Ping path (Optional)**: Specify the command which NetScaler should execute to check your web server. For example, GET /index.html. This is the destination which NetScaler LB service sends the health check queries to. Use a single forward slash in the ping path field so that LB

sends the query to your web server's default home page, whether that default page is named index.html, default.html, or a different name.

- **Response time (Optional)**: How long to wait for a response from the health check (2 - 60 seconds). Default: 5 seconds.
- **Interval time (Optional)**: Amount of time between health checks (1 second - 5 minutes). Default value is set in the global configuration parameter lbrule_health check_time_interval.
- **Healthy threshold (Optional)**: Number of consecutive health check successes that are required before declaring an instance healthy. Default: 2.
- **Unhealthy threshold (Optional)**: Number of consecutive health check failures that are required before declaring an instance unhealthy. Default: 10.

8. Click Add VMs, then select two or more VMs that will divide the load of incoming traffic, and click Apply.

   The new load balancer rule appears in the list.

9. You can repeat these steps to add more load balancer rules for this IP address.

9.6.2. Configuring AutoScale

AutoScaling allows you to scale your back-end services or application VMs up or down seamlessly and automatically according to the conditions you define. With AutoScaling enabled, you can ensure that the number of VMs you are using seamlessly scale up when demand increases, and automatically decreases when demand subsides. Thus it helps you save compute costs by terminating underused VMs automatically and launching new VMs when you need them, without the need for manual intervention.

NetScaler AutoScaling is designed to seamlessly launch or terminate VMs based on user-defined conditions. Conditions for triggering a scaleup or scaledown action can vary from a simple use case like monitoring the CPU usage of a server to a complex use case of monitoring a combination of server's responsiveness and its CPU usage. For example, you can configure AutoScaling to launch an additional VM whenever CPU usage exceeds 80 percent for 15 minutes, or to remove a VM whenever CPU usage is less than 20 percent for 30 minutes.

CloudPlatform uses the NetScaler load balancer to monitor all aspects of a system's health and work in unison with CloudPlatform to initiate scale-up or scale-down actions. The supported NetScaler versions are10.0 and beyond.

Prerequisites

Before you configure an AutoScale rule, consider the following:

- Ensure that the necessary template is prepared before configuring AutoScale. When a VM is deployed by using a template and when it comes up, the application should be up and running.

Note

If the application is not running, the NetScaler device considers the VM as ineffective and continues provisioning the VMs unconditionally until the resource limit is exhausted.

- Deploy the templates you prepared. Ensure that the applications come up on the first boot and is ready to take the traffic. Observe the time requires to deploy the template. Consider this time when you specify the quiet time while configuring AutoScale.
- The AutoScale feature supports the SNMP counters that can be used to define conditions for taking scale up or scale down actions. To monitor the SNMP-based counter, ensure that the SNMP agent is installed in the template used for creating the AutoScale VMs, and the SNMP operations work with the configured SNMP community and port by using standard SNMP managers. For example, see *Section 9.5.1, "Configuring SNMPCommunity String on a RHEL Server"* to configure SNMP on a RHEL machine.
- Ensure that the endpointe.url parameter present in the Global Settings is set to the Management Server API URL. For example, http://10.102.102.22:8080/client/api. In a multi-node Management Server deployment, use the virtual IP address configured in the load balancer for the management server's cluster. Additionally, ensure that the NetScaler device has access to this IP address to provide AutoScale support.

  If you update the endpointe.url, disable the AutoScale functionality of the load balancer rules in the system, then enable them back to reflect the changes. For more information see *Updating an AutoScale Configuration*
- If the API Key and Secret Key are regenerated for an AutoScale user, ensure that the AutoScale functionality of the load balancers that the user participates in are disabled and then enabled to reflect the configuration changes in the NetScaler.
- In an advanced Zone, ensure that at least one VM should be present before configuring a load balancer rule with AutoScale. Having one VM in the network ensures that the network is in implemented state for configuring AutoScale.

Configuration

Specify the following:

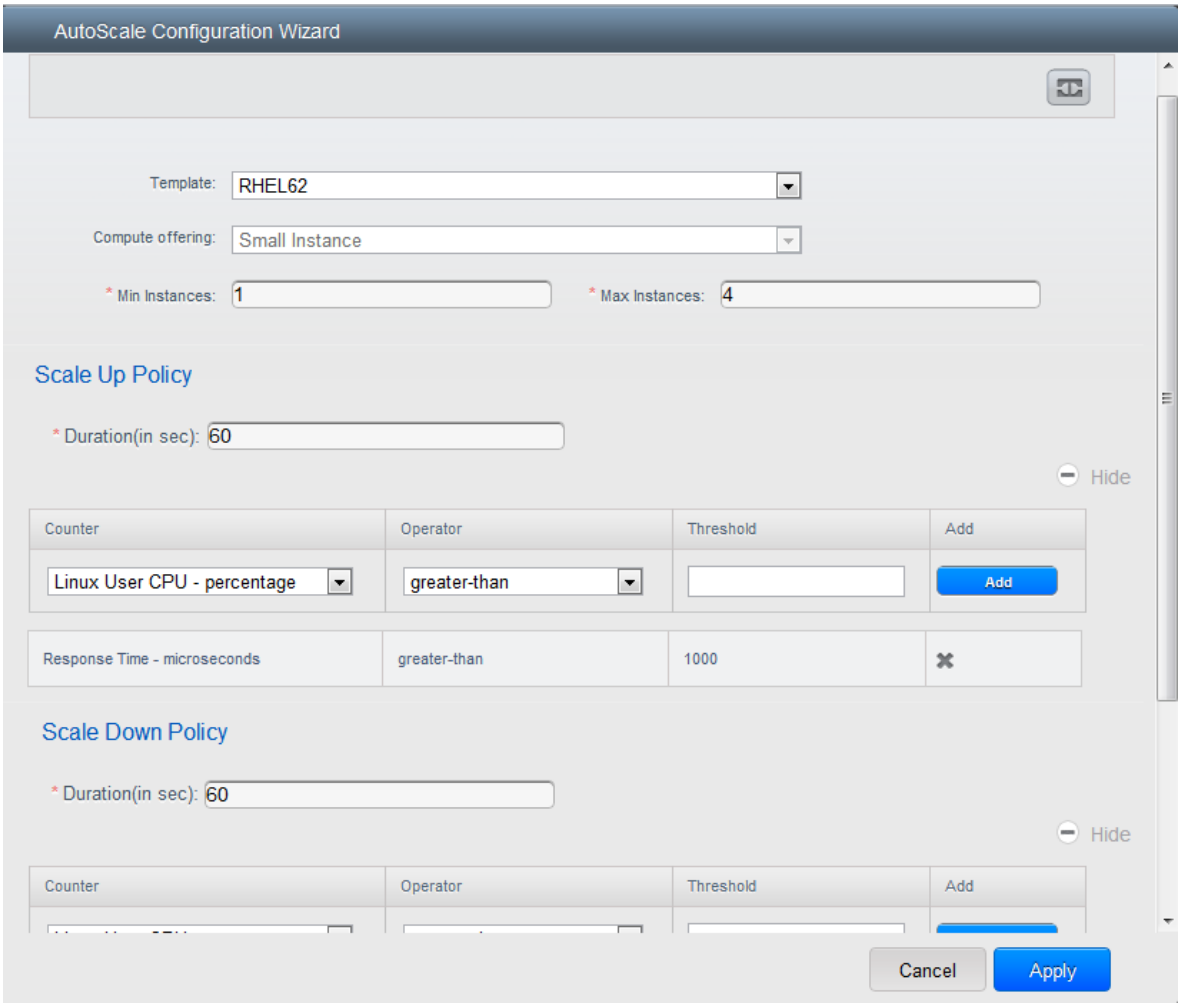

- **Template**: A template consists of a base OS image and application. A template is used to provision the new instance of an application on a scaleup action. When a VM is deployed from a template, the VM can start taking the traffic from the load balancer without any admin intervention. For example, if the VM is deployed for a Web service, it should have the Web server running, the database connected, and so on.
- **Compute offering**: A predefined set of virtual hardware attributes, including CPU speed, number of CPUs, and RAM size, that the user can select when creating a new virtual machine instance. Choose one of the compute offerings to be used while provisioning a VM instance as part of scaleup action.
- **Min Instance**: The minimum number of active VM instances that is assigned to a load balancing rule. The active VM instances are the application instances that are up and serving the traffic, and are being load balanced. This parameter ensures that a load balancing rule has at least the configured number of active VM instances are available to serve the traffic.

Note

If an application, such as SAP, running on a VM instance is down for some reason, the VM is then not counted as part of Min Instance parameter, and the AutoScale feature initiates a scaleup action if the number of active VM instances is below the configured value. Similarly, when an application instance comes up from its earlier down state, this application instance is counted as part of the active instance count and the AutoScale process initiates a scaledown action when the active instance count breaches the Max instance value.

- **Max Instance**: Maximum number of active VM instances that **should be assigned to** a load balancing rule. This parameter defines the upper limit of active VM instances that can be assigned to a load balancing rule.

  Specifying a large value for the maximum instance parameter might result in provisioning large number of VM instances, which in turn leads to a single load balancing rule exhausting the VM instances limit specified at the account or domain level.

Note

If an application, such as SAP, running on a VM instance is down for some reason, the VM is not counted as part of Max Instance parameter. So there may be scenarios where the number of VMs provisioned for a scaleup action might be more than the configured Max Instance value. Once the application instances in the VMs are up from an earlier down state, the AutoScale feature starts aligning to the configured Max Instance value.

Specify the following scale-up and scale-down policies:

- **Duration**: The duration, in seconds, in which the conditions that you specify for a scaleup action to **be true to trigger a scaleup action**. The conditions defined should hold true for the entire duration you specify for an AutoScale action to be invoked.
- **Counter**: The performance counters expose the state of the monitored instances. By default, CloudPlatform offers four performance counters: Three SNMP counters and one NetScaler counter. The SNMP counters are Linux User CPU, Linux System CPU, and Linux CPU Idle. The NetScaler counter is ResponseTime. The root administrator can add additional counters into CloudPlatform by using the CloudStack API.
- **Operator**: The following five relational operators are supported in AutoScale feature: Greater than, Less than, Less than or equal to, Greater than or equal to, and Equal to.
- **Threshold**: Threshold value to be used for the counter. Once the counter defined above breaches the threshold value, the AutoScale feature initiates a scaleup or scaledown action.
- **Add**: Click Add to add the condition.

Additionally, if you want to configure the advanced settings, click Show advanced settings, and specify the following:

- **Polling interval**: Frequency in which the conditions, combination of counter, operator and threshold, are to be evaluated before taking a scale up or down action. The default polling interval is 30 seconds.
- **Quiet Time**: This is the cool down period after an AutoScale action is initiated. The time includes the time taken to complete provisioning a VM instance from its template and the time taken by an application to be ready to serve traffic. This quiet time allows the fleet to come up to a stable state before any action can take place. The default is 300 seconds.
- **Destroy VM Grace Period**: The duration in seconds, after a scaledown action is initiated, to wait before the VM is destroyed as part of scaledown action. This is to ensure graceful close of any pending sessions or transactions being served by the VM marked for destroy. The default is 120 seconds.
- **Security Groups**: (Enabled only for Basic zones.) Security groups provide a way to isolate traffic to the VM instances. A security group is a group of VMs that filter their incoming and outgoing traffic according to a set of rules, called ingress and egress rules. These rules filter network traffic according to the IP address that is attempting to communicate with the VM.
- **Disk Offerings**: A predefined set of disk size for primary data storage.
- **SNMP Community**: The SNMP community string to be used by the NetScaler device to query the configured counter value from the provisioned VM instances. Default is public.
- **SNMP Port**: The port number on which the SNMP agent that run on the provisioned VMs is listening. Default port is 161.
- **User**: This is the user that the NetScaler device use to invoke scaleup and scaledown API calls to the cloud. If no option is specified, the user who configures AutoScaling is applied. Specify another user name to override.
- **Apply**: Click Apply to create the AutoScale configuration.

Disabling and Enabling an AutoScale Configuration

If you want to perform any maintenance operation on the AutoScale VM instances, disable the AutoScale configuration. When the AutoScale configuration is disabled, no scaleup or scaledown action is performed. You can use this downtime for the maintenance activities. To disable the AutoScale configuration, click the Disable AutoScale  button.

The button toggles between enable and disable, depending on whether AutoScale is currently enabled or not. After the maintenance operations are done, you can enable the AutoScale configuration back.

To enable, open the AutoScale configuration page again, then click the Enable AutoScale  button.

Updating an AutoScale Configuration

You can update the various parameters and add or delete the conditions in a scaleup or scaledown rule. Before you update an AutoScale configuration, ensure that you disable the AutoScale load balancer rule by clicking the Disable AutoScale button.

After you modify the required AutoScale parameters, click Apply. To apply the new AutoScale policies, open the AutoScale configuration page again, then click the Enable AutoScale button.

Runtime Considerations

- An administrator should not assign a VM to a load balancing rule which is configured for AutoScale.
- Before a VM provisioning is completed if NetScaler is shutdown or restarted, the provisioned VM cannot be a part of the load balancing rule though the intent was to assign it to a load balancing rule. To workaround, rename the AutoScale provisioned VMs based on the rule name or ID so at any point of time the VMs can be reconciled to its load balancing rule.
- Making API calls outside the context of AutoScale, such as destroyVM, on an autoscaled VM leaves the load balancing configuration in an inconsistent state. Though VM is destroyed from the load balancer rule, NetScaler continues to show the VM as a service assigned to a rule.

9.6.3. Sticky Session Policies for Load Balancer Rules

Sticky sessions are used in Web-based applications to ensure continued availability of information across the multiple requests in a user's session. For example, if a shopper is filling a cart, you need to remember what has been purchased so far. The concept of “stickiness” is also referred to as persistence or maintaining state.

Any load balancer rule defined in CloudPlatform can have a stickiness policy. The policy consists of a name, stickiness method, and parameters. The parameters are name-value pairs or flags, which are defined by the load balancer vendor. The stickiness method could be load balancer-generated cookie, application-generated cookie, or source-based. In the source-based method, the source IP address is used to identify the user and locate the user’s stored data. In the other methods, cookies are used. The cookie generated by the load balancer or application is included in request and response URLs to create persistence. The cookie name can be specified by the administrator or automatically generated. A variety of options are provided to control the exact behavior of cookies, such as how they are generated and whether they are cached.

For the most up to date list of available stickiness methods, see the CloudPlatform UI or call listNetworks and check the SupportedStickinessMethods capability.

For details on how to set a stickiness policy using the UI, see *Section 9.6.1, “Adding a Load Balancer Rule”*.

9.6.4. Health Checks for Load Balancer Rules

(NetScaler load balancer only; requires NetScaler version 10.0 and beyond)

Health checks are used in load balanced applications to ensure that requests are forwarded only to running, available services. When creating a load balancer rule, you can specify a health check policy. This is in addition to specifying the stickiness policy, algorithm, and other load balancer rule options. You can configure one health check policy per load balancer rule.

Any load balancer rule defined on a NetScaler load balancer in CloudPlatform can have a health check policy. The policy consists of a ping path, thresholds to define "healthy" and "unhealthy" states, health check frequency, and timeout wait interval.

When a health check policy is in effect, the load balancer will stop forwarding requests to any resources that are found to be unhealthy. If the resource later becomes available again, the periodic health check will discover it, and the resource will once again be added to the pool of resources that

can receive requests from the load balancer. At any given time, the most recent result of the health check is displayed in the UI. For any VM that is attached to a load balancer rule with a health check configured, the state will be shown as UP or DOWN in the UI depending on the result of the most recent health check.

You can delete or modify existing health check policies.

To configure how often the health check is performed by default, use the global configuration setting healthcheck.update.interval (default value is 600 seconds). You can override this value for an individual health check policy.

For details on how to set a health check policy using the UI, see *Section 9.6.1, "Adding a Load Balancer Rule"*.

9.7. Global Server Load Balancing

CloudPlatform supports Global Server Load Balancing (GSLB) functionalities to provide business continuity by load balancing traffic to an instance on active zones only in case of zone failures. CloudPlatform achieves this by extending its functionality of integrating with NetScaler Application Delivery Controller (ADC), which also provides various GSLB capabilities, such as disaster recovery and load balancing. The DNS redirection technique is used to achieve GSLB in CloudPlatform.

In order to support this functionality, region level services and service provider are introduced. A new service 'GSLB' is introduced as a region level service. The GSLB service provider is introduced that will provider the GSLB service. Currently, NetScaler is the supported GSLB provider in CloudPlatform. GSLB functionality works in an Active-Active data center environment.

9.7.1. Configuring GSLB

A GSLB deployment is the logical collection of GSLB virtual server, GSLB service, LB virtual server, service, domain, and ADNS service. To create a GSLB site, you must configure load balancing in the zone. You must create GSLB servers and GSLB services for each site. You must bind GSLB services to GSLB servers. You must then create an ADNS service that provides the IP address of the best performing site to the client's request. A GSLB server is an entity that performs load balancing for the domains bound to it by returning the IP address of the best GSLB service. A GSLB service is a representation of the load balancing/content switching server. An LB server load balances incoming traffic by identifying the best server, then directs traffic to the corresponding service. It can also load-balance external DNS name servers. Services are entities that represent the servers. The domain is the domain name for which the system is the authoritative DNS server. By creating an ADNS service, the system can be configured as an authoritative DNS server.

To configure GSLB in your cloud environment, as a cloud administrator you must first configure a standard load balancing setup for each zone. This enables to balance load across different servers in each zone in the region. Then, configure both the NetScaler appliances that you plan to add to each zone as authoritative DNS (ADNS) servers.

Next, as a domain administrator or user, create a GSLB site for each zone, configure GSLB virtual servers for each site, create GLSB services, and bind the GSLB services to the GSLB virtual servers. Finally, bind the domain to the GSLB virtual servers. The GSLB configurations on the two appliances at the two different sites are identical, although each sites load-balancing configuration is specific to that site.

As per the example given above, the administrator of xyztelco is the one who sets up GSLB. Perform steps *1* through *b* as a cloud administrator. As a domain administrator or user when you create a GSLB rule and assign load balancer rules on the CloudPlatform side, CloudPlatform orchestrates what is given in *c* through *g*.

1. For the **`cloud.dns.name`** parameter in **Global Settings**, specify the DNS name of your cloud that tenants will use when creating GSLB rules. For example, **`xyztelco.com`**.
2. On the NetScaler side, configure GSLB as given in *Configuring Global Server Load Balancing (GSLB)*[1]:
   a. Configure a standard load balancing setup.
   b. Configure Authoritative DNS, as explained in *Configuring an Authoritative DNS Service*[2].
   c. Configure a GSLB site with site name formed from the domain name details.

      Configure a GSLB site with the site name formed from the domain name.

      As per the example given above, the site names are A.xyztelco.com and B.xyztelco.com.

      For more information, see *Configuring a Basic GSLB Site*[3].
   d. Configure a GSLB virtual server.

      For more information, see *Configuring a GSLB Virtual Server*[4].
   e. Configure a GSLB service for each virtual server.

      For more information, see *Configuring a GSLB Service*[5].
   f. Bind the GSLB services to the GSLB virtual server.

      For more information, see *Binding GSLB Services to a GSLB Virtual Server*[6].
   g. Bind domain name to GSLB virtual server. Domain name is obtained from the domain details.

      For more information, see *Binding a Domain to a GSLB Virtual Server*[7].
3. In each zone that are participating in GSLB, add GSLB-enabled NetScaler device.

   For more information, see *Section 9.7.1.2, "Enabling GSLB in NetScaler"*.

On CloudPlatform side, perform the following as a domain administrator or user:

1. Add a GSLB rule on both the sites.

   See *Section 9.7.1.3, "Adding a GSLB Rule"*.
2. Assign load balancer rules.

   See *Section 9.7.1.4, "Assigning Load Balancing Rules to GSLB"*.

9.7.1.1. Prerequisites and Guidelines

- The GSLB functionality is supported both Basic and Advanced zones.

[1] http://support.citrix.com/proddocs/topic/netscaler-traffic-management-10-map/ns-gslb-config-con.html
[2] http://support.citrix.com/proddocs/topic/netscaler-traffic-management-10-map/ns-gslb-config-adns-svc-tsk.html
[3] http://support.citrix.com/proddocs/topic/netscaler-traffic-management-10-map/ns-gslb-config-basic-site-tsk.html
[4] http://support.citrix.com/proddocs/topic/netscaler-traffic-management-10-map/ns-gslb-config-vsvr-tsk.html
[5] http://support.citrix.com/proddocs/topic/netscaler-traffic-management-10-map/ns-gslb-config-svc-tsk.html
[6] http://support.citrix.com/proddocs/topic/netscaler-traffic-management-10-map/ns-gslb-bind-svc-vsvr-tsk.html
[7] http://support.citrix.com/proddocs/topic/netscaler-traffic-management-10-map/ns-gslb-bind-dom-vsvr-tsk.html

- GSLB is added as a new network service.
- GSLB service provider can be added to a physical network in a zone.
- When users have VMs deployed in multiple availability zones which are GSLB enabled, they can use the GSLB functionality to load balance traffic across the VMs in multiple zones.
- The users can use GSLB to load balance across the VMs across zones in a region only if the admin has enabled GSLB in that region.
- The users can load balance traffic across the availability zones in the same region or different regions.
- The admin can configure DNS name for the entire cloud.
- The users can specify an unique name across the cloud for a globally load balanced service. The provided name is used as the domain name under the DNS name associated with the cloud.

  The user-provided name along with the admin-provided DNS name is used to produce a globally resolvable FQDN for the globally load balanced service of the user. For example, if the admin has configured xyztelco.com as the DNS name for the cloud, and user specifies 'foo' for the GSLB virtual service, then the FQDN name of the GSLB virtual service is foo.xyztelco.com.
- While setting up GSLB, users can select a load balancing method, such as round robin, for using across the zones that are part of GSLB.
- The user shall be able to set weight to zone-level virtual server. Weight shall be considered by the load balancing method for distributing the traffic.
- The GSLB functionality shall support session persistence, where series of client requests for particular domain name is sent to a virtual server on the same zone.

  Statistics is collected from each GSLB virtual server.

9.7.1.2. Enabling GSLB in NetScaler

In each zone, add GSLB-enabled NetScaler device for load balancing.

1. Log in as administrator to the CloudPlatform UI.
2. In the left navigation bar, click Infrastructure.
3. In Zones, click View More.
4. Choose the zone you want to work with.
5. Click the Physical Network tab, then click the name of the physical network.
6. In the Network Service Providers node of the diagram, click Configure.

   You might have to scroll down to see this.
7. Click NetScaler.
8. Click Add NetScaler device and provide the following:

   For NetScaler:

   - **IP Address**: The IP address of the NetScaler appliance.

- **Username/Password**: The authentication credentials to access the device. CloudPlatform uses these credentials to access the device.
- **Type**: The type of device that is being added. It could be F5 Big IP Load Balancer, NetScaler VPX, NetScaler MPX, or NetScaler SDX. For a comparison of the NetScaler types, see the CloudPlatform Administration Guide.
- **Public interface**: Interface of device that is configured to be part of the public network.
- **Private interface**: Interface of device that is configured to be part of the private network.
- **GSLB service**: Select this option.
- **GSLB service Public IP**: The public IP address of the NAT translator for a GSLB service that is on a private network.
- **GSLB service Private IP**: The private IP of the GSLB service.
- **Number of Retries**. Number of times to attempt a command on the device before considering the operation failed. Default is 2.
- **Capacity**: The number of networks the device can handle.
- **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.

9. Click OK.

9.7.1.3. Adding a GSLB Rule

1. Log in to the CloudPlatform UI as a domain administrator or user.
2. In the left navigation pane, click Region.
3. Select the region for which you want to create a GSLB rule.
4. In the Details tab, click View GSLB.
5. Click Add GSLB.

   The Add GSLB page is displayed as follows:

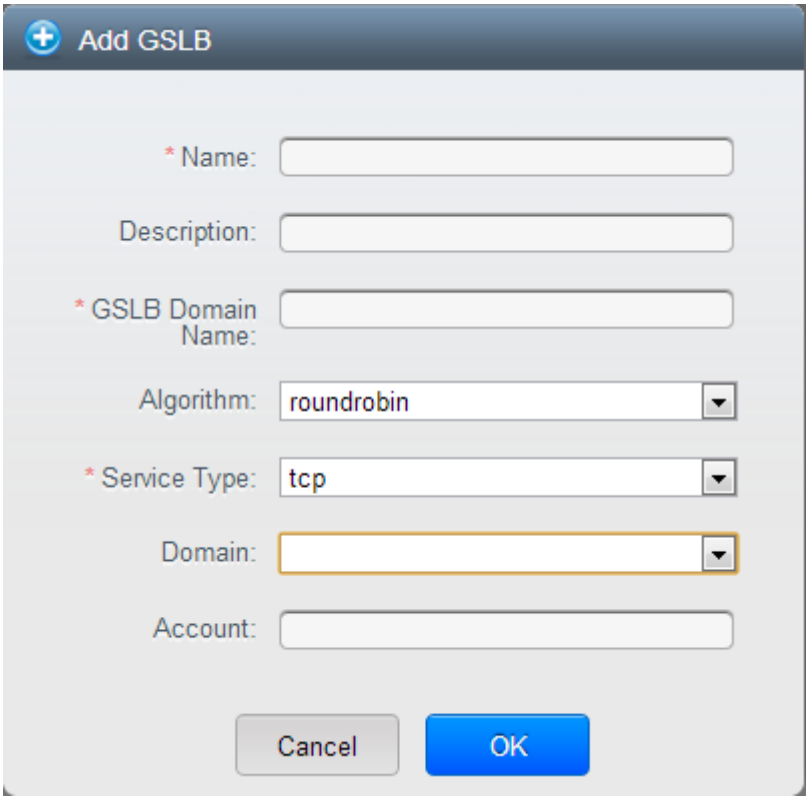

6. Specify the following:
   - **Name**: Name for the GSLB rule.
   - **Description**: (Optional) A short description of the GSLB rule that can be displayed to users.
   - **GSLB Domain Name**: A preferred domain name for the service.
   - **Algorithm**: (Optional) The algorithm to use to load balance the traffic across the zones. The options are Round Robin, Least Connection, and Proximity.
   - **Service Type**: The transport protocol to use for GSLB. The options are TCP and UDP.
   - **Domain**: (Optional) The domain for which you want to create the GSLB rule.
   - **Account**: (Optional) The account on which you want to apply the GSLB rule.
7. Click OK to confirm.

9.7.1.4. Assigning Load Balancing Rules to GSLB

1. Log in to the CloudPlatform UI as a domain administrator or user.
2. In the left navigation pane, click Region.
3. Select the region for which you want to create a GSLB rule.
4. In the Details tab, click View GSLB.
5. Select the desired GSLB.
6. Click view assigned load balancing.

7. Click assign more load balancing.
8. Select the load balancing rule you have created for the zone.
9. Click OK to confirm.

9.8. Using Multiple Guest Networks

In zones that use advanced networking, additional networks for guest traffic may be added at any time after the initial installation. You can also customize the domain name associated with the network by specifying a DNS suffix for each network.

A VM's networks are defined at VM creation time. A VM cannot add or remove networks after it has been created, although the user can go into the guest and remove the IP address from the NIC on a particular network.

Each VM has just one default network. The virtual router's DHCP reply will set the guest's default gateway as that for the default network. Multiple non-default networks may be added to a guest in addition to the single, required default network. The administrator can control which networks are available as the default network.

Additional networks can either be available to all accounts or be assigned to a specific account. Networks that are available to all accounts are zone-wide. Any user with access to the zone can create a VM with access to that network. These zone-wide networks provide little or no isolation between guests.Networks that are assigned to a specific account provide strong isolation.

9.8.1. Adding an Additional Guest Network

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. Click Add guest network. 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 user-visible.
   - **Zone**. The name of the zone this network applies to. Each zone is a broadcast domain, and therefore each zone has a different IP range for the guest network. The administrator must configure the IP range for each zone.
   - **Network offering**: If the administrator has configured multiple network offerings, select the one you want to use for this network.
   - **Guest Gateway**: The gateway that the guests should use.
   - **Guest Netmask**: The netmask in use on the subnet the guests will use.
4. Click Create.

9.8.2. Reconfiguring Networks in VMs

CloudPlatform provides you the ability to move VMs between networks and reconfigure a VM's network. You can remove a VM from a network and add to a new network. You can also change the default network of a virtual machine. With this functionality, hybrid or traditional server loads can be accommodated with ease.

This feature is supported on XenServer, VMware, and KVM hypervisors.

9.8.2.1. Prerequisites

For adding or removing networks to work, ensure that vm-tools are running on the guest VMs on VMware host.

9.8.2.2. Adding a Network

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, click Instances.
3. Choose the VM that you want to work with.
4. Click the NICs tab.
5. Click Add network to VM.

   The Add network to VM dialog is displayed.
6. In the drop-down list, select the network that you would like to add this VM to.

   A new NIC is added for this network. You can view the following details in the NICs page:

   - ID
   - Network Name
   - Type
   - IP Address
   - Gateway
   - Netmask
   - Is default

9.8.2.3. Removing a Network

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, click Instances.
3. Choose the VM that you want to work with.
4. Click the NICs tab.
5. Locate the NIC you want to remove.
6. Click Remove NIC button. 
7. Click Yes to confirm.

9.8.2.4. Selecting the Default Network

1. Log in to the CloudPlatform UI as an administrator or end user.

2. In the left navigation, click Instances.
3. Choose the VM that you want to work with.
4. Click the NICs tab.
5. Locate the NIC you want to work with.
6. Click the Set default NIC button. 
7. Click Yes to confirm.

9.9. Guest IP Ranges

The IP ranges for guest network traffic are set on a per-account basis by the user. This allows the users to configure their network in a fashion that will enable VPN linking between their guest network and their clients.

In shared networks in Basic zone and Security Group-enabled Advanced networks, you will have the flexibility to add multiple guest IP ranges from different subnets. You can add or remove one IP range at a time. For more information, see *Section 9.15, "Multiple Subnets in Shared Network"*.

9.10. Acquiring a New IP Address

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. Click the name of the network where you want to work with.
4. Click View IP Addresses.
5. Click Acquire New IP, and click Yes in the confirmation dialog.

   You are prompted for confirmation because, typically, IP addresses are a limited resource. Within a few moments, the new IP address should appear with the state Allocated. You can now use the IP address in port forwarding or static NAT rules.

9.11. Releasing an IP Address

When the last rule for an IP address is removed, you can release that IP address. The IP address still belongs to the VPC; however, it can be picked up for any guest network again.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. Click the name of the network where you want to work with.
4. Click View IP Addresses.
5. Click the IP address you want to release.
6. Click the Release IP button. 

9.12. Reserving Public IP Addresses and VLANs for Accounts

CloudPlatform provides you the ability to reserve a set of public IP addresses and VLANs exclusively for an account. During zone creation, you can continue defining a set of VLANs and multiple public IP ranges. This feature extends the functionality to enable you to dedicate a fixed set of VLANs and guest IP addresses for a tenant.

Note that if an account has consumed all the VLANs and IPs dedicated to it, the account can acquire two more resources from the system. CloudPlatform provides the root admin with two configuration parameter to modify this default behavior—use.system.public.ips and use.system.guest.vlans. These global parameters enable the root admin to disallow an account from acquiring public IPs and guest VLANs from the system, if the account has dedicated resources and these dedicated resources have all been consumed. Both these configurations are configurable at the account level.

This feature provides you the following capabilities:

- Reserve a VLAN range and public IP address range from an Advanced zone and assign it to an account
- Disassociate a VLAN and public IP address range from an account
- View the number of public IP addresses allocated to an account
- Check whether the required range is available and is conforms to account limits.

  The maximum IPs per account limit cannot be superseded.

9.12.1. Dedicating IP Address Ranges to an Account

1. Log in to the CloudPlatform UI as administrator.
2. In the left navigation bar, click Infrastructure.
3. In Zones, click View All.
4. Choose the zone you want to work with.
5. Click the Physical Network tab.
6. In the Public node of the diagram, click Configure.
7. Click the IP Ranges tab.

   You can either assign an existing IP range to an account, or create a new IP range and assign to an account.
8. To assign an existing IP range to an account, perform the following:
   a. Locate the IP range you want to work with.
   b. Click Add Account  button.

      The Add Account dialog is displayed.
   c. Specify the following:
      - **Account**: The account to which you want to assign the IP address range.

- **Domain**: The domain associated with the account.

To create a new IP range and assign an account, perform the following:

a. Specify the following:

- **Gateway**
- **Netmask**
- **VLAN**
- **Start IP**
- **End IP**
- **Account**: Perform the following:

i. Click Account.

The Add Account page is displayed.

ii. Specify the following:

- **Account**: The account to which you want to assign an IP address range.
- **Domain**: The domain associated with the account.

iii. Click OK.

b. Click Add.

9.12.2. Dedicating VLAN Ranges to an Account

1. After the CloudPlatform Management Server is installed, log in to the CloudPlatform UI as administrator.
2. In the left navigation bar, click Infrastructure.
3. In Zones, click View All.
4. Choose the zone you want to work with.
5. Click the Physical Network tab.
6. In the Guest node of the diagram, click Configure.
7. Select the Dedicated VLAN Ranges tab.
8. Click Dedicate VLAN Range.

   The Dedicate VLAN Range dialog is displayed.
9. Specify the following:

   - **VLAN Range**: The VLAN range that you want to assign to an account.
   - **Account**: The account to which you want to assign the selected VLAN range.

- **Domain**: The domain associated with the account.

9.13. IP Reservation in Isolated Guest Networks

In isolated guest networks, a part of the guest IP address space can be reserved for non-CloudPlatform VMs or physical servers. To do so, you configure a range of Reserved IP addresses by specifying the CIDR when a guest network is in Implemented state. If your customers wish to have non-CloudPlatform controlled VMs or physical servers on the same network, they can share a part of the IP address space that is primarily provided to the guest network.

In an Advanced zone, an IP address range or a CIDR is assigned to a network when the network is defined. The CloudPlatform virtual router acts as the DHCP server and uses CIDR for assigning IP addresses to the guest VMs. If you decide to reserve CIDR for non-CloudPlatform purposes, you can specify a part of the IP address range or the CIDR that should only be allocated by the DHCP service of the virtual router to the guest VMs created in CloudPlatform. The remaining IPs in that network are called Reserved IP Range. When IP reservation is configured, the administrator can add additional VMs or physical servers that are not part of CloudPlatform to the same network and assign them the Reserved IP addresses. CloudPlatform guest VMs cannot acquire IPs from the Reserved IP Range.

9.13.1. IP Reservation Considerations

Consider the following before you reserve an IP range for non-CloudPlatform machines:

- IP Reservation is supported only in Isolated networks and VPC tiers.
- IP Reservation can be applied only when the network is in Implemented state.
- No IP Reservation is done by default.
- Guest VM CIDR you specify must be a subset of the network CIDR.
- Specify a valid Guest VM CIDR. IP Reservation is applied only if no active IPs exist outside the Guest VM CIDR.

  You cannot apply IP Reservation if any VM is alloted with an IP address that is outside the Guest VM CIDR.
- To reset an existing IP Reservation, apply IP reservation by specifying the value of network CIDR in the CIDR field.

  For example, the following table describes three scenarios of guest network creation:

Case	CIDR	Network CIDR	Reserved IP Range for Non-CloudPlatform VMs	Description
1	10.1.1.0/24	None	None	No IP Reservation.
2	10.1.1.0/26	10.1.1.0/24	10.1.1.64 to 10.1.1.254	IP Reservation configured by the UpdateNetwork API with guestvmcidr= 10.1.1.0/26 or

Case	CIDR	Network CIDR	Reserved IP Range for Non-CloudPlatform VMs	Description
				enter 10.1.1.0/26 in the CIDR field in the UI.
3	10.1.1.0/24	None	None	Removing IP Reservation by the UpdateNetwork API with guestvmcidr= 10.1.1.0/24 or enter 10.1.1.0/24 in the CIDR field in the UI.

9.13.2. Limitations

- The IP Reservation is not supported if active IPs that are found outside the Guest VM CIDR.
- Upgrading network offering which causes a change in CIDR (such as upgrading an offering with no external devices to one with external devices) IP Reservation becomes void if any. Reconfigure IP Reservation in the new re-implemeted network.

9.13.3. Best Practices

Apply IP Reservation to the guest network as soon as the network state changes to "Implemented". If you apply reservation soon after the first guest VM is deployed, lesser conflicts occurs while applying reservation.

You can use persistent network offering to create your network. When you do so, your network will be in the "Implemented" state and no VMs will have deployed in the network.

9.13.4. Reserving an IP Range

1. Log-in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation bar, click **Network**.
3. On the right side panel, in the table that lists networks, click the name of the network you want to modify.
4. Under the **Details** tab, click the **Edit** icon.

   The CIDR field changes to editable one.
5. In the **CIDR** field, specify the Guest VM CIDR.
6. Click **Apply**.

   Wait for the update to complete. The Network CIDR and the Reserved IP Range are displayed on the **Details** page.

9.14. Configuring Multiple IP Addresses on a Single NIC

CloudPlatform provides you the ability to associate multiple private IP addresses per guest VM NIC. In addition to the primary IP, you can assign additional IPs to the guest VM NIC. This feature is supported on all the network configurations—Basic, Advanced, and VPC. Security Groups, Static NAT and Port forwarding services are supported on these additional IPs.

As always, you can specify an IP from the guest subnet; if not specified, an IP is automatically picked up from the guest VM subnet. You can view the IPs associated with for each guest VM NICs on the UI. You can apply NAT on these additional guest IPs by using network configuration option in the CloudPlatform UI. You must specify the NIC to which the IP should be associated.

This feature is supported on XenServer, KVM, and VMware hypervisors. Note that Basic zone security groups are not supported on VMware.

9.14.1. Use Cases

Some of the use cases are described below:

- Network devices, such as firewalls and load balancers, generally work best when they have access to multiple IP addresses on the network interface.
- Moving private IP addresses between interfaces or instances. Applications that are bound to specific IP addresses can be moved between instances.
- Hosting multiple SSL Websites on a single instance. You can install multiple SSL certificates on a single instance, each associated with a distinct IP address.

9.14.2. Guidelines

To prevent IP conflict, configure different subnets when multiple networks are connected to the same VM.

9.14.3. Assigning Additional IPs to a VM

1. Log in to the CloudPlatform UI.
2. In the left navigation bar, click Instances.
3. Click the name of the instance you want to work with.
4. In the Details tab, click NICs.
5. Click View Secondary IPs.
6. Click Acquire New Secondary IP, and click Yes in the confirmation dialog.

   You need to configure the IP on the guest VM NIC manually. CloudPlatform will not automatically configure the acquired IP address on the VM. Ensure that the IP address configuration persist on VM reboot.

   Within a few moments, the new IP address should appear with the state Allocated. You can now use the IP address in Port Forwarding or StaticNAT rules.

9.14.4. Port Forwarding and StaticNAT Services Changes

Because multiple IPs can be associated per NIC, you are allowed to select a desired IP for the Port Forwarding and StaticNAT services. The default is the primary IP. To enable this functionality, an extra

optional parameter 'vmguestip' is added to the Port forwarding and StaticNAT APIs (enableStaticNat, createIpForwardingRule) to indicate on what IP address NAT need to be configured. If vmguestip is passed, NAT is configured on the specified private IP of the VM. if not passed, NAT is configured on the primary IP of the VM.

9.15. Multiple Subnets in Shared Network

CloudPlatform provides you with the flexibility to add guest IP ranges from different subnets in Basic zones and security groups-enabled Advanced zones. For security groups-enabled Advanced zones, it implies multiple subnets can be added to the same VLAN. With the addition of this feature, you will be able to add IP address ranges from the same subnet or from a different one when IP address are exhausted. This would in turn allows you to employ higher number of subnets and thus reduce the address management overhead. You can delete the IP ranges you have added.

9.15.1. Prerequisites and Guidelines

- This feature can only be implemented:
  - on IPv4 addresses
  - if virtual router is the DHCP provider
  - on KVM, xenServer, and VMware hypervisors
- Manually configure the gateway of the new subnet before adding the IP range.
- CloudPlatform supports only one gateway for a subnet; overlapping subnets are not currently supported

9.15.2. Adding Multiple Subnets to a Shared Network

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Infrastructure.
3. On Zones, click View More, then click the zone to which you want to work with..
4. Click Physical Network.
5. In the Guest node of the diagram, click Configure.
6. Click Networks.
7. Select the networks you want to work with.
8. Click View IP Ranges.
9. Click Add IP Range.

   The Add IP Range dialog is displayed, as follows:

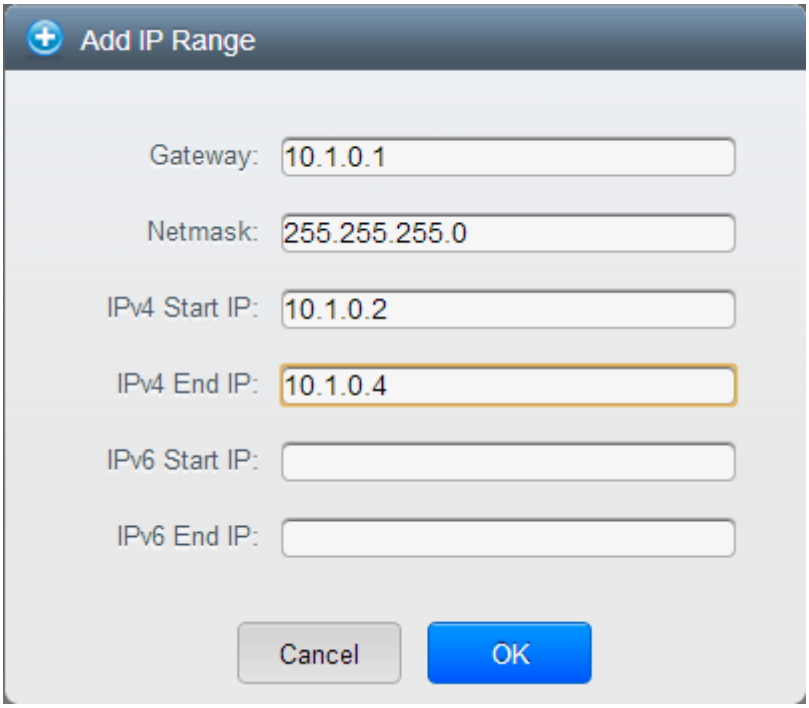

10. Specify the following:

    All the fields are mandatory.

    - **Gateway**: The gateway for the tier you create. Ensure that the gateway is within the Super CIDR range that you specified while creating the VPC, and is not overlapped with the CIDR of any existing tier within the VPC.
    - **Netmask**: The netmask for the tier you create.

      For example, if the VPC CIDR is 10.0.0.0/16 and the network tier CIDR is 10.0.1.0/24, the gateway of the tier is 10.0.1.1, and the netmask of the tier is 255.255.255.0.
    - **Start IP/ End IP**: A range of IP addresses that are accessible from the Internet and will be allocated to guest VMs. Enter the first and last IP addresses that define a range that CloudPlatform can assign to guest VMs .

11. Click OK.

9.16. Portable IPs

9.16.1. About Portable IP

Portable IPs in CloudPlatform are region-level pool of IPs, which are elastic in nature, that can be transferred across geographically separated zones. As an administrator, you can provision a pool of portable public IPs at region level and are available for user consumption. The users can acquire portable IPs if admin has provisioned portable IPs at the region level they are part of. These IPs can be use for any service within an advanced zone. You can also use portable IPs for EIP services in basic zones.

The salient features of Portable IP are as follows:

- IP is statically allocated

- IP need not be associated with a network
- IP association is transferable across networks
- IP is transferable across both Basic and Advanced zones
- IP is transferable across VPC, non-VPC isolated and shared networks
- Portable IP transfer is available only for static NAT.

Guidelines

Before transferring to another network, ensure that no network rules (Firewall, Static NAT, Port Forwarding, and so on) exist on that portable IP.

9.16.2. Configuring Portable IPs

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, click Regions.
3. Choose the Regions that you want to work with.
4. Click View Portable IP.
5. Click Portable IP Range.

   The Add Portable IP Range window is displayed.
6. Specify the following:
   - **Start IP/ End IP**: A range of IP addresses that are accessible from the Internet and will be allocated to guest VMs. Enter the first and last IP addresses that define a range that CloudPlatform can assign to guest VMs.
   - **Gateway**: The gateway in use for the Portable IP addresses you are configuring.
   - **Netmask**: The netmask associated with the Portable IP range.
   - **VLAN**: The VLAN that will be used for public traffic.
7. Click OK.

9.16.3. Acquiring a Portable IP

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. Click the name of the network where you want to work with.
4. Click View IP Addresses.
5. Click Acquire New IP.

   The Acquire New IP window is displayed.
6. Specify whether you want cross-zone IP or not.

7. Click Yes in the confirmation dialog.

   Within a few moments, the new IP address should appear with the state Allocated. You can now use the IP address in port forwarding or static NAT rules.

9.16.4. Transferring Portable IP

Portable IP is transferred from one network to another only if Static NAT is enabled. However, when a portable IP is associated with a network, you can use it for any service in the network.

To transfer a portable IP across the networks:

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. Click the name of the network where you want to work with.
4. Click View IP Addresses.
5. Click the Portable IP you want to work with.

   If none, acquire one,as explained in *Section 9.16.3, “Acquiring a Portable IP”*.
6. Click Enable Static NAT

   The Select VM for Static NAT page is displayed.
7. Select the desired VM.
8. Specify which IP to be replaced with for the Static NAT service.

   The VM can belong to any network owned by you.
9. Click Apply.

9.17. Static NAT

A static NAT rule maps a public IP address to the private IP address of a VM in order to allow Internet traffic into the VM. The public IP address always remains the same, which is why it is called “static” NAT. This section tells how to enable or disable static NAT for a particular IP address.

9.17.1. Enabling or Disabling Static NAT

If port forwarding rules are already in effect for an IP address, you cannot enable static NAT to that IP.

If a guest VM is part of more than one network, static NAT rules will function only if they are defined on the default network.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. Click the name of the network where you want to work with.
4. Click View IP Addresses.
5. Click the IP address you want to work with.

6. Click **Static NAT**. .

   The button toggles between Enable and Disable, depending on whether static NAT is currently enabled for the IP address.

7. If you are enabling static NAT, a dialog appears where you can choose the destination VM and click Apply.

9.18. IP Forwarding and Firewalling

By default, all incoming traffic to the public IP address is rejected. All outgoing traffic from the guests is also blocked by default.

To allow outgoing traffic, follow the procedure in *Section 9.18.1, “Egress Firewall Rules in an Advanced Zone”*.

To allow incoming traffic, users may set up firewall rules and/or port forwarding rules. For example, you can use a firewall rule to open a range of ports on the public IP address, such as 33 through 44. Then use port forwarding rules to direct traffic from individual ports within that range to specific ports on user VMs. For example, one port forwarding rule could route incoming traffic on the public IP's port 33 to port 100 on one user VM's private IP. For more information, see *Section 9.18.2, “Firewall Rules”* and *Section 9.18.3, “Port Forwarding”*.

9.18.1. Egress Firewall Rules in an Advanced Zone

The egress traffic originates from a private network to a public network, such as the Internet. By default, the egress traffic is blocked in default network offerings, so no outgoing traffic is allowed from a guest network to the Internet. However, you can control the egress traffic in an Advanced zone by creating egress firewall rules. When an egress firewall rule is applied, the traffic specific to the rule is allowed and the remaining traffic is blocked. When all the firewall rules are removed the default policy, Block, is applied.

9.18.1.1. Prerequisites and Guidelines

Consider the following scenarios to apply egress firewall rules:

- Egress firewall rules are supported on Juniper SRX and virtual router.
- The egress firewall rules are not supported on shared networks.
- Allow the egress traffic from specified source CIDR. The Source CIDR is part of guest network CIDR.
- Allow the egress traffic with protocol TCP,UDP,ICMP, or ALL.
- Allow the egress traffic with protocol and destination port range. The port range is specified for TCP, UDP or for ICMP type and code.
- The default policy is Allow for the new network offerings, whereas on upgrade existing network offerings with firewall service providers will have the default egress policy Deny.

9.18.1.2. Configuring an Egress Firewall Rule

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.

3. In Select view, choose Guest networks, then click the Guest network you want.
4. To add an egress rule, click the Egress rules tab and fill out the following fields to specify what type of traffic is allowed to be sent out of VM instances in this guest network:

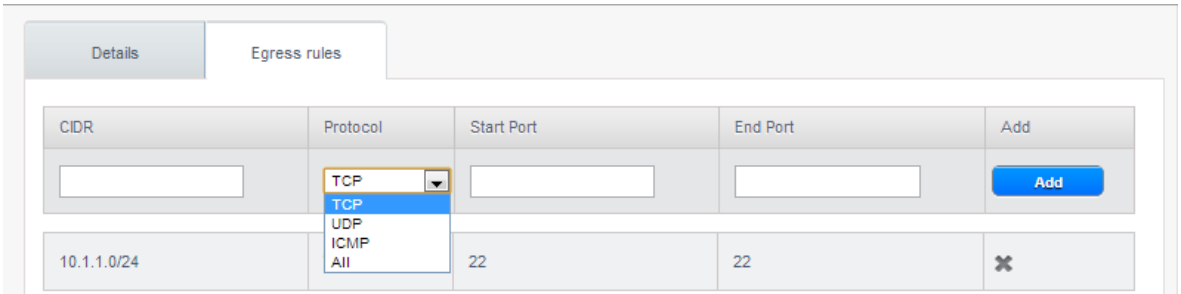

- **CIDR**: (Add by CIDR only) To send traffic only to the IP addresses within a particular address block, enter a CIDR or a comma-separated list of CIDRs. The CIDR is the base IP address of the destination. For example, 192.168.0.0/22. To allow all CIDRs, set to 0.0.0.0/0.
- **Protocol**: The networking protocol that VMs uses to send outgoing traffic. The TCP and UDP protocols are typically used for data exchange and end-user communications. The ICMP protocol is typically used to send error messages or network monitoring data.
- **Start Port, End Port**: (TCP, UDP only) A range of listening ports that are the destination for the outgoing traffic. If you are opening a single port, use the same number in both fields.
- **ICMP Type, ICMP Code**: (ICMP only) The type of message and error code that are sent.

5. Click Add.

9.18.1.3. Configuring the Default Egress Policy

The default egress policy for Isolated guest network is configured by using Network offering. Use the create network offering option to determine whether the default policy should be block or allow all the traffic to the public network from a guest network. Use this network offering to create the network. If no policy is specified, by default all the traffic is allowed from the guest network that you create by using this network offering.

You have two options: Allow and Deny.

Allow

If you select Allow for a network offering, by default egress traffic is allowed. However, when an egress rule is configured for a guest network, rules are applied to block the specified traffic and rest are allowed. If no egress rules are configured for the network, egress traffic is accepted.

Deny

If you select Deny for a network offering, by default egress traffic for the guest network is blocked. However, when an egress rules is configured for a guest network, rules are applied to allow the specified traffic. While implementing a guest network, CloudPlatform adds the firewall egress rule specific to the default egress policy for the guest network.

This feature is supported only on virtual router and Juniper SRX.

1. Create a network offering with your desirable default egress policy:
   a. Log in with admin privileges to the CloudPlatform UI.

b. In the left navigation bar, click Service Offerings.

c. In Select Offering, choose Network Offering.

d. Click Add Network Offering.

e. In the dialog, make necessary choices, including firewall provider.

f. In the Default egress policy field, specify the behaviour.

g. Click OK.

2. Create an isolated network by using this network offering.

   Based on your selection, the network will have the egress public traffic blocked or allowed.

9.18.2. Firewall Rules

By default, all incoming traffic to the public IP address is rejected by the firewall. To allow external traffic, you can open firewall ports by specifying firewall rules. You can optionally specify one or more CIDRs to filter the source IPs. This is useful when you want to allow only incoming requests from certain IP addresses.

You cannot use firewall rules to open ports for an elastic IP address. When elastic IP is used, outside access is instead controlled through the use of security groups. See *Section 9.4.4, “Adding a Security Group”*.

In an advanced zone, you can also create egress firewall rules by using the virtual router. For more information, see *Section 9.18.1, “Egress Firewall Rules in an Advanced Zone”*.

Firewall rules can be created using the Firewall tab in the Management Server UI. This tab is not displayed by default when CloudPlatform is installed. To display the Firewall tab, the CloudPlatform administrator must set the global configuration parameter firewall.rule.ui.enabled to "true."

To create a firewall rule:

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. Click the name of the network where you want to work with.
4. Click View IP Addresses.
5. Click the IP address you want to work with.
6. Click the Configuration tab and fill in the following values.

   - **Source CIDR**. (Optional) To accept only traffic from IP addresses within a particular address block, enter a CIDR or a comma-separated list of CIDRs. Example: 192.168.0.0/22. Leave empty to allow all CIDRs.
   - **Protocol**. The communication protocol in use on the opened port(s).
   - **Start Port and End Port**. The port(s) you want to open on the firewall. If you are opening a single port, use the same number in both fields

- **ICMP Type and ICMP Code**. Used only if Protocol is set to ICMP. Provide the type and code required by the ICMP protocol to fill out the ICMP header. Refer to ICMP documentation for more details if you are not sure what to enter

7. Click Add.

9.18.3. Port Forwarding

A port forward service is a set of port forwarding rules that define a policy. A port forward service is then applied to one or more guest VMs. The guest VM then has its inbound network access managed according to the policy defined by the port forwarding service. You can optionally specify one or more CIDRs to filter the source IPs. This is useful when you want to allow only incoming requests from certain IP addresses to be forwarded.

A guest VM can be in any number of port forward services. Port forward services can be defined but have no members. If a guest VM is part of more than one network, port forwarding rules will function only if they are defined on the default network

You cannot use port forwarding to open ports for an elastic IP address. When elastic IP is used, outside access is instead controlled through the use of security groups. See Security Groups.

A Virtual Router can forward only TCP or UDP traffic. You can create the following port forwarding rules in a Virtual Router:

- A single public port that forwards traffic to an identical single, private port. For example, 3389/TCP that forwards traffic to 3389/TCP and 22/TCP that forwards traffic to 22/TCP.
- A single public port that forwards traffic to a non-identical single private port. For example, 53389/TCP that forwards traffic to 3389/TCP and 50022/TCP that forwards traffic to 22/TCP.
- A range of public ports that forward traffic to an identical range of private ports. For example, 3001-3005/TCP that forward traffic to 3001-3005/TCP.

Note

A range of public ports cannot forward traffic to a non-identical range of private ports. For example, 3001-3005/TCP cannot forward traffic to 4001-4005/TCP.

If you want to forward IPSec traffic through a Virtual Router, you must enable NAT-T or IPSec-over-UDP on the VPN endpoints of the Virtual Router.

To set up port forwarding:

1. Log in to the CloudPlatform UI as an administrator or end user.
2. If you have not already done so, add a public IP address range to a zone in CloudPlatform. See Adding a Zone and Pod in the Installation Guide.
3. Add one or more VM instances to CloudPlatform.
4. In the left navigation bar, click Network.
5. Click the name of the guest network where the VMs are running.

6. Choose an existing IP address or acquire a new IP address. See *Section 9.10, "Acquiring a New IP Address"*. Click the name of the IP address in the list.
7. Click the Configuration tab.
8. In the Port Forwarding node of the diagram, click View All.
9. Fill in the following:
   - **Public Port**. The port to which public traffic will be addressed on the IP address you acquired in the previous step.
   - **Private Port**. The port on which the instance is listening for forwarded public traffic.
   - **Protocol**. The communication protocol in use between the two ports
10. Click Add.

9.19. IP Load Balancing

The user may choose to associate the same public IP for multiple guests. CloudPlatform implements a TCP-level load balancer with the following policies.

- Round-robin
- Least connection
- Source IP

This is similar to port forwarding but the destination may be multiple IP addresses.

9.20. DNS and DHCP

The Virtual Router provides DNS and DHCP services to the guests. It proxies DNS requests to the DNS server configured on the Availability Zone.

9.21. Virtual Private Network (VPN)

CloudPlatform account owners can create virtual private networks (VPN) to access their virtual machines. they can create two types of VPNs - Remote Access VPN and Site-to-Site VPN. For more information on Site-to-Site VPN, see *Section 9.21.4, "Setting Up a Site-to-Site VPN Connection"*

If the guest network is instantiated from a network offering that offers the Remote Access VPN service, the virtual router (based on the System VM) is used to provide the service. CloudPlatform provides a L2TP-over-IPsec-based remote access VPN service to guest virtual networks. Since each network gets its own virtual router, VPNs are not shared across the networks. CloudPlatform supports the VPN clients that use L2TP-over-IPsec protocol. You can use VPN clients native to Linux with OpenSwan, MAC OSx, Windows 2003, Windows Xp, and Windows 7 to connect to the guest network. The account owner can create and manage users for their VPN. CloudPlatform does not use its account database for this purpose but uses a separate table. The VPN user database is shared across all the VPNs created by the account owner. All VPN users get access to all VPNs created by the account owner.

Note

CloudPlatform does not advertise remote subnets when configured with a client VPN. You must set up the subnet routing in guest VMs manually.

Make sure that not all traffic goes through the VPN. That is, the route installed by the VPN should be only for the guest network and not for all traffic.

- **Road Warrior / Remote Access**. Users want to be able to connect securely from a home or office to a private network in the cloud. Typically, the IP address of the connecting client is dynamic and cannot be preconfigured on the VPN server.
- **Site to Site**. In this scenario, two private subnets are connected over the public Internet with a secure VPN tunnel. The cloud user's subnet (for example, an office network) is connected through a gateway to the network in the cloud. The address of the user's gateway must be preconfigured on the VPN server in the cloud. Note that although L2TP-over-IPsec can be used to set up Site-to-Site VPNs, this is not the primary intent of this feature.

9.21.1. Configuring Remote Access VPN

To set up VPN for the cloud:

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, click Global Settings.
3. Set the following global configuration parameters.
    - remote.access.vpn.client.ip.range – The range of IP addresses to be allocated to remote access VPN clients. The first IP in the range is used by the VPN server.
    - remote.access.vpn.psk.length – Length of the IPSec key.
    - remote.access.vpn.user.limit – Maximum number of VPN users per account.

Note

To make the updated global configuration parameters effective on the existing Remote Access VPNs that are in the enabled state, you must first disable these VPNs and then enable them. After you do this, you must restart Management Server.

In the case of a Remote Access VPN, the routing information will not be propagated to the client. After the VPN is connected, you can verify and update the client's route table, if necessary.

To enable VPN for a particular network:

1. Log in as a user or administrator to the CloudPlatform UI.
2. In the left navigation, click Network.

3. Click the name of the network you want to work with.
4. Click View IP Addresses.
5. Click one of the displayed IP address names.
6. Click the Enable VPN button. 

   The IPsec key is displayed in a popup window.

9.21.2. Using Remote Access VPN with Windows

The procedure to use VPN varies by Windows version. Generally, the user must edit the VPN properties and make sure that the default route is not the VPN. The following steps are for Windows L2TP clients on Windows Vista. The commands should be similar for other Windows versions.

1. Log in to the CloudPlatform UI and click on the source NAT IP for the account. The VPN tab should display the IPsec preshared key. Make a note of this and the source NAT IP. The UI also lists one or more users and their passwords. Choose one of these users, or, if none exists, add a user and password.
2. On the Windows box, go to Control Panel, then select Network and Sharing center. Click Setup a connection or network.
3. In the next dialog, select No, create a new connection.
4. In the next dialog, select Use my Internet Connection (VPN).
5. In the next dialog, enter the source NAT IP from step *1* and give the connection a name. Check Don't connect now.
6. In the next dialog, enter the user name and password selected in step *1*.
7. Click Create.
8. Go back to the Control Panel and click Network Connections to see the new connection. The connection is not active yet.
9. Right-click the new connection and select Properties. In the Properties dialog, select the Networking tab.
10. In Type of VPN, choose L2TP IPsec VPN, then click IPsec settings. Select Use preshared key. Enter the preshared key from step *1*.
11. The connection is ready for activation. Go back to Control Panel -> Network Connections and double-click the created connection.
12. Enter the user name and password from step *1*.

9.21.3. Using Remote Access VPN with Mac OS X

First, be sure you've configured the VPN settings in your CloudPlatform install. This section is only concerned with connecting via Mac OS X to your VPN.

Note, these instructions were written on Mac OS X 10.7.5. They may differ slightly in older or newer releases of Mac OS X.

1. On your Mac, open System Preferences and click Network.

2. Make sure Send all traffic over VPN connection is not checked.
3. If your preferences are locked, you'll need to click the lock in the bottom left-hand corner to make any changes and provide your administrator credentials.
4. You will need to create a new network entry. Click the plus icon on the bottom left-hand side and you'll see a dialog that says "Select the interface and enter a name for the new service." Select VPN from the Interface drop-down menu, and "L2TP over IPSec" for the VPN Type. Enter whatever you like within the "Service Name" field.
5. You'll now have a new network interface with the name of whatever you put in the "Service Name" field. For the purposes of this example, we'll assume you've named it "CloudStack." Click on that interface and provide the IP address of the interface for your VPN under the Server Address field, and the user name for your VPN under Account Name.
6. Click Authentication Settings, and add the user's password under User Authentication and enter the pre-shared IPSec key in the Shared Secret field under Machine Authentication. Click OK.
7. You may also want to click the "Show VPN status in menu bar" but that's entirely optional.
8. Now click "Connect" and you will be connected to the CloudStack VPN.

9.21.4. Setting Up a Site-to-Site VPN Connection

A Site-to-Site VPN connection helps you establish a secure connection from an enterprise datacenter to the cloud infrastructure. This allows users to access the guest VMs by establishing a VPN connection to the virtual router of the account from a device in the datacenter of the enterprise. You can also establish a secure connection between two VPC setups or high availability zones in your environment. Having this facility eliminates the need to establish VPN connections to individual VMs.

The difference from Remote VPN is that Site-to-site VPNs connects entire networks to each other, for example, connecting a branch office network to a company headquarters network. In a site-to-site VPN, hosts do not have VPN client software; they send and receive normal TCP/IP traffic through a VPN gateway.

The supported endpoints on the remote datacenters are:

- Cisco ISR with IOS 12.4 or later
- Juniper J-Series routers with JunOS 9.5 or later
- CloudPlatform virtual routers

Note

In addition to the specific Cisco and Juniper devices listed above, the expectation is that any Cisco or Juniper device running on the supported operating systems are able to establish VPN connections.

To set up a Site-to-Site VPN connection, perform the following:

1. Create a Virtual Private Cloud (VPC).

See *Section 9.24, "Configuring a Virtual Private Cloud"*.

2. Create a VPN Customer Gateway.
3. Create a VPN gateway for the VPC that you created.
4. Create VPN connection from the VPC VPN gateway to the customer VPN gateway.

9.21.4.1. Creating and Updating a VPN Customer Gateway

Note

A VPN customer gateway can be connected to only one VPN gateway at a time.

To add a VPN Customer Gateway:

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPN Customer Gateway.
4. Click Add VPN Customer Gateway.

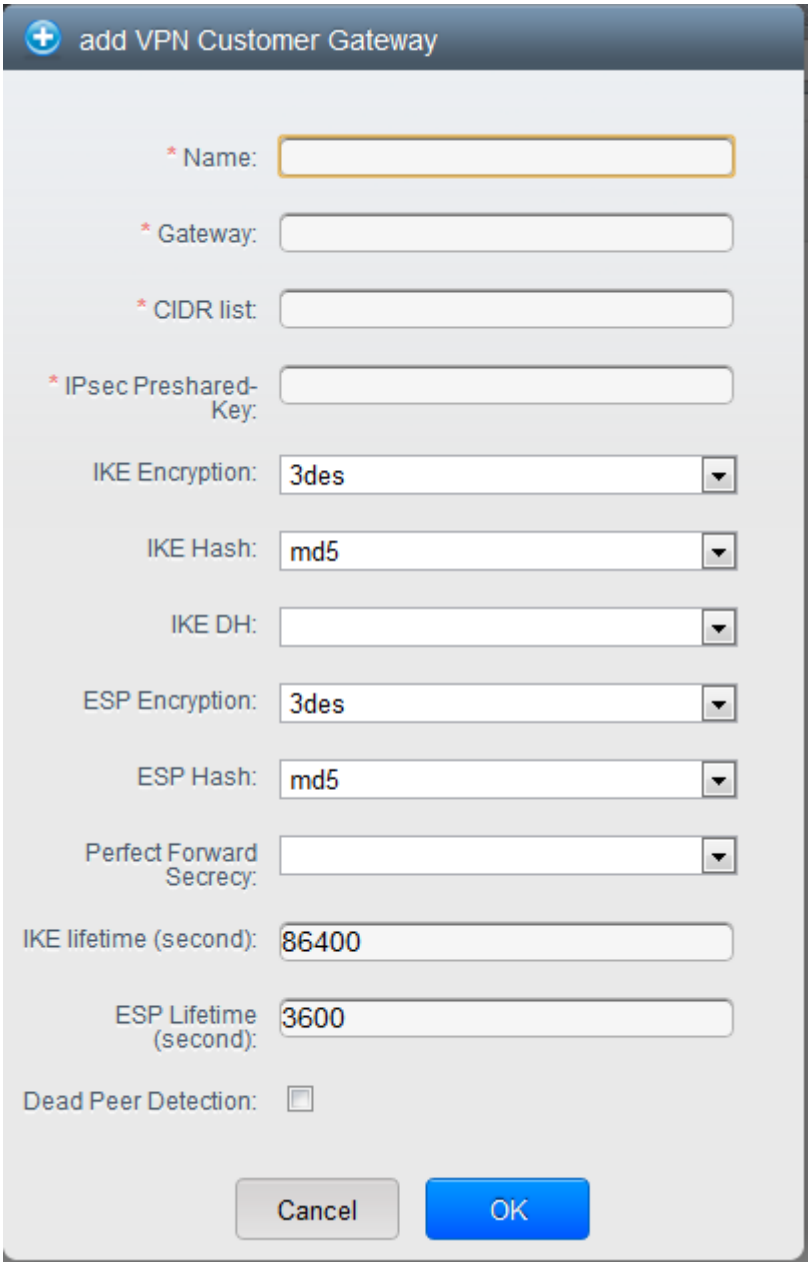

Provide the following information:

- **Name**: A unique name for the VPN customer gateway you create.
- **Gateway**: The IP address for the remote gateway.
- **CIDR list**: The guest CIDR list of the remote subnets. Enter a CIDR or a comma-separated list of CIDRs. Ensure that a guest CIDR list is not overlapped with the VPC's CIDR, or another guest CIDR. The CIDR must be RFC1918-compliant.
- **IPsec Preshared Key**: Preshared keying is a method where the endpoints of the VPN share a secret key. This key value is used to authenticate the customer gateway and the VPC VPN gateway to each other.

Note

The IKE peers (VPN end points) authenticate each other by computing and sending a keyed hash of data that includes the Preshared key. If the receiving peer is able to create the same hash independently by using its Preshared key, it knows that both peers must share the same secret, thus authenticating the customer gateway.

- **IKE Encryption**: The Internet Key Exchange (IKE) policy for phase-1. The supported encryption algorithms are AES128, AES192, AES256, and 3DES. Authentication is accomplished through the Preshared Keys.

Note

The phase-1 is the first phase in the IKE process. In this initial negotiation phase, the two VPN endpoints agree on the methods to be used to provide security for the underlying IP traffic. The phase-1 authenticates the two VPN gateways to each other, by confirming that the remote gateway has a matching Preshared Key.

- **IKE Hash**: The IKE hash for phase-1. The supported hash algorithms are SHA1 and MD5.
- **IKE DH**: A public-key cryptography protocol which allows two parties to establish a shared secret over an insecure communications channel. The 1536-bit Diffie-Hellman group is used within IKE to establish session keys. The supported options are None, Group-5 (1536-bit) and Group-2 (1024-bit).
- **ESP Encryption**: Encapsulating Security Payload (ESP) algorithm within phase-2. The supported encryption algorithms are AES128, AES192, AES256, and 3DES.

Note

The phase-2 is the second phase in the IKE process. The purpose of IKE phase-2 is to negotiate IPSec security associations (SA) to set up the IPSec tunnel. In phase-2, new keying material is extracted from the Diffie-Hellman key exchange in phase-1, to provide session keys to use in protecting the VPN data flow.

- **ESP Hash**: Encapsulating Security Payload (ESP) hash for phase-2. Supported hash algorithms are SHA1 and MD5.
- **Perfect Forward Secrecy**: Perfect Forward Secrecy (or PFS) is the property that ensures that a session key derived from a set of long-term public and private keys will not be compromised. This property enforces a new Diffie-Hellman key exchange. It provides the keying material

that has greater key material life and thereby greater resistance to cryptographic attacks. The available options are None, Group-5 (1536-bit) and Group-2 (1024-bit). The security of the key exchanges increase as the DH groups grow larger, as does the time of the exchanges.

Note

When PFS is turned on, for every negotiation of a new phase-2 SA the two gateways must generate a new set of phase-1 keys. This adds an extra layer of protection that PFS adds, which ensures if the phase-2 SA's have expired, the keys used for new phase-2 SA's have not been generated from the current phase-1 keying material.

- **IKE Lifetime (seconds)**: The phase-1 lifetime of the security association in seconds. Default is 86400 seconds (1 day). Whenever the time expires, a new phase-1 exchange is performed.
- **ESP Lifetime (seconds)**: The phase-2 lifetime of the security association in seconds. Default is 3600 seconds (1 hour). Whenever the value is exceeded, a re-key is initiated to provide a new IPsec encryption and authentication session keys.
- **Dead Peer Detection**: A method to detect an unavailable Internet Key Exchange (IKE) peer. Select this option if you want the virtual router to query the liveliness of its IKE peer at regular intervals. It's recommended to have the same configuration of DPD on both side of VPN connection.

5. Click OK.

Updating and Removing a VPN Customer Gateway

You can update a customer gateway either with no VPN connection, or related VPN connection is in error state.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPN Customer Gateway.
4. Select the VPN customer gateway you want to work with.
5. To modify the required parameters, click the Edit VPN Customer Gateway button 
6. To remove the VPN customer gateway, click the Delete VPN Customer Gateway button 
7. Click OK.

9.21.4.2. Creating a VPN gateway for the VPC

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.

All the VPCs that you have created for the account is listed in the page.

4. Click the Configure button of the VPC to which you want to deploy the VMs.

   The VPC page is displayed where all the tiers you created are listed in a diagram.

   For each tier, the following options are displayed:

   - Internal LB
   - Public LB IP
   - Static NAT
   - Virtual Machines
   - CIDR

   The following router information is displayed:

   - Private Gateways
   - Public IP Addresses
   - Site-to-Site VPNs
   - Network ACL Lists

5. Select Site-to-Site VPN.

   If you are creating the VPN gateway for the first time, selecting Site-to-Site VPN prompts you to create a VPN gateway.

6. In the confirmation dialog, click Yes to confirm.

   Within a few moments, the VPN gateway is created. You will be prompted to view the details of the VPN gateway you have created. Click Yes to confirm.

   The following details are displayed in the VPN Gateway page:

   - IP Address
   - Account
   - Domain

9.21.4.3. Creating a VPN Connection

Note

CloudPlatform supports creating up to 8 VPN connections.

1. Log in to the CloudPlatform UI as an administrator or end user.

2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPCs that you create for the account are listed in the page.
4. Click the Configure button of the VPC to which you want to deploy the VMs.

   The VPC page is displayed where all the tiers you created are listed in a diagram.
5. Click the Settings icon.

   For each tier, the following options are displayed:

   - Internal LB
   - Public LB IP
   - Static NAT
   - Virtual Machines
   - CIDR

   The following router information is displayed:

   - Private Gateways
   - Public IP Addresses
   - Site-to-Site VPNs
   - Network ACL Lists
6. Select Site-to-Site VPN.

   The Site-to-Site VPN page is displayed.
7. From the Select View drop-down, ensure that VPN Connection is selected.
8. Click Create VPN Connection.

   The Create VPN Connection dialog is displayed:

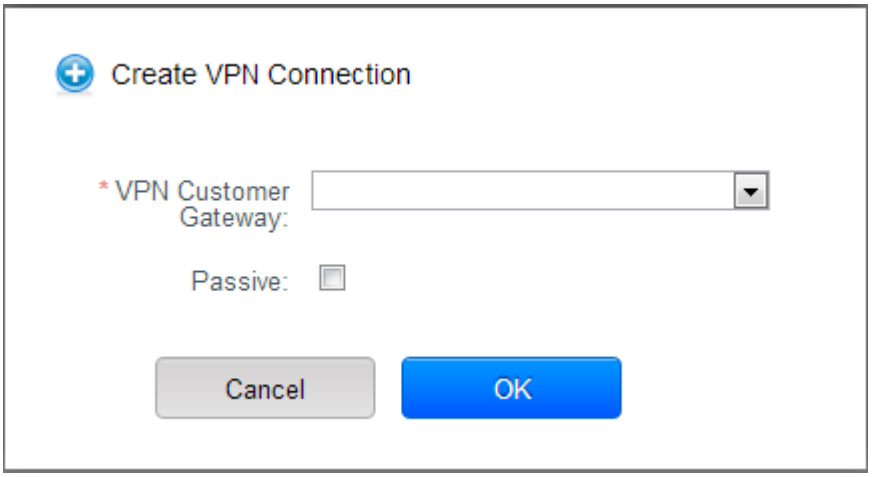

9. Select the desired customer gateway.
10. Select Passive if you want to establish a connection between two VPC virtual routers.

    If you want to establish a connection between two VPC virtual routers, select Passive only on one of the VPC virtual routers, which waits for the other VPC virtual router to initiate the connection. Do not select Passive on the VPC virtual router that initiates the connection.
11. Click OK to confirm.

    Within a few moments, the VPN Connection is displayed.

    The following information on the VPN connection is displayed:

    - IP Address
    - Gateway
    - State
    - IPSec Preshared Key
    - IKE Policy
    - ESP Policy

9.21.4.4. Site-to-Site VPN Connection Between VPC Networks

CloudPlatform provides you with the ability to establish a site-to-site VPN connection between CloudPlatform virtual routers. To achieve that, add a passive mode Site-to-Site VPN. With this functionality, users can deploy applications in multiple Availability Zones or VPCs, which can communicate with each other by using a secure Site-to-Site VPN Tunnel.

This feature is supported on all the hypervisors.

1. Create two VPCs. For example, VPC A and VPC B.

   For more information, see *Section 9.24, "Configuring a Virtual Private Cloud"*.
2. Create VPN gateways on both the VPCs you created.

   For more information, see *Section 9.21.4.2, "Creating a VPN gateway for the VPC"*.
3. Create VPN customer gateway for both the VPCs.

   For more information, see *Section 9.21.4.1, "Creating and Updating a VPN Customer Gateway"*.
4. Enable a VPN connection on VPC A in passive mode.

   For more information, see *Section 9.21.4.3, "Creating a VPN Connection"*.

   Ensure that the customer gateway is pointed to VPC B. The VPN connection is shown in the Disconnected state.
5. Enable a VPN connection on VPC B.

   Ensure that the customer gateway is pointed to VPC A. Because virtual router of VPC A, in this case, is in passive mode and is waiting for the virtual router of VPC B to initiate the connection, VPC B virtual router should not be in passive mode.

The VPN connection on VPC B is directly connected to VPC A's customer gateway.

Creating VPN connection on both the VPCs initiates a VPN connection. Wait for few seconds. The default is 30 seconds for both the VPN connections to show the Connected state.

9.21.4.5. Restarting and Removing a VPN Connection

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPCs that you have created for the account is listed in the page.
4. Click the Configure button of the VPC to which you want to deploy the VMs.

   The VPC page is displayed where all the tiers you created are listed in a diagram.
5. Click the Settings icon.

   For each tier, the following options are displayed:

   - Internal LB
   - Public LB IP
   - Static NAT
   - Virtual Machines
   - CIDR

   The following router information is displayed:

   - Private Gateways
   - Public IP Addresses
   - Site-to-Site VPNs
   - Network ACL Lists
6. Select Site-to-Site VPN.

   The Site-to-Site VPN page is displayed.
7. From the Select View drop-down, ensure that VPN Connection is selected.

   All the VPN connections you created are displayed.
8. Select the VPN connection you want to work with.

   The Details tab is displayed.
9. To remove a VPN connection, click the Delete VPN connection button

   To restart a VPN connection, click the Reset VPN connection button present in the Details tab.

9.22. Isolation in Advanced Zone Using Private VLAN

Isolation of guest traffic in shared networks can be achieved by using Private VLANs (PVLAN). PVLANs provide Layer 2 isolation between ports within the same VLAN. In a PVLAN-enabled shared network, a user VM cannot reach other user VM though they can reach the DHCP server and gateway, this would in turn allow users to control traffic within a network and help them deploy multiple applications without communication between application as well as prevent communication with other users' VMs.

- Isolate VMs in a shared networks by using Private VLANs.
- Supported on KVM, XenServer, and VMware hypervisors
- PVLAN-enabled shared network can be a part of multiple networks of a guest VM.

9.22.1. About Private VLAN

In an Ethernet switch, a VLAN is a broadcast domain where hosts can establish direct communication with each another at Layer 2. Private VLAN is designed as an extension of VLAN standard to add further segmentation of the logical broadcast domain. A regular VLAN is a single broadcast domain, whereas a private VLAN partitions a larger VLAN broadcast domain into smaller sub-domains. A sub-domain is represented by a pair of VLANs: a Primary VLAN and a Secondary VLAN. The original VLAN that is being divided into smaller groups is called Primary, which implies that all VLAN pairs in a private VLAN share the same Primary VLAN. All the secondary VLANs exist only inside the Primary. Each Secondary VLAN has a specific VLAN ID associated to it, which differentiates one sub-domain from another.

Three types of ports exist in a private VLAN domain, which essentially determine the behaviour of the participating hosts. Each ports will have its own unique set of rules, which regulate a connected host's ability to communicate with other connected host within the same private VLAN domain. Configure each host that is part of a PVLAN pair can be by using one of these three port designation:

- **Promiscuous**: A promiscuous port can communicate with all the interfaces, including the community and isolated host ports that belong to the secondary VLANs. In Promiscuous mode, hosts are connected to promiscuous ports and are able to communicate directly with resources on both primary and secondary VLAN. Routers, DHCP servers, and other trusted devices are typically attached to promiscuous ports.
- **Isolated VLANs**: The ports within an isolated VLAN cannot communicate with each other at the layer-2 level. The hosts that are connected to Isolated ports can directly communicate only with the Promiscuous resources. If your customer device needs to have access only to a gateway router, attach it to an isolated port.
- **Community VLANs**: The ports within a community VLAN can communicate with each other and with the promiscuous ports, but they cannot communicate with the ports in other communities at the layer-2 level. In a Community mode, direct communication is permitted only with the hosts in the same community and those that are connected to the Primary PVLAN in promiscuous mode. If your customer has two devices that need to be isolated from other customers' devices, but to be able to communicate among themselves, deploy them in community ports.

For further reading:

- *Understanding Private VLANs*[8]

[8] http://www.cisco.com/en/US/docs/switches/lan/catalyst3750/software/release/12.2_25_see/configuration/guide/swpvlan.html#wp1038379

- *Cisco Systems' Private VLANs: Scalable Security in a Multi-Client Environment*[9]
- *Private VLAN (PVLAN) on vNetwork Distributed Switch - Concept Overview (1010691)*[10]

9.22.2. Prerequisites

- Use a PVLAN supported switch.

  See *Private VLAN Catalyst Switch Support Matrix*[11]for more information.

- All the layer 2 switches, which are PVLAN-aware, are connected to each other, and one of them is connected to a router. All the ports connected to the host would be configured in trunk mode. Open Management VLAN, Primary VLAN (public) and Secondary Isolated VLAN ports. Configure the switch port connected to the router in PVLAN promiscuous trunk mode, which would translate an isolated VLAN to primary VLAN for the PVLAN-unaware router.

  Note that only Cisco Catalyst 4500 has the PVLAN promiscuous trunk mode to connect both normal VLAN and PVLAN to a PVLAN-unaware switch. For the other Catalyst PVLAN support switch, connect the switch to upper switch by using cables, one each for a PVLAN pair.

- Configure private VLAN on your physical switches out-of-band.

- Before you use PVLAN on XenServer and KVM, enable Open vSwitch (OVS).

Note

OVS on XenServer and KVM does not support PVLAN natively. Therefore, CloudPlatform managed to simulate PVLAN on OVS for XenServer and KVM by modifying the flow table.

9.22.3. Creating a PVLAN-Enabled Guest Network

1. Log in to the CloudPlatform UI as administrator.
2. In the left navigation, choose Infrastructure.
3. On Zones, click View More.
4. Click the zone to which you want to add a guest network.
5. Click the Physical Network tab.
6. Click the physical network you want to work with.
7. On the Guest node of the diagram, click Configure.
8. Click the Network tab.
9. Click Add guest network.

[9] http://tools.ietf.org/html/rfc5517
[10] http://kb.vmware.com
[11] http://www.cisco.com/en/US/products/hw/switches/ps708/products_tech_note09186a0080094830.shtml

The Add guest network window is displayed.

10. Specify the following:

- **Name**: The name of the network. This will be visible to the user.
- **Description**: The short description of the network that can be displayed to users.
- **VLAN ID**: The unique ID of the VLAN.
- **Secondary Isolated VLAN ID**: The unique ID of the Secondary Isolated VLAN.

  For the description on Secondary Isolated VLAN, see *Section 9.22.1, "About Private VLAN"*.
- **Scope**: The available scopes are Domain, Account, Project, and All.
  - **Domain**: Selecting Domain limits the scope of this guest network to the domain you specify. The network will not be available for other domains. If you select Subdomain Access, the guest network is available to all the sub domains within the selected domain.
  - **Account**: The account for which the guest network is being created for. You must specify the domain the account belongs to.
  - **Project**: The project for which the guest network is being created for. You must specify the domain the project belongs to.
  - **All**: The guest network is available for all the domains, account, projects within the selected zone.
- **Network Offering**: If the administrator has configured multiple network offerings, select the one you want to use for this network.
- **Gateway**: The gateway that the guests should use.
- **Netmask**: The netmask in use on the subnet the guests will use.
- **IP Range**: A range of IP addresses that are accessible from the Internet and are assigned to the guest VMs.
- **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.

11. Click OK to confirm.

9.23. About Inter-VLAN Routing

Inter-VLAN Routing is the capability to route network traffic between VLANs. This feature enables you to build Virtual Private Clouds (VPC), an isolated segment of your cloud, that can hold multi-tier applications. These tiers are deployed on different VLANs that can communicate with each other. You provision VLANs to the tiers your create, and VMs can be deployed on different tiers. The VLANs are connected to a virtual router, which facilitates communication between the VMs. In effect, you can segment VMs by means of VLANs into different networks that can host multi-tier applications, such as Web, Application, or Database. Such segmentation by means of VLANs logically separate application VMs for higher security and lower broadcasts, while remaining physically connected to the same device.

This feature is supported on XenServer and VMware hypervisors.

The major advantages are:

- The administrator can deploy a set of VLANs and allow users to deploy VMs on these VLANs. A guest VLAN is randomly alloted to an account from a pre-specified set of guest VLANs. All the VMs of a certain tier of an account reside on the guest VLAN allotted to that account.

  Note

  A VLAN allocated for an account cannot be shared between multiple accounts.

- The administrator can allow users create their own VPC and deploy the application. In this scenario, the VMs that belong to the account are deployed on the VLANs allotted to that account.
- Both administrators and users can create multiple VPCs. The guest network NIC is plugged to the VPC virtual router when the first VM is deployed in a tier.
- The administrator can create the following gateways to send to or receive traffic from the VMs:
  - **VPN Gateway**: For more information, see *Section 9.21.4.2, "Creating a VPN gateway for the VPC"*.
  - **Public Gateway**: The public gateway for a VPC is added to the virtual router when the virtual router is created for VPC. The public gateway is not exposed to the end users. You are not allowed to list it, nor allowed to create any static routes.
  - **Private Gateway**: For more information, see *Section 9.24.5, "Adding a Private Gateway to a VPC"*.
- Both administrators and users can create various possible destinations-gateway combinations. However, only one gateway of each type can be used in a deployment.

  For example:

  - **VLANs and Public Gateway**: For example, an application is deployed in the cloud, and the Web application VMs communicate with the Internet.
  - **VLANs, VPN Gateway, and Public Gateway**: For example, an application is deployed in the cloud; the Web application VMs communicate with the Internet; and the database VMs communicate with the on-premise devices.
- The administrator can define Access Control List (ACL) on the virtual router to filter the traffic among the VLANs or between the Internet and a VLAN. You can define ACL based on CIDR, port range, protocol, type code (if ICMP protocol is selected) and Ingress/Egress type.

The following figure shows the possible deployment scenarios of a Inter-VLAN setup:

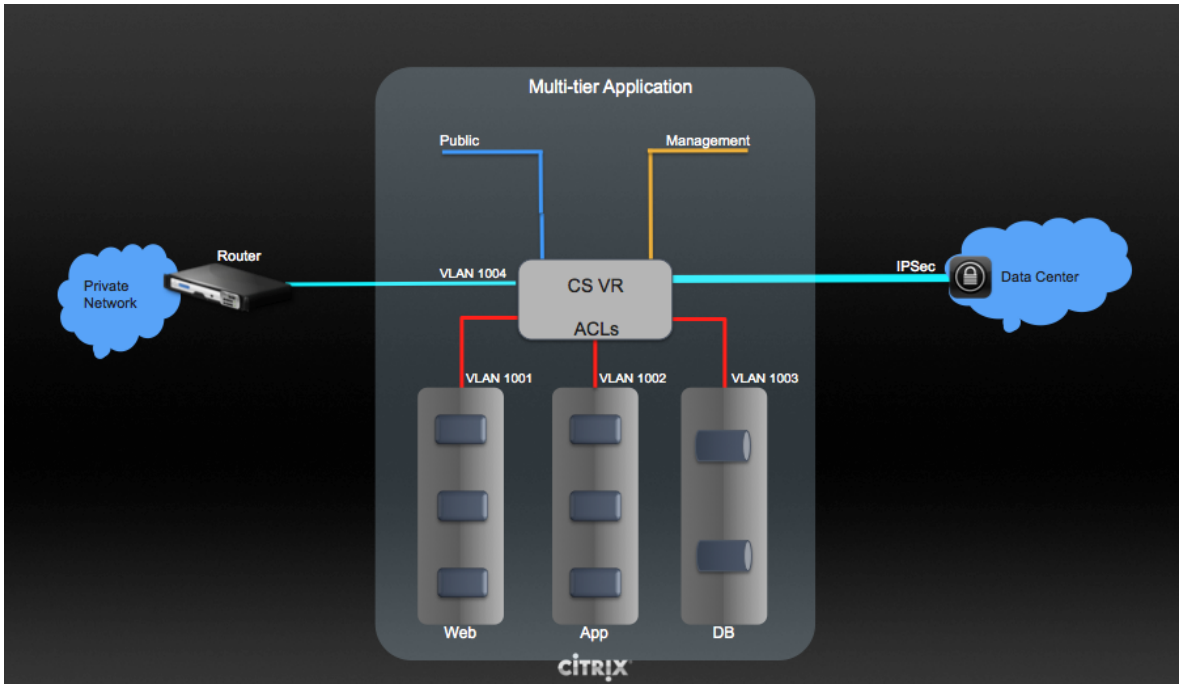

To set up a multi-tier Inter-VLAN deployment, see *Section 9.24, "Configuring a Virtual Private Cloud"*.

9.24. Configuring a Virtual Private Cloud

Hyper-V is not supported in VPC.

9.24.1. About Virtual Private Clouds

CloudPlatform Virtual Private Cloud is a private, isolated part of CloudPlatform. A VPC can have its own virtual network topology that resembles a traditional physical network. You can launch VMs in the virtual network that can have private addresses in the range of your choice, for example: 10.0.0.0/16. You can define network tiers within your VPC network range, which in turn enables you to group similar kinds of instances based on IP address range.

For example, if a VPC has the private range 10.0.0.0/16, its guest networks can have the network ranges 10.0.1.0/24, 10.0.2.0/24, 10.0.3.0/24, and so on.

Major Components of a VPC:

A VPC is comprised of the following network components:

- **VPC**: A VPC acts as a container for multiple isolated networks that can communicate with each other via its virtual router.
- **Network Tiers**: Each tier acts as an isolated network with its own VLANs and CIDR list, where you can place groups of resources, such as VMs. The tiers are segmented by means of VLANs. The NIC of each tier acts as its gateway.
- **Virtual Router**: A virtual router is automatically created and started when you create a VPC. The virtual router connect the tiers and direct traffic among the public gateway, the VPN gateways, and the NAT instances. For each tier, a corresponding NIC and IP exist in the virtual router. The virtual router provides DNS and DHCP services through its IP.

- **Public Gateway**: The traffic to and from the Internet routed to the VPC through the public gateway. In a VPC, the public gateway is not exposed to the end user; therefore, static routes are not support for the public gateway.
- **Private Gateway**: All the traffic to and from a private network routed to the VPC through the private gateway. For more information, see *Section 9.24.5, "Adding a Private Gateway to a VPC"*.
- **VPN Gateway**: The VPC side of a VPN connection.
- **Site-to-Site VPN Connection**: A hardware-based VPN connection between your VPC and your datacenter, home network, or co-location facility. For more information, see *Section 9.21.4, "Setting Up a Site-to-Site VPN Connection"*.
- **Customer Gateway**: The customer side of a VPN Connection. For more information, see *Section 9.21.4.1, "Creating and Updating a VPN Customer Gateway"*.
- **NAT Instance**: An instance that provides Port Address Translation for instances to access the Internet via the public gateway. For more information, see *Section 9.24.10, "Enabling or Disabling Static NAT on a VPC"*.

Network Architecture in a VPC

In a VPC, the following four basic options of network architectures are present:

- VPC with a public gateway only
- VPC with public and private gateways
- VPC with public and private gateways and site-to-site VPN access
- VPC with a private gateway only and site-to-site VPN access

Connectivity Options for a VPC

You can connect your VPC to:

- The Internet through the public gateway.
- The corporate datacenter by using a site-to-site VPN connection through the VPN gateway.
- Both the Internet and your corporate datacenter by using both the public gateway and a VPN gateway.

VPC Network Considerations

Consider the following before you create a VPC:

- A VPC, by default, is created in the enabled state.
- A VPC can be created in Advance zone only, and can't belong to more than one zone at a time.
- The default number of VPCs an account can create is 20. However, you can change it by using the max.account.vpcs global parameter, which controls the maximum number of VPCs an account is allowed to create.
- The default number of tiers an account can create within a VPC is 3. You can configure this number by using the vpc.max.networks parameter.

- Each tier should have an unique CIDR in the VPC. Ensure that the tier's CIDR should be within the VPC CIDR range.
- A tier belongs to only one VPC.
- All network tiers inside the VPC should belong to the same account.
- When a VPC is created, by default, a SourceNAT IP is allocated to it. The Source NAT IP is released only when the VPC is removed.
- A public IP can be used for only one purpose at a time. If the IP is a sourceNAT, it cannot be used for StaticNAT or port forwarding.
- The instances can only have a private IP address that you provision. To communicate with the Internet, enable NAT to an instance that you launch in your VPC.
- Only new networks can be added to a VPC. The maximum number of networks per VPC is limited by the value you specify in the vpc.max.networks parameter. The default value is three.
- The load balancing service can be supported by only one tier inside the VPC.
- If an IP address is assigned to a tier:
  - That IP can't be used by more than one tier at a time in the VPC. For example, if you have tiers A and B, and a public IP1, you can create a port forwarding rule by using the IP either for A or B, but not for both.
  - That IP can't be used for StaticNAT, load balancing, or port forwarding rules for another guest network inside the VPC.
- Remote access VPN is not supported in VPC networks.

9.24.2. Adding a Virtual Private Cloud

When creating the VPC, you simply provide the zone and a set of IP addresses for the VPC network address space. You specify this set of addresses in the form of a Classless Inter-Domain Routing (CIDR) block.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.
4. Click Add VPC. The Add VPC page is displayed as follows:

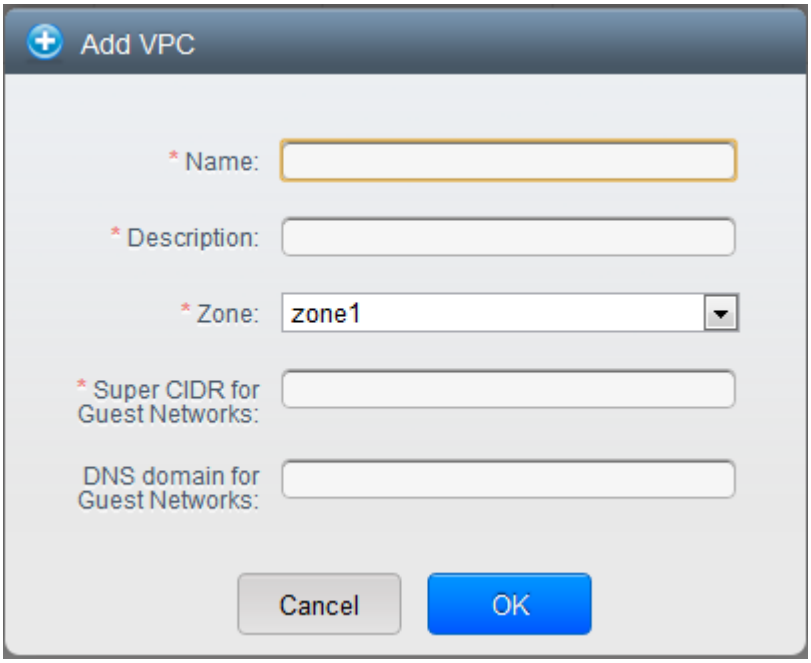

Provide the following information:

- **Name**: A short name for the VPC that you are creating.
- **Description**: A brief description of the VPC.
- **Zone**: Choose the zone where you want the VPC to be available.
- **Super CIDR for Guest Networks**: Defines the CIDR range for all the tiers (guest networks) within a VPC. When you create a tier, ensure that its CIDR is within the Super CIDR value you enter. The CIDR must be RFC1918 compliant.
- **DNS domain for Guest Networks**: If you want to assign a special domain name, specify the DNS suffix. This parameter is applied to all the tiers within the VPC. That implies, all the tiers you create in the VPC belong to the same DNS domain. If the parameter is not specified, a DNS domain name is generated automatically.

9.24.3. Adding Tiers

Tiers are distinct locations within a VPC that act as isolated networks, which do not have access to other tiers by default. Tiers are set up on different VLANs that can communicate with each other by using a virtual router. Tiers provide inexpensive, low latency network connectivity to other tiers within the VPC.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPC that you have created for the account is listed in the page.

Note

The end users can see their own VPCs, while root and domain admin can see any VPC they are authorized to see.

4. Click the Configure button of the VPC for which you want to set up tiers.
5. Click Create network.

   The Add new tier dialog is displayed, as follows:

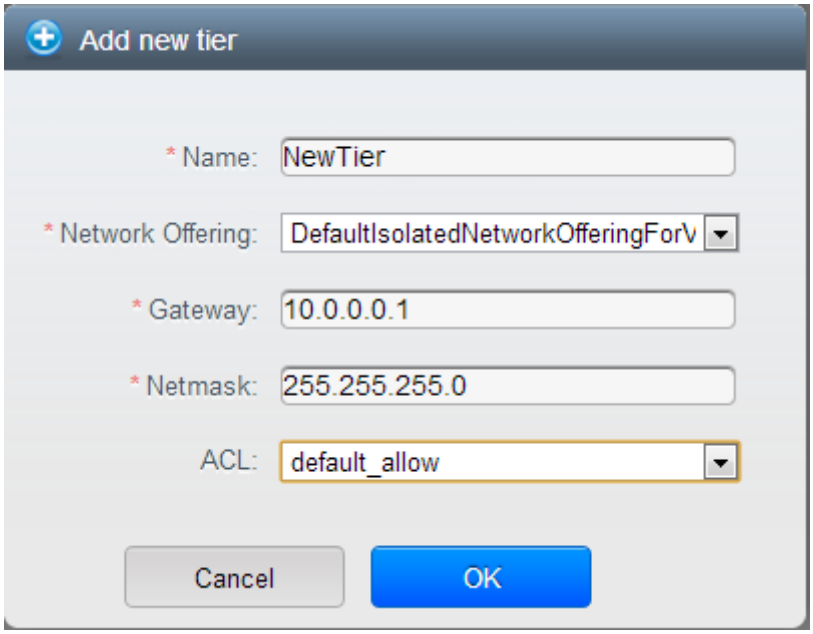

   If you have already created tiers, the VPC diagram is displayed. Click Create Tier to add a new tier.

6. Specify the following:

   All the fields are mandatory.

   - **Name**: A unique name for the tier you create.
   - **Network Offering**: The following default network offerings are listed: Internal LB, DefaultIsolatedNetworkOfferingForVpcNetworksNoLB, DefaultIsolatedNetworkOfferingForVpcNetworks

     In a VPC, only one tier can be created by using LB-enabled network offering.
   - **Gateway**: The gateway for the tier you create. Ensure that the gateway is within the Super CIDR range that you specified while creating the VPC, and is not overlapped with the CIDR of any existing tier within the VPC.
   - **VLAN**: The VLAN ID for the tier you create.

     This option is only visible if the network offering you selected is VLAN-enabled.

For more information, see *Section 7.10.3, "Assigning VLANs to Isolated Networks"*.

- **Netmask**: The netmask for the tier you create.

  For example, if the VPC CIDR is 10.0.0.0/16 and the network tier CIDR is 10.0.1.0/24, the gateway of the tier is 10.0.1.1, and the netmask of the tier is 255.255.255.0.

7. Click OK.

8. Continue with configuring access control list for the tier.

9.24.4. Configuring Network Access Control List

Define Network Access Control List (ACL) on the VPC virtual router to control incoming (ingress) and outgoing (egress) traffic between the VPC tiers, and the tiers and Internet. By default, all incoming traffic to the guest networks is blocked and all outgoing traffic from guest networks is allowed, once you add an ACL rule for outgoing traffic, then only outgoing traffic specified in this ACL rule is allowed, the rest is blocked. To open the ports, you must create a new network ACL. The network ACLs can be created for the tiers only if the NetworkACL service is supported.

9.24.4.1. About Network ACL Lists

In CloudPlatform terminology, Network ACL is a group of Network ACL items. Network ACL items are nothing but numbered rules that are evaluated in order, starting with the lowest numbered rule. These rules determine whether traffic is allowed in or out of any tier associated with the network ACL. You need to add the Network ACL items to the Network ACL, then associate the Network ACL with a tier. Network ACL is associated with a VPC and can be assigned to multiple VPC tiers within a VPC. A Tier is associated with a Network ACL at all the times. Each tier can be associated with only one ACL.

The default Network ACL is used when no ACL is associated. Default behavior is all the incoming traffic is blocked and outgoing traffic is allowed from the tiers. Default network ACL cannot be removed or modified. Contents of the default Network ACL is:

Rule	Protocol	Traffic type	Action	CIDR
1	All	Ingress	Deny	0.0.0.0/0
2	All	Egress	Allow	0.0.0.0/0

9.24.4.2. Creating ACL Lists

1. Log in to the CloudPlatform UI as an administrator or end user.

2. In the left navigation, choose Network.

3. In the Select view, select VPC.

   All the VPCs that you have created for the account is listed in the page.

4. Click the Configure button of the VPC.

   For each tier, the following options are displayed:

   - Internal LB
   - Public LB IP
   - Static NAT

- Virtual Machines
- CIDR

The following router information is displayed:

- Private Gateways
- Public IP Addresses
- Site-to-Site VPNs
- Network ACL Lists

5. Select Network ACL Lists.

   The following default rules are displayed in the Network ACLs page: default_allow, default_deny.

6. Click Add ACL Lists, and specify the following:

   - **ACL List Name**: A name for the ACL list.
   - **Description**: A short description of the ACL list that can be displayed to users.

9.24.4.3. Creating an ACL Rule

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPCs that you have created for the account is listed in the page.

4. Click the Configure button of the VPC.
5. Select Network ACL Lists.

   In addition to the custom ACL lists you have created, the following default rules are displayed in the Network ACLs page: default_allow, default_deny.

6. Select the desired ACL list.
7. Select the ACL List Rules tab.

   To add an ACL rule, fill in the following fields to specify what kind of network traffic is allowed in the VPC.

   - **Rule Number**: The order in which the rules are evaluated.
   - **CIDR**: The CIDR acts as the Source CIDR for the Ingress rules, and Destination CIDR for the Egress rules. To accept traffic only from or to the IP addresses within a particular address block, enter a CIDR or a comma-separated list of CIDRs. The CIDR is the base IP address of the incoming traffic. For example, 192.168.0.0/22. To allow all CIDRs, set to 0.0.0.0/0.
   - **Action**: What action to be taken. Allow traffic or block.
   - **Protocol**: The networking protocol that sources use to send traffic to the tier. The TCP and UDP protocols are typically used for data exchange and end-user communications. The ICMP

protocol is typically used to send error messages or network monitoring data. All supports all the traffic. Other option is Protocol Number.

- **Start Port**, **End Port** (TCP, UDP only): A range of listening ports that are the destination for the incoming traffic. If you are opening a single port, use the same number in both fields.
- **Protocol Number**: The protocol number associated with IPv4. For more information, see *Protocol Numbers*[12].
- **ICMP Type**, **ICMP Code** (ICMP only): The type of message and error code that will be sent.
- **Traffic Type**: The type of traffic: Incoming or outgoing.

8. Click Add. The ACL rule is added.

   You can edit the tags assigned to the ACL rules and delete the ACL rules you have created. Click the appropriate button in the Details tab.

9.24.4.4. Creating a Tier with Custom ACL List

1. Create a VPC.
2. Create a custom ACL list.
3. Add ACL rules to the ACL list.
4. Create a tier in the VPC.

   Select the desired ACL list while creating a tier.
5. Click OK.

9.24.4.5. Assigning a Custom ACL List to a Tier

1. Create a VPC.
2. Create a tier in the VPC.
3. Associate the tier with the default ACL rule.
4. Create a custom ACL list.
5. Add ACL rules to the ACL list.
6. Select the tier for which you want to assign the custom ACL.
7. Click the Replace ACL List icon. 

   The Replace ACL List dialog is displayed.
8. Select the desired ACL list.
9. Click OK.

[12] http://www.iana.org/assignments/protocol-numbers/protocol-numbers.xml

9.24.5. Adding a Private Gateway to a VPC

A private gateway can be added by the root admin only. The VPC private network has 1:1 relationship with the NIC of the physical network. You can configure multiple private gateways to a single VPC. No gateways with duplicated VLAN and IP are allowed in the same data center.

1. Log in to the CloudPlatform UI using the root administrator privileges.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPCs that you have created for the account is listed in the page.
4. Click the Configure button of the VPC to which you want to configure load balancing rules.

   The VPC page is displayed where all the tiers you created are listed in a diagram.
5. Click the Settings icon.

   The following options are displayed.

   - Internal LB
   - Public LB IP
   - Static NAT
   - Virtual Machines
   - CIDR

   The following router information is displayed:

   - Private Gateways
   - Public IP Addresses
   - Site-to-Site VPNs
   - Network ACL Lists
6. Select Private Gateways.

   The Gateways page is displayed.
7. Click Add new gateway:

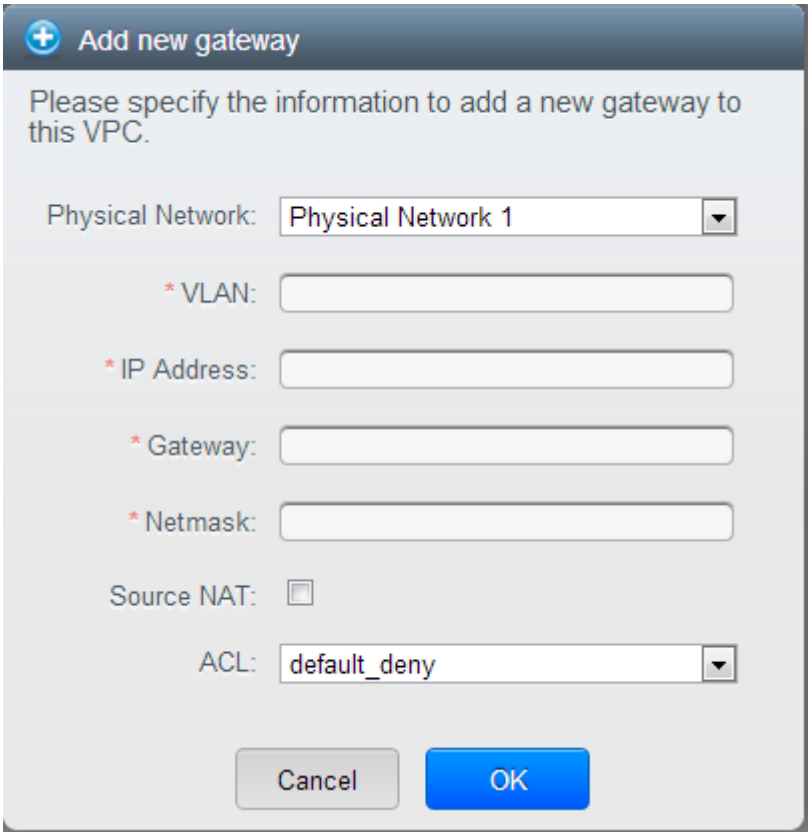

8. Specify the following:

   - **Physical Network**: The physical network you have created in the zone.
   - **IP Address**: The IP address associated with the VPC gateway.
   - **Gateway**: The gateway through which the traffic is routed to and from the VPC.
   - **Netmask**: The netmask associated with the VPC gateway.
   - **VLAN**: The VLAN associated with the VPC gateway.
   - **Source NAT**: Select this option to enable the source NAT service on the VPC private gateway.

     See *Section 9.24.5.1, "Source NAT on Private Gateway"*.
   - **ACL**: Controls both ingress and egress traffic on a VPC private gateway. By default, all the traffic is blocked.

     See *Section 9.24.5.2, "ACL on Private Gateway"*.

   The new gateway appears in the list. You can repeat these steps to add more gateway for this VPC.

9.24.5.1. Source NAT on Private Gateway

You might want to deploy multiple VPCs with the same super CIDR and guest tier CIDR. Therefore, multiple guest VMs from different VPCs can have the same IPs to reach a enterprise data center through the private gateway. In such cases, a NAT service need to be configured on the private

gateway to avoid IP conflicts. If Source NAT is enabled, the guest VMs in VPC reaches the enterprise network via private gateway IP address by using the NAT service.

The Source NAT service on a private gateway can be enabled while adding the private gateway. On deletion of a private gateway, source NAT rules specific to the private gateway are deleted.

To enable source NAT on existing private gateways, delete them and create afresh with source NAT.

9.24.5.2. ACL on Private Gateway

The traffic on the VPC private gateway is controlled by creating both ingress and egress network ACL rules. The ACLs contains both allow and deny rules. As per the rule, all the ingress traffic to the private gateway interface and all the egress traffic out from the private gateway interface are blocked.

You can change this default behaviour while creating a private gateway. Alternatively, you can do the following:

1. In a VPC, identify the Private Gateway you want to work with.
2. In the Private Gateway page, do either of the following:
   - Use the Quickview. See *3*.
   - Use the Details tab. See *4* through .
3. In the Quickview of the selected Private Gateway, click Replace ACL, select the ACL rule, then click OK
4. Click the IP address of the Private Gateway you want to work with.
5. In the Detail tab, click the Replace ACL button. 

   The Replace ACL dialog is displayed.
6. select the ACL rule, then click OK.

   Wait for few seconds. You can see that the new ACL rule is displayed in the Details page.

9.24.5.3. Creating a Static Route

CloudPlatform enables you to specify routing for the VPN connection you create. You can enter one or CIDR addresses to indicate which traffic is to be routed back to the gateway.

1. In a VPC, identify the Private Gateway you want to work with.
2. In the Private Gateway page, click the IP address of the Private Gateway you want to work with.
3. Select the Static Routes tab.
4. Specify the CIDR of destination network.
5. Click Add.

   Wait for few seconds until the new route is created.

9.24.5.4. Blacklisting Routes

CloudPlatform enables you to block a list of routes so that they are not assigned to any of the VPC private gateways. Specify the list of routes that you want to blacklist in the **`blacklisted.routes`** global parameter. Note that the parameter update affects only new static route creations. If you block an existing static route, it remains intact and continue functioning. You cannot add a static route if the route is blacklisted for the zone.

9.24.6. Deploying VMs to the Tier

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPCs that you have created for the account is listed in the page.
4. Click the Configure button of the VPC to which you want to deploy the VMs.

   The VPC page is displayed where all the tiers you have created are listed.
5. Click Virtual Machines tab of the tier to which you want to add a VM.

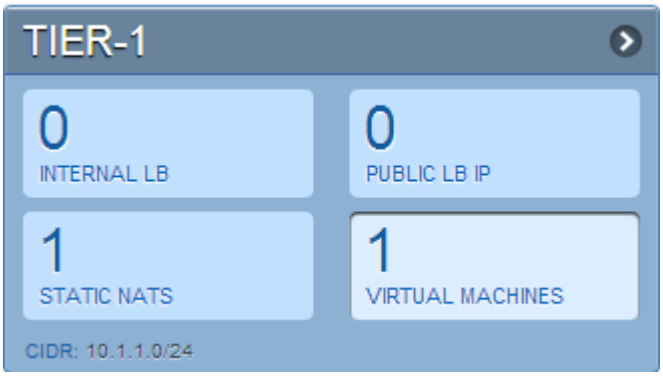

   The Add Instance page is displayed.

   Follow the on-screen instruction to add an instance. For information on adding an instance, see the Installation Guide.

9.24.7. Deploying VMs to VPC Tier and Shared Networks

CloudPlatform allows you deploy VMs on a VPC tier and one or more shared networks. With this feature, VMs deployed in a multi-tier application can receive services offered by a service provider over the shared network. One example of such a service is monitoring service.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Instances.
3. Select the VM you want to work with.
4. Click Add Instance.
5. Select a zone.
6. Select a template or ISO, then follow the steps in the wizard.

7. Ensure that the hardware you have allows starting the selected service offering.
8. Under Networks, select networks for the VM you are launching.

   You can deploy a VM to a VPC tier and multiple shared networks.

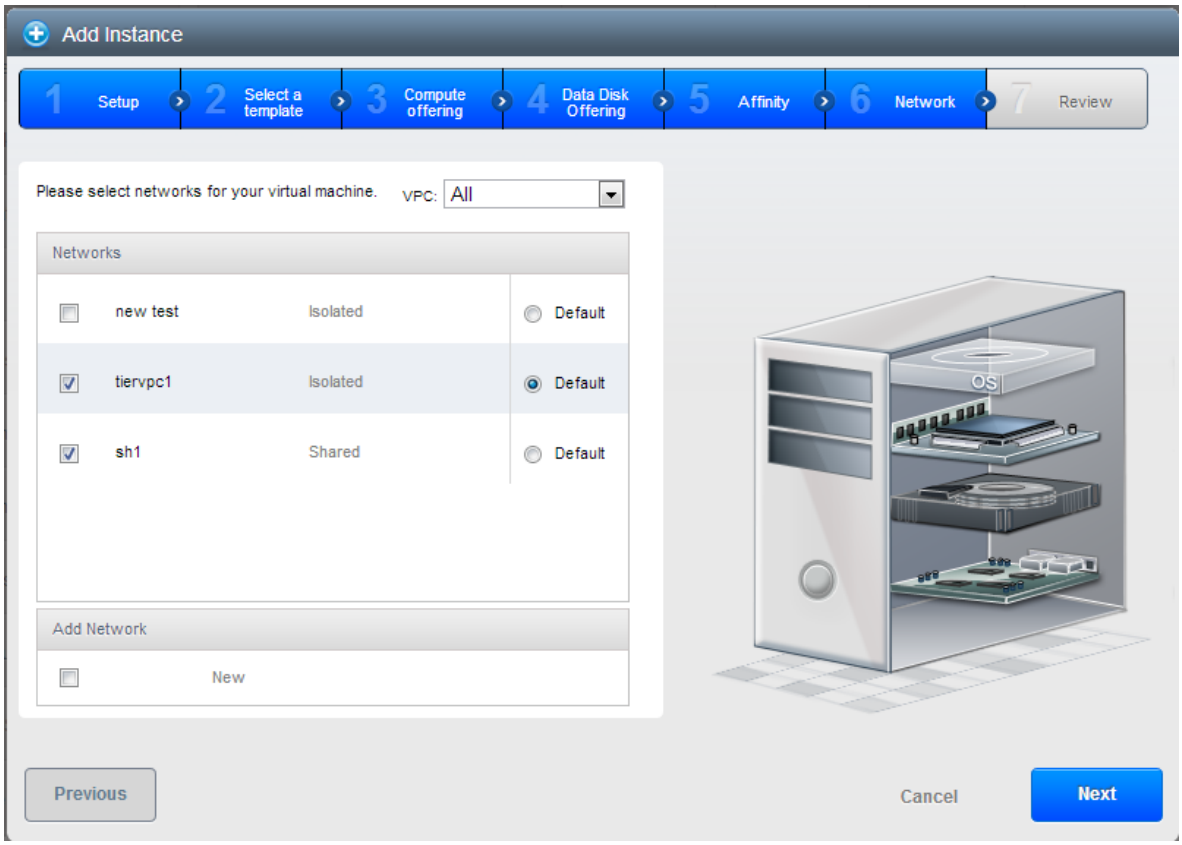

9. Click Next, review the configuration and click Launch.

   Your VM will be deployed to the selected VPC tier and shared network.

9.24.8. Acquiring a New IP Address for a VPC

When you acquire an IP address, all IP addresses are allocated to VPC, not to the guest networks within the VPC. The IPs are associated to the guest network only when the first port-forwarding, load balancing, or Static NAT rule is created for the IP or the network. IP can't be associated to more than one network at a time.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPCs that you have created for the account is listed in the page.
4. Click the Configure button of the VPC to which you want to deploy the VMs.

   The VPC page is displayed where all the tiers you created are listed in a diagram.

   The following options are displayed.

- Internal LB
- Public LB IP
- Static NAT
- Virtual Machines
- CIDR

The following router information is displayed:

- Private Gateways
- Public IP Addresses
- Site-to-Site VPNs
- Network ACL Lists

5. Select IP Addresses.

   The Public IP Addresses page is displayed.

6. Click Acquire New IP, and click Yes in the confirmation dialog.

   You are prompted for confirmation because, typically, IP addresses are a limited resource. Within a few moments, the new IP address should appear with the state Allocated. You can now use the IP address in port forwarding, load balancing, and static NAT rules.

9.24.9. Releasing an IP Address Alloted to a VPC

The IP address is a limited resource. If you no longer need a particular IP, you can disassociate it from its VPC and return it to the pool of available addresses. An IP address can be released from its tier, only when all the networking (port forwarding, load balancing, or StaticNAT) rules are removed for this IP address. The released IP address will still belongs to the same VPC.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPCs that you have created for the account is listed in the page.

4. Click the Configure button of the VPC whose IP you want to release.

   The VPC page is displayed where all the tiers you created are listed in a diagram.

   The following options are displayed.

   - Internal LB
   - Public LB IP
   - Static NAT
   - Virtual Machines

- CIDR

The following router information is displayed:

- Private Gateways
- Public IP Addresses
- Site-to-Site VPNs
- Network ACL Lists

5. Select Public IP Addresses.

   The IP Addresses page is displayed.

6. Click the IP you want to release.

7. In the Details tab, click the Release IP button 

9.24.10. Enabling or Disabling Static NAT on a VPC

A static NAT rule maps a public IP address to the private IP address of a VM in a VPC to allow Internet traffic to it. This section tells how to enable or disable static NAT for a particular IP address in a VPC.

If port forwarding rules are already in effect for an IP address, you cannot enable static NAT to that IP.

If a guest VM is part of more than one network, static NAT rules will function only if they are defined on the default network.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPCs that you have created for the account is listed in the page.

4. Click the Configure button of the VPC to which you want to deploy the VMs.

   The VPC page is displayed where all the tiers you created are listed in a diagram.

   For each tier, the following options are displayed.

   - Internal LB
   - Public LB IP
   - Static NAT
   - Virtual Machines
   - CIDR

   The following router information is displayed:

   - Private Gateways

- Public IP Addresses
- Site-to-Site VPNs
- Network ACL Lists

5. In the Router node, select Public IP Addresses.

   The IP Addresses page is displayed.

6. Click the IP you want to work with.

7. In the Details tab,click **Static NAT**. The button toggles between Enable and Disable, depending on whether static NAT is currently enabled for the IP address.

8. If you are enabling static NAT, a dialog appears as follows:

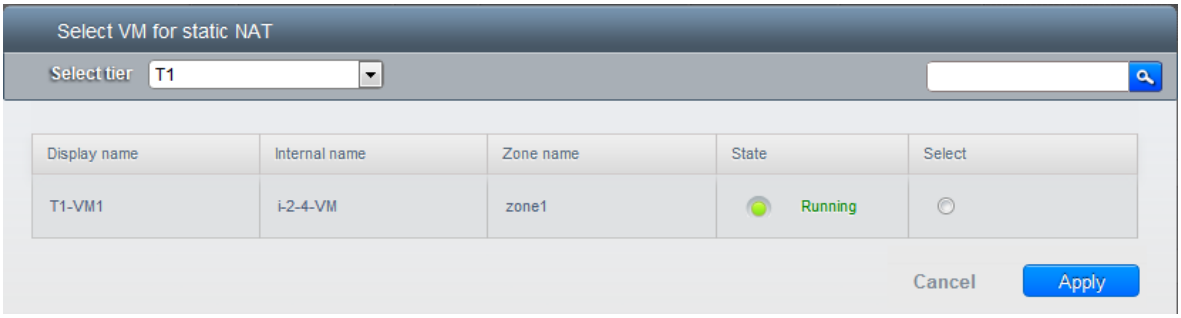

9. Select the tier and the destination VM, then click Apply.

9.24.11. Adding Load Balancing Rules on a VPC

In a VPC, you can configure two types of load balancing—public LB and internal LB. External LB is nothing but a LB rule created to redirect the traffic received at a public IP of the VPC virtual router. The traffic is load balanced within a tier based on your configuration. Citrix NetScaler and VPC virtual router are supported for public LB. When you use internal LB service, traffic received at a tier is load balanced across different VMs within that tier. For example, traffic reached at Web tier is redirected to another VM in that tier. External load balancing devices are not supported for internal LB. The service is provided by a internal LB VM configured on the target tier.

9.24.11.1. Load Balancing Public Traffic (Public LB)

A CloudPlatform user or administrator may create load balancing rules that balance traffic received at a public IP to one or more VMs that belong to a network tier that provides load balancing service in a VPC. A user creates a rule, specifies an algorithm, and assigns the rule to a set of VMs within a tier.

9.24.11.1.1. Enabling NetScaler as the LB Provider on a VPC Tier

1. Add and enable Netscaler VPX in dedicated mode.

   Netscaler can be used in a VPC environment only if it is in dedicated mode.

2. Create a network offering, as given in *Section 9.24.11.1.2, "Creating a Network Offering for Public LB"*.

3. Create a VPC with Netscaler as the Public LB provider.

For more information, see *Section 9.24.2, "Adding a Virtual Private Cloud"*.

4. For the VPC, acquire an IP.
5. Create an public load balancing rule and apply, as given in *Section 9.24.11.1.3, "Creating a Public LB Rule"*.

9.24.11.1.2. Creating a Network Offering for Public LB

To have public LB support on VPC, create a network offering as follows:

1. Log in to the CloudPlatform UI as a user or admin.
2. From the Select Offering drop-down, choose Network Offering.
3. Click Add Network Offering.
4. In the dialog, make the following choices:
   - **Name**: Any desired name for the network offering.
   - **Description**: A short description of the offering that can be displayed to users.
   - **Network Rate**: Allowed rate of data transfer in MB per second (megabits per second).
   - **Traffic Type**: The type of network traffic that will be carried on the network.
   - **Guest Type**: Choose whether the guest network is isolated or shared.
   - **Persistent**: Indicate whether the guest network is persistent or not. The network that you can provision without having to deploy a VM on it is termed persistent network.
   - **VPC**: This option indicate whether the guest network is Virtual Private Cloud-enabled. A Virtual Private Cloud (VPC) is a private, isolated part of CloudPlatform. A VPC can have its own virtual network topology that resembles a traditional physical network. For more information on VPCs, see *Section 9.24.1, "About Virtual Private Clouds"*.
   - **Specify VLAN**: (Isolated guest networks only) Indicate whether a VLAN should be specified when this offering is used.
   - **Supported Services**: Select Load Balancer. Use Netscaler or VpcVirtualRouter.
   - **Load Balancer Type**: Select Public LB from the drop-down.
   - **LB Isolation**: Select Dedicated if Netscaler is used as the public LB provider.
   - **System Offering**: Choose the system service offering that you want virtual routers to use in this network.
   - **Conserve mode**: Indicate whether to use conserve mode. In this mode, network resources are allocated only when the first virtual machine starts in the network.
   - **Disk Read Rate**: Allowed rate of disk read in bits per second (byte per second).
   - **Disk Write Rate**: Allowed rate of disk write in bits per second (byte per second).
5. Click OK and the network offering is created.

9.24.11.1.3. Creating a Public LB Rule

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPCs that you have created for the account is listed in the page.
4. Click the Configure button of the VPC, for which you want to configure load balancing rules.

   The VPC page is displayed where all the tiers you created listed in a diagram.

   For each tier, the following options are displayed:

   - Internal LB
   - Public LB IP
   - Static NAT
   - Virtual Machines
   - CIDR

   The following router information is displayed:

   - Private Gateways
   - Public IP Addresses
   - Site-to-Site VPNs
   - Network ACL Lists
5. In the Router node, select Public IP Addresses.

   The IP Addresses page is displayed.
6. Click the IP address for which you want to create the rule, then click the Configuration tab.
7. In the Load Balancing node of the diagram, click View All.
8. Select the tier to which you want to apply the rule.
9. Specify the following:

   - **Name**: A name for the load balancer rule.
   - **Public Port**: The port that receives the incoming traffic to be balanced.
   - **Private Port**: The port that the VMs will use to receive the traffic.
   - **Algorithm**. Choose the load balancing algorithm you want CloudPlatform to use. CloudPlatform supports the following well-known algorithms:

     - Round-robin
     - Least connections

    - Source
- **Stickiness**. (Optional) Click Configure and choose the algorithm for the stickiness policy. See Sticky Session Policies for Load Balancer Rules.
- **Add VMs**: Click Add VMs, then select two or more VMs that will divide the load of incoming traffic, and click Apply.

The new load balancing rule appears in the list. You can repeat these steps to add more load balancing rules for this IP address.

9.24.11.2. Load Balancing Tier-to-Tier traffic (Internal LB)

CloudPlatform supports sharing workload across different tiers within your VPC. Assume that multiple tiers are set up in your environment, such as Web tier and Application tier. Traffic to each tier is balanced on the VPC virtual router on the public side, as explained in *Section 9.24.11, "Adding Load Balancing Rules on a VPC"*. If you want the traffic coming from the Web tier to the Application tier to be balanced, use the internal load balancing feature offered by CloudPlatform.

9.24.11.2.1. How Does Internal LB Work in VPC?

In this figure, a public LB rule is created for the public IP 72.52.125.10 with public port 80 and private port 81. The LB rule, created on the VPC virtual router, is applied on the traffic coming from the Internet to the VMs on the Web tier. On the Application tier two internal load balancing rules are created. An internal LB rule for the guest IP 10.10.10.4 with load balancer port 23 and instance port 25 is configured on the VM, InternalLBVM1. Another internal LB rule for the guest IP 10.10.10.4 with load balancer port 45 and instance port 46 is configured on the VM, InternalLBVM1. Another internal LB rule for the guest IP 10.10.10.6, with load balancer port 23 and instance port 25 is configured on the VM, InternalLBVM2.

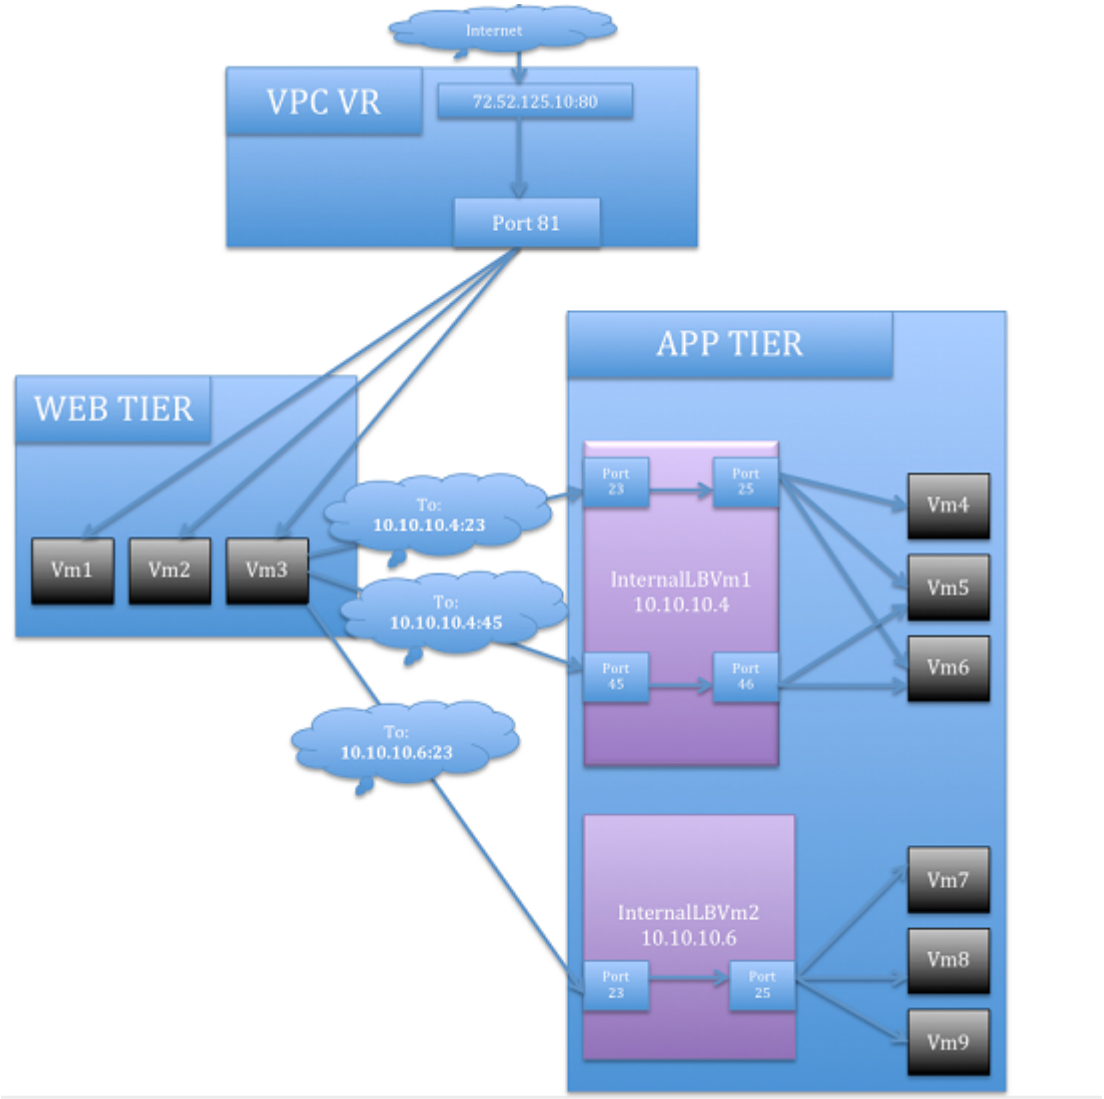

9.24.11.2.2. Enabling Internal LB on a VPC Tier

1. Create a network offering, as given in *Section 9.24.11.2.4, "Creating an Internal LB Rule"*.
2. Create an internal load balancing rule and apply, as given in *Section 9.24.11.2.4, "Creating an Internal LB Rule"*.

9.24.11.2.3. Creating a Network Offering for Internal LB

To have internal LB support on VPC, either use the default offering, DefaultIsolatedNetworkOfferingForVpcNetworksWithInternalLB, or create a network offering as follows:

1. Log in to the CloudPlatform UI as a user or admin.
2. From the Select Offering drop-down, choose Network Offering.
3. Click Add Network Offering.
4. In the dialog, make the following choices:
   - **Name**: Any desired name for the network offering.

- **Description**: A short description of the offering that can be displayed to users.
- **Network Rate**: Allowed rate of data transfer in MB per second (megabits per second).
- **Traffic Type**: The type of network traffic that will be carried on the network.
- **Guest Type**: Choose whether the guest network is isolated or shared.
- **Persistent**: Indicate whether the guest network is persistent or not. The network that you can provision without having to deploy a VM on it is termed persistent network.
- **VPC**: This option indicate whether the guest network is Virtual Private Cloud-enabled. A Virtual Private Cloud (VPC) is a private, isolated part of CloudPlatform. A VPC can have its own virtual network topology that resembles a traditional physical network. For more information on VPCs, see *Section 9.24.1, "About Virtual Private Clouds"*.
- **Specify VLAN**: (Isolated guest networks only) Indicate whether a VLAN should be specified when this offering is used.
- **Supported Services**: Select Load Balancer. Select **`InternalLbVM`** from the provider list.
- **Load Balancer Type**: Select Internal LB from the drop-down.
- **System Offering**: Choose the system service offering that you want virtual routers to use in this network.
- **Conserve mode**: Indicate whether to use conserve mode. In this mode, network resources are allocated only when the first virtual machine starts in the network.
- **Disk Read Rate**: Allowed rate of disk read in bits per second (byte per second).
- **Disk Write Rate**: Allowed rate of disk write in bits per second (byte per second).

5. Click OK and the network offering is created.

9.24.11.2.4. Creating an Internal LB Rule

When you create the Internal LB rule and applies to a VM, an Internal LB VM, which is responsible for load balancing, is created. You can view the created Internal LB VM in the Instances page if you navigate to **Infrastructure** > **Zones** > <*zone_ name*> > <*physical_network_name*> > **Network Service Providers** > **Internal LB VM**.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPCs that you have created for the account is listed in the page.
4. Locate the VPC for which you want to configure internal LB, then click Configure.

   The VPC page is displayed where all the tiers you created listed in a diagram.
5. Locate the Tier for which you want to configure an internal LB rule, click Internal LB.

   In the Internal LB page, click Add Internal LB.

6. In the dialog, specify the following:
   - **Name**: A name for the load balancer rule.
   - **Description**: A short description of the rule that can be displayed to users.
   - **Source IP Address**: The source IP from which traffic originates. The IP is acquired from the CIDR of that particular tier on which you want to create the Internal LB rule.

     For every Source IP, a new Internal LB VM is created for load balancing.
   - **Source Port**: The port associated with the source IP. Traffic on this port is load balanced.
   - **Instance Port**: The port of the internal LB VM.
   - **Algorithm**. Choose the load balancing algorithm you want CloudPlatform to use. CloudPlatform supports the following well-known algorithms:
     - Round-robin
     - Least connections
     - Source

9.24.12. Configuring Remote Access VPN in VPC

On enabling Remote Access VPN on a VPC, any VPN client present outside the VPC can access VMs present in the VPC by using the Remote VPN connection. The VPN client can be present anywhere except inside the VPC on which the user enabled the Remote Access VPN service.

9.24.12.1. Enabling or Disabling Remote Access VPN for a VPC

1. Log in as a user or administrator to the CloudPlatform UI.
2. In the left navigation, click Network.
3. In the Select view, select VPC.

   All the VPCs that you have created for the account is listed in the page.
4. Click the Configure button of the VPC.

   For each tier, the following options are displayed:
   - Internal LB
   - Public LB IP
   - Static NAT
   - Virtual Machines
   - CIDR

   The following router information is displayed:
   - Private Gateways
   - Public IP Addresses

- Site-to-Site VPNs
- Network ACL Lists

5. In the Router node, select Public IP Addresses.

   The IP Addresses page is displayed.

6. Click Source NAT IP address.

7. Click the Enable VPN button. 

   Click OK to confirm. The IPsec key is displayed in a pop-up window.

   To disable, click the Disable VPN button. 

8. Now, continue with *Section 9.24.12.2, "Adding Remote Access VPN Users"*.

9.24.12.2. Adding Remote Access VPN Users

1. Click the Source NAT IP.
2. Select the VPN tab.
3. Add the username and the corresponding password of the user you wanted to add.
4. Click Add.
5. Repeat the same steps to add the VPN users.

9.24.12.3. Removing Remote Access VPN Users

1. Click the Source NAT IP.
2. Select the VPN tab.
3. Locate the user you want to remove.
4. Click Delete User. 
5. Repeat the same steps to remove the VPN users.

9.24.13. Adding a Port Forwarding Rule on a VPC

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPCs that you have created for the account is listed in the page.

4. Click the Configure button of the VPC to which you want to deploy the VMs.

   The VPC page is displayed where all the tiers you created are listed in a diagram.

For each tier, the following options are displayed:

- Internal LB
- Public LB IP
- Static NAT
- Virtual Machines
- CIDR

The following router information is displayed:

- Private Gateways
- Public IP Addresses
- Site-to-Site VPNs
- Network ACL Lists

5. In the Router node, select Public IP Addresses.

   The IP Addresses page is displayed.

6. Click the IP address for which you want to create the rule, then click the Configuration tab.
7. In the Port Forwarding node of the diagram, click View All.
8. Select the tier to which you want to apply the rule.
9. Specify the following:

   - **Public Port**: The port to which public traffic will be addressed on the IP address you acquired in the previous step.
   - **Private Port**: The port on which the instance is listening for forwarded public traffic.
   - **Protocol**: The communication protocol in use between the two ports.
     - TCP
     - UDP
   - **Add VM**: Click Add VM. Select the name of the instance to which this rule applies, and click Apply.

     You can test the rule by opening an SSH session to the instance.

9.24.14. Removing Tiers

You can remove a tier from a VPC. A removed tier cannot be revoked. When a tier is removed, only the resources of the tier are expunged. All the network rules (port forwarding, load balancing and staticNAT) and the IP addresses associated to the tier are removed. The IP address still be belonging to the same VPC.

1. Log in to the CloudPlatform UI as an administrator or end user.

2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPC that you have created for the account is listed in the page.
4. Click the Configure button of the VPC for which you want to set up tiers.

   The Configure VPC page is displayed. Locate the tier you want to work with.
5. Select the tier you want to remove.
6. In the Network Details tab, click the Delete Network button. 

   Click Yes to confirm. Wait for some time for the tier to be removed.

9.24.15. Editing, Restarting, and Removing a Virtual Private Cloud

Note

Ensure that all the tiers are removed before you remove a VPC.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation, choose Network.
3. In the Select view, select VPC.

   All the VPCs that you have created for the account is listed in the page.
4. Select the VPC you want to work with.
5. In the Details tab, click the Remove VPC button 

   You can remove the VPC by also using the remove button in the Quick View.

   You can edit the name and description of a VPC. To do that, select the VPC, then click the Edit button. 

   To restart a VPC, select the VPC, then click the Restart button. 

9.25. Persistent Networks

The network that you can provision without having to deploy any VMs on it is called a persistent network. A persistent network can be part of a VPC or a non-VPC environment.

When you create other types of network, a network is only a database entry until the first VM is created on that network. When the first VM is created, a VLAN ID is assigned and the network is provisioned. Also, when the last VM is destroyed, the VLAN ID is released and the network is no

longer available. With the addition of persistent network, you will have the ability to create a network in CloudPlatform in which physical devices can be deployed without having to run any VMs. Additionally, you can deploy physical devices on that network.

One of the advantages of having a persistent network is that you can create a VPC with a tier consisting of only physical devices. For example, you might create a VPC for a three-tier application, deploy VMs for Web and Application tier, and use physical machines for the Database tier. Another use case is that if you are providing services by using physical hardware, you can define the network as persistent and therefore even if all its VMs are destroyed the services will not be discontinued.

9.25.1. Persistent Network Considerations

- Persistent network is designed for Isolated networks.
- All default network offerings are non-persistent.
- A network offering cannot be editable because changing it affects the behavior of the existing networks that were created using this network offering.
- When you create a guest network, the network offering that you select defines the network persistence. This in turn depends on whether persistent network is enabled in the selected network offering.
- An existing network can be made persistent by changing its network offering to an offering that has the Persistent option enabled. While setting this property, even if the network has no running VMs, the network is provisioned.
- An existing network can be made non-persistent by changing its network offering to an offering that has the Persistent option disabled. If the network has no running VMs, during the next network garbage collection run the network is shut down.
- When the last VM on a network is destroyed, the network garbage collector checks if the network offering associated with the network is persistent, and shuts down the network only if it is non-persistent.

9.25.2. Creating a Persistent Guest Network

To create a persistent network, perform the following:

1. Create a network offering with the Persistent option enabled.

   See *Section 5.1, “Creating a New Network Offering”*.

2. Select Network from the left navigation pane.
3. Select the guest network that you want to offer this network service to.
4. Click the Edit button.
5. From the Network Offering drop-down, select the persistent network offering you have just created.
6. Click OK.

Working with Templates

A template is a reusable configuration for virtual machines. When users launch VMs, they can choose from a list of templates in CloudPlatform.

Specifically, a template is a virtual disk image that includes one of a variety of operating systems, optional additional software such as office applications, and settings such as access control to determine who can use the template. Each template is associated with a particular type of hypervisor, which is specified when the template is added to CloudPlatform.

CloudPlatform ships with a default template. In order to present more choices to users, CloudPlatform administrators and users can create templates and add them to CloudPlatform.

For conceptual information about templates in CloudPlatform , refer to **Chapter 10: About Templates in CloudPlatform** in the *Citrix CloudPlatform 4.5 (powered by Apache CloudStack) Concepts Guide*.

10.1. Creating Templates: Overview

CloudPlatform ships with a default template for the CentOS operating system. There are a variety of ways to add more templates. Administrators and end users can add templates. The typical sequence of events is:

1. Launch a VM instance that has the operating system you want. Make any other desired configuration changes to the VM.
2. Stop the VM.
3. Convert the volume into a template.

There are other ways to add templates to CloudPlatform. For example, you can take a snapshot of the VM's volume and create a template from the snapshot, or import a VHD from another system into CloudPlatform.

The various techniques for creating templates are described in the next few sections.

10.2. Requirements for Templates

- For XenServer, install PV drivers / Xen tools on each template that you create. This will enable live migration, clean guest shutdown, and dynamic scaling.
- For vSphere, install VMware Tools on each template that you create. This will enable dynamic scaling and console view to work properly.
- For Hyper-V, a default UI-enabled CentOS template is provided for guest VMs. For any other templates, you can install whatever packages you need and create a template similar to the way you do for any other hypervisor.
- For Linux VMs that run on KVM, the drivers are built into the OS.
- Do not install **`cloudservice.exe`** on the Windows OS templates which are not password enabled.

10.3. Best Practices for Templates

If you plan to use large templates (100 GB or larger), be sure you have a 10-gigabit network to support the large templates. A slower network can lead to timeouts and other errors when large templates are used.

10.4. Creating a Template from an Existing Virtual Machine

Once you have at least one VM set up in the way you want, you can use it as the prototype for other VMs.

1. Create and start a virtual machine using any of the techniques given in *Section 6.2, "Creating Virtual Machines "*.
2. Make any desired configuration changes on the running VM, then click Stop.
3. Wait for the VM to stop. When the status shows Stopped, go to the next step.
4. Click Create Template and provide the following:
   - **Name and Display Text**. These will be shown in the UI, so choose something descriptive.
   - **OS Type**. This helps CloudPlatform and the hypervisor perform certain operations and make assumptions that improve the performance of the guest. Select one of the following.
     - If the operating system of the stopped VM is listed, choose it.
     - If the OS type of the stopped VM is not listed, choose Other.
     - If you want to boot from this template in PV mode, choose Other PV (32-bit) or Other PV (64-bit). This choice is available only for XenServer:

       Note

       Generally you should not choose an older version of the OS than the version in the image. For example, choosing CentOS 5.4 to support a CentOS 6.2 image will in general not work. In those cases you should choose Other.

   - **Public**. Choose Yes to make this template accessible to all users of this CloudPlatform installation. The template will appear in the Community Templates list.
   - **Password Enabled**. Choose Yes if your template has the CloudPlatform password change script installed. See *Section 10.12, "Adding Password Management to Your Templates"*.
5. Click Add.

The new template will be visible in the Templates section when the template creation process has been completed. The template is then available when creating a new VM.

10.5. Creating a Template from a Virtual Machine that is Stopped

You can create a template based on the root volume of the virtual machine that you have stopped.

1. In the CloudPlatform UI, on the left-side panel, click **Instances**.
2. On the right-side panel, identify the VM based on which you want to create the volume and ensure that its state is displayed as **Stopped**.

3. On the left-side panel, click **Storage**.
4. On the right-side panel, in the **Name** column, click the name of the root volume of the VM that you stopped.
5. In the **Details** page, click the **Create template** icon.
6. In the **Create template** panel, enter the following:
   - **Name**: Enter the name that you want to display on the UI for this template.
   - **Description**: Enter a brief description that you want to display on the UI for this template.
   - **OS Type**: Select the OS type. For more information on the OS type, refer to *Section 10.4, "Creating a Template from an Existing Virtual Machine"*.
   - **Public**: Select this check box for making this template accessible to all users. A pubic template will appear in the Community Templates list.
   - **Password Enabled**: Select this check box if your template has the CloudPlatform password change script installed.
   - **Featured**: Select this check box if you want this template to be more prominent for users to select. A featured template will appear in the Featured Templates list.
   - **Dynamically Scalable**: Select this check box if you have XenServer/VMware tools installed on this template and if the template supports dynamic scaling.
7. Click **OK**.

10.6. Creating a Template from a Snapshot

Note

Not supported by Oracle VM and Hyper-V.

If you do not want to stop the VM to use the Create Template menu item, as described in *Section 10.4, "Creating a Template from an Existing Virtual Machine"*, you can create a template directly from any snapshot through the CloudPlatform UI.

10.7. Uploading Templates

Note

If you are uploading a template that was created using vSphere Client, be sure the OVA file does not contain an ISO. If it does, the deployment of VMs from the template will fail.

Templates are uploaded based on a URL. HTTP is the supported access protocol. Templates are frequently large files. You can optionally gzip them to decrease upload time.

To upload a template:

1. In the CloudPlatform UI, on the left navigation bar, click **Templates**.
2. On the right side panel, click **Register Template**.
3. In the **Register Template** dialog box, provide the following details:
   - **Name and Description**. Enter a unique name to identify the template. Also, provide a description about the template. These will be displayed on the UI. So, you must use descriptive name and description..
   - **URL**. The Management Server will download the file from the specified URL, such as http://my.web.server/filename.vhd.gz.
   - **Zone**. Choose the zone where you want the template to be available. If your CloudPlatform deployment includes multiple zones running the same hypervisor (the one selected in the Hypervisor field later in this dialog box), and you want the template to be available in all of them, choose All Zones.
   - **OS Type**: This helps CloudPlatform and the hypervisor perform certain operations and make assumptions that improve the performance of the guest. Select one of the following:
     - If the operating system of the stopped VM is listed, choose it.
     - If the OS type of the stopped VM is not listed, choose Other.

     

     Note

     Generally you should not choose an older version of the OS than the version in the image. For example, choosing CentOS 5.4 to support a CentOS 6.2 image will in general not work. In those cases you should choose Other.
   - **Hypervisor**: The supported hypervisors are listed. Select the desired one.
   - **Format**. The format of the template upload file, such as VHD or OVA.
   - **Password Enabled**. Choose Yes if your template has the CloudPlatform password change script installed. See Adding Password Management to Your Templates
   - **Extractable**. Select **Yes** if the template and its derivatives (templates and the volumes created from this template) are available for extraction. If this option is selected, end users can download the full image of a template and its derivatives.
   - **Public**. Choose Yes to make this template accessible to all users of this CloudPlatform installation. The template will appear in the Community Templates list.
   - **Featured**. Choose Yes if you would like this template to be more prominent for users to select. The template will appear in the Featured Templates list. Only an administrator can make a template Featured.

- **Dynamically Scalable**: Select this check box if you have XenServer/VMware tools installed on this template and if the template supports dynamic scaling.
- **Routing**: Select this check box if you want to use this template for deploying routers.

10.8. Exporting Templates

End users and Administrators may export templates from the CloudPlatform. Navigate to the template in the UI and choose the Download function from the Actions menu.

10.9. Creating a Windows Template

Windows templates must be prepared with Sysprep before they can be provisioned on multiple machines. Sysprep allows you to create a generic Windows template and avoid any possible SID conflicts.

Note

(XenServer) Windows VMs running on XenServer require PV drivers, which may be provided in the template or added after the VM is created. The PV drivers are necessary for essential management functions such as mounting additional volumes and ISO images, live migration, and graceful shutdown. Do not confuse with the OS type, Windows PV. Choose the OS type that most closely resembles the underlying OS.

Similarly, VMs running on vSphere require VMWare Tools, which may be provided in the template or added after the VM is created. Installing VMWare Tools improves network performance, storage performance, and overall performance of VMs. Also, VMWare tools ensure that the console functions seamlessly. This helps resolve the improper functioning of mouse, which you may experience after you connect to the VMs using a VNC (Virtual Network Computing) client.

An overview of the procedure is as follows:

1. Upload your Windows ISO.

   For more information, see *Section 10.14, “Adding an ISO”*.

2. Create a VM Instance with this ISO.

   For more information, see *Section 6.2, “Creating Virtual Machines ”*.

3. Follow the steps in Sysprep for Windows Server 2008 R2 (below) or Sysprep for Windows Server 2003 R2, depending on your version of Windows Server

4. The preparation steps are complete. Now you can actually create the template as described in Creating the Windows Template.

[1] http://www.microsoft.com/en-us/download/details.aspx?id=9085

10.9.1. System Preparation for Windows Server 2008 R2

For Windows 2008 R2, you run Windows System Image Manager to create a custom sysprep response XML file. Windows System Image Manager is installed as part of the Windows Automated Installation Kit (AIK). Windows AIK can be downloaded from *Microsoft Download Center*[1].

Use the following steps to run sysprep for Windows 2008 R2:

Note

The steps outlined here are derived from the excellent guide by Charity Shelbourne, originally published at *Windows Server 2008 Sysprep Mini-Setup.*[2]

1. Download and install the Windows AIK

   

   Note

   Windows AIK should not be installed on the Windows 2008 R2 VM you just created. Windows AIK should not be part of the template you create. It is only used to create the sysprep answer file.

2. Copy the install.wim file in the \sources directory of the Windows 2008 R2 installation DVD to the hard disk. This is a very large file and may take a long time to copy. Windows AIK requires the WIM file to be writable.
3. Start the Windows System Image Manager, which is part of the Windows AIK.
4. In the Windows Image pane, right click the Select a Windows image or catalog file option to load the install.wim file you just copied.
5. Select the Windows 2008 R2 Edition.

   You may be prompted with a warning that the catalog file cannot be opened. Click Yes to create a new catalog file.
6. In the Answer File pane, right click to create a new answer file.
7. Generate the answer file from the Windows System Image Manager using the following steps:

   a. The first page you need to automate is the Language and Country or Region Selection page. To automate this, expand Components in your Windows Image pane, right-click and add the Microsoft-Windows-International-Core setting to Pass 7 oobeSystem. In your Answer File pane, configure the InputLocale, SystemLocale, UILanguage, and UserLocale with the appropriate settings for your language and country or region. Should you have a question

[2] http://blogs.technet.com/askcore/archive/2008/10/31/automating-the-oobe-process-during-windows-server-2008-sysprep-mini-setup.aspx

about any of these settings, you can right-click on the specific setting and select Help. This will open the appropriate CHM help file with more information, including examples on the setting you are attempting to configure.

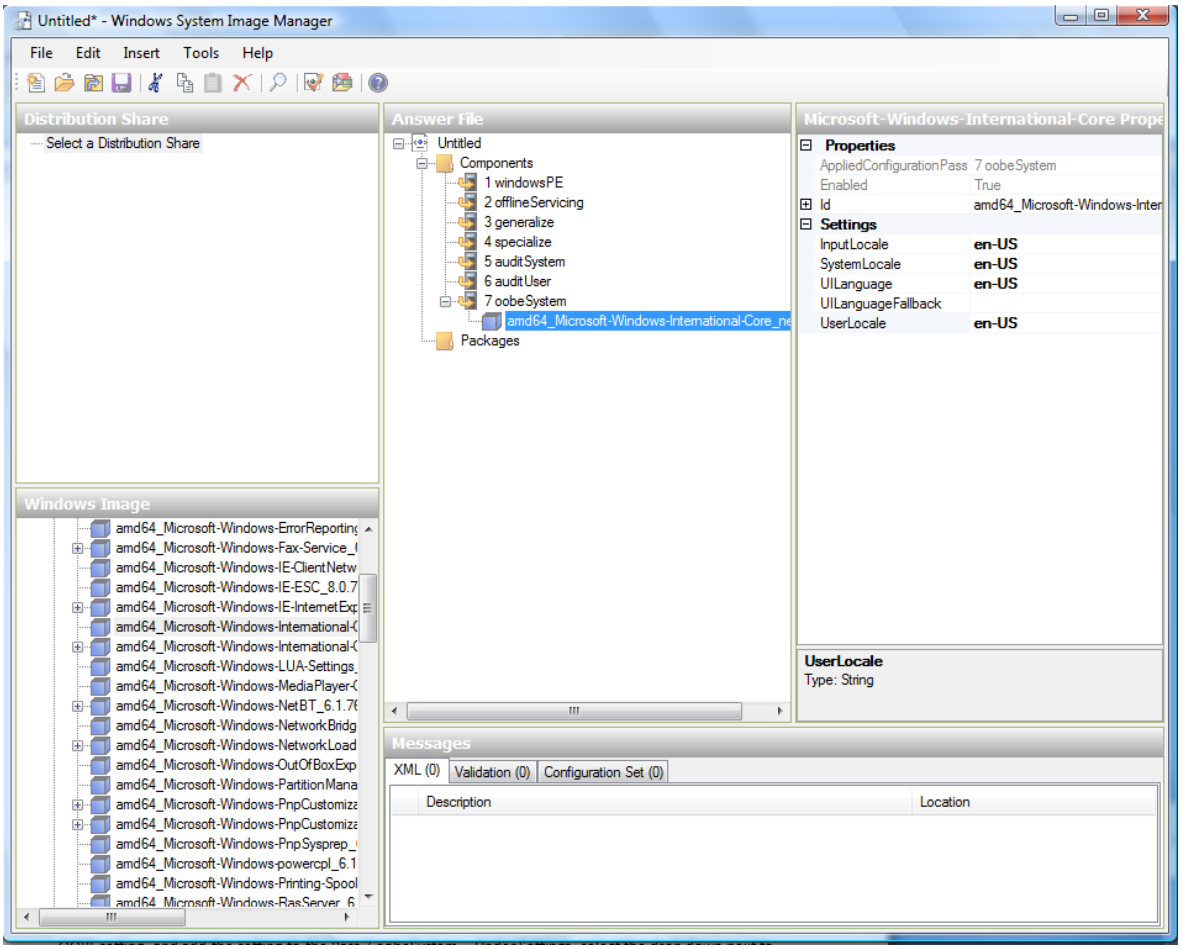

b. You need to automate the Software License Terms Selection page, otherwise known as the End-User License Agreement (EULA). To do this, expand the Microsoft-Windows-Shell-Setup component. High-light the OOBE setting, and add the setting to the Pass 7 oobeSystem. In Settings, set HideEULAPage true.

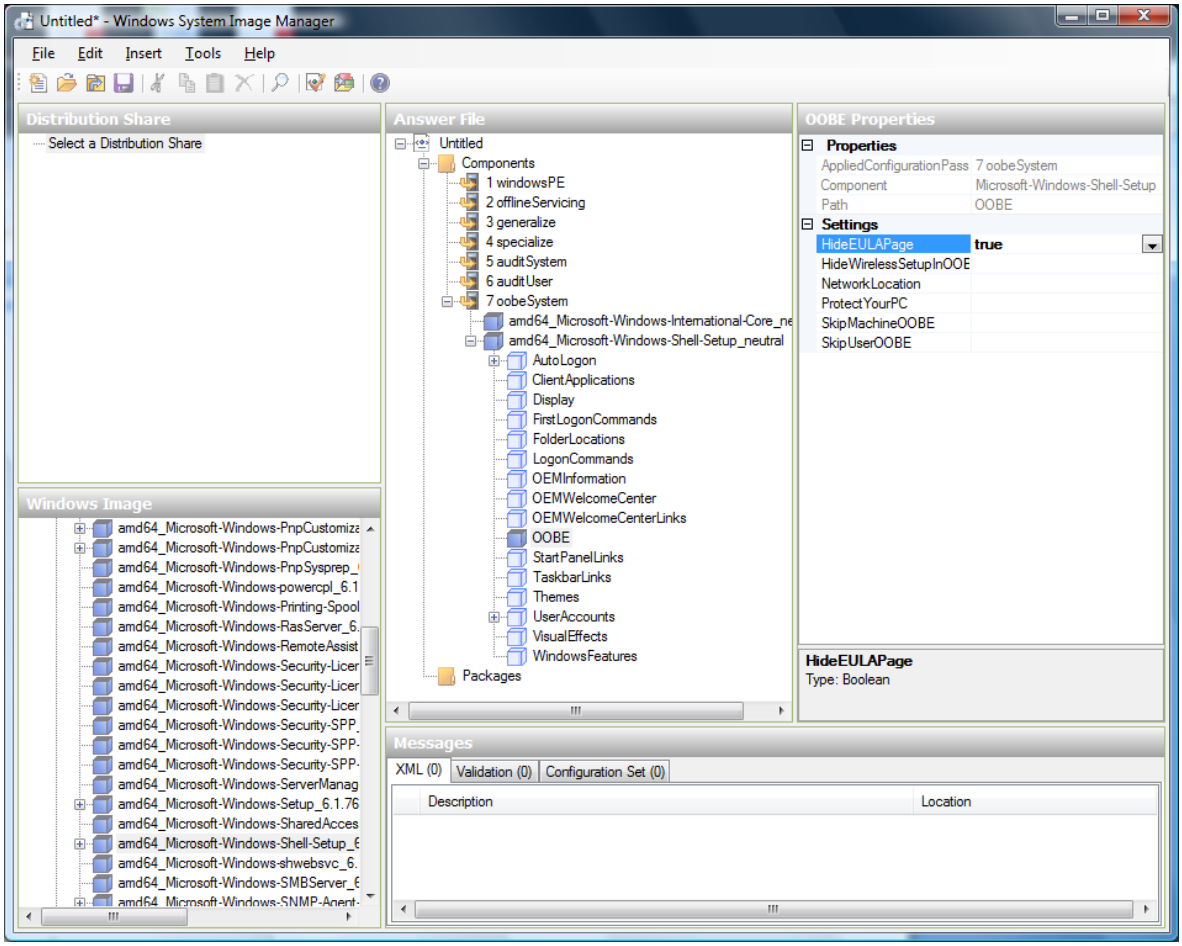

c. Make sure the license key is properly set. If you use MAK key, you can just enter the MAK key on the Windows 2008 R2 VM. You need not input the MAK into the Windows System Image Manager. If you use KMS host for activation you need not enter the Product Key. Details of Windows Volume Activation can be found at *http://technet.microsoft.com/en-us/library/bb892849.aspx*

d. You need to automate is the Change Administrator Password page. Expand the Microsoft-Windows-Shell-Setup component (if it is not still expanded), expand UserAccounts, right-click on AdministratorPassword, and add the setting to the Pass 7 oobeSystem configuration pass of your answer file. Under Settings, specify a password next to Value.

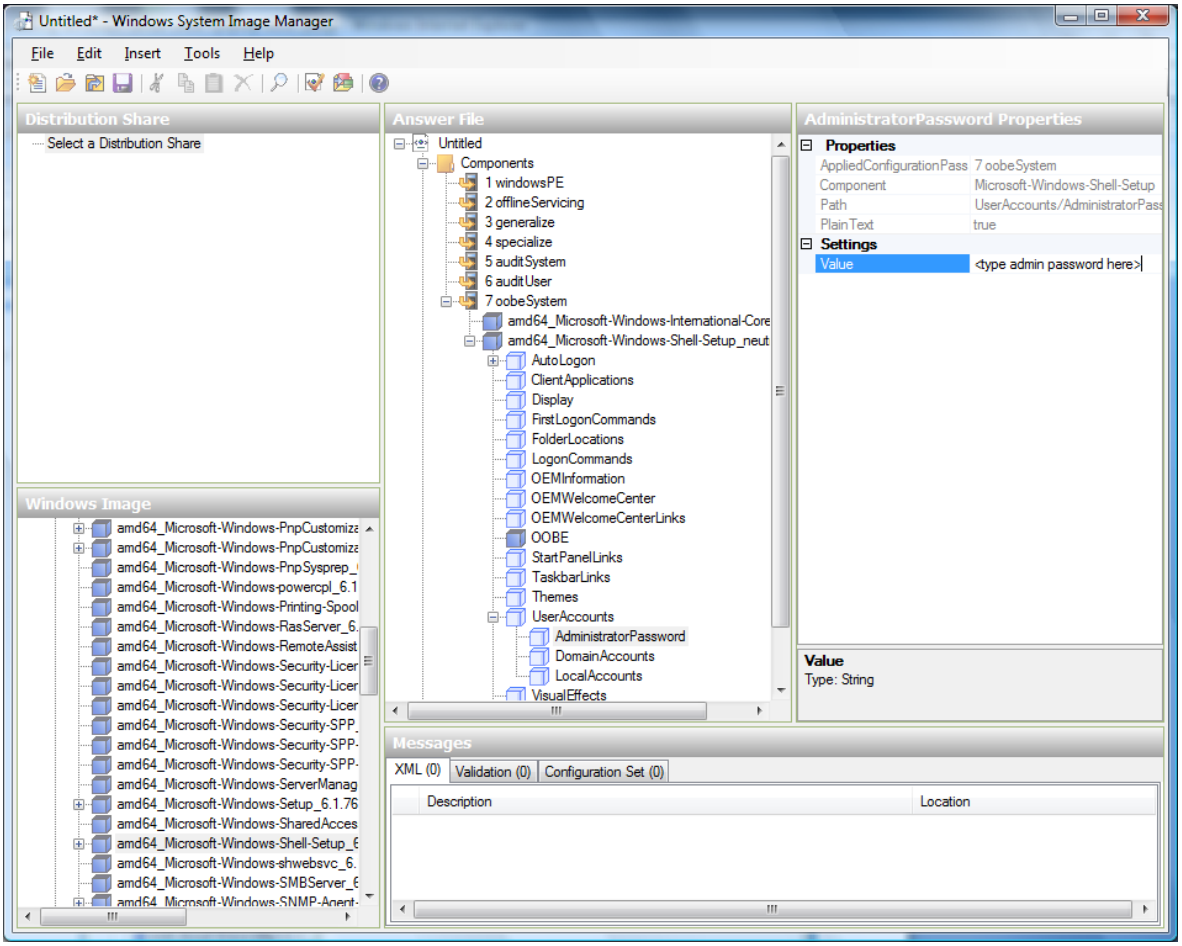

You may read the AIK documentation and set many more options that suit your deployment. The steps above are the minimum needed to make Windows unattended setup work.

8. Save the answer file as unattend.xml. You can ignore the warning messages that appear in the validation window.

9. Copy the unattend.xml file into the c:\windows\system32\sysprep directory of the Windows 2008 R2 Virtual Machine

10. Once you place the unattend.xml file in c:\windows\system32\sysprep directory, you run the sysprep tool as follows:

```
cd c:\Windows\System32\sysprep
sysprep.exe /oobe /generalize /shutdown
```

The Windows 2008 R2 VM will automatically shut down after sysprep is complete.

10.9.2. System Preparation for Windows Server 2003 R2

Earlier versions of Windows have a different sysprep tool. Follow these steps for Windows Server 2003 R2.

1. Extract the content of \support\tools\deploy.cab on the Windows installation CD into a directory called c:\sysprep on the Windows 2003 R2 VM.

2. Run c:\sysprep\setupmgr.exe to create the sysprep.inf file.

a. Select Create New to create a new Answer File.

b. Enter “Sysprep setup” for the Type of Setup.

c. Select the appropriate OS version and edition.

d. On the License Agreement screen, select “Yes fully automate the installation”.

e. Provide your name and organization.

f. Leave display settings at default.

g. Set the appropriate time zone.

h. Provide your product key.

i. Select an appropriate license mode for your deployment

j. Select “Automatically generate computer name”.

k. Type a default administrator password. If you enable the password reset feature, the users will not actually use this password. This password will be reset by the instance manager after the guest boots up.

l. Leave Network Components at “Typical Settings”.

m. Select the “WORKGROUP” option.

n. Leave Telephony options at default.

o. Select appropriate Regional Settings.

p. Select appropriate language settings.

q. Do not install printers.

r. Do not specify “Run Once commands”.

s. You need not specify an identification string.

t. Save the Answer File as c:\sysprep\sysprep.inf.

3. Run the following command to sysprep the image:

```
c:\sysprep\sysprep.exe -reseal -mini -activated
```

After this step the machine will automatically shut down

10.10. Importing Amazon Machine Images

The following procedures describe how to import an Amazon Machine Image (AMI) into CloudPlatform when using the XenServer hypervisor.

Assume you have an AMI file and this file is called CentOS_6.2_x64. Assume further that you are working on a CentOS host. If the AMI is a Fedora image, you need to be working on a Fedora host initially.

You need to have a XenServer host with a file-based storage repository (either a local ext3 SR or an NFS SR) to convert to a VHD once the image file has been customized on the Centos/Fedora host.

Note

When copying and pasting a command, be sure the command has pasted as a single line before executing. Some document viewers may introduce unwanted line breaks in copied text.

To import an AMI:

1. Set up loopback on the image file:

```
# mkdir -p /mnt/loop/centos62
# mount -o loop  CentOS_6.2_x64 /mnt/loop/centos54
```

2. Install the kernel-xen package into the image. This downloads the PV kernel and ramdisk to the image.

```
# yum -c /mnt/loop/centos54/etc/yum.conf --installroot=/mnt/loop/centos62/ -y install
 kernel-xen
```

3. Create a grub entry in /boot/grub/grub.conf.

```
# mkdir -p /mnt/loop/centos62/boot/grub
# touch /mnt/loop/centos62/boot/grub/grub.conf
# echo "" > /mnt/loop/centos62/boot/grub/grub.conf
```

4. Determine the name of the PV kernel that has been installed into the image.

```
# cd /mnt/loop/centos62
# ls lib/modules/
2.6.16.33-xenU  2.6.16-xenU  2.6.18-164.15.1.el5xen  2.6.18-164.6.1.el5.centos.plus
 2.6.18-xenU-ec2-v1.0  2.6.21.7-2.fc8xen  2.6.31-302-ec2
# ls boot/initrd*
boot/initrd-2.6.18-164.6.1.el5.centos.plus.img boot/initrd-2.6.18-164.15.1.el5xen.img
# ls boot/vmlinuz*
boot/vmlinuz-2.6.18-164.15.1.el5xen  boot/vmlinuz-2.6.18-164.6.1.el5.centos.plus  boot/
vmlinuz-2.6.18-xenU-ec2-v1.0  boot/vmlinuz-2.6.21-2952.fc8xen
```

   Xen kernels/ramdisk always end with "xen". For the kernel version you choose, there has to be an entry for that version under lib/modules, there has to be an initrd and vmlinuz corresponding to that. Above, the only kernel that satisfies this condition is 2.6.18-164.15.1.el5xen.

5. Based on your findings, create an entry in the grub.conf file. Below is an example entry.

```
default=0
timeout=5
hiddenmenu
title CentOS (2.6.18-164.15.1.el5xen)
        root (hd0,0)
        kernel /boot/vmlinuz-2.6.18-164.15.1.el5xen ro root=/dev/xvda
        initrd /boot/initrd-2.6.18-164.15.1.el5xen.img
```

6. Edit etc/fstab, changing “sda1” to “xvda” and changing “sdb” to “xvdb”.

```
# cat etc/fstab
/dev/xvda  /         ext3    defaults        1 1
/dev/xvdb  /mnt      ext3    defaults        0 0
none       /dev/pts  devpts  gid=5,mode=620  0 0
none       /proc     proc    defaults        0 0
none       /sys      sysfs   defaults        0 0
```

7. Enable login via the console. The default console device in a XenServer system is xvc0. Ensure that etc/inittab and etc/securetty have the following lines respectively:

```
# grep xvc0 etc/inittab
co:2345:respawn:/sbin/agetty xvc0 9600 vt100-nav
# grep xvc0 etc/securetty
xvc0
```

8. Ensure the ramdisk supports PV disk and PV network. Customize this for the kernel version you have determined above.

```
# chroot /mnt/loop/centos54
# cd /boot/
# mv initrd-2.6.18-164.15.1.el5xen.img initrd-2.6.18-164.15.1.el5xen.img.bak
# mkinitrd -f /boot/initrd-2.6.18-164.15.1.el5xen.img --with=xennet --preload=xenblk --
omit-scsi-modules 2.6.18-164.15.1.el5xen
```

9. Change the password.

```
# passwd
Changing password for user root.
New UNIX password:
Retype new UNIX password:
passwd: all authentication tokens updated successfully.
```

10. Exit out of chroot.

```
# exit
```

11. Check etc/ssh/sshd_config for lines allowing ssh login using a password.

```
# egrep "PermitRootLogin|PasswordAuthentication" /mnt/loop/centos54/etc/ssh/sshd_config
PermitRootLogin yes
PasswordAuthentication yes
```

12. If you need the template to be enabled to reset passwords from the CloudPlatform UI or API, install the password change script into the image at this point. See *Section 10.12, “Adding Password Management to Your Templates”*.

13. Unmount and delete loopback mount.

```
# umount /mnt/loop/centos54
# losetup -d /dev/loop0
```

14. Copy the image file to your XenServer host's file-based storage repository. In the example below, the Xenserver is "xenhost". This XenServer has an NFS repository whose uuid is a9c5b8c8-536b-a193-a6dc-51af3e5ff799.

```
# scp CentOS_6.2_x64 xenhost:/var/run/sr-mount/a9c5b8c8-536b-a193-a6dc-51af3e5ff799/
```

15. Log in to the Xenserver and create a VDI the same size as the image.

```
[root@xenhost ~]# cd /var/run/sr-mount/a9c5b8c8-536b-a193-a6dc-51af3e5ff799
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]#  ls -lh CentOS_6.2_x64
-rw-r--r-- 1 root root 10G Mar 16 16:49 CentOS_6.2_x64
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]# xe vdi-create virtual-size=10GiB sr-
uuid=a9c5b8c8-536b-a193-a6dc-51af3e5ff799 type=user name-label="Centos 6.2 x86_64"
cad7317c-258b-4ef7-b207-cdf0283a7923
```

16. Import the image file into the VDI. This may take 10–20 minutes.

```
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]# xe vdi-import
 filename=CentOS_6.2_x64 uuid=cad7317c-258b-4ef7-b207-cdf0283a7923
```

17. Locate a the VHD file. This is the file with the VDI's UUID as its name. Compress it and upload it to your web server.

```
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]# bzip2 -c cad7317c-258b-4ef7-b207-
cdf0283a7923.vhd > CentOS_6.2_x64.vhd.bz2
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]# scp CentOS_6.2_x64.vhd.bz2
 webserver:/var/www/html/templates/
```

10.11. Converting a Hyper-V VM to a Template

To convert a Hyper-V VM to a XenServer-compatible CloudPlatform template, you will need a standalone XenServer host with an attached NFS VHD SR. Use whatever XenServer version you are using with CloudPlatform, but use XenCenter 5.6 FP1 or SP2 (it is backwards compatible to 5.6). Additionally, it may help to have an attached NFS ISO SR.

For Linux VMs, you may need to do some preparation in Hyper-V before trying to get the VM to work in XenServer. Clone the VM and work on the clone if you still want to use the VM in Hyper-V. Uninstall Hyper-V Integration Components and check for any references to device names in /etc/fstab:

1. From the linux_ic/drivers/dist directory, run make uninstall (where "linux_ic" is the path to the copied Hyper-V Integration Components files).
2. Restore the original initrd from backup in /boot/ (the backup is named *.backup0).
3. Remove the "hdX=noprobe" entries from /boot/grub/menu.lst.
4. Check /etc/fstab for any partitions mounted by device name. Change those entries (if any) to mount by LABEL or UUID. You can get that information with the blkid command.

The next step is make sure the VM is not running in Hyper-V, then get the VHD into XenServer. There are two options for doing this.

Option one:

1. Import the VHD using XenCenter. In XenCenter, go to Tools>Virtual Appliance Tools>Disk Image Import.

2. Choose the VHD, then click Next.
3. Name the VM, choose the NFS VHD SR under Storage, enable "Run Operating System Fixups" and choose the NFS ISO SR.
4. Click Next, then Finish. A VM should be created.

Option two:

1. Run XenConvert, under From choose VHD, under To choose XenServer. Click Next.
2. Choose the VHD, then click Next.
3. Input the XenServer host info, then click Next.
4. Name the VM, then click Next, then Convert. A VM should be created.

Once you have a VM created from the Hyper-V VHD, prepare it using the following steps:

1. Boot the VM, uninstall Hyper-V Integration Services, and reboot.
2. Install XenServer Tools, then reboot.
3. Prepare the VM as desired. For example, run sysprep on Windows VMs. See *Section 10.9, "Creating a Windows Template"*.

Either option above will create a VM in HVM mode. This is fine for Windows VMs, but Linux VMs may not perform optimally. Converting a Linux VM to PV mode will require additional steps and will vary by distribution.

1. Shut down the VM and copy the VHD from the NFS storage to a web server; for example, mount the NFS share on the web server and copy it, or from the XenServer host use sftp or scp to upload it to the web server.
2. In CloudPlatform, create a new template using the following values:
   - URL. Give the URL for the VHD
   - OS Type. Use the appropriate OS. For PV mode on CentOS, choose Other PV (32-bit) or Other PV (64-bit). This choice is available only for XenServer.
   - Hypervisor. XenServer
   - Format. VHD

The template will be created, and you can create instances from it.

10.12. Adding Password Management to Your Templates

CloudPlatform provides an optional password reset feature that allows users to set a temporary admin or root password as well as reset the existing admin or root password from the CloudPlatform UI.

To enable the Reset Password feature, you will need to download an additional script to patch your template. When you later upload the template into CloudPlatform, you can specify whether reset admin/root password feature should be enabled for this template.

The password management feature works always resets the account password on instance boot. The script does an HTTP call to the virtual router to retrieve the account password that should be set. As long as the virtual router is accessible the guest will have access to the account password that should

be used. When the user requests a password reset the management server generates and sends a new password to the virtual router for the account. Thus an instance reboot is necessary to effect any password changes.

If the script is unable to contact the virtual router during instance boot it will not set the password but boot will continue normally.

10.12.1. Linux OS Installation

Use the following steps to begin the Linux OS installation:

1. Download the script file **`cloud-set-guest-password`**:
   - Linux: *http://download.cloud.com/templates/4.2/bindir/cloud-set-guest-password.in*
   - Windows: *http://sourceforge.net/projects/cloudstack/files/Password%20Management%20Scripts/CloudInstanceManager.msi/download*
2. Copy this file to **`/etc/init.d`**.

   On some Linux distributions, copy the file to **`/etc/rc.d/init.d`**.
3. Run the following command to make the script executable:

   ```
   chmod +x /etc/init.d/cloud-set-guest-password
   ```
4. Depending on the Linux distribution, continue with the appropriate step.

10.12.2. Windows OS Installation

Download the installer, CloudInstanceManager.msi, from the *Download page*[3] and run the installer in the newly created Windows VM.

10.13. Deleting Templates

Templates may be deleted. In general, when a template spans multiple Zones, only the copy that is selected for deletion will be deleted; the same template in other Zones will not be deleted. The provided CentOS template is an exception to this. If the provided CentOS template is deleted, it will be deleted from all Zones.

When templates are deleted, the VMs instantiated from them will continue to run. However, new VMs cannot be created based on the deleted template.

10.14. Adding an ISO

To make additional operating system or other software available for use with guest VMs, you can add an ISO. The ISO is typically thought of as an operating system image, but you can also add ISOs for other types of software, such as desktop applications that you want to be installed as part of a template.

1. Log in to the CloudPlatform UI as an administrator or end user.
2. In the left navigation bar, click Templates.

[3] http://cloudstack.org/download.html

3. In Select View, choose ISOs.
4. Click Add ISO.
5. In the Add ISO screen, provide the following:
   - **Name**: Short name for the ISO image. For example, CentOS 6.2 64-bit.
   - **Description**: Display test for the ISO image. For example, CentOS 6.2 64-bit.
   - **URL**: The URL that hosts the ISO image. The Management Server must be able to access this location via HTTP. If needed you can place the ISO image directly on the Management Server
   - **Zone**: Choose the zone where you want the ISO to be available, or All Zones to make it available throughout CloudPlatform.
   - **Bootable**: Whether or not a guest could boot off this ISO image. For example, a CentOS ISO is bootable, a Microsoft Office ISO is not bootable.
   - **OS Type**: This helps CloudPlatform and the hypervisor perform certain operations and make assumptions that improve the performance of the guest. Select one of the following.
     - If the operating system of your desired ISO image is listed, choose it.
     - If the OS Type of the ISO is not listed or if the ISO is not bootable, choose Other.
     - (XenServer only) If you want to boot from this ISO in PV mode, choose Other PV (32-bit) or Other PV (64-bit)
     - (KVM only) If you choose an OS that is PV-enabled, the VMs created from this ISO will have a SCSI (virtio) root disk. If the OS is not PV-enabled, the VMs will have an IDE root disk. The PV-enabled types are:

Fedora 13	Fedora 12	Fedora 11
Fedora 10	Fedora 9	Other PV
Debian GNU/Linux	CentOS 5.3	CentOS 5.4
CentOS 5.5	Red Hat Enterprise Linux 5.3	Red Hat Enterprise Linux 5.4
Red Hat Enterprise Linux 5.5	Red Hat Enterprise Linux 6	

Note

It is not recommended to choose an older version of the OS than the version in the image. For example, choosing CentOS 5.4 to support a CentOS 6.2 image will usually not work. In these cases, choose Other.

   - **Extractable**: Choose Yes if the ISO should be available for extraction.
   - **Public**: Choose Yes if this ISO should be available to other users.
   - **Featured**: Choose Yes if you would like this ISO to be more prominent for users to select. The ISO will appear in the Featured ISOs list. Only an administrator can make an ISO Featured.

6. Click OK.

   The Management Server will download the ISO. Depending on the size of the ISO, this may take a long time. The ISO status column will display Ready once it has been successfully downloaded into secondary storage. Clicking Refresh updates the download percentage.

7. **Important**: Wait for the ISO to finish downloading. If you move on to the next task and try to use the ISO right away, it will appear to fail. The entire ISO must be available before CloudPlatform can work with it.

10.15. Attaching an ISO to a VM

1. In the left navigation, click Instances.
2. Choose the virtual machine you want to work with.
3. Click the Attach ISO button. 
4. In the Attach ISO dialog box, select the desired ISO.
5. Click OK.

10.16. Changing a VM's Base Image

Every VM is created from a base image, which is a template or ISO which has been created and stored in CloudPlatform. Both cloud administrators and end users can create and modify templates, ISOs, and VMs.

In CloudPlatform, you can change an existing VM's base image from one template to another, or from one ISO to another. (You can not change from an ISO to a template, or from a template to an ISO).

For example, suppose there is a template based on a particular operating system, and the OS vendor releases a software patch. The administrator or user naturally wants to apply the patch and then make sure existing VMs start using it. Whether a software update is involved or not, it's also possible to simply switch a VM from its current template to any other desired template.

To change a VM's base image, call the restoreVirtualMachine API command and pass in the virtual machine ID and a new template ID. The template ID parameter may refer to either a template or an ISO, depending on which type of base image the VM was already using (it must match the previous type of image). When this call occurs, the VM's root disk is first destroyed, then a new root disk is created from the source designated in the template ID parameter. The new root disk is attached to the VM, and now the VM is based on the new template.

You can also omit the template ID parameter from the restoreVirtualMachine call. In this case, the VM's root disk is destroyed and recreated, but from the same template or ISO that was already in use by the VM.

Working with System Virtual Machines

CloudPlatform uses several types of system virtual machines to perform tasks in the cloud. In general CloudPlatform manages these system VMs and creates, starts, and stops them as needed based on scale and immediate needs. However, the administrator should be aware of them and their roles to assist in debugging issues.

For conceptual information about system VM template, console proxy, and virtual router, refer to **Chapter 9: About Virtual Machines in CloudPlatform and Chapter 10: About Templates in CloudPlatform** in the *Citrix CloudPlatform 4.5 (powered by Apache CloudStack) Concepts Guide*.

11.1. Configuring the Virtual Router

You can set the following:

- IP range
- Supported network services
- Default domain name for the network serviced by the virtual router
- Gateway IP address
- How often CloudPlatform fetches network usage statistics from CloudPlatform virtual routers. If you want to collect traffic metering data from the virtual router, set the global configuration parameter router.stats.interval. If you are not using the virtual router to gather network usage statistics, set it to 0.

11.2. Upgrading a Virtual Router with System Service Offerings

When CloudPlatform creates a virtual router, it uses default settings which are defined in a default system service offering. All the virtual routers in a single guest network use the same system service offering. You can upgrade the capabilities of the virtual router by creating and applying a custom system service offering.

1. Define your custom system service offering. See *Section 2.4, “Creating a New System Service Offering”*. In System VM Type, choose Domain Router.
2. Associate the system service offering with a network offering. See *Section 5.1, “Creating a New Network Offering”*
3. Apply the network offering to the network where you want the virtual routers to use the new system service offering. If this is a new network, follow the steps in Adding an Additional Guest Network on page 66. To change the service offering for existing virtual routers, follow the steps in *Section 5.2, “Changing the Network Offering on a Guest Network”*.

11.3. Best Practices for Virtual Routers

- WARNING: Restarting a virtual router from a hypervisor console deletes all the iptables rules. To work around this issue, stop the virtual router and start it from the CloudPlatform UI.
- WARNING: Do not use the destroyRouter API when only one router is available in the network, because restartNetwork API with the cleanup=false parameter can't recreate it later. If you want to

destroy and recreate the single router available in the network, use the restartNetwork API with the cleanup=true parameter.

11.4. Service Monitoring Tool for Virtual Router

The network service daemons running on the CloudPlatform virtual routers can be monitored by using the Service Monitoring tool. If a service goes down, the tool automatically restarts the service, and if that does not help bringing up the service, an event log is written to **`/var/log/messages`** indicating the failure. A new global parameter, **`network.router.enableservicemonitoring`**, has been introduced to control this feature. The default value is true, implies, monitoring is enabled. On changing this parameter, you need to restart the Management Server and the virtual router.

Monitoring tool can help to start a VR service, which is crashed due to an unexpected reason. For example:

- The services crashed due to the defects in the source code.
- The services that are terminated by the OS when memory or CPU is not sufficiently available for the service.

Note

Only those services with daemons are monitored. The services that are failed due to errors in the service/daemon configuration file cannot be restarted by the Monitoring tool.

VPC networks are not supported.

The following services are monitored in a VR:

- DNS
- DHCP
- SSH
- Apache Web Server
- Load balancing service

The following networks are supported:

- Isolated Networks
- Shared Networks in Advanced zone
- Shared Networks in Basic zone

Note

VPC networks are not supported.

This feature is supported on all the hypervisors.

11.5. Enhanced Upgrade for Virtual Routers

Upgrading VR is made flexible. The CloudPlatform administrators will be able to control the sequence of the VR upgrades. The sequencing is based on Infrastructure hierarchy, such as by Cluster, Pod, or Zone, and Administrative (Account) hierarchy, such as by Tenant or Domain. As an administrator, you can also determine when a particular customer service, such as VR, is upgraded within a specified upgrade interval. Upgrade operation is enhanced to increase the upgrade speed by allowing as many upgrade operations in parallel as possible.

During the entire duration of the upgrade, users cannot launch new services or make changes to an existing service.

Additionally, using multiple versions of VRs in a single instance is supported. In the Details tab of a VR, you can view the version and whether it requires upgrade. During the Management Server upgrade, CloudPlatform checks whether VR is at the latest version before performing any operation on the VR. To support this, a new global parameter, *router.version.check*, has been added. This parameter is set to true by default, which implies minimum required version is checked before performing any operation. No operation is performed if the VR is not at the required version. Services of the older version VR continue to be available, but no further operations can be performed on the VR until it is upgraded to the latest version. This will be a transient state until the VR is upgraded. This will ensure that the availability of VR services and VR state is not impacted due to the Management Server upgrade.

The following service will be available even if the VR is not upgraded. However, no changes for any of the services can be sent to the VR, until it is upgraded:

- SecurityGroup
- UserData
- DHCP
- DNS
- LB
- Port Forwarding
- VPN
- Static NAT
- Source NAT
- Firewall
- Gateway

- NetworkACL

11.5.1. Supported Virtual Routers

- VR
- VPC VR
- Redundant VR

11.5.2. Upgrading Virtual Routers

1. Download the latest System VM template.
2. By using the prepareTemplate API, download the latest System VM to all the primary storage pools.
3. Upgrade the Management Server.
4. Upgrade CPVM and SSVM either from the UI or by using the following script:

   To upgrade the system VMs using the UI, you must stop and start the system VMs using the **`cloudstack-sysvmadm`** script.

   ```
   # cloudstack-sysvmadm -d <IP address> -u cloud -p -s
   ```

   <IP Address> is the IP address of the cloud database server. If you have not specified this, it will display as root.

   Even when the VRs are still on older versions, existing services will continue to be available to the VMs. The Management Server cannot perform any operations on the VRs until they are upgraded.
5. Selectively upgrade the VRs:
   a. Log in to the CloudPlatform UI as the root administrator.
   b. In the left navigation, choose Infrastructure.
   c. On Virtual Routers, click View More.

      All the VRs are listed in the Virtual Routers page.
   d. In Select View drop-down, select desired grouping based on your requirement.

      You can use either of the following:

      - Group by zone
      - Group by pod
      - Group by cluster
      - Group by account
   e. Click the group which has the VRs to be upgraded.

      For example, if you have selected Group by zone, select the name of the desired zone.

f. Click the Upgrade button to upgrade all the VRs. 

g. Click OK to confirm.

11.6. Setting a Random System VM Password

You can log in to system virtual machines, just like any other VM, by using the procedure in *Section 6.3, "Accessing Virtual Machines "*. When you log in to a system VM, you will need to provide a password. For added security, it is recommended that you create a randomized system VM password rather than accepting the default that is provided when you install CloudPlatform.

To generate a random system VM password:

1. Set the global configuration setting system.vm.random.password to true.
2. Restart the Management Server.
3. To obtain the new password, view the global configuration setting system.vm.password.

The new password will remain the same even if you restart the Management Server again.

If you ever need to change the password, manually edit the CloudPlatform database. In the configuration table, remove the entry system.vm.password. The next time you restart the Management Server, assuming the value of the global configuration setting system.vm.random.password is still true, a new random password will be generated and stored in the database.

11.7. Secure Connections for CloudPlatform System VMs

CloudPlatform System VMs, such as console proxy and Secondary storage VMs, host HTTP connections. You can secure these connections by enabling SSL communication on them. For this, CloudPlatform has a built-in mechanism to use one SSL certificate for all the instances of system VMs. To achieve this, you need the following:

- A software that runs a wildcard DNS service.
- A wildcard certificate for this domain name.
  - Public certificate of root CA in PEM format
  - Public certificate(s) of intermediate CA(s) (if any) in PEM format
  - Wildcard domain certificate in PEM format
  - Private key in PKCS8 format

Note

Self-signed certificates are not supported.

- A domain, which can run a DNS service that is capable of resolving queries for addresses of the form aaa-bbb-ccc-ddd.yourdomain.com to an IPv4 IP address in the form aaa.bbb.ccc.ddd, for example, 202.8.44.1.

11.7.1. Replacing realhostip.com with Your Own Domain Name

To replace realhostip.com with your own domain name, you must do the following:

- Setup Your Domain Name in a DNS Server.
- Get a Signed Certificate from CA.
- Modify Global Configuration Parameters.
- Upload Custom Certificates.

11.7.1.1. Setting up Your Domain Name in a DNS Server

You must require a publicly resolvable DNS server for your domain that you want to use to replace realhostip.com. Use the following procedure to setup your own domain in your DNS server:

Note

You can set up your domain in any DNS server. However, the following procedure uses BIND as the DNS server.

In the following procedure, **yourhostip.com** is used as the new domain name. **realhostip.com** is configured in such a way that it converts every public IP address that is entered into CloudPlatform to a DNS name. For example, 77.88.99.11 maps to 77-88-99-11.realhostip.com. This is required for SSL. When the browser connects to this name, it matches the wildcard cert **`*.realhostip.com`**. Therefore, you must use a similar set up for your environment.

1. Set up your zone in your DNS server. In BIND 9, it appears as the following:

```
zone "yourhostip.com" IN { type master;
file "yourhostip.com.zone"; allow-update { none; };
};
```

2. Populate an A record for every public IP, which you have entered into CloudPlatform, that the console proxy can allocate. For example, you have a range from 55.66.77.100 to 55.66.77.200. In your zone file, you require a range similar to the following:

55-66-77-100	IN	A	55.66.77.100
55-66-77-101	IN	A	55.66.77.101
55-66-77-102	IN	A	55.66.77.102
55-66-77-103	IN	A	55.66.77.103
55-66-77-200	IN	A	55.66.77.200

11.7.1.2. Getting a Signed Certificate from CA

A signed wildcard certificate is required for your domain. You can obtain this certificate from any CA (for example, VeriSign). When you receive the certificate, ensure the following:

- Public certificate of root CA in PEM format.
- Public certificate(s) of intermediate CA(s) (if any) in PEM format.
- Wildcard domain certificate in PEM format.
- Private key in PKCS8 format.

11.7.1.3. Changing the Console Proxy SSL Certificate and Domain

If the administrator prefers, it is possible for the URL of the customer's console session to show a domain other than realhostip.com. The administrator can customize the displayed domain by selecting a different domain and uploading a new SSL certificate and private key. The domain must run a DNS service that is capable of resolving queries for addresses of the form aaa-bbb-ccc-ddd.your.domain to an IPv4 IP address in the form aaa.bbb.ccc.ddd, for example, 202.8.44.1. To change the console proxy domain, SSL certificate, and private key:

1. Set up dynamic name resolution or populate all possible DNS names in your public IP range into your existing DNS server with the format aaa-bbb-ccc-ddd.company.com -> aaa.bbb.ccc.ddd.
2. Generate the private key and certificate signing request (CSR). When you are using openssl to generate private/public key pairs and CSRs, for the private key that you are going to paste into the CloudPlatform UI, be sure to convert it into PKCS#8 format.
    a. Generate a new 2048-bit private key

    ```
    openssl genrsa -des3 -out yourprivate.key 2048
    ```

    b. Generate a new certificate CSR

    ```
    openssl req -new -key yourprivate.key -out yourcertificate.csr
    ```

    c. Head to the website of your favorite trusted Certificate Authority, purchase an SSL certificate, and submit the CSR. You should receive a valid certificate in return

    d. Convert your private key format into PKCS#8 encrypted format.

    ```
    openssl pkcs8 -topk8 -in yourprivate.key -out yourprivate.pkcs8.encrypted.key
    ```

    e. Convert your PKCS#8 encrypted private key into the PKCS#8 format that is compliant with CloudPlatform

    ```
    openssl pkcs8 -in yourprivate.pkcs8.encrypted.key -out yourprivate.pkcs8.key
    ```

3. In the Update SSL Certificate screen of the CloudPlatform UI, paste the following:
    - The certificate you've just generated.
    - The private key you've just generated.
    - The desired new domain name; for example, company.com

4. The desired new domain name; for example, company.com

   This stops all currently running console proxy VMs, then restarts them with the new certificate and key. Users might notice a brief interruption in console availability.

The Management Server generates URLs of the form "aaa-bbb-ccc-ddd.company.com" after this change is made. The new console requests will be served with the new DNS domain name, certificate, and key.

11.7.1.4. Updating Global Configuration Parameters

Update the following Global Configuration parameters:

- Ensure that the value for **`secstorage.encrypt.copy`** is **`true`**. By default, this value is set to **`false`**.

  `secstorage.encrypt.copy = true`

- Provide your domain name as the value for **`secstorage.ssl.cert.domain`**. For example, if your domain name is xyz.com, update the value for this parameter as follows:

  `secstorage.ssl.cert.domain = xyz.com`

- Provide your domain name as the value for **`consoleproxy.url.domain`**. For example, if your domain name is yourdomain.com, update the value for this parameter as follows:

  `consoleproxy.url.domain = *.yourdomain.com`

  

  Note

  Ensure that '*.' is added before the domain name.

- Restart Management Server after you update the Global Configuration parameters.

11.7.1.5. Uploading Custom Certificates using API

As part of uploading custom certificates, you can upload root and intermediate certificates using the API. Then, you can upload the SSL certificate through the CloudPlatform UI.

You must ensure the following:

- All System VMs must be in the "UP" state.
- The ID and the name must be unique for each upload.
- Typically, the session key and the domain suffix would be the same.

Note

Provide the full certificate path for the System VMs if you are using a certificate from an intermediate CA. The certificate path begins with the certificate of that certifying entity, and each certificate in the chain is signed by the entity identified by the next certificate in the chain. The chain terminates with a root CA certificate. For browsers to trust the site's certificate, you must specify the full chain: site certificate, intermediate CA, and root CA. Use the uploadCustomCertificate API calls for each level of the chain. The certificate and private key parameters need to have the full text in PEM encoded format. For example: **`'certificate':'-----BEGIN CERTIFICATE----- \nMIIDYTCCAkmgAwIBAgIQCgEBAQAAAnwasdfKasd`**

1. As the first step, you must upload **Root Certificate** as follows:

   ```
   http://IPAddress:8096/client/api?
   command=uploadCustomCertificate&idID&name=NAME&domainsuffix=DOMAIN_SUFFIX
                       &certificate=URL_ENCODED_CERTIFICATE
   ```

   ID is always 1 for the root certificate.

2. Optionally, you can upload **Intermediate Certificate**s that are part of the certificate chain.

   `http://IPAddress:8096/client/api?command=uploadCustomCertificate&id=ID&name=NAME&domainsuffix=DOMAIN_SUFFIX&certificate=URL_ENCODED_CERTIFICATE`

   ID must always be incremented. Repeat this step to upload all the intermediate certificates.

3. To upload the SSL Certificate using the CloudPlatform UI, do the following:

   a. Log-in to the CloudPlatform UI.

   b. On the left navigation bar, click **Infrastructure**.

   c. On the right side panel, click **SSL Certificate**.

   d. In the SSL Certificate panel, enter the details of the certificate, PKCS#8 Private Key, and the DNS Domain Suffix and click **OK**.

      The "Update SSL certificate succeeded" message appears.

      

      Note

      Certificate and Key should not have the URL encoded. SSVM and CPVM are re-booted to synchronize with the certificates.

11.7.1.6. Verification

You can do the following to ensure that the configuration is successful:

- **Console Proxy**: Check the console view of user VMs and ensure that it is working. The console must show the embedded iframe's source URL with the HTTPS protocol as configured.
- **Copy Template**: Ensure that you can copy a template from one zone to another.
- **Download Template/ISO/Volume**: The download URL must show the HTTPS protocol as configured and the download must succeed.

11.7.2. Load Balancing Console Proxy VMs

You can load balance console proxy VMs with an external load balancer such as Citrix NetScaler only if your cloud infrastructure contains a single zone. Load balancing console proxy VMs does not work if you have configured multiple zones as part of your cloud infrastructure.

To load balance console proxy VMs, do the following:

1. On an external LB device, such as Citrix Netscaler, configure an LB rule.
2. Map the LB rule to one of the public IPs from the public IP pool of Console Proxy.
3. In CloudPlatform, set the *`consoleproxy.url.domain`* parameter to *`xyz.yourdomain.com`* to perform LB on Console Proxy VMS.
4. Configure DNS server to resolve the specific domain name, xyz.yourdomain.com, to the LB public IP you have configured.
5. Restart the Management Server for the new settings to take effect.

   CloudPlatform sends a request as given below :

   ```
   # wget https://xyz.yourdomain.com/ajax?token=token
   ```

   When you open a Console Proxy VM, CloudPlatform sends the request to xyz.yourdomain.com, which is internally mapped to the public IP of the LB rule on the DNS server. DNS server forwards this request to the LB Public IP. When the external LB device receives, request is internally load balanced and forwarded to associated Console Proxy VMs.

Accounts

12.1. Accounts, Users, and Domains

Accounts

An account typically represents a customer of the service provider or a department in a large organization. Multiple users can exist in an account.

Domains

Accounts are grouped by domains. Domains usually contain multiple accounts that have some logical relationship to each other and a set of delegated administrators with some authority over the domain and its subdomains. For example, a service provider with several resellers could create a domain for each reseller.

For each account created, the Cloud installation creates three different types of user accounts: root administrator, domain administrator, and user.

Users

Users are like aliases in the account. Users in the same account are not isolated from each other, but they are isolated from users in other accounts. Most installations need not surface the notion of users; they just have one user per account. The same user cannot belong to multiple accounts.

Username is unique in a domain across accounts in that domain. The same username can exist in other domains, including sub-domains. Domain name can repeat only if the full pathname from root is unique. For example, you can create root/d1, as well as root/foo/d1, and root/sales/d1.

Administrators are accounts with special privileges in the system. There may be multiple administrators in the system. Administrators can create or delete other administrators, and change the password for any user in the system.

Domain Administrators

Domain administrators can perform administrative operations for users who belong to that domain. Domain administrators do not have visibility into physical servers or other domains.

Root Administrator

Root administrators have complete access to the system, including managing templates, service offerings, customer care administrators, and domains

Resource Ownership

Resources belong to the account, not individual users in that account. For example, billing, resource limits, and so on are maintained by the account, not the users. A user can operate on any resource in the account provided the user has privileges for that operation. The privileges are determined by the role. A root administrator can change the ownership of any virtual machine from one account to any other account by using the assignVirtualMachine API. A domain or sub-domain administrator can do the same for VMs within the domain from one account to any other account in the domain or any of its sub-domains.

12.1.1. Dedicating Resources to Accounts and Domains

The root administrator can dedicate resources to a specific domain or account that needs private infrastructure for additional security or performance guarantees. A zone, pod, cluster, or host can be reserved by the root administrator for a specific domain or account. Only users in that domain or its subdomain may use the infrastructure. For example, only users in a given domain can create guests in a zone dedicated to that domain.

There are several types of dedication available:

- Explicit dedication. A zone, pod, cluster, or host is dedicated to an account or domain by the root administrator during initial deployment and configuration.
- Strict implicit dedication. A host will not be shared across multiple accounts. For example, strict implicit dedication is useful for deployment of certain types of applications, such as desktops, where no host can be shared between different accounts without violating the desktop software's terms of license.
- Preferred implicit dedication. The VM will be deployed in dedicated infrastructure if possible. Otherwise, the VM can be deployed in shared infrastructure.

12.1.1.1. How to Dedicate a Zone, Cluster, Pod, or Host to an Account or Domain

For explicit dedication: When deploying a new zone, pod, cluster, or host, the root administrator can click the Dedicated checkbox, then choose a domain or account to own the resource.

To explicitly dedicate an existing zone, pod, cluster, or host: log in as the root admin, find the resource in the UI, and click the **Dedicate** button.

For implicit dedication: The administrator creates a compute service offering and in the Deployment Planner field, chooses ImplicitDedicationPlanner. Then in Planner Mode, the administrator specifies either Strict or Preferred, depending on whether it is permissible to allow some use of shared resources when dedicated resources are not available. Whenever a user creates a VM based on this service offering, it is allocated on one of the dedicated hosts.

12.1.1.2. How to Use Dedicated Hosts

To use an explicitly dedicated host, use the explicit-dedicated type of affinity group. For example, when creating a new VM, an end user can choose to place it on dedicated infrastructure. This operation will succeed only if some infrastructure has already been assigned as dedicated to the user's account or domain.

12.1.1.3. Behavior of Dedicated Hosts, Clusters, Pods, and Zones

The administrator can live migrate VMs away from dedicated hosts if desired, whether the destination is a host reserved for a different account/domain or a host that is shared (not dedicated to any particular account or domain). CloudPlatform will generate an alert, but the operation is allowed.

Dedicated hosts can be used in conjunction with host tags. If both a host tag and dedication are requested, the VM will be placed only on a host that meets both requirements. If there is no dedicated resource available to that user that also has the host tag requested by the user, then the VM will not deploy.

If you delete an account or domain, any hosts, clusters, pods, and zones that were dedicated to it are freed up. They will now be available to be shared by any account or domain, or the administrator may choose to re-dedicate them to a different account or domain.

System VMs and virtual routers affect the behavior of host dedication. System VMs and virtual routers are owned by the CloudPlatform system account, and they can be deployed on any host. They do not adhere to explicit dedication. The presence of system vms and virtual routers on a host makes it unsuitable for strict implicit dedication. The host can not be used for strict implicit dedication, because the host already has VMs of a specific account (the default system account). However, a host with system VMs or virtual routers can be used for preferred implicit dedication.

12.2. Using an LDAP Server for User Authentication

CloudPlatform supports authentication through a Lightweight Directory Access Protocol (LDAP) server, such as Microsoft Active Directory or ApacheDS. You can add LDAP associations to CloudPlatform so users can log in by using credentials based on your existing authentication scheme. Additionally, the simplified LDAP authentication mechanism in CloudPlatform allows you to import users directly from the configured LDAP Group. LDAP users are authenticated without creating individual users in CloudPlatform.

To use LDAP for authentication of CloudPlatform users, you must do the following steps:

1. Install Microsoft Active Directory
2. Set up an OU within Active Directory to house the users you want to authenticate into CloudPlatform.
3. Add a working LDAP server.

   See *Section 12.2.1, “Configuring an LDAP Server”*.
4. Configure the LDAP attributes.

   See *Section 12.2.1.3, “Configuring LDAP Attributes in CloudPlatform”*.
5. Import users from the LDAP group.

   See *Section 12.2.2, “Importing LDAP Users to CloudPlatform”*.
6. To confirm authentication, log in to CloudPlatform UI as one of the LDAP user you have imported.

12.2.1. Configuring an LDAP Server

You can configure CloudPlatform to authenticate user access with a LDAP server. To set up LDAP authentication, provide the following:

- Hostname or IP address and listening port of the LDAP server.
- The LDAP global parameters.
- Base directory and query filter.
- Search user DN credentials, which give CloudPlatform permission to search on the LDAP server.
- SSL keystore and password, if SSL is used.

12.2.1.1. LDAP Configuration Considerations

The users should have an email address set against their Active Directory account. If this is not set, the user account with no email address as well as other user accounts within the same OU are not displayed in the CloudPlatform UI. If you encounter this issues, find the user who does not have the email address attribute and remove them from the Active Directory group you might have set as your

`Ldap.search.group.principle value`. Notice that the remaining users are displayed provided that none of these are missing the email attribute.

To establish higher security, LDAP server should be TLS-enabled. In this case, CloudPlatform Management Server use TLS to connect with LDAP server. We recommend you to not use SSL.

12.2.1.2. Adding an LDAP Server

1. Log in to the CloudPlatform UI.
2. From the left navigational bar, click Global Settings.
3. From the Select view drop down, select LDAP Configuration.
4. Click Configure LDAP.

   The Configure LDAP dialog is displayed.

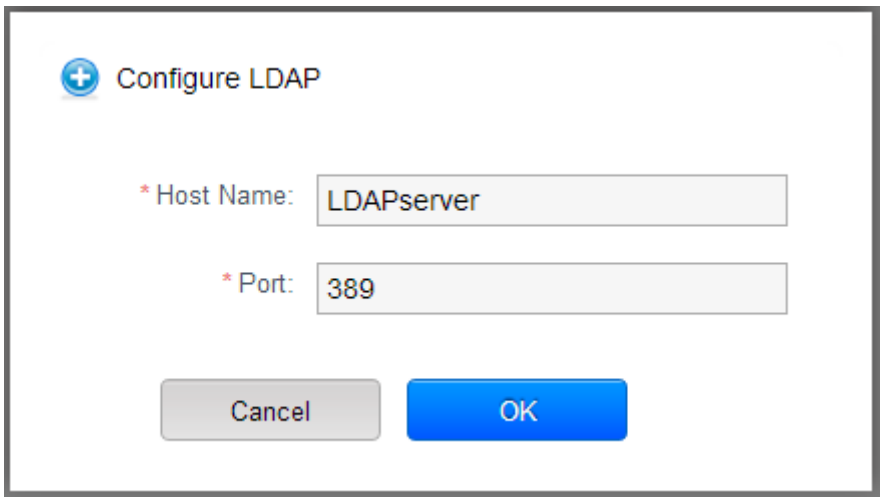

5. Specify the following:
   - **Hostname**: Hostname or IP address of the LDAP server.
   - **Port**: The Listening port of the LDAP server.

     The port numbers for LDAP connections are:

     - 389 for unsecured LDAP connections. This is the default value.
     - 636 for secure LDAP connections.
     - 3268 for Microsoft unsecure LDAP connections.
     - 3269 for Microsoft secure LDAP connections.
     - 10389 for ApacheDS.
6. Click OK.

12.2.1.3. Configuring LDAP Attributes in CloudPlatform

CloudPlatform provides the following global LDAP configuration parameters. You can locate them by searching for ldap in the Global settings.

- *ldap basedn*: Defines the location of the users. This is usually derived from the **binddn**. Remove the user name from **bind dn** and specify the group where users are located. The entire subtree under the **binddn** will be searched for user accounts.

  ```
  cn=users,dc=<sub-domain>,dc=<domain>,dc=com
  ```

  For example:

  ```
  ou=WWUsers,dc=uklab,dc=internal
  ```

- *ldap bind password*: The password used in association with the administrator bind DN. This is used for querying the LDAP directory. If this is left blank along with bind principal then anonymous binding is used.

  For example:

  ```
  password
  ```

- *ldap bind principal*: The principle to bind to the LDAP server for creating the system context. The value is frequently the DN (Distinguished Name) of the user entry with the user ID. If this field is left blank along with the bind password then anonymous binding is used.

  For example:

  ```
  Administrator
  ```

- *ldap email attribute*: The attribute that your LDAP directory uses to hold the user's e-mail address. Default attribute name is *mail*.

  For example:

  ```
  Mail
  ```

- *ldap firstname attribute*: The attribute that your LDAP directory uses to hold the first name of the user. Default is *cn*.

  For example:

  ```
  givenName
  ```

- *ldap group object*: The attribute that sets the object types for groups.

  For example:

  ```
  user
  ```

- *ldap group user uniquemember*: The attribute that your LDAP directory uses to hold the unique members of the group.

- *ldap lastname attribute*: The attribute that your LDAP directory uses to hold the last name of the user.

  For example:

```
sn
```

- *ldap search group principle*: Sets the principle of the group that the LDAP users must be part of.

  For example:

  ```
  cn=CloudPlatform,ou=WWUsers,dc=uklab,dc=internal
  ```

- *ldap trust store*: Sets the path to the trust store to be used for secure connections. You can use the trust store to install CA certificates and client certificates.

  For example:

  ```
  \lib\security\cacerts
  ```

  If you are not using LDAP, leave it as blank.

- *ldap trust store password*: Sets the password for the trust store. Password protects the trust store.

  If you are not using LDAP, leave it as blank.

- *ldap user object*: The object type of user accounts within LDAP. The default is *inetOrgperson*.

  For example:

  ```
  user
  ```

12.2.1.4. LDAP Attributes

The default LDAP authentication mechanism used is Active Directory. However, you can change the default values based on the LDAP server that you are using:

LDAP Attribute	OpenLDAP	Active Directory
ldap.user.object	inetOrgPerson	user
ldap.username.attribute	uid	sAMAccountName
ldap.group.object	groupOfUniqueNames	group
ldap.group.user.uniquemember	member	uniquemember
(optional) ldap.search.group.principal	customer-specified	customer-specified

12.2.1.5. Removing an LDAP Configuration

1. Log in to the CloudPlatform.
2. From the left navigational bar, click Global Settings.
3. From the Select view drop down, select LDAP Configuration.
4. In the Quick View, click Remove LDAP.

Alternatively, you can click Remove LDAP in the LDAP Configuration Details page.

12.2.2. Importing LDAP Users to CloudPlatform

You can import LDAP users without creating users individually in CloudPlatform. After users are imported to CloudPlatform they can by authenticated to use CloudPlatform by using their LDAP credentials. You can either directly select users from the group used in the **`Ldap.search.group.principle`** value. Alternatively, specify the search string or user attributes for CloudPlatform to search the LDAP directory tree based on the query strings.

1. Log in to the CloudPlatform UI as an administrator.
2. In the left navigation bar, click Accounts.

   The Account page is displayed.
3. In the Account page, click Add LDAP Users.

   The Add LDAP User screen lists all the users associated with the LDAP server you have configured.
4. In the Add LDAP Account screen, perform either of the following:
   - Manually select the users from the user list.

     CloudPlatform displays all the users from the LDAP group configured.
   - Specify the search string or desired user attribute, then import users.
     - Domain
     - Account
     - User type
     - Timezone
     - Network Domain
     - LDAP Group
5. Click Add.

12.2.3. Configuring CloudPlatform to Use Global Catalog

A global catalog holds information about all the objects within a forest. You can set up CloudPlatform Active Directory authentication to make use of an Active Directory Global Catalog to query more than one domain. To query a global catalog you must first setup CloudPlatform to point at a valid global catalog server within the environment and use port 3268 instead of 389 as this will query the global catalog and not just the local domain the global catalog server is joined to.

Using Projects to Organize Users and Resources

13.1. Overview of Projects

Projects are used to organize people and resources. CloudPlatform users within a single domain can group themselves into project teams so they can collaborate and share virtual resources such as VMs, snapshots, templates, data disks, and IP addresses. CloudPlatform tracks resource usage per project as well as per user, so the usage can be billed to either a user account or a project. For example, a private cloud within a software company might have all members of the QA department assigned to one project, so the company can track the resources used in testing while the project members can more easily isolate their efforts from other users of the same cloud

You can configure CloudPlatform to allow any user to create a new project, or you can restrict that ability to just CloudPlatform administrators. Once you have created a project, you become that project's administrator, and you can add others within your domain to the project. CloudPlatform can be set up either so that you can add people directly to a project, or so that you have to send an invitation which the recipient must accept. Project members can view and manage all virtual resources created by anyone in the project (for example, share VMs). A user can be a member of any number of projects and can switch views in the CloudPlatform UI to show only project-related information, such as project VMs, fellow project members, project-related alerts, and so on.

The project administrator can pass on the role to another project member. The project administrator can also add more members, remove members from the project, set new resource limits (as long as they are below the global defaults set by the CloudPlatform administrator), and delete the project. When the administrator removes a member from the project, resources created by that user, such as VM instances, remain with the project. This brings us to the subject of resource ownership and which resources can be used by a project.

Resources created within a project are owned by the project, not by any particular CloudPlatform account, and they can be used only within the project. A user who belongs to one or more projects can still create resources outside of those projects, and those resources belong to the user's account; they will not be counted against the project's usage or resource limits. You can create project-level networks to isolate traffic within the project and provide network services such as port forwarding, load balancing, VPN, and static NAT. A project can also make use of certain types of resources from outside the project, if those resources are shared. For example, a shared network or public template is available to any project in the domain. A project can get access to a private template if the template's owner will grant permission. A project can use any service offering or disk offering available in its domain; however, you can not create private service and disk offerings at the project level..

13.2. Configuring Projects

Before CloudPlatform users start using projects, the CloudPlatform administrator must set up various systems to support them, including membership invitations, limits on project resources, and controls on who can create projects.

13.2.1. Setting Up Invitations

CloudPlatform can be set up either so that project administrators can add people directly to a project, or so that it is necessary to send an invitation which the recipient must accept. The invitation can be sent by email or through the user's CloudPlatform account. If you want administrators to use invitations to add members to projects, turn on and set up the invitations feature in CloudPlatform.

1. Log in as administrator to the CloudPlatform UI.
2. In the left navigation, click Global Settings.
3. In the search box, type project and click the search button.
4. In the search results, you can see a few other parameters you need to set to control how invitations behave. The table below shows global configuration parameters related to project invitations. Click the edit button to set each parameter.

Configuration Parameters	Description
project.invite.required	Set to true to turn on the invitations feature.
project.email.sender	The email address to show in the From field of invitation emails.
project.invite.timeout	Amount of time to allow for a new member to respond to the invitation.
project.smtp.host	Name of the host that acts as an email server to handle invitations.
project.smtp.password	(Optional) Password required by the SMTP server. You must also set project.smtp.username and set project.smtp.useAuth to true.
project.smtp.port	SMTP server's listening port.
project.smtp.useAuth	Set to true if the SMTP server requires a username and password.
project.smtp.username	(Optional) User name required by the SMTP server for authentication. You must also set project.smtp.password and set project.smtp.useAuth to true..

5. Restart the Management Server:

```
service cloud-management restart
```

13.2.2. Setting Resource Limits for Projects

The CloudPlatform administrator can set global default limits to control the amount of resources that can be owned by each project in the cloud. This serves to prevent uncontrolled usage of resources such as snapshots, IP addresses, and virtual machine instances. Domain administrators can override these resource limits for individual projects with their domains, as long as the new limits are below the global defaults set by the CloudPlatform root administrator. The root administrator can also set lower resource limits for any project in the cloud

13.2.3. Setting Project Creator Permissions

You can configure CloudPlatform to allow any user to create a new project, or you can restrict that ability to just CloudPlatform administrators.

1. Log in as administrator to the CloudPlatform UI.
2. In the left navigation, click Global Settings.

3. In the search box, type allow.user.create.projects.
4. Click the edit button to set the parameter.

allow.user.create.projects	Set to true to allow end users to create projects. Set to false if you want only the CloudPlatform root administrator and domain administrators to create projects.

5. Restart the Management Server.

```
# service cloud-management restart
```

13.3. Creating a New Project

CloudPlatform administrators and domain administrators can create projects. If the global configuration parameter allow.user.create.projects is set to true, end users can also create projects.

1. Log in as administrator to the CloudPlatform UI.
2. In the left navigation, click Projects.
3. In Select view, click Projects.
4. Click New Project.
5. Give the project a name and description for display to users, then click Create Project.
6. A screen appears where you can immediately add more members to the project. This is optional. Click Next when you are ready to move on.
7. Click Save.

13.4. Adding Members to a Project

New members can be added to a project by the project's administrator, the domain administrator of the domain where the project resides or any parent domain, or the CloudPlatform root administrator. There are two ways to add members in CloudPlatform, but only one way is enabled at a time:

- If invitations have been enabled, you can send invitations to new members.
- If invitations are not enabled, you can add members directly through the UI.

13.4.1. Sending Project Membership Invitations

Use these steps to add a new member to a project if the invitations feature is enabled in the cloud. If the invitations feature is not turned on, use the procedure in Adding Project Members From the UI.

1. Log in to the CloudPlatform UI.
2. In the left navigation, click Projects.
3. In Select View, choose Projects.
4. Click the name of the project you want to work with.
5. Click the Invitations tab.

6. In Add by, select one of the following:

    a. Account – The invitation will appear in the user's Invitations tab in the Project View. See Using the Project View.

    b. Email – The invitation will be sent to the user's email address. Each emailed invitation includes a unique code called a token which the recipient will provide back to CloudPlatform when accepting the invitation. Email invitations will work only if the global parameters related to the SMTP server have been set.

7. Type the user name or email address of the new member you want to add, and click Invite. Type the CloudPlatform user name if you chose Account in the previous step. If you chose Email, type the email address. You can invite only people who have an account in this cloud within the same domain as the project. However, you can send the invitation to any email address.

8. To view and manage the invitations you have sent, return to this tab. When an invitation is accepted, the new member will appear in the project's Accounts tab.

13.4.2. Adding Project Members From the UI

The steps below tell how to add a new member to a project if the invitations feature is not enabled in the cloud. If the invitations feature is enabled cloud,use the procedure in the Sending Project Membership Invitations section.

1. Log in to the CloudPlatform UI.

2. In the left navigation, click Projects.

3. In Select View, choose Projects.

4. Click the name of the project you want to work with.

5. Click the Accounts tab. The current members of the project are listed.

6. Type the account name of the new member you want to add, and click Add Account. You can add only people who have an account in this cloud and within the same domain as the project.

13.5. Accepting a Membership Invitation

If you have received an invitation to join a CloudPlatform project, and you want to accept the invitation, follow these steps:

1. Log in to the CloudPlatform UI.

2. In the left navigation, click Projects.

3. In Select View, choose Invitations.

4. If you see the invitation listed onscreen, click the Accept button.

    Invitations listed on screen were sent to you using your CloudPlatform account name.

5. If you received an email invitation, click the Enter Token button, and provide the project ID and unique ID code (token) from the email.

13.6. Suspending or Deleting a Project

When a project is suspended, it retains the resources it owns, but they can no longer be used. No new resources or members can be added to a suspended project.

When a project is deleted, its resources are destroyed, and member accounts are removed from the project. The project's status is shown as Disabled pending final deletion.

A project can be suspended or deleted by the project administrator, the domain administrator of the domain the project belongs to or of its parent domain, or the CloudPlatform root administrator.

1. Log in to the CloudPlatform UI.
2. In the left navigation, click Projects.
3. In Select View, choose Projects.
4. Click the name of the project.
5. Click one of the buttons:

   To delete, use 

   To suspend, use 

13.7. Using the Project View

If you are a member of a project, you can use CloudPlatform's project view to see project members, resources consumed, and more. The project view shows only information related to one project. It is a useful way to filter out other information so you can concentrate on a project status and resources.

1. Log in to the CloudPlatform UI.
2. Click Project View.
3. The project dashboard appears, showing the project's VMs, volumes, users, events, network settings, and more. From the dashboard, you can:

   - Click the Accounts tab to view and manage project members. If you are the project administrator, you can add new members, remove members, or change the role of a member from user to admin. Only one member at a time can have the admin role, so if you set another user's role to admin, your role will change to regular user.
   - (If invitations are enabled) Click the Invitations tab to view and manage invitations that have been sent to new project members but not yet accepted. Pending invitations will remain in this list until the new member accepts, the invitation timeout is reached, or you cancel the invitation.

System Reliability and High Availability

14.1. HA for Management Server

The CloudPlatform Management Server should be deployed in a multi-node configuration such that it is not susceptible to individual server failures. The Management Server itself (as distinct from the MySQL database) is stateless and may be placed behind a load balancer.

Normal operation of Hosts is not impacted by an outage of all Management Serves. All guest VMs will continue to work.

When the Management Server is down, no new VMs can be created, and the end user and admin UI, API, dynamic load distribution, and HA will cease to work.

14.2. HA-Enabled Virtual Machines

The user can specify a virtual machine as HA-enabled. By default, all virtual router VMs and Elastic Load Balancing VMs are automatically configured as HA-enabled. When an HA-enabled VM crashes, CloudPlatform detects the crash and restarts the VM automatically within the same Availability Zone. HA is never performed across different Availability Zones. CloudPlatform has a conservative policy towards restarting VMs and ensures that there will never be two instances of the same VM running at the same time. The Management Server attempts to start the VM on another Host in the same cluster.

Note

On VMware deployments, enable native HA. Because CloudPlatform relies on native HA for VMware, Offer HA field in the Create Compute Offering dialog is ignored.

HA features work with iSCSI or NFS primary storage. HA with local storage is not supported.

14.3. Dedicated HA Hosts

One or more hosts can be designated for use only by HA-enabled VMs that are restarting due to a host failure. Setting up a pool of such dedicated HA hosts as the recovery destination for all HA-enabled VMs is useful to:

- Make it easier to determine which VMs have been restarted as part of the CloudPlatform high-availability function. If a VM is running on a dedicated HA host, then it must be an HA-enabled VM whose original host failed. (With one exception: It is possible for an administrator to manually migrate any VM to a dedicated HA host.).
- Keep HA-enabled VMs from restarting on hosts which may be reserved for other purposes.

The dedicated HA option is set through a special host tag when the host is created. To allow the administrator to dedicate hosts to only HA-enabled VMs, set the global configuration variable ha.tag to the desired tag (for example, "ha_host"), and restart the Management Server. Enter the value in the Host Tags field when adding the host(s) that you want to dedicate to HA-enabled VMs.

Note

If you set ha.tag, be sure to actually use that tag on at least one host in your cloud. If the tag specified in ha.tag is not set for any host in the cloud, the HA-enabled VMs will fail to restart after a crash.

14.4. Primary Storage Outage and Data Loss

When a primary storage outage occurs, all hosts in that cluster are rebooted. This ensures that affected VMs running on the hypervisor are appropriately marked as stopped. Guests that are marked for HA will be restarted as soon as practical when the primary storage comes back on line. With NFS, the hypervisor may allow the virtual machines to continue running depending on the nature of the issue. For example, an NFS hang will cause the guest VMs to be suspended until storage connectivity is restored. Primary storage is not designed to be backed up. Individual volumes in primary storage can be backed up using snapshots.

Note

If there are multiple primary storage servers in a cluster and only one goes down, VMs using a healthy primary storage will also be affected, because all hosts are rebooted.

14.5. Secondary Storage Outage and Data Loss

For a Zone that has only one secondary storage server, a secondary storage outage will have feature level impact to the system but will not impact running guest VMs. It may become impossible to create a VM with the selected template for a user. A user may also not be able to save snapshots or examine/restore saved snapshots. These features will automatically be available when the secondary storage comes back online.

Secondary storage data loss will impact recently added user data including templates, snapshots, and ISO images. Secondary storage should be backed up periodically.Multiple secondary storage servers can be provisioned within each zone to increase the scalability of the system.

14.6. Database High Availability

To help ensure high availability of the databases that store the internal data for CloudPlatform, you can set up database high availability. This covers both the main CloudPlatform database and the Usage database. High availability is achieved using the MySQL connector parameters and two-way high availability. Tested with MySQL 5.1.

14.6.1. How to Set Up Database High Availability

Database high availability in CloudPlatform is provided using the MySQL high availability capabilities. The steps to set up high availability can be found in the MySQL documentation (links are provided

below). It is suggested that you set up two-way high availability, which involves two database nodes. In this case, for example, you might have node1 and node2. In Asynchronous high availability configuration, not more than two database nodes are supported.

References:

- *http://dev.mysql.com/doc/refman/5.0/en/replication-howto.html*
- *https://wikis.oracle.com/display/CommSuite/MySQL+High+Availability+and+Replication+Information+For+Calendar+Server*

14.6.2. Database High Availability Considerations

- To clean up bin log files automatically, perform the following configuration in the **`my.cnf`** file of each MySQL server:

  expire_logs_days=10 : Number of days to keep the log files.

  max_binlog_size=100M: The maximum size of each log file.

- To change the bin log files location, use the log-bin=/var/lib/mysql/binlog/bin-log property in the **`my.cnf`** file of each MySQL server.

- If two Management Servers happen to connect to two different database HA (split brain), modify the following properties in db.properties of the slave:

  - auto_increment_increment = 10: Instruct the MySQL node to auto increment values by 10 instead of default value 1.

  - auto_increment_offset = 2: Instruct the MySQL node what is the starting point of the auto increment column value to be start with. The second property is relevant only when split brain occurs on a fresh setup.

14.6.3. Configuring Database High Availability

To control the database high availability behavior, use the following configuration settings in the file /etc/cloudstack/management/db.properties.

Required Settings

Be sure you have set the following in db.properties:

- db.ha.enabled: set to true if you want to use the high availability feature.

  Example: db.ha.enabled=true

- db.cloud.slaves: set to a comma-delimited set of slave hosts for the cloud database. This is the list of nodes set up with high availability. The master node is not in the list, since it is already mentioned elsewhere in the properties file.

  Example: db.cloud.slaves=node2,node3,node4

- db.usage.slaves: set to a comma-delimited set of slave hosts for the usage database. This is the list of nodes set up with high availability. The master node is not in the list, since it is already mentioned elsewhere in the properties file.

  Example: db.usage.slaves=node2,node3,node4

Optional Settings

The following settings must be present in db.properties, but you are not required to change the default values unless you wish to do so for tuning purposes:

- db.cloud.secondsBeforeRetryMaster: The number of seconds the MySQL connector should wait before trying again to connect to the master after the master went down. Default is 1 hour. The retry might happen sooner if db.cloud.queriesBeforeRetryMaster is reached first.

  Example: db.cloud.secondsBeforeRetryMaster=3600

- db.cloud.queriesBeforeRetryMaster: The minimum number of queries to be sent to the database before trying again to connect to the master after the master went down. Default is 5000. The retry might happen sooner if db.cloud.secondsBeforeRetryMaster is reached first.

  Example: db.cloud.queriesBeforeRetryMaster=5000

- db.cloud.initialTimeout: Initial time the MySQL connector should wait before trying again to connect to the master. Default is 3600.

  Example: db.cloud.initialTimeout=3600

14.6.4. Asynchronous Configuration for Database High Availability

The MySQL configuration to support DB HA in MySQL server is goes into the **`/etc/my.cnf`** file and configuration slightly varies between master and slave.

14.6.4.1. Master Configuration

```
[mysqld]
datadir=/var/lib/mysql
socket=/var/lib/mysql/mysql.sock
# Disabling symbolic-links is recommended to prevent assorted security risks
symbolic-links=0
# Settings user and group are ignored when systemd is used.
# If you need to run mysqld under a different user or group,
# customize your systemd unit file for mysqld according to the
# instructions in http://fedoraproject.org/wiki/Systemd
server-id=1
default-storage-engine=InnoDB
character-set-server=utf8
transaction-isolation=READ-COMMITTED
log-bin=mysql-bin
innodb_flush_log_at_trx_commit=1
sync_binlog=1
binlog-format=ROW
#Bin logs cleanup configuration
expire_log_days=10
max_binlog_size=100M

[mysqld_safe]
log-error=/var/log/mysqld.log
pid-file=/var/run/mysqld/mysqld.pid
```

14.6.4.2. Slave Configuration

```
[mysqld]
datadir=/var/lib/mysql
socket=/var/lib/mysql/mysql.sock
# Disabling symbolic-links is recommended to prevent assorted security risks
symbolic-links=0
# Settings user and group are ignored when systemd is used.
# If you need to run mysqld under a different user or group,
# customize your systemd unit file for mysqld according to the
# instructions in http://fedoraproject.org/wiki/Systemd
server-id=2
default-storage-engine = InnoDB
character-set-server = utf8
transaction-isolation = READ-COMMITTED
log-bin=mysql-bin
innodb_flush_log_at_trx_commit=1
sync_binlog=1
binlog-format=ROW
#Parameters to solve split brain problem
auto_increment_increment=10 auto_increment_offset=2
#Bin logs cleanup configuration
expire_log_days=10
max_binlog_size=100M
[mysqld_safe]
log-error=/var/log/mysqld.log
pid-file=/var/run/mysqld/mysqld.pid
```

14.6.5. Limitations on Database High Availability

The following limitations exist in the current implementation of this feature.

- Slave hosts can not be monitored through CloudPlatform. You will need to have a separate means of monitoring.
- Events from the database side are not integrated with the CloudPlatform Management Server events system.
- MySQL 5.1 supports only Asynchronous high availability; therefore, there is a chance of data inconsistency while server is down.

14.6.6. Master-Master High Availability Usecase

DEMO-CUSTOMER plans to deploy multiple CloudPlatform implementations over several regions. Each region contains two datacenters. One key requirement is to make sure that critical infrastructure systems are always highly-available over both datacenters in a specific region. Both Citrix CloudPlatform and Citrix CloudPortal Business Manager (CPBM) are using MySQL for hosting the databases.

To make the critical MySQL database highly available, setup a MySQL Master - Master / Slave topology. CloudPlatform supports MySQL Master-Master Replication. Master – Master replication provides automatic failover, if one of the MySQL nodes goes down or becomes otherwise unavailable unexpectedly.

The Master-Master replication ensures that writing to the MySQL databases are always performed synchronously on both master nodes. The additional slave hosts in the second datacenter, monitors the databases of the master hosts and replicates database changes asynchronous immediately. The ClouPlatform management servers can be configured to work with a MySQL Master-Master cluster. In this configuration, CloudPlatform will always connect to the primary database node. If the primary

node goes down, CloudPlatform will immediately and automatically connect to the available MySQL Master node without interruption.

Replication never replaces the need for a consistent backup. There is always a potential risk that a database, could get into an inconsistent state. In this case, only a restore of a consistent database backup will bring the platform online again.

This section describes how to configure a pair of MySQL Master-Master hosts for the main datacenter, and also explains how to add additional slave MySQL hosts for the second datacenter within a given region. With such a configuration in place, CloudPlatform will continue to work even if one of the datacenters goes down unexpectedly.

14.6.6.1. Environment

This section assumes the use of the DEMO-CUSTOMER region EMEA with the datacenters Primary and datacenter Secondary as examples. Each datacenter represents a Citrix CloudPlatform zone and the whole Region will represent a Citrix CloudPlatform Region. The DEMO-CUSTOMER datacenters are connected over a high performance and low latency WAN connections (20Gbit/s, less than 2ms). This allows to run the CloudPlatform Management Servers active / active across the datacenters. Though you are allowed to run multiple MySQL servers in a Master / Master configuration across datacenters, it´s not recommended. An interruption in the datacenter-to-datacenter connection would immediately lead to a “split-brain” situation. Each additional Master node also increases the risk of inconsistent data and the restore process becomes more challenging. Instead, it´s recommended to run MySQL Master / Master nodes only in the main datacenter. The CloudPlatform Management Servers in both datacenters will connect to the MySQL primary Master node, which is specified in the CloudPlatform configuration.

14.6.6.2. Operation Mode

In general, the platform can run in three different operational modes:

- *Section 14.6.6.2.1, “Normal Operation Mode”*
- *Section 14.6.6.2.2, “Primary Master Failover Mode”*
- *Section 14.6.6.2.3, “Datacenter Failover Mode”*

14.6.6.2.1. Normal Operation Mode

All components are reachable. The CloudPlatform servers from both sides are connected to the primary Master SITE01-MYSQL01, and SITE01-MYSQL02 replicates the changes from SITE01-MYSQL01. The Slave node of the second datacenter (SITE02-MYSQLL01) replicates the changes from the node SITE01-MYSQL02:

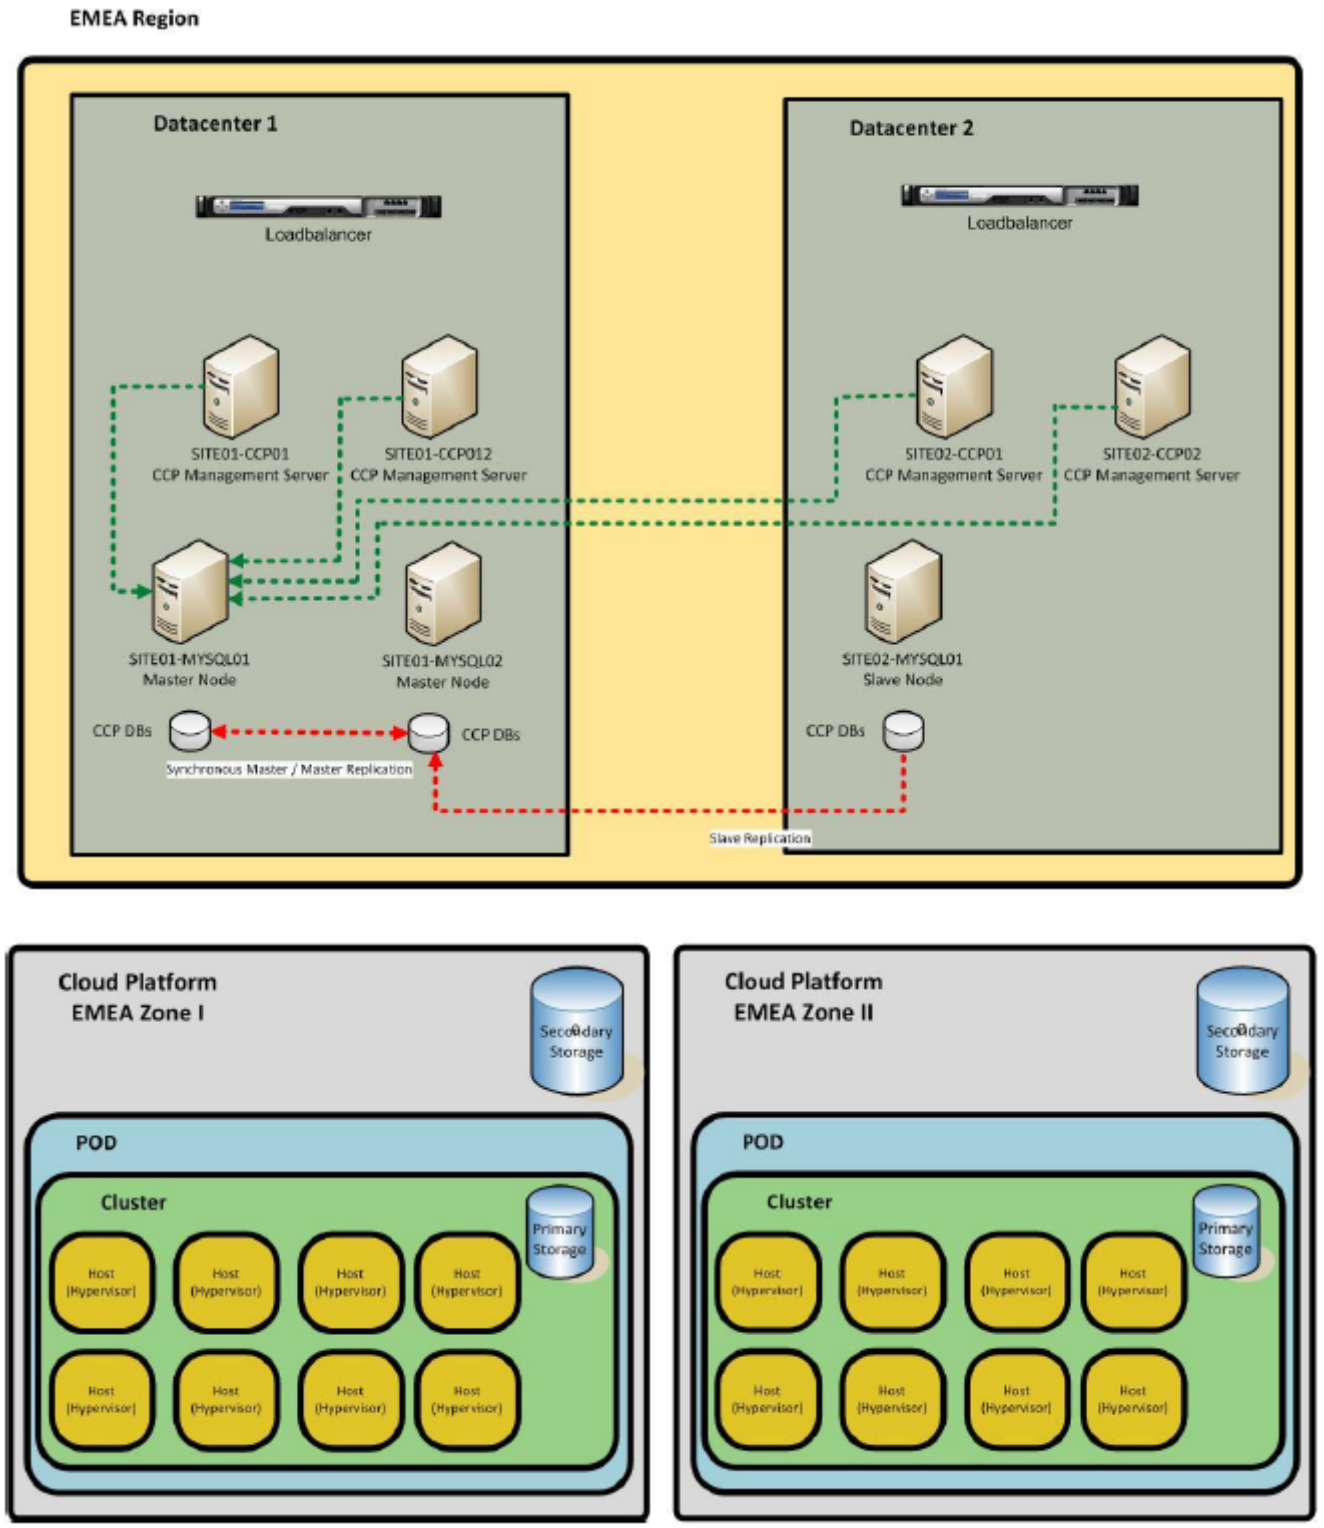

14.6.6.2.2. Primary Master Failover Mode

If the primary Master MySQL node goes down unexpectedly, CloudPlatform automatically switches the connection automatically and without interruption to the secondary master node SITE01-MYSQL02. The Slave node of the second datacenter, SITE02-MYSQL02, continues to replicate database changes from SITE01-MYSQL02.

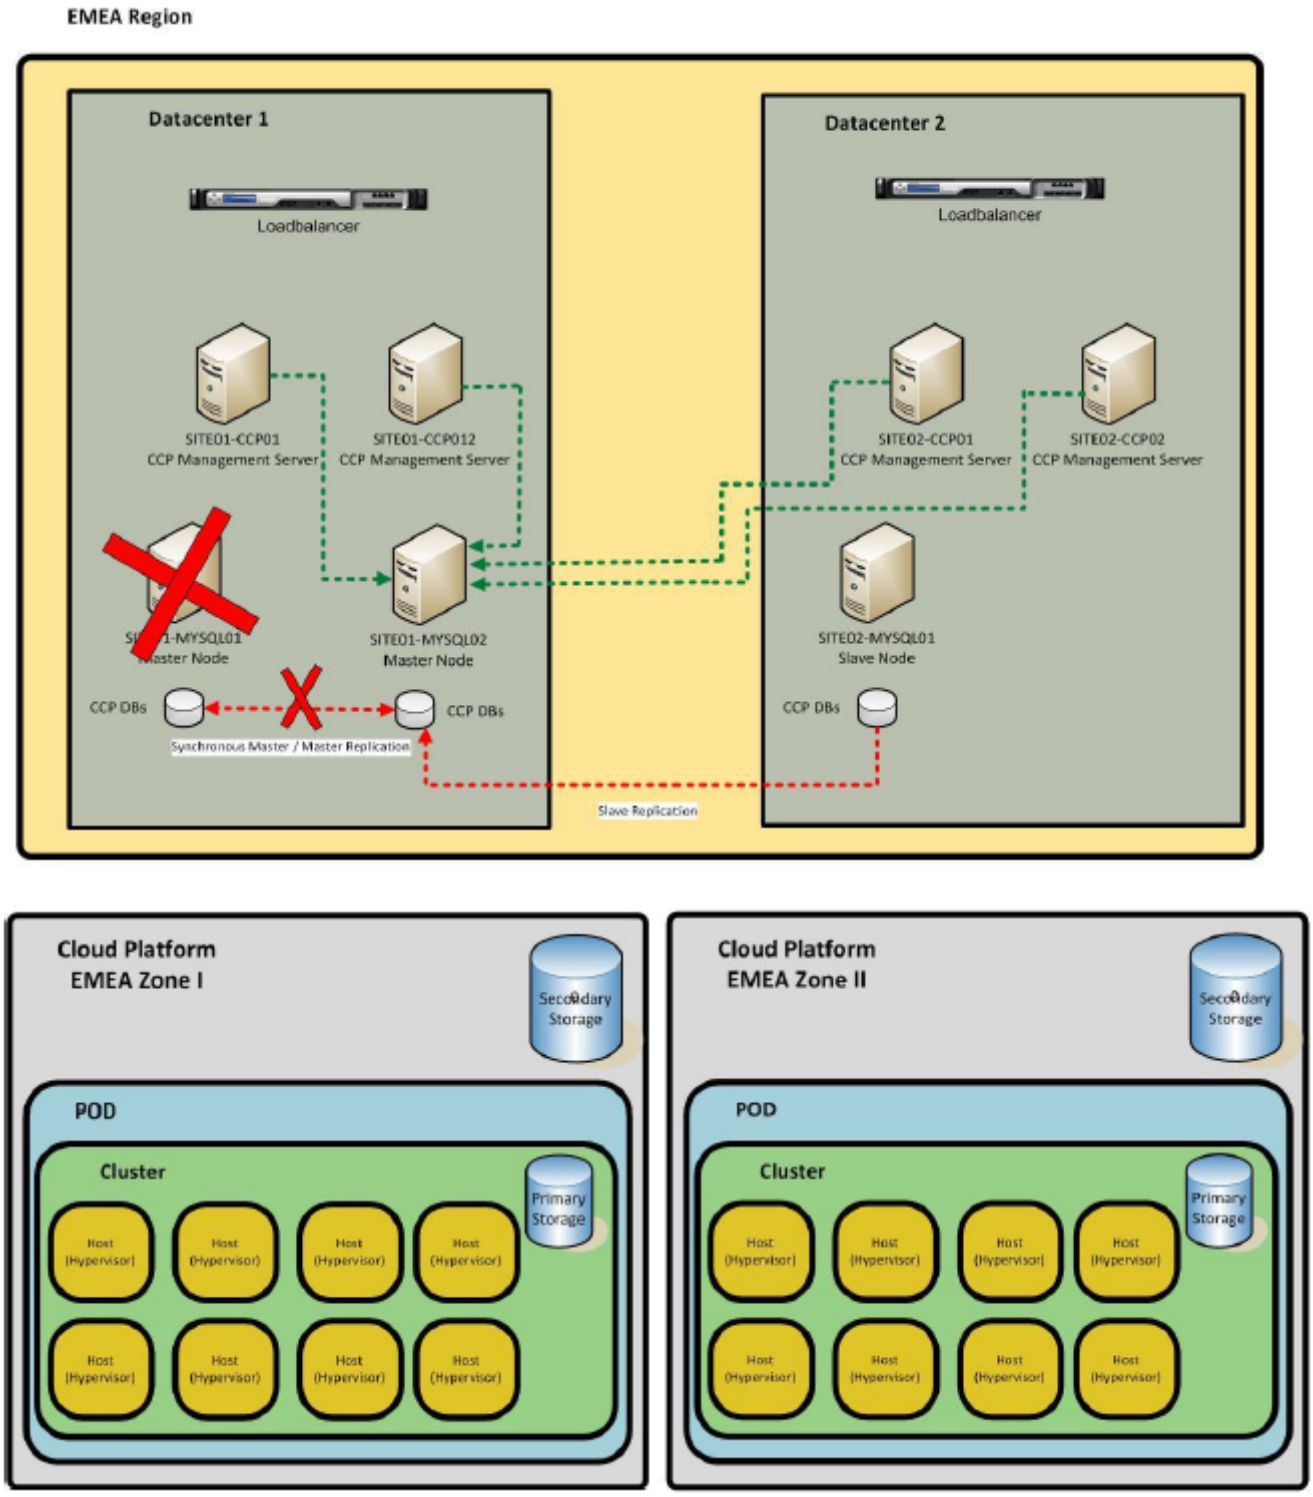

14.6.6.2.3. Datacenter Failover Mode

If the entire Master-Master cluster becomes unreachable, or the primary datacenter is down, manual steps are required to reconfigure the environment to use the MySQL Slave node in the secondary datacenter. The MySQL Slave node must be made the new Master. Additionally, you must restart the CloudPlatform Management Services. The Management Servers of the second datacenter will then continue to manage the zone of the second datacenter.

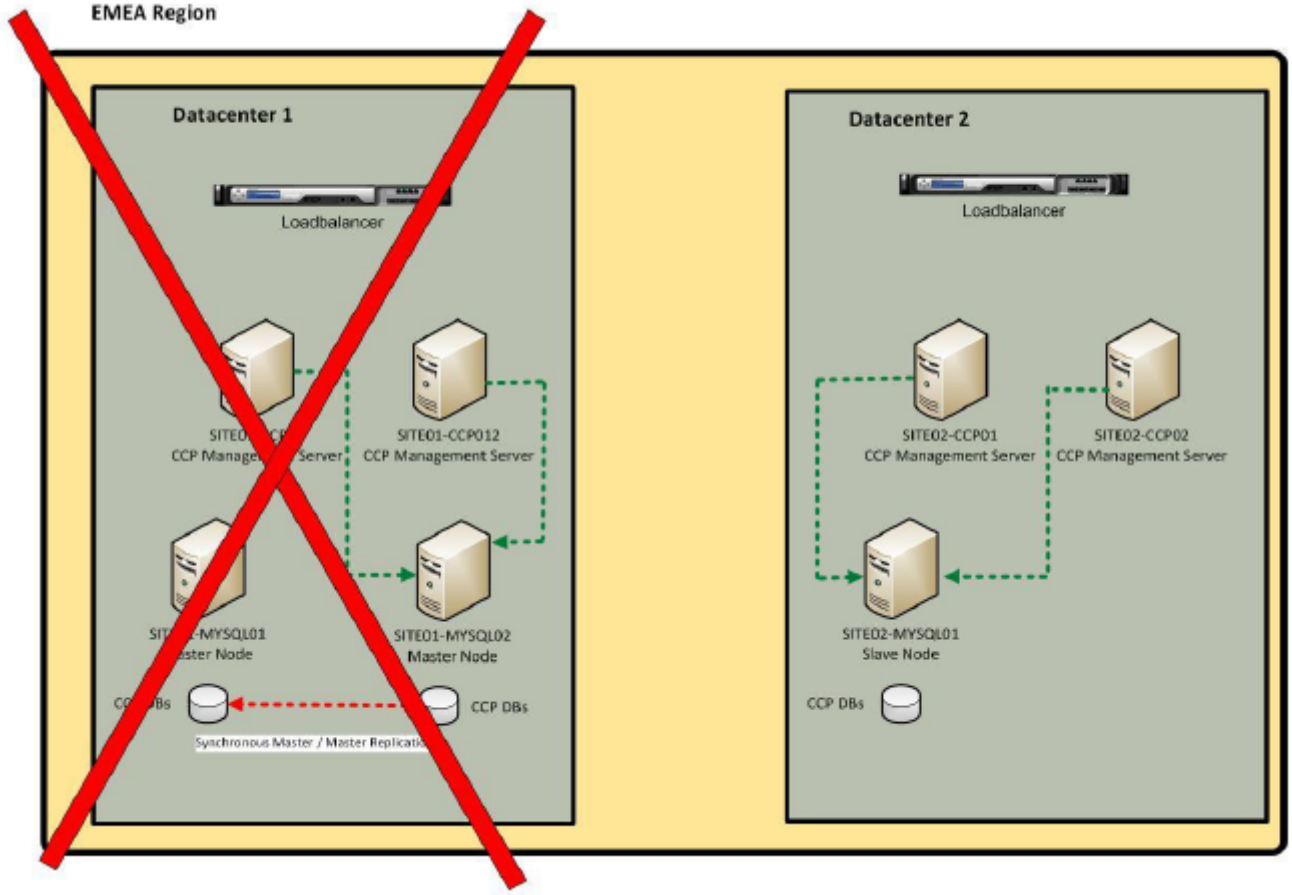

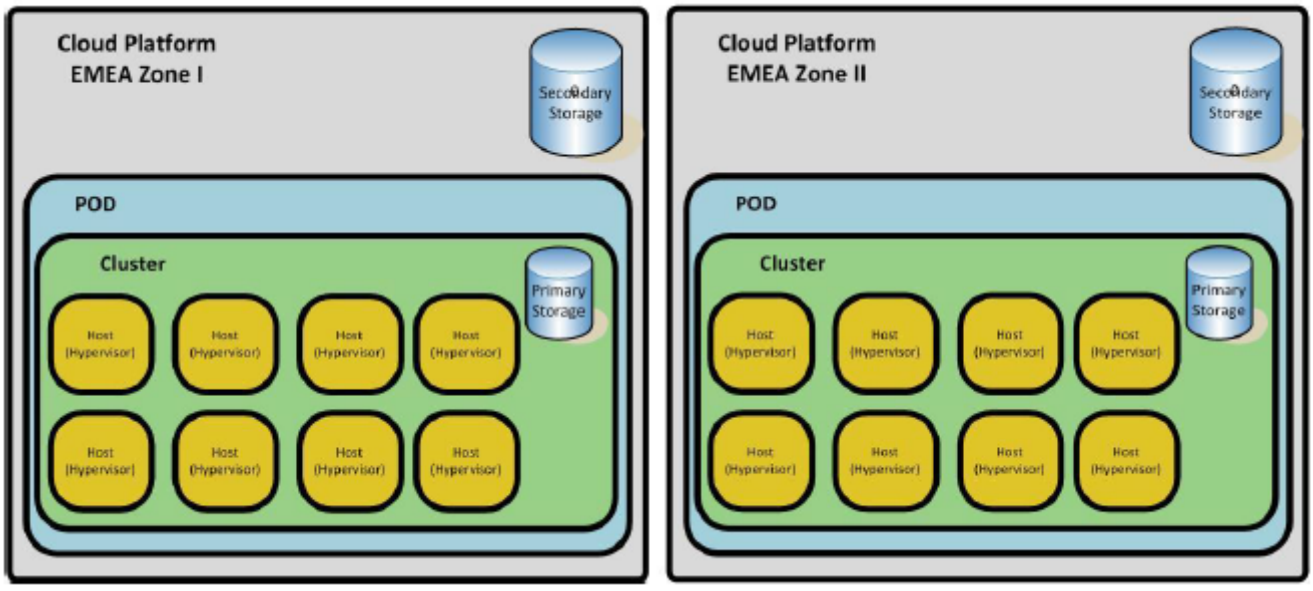

14.6.6.3. Installing MySQL on CloudPlatform

The following versions are used in the example configuration:

- Citrix Cloud Platform 4.3
- MySQL Server 5.1.73

All systems are installed on CentOS 6.5 operating system.

14.6.6.4. Prerequisites and Guidelines

The assumption is that CloudPlatform 4.3 is already installed on the following nodes:

Datacenter Primary, SITE01:

SITE01-CCP01 and SITE01-CCP02

Datacenter Secondary, SITE02:

SITE02-CCP01 and SITE02-CCP02

All CloudPlatform nodes are currently connected to the main database node SITE01-MYSQL01. MySQL has been installed as part of the CloudPlatform installation. The other MYSQL nodes which are used for the database replication are not yet installed, just the OS is prepared and the basic configuration is done. Due to limited computing resources, only two CloudPlatform servers are used, one for each site.

The Basic configuration includes:

- Static IP configured
- SElinux disabled
- Paravirtualized drivers installed (VMware Tools, XenServer Tools)
- NTP configured and time is synced between all hosts
- DNS working and all hosts are resolvable
- MySQL Port 3306 opened between all MySQL and CloudPlatform nodes
- CloudPlatform 4.3 installation tar ball available and copied to all MySQL nodes under /tmp

14.6.6.5. Installing MySQL

This section explains how to install MySQL on the systems SITE01-MYSQL02 and SITE02-MYSQL01. You are recommended to use the MySQL installer that ships with the CloudPlatform installation package. To guarantee that the replication works as expected, the same MySQL versions are required across all nodes.

1. On the nodes SITE01-MYSQL02 and SITE02-MYSQL02, extract the CloudPlatform tar ball:

   ```
   # tar xvf CloudPlatform-4.3.0.0-rhel6.4.tar.gz
   ```

2. On the nodes SITE01-MYSQL01 and SITE01-MYSQL02 launch the CloudPlatform installer (install.sh) and select Option D to install the database server from distribution's repository:

3. Set the MySQL service to start on system-boot:

   ```
   # Chkconfig mysqld on
   ```

4. On the nodes SITE01-MYSQL01 and SITE01-MYSQL02 customize the **`my.cnf`** file with all applicable changes. To guarantee a working failover, all servers should have consistent settings.

   Example configuration:

```
[mysqld]
datadir=/var/lib/mysql
socket=/var/lib/mysql/mysql.sock
user=mysql
# Disabling symbolic-links is recommended to prevent assorted security risks
symbolic-links=0

server-id=1

innodb_rollback_on_timeout=1
innodb_lock_wait_timeout=600
default-time-zone='+00:00'
max_connections=1000
log-bin=mysql-bin
binlog-format = 'ROW'

innodb_buffer_pool_size=5500


[mysqld_safe]
log-error=/var/log/mysqld.log
pid-file=/var/run/mysqld/mysqld.pid
```

Ensure that you use the appropriate settings, which were set during the installation of the first MySQL node. Note down the value for *`server-id`*, which has to be unique across all hosts in a replication configuration.

5. Restart the mysqld service on the nodes SITE01-MYSQL01 and SITE01-MYSQL02:

```
# Service mysqld restart
```

14.6.6.6. Configuring Multi-Master Replication

This section describes how to configure the Master-Master replication between the nodes SITE01-MYSQL01 and SITE01-MYSQL02. In general, the MySQL multi-master replication is configured by setting up a two-way slave replication, there is no dedicated Master-Master replication.

- *Section 14.6.6.6.1, “Configuring MySQL Nodes”*
- *Section 14.6.6.6.2, “Creating Replica Users”*
- *Section 14.6.6.6.3, “Transferring Database”*
- *Section 14.6.6.6.4, “Configuring Replication for SITE01-MYSQL01”*
- *Section 14.6.6.6.5, “Configuring Replication for SITE01-MYSQL02”*

14.6.6.6.1. Configuring MySQL Nodes

Add all replication-required values to the **`my.cnf`** configuration file on both hosts. Important values are given below:

server-id

The unique id used for replication. In basic configuration, the following IDs are used for the nodes:

- SITE01-MYSQL01 = server-id 11

- SITE01-MYSQL02 = server-id 12
- SITE02-MYSQL01 = server-id 21

auto_increment_increment = 10

Specifies that the MySQL node that auto increment values to be incremented by 10 instead of default value 1. This setting is used to overvcome split brain situations.

- log-bin=mysql-bin
- binlog-format = 'ROW'

These two values activate row-based binary logging for replication. This is a mandatory operation.

expire_logs_days=10

The number of days to keep the bin log files. It enables MySQL to automatically clean up the obsolete bin log files and assist to avoid situations where disks run out of space.

max_binlog_size=100M

The maximum size of each bin log file.

max_connections=1200

The maximum number of client connections, implies the number of CloudPlatform hosts multiplied with 350. As you have four CloudPlatform Management Servers, two each in datacenter Primary and datacenter secondary, the total number of connections should be set to 1200.

innodb_flush_log_at_trx_commit=1/ sync_binlog=1

For the maximum reliability and consistency in a replication setup using InnoDB with transactions, set these two values. Both values have impact on the performance on the MySQL server but also lowers the risk of data inconsistency / loss.

To configure:

1. Configure **`my.cnf`** of SITE01-MYSQL01:

   An example configuration:

```
[mysqld]
datadir=/var/lib/mysql
socket=/var/lib/mysql/mysql.sock
# Disabling symbolic-links is recommended to prevent assorted security risks
symbolic-links=0
# unique ID for each MySQL Node that is part of the replication group
server-id=11
innodb_rollback_on_timeout=1
innodb_lock_wait_timeout=600
default-time-zone='+00:00'
default-storage-engine=InnoDB
# maximum number of client connections (number of CloudPlatform hosts multiplied with
 350)
max_connections=1200
# set the innodb_buffer_pool_size (recommended is 70% of system RAM)
Innodb_buffer_pool_size=5500m
default-storage-engine=InnoDB
character-set-server=utf8
transaction-isolation=READ-COMMITTED
```

```
Page 13
# activation of row based Binary logging for replication
log-bin=mysql-bin
binlog-format=ROW
# commit configuration for log files
innodb_flush_log_at_trx_commit=1
sync_binlog=1
#Bin logs cleanup configuration
expire_logs_days=10
max_binlog_size=100M
[mysqld_safe]
log-error=/var/log/mysqld.log
pid-file=/var/run/mysqld/mysqld.pid
```

2. Configuration **`my.cnf`** of SITE01-MYSQL02:

   An example configuration:

```
[mysqld]
datadir=/var/lib/mysql
socket=/var/lib/mysql/mysql.sock
# Disabling symbolic-links is recommended to prevent assorted security risks
symbolic-links=0
# unique ID for each MySQL Node that is part of the replication group
server-id=12
innodb_rollback_on_timeout=1
innodb_lock_wait_timeout=600
default-time-zone='+00:00'
default-storage-engine=InnoDB
# maximum number of client connections (number of CloudPlatform hosts multiplied with
 350)
max_connections=1200
default-storage-engine=InnoDB
character-set-server=utf8
transaction-isolation=READ-COMMITTED
Page 14
# activation of row based Binary logging for replication
log-bin=mysql-bin
binlog-format=ROW
# commit configuration for log files
innodb_flush_log_at_trx_commit=1
sync_binlog=1
#Bin logs cleanup configuration
expire_logs_days=10
max_binlog_size=100M
#Parameters to solve split brain problem
auto_increment_increment=10
auto_increment_offset=2
[mysqld_safe]
log-error=/var/log/mysqld.log
pid-file=/var/run/mysqld/mysqld.pid
```

3. Copy and paste the content of the configuration files to the appropriate MySQL server and adjust the value for *`server-id`* if needed.

4. Restart the mysqld service.

   Restarting the mysqld service on the node SITE01-MYSQL01 will interrupt the database connection from the CloudPlatform nodes. The VMs will continue to work but the Management Server will not be available. It is recommended to stop the CloudPlatform services, cloudstack-management and cloudstack-usage, on all the CloudPlatform hosts first.

14.6.6.6.2. Creating Replica Users

The next step is to create MySQL users with the right replication slave. MySQL will use these accounts to connect to the Master nodes and replicate the database changes.

You need two replication users, one for replicating from SITE01-MYSQL01 and one for replicating from SITE01-MYSQL02. Create both users just on the node MYSQL01-SITE01 because you may dump the MySQL database to have the MySQL users consistent across both hosts.

1. On SITE01-MYSQL01, run the following mysql commands:

```
# GRANT REPLICATION SLAVE ON *.* to 'replication'@'SITE01-MYSQL01.inf.local'
  IDENTIFIED BY 'DEMO-CUSTOMER123';
```

```
# GRANT REPLICATION SLAVE ON *.* to 'replication'@'SITE01-MYSQL02.inf.local'
  IDENTIFIED BY 'DEMO-CUSTOMER123';
```

2. Update the MySQL user and privileges:

   On SITE01-MYSQL01 execute the following command:

```
Flush privileges;
```

3. Check if the login works:

   a. Try to connect to the MySQL node from SITE01-MYSQL01 with the user replication:

```
# mysql –h SITE01-MYSQL01.inf.local –u replication
```

   b. Try to connect to the MySQL node from SITE01-MYSQL02 with the user replication:

```
# mysql –h SITE01-MYSQL02.inf.local –u replication
```

14.6.6.6.3. Transferring Database

A copy of all databases (mysql, cloud, cloud_usage) are transferred to the server SITE01-MYSQL02. MySQL replication cannot replicate whole databases, only changes can be replicated.

1. Lock the tables to prepare for a consistent database dump.

   On SITE01-MYSQL01 run the following:

```
# flush tables with read lock;
```

2. Display the current Master relay log file and position.

   On SITE01-MYSQL01 run the following:

```
# show master status;
```

   Write the relay log file and location down, as they will be needed later for the initial replication.

3. Create the dump of all databases:

Now don't close the MySQL shell, the read lock is valid only for the current session. While the MySQL shell is still open, open a second command line window to create the SQL dump.

On SITE01-MYSQL01, run the following CLI commands:

```
# cd /tmp
```

```
# mysqldump –u root –p –all-databases > SITE01-MYSQL01-DBs.sql
```

The warning message is displayed: Skipping the data of table mysql.event is normal, server related events will not backed up per default as they only valid for that particular server.

4. Release the table lock:

   On SITE01-MYSQL01, run the following:

   ```
   # unlock tables;
   ```

5. Transfer the database dump to SITE01-MYSQL02:

   On SITE01-MYSQL01, run the following:

   ```
   # scp SITE01-MYSQL01-DBs.sql root@SITE01-MYSQL02.inf.local:/tmp
   ```

6. Import the database dump.

   On SITE01-MYSQL02, run the following:

   ```
   # mysql –u root –p < SITE01-MYSQL01-DBs.sql
   ```

7. Check if the databases are now present on SITE01-MYSQL02.

   On SITE01-MYSQL02, run the following:

   ```
   # show databases;
   ```

14.6.6.6.4. Configuring Replication for SITE01-MYSQL01

The next step is to configure the node SITE01-MYSQL02 to replicate all databases from the primary Master SITE01-MYSQL01

1. Set the replication target.

   On SITE01-MYSQL02, run the following:

   ```
   # CHANGE MASTER TO MASTER_HOST='SITE01-MYSQL01.inf.local’,
   MASTER_USER='replication', MASTER_PASSWORD='DEMO-CUSTOMER123',
   MASTER_LOG_FILE='mysql-bin.000004', MASTER_LOG_POS=973;
   ```

2. Start the replication of SITE01-MYSQL02.

   On SITE01-MYSQL02, run the following:

```
# start slave;
```

3. Check the status of the replication.

   On SITE01-MYSQL02, run the following:

   ```
   # show slave status;
   ```

   Consider the following values:

   Slave_IO_State (Waiting for master to send event) : The slave host can connect to the master.

   Slave_IO_Running : Replication is currently active.

   Seconds_Behind_Master (0) : The databases of both nodes are currently synchronized.

   LAST_IO_ERROR (0) : No consistency problems are found.

14.6.6.6.5. Configuring Replication for SITE01-MYSQL02

The next step is to configure the node SITE01-MYSQL01 to replicate all databases from the secondary Master SITE01-MYSQL02.

1. Set a read lock on all tables of SITE01-MYSQL02 On SITE01-MYSQL02.

   Run the following:

   ```
   # flush tables with read lock;
   ```

2. Display the current Master relay log file and position.

   On SITE01-MYSQL02 run the following:

   ```
   # show master status;
   ```

   Write the relay log file and position down, as they are needed later for the initial replication.

3. Set the replication target On SITE01-MYSQL02, run the following:

   ```
   # CHANGE MASTER TO MASTER_HOST='SITE01-MYSQL01.inf.local’,
   MASTER_USER='replication', MASTER_PASSWORD='DEMO-CUSTOMER123',
   MASTER_LOG_FILE='mysql-bin.000004', MASTER_LOG_POS=23311687;
   ```

4. Start the replication of SITE01-MYSQL01.

   On SITE01-MYSQL01, run the following:

   ```
   # start slave;
   ```

5. Check the status of the replication.

   On SITE01-MYSQL01, run the following:

```
# show slave status;
```

Consider the following values:

Slave_IO_State (Waiting for the Master to send event) : The slave host can connect to the master.

Slave_IO_Running : Replication is currently active.

Seconds_Behind_Master (0) : The databases of both nodes are currently synchronized.

LAST_IO_ERROR (0) : No consistency problems are found.

The configuration of the bi-directional Master-Master replication is now done. Each node replicates the database changes from the other.

14.6.6.7. Configuring CloudPlatform to Use Database High Availability

The last step is to configure the CloudPlatform Hosts to use the database HA feature. All related settings are done in the configuration file db.properties under /etc/cloudstack/management/. These settings must be done for each CloudPlatform node that connects to the MySQL Master-Master Cluster. In the example configuration, we just have two CloudPlatform nodes (INF-CCP01 at SITE01 and INF-CCP02 at SITE02) The node INF-CCP01 has the IP-Address 10.0.0.21 and INF-CCP02 has the IP-Address 10.0.0.22.

- *Section 14.6.6.7.1, “Configuring INF-CCP01”*
- *Section 14.6.6.7.2, “Configuring INF-CCP02”*

14.6.6.7.1. Configuring INF-CCP01

1. Modify the **`db.properties`** file:

   a. On INF-CCP01, run the following CLI command:

   ```
   # vi /etc/cloudstack/management.db.properties
   ```

   b. Set the following values:

   ```
   db.cloud.slaves=SITE01-MYSQL02.inf.local
   db.cloud.host=SITE01-MYSQL01.inf.local
   db.usage.host=SITE01-MYSQL01.inf.local
   db.usage.slaves=SITE01-MYSQL02.inf.local
   db.ha.enabled=true
   ```

2. Restart CloudPlatform services:

   On INF-CCP01, run the following CLI commands:

   ```
   # service cloudstack-management restart
   # service cloudstack-usage restart
   ```

Example **`db.properties`** configuration file for INF-CCP01:

```
db.usage.maxActive=100
db.cloud.name=cloud
```

```
db.cloud.password=ENC(IoYAL1db4qcAQDbj4jfiQn4zA+AgVyiy)
db.usage.maxWait=10000
Page 24
db.usage.maxIdle=30
db.cloud.autoReconnectForPools=true
db.cloud.trustStore=
db.awsapi.host=localhost
db.cloud.port=3306
db.cloud.testOnBorrow=true
db.usage.name=cloud_usage
db.cloud.poolPreparedStatements=false
db.cloud.maxIdle=30
db.ha.enabled=true
db.simulator.maxActive=250
db.usage.port=3306
db.cloud.url.params=prepStmtCacheSize=517&cachePrepStmts=true
db.cloud.keyStorePassword=
db.cloud.queriesBeforeRetryMaster=5000
db.usage.failOverReadOnly=false
db.cloud.secondsBeforeRetryMaster=100
db.awsapi.username=cloud
db.cloud.username=cloud
cluster.node.IP=10.1.20.21
db.cloud.initialTimeout=3600
cluster.servlet.port=9090
db.cloud.slaves=SITE01-MYSQL02.inf.local
db.cloud.host=SITE01-MYSQL01.inf.local
db.simulator.maxIdle=30
db.cloud.maxActive=250
db.usage.host=SITE01-MYSQL01.inf.local
db.usage.slaves=SITE01-MYSQL02.inf.local
region.id=1
db.simulator.port=3306
db.cloud.testWhileIdle=true
db.usage.password=ENC(KTnGDzBgQctA7UoB1rqRzqO6P598reph))
db.usage.initialTimeout=3600
db.cloud.reconnectAtTxEnd=true
db.cloud.maxWait=10000
db.usage.autoReconnect=true
db.cloud.timeBetweenEvictionRunsMillis=40000
db.cloud.trustStorePassword=
db.cloud.minEvictableIdleTimeMillis=240000
db.cloud.validationQuery=SELECT 1
db.cloud.useSSL=false
db.simulator.password=cloud
db.cloud.keyStore=
Page 25
db.usage.reconnectAtTxEnd=true
db.simulator.autoReconnect=true
db.simulator.username=cloud
db.ha.loadBalanceStrategy=com.cloud.utils.db.StaticStrategy
db.awsapi.name=cloudbridge
db.cloud.failOverReadOnly=false
db.usage.username=cloud
db.awsapi.password=cloud
db.cloud.encrypt.secret=ENC(ZnHZLPc3M815TxU9rn5nJOJLyGePk9dc)
db.cloud.encryption.type=file
db.simulator.maxWait=10000
# awsapi database settings
# Simulator database settings
db.cloud.autoReconnect=true
db.awsapi.port=3306
db.simulator.host=localhost
db.usage.url.params=
db.simulator.name=simulator
db.usage.queriesBeforeRetryMaster=5000
db.usage.secondsBeforeRetryMaster=3600
```

```
db.usage.autoReconnectForPools=true
```

14.6.6.7.2. Configuring INF-CCP02

1. Modify the db.properties file.

    a. On INF-CCP02, run the following CLI command:

    ```
    # vi /etc/cloudstack/management.db.properties
    ```

    b. Set the following values:

    ```
    db.cloud.slaves=SITE01-MYSQL02.inf.local
    db.cloud.host=SITE01-MYSQL01.inf.local
    db.usage.host=SITE01-MYSQL01.inf.local
    db.usage.slaves=SITE01-MYSQL02.inf.local
    db.ha.enabled=true
    ```

2. Restart CloudPlatform services On INF-CCP02, run the following CLI commands:

    ```
    # service cloudstack-management restart
    # service cloudstack-usage resta
    ```

Example **`db.properties`** configuration file for INF-CCP02:

```
db.usage.maxActive=100
db.cloud.name=cloud
db.cloud.password=ENC(IoYAL1db4qcAQDbj4jfiQn4zA+AgVyiy)
db.usage.maxWait=10000
db.usage.maxIdle=30
db.cloud.autoReconnectForPools=true
db.cloud.trustStore=
db.awsapi.host=localhost
db.cloud.port=3306
db.cloud.testOnBorrow=true
db.usage.name=cloud_usage
db.cloud.poolPreparedStatements=false
db.cloud.maxIdle=30
db.ha.enabled=true
db.simulator.maxActive=250
db.usage.port=3306
db.cloud.url.params=prepStmtCacheSize=517&cachePrepStmts=true
db.cloud.keyStorePassword=
db.cloud.queriesBeforeRetryMaster=5000
db.usage.failOverReadOnly=false
db.cloud.secondsBeforeRetryMaster=100
db.awsapi.username=cloud
db.cloud.username=cloud
cluster.node.IP=10.1.20.22
db.cloud.initialTimeout=3600
cluster.servlet.port=9090
db.cloud.slaves=SITE01-MYSQL02.inf.local
db.cloud.host=SITE01-MYSQL01.inf.local
db.simulator.maxIdle=30
db.cloud.maxActive=250
db.usage.host=SITE01-MYSQL01.inf.local
db.usage.slaves=SITE01-MYSQL02.inf.local
region.id=1
db.simulator.port=3306
db.cloud.testWhileIdle=true
db.usage.password=ENC(KTnGDzBgQctA7UoB1rqRzqO6P598reph))
```

```
db.usage.initialTimeout=3600
db.cloud.reconnectAtTxEnd=true
Page 27
db.cloud.maxWait=10000
db.usage.autoReconnect=true
db.cloud.timeBetweenEvictionRunsMillis=40000
db.cloud.trustStorePassword=
db.cloud.minEvictableIdleTimeMillis=240000
db.cloud.validationQuery=SELECT 1
db.cloud.useSSL=false
db.simulator.password=cloud
db.cloud.keyStore=
db.usage.reconnectAtTxEnd=true
db.simulator.autoReconnect=true
db.simulator.username=cloud
db.ha.loadBalanceStrategy=com.cloud.utils.db.StaticStrategy
db.awsapi.name=cloudbridge
db.cloud.failOverReadOnly=false
db.usage.username=cloud
db.awsapi.password=cloud
db.cloud.encrypt.secret=ENC(ZnHZLPc3M815TxU9rn5nJOJLyGePk9dc)
db.cloud.encryption.type=file
db.simulator.maxWait=10000
# awsapi database settings
# Simulator database settings
db.cloud.autoReconnect=true
db.awsapi.port=3306
db.simulator.host=localhost
db.usage.url.params=
db.simulator.name=simulator
db.usage.queriesBeforeRetryMaster=5000
db.usage.secondsBeforeRetryMaster=3600
db.usage.autoReconnectForPools=true
```

14.6.6.8. Failover Testing

Test failover before releasing the platform into production:

1. Make sure that the CloudPlatform services are started On INF-CCP01, run the following CLI commands:

```
# service cloudstack-management status
# service cloudstack-usage status
```

2. On INF-CCP02, run the following CLI commands:

```
# service cloudstack-management status
# service cloudstack-usage status
```

3. Check the database connections.

   a. On SITE01-MYSQL01, run the following :

```
# show processlist;
```

   b. On SITE01-MYSQL02, run the following:

```
# show processlist;
```

Currently, all the CloudPlatform hosts are currently connected only to the primary Master SITE01-MYSQL01. The secondary Master is currently connected only to the primary master to check for datanase changes.

4. Make sure that the CloudPlatform UI is reachable
5. Simulate an unexpected outage of the primary master. For example, perform a force shutdown of the VM.
6. Check the logs on the CloudPlatform hosts.

   On INF-CCP01 and INF-CCP02, run the following cli commands:

   ```
   # tail –f /var/log/cloudstack/management/management-server.log
   ```

   CloudPlatform loses the connection to the primary Master MySQL node but immediately switches to the slave Master node and everything continues to work.

   The CloudPlatform hosts will frequently try to reconnect to the failed Master. The time frame is controlled by the following values:

   ```
   db.cloud.secondsBeforeRetryMaster=100
   db.cloud.initialTimeout=3600
   db.cloud.queriesBeforeRetryMaster=5000
   db.usage.queriesBeforeRetryMaster=5000
   db.usage.secondsBeforeRetryMaster=3600
   ```

7. Bring the primary Master back online and check the replication state.

   a. On SITE01-MYSQL01 run the following:

      ```
      # show slave status;
      ```

   b. On SITE02-MYSQL02 run the following:

      ```
      # show slave status;
      ```

As you can see, the replication is re-established and no errors were found. This indicates that no inconsistencies were found, and no additional steps are needed. If you want to force the CloudPlatform hosts to reconnect to the primary Master immediately, you have to restart the cloudstack-management and cloudstack-usage service on all CloudPlatform nodes. Otherwise, the CloudPlatform nodes will automatically failback if the reconnect values specified in the config file are expired.

14.6.6.9. Recovering an Inconsistent Master

Due the nature of the asynchronous replication that is used in MySQL 5.1, replication can be broken after the failed Master comes back online. In such a situation, you will see an error message like:

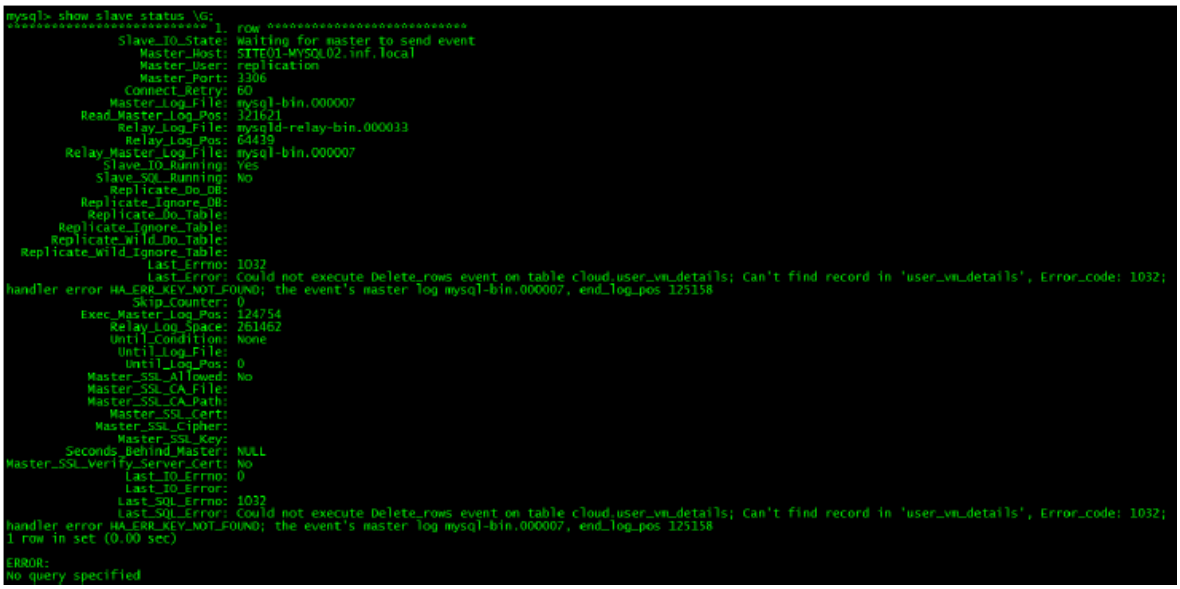

The only safe way to recover from such an inconsistent state is to reestablish the replication with a new and consistent state of the database. That means, you have to create a dump from the consistent database of the current active node, in this example SITE01-MYSQL02.

Do not use SQL_SLAVE_SKIP_COUNTER to resolve database inconsistencies. This will not fix the inconsistent state, rather it just drops the database entry that is causing the conflict. Doing this can lead to serious damage of the cloud databases.

Perform the following to recover an inconsistent primary master:

1. Shutdown all CloudPlatform services.

   To avoid the CloudPlatform nodes are falling back to the inconsistent Master node, stop the cloudstack-management and cloudstack-usage services.

2. Stop the replication on both nodes.

   On SITE01-MYSQL01 and SITE01-MYSQL02, run the following:

   ```
   stop slave;
   ```

3. Create a database dump of the databases, cloud and cloud_portal.

   On SITE01-MYSQL02, run the following CLI commands:

   ```
   # mysqldump –master-data –u root –p cloud > SITE01-MYSQL02-consistent-cloud-db.sql
   # mysqldump –master-data –u root –p cloud_portal > SITE01-MYSQL02-consistent-
   cloud_portal-db.sql
   ```

4. Transfer the database dumps to the Host SITE01-MYSQL01.

   On SITE01-MYSQL02, execute the following CLI command:

   ```
   # scp SITE01_MYSQL02-consistent-cloud-db.sql SITE01-MYSQL02-consistent-cloud_usage-db.sql
   root@SITE01-MYSQL01.inf.local:/tmp
   ```

5. Restore the databases.

On SITE01-MYSQL01, execute the following CLI commands:

```
# mysql –u root –p cloud < SITE01-MYSQL02-consistent-cloud-db.sql
# mysql –u root –p cloud_portal < SITE01-MYSQL02-consitent-cloud_portal-db.sql
```

6. Restart the replication.

   On SITE01-MYSQL01, run the following:

   ```
   # start slave;
   ```

7. Check the replication state.

   As you can see, the SLAVE replication from SITE01-MYSQL01 to SITE01-MYSQL02 is now re-established. The last step is to fix the replication in the other direction, from SITE01-MYSQL02 to SITE01-MYSQL02.

8. Show the current relay-log file and log-position of SITE01-MYSQL01.

   On SITE01-MYSQL01, run the following:

   ```
   # flush tables with read lock;
   # show master status;
   ```

9. Configure the replication to start at the corresponding position.

   On SITE01-MYSQL02, run the following:

   ```
   # change master to master_log_file='mysql-bin.00020', master_log_pos=15457519
   # start slave;
   ```

10. Check the state of the replication:

    As you can see, the Slave replication from SITE02-MYSQL01 to SITE01-MYSQL01 is now re-established as well. The restore process is now finished.

14.6.6.10. Integrating Additional MySQL Slave Hosts

As already mentioned is, it´s important for DEMO-CUSTOMER to make sure that critical components are high-available across datacenters. The service must continue to work even if the primary datacenter faces an outage. This section will describe the steps to add an additional MySQL Slave node to the infrastructure. That host will run on the second datacenter and replicates all DB changes from the Master-Master Cluster at the primary datacenter. The two nodes SITE01-MYSQL01 and SITE02-MYSQL02 are already configured, each node replicates db changes to the other node. We will now configure a third node “SITE02-MYSQL01”. That node will monitor the DBs of the node SITE01-MYSQL02 and replicates changes asynchronous

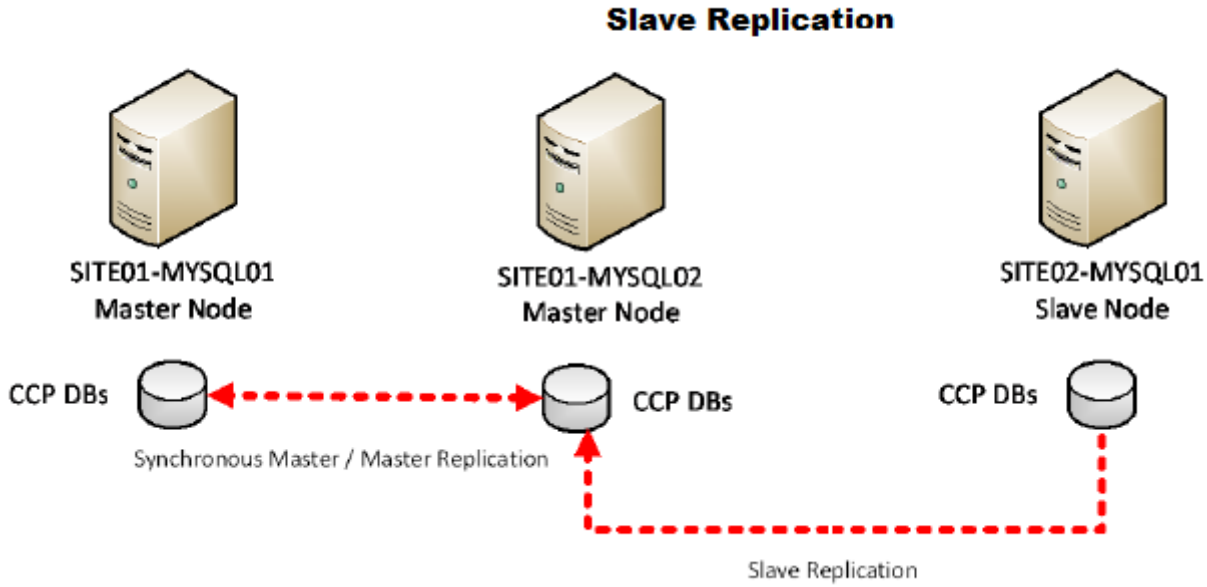

You are in effect using a replication chain. The advantage of using SITE01-MYSQL02 as the source for the replication to the second datacenter is, if the primary Master that is used by CloudPlatform goes down, the automatic failover will switch the CloudPlatform hosts to SITE01-MYSQL02, and SITE02-MYSQL01 will continue to replicate without interruption. The database replication to the second datacenter is still active, no manual steps are required. If SITE01-MYSQL02 goes down unexpectedly before the original master SITE01-MYSQL01 is restored, the Slave node in the second datacenter has the latest database version and can promoted to the new Master.

- *Section 14.6.6.10.1, "Installing MySQL on Slave Node"*
- *Section 14.6.6.10.2, "Preparing the Slave Node"*
- *Section 14.6.6.10.3, "Creating a Replication User"*
- *Section 14.6.6.10.4, "Transferring Database"*
- *Section 14.6.6.10.5, "Configuring Replication for SITE02-MYSQL01"*

14.6.6.10.1. Installing MySQL on Slave Node

The first step is to install MySQL 5.1 on the slave node. Note: It is recommended to use the MySQL installer that ships with the CloudPlatform installation. To guarantee that the replication work as expected, the same MySQL versions needed across all nodes.

1. On the node SITE02-MYSQL01 extract the CloudPlatform tar ball:

```
# tar xvf CloudPlatform-4.3.0.0-rhel6.4.tar.gz
```

2. Launch the CloudPlatform installer and select the option "D" to install MySQL 5.1.73.
3. Set the mysql service to start on system-boot:

```
# chkconfig mysqld on
```

4. Make the necessary changes to the **`my.cnf`** configuration file.

To guarantee a working failover, all servers have consistent settings. Example configuration:

```
[mysqld]
datadir=/var/lib/mysql
socket=/var/lib/mysql/mysql.sock
user=mysql
# Disabling symbolic-links is recommended to prevent assorted security risks
symbolic-links=0

server-id=21

innodb_rollback_on_timeout=1
innodb_lock_wait_timeout=600
default-time-zone='+00:00'
max_connections=1000
log-bin=mysql-bin
binlog-format = 'ROW'

#innodb_buffer_pool_size=5500


[mysqld_safe]
log-error=/var/log/mysqld.log
pid-file=/var/run/mysqld/mysqld.pid
```

This is just an example configuration. Use the appropriate settings, which were set during the installation of the first MySQL node. Note down the value for *`server-id`*, which has to be unique across all host in a replication configuration.

5. Restart the mysqld service on the node SITE02-MYSQL01:

```
# service mysqld restart
```

14.6.6.10.2. Preparing the Slave Node

The next step is to prepare the Slave node for the replication.

1. Add the following lines to the **`my.cnf`** configuration file

```
#Added for MySQL Replication
server-id=21
relay_log=mysql-relay-bin
innodb_flush_log_at_trx_commit=1
sync_binlog=1
read-only=1
```

2. Restart the mysqld daemon.

```
# service mysqld restart
```

14.6.6.10.3. Creating a Replication User

Create a MySQL users with the right replication Slave. The MySQL Slave node SITE02-MYSQL01 will connect with that account to the node SITE01-MYSQL02 to replicate the database changes.

1. On SITE01-MYSQL02, run the following:

```
# GRANT REPLICATION SLAVE ON *.* to 'replication'@'SITE02-MYSQL01.inf.local'
 IDENTIFIED BY 'DEMO-CUSTOMER123';
```

2. Update the MySQL user and privileges:

   On SITE01-MYSQL02 run the following:

   ```
   # flush privileges;
   ```

3. Check if the login works:

   Try to connect to the MySQL node from SITE02-MYSQL01 with the user replication:

   ```
   # mysql –h SITE01-MYSQL02.inf.local –u replication -p
   ```

14.6.6.10.4. Transferring Database

Create a consistent backup of all databases (mysql, cloud, cloud_usage) from the SITE01-MYSQL02 noce and restore them at the Slave node, SITE02-MYSQL01.

1. Lock the tables to prepare for a consistent database dump.

   On SITE01-MYSQL02 run the following:

   ```
   # flush tables with read lock;
   ```

2. Display the current Master relay log file and position.

   On SITE01-MYSQL01 run the following:

   ```
   # show master status;
   ```

   Write the relay log file and position down as they are later needed for the initial replication.

3. Create the dump of all databases.

   Don't close the MySQL shell because the read lock is always only valid for the current session. While the MySQL shell is still open, open a second command line window and create the SQL dump.

   On SITE01-MYSQL03, run the following CLI commands

   ```
   # cd /tmp
   # mysqldump –u root –p –all-databases > SITE01-MYSQL02-DBs.sql
   ```

   You see the warning message: Skipping the data of table mysql.event is normal, server related events will not backed up per default as they only valid for that particular server.

4. Release the table Lock On SITE01-MYSQL02, run the following CLI command:

   ```
   # unlock tables;
   ```

5. Transfer the database dump to SITE02-MYSQL01

   On SITE01-MYSQL02, run the following CLI command:

   ```
   # scp SITE01-MYSQL02-DBs.sql root@SITE02-MYSQL01.inf.local:/tmp
   ```

6. Import the database dump.

   On SITE02-MYSQL01, run the following CLI command:

   ```
   # mysql –u root –p < SITE01-MYSQL02-DBs.sql
   ```

   If you get an error message, such as ERROR 1449 The user specified as a definer does not exist, log in to the mysql shell, run flush privileges and import commands again.

7. Check if the databases are now present.

   On SITE02-MYSQL01 On SITE02-MYSQL01, run the following:

   ```
   # show databases;
   ```

14.6.6.10.5. Configuring Replication for SITE02-MYSQL01

Set the replication target.

1. On SITE02-MYSQL01, run the following:

   ```
   #CHANGE MASTER TO MASTER_HOST='SITE01-MYSQL02.inf.local’,
   MASTER_USER='replication', MASTER_PASSWORD='DEMO-CUSTOMER123',
   MASTER_LOG_FILE='mysql-bin.000004', MASTER_LOG_POS=973;
   ```

2. Start the replication of SITE01-MYSQL01

   On SITE01-MYSQL01, run the following:

   ```
   # start slave;
   ```

3. Check the status of the replication.

   On SITE02-MYSQL01, run the following:

   ```
   # show slave status;
   ```

   Important are the following values:

   Slave_IO_State (Waiting for master to send event): Slave host can connect to the master

   Slave_IO_Running: Replication is currently active

   Seconds_Behind_Master (0): the databases of both nodes are currently in sync

   LAST_IO_ERROR (0): no consistency problems were found

The slave node of the second datacenter now replicates all database changes.

14.6.6.11. Failover Switching to the Second Datacenter

In an unlikely event that either the Master-Master cluster or the entire primary datacenter becomes unavailable, the Slave node SITE02-MYSQL01 of the second datacenter can be promoted to Master, to bring the platform back online. It´s necessary to reconfigure the remaining CloudPlatform hosts to use the Slave host.

In general, the steps are:

- Promote the slave host to the new Master.
- Point the CloudPlatform hosts to the new Master.
- Restart the CloudPlatform services.

The following steps simulate the outage of the primary datacenter, meaning the entire MySQL Master-Master cluster and the CloudPlatform hosts of the main datacenter are down. If CloudPlatform cannot reach a MySQL server that is specified in the db.properties, the CloudPlatform service performs a self-fence to avoid data inconsistencies. For the following steps this means that the CloudPlatform has attempted to re-start is would normally be in a “hang” status.

- *Section 14.6.6.11.1, “Promoting a Slave to Master”*
- *Section 14.6.6.11.2, “Pointing CloudPlatform to the New Master”*

14.6.6.11.1. Promoting a Slave to Master

1. Disable the read-only mode of the slave node.

   On the node SITE02-MYSQL01, open the **`my.cnf`** file and set read-only to “0”:

2. Restart the mysqld service to apply the change:

   ```
   # service mysqld restart
   ```

3. Clear the slave replication configuration.

   On the node SITE02-MYSQL01, run the following:

   ```
   # stop slave;
   # reset slave;
   ```

   The slave is now working as a Master and ready to accept connections.

14.6.6.11.2. Pointing CloudPlatform to the New Master

The CloudPlatform hosts of the second datacenter must now re-configured to use the new database.

1. Change the target database.

   On each of the CloudPlatform nodes:

   a. Edit the **`db.properties`** file under **`/etc/cloudstack/management/db.properties`**

   b. Set the following values:

   ```
   db.cloud.host=SITE02-MYSQL01.inf.local
                    db.usage.host=SITE02-MYSQL01.inf.local
   ```

2. Restart the CloudPlatform services.

   The CloudPlatform nodes are now back online again:

A Failback from a failed datacenter or Master-Master cluster follows the same steps as already described.

In general, you have to perform the following steps:

1. Create a consistent dump of the databases, cloud and cloud_portal, on the node SITE02-MYSQL01.
2. Restore the dump on the nodes SITE01-MYSQL01 and SITE01-MYSQL02.
3. Re-establish the replication between the two nodes SITE01-MYSQL01 and SITE01-MYSQL02 in both directions.
4. Re-establish the slave replication from node SITE02-MYSQL01 to SITE01-MYSQL02 and set it back to read-only.
5. Reconfigure the CloudPlatform nodes to use SITE01-MYSQL01 as primary master and SITE01-MYSQL02 as slave.
6. Restart the CloudPlatform services.

14.7. Limiting the Rate of API Requests

You can limit the rate at which API requests can be placed for each account. This is useful to avoid malicious attacks on the Management Server, prevent performance degradation, and provide fairness to all accounts.

If the number of API calls exceeds the threshold, an error message is returned for any additional API calls. The caller will have to retry these API calls at another time.

14.7.1. Configuring the API Request Rate

To control the API request rate, use the following global configuration settings:

- api.throttling.enabled - Enable/Disable API throttling. By default, this setting is false, so API throttling is not enabled.
- api.throttling.interval (in seconds) - Time interval during which the number of API requests is to be counted. When the interval has passed, the API count is reset to 0.
- api.throttling.max - Maximum number of APIs that can be placed within the api.throttling.interval period.
- api.throttling.cachesize - Cache size for storing API counters. Use a value higher than the total number of accounts managed by the cloud. One cache entry is needed for each account, to store the running API total for that account.

14.7.2. Limitations on API Throttling

The following limitations exist in the current implementation of this feature.

Note

Even with these limitations, CloudPlatform is still able to effectively use API throttling to avoid malicious attacks causing denial of service.

- In a deployment with multiple Management Servers, the cache is not synchronized across them. In this case, CloudPlatform might not be able to ensure that only the exact desired number of API requests are allowed. In the worst case, the number of API calls that might be allowed is (number of Management Servers) * (api.throttling.max).
- The API commands resetApiLimit and getApiLimit are limited to the Management Server where the API is invoked.

Managing the Cloud

This chapter describes how you can mange the cloud to ensure its optimal performance.

15.1. Reporting CPU Sockets

CloudPlatform manages different types of hosts that contains one or more physical CPU sockets. CPU socket is considered as a unit of measure used for licensing and billing cloud infrastructure. CloudPlatform provides both UI and API support to collect the CPU socket statistics for billing purpose. The Infrastructure tab has a new tab for CPU sockets. You can view the statistics for CPU sockets managed by CloudPlatform, which in turn reflects the size of the cloud. The CPU Socket page will give you the number of hosts and sockets used for each hypervisor type.

Cores per socket functionality allows you to specify the number of cores should be created per socket/ CPU when a VM is created. You can specify the value for cores per socket only while registering a template. For example:

```
http://<managementip>:8080/client/api?command=registerTemplate&response=
json&sessionkey=jWSl&name=Windows7&displayText=Windows7&url=
http%3A%2F%2F10.100.100.200%2Fcloudstack%2Ftemplates%2Fmyvol.vhd&zoneid
=-1&format=VHD&isextractable=false&passwordEnabled=false&isdynamicallyscalable
=false&osTypeId=80fc1f21b&hypervisor=Hyperv&ispublic=true&isfeatured
=true&isrouting=false&details[0].cpu.corespersocket=2
```

Note down the details parameter while registering the template. A template which is already registered can also be updated with cores per socket details. Example api is:

```
http://<management ip>:8080/client/api?command=addResourceDetail&response
=json&sessionkey=<mykey>&resourceid=a1935bea5&details[0].key
=cpu.corespersocket&details[0].value=4&resourcetype=template
```

Multiple cores per socket, as per the template details, will be created for any VM created from the template. Ensure that the service offering you create matches the same number of cores and use the same service offering and template while creating a VM.

Locate the CPU sockets details in the CloudPlatform UI as follows:

1. Log in to the CloudPlatform UI.
2. In the left navigation bar, click Infrastructure.
3. On CPU Sockets, click View all.

   The CPU Socket page is displayed. The page shows the number of hosts and CPU sockets based on hypervisor types.

   CPU sockets are displayed for XenServer version 6.2 and beyond, KVM, Hyper-V and VMware hypervisors.

   This feature is not available for XenServer versions prior to 6.2 as they don't support retrieving CPU socket information. Additionally, this feature is not supported for Baremetal.

15.2. Using Tags to Organize Resources in the Cloud

A tag is a key-value pair that stores metadata about a resource in the cloud. Tags are useful for categorizing resources. For example, you can tag a user VM with a value that indicates the user's city

of residence. In this case, the key would be "city" and the value might be "Toronto" or "Tokyo." You can then request CloudPlatform to find all resources that have a given tag; for example, VMs for users in a given city.

You can tag a user virtual machine, volume, snapshot, guest network, template, ISO, firewall rule, port forwarding rule, public IP address, security group, load balancer rule, project, VPC, network ACL, or static route. You can not tag a remote access VPN.

You can work with tags through the UI or through the API commands createTags, deleteTags, and listTags. You can define multiple tags for each resource. There is no limit on the number of tags you can define. Each tag can be up to 255 characters long. Users can define tags on the resources they own, and administrators can define tags on any resources in the cloud.

An optional input parameter, "tags," exists on many of the list* API commands. The following example shows how to use this new parameter to find all the volumes having tag region=canada OR tag city=Toronto:

```
command=listVolumes
    &listAll=true
    &tags[0].key=region
    &tags[0].value=canada
    &tags[1].key=city
    &tags[1].value=Toronto
```

The following API commands have the "tags" input parameter:

- listVirtualMachines
- listVolumes
- listSnapshots
- listNetworks
- listTemplates
- listIsos
- listFirewallRules
- listPortForwardingRules
- listPublicIpAddresses
- listSecurityGroups
- listLoadBalancerRules
- listProjects
- listVPCs
- listNetworkACLs
- listStaticRoutes

15.3. Setting Configuration Parameters

15.3.1. About Configuration Parameters

CloudPlatform provides a variety of settings you can use to set limits, configure features, and enable or disable features in the cloud. Once your Management Server is running, you might need to set some of these configuration parameters, depending on what optional features you are setting up. You can set default values at the global level, which will be in effect throughout the cloud unless you override them at a lower level. You can make local settings, which will override the global configuration parameter values, at the level of an account, zone, cluster, or primary storage.

The documentation for each CloudPlatform feature should direct you to the names of the applicable parameters. The following table shows a few of the more useful parameters.

Field	Value
management.network.cidr	A CIDR that describes the network that the management CIDRs reside on. This variable must be set for deployments that use vSphere. It is recommended to be set for other deployments as well. Example: 192.168.3.0/24.
xen.setup.multipath	For XenServer nodes, this is a true/false variable that instructs CloudStack to enable iSCSI multipath on the XenServer Hosts when they are added. The default value is 'false'. Set it to true if you would like CloudStack to enable multipath. If this is true for a NFS-based deployment multipath will still be enabled on the XenServer host. However, this does not impact NFS operation and is harmless.
secstorage.allowed.internal.sites	This is used to protect your internal network from rogue attempts to download arbitrary files using the template download feature. This is a comma-separated list of CIDRs. If a requested URL matches any of these CIDRs the Secondary Storage VM will use the private network interface to fetch the URL. Other URLs will go through the public interface. We suggest you set this to 1 or 2 hardened internal machines where you keep your templates. For example, set it to 192.168.1.66/32.
use.local.storage	Determines whether CloudStack will use storage that is local to the Host for data disks, templates, and snapshots. By default CloudStack will not use this storage. You should change this to true if you want to use local storage and you understand the reliability and feature drawbacks to choosing local storage.
host	This is the IP address of the Management Server. If you are using multiple Management Servers you should enter a load balanced IP address that is reachable via the private network.

Field	Value
default.page.size	Maximum number of items per page that can be returned by a CloudStack API command. The limit applies at the cloud level and can vary from cloud to cloud. You can override this with a lower value on a particular API call by using the page and page size API command parameters. For more information, see the Developer's Guide. Default: 500.
ha.tag	The label you want to use throughout the cloud to designate certain hosts as dedicated HA hosts. These hosts will be used only for HA-enabled VMs that are restarting due to the failure of another host. For example, you could set this to ha_host. Specify the ha.tag value as a host tag when you add a new host to the cloud.

15.3.2. Setting Global Configuration Parameters

Use the following steps to set global configuration parameters. These values will be the defaults in effect throughout your CloudPlatform deployment.

1. Log in to the UI as administrator.
2. In the left navigation bar, click Global Settings.
3. In Select View, choose one of the following:
   - Global Settings. This displays a list of the parameters with brief descriptions and current values.
   - LDAP Configuration. Displays the LDAP details that you have configured. Also, you can click **Configure LDAP** to create a new LDAP configuration.
   - Hypervisor Capabilities. This displays a list of hypervisor versions with the maximum number of guests supported for each.
4. Use the search box to narrow down the list to those you are interested in.
5. In the Actions column, click the Edit icon to modify a value. If you are viewing Hypervisor Capabilities, you must click the name of the hypervisor first to display the editing screen.

15.3.3. Setting Local Configuration Parameters

Use the following steps to set local configuration parameters for an account, zone, cluster, or primary storage. These values will override the global configuration settings.

1. Log in to the UI as administrator.
2. In the left navigation bar, click Infrastructure or Accounts, depending on where you want to set a value.
3. Find the name of the particular resource that you want to work with. For example, if you are in Infrastructure, click View All on the Zones, Clusters, or Primary Storage area.
4. Click the name of the resource where you want to set a limit.

5. Click the Settings tab.
6. Use the search box to narrow down the list to those you are interested in.
7. In the Actions column, click the Edit icon to modify a value.

15.3.4. Granular Global Configuration Parameters

The following global configuration parameters have been made more granular. The parameters are listed under three different scopes: account, cluster, and zone.

Field	Field	Value
account	remote.access.vpn.client.iprange	The range of IPs to be allocated to remotely access the VPN clients. The first IP in the range is used by the VPN server.
account	allow.public.user.templates	If false, users will not be able to create public templates.
account	use.system.public.ips	If true and if an account has one or more dedicated public IP ranges, IPs are acquired from the system pool after all the IPs dedicated to the account have been consumed.
account	use.system.guest.vlans	If true and if an account has one or more dedicated guest VLAN ranges, VLANs are allocated from the system pool after all the VLANs dedicated to the account have been consumed.
cluster	cluster.storage.allocated.capacity.notificationthreshold	The percentage, as a value between 0 and 1, of allocated storage utilization above which alerts are sent that the storage is below the threshold.
cluster	cluster.storage.capacity.notificationthreshold	The percentage, as a value between 0 and 1, of storage utilization above which alerts are sent that the available storage is below the threshold.
cluster	cluster.cpu.allocated.capacity.notificationthreshold	The percentage, as a value between 0 and 1, of cpu utilization above which alerts are sent that the available CPU is below the threshold.
cluster	cluster.memory.allocated.capacity.notificationthreshold	The percentage, as a value between 0 and 1, of memory utilization above which alerts

Field	Field	Value
		are sent that the available memory is below the threshold.
cluster	cluster.cpu.allocated.capacity.disablethreshold	The percentage, as a value between 0 and 1, of CPU utilization above which allocators will disable that cluster from further usage. Keep the corresponding notification threshold lower than this value to be notified beforehand.
cluster	cluster.memory.allocated.capacity.disablethreshold	The percentage, as a value between 0 and 1, of memory utilization above which allocators will disable that cluster from further usage. Keep the corresponding notification threshold lower than this value to be notified beforehand.
cluster	cpu.overprovisioning.factor	Used for CPU over-provisioning calculation; the available CPU will be the mathematical product of actualCpuCapacity and cpu.overprovisioning.factor.
cluster	mem.overprovisioning.factor	Used for memory over-provisioning calculation.
cluster	vmware.reserve.cpu	Specify whether or not to reserve CPU when not over-provisioning; In case of CPU over-provisioning, CPU is always reserved.
cluster	vmware.reserve.mem	Specify whether or not to reserve memory when not over-provisioning; In case of memory over-provisioning memory is always reserved.
zone	pool.storage.allocated.capacity.disablethreshold	The percentage, as a value between 0 and 1, of allocated storage utilization above which allocators will disable that pool because the available allocated storage is below the threshold.
zone	pool.storage.capacity.disablethreshold	The percentage, as a value between 0 and 1, of storage utilization above which allocators will disable the pool

Field	Field	Value
		because the available storage capacity is below the threshold.
zone	storage.overprovisioning.factor	Used for storage over-provisioning calculation; available storage will be the mathematical product of actualStorageSize and storage.overprovisioning.factor.
zone	network.throttling.rate	Default data transfer rate in megabits per second allowed in a network.
zone	guest.domain.suffix	Default domain name for VMs inside a virtual networks with a router.
zone	router.template.xen	Name of the default router template on Xenserver.
zone	router.template.kvm	Name of the default router template on KVM.
zone	router.template.vmware	Name of the default router template on VMware.
zone	enable.dynamic.scale.vm	Enable or diable dynamically scaling of a VM.
zone	use.external.dns	Bypass internal DNS, and use the external DNS1 and DNS2
zone	blacklisted.routes	Routes that are blacklisted cannot be used for creating static routes for a VPC Private Gateway.

15.4. Changing the Database Configuration

The CloudPlatform Management Server stores database configuration information, for example host name, port, credentials, in the file /etc/cloudstack/management/db.properties. To effect a change, edit this file on each Management Server, then restart the Management Server.

15.5. Administrator Alerts

CloudPlatform provides alerts to help with the management of the cloud. Alerts are notices to an administrator that an error has occurred in the cloud. Alerts are displayed on the Dashboard in the CloudPlatform UI, and can also be sent to an email address, an external SNMP manager, or an external syslog manager.

You can find the MIB file for the SNMP alerts in the CloudPlatform tar ball at the root.

Alerts are generated under the following circumstances:

- The Management Server cluster runs low on CPU, memory, or storage resources
- The Management Server loses heartbeat from a Host for more than 3 minutes

- The Host cluster runs low on CPU, memory, or storage resources

For a list of CloudPlatform alerts, see the Alerts section. For the most up-to-date list, call the listAlerts API.

Note

In addition to alerts, CloudPlatform also generates events. Unlike alerts, which indicate issues of concern, events track all routine user and administrator actions in the cloud. For example, every time a guest VM starts, this creates an associated event. Events are stored in the Management Server's database.

15.5.1. Customizing Alerts with Global Configuration Settings

To exercise some control over how alerts behave, you can use the global configuration settings. You can configure recipient and sender email addresses, SMTP server and authentication, timeouts, frequency intervals, and more. To access these settings through the CloudPlatform UI, go to the Global Settings screen (click the Global Settings button in the left navbar) and type "alert" in the search box.

The following table shows some of the more useful alert configuration settings.

Configuration Variable	Description
alert.email.addresses	One or more email addresses to which alerts will be sent. There are several companion settings for the SMTP host, From: address, etc.
alert.purge.delay	A useful tuning parameter. Alerts older than the specified number days will be deleted, freeing up resources. If you want to keep alerts forever, you can set this to 0.
alert.wait	Another useful tuning parameter, this one controls the sensitivity of the alerting mechanism. CloudPlatform will wait the specified number of seconds before generating an alert on a disconnected agent. By setting this value high, you can potentially reduce the number of alerts, allowing for issues to self-correct. However, this should be used with caution, as it can also obscure issues and delay the fix.

15.5.2. Sending Alerts to External SNMP and Syslog Managers

In addition to showing administrator alerts on the Dashboard in the CloudPlatform UI and sending them in email, CloudPlatform can also send the same alerts to external SNMP or Syslog management software. This is useful if you prefer to use an SNMP or Syslog manager to monitor your cloud.

The alerts which can be sent are listed in the Alerts section. *Appendix B, Alerts*. You can also display the most up to date list by calling the API command listAlerts.

15.5.2.1. SNMP Alert Details

The supported protocol is SNMP version 2.

Each SNMP trap contains the following information: message, podId, dataCenterId, clusterId, and generationTime.

15.5.2.2. Syslog Alert Details

CloudPlatform generates a syslog message for every alert. Each syslog message incudes the fields alertType, message, podId, dataCenterId, and clusterId, in the following format. If any field does not have a valid value, it will not be included.

```
Date severity_level Management_Server_IP_Address/Name  alertType:: value dataCenterId:: value
  podId:: value  clusterId:: value  message:: value
```

For example:

```
Mar  4 10:13:47    WARN    localhost    alertType:: managementNode message:: Management
 server node 127.0.0.1 is up
```

15.5.2.3. Configuring SNMP and Syslog Managers

To configure one or more SNMP managers or Syslog managers to receive alerts from CloudPlatform:

1. For an SNMP manager, install the CloudPlatform MIB file, which is available in the root directory inside the tarball for CloudPlatform, on your SNMP manager system. This maps the SNMP OIDs to trap types that can be more easily read by users. The file must be publicly available. For more information on how to install this file, consult the documentation provided with the SNMP manager.

2. Edit the file /etc/cloudstack/management/log4j-cloud.xml.

   ```
   # vi /etc/cloudstack/management/log4j-cloud.xml
   ```

3. Add an entry using the syntax shown below. Follow the appropriate example depending on whether you are adding an SNMP manager or a Syslog manager. To specify multiple external managers, separate the IP addresses and other configuration values with commas (,).

   

   Note

   The recommended maximum number of SNMP or Syslog managers is 20 for each.

   The following example shows how to configure two SNMP managers at IP addresses 10.1.1.1 and 10.1.1.2. Substitute your own IP addresses, ports, and communities. Do not change the other values (name, threshold, class, and layout values).

   ```
   <appender name="SNMP" class="org.apache.cloudstack.alert.snmp.SnmpTrapAppender">
     <param name="Threshold" value="WARN"/>  <!-- Do not edit. The alert feature assumes
    WARN. -->
     <param name="SnmpManagerIpAddresses" value="10.1.1.1,10.1.1.2"/>
     <param name="SnmpManagerPorts" value="162,162"/>
     <param name="SnmpManagerCommunities" value="public,public"/>
   ```

```
  <layout class="org.apache.cloudstack.alert.snmp.SnmpEnhancedPatternLayout"> <!-- Do not
 edit -->
    <param name="PairDelimeter" value="//"/>
    <param name="KeyValueDelimeter" value="::"/>
  </layout>
</appender>
```

For sending the syslog alerts to non-default port in Goleta, you must specify the IP addresses as follows:

```
<param name="SnmpManagerIpAddresses" value="10.1.1.1:897,10.12.1.3:977"/>
```

The following example shows how to configure two Syslog managers at IP addresses 10.1.1.1 and 10.1.1.2. Substitute your own IP addresses. You can set Facility to any syslog-defined value, such as LOCAL0 - LOCAL7. Do not change the other values.

```
<appender name="ALERTSYSLOG">
  <param name="Threshold" value="WARN"/>
  <param name="SyslogHosts" value="10.1.1.1,10.1.1.2"/>
  <param name="Facility" value="LOCAL6"/>
  <layout>
    <param name="ConversionPattern" value=""/>
  </layout>
</appender>
```

4. If your cloud has multiple Management Server nodes, repeat these steps to edit log4j-cloud.xml on every instance.

5. If you have made these changes while the Management Server is running, wait a few minutes for the change to take effect.

Troubleshooting: If no alerts appear at the configured SNMP or Syslog manager after a reasonable amount of time, it is likely that there is an error in the syntax of the <appender> entry in log4j-cloud.xml. Check to be sure that the format and settings are correct.

You can modify the **`log4j-cloud.xml`** file for sending the Management Server logs to Syslog manager. You can modify the appender entry in the **`log4j-cloud.xml`** file as follows:

```
<appender name="SYSLOG" class="org.apache.log4j.net.SyslogAppender">
        <param name="Threshold" value="WARN"/>
         <param name="SyslogHost" value="localhost"/>
        <param name="Facility" value="LOCAL6"/>
        <layout class="org.apache.log4j.PatternLayout">
        <param name="ConversionPattern" value="%-5p [%c{3}] (%t:%x) %m%n"/>
    </layout>
</appender>
```

An **`appender`** entry supports a syslog host. If you want to add more syslog hosts, you must add **`appender`** entries for each of them. Then, you must add these appenders to the **`logger`** entry in the log4j-cloud.xml file.

```
<logger name="org.apache.cloudstack.alerts" additivity="false">
      <level value="WARN"/>
      <appender-ref ref="SYSLOG"/>
      <appender-ref ref="CONSOLE"/>
      <appender-ref ref="FILE"/>
      <appender-ref ref="SNMP"/>
```

```
      <appender-ref ref="ALERTSYSLOG"/>
</logger>
```

15.5.2.4. Deleting an SNMP or Syslog Manager

To remove an external SNMP manager or Syslog manager so that it no longer receives alerts from CloudPlatform, remove the corresponding entry from the file /etc/cloudstack/management/log4j-cloud.xml.

15.6. Customizing the Network Domain Name

The root administrator can optionally assign a custom DNS suffix at the level of a network, account, domain, zone, or entire CloudPlatform installation, and a domain administrator can do so within their own domain. To specify a custom domain name and put it into effect, follow these steps.

1. Set the DNS suffix at the desired scope

    - At the network level, the DNS suffix can be assigned through the UI when creating a new network, as described in *Section 9.8.1, "Adding an Additional Guest Network"* or with the updateNetwork command in the CloudPlatform API.
    - At the account, domain, or zone level, the DNS suffix can be assigned with the appropriate CloudPlatform API commands: createAccount, editAccount, createDomain, editDomain, createZone, or editZone.
    - At the global level, use the configuration parameter guest.domain.suffix. You can also use the CloudPlatform API command updateConfiguration. After modifying this global configuration, restart the Management Server to put the new setting into effect.

2. To make the new DNS suffix take effect for an existing network, call the CloudPlatform API command updateNetwork. This step is not necessary when the DNS suffix was specified while creating a new network.

The source of the network domain that is used depends on the following rules.

- For all networks, if a network domain is specified as part of a network's own configuration, that value is used.
- For an account-specific network, the network domain specified for the account is used. If none is specified, the system looks for a value in the domain, zone, and global configuration, in that order.
- For a domain-specific network, the network domain specified for the domain is used. If none is specified, the system looks for a value in the zone and global configuration, in that order.
- For a zone-specific network, the network domain specified for the zone is used. If none is specified, the system looks for a value in the global configuration.

15.7. Stopping and Restarting the Management Server

The root administrator will need to stop and restart the Management Server from time to time.

For example, after changing a global configuration parameter, a restart is required. If you have multiple Management Server nodes, restart all of them to put the new parameter value into effect consistently throughout the cloud..

To stop the Management Server, issue the following command at the operating system prompt on the Management Server node:

```
# service cloudstack-management stop
```

To start the Management Server:

```
# service cloudstack-management start
```

Working with Usage

The Usage Server is an optional, separately-installed part of CloudPlatform that provides aggregated usage records which you can use to create billing integration for CloudPlatform. The Usage Server works by taking data from the events log and creating summary usage records that you can access using the listUsageRecords API call.

The usage records show the amount of resources, such as VM run time or template storage space, consumed by guest instances.

The Usage Server runs at least once per day. It can be configured to run multiple times per day.

In addition to this, CloudPlatform 4.5 supports Citrix Insight Service to provide you with data analytics specific to your environment.

16.1. listUsageRecords API Usage Types

The following table describes the usage types that **`listUsageRecords`** API uses:

Usage Type	Description
RUNNING_VM	Track the total running time of a VM per usage record period. If the VM is upgraded during the usage period, you will get a separate usage record for the upgraded VM.
ALLOCATED_VM	Tracks the duration of existence of a VM. It is calculated from the time the VM is created until the time it is destroyed. This usage type helps determine the usage for specific templates such as Windows-based templates.
IP_ADDRESS	Tracks the public IP Address owned by the account.
NETWORK_BYTES_SENT	Tracks the total number of bytes sent by all the VMs for an account.
NETWORK_BYTES_RECEIVED	Tracks the total number of bytes received by all the VMs for an account.
VOLUME	Tracks the duration of existence of a volume. It is calculated from the time the volume is created until the time it is destroyed.
TEMPLATE	Tracks the duration of existence of a template (either created from a snapshot or uploaded to the cloud). It is calculated from the time the template is created until the time it is destroyed. Also, this usage type returns the size of the template.
ISO	Tracks the usage duration of an ISO. It is calculated from the time the ISO is uploaded to the cloud until the time it is removed from the cloud. Also, this usage type returns the size of the ISO.

Usage Type	Description
SNAPSHOT	Tracks the duration of existence of a snapshot. It is calculated from the time the snapshot is created until the time it is destroyed.
LOAD_BALANCER_POLICY	Tracks the duration of existence of a load balancer policy. It is calculated from the time the load balancer policy is created until the time it is removed.
PORT_FORWARDING_RULE	Tracks the duration of existence of a port forwarding rule. It is calculated from the time the port forwarding rule is created until the time it is removed.
NETWORK_OFFERING	Tracks the duration of existence of a network offering. It is calculated from the time the network offering is assigned until the time it is removed.
VPN_USERS	Tracks the duration for which the VPN user is activated. It is calculated from the time the VPN user record is created until the time it is removed.

16.2. Configuring the Usage Server

To configure the usage server:

1. Be sure the Usage Server has been installed. This requires extra steps beyond just installing the CloudPlatform software. See Installing the Usage Server (Optional) in the Advanced Installation Guide.
2. Log in to the CloudPlatform UI as administrator.
3. Click Global Settings.
4. In Search, type usage. Find the configuration parameter that controls the behavior you want to set. See the table below for a description of the available parameters.
5. In Actions, click the Edit icon.
6. Type the desired value and click the Save icon.
7. Restart the Management Server (as usual with any global configuration change) and also the Usage Server:

```
# service cloud-management restart
# service cloud-usage restart
```

The following table shows the global configuration settings that control the behavior of the Usage Server.

Parameter Name	Description
enable.usage.server	Whether the Usage Server is active.
usage.aggregation.timezone	Time zone of usage records. Set this if the usage records and daily job execution are in different

Parameter Name	Description

	time zones. For example, with the following settings, the usage job will run at PST 00:15 and generate usage records for the 24 hours from 00:00:00 GMT to 23:59:59 GMT:
	
	`usage.stats.job.exec.time = 00:15`
	`usage.execution.timezone = PST`
	`usage.aggregation.timezone = GMT`
	
	Default: GMT

usage.execution.timezone	The time zone of usage.stats.job.exec.time. Default: The time zone of the management server.
usage.sanity.check.interval	The number of days between sanity checks. Set this in order to periodically search for records with erroneous data before issuing customer invoices. For example, this checks for VM usage records created after the VM was destroyed, and similar checks for templates, volumes, and so on. It also checks for usage times longer than the aggregation range. If any issue is found, the alert ALERT_TYPE_USAGE_SANITY_RESULT = 21 is sent.
usage.stats.job.aggregation.range	The time period in minutes between Usage Server processing jobs. For example, if you set it to 1440, the Usage Server will run once per day. If you set it to 600, it will run every ten hours. In general, when a Usage Server job runs, it processes all events generated since usage was last run. There is special handling for the case of 1440 (once per day). In this case the Usage Server does not necessarily process all records since Usage was last run. CloudPlatform assumes that you require processing once per day for the previous, complete day's records. For example, if the current day is October 7, then it is assumed you would like to process records for October 6, from midnight to midnight. CloudPlatform assumes this "midnight to midnight" is relative to the usage.execution.timezone. Default: 1440
usage.stats.job.exec.time	The time when the Usage Server processing will start. It is specified in 24-hour format (HH:MM) in the time zone of the server, which should be GMT. For example, to start the Usage job at 10:30 GMT, enter "10:30".

Parameter Name	Description
	If usage.stats.job.aggregation.range is also set, and its value is not 1440, then its value will be added to usage.stats.job.exec.time to get the time to run the Usage Server job again. This is repeated until 24 hours have elapsed, and the next day's processing begins again at usage.stats.job.exec.time. Default: 00:15.

For example, suppose that your server is in GMT, your user population is predominantly in the East Coast of the United States, and you would like to process usage records every night at 2 AM local (EST) time. Choose these settings:

- enable.usage.server = true
- usage.execution.timezone = America/New_York
- usage.stats.job.exec.time = 07:00. This will run the Usage job at 2:00 AM EST. Note that this will shift by an hour as the East Coast of the U.S. enters and exits Daylight Savings Time.
- usage.stats.job.aggregation.range = 1440

With this configuration, the Usage job will run every night at 2 AM EST and will process records for the previous day's midnight-midnight as defined by the EST (America/New_York) time zone.

Note

Because the special value 1440 has been used for usage.stats.job.aggregation.range, the Usage Server will ignore the data between midnight and 2 AM. That data will be included in the next day's run.

16.3. Setting Usage Limits

CloudPlatform provides several administrator control points for capping resource usage by users. Some of these limits are global configuration parameters. Others are applied at the ROOT domain and may be overridden on a per-account basis.

Aggregate limits may be set on a per-domain basis. For example, you may limit a domain and all sub domains to the creation of 100 VMs.

There are two types of limits you can set. First, you can set limits based on the resource count, that is, restricting a user or domain on the basis of the number of VMs, volumes, or snapshots used.

In addition, CloudPlatform supports the customization model—need-basis usage, such as large VM or small VM. The resource types are broadly classified as CPU, RAM, Primary Storage, and Secondary Storage. The root administrator can impose resource usage limits by the following resource types for Domains, Projects, and Accounts:

- CPUs

- Memory (RAM)
- Primary Storage (Volumes)
- Secondary Storage (Snapshots, Templates, ISOs)

To control the behaviour of the needs-based usage feature, use the following configuration parameters:

Parameter Name	Description
max.account.cpus	Maximum number of CPU cores that can be used for an account. Default is 40.
max.account.ram (MB)	Maximum RAM that can be used for an account. Default is 40960.
max.account.primary.storage (GB)	Maximum primary storage space that can be used for an account. Default is 20*10.
max.account.secondary.storage (GB)	Maximum secondary storage space that can be used for an account. Default is 20*20.
max.project.cpus	Maximum number of CPU cores that can be used for an account. Default is 40.
max.project.ram (MB)	Maximum RAM that can be used for an account. Default is 40960.
max.project.primary.storage (GB)	Maximum primary storage space that can be used for an account. Default is 20*10.
max.project.secondary.storage (GB)	Maximum secondary storage space that can be used for an account. Default is 20*20.
max.project.network.rate (Mbps)	Maximum network rate that can be used for an account. Default is 200.

16.3.1. Globally Configured Limits

In a zone, the guest virtual network has a 24 bit CIDR by default. This limits the guest virtual network to 254 running instances. It can be adjusted as needed, but this must be done before any instances are created in the zone. For example, 10.1.1.0/22 would provide for ~1000 addresses.

The following table lists limits set in the Global Configuration:

Parameter Name	Definition
max.account.public.ips	Number of public IP addresses that can be owned by an account
max.account.snapshots	Number of snapshots that can exist for an account
max.account.templates	Number of templates that can exist for an account
max.account.user.vms	Number of virtual machine instances that can exist for an account
max.account.volumes	Number of disk volumes that can exist for an account
max.template.iso.size	Maximum size for a downloaded template or ISO in GB
max.volume.size.gb	Maximum size for a volume in GB
network.throttling.rate	The default data transfer rate in megabits per second allowed in network.
snapshot.max.hourly	Maximum recurring hourly snapshots to be retained for a volume. If the limit is reached, early snapshots from the start of the hour are deleted so that newer ones can be saved. This limit does not apply to manual snapshots. If set to 0, recurring hourly snapshots can not be scheduled
snapshot.max.daily	Maximum recurring daily snapshots to be retained for a volume. If the limit is reached, snapshots from the start of the day are deleted so that newer ones can be saved. This limit does not apply to manual snapshots. If set to 0, recurring daily snapshots can not be scheduled
snapshot.max.weekly	Maximum recurring weekly snapshots to be retained for a volume. If the limit is reached, snapshots from the beginning of the week are deleted so that newer ones can be saved. This limit does not apply to manual snapshots. If set to 0, recurring weekly snapshots can not be scheduled
snapshot.max.monthly	Maximum recurring monthly snapshots to be retained for a volume. If the limit is reached, snapshots from the beginning of the month are deleted so that newer ones can be saved. This limit does not apply to manual snapshots. If set to 0, recurring monthly snapshots can not be scheduled.

To modify global configuration parameters, use the global configuration screen in the CloudPlatform UI. See Setting Global Configuration Parameters

16.3.2. Default Account Resource Limits

You can limit resource use by accounts. The default limits are set by using global configuration parameters, and they affect all accounts within a cloud. The relevant parameters are those beginning with max.account, for example: max.account.snapshots.

To override a default limit for a particular account, set a per-account resource limit.

1. Log in to the CloudPlatform UI.
2. In the left navigation tree, click Accounts.
3. Select the account you want to modify. The current limits are displayed. A value of -1 shows that there is no limit in place.
4. Click the Edit button. 

16.3.3. Per-Domain Limits

CloudPlatform allows the configuration of limits on a domain basis. With a domain limit in place, all users still have their account limits. They are additionally limited, as a group, to not exceed the resource limits set on their domain. Domain limits aggregate the usage of all accounts in the domain as well as all accounts in all subdomains of that domain. Limits set at the root domain level apply to the sum of resource usage by the accounts in all domains and sub-domains below that root domain.

To set a domain limit:

1. Log in to the CloudPlatform UI.
2. In the left navigation tree, click Domains.
3. 3. Select the domain you want to modify. The current domain limits are displayed. A value of -1 shows that there is no limit in place.
4. Click the Edit button 

CloudPlatform API

The CloudPlatform API is a low level API that has been used to implement the CloudPlatform web UIs. It is also a good basis for implementing other popular APIs such as EC2/S3 and emerging DMTF standards.

Many CloudPlatform API calls are asynchronous. These will return a Job ID immediately when called. This Job ID can be used to query the status of the job later. Also, status calls on impacted resources will provide some indication of their state.

The API has a REST-like query basis and returns results in XML or JSON.

See the Developer's Guide and the API Reference.

17.1. Provisioning and Authentication API

CloudPlatform expects that a customer will have their own user provisioning infrastructure. It provides APIs to integrate with these existing systems where the systems call out to CloudPlatform to add/ remove users..

CloudPlatform supports pluggable authenticators. By default, CloudPlatform assumes it is provisioned with the user's password, and as a result authentication is done locally. However, external authentication is possible as well. For example, see Using an LDAP Server for User Authentication .

17.2. Allocators

CloudPlatform enables administrators to write custom allocators that will choose the Host to place a new guest and the storage host from which to allocate guest virtual disk images.

17.3. User Data and Meta Data

CloudPlatform provides API access to attach up to 32KB of user data to a deployed VM. Deployed VMs also have access to instance metadata via the virtual router.

User data can be accessed once the IP address of the virtual router is known. Once the IP address is known, use the following steps to access the user data:

1. Run the following command to find the virtual router.

```
# cat /var/lib/dhclient/dhclient-eth0.leases | grep dhcp-server-identifier | tail -1
```

2. Access user data by running the following command using the result of the above command

```
# curl http://10.1.1.1/latest/user-data
```

Meta Data can be accessed similarly, using a URL of the form http://10.1.1.1/latest/meta-data/ {metadata type}. (For backwards compatibility, the previous URL http://10.1.1.1/latest/{metadata type} is also supported.) For metadata type, use one of the following:

- service-offering. A description of the VMs service offering
- availability-zone. The Zone name
- local-ipv4. The guest IP of the VM

- local-hostname. The hostname of the VM
- public-ipv4. The first public IP for the router. (E.g. the first IP of eth2)
- public-hostname. This is the same as public-ipv4
- instance-id. The instance name of the VM

Tuning

This section provides tips on how to improve the performance of your cloud.

18.1. Performance Monitoring

Host and guest performance monitoring is available to end users and administrators. This allows the user to monitor their utilization of resources and determine when it is appropriate to choose a more powerful service offering or larger disk.

18.2. Increase Management Server Maximum Memory

If the Management Server is subject to high demand, the default maximum JVM memory allocation can be insufficient. To increase the memory:

1. Edit the Tomcat configuration file:

   ```
   /etc/cloud/management/tomcat6.conf
   ```

2. Change the command-line parameter -XmxNNNm to a higher value of N.

   For example, if the current value is -Xmx128m, change it to -Xmx1024m or higher.

3. To put the new setting into effect, restart the Management Server.

   ```
   # service cloud-management restart
   ```

For more information about memory issues, see "FAQ: Memory" at *Tomcat Wiki.*[1]

18.3. Set Database Buffer Pool Size

It is important to provide enough memory space for the MySQL database to cache data and indexes:

1. Edit the Tomcat configuration file:

   ```
   /etc/my.cnf
   ```

2. 2. Insert the following line in the [mysqld] section, below the datadir line. Use a value that is appropriate for your situation. We recommend setting the buffer pool at 40% of RAM if MySQL is on the same server as the management server or 70% of RAM if MySQL has a dedicated server. The following example assumes a dedicated server with 1024M of RAM.

   ```
   innodb_buffer_pool_size=700M
   ```

3. Restart the MySQL service.

   ```
   # service mysqld restart
   ```

[1] http://wiki.apache.org/tomcat/FAQ/Memory

For more information about the buffer pool, see "The InnoDB Buffer Pool" at *MySQL Reference Manual*[2].

18.4. Set and Monitor Total VM Limits per Host

The CloudPlatform administrator should monitor the total number of VM instances in each cluster, and disable allocation to the cluster if the total is approaching the maximum that the hypervisor can handle. Be sure to leave a safety margin to allow for the possibility of one or more hosts failing, which would increase the VM load on the other hosts as the VMs are automatically redeployed. Consult the documentation for your chosen hypervisor to find the maximum permitted number of VMs per host, then use CloudPlatform global configuration settings to set this as the default limit. Monitor the VM activity in each cluster at all times. Keep the total number of VMs below a safe level that allows for the occasional host failure. For example, if there are N hosts in the cluster, and you want to allow for one host in the cluster to be down at any given time, the total number of VM instances you can permit in the cluster is at most (N-1) * (per-host-limit). Once a cluster reaches this number of VMs, use the CloudPlatform UI to disable allocation of more VMs to the cluster.

18.5. Configure XenServer dom0 Memory

Note

The following configuration applies to the versions lower than XenServer 6.2.0

Configure the XenServer dom0 settings to allocate more memory to dom0. This can enable XenServer to handle larger numbers of virtual machines. We recommend 2940 MB of RAM for XenServer dom0. For instructions on how to do this, see *Citrix Knowledgebase Article*[3].The article refers to XenServer 5.6, but the same information applies to XenServer 6

[2] http://dev.mysql.com/doc/refman/5.5/en/innodb-buffer-pool.html
[3] http://support.citrix.com/article/CTX126531

Troubleshooting

19.1. Configuring Citrix Insight Service for CloudPlatform

Citrix Insight Service (CIS) is an utility to collect log files required for diagnostics and analysis of issues. CIS enables you to automatically detects configuration issues from customer database dump and logs. Additionally, CIS provides you with product insights, feature usage and telemetry.

You can run CIS service as an utility as explained in *Section 19.1.2, "CIS Workflow"*.

19.1.1. CIS Analytics Overview

The following types of analytics have been provided for CloudPlatform by the CIS tool:

- Insights: These are the aggregate of all the file bundles uploaded to CIS. As name suggests, insights provide in depth distribution data. For example, distribution of guest networks and API use.
- Content plug-ins: These are specific to a single file bundle, and provide information related to that specific environment. For example, resource usage and upgrade history.
- Diagnostic plug-ins: These are nothing but health check plug-ins, which provide information specific to a single file bundle. For example, alerts related to a specific condition based on the analysis.

19.1.2. CIS Workflow

1. Run the utility from **`/usr/share/cloudstack-management/util`**.

   ```
   # ./CloudStack-bugtool
   ```

   For more information on usage, see *Section 19.1.3, "CIS Utility Usage"*.

   The output is the instrumentation bundle. The zip file is created under **`/tmp`** in following format: **`cloud-bugtool_<DATE and TIME>.<RANDOM STRING>.zip`**

2. Collect **`cloud-bugtool_<DATE and TIME>.<RANDOM STRING>.zip`**
3. Upload the file to *Citrix Insight Services*[1].

   CIS automatically detects issues and recommends resolutions.

19.1.3. CIS Utility Usage

The location of the CloudStack-bugtool utility in CloudPlatform Management server is **`/usr/share/cloudstack-management/util/cloud-bugtool`**.

Usage: **`./CloudStack-bugtool`**

The utility takes the following arguments:

- -f --full: Full mode, collects everything.
- -m --minimal: Collects only system properties and the latest CloudPlatform log files.

[1] https://cis.citrix.com

- -d --nodb: Do not include cloud database dump
- h --help: Display the help options.

The utility operates on the following modes:

- NORMAL: This is the latest mode. Considers system information, cloud database, latest CloudPlatfrom log files, and latest system log files.
- MINIMAL: Considers system information and latest CloudPlatform log files.
- FULL: Considers system information, cloud database, all the CloudPlatform log files, all the system log files.

19.2. Events

An event is essentially a significant or meaningful change in the state of both virtual and physical resources associated with a cloud environment. Events are used by monitoring systems, usage and billing systems, or any other event-driven workflow systems to discern a pattern and make the right business decision. In CloudPlatform an event could be a state change of virtual or psychical resources, an action performed by an user (action events), or policy based events (alerts).

19.2.1. Event Logs

There are two types of events logged in the CloudPlatform Event Log. Standard events log the success or failure of an event and can be used to identify jobs or processes that have failed. There are also long running job events. Events for asynchronous jobs log when a job is scheduled, when it starts, and when it completes. Other long running synchronous jobs log when a job starts, and when it completes. Long running synchronous and asynchronous event logs can be used to gain more information on the status of a pending job or can be used to identify a job that is hanging or has not started. The following sections provide more information on these events.

19.2.2. Event Notification

Event notification framework provides a means for the Management Server components to publish and subscribe to CloudPlatform events. Event notification is achieved by implementing the concept of event bus abstraction in the Management Server. An event bus is introduced in the Management Server that allows the CloudPlatformcomponents and extension plug-ins to subscribe to the events by using the Advanced Message Queuing Protocol (AMQP) client. In CloudPlatform, a default implementation of event bus is provided as a plug-in that uses the RabbitMQ AMQP client. The AMQP client pushes the published events to a compatible AMQP server. Therefore all the CloudPlatform events are published to an exchange in the AMQP server.

A new event for state change, resource state change, is introduced as part of Event notification framework. Every resource, such as user VM, volume, NIC, network, public IP, snapshot, and template, is associated with a state machine and generates events as part of the state change. That implies that a change in the state of a resource results in a state change event, and the event is published in the corresponding state machine on the event bus. All the CloudPlatform events (alerts, action events, usage events) and the additional category of resource state change events, are published on to the events bus.

Use Cases

The following are some of the use cases:

- Usage or Billing Engines: A third-party cloud usage solution can implement a plug-in that can connects to CloudPlatform to subscribe to CloudPlatform events and generate usage data. The usage data is consumed by their usage software.
- AMQP plug-in can place all the events on the a message queue, then a AMQP message broker can provide topic-based notification to the subscribers.
- Publish and Subscribe notification service can be implemented as a pluggable service in CloudPlatform that can provide rich set of APIs for event notification, such as topics-based subscription and notification. Additionally, the pluggable service can deal with multi-tenancy, authentication, and authorization issues.

Configuration

As a CloudPlatform administrator, perform the following one-time configuration to enable event notification framework. At run time no changes can control the behaviour.

1. Create a directory, **`META-INF/cloudstack/core/`**, under **`/etc/cloudstack/management/`**.
2. Create a new configuration file, **`spring-event-bus-context.xml`** in the **`/etc/cloudstack/management/META-INF/cloudstack/core`** directory.
3. Define a bean named **`eventNotificationBus`** as follows:
   - name : Specify a name for the bean.
   - server : The name or the IP address of the RabbitMQ AMQP server.
   - port : The port on which RabbitMQ server is running.
   - username : The user name associated with the account to access the RabbitMQ server.
   - password : The password associated with the user name of the account to access the RabbitMQ server.
   - exchange : The exchange name on the RabbitMQ server where CloudPlatform events are published.

   An example bean is given below:

```
<bean id="eventNotificationBus"
 class="org.apache.cloudstack.mom.rabbitmq.RabbitMQEventBus">
    <property name="name" value="eventNotificationBus"/>
    <property name="server" value="127.0.0.1"/>
    <property name="port" value="5672"/>
    <property name="username" value="guest"/>
    <property name="password" value="guest"/>
   <property name="exchange" value="cloudstack-events"/>
   </bean>
```

   The **`eventNotificationBus`** bean represents the **`org.apache.cloudstack.mom.rabbitmq.RabbitMQEventBus`** class.

   If you want to use encrypted values for the user name and password, you have to include a bean to pass those as variables from a credentials file.

   An example is given below:

```
  <beans xmlns="http://www.springframework.org/schema/beans"
      xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xmlns:context="http://www.springframework.org/schema/context"
      xmlns:aop="http://www.springframework.org/schema/aop"
      xsi:schemaLocation="http://www.springframework.org/schema/beans
        http://www.springframework.org/schema/beans/spring-beans-3.0.xsd
        http://www.springframework.org/schema/aop http://www.springframework.org/
schema/aop/spring-aop-3.0.xsd
        http://www.springframework.org/schema/context
        http://www.springframework.org/schema/context/spring-context-3.0.xsd">

      <bean id="eventNotificationBus"
 class="org.apache.cloudstack.mom.rabbitmq.RabbitMQEventBus">
        <property name="name" value="eventNotificationBus"/>
        <property name="server" value="10.146.0.128"/>
        <property name="port" value="55672"/>
        <property name="username" value="$ {username}"/<
        <property name="password" value="password" value="${password}"/>
        <property name="exchange" value="cloudstack-events"/>
        </bean>

    <bean id="environmentVariablesConfiguration"
 class="org.jasypt.encryption.pbe.config.EnvironmentStringPBEConfig">
       <property name="algorithm" value="PBEWithSHA1AndDESede"/>
       <property name="passwordEnvName" value="APP_ENCRYPTION_PASSWORD"/>
    </bean>

    <bean id="configurationEncryptor"
 class="org.jasypt.encryption.pbe.StandardPBEStringEncryptor">
        <property name="config" ref="environmentVariablesConfiguration" />
    </bean>

    <bean id="propertyConfigurer"
 class="org.jasypt.spring3.properties.EncryptablePropertyPlaceholderConfigurer">
      <constructor-arg ref="configurationEncryptor" />
      <property name="location" value="classpath:/cred.properties" />
    </bean>
  </beans>
```

Create a new file in the same directory called **`cred.properties`** and the specify the values for user name and password as jascrypt encrypted strings.

An example, with guest as values for both fields:

```
username=nh2XrM7jWHMG4VQK18iiBQ==
password=nh2XrM7jWHMG4VQK18iiBQ==
```

The value specified for user name and password are the jasypt encrypted strings for the plaint text string, guest, for the user name and password.

4. Restart the Management Server.

19.2.3. Standard Events

The events log records three types of standard events.

- INFO. This event is generated when an operation has been successfully performed.
- WARN. This event is generated in the following circumstances.
  - When a network is disconnected while monitoring a template download.

  - When a template download is abandoned.
  - When an issue on the storage server causes the volumes to fail over to the mirror storage server.
- ERROR. This event is generated when an operation has not been successfully performed

19.2.4. Configuring AMQP-Based Event Bus

CloudPlatform supports event bus with AMQP-based middleware. In AMQP, middleware notion of routing keys and binding keys defines publish and subscribe semantics.

19.2.4.1. Routing Key

CloudPlatform integration with the AMQP middleware follows the below routing key convention for the routing message for each of the events that are published to the AMQP server.

```
Event-source.event-category.event-type.resource-type.resource-uuid
```

19.2.4.2. Event Source

Event source is the component or entity that publishes the message. In the current releases of CloudPlatform only Management Server publishes the events. However, CloudPlatform plug-ins are capable of publishing their own events.

All the events generated by the CloudPlatform Management Server come with **`management-server`** string value as the event source in the routing key of the message published on the AMQP servers.

19.2.4.3. Event Category

Event category in the routing key defines the category of the events specific to a publisher. For the Management Server, following are the possible event categories and corresponding values as seen in the routing key.

- **Action Events**: The category of events that are generated due to user/admin user actions. For example, user starting a VM or creating a network. The string value seen in the routing key for all the action events would be “Action Event”.
- **Alert Events**: The category of events that are generated by the Management Server as per polices set. For example, system running out of free public IPs in the public IP pool . The string value seen in the routing key for the alert events would be “AlertEvent”.
- **Usage Events**: The category of events generated by the Management Server for events related to consumption of a resource. For example, acquiring a public IP would generate usage event indicating that an IP is acquired. The string value seen in the routing key for usage events would be “UageEvent”.
- **Resource state event**: The category of events related to state changes to a resource either by user or system actions. For example, migrating a VM is a state change of the VM resource. The Management Server would generate a state change event for the VM migration operation. The string value seen in the routing key for resource state change events would be “ResourceStateEvent”.
- **Asyn job event**: The category of events related to job scheduling performed by CloudPlatform. There are events generated whenever an async job is submitted, scheduled, or completed. The string value seen in the routing key for async job events would be “AsyncJobevent".

19.2.4.4. Event Type

Event type in the routing key indicates the specific event type with in the category of the events. The following are the possible event types for both the "ActionEvent" and "UsageEvent" category of events.

Events		
"VM-CREATE"	"VOLUME-MIGRATE"	"PROJECT-UPDATE"
"VM-DESTROY"	"VOLUME-RESIZE"	"PROJECT-DELETE"
"VM-START"	"VOLUME-DETAIL-UPDATE"	"PROJECT-ACTIVATE"
"VM-STOP"	"VOLUME-DETAIL-ADD"	"PROJECT-SUSPEND"
"VM-REBOOT"	"VOLUME-DETAIL-REMOVE"	"PROJECT-ACCOUNT-ADD"
"VM-UPDATE"	"DOMAIN-CREATE"	"PROJECT-INVITATION-UPDATE"
"VM-UPGRADE"	"DOMAIN-DELETE"	"PROJECT-INVITATION-REMOVE"
"VM-DYNAMIC-SCALE"	"DOMAIN-UPDATE"	"PROJECT-ACCOUNT-REMOVE"
"VM-RESETPASSWORD"	"SNAPSHOT-CREATE"	"NETWORK-ELEMENT-CONFIGURE"
"VM-RESETSSHKEY"	"SNAPSHOT-DELETE"	"PHYSICAL-NETWORK-CREATE"
"VM-MIGRATE"	"SNAPSHOTPOLICY-CREATE"	"PHYSICAL-NETWORK-DELETE"
"VM-MOVE"	"SNAPSHOTPOLICY-UPDATE"	"PHYSICAL-NETWORK-UPDATE"
"VM-RESTORE"	"SNAPSHOTPOLICY-DELETE"	"SERVICE-PROVIDER-CREATE"
"ROUTER-CREATE"	"ISO-CREATE"	"SERVICE-PROVIDER-DELETE"
"ROUTER-DESTROY"	"ISO-DELETE"	"SERVICE-PROVIDER-UPDATE"
"ROUTER-START"	"ISO-COPY"	"TRAFFIC-TYPE-CREATE"
"ROUTER-STOP"	"ISO-ATTACH"	"TRAFFIC-TYPE-DELETE"
"ROUTER-REBOOT"	"ISO-DETACH"	"TRAFFIC-TYPE-UPDATE"
"ROUTER-HA"	"ISO-EXTRACT"	"PHYSICAL-LOADBALANCER-ADD"
"ROUTER-UPGRADE"	"ISO-UPLOAD"	"PHYSICAL-LOADBALANCER-DELETE"
"PROXY-CREATE"	"SSVM-CREATE"	
"PROXY-DESTROY"	"SSVM-DESTROY"	
"PROXY-START"	"SSVM-START"	
"PROXY-STOP"	"SSVM-STOP"	
"PROXY-REBOOT"	"SSVM-REBOOT"	
"PROXY-HA"	"SSVM-HA"	
"VNC-CONNECT"		

Events		
"VNC-DISCONNECT"	"SERVICE-OFFERING-CREATE"	"PHYSICAL-LOADBALANCER-CONFIGURE"
"NET-IPASSIGN"	"SERVICE-OFFERING-EDIT"	"SWITCH-MGMT-ADD"
"NET-IPRELEASE"	"SERVICE-OFFERING-DELETE"	"SWITCH-MGMT-DELETE"
"PORTABLE-IPASSIGN"	"DISK-OFFERING-CREATE"	"SWITCH-MGMT-CONFIGURE"
"PORTABLE-IPRELEASE"	"DISK-OFFERING-EDIT"	"SWITCH-MGMT-ENABLE"
"NET-RULEADD"	"DISK-OFFERING-DELETE"	"SWITCH-MGMT-DISABLE"
"NET-RULEDELETE"	"NETWORK-OFFERING-CREATE"	"PHYSICAL-FIREWALL-ADD"
“NET-RULEMODIFY"	"NETWORK-OFFERING-ASSIGN"	"PHYSICAL-FIREWALL-DELETE"
"NETWORK-CREATE"	"NETWORK-OFFERING-EDIT"	"PHYSICAL-FIREWALL-CONFIGURE"
"NETWORK-DELETE"	"NETWORK-OFFERING-REMOVE"	"VPC-CREATE"
"NETWORK-UPDATE"	"NETWORK-OFFERING-DELETE"	"VPC-UPDATE"
"FIREWALL-OPEN"	"POD-CREATE"	"VPC-DELETE"
"FIREWALL-CLOSE"	"POD-EDIT"	"VPC-RESTART"
"NIC-CREATE"	"POD-DELETE"	"NETWORK-ACL-CREATE"
"NIC-DELETE"	"ZONE-CREATE"	"NETWORK-ACL-DELETE"
"NIC-UPDATE"	"ZONE-EDIT"	"NETWORK-ACL-REPLACE"
"NIC-DETAIL-ADD"	"ZONE-DELETE"	"NETWORK-ACL-ITEM-CREATE"
"NIC-DETAIL-UPDATE"	"VLAN-IP-RANGE-CREATE"	"NETWORK-ACL-ITEM-UPDATE"
"NIC-DETAIL-REMOVE"	"VLAN-IP-RANGE-DELETE"	"NETWORK-ACL-ITEM-DELETE"
"LB-ASSIGN-TO-RULE"	"VLAN-IP-RANGE-DEDICATE"	"VPC-OFFERING-CREATE"
"LB-REMOVE-FROM-RULE"	"VLAN-IP-RANGE-RELEASE"	"VPC-OFFERING-UPDATE"
"LB-CREATE"	"STORAGE-IP-RANGE-CREATE"	"VPC-OFFERING-DELETE"
"LB-DELETE"	"STORAGE-IP-RANGE-DELETE"	"PRIVATE-GATEWAY-CREATE"
"LB-STICKINESSPOLICY-CREATE"	"STORAGE-IP-RANGE-UPDATE"	"PRIVATE-GATEWAY-DELETE"
"LB-STICKINESSPOLICY-DELETE"		"STATIC-ROUTE-CREATE"
"LB-HEALTHCHECKPOLICY-CREATE"		
"LB-HEALTHCHECKPOLICY-DELETE"		

Events		
"LB-UPDATE"	"CONFIGURATION-VALUE-EDIT"	"STATIC-ROUTE-DELETE"
"GLOBAL-LB-ASSIGN"	"SG-AUTH-INGRESS"	"CREATE_TAGS"
"GLOBAL-LB-REMOVE"	"SG-REVOKE-INGRESS"	"DELETE_TAGS"
"GLOBAL-LB-CREATE"	"SG-AUTH-EGRESS"	"CREATE_RESOURCE_DETAILS"
"GLOBAL-LB-DELETE"	"SG-REVOKE-EGRESS"	"DELETE_RESOURCE_DETAILS"
"GLOBAL-LB-UPDATE"	"SG-CREATE"	"VMSNAPSHOT-CREATE"
"ACCOUNT-ENABLE"	"SG-DELETE"	"VMSNAPSHOT-DELETE"
"ACCOUNT-DISABLE"	"SG-ASSIGN"	"VMSNAPSHOT-REVERTTO"
"ACCOUNT-CREATE"	"SG-REMOVE"	"PHYSICAL-NVPCONTROLLER-ADD"
"ACCOUNT-DELETE"	"HOST-RECONNECT"	"PHYSICAL-NVPCONTROLLER-DELETE"
"ACCOUNT-UPDATE"	"MAINT-CANCEL"	"PHYSICAL-NVPCONTROLLER-CONFIGURE"
"ACCOUNT-MARK-DEFAULT-ZONE"	"MAINT-CANCEL-PS"	"COUNTER-CREATE"
"USER-LOGIN"	"MAINT-PREPARE"	"COUNTER-DELETE"
"USER-LOGOUT"	"MAINT-PREPARE-PS"	"CONDITION-CREATE"
"USER-CREATE"	"VPN-REMOTE-ACCESS-CREATE"	"CONDITION-DELETE"
"USER-DELETE"	"VPN-REMOTE-ACCESS-DESTROY"	"AUTOSCALEPOLICY-CREATE"
"USER-DISABLE"	"VPN-USER-ADD"	"AUTOSCALEPOLICY-UPDATE"
"USER-UPDATE"	"VPN-USER-REMOVE"	"AUTOSCALEPOLICY-DELETE"
"USER-ENABLE"	"VPN-S2S-VPN-GATEWAY-CREATE"	"AUTOSCALEVMPROFILE-CREATE"
"USER-LOCK"	"VPN-S2S-VPN-GATEWAY-DELETE"	"AUTOSCALEVMPROFILE-DELETE"
"REGISTER-SSH-KEYPAIR"	"VPN-S2S-CUSTOMER-GATEWAY-CREATE"	"AUTOSCALEVMPROFILE-UPDATE"
"REGISTER-USER-KEY"	"VPN-S2S-CUSTOMER-GATEWAY-DELETE"	"AUTOSCALEVMGROUP-CREATE"
"TEMPLATE-CREATE"	"VPN-S2S-CUSTOMER-GATEWAY-UPDATE"	
"TEMPLATE-DELETE"	"VPN-S2S-CONNECTION-CREATE"	
"TEMPLATE-UPDATE"	"VPN-S2S-CONNECTION-DELETE"	
"TEMPLATE-DOWNLOAD-START"		
"TEMPLATE-DOWNLOAD-SUCCESS"		
"TEMPLATE-DOWNLOAD-FAILED"		

Events		
"TEMPLATE-COPY" "TEMPLATE-EXTRACT" "TEMPLATE-UPLOAD" "TEMPLATE-CLEANUP" "VOLUME-CREATE" "VOLUME-DELETE" "VOLUME-ATTACH" "VOLUME-DETACH" "VOLUME-EXTRACT" "VOLUME-UPLOAD"	"VPN-S2S-CONNECTION-RESET" "NETWORK-RESTART" "UPLOAD-CUSTOM-CERTIFICATE" "STATICNAT-ENABLE" "STATICNAT-DISABLE" "ZONE-VLAN-ASSIGN" "ZONE-VLAN-RELEASE" "PROJECT-CREATE"	"AUTOSCALEVMGROUP-DELETE" "AUTOSCALEVMGROUP-UPDATE" "AUTOSCALEVMGROUP-ENABLE" "AUTOSCALEVMGROUP-DISABLE" "PHYSICAL-DHCP-ADD" "PHYSICAL-DHCP-DELETE" "PHYSICAL-PXE-ADD" "PHYSICAL-PXE-DELETE" "AG-CREATE" "AG-DELETE" "AG-ASSIGN" "AG-REMOVE" "VM-AG-UPDATE" "INTERNALLBVM-START" "INTERNALLBVM-STOP" "HOST-RESERVATION-RELEASE" "GUESTVLANRANGE-DEDICATE" "GUESTVLANRANGE-RELEASE" "PORTABLE-IP-RANGE-CREATE" "PORTABLE-IP-RANGE-DELETE" "PORTABLE-IP-TRANSFER" "DEDICATE-RESOURCE" "DEDICATE-RESOURCE-RELEASE" "VM-RESERVATION-CLEANUP"

Events		
		"UCS-ASSOCIATEPROFILE" "UCS-TEMPLATEASSOCIATION" "UCS-DISASSOCIATEPROFILE" "UCS-REFRESHBLADES"

19.2.4.5. Resource Type

The resource type in the routing key indicates the resources for which event is being generated. The following are possible values in the routing key for resource type:

Events		
“VirtualMachine” “VirtualRouter” “ConsoleProxy” “VNC” “Network” “LoadBalancer” "Account" "User" "VirtualMachineTemplate"	"Volume" "Domain" "Snapshot" "Iso" "SecondaryStorageVmServiceOffering" "DiskOffering" "NetworkOffering" "Pod" "DataCenter"	"Vlan" "Host" "Project" "PhysicalNetwork" "Vpc" "RemoteAccessVpn" "StaticNat" "SecurityGroup"

19.2.4.6. Binding Key

CloudPlatform integration with AMQP middleware follows the binding key convention given below to selectively subscribe to the events published on AMQP server middleware.

```
Event--#source.event--#category.event--#type.resource--#type.resource--#uuid
```

CloudPlatform supports binding key to have a wild card for multiple values in the above tuple key.

Example Binding Keys

Here are a list of examples:

- *.*.*.*.* : This binding indicates a match against all the possible events, irrespective of event source, event category, event type, resource type, and resource UUID.
- *.*.*.VirtualMachine.* : This binding key indicates a match against all the possible events corresponding to all VMs irrespective of event source, event category, event type, and VM UUID.
- *.ActionEvents.*.*.* : This binding key indicates a match against all the possible Action Events irrespective of event source, event type, resource type, and resource UUID.

You can subscribe to additional fine grained events. For example, *.ActionEvents.VM-REBOOT.VirtualMachine. 172dd6f8-7d36-44ef-9a03- 3b342034aaa4. This binding key matches all the action events corresponding to VM reboot operation for the VM with the UUID,172dd6f8-7d36-44ef-9a03- 3b342034aaa4.

19.2.5. Long Running Job Events

The events log records three types of standard events.

- INFO. This event is generated when an operation has been successfully performed.
- WARN. This event is generated in the following circumstances.
  - When a network is disconnected while monitoring a template download.
  - When a template download is abandoned.
  - When an issue on the storage server causes the volumes to fail over to the mirror storage server.
- ERROR. This event is generated when an operation has not been successfully performed

19.2.6. Event Log Queries

Database logs can be queried from the user interface. The list of events captured by the system includes:

- Virtual machine creation, deletion, and on-going management operations
- Virtual router creation, deletion, and on-going management operations
- Template creation and deletion
- Network/load balancer rules creation and deletion
- Storage volume creation and deletion
- User login and logout

19.2.7. Deleting and Archiving Events and Alerts

CloudPlatform provides you the ability to delete or archive the existing alerts and events that you no longer want to implement. You can regularly delete or archive any alerts or events that you cannot, or do not want to resolve from the database.

You can delete or archive individual alerts or events either directly by using the Quickview or by using the Details page. If you want to delete multiple alerts or events at the same time, you can use the respective context menu. You can delete alerts or events by category for a time period. For example, you can select categories such as **USER.LOGOUT**, **VM.DESTROY**, **VM.AG.UPDATE**, **CONFIGURATION.VALUE.EDI**, and so on. You can also view the number of events or alerts archived or deleted.

In order to support the delete or archive alerts, the following global parameters have been added:

- **alert.purge.delay**: The alerts older than specified number of days are purged. Set the value to 0 to never purge alerts automatically.
- **alert.purge.interval**: The interval in seconds to wait before running the alert purge thread. The default is 86400 seconds (one day).

Note

Archived alerts or events cannot be viewed in the UI or by using the API. They are maintained in the database for auditing or compliance purposes.

19.2.7.1. Permissions

Consider the following:

- The root admin can delete or archive one or multiple alerts or events.
- The domain admin or end user can delete or archive one or multiple events.

19.2.7.2. Procedure

1. Log in as administrator to the CloudPlatform UI.
2. In the left navigation, click Events.
3. Perform either of the following:
   - To archive events, click Archive Events, and specify event type and date.
   - To archive events, click Delete Events, and specify event type and date.
4. Click OK.

19.3. Working with Server Logs

The CloudPlatform Management Server logs all web site, middle tier, and database activities for diagnostics purposes in /var/log/cloudstack/management/. The CloudPlatform logs a variety of error messages. We recommend this command to find the problematic output in the Management Server log:.

Note

When copying and pasting a command, be sure the command has pasted as a single line before executing. Some document viewers may introduce unwanted line breaks in copied text.

```
        grep -i -E 'exception|unable|fail|invalid|leak|warn|error' /var/log/cloudstack/
management/management-server.log
```

The CloudPlatform processes requests with a Job ID. If you find an error in the logs and you are interested in debugging the issue you can grep for this job ID in the management server log. For example, suppose that you find the following ERROR message:

```
        2010-10-04 13:49:32,595 ERROR [cloud.vm.UserVmManagerImpl] (Job-Executor-11:job-1076)
 Unable to find any host for [User|i-8-42-VM-untagged]
```

Note that the job ID is 1076. You can track back the events relating to job 1076 with the following grep:

```
        grep "job-1076)" management-server.log
```

The CloudPlatform Agent Server logs its activities in /var/log/cloudstack/agent/.

19.4. Getting the Exact CloudPlatform Version with cloudstack-sccs

CloudPlatform provides a command you can use to obtain the exact software version and patch that is currently installed and running, by identifying the most recently committed software code that is included in the install. This information can help with troubleshooting issues. Just log in as root on any compute node where CloudPlatform has been installed and run the following command:

```
# cloudstack-sccs
```

The output is a git commit hash. For example:

```
745811fc01bf93b6517e3ff1160bb8570034860b
```

19.5. Log Collection Utility cloud-bugtool

CloudPlatform provides a command-line utility called cloud-bugtool to make it easier to collect the logs and other diagnostic data required for troubleshooting. This is especially useful when interacting with Citrix Technical Support.

You can use cloud-bugtool to collect the following:

- Basic system and environment information and network configuration including IP addresses, routing, and name resolver settings
- Information about running processes
- Management Server logs
- System logs in /var/log/
- Dump of the cloud database

Warning

cloud-bugtool collects information which might be considered sensitive and confidential. Using the **--nodb** option to avoid the cloud database can reduce this concern, though it is not guaranteed to exclude all sensitive data.

19.5.1. Using cloud-bugtool

1. Log in as root on any compute node where CloudPlatform has been installed.
2. To gather all the possible troubleshooting data, run cloud-bugtool with no arguments:

```
# cloud-bugtool
```

The output is written to a .zip file. The location of this file is displayed on the console.

You can also use command-line options to specify which data will be collected:

- **`-f, --full`** : Collects all the data.
- **`-m, --minimal`** : Collects only system properties and the most recent log files.
- **`-d, --nodb`** : Does not collect the cloud database dump.

To display the current list of options on standard output, run cloud-bugtool with the command-line option **`-h`** or **`--help`**.

19.6. Data Loss on Exported Primary Storage

Symptom

Loss of existing data on primary storage which has been exposed as a Linux NFS server export on an iSCSI volume.

Cause

It is possible that a client from outside the intended pool has mounted the storage. When this occurs, the LVM is wiped and all data in the volume is lost

Solution

When setting up LUN exports, restrict the range of IP addresses that are allowed access by specifying a subnet mask. For example:

```
echo “/export 192.168.1.0/24(rw,async,no_root_squash)” > /etc/exports
```

Adjust the above command to suit your deployment needs.

More Information

See the export procedure in the "Secondary Storage" section of the CloudPlatform Installation Guide

19.7. Recovering a Lost Virtual Router

Symptom

A virtual router is running, but the host is disconnected. A virtual router no longer functions as expected.

Cause

The Virtual router is lost or down.

Solution

If you are sure that a virtual router is down forever, or no longer functions as expected, destroy it. You must create one afresh while keeping the backup router up and running (it is assumed this is in a redundant router setup):

- Force stop the router. Use the stopRouter API with forced=true parameter to do so.
- Before you continue with destroying this router, ensure that the backup router is running. Otherwise the network connection will be lost.
- Destroy the router by using the destroyRouter API.

Recreate the missing router by using the restartNetwork API with cleanup=false parameter. For more information about redundant router setup, see Creating a New Network Offering.

For more information about the API syntax, see the API Reference.

19.8. Maintenance mode not working on vCenter

Symptom

Host was placed in maintenance mode, but still appears live in vCenter.

Cause

The CloudPlatform administrator UI was used to place the host in scheduled maintenance mode. This mode is separate from vCenter's maintenance mode.

Solution

Use vCenter to place the host in maintenance mode.

More Information

See *Section 7.2, “Scheduled Maintenance and Maintenance Mode for Hosts ”*

19.9. Unable to deploy VMs from uploaded vSphere template

Symptom

When attempting to create a VM, the VM will not deploy.

Cause

If the template was created by uploading an OVA file that was created using vSphere Client, it is possible the OVA contained an ISO image. If it does, the deployment of VMs from the template will fail.

Solution

Remove the ISO and re-upload the template.

19.10. Unable to power on virtual machine on VMware

Symptom

Virtual machine does not power on. You might see errors like:

- Unable to open Swap File
- Unable to access a file since it is locked
- ¬ Unable to access Virtual machine configuration

Cause

A known issue on VMware machines. ESX hosts lock certain critical virtual machine files and file systems to prevent concurrent changes. Sometimes the files are not unlocked when the virtual machine is powered off. When a virtual machine attempts to power on, it can not access these critical files, and the virtual machine is unable to power on.

Solution

See the following:

VMware Knowledge Base Article[2]

19.11. Load balancer rules fail after changing network offering

Symptom

After changing the network offering on a network, load balancer rules stop working.

Cause

Load balancing rules were created while using a network service offering that includes an external load balancer device such as NetScaler, and later the network service offering changed to one that uses the CloudPlatform virtual router.

Solution

Create a firewall rule on the virtual router for each of your existing load balancing rules so that they continue to function.

[2] http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=10051/

Appendix A. Event Types

Following are the event types associated with CloudPlatform

Types	Events
VM	VM.CREATE VM.DESTROY VM.START VM.STOP VM.REBOOT VM.UPDATE VM.UPGRADE VM.DYNAMIC.SCALE VM.RESETPASSWORD VM.RESETSSHKEY VM.MIGRATE VM.MOVE VM.RESTORE VM.EXPUNGE
Domain Router	ROUTER.CREATE ROUTER.DESTROY ROUTER.START ROUTER.STOP ROUTER.REBOOT ROUTER.HA ROUTER.UPGRADE
Console proxy	PROXY.CREATE PROXY.DESTROY PROXY.START PROXY.STOP PROXY.REBOOT PROXY.HA
VNC Console Events	VNC.CONNECT

Types	Events
	VNC.DISCONNECT
Network Events	NET.IPASSIGN NET.IPRELEASE PORTABLE.IPASSIGN PORTABLE.IPRELEASE NET.RULEADD NET.RULEDELETE NET.RULEMODIFY NETWORK.CREATE NETWORK.DELETE NETWORK.UPDATE FIREWALL.OPEN FIREWALL.CLOSE
NIC Events	NIC.CREATE NIC.DELETE NIC.UPDATE NIC.DETAIL.ADD NIC.DETAIL.UPDATE NIC.DETAIL.REMOVE
Load Balancers	LB.ASSIGN.TO.RULE LB.REMOVE.FROM.RULE LB.CREATE LB.DELETE LB.STICKINESSPOLICY.CREATE LB.STICKINESSPOLICY.DELETE LB.HEALTHCHECKPOLICY.CREATE LB.HEALTHCHECKPOLICY.DELETE LB.UPDATE LB.STICKINESSPOLICY.UPDATE LB.HEALTHCHECKPOLICY.UPDATE

Types	Events
	LB.CERT.UPLOAD LB.CERT.DELETE LB.CERT.ASSIGN LB.CERT.REMOVE
Global Load Balancer rules	GLOBAL.LB.ASSIGN GLOBAL.LB.REMOVE GLOBAL.LB.CREATE GLOBAL.LB.DELETE GLOBAL.LB.UPDATE
Account events	ACCOUNT.ENABLE ACCOUNT.DISABLE ACCOUNT.CREATE ACCOUNT.DELETE ACCOUNT.UPDATE ACCOUNT.MARK.DEFAULT.ZONE
UserVO Events	USER.LOGIN USER.LOGOUT USER.CREATE USER.DELETE USER.DISABLE USER.UPDATE USER.ENABLE USER.LOCK
Registering SSH keypair events	REGISTER.SSH.KEYPAIR
Register for user API and secret keys	REGISTER.USER.KEY
Template Events	TEMPLATE.CREATE TEMPLATE.DELETE TEMPLATE.UPDATE TEMPLATE.DOWNLOAD.START TEMPLATE.DOWNLOAD.SUCCESS TEMPLATE.DOWNLOAD.FAILED

Types	Events
	TEMPLATE.COPY
	TEMPLATE.EXTRACT
	TEMPLATE.UPLOAD
	TEMPLATE.CLEANUP
Volume Events	VOLUME.CREATE
	VOLUME.DELETE
	VOLUME.ATTACH
	VOLUME.DETACH
	VOLUME.EXTRACT
	VOLUME.UPLOAD
	VOLUME.MIGRATE
	VOLUME.RESIZE
	VOLUME.DETAIL.UPDATE
	VOLUME.DETAIL.ADD
	VOLUME.DETAIL.REMOVE
	VOLUME.UPDATE
Domains	DOMAIN.CREATE
	DOMAIN.DELETE
	DOMAIN.UPDATE
Snapshots	SNAPSHOT.CREATE
	SNAPSHOT.DELETE
	SNAPSHOTPOLICY.CREATE
	SNAPSHOTPOLICY.UPDATE
	SNAPSHOTPOLICY.DELETE
	SNAPSHOT.REVERT
ISO	ISO.CREATE
	ISO.DELETE
	ISO.COPY
	ISO.ATTACH
	ISO.DETACH
	ISO.EXTRACT

Types	Events
	ISO.UPLOAD
SSVM	SSVM.CREATE SSVM.DESTROY SSVM.START SSVM.STOP SSVM.REBOOT SSVM.HA
Service Offerings	SERVICE.OFFERING.CREATE SERVICE.OFFERING.EDIT SERVICE.OFFERING.DELETE
Disk Offerings	DISK.OFFERING.CREATE DISK.OFFERING.EDIT DISK.OFFERING.DELETE
Network offerings	NETWORK.OFFERING.CREATE NETWORK.OFFERING.ASSIGN NETWORK.OFFERING.EDIT NETWORK.OFFERING.REMOVE NETWORK.OFFERING.DELETE
Pods	POD.CREATE POD.EDIT POD.DELETE
Zones	ZONE.CREATE ZONE.EDIT ZONE.DELETE
VLANs/IP ranges	VLAN.IP.RANGE.CREATE VLAN.IP.RANGE.DELETE VLAN.IP.RANGE.DEDICATE VLAN.IP.RANGE.RELEASE STORAGE.IP.RANGE.CREATE STORAGE.IP.RANGE.DELETE STORAGE.IP.RANGE.UPDATE

Types	Events
Configuration Table	CONFIGURATION.VALUE.EDIT
Security Groups	SG.AUTH.INGRESS SG.REVOKE.INGRESS SG.AUTH.EGRESS SG.REVOKE.EGRESS SG.CREATE SG.DELETE SG.ASSIGN SG.REMOVE
Host	HOST.RECONNECT
Maintenance	MAINT.CANCEL MAINT.CANCEL.PS MAINT.PREPARE MAINT.PREPARE.PS
VPN	VPN.REMOTE.ACCESS.CREATE VPN.REMOTE.ACCESS.DESTROY VPN.USER.ADD VPN.USER.REMOVE VPN.S2S.VPN.GATEWAY.CREATE VPN.S2S.VPN.GATEWAY.DELETE VPN.S2S.CUSTOMER.GATEWAY.CREATE VPN.S2S.CUSTOMER.GATEWAY.DELETE VPN.S2S.CUSTOMER.GATEWAY.UPDATE VPN.S2S.CONNECTION.CREATE VPN.S2S.CONNECTION.DELETE VPN.S2S.CONNECTION.RESET VPN.REMOTE.ACCESS.UPDATE VPN.S2S.VPN.GATEWAY.UPDATE VPN.S2S.CONNECTION.UPDATE
Network	NETWORK.RESTART NET.IPUPDATE

Types	Events
	FIREWALL.UPDATE FIREWALL.EGRESS.OPEN FIREWALL.EGRESS.CLOSE FIREWALL.EGRESS.UPDATE
Custom certificates	UPLOAD.CUSTOM.CERTIFICATE
OneToOnenat	STATICNAT.ENABLE STATICNAT.DISABLE ZONE.VLAN.ASSIGN ZONE.VLAN.RELEASE
Projects	PROJECT.CREATE PROJECT.UPDATE PROJECT.DELETE PROJECT.ACTIVATE PROJECT.SUSPEND PROJECT.ACCOUNT.ADD PROJECT.INVITATION.UPDATE PROJECT.INVITATION.REMOVE PROJECT.ACCOUNT.REMOVE
Network as a Service	NETWORK.ELEMENT.CONFIGURE
Physical Network Events	PHYSICAL.NETWORK.CREATE PHYSICAL.NETWORK.DELETE PHYSICAL.NETWORK.UPDATE
Physical Network Service Provider Events	SERVICE.PROVIDER.CREATE SERVICE.PROVIDER.DELETE SERVICE.PROVIDER.UPDATE
Physical Network Traffic Type Events	TRAFFIC.TYPE.CREATE TRAFFIC.TYPE.DELETE TRAFFIC.TYPE.UPDATE
External network device events	PHYSICAL.LOADBALANCER.ADD PHYSICAL.LOADBALANCER.DELETE PHYSICAL.LOADBALANCER.CONFIGURE
External switch management device events	SWITCH.MGMT.ADD

Types	Events
For example: Cisco Nexus 1000v Virtual Supervisor Module.	SWITCH.MGMT.DELETE SWITCH.MGMT.CONFIGURE SWITCH.MGMT.ENABLE SWITCH.MGMT.DISABLE PHYSICAL.FIREWALL.ADD PHYSICAL.FIREWALL.DELETE PHYSICAL.FIREWALL.CONFIGURE
VPC	VPC.CREATE VPC.UPDATE VPC.DELETE VPC.RESTART
Network ACL	NETWORK.ACL.CREATE NETWORK.ACL.DELETE NETWORK.ACL.REPLACE NETWORK.ACL.ITEM.CREATE NETWORK.ACL.ITEM.UPDATE NETWORK.ACL.ITEM.DELETE NETWORK.ACL.UPDATE
VPC offerings	VPC.OFFERING.CREATE VPC.OFFERING.UPDATE VPC.OFFERING.DELETE
Private gateway	PRIVATE.GATEWAY.CREATE PRIVATE.GATEWAY.DELETE
Static routes	STATIC.ROUTE.CREATE STATIC.ROUTE.DELETE
Tag-related events	CREATE_TAGS DELETE_TAGS
Meta data-related events	CREATE_RESOURCE_DETAILS DELETE_RESOURCE_DETAILS
VM snapshot events	VMSNAPSHOT.CREATE VMSNAPSHOT.DELETE

Types	Events
	VMSNAPSHOT.REVERTTO
External network device events	PHYSICAL.NVPCONTROLLER.ADD PHYSICAL.NVPCONTROLLER.DELETE PHYSICAL.NVPCONTROLLER.CONFIGURE PHYSICAL.OVSCONTROLLER.ADD PHYSICAL.OVSCONTROLLER.DELETE
AutoScale	COUNTER.CREATE COUNTER.DELETE CONDITION.CREATE CONDITION.DELETE AUTOSCALEPOLICY.CREATE AUTOSCALEPOLICY.UPDATE AUTOSCALEPOLICY.DELETE AUTOSCALEVMPROFILE.CREATE AUTOSCALEVMPROFILE.DELETE AUTOSCALEVMPROFILE.UPDATE AUTOSCALEVMGROUP.CREATE AUTOSCALEVMGROUP.DELETE AUTOSCALEVMGROUP.UPDATE AUTOSCALEVMGROUP.ENABLE AUTOSCALEVMGROUP.DISABLE PHYSICAL.DHCP.ADD PHYSICAL.DHCP.DELETE PHYSICAL.PXE.ADD PHYSICAL.PXE.DELETE AG.CREATE AG.DELETE AG.ASSIGN AG.REMOVE VM.AG.UPDATE

Types	Events
	INTERNALLBVM.START INTERNALLBVM.STOP HOST.RESERVATION.RELEASE BAREMETAL.RCT.ADD
Dedicated guest vlan range	GUESTVLANRANGE.DEDICATE GUESTVLANRANGE.RELEASE PORTABLE.IP.RANGE.CREATE PORTABLE.IP.RANGE.DELETE PORTABLE.IP.TRANSFER
Dedicated Resources	DEDICATE.RESOURCE DEDICATE.RESOURCE.RELEASE VM.RESERVATION.CLEANUP
Bare Metal Hosts	UCS.ASSOCIATEPROFILE
Object store migration	MIGRATE.PREPARE.SS
Alert generation	ALERT.GENERATE
OpenDaylight	PHYSICAL.ODLCONTROLLER.ADD PHYSICAL.ODLCONTROLLER.DELETE PHYSICAL.ODLCONTROLLER.CONFIGURE
Guest OS-related events	GUEST.OS.ADD GUEST.OS.REMOVE GUEST.OS.UPDATE GUEST.OS.MAPPING.ADD GUEST.OS.MAPPING.REMOVE GUEST.OS.MAPPING.UPDATE NIC.SECONDARY.IP.ASSIGN NIC.SECONDARY.IP.UNASSIGN NIC.SECONDARY.IP.CONFIGURE
External network mapping events	PHYSICAL.NUAGE.VSD.ADD PHYSICAL.NUAGE.VSD.DELETE

Appendix B. Alerts

The following is the list of alert type numbers. The current alerts can be found by calling listAlerts.

```
MEMORY = 0 // Available Memory below configured threshold
```

```
CPU = 1 // Unallocated CPU below configured threshold
```

```
STORAGE =2 // Available Storage below configured threshold
```

```
STORAGE_ALLOCATED = 3 // Remaining unallocated Storage is below configured threshold
```

```
PUBLIC_IP = 4 // Number of unallocated virtual network public IPs is below configured
 threshold
```

```
PRIVATE_IP = 5 // Number of unallocated private IPs is below configured threshold
```

```
SECONDARY_STORAGE = 6 //  Available Secondary Storage in availability zone is below
 configured threshold
```

```
HOST = 7 // Host related alerts like host disconnected
```

```
USERVM = 8 // User VM stopped unexpectedly
```

```
DOMAIN_ROUTER = 9 // Domain Router VM stopped unexpectedly
```

```
CONSOLE_PROXY = 10 // Console Proxy VM stopped unexpectedly
```

```
ROUTING = 11 // Lost connection to default route (to the gateway)
```

```
STORAGE_MISC = 12 // Storage issue in system VMs
```

```
USAGE_SERVER = 13 // No usage server process running
```

```
MANAGMENT_NODE = 14 // Management network CIDR is not configured originally
```

```
DOMAIN_ROUTER_MIGRATE = 15 // Domain Router VM Migration was unsuccessful
```

```
CONSOLE_PROXY_MIGRATE = 16 // Console Proxy VM Migration was unsuccessful
```

```
USERVM_MIGRATE = 17 // User VM Migration was unsuccessful
```

```
VLAN = 18 // Number of unallocated VLANs is below configured threshold in availability zone
```

```
SSVM = 19 // SSVM stopped unexpectedly
```

```
USAGE_SERVER_RESULT = 20 // Usage job failed
```

```
STORAGE_DELETE = 21 // Failed to delete storage pool
```

```
UPDATE_RESOURCE_COUNT = 22 // Failed to update the resource count
```

```
USAGE_SANITY_RESULT = 23 // Usage Sanity Check failed
```

```
DIRECT_ATTACHED_PUBLIC_IP = 24 // Number of unallocated shared network IPs is low in
 availability zone
```

```
LOCAL_STORAGE = 25 // Remaining unallocated Local Storage is below configured threshold
```

```
RESOURCE_LIMIT_EXCEEDED = 26 //Generated when the resource limit exceeds the limit. Currently
 used for recurring snapshots only
```

```
ALERT_TYPE_SYNC =  //Out-of-context, risky state transitions of VMs.
```

Appendix C. Time Zones

CloudPlatform API accepts the following time zone identifiers. There are several places that have a time zone as a required or optional parameter. These include scheduling recurring snapshots, creating a user, and specifying the usage time zone in the Configuration table.

Etc/GMT+12	Atlantic/Cape_Verde	Asia/Karachi
Etc/GMT+11	America/St_Johns	Asia/Kolkata
Pacific/Samoa	America/Araguaina	Asia/Bangkok
Pacific/Honolulu	America/Argentina/ Buenos_Aires	Asia/Shanghai
US/Alaska	America/Cayenne	Asia/Kuala_Lumpur
America/Los_Angeles	America/Godthab	Australia/Perth
Mexico/BajaNorte	America/Montevideo	Asia/Taipei
US/Arizona	Etc/GMT+2	Asia/Tokyo
US/Mountain	Atlantic/Azores	Asia/Seoul
America/Chihuahua	Africa/Casablanca	Australia/Adelaide
America/Chicago	Etc/UTC	Australia/Darwin
America/Costa_Rica	Atlantic/Reykjavik	Australia/Brisbane
America/Mexico_City	Europe/London	Australia/Canberra
Canada/Saskatchewan	CET	Pacific/Guam
America/Bogota	Europe/Bucharest	Pacific/Auckland
America/New_York	Africa/Johannesburg	
America/Caracas	Asia/Beirut	
America/Asuncion	Africa/Cairo	
America/Cuiaba	Asia/Jerusalem	
America/Halifax	Europe/Minsk	
America/La_Paz	Europe/Moscow	
America/Santiago	Africa/Nairobi	

Appendix D. Guest Operating Systems that CloudPlatform Supports

CloudPlatform supports the following guest operating systems:

For more information on the operating systems that the hypervisors support, refer to the *CloudPlatform (powered by Apache CloudStack) Version 4.5 Hypervisor Configuration Guide*.

Apple Mac OS X 10.6 (32-bit)	Oracle Enterprise Linux 6.0 (32-bit)	SUSE Linux Enterprise Server 10 SP3 (32-bit)
Apple Mac OS X 10.6 (64-bit)	Oracle Enterprise Linux 6.0 (64-bit)	SUSE Linux Enterprise Server 10 SP3 (64-bit)
Apple Mac OS X 10.7 (32-bit)	Oracle Enterprise Linux 6.1 (32-bit)	SUSE Linux Enterprise Server 10 SP4 (32-bit)
Apple Mac OS X 10.7 (64-bit)	Oracle Enterprise Linux 6.1 (64-bit)	SUSE Linux Enterprise Server 10 SP4 (64-bit)
Asianux 3(32-bit)	Oracle Enterprise Linux 6.2 (32-bit)	SUSE Linux Enterprise Server 11 (32-bit)
Asianux 3(64-bit)	Oracle Enterprise Linux 6.2 (64-bit)	SUSE Linux Enterprise Server 11 (64-bit)
CentOS 4.5 (32-bit)	Oracle Enterprise Linux 6.3 (32-bit)	SUSE Linux Enterprise Server 11 SP1 (32-bit)
CentOS 4.6 (32-bit)	Oracle Enterprise Linux 6.3 (64-bit)	SUSE Linux Enterprise Server 11 SP1 (64-bit)
CentOS 4.7 (32-bit)	Oracle Enterprise Linux 6.4 (32-bit)	SUSE Linux Enterprise Server 11 SP2 (32-bit)
CentOS 4.8 (32-bit)	Oracle Enterprise Linux 6.4 (64-bit)	SUSE Linux Enterprise Server 11 SP2 (64-bit)
CentOS 5.0 (32-bit)	Oracle Enterprise Linux 6.5 (32-bit)	SUSE Linux Enterprise Server 11 SP3 (32-bit)
CentOS 5.0 (64-bit)	Oracle Enterprise Linux 6.5 (64-bit)	SUSE Linux Enterprise Server 11 SP3 (64-bit)
CentOS 5.1 (32-bit)	OS/2	SUSE Linux Enterprise Server 12 (32-bit)
CentOS 5.1 (64-bit)	Other (32-bit)	SUSE Linux Enterprise Server 12 (64-bit)
CentOS 5.10 (32-bit)	Other (64-bit)	SUSE Linux Enterprise Server 9 SP4 (32-bit)
CentOS 5.10 (64-bit)	Other 2.6x Linux (32-bit)	Ubuntu 10.04 (32-bit)
CentOS 5.2 (32-bit)	Other 2.6x Linux (64-bit)	Ubuntu 10.04 (64-bit)
CentOS 5.2 (64-bit)	Other CentOS (32-bit)	
CentOS 5.3 (32-bit)		
CentOS 5.3 (64-bit)		
CentOS 5.4 (32-bit)		
CentOS 5.4 (64-bit)		
CentOS 5.5 (32-bit)		
CentOS 5.5 (64-bit)		
CentOS 5.6 (32-bit)		

CentOS 5.6 (64-bit)
CentOS 5.7 (32-bit)
CentOS 5.7 (64-bit)
CentOS 5.8 (32-bit)
CentOS 5.8 (64-bit)
CentOS 5.9 (32-bit)
CentOS 5.9 (64-bit)
CentOS 6.0 (32-bit)
CentOS 6.0 (64-bit)
CentOS 6.1 (32-bit)
CentOS 6.1 (64-bit)
CentOS 6.2 (32-bit)
CentOS 6.2 (64-bit)
CentOS 6.3 (32-bit)
CentOS 6.3 (64-bit)
CentOS 6.4 (32-bit)
CentOS 6.4 (64-bit)
CentOS 6.5 (32-bit)
CentOS 6.5 (64-bit)
Debian GNU/Linux 4(32-bit)
Debian GNU/Linux 4(64-bit)
Debian GNU/Linux 5(64-bit)
Debian GNU/Linux 5.0 (64-bit)
Debian GNU/Linux 6(32-bit)
Debian GNU/Linux 6(64-bit)
Debian GNU/Linux 7(32-bit)
Debian GNU/Linux 7(64-bit)
DOS
Fedora 10
Fedora 11
Other CentOS (64-bit)
Other Linux (32-bit)
Other Linux (64-bit)
Other PV (32-bit)
Other PV (64-bit)
Other SUSE Linux(32-bit)
Other SUSE Linux(64-bit)
Other Ubuntu (32-bit)
Other Ubuntu (64-bit)
Red Hat Enterprise Linux 2
Red Hat Enterprise Linux 3(32-bit)
Red Hat Enterprise Linux 3(64-bit)
Red Hat Enterprise Linux 4(64-bit)
Red Hat Enterprise Linux 4.5 (32-bit)
Red Hat Enterprise Linux 4.6 (32-bit)
Red Hat Enterprise Linux 4.7 (32-bit)
Red Hat Enterprise Linux 4.8 (32-bit)
Red Hat Enterprise Linux 5.0 (32-bit)
Red Hat Enterprise Linux 5.0 (64-bit)
Red Hat Enterprise Linux 5.1 (32-bit)
Red Hat Enterprise Linux 5.1 (64-bit)
Red Hat Enterprise Linux 5.10 (32-bit)
Red Hat Enterprise Linux 5.10 (64-bit)
Ubuntu 10.10 (32-bit)
Ubuntu 10.10 (64-bit)
Ubuntu 11.04 (32-bit)
Ubuntu 11.04 (64-bit)
Ubuntu 12.04 (32-bit)
Ubuntu 12.04 (64-bit)
Ubuntu 14.04 (32-bit)
Ubuntu 14.04 (64-bit)
Ubuntu 8.04 (32-bit)
Ubuntu 8.04 (64-bit)
Ubuntu 8.10 (32-bit)
Ubuntu 8.10 (64-bit)
Ubuntu 9.04 (32-bit)
Ubuntu 9.04 (64-bit)
Ubuntu 9.10 (32-bit)
Ubuntu 9.10 (64-bit)
Windows 2000 Advanced Server
Windows 2000 Professional
Windows 2000 Server
Windows 2000 Server SP4 (32-bit)
Windows 3.1
Windows 7 (32-bit)
Windows 7 (64-bit)
Windows 8 (32-bit)
Windows 8 (64-bit)
Windows 8.1 (32-bit)
Windows 8.1 (64-bit)
Windows 95
Windows 98

Fedora 12

Fedora 13

Fedora 8

Fedora 9

FreeBSD (32-bit)

FreeBSD (64-bit)

FreeBSD 10 (32-bit)

FreeBSD 10 (64-bit)

Microsoft Small Bussiness Server 2003

Novell Netware 5.1

Novell Netware 6.x

Open Enterprise Server

Oracle Enterprise Linux 5.0 (32-bit)

Oracle Enterprise Linux 5.0 (64-bit)

Oracle Enterprise Linux 5.1 (32-bit)

Oracle Enterprise Linux 5.1 (64-bit)

Oracle Enterprise Linux 5.10 (32-bit)

Oracle Enterprise Linux 5.10 (64-bit)

Oracle Enterprise Linux 5.2 (32-bit)

Oracle Enterprise Linux 5.2 (64-bit)

Oracle Enterprise Linux 5.3 (32-bit)

Oracle Enterprise Linux 5.3 (64-bit)

Oracle Enterprise Linux 5.4 (32-bit)

Red Hat Enterprise Linux 5.2 (32-bit)

Red Hat Enterprise Linux 5.2 (64-bit)

Red Hat Enterprise Linux 5.3 (32-bit)

Red Hat Enterprise Linux 5.3 (64-bit)

Red Hat Enterprise Linux 5.4 (32-bit)

Red Hat Enterprise Linux 5.4 (64-bit)

Red Hat Enterprise Linux 5.5 (32-bit)

Red Hat Enterprise Linux 5.5 (64-bit)

Red Hat Enterprise Linux 5.6 (32-bit)

Red Hat Enterprise Linux 5.6 (64-bit)

Red Hat Enterprise Linux 5.7 (32-bit)

Red Hat Enterprise Linux 5.7 (64-bit)

Red Hat Enterprise Linux 5.8 (32-bit)

Red Hat Enterprise Linux 5.8 (64-bit)

Red Hat Enterprise Linux 5.9 (32-bit)

Red Hat Enterprise Linux 5.9 (64-bit)

Red Hat Enterprise Linux 6.0 (32-bit)

Red Hat Enterprise Linux 6.0 (64-bit)

Red Hat Enterprise Linux 6.1 (32-bit)

Windows NT 4

Windows PV

Windows Server 2003 DataCenter Edition(32-bit)

Windows Server 2003 Enterprise Edition(32-bit)

Windows Server 2003 Enterprise Edition(64-bit)

Windows Server 2003 Standard Edition(32-bit)

Windows Server 2003 Standard Edition(64-bit)

Windows Server 2003 Web Edition

Windows Server 2008 (32-bit)

Windows Server 2008 (64-bit)

Windows Server 2008 R2 (64-bit)

Windows Server 2012 (64-bit)

Windows Server 2012 R2 (64-bit)

Windows Vista (32-bit)

Windows Vista (64-bit)

Windows XP (32-bit)

Windows XP (64-bit)

Windows XP SP2 (32-bit)

Windows XP SP3 (32-bit)

Oracle Enterprise Linux 5.4 (64-bit)	Red Hat Enterprise Linux 6.1 (64-bit)
Oracle Enterprise Linux 5.5 (32-bit)	Red Hat Enterprise Linux 6.2 (32-bit)
Oracle Enterprise Linux 5.5 (64-bit)	Red Hat Enterprise Linux 6.2 (64-bit)
Oracle Enterprise Linux 5.6 (32-bit)	Red Hat Enterprise Linux 6.3 (32-bit)
Oracle Enterprise Linux 5.6 (64-bit)	Red Hat Enterprise Linux 6.3 (64-bit)
Oracle Enterprise Linux 5.7 (32-bit)	Red Hat Enterprise Linux 6.4 (32-bit)
Oracle Enterprise Linux 5.7 (64-bit)	Red Hat Enterprise Linux 6.4 (64-bit)
Oracle Enterprise Linux 5.8 (32-bit)	Red Hat Enterprise Linux 6.5 (32-bit)
Oracle Enterprise Linux 5.8 (64-bit)	Red Hat Enterprise Linux 6.5 (64-bit)
Oracle Enterprise Linux 5.9 (32-bit)	SCO OpenServer 5
	SCO UnixWare 7
Oracle Enterprise Linux 5.9 (64-bit)	Sun Solaris 10(32-bit)
	Sun Solaris 10(64-bit)
	Sun Solaris 11 (32-bit)
	Sun Solaris 11 (64-bit)
	Sun Solaris 8(Experimental)
	Sun Solaris 9(Experimental)
	SUSE Linux Enterprise 10(32-bit)
	SUSE Linux Enterprise 10(64-bit)
	SUSE Linux Enterprise 8(32-bit)
	SUSE Linux Enterprise 8(64-bit)
	SUSE Linux Enterprise 9(32-bit)

	SUSE Linux Enterprise 9(64-bit) SUSE Linux Enterprise Server 10 SP1 (32-bit) SUSE Linux Enterprise Server 10 SP1 (64-bit) SUSE Linux Enterprise Server 10 SP2 (32-bit) SUSE Linux Enterprise Server 10 SP2 (64-bit)

Appendix E. Hypervisor Feature Support Matrix

CloudPlatform supports the following hypervisors:

- XenServer
- KVM
- VMWare ESXi
- Microsoft Hyper-V
- Linux Containers (LXC)
- Bare Metal

The tables in the following sections list the CloudPlatform features that these hypervisors support.

E.1. Compute

The following table displays the hypervisor support for the features associated with compute:

Feature	XenServer	KVM	VMWare ESXi	Microsoft Hyper-V	LXC	Bare Metal
Host-to-Host manual live migration	Yes	Yes (2.2.13)	Yes	Yes	No	NA
Multi-core per socket for VMs	Yes - Only for HVM	No	No	No	No	NA
Hyper-threading for VMs	Yes - Only for HVM	TBD	Yes	No	No	NA
Affinity and Anti-affinity Groups	Yes	Yes	Yes	No	Yes	Yes
VM reset on boot	Yes	Yes	Yes	No	Yes	No
Dynamic Resource Scaling	Yes	No	Yes	No	No	No
CPU and Memory Overprovision--ing	Yes	Yes	Yes	No	Yes	No

E.2. Storage

The following table displays the hypervisor support for the features associated with storage:

Feature	XenServer	KVM	VMWare ESXi	Microsoft Hyper-V	LXC	Bare Metal
iSCSI	Yes	Yes	Yes	No	No	NA
FC	Yes	Yes	Yes	No	No	NA
Local Disk	Yes	Yes	Yes	Yes	Yes (Only ROOT Disk)	Yes
NFS	Yes	Yes	Yes	No	Yes (Only ROOT Disk)	NA
Local Disk Snapshot	Yes	Yes	Yes	No	No	NA
Local Disk Data Disk	No	No	No	No	No	NA
Ceph	No	Yes	No	No	Yes (Only Data Disk)	No
Volume Snapshots	Yes - Delta	Yes	Yes - Full	No	No	No
Zone wide primary storage	Yes	Yes	Yes	No	Yes	No
VM Snapshots	Yes	No	Yes	No	No	No
Object Store	Yes	Yes	Yes	No	No	No
Resize Disk Volume	Yes - Online Grow/ No Shrink	Yes; CLVM: Online Grow/ Shrink; QCOW2: Online Grow	Yes - Online Grow	No	No	No
Storage Migration	TBD	NFS to NFS	NFS to iSCSI	No	No	No
Live Storage Migration	Yes	No	Yes	No	No	No
SMB/CIFS	No	No	No	Yes	No	No
PVSCSI	No	No	No	No	No	No

E.3. Networking

The following tables display the hypervisor support for the features associated with Networking:

Feature	XenServer	KVM	VMWare ESXi	Microsoft Hyper-V	LXC	Bare Metal
Network Throttling	Yes	No	Yes	No	No	No
VPC / nTier Apps 2.0	Yes	Yes	Yes	No	Yes	No
Open vSwitch	Yes	Yes	No	NA	No	No

CloudPlatform supports two types of networking – Basic and Advanced. They reflect the two zone types that you can create in CloudPlatform.

E.4. Basic Zone Networking

The following table displays the hypervisor support for Basic Zone Networking and associated features:

Feature	XenServer	KVM	VMWare ESXi	Microsoft Hyper-V	LXC	Bare Metal
Basic Zone	Yes	Yes	No	No	Yes	Yes
Basic Zone with Elastic IP and Elastic Load Balancing	Yes	Yes	No	No	Yes	Yes
Security Groups	Yes	Yes	No	No	Yes	No

E.5. Advanced Zone Networking

The following tables display the hypervisor support for Advanced Zone Networking and associated features:

Advanced Zone Networking with isolated networks:

Feature	XenServer	KVM	VMWare ESXi	Microsoft Hyper-V	LXC	Bare Metal
Advanced Zone	Yes	Yes	Yes	Yes	Yes	Yes
Security Group	No	Yes	No	Yes	Yes	No
Firewall Rules (by Virtual Router, Cisco ASA 1000 via VNCM or Juniper SRX)	Yes	Yes	Yes	Yes	Yes	Yes
Port Forwarding (by Virtual	Yes	Yes	Yes	Yes	Yes	Yes

Feature	XenServer	KVM	VMWare ESXi	Microsoft Hyper-V	LXC	Bare Metal
Router, Cisco ASA 1000 via VNCM or Juniper SRX)						
Load Balancing (by Virtual Router, Citrix NetScaler or F5 BigIP) - Stickiness rules	Yes	Yes	Yes	Yes	Yes	Yes
Load Balancing (by Virtual Router, Citrix NetScaler or F5 BigIP) - Health checks (only by Citrix NetScaler)	Yes	Yes	Yes	Yes	Yes	Yes
Load Balancing (by Virtual Router, Citrix NetScaler or F5 BigIP) - AutoScale (only by Citrix NetScaler)	Yes	Yes	Yes	Yes	Yes	Yes
VPN (client to site)	Yes	Yes	Yes	Yes	Yes	Yes
DNS	Yes	Yes	Yes	Yes	Yes	Yes
DHCP	Yes	Yes	Yes	Yes	Yes	Yes
UserData	Yes	Yes	Yes	Yes	Yes	Yes
Source NAT	Yes	Yes	Yes	Yes	Yes	Yes
Static NAT	Yes	Yes	Yes	Yes	Yes	Yes
Network Usage monitoring	Yes	Yes	Yes	Yes	Yes	Yes
Multiple IPs per NIC	Yes	Yes	Yes	No	Yes	No
Redundant Router (not on Hyper-V)	Yes	Yes	Yes	No	Yes	Yes

Advanced Zone Networking with shared networks:

Feature	XenServer	KVM	VMWare ESXi	Microsoft Hyper-V	LXC	Bare Metal
DHCP	Yes	Yes	Yes	Yes	Yes	No
DNS	Yes	Yes	Yes	Yes	Yes	No
UserData	Yes	Yes	Yes	Yes	Yes	No
External appliance Load Balancing	Yes	Yes	Yes	Yes	Yes	No
IPv6	No	Yes	No	No	Yes	No
PVLAN isolation	Yes	Yes	Yes–except Nexus 1000v	No	Yes	No

Advanced Zone Networking with Virtual Private Cloud (VPC):

Feature	XenServer	KVM	VMWare ESXi	Microsoft Hyper-V	LXC	Bare Metal
Network ACLs	Yes	Yes	Yes	No	Yes	No
Multiple Tiers	Yes	Yes	Yes	No	Yes	No
Site to Site VPN	Yes	Yes	Yes	No	Yes	No
LB Between Tiers	Yes	Yes	Yes	No	Yes	No
LB using Citrix NetScaler	Yes	Yes	Yes	No	Yes	No
Public IP to Tier LB	Yes	Yes	Yes	No	Yes	No
Source NAT	Yes	Yes	Yes	No	Yes	No
Static NAT	Yes	Yes	Yes	No	Yes	No
Port Forwarding	Yes	Yes	Yes	No	Yes	No
DNS	Yes	Yes	Yes	No	Yes	No
DHCP	Yes	Yes	Yes	No	Yes	No
UserData	Yes	Yes	Yes	No	Yes	No
Network Usage monitoring	Yes	Yes	Yes	No	Yes	No

E.6. Virtual Machine (VM) Operations

The following table displays the hypervisor support for various VM operations:

Feature	XenServer	KVM	VMWare ESXi	Microsoft Hyper-V	LXC	Bare Metal
VM Start, Stop, reboot	Yes	Yes	Yes	Yes	Yes	Yes
Live VM Memory and CPU resize	Yes	No	Yes	No	No	No
Stopped VM Memory and CPU resize	Yes	Yes	Yes	Yes	Yes	Yes
Reset VM	Yes	Yes	Yes	No	Yes	No
Add or remove Network (while VM is running)	Yes	Yes	Yes	No	Yes	No
Anti-Affinity	Yes	Yes	Yes	No	Yes	No
Deploy VM from ISO	Yes	Yes	Yes	Yes	No	Yes
HA	Yes - CS	Yes - CS	Yes - Native	No	Yes - CS	No
Work Load Balancing	No	No	DRS	No	No	No
Password-protected VNC Access	NA	Yes	Yes	NA	No	No

E.7. Features that All Hypervisors Support

The following features work with all hypervisors that CloudPlatform supports:

Feature	XenServer	KVM	VMWare ESXi	Microsoft Hyper-V	LXC	Bare Metal
LDAP/AD Integration for user authentication	Yes	Yes	Yes	Yes	Yes	Yes
Projects	Yes	Yes	Yes	Yes	Yes	Yes
Domains	Yes	Yes	Yes	Yes	Yes	Yes
Accounts	Yes	Yes	Yes	Yes	Yes	Yes
Domain quotas and Limits	Yes	Yes	Yes	Yes	Yes	Yes
Account quotas and limits	Yes	Yes	Yes	Yes	Yes	Yes

Feature	XenServer	KVM	VMWare ESXi	Microsoft Hyper-V	LXC	Bare Metal
Dedicated Resources (Zone, Pod, Cluster, Host)	Yes	Yes	Yes	Yes	Yes	Yes
Dedicated IP Range, Public VLANs (to account or domain)	Yes	Yes	Yes	Yes	Yes	Yes
Template and ISO Upload	Yes	Yes	Yes	Yes	Yes	Yes
Template and ISO Download	Yes	Yes	Yes	Yes	Yes	Yes
Citrix CloudPlatform Events and Alerts	Yes	Yes	Yes	Yes	Yes	Yes
Custom Compute Offerings	Yes	Yes	Yes	Yes	Yes	Yes
Host and Storage Maintenance mode	Yes	Yes	Yes	Yes	Yes	Yes

E.8. Hypervisor Support for External Devices

The following table displays the hypervisor support for various VM operations:

Feature	XenServer	KVM	VMWare ESXi	Microsoft Hyper-V	LXC	Bare Metal
NetScaler VPX	Yes	Yes	Yes	No	No	No

Appendix F. Deployment Behaviour

The following table shows the type of deployment is supported for hypervisors and VR configuration.

Version 4.5	XenServer	KVM	VMware	Hyper-V
VM deployment	Serial	Serial	Parallel with link-clone	Serial
VR deployment	Serial	Serial	Serial	Serial

Index