



CloudStack Administration Guide

For CloudStack Version 3.0.0 – 3.0.2

Revised August 16, 2012 4:41 PM

© 2011, 2012 Citrix Systems, Inc. All rights reserved. Specifications are subject to change without notice. Citrix Systems, Inc., the Citrix logo, Citrix XenServer, Citrix XenCenter, and CloudStack are trademarks or registered trademarks of Citrix Systems, Inc. All other brands or products are trademarks or registered trademarks of their respective holders.

Contents

What's In This Guide	12
What Is CloudStack?	12
What Can CloudStack Do?	13
Deployment Architecture Overview	14
Management Server Overview	14
Networking Overview	16
User Services Overview.....	16
Service Offerings, Disk Offerings, Network Offerings, and Templates	16
Accounts, Users, and Domains	17
Using an LDAP Server for User Authentication.....	17
Logging In to the CloudStack UI	21
End User's UI Overview.....	21
Root Administrator's UI Overview	21
Logging In as the Root Administrator	21
Changing the Root Password	22
Provisioning Cloud Infrastructure	23
About Zones.....	23
About Pods.....	25
About Clusters	25
About Physical Networks	26
Basic Zone Network Traffic Types.....	26
Basic Zone Guest IP Addresses	27
Advanced Zone Network Traffic Types	27
Advanced Zone Guest IP Addresses.....	28

Advanced Zone Public IP Addresses	28
System Reserved IP Addresses	28
Providing Services for Users	30
About Physical Networks	30
Configurable Characteristics of Physical Networks.....	30
About Virtual Networks	31
Isolated Networks.....	31
Shared Networks	31
Runtime Allocation of Virtual Network Resources	31
Network Service Providers.....	32
Supported Network Service Providers.....	32
Network Offerings	33
Creating a New Network Offering	34
Compute and Disk Service Offerings.....	35
Creating a New Compute Offering.....	36
Creating a New Disk Offering.....	37
Modifying or Deleting a Service Offering.....	37
System Service Offerings	38
Creating a New System Service Offering	38
Working With Virtual Machines.....	40
About Working With Virtual Machines.....	40
Best Practices for Virtual Machines	40
VM Lifecycle.....	41
Creating VMs	42
Accessing VMs	42
Stopping and Starting VMs	43

Changing the VM Name, OS, or Group	43
Changing the Service Offering for a VM	43
Moving VMs Between Hosts (Manual Live Migration)	44
Deleting VMs.....	45
Using Projects to Organize Users and Resources	46
Setting Up Invitations	46
Configuring Projects.....	47
Setting Resource Limits for Projects	47
Setting Project Creator Permissions	49
Creating a New Project	49
Adding Members to a Project.....	50
Sending Project Membership Invitations.....	50
Adding Project Members From the UI	51
Accepting a Membership Invitation.....	51
Removing a Member From a Project	51
Suspending or Deleting a Project.....	52
Using the Project View.....	52
Working with Hosts	54
About Hosts	54
Adding Hosts.....	54
Scheduled Maintenance and Maintenance Mode for Hosts	55
vCenter and Maintenance Mode.....	55
XenServer and Maintenance Mode	55
Disabling and Enabling Zones, Pods, and Clusters	56
Removing Hosts	57
Removing XenServer and KVM Hosts	57

Removing vSphere Hosts	57
Re-Installing Hosts	57
Maintaining Hypervisors on Hosts	57
Changing Host Password.....	58
Host Allocation.....	58
Over-Provisioning and Service Offering Limits	58
VLAN Provisioning.....	59
Managing Networks and Traffic.....	60
Guest Traffic.....	60
Networking in a Pod	61
Networking in a Zone.....	62
Basic Zone Physical Network Configuration	64
About Guest IP Addresses in a Basic Zone	64
Advanced Zone Physical Network Configuration	64
Configure Guest Traffic in an Advanced Zone	64
Configure Public Traffic in an Advanced Zone	65
Using Multiple Guest Networks	66
Adding an Additional Guest Network	66
Changing the Network Offering on a Guest Network	67
Security Groups.....	67
About Security Groups.....	67
Enabling Security Groups	68
Adding a Security Group	68
Adding Ingress and Egress Rules to a Security Group	68
External Firewalls and Load Balancers.....	70
About Using a NetScaler Load Balancer.....	70

Initial Setup of External Firewalls and Load Balancers	71
Ongoing Configuration of External Firewalls and Load Balancers	72
Load Balancer Rules	72
Adding a Load Balancer Rule	72
Sticky Session Policies for Load Balancer Rules	73
Guest IP Ranges	73
Acquiring a New IP Address	73
Releasing an IP Address	74
Static NAT	74
Enabling or Disabling Static NAT	74
IP Forwarding and Firewalling	75
Firewall Rules	75
Port Forwarding	76
IP Load Balancing	77
DNS and DHCP	77
VPN	77
Configuring VPN	78
Using VPN with Windows	78
Using VPN with Mac OS X	79
Working With Storage	80
Primary Storage	80
About Primary Storage	80
System Requirements for Primary Storage	80
Best Practices for Primary Storage	81
Runtime Behavior of Primary Storage	81
Hypervisor Support for Primary Storage	82

Storage Tags.....	83
Maintenance Mode for Primary Storage	83
Secondary Storage	83
About Secondary Storage	83
System Requirements for Secondary Storage	84
Best Practices for Secondary Storage	84
Secondary Storage VM.....	84
Changing the Secondary Storage IP Address	84
Changing Secondary Storage Servers	85
Using Swift for Secondary Storage.....	85
Working with Volumes	85
Creating a New Volume	86
Uploading an Existing Volume to a Virtual Machine	86
Attaching a Volume	87
Detaching and Moving Volumes	88
VM Storage Migration	88
Resizing Volumes	89
Volume Deletion and Garbage Collection.....	90
Working with ISOs.....	90
Adding an ISO.....	90
Attaching an ISO to a VM.....	92
Working with Templates.....	92
Creating Templates: Overview.....	93
Requirements for Templates	93
Best Practices for Templates	93
The Default Template	93

Private and Public Templates	94
Creating a Template from an Existing Virtual Machine	94
Creating a Template From a Snapshot.....	95
Uploading Templates	95
Exporting Templates	96
Creating a Windows Template.....	96
Importing AMIs	102
Creating an Ubuntu 10.04 LTS Template for XenServer	104
Converting a Hyper-V VM to a Template.....	106
Adding Password Management to Your Templates	107
Deleting Templates	108
Working with Snapshots	108
Automatic Snapshot Creation and Retention	109
Incremental Snapshots and Backup.....	109
Volume Status.....	109
Snapshot Restore	110
Runtime Considerations	110
Working with System Virtual Machines.....	111
The System VM Template.....	111
Multiple System VM Support for VMware	111
Console Proxy	111
Changing the Console Proxy SSL Certificate and Domain	112
Virtual Router	113
Configuring the Virtual Router.....	113
Upgrading a Virtual Router with System Service Offerings	114
Best Practices for Virtual Routers	114

Secondary Storage VM.....	114
System Reliability and High Availability	115
HA for Management Server	115
HA for Hosts	115
Primary Storage Outage and Data Loss	115
Secondary Storage Outage and Data Loss	115
HA-Enabled Virtual Machines	116
Managing the Cloud.....	117
Setting Global Configuration Parameters	117
Changing the Database Configuration	117
Administrator Alerts	117
Customizing the Network Domain Name	118
Stopping and Restarting the Management Server	118
Working with Usage.....	120
Configuring the Usage Server	120
Setting Usage Limits.....	122
Globally Configured Limits	122
Default Account Resource Limits	124
Per-Domain Limits	124
CloudStack API	125
Provisioning and Authentication API	125
Allocators	125
User Data and Meta Data	125
Tuning	127
Performance Monitoring	127
Increase Management Server Maximum Memory	127

Set Database Buffer Pool Size	127
Set and Monitor Total VM Limits per Host	128
Configure XenServer dom0 Memory	128
Troubleshooting.....	129
Event Logs	129
Standard Events	129
Long Running Job Events	129
Event Log Queries	130
Event Types.....	131
Alerts.....	132
Working with Server Logs	132
Data Loss on Exported Primary Storage	133
Recovering a Lost Virtual Router	133
Maintenance mode not working on vCenter	134
Unable to deploy VMs from uploaded vSphere template	134
Unable to power on virtual machine on VMware	135
Load balancer rules fail after changing network offering	135
Contacting Support	136
Appendix A—Time Zones.....	137

What's In This Guide

If you have already installed CloudStack or you want to learn more about the ongoing operation and maintenance of a CloudStack-powered cloud, this guide is for you. With the procedures in this Administration Guide, you can start using, configuring, and managing the ongoing operation of your cloud.

What Is CloudStack?

CloudStack™ is an open source software platform that pools computing resources to build public, private, and hybrid Infrastructure as a Service (IaaS) clouds. CloudStack manages the network, storage, and compute nodes that make up a cloud infrastructure. Use CloudStack to deploy, manage, and configure cloud computing environments.

Typical users are service providers and enterprises. With CloudStack, you can:

- Set up an on-demand, elastic cloud computing service. Service providers can sell self-service virtual machine instances, storage volumes, and networking configurations over the Internet.
- Set up an on-premise private cloud for use by employees. Rather than managing virtual machines in the same way as physical machines, with CloudStack an enterprise can offer self-service virtual machines to users without involving IT departments.

Who Should Read This

If you are new to CloudStack or you want to learn more about concepts before installing and running CloudStack, read this overview.

If you just want to get started, see the Basic Installation Guide.



What Can CloudStack Do?

Multiple Hypervisor Support

CloudStack works with a variety of hypervisors. A single cloud deployment can contain multiple hypervisor implementations. You have the complete freedom to choose the right hypervisor for your workload.

CloudStack is designed to work with open source Xen and KVM hypervisors as well as enterprise-grade hypervisors such as VMware vSphere, and Citrix XenServer.

Massively Scalable Infrastructure Management

CloudStack can manage tens of thousands of servers installed in multiple geographically distributed datacenters. The centralized management server scales linearly, eliminating the need for intermediate cluster-level management servers. No single component failure can cause cloud-wide outage. Periodic maintenance of the management server can be performed without affecting the functioning of virtual machines running in the cloud.

Automatic Configuration Management

CloudStack automatically configures each guest virtual machine's networking and storage settings.

CloudStack internally manages a pool of virtual appliances to support the cloud itself. These appliances offer services such as firewalling, routing, DHCP, VPN access, console proxy, storage access, and storage replication. The extensive use of virtual appliances greatly simplifies the installation, configuration, and on-going management of a cloud deployment.

Graphical User Interface

CloudStack offers an administrator's Web interface, used for provisioning and managing the cloud, as well as an end-user's Web interface, used for running VMs and managing VM templates. The UI can be customized to reflect the desired service provider or enterprise look and feel.

API and Extensibility

CloudStack provides an API that gives programmatic access to all the management features available in the UI. The API is maintained and documented. This API enables the creation of command line tools and new user interfaces to suit particular needs. See the Developer's Guide and API Reference, both available at http://docs.cloud.com/CloudStack_Documentation.

CloudStack's pluggable allocation architecture allows the creation of new types of allocators for the selection of storage and hosts. See the Allocator Implementation Guide (http://docs.cloud.com/CloudStack_Documentation/Allocator_Implementation_Guide).


High Availability

CloudStack has a number of features to increase the availability of the system. The Management Server itself may be deployed in a multi-node installation where the servers are load balanced. MySQL may be configured to use replication to provide for a manual failover in the event of database loss. For the Hosts, CloudStack supports NIC bonding and the use of separate networks for storage as well as iSCSI Multipath.

Deployment Architecture Overview

A CloudStack installation consists of two parts: the Management Server and the cloud infrastructure that it manages. When you set up and manage a CloudStack cloud, you provision resources such as hosts, storage devices, and IP addresses into the Management Server, and the Management Server manages those resources.

The minimum installation consists of one machine running the CloudStack Management Server and another machine to act as the cloud infrastructure (in this case, a very simple infrastructure consisting of one host running hypervisor software).



Simplified view of a basic deployment

A more full-featured installation consists of a highly-available multi-node Management Server installation and up to thousands of hosts using any of several advanced networking setups. For information about deployment options, see the Advanced Installation Guide.

Management Server Overview

The Management Server:


- Provides the web user interface for the administrator and a reference user interface for end users.
- Provides the APIs for CloudStack.
- Manages the assignment of guest VMs to particular hosts.
- Manages the assignment of public and private IP addresses to particular accounts.
- Manages the allocation of storage to guests as virtual disks.

- Manages snapshots, templates, and ISO images, possibly replicating them across data centers.
- Provides a single point of configuration for the cloud.

For additional options, including how to set up a multi-node management server installation, see the Advanced Installation Guide.

The cloud infrastructure is organized as follows:

- **Zone:** Typically, a zone is equivalent to a single datacenter. A zone consists of one or more pods and secondary storage. See About Zones on page 23.
- **Pod:** Typically, one rack of hardware that includes a layer-2 switch and one or more clusters. See About Pods on page 25.
- **Cluster:** A cluster consists of one or more hosts and primary storage. See About Clusters on page 25.
- **Host:** A single compute node within a cluster. The hosts are where the actual cloud services run in the form of guest virtual machines. See About Hosts on page 54.
- **Primary storage** is associated with a cluster, and it stores the disk volumes for all the VMs running on hosts in that cluster. See About Primary Storage on page 80.
- **Secondary storage** is associated with a zone, and it stores templates, ISO images, and disk volume snapshots. See About Secondary Storage on page 83.



Nested organization of a zone

Networking Overview

CloudStack offers two types of networking scenario:

- **Basic.** For AWS-style networking. Provides a single network where guest isolation can be provided through layer-3 means such as security groups (IP address source filtering).
- **Advanced.** For more sophisticated topologies. This network model provides the most flexibility in defining guest networks.

For more on networking, see:

- About Physical Networks on page 26
- Providing Services for Users on page 30
- Network Setup in the Advanced Installation Guide

User Services Overview

In addition to the physical and logical infrastructure of your cloud, and the CloudStack software and servers, you also need a layer of user services so that people can actually make use of the cloud. This means not just a user UI, but a set of options and resources that users can choose from, such as templates for creating virtual machines, disk storage, and more. If you are running a commercial service, you will be keeping track of what services and resources users are consuming and charging them for that usage. Even if you do not charge anything for people to use your cloud – say, if the users are strictly internal to your organization, or just friends who are sharing your cloud – you can still keep track of what services they use and how much of them.

Service Offerings, Disk Offerings, Network Offerings, and Templates

A user creating a new instance can make a variety of choices about its characteristics and capabilities. CloudStack provides several ways to present users with choices when creating a new instance:

- Service Offerings, defined by the CloudStack administrator, provide a choice of CPU speed, number of CPUs, RAM size, tags on the root disk, and other choices. See [Creating a New Compute Offering](#) on page 36.
- Disk Offerings, defined by the CloudStack administrator, provide a choice of disk size for primary data storage. See [Creating a New Disk Offering](#) on page 37.
- Network Offerings, defined by the CloudStack administrator, describe the feature set that is available to end users from the virtual router or external networking devices on a given guest network. See [Network Offerings](#) on page 33.

- Templates, defined by the CloudStack administrator or by any CloudStack user, are the base OS images that the user can choose from when creating a new instance. For example, CloudStack includes CentOS as a template. See [Working with Templates](#) on page 92.

In addition to these choices that are provided for users, there is another type of service offering which is available only to the CloudStack root administrator, and is used for configuring virtual infrastructure resources. For more information, see [Upgrading a Virtual Router with System Service Offerings](#) on page 114.

Accounts, Users, and Domains

An account typically represents a customer of the service provider or a department in a large organization. Multiple users can exist in an account. 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.

Accounts are grouped by domains. Domains usually contain 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.

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. Root administrators have complete access to the system, including managing templates, service offerings, customer care administrators, and domains. 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.

Using an LDAP Server for User Authentication

You can use an external LDAP server such as Microsoft Active Directory or ApacheDS to authenticate CloudStack end-users. Just map CloudStack accounts to the corresponding LDAP accounts using a query filter. The query filter is written using the query syntax of the particular LDAP server, and can include special wildcard characters provided by CloudStack for matching common values such as the user's email address and name. CloudStack will search the external LDAP directory tree starting at a specified base directory and return the distinguished name (DN) and password of the matching user. This information along with the given password is used to authenticate the user.

To set up LDAP authentication in CloudStack, call the CloudStack API command `ldapConfig` and provide the following:

- Hostname or IP address and listening port of the LDAP server
- Base directory and query filter
- Search user DN credentials, which give CloudStack permission to search on the LDAP server
- SSL keystore and password, if SSL is used

Example LDAP Configuration Commands

To understand the examples in this section, you need to know the basic concepts behind calling the CloudStack API, which are explained in the [Developer's Guide](#).

The following shows an example invocation of `ldapConfig` with an ApacheDS LDAP server.

```
http://127.0.0.1:8080/client/api?command=ldapConfig&hostname=127.0.0.1&searchbase=ou%3Dtesting%2Co%3Dproject&queryfilter=%28%26%28uid%3D%25u%29%29&binddn=cn%3DJohn+Singh%2Co%3Dtesting%2Co%3Dproject&bindpass=secret&port=10389&ssl=true&truststore=C%3A%2Fcompany%2Finfo%2Ftrusted.ks&truststorepass=secret&response=json&apiKey=YourAPIKey&signature=YourSignatureHash
```

The command must be URL-encoded. Here is the same example without the URL encoding:

```
http://127.0.0.1:8080/client/api?command=ldapConfig
&hostname=127.0.0.1
&searchbase=ou=testing,o=project
&queryfilter=((&(%uid=%u))
&binddn=cn=John+Singh,ou=testing,o=project
&bindpass=secret
&port=10389
&ssl=true
&truststore=C:/company/info/trusted.ks
&truststorepass=secret
&response=json
&apiKey=YourAPIKey&signature=YourSignatureHash
```

The following shows a similar command for Active Directory. Here, the search base is the testing group within a company, and the users are matched up based on email address.

```
http://10.147.29.101:8080/client/api?command=ldapConfig&hostname=10.147.28.250&searchbase=OU%3Dtesting%2CDC%3Dcompany&queryfilter=%28%26%28mail%3D%25e%29%29
&binddn=CN%3DAdministrator%2COU%3Dtesting%2CDC%3Dcompany&bindpass=1111_aaaa&port=389&response=json&apiKey=YourAPIKey&signature=YourSignatureHash
```

The next few sections explain some of the concepts you will need to know when filling out the `ldapConfig` parameters.

Search Base

An LDAP query is relative to a given node of the LDAP directory tree, called the search base. The search base is the distinguished name (DN) of a level of the directory tree below which all users can be found. The users can be in the immediate base directory or in some subdirectory. The search base may be equivalent to the organization, group, or domain name. The syntax for writing a DN varies depending on which LDAP server you are using. A full discussion of distinguished names is outside the scope of our documentation. The following table shows some examples of search bases to find users in the testing department.

LDAP Server	Example Search Base DN
ApacheDS	ou=testing,o=project
Active Directory	OU=testing, DC=company

Query Filter

The query filter is used to find a mapped user in the external LDAP server. The query filter should uniquely map the Cloudstack user to LDAP user for a meaningful authentication. For more information about query filter syntax, consult the documentation for your LDAP server.

The CloudStack query filter wildcards are:

Query Filter Wildcard	Description
%u	User name
%e	Email address
%n	First and last name

The following examples assume you are using Active Directory, and refer to user attributes from the Active Directory schema.

If the CloudStack user name is the same as the LDAP user ID:

```
(uid=%u)
```

If the CloudStack user name is the LDAP display name:

```
(displayName=%u)
```

To find a user by email address:

```
(mail=%e)
```

Search User Bind DN

The bind DN is the user on the external LDAP server permitted to search the LDAP directory within the defined search base. When the DN is returned, the DN and passed password are used to authenticate the Cloudstack user with an LDAP bind. A

full discussion of bind DN is outside the scope of our documentation. The following table shows some examples of bind DNs.

LDAP Server	Example Bind DN
ApacheDS	cn=Administrator,dc=testing,ou=project,ou=org
Active Directory	CN=Administrator, OU=testing, DC=company, DC=com

SSL Keystore Path and Password

If the LDAP server requires SSL, you need to enable it in the `ldapConfig` command by setting the parameters `ssl`, `truststore`, and `truststorepass`. Before enabling SSL for `ldapConfig`, you need to get the certificate which the LDAP server is using and add it to a trusted keystore. You will need to know the path to the keystore and the password.

Logging In to the CloudStack UI

CloudStack provides a web-based UI that can be used by both administrators and end users. The appropriate version of the UI is displayed depending on the credentials used to log in. The UI is available in popular browsers including IE7, IE8, IE9, Firefox 3.5+, Firefox 4, Safari 4, and Safari 5. The URL is: (substitute your own management server IP address)

```
http://<management-server-ip-address>:8080/client
```

For more guidance about the choices that appear when you log in to this UI, see *Logging In as the Root Administrator* on page 21.

End User's UI Overview

The CloudStack UI helps users of cloud infrastructure to view and use their cloud resources, including virtual machines, templates and ISOs, data volumes and snapshots, guest networks, and IP addresses. If the user is a member or administrator of one or more CloudStack projects, the UI can provide a project-oriented view.

Root Administrator's UI Overview

The CloudStack UI helps the CloudStack administrator provision, view, and manage the cloud infrastructure, domains, user accounts, projects, and configuration settings. The first time you start the UI after a fresh Management Server installation, you can choose to follow a guided tour to provision your cloud infrastructure. On subsequent logins, the dashboard of the logged-in user appears. The various links in this screen and the navigation bar on the left provide access to a variety of administrative functions. The root administrator can also use the UI to perform all the same tasks that are present in the end-user's UI.

Logging In as the Root Administrator

After the Management Server software is installed and running, you can run the CloudStack user interface. To log in as the root administrator:

1. Open your favorite Web browser and go to this URL. Substitute the IP address of your own Management Server:

```
http://<management-server-ip-address>:8080/client
```

On a fresh Management Server installation, a guided tour splash screen appears. On later visits, you'll see a login screen where you can enter a user ID and password and proceed to your Dashboard.

2. If you see the first-time splash screen, choose one of the following.
 - **Continue with basic setup.** Choose this if you're just trying CloudStack, and you want a guided walkthrough of the simplest possible configuration so that you can get started using CloudStack right away. We'll help you set

up a cloud with the following features: a single machine that runs CloudStack software and uses NFS to provide storage; a single machine running VMs under the XenServer hypervisor; and a shared public network.

The prompts in this guided tour should give you all the information you need, but if you want just a bit more detail, you can follow along in the CloudStack Basic Installation Guide.


- **I have used CloudStack before.** Choose this if you have already gone through a design phase and planned a more sophisticated CloudStack deployment, or you are ready to start scaling up a trial cloud that you set up earlier with the basic setup screens. In the Administrator UI, you can start using the more powerful features of CloudStack, such as advanced VLAN networking, high availability, additional network elements such as load balancers and firewalls, and support for multiple hypervisors including Citrix XenServer, KVM, and VMware vSphere.

The root administrator Dashboard appears.

3. If it was not already done during CloudStack installation, you should set a new root administrator password. If you chose basic setup, you'll be prompted to create a new password right away. If you chose experienced user, use the steps in Changing the Root Password on page 22.

Changing the Root Password

During CloudStack installation, you are logging in as the root administrator. This account manages the CloudStack deployment, including physical infrastructure. The root administrator can modify configuration settings to change basic functionality, create or delete user accounts, and take many actions that should be performed only by an authorized person. Please change the default password (which is "password") to a new, unique value.

1. Log in to the CloudStack UI using the current root user ID and password. The default is admin, password.
2. Click Accounts.
3. Click the admin account name.
4. Click View Users.
5. Click the admin user name.
6. Click the Change Password button. 
7. Type the new password, and click OK.


Provisioning Cloud Infrastructure

After the Management Server is installed and running, you can add the compute resources for it to manage.

To provision the cloud infrastructure, or to scale it up at any time, follow these procedures, all of which are available in the Advanced Installation Guide:

1. Add zones and pods
2. Configure the physical network
3. Add clusters
4. Add hosts
5. Add primary storage
6. Add secondary storage

When you have finished these steps, you will have a deployment with the following basic structure:



Conceptual view of a basic deployment

Your actual deployment can have multiple management servers and zones.


About Zones

A zone is the largest organizational unit within a CloudStack deployment. A zone typically corresponds to a single datacenter, although it is permissible to have multiple zones in a datacenter. The benefit of organizing infrastructure into

zones is to provide physical isolation and redundancy. For example, each zone can have its own power supply and network uplink, and the zones can be widely separated geographically (though this is not required).

A zone consists of:

- One or more pods. Each pod contains one or more clusters of hosts and one or more primary storage servers.
- Secondary storage, which is shared by all the pods in the zone.



A simple zone

Zones are visible to the end user. When a user starts a guest VM, the user must select a zone for their guest. Users might also be required to copy their private templates to additional zones to enable creation of guest VMs in those zones from their templates.

Zones may be public or private. Public zones are visible to all users. This means that any user may create a guest in that zone. Private zones are reserved for a specific domain. Only users in that domain or its subdomains may create guests in that zone.

Hosts in the same zone are directly accessible to each other without having to go through a firewall. Hosts in different zones can access each other through statically configured VPN tunnels.

For each zone, the administrator must decide the following.

- How many pods to place in a zone.
- How many clusters to have per pod.
- How many hosts to place in each cluster.


- How many primary storage servers to place in each cluster and total capacity for the storage servers.
- How much secondary storage to deploy in a zone.

About Pods

A pod often represents a single rack. Hosts in the same pod are in the same subnet.

A pod is the second-largest organizational unit within a CloudStack deployment. Pods are contained within zones. Each zone can contain one or more pods.

A pod consists of one or more clusters of hosts and one or more primary storage servers.



A simple pod


Pods are not visible to the end user.

About Clusters

A cluster provides a way to group hosts. To be precise, a cluster is a XenServer server pool, a set of KVM servers, or a VMware cluster preconfigured in vCenter. The hosts in a cluster all have identical hardware, run the same hypervisor, are on the same subnet, and access the same shared primary storage. Virtual machine instances (VMs) can be live-migrated from one host to another within the same cluster, without interrupting service to the user.

A cluster is the third-largest organizational unit within a CloudStack deployment. Clusters are contained within pods, and pods are contained within zones. Size of the cluster is limited by the underlying hypervisor, although the CloudStack recommends less in most cases; see the Best Practices section in the Installation Guide.

A cluster consists of one or more hosts and one or more primary storage servers.



A simple cluster

CloudStack allows multiple clusters in a cloud deployment.

Every VMware cluster is managed by a vCenter server. Administrator must register the vCenter server with CloudStack. There may be multiple vCenter servers per zone. Each vCenter server may manage multiple VMware clusters.

Even when local storage is used, clusters are still required. There is just one host per cluster.

About Physical Networks

Part of adding a zone is setting up the physical network. One or (in an advanced zone) more physical networks can be associated with each zone. The network corresponds to a NIC on the hypervisor host. Each physical network can carry one or more types of network traffic. The choices of traffic type for each network vary depending on whether you are creating a zone with basic networking or advanced networking.

Basic Zone Network Traffic Types

When basic networking is used, there can be only one physical network in the zone. That physical network carries three traffic types:

- **Guest.** When end users run VMs, they generate guest traffic. The guest VMs communicate with each other over a network that can be referred to as the guest network. Each pod in a basic zone is a broadcast domain, and therefore each pod has a different IP range for the guest network. The administrator must configure the IP range for each pod.
- **Management.** When CloudStack's internal resources communicate with each other, they generate management traffic. This includes communication between hosts, system VMs (VMs used by CloudStack to perform various tasks in the cloud), and any other component that communicates directly with the CloudStack Management Server. You must configure the IP range for the system VMs to use.
- **Storage.** Traffic such as VM templates and snapshots, which is sent between the secondary storage VM and secondary storage servers. CloudStack uses a separate Network Interface Controller (NIC) named storage NIC for

We strongly recommend the use of separate NICs for management traffic and guest traffic.

storage network traffic. Use of a storage NIC that always operates on a high bandwidth network allows fast template and snapshot copying. You must configure the IP range to use for the storage network.

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.

Basic Zone Guest IP Addresses

When basic networking is used, CloudStack will assign IP addresses in the CIDR of the pod to the guests in that pod. The administrator must add a Direct IP range on the pod for this purpose. These IPs are in the same VLAN as the hosts.

If the administrator changes the guest traffic CIDR at any time, the existing VMs continue to use the old CIDR. The new CIDR affects only VMs created from that point forward.

Advanced Zone Network Traffic Types

When advanced networking is used, there can be multiple physical networks in the zone. Each physical network can carry one or more traffic types, and you need to let CloudStack know which type of network traffic you want each network to carry. The traffic types in an advanced zone are:

- **Guest.** When end users run VMs, they generate guest traffic. The guest VMs communicate with each other over a network that can be referred to as the guest network. This network can be isolated or shared. In an isolated guest network, the administrator needs to reserve VLAN ranges to provide isolation for each CloudStack account's network (potentially a large number of VLANs). In a shared guest network, all guest VMs share a single network.
- **Management.** When CloudStack's internal resources communicate with each other, they generate management traffic. This includes communication between hosts, system VMs (VMs used by CloudStack to perform various tasks in the cloud), and any other component that communicates directly with the CloudStack Management Server. You must configure the IP range for the system VMs to use.
- **Public.** Public traffic is generated when VMs in the cloud access the Internet. Publicly accessible IPs must be allocated for this purpose. End users can use the CloudStack UI to acquire these IPs to implement NAT between their guest network and the public network, as described in [Acquiring a New IP Address](#) on page 73.
- **Storage.** Traffic such as VM templates and snapshots, which is sent between the secondary storage VM and secondary storage servers. CloudStack uses a separate Network Interface Controller (NIC) named storage NIC for storage network traffic. Use of a storage NIC that always operates on a high bandwidth network allows fast template and snapshot copying. You must configure the IP range to use for the storage network.

These traffic types can each be on a separate physical network, or they can be combined with certain restrictions. When you use the Add Zone wizard in the UI to create a new zone, you are guided into making only valid choices.

Advanced Zone Guest IP Addresses

When advanced networking is used, the administrator can create additional networks for use by the guests. These networks can span the zone and be available to all accounts, or they can be scoped to a single account, in which case only the named account may create guests that attach to these networks. The networks are defined by a VLAN ID, IP range, and gateway. The administrator may provision thousands of these networks if desired.

If the administrator changes the guest traffic CIDR at any time, the existing VMs continue to use the old CIDR. The new CIDR affects only guest networks and VMs created from that point forward.

Advanced Zone Public IP Addresses

CloudStack provisions one public IP address per account for use as the source NAT IP address. If a Juniper SRX firewall is used, CloudStack can instead use a single public IP address as an interface NAT IP for all accounts, reducing the number of IP addresses consumed. Users may request additional public IP addresses. The administrator must configure one or more ranges of public IP addresses for use by CloudStack. These IP addresses could be RFC1918 addresses in private clouds.

System Reserved IP Addresses

In each zone, you need to configure a range of reserved IP addresses for the management network. This network carries communication between the CloudStack Management Server and various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP.

The reserved IP addresses must be unique across the cloud. You cannot, for example, have a host in one zone which has the same private IP address as a host in another zone.

The hosts in a pod are assigned private IP addresses. These are typically RFC1918 addresses. The Console Proxy and Secondary Storage system VMs are also allocated private IP addresses in the CIDR of the pod that they are created in.

The management network IP addresses are in the same subnet as the compute nodes where hypervisors and the Management Server run. You therefore need to make sure computing servers and Management Servers use IP addresses outside of the System Reserved IP range. For example, suppose the System Reserved IP range starts at 192.168.154.2 and ends at 192.168.154.7. CloudStack can use .2 to .7 for System VMs. This leaves the rest of the pod CIDR, from .8 to .254, for the Management Server and hypervisor hosts.

In all zones

Provide private IPs for the system in each pod and provision them in CloudStack.

For KVM and XenServer, the recommended number of private IPs per pod is one per host. If you expect a pod to grow, add enough private IPs now to accommodate the growth.

In a zone that uses advanced networking

For vSphere with advanced networking, we recommend provisioning enough private IPs for your total number of customers, plus enough for the required CloudStack System VMs. Typically, about 10 additional IPs are required for the System VMs. For more information about System VMs, see *Working with System Virtual Machines* in the Administrator's Guide.

When advanced networking is being used, the number of private IP addresses available in each pod varies depending on which hypervisor is running on the nodes in that pod. Citrix XenServer and KVM use link-local addresses, which in theory provide more than 65,000 private IP addresses within the address block. As the pod grows over time, this should be more than enough for any reasonable number of hosts as well as IP addresses for guest virtual routers. VMWare ESXi, by contrast uses any administrator-specified subnetting scheme, and the typical administrator provides only 255 IPs per pod. Since these are shared by physical machines, the guest virtual router, and other entities, it is possible to run out of private IPs when scaling up a pod whose nodes are running ESXi.

To ensure adequate headroom to scale private IP space in an ESXi pod that uses advanced networking, use one or more of the following techniques:

- Specify a larger CIDR block for the subnet. A subnet mask with a /20 suffix will provide more than 4,000 IP addresses.
- Create multiple pods, each with its own subnet. For example, if you create 10 pods and each pod has 255 IPs, this will provide 2,550 IP addresses.

Providing Services for Users

People using cloud infrastructure have a variety of needs and preferences when it comes to the networking services provided by the cloud. As a CloudStack administrator, you can do the following things to set up networking for your users:

- Set up physical networks in zones (see the Advanced Installation Guide)
- Set up several different providers for the same service on a single physical network (for example, both Cisco and Juniper firewalls)
- Bundle different types of network services into network offerings, so users can choose the desired network services for any given virtual machine (see Network Offerings on page 33)
- Add new network offerings as time goes on so end users can upgrade to a better class of service on their network
- Provide more ways for a network to be accessed by a user, such as through a project of which the user is a member (see Using Projects to Organize Users and Resources on page 46)

About Physical Networks

A physical network is the actual network hardware and wiring in a zone. A zone can have multiple physical networks. An administrator can:

- Add/Remove/Update physical networks in a zone
- Configure VLANs on the physical network
- Configure a name so the network can be recognized by hypervisors
- Configure the service providers (firewalls, load balancers, etc.) available on a physical network
- Configure the IP addresses trunked to a physical network
- Specify what type of traffic is carried on the physical network, as well as other properties like network speed

Configurable Characteristics of Physical Networks

CloudStack provides configuration settings you can use to set up a physical network in a zone, including:

- What type of network traffic it carries (guest, public, management, storage)
- VLANs
- Unique name that the hypervisor can use to find that particular network
- Enabled or disabled. When a network is first set up, it is disabled – not in use yet. The administrator sets the physical network to enabled, and it begins to be used. The administrator can later disable the network again, which prevents

any new virtual networks from being created on that physical network; the existing network traffic continues even though the state is disabled.

- Speed
- Tags, so network offerings can be matched to physical networks
- Isolation method

About Virtual Networks

A virtual network is a logical construct that enables multi-tenancy on a single physical network. In CloudStack, a virtual network can be shared or isolated.

Isolated Networks

An isolated network can be accessed only by virtual machines of a single account. Isolated networks have the following properties.

- Resources such as VLAN are allocated and garbage collected dynamically.
- There is one network offering for the entire network.
- The network offering can be upgraded or downgraded but it is for the entire network.

Shared Networks

A shared network can be accessed by virtual machines that belong to many different accounts. Network Isolation on shared networks is accomplished using techniques such as security groups (supported only in basic zones in CloudStack 3.0.0 - 3.0.2).

- Shared Networks are created by the administrator
- Shared Networks can be designated to a certain domain
- Shared Network resources such as VLAN and physical network that it maps to are designated by the administrator
- Shared Networks are isolated by security groups
- Public Network is a shared network that is not shown to the end users.

Runtime Allocation of Virtual Network Resources

When you define a new virtual network, all your settings for that network are stored in CloudStack. The actual network resources are activated only when the first virtual machine starts in the network. When all virtual machines have left the virtual network, the network resources are garbage collected so they can be allocated again. This helps to conserve network resources.

Network Service Providers

A service provider (also called a network element) is hardware or virtual appliance that makes a network service possible; for example, a firewall appliance can be installed in the cloud to provide firewall service. On a single network, multiple providers can provide the same network service. For example, a firewall service may be provided by Cisco or Juniper devices in the same physical network.

For the most up-to-date list of which network service providers CloudStack supports, see the CloudStack UI or call `listNetworkServiceProviders`.

You can have multiple instances of the same service provider in a network (say, more than one Juniper SRX device).

If different providers are set up to provide the same service on the network, the administrator can create *network offerings* so users can specify which network service provider they prefer (along with the other choices offered in network offerings). Otherwise, CloudStack will choose which provider to use whenever the service is called for.

Supported Network Service Providers

CloudStack ships with an internal list of the supported service providers, and you can choose from this list when creating a network offering.

	Virtual Router	Citrix NetScaler	Juniper SRX	F5 BigIP	Host based (KVM/Xen)
Remote Access VPN	Yes	No	No	No	No
DNS/DHCP/User Data	Yes	No	No	No	No
Firewall	Yes	No	Yes	No	No
Load Balancing	Yes	Yes	No	Yes	No
Elastic IP	No	Yes	No	No	No
Elastic LB	No	Yes	No	No	No
Source NAT	Yes	No	Yes	No	No
Static NAT	Yes	Yes	Yes	No	No
Port Forwarding	Yes	No	Yes	No	No

Network Offerings

A network offering is a named set of network services, such as:

- DHCP
- DNS
- Source NAT
- Static NAT
- Port Forwarding
- Load Balancing
- Firewall
- VPN

For the most up-to-date list of which network services CloudStack supports, see the CloudStack UI or call `listNetworkServices`.

- (Optional) Name one of several available providers to use for a given service, such as Juniper for the firewall
- (Optional) Network tag to specify which physical network to use

When creating a new VM, the user chooses one of the available network offerings, and that determines which network services the VM can use.

The CloudStack administrator can create any number of custom network offerings, in addition to the default network offerings provided by CloudStack. By creating multiple custom network offerings, you can set up your cloud to offer different classes of service on a single multi-tenant physical network. For example, while the underlying physical wiring may be the same for two tenants, tenant A may only need simple firewall protection for their website, while tenant B may be running a web server farm and require a scalable firewall solution, load balancing solution, and alternate networks for accessing the database backend.

When creating a new virtual network, the CloudStack administrator chooses which network offering to enable for that network. Each virtual network is associated with one network offering. A virtual network can be upgraded or downgraded by changing its associated network offering. If you do this, be sure to reprogram the physical network to match.

CloudStack also has internal network offerings for use by CloudStack system VMs. These network offerings are not visible to users but can be modified by administrators.

WARNING

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 CloudStack 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.

Creating a New Network Offering

1. Log in with admin privileges to the CloudStack UI.
2. In the left navigation bar, click Service Offerings.
3. In Select Offering, choose Network Offering.
4. Click Add Network Offering.
5. In the dialog, make the following choices:
 - **Name.** Any desired name for the network offering.
 - **Description.** A short description of the offering that can be displayed to users.
 - **Network Rate.** Allowed data transfer rate in MB 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. For a description of these terms, see About Virtual Networks on page 31.

- **Specify VLAN.** (Isolated guest networks only) Indicate whether a VLAN should be specified when this offering is used.
- **Supported Services.** Select one or more of the possible network services. For some services, you must also choose the service provider; for example, if you select Load Balancer, you can choose the CloudStack virtual router or any other load balancers that have been configured in the cloud. Depending on which services you choose, additional fields may appear in the rest of the dialog box.
- **System Offering.** If the service provider for any of the services selected in Supported Services is a virtual router, the System Offering field appears. Choose the system service offering that you want virtual routers to use in this network. For example, if you selected Load Balancer in Supported Services and selected a virtual router to provide load balancing, the System Offering field appears so you can choose between the CloudStack default system service offering and any custom system service offerings that have been defined by the CloudStack root administrator. For more information, see System Service Offerings on page 38.
- **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.
- **Tags.** Network tag to specify which physical network to use.

6. Click Add.

Compute and Disk Service Offerings

A service offering is a set of virtual hardware features such as CPU core count and speed, memory, and disk size. The CloudStack administrator can set up various offerings, and then end users choose from the available offerings when they create a new VM.

A service offering includes the following elements:

- CPU, memory, and network resource guarantees.
- How resources are metered.
- How the resource usage is charged.
- How often the charges are generated.

For example, one service offering might allow users to create a virtual machine instance that is equivalent to a 1 GHz Intel® Core™ 2 CPU, with 1 GB memory at \$0.20/hour, with network traffic metered at \$0.10/GB. Based on the user's selected offering, CloudStack emits usage records that can be integrated with billing systems.

CloudStack separates service offerings into compute offerings and disk offerings. The computing service offering specifies:

- Guest CPU
- Guest RAM
- Guest Networking type (virtual or direct)
- Tags on the root disk

The disk offering specifies:

- Disk size (optional). An offering without a disk size will allow users to pick their own.
- Tags on the data disk

Creating a New Compute Offering

1. Log in with admin privileges to the CloudStack UI.
2. In the left navigation bar, click Service Offerings.
3. In Select Offering, choose Compute Offerings.
4. Click Add Compute Offering.
5. In the dialog, make the following choices:
 - **Name.** Any desired name for the service offering.
 - **Description.** A short description of the offering that can be displayed to users.
 - **Storage type.** The type of disk that should be allocated to the guest. Local allocates from storage attached to the hypervisor host directly. Shared allocates from storage accessible via NFS.
 - **# of CPU cores.** The number of cores which should be allocated to an instance with this offering.
 - **CPU (in MHz).** The CPU speed of the cores that the instance is allocated. For example, “2000” would provide for a 2 GHz clock.
 - **Memory (in MB).** The amount of memory in megabytes that the instance should be allocated. For example, “2048” would provide for a 2 GB RAM allocation.
 - **Network Rate.** Allowed data transfer rate in MB per second.
 - **Offer HA.** If yes, the user will be able to choose a VM to be monitored and as highly available as possible.
 - **Storage Tags.** The tags that should be associated with the primary storage for this disk.
 - **Host Tags.** (Optional) Any tags that you use to organize your hosts.

- **CPU cap.** Whether to cap users at their purchased level of CPU usage even if spare capacity is available.
- **Public.** Indicate whether the service offering should be available all domains or only some domains. Choose Yes to make it available to all domains. Choose No to limit the scope to a subdomain; CloudStack will then prompt for the subdomain's name.

6. Click Add.

Creating a New Disk Offering

1. Log in to the CloudStack UI as administrator.
2. Click Service Offerings.
3. In Select Offering, choose Disk Offerings.
4. Click Add Disk Offering.
5. Make the following choices, and click OK.
 - **Name and Description.** Give the offering a meaningful name and description that will be shown to users to help them select between various disk offerings.
 - **Custom Disk Size.** If checked, the user can set their own disk size. If not checked, the root administrator must define a value in Disk Size.
 - **Disk Size.** Appears only if Custom Disk Size is not selected. Define the volume size in GB.
 - **Storage Tags (optional).** The tags that should be associated with the primary storage for this disk. Tags are a comma separated list of attributes of the storage. For example "ssd,blue". Tags are also added on Primary Storage. CloudStack matches tags on a disk offering to tags on the storage. If a tag is present on a disk offering that tag (or tags) must also be present on Primary Storage for the volume to be provisioned. If no such primary storage exists, allocation from the disk offering will fail.
 - **Public.** Should the offering be available all domains or only some domains. Choose Yes to make it available to all domains. Choose No to limit the scope to a subdomain; CloudStack will then prompt for the subdomain's name.

Modifying or Deleting a Service Offering

Service offerings cannot be changed once created. This applies to both compute offerings and disk offerings.

A service offering can be deleted. If it is no longer in use, it is deleted immediately and permanently. If the service offering is still in use, it will remain in the database until all the virtual machines referencing it have been deleted. After deletion by the administrator, a service offering will not be available to end users that are creating new instances.

System Service Offerings

System service offerings provide a choice of CPU speed, number of CPUs, tags, and RAM size, just as other service offerings do. But rather than being used for virtual machine instances and exposed to users, system service offerings are used to change the default properties of virtual routers, console proxies, and other system VMs. System service offerings are visible only to the CloudStack root administrator. CloudStack provides default system service offerings. The CloudStack root administrator can create additional custom system service offerings.

When CloudStack creates a virtual router for a guest network, it uses default settings which are defined in the system service offering associated with the network offering. You can upgrade the capabilities of the virtual router by applying a new network offering that contains a different system service offering. All virtual routers in that network will begin using the settings from the new service offering.

Creating a New System Service Offering

1. Log in with admin privileges to the CloudStack UI.
2. In the left navigation bar, click Service Offerings.
3. In Select Offering, choose System Offerings.
4. Click Add System Service Offering.
5. In the dialog, make the following choices:
 - **Name.** Any desired name for the service offering.
 - **Description.** A short description of the offering that can be displayed to administrators.
 - **System VM Type.** Select the type of system virtual machine that this offering is intended to support.
 - **Storage type.** The type of disk that should be allocated. Local allocates from storage attached directly to the host where the system VM is running. Shared allocates from storage accessible via NFS.
 - **# of CPU cores.** The number of cores which should be allocated to a system VM with this offering.
 - **CPU (in MHz).** The CPU speed of the cores that the system VM is allocated. For example, “2000” would provide for a 2 GHz clock.
 - **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.** Allowed data transfer rate in MB per second.
 - **Offer HA.** If yes, the administrator can choose to have the system VM be monitored and as highly available as possible.
 - **Storage Tags.** The tags that should be associated with the primary storage used by the system VM.

- **Host Tags.** (Optional) Any tags that you use to organize your hosts.
- **CPU cap.** Whether to limit the level of CPU usage even if spare capacity is available.
- **Public.** Indicate whether the service offering should be available all domains or only some domains. Choose Yes to make it available to all domains. Choose No to limit the scope to a subdomain; CloudStack will then prompt for the subdomain's name.

6. Click Add.

Working With Virtual Machines

About Working With Virtual Machines

CloudStack provides administrators with complete control over the lifecycle of all guest VMs executing in the cloud.

CloudStack provides several guest management operations for end users and administrators. VMs may be stopped, started, rebooted, and destroyed.

Guest VMs have a name and group. VM names and groups are opaque to CloudStack and are available for end users to organize their VMs. Each VM can have three names for use in different contexts. Only two of these names can be controlled by the user:

- Instance name – a unique, immutable ID that is generated by CloudStack and can not be modified by the user. This name conforms to the requirements in IETF RFC 1123.
- Display name – the name displayed in the CloudStack web UI. Can be set by the user. Defaults to instance name.
- Name – host name that the DHCP server assigns to the VM. Can be set by the user. Defaults to instance name.

Guest VMs can be configured to be Highly Available (HA). An HA-enabled VM is monitored by the system. If the system detects that the VM is down, it will attempt to restart the VM, possibly on a different host. For more information, see HA-Enabled Virtual Machines on page 116.

Each new VM is allocated one public IP address. When the VM is started, CloudStack automatically creates a static NAT between this public IP address and the private IP address of the VM.

If elastic IP is in use (with the NetScaler load balancer), the IP address initially allocated to the new VM is not marked as elastic. The user must replace the automatically configured IP with a specifically acquired elastic IP, and set up the static NAT mapping between this new IP and the guest VM's private IP. The VM's original IP address is then released and returned to the pool of available public IPs.

CloudStack cannot distinguish a guest VM that was shut down by the user (such as with the “shutdown” command in Linux) from a VM that shut down unexpectedly. If an HA-enabled VM is shut down from inside the VM, CloudStack will restart it. To shut down an HA-enabled VM, you must go through the CloudStack UI or API.

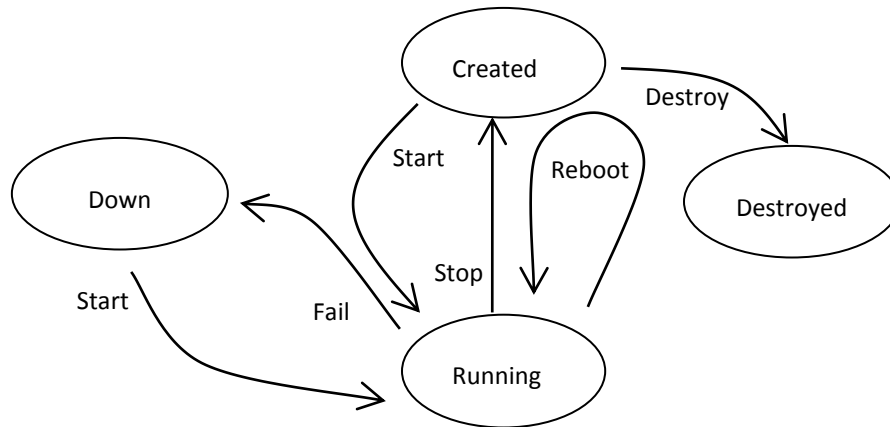
Best Practices for Virtual Machines

The CloudStack 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 CloudStack 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 $(N-1) * (\text{per-host-limit})$. Once a cluster reaches this number of VMs, use the CloudStack UI to disable allocation of more VMs to the cluster.

VM Lifecycle

Virtual machines can be in the following states:



Once a virtual machine is destroyed, it cannot be recovered. All the resources used by the virtual machine will be reclaimed by the system. This includes the virtual machine's IP address.

A stop will attempt to gracefully shut down the operating system, which typically involves terminating all the running applications. If the operation system cannot be stopped, it will be forcefully terminated. This has the same effect as pulling the power cord to a physical machine.

A reboot is a stop followed by a start.

CloudStack preserves the state of the virtual machine hard disk until the machine is destroyed.

A running virtual machine may fail because of hardware or network issues. A failed virtual machine is in the down state.

The system places the virtual machine into the down state if it does not receive the heartbeat from the hypervisor for three minutes.

The user can manually restart the virtual machine from the down state.

The system will start the virtual machine from the down state automatically if the virtual machine is marked as HA-enabled.

Creating VMs

Virtual machines 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.

To create a VM from a template:

1. Log in to the CloudStack UI as an administrator or user.
2. In the left navigation bar, click Instances.
3. Click Add Instance.
4. Select a template, then follow the steps in the wizard. (For more information about how the templates came to be in this list, see [Working with Templates](#) on page 92.)
5. Be sure that the hardware you have allows starting the selected service offering.
6. Click Submit and your VM will be created and started.

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 CloudStack management functions such as mounting additional volumes and ISO images, live migration, and graceful shutdown.


To create a VM from an ISO:

1. Log in to the CloudStack UI as an administrator or user.
2. In the left navigation bar, click Instances.
3. Click Add Instance.
4. Select ISO Boot, and follow the steps in the wizard.
5. Click Submit and your VM will be created and started.
6. (Oracle VM only) After ISO installation, the installer reboots into the operating system. Due to a known issue in OVM, the reboot will place the VM in the Stopped state. In the CloudStack UI, detach the ISO from the VM (so that the VM will not boot from the ISO again), then click the Start button to restart the VM.

Accessing VMs

Any user can access their own virtual machines. The administrator can access all VMs running in the cloud.

To access a VM through the CloudStack UI:

1. Log in to the CloudStack UI as a user or admin.
2. Click Instances, then click the name of a running VM.
3. Click the View Console button. 

To access a VM directly over the network:



1. The VM must have some port open 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 what security group you picked when creating the VM. In other cases, you can open a port by setting up a port forwarding policy. See [IP Forwarding and Firewalling](#) on page 75.
2. If a port is open but you can not access the VM using ssh, it's possible that ssh is not already enabled on the VM. This will depend on whether ssh is enabled in the template you picked when creating the VM. Access the VM through the CloudStack UI and enable ssh on the machine using the commands for the VM's operating system.
3. If the network has an external firewall device, you will need to create a firewall rule to allow access. See [IP Forwarding and Firewalling](#) on page 75.

Stopping and Starting VMs

Once a VM instance is created (see [Creating VMs](#) on page 42) you can stop, restart, or delete it as needed. In the CloudStack UI, click [Instances](#), select the VM, and use the [Stop](#), [Start](#), [Reboot](#), and [Destroy](#) links.

Changing the VM Name, OS, or Group

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

1. Log in to the CloudStack UI as an administrator or user.
2. In the left navigation, click [Instances](#).
3. Select the VM that you want to modify.
4. Click the Stop button to stop the VM. 
5. Click Edit. 
6. Make the desired changes to the following:
 - 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](#).

Changing the Service Offering for a VM

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 CloudStack UI as an administrator or user.
2. In the left navigation, click Instances.
3. Choose the VM that you want to work with.

4. Click the Stop Instance button. 

5. Click the Change Service button. 

The Change service dialog box is displayed.


6. Select the offering you want.
7. Click OK.

Moving VMs Between Hosts (Manual Live Migration)

The CloudStack administrator 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, and can be done under the following conditions:

- The root administrator is logged in. Domain admins and users can not perform manual live migration of VMs.
- The VM is running. Stopped VMs can not be live migrated.
- The destination host must be in the same cluster as the original host.
- The VM must not be using local disk storage.
- The destination host must have enough available capacity. If not, the VM will remain in the "migrating" state until memory becomes available.
- (OVM) If the VM is running on the OVM hypervisor, it must not have an ISO attached. Live migration of a VM with attached ISO is not supported in OVM.

To manually live migrate a virtual machine:


1. Log in to the CloudStack UI as the CloudStack root administrator.
2. Go to Instances.
3. Choose the VM that you want to migrate.
4. Click the Migrate Instance button. 
5. From the list of hosts, choose the one to which you want to move the VM.
6. Click OK.

Deleting VMs

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

Administrators can delete any virtual machines.

To delete a virtual machine:

1. Log in to the CloudStack UI as an administrator or user.
2. Go to Instances.
3. Choose the VM that you want to delete.
4. Click the Destroy Instance button. 

Using Projects to Organize Users and Resources

CloudStack 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. CloudStack 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 CloudStack to allow any user to create a new project, or you can restrict that ability to just CloudStack administrators. Once you have created a project, you become that project's administrator, and you can add others within your domain to the project. CloudStack 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 CloudStack 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 CloudStack 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 CloudStack 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.

Setting Up Invitations

CloudStack 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 CloudStack account. If you want administrators to use invitations to add members to projects, turn on and set up the invitations feature in CloudStack.

1. Log in as administrator to the CloudStack UI.

2. In the left navigation, click Global Settings.
3. In the search box, type project and click the search button.
4. Set project.invite.required to true. This enables the invitation feature.
5. In the search results, you will 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.



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.

6. Restart the Management Server.

```
# service cloud-management restart
```

Configuring Projects

Before CloudStack users start using projects, the CloudStack administrator must set up various systems to support them, including membership invitations, limits on project resources, and controls on who can create projects.

Setting Resource Limits for Projects

The CloudStack 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 CloudStack root administrator. The root administrator can also set lower resource limits for any project in the cloud.

Setting the Global Project Resource Limits

1. Log in as administrator to the CloudStack UI.
2. In the left navigation, click Global Settings.
3. In the search box, type max.project and click the search button.
4. In the search results, you will see the parameters you can use to set per-project maximum resource amounts that apply to all projects in the cloud. No project can have more resources, but an individual project can have lower limits. Click the edit button to set each parameter.



max.project.public.ips	Maximum number of public IP addresses that can be owned by any project in the cloud. See About Public IP Addresses on page 65.
max.project.snapshots	Maximum number of snapshots that can be owned by any project in the cloud. See Working with Snapshots on page 108.
max.project.templates	Maximum number of templates that can be owned by any project in the cloud. See Working with Templates on page 92.
max.project.uservms	Maximum number of guest virtual machines that can be owned by any project in the cloud. See Working With Virtual Machines on page 40.
max.project.volumes	Maximum number of data volumes that can be owned by any project in the cloud. See Working with Volumes on page 85.

5. Restart the Management Server.

```
# service cloud-management restart
```

Setting Per-Project Resource Limits

The CloudStack root administrator or the domain administrator of the domain where the project resides can set new resource limits for an individual project. The project owner can set resource limits only if the owner is also a domain or root administrator.

The new limits must be below the global default limits set by the CloudStack administrator (as described in Setting Resource Limits for Projects on page 47). If the project already owns more of a given type of resource than the new maximum, the resources are not affected; however, the project can not add any new resources of that type until the total drops below the new limit.

1. Log in to the CloudStack 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 Resources tab. This tab lists the current maximum amount that the project is allowed to own for each type of resource.
6. Type new values for one or more resources.
7. Click Apply.

Setting Project Creator Permissions

You can configure CloudStack to allow any user to create a new project, or you can restrict that ability to just CloudStack administrators.

1. Log in as administrator to the CloudStack 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.



<code>allow.user.create.projects</code>	Set to true to allow end users to create projects. Set to false if you want only the CloudStack root administrator and domain administrators to create projects.
---	--

5. Restart the Management Server.

```
# service cloud-management restart
```

Creating a New Project

CloudStack 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 to the CloudStack 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.

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 CloudStack root administrator. There are two ways to add members in CloudStack, 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.

Sending Project Membership Invitations

Use these steps to add a new member to a project if the invitations feature is enabled in the cloud (as described in Setting Up Invitations on page 46). If the invitations feature is not turned on, use the procedure in Adding Project Members From the UI on page 51 instead.

1. Log in to the CloudStack 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:
 - Account – The invitation will appear in the user's Invitations tab in the Project View. See Using the Project View on page 52.
 - 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 CloudStack when accepting the invitation. Email invitations will work only if the global parameters related to the SMTP server have been set. See Setting Up Invitations on page 46.
7. Type the user name or email address of the new member you want to add, and click Invite. Type the CloudStack 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.

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 (as described in Setting Up Invitations on page 46), use the procedure in Sending Project Membership Invitations on page 50 instead.

1. Log in to the CloudStack UI.
2. In the left navigation, click Projects.
3. In Select view, click 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.

Accepting a Membership Invitation

If you have received an invitation to join a CloudStack project, and you want to accept the invitation, follow these steps.

1. Log in to the CloudStack 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 CloudStack 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.

Removing a Member From a Project

When a member is removed from a project, the member's resources continue to be owned by the project. The former project member can not create any new resources within the project or use any of the project's existing resources.

A member of a project can be removed by the project administrator, the domain administrator of the domain the project belongs to or of its parent domain, or the CloudStack root administrator.

1. Log in to the CloudStack UI.
2. In the left navigation, click Projects.
3. In Select view, click Projects.

4. Click the name of the project.
5. Click the Accounts tab.
6. Click the name of the member.
7. Click the Delete button. 

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 CloudStack root administrator.

1. Log in to the CloudStack UI.
2. In the left navigation, click Projects.
3. In Select view, click Projects.
4. Click the name of the project.
5. Click one of the buttons:



Suspend project



Delete project

Using the Project View

If you are a member of a project, you can use CloudStack'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 CloudStack UI as an administrator or user.
2. Click Project View.
3. In Select Project, choose the name of the project you want to view. This dialog lists all projects of which you are a member.

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.
- Click the Resources tab to see the current resource limits for the project. If you are the project administrator, you can lower the resource limits.
- In Networking and Security, click Manage Resources to acquire IP addresses, set up load balancing and port forwarding rules.
- Click Default View to close the project view.

Working with Hosts

This section gives concepts and technical details about hosts in CloudStack. For information about how to install and configure hosts through the CloudStack UI, see the Advanced Installation Guide.

About Hosts

A host is a single computer. Hosts provide the computing resources that run the guest virtual machines. Each host has hypervisor software installed on it to manage the guest VMs. For example, a Linux KVM-enabled server, a Citrix XenServer server, and an ESXi server are hosts.

The host is the smallest organizational unit within a CloudStack deployment. Hosts are contained within clusters, clusters are contained within pods, and pods are contained within zones.

Hosts in a CloudStack deployment:

- Provide the CPU, memory, storage, and networking resources needed to host the virtual machines
- Interconnect using a high bandwidth TCP/IP network and connect to the Internet
- May reside in multiple data centers across different geographic locations
- May have different capacities (different CPU speeds, different amounts of RAM, etc.), although the hosts within a cluster must all be homogeneous

CloudStack automatically detects the amount of CPU and memory resources provided by the hosts.

Hosts are not visible to the end user. An end user cannot determine which host their guest has been assigned to.

For a host to function in CloudStack, you must do the following:

- Install hypervisor software on the host
- Assign an IP address to the host
- Ensure the host is connected to the CloudStack Management Server

Adding Hosts

Additional hosts can be added at any time to provide more capacity for guest VMs. For requirements and instructions, see the Advanced Installation Guide.

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.

vCenter and Maintenance Mode

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

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

When the CloudStack 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 CloudStack to reactivate it.

- b. Then use CloudStack's administrator UI to cancel the CloudStack 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.

XenServer and Maintenance Mode

For XenServer, you can take a server offline temporarily by using the Maintenance Mode feature in XenCenter. When you place a server into Maintenance Mode, all running VMs are automatically migrated from it to another host in the same pool. If the server is the pool master, a new master will also be selected for the pool. While a server is Maintenance Mode, you cannot create or start any VMs on it.

To place a server in Maintenance Mode

1. In the **Resources** pane, select the server, then do one of the following:
 - Right-click, then click **Enter Maintenance Mode** on the shortcut menu.
 - On the Server menu, click **Enter Maintenance Mode**.
2. Click **Enter Maintenance Mode**.

The server's status in the Resources pane shows when all running VMs have been successfully migrated off the server.



To take a server out of Maintenance Mode

1. In the **Resources** pane, select the server, then do one of the following:
 - Right-click, then click **Exit Maintenance Mode** on the shortcut menu.
 - On the **Server** menu, click **Exit Maintenance Mode**.
2. Click **Exit Maintenance Mode**.

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.

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

1. Log in to the CloudStack 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. 

Removing Hosts

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

Removing XenServer and 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 Scheduled Maintenance and Maintenance Mode for Hosts on page 55).
2. For KVM, 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.

Removing vSphere Hosts

To remove this type of host, first place it in maintenance mode, as described in Scheduled Maintenance and Maintenance Mode for Hosts on page 55. Then use CloudStack to remove the host. CloudStack will not direct commands to a host that has been removed using CloudStack. However, the host may still exist in the vCenter cluster.

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. See Removing Hosts on page 54.

Maintaining Hypervisors on Hosts

When running hypervisor software on hosts, be sure all the hotfixes provided by the hypervisor vendor are applied. Track the release of hypervisor patches through your hypervisor vendor's support channel, and apply patches as soon as possible after they are released.

CloudStack will not track or notify you of required hypervisor patches.

It is essential that your hosts are completely up to date with the provided hypervisor patches. The hypervisor vendor is likely to refuse to support any system that is not up to date with patches.

WARNING

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

(XenServer) For more information, see [Highly Recommended Hotfixes for XenServer](#) in the CloudStack Knowledge Base.

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 CloudStack will not match. Operations on the cluster will fail until the two passwords match.
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 hostname "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 set password='password' where id=5 or id=10 or id=12;
```

Host Allocation

The system automatically picks the most appropriate host to run each virtual machine. End users may specify the zone in which the virtual machine will be created. End users do not have control over which host will run the virtual machine instance.

CloudStack administrators can specify that certain hosts should have a preference for particular types of guest instances. For example, an administrator could state that a host should have a preference to run Windows guests. The default host allocator will attempt to place guests of that OS type on such hosts first. If no such host is available, the allocator will place the instance wherever there is sufficient physical capacity.

Over-Provisioning and Service Offering Limits

CloudStack performs CPU over-provisioning based on an over-provisioning ratio configured by the administrator. This is defined by the `cpu.overprovisioning.factor` global configuration variable.

CPU over-provisioning allows the sum total of the gigahertz of CPU speed allocated to guests to exceed the physically available gigahertz. For example, if a Host had 2 cores at 2 GHz each, it would have 4 GHz total. With a CPU over provisioning factor of 1.5, the CloudStack would allocate VMs up to 6 GHz total on the Host.

Service offerings 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.

CloudStack does not perform memory over-provisioning.

VLAN Provisioning

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

CloudStack 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.

Managing Networks and Traffic

In a CloudStack cloud, guest VMs can communicate with each other using shared infrastructure with the security and user perception that the guests have a private LAN.

The CloudStack virtual router is the main component providing networking features for guest traffic.

Guest Traffic

A network can carry guest traffic only between VMs within one zone. Virtual machines in different zones cannot communicate with each other using their IP addresses; they must communicate with each other by routing through a public IP address.

Figure 1 illustrates a typical guest traffic setup.




Figure 1 Guest Traffic Setup

The Management Server automatically creates a virtual router for each network. A virtual router is a special virtual machine that runs on the hosts. Each virtual router has three network interfaces. Its eth0 interface serves as the gateway for the guest traffic and has the IP address of 10.1.1.1. Its eth1 interface is used by the system to configure the virtual router. Its eth2 interface is assigned a public IP address for public traffic.

The virtual router provides DHCP and will automatically assign an IP address for each guest VM within the IP range assigned for the network. The user can manually reconfigure guest VMs to assume different IP addresses.

Source NAT is automatically configured in the virtual router to forward outbound traffic for all guest VMs.

Networking in a Pod

Figure 2 illustrates network setup within a single pod. The hosts are connected to a pod-level switch. At a minimum, the hosts should have one physical uplink to each switch. Bonded NICs are supported as well. The pod-level switch is a pair of redundant gigabit switches with 10 G uplinks.

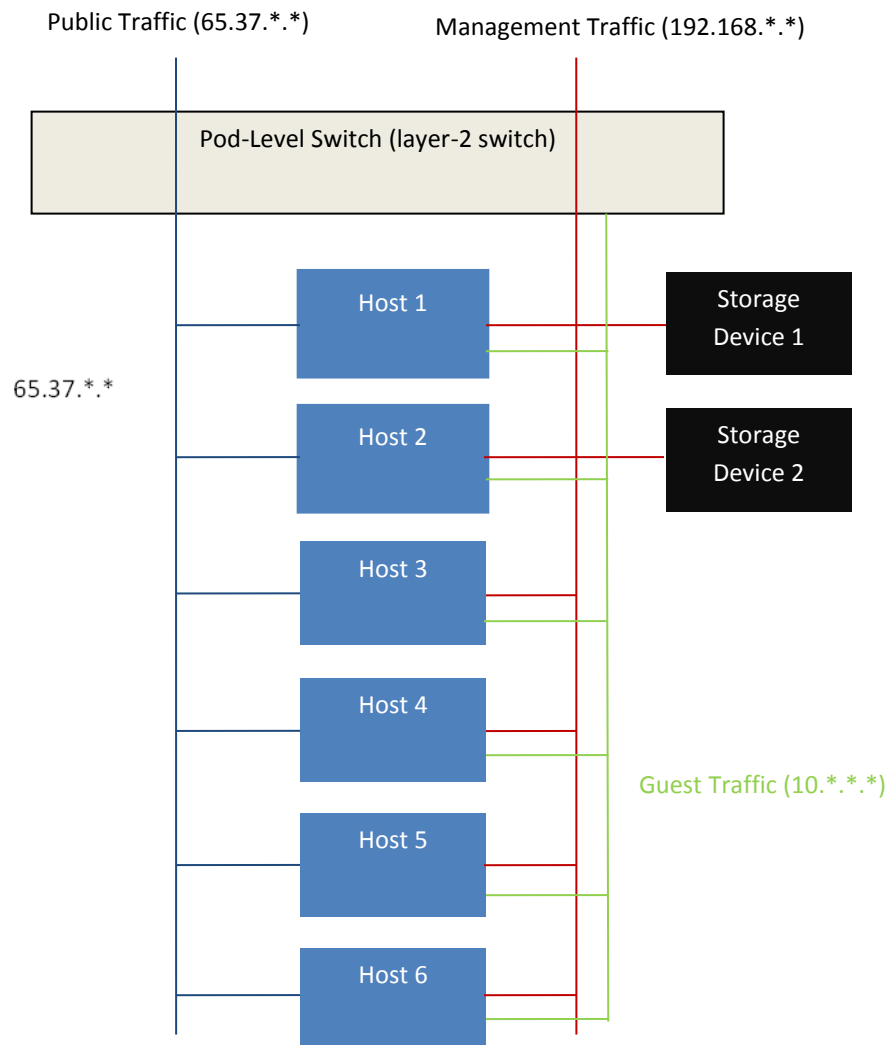


Figure 2 Network Setup within a Single Pod – Logical View

Servers are connected as follows:

- Storage devices are connected to only the network that carries management traffic.
- Hosts are connected to networks for both management traffic and public traffic.
- Hosts are also connected to one or more networks carrying guest traffic.

We recommend the use of multiple physical Ethernet cards to implement each network interface as well as redundant switch fabric in order to maximize throughput and improve reliability.

Networking in a Zone

Figure 3 illustrates the network setup within a single zone.

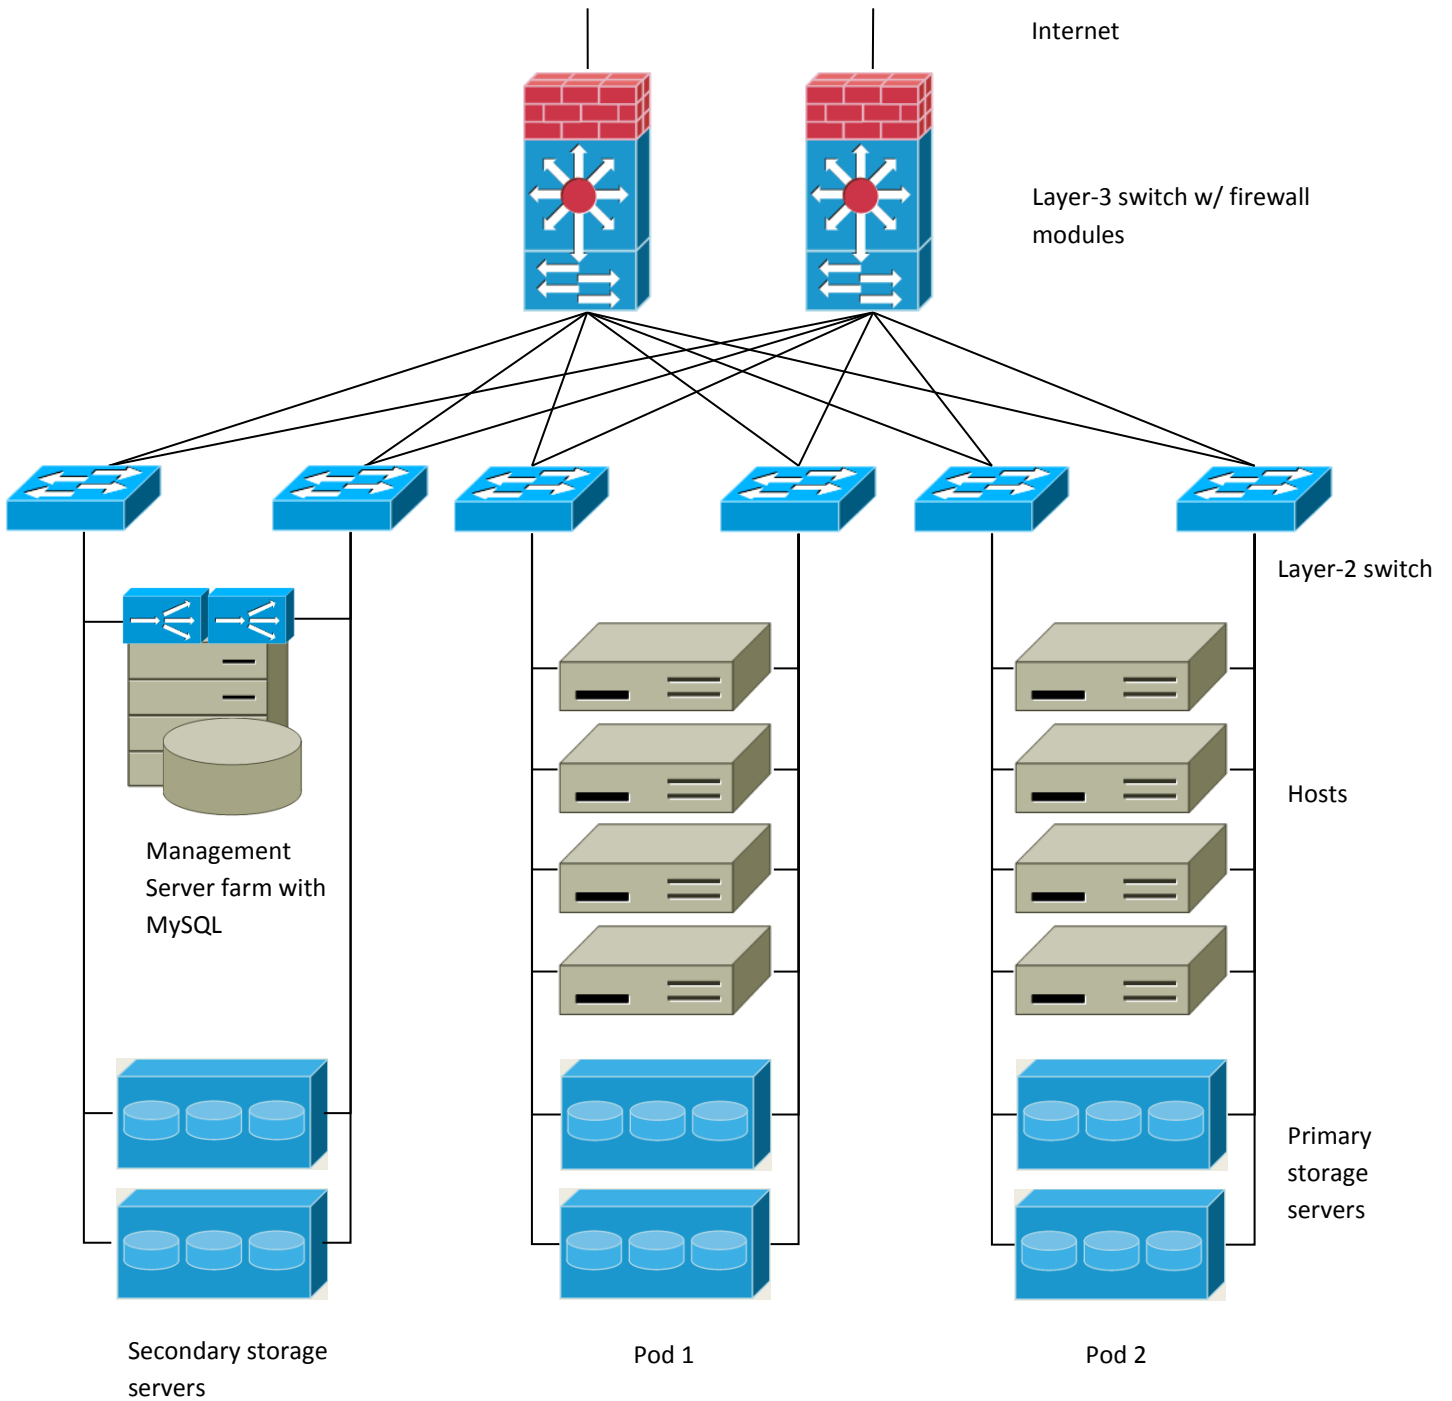


Figure 3 Network setup in a zone

A firewall for management traffic operates in the NAT mode. The network typically is assigned IP addresses in the 192.168.0.0/16 Class B private address space. Each pod is assigned IP addresses in the 192.168.*.0/24 Class C private address space.

Each zone has its own set of public IP addresses. Public IP addresses from different zones do not overlap.

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 CloudStack, you set up the guest network through the Add Zone screens (see the Advanced Installation Guide).

About Guest IP Addresses in a Basic Zone

When basic networking is used, CloudStack will assign IP addresses in the CIDR of the pod to the guests in that pod. The administrator must add a Direct IP range on the pod for this purpose. These IPs are in the same VLAN as the hosts.

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.

Configure Guest Traffic in an Advanced Zone

In a zone that uses advanced networking, you need to configure at least one guest network to carry traffic that is generated by guest VMs.

About Guest IP Addresses in an Advanced Zone

When advanced networking is used, the administrator can create additional networks for use by the guests. These networks can span the zone and be available to all accounts, or they can be scoped to a single account, in which case only the named account may create guests that attach to these networks. The networks are defined by a VLAN ID, IP range, and gateway. The administrator may provision thousands of these networks if desired.

Configuring Guest Traffic in an Advanced Zone

These steps assume you have already logged in to the CloudStack UI (see page 21). 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 network.
4. Provide the following information:
 - **Name.** The name of the network. This will be user-visible.

- **Description:** The description of the network. This will be user-visible.
- **VLAN ID:** Enter an administrator-configured VLAN ID so you can create different networks for use by different VM users in the zone.
- **Scope:** Choose account-specific or domain-specific if you would like to make the network accessible to only a single account or domain. Choose zone-wide if all accounts with access to the zone should be able to access the network.
- **Domain/Account:** If Scope is account-specific, enter the domain and account name for the account.
- **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.
- **Start IP/End IP:** Enter the first and last IP addresses that define a range that CloudStack can assign to guests. If one NIC is used, these IPs should be in the same CIDR as the pod CIDR. If multiple NICs are used, they may be in a different subnet.
- **Network Domain:** (Optional) If you want to assign a special domain name to this network, specify the DNS suffix.

5. Click OK.

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.

About Public IP Addresses

CloudStack provisions one public IP address per account for use as the source NAT IP address. If a Juniper SRX firewall is used, CloudStack can instead use a single public IP address as an interface NAT IP for all accounts, reducing the number of IP addresses consumed. Users may request additional public IP addresses. The administrator must configure one or more ranges of public IP addresses for use by CloudStack. These IP addresses could be RFC1918 addresses in private clouds.

Adding IP Addresses for the Public Network

These steps assume you have already logged in to the CloudStack UI (see page 21).

1. In the left navigation, choose Infrastructure. In Zones, click View More, then click the desired zone
2. Click the Network tab.
3. In the Public node of the diagram, click Configure.
4. Click the IP Ranges tab.

5. Enter the following details.

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

6. Click Add.

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 controls 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.

Adding an Additional Guest Network


1. Log in to the CloudStack UI as an administrator or end user (see page 21).
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.
 - **Description.** The description of the network. This will be user-visible.
 - **Network offering.** If the administrator has configured multiple network offerings, select the one you want to use for this network.

- **Pod.** The name of the pod this network applies to. Each pod in a basic zone is a broadcast domain, and therefore each pod has a different IP range for the guest network. The administrator must configure the IP range for each pod.
- **VLAN ID.** The VLAN tag for this network.
- **Gateway.** The gateway that the guests should use.
- **Netmask.** The netmask in use on the subnet the guests will use.
- **Start IP/End IP.** Enter the first and last IP addresses that define a range that CloudStack can assign to guests. We strongly recommend the use of multiple NICs. If multiple NICs are used, they may be in a different subnet. If one NIC is used, these IPs should be in the same CIDR as the pod CIDR.

4. Click Create.

Changing the Network Offering on a Guest Network

A user or administrator can change the network offering that is associated with an existing guest network.

1. Log in to the CloudStack UI as an administrator or end user (see page 21).
2. If you are changing from a network offering that uses the CloudStack virtual router to one that uses external devices as network service providers, you must first stop all the VMs on the network. See *Stopping and Starting VMs* on page 43. Then return here and continue to the next step.
3. In the left navigation, choose Networks.
4. Click the name of the network you want to modify.
5. Click the Edit button. 
6. In Network Offering, choose the new network offering, then click Apply.
7. A prompt appears asking whether you want to keep the existing CIDR. This is to let you know that if you change the network offering, the CIDR will be affected. Choose No to proceed with the change.
8. Wait for the update to complete. Don't try to restart VMs until after the network change is complete.
9. If you stopped any VMs in step 2, restart them.

Security Groups

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 CloudStack 3.0.0 - 3.0.2, security groups are supported only in zones that use basic networking.

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

Each CloudStack 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 CloudStack 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.

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 Basic Zone Configuration in the Advanced Installation Guide.

Adding a Security Group

1. Log in to the CloudStack UI as a user or administrator.
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.
7. The new security group appears in the Security Groups Details tab.
8. To make the security group useful, continue to Adding Ingress and Egress Rules to a Security Group on page 68.

Adding Ingress and Egress Rules to a Security Group

1. Log in to the CloudStack UI as a user or administrator.
2. In the left navigation, choose Network.
3. In Select view, choose Security Groups, then click the security group you want to work with.

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 CloudStack 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 CloudStack 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 the same name you used in step 7.

The following example allows inbound HTTP access from anywhere.

Protocol	Start Port	End Port	CIDR	Add
TCP	80	80	0.0.0.0/0	Add

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 CloudStack 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 CloudStack 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.

6. Click Add.

External Firewalls and Load Balancers

CloudStack is capable of replacing its Virtual Router with an external Juniper SRX device and an optional external NetScaler or F5 load balancer for gateway and load balancing services. In this case, the VMs use the SRX as their gateway.

For information about how to add external firewalls and load balancers to CloudStack, see the Advanced Installation Guide.

About Using a NetScaler Load Balancer

Citrix NetScaler is supported as an external network element for load balancing in zones that use advanced networking (also called advanced zones). Set up an external load balancer when you want to provide load balancing through means other than CloudStack's provided virtual router.

The NetScaler can be set up in direct (outside the firewall) mode. It must be added before any load balancing rules are deployed on guest VMs in the zone.

The functional behavior of the NetScaler with CloudStack is the same as described in the CloudStack documentation for using an F5 external load balancer. The only exception is that the F5 supports routing domains, and NetScaler does not. NetScaler can not yet be used as a firewall.

The Citrix NetScaler comes in three varieties. The following table summarizes how these variants are treated in CloudStack.

NetScaler ADC Type	Description of Capabilities	CloudStack 3.0 Supported Features
MPX	Physical appliance. Capable of deep packet inspection. Can act as application firewall and load balancer.	In advanced zones, load balancer functionality fully supported without limitation. In basic zones, static NAT, elastic IP (EIP), and elastic load balancing (ELB) are also provided.
VPX	Virtual appliance. Can run as VM on XenServer, ESXi, and Hyper-V hypervisors. Same functionality as MPX.	Supported only on ESXi. Same functional support as for MPX. CloudStack will treat VPX and MPX as the same device type.
SDX	Physical appliance. Can create multiple fully isolated VPX instances on a single appliance to support multi-tenant usage.	Cloudstack will dynamically provision, configure, and manage the lifecycle of VPX instances on the SDX. Provisioned instances are added into CloudStack automatically – no manual configuration by the administrator is required. Once a VPX instance is added into CloudStack, it is treated the same as a VPX on an ESXi host.

Initial Setup of External Firewalls and Load Balancers

When the first VM is created for a new account, CloudStack 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).

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, CloudStack 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 CloudStack database.

Load Balancer Rules

A CloudStack user or administrator may create load balancing rules that balance traffic received at a public IP to one or more VMs. A user creates a rule, specifies an algorithm, and assigns the rule to a set of VMs.

Adding a Load Balancer Rule

1. Log in to the CloudStack UI as an administrator or user.
2. In the left navigation bar, click 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.
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.

WARNING

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 CloudStack 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.

- **Algorithm.** Choose the load balancing algorithm you want CloudStack to use. CloudStack 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 Sticky Session Policies for Load Balancer Rules on page 73.

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

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 CloudStack 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 CloudStack UI or call `listNetworks` and check the `SupportedStickinessMethods` capability.

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.

Acquiring a New IP Address

The administrator allocates a set of publicly accessible IPs for use by VMs that require access to a public network. Administrators or end users can use the CloudStack UI to acquire these IPs to implement NAT between their guest network and the public network.


1. Log in to the CloudStack UI as an administrator or user.
2. In the left navigation bar, click Network.
3. Click the name of the guest network 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.

Releasing an IP Address

IP addresses are a limited resource. If you no longer need a particular IP, you can disassociate it from its network and return it to the pool of available addresses.

1. Log in to the CloudStack UI as an administrator or user.
2. In the left navigation bar, click Network.
3. Click the name of the guest network 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. 

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.

Enabling or Disabling Static NAT

If port forwarding rules are already in effect for an IP address, you can not 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 CloudStack UI as an administrator or user.
2. In the left navigation bar, click Network.
3. Click the name of the guest network you want to work with.
4. Click View IP Addresses.
5. Click the IP address you want to work with.

6. Click the Static NAT button. 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.

IP Forwarding and Firewalling

By default, all incoming traffic to the public IP address is rejected. All outgoing traffic from the guests is translated via NAT to the public IP address and is allowed.

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 the steps to implement these rules, see Firewall Rules on page 75 and Port Forwarding on page 76.

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 can not 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 Security Groups on page 67.

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

To create a firewall rule:

1. Log in to the CloudStack UI as a user or administrator.
2. In the left navigation bar, click Network.
3. Click the name of the network 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.

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 can not 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 on page 67.

To set up port forwarding:

1. Log in to the CloudStack UI.
2. If you have not already done so, add a public IP address range to a zone in CloudStack. See Adding a Zone and Pod in the Installation Guide.
3. Add one or more VM instances to CloudStack.
4. In the left navigation bar, click Network.
5. Click the name of the guest network where the VMs are running.
6. Click View IP Addresses.
7. Choose an existing IP address or acquire a new IP address. (See Acquiring a New IP Address on page 73.) Click the name of the IP address in the list.
8. Click the Configuration tab.
9. In the Port Forwarding node of the diagram, click View All.
10. 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.

11. Click Add VM. Choose the name of the instance to which this rule applies.

12. Test by opening an ssh session to the instance.

IP Load Balancing

The user may choose to associate the same public IP for multiple guests. CloudStack 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.

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.

VPN

CloudStack account owners can create virtual private networks (VPN) to access their virtual machines. 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.

CloudStack 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. VPN clients native to Windows, Mac OS X and iOS can be used to connect to the guest networks. The account owner can create and manage users for their VPN. CloudStack 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.

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.

Configuring VPN

To set up VPN for the cloud:

1. Log in as administrator to the CloudStack UI.
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.

To enable VPN for a particular network:

1. Log in as a user or administrator to the CloudStack 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.

Using 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 Cloudstack 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 -> 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.

Using VPN with Mac OS X

In Mac OS X, in Network Preferences – Advanced, make sure Send all traffic over VPN connection is not checked.

Working With Storage

CloudStack 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.

There is no ephemeral storage in CloudStack. All volumes on all nodes are persistent.

Primary Storage

This section gives concepts and technical details about CloudStack primary storage. For information about how to install and configure primary storage through the CloudStack UI, see the Advanced Installation Guide.

About Primary Storage

Primary storage is associated with a cluster, and it stores the disk volumes for all the VMs running on hosts in that cluster. You can add multiple primary storage servers to a cluster. At least one is required. It is typically located close to the hosts for increased performance. CloudStack manages the allocation of guest virtual disks to particular primary storage devices.

CloudStack is designed to work with all standards-compliant iSCSI and NFS servers that are supported by the underlying hypervisor, including, for example:

- Dell EqualLogic™ for iSCSI
- Network Appliances filers for NFS and iSCSI
- Scale Computing for NFS

The use of the Cluster Logical Volume Manager (CLVM) for KVM is not officially supported with CloudStack 3.0.x.

System Requirements for Primary Storage

Hardware requirements:

- Any standards-compliant iSCSI or NFS server that is supported by the underlying hypervisor.
- The storage server should be a machine with a large number of disks. The disks should ideally be managed by a hardware RAID controller.
- Minimum required capacity depends on your needs.

WARNING

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

When setting up primary storage, follow these restrictions:

- Primary storage cannot be added until a host has been added to the cluster.

- If you do not provision shared storage for primary storage, you will not be able to create additional volumes.
- If you do not provision shared primary storage, you must set the global configuration parameter `system.vm.local.storage.required` to true, or else you will not be able to start VMs.

Best Practices for Primary Storage

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

Runtime Behavior of Primary Storage

Root volumes are created automatically when a virtual machine is created. Root volumes are deleted when the VM is destroyed. Data volumes can be created and dynamically attached to VMs (although, when the OVM hypervisor is used, the VM must be stopped before an additional volume can be attached). Data volumes are not deleted when VMs are destroyed.

Administrators should monitor the capacity of primary storage devices and add additional primary storage as needed. See the Advanced Installation Guide.

Administrators add primary storage to the system by creating a CloudStack storage pool. Each storage pool is associated with a cluster.

Hypervisor Support for Primary Storage

The following table shows storage options and parameters for different hypervisors.

	VMware vSphere	Citrix XenServer	KVM	Oracle VM (OVM)
Format for Disks, Templates, and Snapshots	VMDK	VHD	QCOW2	RAW
iSCSI support	VMFS	Clustered LVM	Yes, via Shared Mountpoint	Yes, via OCFS2
Fiber Channel support	VMFS	Yes, via Existing SR	Yes, via Shared Mountpoint	No
NFS support	Y	Y	Y	Y
Local storage support	Y	Y	Y	N
Storage over-provisioning	NFS and iSCSI	NFS	NFS	No

XenServer uses a clustered LVM system to store VM images on iSCSI and Fiber Channel volumes and does not support over-provisioning in the hypervisor. The storage server itself, however, can support thin-provisioning. As a result the CloudStack can still support storage over-provisioning by running on thin-provisioned storage volumes.

KVM supports "Shared Mountpoint" storage. A shared mountpoint is a file system path local to each server in a given cluster. The path must be the same across all Hosts in the cluster, for example /mnt/primary1. This shared mountpoint is assumed to be a clustered filesystem such as OCFS2. In this case the CloudStack does not attempt to mount or unmount the storage as is done with NFS. The CloudStack requires that the administrator insure that the storage is available.

Oracle VM (OVM) supports both iSCSI and NFS storage. When iSCSI is used with OVM, the CloudStack administrator is responsible for setting up iSCSI on the host, including re-mounting the storage after the host recovers from a failure such as a network outage.

With hypervisors other than OVM, CloudStack takes care of mounting the iSCSI target on the host whenever it discovers a connection with an iSCSI server and unmounting the target when it discovers the connection is down.

With NFS storage, CloudStack manages the overprovisioning. In this case the global configuration parameter `storage.overprovisioning.factor` controls the degree of overprovisioning. This is independent of hypervisor type.

Local storage is an option for primary storage for vSphere, XenServer, KVM, and OVM. When the local disk option is enabled, a local disk storage pool is automatically created on each host. To use local storage for the System Virtual Machines (such as the Virtual Router), set `system.vm.use.local.storage` to `true` in global configuration.

CloudStack supports multiple primary storage pools in a Cluster. For example, you could provision 2 NFS servers in primary storage. Or you could provision 1 iSCSI LUN initially and then add a second iSCSI LUN when the first approaches capacity.

Storage Tags

Storage may be "tagged". A tag is a text string attribute associated with primary storage, a Disk Offering, or a Service Offering. Tags allow administrators to provide additional information about the storage. For example, that is a "SSD" or it is "slow". Tags are not interpreted by CloudStack. They are matched against tags placed on service and disk offerings. CloudStack requires all tags on service and disk offerings to exist on the primary storage before it allocates root or data disks on the primary storage. Service and disk offering tags are used to identify the requirements of the storage that those offerings have. For example, the high end service offering may require "fast" for its root disk volume.

The interaction between tags, allocation, and volume copying across clusters and pods can be complex. To simplify the situation, use the same set of tags on the primary storage for all clusters in a pod. Even if different devices are used to present those tags, the set of exposed tags can be the same.

Maintenance Mode for Primary Storage

Primary storage may be placed into maintenance mode. This is useful, for example, to replace faulty RAM in a storage device. Maintenance mode for a storage device will first stop any new guests from being provisioned on the storage device. Then it will stop all guests that have any volume on that storage device. When all such guests are stopped the storage device is in maintenance mode and may be shut down. When the storage device is online again you may cancel maintenance mode for the device. The CloudStack will bring the device back online and attempt to start all guests that were running at the time of the entry into maintenance mode.

Secondary Storage

This section gives concepts and technical details about CloudStack secondary storage. For information about how to install and configure secondary storage through the CloudStack UI, see the Advanced Installation Guide.

About Secondary Storage

Secondary storage is associated with a zone, and it stores the following:

- Templates – OS images that can be used to boot VMs and can include additional configuration information, such as installed applications
- ISO images – disc images containing data or bootable media for operating systems
- Disk volume snapshots – saved copies of VM data which can be used for data recovery or to create new templates

The items in secondary storage are available to all hosts in the zone.

System Requirements for Secondary Storage

- NFS storage appliance or Linux NFS server
- (Optional) OpenStack Object Storage (Swift)
- 100GB minimum capacity
- A secondary storage device must be located in the same zone as the guest VMs it serves.
- Each secondary storage server must be available to all hosts in the zone.

WARNING

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

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

Secondary Storage VM

In addition to the hosts, CloudStack's Secondary Storage VM mounts and writes to secondary storage.

Submissions to secondary storage go through the Secondary Storage VM. The Secondary Storage VM can retrieve templates and ISO images from URLs using a variety of protocols.

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.

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.

```
# mysql -p
mysql> use cloud;
mysql> select id from host where type = 'SecondaryStorage';
mysql> update host_details set value = 'nfs://192.168.160.20/export/mike-ss1'
  where host_id = HOSTID and name = 'orig.url';
```

```
mysql> update host set name = 'nfs://192.168.160.20/export/mike-ss1' where type  
= 'SecondaryStorage' and id = #;  
mysql> update host set url = 'nfs://192.168.160.20/export/mike-ss1' where type  
= 'SecondaryStorage' and id = #;  
mysql> update host set guid = 'nfs://192.168.160.20/export/mike-ss1' where type  
= 'SecondaryStorage' and id = #;
```

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

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. Use the procedure above to change the IP address for secondary storage if required.
5. Start the Management Server(s).

Using Swift for Secondary Storage

CloudStack supports OpenStack Object Storage (Swift, <http://swift.openstack.org>) for secondary storage. When using Swift, you configure Swift storage for the entire CloudStack, then set up NFS secondary storage for each zone as usual. The NFS storage in each zone acts as a staging area through which all templates and other secondary storage data pass before being forwarded to Swift. The Swift storage acts as a cloud-wide resource, making templates and other data available to any zone in the cloud. There is no hierarchy in the Swift storage, just one Swift container per storage object. Any secondary storage in the whole cloud can pull a container from Swift at need. It is not necessary to copy templates and snapshots from one zone to another, as would be required when using zone NFS alone. Everything is available everywhere.

Swift storage must be set up before you add NFS secondary storage to zones. This is accomplished through some additional configuration steps on a fresh Management Server installation, before you add the first zone. The procedure is described in [Adding a Zone in the Advanced Installation Guide](#).

Working with Volumes

A volume provides storage to a guest VM. The volume can provide for a root disk or an additional data disk. CloudStack 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 (e.g. XenServer) may not be attached to a guest that is using another hypervisor type (e.g. vSphere, KVM, OVM). This is because the different hypervisors use different disk image formats.

CloudStack 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 (e.g. 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.

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 CloudStack administrators and users can add volumes to VM instances. When you create a new volume, it is stored as an entity in CloudStack, but the actual storage resources are not allocated on the physical storage device until you attach the volume. This optimization allows the CloudStack to provision the volume nearest to the guest that will use it when the first attachment is made.

1. Log in to the CloudStack UI as an administrator or user.
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 CloudStack, but the volume is not yet ready for use.

5. To start using the volume, continue to Attaching a Volume on page 87.

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 can not 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 on page 122.

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, CloudStack will use this value to verify that no data corruption has occurred.
2. Log in to the CloudStack UI as an administrator or user.
3. In the left navigation, 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
VMware	OVA
KVM	QCOW2
OVM	RAW

- **URL.** The secure HTTP or HTTPS URL that CloudStack 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. Click Instances - Volumes, find the name you specified in step 5, and make sure the status is Uploaded.

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 CloudStack UI as an administrator or user.
2. In the left navigation bar, 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 popup, 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.

If the VM is running in the OVM hypervisor, the VM must be stopped before a new volume can be attached to it.

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


Detaching and Moving Volumes

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

This procedure is different from moving disk volumes from one storage pool to another. See VM Storage Migration on page 88.

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.

If the destination VM is running in the OVM hypervisor, the VM must be stopped before a new volume can be attached to it.

1. Log in to the CloudStack UI as an administrator or user.
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 Attaching a Volume on page 87.

VM Storage Migration

Supported in XenServer, KVM, and VMware.

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.

This procedure is different from moving disk volumes from one VM to another. See Detaching and Moving Volumes on page 88.

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.

Migrating a Data Disk Volume to a New Storage Pool

1. Log in to the CloudStack UI as an administrator.
2. Detach the data disk from the VM. See Detaching and Moving Volumes on page 88 (but skip the “reattach” step at the end. You will do that after migrating to new storage.)
3. Call the CloudStack API command `migrateVolume` and pass in the volume ID and the ID of any storage pool in the zone.
4. Watch for the volume status to change to `Migrating`, then back to `Ready`.
5. Attach the volume to any desired VM running in the same cluster as the new storage server. See Attaching a Volume on page 87.

Migrating a VM Root Volume to a New Storage Pool

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 CloudStack UI as administrator.
2. Detach any data disks from the VM. See Detaching and Moving Volumes on page 88.
3. Stop the VM.
4. Call the CloudStack API command `migrateVirtualMachine` with the ID of the VM to migrate and the IDs of a destination host and destination storage pool in the same zone.
5. Watch for the VM status to change to `Migrating`, then back to `Stopped`.
6. Restart the VM.

Resizing Volumes

CloudStack does not provide the ability to resize root disks or data disks; the disk size is fixed based on the template used to create the VM. However, the tool [VHD Resizer](http://vmtoolkit.com/files/folders/converters/entry87.aspx) (<http://vmtoolkit.com/files/folders/converters/entry87.aspx>), while not officially supported by Cloud.com or Citrix, might provide a workaround. To increase disk size with VHD Resizer:

1. Get the VHD from the secondary storage.
2. Import it into VHD Resizer.
3. Resize the VHD.
4. Upload the new VHD.
5. Create a new VM.
6. On a Linux guest, extend the file system to reflect the new disk size. Manually resize your partitions using the operating system's utilities. For example, use `resize2fs` or use LVM utilities to add partition space.
7. Take a snapshot, then create a new template from that snapshot.

For more information, see [How to Resize a Provisioning Server 5 Virtual Disk](http://support.citrix.com/article/CTX118608) at the Citrix Knowledge Center (<http://support.citrix.com/article/CTX118608>).

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.

Working with ISOs

CloudStack supports ISOs and their attachment to guest VMs. An ISO is a read-only file that has an ISO/CD-ROM style file system. Users can upload their own ISOs and mount them on their guest VMs.

ISOs are uploaded based on a URL. HTTP is the supported protocol. Once the ISO is available via HTTP specify an upload URL such as `http://my.web.server/filename.iso`.

ISOs may be public or private, like templates.

ISOs are not hypervisor-specific. That is, a guest on vSphere can mount the exact same image that a guest on KVM can mount.

ISO images may be stored in the system and made available with a privacy level similar to templates. ISO images are classified as either bootable or not bootable. A bootable ISO image is one that contains an OS image (E.g. An Ubuntu 10.04 installation CD). CloudStack allows a user to boot a guest VM off of an ISO image. Users can also attach ISO images to guest VMs. For example, this enables installing PV drivers into Windows. ISO images are not hypervisor-specific.

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 CloudStack UI as an administrator or user.
2. In the left navigation bar, click Templates.

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. (E.g. CentOS 6.2 64 bit)
 - **Description.** Display test for the ISO image. (E.g. 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 CloudStack.
 - **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 CloudStack 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). This choice is not valid for Ubuntu. To create an Ubuntu PV template, see [Creating an Ubuntu 10.04 LTS Template for XenServer](#) on page 104.
 - (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:
 - Ubuntu 10.04
 - Ubuntu 9
 - Ubuntu 8.10
 - Fedora 13
 - Fedora 12
 - Fedora 11
 - Fedora 10
 - Fedora 9
 - 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

Debian GNU/Linux

Other PV

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 CloudStack can work with it.

Attaching an ISO to a VM

A virtual machine is created using either a template or an ISO image. Once the VM exists, you can attach an ISO to it.

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.

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 CloudStack.

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 CloudStack.

CloudStack ships with a default template. In order to present more choices to users, CloudStack administrators and users can create templates and add them to CloudStack.

Creating Templates: Overview

CloudStack 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 CloudStack. 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 CloudStack.

The various techniques for creating templates are described in the next few sections.

Requirements for Templates

- For XenServer, install PV drivers / Xen tools on each template that you create. This will enable live migration and clean guest shutdown.
- For vSphere, install VMware Tools on each template that you create. This will enable console view to work properly.

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.

The Default Template

CloudStack includes a CentOS template. This template is downloaded by the Secondary Storage VM after the primary and secondary storage are configured. You can use this template in your production deployment or you can delete it and use custom templates.

The root password for the default template is "password".

A default template is provided for each of XenServer, KVM, and vSphere. The templates that are downloaded depend on the hypervisor type that is available in your cloud. Each template is approximately 2.5 GB physical size.

The default template includes the standard iptables rules, which will block most access to the template excluding ssh.

```
# iptables --list
Chain INPUT (policy ACCEPT)
```

```

target      prot opt source      destination
RH-Firewall-1-INPUT  all  --  anywhere    anywhere

Chain FORWARD (policy ACCEPT)
target      prot opt source      destination
RH-Firewall-1-INPUT  all  --  anywhere    anywhere

Chain OUTPUT (policy ACCEPT)
target      prot opt source      destination

Chain RH-Firewall-1-INPUT (2 references)
target      prot opt source      destination
ACCEPT      all  --  anywhere    anywhere
ACCEPT      icmp --  anywhere    anywhere    icmp any
ACCEPT      esp  --  anywhere    anywhere
ACCEPT      ah   --  anywhere    anywhere
ACCEPT      udp  --  anywhere    224.0.0.251  udp dpt:mdns
ACCEPT      udp  --  anywhere    anywhere    udp dpt:ipp
ACCEPT      tcp  --  anywhere    anywhere    tcp dpt:ipp
ACCEPT      all  --  anywhere    anywhere    state RELATED,ESTABLISHED
ACCEPT      tcp  --  anywhere    anywhere    state NEW tcp dpt:ssh
REJECT      all  --  anywhere    anywhere    reject-with icmp-host-prohibited

```

Private and Public Templates

When a user creates a template, it can be designated private or public.

Private templates are only available to the user who created them. By default, an uploaded template is private.

When a user marks a template as “public,” the template becomes available to all users in all accounts in the user's domain, as well as users in any other domains that have access to the Zone where the template is stored. This depends on whether the Zone, in turn, was defined as private or public. A private Zone is assigned to a single domain, and a public Zone is accessible to any domain. If a public template is created in a private Zone, it is available only to users in the domain assigned to that Zone. If a public template is created in a public Zone, it is available to all users in all domains.

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 in [Creating VMs on page 42](#).
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 CloudStack 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 CloudStack installation. The template will appear in the Community Templates list. See Private and Public Templates on page 94.
- **Password Enabled.** Choose Yes if your template has the CloudStack password change script installed. See Adding Password Management to Your Templates on page 107.

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.

Creating a Template From a Snapshot

(Not supported with OVM)

If you do not want to stop the VM in order to use the Create Template menu item (as described in Creating a Template from an Existing Virtual Machine on page 94), you can create a template directly from any snapshot through the CloudStack UI.

Uploading Templates

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 times.

WARNING

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.

To upload a template:

1. In the left navigation bar, click Templates.
2. Click Create Template.
3. Provide the following:
 - **Name and Display Text.** These will be shown in the UI, so choose something descriptive.
 - **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, or All Zones to make it available throughout CloudStack.
 - **OS Type:** This helps CloudStack 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 use Ubuntu 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.

 - **Hypervisor.**
 - **Format.** The format of the template upload file, such as VHD or OVA.
 - **Password Enabled.** Choose Yes if your template has the CloudStack password change script installed. See Adding Password Management to Your Templates on page 107.
 - **Public.** Choose Yes to make this template accessible to all users of this CloudStack installation. The template will appear in the Community Templates list. See Private and Public Templates on page 94.
 - **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.

Exporting Templates

End users and Administrators may export templates from the CloudStack. Navigate to the template in the UI and choose the Download function from the Actions menu.

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.

An overview of the procedure is as follows:

1. Upload your Windows ISO. (If you need help, see Adding an ISO on page 90.)
2. Create a VM Instance with this ISO. (If you need help, see Creating VMs on page 42.)
3. Follow the steps in Sysprep for Windows Server 2008 R2 (below) or Sysprep for Windows Server 2003 R2 on page 100, 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 on page 101.

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 CloudStack management functions such as mounting additional volumes and ISO images, live migration, and graceful shutdown.

Sysprep 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 the Microsoft Download Center at the following location:

<http://www.microsoft.com/downloads/details.aspx?FamilyID=94bb6e34-d890-4932-81a5-5b50c657de08&DisplayLang=en>

Use the following steps to run sysprep for Windows 2008 R2.¹

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 “Select a Windows image or catalog file” 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.

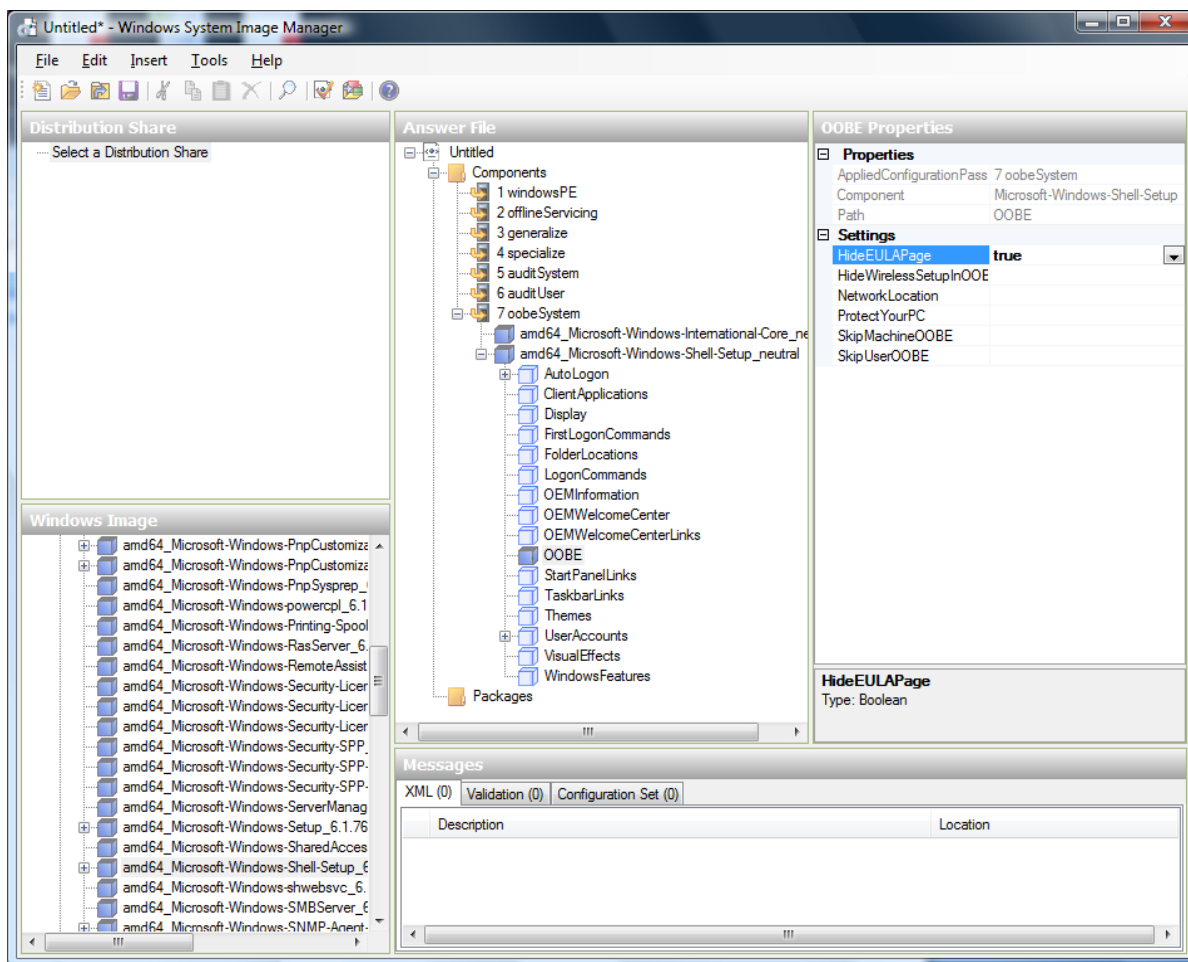
¹ The steps outlined here are derived from the excellent guide by Charity Shelbourne, originally published at the following URL: <http://blogs.technet.com/askcore/archive/2008/10/31/automating-the-oobe-process-during-windows-server-2008-sysprep-mini-setup.aspx>

7. Generate the answer file from the Windows System Image Manager using the following steps.

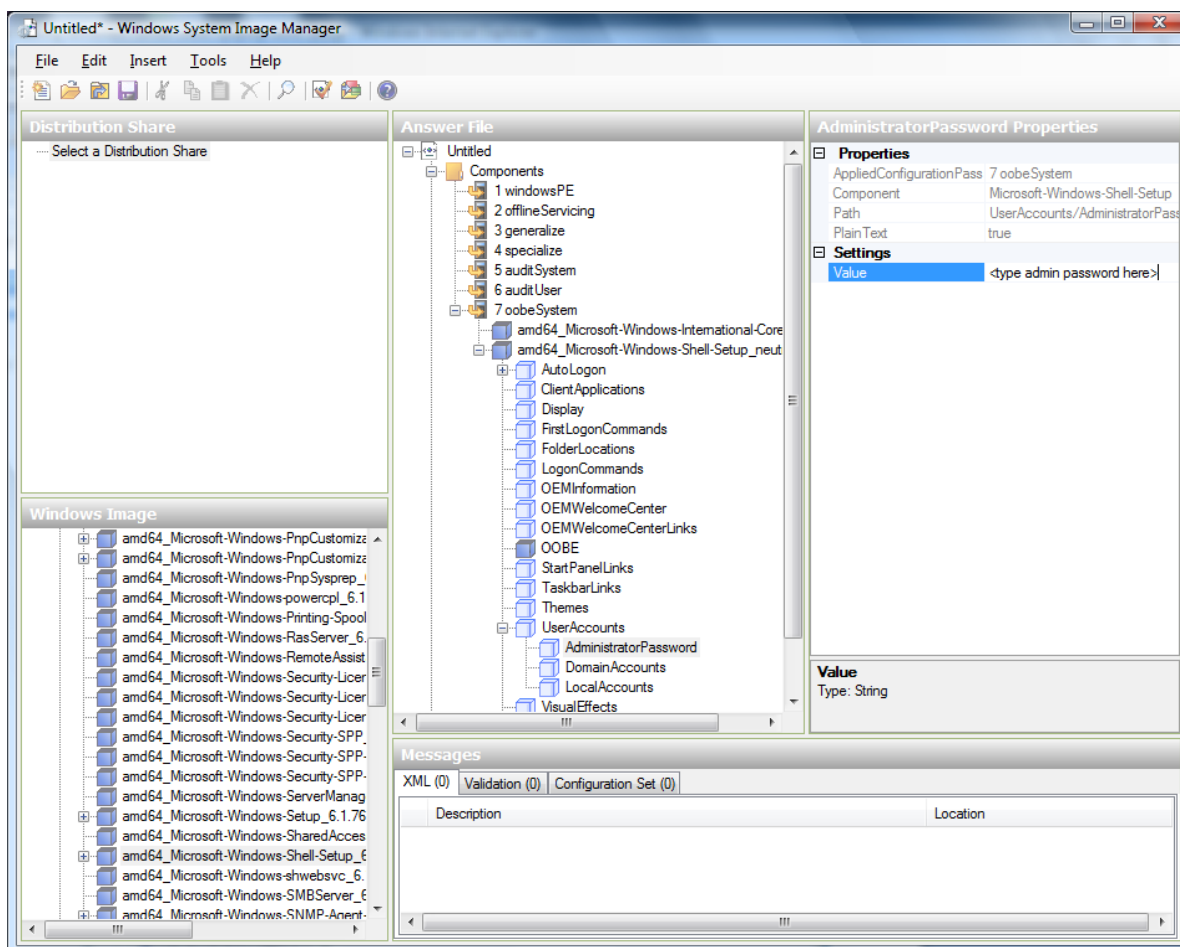
- a. The first page you need to automate is the Language and Country or Region Selection page. To automate this, expand Components in your Windows Image pane, right-click and add the Microsoft-Windows-International-Core setting to Pass 7 oobeSystem. In your Answer File pane, configure the InputLocale, SystemLocale, UILanguage, and UserLocale with the appropriate settings for your language and country or region. Should you have a question 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 to 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 here: <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

Sysprep 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:
4. `c:\sysprep\sysprep.exe -reseal -mini -activated`
5. After this step the machine will automatically shut down.

Creating the Windows Template

Once your VM has shut down, you can create a template.

1. Make sure the VM status is Stopped.
2. Click on Instances and find your VM. Click on it.
3. Click View Volumes and find the root disk. Click on it.
4. Click the Create Template button. This button appears only on stopped VMs.



Importing AMIs

The following procedures describe how to import an Amazon Machine Image (AMI) into CloudStack 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.

Note: 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.

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.

1. Set up loopback on 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 CloudStack UI or API, install the password change script into the image at this point. See [Adding Password Management to Your Templates](#) on page 107.

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/
```

Creating an Ubuntu 10.04 LTS Template for XenServer

This section tells how to create an Ubuntu 10.04 LTS template so that you can create Ubuntu VM instances on the XenServer hypervisor.

1. In XenServer, create a running Ubuntu 10.04 VM by following these steps. (Copied from [Installing Ubuntu 10.04 LTS](http://community.citrix.com/display/xs/Installing+Ubuntu+Server+10.04+%2832bit+and+64bit%29+LTS) on the Citrix Developer Network. Check the website for the most up-to-date information: <http://community.citrix.com/display/xs/Installing+Ubuntu+Server+10.04+%2832bit+and+64bit%29+LTS>.)
 - a. Copy `makeubuntu.sh` script to the Pool Master (download the script from the Citrix Developer Network link above).

- b. Execute `makeubuntu.sh` script to create Ubuntu Templates.
- c. Create an Ubuntu VM with the new templates.
- d. Perform install and reboot.

2. Perform the following tests.

- a. Make sure the VM is booted with one NIC (`eth0`).
- b. Open the file `/etc/network/interfaces` and be sure that `eth0` is set to use DHCP.

3. Stop the Ubuntu VM.

In the next few steps, you will copy the virtual machine's virtual hard disk (VHD) to a web server.

4. From the XenServer command line, list the VMs with the following command. Note the UUID of your Ubuntu VM.

```
# xe vm-list
```

5. List the virtual block devices (VBDs) with the following command, passing in the VM UUID you discovered in the previous step. Note the VDI UUID for the VBD.

```
# xe vbd-list <your Ubuntu VM UUID>
```

6. Navigate to the mount point for primary storage to find the VHD file. The name of the VHD file is the VDI UUID you discovered in the previous step.

7. Copy the VHD file to a webserver.

8. In the CloudStack UI, click Add Template.

- a. Select the URL of the VHD file on the web server as the location.
- b. For the guest OS type, select Ubuntu if you are running XenServer 5.6 FP1 or greater (for earlier XenServer versions, select CentOS 5.4 x64). Alternatively, if you want Ubuntu to boot in PV mode, select Other PV (64-bit).

9. Start a new VM from the template.

10. Make sure the VM was able to get an IP address. If not, follow these troubleshooting steps:

- a. Start a CentOS 5.3 x64 VM.
- b. On the CentOS VM, run this command to find the location of the DHCP client script.

```
# which dhclient
```

The location returned should be `/sbin/modified-dhclient/dhclient`.

- c. On the Ubuntu VM, create a new folder.

```
# mkdir /sbin/modified-dhclient
```

- d. Copy the `dhclient` script from the CentOS VM to the Ubuntu VM at `/sbin/modified-dhclient/dhclient`.
- e. Add the new folder to the front of your VM's path.
- f. Log out of the VM and log in again.

Converting a Hyper-V VM to a Template

To convert a Hyper-V VM to a XenServer-compatible CloudStack template, you will need a standalone XenServer host with an attached NFS VHD SR. Use whatever XenServer version you are using with CloudStack, 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 (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 [Creating a Windows Template](#) on page 96).

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 CloudStack, 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.

Adding Password Management to Your Templates

CloudStack 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 CloudStack 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 CloudStack, 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.

Window OS Installation

Download the installer, CloudInstanceManager.msi, from <http://cloudstack.org/download.html> and run the installer in the newly created Windows VM.

Linux OS Installation

Use the following steps to begin the Linux OS installation.

1. Download the script file `cloud-set-guest-password` from the CloudStack community on the Web:

- Linux:
<http://cloudstack.org/dl/cloud-set-guest-password>
- 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.

- Fedora, CentOS/RHEL, and Debian.** Run “chkconfig --add cloud-set-guest-password”.
- Ubuntu.** If you are using Ubuntu 11.04, start by creating a directory called /var/lib/dhcp3 on your Ubuntu machine (works around a known issue with this version of Ubuntu). On all Ubuntu versions: Run “sudo update-rc.d cloud-set-guest-password defaults 98”. To test, run “mkpasswd” and check that it is generating a new password. If the “mkpasswd” command does not exist, run “sudo apt-get install whois” (or sudo apt-get install mkpasswd, depending on your Ubuntu version) and repeat.

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.

Working with Snapshots

(Supported for the following hypervisors: XenServer, VMware vSphere, and KVM)

CloudStack supports snapshots of disk volumes. Snapshots are a point-in-time capture of virtual machine disks. Memory and CPU states are not captured. If you are using the Oracle VM hypervisor, you can not take snapshots, since OVM does not support them.

Snapshots may be taken for volumes, including both root and data disks (except when the OVM hypervisor is used, which does not support snapshots). The administrator places a limit on the number of stored snapshots per user. Users can create new volumes from the snapshot for recovery of particular files and they can create templates from snapshots to boot from a restored disk.

Users can create snapshots manually or by setting up automatic recurring snapshot policies. Users can also create disk volumes from snapshots, which may be attached to a VM like any other disk volume. Snapshots of both root disks and data

disks are supported. However, CloudStack does not currently support booting a VM from a recovered root disk. A disk recovered from snapshot of a root disk is treated as a regular data disk; the data on recovered disk can be accessed by attaching the disk to a VM.

A completed snapshot is copied from primary storage to secondary storage, where it is stored until deleted or purged by newer snapshots.

Automatic Snapshot Creation and Retention

(Supported for XenServer, VMware vSphere, and KVM)

Users can set up a recurring snapshot policy to automatically create multiple snapshots of a disk at regular intervals. Snapshots can be created on an hourly, daily, weekly, or monthly intervals. One snapshot policy can be set up per disk volume. For example, a user can set up a daily snapshot at 02:30.

With each snapshot schedule, users can also specify the number of scheduled snapshots to be retained. Older snapshots that exceed the retention limit are automatically deleted. This user-defined limit must be equal to or lower than the global limit set by the CloudStack administrator (see Globally Configured Limits on page 122). The limit applies only to those snapshots that are taken as part of an automatic recurring snapshot policy. Additional manual snapshots can be created and retained.

Incremental Snapshots and Backup

Snapshots are created on primary storage where a disk resides. After a snapshot is created, it is immediately backed up to secondary storage and removed from primary storage for optimal utilization of space on primary storage.

CloudStack does incremental backups for some hypervisors. When incremental backups are supported, every N backup is a full backup.

	VMware vSphere	Citrix XenServer	KVM
Support incremental backup	N	Y	N

Volume Status

When a snapshot operation is triggered by means of a recurring snapshot policy, a snapshot is skipped if a volume has remained inactive since its last snapshot was taken. A volume is considered to be inactive if it is either detached or attached to a VM that is not running. CloudStack ensures that at least one snapshot is taken since the volume last became inactive.

When a snapshot is taken manually, a snapshot is always created regardless of whether a volume has been active or not.

Snapshot Restore

There are two paths to restoring snapshots. Users can create a volume from the snapshot. The volume can then be mounted to a VM and files recovered as needed. Alternatively, a template may be created from the snapshot of a root disk. The user can then boot a VM from this template to effect recovery of the root disk.

Runtime Considerations

- Snapshots not only consume space in secondary storage, but can take up significant CPU cycles and network bandwidth as the snapshots are moved between primary and secondary storage. This is something to be factored in for capacity planning and end-user pricing of snapshot operations.
- (VMware vSphere only) When attaching a volume that is created from a snapshot based on a root volume as a data disk, depending on how the root volume was previously partitioned by a guest VM, the guest operating system of the new host VM might not recognize all the partitions. Users will have to find corresponding tools if they want to fully work on all the partitions inside the volume. CloudStack does not provide such tools.
- (VMware vSphere only) When deploying a new VM based on the template created from a snapshot based on a data volume, the VM might not be bootable and usable.

Working with System Virtual Machines

CloudStack uses several types of system virtual machines to perform tasks in the cloud. In general CloudStack 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.

The System VM Template

The System VMs come from a single template. The System VM has the following characteristics:

- Debian 6.0 ("Squeeze"), 2.6.32 kernel with the latest security patches from the Debian security APT repository
- Has a minimal set of packages installed thereby reducing the attack surface
- 32-bit for enhanced performance on Xen/VMWare
- pvops kernel with Xen PV drivers, KVM virtio drivers, and VMware tools for optimum performance on all hypervisors
- Xen tools inclusion allows performance monitoring
- Latest versions of HAProxy, iptables, IPsec, and Apache from debian repository ensures improved security and speed
- Latest version of JRE from Sun/Oracle ensures improved security and speed

Multiple System VM Support for VMware

Every CloudStack zone has single System VM for template processing tasks such as downloading templates, uploading templates, and uploading ISOs. In a zone where VMware is being used, additional System VMs can be launched to process VMware-specific tasks such as taking snapshots and creating private templates. The CloudStack management server launches additional System VMs for VMware-specific tasks as the load increases. The management server monitors and weights all commands sent to these System VMs and performs dynamic load balancing and scaling-up of more System VMs.

Console Proxy

The Console Proxy is a type of System Virtual Machine that has a role in presenting a console view via the web UI. It connects the user's browser to the VNC port made available via the hypervisor for the console of the guest. Both the administrator and end user web UIs offer a console connection.

Clicking on a console icon brings up a new window. The AJAX code downloaded into that window refers to the public IP address of a console proxy VM. There is exactly one public IP address allocated per console proxy VM. The AJAX application connects to this IP. The console proxy then proxies the connection to the VNC port for the requested VM on the Host hosting the guest.

Note: The hypervisors will have many ports assigned to VNC usage so that multiple VNC sessions can occur simultaneously.

There is never any traffic to the guest virtual IP, and there is no need to enable VNC within the guest.

The console proxy VM will periodically report its active session count to the Management Server. The default reporting interval is five seconds. This can be changed through standard Management Server configuration with the parameter `consoleproxy.loadscan.interval`.

Assignment of guest VM to console proxy is determined by first determining if the guest VM has a previous session associated with a console proxy. If it does, the Management Server will assign the guest VM to the target Console Proxy VM regardless of the load on the proxy VM. Failing that, the first available running Console Proxy VM that has the capacity to handle new sessions is used.

Console proxies can be restarted by administrators but this will interrupt existing console sessions for users.

The console viewing functionality uses a dynamic DNS service under the domain name `realhostip.com` to assist in providing SSL security to console sessions. The console proxy is assigned a public IP address. In order to avoid browser warnings for mismatched SSL certificates, the URL for the new console window is set to the form of `https://aaa-bbb-ccc-ddd.realhostip.com`. Customers will see this URL during console session creation. CloudStack includes the `realhostip.com` SSL certificate in the console proxy VM. Of course, CloudStack cannot know about DNS A records for our customers' public IPs prior to shipping the software. CloudStack therefore runs a dynamic DNS server that is authoritative for the `realhostip.com` domain. It maps the `aaa-bbb-ccc-ddd` part of the DNS name to the IP address `aaa.bbb.ccc.ddd` on lookups. This allows the browser to correctly connect to the console proxy's public IP, where it then expects and receives a SSL certificate for `realhostip.com`, and SSL is set up without browser warnings.

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 CloudStack 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 CloudStack.

```
openssl pkcs8 -in yourprivate.pkcs8.encrypted.key -out yourprivate.pkcs8.key
```

3. In the Update SSL Certificate screen of the CloudStack UI, paste the following:

- Certificate from step 1(c).
- Private key from step 1(e).
- The desired new domain name; for example, company.com.

4. Click Add to put the changes into effect.

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 will generate URLs of the form "aaa-bbb-ccc-ddd.company.com" after this change is made. New console requests will be served with the new DNS domain name, certificate, and key.

Virtual Router

The virtual router is a type of System Virtual Machine. The virtual router is one of the most frequently used service providers in CloudStack. The end user has no direct access to the virtual router. Users can ping the virtual router and take actions that affect it (such as setting up port forwarding), but users do not have SSH access into the virtual router.

There is no mechanism for the administrator to log in to the virtual router. Virtual routers can be restarted by administrators, but this will interrupt public network access and other services for end users.

A basic test in debugging networking issues is to attempt to ping the virtual router from a guest VM.

Some of the characteristics of the virtual router are determined by its associated system service offering.

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 CloudStack fetches network usage statistics from CloudStack 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.

Upgrading a Virtual Router with System Service Offerings

When CloudStack creates a virtual router, it uses default settings which are defined in a default system service offering (see System Service Offerings on page 38). 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 Creating a New System Service Offering on page 38. In System VM Type, choose Domain Router.
2. Associate the system service offering with a network offering. See Creating a New Network Offering on page 34.
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 Changing the Network Offering on a Guest Network on page 67.

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 CloudStack 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.

Secondary Storage VM

The secondary storage VM provides a background task that takes care of a variety of secondary storage activities: downloading a new template to a Zone, copying templates between Zones, and snapshot backups.

The administrator can log in to the secondary storage VM if needed. The procedure for this is documented in the Troubleshooting section of the Installation Guide.

System Reliability and High Availability

HA for Management Server

The CloudStack 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 Servers. 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.

HA for Hosts

When hosts are down, CloudStack will restart impacted HA-enabled VMs automatically, assuming that other hosts have sufficient resources available. When the host comes back online it will be marked as available and newly started VMs may be allocated to it. VMs previously migrated from it will not be migrated back. VMs that were running on it but did not have HA enabled will not be started automatically.

The user will not lose the virtual machine disk image during a host outage. However, the guest OS may perceive its disk image as corrupt (and needing fsck or equivalent) on restart.

Primary Storage Outage and Data Loss

When a primary storage outage occurs the hypervisor immediately stops all VMs stored on that storage device. 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.

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.

HA-Enabled Virtual Machines

The user can specify a virtual machine as HA-enabled. All virtual router VMs and system VMs are automatically configured as HA-enabled. When an HA-enabled VM crashes, CloudStack detects the crash and restarts the VM automatically within the same Availability Zone. HA is never performed across different Availability Zones. CloudStack 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.

HA features work with iSCSI or NFS primary storage. HA with local storage is not supported.

Managing the Cloud

Setting Global Configuration Parameters

CloudStack provides parameters that you can set to control many aspects of the cloud. When CloudStack is first installed, and periodically thereafter, you might need to modify these settings.

1. Log in to the CloudStack 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.
 - 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. 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.

Changing the Database Configuration

The CloudStack Management Server stores database configuration information (e.g., hostname, port, credentials) in the file `/etc/cloud/management/db.properties`. To effect a change, edit this file on each Management Server, then restart the Management Server.

Administrator Alerts

The system provides alerts and events to help with the management of the cloud. Alerts are notices to an administrator, generally delivered by e-mail, notifying the administrator that an error has occurred in the cloud. Alert behavior is configurable.

Events track all of the user and administrator actions in the cloud. For example, every guest VM start creates an associated event. Events are stored in the Management Server's database.

Emails will be sent to administrators 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

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 CloudStack 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 (see Adding an Additional Guest Network on page 66) or with the `updateNetwork` command in the CloudStack API.
 - At the account, domain, or zone level, the DNS suffix can be assigned with the appropriate CloudStack 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 CloudStack 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 CloudStack 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.

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 cloud-management stop
```

To start the Management Server:

```
# service cloud-management start
```

To perform a stop and start in one command:

```
# service cloud-management restart
```

Working with Usage

The Usage Server is an optional, separately-installed part of CloudStack that provides aggregated usage records which you can use to create billing integration for CloudStack. 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.

Configuring the Usage Server

1. Be sure the Usage Server has been installed. This requires extra steps beyond just installing the CloudStack software. See [Installing the Usage Server \(Optional\)](#) in the Advanced Installation Guide.
2. Log in to the CloudStack 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	Parameter Definition
enable.usage.server	Whether the Usage Server is active.
usage.aggregation.timezone	<p>Time zone of usage records. Set this if the usage records and daily job execution are in different 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:</p> <pre>usage.stats.job.exec.time = 00:15 usage.execution.timezone = PST usage.aggregation.timezone = GMT</pre> <p>Valid values for the time zone are specified in Appendix A—Time Zones on page 137. Default: GMT.</p>
usage.execution.timezone	<p>The time zone of usage.stats.job.exec.time. Valid values for the time zone are specified in Appendix A—Time Zones on page 137.</p> <p>Default: The time zone of the management server.</p>
usage.sanity.check.interval	<p>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 <code>ALERT_TYPE_USAGE_SANITY_RESULT = 21</code> is sent.</p>
usage.stats.job.aggregation.range	<p>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.</p> <p>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. CloudStack 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. CloudStack assumes this "midnight to midnight" is relative to the <code>usage.execution.timezone</code>.</p> <p>Default: 1440</p>

usage.stats.job.exec.time	<p>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”.</p> <p>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.</p> <p>Default: 00:15.</p>
---------------------------	---

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.

Setting Usage Limits

CloudStack 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 subdomains to the creation of 100 VMs.

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	Default data transfer rate in megabits per second allowed per user (supported on XenServer)
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 CloudStack UI. See [Setting Global Configuration Parameters](#) on page 117.

Default Account Resource Limits

You can limit resource use by accounts. The default limits are set using global configuration parameters, and they affect all accounts within a cloud. The relevant parameters are those beginning with `max.account` (`max.account.snapshots`, etc.).


To override a default limit for a particular account, set a per-account resource limit.

1. Log in to the CloudStack 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. 

Per-Domain Limits

CloudStack 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 CloudStack UI.
2. In the left navigation tree, click **Domains**.
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. 

CloudStack API

The CloudStack API is a low level API that has been used to implement the CloudStack web UIs. It is also a good basis for implementing other popular APIs such as EC2/S3 and emerging DMTF standards.

Many CloudStack 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 API Reference at http://docs.cloudstack.org/CloudStack_Documentation.

Provisioning and Authentication API

CloudStack 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 CloudStack to add/remove users.

CloudStack supports pluggable authenticators. By default, CloudStack 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 on page 17.

Allocators

CloudStack 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.

User Data and Meta Data

CloudStack provides API access to attach 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.

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.

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 `-XmxNNMm` 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" in the Tomcat Wiki at <http://wiki.apache.org/tomcat/FAQ/Memory>.

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 MySQL configuration file `/etc/my.cnf`.
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
```

For more information about the buffer pool, see "The InnoDB Buffer Pool" in the MySQL Reference Manual at <http://dev.mysql.com/doc/refman/5.5/en/innodb-buffer-pool.html>.

Set and Monitor Total VM Limits per Host

The CloudStack 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 CloudStack 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 $(N-1) * (\text{per-host-limit})$. Once a cluster reaches this number of VMs, use the CloudStack UI to disable allocation of more VMs to the cluster.

Configure XenServer dom0 Memory

Configure the XenServer dom0 settings to allocate more memory to dom0. This can enable XenServer to handle larger numbers of virtual machines. We recommend 2940 MB of RAM for XenServer dom0. For instructions on how to do this, see <http://support.citrix.com/article/CTX126531>. The article refers to XenServer 5.6, but the same information applies to XenServer 6.

Troubleshooting

This section describes how to diagnose and remedy runtime issues.

Event Logs

There are two types of events logged in the CloudStack 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.

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.

Long Running Job Events

In addition to the three standard event types, the events log also records the three following events for long running jobs.

- **SCHEDULED.** (Asynchronous jobs only) This event is generated when an asynchronous job is submitted.
- **STARTED.** This event is generated when a job begins execution.
- **COMPLETED.** This event is generated when a job is completed.

Both the Started and Completed events are logged for all long running job types. The Scheduled event is only logged for asynchronous events. When an action is initiated synchronously or as part of another asynchronous job, the Scheduled event won't be logged.

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

Event Types

The following is the full list of event types.

VM.CREATE	USER.DISABLE	DISK.OFFERING.CREATE
VM.DESTROY	TEMPLATE.CREATE	DISK.OFFERING.EDIT
VM.START	TEMPLATE.DELETE	DISK.OFFERING.DELETE
VM.STOP	TEMPLATE.UPDATE	NETWORK.OFFERING.CREATE
VM.REBOOT	TEMPLATE.COPY	NETWORK.OFFERING.EDIT
VM.UPGRADE	TEMPLATE.DOWNLOAD.START	NETWORK.OFFERING.DELETE
VM.RESETPASSWORD	TEMPLATE.DOWNLOAD.SUCCESS	POD.CREATE
ROUTER.CREATE	TEMPLATE.DOWNLOAD.FAILED	POD.EDIT
ROUTER.DESTROY	TEMPLATE.EXTRACT	POD.DELETE
ROUTER.START	TEMPLATE.UPLOAD	ZONE.CREATE
ROUTER.STOP	TEMPLATE.CLEANUP	ZONE.EDIT
ROUTER.REBOOT	VOLUME.CREATE	ZONE.DELETE
ROUTER.HA	VOLUME.DELETE	VLAN.IP.RANGE.CREATE
PROXY.CREATE	VOLUME.ATTACH	VLAN.IP.RANGE.DELETE
PROXY.DESTROY	VOLUME.DETACH	CONFIGURATION.VALUE.EDIT
PROXY.START	VOLUME.EXTRACT	SG.AUTH.INGRESS
PROXY.STOP	VOLUME.UPLOAD	SG.REVOKE.INGRESS
PROXY.REBOOT	SERVICEOFFERING.CREATE	HOST.RECONNECT
PROXY.HA	SERVICEOFFERING.UPDATE	MAINT.CANCEL
VNC.CONNECT	SERVICEOFFERING.DELETE	MAINT.CANCEL.PS
VNC.DISCONNECT	DOMAIN.CREATE	MAINT.PREPARE
NET.IPASSIGN	DOMAIN.DELETE	MAINT.PREPARE.PS
NET.IPRELEASE	DOMAIN.UPDATE	VPN.REMOTE.ACCESS.CREATE
NET.RULEADD	SNAPSHOT.CREATE	VPN.REMOTE.ACCESS.DESTROY
NET.RULEDELETE	SNAPSHOT.DELETE	VPN.USER.ADD
NET.RULEMODIFY	SNAPSHOTPOLICY.CREATE	VPN.USER.REMOVE
NETWORK.CREATE	SNAPSHOTPOLICY.UPDATE	NETWORK.RESTART
NETWORK.DELETE	SNAPSHOTPOLICY.DELETE	UPLOAD.CUSTOM.CERTIFICATE
LB.ASSIGN.TO.RULE	ISO.CREATE	STATICNAT.ENABLE
LB.REMOVE.FROM.RULE	ISO.DELETE	STATICNAT.DISABLE
LB.CREATE	ISO.COPY	SSVM.CREATE
LB.DELETE	ISO.ATTACH	SSVM.DESTROY
LB.UPDATE	ISO.DETACH	SSVM.START
USER.LOGIN	ISO.EXTRACT	SSVM.STOP
USER.LOGOUT	ISO.UPLOAD	SSVM.REBOOT
USER.CREATE	SERVICE.OFFERING.CREATE	SSVM.HA
USER.DELETE	SERVICE.OFFERING.EDIT	
USER.UPDATE	SERVICE.OFFERING.DELETE	

Alerts

The following is the list of alert type numbers.

```
MEMORY = 0
CPU = 1
STORAGE = 2
STORAGE_ALLOCATED = 3
PUBLIC_IP = 4
PRIVATE_IP = 5
HOST = 6
USERVM = 7
DOMAIN_ROUTER = 8
CONSOLE_PROXY = 9
ROUTING = 10 // lost connection to default route (to the gateway)
STORAGE_MISC = 11 // lost connection to default route (to the gateway)
USAGE_SERVER = 12 // lost connection to default route (to the gateway)
MANAGMENT_NODE = 13 // lost connection to default route (to the gateway)
DOMAIN_ROUTER_MIGRATE = 14
CONSOLE_PROXY_MIGRATE = 15
USERVM_MIGRATE = 16
VLAN = 17
SSVM = 18
USAGE_SERVER_RESULT = 19
STORAGE_DELETE = 20;
UPDATE_RESOURCE_COUNT = 21; //Generated when we fail to update the resource count
USAGE_SANITY_RESULT = 22;
DIRECT_ATTACHED_PUBLIC_IP = 23;
LOCAL_STORAGE = 24;
RESOURCE_LIMIT_EXCEEDED = 25; //Generated when the resource limit exceeds the limit. Currently
used for recurring snapshots only
```

Working with Server Logs

The CloudStack Management Server logs all web site, middle tier, and database activities for diagnostics purposes in `/var/log/cloud/management/`.

The CloudStack logs a variety of error messages. We recommend this command to find the problematic output in the Management Server log:

```
grep -i -E 'exception|unable|fail|invalid|leak|warn|error'
/var/log/cloud/management/management-server.log
```

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.

The CloudStack 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 CloudStack Agent Server logs its activities in `/var/log/cloud/agent/`.

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 CloudStack Installation Guide.

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):

1. Force stop the router. Use the stopRouter API with forced=true parameter to do so.
2. Before you continue with destroying this router, ensure that the backup router is running. Otherwise the network connection will be lost.
3. 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](#) on page 34.

For more information about the API syntax, see the API Reference at http://docs.cloudstack.org/CloudStack_Documentation/API_Reference%3A_CloudStack.

Maintenance mode not working on vCenter

Symptom

Host was placed in maintenance mode, but still appears live in vCenter.

Cause

The CloudStack 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 [Scheduled Maintenance and Maintenance Mode for Hosts](#) on page 55.

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.

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 in VMware. 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 http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=10051.

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 CloudStack 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.

Contacting Support

Open-source community

A variety of channels are available for getting help with CloudStack, from forums to IRC chat and more. For details, see <http://cloudstack.org/discuss/>.

Commercial customers

The CloudStack support team is available to help commercial customers plan and execute their installations. To contact the support team, log in to the support portal at <https://na6.salesforce.com/sserv/login.jsp?orgId=00D80000000LWom> using the account credentials you received when you purchased your support contract.

Appendix A—Time Zones

The following time zone identifiers are accepted by CloudStack. 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	America/La_Paz	Asia/Jerusalem
Etc/GMT+11	America/Santiago	Europe/Minsk
Pacific/Samoa	America/St_Johns	Europe/Moscow
Pacific/Honolulu	America/Araguaina	Africa/Nairobi
US/Alaska	America/Argentina/Buenos_Aires	Asia/Karachi
America/Los_Angeles	America/Cayenne	Asia/Kolkata
Mexico/BajaNorte	America/Godthab	Asia/Bangkok
US/Arizona	America/Montevideo	Asia/Shanghai
US/Mountain	Etc/GMT+2	Asia/Kuala_Lumpur
America/Chihuahua	Atlantic/Azores	Australia/Perth
America/Chicago	Atlantic/Cape_Verde	Asia/Taipei
America/Costa_Rica	Africa/Casablanca	Asia/Tokyo
America/Mexico_City	Etc/UTC	Asia/Seoul
Canada/Saskatchewan	Atlantic/Reykjavik	Australia/Adelaide
America/Bogota	Europe/London	Australia/Darwin
America/New_York	CET	Australia/Brisbane
America/Caracas	Europe/Bucharest	Australia/Canberra
America/Asuncion	Africa/Johannesburg	Pacific/Guam
America/Cuiaba	Asia/Beirut	Pacific/Auckland
America/Halifax	Africa/Cairo	