Apache CloudStack 4.1.0

CloudStack Administrator's Guide

cloudstack™

open source cloud computing

Apache CloudStack

Apache CloudStack 4.1.0 CloudStack Administrator's Guide

作者　　　　　　　　　　　　　　Apache CloudStack

Apache CloudStack is an effort undergoing incubation at The Apache Software Foundation (ASF).

Incubation is required of all newly accepted projects until a further review indicates that the infrastructure, communications, and decision making process have stabilized in a manner consistent with other successful ASF projects. While incubation status is not necessarily a reflection of the completeness or stability of the code, it does indicate that the project has yet to be fully endorsed by the ASF.

CloudStack管理指南

概念

1.1. 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.

1.2. CloudStack能做什么?

多种Hypervisor支持

CloudStack works with a variety of hypervisors, and a single cloud deployment can contain multiple hypervisor implementations. The current release of CloudStack supports pre-packaged enterprise solutions like Citrix XenServer and VMware vSphere, as well as KVM or Xen running on Ubuntu or CentOS.

大规模可扩展的管理架构

CloudStack可以管理数万台服务器；这些服务器可以部署在不同地域的数据中心里．处于中心位置的管理服务器可以线性扩展，这样就消除了对中间层集群级别管理服务器的依赖．任何一个组件失效不会导致云平台的服务暂停.对于管理服务器的定期维护不会对云平台中正在运行的虚拟机造成影响.

自动化配置管理

CloudStack会对客户虚拟机的网络和存储进行自动化配置.

CloudStack内部提供的虚拟设备池用来支持云平台自身功能．这些虚拟设备可以提供的服务有防火墙，路由，DHCP，VPN访问，控制台代理，存储访问以及存储备份等．虚拟设备的大量使用简化了安装，配置和持续的云平台部署管理流程.

图形用户界面

CloudStack提供了管理员Web接口，用来供应和管理整个云平台；同时也提供了类似最终用户的Web接口,用来管理运行中的虚机和模板．UI可以根据服务提供商的需求或企业的Web风格进行定制化.

API及其扩展性

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 Apache CloudStack Guides[1] and Apache CloudStack API Reference[2] respectively.

CloudStack 可插拨的allocation架构允许对选择的存储和主机创建新的allocator类型. 参见Allocator实现指导 (http://docs.cloudstack.org/CloudStack_Documentation/Allocator_Implementation_Guide).

高可用性

CloudStack平台有很多功能来增加系统的可用性．管理服务器自身可以在前端负载均衡的前提下部署在多个节点上．MySQL可以配置使用备份来提供在数据库丢失情况下的手工故障恢复．对于主机，CloudStack平台提供网卡绑定及为存储使用单独网络,这类似于iSCSI的多路径.

1.3. 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 production 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). In its smallest deployment, a single machine can act as both the Management Server and the hypervisor host (using the KVM hypervisor).

[1] http://cloudstack.apache.org/docs/en-US/index.html

[2] http://cloudstack.apache.org/docs/api/index.html

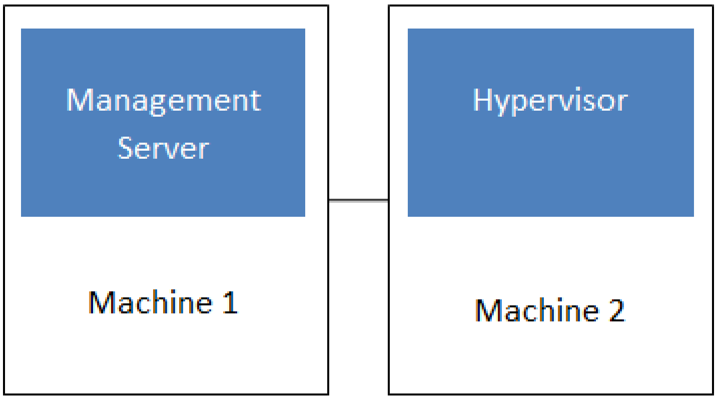

Simplified view of a basic deployment

A more full-featured installation consists of a highly-available multi-node Management Server installation and up to tens of thousands of hosts using any of several advanced networking setups. For information about deployment options, see the "Choosing a Deployment Architecture" section of the $PRODUCT; Installation Guide.

1.3.1. 管理服务器概述

管理服务器是CloudStack软件用来管理云环境的所有资源. 通过UI或API与管理服务器交互，你就可以配置并管理你的云基础架构.

一个管理服务器运行在专属的服务器或虚机里. 它控制虚拟机在主机上的分配，并且分配存储和IP地址到虚拟机实例. 管理服务器运行在一个Tomcat容器里并通过MySQL数据库进行持久化.

机器必须符合系统需求，在系统需求里有相关描述.

管理服务器:

- 为管理员提供一个Web用户接口并且为最终用户提供一个引用的用户接口.
- 为CloudStack提供API.
- 管理客户虚拟机到特定的主机分配.
- 管理公共IP及私有IP地址到账号的分配.
- 管理客户的存储作为虚拟磁盘的分配.
- 管理快照，模板，和ISO映像，并且可以在多个数据中心复制.
- 提供整个云环境的集中式配置.

1.3.2. Cloud Infrastructure Overview

The Management Server manages one or more zones (typically, datacenters) containing host computers where guest virtual machines will run. 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.
- Pod: A pod is usually one rack of hardware that includes a layer-2 switch and one or more clusters.

- Cluster: A cluster consists of one or more hosts and primary storage.

- Host: A single compute node within a cluster. The hosts are where the actual cloud services run in the form of guest virtual machines.

- Primary storage is associated with a cluster, and it stores the disk volumes for all the VMs running on hosts in that cluster.

- Secondary storage is associated with a zone, and it stores templates, ISO images, and disk volume snapshots.

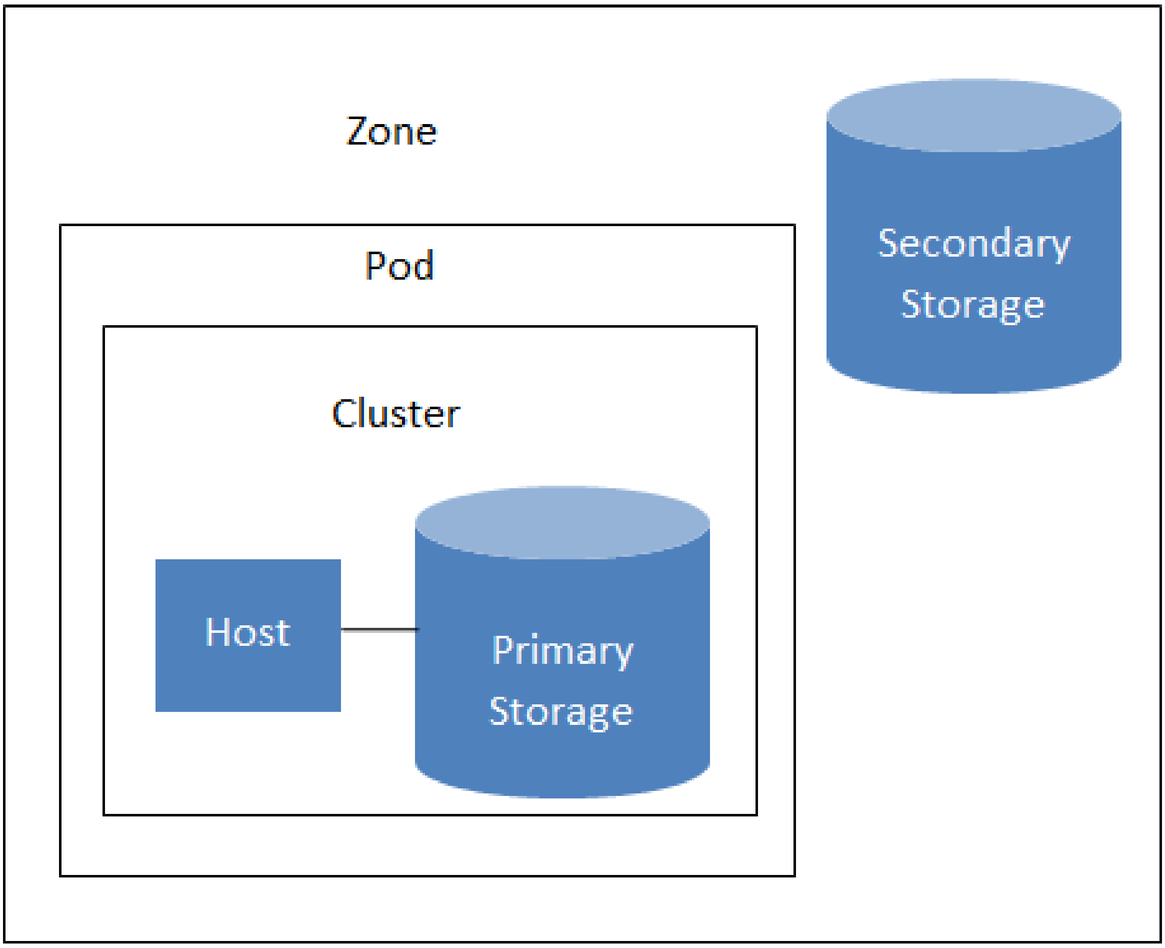

Nested organization of a zone

More Information

For more information, see documentation on cloud infrastructure concepts.

1.3.3. 网络概述

CloudStack提供两种类型的网络应用场景:

- 基本网络. 类似于AWS类型的网络. 提供一个单一网络, 在这个网络里客户通过提供的三层方式进行隔离,比如借安全组方式(源IP地址过滤).

- 高级网络. 为更复杂的网络拓扑设计. 网络模型提供了更为灵活的客户网络定义.

更详细的信息, 请参数网络设置.

云基础设施概念

2.1. About Regions

To increase reliability of the cloud, you can optionally group resources into multiple geographic regions. A region is the largest available organizational unit within a CloudStack deployment. A region is made up of several availability zones, where each zone is roughly equivalent to a datacenter. Each region is controlled by its own cluster of Management Servers, running in one of the zones. The zones in a region are typically located in close geographical proximity. Regions are a useful technique for providing fault tolerance and disaster recovery.

By grouping zones into regions, the cloud can achieve higher availability and scalability. User accounts can span regions, so that users can deploy VMs in multiple, widely-dispersed regions. Even if one of the regions becomes unavailable, the services are still available to the end-user through VMs deployed in another region. And by grouping communities of zones under their own nearby Management Servers, the latency of communications within the cloud is reduced compared to managing widely-dispersed zones from a single central Management Server.

Usage records can also be consolidated and tracked at the region level, creating reports or invoices for each geographic region.

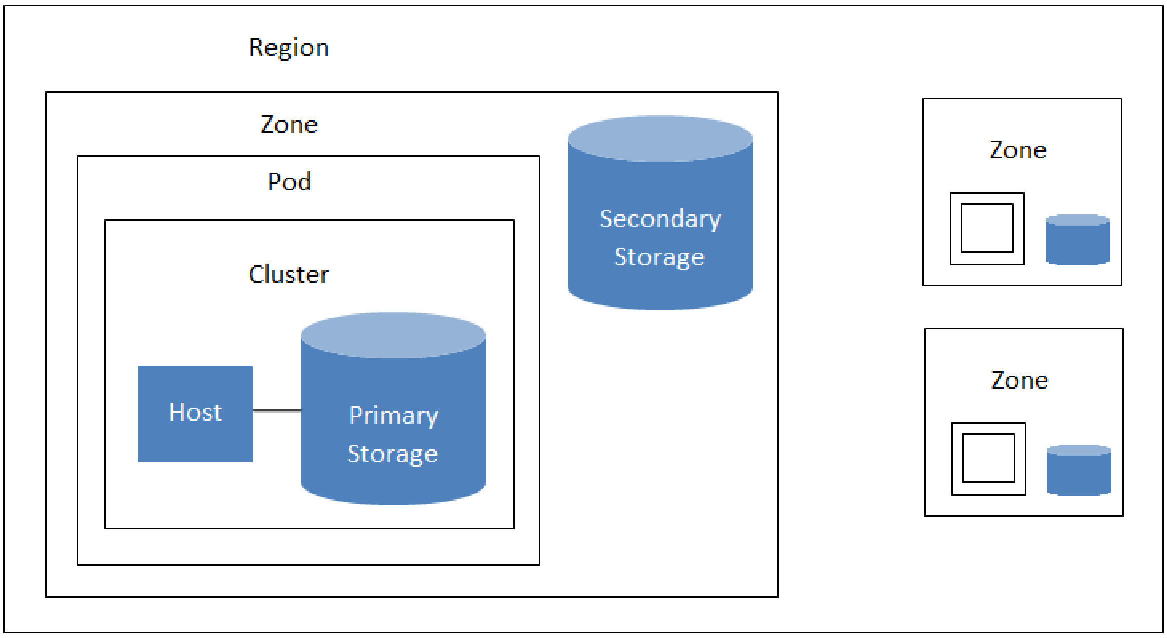

A region with multiple zones

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

2.2. 关于资源域

A zone is the second 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).

一个资源域包括:

- 一个或多个提供点。每个提供点包含一个或多个宿主机集群和iygehuoduoge主存储服务器。
- 二级存储是在资源域下的所有提供点共享的。

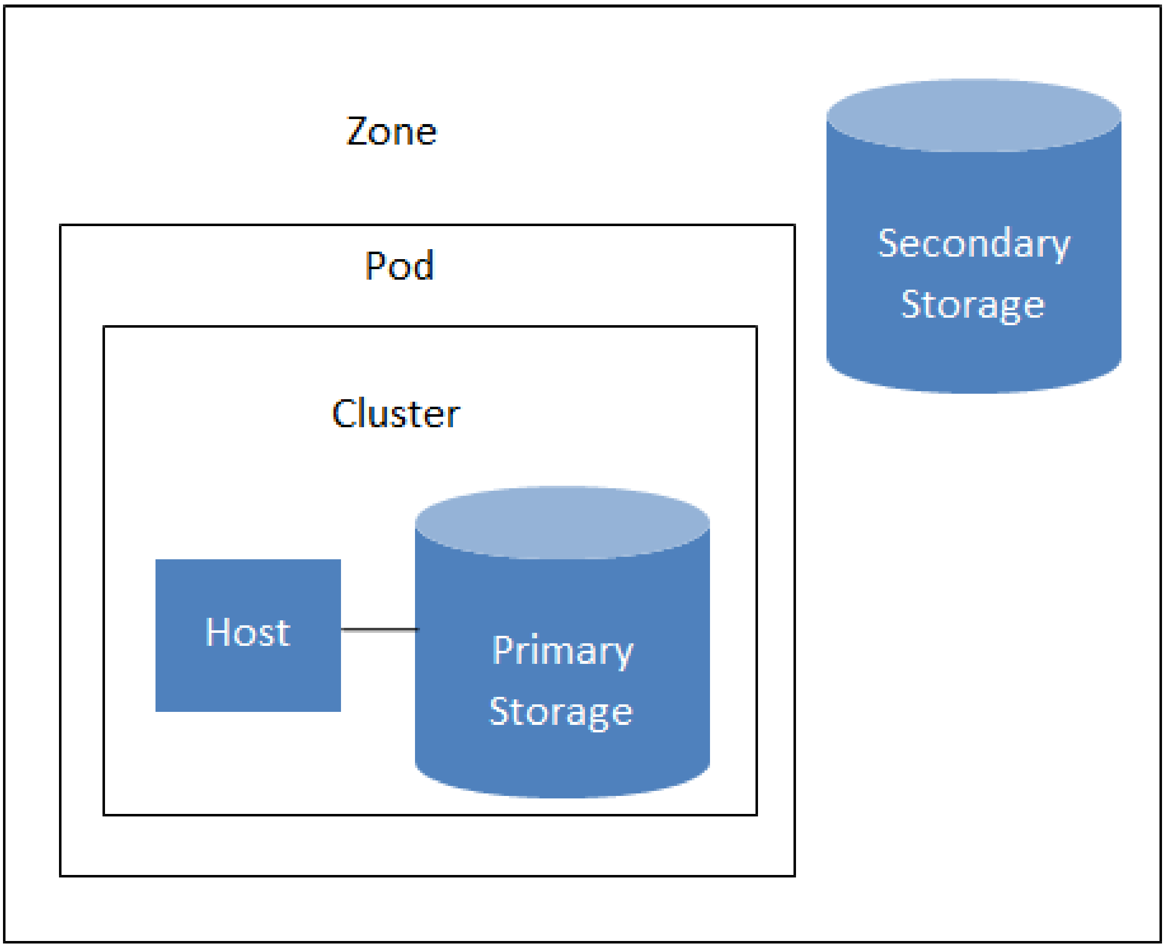

Nested organization of a zone

资源与对用户使可见的。当用户运行一个来宾虚拟机，他必须选择该虚拟机运行在哪个资源域上。当来宾虚拟机需要运行在额外的资源域时，用户可能需要拷贝私有模板向这些资源域拷贝私有模板。

资源域可以是公有也可以是私有的。公共资源与对所有用户可见。这意味着所用用户都可以在上面创建来宾虚拟机。私有资源域只对特定的鱼保留。只有在这个域或其子域的用户才能创建来宾虚拟机。

同一个资源域下的宿主机是不需要穿过防火墙的互连的机器。不同资源域的宿主机可以通过静态配置的vpn通道互相访问。

管理员必须决定每个资源域以下内容。

- 资源与中有多少提供点。
- 每个提供点中有多少集群。

- 每个集群中有几台宿主机。

- 每个集群中有几个主存储，且总的存储量要多大。

- 每个资源域下要有多少二级存储。

当您添加一个新的区域时，系统会提示您配置区域的物理网络，并添加提供点，集群，主机，主存储和二级存储。

2.3. 关于POD

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. Pods are not visible to the end user.

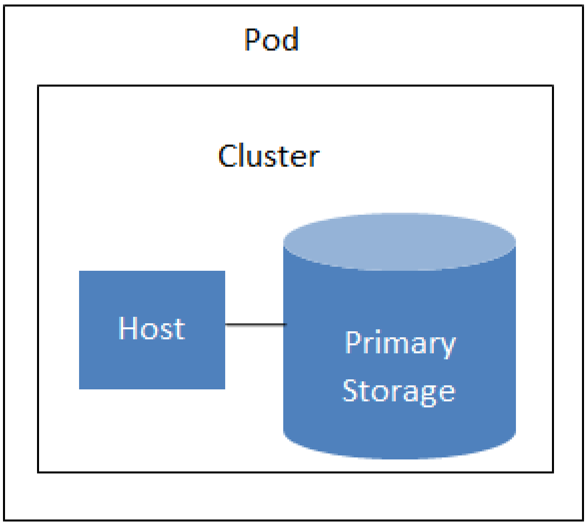

A simple pod

2.4. 关于集群

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.

集群在本产品中是第三大组织单位；\n部署\n集群隶属于pod之下，而pod隶属于zone之下。集群的大小取决于下层虚拟机软件。大多数情况下基本无建议；详见最佳实践

集群由一个或多个宿主机和一个或多个主要存储服务器构成。

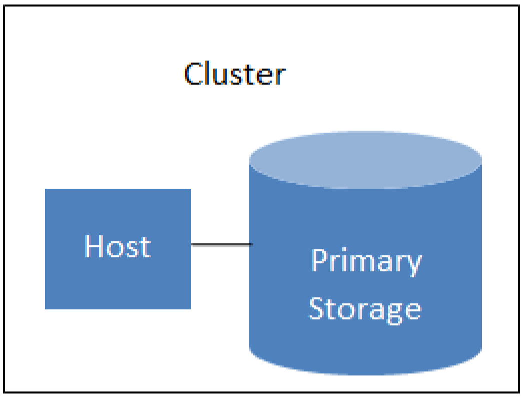

A simple cluster

本产品允许在云部署中有多个集群

Even when local storage is used exclusively, clusters are still required organizationally, even if there is just one host per cluster.

当使用VMware时，每个VMware集群都被vCenter 服务器管理。管理员必须在本产品中登记vCenter。每个zone下可以有多个vCenter服务器。每个vCenter服务器可能管理多个VMware集群。

2.5. 关于宿主机

宿主机就是个独立的计算机。宿主机运行来宾虚拟机并提供其相应的计算资源。每个宿主机都装有虚拟机软件来运行来宾虚拟机。比如一个开启了kvm支持的服务器，一个思杰XenServer服务器，或者一个ESXi服务器都可以作为宿主机。

宿主机在CloudStack部署中属于最小的组织单元。宿主机包含于集群中，集群有属于提供点，而区域中包含提供点（就是在逻辑概念上zone>pod>cluster>host） 。

CloudStack部署中的宿主机:

- Provide the CPU, memory, storage, and networking resources needed to host the virtual machines
- 通过高带宽TCP/IP网络并连接到因特网
- 可能在不同地理位置有多个数据中中心。
- 虽说包含在集群中的宿主机必须是同质的（使用相同的虚拟机软件）但是他们可以具有不同的计算能力（不同的CPU速度，不同的内存数量等等）

新增的宿主机可以随时添加以提供更多资源给来宾虚拟机

CloudStack自动探测宿主机的cpu数量和内存资源。

宿主机对终端用户不可见。终端用户不能决定他们的虚拟机被分配到哪台宿主机。

如果您想让宿主机在CloudStack上正常运行，你必须作如下步骤:

- 在宿主机上安装虚拟机软件
- 为宿主机分配IP（固定IP）
- 确定宿主机已经连接到CloudStack的管理服务器

2.6. 关于主存储

主存储是和群集有关联的，它为所有在那个群集里运行在主机上的虚拟机储存磁盘卷。你能添加多个主存储服务器给群集。至少一个是必须的。为了提高性能它的位置最好是接近主机放置。

CloudStack被设计和标准iSCSI或者NFS 服务器一起工作，这些被底层的虚拟机平台支持，包括，例如：

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

如果你打算使用本地磁盘当你安装的时候，你可以跳过去安装辅助存储。

2.7. 关与辅助存储

辅助存储是和区域关联，它存储如下事物：

- 模板— 操作系统镜像能够用来启动虚拟机并且可一包含额外配置信息，例如被按装的应用程序。
- ISO 镜像—磁盘镜像包含数据或者操作系统引导媒体。
- 磁盘卷快照—被保存的虚拟机数据复制品能够被用来做数据恢复或者建立新的模板。

这些在基于区域的NFS辅助存储中的数据是对所有在这个区域内的主机有效的。

为了让这些在辅助存储中的数据对所有在云中的主机有效，你可以另外添加OpenStack对象存储(Swift，swift.openstack.org[1]) 给基于区域的NFS辅助存储。当使用Swift时，你配置Swift存储给整个CloudStack，然后照常设置NFS 辅助存储给每个区域。在每个区域的NFS存储扮演了一个代转区，所有的模板和其他辅助存储的数据在转向Swift前将通过它。Swift 存储扮演了一个广泛云的资源,使得模板和其它数据是有效的对任何在云中的区域。Swift 存储中没有分级，每个存储对象只有一个Swift容器。任何在整个云中的辅助存储能够拖一个容器从Swift中在需要的时候。这样就不用拷贝模板和快照从一个区域到另外一个区域了，如果单独使用区域NFS的话，还是需要的。任何事情都是有效的在任何地方。

2.8. 关于物理网络

设置物理网络是添加区域步骤中的一个部分。每个区域可以分配一个（极限与高级资源域）或者多个物理网络。这个网络对应宿主机的一个网卡。每个物理网络可以承载一种或多种网络流量。每个网络流量的类型的选项由你选择的是基本网络与还是高级网络与而不同。

物理网络是连接到资源与的真实网络硬件。一个资源与能有多个物理网络。管理员能做一下操作：

- 添加/删除/更新 域中的物理网络
- 在物理网络上设置VLAN
- 通过设置名字使网络能被虚拟机软件识别
- 设置在物理网络上能够提供的服务（防火墙，负载均衡器，等等）
- 设置能直连到物理网络的IP地址

[1] http://swift.openstack.org

- 指定物理网络承载的流量类型还有其他类似网络速度之类的属性

2.8.1. 基本区域网络流量类型

当使用基本网络是，这里只有一个物理网卡在区域中。这个物理网卡承载以下类型流量：

- 来宾。当终端用户运行虚拟，他们产生来宾流量。来宾虚拟机和其他虚拟机通信在网络上的流量，归功来宾网络。每一个pod中的基本区域就是一个广播域，因此每个一个pod中的来宾网络拥有不同的ip范围。管理员必须为每一个pod配置ip范围。
- 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.

> **注意**
>
> 我们强烈要求管理和来宾流量使用独立的网卡

- 公共。云中的虚拟机访问internet时产生公共流量，基于这个原因必须分配可供访问的ip地址。终端用户可以使用CloudStack UI获得一个ip，用来构建来宾网络和公共网络的nat。定义为：获取新的ip地址。
- Storage. While labeled "storage" this is specifically about secondary storage, and doesn't affect traffic for primary storage. This includes 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.

在基本区域中，配置物理网络是相当简单的。在大多数情况下，你只需要配置一个来宾网络承载所有来宾虚拟机流量。如果你使用NetScaler负载平衡器，开启弹性的IP和弹性负载平衡(EIP和ELB)的功能，你还必须配置一个网络承载公共流量。 当你通过UI添加一个新的区域， CloudStack负责提出必要的网络配置的步骤。

2.8.2. 基本区域宾客 IP 地址

当基本联网方式被使用， CloudStack 将在POD CIDR的IP 地址分配给该POD中的宾客。 管理员必须在此POD增加一个直接IP 范围用于此目的。 这些IP 和主机位于同样的VLAN。

2.8.3. 高级区域网络流量类型

当使用高级网路时，在区域包含多种物理网络。每一个物理网络能够承载一种或者多种类型的流量。要让 CloudStack 知道每种网络承载那种类型的流量。高级区域包含的流量类型：

- 来宾。当用户允许VM时，他们产生来宾流量。来宾VM通过来宾网络进相关通讯。这个网络可以是隔离或者共享的；在隔离的来宾网络中，管理员需要为每一个 CloudStack 中隔离网络分配 VLAN 范围；帐户的网络(潜在的大量的VLAN)。在一个共享来宾网络中，所有来宾VM共享一个网络。

- 管理。当 CloudStack最为内部资源和其他通信时，他们产生管理流量。包括主机，系统vm(在云中，被用于CloudStack 执行大量任务的虚拟机)之间的通信， 其他组件和CloudStack 管理服务器的直接通信。你必须为系统vm配置一个ip范围。

- 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” in the Administration Guide.

- Storage. While labeled "storage" this is specifically about secondary storage, and doesn't affect traffic for primary storage. This includes 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.

2.8.4. 高级区域宾客的IP地址

使用高级的网络时，管理员可以创建额外的网络供客人使用。这些网络可以跨越区域，并提供给所有帐户，或者他们可以到一个单一的帐户范围内，在这种情况下，只有指定的帐户可以创建连接到这些网络的宾客。网络被定义为一个VLAN ID，IP范围和网关。如果需要的话，系统管理员可能会提供成千上万的网络。

2.8.5. Advanced Zone Public IP Addresses

使用高级的网络时，管理员可以创建额外的网络供客人使用。这些网络可以跨越区域，并提供给所有帐户，或者他们可以到一个单一的帐户范围内，在这种情况下，只有指定的帐户可以创建连接到这些网络的宾客。网络被定义为一个VLAN ID，IP范围和网关。如果需要的话，系统管理员可能会提供成千上万的网络。

2.8.6. 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.

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 zones 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 both 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.

第 3

帐号

3.1. 账户，用户，域

å¸#æ#·

一个账号通常代表一个客户的服务提供者或一个大组织中的一个部门。一个账户下可有多个用户。

å##

域通常包含多个账户。域中经常包含多个账户，这些账户间存在一些逻辑上关系和一系列该域和其子域下的委派的管理员（这段的意思就是说在逻辑上域下可以有管理员，子域下也可以有管理员）。比如，一个服务提供商可有多个分销商这样的服务提供商就能为每一个分销商创建一个域

对于每个账户的创建，CloudStack 创建了三种不同类型的用户账户：根管理员，域管理员，普通用户。

普通用户

用户就像是账户的别名。在同一账户下的用户彼此之间并非隔离的。但是他们与不同账户下的用户是相互隔离的。大多数安装不需要表面的用户的概念；他们只是每一个帐户的用户。相同的用户不能属于多个帐户。

多个账户中的用户名在域中应该是唯一的。相同的用户名能在其他的域中存在，包括子域。域名只有在全路径名唯一的时候才能重复。比如，你能创建一个root/d1，也可以创建root/foo/d1和root/sales/d1。

管理员在系统中是拥有特权的账户。可能有多个管理员在系统中，管理员能创建删除其他管理员，并且修改系统中任意用户的密码。

域管理员

域管理员能在它所属域中执行和管理员相同的操作。域管理员在物理服务器或其他域中不可见。

根管理员

根管理员拥有系统完全访问权限，包括管理模板，服务方案，客户服务管理员和域。

资源属于帐户而不是个人用户。例如,帐单、资源限制,等等都是由该帐户维护,而不是用户维护。一个用户可以操作任何资源帐户提供了用户权限的操作。特权由角色决定。

3.2. 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

3.2.1. LDAP配置命令样例

为了明白本章节的样例，你需要知道调用CloudStack API的基本概念，这在开发者文档中有讲解.

以下展示了通过ApacheDS LDAP服务器调用 ldapConfig命令的样例

```
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%2Cou
%3Dtesting%2Co%project&bindpass=secret&port=10389&ssl=true&truststore=C%3A%2Fcompany%2Finfo
%2Ftrusted.ks&truststorepass=secret&response=json&apiKey=YourAPIKey&signature=YourSignatureHash
```

命令调用的URL必须进行编码. 这里是一个例子没有进行URL编码:

```
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
```

以下展示了与活动目录类似的命令. 这里，搜索是基于一个公司的测试组，用户是基于邮件地址进行匹配查找.

```
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
```

接下来的几个章节解释了你需要知道的通过ldapConfig参数进行过滤的一些概念.

3.2.2. 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

LDAP Server	Example Search Base DN
Active Directory	OU=testing, DC=company

3.2.3. 查询过滤器

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:

查询过滤器通配符	è¯´æ##
%u	用户名
%e	邮件地址
%n	姓名

下面的例子假定你使用活动目录，并从活动目录的概要里获得用户属性.

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)
```

利用邮件地址来查找一个用户:

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

3.2.4. 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 DNs 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

3.2.5. SSL Keystore路径和密码

如果LDAP 服务器要求SSL， 你需要在ldapConfig命令中通过设置参数ssl，truststore和truststorepass使其生效。在使SSL 对ldapConfig 生效之前，你需要得到LDAP服务器在使用的证书并把它加到被信任的键存储中。你将需要知道到键存储和密码的路径。

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.

4.1. 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.

- Disk Offerings, defined by the CloudStack administrator, provide a choice of disk size for primary data storage. See Creating a New Disk Offering.

- 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.

- 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.

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.

第 5

用户界面

5.1. 登陆到用户界面

CloudStack提供一个基于web的用户界面能够被管理员和终端用户使用。适当的用户界面版本被展现依赖于登陆时使用的凭证。用户界面是适用于大多数流行的浏览器包括IE7,IE8,IE9,Firefox3.5+,Firefox4,Safari4,和Safari5。URL是:(用你自己的管理控制服务器IP地址代替)

```
http://<管理控制-服务器-ip-地址>:8080/client
```

初次登陆管理控制服务器时，一个向导启动画面会显现。你将看到登陆界面当你执行下面的过程在你的控制面板上:

用户名

你的帐户的用户标识。默认的用户名是admin。

密码

相关用户标识的密码。默认用户（admin）的密码是password。

域

如果你是一个root用户，不用填写域这个字段。

如果你是一个子域的用户，输入完全路径在域字段，不包括根域。

例如，假设多个层级被建立在根域下，像Comp1/hr，在Comp1域的用户应该输入Comp1在域字段，在Comp1/sales域的用户应该输入Comp1/sales。

更多关于当你登陆这个界面时选项的指导，参照作为根管理员登陆。

5.1.1. 最终用户界面概览

CloudStack 用户界面帮助云基础设施的用户查看和使用他们的云资源，包括虚拟机、模板和ISO、数据卷和快照、宾客网络，以及IP 地址。如果用户是一个或多个CloudStack 项目的成员或管理员，用户界面能提供一个面向项目的视图。

5.1.2. 根管理员界面的概述

CloudStack 界面帮助 CloudStack 管理员配置，查看和管理云的基础设施，用户域，账号，项目，参数设置。当一个全新的管理服务器安装完成后，第一次启动界面的时候，可以选择根随引导步骤配置云的基础设施。当再次登录时，会显示当前登录用户的仪表板。在这个页面有很多的连接，可以通过左边的导航栏访问各种管理功能。根管理员也可以使用界面像最终用户一样来执行所有的功能。

5.1.3. 作为根管理员登录

在管理服务器软件安装并且运行后，你可以运行 CloudStack 的用户界面. 在这里通过UI,可以供给，查看并管理你的云基础架构.

1. 打开你自己喜欢的浏览器并访问这个URL. 请把IP地址替换成你自己的管理服务器的IP.

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

After logging into a fresh Management Server installation, a guided tour splash screen appears. On later visits, you’ll be taken directly into the Dashboard.

2. 如果你看到第一次的向导屏幕，可以选择下面步骤之一进行.

   - 继续简单设置. 如果你只是简单试用CloudStack,并且你想通过一个配置向导尽可能简单快速的开始，请选择这项. 我们将帮助你建立一个有下列功能的云环境: 一个单独的机器运行CloudStack并通过NFS提供存储; 一个单独的机器提供在XenServer或KVM上运行虚拟机; 以及一个共享的公共网络.

     安装向导的提示应该给你需要的所有信息. 但如果你需要更多的详细信息，你可以按照试用安装向导进行.

   - I have used CloudStack before. Choose this if you have already gone through a design phase and planned a more sophisticated 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.

     根管理员的仪表盘出现了.

3. 你应该为根管理员设置一个新的密码. 如果你选择简单设置，将会提示你立即创建一个新的密码. 如果你选择有经验的用户，请选择第 5.1.4 节 “修改Root口令” 里的步骤.

警告

你正作为根管理员登入. 这个账号管理@PRODUCT;的部署，包括物理架构. 根管理员可以更改配置以更改基本的功能，创建或删除用户账号，以及其它许多只有被鉴权的用户执行的操作. 请更改默认的密码,确保其唯一性和安全性.

5.1.4. 修改Root口令

During installation and ongoing cloud administration, you will need to log in to the UI as the root administrator. The root administrator 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. When first installing CloudStack, be sure to change the default password to a new, unique value.

1. 打开你自己喜欢的浏览器并访问这个URL. 请把IP地址替换成你自己的管理服务器的IP.

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

2. 使用当前root用户的ID和口令登录UI。缺省为admin/pawword。

3. 点击账户

4. 点击管理员账户名

5. 点击查看用户

6. 点击管理员用户名

7. Click the Change Password button. 🔑

8. 键入新密码，然后点击确认

5.2. Using SSH Keys for Authentication

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

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

5.2.1. Creating an Instance Template that Supports SSH Keys

Create a instance template that supports SSH Keys.

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

   For more information on creating a new instance, see

2. Download the cloudstack script from The SSH Key Gen Script[1]to the instance you have created.

   ```
   wget http://downloads.sourceforge.net/project/cloudstack/SSH%20Key%20Gen%20Script/cloud-set-guest-sshkey.in?r=http%3A%2F%2Fsourceforge.net%2Fprojects%2Fcloudstack%2Ffiles%2FSSH%2520Key%2520Gen%2520Script%2F&ts=1331225219&use_mirror=iweb
   ```

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

   ```
   cp cloud-set-guest-sshkey.in /etc/init.d/
   ```

4. Give the necessary permissions on the script:

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

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

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

6. Stop the instance.

[1] http://sourceforge.net/projects/cloudstack/files/SSH%20Key%20Gen%20Script/

5.2.2. Creating the SSH Keypair

You must make a call to the createSSHKeyPair api method. You can either use the CloudStack Python API library or the curl commands to make the call to the cloudstack api.

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

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

1. Run the following curl command:

```
curl --globoff "http://localhost:8096/?command=createSSHKeyPair&name=keypair-
doc&account=admin&domainid=5163440e-c44b-42b5-9109-ad75cae8e8a2"
```

   The output is something similar to what is given below:

```
<?xml version="1.0" encoding="ISO-8859-1"?><createsshkeypairresponse
 cloud-stack-version="3.0.0.20120228045507"><keypair><name>keypair-doc</
name><fingerprint>f6:77:39:d5:5e:77:02:22:6a:d8:7f:ce:ab:cd:b3:56</fingerprint><privatekey>-----BEGIN RSA
 PRIVATE KEY-----
MIICXQIBAAKBgQCSydmnQ67jP61NoXdX3noZjQdrMAWNQZ7y5SrEu4wDxplvhYci
dXYBeZVwakDVsU2MLG1/K+wefwefwefwefwefJyKJaogMKn7BperPD6n1wIDAQAB
AoGAdXaJ7uyZKeRDoy6wA0UmF0kSPbMZCR+UTIHNkS/E0/4U+6lhMokmFSHtu
mfDZ1kGGDYhMsdytjDBztljawfawfeawefawfawfawQQDCjEsoRdgkduTy
QpbSGDIa11Jsc+XNDx2fgRinDsxXI/zJYXTKRhS1/LIPHBw/brW8vzxhOlSOrwm7
VvemkkgpAkEAwSeEw394LYZiEVv395ar9MLRVTVLwpo54jC4tsOxQCBlloocK
lYaocpk0yBqqOUSBawfIiDCuLXSdvBo1Xz5ICTM19vgvEp/+kMuECQBzm
nVo8b2Gvyagqt/KEQo8wzH2THghZ1qQ1QRhIeJG2aissEacF6bGB2oZ7Igim5L14
4KR7OeEToyCLC2k+02UCQQCrniSnWKtDVoVqeK/zbB32JhW3Wu11v5p5zUEcd
KfEEuzcCUIxtJYTahJ1pv1FkQ8anpuxjSEDp8x/18bq3
-----END RSA PRIVATE KEY-----
</privatekey></keypair></createsshkeypairresponse>
```

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

```
-----BEGIN RSA PRIVATE KEY-----
MIICXQIBAAKBgQCSydmnQ67jP61NoXdX3noZjQdrMAWNQZ7y5SrEu4wDxplvhYci
dXYBeZVwakDVsU2MLG1/K+wefwefwefwefwefJyKJaogMKn7BperPD6n1wIDAQAB
AoGAdXaJ7uyZKeRDoy6wA0UmF0kSPbMZCR+UTIHNkS/E0/4U+6lhMokmFSHtu
mfDZ1kGGDYhMsdytjDBztljawfawfeawefawfawfawQQDCjEsoRdgkduTy
QpbSGDIa11Jsc+XNDx2fgRinDsxXI/zJYXTKRhS1/LIPHBw/brW8vzxhOlSOrwm7
VvemkkgpAkEAwSeEw394LYZiEVv395ar9MLRVTVLwpo54jC4tsOxQCBlloocK
lYaocpk0yBqqOUSBawfIiDCuLXSdvBo1Xz5ICTM19vgvEp/+kMuECQBzm
nVo8b2Gvyagqt/KEQo8wzH2THghZ1qQ1QRhIeJG2aissEacF6bGB2oZ7Igim5L14
4KR7OeEToyCLC2k+02UCQQCrniSnWKtDVoVqeK/zbB32JhW3Wu11v5p5zUEcd
KfEEuzcCUIxtJYTahJ1pv1FkQ8anpuxjSEDp8x/18bq3
-----END RSA PRIVATE KEY-----
```

3. Save the file.

5.2.3. Creating an Instance

After you save the SSH keypair file, you must create an instance by using the template that you created at 第 5.2.1 节 “ Creating an Instance Template that Supports SSH Keys” . Ensure that you use the same SSH key name that you created at 第 5.2.2 节 “Creating the SSH Keypair” .

注意

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

A sample curl command to create a new instance is:

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

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

5.2.4. Logging In Using the SSH Keypair

To test your SSH key generation is successful, check whether you can log in to the cloud setup.

For exaple, from a Linux OS, run:

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

The -i parameter tells the ssh client to use a ssh key found at ~/.ssh/keypair-doc.

5.2.5. Resetting SSH Keys

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

使用项目来管理用户和资源。

6.1. Overview of Projects

Projects are used to organize people 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..

6.2. 配置项目

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.

6.2.1. 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. 以管理员权限登录CloudStack 界面

2. In the left navigation, click Global Settings.

3. In the search box, type project and click the search button. ✎

4. In the search results, you can see a few other parameters you need to set to control how invitations behave. The table below shows global configuration parameters related to project invitations. Click the edit button to set each parameter.

Configuration Parameters	è¯´æ##
project.invite.required	Set to true to turn on the invitations feature.
project.email.sender	The email address to show in the From field of invitation emails.
project.invite.timeout	Amount of time to allow for a new member to respond to the invitation.
project.smtp.host	Name of the host that acts as an email server to handle invitations.
project.smtp.password	(Optional) Password required by the SMTP server. You must also set project.smtp.username and set project.smtp.useAuth to true.
project.smtp.port	SMTP server’s listening port.
project.smtp.useAuth	Set to true if the SMTP server requires a username and password.
project.smtp.username	(Optional) User name required by the SMTP server for authentication. You must also set project.smtp.password and set project.smtp.useAuth to true..

5. Restart the Management Server:

```
service cloudstack-management restart
```

6.2.2. 设置项目的资源限制

CloudStack管理员可以设置全局默认的限制，控制在云中的每个项目可以拥有的资源量。这是为了防止不受控制的资源使用，如快照，IP地址和虚拟机实例。域管理员可以取代他们域中个别项目的资源限制，只要新的限制是低于CloudStack根管理员设置的默认全局限制。根管理员还可以为云中的任何项目设置较低的资源限制。

6.2.2.1. 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 第 6.2.2 节 “设置项目的资源限制”). 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. 以管理员权限登录CloudStack 界面
2. 在左边导航栏，点击项目.
3. 在下拉列表框选择项目
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.

6.2.2.2. Setting the Global Project Resource Limits

1. 以管理员权限登录CloudStack 界面
2. In the left navigation, click Global Settings.
3. In the search box, type max.projects 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.
max.project.snapshots	Maximum number of snapshots that can be owned by any project in the cloud. See Working with Snapshots.
max.project.templates	Maximum number of templates that can be owned by any project in the cloud. See Working with Templates.
max.project.uservms	Maximum number of guest virtual machines that can be owned by any project in the cloud. See Working With Virtual Machines.
max.project.volumes	Maximum number of data volumes that can be owned by any project in the cloud. See Working with Volumes.

5. 重启管理服务器.

```
# service cloudstack-management restart
```

6.2.3. 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. 以管理员权限登录CloudStack 界面

2. In the left navigation, click Global Settings.

3. In the search box, type allow.user.create.projects.

4. Click the edit button to set the parameter. ✎

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

5. 重启管理服务器.

```
# service cloudstack-management restart
```

6.3. 创建一个新项目

CloudStack （默认）系统管理员和与管理员能创建项目。如果全局变量allow.user.create.projects 设置为true，终端用户也能创建项目。

1. 以管理员权限登录CloudStack 界面

2. 在左边导航栏点击项目。

3. 在选择视图点击项目。

4. 点击新建项目。

5. 为项目命名并添加描述，然后点击创建项目。

6. 会出现一个界面，你可以很快的添加更多成员到项目中，此步可选。当你准备好继续，点击下一步。

7. 点击保存。

6.4. 添加成员到一个项目

新的成员可以由，项目所在的域或任何父域的域管理员，或CloudStack的root管理员，项目的管理员添加到项目中。在CloudStack 中，有两种添加成员的方式，但在一个时刻只有一种可用：

- 如果邀请被设置，你可以给新成员发送邀请。

- 如果邀请没有被设置，你可以通过UI直接添加成员。

6.4.1. Sending Project Membership Invitations

Use these steps to add a new member to a project if the invitations feature is enabled in the cloud as described in 第 6.2.1 节 "Setting Up Invitations" . If the invitations feature is not turned on, use the procedure in Adding Project Members From the UI.

1. 登入到CloudStack UI.

2. 在左边导航栏，点击项目.

3. 在下拉列表框选择项目

4. Click the name of the project you want to work with.

5. Click the Invitations tab.

6. In Add by, select one of the following:

   a. Account — The invitation will appear in the user's Invitations tab in the Project View. See Using the Project View.

   b. Email — The invitation will be sent to the user's email address. Each emailed invitation includes a unique code called a token which the recipient will provide back to CloudStack when accepting the invitation. Email invitations will work only if the global parameters related to the SMTP server have been set. See 第 6.2.1 节 "Setting Up Invitations" .

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.

6.4.2. Adding Project Members From the UI

The steps below tell how to add a new member to a project if the invitations feature is not enabled in the cloud. If the invitations feature is enabled cloud,as described in 第 6.2.1 节 "Setting Up Invitations" , use the procedure in 第 6.4.1 节 "Sending Project Membership Invitations" .

1. 登入到CloudStack UI.

2. 在左边导航栏，点击项目.

3. 在下拉列表框选择项目

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.

6.5. 接受成员身份的邀请

如果你收到一个加入CloudStack项目的邀请，并且你想接受这个邀请，参照下面的步骤:

1. 登入到CloudStack UI.
2. 在左边导航栏，点击项目.
3. 在选择视图中，选择邀请
4. 如果你看到邀请列在屏幕上，点击接受按钮.

   屏幕上列出的邀请是使用你在CloudStack账号名称发给你的.
5. 如果你收到一个邮件邀请，点击权标按钮，然后提供从邮件中获得的项目ID和唯一识别码(权标).

6.6. Suspending or Deleting a Project

When a project is suspended, it retains the resources it owns, but they can no longer be used. No new resources or members can be added to a suspended project.

When a project is deleted, its resources are destroyed, and member accounts are removed from the project. The project’s status is shown as Disabled pending final deletion.

A project can be suspended or deleted by the project administrator, the domain administrator of the domain the project belongs to or of its parent domain, or the CloudStack root administrator.

1. 登入到CloudStack UI.
2. 在左边导航栏，点击项目.
3. 在下拉列表框选择项目
4. Click the name of the project.
5. Click one of the buttons:

   To delete, use 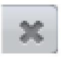

   To suspend, use 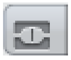

6.7. Using the Project View

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. 登入到CloudStack UI.
2. Click Project View.
3. The project dashboard appears, showing the project’s VMs, volumes, users, events, network settings, and more. From the dashboard, you can:

- Click the Accounts tab to view and manage project members. If you are the project administrator, you can add new members, remove members, or change the role of a member from user to admin. Only one member at a time can have the admin role, so if you set another user’s role to admin, your role will change to regular user.
- (If invitations are enabled) Click the Invitations tab to view and manage invitations that have been sent to new project members but not yet accepted. Pending invitations will remain in this list until the new member accepts, the invitation timeout is reached, or you cancel the invitation.

准备你的云基础设施的步骤

This section tells how to add regions, zones, pods, clusters, hosts, storage, and networks to your cloud. If you are unfamiliar with these entities, please begin by looking through 第 2 章 云基础设施概念.

7.1. 设置步骤概览

管理服务器节点安装并运行之后，你就可以添加计算资源来进行管理了. 要查看CloudStack云架构整体上是如何组织的，请参考 第 1.3.2 节 “Cloud Infrastructure Overview” .

为了提供云基础架构，或者在任何时个需要扩展过规模，请按照下面的步骤进行:

1. Define regions (optional). See 第 7.2 节 “Adding Regions (optional)” .

2. Add a zone to the region. See 第 7.3 节 “创建Zone” .

3. Add more pods to the zone (optional). See 第 7.4 节 “添加一个机架” .

4. Add more clusters to the pod (optional). See 第 7.5 节 “添加集群” .

5. Add more hosts to the cluster (optional). See 第 7.6 节 “Adding a Host” .

6. Add primary storage to the cluster. See 第 7.7 节 “æ·»å# ä¸»å-#å#¨” .

7. Add secondary storage to the zone. See 第 7.8 节 “æ·»å# è¾å#©å-#å#¨” .

8. 初始化并测试新的云环境. 参照第 7.9 节 “初始化和测试” .

当你完成这些步骤以后，你将部署如下的一个基本结构:

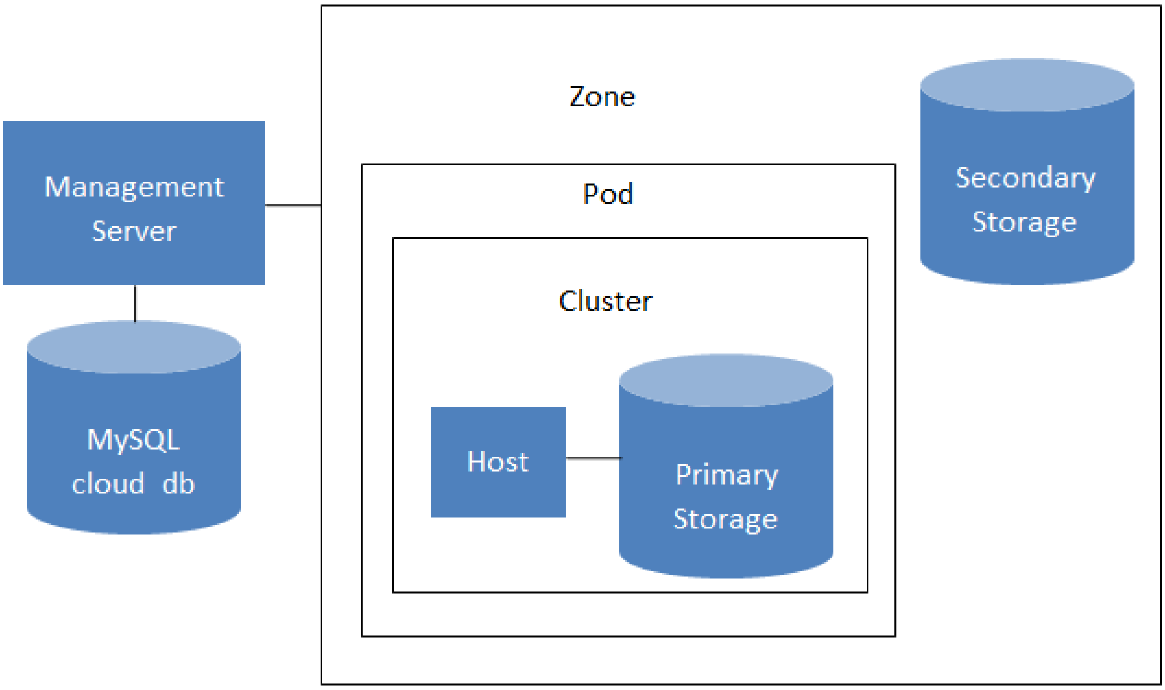

Conceptual view of a basic deployment

7.2. Adding Regions (optional)

Grouping your cloud resources into geographic regions is an optional step when provisioning the cloud. For an overview of regions, see 第 2.1 节 “About Regions” .

7.2.1. The First Region: The Default Region

If you do not take action to define regions, then all the zones in your cloud will be automatically grouped into a single default region. This region is assigned the region ID of 1.

You can change the name or URL of the default region by using the API command updateRegion. For example:

```
http://<IP_of_Management_Server>:8080/client/api?command=updateRegion&id=1&name=Northern&endpoint=http://
<region_1_IP_address_here>:8080/client&apiKey=miVr6X7u6bN_sdahOBpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU%2FcaiK8RAP0O1hU%3D
```

7.2.2. Adding a Region

Use these steps to add a second region in addition to the default region.

1. Each region has its own CloudStack instance. Therefore, the first step of creating a new region is to install the Management Server software, on one or more nodes, in the geographic area where you want to set up the new region. Use the steps in the Installation guide. When you come to the step where you set up the database, use the additional command-line flag -r <region_id> to set a region ID for the new region. The default region is automatically assigned a region ID of 1, so your first additional region might be region 2.

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

2. By the end of the installation procedure, the Management Server should have been started. Be sure that the Management Server installation was successful and complete.

3. Add region 2 to region 1. Use the API command addRegion. (For information about how to make an API call, see the Developer's Guide.)

   ```
   http://<IP_of_region_1_Management_Server>:8080/client/api?
   command=addRegion&id=2&name=Western&endpoint=http://<region_2_IP_address_here>:8080/
   client&apiKey=miVr6X7u6bN_sdahOBpjNejPgEsT35eXq-
   jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU%2FcaiK8RAP0O1hU%3D
   ```

4. Now perform the same command in reverse, adding region 1 to region 2.

   ```
   http://<IP_of_region_2_Management_Server>:8080/client/api?
   command=addRegion&id=1&name=Northern&endpoint=http://<region_1_IP_address_here>:8080/
   client&apiKey=miVr6X7u6bN_sdahOBpjNejPgEsT35eXq-
   jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU%2FcaiK8RAP0O1hU%3D
   ```

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

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

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

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

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

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

6. Remove project accounts. Run these commands on the region 2 database:

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

7. Set the default zone as null:

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

8. Restart the Management Servers in region 2.

7.2.3. Adding Third and Subsequent Regions

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

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

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

2. Once the Management Server is running, add your new region to all existing regions by repeatedly calling the API command addRegion. For example, if you were adding region 3:

```
http://<IP_of_region_1_Management_Server>:8080/client/api?
command=addRegion&id=3&name=Eastern&endpoint=http://<region_3_IP_address_here>:8080/
client&apiKey=miVr6X7u6bN_sdahOBpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU%2FcaiK8RAP0O1hU%3D

http://<IP_of_region_2_Management_Server>:8080/client/api?
command=addRegion&id=3&name=Eastern&endpoint=http://<region_3_IP_address_here>:8080/
client&apiKey=miVr6X7u6bN_sdahOBpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU%2FcaiK8RAP0O1hU%3D
```

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

```
http://<IP_of_region_3_Management_Server>:8080/client/api?
command=addRegion&id=1&name=Northern&endpoint=http://<region_1_IP_address_here>:8080/
```

```
client&apiKey=miVr6X7u6bN_sdahOBpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU%2FcaiK8RAP0O1hU%3D

http://<IP_of_region_3_Management_Server>:8080/client/api?
command=addRegion&id=2&name=Western&endpoint=http://<region_2_IP_address_here>:8080/
client&apiKey=miVr6X7u6bN_sdahOBpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU%2FcaiK8RAP0O1hU%3D
```

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

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

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

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

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

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

5. Remove project accounts. Run these commands on the region 2 database:

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

6. Set the default zone as null:

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

7. Restart the Management Servers in the new region.

7.2.4. Deleting a Region

To delete a region, use the API command removeRegion. Repeat the call to remove the region from all other regions. For example, to remove the 3rd region in a three-region cloud:

```
http://<IP_of_region_1_Management_Server>:8080/client/api?
command=removeRegion&id=3&apiKey=miVr6X7u6bN_sdahOBpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU%2FcaiK8RAP0O1hU%3D

http://<IP_of_region_2_Management_Server>:8080/client/api?
command=removeRegion&id=3&apiKey=miVr6X7u6bN_sdahOBpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU%2FcaiK8RAP0O1hU%3D
```

7.3. 创建Zone

以上安装步骤如果全部完成，登录WEB UI

1. （可选）

   a. 以管理员身份登录进入CloudStack 用户界面。

b. If this is your first time visiting the UI, you will see the guided tour splash screen. Choose “Experienced user.” The Dashboard appears.

c. 在左侧导航栏，点击 全局设置

d. In the search box, type swift.enable and click the search button.

e. Click the edit button and set swift.enable to true.

f. 重启管理服务器.

```
# service cloudstack-management restart
```

g. Refresh the CloudStack UI browser tab and log back in.

2. In the left navigation, choose Infrastructure.

3. On Zones, click View More.

4. (Optional) If you are using Swift storage, click Enable Swift. Provide the following:

   - URL. The Swift URL.

   - Account. The Swift account.

   - Username. The Swift account’s username.

   - Key. The Swift key.

5. Click Add Zone. The zone creation wizard will appear.

6. Choose one of the following network types:

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

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

   For more information about the network types, see 第 2.8 节 “关于物理网络”.

7. The rest of the steps differ depending on whether you chose Basic or Advanced. Continue with the steps that apply to you:

   - 第 7.3.1 节 “基础区域配置”

   - 第 7.3.2 节 “高级资源域配置”

7.3.1. 基础区域配置

1. 你在添加区域向导中选择 ” 基础 “后，点击下一步，你将被询问输入以下细节，接着点击 下一步

   - 名字，区域名字

- dns 1和2 ， 区域中来宾虚拟机的dns服务器，通过你后面添加的公共网络访问dns服务器。区域中的公共ip地址必须有通向已定义dns服务器的路由。

- 内部dns1和内部dns2. 这些dns是被区域中系统vm(这些CloudStack 虚拟机主机：它自己。例如虚拟路由器，console代理和辅助存储虚拟机)使用的。\n系统虚拟机通过管理流量网络接口访问这些dns服务器。pod私有地址必须有通向已定义dns服务器的路由

- hypersior(3.0.1版本中以介绍).选择区域中第一个集群的虚拟化方案。在你完成区域添加后，你可以添加使用不同虚拟化方案的集群。

- 网络方案。 你的选择决定了来宾虚拟机可以使用的网络服务。

ç½#ç»#æ#¹æ¡#	è¯´æ##
DefaultSharedNetworkOfferingWithSGService	如果你打算使用安全组进行来宾流量隔离，选择：（参考 使用安全组控制虚拟机流量）
DefaultSharedNetworkOffering	如果你不需要安全组，选择：
DefaultSharedNetscalerEIPandELBNetworkOffering	你如果在你区域网络中安装了 Citrix NetScaler appliance ， 你打算使用它的弹性ip和弹性负载特性，就选择它。通过EIP and ELB特性，区域中的安全组可以提供1:1静态NAT和负载。

- 网络域。

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

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

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

3. 3. (Introduced in version 3.0.1) Assign a network traffic label to each traffic type on the physical network. These labels must match the labels you have already defined on the hypervisor host. To assign each label, click the Edit button under the traffic type icon. A popup dialog appears where you can type the label, then click OK.

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

4. Click Next.

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

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

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

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

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

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

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

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

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

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

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

   - Netmask. The netmask associated with this IP range.

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

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

7. In a new zone, CloudStack adds the first pod for you. You can always add more pods later. For an overview of what a pod is, see 第 2.3 节 “关于POD”.

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

   - Pod Name. A name for the pod.

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

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

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

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

   - Guest gateway. The gateway that the guests should use.

   - Guest netmask. The netmask in use on the subnet the guests will use.

   - Guest start IP/End IP. Enter the first and last IP addresses that define a range that 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.

9. In a new pod, CloudStack adds the first cluster for you. You can always add more clusters later. For an overview of what a cluster is, see About Clusters.

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

   - Hypervisor. (Version 3.0.0 only; in 3.0.1, this field is read only) Choose the type of hypervisor software that all hosts in this cluster will run. If you choose VMware, additional fields appear so you can give information about a vSphere cluster. For vSphere servers, we recommend creating the cluster of hosts in vCenter and then adding the entire cluster to CloudStack. See Add Cluster: vSphere.

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

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

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

    Before you can configure the host, you need to install the hypervisor software on the host. You will need to know which version of the hypervisor software version is supported by CloudStack and what additional configuration is required to ensure the host will work with CloudStack. To find these installation details, see:

    - Citrix XenServer Installation and Configuration

    - VMware vSphere 安装和配置

    - KVM vSphere Installation and Configuration

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

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

    - Username. The username is root.

    - Password. This is the password for the user named above (from your XenServer or KVM install).

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

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

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

    - Name. The name of the storage device.

    - Protocol. For XenServer, choose either NFS, iSCSI, or PreSetup. For KVM, choose NFS, SharedMountPoint,CLVM, or RBD. For vSphere choose either VMFS (iSCSI or FiberChannel) or NFS. The remaining fields in the screen vary depending on what you choose here.

7.3.2. 高级资源域配置

1. 在添加域向导中选择了高级并且点击下一步之后，你会被要求输入下列信息。然后点击下一步。

    - 名称. 一个区域的名称。

    - DNS 1 和 2. 这些DNS服务器是给在域中的客户虚拟机使用的。这些DNS服务器可以通过稍后添加的公共网络访问。这个域的公共IP 地址必须有一个路由到在这里指定的DNS服务器。

    - 内部 DNS 1 and 内部 DNS 2. 这些DNS 服务器给域中的系统虚拟机使用。(这些系统虚拟机是CloudStack自己使用的，例如虚拟路由，控制代理和辅助存储虚拟机。)这些DNS服务器可以通过系统虚拟机管理网络接口访问。你提供给机架的私有IP地址必须有一个路由到在这里指定的内部DNS服务器。

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

    - 来宾 CIDR. CIDR 描述了使用在来宾虚拟网络的IP地址在这个区域中。例如，10.1.1.0/24。作为一个好的实践你应该设置不同的CIDR给不同的区域。这会更容易去设置VPN在不同的区域之间。

    - Hypervisor. (Introduced in version 3.0.1) 选择hypervisor 给区域的第一个集群。 之后你可以添加集群用不同的hypervisors，在完成添加区域之后。

    - 公共. 一个公共的区域是对所有用户有效的。一个不公开的区域会被安排给一个特别的域。只有在那个域的用户才能被允许建立来宾虚拟机在这个区域里。

2. 选择哪种通信类型

    通信类型是管理，公共，来宾和存储通信。更多的关于类型的信息，覆盖到相应的图标会显示他们的工具提示，或者参考第 2.8.3 节 “高级区域网络流量类型” 。这个界面初始时一个网络已经被配置了。如果你有多个物理网络，你需要添加多个。拖拽或者删除通信类型对一个灰色的网络，它会被激活。你可以移动通信图标从一个网络到另一个；例如，如果默认的通信类型显示在网络1的不匹配你的实际设置，你可以把他们拉下来。你也可以改变网络名按照你期望的。

3. (3.0.1 版本中介绍) 安排一个网络通信标识给每个通信类型在每个物理网络上。这些标识必须匹配你已经定义在hypervisor主机上的标识。安排每个标识，点击编辑按钮在通信类型图标下，一个弹出的对话框显示出来，你可以输入标识，然后点击确定。

    这些通信标识被定义只有hypervisor被选中给第一集群的时候。对于所有其它hypervisor,那些标识能够被配置在区域建立好以后。

4. 点击下一步。

5. 配置IP地址段给公共的Internet通信。输入下面详细信息，然后点击添加。如果有需要，你可以重复这步添加更多的公共Internet IP 段。完成后，点击下一步。

- 网关.网关使用为这些IP地址。
- 子网掩码. 和这个IP地址段相关的子网掩码。
- VLAN. VLAN 会被用到公共通信。
- 开始 IP/结束 IP. 一个IP地址段被安排

6. 在新的区域中，CloudStack添加第一个机架。你可以添加更多的机架在以后。关于什么是机架，请参考第 2.3 节 “关于POD”

   配置第一个机架，输入下面信息，然后点击下一步：

   - 机架名称. 机架的命名.
   - Reserved system gateway. 这个网关是给在机架中的主机使用的。
   - 预留系统网络掩码. 定义给机架的子网前缀。 使用CIDR形式。
   - 开始/结束 预留系统IP. 管理网络中的IP地址段，CloudStack 用来管理各种系统虚拟机，例如辅助存储系统虚拟机,控制代理系统虚拟机和DHCP.更多信息请参考 第 2.8.6 节 “System Reserved IP Addresses” .

7. Specify a range of VLAN IDs to carry guest traffic for each physical network (see VLAN Allocation Example), then click Next.

8. In a new pod, CloudStack adds the first cluster for you. You can always add more clusters later. For an overview of what a cluster is, see 第 2.4 节 “关于集群” .

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

   - Hypervisor. (Version 3.0.0 only; in 3.0.1, this field is read only) Choose the type of hypervisor software that all hosts in this cluster will run. If you choose VMware, additional fields appear so you can give information about a vSphere cluster. For vSphere servers, we recommend creating the cluster of hosts in vCenter and then adding the entire cluster to CloudStack. See Add Cluster: vSphere .
   - Cluster name. Enter a name for the cluster. This can be text of your choosing and is not used by CloudStack.

9. In a new cluster, CloudStack adds the first host for you. You can always add more hosts later. For an overview of what a host is, see 第 2.5 节 “关于宿主机” .

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

   在你能配置主机之前，你需要安装虚拟机管理软件在主机上。你需要知道哪个虚拟机管理软件版本被支持，并且额外的配置信息是需要确认主机是否会和CloudStack好好工作。这些安装详细信息请参考：

   - Citrix XenServer 安装为了 CloudStack

- VMware vSphere 安装和配置

- KVM 安装和配置

配置第一个主机，输入下列信息，然后点击下一步：

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

- Username. Usually root.

- Password. This is the password for the user named above (from your XenServer or KVM install).

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

10. In a new cluster, CloudStack adds the first primary storage server for you. You can always add more servers later. For an overview of what primary storage is, see 第 2.6 节 “关于主存储”.

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

- Name. The name of the storage device.

- Protocol. For XenServer, choose either NFS, iSCSI, or PreSetup. For KVM, choose NFS, SharedMountPoint, CLVM, and RBD. For vSphere choose either VMFS (iSCSI or FiberChannel) or NFS. The remaining fields in the screen vary depending on what you choose here.

NFS	• Server. The IP address or DNS name of the storage device. • Path. The exported path from the server. • Tags (optional). The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. 在一个区（zone）内的多个集群所拥有的主存储标签集必须是完全一致的。例如：如果集群A提供主存储有标签T1和T2，那么在这个区内的所有其它集群提供的主存储也必须有标签T1和T2。
iSCSI	• Server. The IP address or DNS name of the storage device. • Target IQN. The IQN of the target. For example, iqn.1986-03.com.sun:02:01ec9bb549-1271378984.

	• Lun. The LUN number. For example, 3. • Tags (optional). The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. 在一个区（zone）内的多个集群所拥有的主存储标签集必须是完全一致的。例如：如果集群A提供主存储有标签T1和T2，那么在这个区内的所有其它集群提供的主存储也必须有标签T1和T2。
预设置	• Server. The IP address or DNS name of the storage device. • SR Name-Label. 输入已经安装在CloudStack之外的SR名称标识。 • Tags (optional). The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. 在一个区（zone）内的多个集群所拥有的主存储标签集必须是完全一致的。例如：如果集群A提供主存储有标签T1和T2，那么在这个区内的所有其它集群提供的主存储也必须有标签T1和T2。
SharedMountPoint	• 路径. 主存储被挂载到每个主机上的路径。例如，"/mnt/primary". • Tags (optional). The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. 在一个区（zone）内的多个集群所拥有的主存储标签集必须是完全一致的。例如：如果集群A提供主存储有标签T1和T2，那么在这个区内的所有其它集群提供的主存储也必须有标签T1和T2。
VMFS	• 服务器. vCenter服务器的IP地址或者是DNS名称。 • 路径. 一个数据中心和数据存储合并的名称。格式是 "/" 数据中心名 "/" 数据存储名。例如，"/cloud.dc.VM/cluster1datastore". • Tags (optional). The comma-separated list of tags for this storage device.

	It should be an equivalent set or superset of the tags on your disk offerings. 在一个区（zone）内的多个集群所拥有的主存储标签集必须是完全一致的。例如：如果集群A提供主存储有标签T1和T2，那么在这个区内的所有其它集群提供的主存储也必须有标签T1和T2。

11. 在这个新区域中，CloudStack添加第一个辅助存储服务器给你。对于辅助存储概览，请参考 第 2.7 节 “关与辅助存储” 。

    在你填写这个界面之前，你需要提前准备好辅助存储通过设置NFS共享并且安装最新版本的CloudStack系统虚拟机模板。参考添加辅助存储。

    - NFS Server. The IP address of the server or fully qualified domain name of the server.

    - Path. The exported path from the server.

12. 点击启动。

7.4. 添加一个机架

When you created a new zone, CloudStack adds the first pod for you. You can add more pods at any time using the procedure in this section.

1. Log in to the CloudStack UI. See 第 5.1 节 “登陆到用户界面” .

2. In the left navigation, choose Infrastructure. In Zones, click View More, then click the zone to which you want to add a pod.

3. Click the Compute and Storage tab. In the Pods node of the diagram, click View All.

4. Click Add Pod.

5. Enter the following details in the dialog.

    - Name. The name of the pod.

    - Gateway. The gateway for the hosts in that pod.

    - Netmask. The network prefix that defines the pod's subnet. Use CIDR notation.

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

6. 点击 确定。

7.5. 添加集群

你需要告诉CloudStack 它要管理的主机。主机存在于集群中，所以在你开始加入主机到云中之前，你必须增加至少一个集群。

7.5.1. Add Cluster: KVM or XenServer

These steps assume you have already installed the hypervisor on the hosts and logged in to the CloudStack UI.

1. In the left navigation, choose Infrastructure. In Zones, click View More, then click the zone in which you want to add the cluster.

2. 点击计算标签。

3. In the Clusters node of the diagram, click View All.

4. Click Add Cluster.

5. Choose the hypervisor type for this cluster.

6. Choose the pod in which you want to create the cluster.

7. Enter a name for the cluster. This can be text of your choosing and is not used by CloudStack.

8. 点击 确定。

7.5.2. Add Cluster: vSphere

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

For vSphere servers, we recommend creating the cluster of hosts in vCenter and then adding the entire cluster to CloudStack. Follow these requirements:

- Do not put more than 8 hosts in a vSphere cluster

- Make sure the hypervisor hosts do not have any VMs already running before you add them to CloudStack.

To add a vSphere cluster to CloudStack:

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

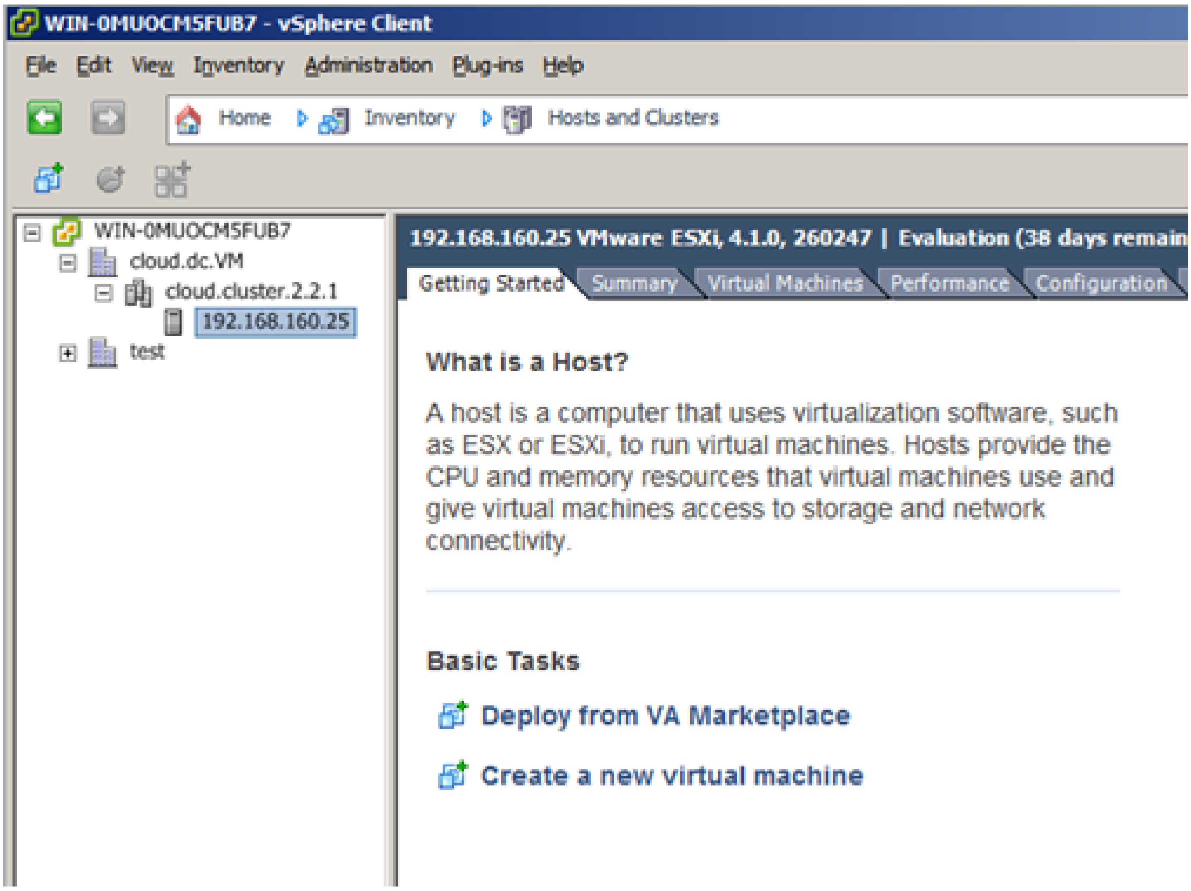
WIN-0MUOCM5FUB7 - vSphere Client
File Edit View Inventory Administration Plug-ins Help
Home
Inventory
Hosts and Clusters
WIN-0MUOCM5FUB7
cloud.dc.VM
cloud.cluster.2.2.1
192.168.160.25
test
192.168.160.25 VMware ESXi, 4.1.0, 260247 | Evaluation (38 days remain
Getting Started
Summary
Virtual Machines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

What is a Host?

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

Basic Tasks

Deploy from VA Marketplace
Create a new virtual machine

2. Log in to the UI.

3. In the left navigation, choose Infrastructure. In Zones, click View More, then click the zone in which you want to add the cluster.

4. Click the Compute tab, and click View All on Pods. Choose the pod to which you want to add the cluster.

5. Click View Clusters.

6. Click Add Cluster.

7. In Hypervisor, choose VMware.

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

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

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

   - vCenter Username. Enter the username that CloudStack should use to connect to vCenter. This user must have all administrative privileges.

   - vCenter Password. Enter the password for the user named above

   - vCenter Datacenter. Enter the vCenter datacenter that the cluster is in. For example, "cloud.dc.VM".

- 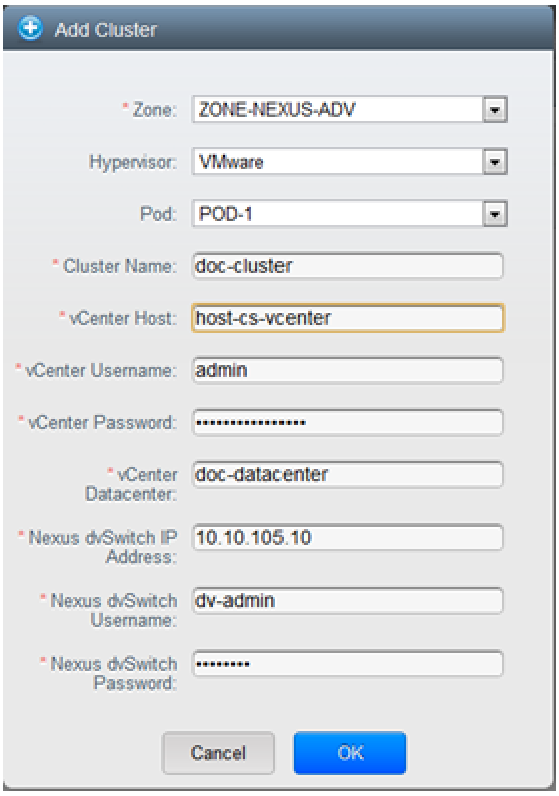

  Add Cluster

  * Zone: ZONE-NEXUS-ADV

  Hypervisor: VMware

  Pod: POD-1

  * Cluster Name: doc-cluster

  * vCenter Host: host-cs-vcenter

  * vCenter Username: admin

  * vCenter Password: ••••••••••••••••

  * vCenter Datacenter: doc-datacenter

  * Nexus dvSwitch IP Address: 10.10.105.10

  * Nexus dvSwitch Username: dv-admin

  * Nexus dvSwitch Password: ••••••••

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

7.6. Adding a Host

1. Before adding a host to the CloudStack configuration, you must first install your chosen hypervisor on the host. CloudStack can manage hosts running VMs under a variety of hypervisors.

   The CloudStack Installation Guide provides instructions on how to install each supported hypervisor and configure it for use with CloudStack. See the appropriate section in the Installation Guide for information about which version of your chosen hypervisor is supported, as well as crucial additional steps to configure the hypervisor hosts for use with CloudStack.

   > **警告**
   >
   > Be sure you have performed the additional CloudStack-specific configuration steps described in the hypervisor installation section for your particular hypervisor.

2. Now add the hypervisor host to CloudStack. The technique to use varies depending on the hypervisor.

   - 第 7.6.1 节 “Adding a Host (XenServer or KVM)”
   - 第 7.6.2 节 “增加一台主机(vSphere)”

7.6.1. Adding a Host (XenServer or KVM)

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

7.6.1.1. Requirements for XenServer and KVM Hosts

> **警告**
>
> Make sure the hypervisor host does not have any VMs already running before you add it to CloudStack.

Configuration requirements:

- Each cluster must contain only hosts with the identical hypervisor.
- For XenServer, do not put more than 8 hosts in a cluster.
- For KVM, do not put more than 16 hosts in a cluster.

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

7.6.1.1.1. XenServer Host Additional Requirements

If network bonding is in use, the administrator must cable the new host identically to other hosts in the cluster.

For all additional hosts to be added to the cluster, run the following command. This will cause the host to join the master in a XenServer pool.

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

> **注意**
>
> 当拷贝粘贴一条命令，确保在运行前粘贴的命令在一行上．一些文档查看器可能会在拷贝时引入不希望的换行符．

With all hosts added to the XenServer pool, run the cloud-setup-bond script. This script will complete the configuration and setup of the bonds on the new hosts in the cluster.

1. Copy the script from the Management Server in /usr/lib64/cloud/common/scripts/vm/hypervisor/xenserver/cloud-setup-bonding.sh to the master host and ensure it is executable.

2. 运行脚本。

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

7.6.1.1.2. KVM Host Additional Requirements

- If shared mountpoint storage is in use, the administrator should ensure that the new host has all the same mountpoints (with storage mounted) as the other hosts in the cluster.

- Make sure the new host has the same network configuration (guest, private, and public network) as other hosts in the cluster.

- If you are using OpenVswitch bridges edit the file agent.properties on the KVM host and set the parameter network.bridge.type to openvswitch before adding the host to CloudStack

7.6.1.2. Adding a XenServer or KVM Host

- If you have not already done so, install the hypervisor software on the host. You will need to know which version of the hypervisor software version is supported by CloudStack and what additional configuration is required to ensure the host will work with CloudStack. To find these installation details, see the appropriate section for your hypervisor in the CloudStack Installation Guide.

- 以管理员身份登录进入CloudStack 用户界面。

- In the left navigation, choose Infrastructure. In Zones, click View More, then click the zone in which you want to add the host.

- Click the Compute tab. In the Clusters node, click View All.

- Click the cluster where you want to add the host.

- Click View Hosts.

- Click Add Host.

- Provide the following information.

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

  - Username. Usually root.

  - Password. This is the password for the user from your XenServer or KVM install).

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

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

- Repeat for additional hosts.

7.6.2. 增加一台主机(vSphere)

对于 vSphere 服务器，我们建议在vCenter 中创建主机的集群并把整个集群加入到 CloudStack 中。见 Add Cluster: vSphere。

7.7. æ·»å# ä¸»å-#å#¨

7.7.1. 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.

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 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.

7.7.2. Adding Primary Stroage

当你建立一个新的区域的时候，主存储作为过程的一部分被添加。你也可以添加主存储在任何时候，例如当添加一个新的群集或者添加更多的主机到一个存在的群集的时候。

警告

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

1. Log in to the CloudStack UI (see 第 5.1 节 "登陆到用户界面").

2. 在左边的导航栏，选择基础架构。在区域中，点击查看全部，然后点击你想添加主存储的那个区域。

3. 点击计算标签。

4. 在图的主存储节点，点击查看所有。

5. 点击添加主存储。

6. 在对话框中提供下面的信息。

   - Pod. The pod for the storage device.

   - Cluster. The cluster for the storage device.

   - Name. The name of the storage device.

   - Protocol. For XenServer, choose either NFS, iSCSI, or PreSetup. For KVM, choose NFS or SharedMountPoint. For vSphere choose either VMFS (iSCSI or FiberChannel) or NFS.

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

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

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

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

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

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

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

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

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

在一个区（zone）内的多个集群所拥有的主存储标签集必须是完全一致的。例如：如果集群A提供主存储有标签T1和T2，那么在这个区内的所有其它集群提供的主存储也必须有标签T1和T2。

7. 点击 确定。

7.8. æ·»å# è¾ å#©å-#å#¨

7.8.1. System Requirements for Secondary Storage

- NFS storage appliance or Linux NFS server

- (Optional) OpenStack Object Storage (Swift) (see http://swift.openstack.org)

- 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.

7.8.2. 增加二级存储

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

警告

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

1. If you are going to use Swift for cloud-wide secondary storage, you must add the Swift storage to CloudStack before you add the local zone secondary storage servers. See 第 7.3 节 “创建Zone” .

2. To prepare for local zone secondary storage, you should have created and mounted an NFS share during Management Server installation. See Preparing NFS Shares in the Installation Guide.

3. Make sure you prepared the system VM template during Management Server installation. See Prepare the System VM Template in the Installation Guide.

4. Now that the secondary storage server for per-zone storage is prepared, add it to CloudStack. Secondary storage is added as part of the procedure for adding a new zone. See 第 7.3 节 “创建Zone” .

7.9. 初始化和测试

所有内容配置好之后，执行初始化。根据你的网络速度，大概花费30分钟或更长时间。当初始化成功执行之后，管理面板会在CloudStack UI中显示。

1. 系统验证完成后，在左边的导航栏，选择模版，点击‘ CentOS 5.5 (64bit) no Gui (KVM) template’，在下载完成前不要做下一步操作。

2. 进入‘实例’标签，查看我的实例。

3. 点击添加实例，进入导航。

    a. 选择你添加的区域。

    b. 在选择模版中，选择要在虚拟机中使用的模版，如果是新安装的系统，只有CentOS模版供选择。

    c. 选择计算方案，确定硬件环境允许启动你选择的计算方案。

    d. 在磁盘方案中，如果需要，添加另一个数据磁盘。第二个卷可能不会挂载在虚拟机上。例如：在XenServer的Linux下重启后你会看到/dev/xvdb。如果使用PV-enabled 的操作系统内核，则不需要重启。

    e. 在默认网络，为客户虚拟机选择主要网络。在试用安装中，在这里只有一个选项。

    f. 任意的为虚拟机添加名称和组。

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

4. 使用虚拟机，点击查看按钮

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

恭喜你，你成功完成了 CloudStack的安装

如果你决定增加你的部署，你可以增加更多的主机，主存储，区域，机架和集群

服务提供

In this chapter we discuss compute, disk, and system service offerings. Network offerings are discussed in the section on setting up networking for users.

8.1. 计算和磁盘服务提供方案

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，内存和网络资源保障
- 资源是如何计费的
- 资源使用是如何收费的
- 计费信息多久更新一次

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:

- 来宾CPU
- 来宾内存
- 来宾网络类型（虚拟网络还是直连网络）
- 根磁盘的标签

磁盘提供方案包括的内容：

- 磁盘大小（可选）。未指定批判大小的磁盘提供方案将允许用户自定义磁盘大小
- 数据盘标签

8.1.1. 创建计算资源提供方案

创建计算资源提供方案：

1. 以管理员权限登录CloudStack 用户界面。
2. 在左侧导航栏点击 服务提供方案。
3. 在下拉列表中选择 计算资源提供方案。
4. 点击添加计算资源提供方案。
5. 在对话框中填选以下内容：
    - 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. 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. 点击添加。

8.1.2. 创建一个新的磁盘方案

为了创建一个新的磁盘方案:

1. 使用管理员权限登入CloudStack UI.

2. 在左边导航栏，点击服务方案.

3. 在选择方案中，选择磁盘方案.

4. 点击添加磁盘方案.

5. 在对话框中，做出如下选择:

   - 名称. 任何你想要的磁盘方案名称.

   - 描述. 一个简单的描述，将会显示给最终用户

   - 自定义磁盘大小. 如果勾选，用户可以自己设置磁盘大小. 如果没有勾选，根管理员必须定义一个磁盘大小的数值.

   - 磁盘大小. 只有当自定义磁盘大小没选中时才会出现. 其定义的卷大小以GB为单位.

   - (可选)存储标签. 这个标签应与这个磁盘的主存储相关联. 标签以逗号分隔存储的属性列表. 比如 "ssd,blue". 标签所被添加在主存储上. CloudStack 通过标签来匹配存储和磁盘方案. 如果一个标签出现在磁盘方案里，那这个标签也必须出现在将要分配这个卷的主存储上. 如果这样的主存储不存在，从这个磁盘方案中进行分配将失败.

- 公有. 显示这个磁盘方案对所有的域都可见还是部分域可见. 选择是以便其对所有的域都可见. 选择不以限制其只对子域可见；CloudStack接着会提示子域的名字.

6. 点击添加.

8.1.3. 修改或删除服务提供

服务提供一旦被创建，不能改变。这也适用于计算提供和磁盘提供。

一个服务提供可以被删除。如果不再使用时，它立即并永久删除。如果服务提供仍然是在使用中，它会保留在数据库中，直到所有的虚拟机引用已被删除。管理员删除后，服务提供将不提供给创建新的实例的最终用户。

8.2. 系统服务方案

我们假设系统的服务提供了一个选择的CPU速度，CPU的数量，标签和RAM大小，就像其他的服务一样。但系统提供的服务，而不是用于虚拟机实例向用户公开，系统服务器用于更改虚拟路由器默熟悉，主机代理，和其他系统的虚拟机。系统服务器方案提供的服务是可见的CloudStackroot管理员。 CloudStack提供了默认的系服务方案。 ，CloudStackroot管理员可以创建额外的自定义系统服务方案

当CloudStack为来宾网络创建了一个虚拟路由器，它使用在系统方案中和网络方案关联的默认设置，你可以使用包含不同网络方案的系统方案来升级你虚拟路由器的性能。在改网络中的路由器将使用新的网络方案设置。

8.2.1. Creating a New System Service Offering

为了创建一个新的磁盘方案:

1. 使用管理员权限登入CloudStack UI.

2. 在左边导航栏，点击服务方案.

3. In Select Offering, choose System Offering.

4. Click Add System Service Offering.

5. 在对话框中，做出如下选择:

   - 名称. 任何你想要的磁盘方案名称.
   - 描述. 一个简单的描述，将会显示给最终用户
   - 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.

- 公有. 显示这个磁盘方案对所有的域都可见还是部分域可见. 选择是以便其对所有的域都可见. 选择不以限制其只对子域可见; CloudStack接着会提示子域的名字.

6. 点击添加。

8.3. Network Throttling

Network throttling is the process of controlling the network access and bandwidth usage based on certain rules. CloudStack controls this behaviour of the guest networks in the cloud by using the network rate parameter. This parameter is defined as the default data transfer rate in Mbps (Megabits Per Second) allowed in a guest network. It defines the upper limits for network utilization. If the current utilization is below the allowed upper limits, access is granted, else revoked.

You can throttle the network bandwidth either to control the usage above a certain limit for some accounts, or to control network congestion in a large cloud environment. The network rate for your cloud can be configured on the following:

- ç½#ç»#æ#¹ æ¡#

- æ##å#¡æ#¹ æ¡#

- Global parameter

If network rate is set to NULL in service offering, the value provided in the vm.network.throttling.rate global parameter is applied. If the value is set to NULL for network offering, the value provided in the network.throttling.rate global parameter is considered.

For the default public, storage, and management networks, network rate is set to 0. This implies that the public, storage, and management networks will have unlimited bandwidth by default. For default guest networks, network rate is set to NULL. In this case, network rate is defaulted to the global parameter value.

The following table gives you an overview of how network rate is applied on different types of networks in CloudStack.

ç½#ç»#	Network Rate Is Taken from
Guest network of Virtual Router	Guest Network Offering
Public network of Virtual Router	Guest Network Offering
Storage network of Secondary Storage VM	System Network Offering
Management network of Secondary Storage VM	System Network Offering

ç½#ç»#	Network Rate Is Taken from
Storage network of Console Proxy VM	System Network Offering
Management network of Console Proxy VM	System Network Offering
Storage network of Virtual Router	System Network Offering
Management network of Virtual Router	System Network Offering
Public network of Secondary Storage VM	System Network Offering
Public network of Console Proxy VM	System Network Offering
Default network of a guest VM	Compute Offering
Additional networks of a guest VM	Corresponding Network Offerings

A guest VM must have a default network, and can also have many additional networks. Depending on various parameters, such as the host and virtual switch used, you can observe a difference in the network rate in your cloud. For example, on a VMware host the actual network rate varies based on where they are configured (compute offering, network offering, or both); the network type (shared or isolated); and traffic direction (ingress or egress).

The network rate set for a network offering used by a particular network in CloudStack is used for the traffic shaping policy of a port group, for example: port group A, for that network: a particular subnet or VLAN on the actual network. The virtual routers for that network connects to the port group A, and by default instances in that network connects to this port group. However, if an instance is deployed with a compute offering with the network rate set, and if this rate is used for the traffic shaping policy of another port group for the network, for example port group B, then instances using this compute offering are connected to the port group B, instead of connecting to port group A.

The traffic shaping policy on standard port groups in VMware only applies to the egress traffic, and the net effect depends on the type of network used in CloudStack. In shared networks, ingress traffic is unlimited for CloudStack, and egress traffic is limited to the rate that applies to the port group used by the instance if any. If the compute offering has a network rate configured, this rate applies to the egress traffic, otherwise the network rate set for the network offering applies. For isolated networks, the network rate set for the network offering, if any, effectively applies to the ingress traffic. This is mainly because the network rate set for the network offering applies to the egress traffic from the virtual router to the instance. The egress traffic is limited by the rate that applies to the port group used by the instance if any, similar to shared networks.

For example:

Network rate of network offering = 10 Mbps

Network rate of compute offering = 200 Mbps

In shared networks, ingress traffic will not be limited for CloudStack, while egress traffic will be limited to 200 Mbps. In an isolated network, ingress traffic will be limited to 10 Mbps and egress to 200 Mbps.

8.4. Changing the Default System Offering for System VMs

You can manually change the system offering for a particular System VM. Additionally, as a CloudStack administrator, you can also change the default system offering used for System VMs.

1. Create a new system offering.

   For more information, see Creating a New System Service Offering.

2. 备份数据库:

   ```
   mysqldump -u root -p cloud | bzip2 > cloud_backup.sql.bz2
   ```

3. Open an MySQL prompt:

   ```
   mysql -u cloud -p cloud
   ```

4. Run the following queries on the cloud database.

   a. In the disk_offering table, identify the original default offering and the new offering you want to use by default.

      Take a note of the ID of the new offering.

      ```
      select id,name,unique_name,type from disk_offering;
      ```

   b. For the original default offering, set the value of unique_name to NULL.

      ```
      # update disk_offering set unique_name = NULL where id = 10;
      ```

      Ensure that you use the correct value for the ID.

   c. For the new offering that you want to use by default, set the value of unique_name as follows:

      For the default Console Proxy VM (CPVM) offering,set unique_name to 'Cloud.com-ConsoleProxy'. For the default Secondary Storage VM (SSVM) offering, set unique_name to 'Cloud.com-SecondaryStorage'. For example:

      ```
      update disk_offering set unique_name = 'Cloud.com-ConsoleProxy' where id = 16;
      ```

5. Restart CloudStack Management Server. Restarting is required because the default offerings are loaded into the memory at startup.

   ```
   service cloudstack-management restart
   ```

6. Destroy the existing CPVM or SSVM offerings and wait for them to be recreated. The new CPVM or SSVM are configured with the new offering.

Setting Up Networking for Users

9.1. Overview of Setting Up Networking 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
- 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
- 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

9.2. 关于虚拟网络

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.

9.2.1. 隔离的网络

一个隔离的网络可以访问虚拟机的单一账户。隔离的网络具有下列性质。

- 如VLAN等资源被动态分配和垃圾收集
- 有一个用于整个网络的网络提供
- 网络提供可升级或降级，但它是用于整个网络的

9.2.2. 共享网络

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).

- 管理员创建共享网络
- 在一个确定的网络中设计共享网络
- 共享网络资源如VLAN和物理网络，它映射到指定的管理员
- 共享网络通过安全组实现隔离
- 公网网络作为一个共享网络不会展示给终端用户

9.2.3. 虚拟网络资源的运行时分配

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.

9.3. 网络服务提供方案

> **注意**
>
> 与查看最新的网络服务提供者支持列表请见CloudStack用户界面或者条用API listNetworkServiceProviders。

服务提供者（也称为网络元件）是指通过硬件或虚拟应用来实现网络应用。比如，防火墙应用可以安装在云端来提供防火墙服务。在独立网络中多个提供者能提供相同的网络服务。比如，可以通过思科或者Juniper的设备在同一个物理网络中提供防火墙服务。

在一个网络中你可以多个实例使用相同的服务提供者（也可以使用多个Juniper SRX设备）

如果不同提供者被设置在网络中提供相同服务，管理员可以通过创建网络提供方案，因此用户能够自己制定使用哪个物理网络提供者（要遵从网络提供方案中的其他选项）。否则CloudStack 会在服务被需求的时候选择使用哪个提供者。

支持的网络服务提供者

CloudStack已经预置了一些内置的服务提供者支持列表。你能在创建网络提供方案的时候你能从这列表中选择。

	è##æ##è·¯ç#±å#¨	Citrix NetScaler	Juniper SRX	F5 BigIP	Host based (KVM/Xen)
Remote Access VPN	是	否	否	否	否
DNS/DHCP/User Data	是	否	否	否	否
é#² ç#«å¢#	是	否	是	否	否
è´#è½½å¹³è¡¡	是	是	否	是	否
å¼¹æ#§ IP	否	是	否	否	否
å¼¹æ#§è´#è½½å¹³è¡¡å#¨	否	是	否	否	否
æº# NAT	是	否	是	否	否
静态 NAT	是	是	是	否	否
ç«¯å#£è½¬å##	是	否	是	否	否

9.4. ç½#ç»#æ#¹ æ¡#

注意

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

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

- DHCP
- DNS
- æº# NAT
- é##æ## NAT
- ç«¯å#£è½¬å##
- è´#è½½å¹³è¡¡
- é#² ç#«å¢#
- VPN
- 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.

注意

如果你创建的负载均衡规则同时使用的网络提供方案使用到了外部的负载均衡器设备（比如NetScaler），随后又改变为使用虚拟路由器的网络提供方案,你必须在虚拟路由器上为每个已存在的负载聚恒规则创建一个防火墙规则以使它们继续生效。

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.

9.4.1. Creating a New Network Offering

To create a 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. 在对话框中，做出如下选择：

   - 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.

   - Guest Type. Choose whether the guest network is isolated or shared.

     For a description of this term, see 第 9.2 节 “关于虚拟网络”.

   - Persistent. Indicate whether the guest network is persistent or not. The network that you can provision without having to deploy a VM on it is termed persistent network. For more information, see 第 15.20 节 “Persistent Networks”.

   - Specify VLAN. (Isolated guest networks only) Indicate whether a VLAN should be specified when this offering is used.

   - VPC. This option indicate whether the guest network is Virtual Private Cloud-enabled. A Virtual Private Cloud (VPC) is a private, isolated part of CloudStack. A VPC can have its own virtual network topology that resembles a traditional physical network. For more information on VPCs, see 第 15.19.1 节 “About Virtual Private Clouds”.

   - 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.

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

æ#¯æ##ç##æ##å#¡	è¯´æ##	Isolated	å·²å ±äº«
DHCP	For more information, see 第 15.16 节 “DNS和DHCP”.	Supported	Supported

æ#¯æ##ç##æ##å#¡	è¯´æ##	Isolated	å·²å±äº«
DNS	For more information, see 第 15.16 节 “DNS和 DHCP” .	Supported	Supported
è´#è½½å¹³è¡¡å#¨	If you select Load Balancer, you can choose the CloudStack virtual router or any other load balancers that have been configured in the cloud.	Supported	Supported
é#² ç#«å¢#	For more information, see the Administration Guide.	Supported	Supported
æº# NAT	If you select Source NAT, you can choose the CloudStack virtual router or any other Source NAT providers that have been configured in the cloud.	Supported	Supported
é##æ## NAT	If you select Static NAT, you can choose the CloudStack virtual router or any other Static NAT providers that have been configured in the cloud.	Supported	Supported
ç«¯å#£è½¬å##	If you select Port Forwarding, you can choose the CloudStack virtual router or any other Port Forwarding providers that have been configured in the cloud.	Supported	Not Supported
VPN	For more information, see 第 15.17 节 “VPN 虚拟专用网” .	Supported	Not Supported

æ#¯æ##ç##æ##å#¡	è¯´æ##	Isolated	å·²å±äº«
User Data	For more information, see 第 20.3 节 “User Data and Meta Data” .	Not Supported	Supported
Network ACL	For more information, see 第 15.19.4 节 “Configuring Access Control List” .	Supported	Not Supported
å®#å¨ç»#	For more information, see 第 15.7.2 节 “Adding a Security Group” .	Not Supported	Supported

- 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 第 8.2 节 “系统服务方案” .

- Redundant router capability. Available only when Virtual Router is selected as the Source NAT provider. Select this option if you want to use two virtual routers in the network for uninterrupted connection: one operating as the master virtual router and the other as the backup. The master virtual router receives requests from and sends responses to the user’ s VM. The backup virtual router is activated only when the master is down. After the failover, the backup becomes the master virtual router. CloudStack deploys the routers on different hosts to ensure reliability if one host is down.

- 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. When conservative mode is off, the public IP can only be used for a single service. For example, a public IP used for a port forwarding rule cannot be used for defining other services, such as StaticNAT or load balancing. When the conserve mode is on, you can define more than one service on the same public IP.

注意

If StaticNAT is enabled, irrespective of the status of the conserve mode, no port forwarding or load balancing rule can be created for the IP. However, you can add the firewall rules by using the createFirewallRule command.

- Tags. Network tag to specify which physical network to use.

6. 点击添加.

使用虚拟机

10.1. 关于虚拟机的使用

CloudStack可以让管理员对运行在云中所有客户虚拟机进行完整的生命周期管理。CloudStack 提供给最终用户和管理员一些客户的管理操作。可以停止，启动，重启和删除虚拟机

客户虚拟机可以拥有名称和组。虚拟机的名字和组对于CloudStack 是透明的，方便最终用户识别管理自己的虚拟机。每个虚拟机可以拥有不同名称的3个名字。但只有其中的两个名字可以让用户来管理。

- 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

客户虚拟机可以设定为高可用性(HA)。启用HA的虚拟机会被系统监控。如果系统检测到这个虚拟机关机了，会试图重启这个虚拟机，一般会换一台主机。更多信息，请参阅启用虚拟机高可用性一章。

每个新虚拟机会被分配一个公共IP地址。当虚拟机启动的时候，CloudStack 会自动为虚拟机的内外网IP地址之间创建一个静态NAT。

如果使用了动态IP (配合NetScaler负载均衡设备)，新虚拟机在最初分配IP地址的时候不会标记为动态的。用户必须设定为动态IP来替代自动配置的IP地址，然后设定静态NAT来映射新IP和客户虚拟机的私有IP。虚拟机原有的IP会被释放然后重新回到可用的公共IP地址池中。

CloudStack 不能分辨客户虚拟机关机是由用户操作的(如在Linux下使用“shutdown”命令) 还是非预期的关机。如果开启HA功能的虚拟机从内部关机， CloudStack 仍然会重启它。如果想关闭有HA功能的虚拟机，必须要通过 CloudStack 的UI或 API。

10.2. 虚拟机最佳实践

CloudStack管理员应该监视每个集群中的虚拟机实例的总数，并禁止分配给集群接近最大的虚拟机管理程序可以处理。一定要留一个安全余量，允许一个或多个主机出现故障，这会增加其他主机上的虚拟机负载，增加自动重新部署的可能性。查询你选择虚拟机管理程序的文档，确认每个主机运行的最大虚拟机数量。然后使用CloudStack的全局配置设置为默认的限制。每时每刻监控集群上的活动虚拟机，保持主机故障后，虚拟机的总数仍旧属于一个安全的水平。例如，如果有N个主机在集群中，你想允许在集群中的一台主机任何时间被关机下线。你被允许在集群中部署的虚拟机主机的最大数量为(N-1)) *(每个主机的限制)。集群一旦达到这个数量的虚拟机，使用CloudStack UI禁用分配更多的虚拟机的集群。

10.3. VM Lifecycle

Virtual machines can be in the following states:

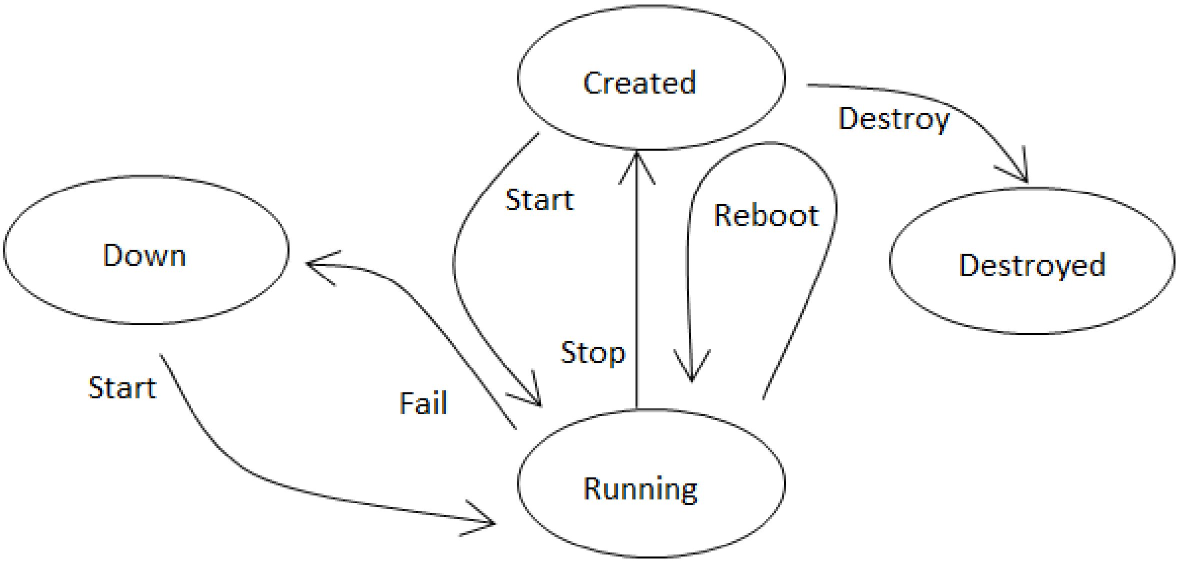

Once a virtual machine is destroyed, it cannot be recovered. All the resources used by the virtual machine will be reclaimed by the system. This includes the virtual machine’s IP address.

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.

10.4. 创建虚拟机实例

通常建议使用模板创建虚拟机实例。除此以外，用户也可以不使用模板而创建一个没有操作系统的空白虚拟机。用户可以为空白虚拟机附加ISO文件，并通过虚拟的CD/DVD-ROM设备安装操作系统。

注意

You can create a VM without starting it. You can determine whether the VM needs to be started as part of the VM deployment. A request parameter, startVM, in the deployVm API provides this feature. For more information, see the Developer's Guide

通过模板创建虚拟机实例

1. 使用管理员或普通用户，登陆到 CloudStack 。

2. 在左边的任务栏，点击 实例 。

3. 点击添加实例

4. Select a zone.

5. Select a template, then follow the steps in the wizard. For more information about how the templates came to be in this list, see 第 12 章 使用模板.

6. 请确认你的硬件设备允许运行所选的服务。

7. 点击提交，您的虚拟机将会被创建并启动。

> **注意**
>
> 因为安全的原因， 虚拟机的网络名称只有root可以看到 管理。

通过ISO文件创建虚拟机实例

> **注意**
>
> (XenServer) 运行在XenServer上的Windows虚拟机需要安装PV驱动，PV驱动可以 通过模板安装或在虚拟机创建后再安装。主要的管理功能 会用到PV驱动，例如添加额外的磁盘和ISO镜像，在线迁移，和Graceful 关机。

1. 使用管理员或普通用户，登陆到 CloudStack 。

2. 在左边的任务栏，点击 实例 。

3. 点击添加实例

4. Select a zone.

5. 选择ISO启动，然后按照向导的步骤进行操作。

6. 点击提交，您的虚拟机将会被创建并启动。

10.5. Accessing VMs

任何用户可以访问他们自己的虚拟机。管理员能够访问在云中运行的所有虚拟机。

To access a VM through the CloudStack UI:

1. 以管理员或用户身份登陆至CloudStack用户界面。

2. Click Instances, then click the name of a running VM.

3. Click the View Console button

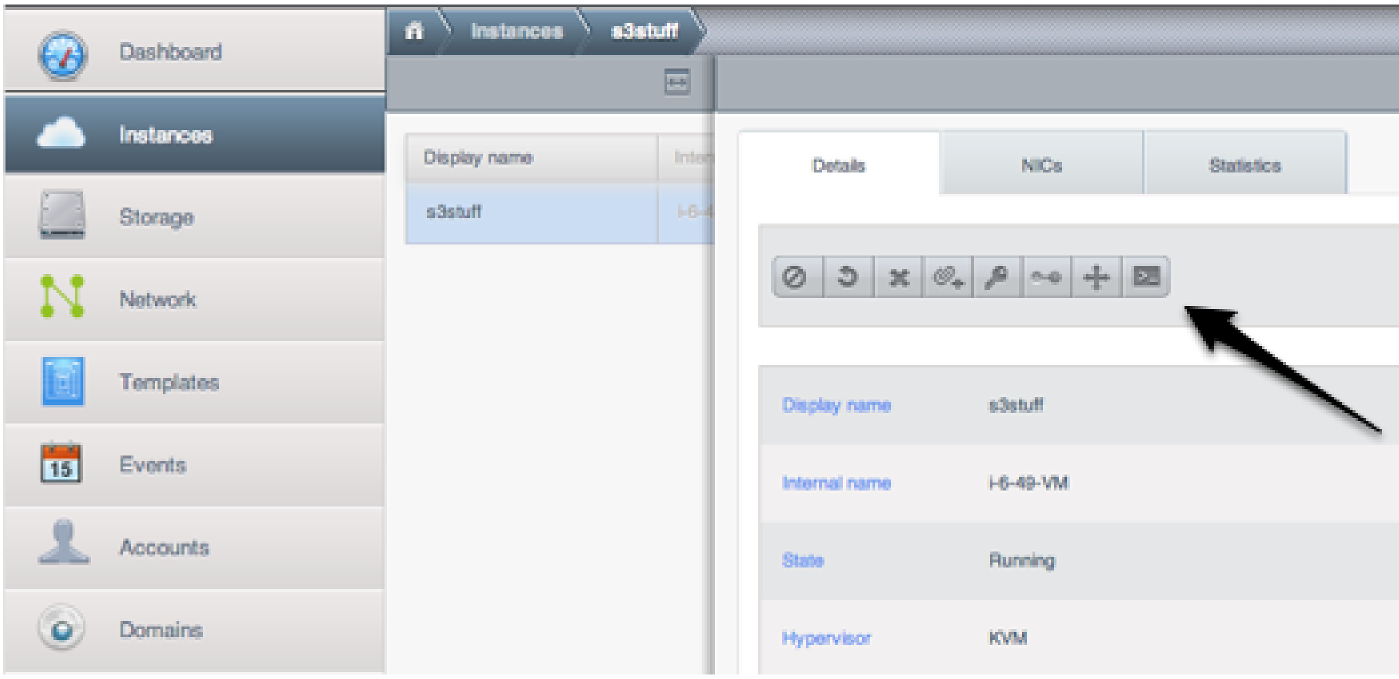

To access a VM directly over the network:

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 第 15.14 节 "IP转发及防火墙".

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 第 15.14 节 "IP转发及防火墙".

10.6. 停止和启动虚拟机

Once a VM instance is created, 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.

10.7. 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.

To access a VM through the CloudStack UI:

1. 以管理员或用户身份登陆至CloudStack用户界面。

2. 在左侧的导航菜单中，点击实例。

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:

7. Display name: Enter a new display name if you want to change the name of the VM.

8. OS Type: Select the desired operating system.

9. Group: Enter the group name for the VM.

10. Click Apply.

10.8. 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. 以管理员或用户身份登陆至CloudStack用户界面。

2. 在左侧的导航菜单中，点击实例。

3. Choose the VM that you want to work with.

4. Click the Stop button to stop the VM. ⊘

5. Click the Change Service button. ⊶

   The Change service dialog box is displayed.

6. Select the offering you want to apply to the selected VM.

7. 点击 确定。

10.9. 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.

To manually live migrate a virtual machine

1. 以管理员或普通用户身份登录进入CloudStack用户界面。

2. 在左侧的导航菜单中，点击实例。

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. 点击 确定。

10.10. 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. 以管理员或用户身份登陆至CloudStack用户界面。

2. 在左侧的导航菜单中，点击实例。

3. Choose the VM that you want to delete.

4. Click the Destroy Instance button. ✖

10.11. 使用iso文件

CloudStack支持iso并加载到客户vm中。一个iso文件是包含iso/cd-rom格式文件系统的只读文件。用户可以上传自己的iso文件， 并挂载到vm中

iso文件通过url方法上传。HTTP支持这个这个协议。你可以通过以下方式访问通过http方式发布的iso.比如http://my.web.server/filename.iso.

iso文件就像模板一样可以是共有或者私有. iso不是一个定制的hypersior,用户可以在vSphere和KVM已同样的方式挂载

iso镜像存储在系统中，配置和模板系统的相似的隐私级别。iso镜像被分为可引导和不可引导。一个可引导iso镜像包含一个OS镜像。CloudStack允许在不使用iso镜像时候启动客户vm.\n用户也可以在客户vm中挂载iso镜像。例如，在windows中安装pv驱动。iso镜像不是一个定制的hypervisor

10.11.1. 添加一个ISO

为了制造额外的操作系统或者其它有效的软件为了guest VMs使用，你能够添加ISO。最典型的ISO被认为是一个操作系统，但是你也能添加软件类型的ISO，例如你想作为模板被安装的桌面应用。

1. 登陆到CloudStack界面以管理员或者终端用户。

2. 在左边的导航栏，点击模板。

3. 在选择视图中，选择ISO。

4. 点击添加ISO。

5. 在添加ISO屏幕中，提供下列信息：

   - 名称:ISO 镜像的简称。例如，CentOS6.2 64-bit。

- 说明:对于ISO镜像的说明。例如，CentOS6.2 64-bit。
- URL:URL是ISO镜像的主机。管理控制服务器必须能够访问这个地址通过HTTP。如果有需要你可以直接放置ISO镜像在管理控制服务器。
- 区域:选择你想让ISO有效的区域，或者ALL Zones可以让它在整个CloudStack可用。
- 可启动:是否一个guest能够引导这个ISO image。例如，一个CentOS ISO 是可启动的，一个Microsoft Office ISO 是不可启动的。
- 操作系统类型:这帮助CloudStack和虚拟化平台执行正常的操作并且假设能够提高guest性能。选择下列中的一个。
  - 如果你需要的操作系统在ISO镜像被列出来，请选择它。
  - 如果ISO操作系统类型没有被列出或者ISO是不可引导的，选择Other。
  - (仅针对XenServer) 如果你想从这个ISO引导以PV 模式，选择 Other PV (32-bit) or Other PV(64-bit)
  - (仅针对KVM) 如果你选择一个操作系统是PV-enabled， 从这个ISO创建的虚拟机会有一个SCSI(virtio)引导磁盘。如果这个操作系统没有PV-enabled,虚拟机将有一个IDE引导磁盘。PV-enabled 类型是:

Fedora 13	Fedora 12	Fedora 11
Fedora 10	Fedora 9	Other PV
Debian GNU/Linux	CentOS 5.3	CentOS 5.4
CentOS 5.5	Red Hat Enterprise Linux 5.3	Red Hat Enterprise Linux 5.4
Red Hat Enterprise Linux 5.5	Red Hat Enterprise Linux 6	

> **注意**
>
> 不建议选择一个比镜像操作系统老的版本。比如，选择CentOS 5.4 去支持一个CentOS6.2的镜像通常会导致不能工作。在这种情况下，去选择Other。

- Extractable: Choose Yes if the ISO should be available for extraction.
- Public: Choose Yes if this ISO should be available to other users.
- 精选: 如果你想让这个ISO更显著为了让用户选择的话，请选中这个选项。这个ISO将会显示到精选ISOs列表中。只有管理员可以制作一个ISO精选。

6. 点击确定。

   管理控制服务器将下载那个ISO。依赖与那个ISO的大小，下载过程可能会用很长时间。ISO 状态栏将会显示Ready当它已经成功被下载到第二存储上的时候。点击刷新更新下载率。

7. 要点：等ISO下载完成的过程中。如果你想执行下一个任务并且试着正常使用这个ISO，这将会失败。完整的ISO必须是有效的在CloudStack能够使用它之前。

10.11.2. 附加ISO到虚拟机

1. 在左侧的导航菜单中，点击实例。

2. 选择要使用的虚拟机。

3. Click the Attach ISO button.

4. 在附加ISO对话框中，选择所需的ISO。

5. 点击 确定。

Working With Hosts

11.1. Adding Hosts

Additional hosts can be added at any time to provide more capacity for guest VMs. For requirements and instructions, see 第 7.6 节 "Adding a Host" .

11.2. 主机定期维护和维护模式

您可以将主机置于维护模式。维护模式被激活时，主机将无法接收新的客户虚拟机，并且已经在主机上运行的虚拟机被无缝地迁移到另一台不在维护模式下的主机上。迁移使用实时迁移技术和不中断客户的执行。

11.2.1. 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 may be migrated back to it manually and new VMs can be added.

11.2.2. XenServer和维护模式

使用XenServer的时候,可以在xencenter中对某一台服务器启用维护模式,将该服务器暂时置于离线状态,当你对服务器启用维护模式的时候,所有在该台服务器运行的VM,将自动的迁移到同一个资源池中的其他服务器.如果原来该服务器是资源池的主服务器,一个新的主服务器将自动产生并履职.当一台服务器处于维护模式的时候,不能在其上执行创建VM等操作.

下面来简单描述如何将一个服务器启用维护模式.

1. 在资源栏中,选择一台服务器,然后选择执行如下某一个操作.

   - 右键,在弹出菜单选择 Enter Maintenance Mode

   - On the Server menu, click Enter Maintenance Mode.

2. Click Enter Maintenance Mode.

资源栏中的服务器状态将会显示所有正在该服务器上运行的VM将会成功从该服务器迁移至同资源池中的其他服务器.

下面来简单描述如何将一个服务器退出维护模式.

1. 在资源栏中,选择一台服务器,然后选择执行如下某一个操作.

   - 右键,在弹出菜单选择 Exit Maintenance Mode

   - On the Server menu, click Exit Maintenance Mode.

2. Click Exit Maintenance Mode.

11.3. Disabling and Enabling Zones, Pods, and Clusters

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

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. 点击计算标签。

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.

11.4. 移除主机

主机可以按照需要从云中移除。移除主机的过程取决于hypervisor的类型。

11.4.1. 移除XenServer和KVM宿主机

主机进入维护模式才能在集群中移除。该宿主机上所有的虚拟机已经迁移到其他宿主机上之后才能进入维护模式。从cloud中移除宿主机:

1. 将节点置于维护模式。

   详见 第 11.2 节 “主机定期维护和维护模式”.

2. For KVM, stop the cloudstack-agent service.

3. 点击用户界面中的删除按钮删除节点。

   然后你就可以关闭这台已被移除的宿主机，重新使用它的IP，重装它，等等。

11.4.2. 移除vSphere 主机

如果要移除这种类型的主机，首先将其设置为维护模式，请参照第 11.2 节 “主机定期维护和维护模式”。然后使用 CloudStack 删除主机。CloudStack 将不会直接管理这台已经被CloudStack 删除的主机。然而这台主机可能仍存在于vCenter的集群中。

11.5. 重新安装主机

在主机进入维护模式后，接着移除它，你就可以重新安装主机。如果一个主机宕机或者不能进入维护模式，在重新安装前仍旧应该被移除。

11.6. 维护主机上的虚拟机管理器

在虚拟机管理器软件在主机上执行时，请确定已应用了由虚拟机软件提供商提供的全部补丁。时刻保持虚拟机管理器软件的更新。本产品不会追踪或提醒你虚拟机管理器软件的更新。保持虚拟机管理器软件的补丁最新很重要。虚拟机管理器提供商很可能对于未更新补丁的系统不提供支持。

> **注意**
>
> 缺乏最新补丁更新可能会导致数据和虚拟机丢失。

(XenServer) For more information, see Highly Recommended Hotfixes for XenServer in the CloudStack Knowledge Base[1].

11.7. 更改主机密码

XenServer主机，kvm主机或者vSphere主机的密码可以在数据库中进行修改。需要注意的是在集群中所有节点主机必须拥有同样的密码。

修改节点密码:

1. 验证所有集群主机。

[1] http://docs.cloudstack.org/Knowledge_Base/Possible_VM_corruption_if_XenServer_Hotfix_is_not_Applied/Highly_Recommended_Hotfixes_for_XenServer_5.6_SP2

2. 修改集群所有主机密码。现在主机密码和已知的 CloudStack 密码将不匹配。在二者密码匹配前，所有的对于集群的操作将失败。

3. 得到集群中你打算修改密码主机的主机id列表。你需要通过访问数据库确定这些主机id。比如打算为主机名包含“h”(或者vSphere 集群)的修改密码，执行：

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

4. 将返回一个单独的id.对应这些主机的id记录集。

5. 更新数据库中主机的密码。下面这个例子，我们修改主机id为5,10和12的密码为“password”.

   ```
   mysql> update cloud.host set password='password' where id=5 or id=10 or id=12;
   ```

11.8. 主机分配

系统会选择最合适的宿主机来运新每一台虚拟机。终端用户能指定虚拟机创建在哪个资源域。终端用户无权选择虚拟机运行于哪台宿主机。

CloudStack管理员能指定特定宿主机来有限运行莫特定类型的虚拟机实例。比如，管理员能指定某台宿主机优先运行windows的来宾虚拟机。默认的宿主机分配器会优先将该类型的来宾虚拟机分配到这类宿主机上。如果没有可用的这类宿主机，分配器将会把来宾虚拟机安排在任意满足该虚拟机硬件要求的宿主机上。

垂直分配和水平分配均被支持。垂直分配将优先将虚拟机分配到一台宿主机上（直到不能分配为止）。这能降低电力消耗。水平分配以一种轮训的方式将虚拟机分配到各个宿主机上。这能在一些情况下带来更好的性。CloudStack通过管理员的配置支持CPU超限分配，。超限分配允许管理员配置超过物理CPU能力若干倍的分配量。

CloudStack也支持一种插件似的接口添加新的分配器。这些自定义的分配器能提供任意管理员七万个分配策略。

11.8.1. 超配和服务提供方案的限制

CloudStack 的CPU资源超配是基于管理员设置的超配比例的。这在全局变量 cpu.overprovisioning.factor中定义。

CloudStack 的CPU资源超配是基于管理员设置的超配比例的。这在全局变量 cpu.overprovisioning.factor中定义。

服务提供方案限制（比如1G主频，1核）严格执行核心数。比如，一个使用分配有一个核心服务方案的来宾，会只有一个核心可用，不管主机上其它的活动。

服务方案限制为千兆的，执行CPU资源争用机制。比如，假设一个来宾使用1GHz的服务提供方案，而它所在的主机有一个2GHz的核心，并且该来宾是此主机上唯一运行的实例。这个来宾将可以使用2GHz的cpu资源。当多个来宾打算使用CPU资源时，一个权衡机制将用来分配CPU资源。权重值基于来宾所使用的服务提供方案中指定的时钟频率。来宾被分配的CPU资源将服务提供方案中相应的频率执行。比如，一个使用2GHz服务提供方案的来宾将获得使用1GHz服务方案的来宾CPU资源的两倍，。CloudStack不支持内存超额分配。

11.9. vlan 配置

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.

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

第 12

使用模板

模板相当于虚拟机的重用配置。当用户创建虚拟机时能从CloudStack的模板列表中选择一个。

特殊情况下，模板可以是一个包含一个或多个操作系统的虚拟磁盘镜像，你可以选择性的安装另外的软件，比如office应用并设置访问控制来君顶谁能使用这个模板。每个模板对应一个一般类型的hypervisor，这个hypervisor是在添加入CloudStack时指定的。

CloudStack附带一个默认模板。为了向用户呈现出更多选择CloudStack管理员和用户能创建模板并添加到CloudStack中。

12.1. 创建模板概览

CloudStack默认已经有了一个CentOS的默认模板。有许多添加更多模板的方法，管理员和普通用户均能添加。一般是这样的顺序：

1. 运行一个你想要的操作系统的虚拟机实例，并进行一些你需要的设置。

2. 关闭这虚拟机。

3. 将卷转换为模板。

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.

接下来的几节中将继续描述各种创建模板的技术。

12.2. 针对模板的需求

- 对于 XenServer，在每一个你创建的模板上安装 PV 驱动 / Xen 工具。 这将使动态迁移和干净的宾客关机成为可能。

- 对于 vSphere，在每一个你创建的模板上安装VMware 工具。这将使控制台视图能够正常工作。

12.3. 模板最佳实践

如果你计划使用大的模板（100 GB 或更大），确保你有10g 的网络以支持大的模板。 当大的模板被使用时，较慢的网络可能导致超时及其它错误。

12.4. 默认模版

CloudStack 包含一个CentOS 模版。当主存储和二级存储配置完成后，这个模版会由二级存储虚拟机下载。可以在生产部署中使用这个模版，也可以删除掉它，使用自定义的模版。

默认模版的root用户密码是“password”。

默认模版为XenServer，KVM和vSphere各提供了一个。模版的下载依赖于云中使用了哪种宿主机。每个模版大约是2.5GB的物理大小。

默认模版包括标准的iptables 规则，会阻止除了ssh以外的其他访问。

```
# 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-
```

12.5. 私有模板和公共模板

用户创建模板时可选择模板为公有还是私有。

私有模板只对创建者可用。默认上传的模板都是私有的。

当用户将模板标识为“公有”，该模板将对该用户所处域、能连接至该用户处于区域的其他域用户可用。这取决于zone是设置成公用还是私有。私有区域被分配到一个单一的域，而公共区域能被任何域访问。

12.6. 通过已存在的虚拟机创建模板

当你已经有了一台安妮想法已经配置好的虚拟机，你就能以他为原型创建别的虚拟机。

1. 使用第 10.4 节 “创建虚拟机实例”给出的方法创建并且开启一个虚拟机。

2. 在虚拟机中做好需要的配置然后点击按钮关闭该虚拟机。

3. 等待虚拟机关闭。当虚拟机状态显示为“已停止”，执行下一步。

4. 点击创建模板并填写如下内容:

   - 名字和显示文本. 这些会在网页用户界面显示，所以建议写一些描述信息。

   - 系统类型。这个选项有助于CloudStack 和虚拟机软件执行正确的操作和提高来宾的性能。选择以下其中一个。

     - 如果已停止虚拟机的系统在列表中就选择。

     - 如果已停止虚拟机系统类型不在列表中就选择其他。

     - 如果你打算以PV模式启动该模板，请选择其他PV（32位）或其他PV（64位） 。这个选项只对XenServer有效:

> **注意**
>
> 注意：一般情况下你不应该选择一个比你镜像中系统地的版本。比如选用CentOS 5.4来支持CentOS6.2的镜像会导致启动失败。这种情况下你应该选择其他。

- 公用.选择是使模板对CloudStack installation所有用户可见。该模板会在社区模板中出现。详情请见 第 12.5 节 “私有模板和公共模板”.
- Password Enabled. Choose Yes if your template has the CloudStack password change script installed. See 第 12.13 节 “Adding Password Management to Your Templates”.

5. 点击添加。

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.

12.7. 从一个快照创建一个模板

如果你不想为了使用Create Template 菜单项而停止虚拟机（如在第 12.6 节 “通过已存在的虚拟机创建模板”中描述的）， 你可以通过 CloudStack 用户界面从任何快照直接创建模板。

12.8. Uploading Templates

> **vSphere Templates and ISOs**
>
> If you are uploading a template that was created using vSphere Client, be sure the OVA file does not contain an ISO. If it does, the deployment of VMs from the template will fail.

Templates are uploaded based on a URL. HTTP is the supported access protocol. Templates are frequently large files. You can optionally gzip them to decrease upload times.

To upload a template:

1. 在左边的导航栏，点击模板。

2. Click Register Template.

3. Provide the following:

   - Name and Description. 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.

  > 注意
  >
  > You should not choose an older version of the OS than the version in the image. For example, choosing CentOS 5.4 to support a CentOS 6.2 image will in general not work. In those cases you should choose Other.

- Hypervisor: The supported hypervisors are listed. Select the desired one.
- Format. The format of the template upload file, such as VHD or OVA.
- Password Enabled. Choose Yes if your template has the CloudStack password change script installed. See Adding Password Management to Your Templates
- Extractable. Choose Yes if the template is available for extraction. If this option is selected, end users can download a full image of a template.
- 公用.选择是使模板对CloudStack installation所有用户可见。该模板会在社区模板中出现。详情请见 第 12.5 节 “私有模板和公共模板”.
- 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.

12.9. 导出模板

最终用户和管理员可以从CloudStack 导出模板。导航到用户界面中的模板并选择Actions菜单中的Download功能。

12.10. 创建WINdows 模板

创建Windows模板必须完成公用Sysprep（卷快照）创建。Sysprep（卷快照）允许你创建一个通用的WINdows模板并避免任何可能的SID冲突

> 注意
>
> 在XenServer上运行Windows虚拟机需要PV驱动支持，用来在虚拟机创建后创建模板。PV驱动程序需要基本的管理功能，以及一些额外的卷容量、ISO镜像、实时迁移功能和可以正常关机。

安装步骤概述:

1. 上传WINdows 系统安装ISO文件.

   For more information, see 第 10.11.1 节 “添加一个ISO” .

2. 使用刚上传的Windows ISO创建一个实例。

   For more information, see 第 10.4 节 “创建虚拟机实例” .

3. 按照需求创建WIN 2008 R2或WIN2003 R2 卷快照。

4. 准备步骤完成，现在可以创建WINdows模板了。

12.10.1. System Preparation for Windows Server 2008 R2

For Windows 2008 R2, you run Windows System Image Manager to create a custom sysprep response XML file. Windows System Image Manager is installed as part of the Windows Automated Installation Kit (AIK). Windows AIK can be downloaded from Microsoft Download Center[1].

Use the following steps to run sysprep for Windows 2008 R2:

> 注意
>
> The steps outlined here are derived from the excellent guide by Charity Shelbourne, originally published at Windows Server 2008 Sysprep Mini-Setup.[2]

1. Download and install the Windows AIK

   > 注意
   >
   > Windows AIK should not be installed on the Windows 2008 R2 VM you just created. Windows AIK should not be part of the template you create. It is only used to create the sysprep answer file.

2. Copy the install.wim file in the \sources directory of the Windows 2008 R2 installation DVD to the hard disk. This is a very large file and may take a long time to copy. Windows AIK requires the WIM file to be writable.

3. Start the Windows System Image Manager, which is part of the Windows AIK.

4. In the Windows Image pane, right click the Select a Windows image or catalog file option to load the install.wim file you just copied.

[1] http://www.microsoft.com/en-us/download/details.aspx?id=9085

[2] http://blogs.technet.com/askcore/archive/2008/10/31/automating-the-oobe-process-during-windows-server-2008-sysprep-mini-setup.aspx

5. Select the Windows 2008 R2 Edition.

   You may be prompted with a warning that the catalog file cannot be opened. Click Yes to create a new catalog file.

6. In the Answer File pane, right click to create a new answer file.

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

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

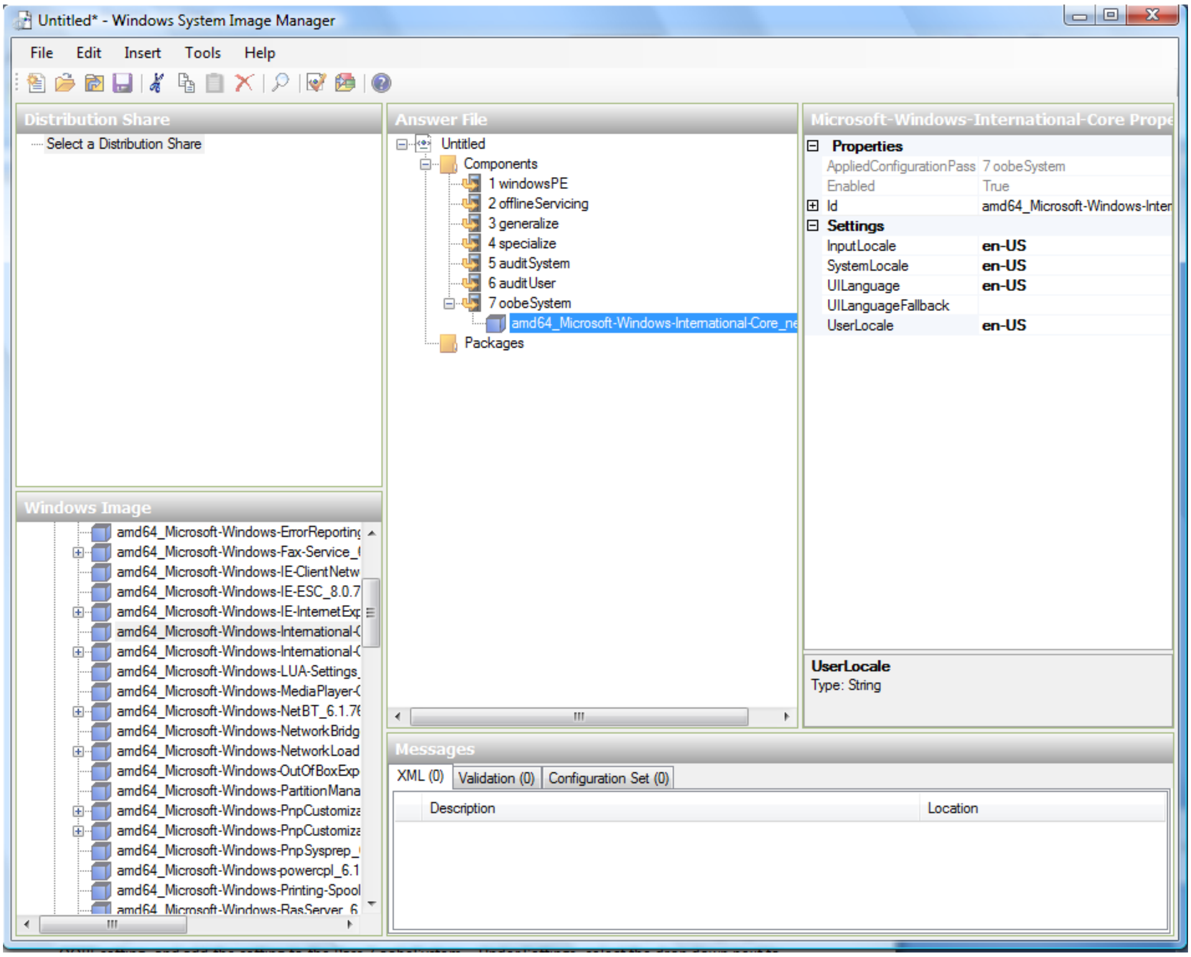

Untitled* - Windows System Image Manager

Answer File: Untitled > Components > 7 oobeSystem > amd64_Microsoft-Windows-International-Core_ne

Microsoft-Windows-International-Core Prope

Properties	
AppliedConfigurationPass	7 oobeSystem
Enabled	True
Id	amd64_Microsoft-Windows-Inter
Settings	
InputLocale	**en-US**
SystemLocale	**en-US**
UILanguage	**en-US**
UILanguageFallback	
UserLocale	**en-US**

UserLocale
Type: String

File
Edit
Insert
Tools
Help
Distribution Share
Select a Distribution Share
Answer File
Untitled
Components
1 windowsPE
2 offlineServicing
3 generalize
4 specialize
5 auditSystem
6 auditUser
7 oobeSystem
Packages
Windows Image
amd64_Microsoft-Windows-ErrorReporting
amd64_Microsoft-Windows-Fax-Service_
amd64_Microsoft-Windows-IE-ClientNetw
amd64_Microsoft-Windows-IE-ESC_8.0.7
amd64_Microsoft-Windows-IE-InternetExp
amd64_Microsoft-Windows-International-
amd64_Microsoft-Windows-International-
amd64_Microsoft-Windows-LUA-Settings_
amd64_Microsoft-Windows-MediaPlayer-
amd64_Microsoft-Windows-NetBT_6.1.76
amd64_Microsoft-Windows-NetworkBridg
amd64_Microsoft-Windows-NetworkLoad
amd64_Microsoft-Windows-OutOfBoxExp
amd64_Microsoft-Windows-PartitionMana
amd64_Microsoft-Windows-PnpCustomiza
amd64_Microsoft-Windows-PnpCustomiza
amd64_Microsoft-Windows-PnpSysprep_
amd64_Microsoft-Windows-powercpl_6.1
amd64_Microsoft-Windows-Printing-Spool
amd64_Microsoft-Windows-RasServer_6
Messages
XML (0)
Validation (0)
Configuration Set (0)
Description
Location

   b. You need to automate the Software License Terms Selection page, otherwise known as the End-User License Agreement (EULA). To do this, expand the Microsoft-Windows-Shell-Setup component. High-light the OOBE setting, and add the setting to the Pass 7 oobeSystem. In Settings, set HideEULAPage true.

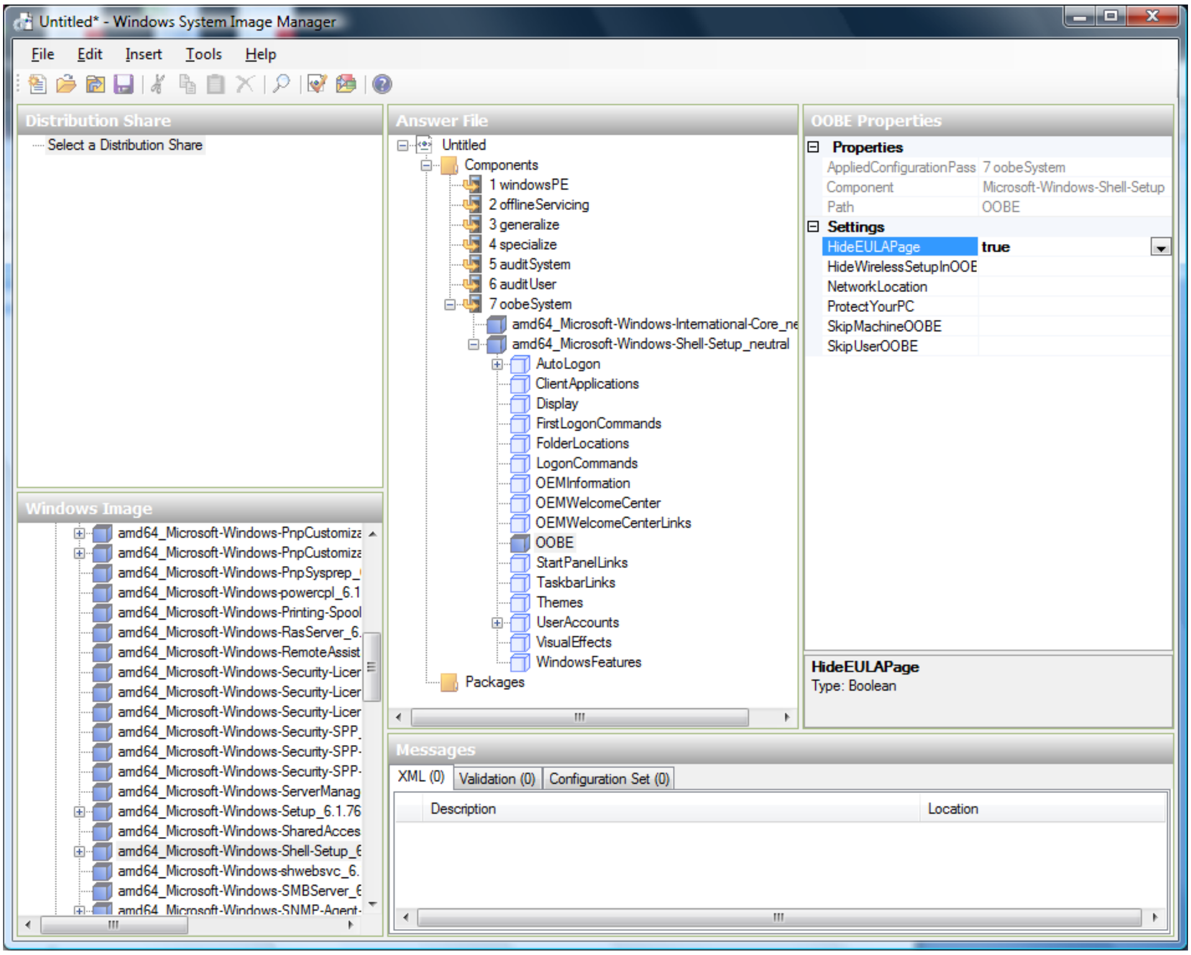
Untitled* - Windows System Image Manager
File
Edit
Insert
Tools
Help
Distribution Share
Select a Distribution Share
Windows Image
amd64_Microsoft-Windows-PnpCustomiza
amd64_Microsoft-Windows-PnpCustomiza
amd64_Microsoft-Windows-PnpSysprep_
amd64_Microsoft-Windows-powercpl_6.1
amd64_Microsoft-Windows-Printing-Spool
amd64_Microsoft-Windows-RasServer_6.
amd64_Microsoft-Windows-RemoteAssist
amd64_Microsoft-Windows-Security-Licer
amd64_Microsoft-Windows-Security-Licer
amd64_Microsoft-Windows-Security-Licer
amd64_Microsoft-Windows-Security-SPP_
amd64_Microsoft-Windows-Security-SPP-
amd64_Microsoft-Windows-Security-SPP-
amd64_Microsoft-Windows-ServerManag
amd64_Microsoft-Windows-Setup_6.1.76
amd64_Microsoft-Windows-SharedAcces
amd64_Microsoft-Windows-Shell-Setup_6
amd64_Microsoft-Windows-shwebsvc_6.
amd64_Microsoft-Windows-SMBServer_6
amd64_Microsoft-Windows-SNMP-Agent-
Answer File
Untitled
Components
1 windowsPE
2 offlineServicing
3 generalize
4 specialize
5 auditSystem
6 auditUser
7 oobeSystem
amd64_Microsoft-Windows-International-Core_ne
amd64_Microsoft-Windows-Shell-Setup_neutral
AutoLogon
ClientApplications
Display
FirstLogonCommands
FolderLocations
LogonCommands
OEMInformation
OEMWelcomeCenter
OEMWelcomeCenterLinks
OOBE
StartPanelLinks
TaskbarLinks
Themes
UserAccounts
VisualEffects
WindowsFeatures
Packages
OOBE Properties
Messages
XML (0)
Validation (0)
Configuration Set (0)
Description
Location

Properties

AppliedConfigurationPass	7 oobeSystem
Component	Microsoft-Windows-Shell-Setup
Path	OOBE

Settings

HideEULAPage	true
HideWirelessSetupInOOE	
NetworkLocation	
ProtectYourPC	
SkipMachineOOBE	
SkipUserOOBE	

HideEULAPage
Type: Boolean

c. Make sure the license key is properly set. If you use MAK key, you can just enter the MAK key on the Windows 2008 R2 VM. You need not input the MAK into the Windows System Image Manager. If you use KMS host for activation you need not enter the Product Key. Details of Windows Volume Activation can be found at http://technet.microsoft.com/en-us/library/bb892849.aspx

d. You need to automate is the Change Administrator Password page. Expand the Microsoft-Windows-Shell-Setup component (if it is not still expanded), expand UserAccounts, right-click on AdministratorPassword, and add the setting to the Pass 7 oobeSystem configuration pass of your answer file. Under Settings, specify a password next to Value.

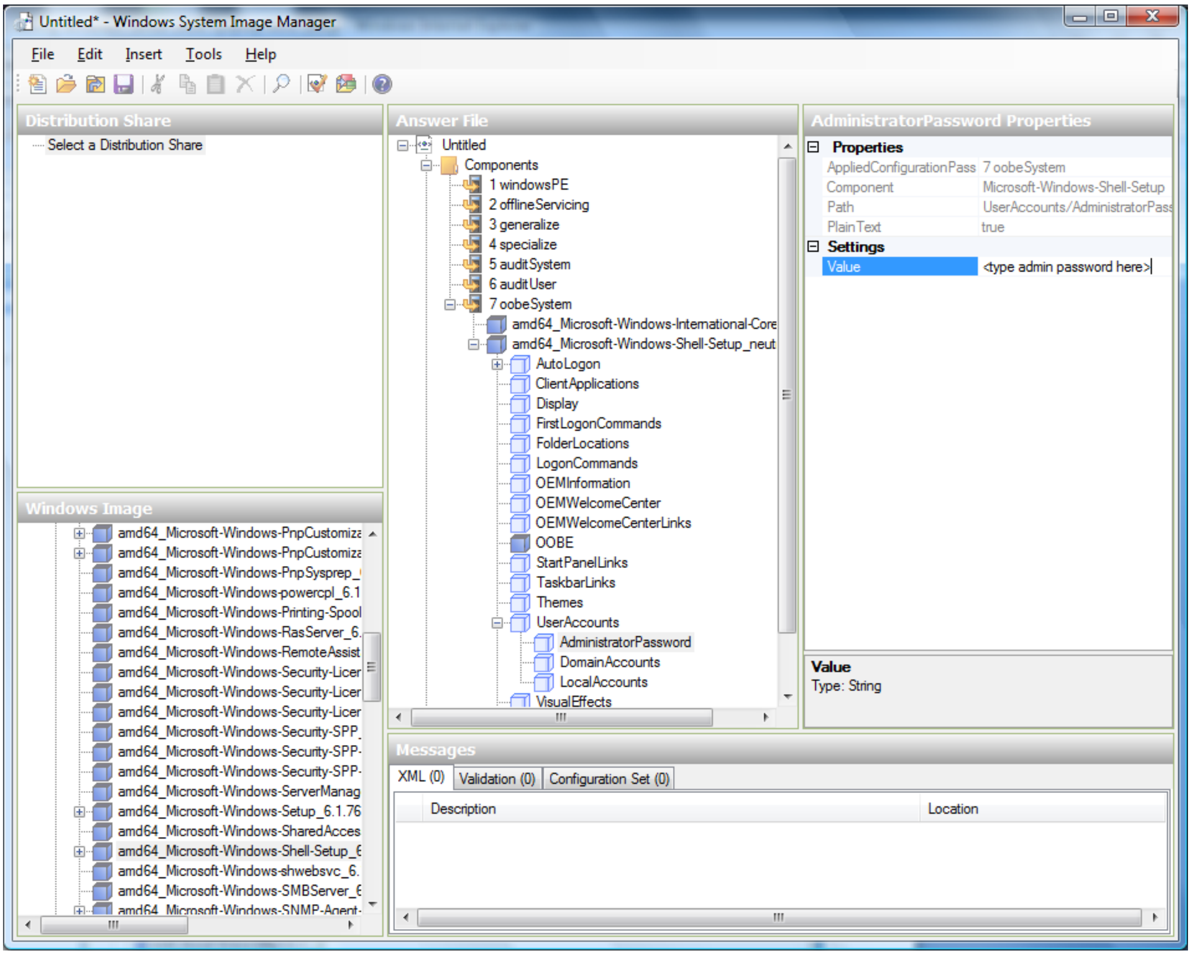
Untitled* - Windows System Image Manager
File Edit Insert Tools Help
Distribution Share
Select a Distribution Share
Answer File
Untitled
Components
1 windowsPE
2 offlineServicing
3 generalize
4 specialize
5 auditSystem
6 auditUser
7 oobeSystem
amd64_Microsoft-Windows-International-Core
amd64_Microsoft-Windows-Shell-Setup_neut
AutoLogon
ClientApplications
Display
FirstLogonCommands
FolderLocations
LogonCommands
OEMInformation
OEMWelcomeCenter
OEMWelcomeCenterLinks
OOBE
StartPanelLinks
TaskbarLinks
Themes
UserAccounts
AdministratorPassword
DomainAccounts
LocalAccounts
VisualEffects
AdministratorPassword Properties
Windows Image
amd64_Microsoft-Windows-PnpCustomiza
amd64_Microsoft-Windows-PnpCustomiza
amd64_Microsoft-Windows-PnpSysprep_
amd64_Microsoft-Windows-powercpl_6.1
amd64_Microsoft-Windows-Printing-Spool
amd64_Microsoft-Windows-RasServer_6.
amd64_Microsoft-Windows-RemoteAssist
amd64_Microsoft-Windows-Security-Licer
amd64_Microsoft-Windows-Security-Licer
amd64_Microsoft-Windows-Security-Licer
amd64_Microsoft-Windows-Security-SPP_
amd64_Microsoft-Windows-Security-SPP-
amd64_Microsoft-Windows-Security-SPP-
amd64_Microsoft-Windows-ServerManag
amd64_Microsoft-Windows-Setup_6.1.76
amd64_Microsoft-Windows-SharedAcces
amd64_Microsoft-Windows-Shell-Setup_6
amd64_Microsoft-Windows-shwebsvc_6.
amd64_Microsoft-Windows-SMBServer_6
amd64_Microsoft-Windows-SNMP-Agent-
Messages
XML (0)
Validation (0)
Configuration Set (0)
Description
Location

Properties

AppliedConfigurationPass	7 oobeSystem
Component	Microsoft-Windows-Shell-Setup
Path	UserAccounts/AdministratorPass
PlainText	true

Settings

Value	<type admin password here>

Value
Type: String

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.

12.10.2. System Preparation for Windows Server 2003 R2

Earlier versions of Windows have a different sysprep tool. Follow these steps for Windows Server 2003 R2.

1. Extract the content of \support\tools\deploy.cab on the Windows installation CD into a directory called c:\sysprep on the Windows 2003 R2 VM.

2. Run c:\sysprep\setupmgr.exe to create the sysprep.inf file.

a. Select Create New to create a new Answer File.

b. Enter “Sysprep setup” for the Type of Setup.

c. Select the appropriate OS version and edition.

d. On the License Agreement screen, select “Yes fully automate the installation”.

e. Provide your name and organization.

f. Leave display settings at default.

g. Set the appropriate time zone.

h. Provide your product key.

i. Select an appropriate license mode for your deployment

j. Select “Automatically generate computer name”.

k. Type a default administrator password. If you enable the password reset feature, the users will not actually use this password. This password will be reset by the instance manager after the guest boots up.

l. Leave Network Components at “Typical Settings”.

m. Select the “WORKGROUP” option.

n. Leave Telephony options at default.

o. Select appropriate Regional Settings.

p. Select appropriate language settings.

q. Do not install printers.

r. Do not specify “Run Once commands”.

s. You need not specify an identification string.

t. Save the Answer File as c:\sysprep\sysprep.inf.

3. Run the following command to sysprep the image:

```
c:\sysprep\sysprep.exe -reseal -mini -activated
```

After this step the machine will automatically shut down

12.11. Importing Amazon Machine Images

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.

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.

注意

当拷贝粘贴一条命令，确保在运行前粘贴的命令在一行上．一些文档查看器可能会在拷贝时引入不希望的换行符．

To import an AMI:

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 第 12.13 节 “Adding Password Management to Your Templates”.

13. Unmount and delete loopback mount.

```
# umount /mnt/loop/centos54
# losetup -d /dev/loop0
```

14. Copy the image file to your XenServer host's file-based storage repository. In the example below, the Xenserver is "xenhost". This XenServer has an NFS repository whose uuid is a9c5b8c8-536b-a193-a6dc-51af3e5ff799.

```
# scp CentOS_6.2_x64 xenhost:/var/run/sr-mount/a9c5b8c8-536b-a193-a6dc-51af3e5ff799/
```

15. Log in to the Xenserver and create a VDI the same size as the image.

```
[root@xenhost ~]# cd /var/run/sr-mount/a9c5b8c8-536b-a193-a6dc-51af3e5ff799
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]#  ls -lh CentOS_6.2_x64
-rw-r--r-- 1 root root 10G Mar 16 16:49 CentOS_6.2_x64
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]# xe vdi-create virtual-size=10GiB sr-
uuid=a9c5b8c8-536b-a193-a6dc-51af3e5ff799 type=user name-label="Centos 6.2 x86_64"
cad7317c-258b-4ef7-b207-cdf0283a7923
```

16. Import the image file into the VDI. This may take 10—20 minutes.

```
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]# xe vdi-import filename=CentOS_6.2_x64
 uuid=cad7317c-258b-4ef7-b207-cdf0283a7923
```

17. Locate a the VHD file. This is the file with the VDI’ s UUID as its name. Compress it and upload it to your web server.

```
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]# bzip2 -c cad7317c-258b-4ef7-b207-cdf0283a7923.vhd >
 CentOS_6.2_x64.vhd.bz2
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]# scp CentOS_6.2_x64.vhd.bz2 webserver:/var/www/html/
templates/
```

12.12. 转换Hyper-v 虚拟机为模板

要转换的Hyper-V虚拟机到的XenServer兼容CloudStack模板，您将需要一个独立的XenServer主机连接的NFS VHD的SR。使用任何 CloudStack 支持的XenServer版本，单必须是XenCenter5.6 FP1或SP2(它是向后兼容5.6)。此外，它可以帮助你瓜子一个NFS ISO SR。

对于linux虚拟机，在你使用XenServer中虚拟机进行工作之前，需要为虚拟化做一些准备工作。克隆虚拟机，并且如果你还打算在hyper-v中使用虚拟机。需要卸载hyper-v集组件，同时检查在/etc/fstab中的相关设备

1. 从linux_ic/驱动器/ dist目录中，运行卸载(其中的“linux_ic”是复制的Hyper-V的集成组件文件的路径)。

2. 在/ boot /(备份是名为*。backup0)从备份中恢复原来的initrd。

3. 从/boot/grub/menu.lst.中移除存在的"hdX=noprobe"

4. Check /etc/fstab for any partitions mounted by device name. Change those entries (if any) to mount by LABEL or UUID. You can get that information with the blkid command.

下一步是确保在Hyper-V的虚拟机没有运行，接着应用VHD到XenServer。这里包含有两个选项。

选项1:

1. 通过xencenter导入VHD，工具 虚拟应用工具 磁盘镜像导入

2. 选择 VHD,点击 下一步

3. 为vm命名，选择存储下面的NFS VHD SR,选中 “运行操作系统修复”，并且选择NFS ISO SR

4. 点击下一步 完成。完成虚拟机创建

Option two:

1. 运行 XenConvert ,选择 VHD,选择 XenServer,点击下一步

2. 选择 VHD,点击 下一步

3. 输入XenServer主机信息，点击下一步

4. Name the VM, then click Next, then Convert. A VM should be created.

一旦你完成从hyper-v VHD到虚拟机的创建，准备使用以下步骤：

1. 启动虚拟机，卸载Hyper-V集成服务，并重新启动

2. 安装XenServer的工具，然后重新启动。

3. Prepare the VM as desired. For example, run sysprep on Windows VMs. See 第 12.10 节 “创建WINdows 模板”.

以上任一选项将创建一个VM在HVM模式。为Windows虚拟机这是很好的，但Linux的虚拟机可能无法达到最佳性能。转换的Linux虚拟机，以PV模式，对于不同的发行版本将需要额外的步骤会 。

1. 关闭虚拟机，从NFS存储拷贝虚拟到一个web服务器；比如，在web服务器上挂载NFS共享然后拷贝它，或者是在其他XenServer逐渐向使用sftp或scp上传到web服务器

2. 在CloudStack使用以下值创建模板

   - URL,提供VHD的url

   - os 类型，使用相应的os。对于centos的pv模式，选择 其他PV(32-位)或者其他PV(64-位)。这个选项只能在Xenserver上使用

   - Hypervisor. XenServer

   - 格式化，VHD

这个模板将被创建，你可以通过它来创建实例

12.13. 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.

12.13.1. Linux系统安装

试用一下步骤开始Linux系统的安装

1. 下载脚本文件cloud-set-guest-password

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

   - Windows: http://sourceforge.net/projects/cloudstack/files/Password%20Management%20Scripts/CloudInstanceManager.msi/download

2. 拷贝本文件到Copy this file to /etc/init.d

   在某些linux发行版拷贝此文件到/etc/rc.d/init.d

3. 执行以下命令使脚本可执行

```
chmod +x /etc/init.d/cloud-set-guest-password （这里注意文件路径）
```

4. 根据不同的Linux发行版，选择适当的步骤继续

   在Fedora、CentOS/RHEL、 和 Debian,执行‘

```
chkconfig --add cloud-set-guest-password
```

12.13.2. Windows OS 安装

从 Download page[3] 下载安装程序CloudInstanceManager.msi，并在新创建的Windows 虚拟机中运行安装程序。

12.14. 删除模板

模板可以被删除。在一般情况下，当一个模板跨越多个区域，只为删除被选中的副本将被删除，在其他区域相同的模板将不会被删除。CentOS的模板是一个例外。如果所提供的CentOS的模板被删除，它从所有区域都将被删除。

删除模板时，从它们中产生的虚拟机实例将继续运行。然而，新的虚拟机不能在被删除模板的基础上创建。

[3] http://cloudstack.org/download.html

使用存储

13.1. Storage Overview

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.

13.2. 主存储

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 Installation Guide.

第 2.6 节 “关于主存储”

13.2.1. 对于主存储最好的实践

- 主存储的速度会影响访问性能。如果可以，选择较小，较高RPM驱动的设备给主存储。

- Ensure that nothing is stored on the server. Adding the server to CloudStack will destroy any existing data

13.2.2. 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. 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.

13.2.3. 虚拟机平台对主存储的支持

下表显示了针对不同虚拟机平台的存储选项和参数。

	VMware vSphere	Citrix XenServer	KVM	
磁盘格式,模板,和快照	VMDK	VHD	QCOW2	
支持 iSCSI	VMFS	集群的逻辑卷(LVM)	可以，通过共享挂载点	
支持 Fiber Channel	VMFS	是，通过现有的存储空间。	可以，通过共享挂载点	
NFS 支持	是	是	是	
本地存储支持	是	是	是	

	VMware vSphere	Citrix XenServer	KVM	
基于供应的存储	NFS 和 iSCSI	NFS	NFS	

XenServer在iSCSI和光纤通道存储卷使用集群的LVM系统，而不支持预留空间特性。但是，存储服务器本身可以支持自动精简供应（thin-provisioning）。所以，其结果是CloudStack在自动精简供应的存储卷上仍然可以实现预留空间的特性。

KVM支持“共享挂载点”存储。在一个给定的集群里，共享挂载点是每个服务器的一个文件系统路径。这个路径对所有集群内的主机都必须是一致的，例如：/mnt/primary1。此共享挂载点假定是一个集群文件系统，比如OCFS2。在此情况下， CloudStack不会像对待NFS那样挂载或者卸载存储。 CloudStack要求管理员保证存储已经可用。

CloudStack管理NFS存储的预留空间特性。在这种情况下，有一个全局参数storage.overprovisioning.factor来控制预留空间的程度。这与虚拟机管理器的类型无关。

Local storage is an option for primary storage for vSphere, XenServer, and KVM. 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支持在一个集群中使用多个主存储池。例如，你能在主存储中准备2个NFS服务器，或者最初可以准备一个iSCSI LUN，然后在第一个容量将满时添加第二个iSCSI LUN。

13.2.4. å-#å#¨æ #ç¾

存储可以被打上标签。 字符串的标签使得主存储，磁盘服务，或计算服务相互关联起来。标签使得管理员可以添加关于存储的额外信息。例如它是"SSD" 或它是"slow"。标签不会被CloudStack 解读。标签是用于在计算服务和磁盘服务上相互匹配使用的。在主存储上分配根和数据磁盘之前，CloudStack需要计算和磁盘服务上的标签都标记于主存储上。计算和磁盘服务的标签用于识别 这些服务拥有存储的要求。例如高级的计算服务可能需要"fast"标签 给其跟磁盘卷

标签，分配，跨集群或机架的卷复制之间的关系是很复杂的。简单的环境就是在一个机架内所有集群的主存储使用相同的标签。即使用这些标签表示不同设备，展现出来的标签组仍可以是一样的。

13.2.5. 主存储的维护模式

主存储可以被设置成维护模式。这是有用的，例如，替换在存储设备中坏的RAM。对存储设备的维护模式将首先停止任何新的来自预处理的来宾，然后停止所有有数据卷的来宾。当所有来宾被停止的时候，这个存储设备就进入维护模式了并且可以关机。当存储设备再一次在线的时候，你可以取消维护模式对这个设备。CloudStack将返回在线状态并且试着启动所有曾在这个设备进入维护模式前运行的来宾机器。

13.3. 二级存储

本节给出了关于CloudStack 二级存储的概念和技术细节。关于通过CloudStack 用户界面如何安装和配置二级存储的信息，见高级安装指南。

第 2.7 节 “关与辅助存储”

13.4. Working With Volumes

卷为用户虚拟机提供存储。卷可以作为根分区或附加硬盘。CloudStack 支持为用户虚拟机添加卷。

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, for example:vSphere, KVM. 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, for example: "/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.

注意

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

13.4.1. 创建新卷

你可以在符合你存储能力的情况下随时向来宾虚拟机添加多个数据卷。CloudStack 的管理员和普通用户都可以向虚拟机实例中添加卷。当你创建了一个新卷，他以一个实体的形式存在于CloudStack但是在你将其附加到实例中之前他并不会被分配实际的物理空间。这种优化会使CloudStack将磁盘空间提供给刚需的用户。

13.4.1.1. Using Local Storage for Data Volumes

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

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

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

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

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

13.4.1.2. To Create a New Volume

1. 以管理员或普通用户身份登录进入CloudStack用户界面。

2. 在左侧导航栏点击存储。

3. 在选择视图选择卷。

4. 点击添加卷来创建一个新卷，填写以下信息后点击确定。

   - 名字。给卷取个独特名字以便于你以后找到他。

   - 可用的资源域。你想让这个个存储在哪个地方有效？这个应该接近要是用这个卷的VM。（就是说你要 在单个资源域内使用这个存储就选择单个资源域，如果此存储要在多个资源与内共享你就选所有资源域）

   - 磁盘方案。选择存储特性。

   新建的存储会在卷列表中显示“已分配”状态。卷数据已经存储到CloudStack了但是他还不能使用。

5. 通过附加卷来开始使用这个卷。

13.4.2. Uploading an Existing Volume to a Virtual Machine

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

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

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

To upload a volume:

1. (Optional) Create an MD5 hash (checksum) of the disk image file that you are going to upload. After uploading the data disk, 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. 在左侧导航栏点击存储。

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.

è##æ##æ#ºç®¡ç##ç¨#åº#	Disk Image Format
XenServer	VHD
VMware	OVA
KVM	QCOW2

- 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 ???, and make sure the status is Uploaded.

13.4.3. 附加卷

您能通过附加一个卷来提供虚拟机的额外磁盘存储。当你第一次创建新卷，或移动已存在的卷到另一台虚拟机，或实在从另一个存储池迁移过来一个卷的时候你才可以附加一个卷。

1. 以管理员身份登录进入 CloudStack 用户界面。

2. 在左侧导航栏点击存储。

3. 在选择视图选择卷。

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

5. 在弹出的实例界面，选择你打算附加卷到哪台虚拟机。你只能看到允许你附加卷的实例。比如，普通用户只能看到他自己创建的实例，而管理员将会有更多的选择。

6. 当卷被附加之后你通过点击实例看到实例名和该实例所附加的卷。

13.4.4. Detaching and Moving Volumes

注意

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

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.

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

1. 以管理员或普通用户身份登录进入CloudStack用户界面。

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 第 13.4.3 节 “附加卷”.

13.4.5. 虚拟机存储迁移

支持xenserver,kvm,vmware

> **注意**
>
> 这个过程有别于移动磁盘卷从一个虚拟机到另外的虚拟机。查看挂载和移动卷

你可以迁移一个虚拟机的根磁盘卷或者附加数据磁盘卷从同一区域中一个存储池到另外一个存储池。

你可以使用磁盘迁移特性完成一些常用的管理目标。如存储池的负载平衡和增加虚拟机的可靠性，帮助他们远离任何存储池常见问题。

13.4.5.1. Migrating a Data Disk Volume to a New Storage Pool

1. 以管理员或普通用户身份登录进入CloudStack用户界面。

2. Detach the data disk from the VM. See Detaching and Moving Volumes 第 13.4.4 节 “Detaching and Moving Volumes” (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 第 13.4.3 节 “附加卷”

13.4.5.2. 迁移虚拟机根卷到新的存储池

迁移根卷的时候虚拟机要先停止，此时用户将不能访问该虚拟机。迁移结束方可重启虚拟机。

1. 以管理员或用户身份登陆至CloudStack用户界面。

2. 取消附加在该虚拟机上的数据磁盘。详见 取消附加并移除卷第 13.4.4 节 “Detaching and Moving Volumes” （但是跳过最后的“重新附加”步骤你会在前几结束之后做这步）。

3. 停止虚拟机。

4. Use 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. 看着虚拟机状态由“迁移中”恢复为“已停止”。

6. 重新启动该虚拟机。

13.4.6. 调整卷大小

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

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

This feature is supported on KVM, XenServer, and VMware hosts. However, shrinking volumes is not supported on VMware hosts.

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

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

To resize a volume:

1. 以管理员或普通用户身份登录进入CloudStack用户界面。
2. 在左侧导航栏点击存储。
3. 在选择视图选择卷。
4. Select the volume name in the Volumes list, then click the Resize Volume button
5. In the Resize Volume pop-up, choose desired characteristics for the storage.

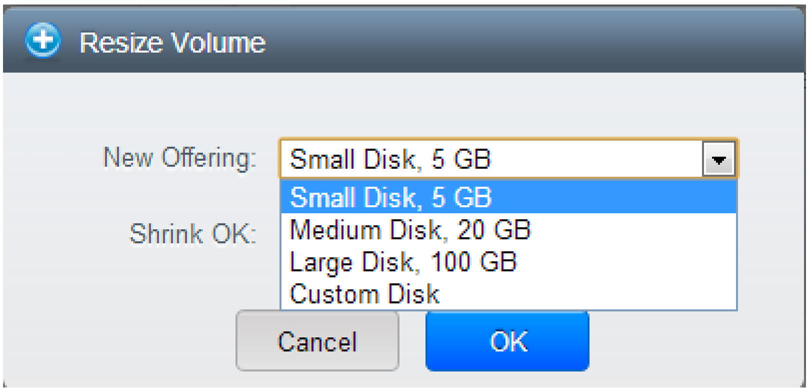

   a. If you select Custom Disk, specify a custom size.

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

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

6. 点击 确定。

13.4.7. 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.

13.5. 使用快照

（支持以下Hypervisor。 XenServer，VMware vSphere，和 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.

Snapshots may be taken for volumes, including both root and data disks. 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.

用户能手动或设置自动循环创建快照策略。用户也能通过快照创建磁盘卷，并向其他盘一样附加到虚拟机中。支持对root和data盘做快照。然而CloudStack当前不支持通过恢复的root盘启动。通过快照恢复的root盘会被当作data盘对待，但是恢复的到的盘能够挂载到虚拟机中。

完整快照会从主存储拷贝到二级存储直到被删除或被新的快照覆盖。

13.5.1. Snapshot Job Throttling

When a snapshot of a virtual machine is requested, the snapshot job runs on the same host where the VM is running or, in the case of a stopped VM, the host where it ran last. If many snapshots are requested for VMs on a single host, this can lead to problems with too many snapshot jobs overwhelming the resources of the host.

To address this situation, the cloud's root administrator can throttle how many snapshot jobs are executed simultaneously on the hosts in the cloud by using the global configuration setting concurrent.snapshots.threshold.perhost. By using this setting, the administrator can better ensure that snapshot jobs do not time out and hypervisor hosts do not experience performance issues due to hosts being overloaded with too many snapshot requests.

Set concurrent.snapshots.threshold.perhost to a value that represents a best guess about how many snapshot jobs the hypervisor hosts can execute at one time, given the current resources of the hosts and the number of VMs running on the hosts. If a given host has more snapshot requests, the additional requests are placed in a waiting queue. No new snapshot jobs will start until the number of currently executing snapshot jobs falls below the configured limit.

The admin can also set job.expire.minutes to place a maximum on how long a snapshot request will wait in the queue. If this limit is reached, the snapshot request fails and returns an error message.

13.5.2. Automatic Snapshot Creation and Retention

（支持以下Hypervisor。 XenServer，VMware vSphere，和 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 interval. 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 第 14.3 节 “Globally Configured Limits” . 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.

13.5.3. 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

13.5.4. 卷状态

快照操作时由一个经常性的快照策略所引发，如果从其上次创建快照后，卷一直处于非活跃状态，快照被跳过。卷被认为是非活跃的，如果它被分离或附加的虚拟机没有运行。CloudStack确保从卷上一次变得不活跃后，至少有一个快照创建。

手动创建快照时，快照总是创建，无论卷活跃与否。

13.5.5. 快照恢复

有两种方式恢复快照。用户能够从快照中创建一个卷。卷可以随后被装载到虚拟机上并且文件根据需要被复原。另一种方式是，模板可以从一个root 盘的快照创建。用户能够从这个模板启动虚拟机从而实际上复原root盘。

第 14

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.

14.1. 配置使用服务器

配置使用服务器之前：

1. 确定你已经安装了使用服务器。这需要你在CloudStack 软件之外的其他步骤。详情请见高级安装指南中的 安装使用服务器（可选）部分。

2. 以管理员身份登录进入CloudStack 用户界面。

3. 点击 全局设置。

4. 在搜索栏输入 usage。找到no需要改变的配置参数。下表是这些参数的详细描述。

5. 在操作栏点击编辑图标。

6. 输入数值点击保存图标。

7. 重新启动管理服务器（通常在改变了全局配置之后都要进行这步）并重启使用服务器。

```
# service cloudstack-management restart
# service cloudstack-usage restart
```

下表流出了全局配置中控制使用服务器的配置项。

变量名	è¯´æ##
enable.usage.server	是否开启使用服务器。

usage.aggregation.timezone	记录使用信息所用的时区。如果使用记录和每日的任务执行在不同时区的话需要设置本参数。例如，使用以下设置会使使用作业在PST 00:15运行并生成从GMT时间0点到23:59:59 的二十四小时使用记录。
	`usage.stats.job.exec.time = 00:15`
	`usage.execution.timezone = PST`
	`usage.aggregation.timezone = GMT`
	Valid values for the time zone are specified in 附录 A, 时区
	默认：GMT

变量名	è¯´æ##
usage.execution.timezone	The time zone of usage.stats.job.exec.time. Valid values for the time zone are specified in 附录 A, 时区 默认时区是管理服务器的时区。
usage.sanity.check.interval	完整性检查时间间隔。设置此值来定期在生成用户账单之前检查出错误的数据。比如，他能检查虚拟机被销毁之后的使用记录和模板，卷等的类似记录。还会检查超过聚合时间的使用时间。如果发生了错误就会发送 ALERT_TYPE_USAGE_SANITY_RESULT = 21 警告。
usage.stats.job.aggregation.range	使用服务器执行任务时间间隔（分钟为单位）。比如，如果你将此值设为1440，使用服务器将每天执行一次。如果你将此值设为600，则会10小时执行一次。一般情况下使用服务器执行任务时会继续在上次的使用统计基础上处理所有事件。 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 usage.execution.timezone. 默认：1440
usage.stats.job.exec.time	The time when the Usage Server processing will start. It is specified in 24-hour format (HH:MM) in the time zone of the server, which should be GMT. For example, to start the Usage job at 10:30 GMT, enter “10:30”. 如果同时设置了 usage.stats.job.aggregation.range参数，并且该参数值不是1440，这个值将被添加到 usage.stats.job.exec.time到时再次运行使用服务器任务。重复此过程，直到24小时已经过去，第二天到达 usage.stats.job.exec.time处理任务开始。 默认值：00:15

比如，假设你服务器时区是GMT，你的用户主要在美国东海岸，而你有打算在当地时间（EST）每天凌晨两点执行使用记录统计。选择这些选项：

- enable.usage.server = true

- usage.execution.timezone = America/New_York
- usage.stats.job.exec.time = 07:00。这将会在东部时间2：00执行使用任务。注意，这会由于东部时间进入退出夏令时间而改变一小时。
- usage.stats.job.aggregation.range = 1440

在这种配置下，使用任务 会在东部时间每天凌晨两点执行，同时会如定义的一样以东部时间统计前一天的“午夜到午夜使用记录。”

注意

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.

14.2. 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.

This section covers the following topics:

14.3. 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

Parameter Name	Definition
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

14.4. Default Account Resource Limits

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

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

1. 登入到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.

14.5. 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. 登陆 CloudStack 用户界面

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

管理网络和流量

在CloudStack ; 来宾vms可以通过安全共享架构和其他人通讯，用户感觉像拥有自己的私有网络一样。CloudStack虚拟路由器是为用户提供网络通讯流量的主要组件

15.1. æ#¥å®¾æµ#é##

一个网络只能为在一个zone中的虚拟机之间传输流量。不同zone中的虚拟机不能使用它们自己的ip地址和对方通信；它们必须通过公共ip地址以路由的方式和对方通信。

This figure illustrates a typical guest traffic setup:

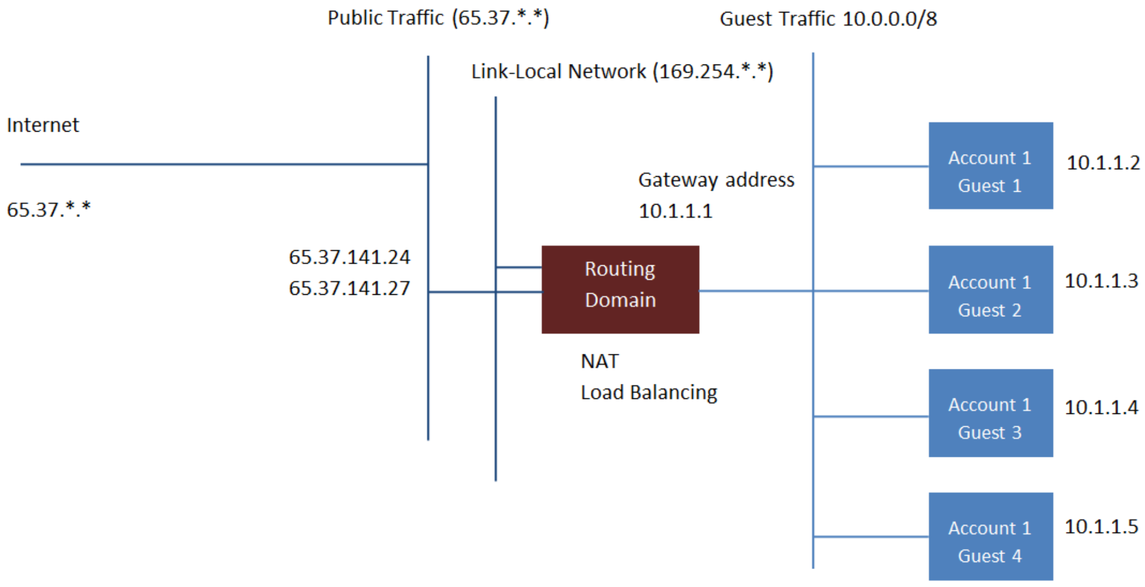

Guest Traffic Setup

管理服务器自动的为每个网络创建一个虚拟路由器。虚拟路由器是运行在主机上的特殊虚拟机，它拥有3个网络接口。eth0接口作为来宾网络的网关，ip地址为10.1.1.1；eth1接口用于系统配置虚拟路由器；eth2接口指定一个公共ip地址传输公共流量。

虚拟路由器提供dhcp服务并且自动将网络保留的ip地址分配一个给虚拟机。用户可以人工的重新配置ip地址

虚拟路由器自动的配置源nat为来宾vm转发外部流量

15.2. 一个POD内的网络

The figure below 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.

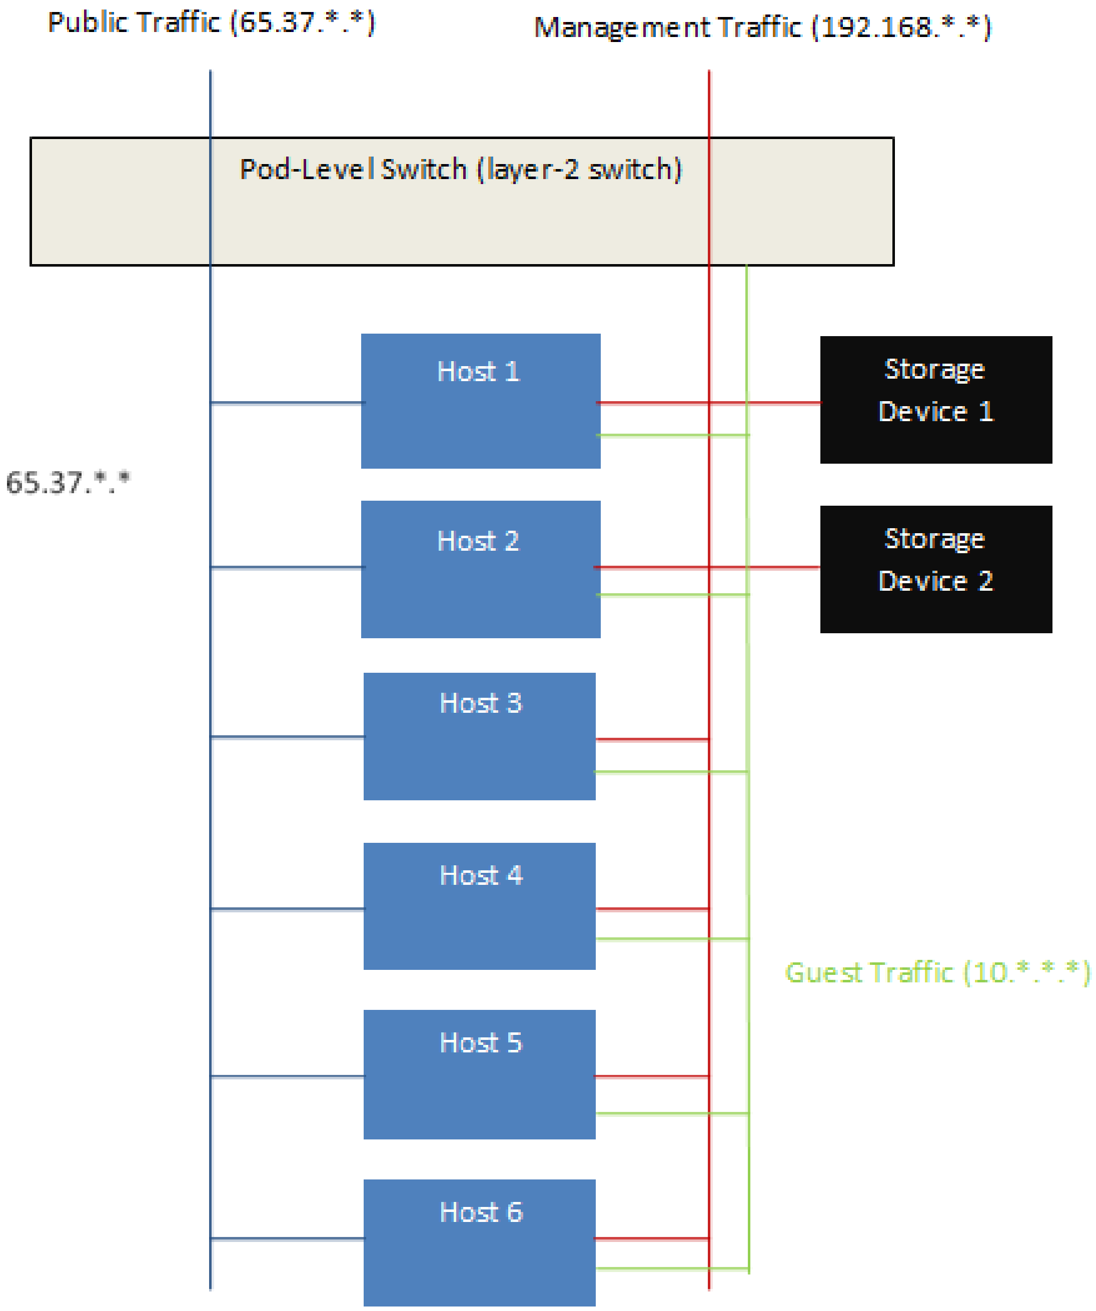

Network Setup within a Single Pod – Logical View

服务器像如下连接:

- 存储设备只连接到携带管理流量的网络.
- 主机连接到管理流量及公共流量的网络.
- 主机也连接到一个或多个携带客户流量的网络.

我们建议使用多个物理网卡来实现每个网络接口，就像冗余的交换机一样，以此来保证最大的吞吐并且改善可靠性.

15.3. 在区域内的联网

The following figure illustrates the network setup within a single zone.

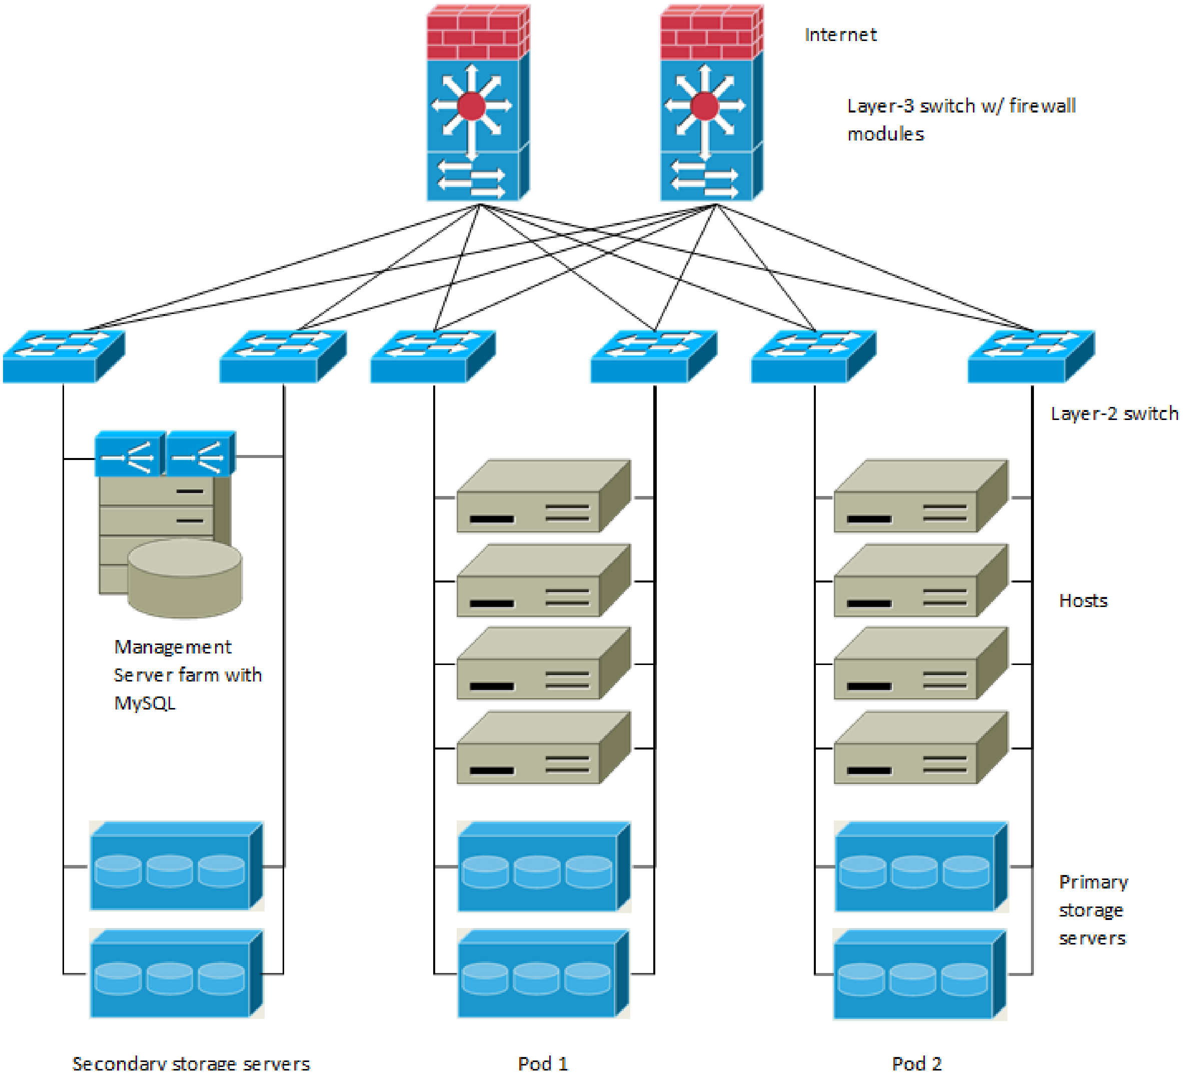
Internet
Layer-3 switch w/ firewall modules
Layer-2 switch
Management Server farm with MySQL
Hosts
Primary storage servers
Secondary storage servers
Pod 1
Pod 2

用于管理流量的防火墙工作在NAT模式。通常是分配给网络中的IP地址192.168.0.0/16 B类私有地址空间。每个POD分配的IP地址192.168.*.0/24 C类私有地址空间。

每个区域都有自己的一套公网IP地址。来自不同区域的公网IP地址不重叠。

15.4. 基本区域物理网络配置

在一个基本网络中，配置物理网络相当直接明了。你只需要配置一个宾客网络以承载由宾客虚拟机产生的流量。当你首次增加一个区域到CloudStack中， 你通过Add Zone 屏设置宾客网络。

15.5. 高级区物理网络配置

在使用高级联网的区中，你需要告诉管理服务器物理网络是如何安装的以独立地传输不同种类的网络流量。

15.5.1. 在高级资源域中设置来宾流量

以下步骤假定你已经登录进入 CloudStack 界面。设置基本的来宾网络:

1. 在左边的导航栏，选择基础架构。在区域数量界面点击查看全部，然后点击你要添加网络的区域名。

2. 点击网络标签。

3. 点击 添加客户网络。

   添加来宾网络窗口显示:

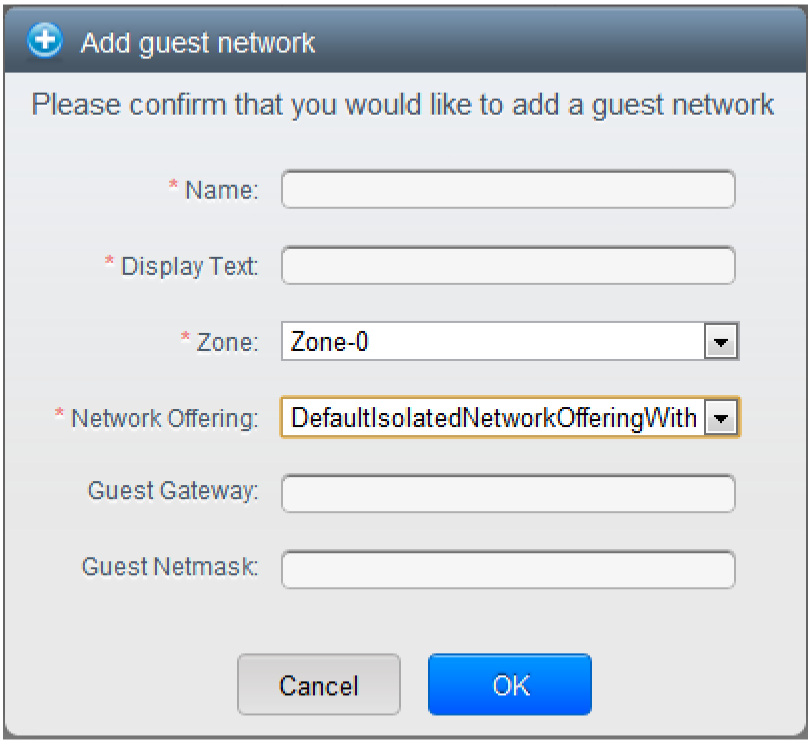

Add guest network

Please confirm that you would like to add a guest network

* Name:

* Display Text:

* Zone: Zone-0

* Network Offering: DefaultIsolatedNetworkOfferingWith

Guest Gateway:

Guest Netmask:

4. 提供以下信息

   - 名称。网络名称。这是对用户可见的。
   - 显示文本:
   - 区域: 这里的区域是指你要配置来宾网络的区域。
   - 网络提供方案: 如果管理员已经配置了多个网络i共方案, 选择你需要的那个。
   - 来宾网络网关: 来宾网络使用的网关
   - 来宾网络子网掩码: 来宾网络使用的子网掩码

5. 点击 确定。

15.5.2. 在高级区中配置公用通信

在一个使用高级网络配置的区中, 你需要配置至少一个用于Internet通信的IP地址范围。

15.6. 使用多个来宾网络

在使用高级网络的资源域里, 附加的来宾网络可以在初始安装后任何时间添加. 你还可以通过指定DNS后缀为每个网络自定义关联的域名.

一个虚机的网络在其创建时定义. 当一个虚机创建以后, 就不能对其网络添加删除, 即便用户可以进入客户虚机删除指定网卡的IP地址.

每一个虚机只有一个默认的网络．在这个默认网络里，虚拟路由器的DHCP响应将设置客户的默认网关．除了单一,必须的默认网络，多个非默认的网络也可以添加到客户虚机里．管理员可以控制哪个网络作为默认的网络．

附加的网络可以给所有账户使用或者分配给特定的账户．对所有账户都可用的网络在整个资源域有效．任何可以访问这个资源域的用户都可以使用这个网络创建虚机．这些资源域一级的网络基本不提供客户之间的隔离．分配给特定帐户的网络提供强制隔离的功能．

15.6.1. 添加另一个虚拟机的网络

1. 作为管理员或最终用户登入到CloudStack UI．
2. 在左边的导航栏里选择网络．
3. 点击添加虚拟机网络。系统给出如下信息：
   - 名称：网络名称。用户可见。
   - 显示文本：网络描述。用户可见。
   - 区域．此网络所应用的区域名。每个区域都是一个广播范围，因此对虚拟机网络来说，每个区域都有一个不同的IP范围。管理员必须为每个区域配置IP范围。
   - 网络提议：如果管理员已经配置了多个网络提议，为此网络从中选择一个。
   - 虚拟机网关：虚拟机应该使用的网关。
   - 虚拟机掩码：虚拟机子网说使用的掩码。
4. 点击创建。

15.6.2. 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.

- 作为管理员或最终用户登入到CloudStack UI．
- 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 第 10.6 节 “停止和启动虚拟机”．
- 在左边的导航栏里选择网络．
- Click the name of the network you want to modify.
- In the Details tab, click Edit.
- In Network Offering, choose the new network offering, then click Apply.
- A prompt is displayed 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.
- Wait for the update to complete. Don’t try to restart VMs until the network change is complete.
- If you stopped any VMs, restart them.

15.7. 安全分组

15.7.1. About Security Groups

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

> 注意
>
> 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.

15.7.2. Adding a Security Group

A user or administrator can define a new security group.

1. 作为管理员或最终用户登入到CloudStack UI.

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. 点击 确定。

   The new security group appears in the Security Groups Details tab.

7. To make the security group useful, continue to Adding Ingress and Egress Rules to a Security Group.

15.7.3. Security Groups in Advanced Zones (KVM Only)

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

Limitations

The following are not supported for this feature:

- Two IP ranges with the same VLAN and different gateway or netmask in security group-enabled shared network.

- Two IP ranges with the same VLAN and different gateway or netmask in account-specific shared networks.

- Multiple VLAN ranges in security group-enabled shared network.

- Multiple VLAN ranges in account-specific shared networks.

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

15.7.4. 启用安全组

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. The administrator can not enable security groups for an existing zone, only when creating a new zone.

15.7.5. Adding Ingress and Egress Rules to a Security Group

1. 作为管理员或最终用户登入到CloudStack UI.

2. In the left navigation, choose Network

3. In Select view, choose Security Groups, then click the security group you want .

4. To add an ingress rule, click the Ingress Rules tab and fill out the following fields to specify what network traffic is allowed into VM instances in this security group. If no ingress rules are specified, then no traffic will be allowed in, except for responses to any traffic that has been allowed out through an egress rule.

   - Add by CIDR/Account. Indicate whether the source of the traffic will be defined by IP address (CIDR) or an existing security group in a 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:

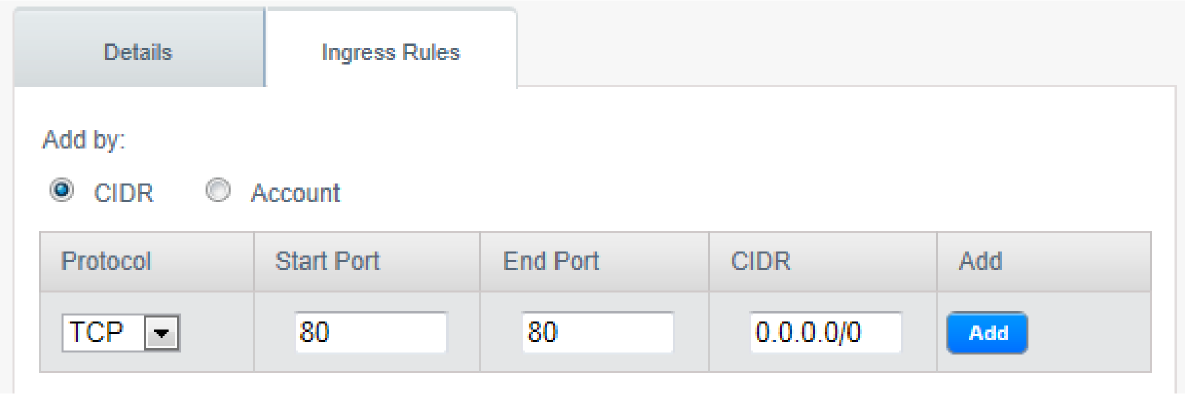

Add by:

● CIDR ○ Account

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. 点击添加。

15.8. 外部防火墙和负载均衡器

CloudStack 能够用外部的Juniper SRX设备和可选的外部NetScaler 或 F5 负载均衡器替换它的虚拟路由器，用于网管和负载均衡服务。在这种情况下，虚拟机使用， 虚拟机使用SRX 作为它们的网关。

15.8.1. 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 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

NetScaler ADC Type	Description of Capabilities	CloudStack Supported Features
		administrator is required. Once a VPX instance is added into CloudStack, it is treated the same as a VPX on an ESXi host.

15.8.2. 配置RHEL服务器上的snmp通信组

SNMP团体字符串是类似的用户ID或密码，提供了访问网络设备，如路由器。此字符串的所有SNMP请求一起发送。如果团体字符串是正确的，设备响应请求的信息。如果团体字符串是不正确的，则丢弃请求没有回应。

NetScaler设备使用SNMP和虚拟机通信。你必须安装SNMP，并且配置SNMP通信组用于netscaler和RHEL机器之间的安全通信。

1. 确保你的RedHat上安装了SNMP,如果没有执行以下命令:

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

2. 编辑 /etc/snmp/snmpd.conf 允许SNMP测试netscaler 设备

   a. 映射团体名到一个安全组(本地和网络，依赖从哪里发来请求):

   注意

   当你编辑以下表格，使用强壮密码替换public

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

   注意

   设置0.0.0.0 允许所有ip查询netscaler服务器

   b. 映射安全名到组名

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

   c. 创建视图授予组以下权限:

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

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

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

3. iptalbes 中放行SNMP

```
iptables -A INPUT -p udp --dport 161 -j ACCEPT
```

4. 启动SNMP服务

```
service snmpd start
```

5. 确保snmp服务随系统自动启动

```
chkconfig snmpd on
```

15.8.3. 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).

15.8.4. 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.

15.8.5. Configuring AutoScale

AutoScaling allows you to scale your back-end services or application VMs up or down seamlessly and automatically according to the conditions you define. With AutoScaling enabled, you can ensure that the number of VMs you are using seamlessly scale up when demand increases, and automatically decreases when demand subsides. Using AutoScaling, you can automatically shut down instances you don't need, or launch new instances, depending on demand.

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

CloudStack uses the NetScaler load balancer to monitor all aspects of a system's health and work in unison with CloudStack to initiate scale-up or scale-down actions.

> 注意
>
> AutoScale is supported on NetScaler Release 10 Build 73.e and beyond.

先决条件

Before you configure an AutoScale rule, consider the following:

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

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

- Deploy the templates you prepared. Ensure that the applications come up on the first boot and is ready to take the traffic. Observe the time requires to deploy the template. Consider this time when you specify the quiet time while configuring AutoScale.

- The AutoScale feature supports the SNMP counters that can be used to define conditions for taking scale up or scale down actions. To monitor the SNMP-based counter, ensure that the SNMP agent is installed in the template used for creating the AutoScale VMs, and the SNMP operations work with the configured SNMP community and port by using standard SNMP managers. For example, see 第 15.8.2 节 "配置RHEL服务器上的snmp通信组" to configure SNMP on a RHEL machine.

- Ensure that the endpointe.url parameter present in the Global Settings is set to the Management Server API URL. For example, http://10.102.102.22:8080/client/api. In a multi-node Management Server deployment, use the virtual IP address configured in the load balancer for the management server's cluster. Additionally, ensure that the NetScaler device has access to this IP address to provide AutoScale support.

  If you update the endpointe.url, disable the AutoScale functionality of the load balancer rules in the system, then enable them back to reflect the changes. For more information see Updating an AutoScale Configuration

- If the API Key and Secret Key are regenerated for an AutoScale user, ensure that the AutoScale functionality of the load balancers that the user participates in are disabled and then enabled to reflect the configuration changes in the NetScaler.

- In an advanced Zone, ensure that at least one VM should be present before configuring a load balancer rule with AutoScale. Having one VM in the network ensures that the network is in implemented state for configuring AutoScale.

云平台配置

Specify the following:

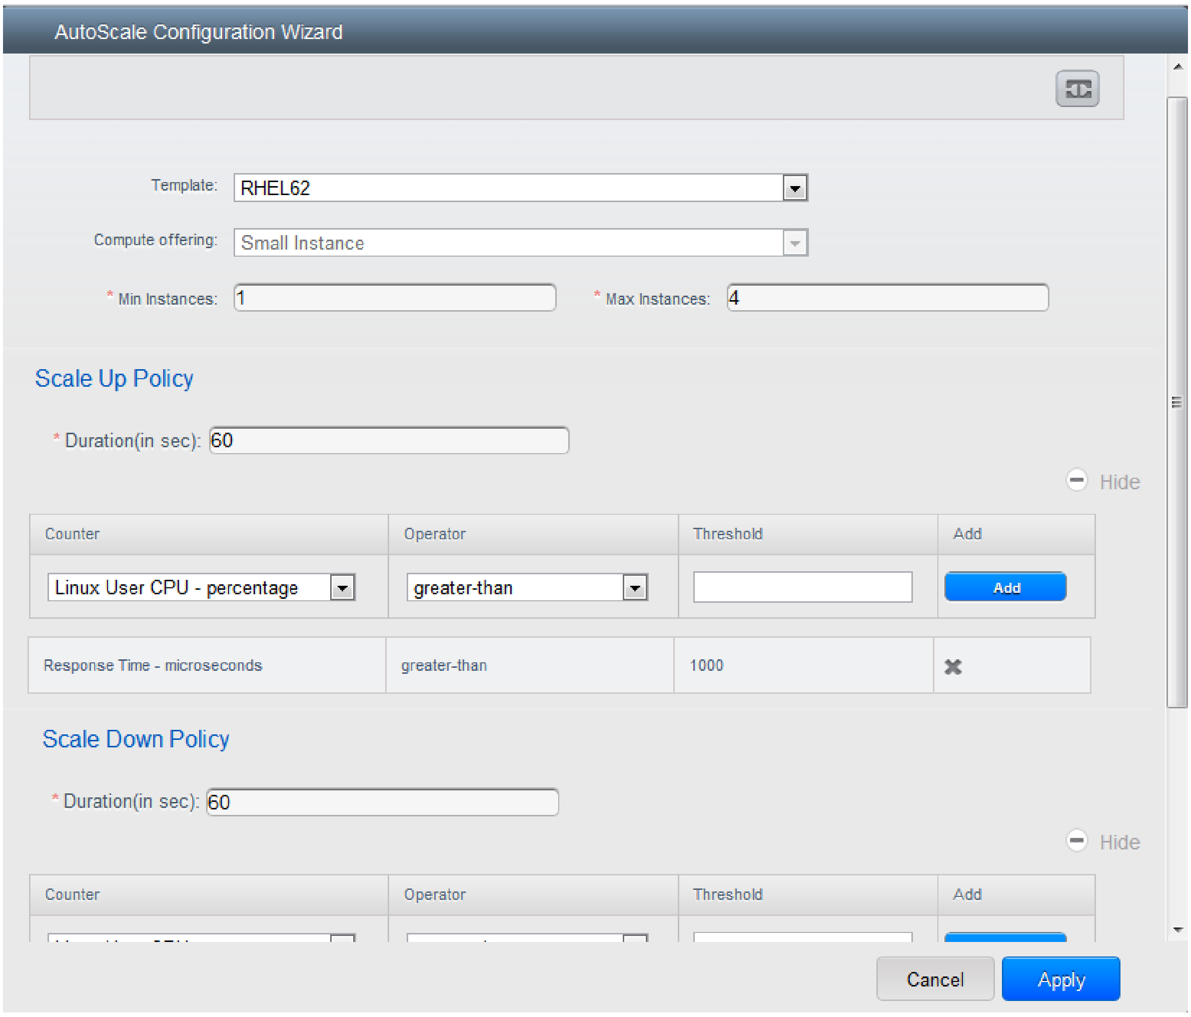

AutoScale Configuration Wizard

Template: RHEL62

Compute offering: Small Instance

* Min Instances: 1

* Max Instances: 4

Scale Up Policy

* Duration(in sec): 60

Counter	Operator	Threshold	Add
Linux User CPU - percentage	greater-than		Add
Response Time - microseconds	greater-than	1000	✖

Scale Down Policy

* Duration(in sec): 60

Counter	Operator	Threshold	Add

- Template: A template consists of a base OS image and application. A template is used to provision the new instance of an application on a scaleup action. When a VM is deployed from a template, the VM can start taking the traffic from the load balancer without any admin intervention. For example, if the VM is deployed for a Web service, it should have the Web server running, the database connected, and so on.

- Compute offering: A predefined set of virtual hardware attributes, including CPU speed, number of CPUs, and RAM size, that the user can select when creating a new virtual machine instance. Choose one of the compute offerings to be used while provisioning a VM instance as part of scaleup action.

- Min Instance: The minimum number of active VM instances that is assigned to a load balancing rule. The active VM instances are the application instances that are up and serving the traffic, and are being load balanced. This parameter ensures that a load balancing rule has at least the configured number of active VM instances are available to serve the traffic.

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

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

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

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

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

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

- Add: Click Add to add the condition.

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

- Polling interval: Frequency in which the conditions, combination of counter, operator and threshold, are to be evaluated before taking a scale up or down action. The default polling interval is 30 seconds.

- Quiet Time: This is the cool down period after an AutoScale action is initiated. The time includes the time taken to complete provisioning a VM instance from its template and the time taken by an application to be ready to serve traffic. This quiet time allows the fleet to come up to a stable state before any action can take place. The default is 300 seconds.

- Destroy VM Grace Period: The duration in seconds, after a scaledown action is initiated, to wait before the VM is destroyed as part of scaledown action. This is to ensure graceful close of any pending sessions or transactions being served by the VM marked for destroy. The default is 120 seconds.

- Security Groups: Security groups provide a way to isolate traffic to the VM instances. A security group is a group of VMs that filter their incoming and outgoing traffic according to a set of rules, called ingress and egress rules. These rules filter network traffic according to the IP address that is attempting to communicate with the VM.

- Disk Offerings: A predefined set of disk size for primary data storage.

- SNMP Community: The SNMP community string to be used by the NetScaler device to query the configured counter value from the provisioned VM instances. Default is public.

- SNMP Port: The port number on which the SNMP agent that run on the provisioned VMs is listening. Default port is 161.

- User: This is the user that the NetScaler device use to invoke scaleup and scaledown API calls to the cloud. If no option is specified, the user who configures AutoScaling is applied. Specify another user name to override.

- Apply: Click Apply to create the AutoScale configuration.

Disabling and Enabling an AutoScale Configuration

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

To disable the AutoScale configuration, click the Disable AutoScale button.

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

Updating an AutoScale Configuration

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

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

Runtime Considerations

- An administrator should not assign a VM to a load balancing rule which is configured for AutoScale.

- Before a VM provisioning is completed if NetScaler is shutdown or restarted, the provisioned VM cannot be a part of the load balancing rule though the intent was to assign it to a load balancing rule. To workaround, rename the AutoScale provisioned VMs based on the rule name or ID so at any point of time the VMs can be reconciled to its load balancing rule.

- Making API calls outside the context of AutoScale, such as destroyVM, on an autoscaled VM leaves the load balancing configuration in an inconsistent state. Though VM is destroyed from the load balancer rule, NetScaler continues to show the VM as a service assigned to a rule.

15.9. 负载均衡规则

本产品用户或管理员能创建来自于公共IP到一个或多个虚拟机之间流量的负载均衡规则。用户创建规则指定算法，然后将规则分配给一个虚拟机集合。

> **注意**
>
> 如果你创建的负载均衡规则同时使用的网络提供方案使用到了外部的负载均衡器设备（比如NetScaler），随后又改变为使用虚拟路由器的网络提供方案，你必须在虚拟路由器上为每个已存在的负载聚恒规则创建一个防火墙规则以使它们继续生效。

15.9.1. Adding a Load Balancer Rule

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. Click the name of the network where you want to load balance the traffic.

4. 点击查看IP地址.

5. Click the IP address for which you want to create the rule, then click the Configuration tab.

6. In the Load Balancing node of the diagram, click View All.

In a Basic zone, you can also create a load balancing rule without acquiring or selecting an IP address. CloudStack internally assign an IP when you create the load balancing rule, which is listed in the IP Addresses page when the rule is created.

To do that, select the name of the network, then click Add Load Balancer tab. Continue with 7.

7. Fill in the following:

   - Name: A name for the load balancer rule.
   - Public Port: The port receiving incoming traffic to be balanced.
   - Private Port: The port that the VMs will use to receive the traffic.
   - Algorithm: Choose the load balancing algorithm you want 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.
   - AutoScale: Click Configure and complete the AutoScale configuration as explained in 第 15.8.5 节 "Configuring AutoScale" .

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.

15.9.2. 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.

15.10. 宾客IP范围

宾客网络流量的IP是由用户以帐号为基础设置的。这允许用户以在他们的宾客网络和他们的客户端之间开通VPN连接的方式配置他们的网络。

15.11. 获得一个新的IP地址

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. 点击你想要与之工作的网络名称.

4. 点击查看IP地址.

5. 点击获得一个新IP，并且在确认的对话框中点击确定.

   你被要求点击确认是因为，通常IP地址是有限的资源. 在稍等片刻之后，新的IP地址将会出现并且状态是已分配. 你现在可以使用这个IP地址进行端口转发或静态NAT规则.

15.12. Releasing an IP Address

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

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. 点击你想要与之工作的网络名称.

4. 点击查看IP地址.

5. Click the IP address you want to release.

6. Click the Release IP button. ✖

15.13. 静态 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.

15.13.1. Enabling or Disabling Static NAT

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

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

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. 点击你想要与之工作的网络名称.

4. 点击查看IP地址.

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.

15.14. IP转发及防火墙

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

To allow outgoing traffic, follow the procedure in 第 15.14.1 节 "Creating Egress Firewall Rules in an Advanced Zone" .

To allow incoming traffic, users may set up firewall rules and/or port forwarding rules. For example, you can use a firewall rule to open a range of ports on the public IP address, such as 33 through 44. Then use port forwarding rules to direct traffic from individual ports within that range to specific ports on user VMs. For example, one port forwarding rule could route incoming traffic on the public IP's port 33 to port 100 on one user VM's private IP. For more information, see 第 15.14.2 节 "防火墙规则" and 第 15.14.3 节 "ç«¯å#£è½¬å##" .

15.14.1. Creating Egress Firewall Rules in an Advanced Zone

注意

The egress firewall rules are supported only on virtual routers.

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

Consider the following scenarios to apply egress firewall rules:

- Allow the egress traffic from specified source CIDR. The Source CIDR is part of guest network CIDR.

- Allow the egress traffic with destination protocol TCP,UDP,ICMP, or ALL.

- Allow the egress traffic with destination protocol and port range. The port range is specified for TCP, UDP or for ICMP type and code.

To configure an egress firewall rule:

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. In Select view, choose Guest networks, then click the Guest network you want.

4. To add an egress rule, click the Egress rules tab and fill out the following fields to specify what type of traffic is allowed to be sent out of VM instances in this guest network:

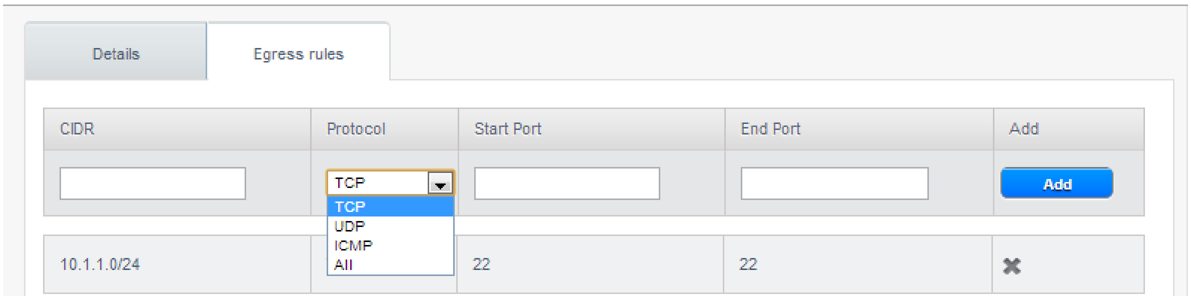

CIDR	Protocol	Start Port	End Port	Add
	TCP			Add
10.1.1.0/24		22	22	✖

- CIDR: (Add by CIDR only) To send traffic only to the IP addresses within a particular address block, enter a CIDR or a comma-separated list of CIDRs. The CIDR is the base IP address of the destination. For example, 192.168.0.0/22. To allow all CIDRs, set to 0.0.0.0/0.

- Protocol: The networking protocol that VMs uses to send outgoing traffic. The TCP and UDP protocols are typically used for data exchange and end-user communications. The ICMP protocol is typically used to send error messages or network monitoring data.

- Start Port, End Port: (TCP, UDP only) A range of listening ports that are the destination for the outgoing traffic. If you are opening a single port, use the same number in both fields.

- ICMP Type, ICMP Code: (ICMP only) The type of message and error code that are sent.

5. 点击添加.

15.14.2. 防火墙规则

默认情况下，防火墙拒绝所有流入公共IP的流量。要是需要允许外部流量，你需要通过制定防火墙规则打开防火墙端口。你可以选择性的制定一个或多个CIDR（无类别域际路由选择，这个说法太晦涩，你可以理解为一个IP网络）来过滤来源IP。这在当你只需要允许特定IP请求时会很有用。

你不能使用防火墙规则打开弹性IP端口。当弹性IP处在使用状态时，外部的通过请求将被安全组管理。详情查看第 15.7.2 节 “Adding a Security Group” 。

In an advanced zone, you can also create egress firewall rules by using the virtual router. For more information, see 第 15.14.1 节 “Creating Egress Firewall Rules in an Advanced Zone” .

防火墙规则能在管理服务器的web UI的防火墙表里创建，这个规则表默认不显示，你需要以管理员身份修改全局变量 firewall.rule.ui.enabled为 "true"才能显示。

如果你想创建防火墙规则:

1. 以用户或管理员身份登录进入WebUI。

2. 在左边的导航，选择网络

3. 选择你要修改的网络名称

4. 点击 查看IP地址

5. 点击你要修改的IP地址

6. 点击设置标签 填入以下值

   - 来源CIDR。（可选）只允许在某个特定地址块的IP流量的话，输入一个CIDR或者一个逗号隔开的CIDR列表。例如：192.168.0.0/22或者192.168.0.0/24,192.168.1.0/24,192.168.2.0/24。留空则为允许所有的CIDR。

   - 协议. 你所要开反复端口所使用的网络协议。

   - 起始端口和结束端口. 你想要在防火墙开放的端口。如果你只打开单个端口，在两个空格中填入相同的端口号。

   - ICMP 类型和编号.。只有在设置ICMP协议时才会用到。提供需要填写的ICMP协议，ICMP头的类型和代码。如果你不知道天什么请参看ICMP文档。（这里推荐一个 http://wenku.baidu.com/view/e235e8ecaeaad1f346933fed.html）

7. 点击添加.

15.14.3. ç«¯å#£è½¬å##

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

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

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

To set up port forwarding:

1. 登陆到CloudStack界面以管理员或者终端用户。

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. Choose an existing IP address or acquire a new IP address. See 第 15.11 节 “获得一个新的IP地址”. Click the name of the IP address in the list.

7. Click the Configuration tab.

8. In the Port Forwarding node of the diagram, click View All.

9. Fill in the following:

- Public Port. The port to which public traffic will be addressed on the IP address you acquired in the previous step.

- Private Port. The port on which the instance is listening for forwarded public traffic.

- Protocol. The communication protocol in use between the two ports

10. 点击添加。

15.15. IP负载均衡

用户可以选择关联到同一个公网IP的多个宾客虚拟机。CloudStack实现了TCP级别的负载平衡器，有一下策略。

- 轮询

- 最少连接

- 源IP

这类似于端口转发，但目标可能会有多个IP地址。

15.16. DNS和DHCP

虚拟路由器为客户机提供DNS和DHCP服服务。它将DNS请求代理到在可用性区中配置的DNS服务器。

15.17. VPN 虚拟专用网

CloudStack account拥有者可以建立 VPN 以便访问他们的虚拟机。如果来宾网络 从一个提供 远程VPN访问服务的网络实例化，虚拟路由(在 System VM上的)将被用于提供服务。\nCloudStack 提供了一个基于 L2TP-over-IPsec-based 协议的远程VPN访问服务 给 guest虚拟网络。 因为每个网络有它自己的 虚拟路由， VPNs 没有在这些网络中被共享。Windows, Mac OS X和iOS的原生(自带) VPN客户端 可以用于访问 来宾网络。account拥有者可以建立并管理他们的VPN用户。CloudStack 不使用自己账户数据库，而是用了另外一个独立的表。VPN用户数据库在同一个account拥有者建立的VPN网络中被共享。 也就是说，同一个account拥有者建立里的全部VPN可以被它的全部VPN用户访问。

注意

确保不是所有的网络流量走VPN。也就是说，用于配置VPN的route不是唯一用于该guest network，也不承担全部的网络流量。

- Road Warrior / Remote Access 公路勇士/远程访问.用户希望可以安全地从家里或者办公室连接到云上的一个 私有网络。特别是连接的客户端的IP地址是动态决定的，不能预先配置到VPN 服务器上。

- 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. For more information, see 第 15.17.4 节 “配置站点到站点的VPN连接”

15.17.1. Configuring VPN

To set up VPN for the cloud:

1. 作为管理员或最终用户登入到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. 点击查看IP地址.

5. Click one of the displayed IP address names.

6. Click the Enable VPN button. 🔒

   The IPsec key is displayed in a popup window.

15.17.2. 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, then select Network and Sharing center. Click Setup a connection or network.

3. In the next dialog, select No, create a new connection.

4. In the next dialog, select Use my Internet Connection (VPN).

5. In the next dialog, enter the source NAT IP from step 1 and give the connection a name. Check Don't connect now.

6. In the next dialog, enter the user name and password selected in step 1.

7. Click Create.

8. Go back to the Control Panel and click Network Connections to see the new connection. The connection is not active yet.

9. Right-click the new connection and select Properties. In the Properties dialog, select the Networking tab.

10. In Type of VPN, choose L2TP IPsec VPN, then click IPsec settings. Select Use preshared key. Enter the preshared key from Step 1.

11. The connection is ready for activation. Go back to Control Panel -> Network Connections and double-click the created connection.

12. Enter the user name and password from Step 1.

15.17.3. 在Mac OS X上使用VPN

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

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

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

2. Make sure Send all traffic over VPN connection is not checked.

3. If your preferences are locked, you'll need to click the lock in the bottom left-hand corner to make any changes and provide your administrator credentials.

4. You will need to create a new network entry. Click the plus icon on the bottom left-hand side and you'll see a dialog that says "Select the interface and enter a name for the new service." Select VPN from the Interface drop-down menu, and "L2TP over IPSec" for the VPN Type. Enter whatever you like within the "Service Name" field.

5. You'll now have a new network interface with the name of whatever you put in the "Service Name" field. For the purposes of this example, we'll assume you've named it "CloudStack." Click on that interface and provide the IP address of the interface for your VPN under the Server Address field, and the user name for your VPN under Account Name.

6. Click Authentication Settings, and add the user's password under User Authentication and enter the pre-shared IPSec key in the Shared Secret field under Machine Authentication. Click OK.

7. You may also want to click the "Show VPN status in menu bar" but that's entirely optional.

8. Now click "Connect" and you will be connected to the CloudStack VPN.

15.17.4. 配置站点到站点的VPN连接

一个站点到站点的VPN连接可以帮助你建立从云基础架构到企业内部数据中心的安全连接.这就允许一个账户从企业内部数据中心的设备连接到此账户启用VPN连接的虚拟路由器上，从而通过VPN连接到该账户的虚拟机.

目前支持的数据中心的终端设备是:

- Cisco ISR IOS 12.4或更新
- Juniper J-系统 路由器 JunOS 9.5 或更新版本

注意

除了上述指定的Cisco和Juniper设备，所期望的是任何Cisco或Juniper的设备在支持的操作系统上都可以建立VPN连接.

为了建立站点到站点的VPN连接，需要执行以下步骤:

1. 创建一个虚拟私有云(VPC).

   参见第 15.19 节 “Configuring a Virtual Private Cloud” .

2. 创建一个VPN客户网关.
3. 为你创建的VPC设定一个VPN网关.
4. 从VPC的VPN网关到客户的VPN网关建立VPN连接.

注意

Appropriate events are generated on the CloudStack UI when status of a Site-to-Site VPN connection changes from connected to disconnected, or vice versa. Currently no events are generated when establishing a VPN connection fails or pending.

15.17.4.1. Creating and Updating a VPN Customer Gateway

注意

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

To add a VPN Customer Gateway:

1. 作为管理员或最终用户登入到CloudStack UI.
2. 在左边的导航栏里选择网络.
3. In the Select view, select VPN Customer Gateway.
4. Click Add site-to-site VPN.

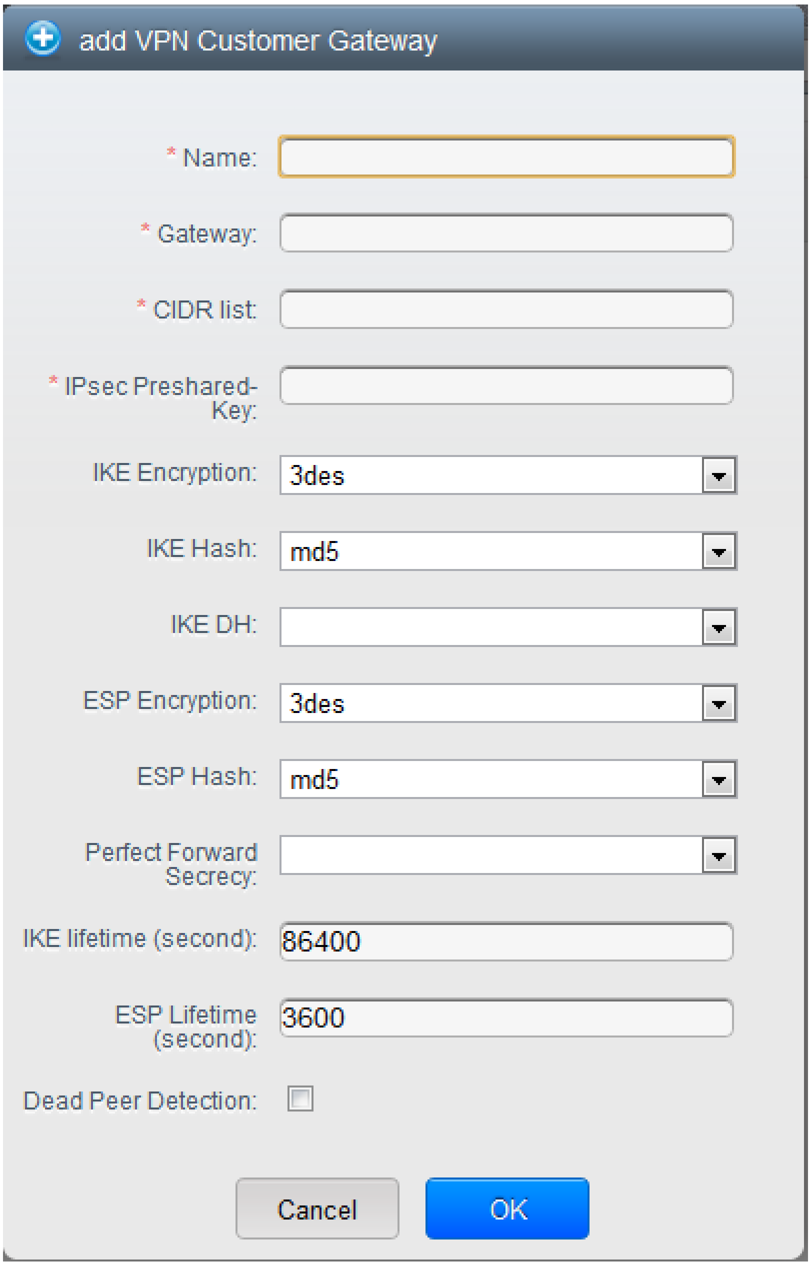

add VPN Customer Gateway

* Name:	
* Gateway:	
* CIDR list:	
* IPsec Preshared-Key:	
IKE Encryption:	3des
IKE Hash:	md5
IKE DH:	
ESP Encryption:	3des
ESP Hash:	md5
Perfect Forward Secrecy:	
IKE lifetime (second):	86400
ESP Lifetime (second):	3600
Dead Peer Detection:	□

填写以下内容。

- Name: A unique name for the VPN customer gateway you create.

- Gateway: The IP address for the remote gateway.

- CIDR list: The guest CIDR list of the remote subnets. Enter a CIDR or a comma-separated list of CIDRs. Ensure that a guest CIDR list is not overlapped with the VPC’ s CIDR, or another guest CIDR. The CIDR must be RFC1918-compliant.

- IPsec Preshared Key: Preshared keying is a method where the endpoints of the VPN share a secret key. This key value is used to authenticate the customer gateway and the VPC VPN gateway to each other.

注意

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

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

注意

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

- IKE Hash: The IKE hash for phase-1. The supported hash algorithms are SHA1 and MD5.

- IKE DH: A public-key cryptography protocol which allows two parties to establish a shared secret over an insecure communications channel. The 1536-bit Diffie-Hellman group is used within IKE to establish session keys. The supported options are None, Group-5 (1536-bit) and Group-2 (1024-bit).

- ESP Encryption: Encapsulating Security Payload (ESP) algorithm within phase-2. The supported encryption algorithms are AES128, AES192, AES256, and 3DES.

注意

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

- ESP Hash: Encapsulating Security Payload (ESP) hash for phase-2. Supported hash algorithms are SHA1 and MD5.

- Perfect Forward Secrecy: Perfect Forward Secrecy (or PFS) is the property that ensures that a session key derived from a set of long-term public and private keys will not be compromised. This property enforces a new Diffie-Hellman key exchange. It provides the keying material that has greater key material life and thereby greater resistance to cryptographic attacks. The available options are None, Group-5 (1536-bit) and Group-2 (1024-bit). The security of the key exchanges increase as the DH groups grow larger, as does the time of the exchanges.

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

- IKE Lifetime (seconds): The phase-1 lifetime of the security association in seconds. Default is 86400 seconds (1 day). Whenever the time expires, a new phase-1 exchange is performed.

- ESP Lifetime (seconds): The phase-2 lifetime of the security association in seconds. Default is 3600 seconds (1 hour). Whenever the value is exceeded, a re-key is initiated to provide a new IPsec encryption and authentication session keys.

- Dead Peer Detection: A method to detect an unavailable Internet Key Exchange (IKE) peer. Select this option if you want the virtual router to query the liveliness of its IKE peer at regular intervals. It' s recommended to have the same configuration of DPD on both side of VPN connection.

5. 点击 确定。

Updating and Removing a VPN Customer Gateway

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

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. In the Select view, select VPN Customer Gateway.

4. Select the VPN customer gateway you want to work with.

5. To modify the required parameters, click the Edit VPN Customer Gateway button ✎

6. To remove the VPN customer gateway, click the Delete VPN Customer Gateway button ✖

7. 点击 确定。

15.17.4.2. Creating a VPN gateway for the VPC

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. vpn连接列表

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

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

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

5. Click the Settings icon.

   The following options are displayed.

   - IP å#°å##
   - 网关
   - 点对点VPN
   - Network ACLs

6. 选择点对点VPN

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

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

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

   The following details are displayed in the VPN Gateway page:

   - IP å#°å##
   - å¸#æ#·
   - å##

15.17.4.3. 新建vpn连接

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. vpn连接列表

   All the VPCs that you create for the account are listed in the page.

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

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

5. Click the Settings icon.

   The following options are displayed.

   - IP地址

- 网关
- 点对点VPN
- Network ASLs

6. 选择点对点VPN

   The Site-to-Site VPN page is displayed.

7. From the Select View drop-down, ensure that VPN Connection is selected.

8. 选择创建vpn连接按钮

   The Create VPN Connection dialog is displayed:

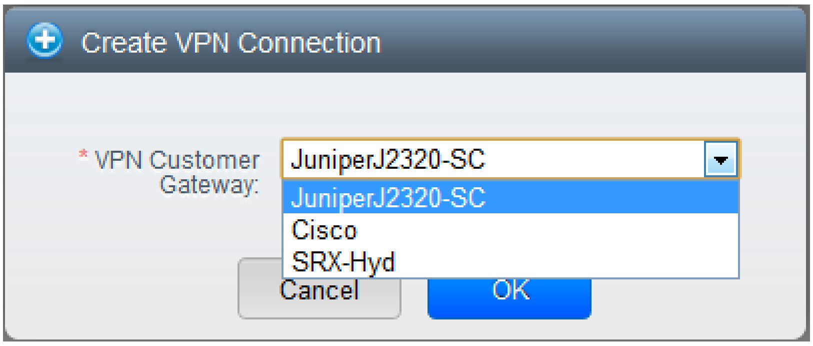

Create VPN Connection

* VPN Customer Gateway: JuniperJ2320-SC

9. Select the desired customer gateway, then click OK to confirm.

   Within a few moments, the VPN Connection is displayed.

   VPN连接信息

   - IP å#°å##
   - ç½#å
     3
   - ç#¶æ##
   - IPSec共享密钥
   - IKE密钥
   - ESP密钥

15.17.4.4. Restarting and Removing a VPN Connection

1. 登陆到CloudStack界面以管理员或者终端用户。

2. 在左边的导航，选择网络

3. vpn连接列表

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

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

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

5. Click the Settings icon.

   The following options are displayed.

   - IP å#°å##
   - 网关
   - 点对点VPN
   - Network ASLs

6. Select Site-to-Site VPN.

   The Site-to-Site VPN page is displayed.

7. From the Select View drop-down, ensure that VPN Connection is selected.

   All the VPN connections you created are displayed.

8. Select the VPN connection you want to work with.

   The Details tab is displayed.

9. To remove a VPN connection, click the Delete VPN connection button ✖

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

15.18. About Inter-VLAN Routing

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

This feature is supported on XenServer and VMware hypervisors.

The major advantages are:

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

注意

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

- The administrator can allow users create their own VPC and deploy the application. In this scenario, the VMs that belong to the account are deployed on the VLANs allotted to that account.

- Both administrators and users can create multiple VPCs. The guest network NIC is plugged to the VPC virtual router when the first VM is deployed in a tier.

- The administrator can create the following gateways to send to or receive traffic from the VMs:

  - VPN Gateway: For more information, see 第 15.17.4.2 节 "Creating a VPN gateway for the VPC" .

  - Public Gateway: The public gateway for a VPC is added to the virtual router when the virtual router is created for VPC. The public gateway is not exposed to the end users. You are not allowed to list it, nor allowed to create any static routes.

  - Private Gateway: For more information, see 第 15.19.5 节 "Adding a Private Gateway to a VPC" .

- Both administrators and users can create various possible destinations-gateway combinations. However, only one gateway of each type can be used in a deployment.

  For example:

  - VLANs and Public Gateway: For example, an application is deployed in the cloud, and the Web application VMs communicate with the Internet.

  - VLANs, VPN Gateway, and Public Gateway: For example, an application is deployed in the cloud; the Web application VMs communicate with the Internet; and the database VMs communicate with the on-premise devices.

- The administrator can define Access Control List (ACL) on the virtual router to filter the traffic among the VLANs or between the Internet and a VLAN. You can define ACL based on CIDR, port range, protocol, type code (if ICMP protocol is selected) and Ingress/ Egress type.

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

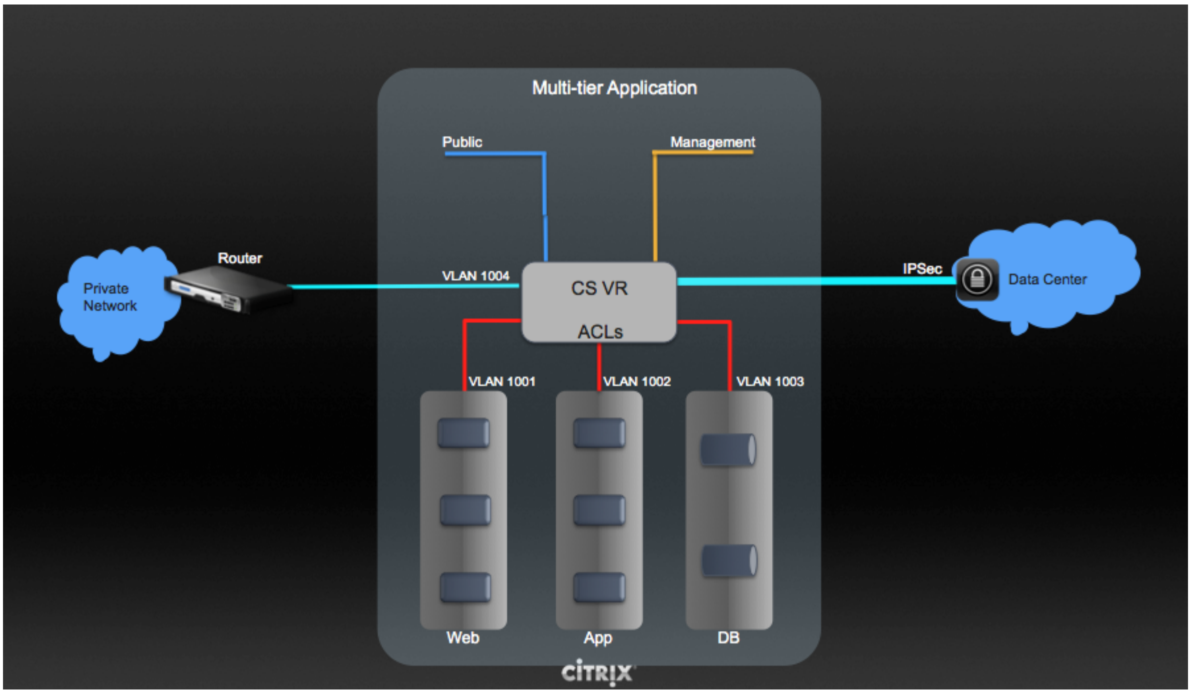

To set up a multi-tier Inter-VLAN deployment, see 第 15.19 节 "Configuring a Virtual Private Cloud" .

15.19. Configuring a Virtual Private Cloud

15.19.1. About Virtual Private Clouds

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

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

Major Components of a VPC:

A VPC is comprised of the following network components:

- VPC: A VPC acts as a container for multiple isolated networks that can communicate with each other via its virtual router.

- Network Tiers: Each tier acts as an isolated network with its own VLANs and CIDR list, where you can place groups of resources, such as VMs. The tiers are segmented by means of VLANs. The NIC of each tier acts as its gateway.

- Virtual Router: A virtual router is automatically created and started when you create a VPC. The virtual router connect the tiers and direct traffic among the public gateway, the VPN gateways, and the NAT instances. For each tier, a corresponding NIC and IP exist in the virtual router. The virtual router provides DNS and DHCP services through its IP.

- Public Gateway: The traffic to and from the Internet routed to the VPC through the public gateway. In a VPC, the public gateway is not exposed to the end user; therefore, static routes are not support for the public gateway.

- Private Gateway: All the traffic to and from a private network routed to the VPC through the private gateway. For more information, see 第 15.19.5 节 "Adding a Private Gateway to a VPC" .

- VPN Gateway: The VPC side of a VPN connection.

- Site-to-Site VPN Connection: A hardware-based VPN connection between your VPC and your datacenter, home network, or co-location facility. For more information, see 第 15.17.4 节 "配置站点到站点的VPN连接" .

- Customer Gateway: The customer side of a VPN Connection. For more information, see 第 15.17.4.1 节 "Creating and Updating a VPN Customer Gateway" .

- NAT Instance: An instance that provides Port Address Translation for instances to access the Internet via the public gateway. For more information, see 第 15.19.9 节 "Enabling or Disabling Static NAT on a VPC" .

Network Architecture in a VPC

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

- VPC with a public gateway only

- VPC with public and private gateways

- VPC with public and private gateways and site-to-site VPN access

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

Connectivity Options for a VPC

You can connect your VPC to:

- The Internet through the public gateway.

- The corporate datacenter by using a site-to-site VPN connection through the VPN gateway.

- Both the Internet and your corporate datacenter by using both the public gateway and a VPN gateway.

VPC Network Considerations

Consider the following before you create a VPC:

- A VPC, by default, is created in the enabled state.

- A VPC can be created in Advance zone only, and can't belong to more than one zone at a time.

- The default number of VPCs an account can create is 20. However, you can change it by using the max.account.vpcs global parameter, which controls the maximum number of VPCs an account is allowed to create.

- The default number of tiers an account can create within a VPC is 3. You can configure this number by using the vpc.max.networks parameter.

- Each tier should have an unique CIDR in the VPC. Ensure that the tier's CIDR should be within the VPC CIDR range.

- A tier belongs to only one VPC.

- All network tiers inside the VPC should belong to the same account.

- When a VPC is created, by default, a SourceNAT IP is allocated to it. The Source NAT IP is released only when the VPC is removed.

- A public IP can be used for only one purpose at a time. If the IP is a sourceNAT, it cannot be used for StaticNAT or port forwarding.

- The instances only have a private IP address that you provision. To communicate with the Internet, enable NAT to an instance that you launch in your VPC.

- Only new networks can be added to a VPC. The maximum number of networks per VPC is limited by the value you specify in the vpc.max.networks parameter. The default value is three.

- The load balancing service can be supported by only one tier inside the VPC.

- If an IP address is assigned to a tier:

  - That IP can't be used by more than one tier at a time in the VPC. For example, if you have tiers A and B, and a public IP1, you can create a port forwarding rule by using the IP either for A or B, but not for both.

  - That IP can't be used for StaticNAT, load balancing, or port forwarding rules for another guest network inside the VPC.

- Remote access VPN is not supported in VPC networks.

15.19.2. Adding a Virtual Private Cloud

When creating the VPC, you simply provide the zone and a set of IP addresses for the VPC network address space. You specify this set of addresses in the form of a Classless Inter-Domain Routing (CIDR) block.

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. vpn连接列表

4. Click Add VPC. The Add VPC page is displayed as follows:

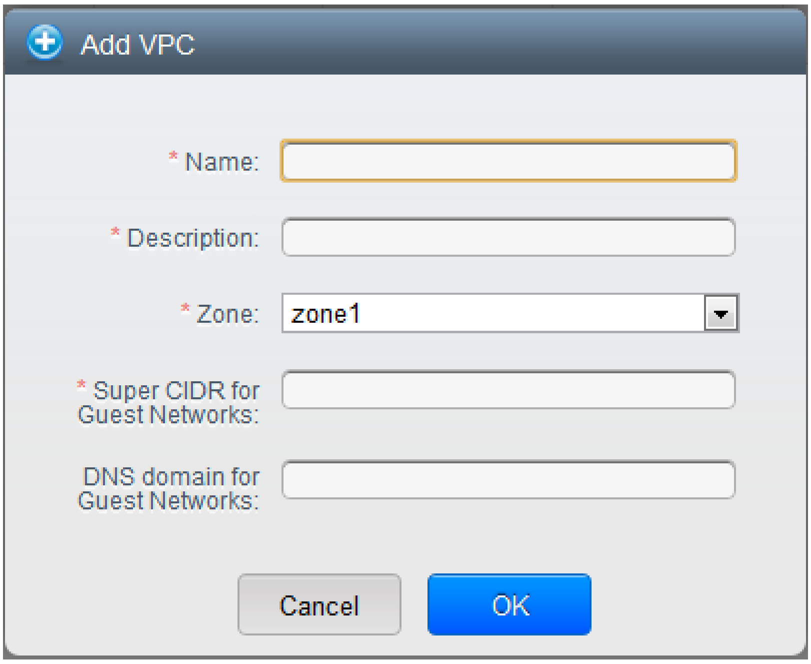
Add VPC
* Name:
* Description:
* Zone:
zone1
* Super CIDR for Guest Networks:
DNS domain for Guest Networks:
Cancel
OK

填写以下内容。

- Name: A short name for the VPC that you are creating.
- Description: A brief description of the VPC.
- Zone: Choose the zone where you want the VPC to be available.
- Super CIDR for Guest Networks: Defines the CIDR range for all the tiers (guest networks) within a VPC. When you create a tier, ensure that its CIDR is within the Super CIDR value you enter. The CIDR must be RFC1918 compliant.
- DNS domain for Guest Networks: If you want to assign a special domain name, specify the DNS suffix. This parameter is applied to all the tiers within the VPC. That implies, all the tiers you create in the VPC belong to the same DNS domain. If the parameter is not specified, a DNS domain name is generated automatically.

15.19.3. Adding Tiers

Tiers are distinct locations within a VPC that act as isolated networks, which do not have access to other tiers by default. Tiers are set up on different VLANs that can communicate with each other by using a virtual router. Tiers provide inexpensive, low latency network connectivity to other tiers within the VPC.

1. 登陆到CloudStack界面以管理员或者终端用户。
2. 在左边的导航，选择网络
3. vpn连接列表

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

注意

The end users can see their own VPCs, while root and domain admin can see any VPC they are authorized to see.

4. Click the Configure button of the VPC for which you want to set up tiers.

   The Add new tier dialog is displayed, as follows:

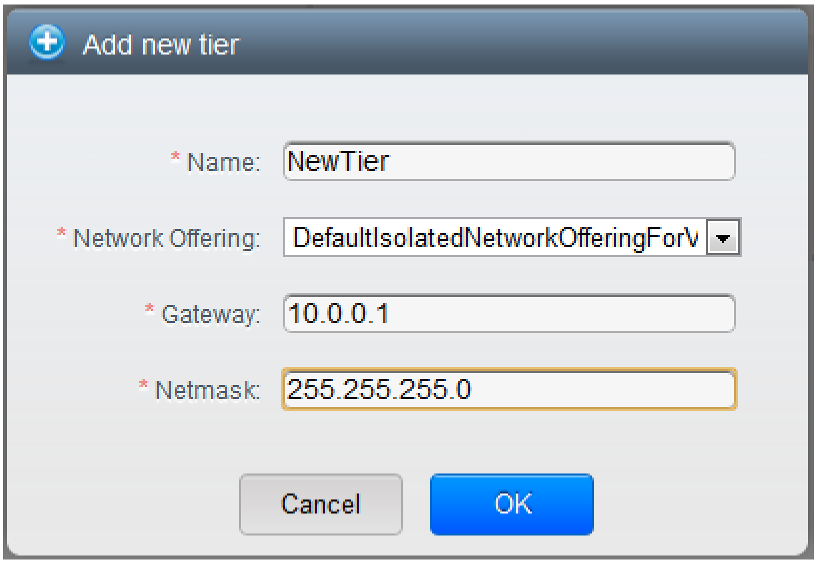

* Name: NewTier

* Network Offering: DefaultIsolatedNetworkOfferingForV

* Gateway: 10.0.0.1

* Netmask: 255.255.255.0

   If you have already created tiers, the VPC diagram is displayed. Click Create Tier to add a new tier.

5. Specify the following:

   All the fields are mandatory.

   - Name: A unique name for the tier you create.

   - Network Offering: The following default network offerings are listed: DefaultIsolatedNetworkOfferingForVpcNetworksNoLB, DefaultIsolatedNetworkOfferingForVpcNetworks

     In a VPC, only one tier can be created by using LB-enabled network offering.

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

   - Netmask: The netmask for the tier you create.

     For example, if the VPC CIDR is 10.0.0.0/16 and the network tier CIDR is 10.0.1.0/24, the gateway of the tier is 10.0.1.1, and the netmask of the tier is 255.255.255.0.

6. 点击 确定。

7. Continue with configuring access control list for the tier.

15.19.4. Configuring Access Control List

Define Network Access Control List (ACL) on the VPC virtual router to control incoming (ingress) and outgoing (egress) traffic between the VPC tiers, and the tiers and Internet. By default, all incoming and outgoing traffic to the guest networks is blocked. To open the ports, you must create a new network ACL. The network ACLs can be created for the tiers only if the NetworkACL service is supported.

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. vpn连接列表

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

4. Click the Settings icon.

   The following options are displayed.

   - IP地址
   - 网关
   - 点对点VPN
   - Network ACLs

5. Select Network ACLs.

   The Network ACLs page is displayed.

6. Click Add Network ACLs.

   To add an ACL rule, fill in the following fields to specify what kind of network traffic is allowed in this tier.

   - CIDR: The CIDR acts as the Source CIDR for the Ingress rules, and Destination CIDR for the Egress rules. To accept traffic only from or to the IP addresses within a particular address block, enter a CIDR or a comma-separated list of CIDRs. The CIDR is the base IP address of the incoming traffic. For example, 192.168.0.0/22. To allow all CIDRs, set to 0.0.0.0/0.

   - Protocol: The networking protocol that sources use to send traffic to the tier. The TCP and UDP protocols are typically used for data exchange and end-user communications. The ICMP protocol is typically used to send error messages or network monitoring data.

   - 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.

   - Select Tier: Select the tier for which you want to add this ACL rule.

   - ICMP Type, ICMP Code (ICMP only): The type of message and error code that will be sent.

- Traffic Type: Select the traffic type you want to apply.
  - Egress: To add an egress rule, select Egress from the Traffic type drop-down box and click Add. This specifies what type of traffic is allowed to be sent out of VM instances in this tier. If no egress rules are specified, all traffic from the tier is allowed out at the VPC virtual router. Once egress rules are specified, only the traffic specified in egress rules and the responses to any traffic that has been allowed in through an ingress rule are allowed out. No egress rule is required for the VMs in a tier to communicate with each other.
  - Ingress: To add an ingress rule, select Ingress from the Traffic type drop-down box and click Add. This specifies what network traffic is allowed into the VM instances in this tier. 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.

> **注意**
>
> By default, all incoming and outgoing traffic to the guest networks is blocked. To open the ports, create a new network ACL.

7. Click Add. The ACL rule is added.

   To view the list of ACL rules you have added, click the desired tier from the Network ACLs page, then select the Network ACL tab.

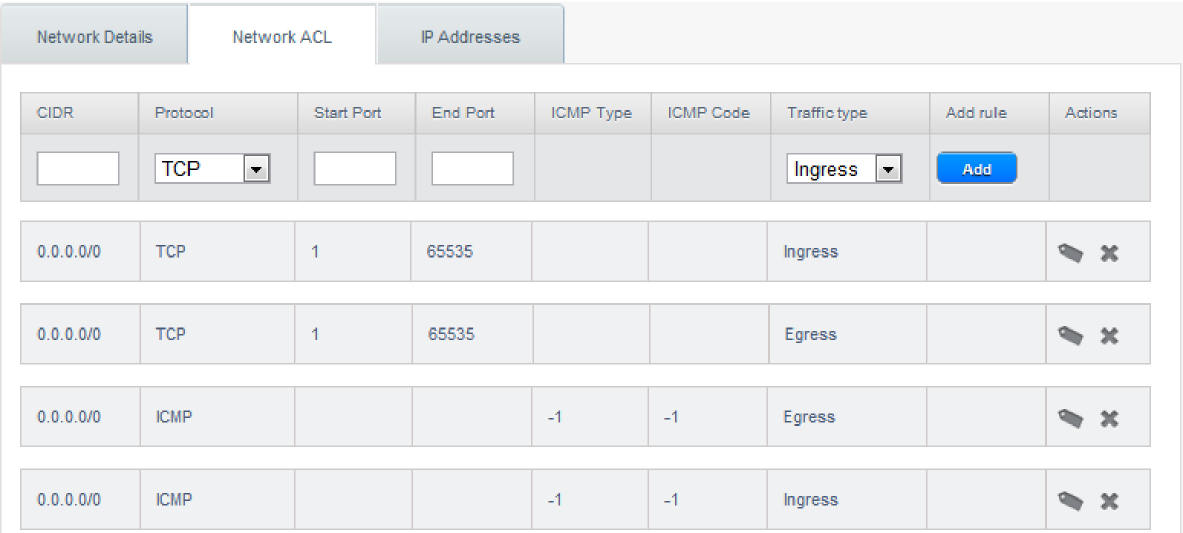

CIDR	Protocol	Start Port	End Port	ICMP Type	ICMP Code	Traffic type	Add rule	Actions
	TCP					Ingress	Add	
0.0.0.0/0	TCP	1	65535			Ingress		✖
0.0.0.0/0	TCP	1	65535			Egress		✖
0.0.0.0/0	ICMP			-1	-1	Egress		✖
0.0.0.0/0	ICMP			-1	-1	Ingress		✖

   You can edit the tags assigned to the ACL rules and delete the ACL rules you have created. Click the appropriate button in the Actions column.

15.19.5. Adding a Private Gateway to a VPC

A private gateway can be added by the root admin only. The VPC private network has 1:1 relationship with the NIC of the physical network. No gateways with duplicated VLAN and IP are allowed in the same data center.

1. 登陆到CloudStack界面以管理员或者终端用户。

2. 在左边的导航，选择网络

3. vpn连接列表

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

4. Click the Configure button of the VPC to which you want to configure load balancing rules.

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

5. Click the Settings icon.

   The following options are displayed.

   - IP å#°å##

   - Private Gateways

   - 点对点VPN

   - Network ACLs

6. Select Private Gateways.

   The Gateways page is displayed.

7. Click Add new gateway:

   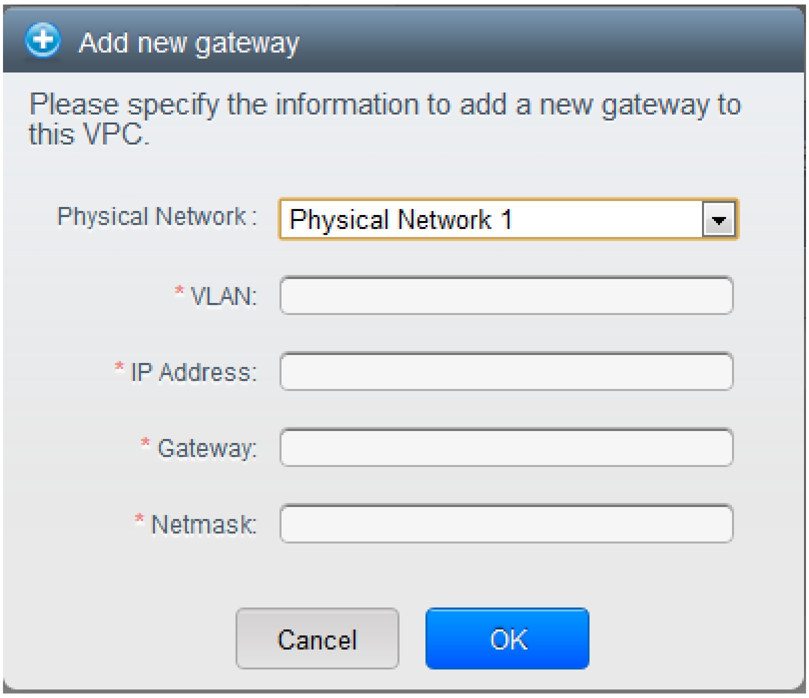

   Add new gateway

   Please specify the information to add a new gateway to this VPC.

   Physical Network : Physical Network 1

   * VLAN:

   * IP Address:

   * Gateway:

   * Netmask:

8. Specify the following:

   - Physical Network: The physical network you have created in the zone.

   - IP Address: The IP address associated with the VPC gateway.

- Gateway: The gateway through which the traffic is routed to and from the VPC.

- Netmask: The netmask associated with the VPC gateway.

- VLAN: The VLAN associated with the VPC gateway.

The new gateway appears in the list. You can repeat these steps to add more gateway for this VPC.

15.19.6. Deploying VMs to the Tier

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. vpn连接列表

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

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

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

5. Click the Add VM button of the tier for which you want to add a VM.

   The Add Instance page is displayed.

   Follow the on-screen instruction to add an instance. For information on adding an instance, see Adding Instances section in the Installation Guide.

15.19.7. Acquiring a New IP Address for a VPC

When you acquire an IP address, all IP addresses are allocated to VPC, not to the guest networks within the VPC. The IPs are associated to the guest network only when the first port-forwarding, load balancing, or Static NAT rule is created for the IP or the network. IP can't be associated to more than one network at a time.

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. vpn连接列表

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

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

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

5. Click the Settings icon.

   The following options are displayed.

   - IP地址

   - 网关

   - 点对点VPN

- Network ACLs

6. Select IP Addresses.

   The IP Addresses page is displayed.

7. 点击获得一个新IP，并且在确认的对话框中点击确定.

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

15.19.8. Releasing an IP Address Alloted to a VPC

The IP address is a limited resource. If you no longer need a particular IP, you can disassociate it from its VPC and return it to the pool of available addresses. An IP address can be released from its tier, only when all the networking (port forwarding, load balancing, or StaticNAT) rules are removed for this IP address. The released IP address will still belongs to the same VPC.

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. vpn连接列表

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

4. Click the Configure button of the VPC whose IP you want to release.

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

5. Click the Settings icon.

   The following options are displayed.

   - IP地址
   - 网关
   - 点对点VPN
   - Network ACLs

6. Select IP Addresses.

   The IP Addresses page is displayed.

7. Click the IP you want to release.

8. In the Details tab, click the Release IP button ✖

15.19.9. Enabling or Disabling Static NAT on a VPC

A static NAT rule maps a public IP address to the private IP address of a VM in a VPC to allow Internet traffic to it. This section tells how to enable or disable static NAT for a particular IP address in a VPC.

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

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

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. vpn连接列表

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

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

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

5. Click the Settings icon.

   The following options are displayed.

   - IP地址
   - 网关
   - 点对点VPN
   - Network ACLs

6. Select IP Addresses.

   The IP Addresses page is displayed.

7. Click the IP you want to work with.

8. In the Details tab,click the Static NAT button. The button toggles between Enable and Disable, depending on whether static NAT is currently enabled for the IP address.

9. If you are enabling static NAT, a dialog appears as follows:

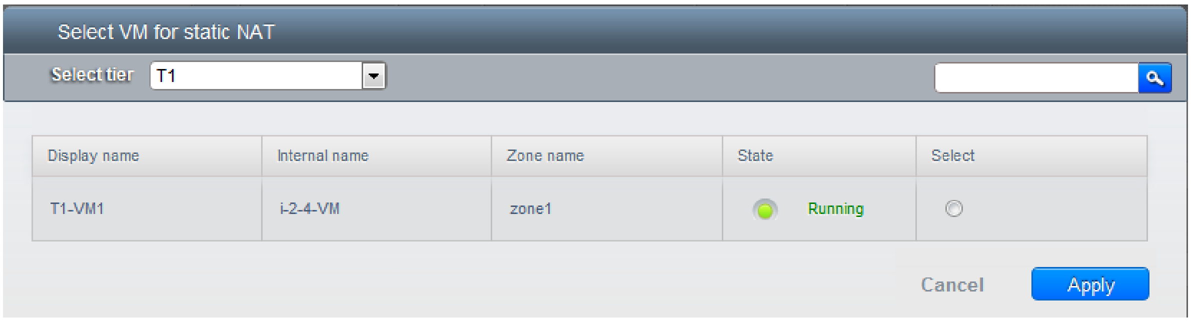

Select VM for static NAT

Select tier T1

Display name	Internal name	Zone name	State	Select
T1-VM1	i-2-4-VM	zone1	● Running	○

10. Select the tier and the destination VM, then click Apply.

15.19.10. Adding Load Balancing Rules on a VPC

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

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. vpn连接列表

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

4. Click the Configure button of the VPC to which you want to configure load balancing rules.

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

5. Click the Settings icon.

   The following options are displayed.

   - IP地址
   - 网关
   - 点对点VPN
   - Network ACLs

6. Select IP Addresses.

   The IP Addresses page is displayed.

7. Click the IP address for which you want to create the rule, then click the Configuration tab.

8. In the Load Balancing node of the diagram, click View All.

9. Select the tier to which you want to apply the rule.

   > **注意**
   >
   > In a VPC, the load balancing service is supported only on a single tier.

10. Specify the following:

    - Name: A name for the load balancer rule.
    - Public Port: The port that receives the incoming traffic to be balanced.
    - Private Port: The port that the VMs will use to receive the traffic.

- Algorithm. Choose the load balancing algorithm you want CloudStack to use. CloudStack supports the following well-known algorithms:
    - 轮询
    - æ##å°#è¿#æ#¥ç®#æ³#
    - æº#ç®#æ³#
- Stickiness. (Optional) Click Configure and choose the algorithm for the stickiness policy. See Sticky Session Policies for Load Balancer Rules.
- Add VMs: Click Add VMs, then select two or more VMs that will divide the load of incoming traffic, and click Apply.

The new load balancing rule appears in the list. You can repeat these steps to add more load balancing rules for this IP address.

15.19.11. Adding a Port Forwarding Rule on a VPC

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. vpn连接列表

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

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

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

5. Click the Settings icon.

    The following options are displayed.

    - IP地址
    - 网关
    - 点对点VPN
    - Network ACLs

6. Choose an existing IP address or acquire a new IP address. Click the name of the IP address in the list.

    The IP Addresses page is displayed.

7. Click the IP address for which you want to create the rule, then click the Configuration tab.

8. In the Port Forwarding node of the diagram, click View All.

9. Select the tier to which you want to apply the rule.

10. Specify the following:

- Public Port: The port to which public traffic will be addressed on the IP address you acquired in the previous step.
- Private Port: The port on which the instance is listening for forwarded public traffic.
- Protocol: The communication protocol in use between the two ports.
  - TCP
  - UDP
- Add VM: Click Add VM. Select the name of the instance to which this rule applies, and click Apply.

  You can test the rule by opening an ssh session to the instance.

15.19.12. Removing Tiers

You can remove a tier from a VPC. A removed tier cannot be revoked. When a tier is removed, only the resources of the tier are expunged. All the network rules (port forwarding, load balancing and staticNAT) and the IP addresses associated to the tier are removed. The IP address still be belonging to the same VPC.

1. 作为管理员或最终用户登入到CloudStack UI.
2. 在左边的导航栏里选择网络.
3. vpn连接列表

   All the VPC that you have created for the account is listed in the page.
4. Click the Configure button of the VPC for which you want to set up tiers.

   The Configure VPC page is displayed. Locate the tier you want to work with.
5. Click the Remove VPC button:

   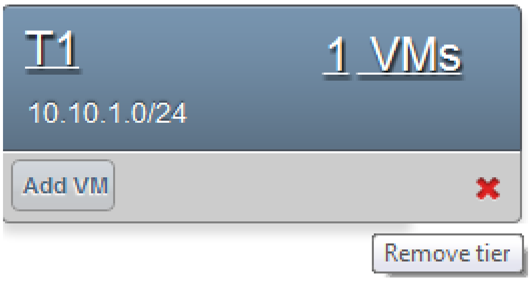

   Wait for some time for the tier to be removed.

15.19.13. Editing, Restarting, and Removing a Virtual Private Cloud

> **注意**
>
> Ensure that all the tiers are removed before you remove a VPC.

1. 作为管理员或最终用户登入到CloudStack UI.

2. 在左边的导航栏里选择网络.

3. vpn连接列表

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

4. Select the VPC you want to work with.

5. To remove, click the Remove VPC button ✖

   You can edit the name and description of a VPC. To do that, select the VPC, then click the Edit button. ✎

   To restart a VPC, select the VPC, then click the Restart button. ↻ i

15.20. Persistent Networks

The network that you can provision without having to deploy any VMs on it is called a persistent network. A persistent network can be part of a VPC or a non-VPC environment.

When you create other types of network, a network is only a database entry until the first VM is created on that network. When the first VM is created, a VLAN ID is assigned and the network is provisioned. Also, when the last VM is destroyed, the VLAN ID is released and the network is no longer available. With the addition of persistent network, you will have the ability to create a network in CloudStack in which physical devices can be deployed without having to run any VMs. Additionally, you can deploy physical devices on that network.

One of the advantages of having a persistent network is that you can create a VPC with a tier consisting of only physical devices. For example, you might create a VPC for a three-tier application, deploy VMs for Web and Application tier, and use physical machines for the Database tier. Another use case is that if you are providing services by using physical hardware, you can define the network as persistent and therefore even if all its VMs are destroyed the services will not be discontinued.

15.20.1. Persistent Network Considerations

- Persistent network is designed for isolated networks.

- All default network offerings are non-persistent.

- A network offering cannot be editable because changing it affects the behavior of the existing networks that were created using this network offering.
- When you create a guest network, the network offering that you select defines the network persistence. This in turn depends on whether persistent network is enabled in the selected network offering.
- An existing network can be made persistent by changing its network offering to an offering that has the Persistent option enabled. While setting this property, even if the network has no running VMs, the network is provisioned.
- An existing network can be made non-persistent by changing its network offering to an offering that has the Persistent option disabled. If the network has no running VMs, during the next network garbage collection run the network is shut down.
- When the last VM on a network is destroyed, the network garbage collector checks if the network offering associated with the network is persistent, and shuts down the network only if it is non-persistent.

15.20.2. Creating a Persistent Guest Network

To create a persistent network, perform the following:

1. Create a network offering with the Persistent option enabled.

   See 第 9.4.1 节 “Creating a New Network Offering” .

2. Select Network from the left navigation pane.

3. Select the guest network that you want to offer this network service to.

4. Click the Edit button.

5. From the Network Offering drop-down, select the persistent network offering you have just created.

6. 点击 确定。

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.

注意

You can configure the system.vm.random.password parameter to create a random system VM password to ensure higher security. If you reset the value for system.vm.random.password to true and restart the Management Server, a random password is generated and stored encrypted in the database. You can view the decrypted password under the system.vm.password global parameter on the CloudStack UI or by calling the listConfigurations API.

16.1. 系统VM模板

系统VMs来自单个模板。该系统VM都具有以下特点:

- Debian 6.0(“Squeeze”),从Debian安全APT存储库安装2.6.32以上内核并安装最新的安全补丁。
- 建议最小化安装, 以降低安全攻击漏洞。
- 基于 Xen/VMWare 32位增强性能版
- 所有虚拟管理程序 Xen PV 驱动, KVM virtio 驱动和VMware tools 在pvops 内核中得到最佳性能
- Xen tools 包含性能监控
- 从Debian 资源库安装最新版的HAProxy、iptables、IPsec和Apache可以有效改善安全与速度
- 从 Sun/Oracle 安装最新版本的JRE也可以有效改善安全性与速度

16.2. 支持VMware的多种系统虚拟机

每个CloudStack 资源域都有单独的系统虚拟机用于处理关于模版的各种任务, 如下载模版, 上传模版, 上传ISO。在使用VMware的资源域中, 会启动额外的系统虚拟机用于执行VMware的特殊任务, 如做快照和创建私有模版。当VMware特殊任务压力增加时, CloudStack 管理服务器会启动额外的系统虚拟机。管理服务器会监控并权衡所有发往这些系统虚拟机的命令, 并执行动态负载均衡和扩展更多的系统虚拟机。

16.3. 控制台代理 (控制台虚拟机)

控制台代理作为系统虚拟机的一种通过网页用户接口为用户呈现一个控制台视图。它通过虚拟机软件为来宾提供的控制台将用户浏览器与vnc端口相连。管理员和终端用户动能通过网页用户接口获得一个控制台连接。

Clicking 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.

注意

虚拟机软件会分配多个vnc端口供vnc会话使用。

不会有任何流量是刘翔来宾虚拟IP的，并且不需要在来宾虚拟机打开vnc。

控制台虚拟机会定时的向管理服务器回报当前活动的会话数。默认报告间隔是五秒钟。可以通过管理服务器的配置参数 consoleproxy.loadscan.interval.更改。

如果来宾虚拟机之前有已经分配的关联控制台代理的会话，控制台代理的分配会由第一次分配的控制台代理决定。如果该来宾虚拟级之前存在已分配的控制台代理，则不论该控制台代理目前负载如何管理服务器都会将该来宾虚拟机分配到目标控制台代理虚拟机。如果失败则会将来宾虚拟机分配到第一个拥有足够资源处理新会话的控制台代理上。

管理员能重启控制台代理，但是这回是已经存在的用户的控制台会话中断。

16.3.1. Using a SSL Certificate for the Console Proxy

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. You 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 the 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.

16.3.2. 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:

   - The certificate you've just generated.

   - The private key you've just generated.

   - The desired new domain name; for example, company.com

4. The desired new domain name; for example, company.com

   This stops all currently running console proxy VMs, then restarts them with the new certificate and key. Users might notice a brief interruption in console availability.

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

16.4. è##æ##è·¯ç#±å#¨

虚拟路由器是一个系统虚拟机。它经常在 CloudStack 服务方案中被使用；终端用户不能够直接访问虚拟路由器。用户可ping和影响它(比如设置端口转发)但是不能通过ssh访问

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..

16.4.1. 配置虚拟路由器

你可以设置一下内容

- IP地址范围

- 支持的网络服务
- 由虚拟路由网络服务提供的默认域名
- 网关IP地址
- 本产品多久从虚拟路由器获取一次网络使用数据。如果你打算手机虚拟路由器的流量计量数据，设置全局变量router.stats.interval。如果你不使用虚拟路怄气收集网络使用数据，设置该值为0

16.4.2. 使用系统计算服务方案升级虚拟路由器

当 CloudStack 创建一个虚拟路由器，它是按照默认的系统计算服务方案进行的设置。参见 第 8.2 节 “系统服务方案”。所有在单独客户网络中的虚拟路由器都使用相同的系统计算服务方案。可以通过新建和使用自定义的系统计算服务方案来提高虚拟路由器的性能。

1. 定义定制的系统计算服务方案，参见第 8.2.1 节 “Creating a New System Service Offering”。在系统虚拟机类型中，选择域路由器。

2. Associate the system service offering with a network offering. See 第 9.4.1 节 “Creating a New Network Offering”.

3. 将网络服务方案应用到使用新系统计算服务方案的虚拟路由器的网络上。如果这是一个新的网络，请根据66页的添加额外客户网络的步骤操作。想要改变已生成的虚拟路由器的计算服务方案，请根据以下步骤操作第 15.6.2 节 “Changing the Network Offering on a Guest Network”。

16.4.3. 虚拟路由器的最佳实践

- 警告：从一个虚拟机管理程序控制台重新启动一台虚拟路由器，将删除所有iptables规则。要解决这个问题，从CloudStack UI停止虚拟路由器和启动。
- 警告：在网络中只有一个路由器可用时，不要使用destroyRouter API，因为restartNetwork API带cleanup=false参数不能随后重新创建它。如果你想销毁并重新创建网络中的单一路由器，使用restartNetwork API 带cleanup=true参数。

16.5. 二级存储虚拟机

除了主机CloudStack的二级存储虚拟机会挂载和往二级存储中写入内容。

通过二级存储虚拟机来提交信息到二级存储。二级存储虚拟机会使用多种协议通过URL来获取模版和ISO镜像文件。

二级存储虚拟机会提供后台任务来负责各种二级存储的活动：将新模版的下载到资源域中，多个资源域之间的模版复制，和快照备份。

管理员如果有需要，可以登录到二级存储虚拟机上。

System Reliability and High Availability

17.1. 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 Serves. All guest VMs will continue to work.

When the Management Server is down, no new VMs can be created, and the end user and admin UI, API, dynamic load distribution, and HA will cease to work.

17.2. Management Server Load Balancing

CloudStack can use a load balancer to provide a virtual IP for multiple Management Servers. The administrator is responsible for creating the load balancer rules for the Management Servers. The application requires persistence or stickiness across multiple sessions. The following chart lists the ports that should be load balanced and whether or not persistence is required.

Even if persistence is not required, enabling it is permitted.

Source Port	Destination Port	å##è®®	Persistence Required?
80 or 443	8080 (or 20400 with AJP)	HTTP (or AJP)	是
8250	8250	TCP	是
8096	8096	HTTP	否

In addition to above settings, the administrator is responsible for setting the 'host' global config value from the management server IP to load balancer virtual IP address. If the 'host' value is not set to the VIP for Port 8250 and one of your management servers crashes, the UI is still available but the system VMs will not be able to contact the management server.

17.3. 开启了高可用特性的虚拟机

用户可以给虚拟机制定开启高可用特性。默认情况下所有的虚拟路由虚拟机负载均衡虚拟机自动开启了高可用特性。当CloudStack检测到开启了高可用特性的虚拟机崩溃时将会在相同的可用资源与中自动重新启动该虚拟机。高可用特性不会跨资源域执行。CloudStack 采用比较保守的方式重启虚拟机，以确使不会同时运行两个相同的实例。管理服务器会尝试在本季群的另一台主机上开启该虚拟机。

高可用特性只在使用iSCSI和NFS做主存储的才可以使用。不支持使用本地存储作为主存储的高可用。

17.4. HA for Hosts

用户可以给虚拟机制定开启高可用特性。默认情况下所有的虚拟路由虚拟机负载均衡虚拟机自动开启了高可用特性。当CloudStack检测到开启了高可用特性的虚拟机崩溃时将会在相同的可用资源与中自动重

新启动该虚拟机。高可用特性不会跨资源域执行。CloudStack 采用比较保守的方式重启虚拟机，以确使不会同时运行两个相同的实例。管理服务器会尝试在本季群的另一台主机上开启该虚拟机。

高可用特性只在使用iSCSI和NFS做主存储的才可以使用。不支持使用本地存储作为主存储的高可用。

17.4.1. Dedicated HA Hosts

One or more hosts can be designated for use only by HA-enabled VMs that are restarting due to a host failure. Setting up a pool of such dedicated HA hosts as the recovery destination for all HA-enabled VMs is useful to:

- Make it easier to determine which VMs have been restarted as part of the CloudStack high-availability function. If a VM is running on a dedicated HA host, then it must be an HA-enabled VM whose original host failed. (With one exception: It is possible for an administrator to manually migrate any VM to a dedicated HA host.).

- Keep HA-enabled VMs from restarting on hosts which may be reserved for other purposes.

The dedicated HA option is set through a special host tag when the host is created. To allow the administrator to dedicate hosts to only HA-enabled VMs, set the global configuration variable ha.tag to the desired tag (for example, "ha_host"), and restart the Management Server. Enter the value in the Host Tags field when adding the host(s) that you want to dedicate to HA-enabled VMs.

注意

If you set ha.tag, be sure to actually use that tag on at least one host in your cloud. If the tag specified in ha.tag is not set for any host in the cloud, the HA-enabled VMs will fail to restart after a crash.

17.5. 主存储故障和数据丢失

当主存储发生故障，中断虚拟机管理程序立即停止该存储设备上存储的所有虚拟机。客户机被标记为，当主存储重新上线时，HA根据实际情况尽快将重新启动。使用NFS时，虚拟机管理程序可能允许虚拟机继续运行，这取决于问题的性质。例如，NFS挂起将导致客户虚拟机暂停，直至恢复存储连接。主存储没有被设计进行备份。在主存储中的单个卷，可以使用快照备份。

17.6. 二级存储的中断和数据丢失

由于一个资源域只有一个二级存储服务器，二级存储的中断将会对系统的一些功能产生影响，但不能影响正在运行的客户虚拟机。可能会让用户无法选择模版来创建虚拟机。用户也可能无法保存快照，检查或恢复已保存的快照。当二级存储恢复连接后，这些功能也就可以自动恢复。

二级存储的数据丢失将会影响最近添加的用户数据，包括模版、快照、和ISO镜像。二级存储应该进行定期备份。为每个资源域提供多个二级存储服务器能够增强系统的可扩展性。

17.7. Limiting the Rate of API Requests

You can limit the rate at which API requests can be placed for each account. This is useful to avoid malicious attacks on the Management Server, prevent performance degradation, and provide fairness to all accounts.

If the number of API calls exceeds the threshold, an error message is returned for any additional API calls. The caller will have to retry these API calls at another time.

17.7.1. Configuring the API Request Rate

To control the API request rate, use the following global configuration settings:

- api.throttling.enabled - Enable/Disable API throttling. By default, this setting is false, so API throttling is not enabled.
- api.throttling.interval (in seconds) - Time interval during which the number of API requests is to be counted. When the interval has passed, the API count is reset to 0.
- api.throttling.max - Maximum number of APIs that can be placed within the api.throttling.interval period.
- api.throttling.cachesize - Cache size for storing API counters. Use a value higher than the total number of accounts managed by the cloud. One cache entry is needed for each account, to store the running API total for that account.

17.7.2. Limitations on API Throttling

The following limitations exist in the current implementation of this feature.

注意

> Even with these limitations, CloudStack is still able to effectively use API throttling to avoid malicious attacks causing denial of service.

- In a deployment with multiple Management Servers, the cache is not synchronized across them. In this case, CloudStack might not be able to ensure that only the exact desired number of API requests are allowed. In the worst case, the number of API calls that might be allowed is (number of Management Servers) * (api.throttling.max).
- The API commands resetApiLimit and getApiLimit are limited to the Management Server where the API is invoked.

第 18

管理云

18.1. Using Tags to Organize Resources in the Cloud

A tag is a key-value pair that stores metadata about a resource in the cloud. Tags are useful for categorizing resources. For example, you can tag a user VM with a value that indicates the user's city of residence. In this case, the key would be "city" and the value might be "Toronto" or "Tokyo." You can then request CloudStack to find all resources that have a given tag; for example, VMs for users in a given city.

You can tag a user virtual machine, volume, snapshot, guest network, template, ISO, firewall rule, port forwarding rule, public IP address, security group, load balancer rule, project, VPC, network ACL, or static route. You can not tag a remote access VPN.

You can work with tags through the UI or through the API commands createTags, deleteTags, and listTags. You can define multiple tags for each resource. There is no limit on the number of tags you can define. Each tag can be up to 255 characters long. Users can define tags on the resources they own, and administrators can define tags on any resources in the cloud.

An optional input parameter, "tags," exists on many of the list* API commands. The following example shows how to use this new parameter to find all the volumes having tag region=canada OR tag city=Toronto:

```
command=listVolumes
    &listAll=true
    &tags[0].key=region
    &tags[0].value=canada
    &tags[1].key=city
    &tags[1].value=Toronto
```

The following API commands have the "tags" input parameter:

- listVirtualMachines
- listVolumes
- listSnapshots
- listNetworks
- listTemplates
- listIsos
- listFirewallRules
- listPortForwardingRules
- listPublicIpAddresses
- listSecurityGroups
- listLoadBalancerRules
- listProjects

- listVPCs
- listNetworkACLs
- listStaticRoutes

18.2. 改变数据库配置

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.

18.3. Changing the Database Password

You may need to change the password for the MySQL account used by CloudStack. If so, you'll need to change the password in MySQL, and then add the encrypted password to /etc/cloud/management/db.properties.

1. Before changing the password, you'll need to stop CloudStack's management server and the usage engine if you've deployed that component.

   ```
   # service cloudstack-management stop
   # service cloudstack-usage stop
   ```

2. Next, you'll update the password for the CloudStack user on the MySQL server.

   ```
   # mysql -u root -p
   ```

   At the MySQL shell, you'll change the password and flush privileges:

   ```
   update mysql.user set password=PASSWORD("newpassword123") where User='cloud';
   flush privileges;
   quit;
   ```

3. The next step is to encrypt the password and copy the encrypted password to CloudStack's database configuration (/etc/cloud/management/db.properties).

   ```
   # java -classpath /usr/share/java/cloud-jasypt-1.8.jar \ org.jasypt.intf.cli.JasyptPBEStringEncryptionCLI
    encrypt.sh \ input="newpassword123" password="`cat /etc/cloud/management/key`" \ verbose=false
   ```

   ### File encryption type

   Note that this is for the file encryption type. If you're using the web encryption type then you'll use password="management_server_secret_key"

4. Now, you'll update /etc/cloud/management/db.properties with the new ciphertext. Open /etc/cloud/management/db.properties in a text editor, and update these parameters:

```
db.cloud.password=ENC(encrypted_password_from_above)
db.usage.password=ENC(encrypted_password_from_above)
```

5. After copying the new password over, you can now start CloudStack (and the usage engine, if necessary).

```
        # service cloudstack-management start
        # service cloudstack-usage start
```

18.4. 管理员告警

系统提供告警和事件用以帮助云的管理。告警通知管理员，一般用邮件发送，提醒管理员云中有错误发生。告警的行为可以进行配置。

会追踪云中所有用户和管理员的操作事件。如，每个客户虚拟机的启动会建立一个对应的事件。每个时间都存储在管理节点的数据库中。

在以下情况，会发送邮件给管理员：

- 管理节点集群中CPU，内存或者存储资源的可用量低。
- 管理节点和主机之间的心跳检查丢失超过3分钟。
- 主机集群中CPU，内存或者存储资源的可用量低。

18.5. 自定义网络域名

根管理员在网络，帐户，域，资源域以及整个CloudStack安装级别可选择的设置DNS后缀，一个域管理员可以在自己的域做这样的设置．要自定义域名并使其有效，请跟随下面的步骤．

1. 在希望的范围内设置DNS后缀

   - 在网络级别，DNS后缀可以通过UI在创建一个新的网络时设置，这些在 第 15.6.1 节 “添加另一个虚拟机的网络” 和CloudStack API的updateNetwork命令都有描述．
   - 在账户，域或资源域级别，DNS后缀可以指派给合适的CloudStack API命令：createAccount，editAccount，createDomain，editDomain，createZone或editZone．
   - 在全局范围内，使用配置参数guest.domain.suffix．你也可以使用CloudStack API命令updateConfiguration．当更改了这个全局配置后，重启管理服务器的服务以便新的设置有效．

2. 为了使你的新DNS后缀对已经存在的网络有效，你需要调用CloudStack API命令updateNetwork．对于DNS后缀已经更改后新建的网络这一步不是必需的．

你使用的网域的源取决于下面的一些规则．

- 对于所有的网络,如果网域作为这个网络自己配置的一部分，那这个网域的值会被使用．
- 对于账户指定的网络，为这个账户指定的网域会被使用．如果没有指定，系统会按照域，资源域，全局配置的顺序查找网域的值．

- 对于域指定的网络，为这个域指定的网域会被使用．如果没有指定，系统会按照 资源域，全局配置的顺序查找网域的值．

- 对于资源域指定的网络，为这个资源域指定的网域会被使用．如果没有指定，系统会在全局配置里查找网域的值．

18.6. 停止重启管理服务器

超级管理需要经常性的关闭和重启管理服务器

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..

停止管理服务器，在操作系统提示符下管理服务器节执行以下命令：

```
# service cloudstack-management stop
```

启动管理服务器：

```
# service cloudstack-management start
```

停止管理服务器：

```
# service cloudstack-management stop
```

Global Configuration Parameters

19.1. 设置全局配置参数

你可以通过设置CloudStack提供的参数控制云的多个方面.CloudStack首次安装后,在此后定期,您可能需要修改这些设置。

1. 使用管理员账号登陆

2. 在左侧导航栏，点击 全局设置

3. 在选择视图中，选择下列操作之一：

   - 全局设置

   - hypervsior 容量。这里显示列出了不同hypervsior版本所支持的最大虚拟机数量。

4. 使用搜索框来缩小列表中你所感兴趣的那些

5. 单击“编辑”图标修改一个值。如果您正在查看的hypervsior容量，您必须单击hypervsior第一次显示编辑屏幕的名称。

19.2. About Global Configuration Parameters

CloudStack provides a variety of settings you can use to set limits, configure features, and enable or disable features in the cloud. Once your Management Server is running, you might need to set some of these global configuration parameters, depending on what optional features you are setting up.

To modify global configuration parameters, use the steps in "Setting Global Configuration Parameters."

The documentation for each CloudStack feature should direct you to the names of the applicable parameters. Many of them are discussed in the CloudStack Administration Guide. The following table shows a few of the more useful parameters.

Field	å#¼
management.network.cidr	A CIDR that describes the network that the management CIDRs reside on. This variable must be set for deployments that use vSphere. It is recommended to be set for other deployments as well. Example: 192.168.3.0/24.
xen.setup.multipath	For XenServer nodes, this is a true/false variable that instructs CloudStack to enable iSCSI multipath on the XenServer Hosts when they are added. This defaults to false. Set it to true if you would

Field	å#¼
	like CloudStack to enable multipath. If this is true for a NFS-based deployment multipath will still be enabled on the XenServer host. However, this does not impact NFS operation and is harmless.
secstorage.allowed.internal.sites	This is used to protect your internal network from rogue attempts to download arbitrary files using the template download feature. This is a comma-separated list of CIDRs. If a requested URL matches any of these CIDRs the Secondary Storage VM will use the private network interface to fetch the URL. Other URLs will go through the public interface. We suggest you set this to 1 or 2 hardened internal machines where you keep your templates. For example, set it to 192.168.1.66/32.
use.local.storage	Determines whether CloudStack will use storage that is local to the Host for data disks, templates, and snapshots. By default CloudStack will not use this storage. You should change this to true if you want to use local storage and you understand the reliability and feature drawbacks to choosing local storage.
host	This is the IP address of the Management Server. If you are using multiple Management Servers you should enter a load balanced IP address that is reachable via the private network.
default.page.size	Maximum number of items per page that can be returned by a CloudStack API command. The limit applies at the

Field	å#¼
	cloud level and can vary from cloud to cloud. You can override this with a lower value on a particular API call by using the page and pagesize API command parameters. For more information, see the Developer's Guide. Default: 500.
ha.tag	The label you want to use throughout the cloud to designate certain hosts as dedicated HA hosts. These hosts will be used only for HA-enabled VMs that are restarting due to the failure of another host. For example, you could set this to ha_host. Specify the ha.tag value as a host tag when you add a new host to the cloud.

第 20

CloudStack API

CloudStack API是用来实现 CloudStack 的一个低层次API；

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[1] and the API Reference[2].

20.1. 自服务和验证API

CloudStack期望客户户有他自己用户的自服务基础架构. 它提供与已有的这些系统集成的API, 通过调用CloudStack的这些API来添加/删除用户.

CloudStack 支持插件类型的验证组件. 默认情况下, CloudStack 假定通过提供用户的密码,作为本地的验证方式. 但外部的验证也是可能的. 例如, 参考使用LDAP服务器来进行用户的验证.

20.2. 分配器

CloudStack 是管理员能够写定制的分配器, 用于选择放置新的客户机的主机和分配客户机虚拟磁盘镜像的存储主机。

20.3. 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

[1] http://docs.cloudstack.org/CloudStack_Documentation/Developer's_Guide%3A_CloudStack
[2] http://docs.cloudstack.org/CloudStack_Documentation/API_Reference%3A_CloudStack

- 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

性能调优

本节提供了如何提高你的云性能的提示。

21.1. 性能监控

终端用户和管理员都能使用宿主机和虚拟机的性能监控。性能监控将允许用户监控他们的资源利用和帮助决定何时将需要选择一个更高的系统性能或者更大的硬盘。

21.2. Increase Management Server Maximum Memory

If the Management Server is subject to high demand, the default maximum JVM memory allocation can be insufficient. To increase the memory:

1. Edit the Tomcat configuration file:

   ```
   /etc/cloud/management/tomcat6.conf
   ```

2. Change the command-line parameter -XmxNNNm to a higher value of N.

   For example, if the current value is -Xmx128m, change it to -Xmx1024m or higher.

3. To put the new setting into effect, restart the Management Server.

   ```
   # service cloudstack-management restart
   ```

For more information about memory issues, see "FAQ: Memory" at Tomcat Wiki.[1]

21.3. Set Database Buffer Pool Size

It is important to provide enough memory space for the MySQL database to cache data and indexes:

1. Edit the 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. 重启MySQL服务.

   ```
   # service mysqld restart
   ```

[1] http://wiki.apache.org/tomcat/FAQ/Memory

For more information about the buffer pool, see "The InnoDB Buffer Pool" at MySQL Reference Manual[2].

21.4. Set and Monitor Total VM Limits per Host

CloudStack管理员应该监视每个集群中的虚拟机实例的总数，并禁止分配给集群接近最大的虚拟机管理程序可以处理。一定要留一个安全余量，允许一个或多个主机出现故障，这会增加其他主机上的虚拟机负载，增加自动重新部署的可能性。查询你选择虚拟机管理程序的文档，来发现的最大许可数量每台主机上的虚拟机，然后使用CloudStack的全局配置设置设置为默认的限制。在每个群集在任何时候都监控虚拟机的活动。保持总的虚拟机数量低于安全水平，允许偶尔主机发生故障。例如，如果有N个集群中的主机，你想允许在集群中的一台主机是在任何给定的时间内，总数量的虚拟机实例，你可以允许集群中最多（N-1 ）*（每个主机的限制）。集群一旦达到这个数字的虚拟机，使用CloudStack UI禁用分配更多的虚拟机的集群。

21.5. 配置XenServer的dom0的内存

配置XenServer dom0的设置，分配更多的内存给dom0。这可以使XenServer能处理大量的虚拟机。我们建议2940 MB RAM给XenServer dom0。对于如何做到这一点的说明，请参阅：思杰知识库文章[3]。文章指的是XenServer 5.6，但相同的信息适用于XenServer 6。

[2] http://dev.mysql.com/doc/refman/5.5/en/innodb-buffer-pool.html

[3] http://support.citrix.com/article/CTX126531

故障排查

22.1. 事件

An event is essentially a significant or meaningful change in the state of both virtual and physical resources associated with a cloud environment. Events are used by monitoring systems, usage and billing systems, or any other event-driven workflow systems to discern a pattern and make the right business decision. In CloudStack an event could be a state change of virtual or psychical resources, an action performed by an user (action events), or policy based events (alerts).

22.1.1. 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..

22.1.2. Event Notification

Event notification framework provides a means for the Management Server components to publish and subscribe to CloudStack events. Event notification is achieved by implementing the concept of event bus abstraction in the Management Server. An event bus is introduced in the Management Server that allows the CloudStackcomponents and extension plug-ins to subscribe to the events by using the Advanced Message Queuing Protocol (AMQP) client. In CloudStack, a default implementation of event bus is provided as a plug-in that uses the RabbitMQ AMQP client. The AMQP client pushes the published events to a compatible AMQP server. Therefore all the CloudStack events are published to an exchange in the AMQP server.

A new event for state change, resource state change, is introduced as part of Event notification framework. Every resource, such as user VM, volume, NIC, network, public IP, snapshot, and template, is associated with a state machine and generates events as part of the state change. That implies that a change in the state of a resource results in a state change event, and the event is published in the corresponding state machine on the event bus. All the CloudStack events (alerts, action events, usage events) and the additional category of resource state change events, are published on to the events bus.

Use Cases

The following are some of the use cases:

- Usage or Billing Engines: A third-party cloud usage solution can implement a plug-in that can connects to CloudStack to subscribe to CloudStack events and generate usage data. The usage data is consumed by their usage software.

- AMQP plug-in can place all the events on the a message queue, then a AMQP message broker can provide topic-based notification to the subscribers.

- Publish and Subscribe notification service can be implemented as a pluggable service in CloudStack that can provide rich set of APIs for event notification, such as topics-based subscription and notification. Additionally, the pluggable service can deal with multi-tenancy, authentication, and authorization issues.

云平台配置

As a CloudStack administrator, perform the following one-time configuration to enable event notification framework. At run time no changes can control the behaviour.

1. Open 'componentContext.xml.

2. Define a bean named eventNotificationBus as follows:

   - name : Specify a name for the bean.

   - server : The name or the IP address of the RabbitMQ AMQP server.

   - port : The port on which RabbitMQ server is running.

   - username : The username associated with the account to access the RabbitMQ server.

   - password : The password associated with the username of the account to access the RabbitMQ server.

   - exchange : The exchange name on the RabbitMQ server where CloudStack events are published.

     A sample bean is given below:

     ```
     <bean id="eventNotificationBus" class="org.apache.cloudstack.mom.rabbitmq.RabbitMQEventBus">
         <property name="name" value="eventNotificationBus"/>
         <property name="server" value="127.0.0.1"/>
         <property name="port" value="5672"/>
         <property name="username" value="guest"/>
         <property name="password" value="guest"/>
        <property name="exchange" value="cloudstack-events"/>
        </bean>
     ```

     The eventNotificationBus bean represents the org.apache.cloudstack.mom.rabbitmq.RabbitMQEventBus class.

3. 重启管理服务器.

22.1.3. Standard Events

The events log records three types of standard events.

- INFO. This event is generated when an operation has been successfully performed.

- WARN. This event is generated in the following circumstances.

  - When a network is disconnected while monitoring a template download.

  - When a template download is abandoned.

  - When an issue on the storage server causes the volumes to fail over to the mirror storage server.

- ERROR. This event is generated when an operation has not been successfully performed

22.1.4. 长时间运行的任务事件

事件日志记录了三种类型的标志事件.

- INFO. 当一个操作成功执行时会产生这样的事件.
- WARN. 这个事件在以下情况下产生.
  - 当一个模板正在下载时检测到网络断开.
  - 当放弃一个模板的下载.
  - 一个存储服务器上的错误导致卷故障无法转移到另一个镜像的存储服务器.
- ERROR. 当一个操作没有被成功执行时会产生这个事件.

22.1.5. 事件日志查询

数据库日志可以从用户界面查询。系统所捕获的事件的列表包括:

- 虚拟机的创建, 删除和持续的管理操作
- 虚拟路由器创建, 删除和持续的管理操作
- 模板的创建和删除
- 网络/负载均衡器规则的创建和删除
- 存储卷的创建和删除
- 用户登录和注销

22.2. Working with Server Logs

The CloudStack Management Server logs all web site, middle tier, and database activities for diagnostics purposes in /var/log/cloudstack/management/. The CloudStack logs a variety of error messages. We recommend this command to find the problematic output in the Management Server log:.

注意

当拷贝粘贴一条命令，确保在运行前粘贴的命令在一行上. 一些文档查看器可能会在拷贝时引入不希望的换行符.

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

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/cloudstack/agent/.

22.3. Data Loss on Exported Primary Storage

症状

Loss of existing data on primary storage which has been exposed as a Linux NFS server export on an iSCSI volume.

原因

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

解决方案

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

22.4. Recovering a Lost Virtual Router

症状

A virtual router is running, but the host is disconnected. A virtual router no longer functions as expected.

原因

The Virtual router is lost or down.

解决方案

If you are sure that a virtual router is down forever, or no longer functions as expected, destroy it. You must create one afresh while keeping the backup router up and running (it is assumed this is in a redundant router setup):

- Force stop the router. Use the stopRouter API with forced=true parameter to do so.
- Before you continue with destroying this router, ensure that the backup router is running. Otherwise the network connection will be lost.
- Destroy the router by using the destroyRouter API.

Recreate the missing router by using the restartNetwork API with cleanup=false parameter. For more information about redundant router setup, see Creating a New Network Offering.

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

22.5. Maintenance mode not working on vCenter

症状

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

原因

The CloudStack administrator UI was used to place the host in scheduled maintenance mode. This mode is separate from vCenter's maintenance mode.

解决方案

Use vCenter to place the host in maintenance mode.

More Information

See 第 11.2 节 "主机定期维护和维护模式"

22.6. 无法从上传的vSphere模板部署虚拟机

症状

当试图创建一个虚拟机，虚拟机将无法部署。

原因

如果模板通过上传OVA文件创建，而OVA文件是使用vSphere Client创建的，可能OVA中包含ISO镜像。如果是的话，从模板部署虚拟机将失败。

解决方案

移除ISO并重新上传模板。

22.7. Unable to power on virtual machine on VMware

症状

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

原因

A known issue on VMware machines. ESX hosts lock certain critical virtual machine files and file systems to prevent concurrent changes. Sometimes the files are not unlocked when the virtual machine is powered off. When a virtual machine attempts to power on, it can not access these critical files, and the virtual machine is unable to power on.

解决方案

See the following:

VMware Knowledge Base Article[1]

22.8. Load balancer rules fail after changing network offering

症状

After changing the network offering on a network, load balancer rules stop working.

原因

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.

解决方案

Create a firewall rule on the virtual router for each of your existing load balancing rules so that they continue to function.

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

附录 A. 时区

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	Etc/GMT+11Etc/GMT+11	Pacific/Samoa
Pacific/Honolulu	US/Alaska	America/Los_Angeles
Mexico/BajaNorte	US/Arizona	US/Mountain
America/Chihuahua	America/Chicago	America/Costa_Rica
America/Mexico_City	Canada/Saskatchewan	America/Bogota
America/New_York	America/Caracas	America/Asuncion
America/Cuiaba	America/Halifax	America/La_Paz
America/Santiago	America/St_Johns	America/Araguaina
America/Argentina/Buenos_Aires	America/Cayenne	America/Godthab
America/Montevideo	Etc/GMT+2	Atlantic/Azores
Atlantic/Cape_Verde	Africa/Casablanca	Etc/UTC
Atlantic/Reykjavik	Europe/London	CET
Europe/Bucharest	Africa/Johannesburg	Asia/Beirut
Africa/Cairo	Asia/Jerusalem	Europe/Minsk
Europe/Moscow	Africa/Nairobi	Asia/Karachi
Asia/Kolkata	Asia/Bangkok	Asia/Shanghai
亚洲/吉隆坡	澳大利亚/珀斯	亚洲/台北
亚洲/东京	亚洲/首尔	澳大利亚/阿德莱德
澳大利亚/达尔文	澳大利亚/布里斯班	澳大利亚/堪培拉
太平洋/关岛	太平洋/奥克兰	

附录 B. Event Types

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

附录 C. Alerts

The following is the list of alert type numbers. The current alerts can be found by calling listAlerts.

```
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
```

附录 D. 更新记录

修订 0-0	Tue May 29 2012	Jessica Tomechak

由publican生成的最初版本