Cloud.com CloudStack Installation Guide

Version 2.2 early access release

November 20, 2010

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

DRAFT

© 2010 Cloud.com, Inc. All rights reserved. Specifications are subject to change without notice. The Cloud.com logo, Cloud.com, CloudStack, Hypervisor Attached Storage, HAS, Hypervisor Aware Network, HAN, and VMSync are trademarks or registered trademarks of Cloud.com, Inc. All other brands or products are trademarks or registered trademarks of their respective holders.

Contents

1 Overview 7

2 Prerequisites 8

3 Choosing a Deployment Architecture 10

3.1 Small-Scale Deployment 10

3.2 Large-Scale Redundant Setup 11

3.3 Separate Storage Network 12

4 Network Setup 15

4.1 VLAN Allocation 15

4.1.1 VLAN Allocation with Virtual Networking 16

4.1.2 VLAN Allocation with Tagged Direct Attached 17

4.1.3 VLAN Allocation with Virtual Networking and Tagged Direct Attached Networking 17

4.2 Layer-3 Switch 18

4.2.1 Example Configuration 18

4.3 Layer-2 Switch 19

4.3.1 Example Configurations 19

4.4 Hardware Firewall 20

4.5 Load Balancer 20

4.6 Additional Topology Requirements 21

5 Storage Setup 22

5.1 Small-Scale Setup 22

5.2 Secondary Storage 22

5.3 Example Configurations 22

5.3.1 Linux NFS on Local Disks and DAS 22

5.3.2 Linux NFS on iSCSI 23

6 Citrix XenServer Installation and Configuration 25

6.1 Username and Password 25

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

6.2 Time Synchronization 25
6.3 Licensing 25
6.3.1 Getting and Deploying a License 25
6.4 Patch for 2.1.2 and Previous Releases 26
6.5 Networking Setup 26
6.5.1 Configuring Public Network with a Dedicated NIC (optional) 26
6.5.2 NIC Bonding (optional) 27
6.5.3 Separate Storage Network (optional) 28
6.6 iSCSI Multipath Setup (optional) 29
7 VMware vSphere Installation and Configuration 30
7.1 ESX/ESXi Host setup 30
7.2 Configure Host Networking 30
7.2.1 Configure Virtual Switch 31
7.2.2 Configure vCenter Management Network 34
7.2.3 Configure NIC Bonding 35
7.3 Configure Storage Multipathing 35
7.4 Add Hosts or Configure Clusters 35
7.4.1 Clusters 35
7.4.2 Standalone Nodes 35
8 KVM Installation and Configuration 36
8.1 Network Configuration 36
8.2 Installing the CloudStack Agent on a Computing Node 36
9 Management Server Installation 38
9.1 Operating System and OS Preparation 38
9.2 Single Node Install (One Management Server) 38
9.2.1 Single Node Database Install 39
9.3 Multinode Install (Multiple Management Servers) 40

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

9.3.1 Install the First Management Server 40
9.3.2 Install the Database 40
9.3.3 Database Replication (Optional) 41
9.3.4 Creating and Initializing the Database 43
9.3.5 OS Configuration for the Management Server 43
9.3.6 Prepare and Start Additional Management Servers 43
10 Prepare Secondary Storage 45
11 Describe Your Deployment 46
11.1 Add a New Zone 47
11.1.1 Adding a Zone 47
11.1.2 Advanced Networking: Adding a Public IP Address Range on the Zone 50
11.1.3 Additional Zones 50
11.1.4 Additional Pods 50
11.1.5 Additional Clusters 50
11.2 Edit Service Offerings (Optional) 51
11.3 Edit Disk Offerings (Optional) 51
11.4 Add Hosts 52
11.5 Add Primary Storage 54
11.5.1 Secondary Storage 55
11.6 SSL 56
12 Initialization and Testing 57
13 Installing the Usage Server (Optional) 58
14 Troubleshooting 59
14.1 Troubleshooting the Secondary Storage VM 59
14.1.1 Running a Diagnostic Script 59
14.1.2 Checking the Log File 59
14.2 Troubleshooting the Console Proxy VM 60

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

15 Contacting Support 61

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

1 Overview

Cloud.com™ CloudStack™ Version 2.2 is designed to work with a wide variety of enterprise-grade and commodity network and storage infrastructure including the following:

- Layer-3 switching at the core and layer-2 switching at the edge. With layer-3 switching at the core, there is no limit on the number of physical servers that can be managed in a cloud.
- 1-GbE and 10-GbE Ethernet NICs and switches
- Redundant network setup with bonded NICs
- NFS and iSCSI storage

The CloudStack consists of two types of nodes:

- **CloudStack Management Server**: The server in this node is the resource manager in the system. It controls allocation of virtual machines to servers in the Computing Node and assigns storage and IP addresses to the virtual machine instances.
- **CloudStack Computing Node**: The servers in this node run the virtual machine instances. Servers are grouped into Zones and Pods. The initial installation has one Zone and one Pod.
    - **Zone**: A Zone consists of multiple Pods. Typically a Zone is a datacenter.
    - **Pod**: A Pod is usually one rack of hardware and will include primary storage and a layer-2 switch.

A small installation may consist of one Management Server and several Computing Nodes. Additional Computing Nodes can be added after the initial installation. The CloudStack Management Server is installed on a RHEL/CentOS 5.4+ system that can be a VM or a dedicated server.

This guide contains detailed information about the following recommended steps for installing the Cloud.com CloudStack.

1. Choosing a deployment architecture
2. Set up networking
3. Set up storage
4. Install Citrix XenServer/VMware vSphere/KVM
5. Install the CloudStack Management Server
6. Prepare secondary storage
7. Describe the deployment
8. Test the deployment

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

2 Prerequisites

The Cloud.com CloudStack has the following hardware and software requirements.

Important: Physical hosts should be selected and racked before starting installation.

	Description	Minimum Requirements
Management Server	Hosts the Cloud.com CloudStack Management Server software.	• 64-bit x86 CPU (more cores results in better performance) • 2 GB of memory • 80 GB of local disk • At least 1 NIC • RHEL/CentOS 5.4+ 64-bit • Statically allocated IP address • Fully qualified domain name as returned by the hostname command
Computing Node	Provides all the CPU and memory resource for allocated guest virtual machines.	• 64-bit x86 CPU (more cores results in better performance) • Hardware virtualization support required • 4 GB of memory • 30 GB of local disk • At least 1 NIC • Statically allocated IP Address • Citrix XenServer 5.6, VMware vSphere 4.1, or RHEL/CentOS 5.5 **Important: The computing server should be certified as compatible by the hypervisor vendor. You can view the Citrix Hardware Compatibility Guide at** http://hcl.xensource.com/**. You can view the VMware Hardware Compatibility Guide at** http://www.vmware.com/resources/compatibility/search.php**. You can view the RHEL Hardware Compatibility Guide at** https://hardware.redhat.com/**.**
vCenter Server	Run VMware vCenter software	• Processor – 2 CPUs 2.0GHz or higher Intel or AMD x86 processors. Processor may be higher if the database runs on the same machine. • Memory – 3GB RAM. RAM requirements may be higher if your database runs on the same machine.

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

		• Disk storage – 2GB. Disk requirements may be higher if your database runs on the same machine. • Microsoft SQL Server 2005 Express disk requirements. The bundled database requires up to 2GB free disk space to decompress the installation archive. • Networking – 1Gbit recommended.
Secondary Storage	Provides storage for templates and snapshots	• NFS storage appliance or Linux NFS server • 100GB minimum capacity
Database Node		• May be co-located with the Management Server • Otherwise requirements identical to Management Server

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

3 Choosing a Deployment Architecture

The architecture used in a deployment will vary depending on the size and purpose of the deployment. This section contains examples of deployment architecture, including a small-scale deployment useful for test and trial deployments and a fully-redundant large-scale setup for production deployments.

3.1 Small-Scale Deployment

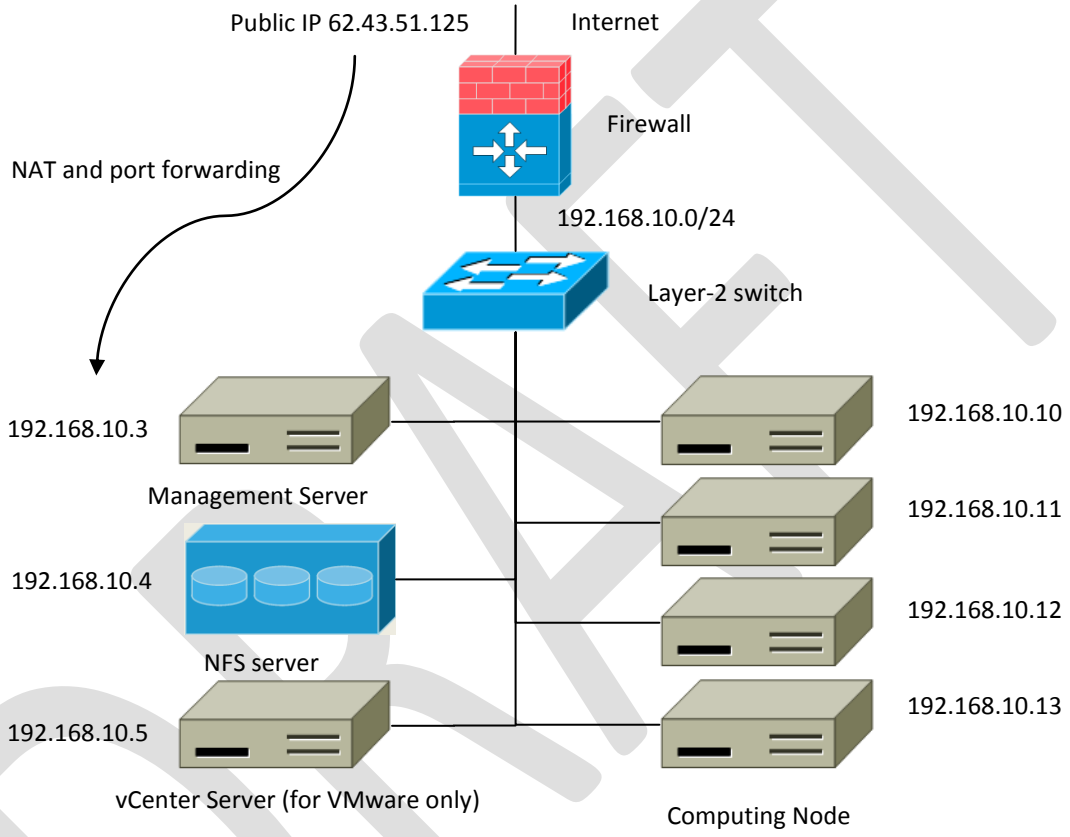

Figure 1 Small-Scale Deployment

Figure 1 illustrates the network architecture of a small-scale Cloud.com CloudStack deployment.

1. A firewall provides a connection to the Internet. The firewall is configured in NAT mode. The firewall forwards HTTP requests and API calls from the Internet to the Management Server. The Management Server resides on the private network.
2. A layer-2 switch connects all physical servers and storage.
3. A single NFS server functions as both the primary and secondary storage.
4. The Management Server is connected to the private network.

© 2010 Cloud.com, Inc. All rights reserved.

3.2 Large-Scale Redundant Setup

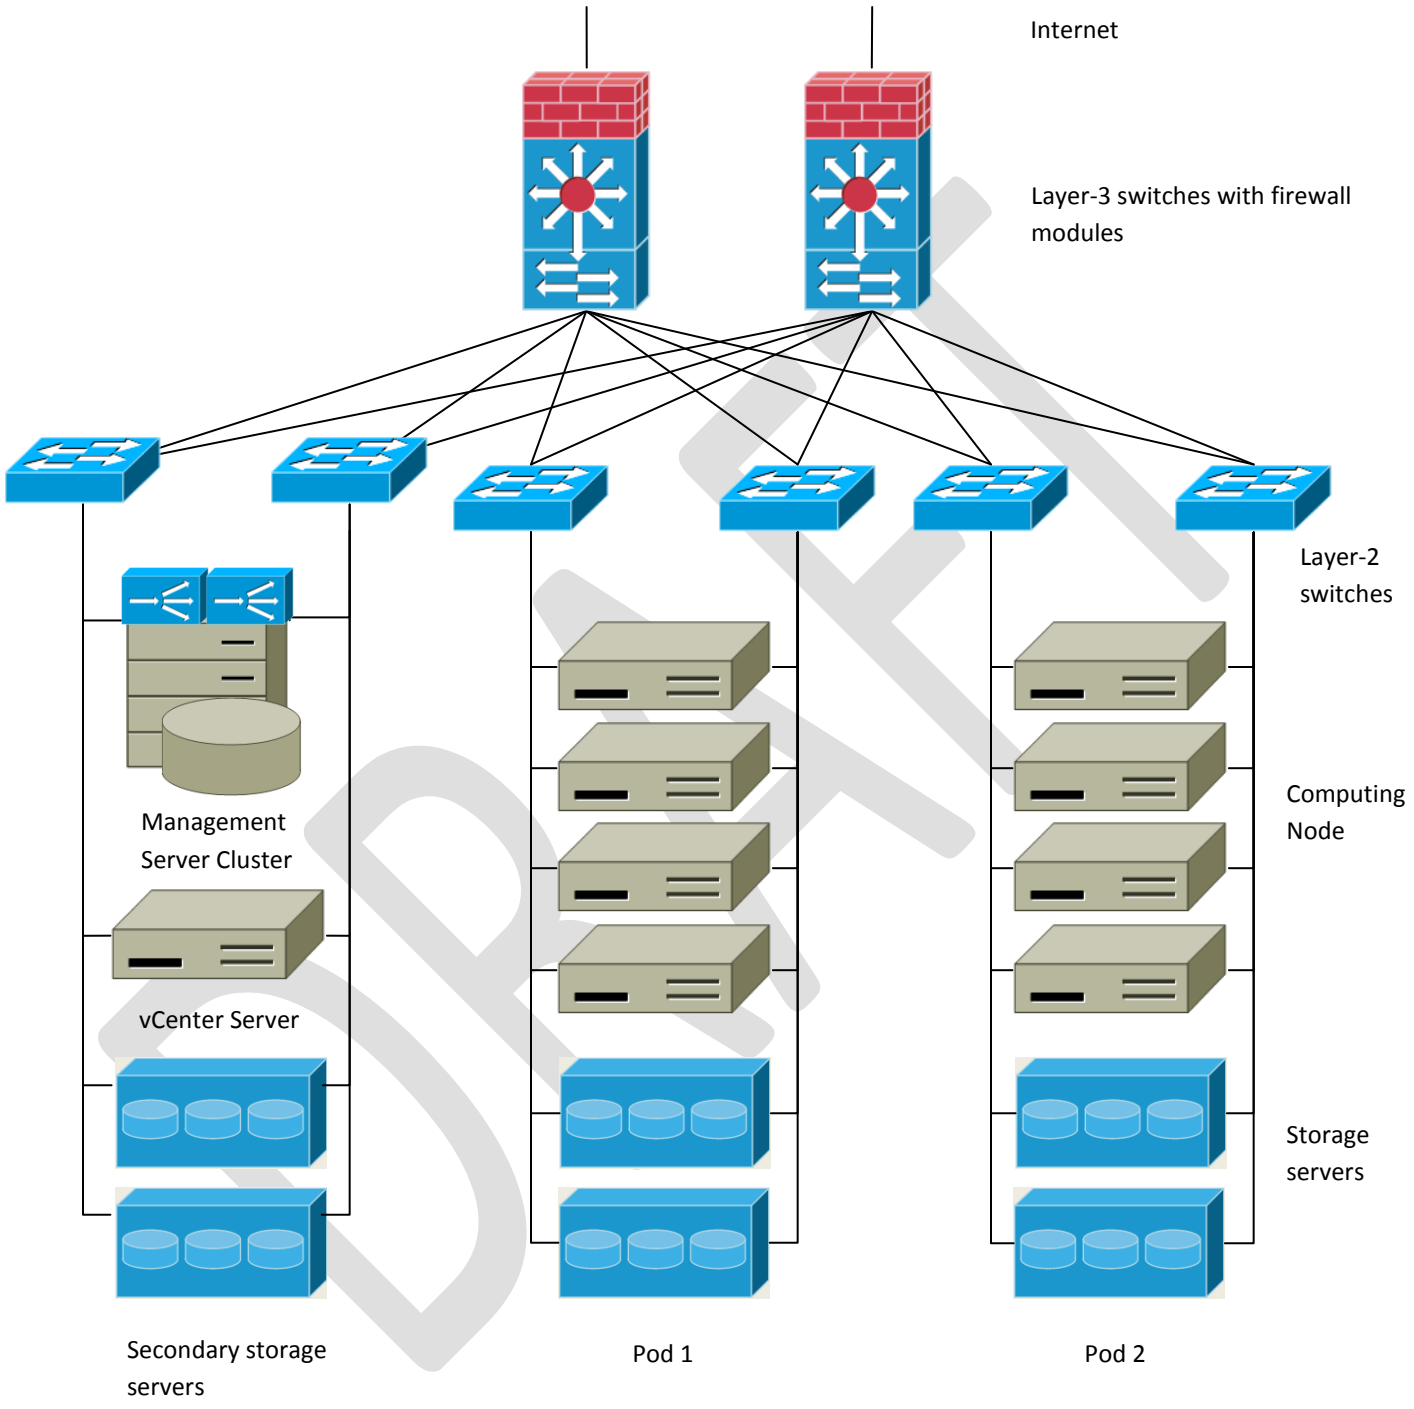

Figure 2 Large-Scale Deployment Architecture

Figure 2 illustrates the network architecture of a large-scale Cloud.com CloudStack deployment.

1. A layer-3 switching layer is at the core of the data center. A router redundancy protocol like VRRP should be deployed. Typically high-end core switches also include firewall modules. Separate firewall appliances may also be used if the layer-3

© 2010 Cloud.com, Inc. All rights reserved.

switch does not have integrated firewall capabilities. The firewalls are configured in NAT mode. The firewalls provide the following functions:

a. Forwards HTTP requests and API calls from the Internet to the Management Server. The Management Server resides on the private network.

b. When the cloud spans multiple availability Zones, the firewalls should enable site-to-site VPN such that servers in different availability Zones can directly reach each other.

2. A layer-2 access switch layer is established for each Pod. Multiple switches can be stacked to increase port count. In either case, redundant pairs of layer-2 switches should be deployed.
3. The Management Server cluster (including front-end load balancers, Management Server nodes, and the MySQL database) is connected to the private network through a pair of load balancers.
4. The secondary storage server is connected to the private network.
5. Each Pod contains storage and computing servers. Each storage and computing server should have redundant NICs connected to separate layer-2 access switches.

3.3 Separate Storage Network

In the Large-Scale Redundant setup described in the previous section, storage traffic can overload the private network. A separate storage network is optional for deployments. Storage protocols such as iSCSI are sensitive to network delays. A separate storage network ensures guest network traffic contention does not impact storage performance.

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

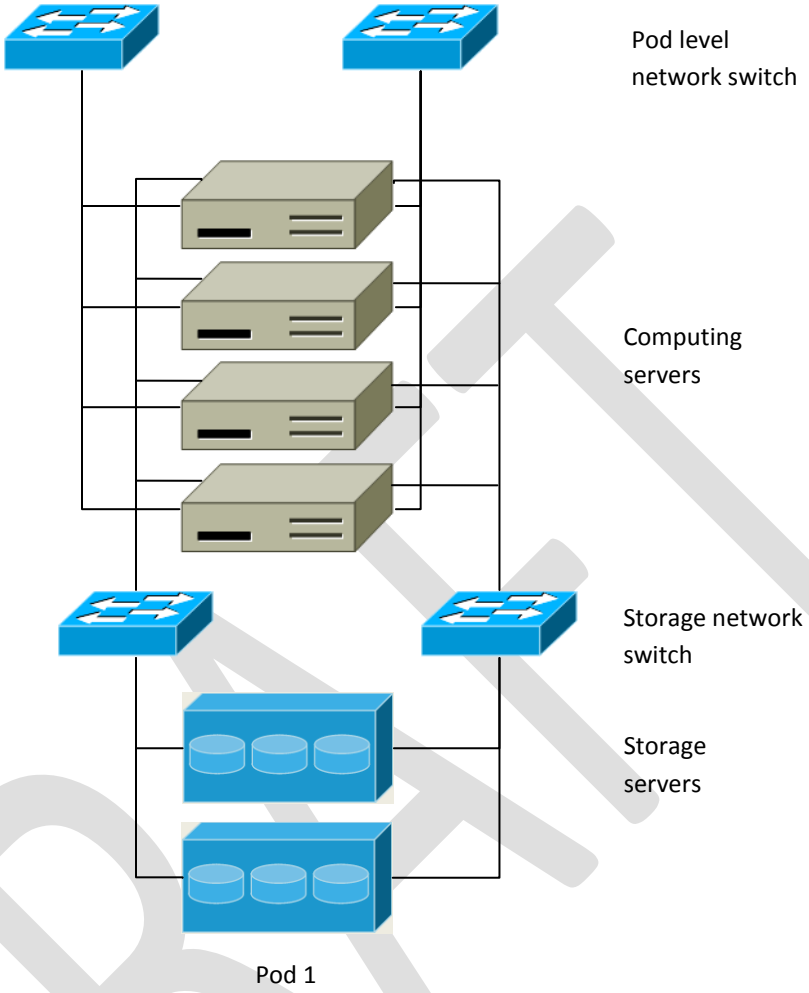

Figure 3 Separate Storage Network

Figure 3 illustrates a setup with a separate storage network. Each server has four NICs, two connected to Pod-level network switches and two connected to storage network switches.

There are two ways to configure the storage network:

1. Bonded NIC and redundant switches can be deployed for NFS. In NFS deployments, redundant switches and bonded NICs still result in one network (one CIDR block+ default gateway address).
2. iSCSI can take advantage of two separate storage networks (two CIDR blocks each with its own default gateway). Multipath iSCSI client can failover and load balance between separate storage networks.

© 2010 Cloud.com, Inc. All rights reserved.

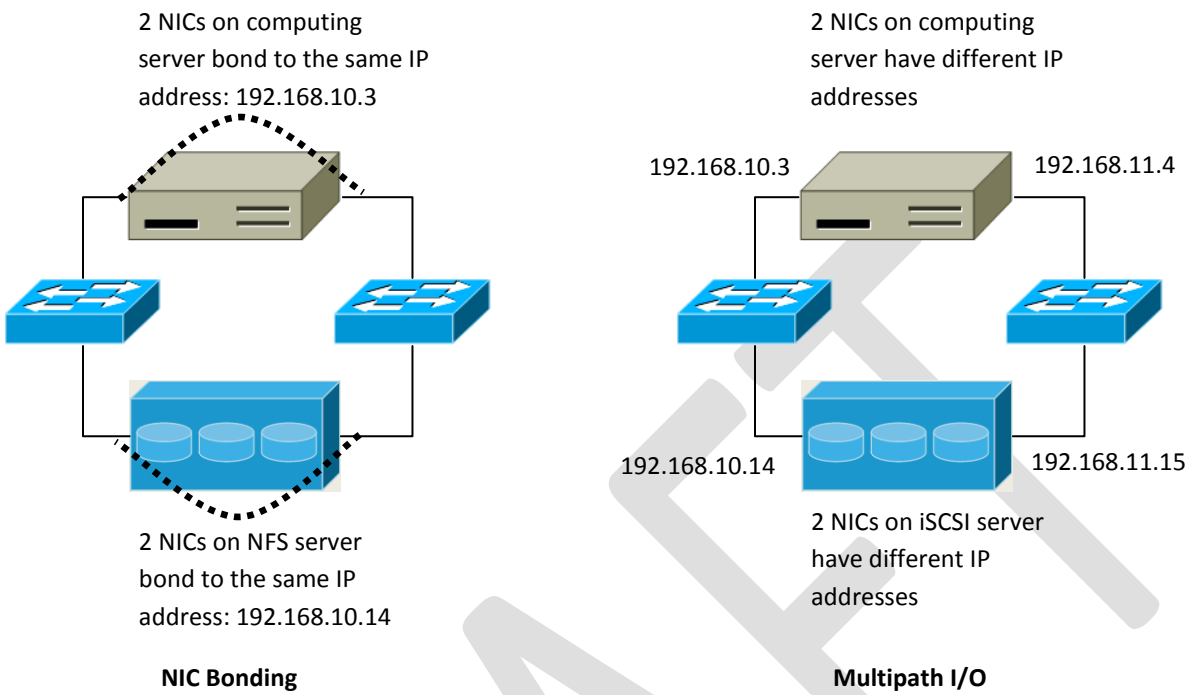

Figure 4 NIC Bonding and Multipath I/O

Figure 4 illustrates the differences between NIC bonding and Multipath I/O (MPIO). NIC bonding configuration involves only one network. MPIO involves 2 separate networks.

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

4 Network Setup

The CloudStack provides three primary networking modes:

- Virtual networking
- Direct attached with tagged VLANs ("Direct tagged")
- Direct attached with untagged VLANs ("Direct untagged")

One distinguishing factor between the modes is whether or not the guest virtual machines are part of an isolated virtual network created by the CloudStack or are attached directly to an existing network in the data center. We call these two modes "Virtual Networking" and "Direct Attached Networking".

With Virtual Networking the guests:

- Use the virtual router as the gateway and firewall
- Have an IP address defined by the Zone's Guest CIDR
- Are always part of an account-specific VLAN

With Direct Tagged Networking the guests:

- Do not use the virtual router as their gateway or firewall
- Have an IP address defined by the Direct IPs provisioned in the CloudStack
- Are on an account-specific, Zone-Wide, tagged VLAN that provides isolation

With Direct Untagged Networking the guests:

- Do not use the virtual router as their gateway or firewall
- Have an IP address defined by the Direct IPs provisioned in the CloudStack
- Are on an untagged VLAN and may have no isolation

The three types of networking may be in use in the same cloud. However, Direct Untagged networking may not be mixed with the other modes in the same Zone. A Zone supports either Direct Untagged networking or a combination of Direct Tagged and Virtual Networking.

The style of networking that a particular guest receives is determined by the Service Offering it is started from.

This chapter discusses network setup that is common to all modes as well as network setup that is specific to one or two of the modes.

4.1 VLAN Allocation

Important: CloudStack networking enables many different deployment styles. Your deployment may not need one more or of the types of VLANs. You should decide what services you want to offer to your users before provisioning VLANs.

VLAN allocation must be considered if Virtual Networking or Tagged Direct Attached Networking is required.

Cloud.com CloudStack is designed to utilize hardware VLANs to isolate guest virtual network traffic. There are three types of VLANs in the CloudStack.

© 2010 Cloud.com, Inc. All rights reserved.

1. **Public VLAN**. A range of VLAN IDs will be reserved for public IP addresses. These VLANs are trunked into every Pod.
2. **Zone VLAN**. A range of VLAN IDs will be reserved for guest virtual networks. These VLANs are trunked into every Pod. One VLAN is allocated per guest virtual network that has active instances.
3. **Direct VLAN.** A range of VLAN IDs will be reserved for assignment to accounts if Tagged Direct Attached networking is required. These VLANs are trunked into every Pod. The administrator assigns these VLANs to customers one at a time; a range is not given to the CloudStack to manage.

Figure 5 illustrates VLAN allocation in an Availability Zone:

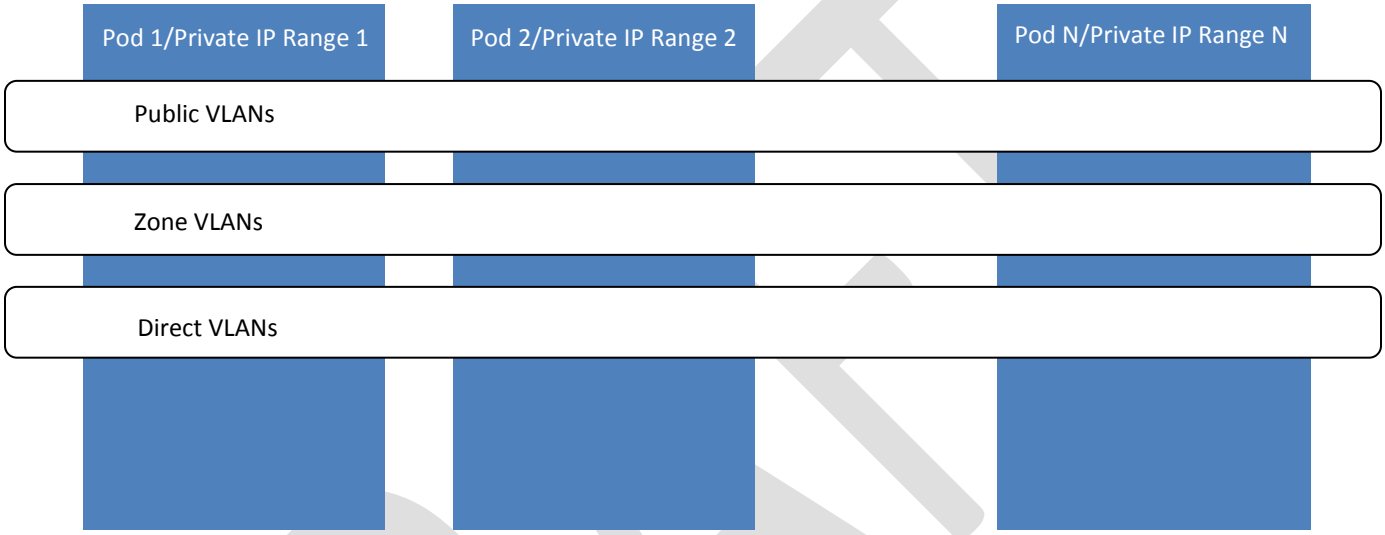

Figure 5 VLAN Allocation in an Availability Zone

The computing and storage servers reside on an Untagged Private Network. The untagged private network in each Pod is mapped to a unique VLAN and private IP range in the layer-3 switch. The layer-3 switch serves as the default gateway for each private network and ensures a packet can be routed from one host to any other host in the same Zone.

4.1.1 VLAN Allocation with Virtual Networking

With Virtual Network VLANs are required for the guests at the Zone level. The following is an example of a VLAN allocation scheme when Virtual Networking is used:

VLAN IDs	Use
< 500	Reserved for administrative purposes and untagged private network for each pod
500-599	Public VLANs
600-999	Zone VLANs

© 2010 Cloud.com, Inc. All rights reserved.

> 1000	Reserved for future use

4.1.2 VLAN Allocation with Tagged Direct Attached

With Tagged Direct Attached networking there is no need for Public VLANs nor Zone VLANs. There is a new need for a VLAN range for the Direct Attached guests.

VLAN IDs	Use
< 300	Reserved for administrative purposes and untagged private network for each pod
300-499	Directed Attached VLANs
> 500	Reserved for future use

4.1.3 VLAN Allocation with Virtual Networking and Tagged Direct Attached Networking

The CloudStack supports deployments that have both Virtual Networking and Tagged Direct Attached guests. In this case it will be necessary to allocate VLANs for both types of guests.

VLAN IDs	Use
< 300	Reserved for administrative purposes and untagged private network for each pod
300-499	Tagged Direct Attached VLANs
500-599	Public VLANs
600-999	Zone VLANs
> 1000	Reserved for future use

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

4.2 Layer-3 Switch

The layer-3 switch is the core switching layer at the Availability Zone level. The layer-3 switch should be programmed as follows:

- If direct tagged or virtual networking is in use, the layer-3 switch trunks public VLANs, Zone VLANs, and Direct Attached VLANs into each Pod.
- The layer-3 switch functions as the gateway for the untagged private network. A separate VLAN is created in the layer-3 switch for each private IP address range. The layer-3 switch should allow packets to flow between private IP ranges.

The "Virtual Network and Direct Tagged" VLAN allocation in this section is used in the configurations described for layer 2 and layer 3 switches. You can adjust VLAN allocation according to your specific needs.

4.2.1 Example Configuration

This section contains an example configuration of specific switch models for Zone-level layer-3 switching. It assumes VLAN management protocols, such as VTP or GVRP, have been disabled. The example scripts must be changed appropriately if you choose to use VTP or GVRP.

4.2.1.1 Dell 62xx

The following steps show how a Dell 62xx is configured for Zone-level layer-3 switching. These steps assume VLAN 201 is used to route untagged private IPs for Pod 1, and Pod 1's layer-2 switch is connected to Ethernet port 1/g1.

Important: Dell 62xx Series switch only supports up to 1024 VLANs.

1. Configure all the VLANs in the database.

```
vlan database
vlan 200-999
exit
```

2. Configure Ethernet port 1/g1.

```
interface ethernet 1/g1
switchport mode general
switchport general pvid 201
switchport general allowed vlan add 201 untagged
switchport general allowed vlan add 300-999 tagged
exit
```

The statements configure Ethernet port 1/g1 as follows:

- VLAN 201 is the native untagged VLAN for port 1/g1.
- Public VLANs (500-599) and Zone VLANs (600-999) and Direct Attached VLANs (300-499) are passed to all the Pod-level layer-2 switches.

4.2.1.2 Cisco 3750

The following steps show how a Cisco 3750 is configured for Zone-level layer-3 switching. These steps assume VLAN 201 is used to route untagged private IPs for Pod 1, and Pod 1's layer-2 switch is connected to GigabitEthernet1/0/1.

© 2010 Cloud.com, Inc. All rights reserved.

1. Setting VTP mode to transparent allows us to utilize VLAN IDs above 1000. Since we only use VLANs up to 999, vtp transparent mode is not strictly required.

```
vtp mode transparent
vlan 200-999
exit
```

2. Configure GigabitEthernet1/0/1.

```
interface GigabitEthernet1/0/1
switchport trunk encapsulation dot1q
switchport mode trunk
switchport trunk native vlan 201
exit
```

The statements configure GigabitEthernet1/0/1 as follows:

- VLAN 201 is the native untagged VLAN for port GigabitEthernet1/0/1.
- Cisco passes all VLANs by default. As a result Public VLANs (500-599), Zone VLANs (600-999), and Direct Attached VLANs (300-499) are passed to all the Pod-level layer-2 switches.

4.3 Layer-2 Switch

The layer-2 switch is the access switching layer inside the Pod.

- It should trunk Public VLANs, Zone VLANs, and Direct Attached VLANs into every computing host.
- It should switch untagged traffic for the private network containing computing and storage hosts. The layer-3 switch will serve as the gateway for the private network.

4.3.1 Example Configurations

This section contains example configurations for specific switch models for Pod-level layer-2 switching. It assumes VLAN management protocols such as VTP or GVRP have been disabled. The scripts must be changed appropriately if you choose to use VTP or GVRP.

4.3.1.1 Dell 62xx

The following steps show how a Dell 62xx is configured for Pod-level layer-2 switching.

1. Configure all the VLANs in the database.

```
vlan database
vlan 300-999
exit
```

2. VLAN 201 is used to route untagged private IP addresses for Pod 1, and Pod 1 is connected to this layer-2 switch.

```
interface range ethernet all
switchport mode general
switchport general allowed vlan add 300-999 tagged
```

© 2010 Cloud.com, Inc. All rights reserved.

```
exit
```

The statements configure all Ethernet ports to function as follows:

- All ports are configured the same way.
- Public VLANs (500-599) , Zone VLANs (600-999) and Direct Attached VLANs (300-499) are passed through all the ports of the layer-2 switch.

4.3.1.2 Cisco 3750

The following steps show how a Cisco 3750 is configured for Pod-level layer-2 switching.

3. Setting VTP mode to transparent allows us to utilize VLAN IDs above 1000. Since we only use VLANs up to 999, vtp transparent mode is not strictly required.

```
vtp mode transparent
vlan 300-999
exit
```

4. Configure all ports to dot1q and set 201 as the native VLAN.

```
interface range GigabitEthernet 1/0/1-24
switchport trunk encapsulation dot1q
switchport mode trunk
switchport trunk native vlan 201
exit
```

By default Cisco passes all VLANs. Cisco switches complain of the native VLAN IDs are different when 2 ports are connected together. That's why we specify VLAN 201 as the native VLAN on the layer-2 switch.

4.4 Hardware Firewall

The hardware firewall is required to serve two purposes:

- Protect the Management Server farm. NAT and port forwarding should be configured to direct traffic from the public Internet to the Management Server farm.
- Route private network traffic between multiple Availability Zones. Site-to-site VPN should be configured between multiple Availability Zones.

To achieve the above purposes you must set up fixed configurations for the firewall. Firewall rules and policies need not change as users are provisioned into the cloud. Any brand of hardware firewall that supports NAT and site-to-site VPN can be used.

4.5 Load Balancer

The 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.

Important: Persistence may still be enabled on source ports which do not require it.

Source Port	Destination Port	Protocol	Persistence Required?

© 2010 Cloud.com, Inc. All rights reserved.

80 or 443	8080	HTTP	Yes
8250	8250	TCP	Yes
8096	8096	HTTP	No

4.6 Additional Topology Requirements

The secondary storage VMs (discussed later) and console proxy VMs connect to the Management Server on port 8250. If you are using multiple Management Servers, the load balanced IP address of the Management Servers on port 8250 must be reachable. Note that you must not expose port 8250 to the public Internet. The secondary storage VM can be located on any Computing Node in the Zone. It uses the private network for its communication.

The secondary storage NFS export is mounted by the secondary storage VM. Secondary storage traffic goes over the private network even if there is a separate storage network. (Primary storage traffic goes over the storage network, if available.) If you choose to place the secondary storage NFS server on the storage network you must make sure there is a route from the private network to the storage network.

The Management Servers communicate with each other to coordinate tasks amongst themselves. This communication uses TCP on ports 8250 and 9090.

The public internet must not be able to access port 8096 on the Management Server.

The Management Servers communicate with the XenServers on ports 22 (ssh) and 80 (HTTP).

The Management Servers communicate with VMware vCenter servers on port 443 (HTTPs).

The Management Servers communicate with the KVM servers on port 22 (ssh).

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

5 Storage Setup

The Cloud.com CloudStack is designed to work with a wide variety of commodity and enterprise-grade storage that supports iSCSI and NFS protocols. Local disk may be used as well.

IMPORTANT: For 2.2 beta, vSphere supports only NFS.

IMPORTANT: For 2.2 beta, KVM supports only NFS and local disk.

5.1 Small-Scale Setup

In a small-scale setup, a single NFS server can function as both primary and secondary storage. The NFS server just needs to export two separate shares, one for primary storage and the other for secondary storage.

5.2 Secondary Storage

The CloudStack is designed to work with any scalable secondary storage system. The only requirement is the secondary storage system supports the NFS protocol.

5.3 Example Configurations

In this section we go through a few examples of how to set up storage to work properly with CloudStack on a few types of NFS and iSCSI storage systems.

5.3.1 Linux NFS on Local Disks and DAS

This section describes how to configure an NFS export on a standard Linux installation. Instructions in this section specifically apply to RHEL/CentOS 5. Steps to setup other distributions may vary.

1. Install the RHEL/CentOS distribution on the storage server.

Important: The storage server should be a machine with a large number of disks. The disks should ideally be managed by a hardware RAID controller. Modern hardware RAID controllers support hot plug functionality independent of the operating system so you can replace faulty disks without impacting the running operating system.

2. If the root volume is more than 2 TB in size, create a smaller boot volume to install RHEL/CentOS. A root volume of 20 GB should be sufficient.
3. After the system is installed, create a directory called /export. This can each be a directory in the root partition itself or a mount point for a large disk volume.
4. If you have more than 16TB of storage on one host, create multiple EXT3 file systems and multiple NFS exports. Individual EXT3 file systems cannot exceed 16TB.
5. After /export directory is created, run the following command to configure it as an NFS export.

```
echo "/export *(rw,async,no_root_squash)" > /etc/exports
```

You may need to adjust the above command to suite your deployment needs.

- **Limiting NFS export**. You can limit the NFS export to a particular subnet by changing * to a more restrictive subnet mask (e.g.,"192.168.1.0/24"). **The limit you place must include the private network(s) and the storage network(s).** If the two are

© 2010 Cloud.com, Inc. All rights reserved.

the same network then one CIDR is sufficient. If you have a separate storage network you must provide separate CIDR's for both or one CIDR that is broad enough to span both.

The following is an example with separate CIDR's:

```
/export 192.168.1.0/24(rw,async,no_root_squash)
10.50.1.0/24(rw,async,no_root_squash)
```

- **Removing the async flag**. The async flag improves performance by allowing the NFS server to respond before writes are committed to the disk. Remove the async flag in your mission critical production deployment.

6. Run the following command to enable NFS service.

```
chkconfig nfs on
```

7. Edit the /etc/sysconfig/nfs file and uncomment the following lines.

```
LOCKD_TCPPORT=32803
LOCKD_UDPPORT=32769
MOUNTD_PORT=892
RQUOTAD_PORT=875
STATD_PORT=662
STATD_OUTGOING_PORT=2020
```

8. Edit the /etc/sysconfig/iptables file and add the following lines at the beginning of the INPUT chain.

```
-A INPUT -m state --state NEW -p udp --dport 111 -j ACCEPT
-A INPUT -m state --state NEW -p tcp --dport 111 -j ACCEPT
-A INPUT -m state --state NEW -p tcp --dport 2049 -j ACCEPT
-A INPUT -m state --state NEW -p tcp --dport 32803 -j ACCEPT
-A INPUT -m state --state NEW -p udp --dport 32769 -j ACCEPT
-A INPUT -m state --state NEW -p tcp --dport 892 -j ACCEPT
-A INPUT -m state --state NEW -p udp --dport 892 -j ACCEPT
-A INPUT -m state --state NEW -p tcp --dport 875 -j ACCEPT
-A INPUT -m state --state NEW -p udp --dport 875 -j ACCEPT
-A INPUT -m state --state NEW -p tcp --dport 662 -j ACCEPT
-A INPUT -m state --state NEW -p udp --dport 662 -j ACCEPT
```

9. Reboot the server.

An NFS share called /export is now setup.

5.3.2 Linux NFS on iSCSI

Use the following steps to set up a Linux NFS server export on an iSCSI volume. These steps apply to RHEL/CentOS 5 distributions.

1. Install iscsiadm.

```
yum install iscsi-initiator-utils
service iscsi start
chkconfig --add iscsi
chkconfig iscsi on
```

2. Discover the iSCSI target.

© 2010 Cloud.com, Inc. All rights reserved.

```
iscsiadm –m discovery –t st –p <iSCSI Server IP address>:3260
```

For example:

```
# iscsiadm -m discovery -t st -p 172.23.10.240:3260
172.23.10.240:3260,1 iqn.2001-05.com.equallogic:0-8a0906-83bcb3401-
16e0002fd0a46f3d-rhel5-test
```

3. Log in.

```
iscsiadm –m node –T <Complete Target Name> –l –p <Group IP>:3260
```

For example:

```
# iscsiadm –m node –l –T iqn.2001-05.com.equallogic:83bcb3401-16e0002fd0a46f3d-
rhel5-test –p 172.23.10.240:3260
```

4. Discover the SCSI disk. For example:

```
# iscsiadm -m session –P3 | grep Attached
Attached scsi disk sdb State: running
```

5. Format the disk as ext3 and mount the volume.

```
mkfs.ext3 /dev/sdb
mkdir -p /export
mount /dev/sdb /export
```

6. Add the disk to /etc/fstab to make sure it gets mounted on boot.

```
/dev/sdb /export ext3 _netdev 0 0
```

Now you can set up /export as an NFS share.

© 2010 Cloud.com, Inc. All rights reserved.

6 Citrix XenServer Installation and Configuration

Citrix XenServer must be installed on the Computing Nodes. Citrix XenServer can be downloaded from the Citrix Website (http://www.citrix.com/lang/English/lp/lp_1688615.asp) and installed using the Citrix XenServer Installation Guide.

Important: All computing nodes must be 64-bit and must support HVM (Intel-VT or AMD-V enabled). All Computing Nodes within a Cluster must be homogenous. That means the CPUs must be of the same type, count, and feature flags. See http://docs.vmd.citrix.com/XenServer/4.0.1/reference/ch02.html for more information on homogenous XenServer hosts.

Important: You must re-install Citrix XenServer if you are going to re-use a host from a previous install. The only exception to this is when following the upgrade procedure from 2.0 GA to 2.1 GA.

Important: The CloudStack requires XenServer 5.6. XenServer 5.5 update 2 is not supported. Xenserver 5.6 FP1 beta is not supported.

Following installation, CloudStack requires the following configuration.

- Username and password
- Time synchronization
- Licensing
- Network Setup
    - Configuring public network with a dedicated NIC (optional)
    - NIC bonding (optional)
- iSCSI Multipath Setup (optional)

The following sections contain information about configuring the XenServer.

6.1 Username and Password

All XenServers in a Pod must have the same username and password as configured in the CloudStack.

6.2 Time Synchronization

The XenServer host must be set to use NTP or another clock synchronization. All Computing Nodes in a Pod must have the same time. You can use the NTP servers provided by Citrix:

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

6.3 Licensing

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

6.3.1 Getting and Deploying a License

© 2010 Cloud.com, Inc. All rights reserved.

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

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

You can install the license with XenCenter or you can contact Cloud.com support for a script that will apply the license to all known Computing Nodes.

6.4 Patch for 2.1.2 and Previous Releases

Special action is required on the XenServer node. On a new XenServer node, the following line should be added to the end of /etc/sysconfig/network, prior to moving forward with the network configuration:

NOZEROCONF=yes

6.5 Networking Setup

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

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

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

The Management Server uses arping to determine which NIC carries the management traffic, and if any other NICs should be bonded with that NIC. You can control this process and provide input to the Management Server through the use of XenServer name labels. The name labels are placed on physical interfaces or bonds and configured in the CloudStack to provide the CloudStack with their purpose. In some simple cases the name labels are not required.

Important: If either of the following is your desired NIC configuration there is no need for further configuration. Continue to the next section, Management Server Installation.

- **A single NIC for public, private, and guest traffic. In this case you must use a tagged VLAN for the public network.**
- **Exactly two NICs to be bonded together (by CloudStack) to carry public, private, and guest traffic. In this case, you should not create the bond. The CloudStack will do so when the host is added to the CloudStack. You must also use a tagged VLAN for the public network.**

6.5.1 Configuring Public Network with a Dedicated NIC (optional)

The CloudStack supports the use of a second NIC (or bonded pair of NICs, described later) for the public network. If bonding is not used the public network can be on any NIC and can be on different NICs on the Computing Nodes in a cluster. For example, the public network can be on eth0 on node A and eth1 on node B. However, the XenServer name-label for the public network must be identical across all computing nodes. The following examples set the network label to “cloud-public”. After the management server is installed and running you must configure it with the name of the chosen network label (E.g. “cloud-public”); this is discussed in Management Server Installation.

© 2010 Cloud.com, Inc. All rights reserved.

When a dedicated NIC is present for the public network, the public network can be implemented using a tagged or untagged VLAN.

If you are using two NICs bonded together to create a public network, see NIC Bonding (optional). If you are using a single dedicated NIC to provide public network access, follow this procedure.

1. Run `xe network-list` and find the public network. This is usually attached to the NIC that is public. Once you find the network make note of its UUID. Call this <UUID-Public>.
2. Run the following command.

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

This procedure should be repeated for each new host that is added to the CloudStack before adding the host.

6.5.2 NIC Bonding (optional)

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

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

All NIC bonding is optional.

XenServer automatically propagates NIC bonds from the master host in a cluster to slave hosts. In an installation the master will be the first host that was added to the cluster and the slave hosts will be all subsequent hosts added to the cluster. All bonds are propagated like this, whether they are for private, public, or storage access. There are several important implications of this:

- You must set bonds on the first host added to a cluster, and not on slaves that are added later. This means that the procedure to add a master host and a slave host are different. You must be aware of this or else your slave hosts will not have the bond established correctly.
- Slave hosts in a cluster must be cabled exactly the same as the master. For example, if eth0 is in the private bond on the master, it must be in the private network for added slave hosts.

6.5.2.1 Private Network Bonding

The CloudStack will automatically create a bond for the private network if it finds two NICs that are attached to the private network. It is assumed that only one NIC will be up, but both NICs are cabled.

Important: For the private network, the CloudStack creates the bond automatically and the administrator must not configure a bond at this time. Only one of the two NICs cabled for the private network should have an IP address.

6.5.2.2 Public Network Bonding

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

© 2010 Cloud.com, Inc. All rights reserved.

Use the following steps to create a bond in XenServer. These steps should be run on only the first host in a cluster.

1. Run the following command.

```
# xe pif-list
```

   This command shows all NICs and their UUIDs. Identify the NICs that you want to use for the public network and determine their UUID's. Call the UUID's of the NICs to be bonded slave1-UUID and slave2-UUID.

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

   Important: This label is significant as the CloudStack looks for a public network by a name you configure. You must use the same name-label for all Computing Nodes in the cloud for the public network.

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

3. The bond-list command will show one or more bonds with output like the following

```
uuid ( RO)      : 8f20d36b-db0d-6aaa-c696-1306c8b0bc14
    master ( RO): dcf22653-4e75-8da8-0a7c-f8f598495207
    slaves ( RO): e3e8997e-2483-d236-335f-ae926776d452; 7bb08b0b-5145-1635-c3a0-
7aa982cd6307
```

   Find the bond that has the slaves with UUIDs that you used in step 3. Take note of the master UUID, highlighted above. Be sure to use the UUID in the master row and not the UUID of the bond itself. Call this the master-uuid. Then

```
# xe pif-plug uuid=[master-uuid]
```

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

Slave hosts that are later added to the same cluster should not have this procedure done. They should be cabled and added to the CloudStack. XenServer will automatically propagate the bonds from the master to the slave.

6.5.3 Separate Storage Network (optional)

You can optionally set up a separate storage network. This can be done using one or two available NICs. With two NICs bonding may be done as above. It is the administrator's responsibility to set up a separate storage network.

Give the storage network a different name-label than what was given for the public network.

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

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

If bonding is done, XenServer will automatically propagate the bond from the master in a cluster to the slave. In this case the bond should be established once on the master and not on the slaves. However, the slaves must be cabled exactly like the master as discussed previously.

© 2010 Cloud.com, Inc. All rights reserved.

If no bonding is done, XenServer will not propagate the network configuration. The administrator must setup and name-label the separate storage network on all hosts (masters and slaves).

6.6 iSCSI Multipath Setup (optional)

XenServer must have multipathing enabled on the host object to use iSCSI multipath. The CloudStack can enable this for you. We recommend that you do not set this parameter on XenServer and instead manage it through the CloudStack. The CloudStack configuration to perform this is discussed in the Management Server section.

Multipath may be enabled across the cloud but disabled on select Computing Nodes. Implement this configuration by using the following steps.

1. Use the XenServer CLI to set multipathing to false on the host object for the specific Compute Nodes where you do not want multipath.
2. Enable multipath in the CloudStack; this will enable it for all other Compute Nodes when they are added to the cloud.

Important: Enabling multipath is only the first step of implementing multipath. The iSCSI target also has to tell the initiator that the LUN is exposed on two separate IP addresses.

Repeat the above Storage and Networking setup steps for each Citrix XenServer Computing Node. Be sure to use the same name-label for the cloud-public across all the Computing Nodes. Multipath configuration is not propagated from master to slave by XenServer.

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

7 VMware vSphere Installation and Configuration

VMware vSphere must be installed on the Computing Nodes. VMware vSphere can be downloaded and purchased from the VMware Website (https://www.vmware.com/tryvmware/index.php?p=vmware-vsphere&lp=1) and installed by following the VMware vSphere Installation Guide.

Important: You must re-install VMware ESXi if you are going to re-use a host from a previous install.

Important: The CloudStack requires VMware vSphere 4.1. VMware vSphere 4.0 is not supported.

Important: All computing nodes must be 64-bit and must support HVM (Intel-VT or AMD-V enabled). All Computing Nodes within a Cluster must be homogenous. That means the CPUs must be of the same type, count, and feature flags

Following installation, CloudStack requires the following configuration.

- ESX/ESXi Host Setup
- Configure Host Networking
    - Configure Virtual Switch
    - Configure vCenter Management Network
    - Configure NIC Bonding (optional)
- Configure Multipath Storage (optional)
- Configure clusters in vCenter and add hosts to them, or add hosts without clusters to vCenter

The following sections contain information about configuring the VMware vSphere.

7.1 ESX/ESXi Host setup

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

7.2 Configure Host Networking

Before an ESX host can be safely added to the CloudStack, proper network configuration is required. To configure ESX/ESXi host, you can use vClient to add it as standalone host to vCenter first. Once you see the host appearing in the vCenter inventory tree, click the host node in the inventory tree, navigate to configuration tab as following:

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

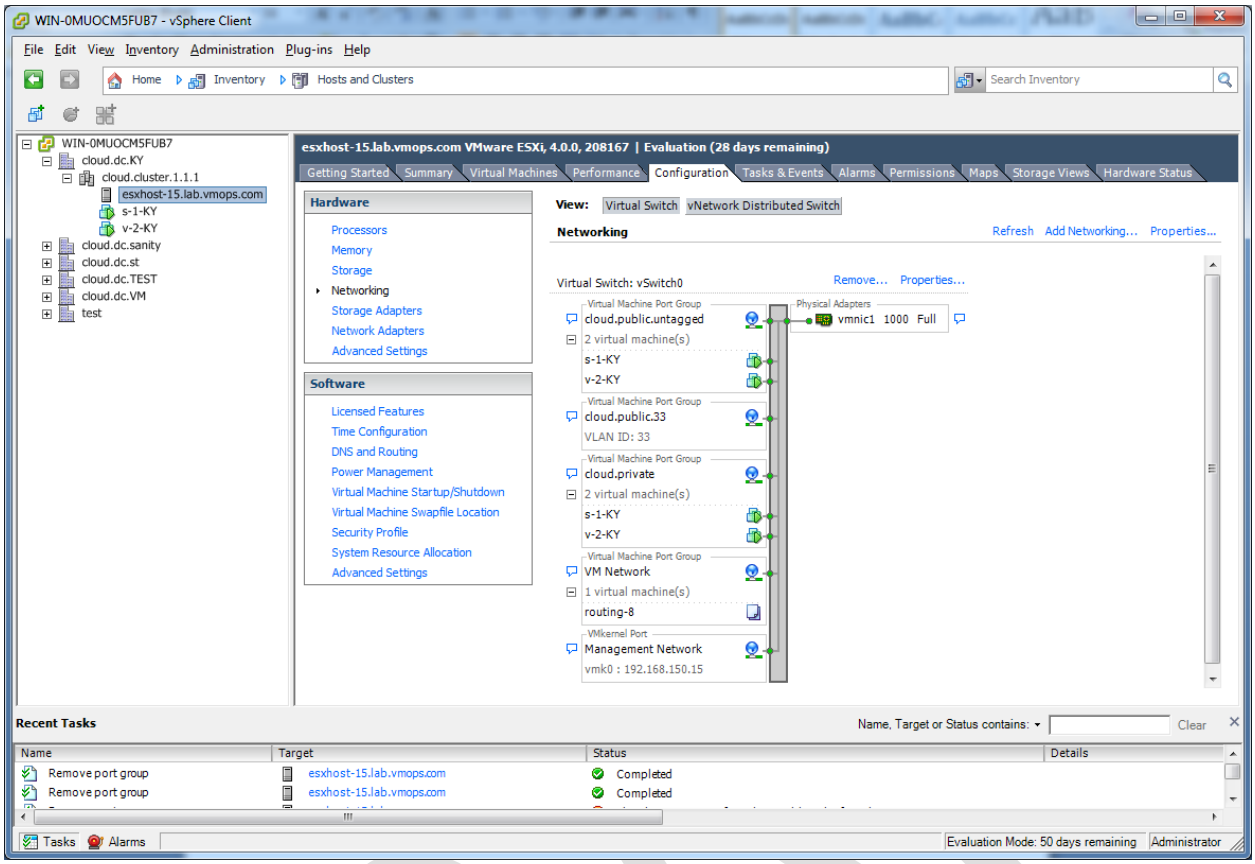

In host configuration tab, click “Hardware/Networking” link to bring up networking configuration page as above.

7.2.1 Configure Virtual Switch

A default virtual switch vSwitch0 is created. CloudStack requires all ESX/ESXi hosts in the cloud to use the same set of virtual switch names. If you change the default virtual switch name, you will need to configure one or more CloudStack configuration variables as well. See section 11.

7.2.1.1 Separating Traffic

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

If you want to separate traffic in this way you should first create and configure vSwitches in vCenter according to the vCenter instructions. Take note of the vSwitch names you have used for each traffic type. You will configure the CloudStack to use these vSwitches in section 11.

© 2010 Cloud.com, Inc. All rights reserved.

7.2.1.2 Increasing Ports

By default a virtual switch on ESX/ESXi hosts is created with 56 ports. We recommend to change it to 4096, the maximum number of ports allowed. To do that, click the "Properties…" link for virtual switch (note this is not the Properties link for Networking.)

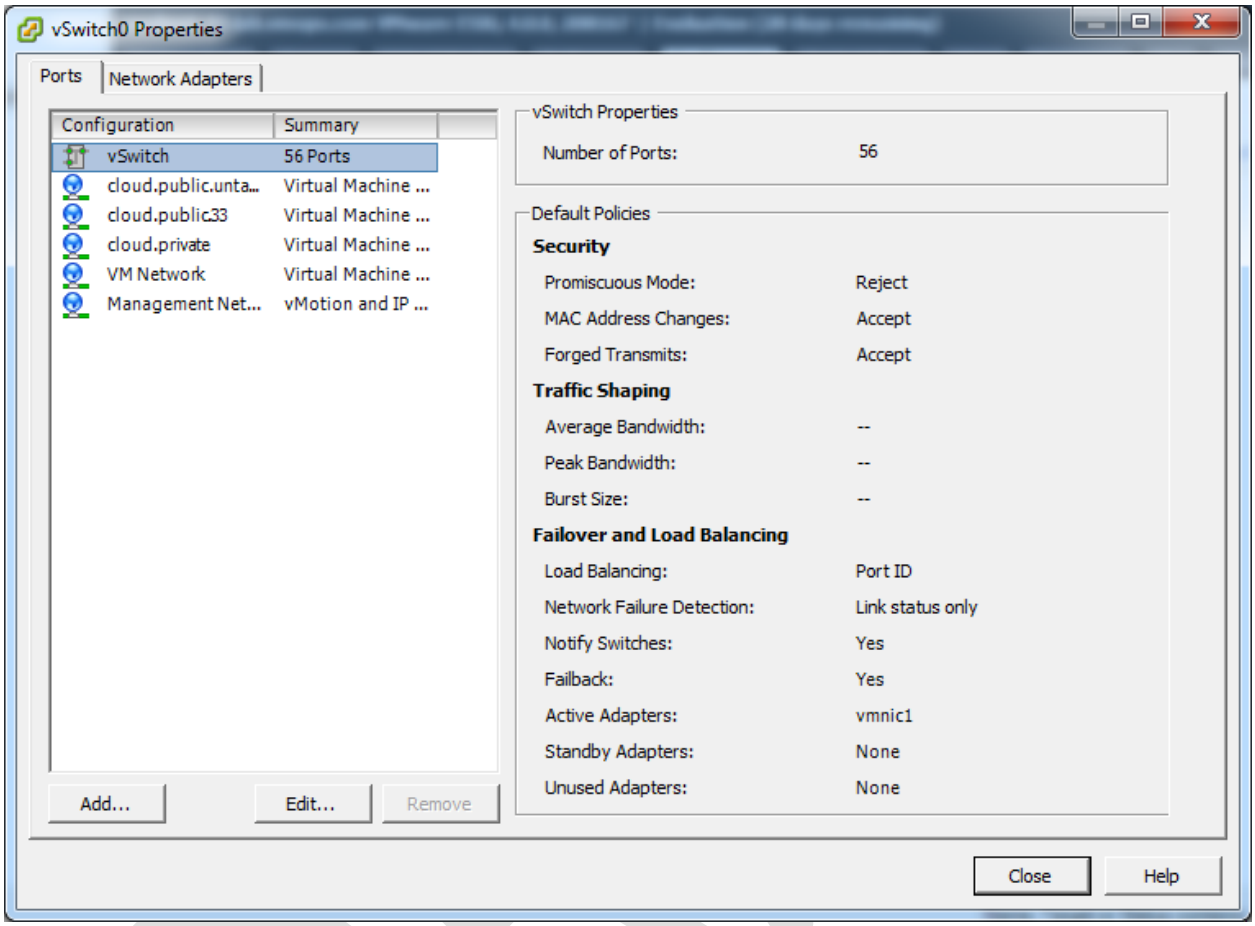

In vSwitch properties dialog, select the vSwitch and click Edit button, you should see the following dialog:

© 2010 Cloud.com, Inc. All rights reserved.

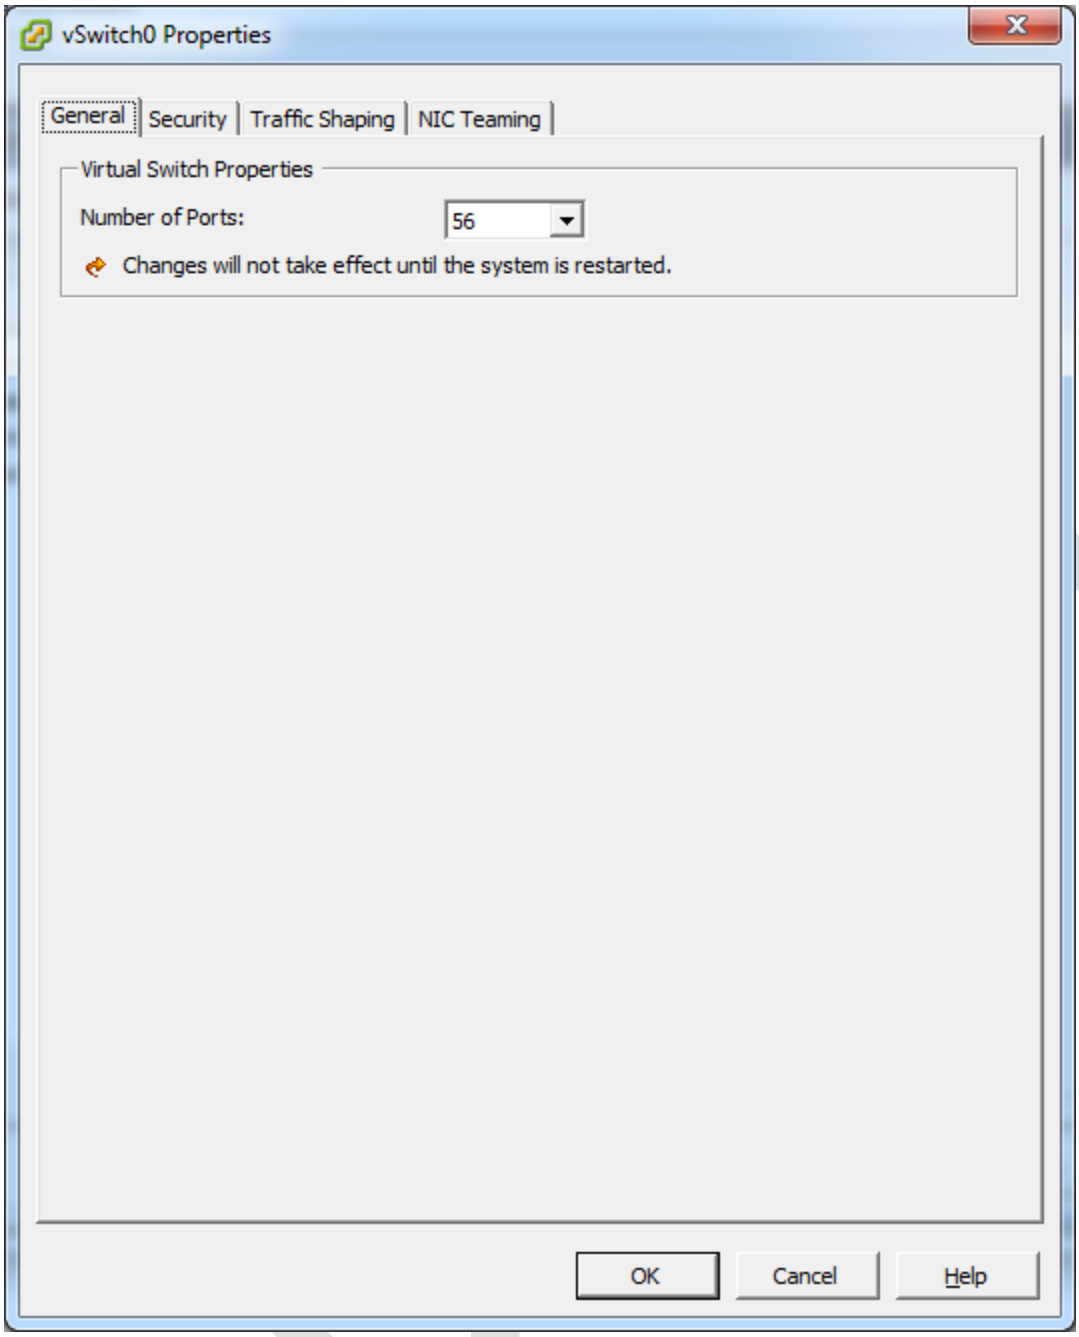

You can change the number of switch ports in above dialog, after you've done that, ESX/ESXi hosts are required to reboot in order for the setting to take effect.

© 2010 Cloud.com, Inc. All rights reserved.

7.2.2 Configure vCenter Management Network

In the vSwitch properties dialog box, you may see a vCenter management network. CloudStack also requires vMotion to be configured properly, select the management network item in the dialog, click edit button

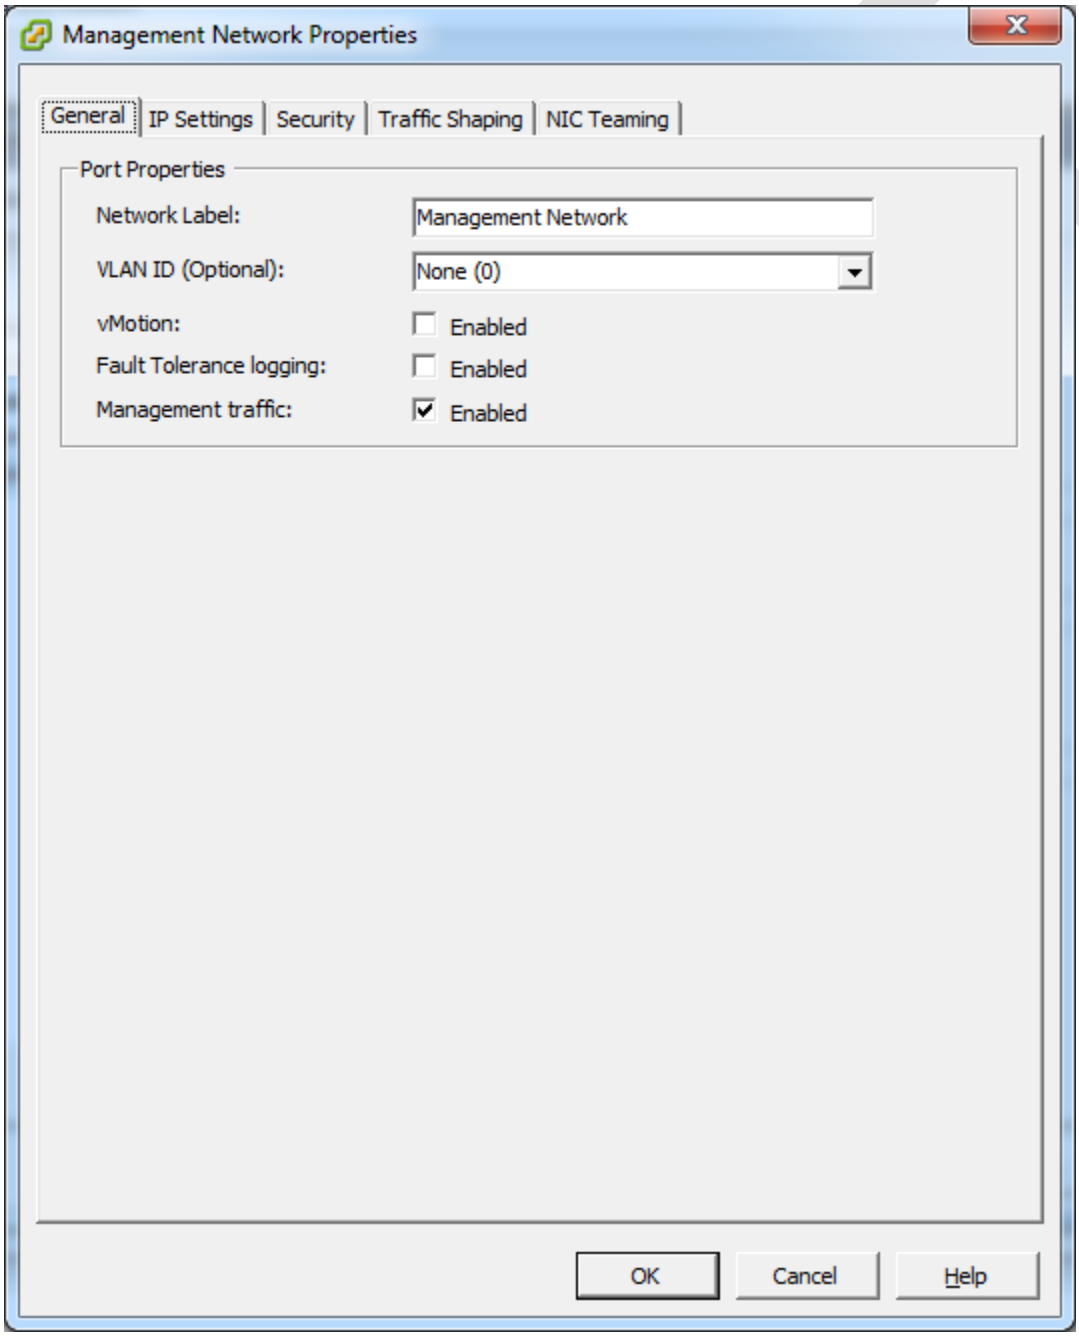

© 2010 Cloud.com, Inc. All rights reserved.

Make sure vMotion and Management traffic options are enabled.

7.2.3 Configure NIC Bonding

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

7.3 Configure Storage Multipathing

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

7.4 Add Hosts or Configure Clusters

Host management for vSphere is done through a combination of vCenter and the CloudStack admin UI. As an administrator you must decide if you would like to use clusters or standalone hosts. Clusters allow for features like live migration. Clusters also required shared storage such as NFS or iSCSI. The CloudStack allows some hosts to be in a cluster and other hosts to be standalone nodes.

7.4.1 Clusters

Use vCenter to create a vCenter cluster and add your desired hosts into the cluster. You will later add the entire cluster into the CloudStack with a single action.

7.4.2 Standalone Nodes

Use vCenter to add standalone nodes into vCenter management. You will later add these nodes into the CloudStack one at a time.

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

8 KVM Installation and Configuration

Important: All computing nodes must be 64-bit and must support HVM (Intel-VT or AMD-V enabled). All Computing Nodes within a Cluster must be homogenous. That means the CPUs must be of the same type, count, and feature flags.

Important: All Computing Nodes in the cloud must be of the same kernel version. For example, all Compute Nodes may be CentOS 5.5 64 bit.

8.1 Network Configuration

By default the CloudStack will use the device that is used for the default route. This device will be placed in a CloudStack-created bridge.

If a system has multiple NICs or bonding is desired the admin may configure the networking on the host. The admin must create a bridge and place the desired device into the bridge. This may be done for each of the public network and the private network. Then the customer should edit /etc/cloud/agent/agent.properties and add values for

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

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

- public.network.device=publicbondbr0

This should be done after the yum install of the software below (Step 2).

8.2 Installing the CloudStack Agent on a Computing Node

Each KVM Computing Node must have the CloudStack Agent installed on it. Install the CloudStack Agent on each Computing Node using the following steps. All commands should be run as root. On Fedora you have to make sure that the hostname resolves in DNS or /etc/hosts.

A Computing Node's host OS must also have a fully qualified domain name. Usually you can correct the lack of a FQDN by editing /etc/hosts. You will also want to make sure that the IP address associated with the hostname is the IP address of the default route table entry (usually eth0). If the hostname resolves to 127.0.0.1 guest networking will not work.

1. Disable SELinux .

   Set SELinux up to be permissive by default. First, set the SELINUX variable in /etc/selinux/config to "permissive". This ensures that MySQL and the Management Server can run properly on system reboot. Then set SELinux to permissive until the system is rebooted:

   ```
   setenforce permissive
   ```

2. Install the CloudStack packages. You should have a file in the form of "CloudStack-NNNN.tar.gz". Untar the file and then run the install.sh script inside it:

   ```
   # tar xzf CloudStack-2.2.0-1-centos.tar.gz
   # cd CloudStack-2.2.0-1-centos
   # ./install.sh
   Setting up the temporary repository...
   ```

© 2010 Cloud.com, Inc. All rights reserved.

```
Cleaning Yum cache...
Loaded plugins: fastestmirror
11 metadata files removed
Welcome to the Cloud.com CloudStack Installer.  What would you like to do?

    M) Install the Management Server
    A) Install the Agent
    S) Install the Usage Monitor
    D) Install the database server
    Q) Quit

    > A
```

Choose “A” to install the Agent software.

The CloudStack Agent is now installed. Later in the installation you will add this host to the CloudStack via the Management Server. This step will configure the Agent on the Computing Node.

Important: A Computing Node must have a statically allocated IP address; host addition will error and fail if a dynamically-assigned address is present.

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

9 Management Server Installation

The Cloud.com Management Server download includes everything you need to get started, except MySQL. This includes the Cloud.com software as well as dependencies. This section describes installing one or more Management Servers with one instance of MySQL, which may be on a different node from the Management Servers. The procedure for the installation is:

1. Prepare the operating system for all Management Servers.
2. Install the first Management Server.
3. Install MySQL.
4. (optional) Install additional Management Servers to create a farm for high availability.

To simplify the installation procedure this document defines two separate installation procedures: one for installing a single Management Server and one for installing multiple Management Servers in a load balanced pool. This document assumes that, in the case of multiple Management Servers, MySQL will be installed on a separate node from the Management Servers.

Important: Upgrade from 2.0 GA to 2.1 GA is supported. Upgrade instructions will be provided in the release notes when the upgrade procedure is complete. No other upgrade scenarios are currently supported.

9.1 Operating System and OS Preparation

The Cloud.com Management Server requires RHEL/CentOS 5.4 64 bit or later. You can download CentOS 64-bit via the following link: http://isoredirect.centos.org/centos/5/isos/x86_64/. The OS must be prepared to host the Management Server using the following steps.

Important: These steps should be done on all Management Servers.

Important: NTP is recommended.

1. Edit the /etc/hosts file to make sure that every Management Server has a fully-qualified host name that resolves to an IP address. Alternatively, you can do this through DNS.
2. Log in to your OS as root. All the following commands should be run as root.
3. Ensure that the SELINUX variable in /etc/selinux/config is set to permissive. This ensures that MySQL and the Management Server can run properly on system reboot.
4. Run the following command.

```
# setenforce permissive
```

5. Make sure that the Management Server can reach the Internet.

```
# ping www.google.com
```

9.2 Single Node Install (One Management Server)

This section describes the procedure for performing a single node install where the Management Server and MySQL are on a single, shared OS instance. If you have multiple Management Servers or if you want to have MySQL on a separate server, see Multinode Install (Multiple Management Servers).

3. Install the CloudStack packages. You should have a file in the form of “CloudStack-NNNN.tar.gz”. Untar the file and then run the install.sh script inside it:

© 2010 Cloud.com, Inc. All rights reserved.

```
# tar xzf CloudStack-2.2.0-1-centos.tar.gz
# cd CloudStack-2.2.0-1-centos
# ./install.sh
Setting up the temporary repository...
Cleaning Yum cache...
Loaded plugins: fastestmirror
11 metadata files removed
Welcome to the Cloud.com CloudStack Installer.  What would you like to do?

    M) Install the Management Server
    A) Install the Agent
    S) Install the Usage Monitor
    D) Install the database server
    Q) Quit

    > M
```

6. Choose "M" to install the Management Server software.

9.2.1 Single Node Database Install

1. Re-run install.sh and choose "D" to install MySQL.

```
# ./install.sh
Setting up the temporary repository...
Cleaning Yum cache...
Loaded plugins: fastestmirror
11 metadata files removed
Welcome to the Cloud.com CloudStack Installer.  What would you like to do?

    A) Install the Agent
    S) Install the Usage Monitor
    D) Install the database server
    U) Upgrade the CloudStack packages installed on this computer
    R) Stop any running CloudStack services and remove the CloudStack packages from
this computer
    Q) Quit

    > D
```

2. Edit the MySQL configuration (/etc/my.cnf) and insert the following lines in the [mysqld] section. You can put these lines below the datadir line. The max_connections parameter should be set to 350 multiplied by the number of Management Servers you are deploying. This example assumes 2 Management Servers.

```
innodb_rollback_on_timeout=1
innodb_lock_wait_timeout=600
max_connections=700
```

3. Restart the MySQL service:

```
# service mysqld restart
```

4. Use the following script to create the cloud user on the database with a password of your choice. Parameters include:

© 2010 Cloud.com, Inc. All rights reserved.

a. **dbpassword**. The password that will be assigned to the cloud user. You can choose to provide no password.

This script deploys the database using the credentials in the deploy-as parameters. By default the MySQL install does not set a password. This is represented in the example below. If you have set a password you should provide it with the root argument, as in "`--deploy-as=root:password`".

```
# cloud-setup-databases cloud:<dbpassword>@localhost --deploy-as=root
```

5. Configure the OS for the Management Server using the following command. This command will set up iptables, sudoers, and start the Management Server.

```
# cloud-setup-management
```

This completes the single node install for the Management Server and database. Continue with Prepare Secondary Storage.

9.3 Multinode Install (Multiple Management Servers)

This section describes installing multiple Management Servers and installing MySQL on a node separate from the Management Servers. If you have just completed the single node install, see Prepare Secondary Storage.

The procedure to install multiple management servers is:

1. Install the first Management Server
2. Install MySQL on a separate node (referred to as the Database Node, below)
3. Set up the MySQL database
4. Install additional Management Servers

9.3.1 Install the First Management Server

1. Install the CloudStack packages. You should have a file in the form of "CloudStack-NNNN.tar.gz". Untar the file and then run the install.sh script inside it:

```
# tar xzf CloudStack-2.2.0-1-centos.tar.gz
# cd CloudStack-2.2.0-1-centos
# ./install.sh
```

2. Choose "M" to install the Management Server software.

9.3.2 Install the Database

1. Log in as root to your Database Node and run the following commands. If you are going to install a replica database then log in to the master.

```
# yum install mysql-server
# chkconfig --level 35 mysqld on
```

2. Edit the MySQL configuration (/etc/my.cnf) and insert the following lines in the [mysqld] section. You can put these lines below the datadir line.

```
innodb_rollback_on_timeout=1
innodb_lock_wait_timeout=600
```

3. Start the MySQL service.

```
# service mysqld start
```

© 2010 Cloud.com, Inc. All rights reserved.

4. Run the following commands to grant access privileges to remote users.

```
# mysql -u root
```

5. Run the following command from the mysql prompt:

```
mysql> GRANT ALL PRIVILEGES ON *.* TO 'root'@'%' WITH GRANT OPTION;
```

6. Restart the MySQL service.

```
# service mysqld restart
```

7. Open the MySQL server port (3306) in the firewall to allow remote clients to connect.

```
# iptables -I INPUT -p tcp --dport 3306 -j ACCEPT
```

8. Edit the /etc/sysconfig/iptables file and add the following lines at the beginning of the INPUT chain.

```
-A INPUT -p tcp --dport 3306 -j ACCEPT
```

9.3.3 Database Replication (Optional)

The CloudStack supports database replication from one MySQL node to another. This is achieved using standard MySQL replication. You may want to do this as insurance against MySQL server or storage loss. MySQL replication is implemented using a master/slave model. The master is the node that the Management Servers are configured to use. The slave is a standby node that receives all write operations from the master and applies them to a local, redundant copy of the database. The following steps are a guide to implementing MySQL replication.

Important: These steps assume that this is a fresh install with no data in the master.

Important: Creating a replica is not a backup solution. You should develop a backup procedure for the MySQL data that is distinct from replication.

1. Edit my.cnf on the master and add the following in the [mysqld] section below datadir.

```
log_bin=mysql-bin
server_id=1
```

   For server_id a common practice is to set it to the last octet of the server's IP address. It must be unique with respect to other servers. Restart the MySQL service:

```
# service mysqld restart
```

2. Create a replication account on the master and give it privileges. We will use the "cloud-repl" user with the password "password". This assumes that master and slave run on the 172.16.1.0/24 network.

```
# mysql -u root
mysql> create user 'cloud-repl'@'172.16.1.%' identified by 'password';
mysql> grant replication slave on *.* TO 'cloud-repl'@'172.16.1.%';
mysql> flush privileges;
mysql> flush tables with read lock;
```

3. Leave the current MySQL session running.
4. In a new shell start a second MySQL session.

© 2010 Cloud.com, Inc. All rights reserved.

5. Retrieve the current position of the database.

```
# mysql -u root
mysql> show master status;
+------------------+----------+--------------+------------------+
| File             | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+------------------+----------+--------------+------------------+
| mysql-bin.000001 |      412 |              |                  |
+------------------+----------+--------------+------------------+
```

6. Note the file and the position that are returned by your instance.
7. Exit from this session.
8. Complete the master setup. Returning to your first session on the master, release the locks and exit MySQL.

```
mysql> unlock tables;
```

9. Install and configure the slave. On the slave server, run the following commands.

```
# yum install mysql-server
# chkconfig mysqld on
```

10. Edit my.cnf and add the following lines in the [mysqld] section below datadir.

```
server_id=2
innodb_rollback_on_timeout=1
innodb_lock_wait_timeout=600
```

11. Restart MySQL.

```
# service mysqld restart
```

12. Instruct the slave to connect to and replicate from the master. Replace the IP address, password, log file, and position with the values you have used in the previous steps.

```
mysql> change master to
    -> master_host='172.16.1.217',
    -> master_user='cloud-repl',
    -> master_password='password',
    -> master_log_file='mysql-bin.000001',
    -> master_log_pos=412;
```

13. Then start replication on the slave.

```
mysql> start slave;
```

14. Optionally, open port 3306 on the slave as was done on the master earlier.

    Important: This is not required for replication to work. But if you choose not to do this you will need to do it when failover to the replica occurs.

© 2010 Cloud.com, Inc. All rights reserved.

9.3.3.1 Failover

This will provide for a replicated database that can be used to implement manual failover for the Management Servers. CloudStack failover from one MySQL instance to another is performed by the administrator. In the event of a database failure you should:

1. Stop the Management Servers (via service cloud-management stop).
2. Change the replica's configuration to be a master and restart it.
3. Ensure that the replica's port 3306 is open to the Management Servers
4. Make a change so that the Management Server uses the new database. The simplest process here is to put the IP address of the new database server into each Management Server's /etc/cloud/management/db.properties.
5. Restart the Management Servers (via service cloud-management start)

9.3.4 Creating and Initializing the Database

Next, create and initialize the database.

1. Return to the root shell on your first Management Server.
2. Use the following script to create the cloud user on the database with a password of your choice. Parameters include:
   a. **dbhost**. The hostname of the database node.
   b. **dbpassword**. The password that will be assigned to the cloud user. You can choose to provide no password.

   This script deploys the database using the credentials in the deploy-as parameters. For example, if you originally installed MySQL with user "root" and password "password", provide `--deploy-as=root:password`.

```
# cloud-setup-databases cloud:<dbpassword>@<dbhost> --deploy-as=root:<rootpassword>
```

9.3.5 OS Configuration for the Management Server

Now run a script that will set up iptables rules and SELinux for use by the Management Server. It will also chkconfig off and start the Management Server.

```
# cloud-setup-management
```

9.3.6 Prepare and Start Additional Management Servers

For your second and subsequent Management Servers you will install the CloudStack, connect it to the database, and set up the OS for the Management Server.

1. Run these commands on each additional Management Server:

```
# tar xzf CloudStack-2.2.0-1-centos.tar.gz
# cd CloudStack-2.2.0-1-centos
# ./install.sh
```

2. Choose "M" to install the Management Server.
3. Configure the database client. Note the absence of the `--deploy-as` argument in this case.

```
# cloud-setup-databases cloud:<dbpassword>@<dbhost>
```

4. Configure the OS and start the Management Server:

© 2010 Cloud.com, Inc. All rights reserved.

```
# cloud-setup-management
```

The Management Server on this node should now be running.

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

10 Prepare Secondary Storage

Secondary storage in all Zones must be seeded with a template that is used for system VMs such as the Virtual Router. For each Secondary storage server you will need to execute the following steps.

1. Mount secondary storage on to your Management Server. This example assumes the path on the secondary storage server is /nfs/share.

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

2. Retrieve and decompress the system VM template. Run the script /usr/lib64/cloud/agent/scripts/storage/secondary/installrtng.sh, which is installed on the Management Server. Run this for each hypervisor type that you expect end users to run in this zone. This will give the system maximum flexibility to run system virtual machines. This process will require approximately 2 GB of free space on the local file system each time it runs.

   The command to run depends on hypervisor type.

   For ESX:

```
# /usr/lib64/cloud/agent/scripts/storage/secondary/installrtng.sh -m /mnt/secondary
-u http://download.cloud.com/releases/2.2.0/systemvm.ova -h vmware -F
```

   For KVM:

```
# /usr/lib64/cloud/agent/scripts/storage/secondary/installrtng.sh -m /mnt/secondary
-u http://download.cloud.com/releases/2.2.0/systemvm.qcow2.bz2 -h kvm -F
```

   For XenServer:

```
# /usr/lib64/cloud/agent/scripts/storage/secondary/installrtng.sh -m /mnt/secondary
-u http://download.cloud.com/releases/2.2.0/systemvm.vhd.bz2 -h xen -F
```

   Each one of these files is large. It may take 30 minutes or more to download and uncompress.

3. When the script has finished running, unmount secondary storage.

```
# umount /mnt/secondary
```

Repeat the above steps for each secondary storage server.

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

11 Describe Your Deployment

Now your Cloud.com Management Server is running. The next step is to tell it about the compute nodes, storage, and network configuration that you have done in the previous sections.

1. Log in to the administrator web UI.

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

   The default credentials are “admin” for user and “password” for password. The domain field should be left blank. A blank domain field is defaulted to the ROOT domain.

2. Using the left navigation tree, click on Configuration then Global Settings.

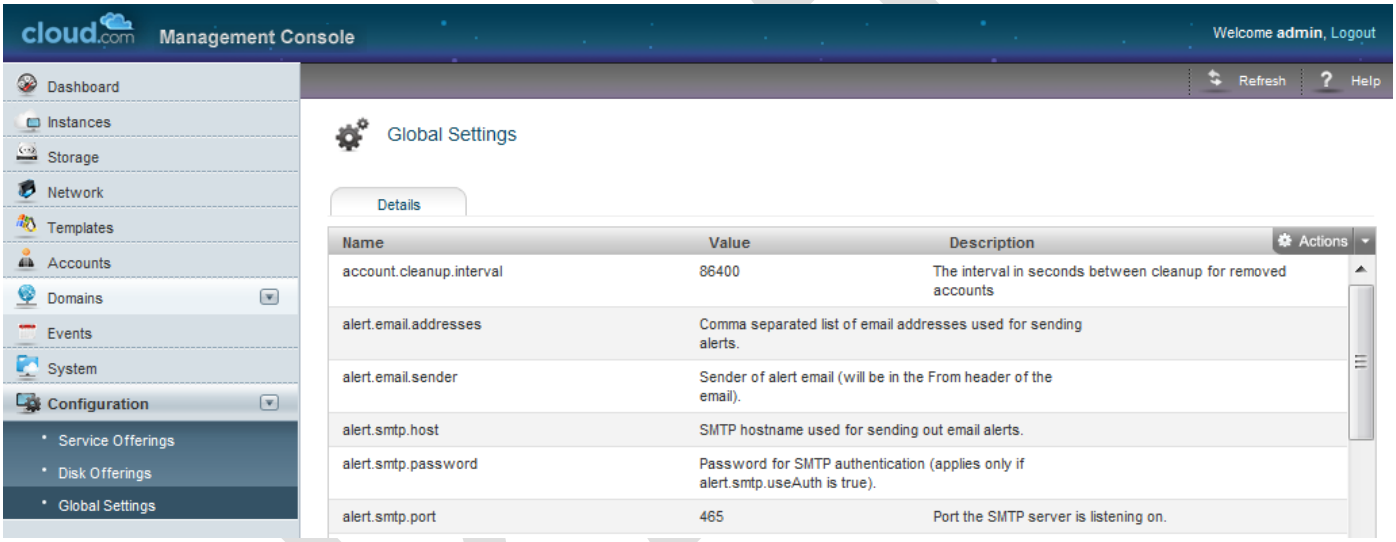

   You may need to edit the following fields.

Field	Value
xen.public.network.device	For XenServer nodes, this is the device name with the name-label that was used for the public network. For example, “cloud-public”. **This is applicable only if you have a second NIC for the public network.** In this case, set the value to the name-label used for the separate public network created in section 7.
xen.setup.multipath	For XenServer nodes, this is a true/false variable that instructs the CloudStack to enable iSCSI multipath on the XenServer Compute Nodes when they are added. This defaults to false. Set it to true if you would the 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.

© 2010 Cloud.com, Inc. All rights reserved.

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	This determines whether or not the CloudStack will use storage that is local to the Compute Node for VHDs. By default the 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. **Important: in 2.2 beta 1 this option is not applicable to vSphere.**
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.
hypervisor.type	Allowed values are "xen", "kvm", and "vmware". This variable controls which hypervisor type will be used to run the system VMs if there are multiple hypervisors in the cloud.
vmware.guest.vswitch	The name of the vSwitch in vCenter that you want to carry the guest-guest traffic. Default is vSwitch0.
vmware.private.vswitch	The name of the vSwitch in vCenter that you want to carry the management and storage traffic. Default is vSwitch0.
vmware.public.vswitch	The name of the vSwitch in vCenter that you want to carry the traffic to/from the public internet. Default is vSwitch0.

There are additional configuration parameters that you may want to set. These are discussed in the Administration Guide. For an initial installation they are not generally necessary.

3. If you changed any of these values you should restart the Management Server now.

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

11.1 Add a New Zone

For 2.2 beta, you should ignore the default Zone and the default Pod in it.

11.1.1 Adding a Zone

© 2010 Cloud.com, Inc. All rights reserved.

Begin by using these steps to edit the default Zone.

1. In the left navigation tree, choose System, then Physical Resources.
2. Click on the + next to the Zone to add a Zone. The Zone creation wizard will appear.
3. Choose your network type.
    a. Basic: this is for use with Direct Untagged networking.
    b. Advanced: this is for use with Direct Tagged or Virtual Networking.
4. Configure the Zone.

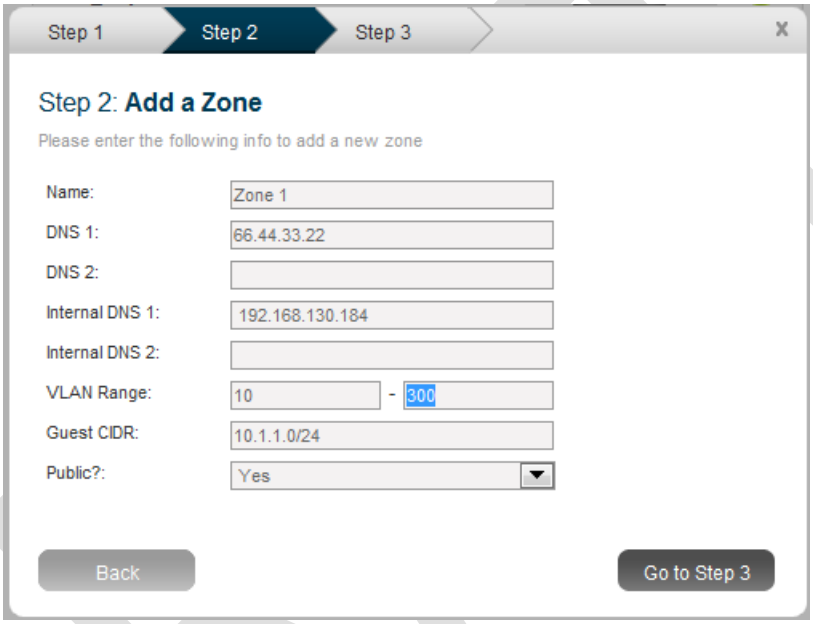

Important: The CloudStack distinguishes between internal and public DNS. Internal DNS is assumed to be capable of resolving internal-only hostnames, such as your NFS server's DNS name. Public DNS is provided to the guest VMs for DNS resolution. You can enter the same DNS server for both types, but if you do so you must make sure that both private and public IP addresses can route to the DNS server. Note that you must provide at least one public DNS server and at least one Internal DNS server.

Enter the following details in the Add Zone dialog.

- **Name**. The name of the Zone.
- **DNS 1 and 2**. These are DNS servers for use by guest VMs in the Zone. These DNS servers will be accessed via the public network you will add later. The public IP addresses for the Zone must have a route to the DNS server named here.
- **Internal DNS 1 and 2**. These are DNS servers for use by system VMs in the Zone. These DNS servers will be accessed via the private network interface of the System VMs. The private IP address you provide for the Pods must have a route to the DNS server named here.
- **Zone VLAN Range (Advanced Only)**. This is the range of Zone VLANs that are used for provided VLANs to the guest networks. It is entered in the format "x-y". (E.g. 800-900). You should set the Zone VLAN fields based on your planned VLAN allocation.

© 2010 Cloud.com, Inc. All rights reserved.

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

  This field should be modified only if you are using Virtual Networking. The CloudStack uses it when VMs with Virtual Networking are created.

- **Public**. Choose if this Zone is public or not. A public Zone is available to all users. A non-public Zone will be assigned to a particular domain. Only users in that domain will be allowed to create Guests in this Zone.

5. Add a Pod for the new Zone.

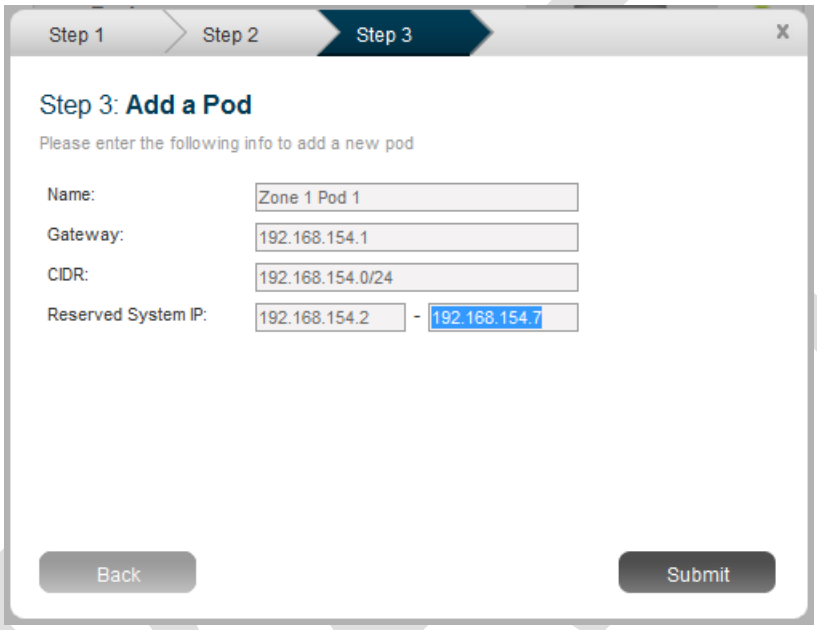

Enter the following details in the Add Pod dialog.

- **Name**. The name of the Pod.
- **Gateway**. The gateway for the hosts in that Pod.
- **CIDR**. The CIDR that encompasses the subnet in the Pod.
- **Reserved System IP**: This is the IP range in the private network that the CloudStack uses to manage Secondary Storage VMs and Console Proxy VMs. These IP addresses are taken from the same subnet as computing servers. You therefore need to make sure computing servers and Management Servers use IP addresses outside of this range. These two values combine to give the system control over a certain IP address range, and leave you in control of allocation for IP addresses in the CIDR but outside of the start and end range. In the screenshot we have start=192.168.154.2 and end=192.168.154.7. These computing servers and Management Servers can use IP addresses .8 to .254 and the CloudStack can use .2 to .7 for System VMs.

  The recommended number of private IPs per Pod is 5 times the number of Clusters. The example above allocates ten IPs as if two Clusters were in place for the Pod.

- **Guest IP Range** (Basic Mode Only). The range of IP addresses that will be available for allocation to guests in this Pod. If one NIC is used these IPs should be in the same CIDR as the Pod CIDR. If multiple NICs are used they may be in a different subnet.
- **Guest Netmask** (Basic Mode Only). The netmask in use on the subnet the guests will use.
- **Guest Gateway** (Basic Mode Only). The gateway that the guests should use.

© 2010 Cloud.com, Inc. All rights reserved.

11.1.2 Advanced Networking: Adding a Public IP Address Range on the Zone

A Public IP address range should be added if you will have guests using Virtual Networking. If you will use Direct Attached only (either tagged or untagged) you should skip this section.

Use the following steps to add a public IP address range.

1. Click on your Zone.
2. Click "Add VLAN". The Add VLAN IP Range dialog displays.

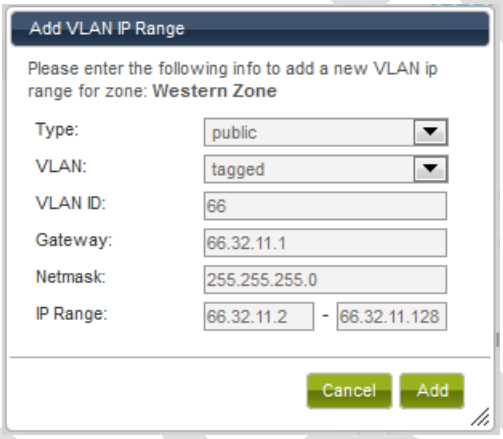

3. Enter the following details in the Add VLAN IP Range dialog.
    - **Type (public or direct).** Choose public.
    - **VLAN (tagged or untagged)**. Choose whether you will be using a tagged VLAN for public network traffic or an untagged VLA. You must choose tagged if you are using a single NIC for all traffic.
    - **VLAN ID**. The VLAN that will be used for public traffic if tagged VLAN was chosen**.**
    - **Gateway**. The gateway in use for these IP addresses.
    - **Netmask**. The netmask associated with this IP range.
    - **IP Range**. This takes a lower and an upper IP address that together form a range. These IP addresses are assumed to be accessible from the Internet and will be allocated for access to guest networks.

11.1.3 Additional Zones

You can add additional Zones as needed. If you choose to add additional Zones, be sure to repeat the installrtng.sh template seeding that you did for the first Zone.

11.1.4 Additional Pods

You can add additional Pods as needed now or after the system is up and running.

11.1.5 Additional Clusters

You can add additional Clusters as needed now or after the system is up and running.

© 2010 Cloud.com, Inc. All rights reserved.

11.2 Edit Service Offerings (Optional)

The service offering defines CPU and RAM for the guests. The CloudStack ships with several default service offerings. You can optionally use the following steps to edit these now or proceed with the defaults.

1. Go to the Configuration tab, and select the Service Offerings section.
2. Add or edit service offerings as needed. Select "Add Service Offering" to add one.

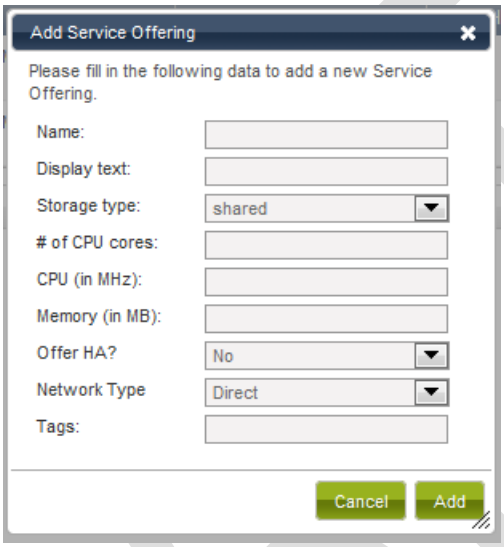

3. Provide the following information to define this service offering.
    - **Name**. The name of the service offering.
    - **Display text**. A short description of the offering.
    - **Storage type**. The type of disk that should be allocated to the guest. Local allocates from storage attached to XenServer directly. Shared allocates from storage accessible via NFS.
    - **# of CPU cores**. The number of cores which should be allocated to an instance with this offering.
    - **CPU (in MHz)**. The CPU speed of the cores that the instance is allocated. For example, “2000” would provide for a 2 GHz clock.
    - **Memory (in MB)**. The amount of memory in megabytes that the instance should be allocated. For example, “2048” would provide for a 2 GB RAM allocation.
    - **Offer HA**. If yes, the user will be able to choose a VM to be monitored and as highly available as possible.
    - **Network Type**. If Public this VM will be part of a Virtual Network. If Direct this VM will be Direct Attached; whether it is tagged or untagged depends on the configuration of the Zone and Pod.
    - **Tags**. The tags that should be associated with the primary storage for this root disk.

11.3 Edit Disk Offerings (Optional)

The disk offering defines the size and characteristics of data disks attached to the guests. The CloudStack ships with several default disk offerings.

© 2010 Cloud.com, Inc. All rights reserved.

1. Go to the Configuration tab, and select the Disk Offerings section. This is located next to the Zones section.
2. Add or edit disk offerings as needed. Select "Add Disk Offering" to add one.

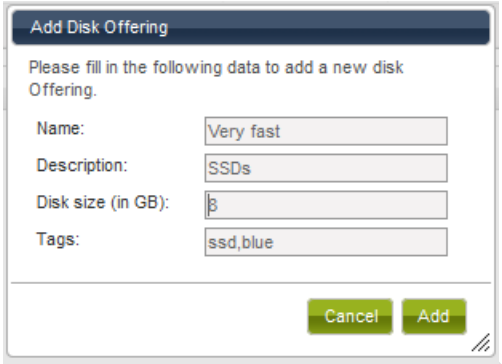

3. Provide the following information to define this disk offering.
    - **Name**. Name of the disk offering (E.g. extra large)
    - **Description**. A short description of the disk offering
    - **Disk size (in GB)**. The size of the disk offering in GB (E.g. 10 is a 10 GB offering)
    - **Tags**. Tags are a comma separated list of attributes of the storage. For example "ssd,blue". Tags are optional. They are also added on Primary Storage. The CloudStack matches tags on a disk offering to tags on the storage. If a tag is present on a disk offering that tag (or tags) must also be present on Primary Storage for the volume to be provisioned. If no such primary storage exists allocation from the disk offering will fail.

11.4 Add Hosts

Now that the offerings are defined, you need to tell the CloudStack about the hosts that it should manage.

The CloudStack supports multiple Clusters per Pod. A Cluster is a XenServer server pool, a set of KVM servers, or a VMware server cluster. Within a Cluster the Compute Nodes may live migrate VMs to and from each other, and all access the same shared storage. We expect that most deployments will have a single Cluster per Pod.

1. Go to the Hosts tab. An empty list of hosts displays.
2. Select Add Host. The Add New Computing Host dialog displays.

© 2010 Cloud.com, Inc. All rights reserved.

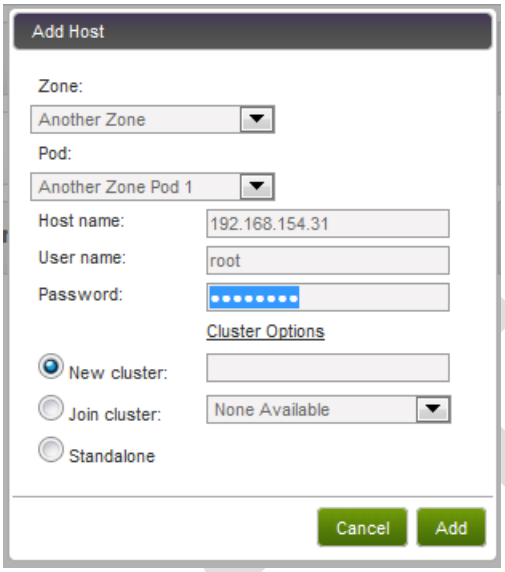

3. Provide the following information in the Add New Computing Host dialog.
   - **Availability Zone**. The Zone to which this host will be allocated.
   - **Pod**. The Pod to which this host will be allocated.
   - **Hostname**. For Xen and KVM this is the DNS name or IP address of the XenServer host or KVM host.

     For vSphere servers this can be either the cluster that you have configured in vCenter or a standalone host. vSphere hosts or clusters require a special syntax for the Hostname. This is a URL of the form vCenter hostname / data center name / { cluster name or hostname }.

     For example, to add all hosts in the cluster "Cluster 1" in data center "Zone 1" as configured in vCenter host vcenter.lab.example.org, you would enter "vcenter.lab.example.org/Zone 1/Cluster 1".

     To add a single standalone host that is not in a cluster in vCenter, extending the example above you would enter "vcenter.lab.example.org/Zone 1/192.168.154.31".

     This vCenter screenshot shows a path to identify a vSphere cluster.

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

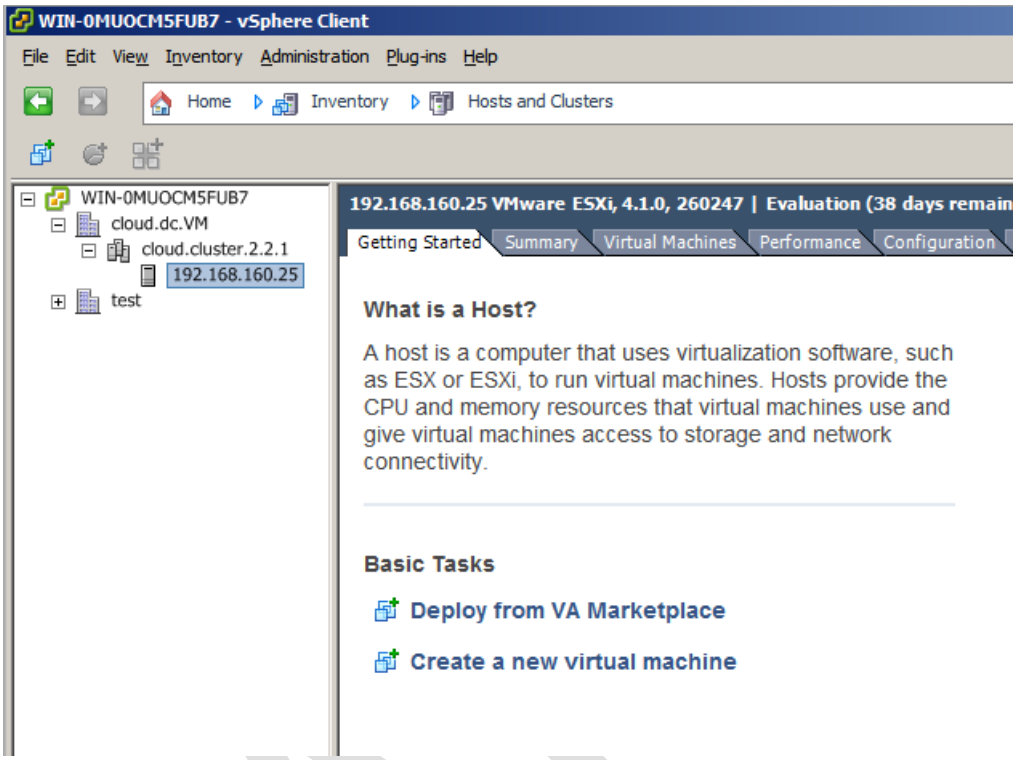

- **Username**. Usually the root user (XenServer or KVM) or Administrator (VMware vCenter).
- **Password**. This is the password for the user named above (from your Citrix Xen Server, KVM, or VMware vSphere install).
- **New Cluster / No Cluster / Existing Cluster**. This is the Cluster that you would like the added Host to join, if any. Choose New Cluster to create a new Cluster in the Pod. If this is your first Host to add to the Pod, and you plan on using shared Primary Storage such as NFS or iSCSI, choose New Cluster and provide a Cluster name. The Cluster name has no particular meaning for the CloudStack.

  If you are using Local Storage you should choose "No Cluster". This will ensure that no Clusters are formed. Clusters are not applicable to Local Storage since live migration is not possible.

  If you are adding a Host into a Pod with an existing Cluster, you can choose either New Cluster (to create a second or subsequent Cluster) or Existing Cluster. Choose Existing Cluster to have this host join the selected Cluster.

  Most customers will choose "New Cluster" for the first Host in the Pod, and "Existing Cluster" for all subsequent Host adds into that Pod.

It may take a minute for the host to be provisioned. It should automatically display in the UI.

Important: when a vCenter cluster is added the CloudStack will show only the cluster. It does not show the individual hosts that are present in the cluster.

11.5 Add Primary Storage

Next you will need to tell the system about the primary and secondary storage devices that are available. If you intend to use only local disk for your installation you can skip to Secondary Storage.

1. Navigate to the CloudStack cluster that you want to add primary storage to. Go to System -> Physical Resources -> Zone -> Pod -> Cluster. Then choose "Add Primary Storage".

© 2010 Cloud.com, Inc. All rights reserved.

Important: Primary storage cannot be added until a Computing Node has been added to the Pod.

Important: if you do not provision shared storage for primary storage, you will not be able to create additional volumes (via Storage > Volumes > Add Volume). Also, if you do not provision shared primary storage, you must have set system.vm.local.storage.required to true in the first part of this section or else you will not be able to start VMs.

2. Click on Add Primary Storage. The Add Primary Storage dialog displays.

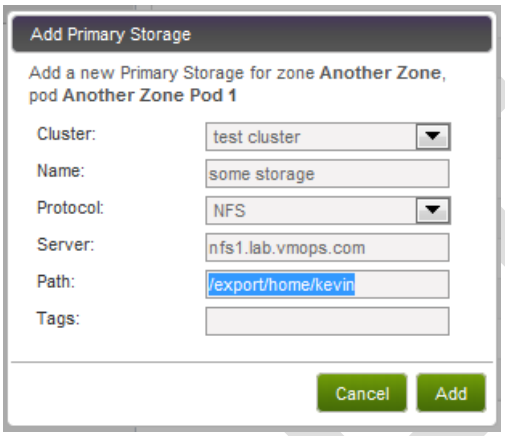

3. Provide the following information in the Add Primary Storage dialog.
   - **Cluster**. The Cluster for the storage device.
   - **Name**. The name of the storage device.
   - **Protocol**. Choose either NFS or iSCSI as determined by your storage device.
   - **Server**. The IP address or DNS name of the server.
   - **Path (NFS Only)**. In NFS this is the exported path from the server
   - **Target IQN (iSCSI Only)**. In iSCSI this is the IQN of the target. For example, iqn.1986-03.com.sun:02:01ec9bb549-1271378984
   - **Lun # (iSCSI Only)**. In iSCSI this is the LUN number. For example, 3.
   - **Tags**. 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. It is an optional field and may be left blank.

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

You can add multiple primary storage servers to a Cluster. At least one is required.

11.5.1 Secondary Storage

You will need to add secondary storage. Secondary Storage is used to store templates, ISOs, and snapshots.

Important: Secondary Storage is always accessed via NFS.

1. Navigate to the Zone you are building. Choose System -> Physical Resources and then click on your Zone.
2. Provide exactly one Secondary Storage device per Zone. Select "Add Secondary Storage".

© 2010 Cloud.com, Inc. All rights reserved.

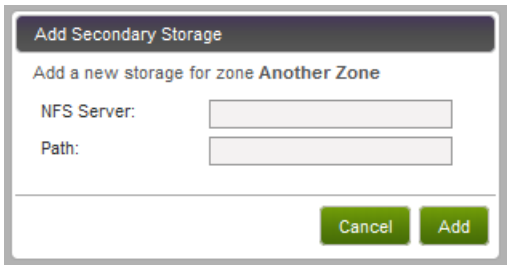

3. Provide the details for Secondary Storage server:
    - **Server**. The IP address of the server.
    - **Path**. The exported path from the server.

11.6 SSL

The CloudStack provides HTTP access in its default installation. There are a number of technologies and sites which choose to implement SSL. As a result we have left the CloudStack to expose HTTP under the assumption that a site will implement its typical practice.

The CloudStack uses Tomcat as its servlet container. For sites that would like the CloudStack to terminate the SSL session, Tomcat's SSL access may be enabled. Tomcat SSL configuration is described at http://tomcat.apache.org/tomcat-6.0-doc/ssl-howto.html.

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

12 Initialization and Testing

At this point you should have one Java process running the Cloud.com software on each Management Server. This is the Management Server process.

After everything is configured the CloudStack will perform its initialization. This can take 30 minutes or more depending on the speed of your network. During this initialization process several things happen:

- The CloudStack will start the Secondary Storage VM and Console Proxy VM from the system VM template downloaded into each Zone. In the Instances tab, System section you will see the status of these VMs listed first as Creating, then as Starting, then as Running. You can click on Refresh is the lower right to update the status.
- After the Secondary Storage VM is running the Management Server will initiate the download of the CentOS template. The Management Server requests that the Secondary Storage VM perform this download. You can go to the Templates tab to check the status of this download. The status will show "Storage agent or storage VM disconnected" until the Secondary Storage VM is running. Then the status will change to show that the download is in progress. You can click Refresh to update the download percentage.
- Once the CentOS template is downloaded it will be uncompressed by the Secondary Storage VM. This is a large file and this operation will take several minutes. The Management Server will then update the template's status to Ready.

If these steps do not work you should see the Troubleshooting section below, or contact support for assistance.

Important: If these steps do not work further testing will not be successful.

Once the CloudStack has performed initialization, use the following steps to try creating a new virtual machine.

1. Click on Launch Test Provisioning Tool to create a new end user account.

   This tool allows you to create end user accounts, domain administrators, and global administrators, as well new domains. We expect most organizations will remove this tool before going to production and replace it with their own provisioning system.
2. Exit from the Test Provisioning Tool.
3. Go to the Instances tab.
4. Click the button to create a new VM and follow the steps in the wizard.
5. To log in as an end user account, go to http://managementserver:8080/client. This URL displays the end user UI or the admin UI based on the access level of the authenticated account.

If you decide to grow your deployment, you can add more Computing Nodes, Primary Storage, Zones, Pods, and Clusters. Repeat the procedures above as needed.

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

13 Installing the Usage Server (Optional)

You can optionally install the Usage Server once the Management Server is configured properly. The Usage Server takes data from the events in the system and enables usage-based billing for accounts.

When multiple Management Servers are present the Usage Server may be installed on any number of them. The Usage Servers will coordinate usage processing. A site that is concerned about availability should install Usage Servers on at least two Management Servers.

Important: The Management Server must be running when Usage Server is installed. The Usage Server must be installed on the same server as a Management Server.

Use the following steps to install the Usage Server.

1. Run ./install.sh.

```
# ./install.sh
Setting up the temporary repository...
Cleaning Yum cache...
Loaded plugins: fastestmirror
11 metadata files removed
Welcome to the Cloud.com CloudStack Installer.  What would you like to do?

    A) Install the Agent
    S) Install the Usage Monitor
    U) Upgrade the CloudStack packages installed on this computer
    R) Stop any running CloudStack services and remove the CloudStack packages from
this computer
    E) Remove the MySQL server (will not remove the MySQL databases)
    Q) Quit

    > S
```

2. Choose “S” to install the Usage Server.
3. Once installed, start the Usage Server with the following command.

```
# service cloud-usage start
```

The Administration Guide discusses further configuration of the Usage Server.

© 2010 Cloud.com, Inc. All rights reserved.

14 Troubleshooting

Many install problems relate to the secondary storage VM. Sample common problems:

- SSVM cannot reach the DNS server
- SSVM cannot reach the Management Server
- SSVM cannot reach the outside world to download templates. It contacts download.cloud.com via HTTP.
- The configured DNS server cannot resolve your internal hostnames. E.g., you entered private-nfs.lab.example.org for secondary storage NFS, but gave a DNS server that your customers use, and that server cannot resolve private-nfs.lab.example.org.

Another common problem is that your VLAN's are not set up correctly.

A quick step to look for errors in the management server log is this:

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

14.1 Troubleshooting the Secondary Storage VM

You can troubleshoot the secondary storage VM either by running a diagnostic script or by checking the log file. The following sections detail each of these methods.

If you have corrected the problem but the template hasn't started to download, restart the cloud service with "service cloud restart". This will restart the default CentOS template download.

Important: To recover a failed SSVM after making changes that fix the root cause of the failure, you must stop the VM first and then start it. A restart merely reboots the VM without resending the configuration, which may have changed.

14.1.1 Running a Diagnostic Script

You can log into the SSVM. To do this you have to find the Compute Node running the SSVM, ssh into it, then ssh into the SSVM's private IP from that host. Once you are logged in, use the following steps to run a diagnostic script.

1. In the admin UI, go to Instances tab, System section.
2. Note the name of the Computing Node hosting the SSVM as shown in the Host column. Also note the private IP of the SSVM as shown in the Private IP column.
3. ssh into the Computing Node using your known user and password.
4. ssh into the private IP of the SSVM with the following.

   ```
   # ssh -i /opt/xensource/bin/id_rsa -p 3922 root@private-ip
   ```

5. Once into the SSVM, run the following diagnostic script:

   ```
   # /usr/local/cloud/systemvm/ssvm-check.sh
   ```

This script will test various aspects of the SSVM and report warnings and errors.

14.1.2 Checking the Log File

© 2010 Cloud.com, Inc. All rights reserved.

You can also check the log file /var/log/cloud/cloud.log for any error messages.

14.2 Troubleshooting the Console Proxy VM

```
Access is denied for console session. Please close the window
```

If you launch the Console Viewer and see this error it most likely means that the Console Proxy VM cannot connect from its private interface to port 8250 on the Management Server (or load balanced Management server pool). Check these things:

- Load balancer has port 8250 open
- All Management Servers have port 8250 open
- There is a network path from the CIDR in the Pod hosting the Console Proxy VM to the load balancer or Management Server
- The "host" global configuration parameter is set to the load balancer if in use

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.

15 Contacting Support

Cloud.com support is available to help you plan and execute your installation. The support team is available at support@cloud.com.

Important: while CloudStack 2.2 is in beta we will give priority to customer requests on the 2.1.x series of releases.

DRAFT

© 2010 Cloud.com, Inc. All rights reserved.