Eucalyptus 4.1.2 Installation Guide

2016-01-26 Eucalyptus Systems

Contents

Installation Overview.....6

Introduction to Eucalyptus.....7

Eucalyptus Overview.....7
Eucalyptus Components.....7
System Requirements.....8

Eucalyptus Installation.....10

Plan Your Installation.....10
Eucalyptus Architecture Overview.....10
Plan Component Placement.....10
Plan Your Hardware.....12
Verify Component Disk Space.....12
Plan Networking Modes.....12
Plan Eucalyptus Features.....16
Prepare the Network.....17
Configure Dependencies.....20
Configure Bridges.....20
Disable the Firewall.....21
Configure SELinux.....21
Configure NTP.....21
Configure an MTA.....22
Enable Packet Routing.....22
Install Eucalyptus.....22
Software Signing.....22
Install Eucalyptus from Release Packages.....23
Install Eucalyptus from Nightly Packages.....25
Configure Eucalyptus.....26
Configure Network Modes.....26
Create Scheduling Policy.....31
Configure Loop Devices.....31
Configure Multi-Cluster Networking.....32
Configure the Firewall.....32
Start Eucalyptus.....33
Start the CLC.....33
Start Walrus.....34
Start the CC.....34
Start the SC.....34

Start the NCs....34
Verify the Startup....35
Register Eucalyptus....35
Register Services....36
Register Walrus....36
Register the CC....36
Register the SC....36
Register the NCs....37
Register Arbitrators....37
Configure the Runtime Environment....39
Generate Administrator Credentials....39
Upload the JSON Configuration File (Edge Mode Networking)....39
Configure Object Storage Gateway....39
Configure the Storage Controller....40
Configure the Imaging Service....54
Configure DNS....56
Configure Node Controller....58
Set Up Security Groups....58
Configure the Load Balancer....59

Eucalyptus Upgrade or Migration....60

Eucalyptus Upgrade....60
Prepare the Configuration File....60
Shutdown Components....61
Upgrade Euca2ools Packages....61
Upgrade Eucalyptus Packages....61
Switch to the Stock DHCP Package....62
Start Eucalyptus....63
Sync Keys on Components....64
Register Services....65
Verify the Components....65
Update the Service Images....66
Upgrade Credentials....66
Dealing with Failed Upgrades....67
Eucalyptus Migration to Edge Networking Mode....69
Eucalyptus Migration to High Availability....69

Find More Information....71

Eucalyptus Installation from Local Package Repository....72

Euca2ools Standalone Installation....73

Tech Preview: Eucalyptus HA Installation....74

Plan Your Installation....74
- Understanding the Eucalyptus HA Architecture....74
- Plan Component Placement....75
- Plan Your Hardware....76
- Verify Component Disk Space....77
- Plan Networking Modes....77
- Plan Eucalyptus Features....81
- Prepare the Network....87

Configure Dependencies....89
- Configure Bridges....89
- Disable the Firewall....90
- Configure SELinux....90
- Configure NTP....91
- Configure an MTA....91
- Enable Packet Routing....91

Install Eucalyptus....92
- Software Signing....92
- Install Eucalyptus from Release Packages....93
- Install Eucalyptus from Nightly Packages....95

Configure Eucalyptus....96
- Configure Network Modes....96
- Configure Loop Devices....101
- Configure Multi-Cluster Networking....102
- Configure the Firewall....102

Start Eucalyptus....104
- Start the CLC Pairs....104
- Start the Walrus Pairs....104
- Start the CC Pairs....104
- Start the SC Pairs....105
- Start the NCs....105
- Verify the Startup....105

Register Eucalyptus....105
- Register the Secondary CLC....106
- Register Services....106
- Register Walrus Pairs....107
- Register the CC Pairs....107
- Register the SC Pairs....107
- Register the NCs....108
- Register Arbitrators....108

Configure the Runtime Environment....110
- Generate Administrator Credentials....110
- Upload the JSON Configuration File (Edge Mode Networking)....110

Configure Object Storage Gateway....111
Configure the Storage Controller....111
Configure DNS....123
Configure Node Controller....125
Configure DRBD....125
Skip Initial Device Synchronization....128
Synchronize Pairs Configuration....128
Set Up Security Groups....129
Configure the Load Balancer....130

Installation Overview

This topic helps you understand, plan for, and install Eucalyptus. If you follow the recommendations and instructions in this guide, you will have a working version of Eucalyptus customized for your specific needs and requirements.

This guide walks you through installations for a few different use cases. You can choose from one of the installation types listed in the following table.

What Do You Want to Do?	Installation Type
Quickly deploy Eucalyptus on one machine	If you have a CentOS 6.6 minimal install and a few IP addresses to spare, try the Faststart script. Run the following command as root: `bash <(curl -Ls hphelion.com/eucalyptus-install)`
Create a development or production environment	*Eucalyptus*
Test a development environment with high availability (technical preview)	*Eucalyptus HA*

We recommend that you read the section you choose in the order presented. There are no shortcuts for installing Eucalyptus, though Eucalyptus FastStart is fairly easy. However, to customize your installation, you have to understand what Eucalyptus is, what the installation requirements are, what your network configuration and restrictions are, and what Eucalyptus components and features are available based on your needs and requirements.

Important: If you are upgrading from a previous version of Eucalyptus, see *Eucalyptus Upgrade*.

Document version: Build 3029 (2016-01-26 12:00:31)

CC-BY-SA, Eucalyptus Systems, Inc.

Introduction to Eucalyptus

Eucalyptus is a Linux-based software architecture that implements scalable private and hybrid clouds within your existing IT infrastructure. Eucalyptus allows you to use your own collections of resources (hardware, storage, and network) using a self-service interface on an as-needed basis.

You deploy a Eucalyptus cloud across your enterprise's on-premise data center. Users access Eucalyptus over your enterprise's intranet. This allows sensitive data to remain secure from external intrusion behind the enterprise firewall.

You can install Eucalyptus on the following Linux distributions:

- CentOS 6
- Red Hat Enterprise Linux 6

Eucalyptus Overview

Eucalyptus was designed to be easy to install and as non-intrusive as possible. The software framework is modular, with industry-standard, language-agnostic communication.

Eucalyptus provides a virtual network overlay that both isolates network traffic of different users and allows two or more clusters to appear to belong to the same Local Area Network (LAN). Also, Eucalyptus offers API compatibility with Amazon's EC2, S3, IAM, ELB, Auto Scaling, and CloudWatch services. This offers you the capability of a hybrid cloud.

Eucalyptus Components

Eucalyptus is comprised of the following components: Cloud Controller (CLC), User-Facing Services (UFS), Object Storage Gateway (OSG), Cluster Controller (CC), Storage Controller (SC), and Node Controller (NC).

A detailed description of each Eucalyptus component follows.

Cloud Controller

In most deployments, the CLC and the UFS are on the same machine. This machine is the entry-point into the cloud for administrators, developers, project managers, and end-users. The CLC handles persistence and is the backend for the UFS.

User-Facing Services

The UFS serve as endpoints for the AWS-compatible services offered by Eucalyptus: EC2 (compute), AS (AutoScaling), CW (CloudWatch), ELB (LoadBalancing), IAM (Euare), and STS (tokens).

Object Storage Gateway

The OSG passes requests to object storage providers and talks to the persistence layer (DB) to authenticate requests. You can use Walrus or Riak CS as the object storage provider.

Cluster Controller

The Cluster Controller (CC) generally executes on a machine that has network connectivity to both the machines running the Node Controllers (NCs) and to the machine running the CLC. CCs gather information about a set of NCs and schedules virtual machine (VM) execution on specific NCs. The CC also manages the virtual machine networks in Managed and Managed (No VLAN) *networking modes*. All NCs associated with a single CC must be in the same subnet.

CC-BY-SA, Eucalyptus Systems, Inc.

Storage Controller

The Storage Controller (SC) provides functionality similar to the Amazon Elastic Block Store (Amazon EBS). The SC can interface with various storage systems. Elastic block storage exports storage volumes that can be attached by a VM and mounted or accessed as a raw block device. EBS volumes persist past VM termination and are commonly used to store persistent data. An EBS volume cannot be shared between VMs and can only be accessed within the same availability zone in which the VM is running. Users can create snapshots from EBS volumes. Snapshots are stored in Walrus and made available across availability zones. Eucalyptus with SAN support lets you use your enterprise-grade SAN devices to host EBS storage within a Eucalyptus cloud.

Node Controller

The Node Controller (NC) executes on any machine that hosts VM instances. The NC controls VM activities, including the execution, inspection, and termination of VM instances. It also fetches and maintains a local cache of instance images, and it queries and controls the system software (host OS and the hypervisor) in response to queries and control requests from the CC. The NC manages the virtual machine networks in Edge *networking mode*. The NC is also responsible for the management of the virtual network endpoint.

System Requirements

To install Eucalyptus, your system must meet the following baseline requirements.

Note: The specific requirements of your Eucalyptus deployment, including the number of physical machines, structure of the physical network, storage requirements, and access to software are ultimately determined by the features you choose for your cloud and the availability of infrastructure required to support those features. For further information, go to *Reference Architectures* page and look at the physical resources recommended for your deployment type.

Compute Requirements

- Physical Machines: All Eucalyptus components must be installed on physical machines, not virtual machines.
- Central Processing Units (CPUs): We recommend that each machine in your Eucalyptus cloud contain either an Intel or AMD processor with a minimum of two, 2GHz cores.
- Operating Systems: Eucalyptus supports the following Linux distributions: CentOS 6 and RHEL 6. Eucalyptus supports only 64-bit architecture.
- Machine Clocks: Each Eucalyptus component machine and any client machine clocks must be synchronized (for example, using NTP). These clocks must be synchronized all the time, not just at installation.
- Machine Access: Verify that all machines in your network allow SSH login, and that root or sudo access is available on each of them.

Storage and Memory Requirements

- Each machine in your network needs a minimum of 30 GB of storage.
- We recommend at least 100GB for Walrus and SC hosts running Linux VMs. We recommend at least 250GB for Walrus and SC hosts running Windows VMs.
- We recommend a range of 50-100GB per NC host running Linux VMs, and at least 250GB per NC host for running Windows VMs. Note that larger available disk space enables greater number of VMs.
- Each machine in your network needs a minimum of 4 GB RAM. However, we recommend more RAM for improved caching.

Network Requirements

- All NCs must have access to a minimum of 1Gb Ethernet network connectivity.
- All Eucalyptus components must have at least one Network Interface Card (NIC) for a base-line deployment. For better network isolation and scale, the CC should have two NICS (one facing the CLC/user network and one facing the NC/VM network). For HA configurations that include network failure resilience, each machine should have one

CC-BY-SA, Eucalyptus Systems, Inc.

extra NIC for each functional NIC (they will be bonded and connected to separate physical network hardware components).

- Some configurations require that machines hosting a CC have two network interfaces, each with a minimum of 1Gb Ethernet.
- For virtual machine traffic isolation, the network ports connecting Ethernet interfaces might need to allow VLAN trunking.
- For *Managed* and *Managed (No VLAN)* modes, Eucalyptus needs two sets of IP addresses.
- For *Edge* mode, Eucalyptus needs at least one existing network.
- The network connecting machines that host Eucalyptus components (except the CC and NC) must support UDP multicast for IP address 228.7.7.3. Note that UDP multicast is not used over the network that connects the CC to the NCs. For information about testing connectivity, see *Verify Connectivity*.

Once you are satisfied that your systems requirements are met, you are ready to plan your Eucalyptus installation.

CC-BY-SA, Eucalyptus Systems, Inc.

Eucalyptus Installation

This section details steps to install Eucalyptus.

To install Eucalyptus, perform the following tasks in the order presented in this section.

Plan Your Installation

In order to get the most out of a Eucalyptus deployment, we recommend that you create a plan that provides a complete set of features, performance, scaling, and resilience characteristics you want in your deployment.

Attention: If you are upgrading from an existing Eucalyptus release, see *Eucalyptus Upgrade*.

To successfully plan for your Eucalyptus installation, you must determine two things:

- **The infrastructure you plan to install Eucalyptus on:** Think about the application workload performance and resource utilization tuning. Think about how many machines you want on your system.
- **The amount of control you plan to give Eucalyptus on your network:** Use your existing architecture and policies to determine the Eucalyptus networking features you want to enable: elastic IPs, security groups, DHCP server, and Layer 2 VM isolation.

This section describes how to evaluate each tradeoff to determine the best choice to make, and how to verify that the resource environment can support the features that are enabled as a consequence of making a choice.

By the end of this section, you should be able to specify how you will deploy Eucalyptus in your environment, any tradeoffs between feature set and flexibility, and where your deployment will integrate with existing infrastructure systems.

Tip: For more help in planning your installation, see the *Eucalyptus Cloud Reference Architectures* page. This page includes use cases and reference architectures for various deployments.

Eucalyptus Architecture Overview

This topics describes the relationship of the components in a Eucalyptus installation.

The cloud components, Cloud Controller (CLC), user-facing services (UFS), and Walrus (WS), communicate with cluster components, the Cluster Controllers (CCs) and Storage Controllers (SCs). The CCs and SCs, in turn, communicate with the Node Controllers (NCs). The networks between machines hosting these components must be able to allow TCP connections between them.

However, if the CCs are on separate network interfaces (one for the network on which the cloud components are hosted and another for the network that NCs use) the CCs will act as software routers between these networks in some networking configurations. So each cluster can use an internal private network for its NCs and the CCs will route traffic from that network to a network shared by the cloud components.

Virtual machines (VMs) run on the machines that host NCs. You can use the CCs as software routers for traffic between clients outside Eucalyptus and VMs. Or the VMs can use the routing framework already in place without CC software routers. However, depending on the layer-2 isolation characteristics of your existing network, you might not be able to implement all of the security features supported by Eucalyptus.

Plan Component Placement

A Eucalyptus deployment is a set of cloud services (Cloud Controller and Walrus) and one or more clusters, each of which contains a Cluster Controller, a Storage Controller, and one or more Node Controllers.

CC-BY-SA, Eucalyptus Systems, Inc.

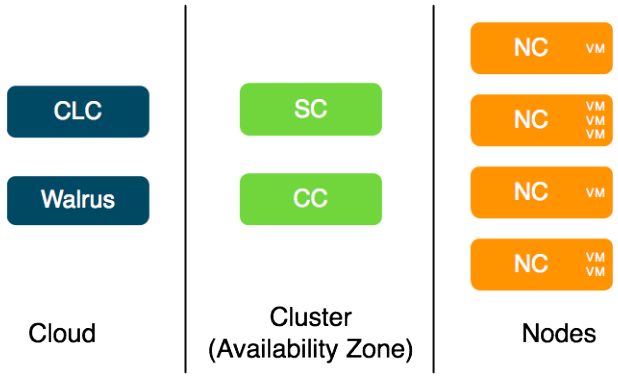

Cloud Components

The main decision for cloud components is whether to install the Cloud Controller (CLC) and Walrus on the same server. If they are on the same server, they operate as separate web services within a single Java environment, and they use a fast-path for inter-service communication. If they are not on the same server, they use SOAP and REST to work together.

Sometimes the key factor for cloud components is not performance, but server cost and data center configuration. If you only have one server available for the cloud, then you have to install the components on the same server.

All components should be in the same data center. They use aggressive time-outs to maintain system responsiveness so separating them over a long-latency, lossy network link will not work.

Cluster Components

The Eucalyptus components deployed in the cluster level of a Eucalyptus deployment are the Cluster Controller (CC) and Storage Controller (SC).

You can install all cluster components on a single machine, or you can distribute them on different machines. The choice of one or multiple machines is dictated by the demands of user workload in terms of external network utilization (CC) and EBS volume access (SC).

Things to consider for CC placement:

- For Managed and Managed (No VLAN) modes, the CC physical machine becomes a software IP gateway between VM instances and the public network. Because of this software routing function, the physical server on which the CC is deployed should have fast, dedicated network access to both the NC network, and the public network.
- For Edge mode, the CC physical machine will not act as a software gateway. Network traffic will be limited to small control messages.
- In all cases, place the CC on a machine that has TCP/IP connectivity to the Eucalyptus front end servers and the NC servers in its cluster.

Things to consider for SC placement:

- The machine on which the SC is deployed must always have TCP/IP connectivity to the CLC. If you are a subscriber and use one of Eucalyptus' provided SAN integration drivers, the SC must also have TCP/IP connectivity to the chosen SAN device. In this case, the SC only sends control messages to the SAN.
- If you do not configure a SAN, the SC requires only TCP/IP connectivity to the NCs in the cluster. The SC will use this TCP/IP connectivity to provide the NCs network access to the dynamic block volumes residing on the SC's storage. SC storage should consist of a fast, reliable disk pool (either local file-system or block-attached storage) so that the SC can create and maintain volumes for the NCs. The capacity of the disk pool should be sufficient to provide the NCs with enough space to accommodate all dynamic block volumes requests from end-users

Node Components

The Node Controllers are the components that comprise the Eucalyptus back-end. All NCs must have network connectivity to whatever hosts their EBS volumes. This host is either a SAN or the SC.

CC-BY-SA, Eucalyptus Systems, Inc.

Plan Your Hardware

This topic describes ways you can install Eucalyptus components on your machines.

You can run Eucalyptus components in any combination on the various physical servers in a data center. You can install the Cloud Controller, Walrus, CC, and SC on one machine, and an NC on one or more machines. Or you can install each component on an independent physical server. This gives each component the most local resource usage.

Often in installation decisions, you must trade deployment simplicity for performance or high-availability. For example, if you place all cloud and cluster components on a single machine, it makes for simple administration. This is because there is only one machine to monitor and control for the Eucalyptus control services. But, each component acts as an independent web service. So if these components share a single machine, the physical resources available to each service could become a performance bottleneck.

Verify Component Disk Space

Eucalyptus components need disk space for log files, databases, buckets, and instances. The following table details the needs of each component. Verify that the machines you plan to install the components on have adequate space.

We recommend that you choose a disk for each Walrus that is large enough to hold all objects and buckets you ever expect to have, including all images that will ever be registered to your system, plus any Amazon S3 application data. For consistent performance, we recommend that you use identical disks for the primary and secondary Walrus.

Tip: We recommend that you use LVM (Logical Volume Manager). If you run out of disk space, LVM allows you to add disks and migrate the data.

Component	Directory	Minimum Size
Cloud Controller (CLC) CLC logging	`/var/lib/eucalyptus/db` `/var/log/eucalyptus`	20GB 2GB
Walrus Walrus logging	`/var/lib/eucalyptus/bukkits` `/var/log/eucalyptus`	250GB 2GB
Storage Controller (SC)	`/var/lib/eucalyptus/volumes` (EBS storage) **Important:** This disk space on the SC is only required if you are not using a SAN driver or if you are using Direct Attached Storage (DAS). For more information, see either *Configure the Storage Controller* or *Configure the Storage Controller* (HA).	250GB
Cluster Controller (CC) CC logging	`/var/lib/eucalyptus/CC` `/var/log/eucalyptus`	5GB 2GB
Node Controller (NC) NC logging	`/var/lib/eucalyptus/instances` `/var/log/eucalyptus`	250GB 2GB

If necessary, create symbolic links to larger filesystems from the above locations. Make sure that the eucalyptus user owns the directories.

Plan Networking Modes

Eucalyptus overlays a virtual network on top of your existing network. In order to do this, Eucalyptus supports three different networking modes: Edge, Managed, and Managed (No VLAN).

CC-BY-SA, Eucalyptus Systems, Inc.

Each mode is designed to allow you to choose an appropriate level of security and flexibility. The purpose of these modes is to direct Eucalyptus to use different network features to manage the virtual networks that connect VMs to each other and to clients external to Eucalyptus.

A Eucalyptus installation must be compatible with local site policies and configurations (e.g., firewall rules). Eucalyptus configuration and deployment interfaces allow a wide range of options for specifying how it should be deployed. However, choosing between these options implies tradeoffs.

Your choice of networking mode depends on the following considerations:

- Do you plan to support elastic IPs and security groups?
- Do you plan to provide your own network DHCP server?

These networking features are described in the following table:

Feature	Description	Mode
Elastic IPs	Eucalyptus instances typically have two IPs associated with them: a private one and a public one. Private IPs are intended for internal communications between instances and are usually only routable within a Eucalyptus cloud. Public IPs are used for external access and are usually routable outside of Eucalyptus cloud. How these addresses are allocated and assigned to instances is determined by a networking mode. The distinction between public and private addresses becomes important in Edge, Managed, and Managed (No VLAN) modes, which support elastic IPs. With elastic IPs the user gains control over a set of static IP addresses. Once allocated to the user, those same IPs can be dynamically associated to running instances, overriding pre-assigned public IPs. This allows users to run well-known services (for example, web sites) within the Eucalyptus cloud and to assign those services fixed IPs that do not change.	Edge Managed Managed (No VLAN)
Security groups	Security groups are sets of networking rules that define the access rules for all VM instances associated with a group. For example, you can specify ingress rules, such as allowing ping (ICMP) or SSH (TCP, port 22) traffic to reach VMs in a specific security group. When you create a VM instance, unless otherwise specified at instance run-time, it is assigned to a default security group that denies incoming network traffic from all sources. Thus, to allow login and usage of a new VM instance you must authorize network access to the default security group with the `euca-authorize` command.	Edge Managed Managed (No VLAN)
VM isolation	Although network traffic between VM instances belonging to a security group is always open, Eucalyptus can enforce isolation of network traffic between different security groups. This isolation is enforced using ebtables (Edge) or VLAN tags (Managed), thus, protecting VMs from possible eavesdropping by VM instances belonging to other security groups.	Edge Managed
DHCP server	Eucalyptus assigns IP addresses to VMs in all modes.	Edge Managed Managed (No VLAN)

If Eucalyptus can control and condition the networks its components use, your deployment will support the full set of API features. However, if Eucalyptus is confined to using an existing network, some of the API features might be disabled. So, understanding and choosing the right networking configuration is an important (and complex) step in deployment planning.

Each networking mode is detailed in the following sections.

CC-BY-SA, Eucalyptus Systems, Inc.

Edge Mode

Edge mode offers the most features of the networking modes. It is designed to integrate into already extant (or straightforward to deploy) underlying network topologies. However, Edge mode can impose constraints in certain environments, in which case you can choose another mode.

In Edge networking mode, the component responsible for implementing Eucalyptus VM networking artifacts is running at the edge of a Eucalyptus deployment: the Node Controller (NC). Eucalyptus provides a stand-alone component called `eucanetd` in each NC. This component dynamically receives changing Eucalyptus networking views and is responsible for configuring the Linux machine on which the NC is running to reflect the latest view.

Edge networking mode integrates with your existing network infrastructure, allowing you to 'tell' Eucalyptus (through configuration parameters for Edge mode) about the existing network, which Eucalyptus then will consume when implementing the networking view.

Edge networking mode integrates with two basic types of pre-existing network setups:

- One flat IP network used to service Eucalyptus component systems, Eucalyptus VM public IPs (elastic IPs), and Eucalyptus VM private IPs.
- Two networks, one for Eucalyptus components and Eucalyptus VM public IPs, and the other for Eucalyptus VM private IPs.

Important: Edge networking mode will not set up the network from scratch as Managed and Managed (No VLAN) modes do. Instead, it integrates with networks that already exist. If the network, netmask, and router don't already exist, you must create them outside of Eucalyptus before configuring Edge.

Edge Mode Requirements

- Each NC must have an interface configured with an IP on a VM public and a VM private network (which can be the same network).
  - There must be unused IP addresses on the VM public network for Eucalyptus to assign VM elastic IPs
  - There must be unused IP addresses on the VM private network for Eucalyptus to assign VM private IPs
- There must be IP connectivity from each NC machine (where `eucanetd` runs) and the CLC, for metadata re-directs for 169.254.169.254 to the active CLC to function.
- There must be a functioning router in place for the private network. This router will be the default gateway for VM instances.
- The private and public networks can be the same network, but they can also be separate networks.
- The Node Controllers (NCs) need a bridge configured on the private network, with the bridge interface itself having been assigned an IP from the network.
- If you're using a public network, the NCs need an interface on the public network as well (if the public and private networks are the same network, then the bridge needs an IP assigned on the network).
- If you run in multi-cluster, each cluster can use the same network as its private network, or they can use separate networks as private networks. If you use separate networks, you need to have a router in place that is configured to route traffic between the networks.
- If you use private addressing only mode, the Cloud Controller machines must have a route back to the VM private network.

Edge Mode Limitations

- All NCs must have an interface on the VM public (Elastic IP) network.
- Global network updates (such as security group rule updates, security group VM membership updates, and elastic IP updates) are be applied through an "eventually consistent" mechanism, as opposed to an "atomic" mechanism. That is, there may be a brief period of time where one NC has the new state implemented but another NC has the previous state implemented.
- Mappings between VM MAC addresses and private IPs are strictly enforced.

CC-BY-SA, Eucalyptus Systems, Inc.

Managed Mode

In Managed mode, Eucalyptus manages the local network of VM instances and provides all networking features Eucalyptus currently supports, including VM network isolation, security groups, elastic IPs, and metadata service.

In Managed mode, you define a large network (usually private, unroutable) from which VM instances will draw their private IP addresses. Eucalyptus maintains a DHCP server with static mappings for each VM instance that is created. When you create a new VM instance, you can specify the name of the security group to which that VM will belong. Eucalyptus then selects a subset of the entire range of IPs, to hand out to other VMs in the same security group.

You can also define a number of security groups, and use those groups to apply network ingress rules to any VM that runs within that network. In this way, Eucalyptus provides functionality similar to Amazon's security groups. In addition, the administrator can specify a pool of public IP addresses that users may allocate, then assign to VMs either at boot or dynamically at run-time. This capability is similar to Amazon's 'elastic IPs'. Eucalyptus administrators that require security groups, elastic IPs, and VM network isolation must use this mode.

Managed mode uses a Virtual LAN (VLAN) to enforce network isolation between instances in different security groups. If your underlying physical network is also using a VLAN, there can be conflicts that prevent instances from being network accessible. So you have to determine if your network between the CC and NCs is VLAN clean (that is, if your VLANs are usable by Eucalyptus). To test if the network is VLAN clean, see VLAN Preparation.

Each VM receives two IP addresses: a public IP address and a private IP address. Eucalyptus maps public IP addresses to private IP addresses. Access control is managed through security groups.

Managed Mode Requirements

- There must be an available range of IP addresses for the virtual subnets. This range must not interfere with the physical network. Typically these IP addresses are selected from the private IP ranges: 192.168.x.x, 10.x.x.x, etc.
- The network between the CC and NCs must be VLAN clean, meaning that all switch ports that Eucalyptus components are connected to will allow and forward VLAN tagged packets.
- Any firewall running on the Cluster Controller must be compatible with the dynamic changes performed by Eucalyptus when working with security groups. (Note that Eucalyptus will flush the 'filter' and 'nat' tables upon boot).
- Any DHCP server on the subnet must be configured not to serve Eucalyptus instances.
- There must be a separate Layer 2 network for each cluster in a multi-cluster setup.

Managed Mode Limitations

- The machine that hosts the CC will be a router in the data path for any VM traffic that is not 'VM private IP to VM private IP, where both VMs are in the same security group'.
- Network switch must be properly configured. For more information, see *Prepare VLAN* or *Prepare VLAN* (for HA).
- In non-HA mode, the machine that hosts the CC is a single point of failure for most VM network communication.

Managed (No VLAN) Mode

In Managed (No VLAN) mode, Eucalyptus fully manages the local VM instance network and provides all of the networking features Eucalyptus currently supports, including security groups, elastic IPs, etc. However, it does not provide VM network isolation.

Without VLAN isolation at the bridge level, it is possible in Managed (No VLAN) mode for a root user on one VM to snoop and/or interfere with the ethernet traffic of other VMs running on the same layer 2 network.

Tip: In Managed (No VLAN) mode, VM isolation is provided by having different security groups on different subnets—this translates into Layer-3 only VM isolation.

Managed (No VLAN) Mode Requirements

- There must be an available range of IP addresses for the virtual subnets. This range must not interfere with the physical network. Typically these IP addresses are selected from the private IP ranges: 192.168.x.x, 10.x.x.x, etc.
- Any firewall running on the Cluster Controller must be compatible with the dynamic changes performed by Eucalyptus when working with security groups. (Note that Eucalyptus will flush the 'filter' and 'nat' tables upon boot).

CC-BY-SA, Eucalyptus Systems, Inc.

- A range of public IP addresses must be available for use by Eucalyptus.
- The CC must have a DHCP server daemon installed that is compatible with ISC DHCP Daemon version 3.0.X.
- If you plan to set up more than one cluster, you need to have a bridge for security groups to span the clusters.

Managed (No VLAN) Mode Limitations

- Limited (Layer-3) VM isolation.

Plan Eucalyptus Features

Before you install Eucalyptus, we recommend that you think about the features you plan to implement with Eucalyptus. These features are detailed in the following sections.

Windows Guest OS Support

This topic details what Eucalyptus needs in order to use Windows as a guest operating system.

- A licensed installation copy (.iso image or CD/DVD disk) of a compatible Windows OS. Eucalyptus currently supports Windows virtual machines created from Windows Server 2003 R2 Enterprise (32/64 bit); Windows Server 2008 SP2, Datacenter (32/64 bit); Windows Server 2008 R2, Datacenter; and Windows 7 Professional.
- A VNC client such as RealVNC or Virtual Manager/Virtual Viewer for initial installation. Subsequent Eucalyptus-hosted Windows instances will use RDP, but the initial installation requires VNC.

For additional Windows-related licensing information, see the following links:

- *http://technet.microsoft.com/en-us/library/dd979803.aspx*
- *http://technet.microsoft.com/en-us/library/dd878528.aspx*
- *http://technet.microsoft.com/en-us/library/dd772269.aspx*

SAN Support

Eucalyptus includes optional, subscription only support for integrating enterprise-grade SAN (Storage Area Network) hardware devices into a Eucalyptus cloud.

SAN support extends the functionality of the Eucalyptus Storage Controller (SC) to provide a high performance data conduit between VMs running in Eucalyptus and attached SAN devices. Eucalyptus dynamically manages SAN storage without the need for the administrator to manually allocate and de-allocate storage, manage snapshots or set up data connections.

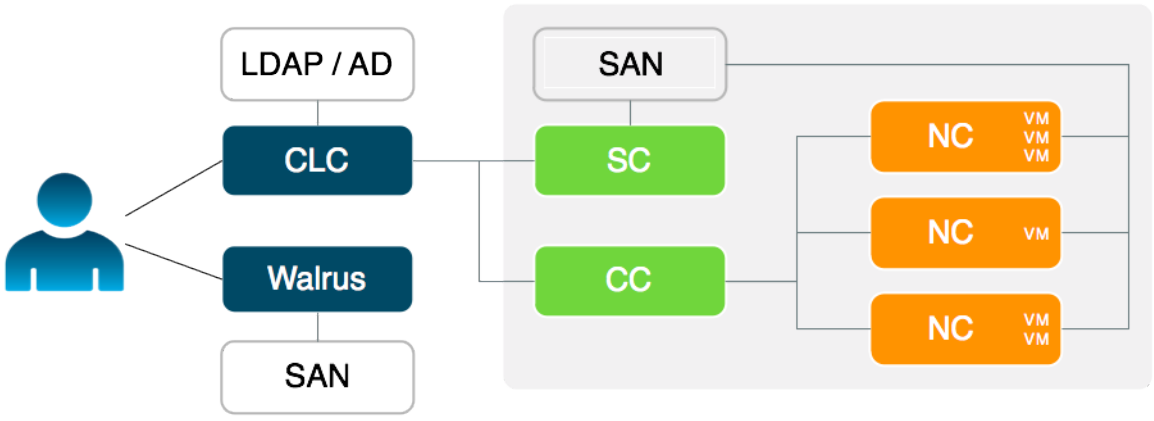

Eucalyptus with SAN support allows you to:

- Integrate Eucalyptus block storage functionality (dynamic block volumes, snapshots, creating volumes from snapshots, etc.) with existing SAN devices
- Link VMs in the Eucalyptus cloud directly to SAN devices, thereby removing I/O communication bottlenecks of the physical hardware host
- Incorporate enterprise-level SAN features (high-speed, large-capacity, reliability) to deliver a production-ready EBS (block storage) solution for the enterprise

CC-BY-SA, Eucalyptus Systems, Inc.

- Attach SAN devices to Eucalyptus deployments on Xen and KVM

To use Eucalyptus with supported SAN storage, you must decide whether administrative access can be provided to Eucalyptus to control the SAN. If this is possible in your environment, Eucalyptus can automatically and dynamically manage SAN storage.

Currently, the Dell Equallogic series of SANs (PS 4000 and PS 6000), NetApp Filer FAS 2000 and FAS 6000 series and EMC VNX are supported. For Dell Equallogic, Eucalyptus requires SSH access to enable automatic provisioning. Eucalyptus will manage NetApp SANs via ONTAPI (version 7.3.3 and above). For EMC, Eucalyptus expects that the EMC NaviSecCLI software will be installed on the Storage Controller host.

SAN Support Prerequisites

Eucalyptus supports the following SAN devices:

- Dell EqualLogic, PS4000 series and PS6000 series (For more information about Dell EqualLogic SANs, go to *http://www.dell.com*)
- NetApp, FAS2000 series and FAS6000 series (For more information about NetApp SANs, go to *http://www.netapp.com*
- EMC VNX Series (For more information about EMC VNX, go to *VNX Family*

For additional information on SAN support for Eucalyptus, contact Eucalyptus Systems, Inc.

Availability Zone Support

Eucalyptus offers the ability to create multiple availability zones. In Eucalyptus, an availability zone is a partition in which there is at least one available cluster.

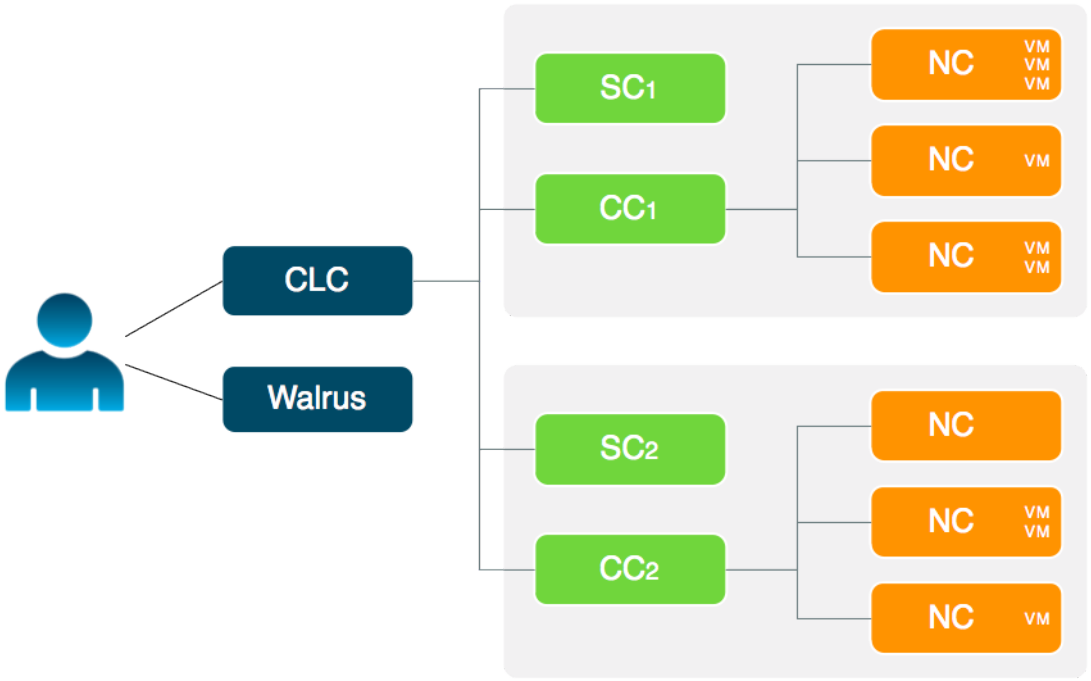

Object Storage

Eucalyptus supports Walrus and Riak CS as its object storage backend. There is no extra planning if you use Walrus. If you use Riak CS, you can use a single Riak CS cluster for several Eucalyptus clouds. Basho (the vendor of RiakCS) recommends five nodes for each Riak CS cluster. This also means that you have to set up and configure a load balancer between the Riak CS nodes and the object storage gateway (OSG).

Prepare the Network

In order for Eucalyptus to function in your local environment, be sure to prepare your network. To prepare your network, perform the tasks listed in this section.

CC-BY-SA, Eucalyptus Systems, Inc.

Reserve Ports

Eucalyptus components use a variety of ports to communicate. The following table lists the all of the important ports used by Eucalyptus.

Port	Description
TCP 5005	DEBUG ONLY: This port is used for debugging Eucalyptus (using the `--debug` flag).
TCP 8080	Port for getting user credentials on the CLC. Configurable with `euca-modify-property`.
TCP 8772	DEBUG ONLY: JMX port. This is disabled by default, and can be enabled with the `--debug` or `--jmx` options for `CLOUD_OPTS`.
TCP 8773	Web services port for the CLC, user-facing services (UFS), object storage gateway (OSG), Walrus SC; also used for external and internal communications by the CLC and Walrus. Configurable with `euca-modify-property`.
TCP 8774	Web services port on the CC. Configured in the `eucalyptus.conf` configuration file
TCP 8775	Web services port on the NC. Configured in the `eucalyptus.conf` configuration file.
TCP 8777	Database port on the CLC
TCP 8779 (or next available port, up to TCP 8849)	jGroups failure detection port on CLC, UFS, OSG, Walrus SC. If port 8779 is available, it will be used, otherwise, the next port in the range will be attempted until an unused port is found.
TCP 8888	The default port for the Eucalyptus Management Console. Configured in the `/etc/eucalyptus-console/console.init` file.
TCP 16514	TLS port on Node Controller, required for node migrations
UDP 7500	Port for diagnostic probing on CLC, UFS, OSG, Walrus SC
UDP 8773	HA membership port
UDP 8778	The bind port used to establish multicast communication
TCP/UDP 53	DNS port on UFS

Verify Connectivity

Verify connectivity between the machines you'll be installing Eucalyptus on. Some Linux distributions provide default TCP/IP firewalling rules that limit network access to machines. Disable these default firewall settings before you install Eucalyptus components to ensure that the components can communicate with one another.

Note: Any firewall running on the CC must be compatible with the dynamic changes performed by Eucalyptus when working with security groups. Eucalyptus will flush the 'filter' and 'nat' tables upon boot.

Verify component connectivity by performing the following checks on the machines that will be running the listed Eucalyptus components.

1. Verify connection from an end-user to the CLC on TCP ports 8443 and 8773
2. Verify connection from an end-user to Walrus on TCP port 8773
3. Verify connection from the CLC, SC, and NC to SC on TCP port 8773
4. Verify connection from the CLC, SC, and NC to Walrus on TCP port 8773
5. Verify connection from Walrus and SC to CLC on TCP port 8777
6. Verify connection from CLC to CC on TCP port 8774
7. Verify connection from CC to NC on TCP port 8775
8. Verify connection from NC to Walrus on TCP port 8773. Or, you can verify the connection from the CC to Walrus on port TCP 8773, and from an NC to the CC on TCP port 8776
9. Verify connection from public IP addresses of Eucalyptus instances (metadata) and CC to CLC on TCP port 8773

CC-BY-SA, Eucalyptus Systems, Inc.

10. Verify TCP connectivity between CLC, Walrus, and SC on TCP port 8779 (or the first available port in range 8779-8849)
11. Verify connection between CLC, Walrus, and SC on UDP port 7500
12. Verify multicast connectivity for IP address 228.7.7.3 between CLC, Walrus, and SC on UDP port 8773
13. If DNS is enabled, verify connection from an end-user and instance IPs to DNS ports
14. If you use tgt (iSCSI open source target) for EBS storage, verify connection from NC to SC on TCP port 3260
15. Test multicast connectivity between each CLC and Walrus, and SC.
    a) Clone the Eucalyptus `deveutils` repository

    ```
    git clone https://github.com/eucalyptus/deveutils
    ```

    b) Run the `network-tomography` tool on the Cloud Controller, Cluster Controller, Storage Controller, and any machines running Walrus, passing a list of IP addresses for each of these machines.

    ```
    cd deveutils/network-tomography
    ./network-tomography 192.168.51.174 192.168.51.196 192.168.51.86
    192.168.51.99
    ```

    This tool may take up to an hour to run. Check the output for reports of packet loss. If there is significant packet loss, ensure that your network is available and multicast enabled.

Prepare VLAN

Managed networking mode requires that switches and routers be “VLAN clean.” This means that switches and routers must allow and forward VLAN tagged packets. If you plan to use the Managed networking mode, you can verify that the network is VLAN clean between machines running Eucalyptus components by performing the following test.

Tip: You only need to read this section if you are using Managed mode. If you aren’t using Managed mode, skip this section.

1. Choose two IP addresses from the subnet you plan to use with Eucalyptus, one VLAN tag from the range of VLANs that you plan to use with Eucalyptus, and the network interface that will connect your planned CC and NC servers. The examples in this section use the IP addresses 192.168.1.1 and 192.168.1.2, VLAN tag 10, and network interface eth3, respectively.
2. On the planned CC server, choose the interface on the local Ethernet and run:

   ```
   vconfig add eth3 10
   ifconfig eth3.10 192.168.1.1 up
   ```

3. On a planned NC server, choose the interface on the local network and run:

   ```
   vconfig add eth3 10
   ifconfig eth3.10 192.168.1.2 up
   ```

4. On the NC, ping the CC:

   ```
   ping 192.168.1.1
   ```

5. On the CC, ping the NC:

   ```
   ping 192.168.1.2
   ```

   - If this VLAN clean test fails, configure your switch to forward VLAN tagged packets. If it is a managed switch, see your switch's documentation to determine how to do this.
   - If the VLAN clean test passes, continue with the following steps to remove the test interfaces.

6. On the CC, remove the test interface by running:

   ```
   vconfig rem eth3.10
   ```

7. On the planned NC, run:

CC-BY-SA, Eucalyptus Systems, Inc.

```
vconfig rem eth3.10
```

Configure Dependencies

Before you install Eucalyptus, make sure you have the following dependencies installed and configured.

Configure Bridges

For Managed (No VLAN) and EDGE modes, you must configure a Linux ethernet bridge on all NCs. This bridge connects your local ethernet adapter to the cluster network. Under normal operation, NCs will attach virtual machine instances to this bridge when the instances are booted.

To configure a bridge in CentOS 6 or RHEL6, you need to create a file with bridge configuration (for example, ifcfg-brX) and modify the file for the physical interface (for example, ifcfg-ethX). The following steps describe how to set up a bridge on both CentOS 6 and RHEL 6. We show examples for configuring bridge devices that either obtain IP addresses using DHCP or statically.

1. Install the `bridge-utils` package.

   ```
   yum install bridge-utils
   ```

2. Go to the `/etc/sysconfig/network-scripts` directory:

   ```
   cd /etc/sysconfig/network-scripts
   ```

3. Open the network script for the device you are adding to the bridge and add your bridge device to it. The edited file should look similar to the following:

   ```
   DEVICE=eth0
   # change the hardware address to match the hardware address your NIC uses
   HWADDR=00:16:76:D6:C9:45
   ONBOOT=yes
   BRIDGE=br0
   NM_CONTROLLED=no
   ```

4. Create a new network script in the `/etc/sysconfig/network-scripts` directory called `ifcfg-br0` or something similar. The br0 is the name of the bridge, but this can be anything as long as the name of the file is the same as the `DEVICE` parameter, and the name is specified correctly in the previously created physical interface configuration (ifcfg-ethX).

   - If you are using DHCP, the configuration will look similar to:

     ```
     DEVICE=br0
     TYPE=Bridge
     BOOTPROTO=dhcp
     ONBOOT=yes
     DELAY=0
     ```

   - If you are using a static IP address, the configuration will look similar to:

     ```
     DEVICE=br0
     TYPE=Bridge
     BOOTPROTO=static
     IPADDR=<static_IP_address>
     NETMASK=<netmask>
     GATEWAY=<gateway>
     ONBOOT=yes
     ```

5. Enter the following command:

   ```
   service network restart
   ```

CC-BY-SA, Eucalyptus Systems, Inc.

Disable the Firewall

If you have existing firewall rules on your hosts, you should disable the firewall in order to install Eucalyptus. You should re-enable it after installation.

Tip: If you do not have a firewall enabled, skip this step.

1. To disable your firewall:
   a) Run the command `system-config-firewall-tui`
   b) Turn off the **Enabled** check box.
2. Repeat on each host that will run a Eucalyptus component: Cloud Controller, Walrus, Cluster Controller, Storage Controller, and Node Controllers.

Configure SELinux

Security-enabled Linux (SELinux) is security feature for Linux that allows you to set access control through policies. Eucalyptus is not compatible with SELinux.

To configure SELinux to allow Eucalyptus access:

1. Open `/etc/selinux/config` and edit the line `SELINUX=enforcing` to `SELINUX=permissive`.
2. Save the file.
3. Run the following command:

```
setenforce 0
```

Configure NTP

Eucalyptus requires that each machine have the Network Time Protocol (NTP) daemon started and configured to run automatically on reboot.

To use NTP:

1. Install NTP on the machines that will host Eucalyptus components.

```
yum install ntp
```

2. Open the `/etc/ntp.conf` file and add NTP servers, as in the following example.

```
server 0.pool.ntp.org
server 1.pool.ntp.org
server 2.pool.ntp.org
```

3. Save and close the file.
4. Configure NTP to run at reboot.

```
chkconfig ntpd on
```

5. Start NTP.

```
service ntpd start
```

6. Synchronize your server.

```
ntpdate -u <your_ntp_server>
```

7. Synchronize your system clock, so that when your system is rebooted, it does not get out of sync.

```
hwclock --systohc
```

8. Repeat on each host that will run a Eucalyptus component.

CC-BY-SA, Eucalyptus Systems, Inc.

Configure an MTA

All machines running the Cloud Controller must run a mail transport agent server (MTA) on port 25. Eucalyptus uses the MTA to deliver or relay email messages to cloud users' email addresses.

You can use Sendmail, Exim, postfix, or something simpler. The MTA server does not have to be able to receive incoming mail.

Many Linux distributions satisfy this requirement with their default MTA. For details about configuring your MTA, go to the documentation for your specific product.

To test your mail relay for localhost, send email to yourself from the terminal using `mail`.

Enable Packet Routing

Edit the `sysctl.conf` on each machine you plan to install the Cluster Controller (CC) component and the Node Controller (NC) on.

In the `sysctl.conf` file, set the following parameters and values:

1. Enable IP forwarding.

```
net.ipv4.ip_forward = 1
```

2. Enable the bridge to forward traffic based on iptables rules.

```
net.bridge.bridge-nf-call-iptables = 1
```

Install Eucalyptus

Eucalyptus installation packages are available for CentOS 6 and RHEL 6. The following sections show installation steps on each supported Linux distribution.

Eucalyptus Subscription allows you access to additional software modules. If you are a subscriber, you will receive an entitlement certificate and a private key that allow you to download Eucalyptus subscription modules. You will also receive a GPG public key to be used to verify the Eucalyptus software's integrity. The files will come in the form of a platform specific package.

Software Signing

We use a number of GPG keys to sign our software packages and package repositories. The necessary public keys are provided with the relevant products and can be used to automatically verify software updates. You can also verify the packages or package repositories manually using the keys on this page.

Use the `rpm --checksig` command on a download file to verify a RPM package for an HP Helion Eucalyptus product. For example:

```
rpm --checksig -v myfilename.rpm
```

Follow the procedure detailed on Debian's *SecureApt* web page to verify a deb package for an HP Helion Eucalyptus product.

Please do not use package signing keys to encrypt email messages.

The following keys are used for signing Eucalyptus software:

c1240596: Eucalyptus Systems, Inc. (release key) <security@eucalyptus.com>

This key is used for signing HP Helion Eucalyptus products released after July 2011 and their updates.

- *Download from Helion Eucalyptus*
- *Download from pgp.mit.edu*
- Fingerprint: `8639 B2D2 11BB 930D 16A2 D654 BE26 4D09 C124 0596`

CC-BY-SA, Eucalyptus Systems, Inc.

0260cf4e: Eucalyptus Systems, Inc. (pre-release key) <security@eucalyptus.com>

This key is used for signing HP Helion Eucalyptus pre-release products due for release after July 2011.

- *Download from Helion Eucalyptus*
- *Download from pgp.mit.edu*
- Fingerprint: `7363 5F5A 9531 308B E83D 3413 8B94 DFB5 0260 CF4E`

9d7b073c: Eucalyptus Systems, Inc. (nightly release key) <security@eucalyptus.com>

This key is used for signing nightly builds of HP Helion Eucalyptus products published after July 2011.

- *Download from Helion Eucalyptus*
- *Download from pgp.mit.edu*
- Fingerprint: `708C DB7C B08A 204E C4CB FA6A 4882 7EF9 9D7B 073C`

Install Eucalyptus from Release Packages

To install Eucalyptus from release packages, perform the tasks listed in this topic.

If you plan to install Eucalyptus HA, we recommend that you install each Eucalyptus component on a separate host. For example, if you are installing CLC, Walrus, CC, and SC, you will install each of these components on a separate host. You will also install each secondary component (the secondary CLC, Walrus, CC, and SC) on a separate host. In this case, you will need eight machines. Each additional cluster needs four more machines for its CCs and SCs. This does not account for NCs, which are not redundant.

1. Configure the Eucalyptus package repository on each host that will run a Eucalyptus component:

   ```
   yum install
   http://downloads.eucalyptus.com/software/eucalyptus/4.1/centos/6/x86_64/eucalyptus-release-4.1.el6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

2. (Optional) If you are a Eucalyptus subscriber, you will receive two rpm package files containing your license for subscription-only components. Install these packages on each host that will run a Eucalyptus component.

   a) Install the license files to access the enterprise repository.

   ```
   yum install eucalyptus-enterprise-license*.noarch.rpm \
   http://subscription.eucalyptus.com/eucalyptus-enterprise-release-4.1-1.el6.noarch.rpm
   ```

   b) Install the actual repository file.

   ```
   yum install eucalyptus-enterprise-release-4.1*.noarch.rpm
   eucalyptus-license-*.rpm
   ```

3. Configure the Euca2ools package repository on each host that will run a Eucalyptus component or Euca2ools:

   ```
   yum install
   http://downloads.eucalyptus.com/software/euca2ools/3.2/centos/6/x86_64/euca2ools-release-3.2.el6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

4. Configure the EPEL package repository on each host that will run a Eucalyptus component or Euca2ools:

   ```
   yum install
   http://downloads.eucalyptus.com/software/eucalyptus/4.1/centos/6/x86_64/epel-release-6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

5. If you are installing on RHEL 6, you must enable the Optional repository in Red Hat Network for each NC, as follows:

   a) Go to *http://rhn.redhat.com* and navigate to the system that will run the NC.

   b) Click **Alter Channel Subscriptions**.

   c) Make sure the **RHEL Server Optional** checkbox is checked.

CC-BY-SA, Eucalyptus Systems, Inc.

d) Click **Change Subscriptions**.

6. a) Install the Eucalyptus node controller software on each planned NC host:

```
yum install eucalyptus-nc
```

b) Remove the default libvirt network. This step allows the eucanetd dhcpd server to start.

```
virsh net-destroy default
virsh net-autostart default --disable
```

c) Check that the KVM device node has proper permissions.

Run the following command:

```
ls -l /dev/kvm
```

Verify the output shows that the device node is owned by user root and group kvm.

```
crw-rw-rw- 1 root kvm 10, 232 Nov 30 10:27 /dev/kvm
```

If your kvm device node does not have proper permissions, you need to reboot your NC host.

7. If you plan to run in *Edge networking mode*, install the package for Edge support on each planned NC host.

```
yum install eucanetd
```

8. On each planned CLC host, install the Eucalyptus cloud controller software.

```
yum install eucalyptus-cloud
```

9. Install the services image package on the machine hosting the primary CLC:

```
yum install eucalyptus-service-image -y
```

This will install worker images for both the Load Balancer and Imaging services.

10. Install the software for the remaining Eucalyptus components. The following example shows most components being installed on the same host. We recommend that you use different hosts for each component:

```
yum install eucalyptus-cc eucalyptus-sc eucalyptus-walrus
```

11. If you are a subscriber and use SAN, run the appropriate command for your device on each machine hosting a CLC:

For EMC SAN:

```
yum install eucalyptus-enterprise-storage-san-emc-libs
```

For EqualLogic SAN:

```
yum install eucalyptus-enterprise-storage-san-equallogic-libs
```

For NetApp SAN:

```
yum install eucalyptus-enterprise-storage-san-netapp-libs
```

12. If you are a subscriber and use SAN, run the appropriate command for your device on each machine hosting a SC:

For EMC SAN:

```
yum install eucalyptus-enterprise-storage-san-emc
```

Important: To use Eucalyptus with EMC SAN support, you must have the `NaviCLI-Linux-64-latest.rpm` package installed on each SC. This package is not supplied with Eucalyptus, please see your SAN vendor if it is not already installed.

For EqualLogic SAN:

```
yum install eucalyptus-enterprise-storage-san-equallogic
```

CC-BY-SA, Eucalyptus Systems, Inc.

For NetApp SAN:

```
yum install eucalyptus-enterprise-storage-san-netapp
```

13. After you have installed Eucalyptus, test multicast connectivity between each CLC and Walrus, and SC.
    a) Install the network tomography package on the Cloud Controller, Cluster Controller, Storage Controller, and any machines running Walrus.

    ```
    yum install --nogpgcheck
    http://downloads.eucalyptus.com/software/tools/centos/6/x86_64/network-tomography-1.0.0-3.el6.x86_64.rpm
    ```

    b) Run the `network-tomography` tool on the Cloud Controller, Cluster Controller, Storage Controller, and any machines running Walrus, passing a list of IP addresses for each of these machines.

    ```
    /usr/bin/network-tomography 192.168.51.174 192.168.51.196 192.168.51.86
    192.168.51.99
    ```

    This tool may take up to an hour to run. Check the output for reports of packet loss. If there is significant packet loss, ensure that your network is available and multicast enabled.

Your installation is complete.

You are now ready to *Configure Eucalyptus*.

Install Eucalyptus from Nightly Packages

To install Eucalyptus from nightly builds, perform the tasks listed in this topic.

Important: Eucalyptus nightly packages are latest Eucalyptus builds. They should be considered unstable/"bleeding edge" software and should not be installed in production. In addition, upgrades from nightlies to released software are not supported.

To install Eucalyptus nightly builds on servers running CentOS 6 or RHEL 6:

1. On all servers, run the following commands:

   ```
   yum install http://downloads.eucalyptus.com/
   software/eucalyptus/nightly/4.1/centos/6/x86_64/eucalyptus-release-4.1.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

2. On all systems that will run either Eucalyptus or Euca2ools, run the following commands:

   ```
   yum install http://downloads.eucalyptus.com/software/
   euca2ools/3.2/centos/6/x86_64/euca2ools-release-3.2.el6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

3. On all servers, enter:

   ```
   yum update
   ```

4. Install Eucalyptus packages. The following example shows most components being installed all on the same server. You can use different servers for each component.

   ```
   yum install eucalyptus-cloud
   yum install eucalyptus-cc eucalyptus-sc eucalyptus-walrus
   ```

   For HA: For Eucalyptus HA, you must install these packages on pairs of systems. For instance, “eucalyptus-cloud” is installed on the primary CLC and the secondary CLC.

5. Install the Imaging Worker image package on the machine hosting the primary CLC:

   ```
   yum install eucalyptus-service-image
   ```

CC-BY-SA, Eucalyptus Systems, Inc.

6. If you would like Load Balancer support in your cloud, you will need to install the Load Balancer image package on the machine hosting the primary CLC:

```
yum install eucalyptus-service-image
```

7. On each planned NC server, install the NC package:

```
yum install eucalyptus-nc
```

8. If you plan to run in *Edge networking mode*, install the package for Edge support on each planned NC host.

```
yum install eucanetd
```

Your installation is complete.

You are now ready to *Configure Eucalyptus*.

Configure Eucalyptus

This topic describes the parameters you need to set in order to launch Eucalyptus for the first time.

The first launch of Eucalyptus is different than a restart of a previously running Eucalyptus deployment in that it sets up the security mechanisms that will be used by the installation to ensure system integrity.

Eucalyptus configuration is stored in a text file, `/etc/eucalyptus/eucalyptus.conf`, that contains key-value pairs specifying various configuration parameters. Eucalyptus reads this file when it launches and when various forms of reset commands are sent it the Eucalyptus components.

Important: Perform the following tasks after you install Eucalyptus software, but before you start the Eucalyptus services.

Configure Network Modes

This section provides detailed configuration instructions for each of the four Eucalyptus networking modes. Eucalyptus requires network connectivity between its clients (end-users) and the cloud components (CC, CLC, and Walrus).

- In Edge mode, most networking configuration is handled through settings in a global Cloud Controller (CLC) property file. For more information, see *Configure for Edge Mode* (or *Configure for Edge Mode* for HA).
- In Managed and Managed (No VLAN) modes, traffic to instances pass through the CC. In these two modes clients must be able to connect to the Cluster Controller (CC).

The `/etc/eucalyptus/eucalyptus.conf` file contains all network-related options in the “Networking Configuration” section. These options use the prefix VNET_. The most commonly used VNET options are described in the following table. The set of networking settings that apply to a cloud varies based on its networking mode. Each setting in this section lists the modes in which it applies. Unless otherwise noted, all of these settings apply only to CCs.

The most commonly used VNET options are described in the following table.

CC-BY-SA, Eucalyptus Systems, Inc.

Option	Description	Modes

`VNET_ADDRESSPERNET`	This option controls how many VM instances can simultaneously be part of an individual user's security group. This option is set to a power of 2 (8, 16, 32, 64, etc.) but it should never be less than 8 and it cannot be larger than: (the total number of available IP addresses - 2).
	
	This option is used with VNET_NETMASK to determine how the IP addresses that are available to VMs are distributed among security groups. VMs within a single security group can communicate directly. Communication between VMs within a security group and clients or VMs in other security groups is controlled by a set of firewall rules. For example, setting
	
	`VNET_NETMASK="255.255.0.0"`
	`VNET_ADDRESSPERNET="32"`
	
	defines a netmask of 255.255.0.0 that uses 16 bits of the IP address to specify a network number. The remaining 16 bits specify valid IP addresses for that network meaning that 2^16 = 65536 IP addresses are assignable on the network. Setting `VNET_ADDRESSPERNET="32"` tells Eucalyptus that each security group can have at most 32 VMs in it (each VM getting its own IP address). Further, it stipulates that at most 2046 security groups can be active at the same time since 65536 / 32 = 2048. Eucalyptus reserves two security groups for its own use.
	
	In addition to subnets at Layer 3, Eucalyptus uses VLANs at Layer 2 in the networking stack to ensure isolation (Managed mode only).
		Managed, Managed (No VLAN)

`VNET_BRIDGE`	On an NC, this is the name of the bridge interface to which instances' network interfaces should attach. A physical interface that can reach the CC must be attached to this bridge. Common setting for KVM is `br0`.	Edge (on NC) Managed (No VLAN)
`VNET_DHCPDAEMON`	The ISC DHCP executable to use. This is set to a distro-dependent value by packaging. The internal default is `/usr/sbin/dhcpd3`.	Edge (on NC) Managed Managed (No VLAN)
`VNET_DHCPUSER`	The user the DHCP daemon runs as on your distribution. For CentOS 6 and RHEL 6, this is typically `root`. Default: `dhcpd`	Managed Managed (No VLAN)
`VNET_DNS`	The address of the DNS server to supply to instances in DHCP responses. Example: `VNET_DNS="173.205.188.129"`	Managed Managed (No VLAN)
`VNET_LOCALIP`	By default the CC automatically determines which IP address to use when setting up tunnels to other CCs. Set this to the IP address that other CCs can use to reach this CC if tunneling does not work.	Managed Managed (No VLAN)

CC-BY-SA, Eucalyptus Systems, Inc.

Option	Description	Modes
`VNET_MACPREFIX`	This option is used to specify a prefix for MAC addresses generated by Eucalyptus for VM instances. The prefix has to be in the form `HH:HH` where H is a hexadecimal digit. Example: `VNET_MACPREFIX="D0:D0"`	Managed, Managed (No VLAN)
`VNET_MODE`	The networking mode in which to run. The same mode must be specified on all CCs and NCs in your cloud. Valid values: `EDGE` `MANAGED`, `MANAGED-NOVLAN`,	All
`VNET_PRIVINTERFACE`	The name of the network interface that is on the same network as the NCs. In Managed and Managed (No VLAN) modes this must be a bridge for instances in different clusters but in the same security group to be able to reach one another with their private addresses. Default: `eth0`	Edge (on NC) Managed
`VNET_PUBINTERFACE`	**On a CC**, this is the name of the network interface that is connected to the "public" network. **On an NC**, this is the name of the network interface that is connected to the same network as the CC. Depending on the hypervisor's configuration this may be a bridge or a physical interface that is attached to the bridge. Default: `eth0`	Edge (on NC) Managed Managed (No VLAN)
`VNET_PUBLICIPS`	A space-separated list of individual and/or hyphenated ranges of public IP addresses to assign to instances. If you do not set a value for this option, all instances will receive only private IP addresses. Example: `VNET_PUBLICIPS= "173.205.188.140-173.205.188.254"` **Tip:** To offer more public IPs, you can span subnets. However you must list each subnet range separately. For example: `"10.133.82.50-10.133.82.254 10.133.83.0-10.133.83.254"`	Managed Managed (No VLAN)
`VNET_SUBNET, VNET_NETMASK`	These options control the internal private network used by instances within Eucalyptus. Eucalyptus assigns a distinct subnet of private IP addresses to each security group. This setting dictates how many addresses each of these subnets should contain. Specify a power of 2 between 16 and 2048. This is directly related, though not equal, to the number of instances that can reside in each security group. Eucalyptus reserves eleven addresses per security group.	Managed, Managed (No VLAN)

Configure for Edge Mode

To configure Eucalyptus for Edge mode, you must edit `eucalyptus.conf` on the Cluster Controller (CC) and Node Controller (NC) hosts. You must also create a JSON file and upload it the Cloud Controller (CLC).

CC-BY-SA, Eucalyptus Systems, Inc.

Configure for Managed Mode

In Managed mode, Eucalyptus manages the local network of VM instances and provides all networking features Eucalyptus currently supports, including VM network isolation, security groups, elastic IPs, and metadata service. Configure each CC to use an Ethernet device that lies within the same broadcast domain as all of its NCs.

Important: In Managed mode, each security group requires a separate subnet and a separate VLAN that Eucalyptus controls and maintains. So the underlying physical network must be "VLAN clean." For more information about VLAN clean, see *Prepare VLAN*.

To configure for Managed mode:

CLC Configuration

No network configuration required.

CC Configuration

Important:

We recommend allowing the CC to act as the gateway for NCs, in Managed mode. To do so, ensure that traffic from all NCs (on private network) is allowed to be masqueraded on the CC, and set the output interface to the public interface of the CC. You can do this using the following `iptables` command:

```
iptables -t nat -A POSTROUTING -s 10.101.104.0/16 -o em1 -j MASQUERADE
```

Where `10.101.104.0/16` is the private network containing all NCs, and `em1` is the public interface set on the CC.

1. Log in to the CC and open the `/etc/eucalyptus/eucalyptus.conf` file.
2. Go to the **Network Configuration** section, uncomment and set the following:

```
VNET_MODE="MANAGED"

VNET_SUBNET="<subnet for instances' private IPs. Example: 192.168.0.0>"
VNET_NETMASK="<your netmask for the vnet_subnet. Example: 255.255.0.0>"
VNET_DNS="<your DNS server's IP>"
VNET_ADDRSPERNET="<# of simultaneous instances per security group>"

VNET_PUBLICIPS="<your_free_public_ip1 your_free_public_ip2 …>"

VNET_LOCALIP="<the IP of the local interface on the cc that is reachable from
 CLC>"

VNET_DHCPDAEMON="<path to DHCP daemon binary. Example: /usr/sbin/dhcpd3>"

VNET_DHCPUSER="<DHCP user name. Example: dhcpd>"
```

3. If your NCs are not reachable from end-users directly and the CC has two (or more) Ethernet devices of which one connects to the client/public network and one connects to the NC network, or the single Ethernet device that the CC uses to connect to both clients and NCs is NOT 'eth0', then you must also uncomment and set:

```
VNET_PRIVINTERFACE="<Ethernet device on same network as NCs. Example: eth1>"

VNET_PUBINTERFACE="<Ethernet device on 'public' network. Example: eth0>"
```

4. Save the file.
5. Repeat on each CC in your system.

Important: Each CC must have the same configuration with the exception of the `VNET_LOCALIP` value, which should be machine-specific. In a multi-cluster configuration, you must set `VNET_PUBLICIPS` identically on all CCs.

CC-BY-SA, Eucalyptus Systems, Inc.

NC Configuration

Important:

We recommend allowing the CC to act as the gateway for NCs, in Managed mode. To do so, ensure that traffic from all NCs (on private network) is allowed to be masqueraded on the CC, and set the output interface to the public interface of the CC. You can do this using the following `iptables` command:

```
iptables -t nat -A POSTROUTING -s 10.101.104.0/16 -o em1 -j MASQUERADE
```

Where `10.101.104.0/16` is the private network containing all NCs, and `em1` is the public interface set on the CC.

1. Log into an NC machine and open the `/etc/eucalyptus/eucalyptus.conf` file.
2. Go to the **Network Configuration** section, uncomment and set the following:

```
VNET_MODE="MANAGED"
VNET_PUBINTERFACE="<Ethernet device/bridge reachable from cc machine. Example:
 eth0>"
```

3. Save the file.
4. Repeat on each NC.

Configure for Managed (No-VLAN) Mode

In Managed (No-VLAN) mode, Eucalyptus does not use VLANs to isolate the network bridges attached to VMs from each other. Configure each CC to use an Ethernet device that lies within the same broadcast domain as all of its NCs.

To configure for Managed (No VLAN) mode:

CLC Configuration

No network configuration required.

CC Configuration

Important: You must set `VNET_PUBLICIPS` identically on all CCs in a multi-cluster configuration.

1. Log in to the CC and open the `/etc/eucalyptus/eucalyptus.conf` file.
2. Go to the **Network Configuration** section, uncomment and set the following:

```
VNET_MODE="MANAGED-NOVLAN"
VNET_SUBNET="[Subnet for VMs private IPs. Example: 192.168.0.0]"
VNET_NETMASK="[Netmask for the vnet_subnet. Example: 255.255.0.0]"
VNET_DNS="[DNS server IP]"
VNET_ADDRSPERNET="[Number of simultaneous instances per security group]"
VNET_PUBLICIPS="[Free public IP 1] [Free public IP 2] ..."
VNET_LOCALIP="[IP address that other CCs can use to reach this CC]"
VNET_DHCPDAEMON="[Path to DHCP daemon binary. Example: /usr/sbin/dhcpd3]"
VNET_DHCPUSER='[DHCP user. Example: dhcpd]"
```

3. If your NCs are not reachable from end-users directly and the CC has two (or more) Ethernet devices of which one connects to the client/public network and one connects to the NC network, or the single Ethernet device that the CC uses to connect to both clients and NCs is NOT ‘eth0’, then you must also uncomment and set:

```
VNET_PRIVINTERFACE="[Ethernet device on same network as NCs. Example: eth1]"

VNET_PUBINTERFACE="[Ethernet device on ‘public’ network. Example: eth0]"
```

4. Save the file.
5. Repeat on each CC in your system.

Important: Each CC must have the same configuration with the exception of the `VNET_LOCALIP` value, which should be machine-specific.

CC-BY-SA, Eucalyptus Systems, Inc.

NC Configuration

1. Log into an NC machine and open the `/etc/eucalyptus/eucalyptus.conf` file.
2. Go to the **Network Configuration** section, uncomment and set the following:

```
VNET_MODE="MANAGED-NOVLAN"
VNET_BRIDGE="[bridge name. Example: br0]"
```

3. Save the file.
4. Repeat on each NC.

Create Scheduling Policy

This topic describes how to set up the Cluster Controller (CC) to choose which Node Controller (NC) to run each new instance.

1. In the CC, open the `/etc/eucalyptus/eucalyptus.conf` file.
2. In the `SCHEDPOLICY=` parameter, set the value to one of the following:

Option	Description
`GREEDY`	When the CC receives a new instance run request, it runs the instance on the first NC in an ordered list of NCs that has capacity to run the instance. At partial capacity with some amount of churn, this policy generally results in a steady state over time where some nodes are running many instances, and some nodes are running few or no instances.
`ROUNDROBIN`	(Default) When the CC receives a new instance run request, it runs the instance on the next NC in an ordered list of NCs that has capacity. The next NC is determined by the last NC to have received an instance. At partial capacity with some amount of churn, this policy generally results in a steady state over time where instances are more evenly distributed across the set of NCs.

3. Save the file.

Configure Loop Devices

In order to start new instances, Eucalyptus needs a sufficient number of loop devices to use for SC and NC components. An SC with insufficient loop devices fails to create new EBS volumes. An NC with insufficient loop devices fails to start new instances.

Eucalyptus installs with a default loop device amount of 256. If you want to change this number, perform the following steps. Otherwise, skip this section.

Tip: We recommend that you err on the side of configuring too many loop devices. Too many loop devices result in a minor amount of memory tie-up and some clutter added to the system's `/dev` directory. Too few loop devices make Eucalyptus unable to use all of a system's resources. We recommend a minimum of 50 loop devices. If you have fewer than 50, the startup script will complain.

1. Log in to the SC server and open the `/etc/eucalyptus/eucalyptus.conf` file.
2. Uncomment the following line:

```
# CREATE_SC_LOOP_DEVICES=256
```

3. Replace 256 with the number of loop devices.
4. Repeat for each SC on your system.
5. Log in to an NC server and open the `/etc/eucalyptus/eucalyptus.conf` file.
6. Uncomment the following line:

```
# CREATE_NC_LOOP_DEVICES=256
```

7. Replace 256 with the number of loop devices.
8. Repeat for each NC on your system.

CC-BY-SA, Eucalyptus Systems, Inc.

Configure Multi-Cluster Networking

Eucalyptus supports multiple clusters within a single Eucalyptus cloud. This topic briefly describes how Eucalyptus manages the networking aspect of a multi-cluster setup.

In Edge networking mode, Eucalyptus does not perform any special configuration for a multi-cluster setup. In Managed and Managed (No VLAN) modes, Eucalyptus sets up Layer 2 Tunneling Protocol (L2TP) between your clusters. This means that virtual machines in the same security group, but distributed across clusters (potentially each in their own broadcast domain), can communicate with one another. Eucalyptus uses the VTun package to handle all L2TP tunnels between clusters. If VTun is installed on each of your CCs, multi-cluster tunneling is automatically handled by each CC.

Depending on the networking mode and network topology, keep the following network configuration considerations in mind.

Managed Mode:	During normal operation, you will see many tunnel interfaces being created and destroyed as virtual networks are constructed and torn down.
Managed (No VLAN) Mode:	In order for VTun tunneling to work in this mode, you must configure each CC with a bridge as its primary, private interface (`VNET_PRIVINTERFACE`). All traffic from nodes in one cluster to nodes in another cluster is routed through the CCs. Each cluster requires that the interface that faces the nodes for the CC (the private interface) be a bridge device for the nodes themselves. You must set `VNET_PUBLICIPS` identically on all CCs in a multi-cluster configuration.
Managed Mode and Managed (No VLAN) Mode:	The CC attempts to auto-discover its list of local IP addresses upon startup, but if the IP that was used to register the CC is not locally available, you can override the CC's notion of 'self' by setting the `VNET_LOCALIP` variable in the `eucalyptus.conf` file.

Important: Note the following:

- You must set VNET_PUBLICIPS identically on all CCs in a multi-cluster configuration.
- To enable tunneling, set `DISABLE_TUNNELING=N` in `eucalyptus.conf` on both CC hosts.
- When L2TP tunneling is enabled in a multi-cluster setup, make sure that you are using different IP ranges for the nodes in each cluster.
- Do not run two CCs in the same broadcast domain with tunneling enabled, as this will potentially lead to a broadcast storm as tunnels start forwarding packets in a loop on your local network. Please disable tunneling by setting `DISABLE_TUNNELING=Y` in `eucalyptus.conf` on both CC hosts.

Configure the Firewall

This topic provides guidelines for restricting network access and managing iptables rules.

Restricting Network Access

This section provides basic guidance on setting up a firewall around your Eucalyptus components. It is not intended to be exhaustive.

On the Cloud Controller (CLC), Walrus, and Storage Controller (SC), allow for the following jGroups traffic:

- TCP connections between CLC, user-facing services (UFS), object storage gateway (OSG), Walrus, and SC on port 8779 (or the first available port in range 8779-8849)
- UDP connections between CLC, UFS, OSG, Walrus, and SC on port 7500
- Multicast connections between CLC, UFS, OSG, Walrus, and SC to IP 228.7.7.3 on UDP port 8773

On the UFS, allow the following connections:

- TCP connections from end-users and instances on ports 8773
- End-user and instance connections to DNS ports

On the CLC, allow the following connections:

CC-BY-SA, Eucalyptus Systems, Inc.

- TCP connections from UFS, CC and Eucalyptus instances (public IPs) on port 8773 (for metadata service)
- TCP connections from UFS, OSG, Walrus, and SC on port 8777

On the CC, make sure that all firewall rules are compatible with the dynamic changes performed by Eucalyptus, described in the section below. Also allow the following connections:

- TCP connections from CLC on port 8774

On OSG, allow the following connections:

- TCP connections from end-users and instances on port 8773
- TCP connections from SC and NC on port 8773

On Walrus, allow the following connections:

- TCP connections from OSG on port 8773

On the SC, allow the following connections:

- TCP connections from CLC and NC on TCP port 8773
- TCP connections from NC on TCP port 3260, if tgt (iSCSI open source target) is used for EBS storage

On the NC, allow the following connections:

- TCP connections from CC on port 8775
- TCP connections from other NCs on port 16514
- DHCP traffic forwarding to VMs
- Traffic forwarding to and from instances' private IP addresses

Managing iptables Rules for the CC

In Managed and Managed (No VLAN) modes, Eucalyptus flushes the CC's iptables rules for both `filter` and `nat`, then it sets the default policy for the `FORWARD` chain in `filter` to `DROP`. At run time, the CC adds and removes rules from `FORWARD` as users add and remove ingress rules from their active security groups. In addition, the `nat` table is configured to allow VMs access to the external network using IP masquerading, and dynamically adds/removes rules in the `nat` table as users assign and unassign public IPs to VMs at instance boot or run-time.

If you have rules you want to apply on the CC, make the following edit on the CC before you start Eucalyptus or while Eucalyptus is stopped:

```
iptables-save > /etc/eucalyptus/iptables-preload
```

Caution: Performing this operation to define special iptables rules that are loaded when Eucalyptus starts could cause Eucalyptus VM networking to fail. We recommend that you only do this if you are completely sure that it will not interfere with the operation of Eucalyptus.

Start Eucalyptus

Start the Eucalyptus components in the order presented in this section.

Make sure that each host you installed a Eucalyptus component on resolves to an IP address. Edit the `/etc/hosts` file if necessary.

Note: Eucalyptus 4.1.2 requires version 7 of the Java Virtual Machine. Make sure that your CLOUD_OPTS settings in the /etc/eucalyptus/eucalyptus.conf file either do not set `--java-home`, or that `--java-home` points to a version 7 JVM. This needs to happen before services are started.

Start the CLC

1. Log in to the Cloud Controller (CLC).
2. Enter the following command to initialize the CLC:

CC-BY-SA, Eucalyptus Systems, Inc.

Note: Make sure that the `eucalyptus-cloud` process is not running prior to executing this command.

```
/usr/sbin/euca_conf --initialize
```

Note: This command might take a minute or more to finish.

3. Enter the following command to start the CLC:

```
service eucalyptus-cloud start
```

Start Walrus

Important: Skip this step if you not using Walrus as your object storage backend, or if you installed Walrus on the same host as the CLC.

To start Walrus:

Log in to the Walrus server and enter the following command:

```
service eucalyptus-cloud start
```

Start the CC

To start the CC:

1. Log in to the CC server and enter the following:

```
service eucalyptus-cc start
```

2. If you have a multi-cluster setup, repeat this step on the CC in each cluster.

Start the SC

Important: If you installed SC on the same host as the CLC, skip this step.

To start the SC:

1. Log in to the SC server and enter the following command:

```
service eucalyptus-cloud start
```

Important: If you are re-installing the SC, please restart the tgt (iSCSI open source target) daemon.

2. If you have a multi-cluster setup, repeat this step on the SC in each cluster.

Start the NCs

To start the Node Controllers, perform the following tasks.

1. Log in to an NC server and enter the following command:

```
service eucalyptus-nc start
```

2. If you are running in Edge networking mode, start the Edge component.

```
service eucanetd start
```

3. Repeat for each NC server.

CC-BY-SA, Eucalyptus Systems, Inc.

Verify the Startup

At this point, all Eucalyptus components are enabled and starting up. Some of these services perform intensive initialization at start-up, particularly the first time they are started. You might have to wait a few minutes until they are fully operational.

One quick way to determine if the components are running is to run netstat on the various hosts and look to see when the service ports are allocated to a process. Specifically, the CLC, Walrus, and the SC allocate ports 8773. The CC listens to port 8774, and the NC uses port 8775.

Verify that everything has started without error. Expected outcomes include:

- The CLC is listening on ports 8443 and 8773
- Walrus is listening on port 8773
- The SC is listening on port 8773
- The CC is listening on port 8774
- The NCs are listening on port 8775
- Log files are being written to `/var/log/eucalyptus/`

Register Eucalyptus

After you start Eucalyptus for the first time, register the Eucalyptus components as described in this section.

Eucalyptus implements a secure protocol for registering separate components so that the overall system can't be tricked into including a component run by an unauthorized administrator or user. You only need to register components the first time Eucalyptus is started after it was installed.

Most registration commands run on the CLC server. NCs, however, are registered on each CC. You must register each NC on every CC for the cluster on which the NC participates.

Note that each registration command will attempt an SSH as root to the remote physical host where the registering component is assumed to be running. The registration command also contacts the component so it must be running at the time of the command is issued. If a password is required to allow SSH access, the command will prompt the user for it.

Except for NCs, each registration command requires four pieces of information:

- The **component** (`--register-XYZ`) you are registering, because this affects where the commands must be executed.
- The **partition** (`--partition`) the component will belong to. The partition is the same thing as availability zone in AWS.
- The **name** (`--component`) you assign to the component, up to 256 characters. This is the name used to identify the component in a human-friendly way. This name is also used when reporting system state changes which require administrator attention. This name must be globally-unique with respect to other component registrations. To ensure this uniqueness, we recommend using a combination of the component type (CLC, SC, CC, etc) and system hostname or IP address when you choose your component names. For example: `clc-eucahost15` or `clc-192.168.0.15`.
- The **IP address** (`--host`) of the service being registered. The host must be specified by IP address to function correctly.

  Note: You must specify public IP addresses.

NCs only have two pieces of information: component name and IP address.

Note: We recommend that you use IP addresses rather than host names when registering Eucalyptus components. If you do use hostnames, the underlying IP address may not be a site-local, any-cast, loopback, link-local, or multicast address.

Note: Once you've registered a Eucalyptus component with a host name, to avoid connectivity issues, do not change the host name's underlying IP address.

CC-BY-SA, Eucalyptus Systems, Inc.

Register Services

To register the user-facing services:

On the CLC server, enter the following command:

```
euca_conf --register-service -T user-api -H <host_url> -N <name>
```

For example:

```
euca_conf --register-service -T user-api -H 10.111.1.135 -N API_135
```

Register Walrus

Important: Skip this step if you are not using Walrus as your object storage backend.

To register Walrus:

On the CLC server, enter the following command:

```
/usr/sbin/euca_conf --register-walrusbackend --partition walrus --host
[walrus_IP_address] --component [walrus_name]
```

The `partition` name for Walrus has to be `walrus`. Like the CLC, the component name is a unique name for this particular component: we recommend a format such as `walrus-[hostname]`.

Register the CC

To register the CC:

1. On the CLC, enter the following command:

   ```
   /usr/sbin/euca_conf --register-cluster --partition [partition_name]
    --host [CC_IP_address] --component [cc_name]
   ```

   We recommend that you set the `partition` name to a descriptive name for the availability zone controlled by the CC. For example: `cluster01`.

   The `component` must be a unique name. We recommend that you use a short-hand name of the hostname or IP address of the machine, like `cc-[hostname]` or `cc-[IP address]`.

2. Repeat for each cluster, replacing the CC name, partition name, CC IP address, and CC name.

Register the SC

To register the SC:

1. On the CLC, enter the following command:

   Note: We recommend that you use IP addresses instead of DNS names when registering Eucalyptus components.

   ```
   /usr/sbin/euca_conf --register-sc --partition [partition_name] --host
   [SC_IP_address]
    --component [SC_name]
   ```

   An SC must have the same `partition` name as the CC in the same cluster. Like the other components, the `component` is a unique name for this particular component: we recommend a format such as `sc-[hostname]`.

CC-BY-SA, Eucalyptus Systems, Inc.

Warning: Newly registered SCs will be in the BROKEN state until they are explicitly configured to use a backend storage provider. The output of the registration for the first SC registered in a partition will look like:

```
SERVICE storage         PARTI00         SC71            BROKEN     37
 http://192.168.51.71:8773/services/Storage
arn:euca:eucalyptus:PARTI00:storage:SC71/
Registered the first storage controller in partition 'PARTI00'.  You must
 choose a storage back end with ``euca-modify-property -p
PARTI00.storage.blockstoragemanager=$BACKEND''
```

This is completely normal and simply indicates that further action must be taken to configure the SC before it will become fully functional. For information about configuring the SC, see *Configure the Runtime Environment->Configure the Storage Controller*

2. Repeat for each cluster, replacing the SC name, partition name, SC IP address, and SC name.

Register the NCs

Important: If you are using host names rather than IP addresses when registering your NCs, ensure that DNS is working properly, or populate /etc/hosts for all nodes in a cluster.

1. On a CC, register all NCs using the following command with the IP address of each NC server:

```
/usr/sbin/euca_conf --register-nodes "[node0_IP_address] ... [nodeN_IP_address]"
```

2. Repeat each cluster in your cloud.

The IP addresses of the NCs are space delimited, as in the following example:

```
/usr/sbin/euca_conf --register-nodes "192.168.71.154 192.168.71.155
 192.168.71.159"
```

For HA: For HA, you must also register the NCs with the secondary CC.

Register Arbitrators

Eucalyptus uses a periodic ICMP echo test to an Arbitrator. This test approximates an end user's ability to access the system. If Eucalyptus determines that it cannot reach the host associated with a registered Arbitrator, all Eucalyptus services operating on that host attempt to failover to the alternate hosts running those services.

For HA: In HA, you can register each Arbitrator service on the primary and secondary CLC and Walrus. If you are using either Managed or Managed (No VLAN) mode, you can also register Arbitrator services on both the primary CC and the secondary CC.

We recommend that you register more than one Arbitrator for each Eucalyptus component. This will allow for normal outages and maintenance. There is no limit on the number of Arbitrators on a CLC and a Walrus. You can only register up to three on a CC.

Register an Arbitrator service on each host that has a cloud component (CLC or Walrus) installed. An Arbitrator is a host-wide component: when an Arbitrator is registered on a host, it is registered with all cloud components enabled on that host. A separate arbitrator has to be registered per each network entity that needs to be monitored from the host.

To register an Arbitrator:

1. Log in to the primary CLC.
2. Enter the following command to register an arbitrator:

CC-BY-SA, Eucalyptus Systems, Inc.

```
/usr/sbin/euca_conf --register-arbitrator --partition [ID]
 --component [ID] --host [target_IP]>
```

where:

- `[ID]` is a globally unique ID that identifies an Arbitrator. Note that you must use the same [ID] as both a partition and component ID.
- `[target_IP]` is the IP of the machine running the Eucalyptus component that will run the Arbitrator.

For example:

```
euca_conf --register-arbitrator --partition EXAMPLE_ARB --component EXAMPLE_ARB
 --host 192.168.1.10
```

3. Repeat for the secondary CLC and for both Walrus servers.
4. Define the gateway for each Arbitrator:

```
/usr/sbin/euca-modify-property -p <ID>.arbitrator.gatewayhost=<gateway>
```

where:

- `<ID>` is the globally unique ID of the registered Arbitrator.
- `<gateway>` is an external hostname or IP address used to approximate connectivity to the end user.

For example:

```
euca-modify-property -p EXAMPLE_ARB.arbitrator.gatewayhost=192.168.1.1
```

5. Repeat for each registered Arbitrator.
6. To register on each CC, log in to the primary CC, and open the `/etc/eucalyptus/eucalyptus.conf` file.
7. Provide a list of Arbitrators (up to three) as values for the `CC_ARBITRATORS` property. For example:

```
CC_ARBITRATORS="192.168.48.11 192.168.48.12"
```

8. Save the file and restart the CC.

```
service eucalyptus-cc restart
```

9. Repeat on the secondary CC.

In the following example, the primary CLC is on <CLC_host_p>, the secondary CLC is on <CLC_host_s>, the primary Walrus is on <Walrus_host_p>, and the secondary Walrus is on <Walrus_host_s>.

```
/usr/sbin/euca_conf --register-arbitrator --host <CLC_host_p>
--component ARB00 --partition ARB00
/usr/sbin/euca_conf --register-arbitrator --host <CLC_host_p>
--component ARB01 --partition ARB01
/usr/sbin/euca_conf --register-arbitrator --host <Walrus_host_p>
--component ARB02 --partition ARB02
/usr/sbin/euca_conf --register-arbitrator --host <Walrus_host_p>
--component ARB03 --partition ARB03
/usr/sbin/euca_conf --register-arbitrator --host <CLC_host_s>
--component ARB04 --partition ARB04
/usr/sbin/euca_conf --register-arbitrator --host <CLC_host_s>
--component ARB05 --partition ARB05
/usr/sbin/euca_conf --register-arbitrator --host <Walrus_host_s>
--component ARB06 --partition ARB06
/usr/sbin/euca_conf --register-arbitrator --host <Walrus_host_s>
--component ARB07 --partition ARB07
```

CC-BY-SA, Eucalyptus Systems, Inc.

Configure the Runtime Environment

After Eucalyptus is installed and registered, perform the tasks in this section to configure the runtime environment.

Generate Administrator Credentials

Now that you have installed Eucalyptus, you're ready to start configuring and using it. To do so, you must generate credentials.

Important: When you run the `euca_conf --get-credentials` command, you are requesting the access and secret keys and an X.509 certificate and key. You cannot retrieve an existing X.509 certificate and key. You can only generate a new pair.

To generate a set of credentials:

1. Generate administrator credentials.

```
/usr/sbin/euca_conf --get-credentials admin.zip

unzip admin.zip
```

2. Source the `eucarc` file.

```
source eucarc
```

You are now able to run Eucalyptus commands.

Upload the JSON Configuration File (Edge Mode Networking)

If you've selected edge mode networking, you will need to upload the JSON configuration file you configured earlier in the installation process.

To upload the JSON file with the Edge mode networking configuration:

Important: This step can only be run after getting your credentials in *Generate Administrator Credentials*.

Run the following command to upload the configuration file to the CLC (with valid eucalyptus/admin credentials sourced):

```
euca-modify-property -f
cloud.network.network_configuration=</path/to/your/json_config_file>
```

Configure Object Storage Gateway

This topic details how to configure OSG for either Walrus or Riak.

You must configure an object storage gateway (OSG) before it will function properly. OSGs that have been registered but not properly configured will be listed in the `BROKEN` state when listed with the euca-describe-services command. For example:

```
SERVICE objectstorage objectstorage osg-192.168.1.16 BROKEN 23
http://192.168.1.16:8773/services/objectstorage
arn:euca:bootstrap:objectstorage:objectstorage:osg-192.168.1.16/
```

Configure OSG for Walrus

1. Enter `walrus` as the storage provider using the `euca-modify-property` command.

```
euca-modify-property -p objectstorage.providerclient=walrus
```

2. Check that the OSG is enabled.

```
euca-describe-services
```

CC-BY-SA, Eucalyptus Systems, Inc.

If the state appears as `DISABLED` or `BROKEN`, check the cloud-*.log files in the `/var/log/eucalyptus` directory. A `DISABLED` state generally indicates that there is a problem with your network or credentials.

Configure OSG for Riak

1. Enter `riakcs` as the storage provider using the `euca-modify-property` command.

   ```
   euca-modify-property -p objectstorage.providerclient=riakcs
   ```

2. Specify the RiakCS/S3 endpoint that you want to use with Eucalyptus. For example:

   ```
   euca-modify-property -p
   objectstorage.s3provider.s3endpoint=riakcs-01.riakcs-cluster.myorg.com
   ```

3. Provide your RiakCS credentials to access your RiakCS installation:

   ```
   euca-modify-property -p objectstorage.s3provider.s3accesskey=<RiakCS access
   key>
   euca-modify-property -p objectstorage.s3provider.s3secretkey=<RiakCS secret
   key>
   ```

4. After successful configuration, check to ensure that the state of the OSG is `ENABLED` by running the `euca-describe-services` command. For example:

   ```
   SERVICE objectstorage objectstorage osg-192.168.1.16 ENABLED 23
   http://192.168.1.16:8773/services/objectstorage
   arn:euca:bootstrap:objectstorage:objectstorage:osg-192.168.1.16/
   ```

   If the state appears as `DISABLED` or `BROKEN`, check the cloud-*.log files in the `/var/log/eucalyptus` directory. A `DISABLED` state generally indicates that there is a problem with your network or credentials.

Configure the Storage Controller

Eucalyptus offers SAN support for Eucalyptus block storage (EBS). Eucalyptus directs the Storage Controller (SC) to manage any supported SAN devices or third-party clients. Eucalyptus automatically creates and tears down volumes, snapshots, and data connections from guest instances. The administrator does not need to pre-allocate volumes or LUNs for Eucalyptus.

Eucalyptus currently offers several backend providers for the Storage Controller:

- Overlay
- DAS
- Equallogic
- Netapp
- EMC-VNX
- Ceph-RBD

You must configure the SC to use one of these options. The SC automatically goes to the `BROKEN` state after being registered with the CLC and will remain in that state until you explicitly configure the SC by telling it which backend storage provider to use.

You can check the state of a storage controller by running `euca-describe-services -E` and note the state and status message of the SC(s). The output for an unconfigured SC will look like:

```
SERVICE storage         PARTI00         SC71            BROKEN     37
http://192.168.51.71:8773/services/Storage
arn:euca:eucalyptus:PARTI00:storage:SC71/
SERVICEEVENT 6c1f7a0a-21c9-496c-bb79-23ddd5749222
arn:euca:eucalyptus:PARTI00:storage:SC71/
SERVICEEVENT 6c1f7a0a-21c9-496c-bb79-23ddd5749222 ERROR
SERVICEEVENT 6c1f7a0a-21c9-496c-bb79-23ddd5749222 Sun Nov 18 22:11:13 PST 2012
SERVICEEVENT 6c1f7a0a-21c9-496c-bb79-23ddd5749222 SC blockstorageamanger not
configured. Found empty or unset manager(unset). Legal values are: das,overlay
```

CC-BY-SA, Eucalyptus Systems, Inc.

Note the error above: `SC blockstoragemanager not configured. Found empty or unset manager(unset). Legal values are: das,overlay`.

This indicates that the SC is not yet configured. It can be configured by setting the `[partition].storage.blockstoragemanager` property to either 'das' or 'overlay'.

If you have installed the Eucalyptus Enterprise packages for your EBS adapter, you will also see additional options in the output line above, and can set the block storage manager to 'netapp','emc-vnx-flare31','emc-vnx', 'equallogic', or 'ceph-rbd' as appropriate.

You can verify that the SC blockstoragemanager is unset using:

```
euca-describe-properties | grep blockstorage
```

To configure SAN support, follow the steps for your desired backend storage device: *Open-Source ISCSI Filesystem-backed*, *Dell Equallogic*, *JBOD*, *Netapp*, *EMC VNX*, or *Enabling Ceph-RBD*.

Configuring the SC to use the local filesystem (Overlay)

This was the default configuration option for the SC in pre-3.2 Eucalyptus. In this configuration the SC itself hosts the volume and snapshots for EBS and stores them as files on the local filesystem. It uses standard linux iSCSI tools to serve the volumes to instances running on NCs.

1. Configure the SC to use the OverlayManager for storage.

   ```
   euca-modify-property -p <partition>.storage.blockstoragemanager=overlay
   ```

   The output of the command should be similar to:

   ```
   PROPERTY PARTI00.storage.blockstoragemanager overlay was <unset>
   ```

2. Verify that the property value is now: 'overlay'

   ```
   euca-describe-properties | grep blockstorage
   ```

Enable Dell Equallogic SANs

1. Configure the SC to use the EquallogicManager for storage.

   ```
   euca-modify-property -p <partition>.storage.blockstoragemanager=equallogic
   ```

   The output of the command should be similar to:

   ```
   PROPERTY PARTI00.storage.blockstoragemanager equallogic was <unset>
   ```

2. Verify that the property value is now: 'equallogic'

   ```
   euca-describe-properties | grep blockstorage
   ```

3. On the CLC, run the following command to verify that the SC is listed; note that it may be in the `BROKEN` state:

   ```
   euca_conf --list-scs
   ```

4. On the CLC, enable SAN support in Eucalyptus by entering your SAN's hostname or IP address, the username, password, and the name of the chap user:

   ```
   euca-modify-property -p [partition_name].storage.sanhost=[SAN_IP_address]
   euca-modify-property -p [partition_name].storage.sanuser=[SAN_admin_user_name]
   euca-modify-property -p
   [partition_name].storage.sanpassword=[SAN_admin_password]
   euca-modify-property -p <partition>.storage.chapuser=[chap_username]
   ```

   If you have multiple management IP addresses for the SAN adapter, provide a comma-delimited list of IP addresses to the `[partition_name].storage.sanhost` property.

5. Set the `<partition>.storage.poolname` property to the name of a storage pool on the Equallogic SAN to be used by Eucalyptus. For example:

CC-BY-SA, Eucalyptus Systems, Inc.

```
euca-modify-property -p
[partition_name].storage.poolname=equallogic-storage-pool
```

Your Equallogic SAN is now ready to use with Eucalyptus.

Enable Direct Attached Storage (JBOD) SANs

Important: Direct Attached Storage still requires that `/var/lib/eucalyptus/volumes` has enough space for locally cached snapshots.

1. Configure the SC to use the (Direct Attached Storage) DASManager for storage.

   ```
   euca-modify-property -p <partition>.storage.blockstoragemanager=das
   ```

   The output of the command should be similar to:

   ```
   PROPERTY PARTI00.storage.blockstoragemanager das was <unset>
   ```

2. Verify that the property value is now: 'das'

   ```
   euca-describe-properties | grep blockstorage
   ```

3. On the CLC, run the following command to verify that the SC is listed; note that it may be in the `BROKEN` state:

   ```
   euca_conf --list-scs
   ```

4. On the CLC, set the DAS device name property. The device name can be either a raw device (/dev/sdX, for example), or the name of an existing Linux LVM volume group.

   ```
   euca-modify-property -p <cluster name>.storage.dasdevice=<device name>
   ```

   For example:

   ```
   euca-modify-property -p cluster0.storage.dasdevice=/dev/sdb
   ```

Your SAN is now ready to use with Eucalyptus.

Enable NetApp SANs

Eucalyptus supports both NetApp Clustered ONTAP and traditional 7-mode SANs. NetApp Vservers and 7-mode Filers are managed by Eucalyptus using NetApp Manageability Software Development Kit (NMSDK) and Data ONTAP APIs. This section covers enabling both NetApp Clustered ONTAP and traditional 7-mode SANs.

Enable NetApp 7-mode SANs

To configure NetApp 7-mode Filer and enable the SAN in Eucalyptus:

1. Verify Data ONTAP version for the 7-mode Filer is 7.3.3 or later.
2. Verify SSL access by typing `secureadmin status`
3. If SSL is marked inactive, enable with `secureadmin setup ssl` and generate a new certificate.
4. Turn on SSL access with `options httpd.admin.ssl.enable on`
5. Enable the iSCSI service on the NetApp device with `option iscsi.enable on` or `option licensed_feature.iscsi.enable on` if you have an embedded license on your array.
6. Turn on the iSCSI service with `iscsi start`
7. Enable the iSCSI service on the NetApp device with `enable iscsi service`
8. Verify that an aggregate with sufficient spare capacity exists.
   - If you have SSH access to the NetApp Filer, enter `aggr show_space`.
   - If an aggregate with spare capacity does not exist, create one using the `aggr create` command.
9. Verify that you have a license for FlexClone installed. At the shell prompt, enter `license` to see the list of all installed licenses.

CC-BY-SA, Eucalyptus Systems, Inc.

10. Verify that administrator account credentials for NetApp Filer are available to be configured in Eucalyptus. If not, create a new administrator account for use by Eucalyptus
11. Configure the SC to use the NetappManager for storage.

```
euca-modify-property -p <partition>.storage.blockstoragemanager=netapp
```

The output of the command should be similar to:

```
PROPERTY <partition>.storage.blockstoragemanager netapp was <unset>
```

12. Verify that the property value is now: 'netapp'

```
euca-describe-properties | grep blockstorage
```

13. On the CLC, run the following command to verify that the SC is listed; note that it may be in the `BROKEN` state:

```
euca_conf --list-scs
```

14. Wait for the SC to transition to the NOTREADY or DISABLED state.
15. On the CLC, enable NetApp SAN support in Eucalyptus by entering the Filer's hostname or IP address, the username and password of the administrator account, and CHAP username.

Note: Eucalyptus uses Challenge Handshake Authentication Protocol (CHAP) for disk operations. The CHAP username can be any value, however it should be unique when sharing a NetApp Filer across multiple Eucalyptus clusters.

Note: CHAP support for NetApp has been added in Eucalyptus 3.3. An SC will not transition to ENABLED state until the CHAP username is configured.

```
euca-modify-property -p <partition>.storage.sanhost=<Filer_IP_address>
euca-modify-property -p <partition>.storage.sanuser=<Filer_admin_username>
euca-modify-property -p <partition>.storage.sanpassword=<Filer_admin_password>
euca-modify-property -p <partition>.storage.chapuser=<Chap_username>
```

16. Wait for the SC to transition to the ENABLED state.

Note: The SC must be in the ENABLED state before configuring the following properties.

17. If no aggregate is set, Eucalyptus will query the NetApp Filer for all available aggregates and use the one that has the highest capacity (free space) by default. To make Eucalyptus use specific aggregate(s) configure the following property:

```
euca-modify-property -p
<partition>.storage.aggregate=<aggregate_1_name,aggregate_2_name,...>
```

If you want Eucalyptus to use the smallest aggregate first configure the following property:

```
euca-modify-property -p <partition>.storage.uselargestaggregate=false
```

18. Set the iSCSI data IP on the ENABLED CLC. This IP is used by NCs to perform disk operations on the Filer.

Note: Filer IP address can be used as the data port IP. If this is not set, Eucalyptus will automatically use the Filer IP address/hostname.

Note: Eucalyptus does not support Multipath I/O for NetApp 7-mode Filers.

```
euca-modify-property -p <partition>.storage.ncpaths=<ip>
```

19. Set the iSCSI data IP on the ENABLED CLC. This IP is used by the SC to perform disk operations on the Filer. The SC connects to the Filer in order to transfer snapshots to Walrus during snapshot operations.

CC-BY-SA, Eucalyptus Systems, Inc.

Note: The Filer IP address can be used as the data port IP. If this is not set, Eucalyptus will automatically use the Filer IP address/hostname.

Note: Eucalyptus does not support Multipath I/O for NetApp 7-mode Filers.

```
euca-modify-property -p <partition>.storage.scpaths=<ip>
```

Your Netapp 7-mode SAN is now ready to use with Eucalyptus.

Enable NetApp Clustered Data ONTAP SAN

Eucalyptus integrates with NetApp Clustered ONTAP SAN by operating against a Vserver. SC must be configured to operate against Vserver contained in the NetApp Clustered ONTAP environment.

For more information on NetApp Clustered Data ONTAP, see *Clustered Data ONTAP 8.1 and 8.1.1: An Introduction*.

To configure NetApp Vserver and enable the SAN in Eucalyptus:

1. Verify Clustered Data ONTAP version for the SAN is 8.1.1 or later.
2. Verify that FlexClone and iSCSI licenses are installed on the SAN.
3. Verify that a Vserver with iSCSI data protocol is available for use by Eucalyptus.
4. Verify that Vserver administration is delegated to a user with administrative privileges for that Vserver. If not, create a new new Vserver administrator account for use by Eucalyptus.
5. Verify that a management (only) Logical Interface (LIF) is configured for the Vserver and an IP address or hostname is assigned to it.
6. Verify that data LIFs are configured on the Vserver.
7. Verify that one or more aggregates with sufficient spare capacity exists.
8. Verify the network connectivity between Eucalyptus components and the Vserver. The SC must be able communicate with the Vserver over both management and data LIFs. The NC must be able to communicate with the Vserver using the data LIFs.
9. Configure the SC to use the NetApp SAN for storage:

   ```
   euca-modify-property -p <partition>.storage.blockstoragemanager=netapp
   ```

   The output of the command should be similar to:

   ```
   PROPERTY <partition>.storage.blockstoragemanager netapp was <unset>
   ```

10. Verify that the property value is now: 'netapp'

    ```
    euca-describe-properties | grep blockstorage
    ```

11. On the CLC, run the following command to verify that the SC is listed; note that it may be in the BROKEN state:

    ```
    euca_conf --list-scs
    ```

12. Wait for the SC to transition to NOTREADY or DISABLED states.
13. On the CLC, enable NetApp SAN support in Eucalyptus by entering the Vserver's hostname or IP address, the username and password of the administrator account, CHAP username and Vserver name.

    Note: Eucalyptus uses Challenge Handshake Authentication Protocol (CHAP) for disk operations. The CHAP username can be any value, however it should be unique when sharing a NetApp Vserver across multiple Eucalyptus clusters.

    Note: CHAP support for NetApp has been added in Eucalyptus 3.3. The SC will not transition to ENABLED state until the CHAP username is configured.

    ```
    euca-modify-property -p <partition>.storage.sanhost=<Vserver_IP_address>
    euca-modify-property -p <partition>.storage.sanuser=<Vserver_admin_username>
    ```

CC-BY-SA, Eucalyptus Systems, Inc.

```
euca-modify-property -p <partition>.storage.sanpassword=<Vserver_admin_password>
euca-modify-property -p <partition>.storage.chapuser=<Chap_username>
```

Note: The following command may fail if tried immediately after configuring the block storage manager. Retry the command a few times, pausing for a few seconds after each retry:

```
euca-modify-property -p <partition>.storage.vservername=<Vserver_name>
```

14. Wait for the SC to transition to ENABLED state.

Note: The SC must be in the ENABLED state before configuring the following properties.

15. If no aggregate is set, Eucalyptus will query the NetApp Vserver for all available aggregates and use the one that has the highest capacity (free space) by default. To make Eucalyptus use specific aggregate(s) configure the following property:

```
euca-modify-property -p <partition>.storage.aggregate=<aggregate_1_name,
      aggregate_2_name,...>
```

If you want Eucalyptus to use the smallest aggregate first configure the following property:

```
euca-modify-property -p <partition>.storage.uselargestaggregate=false
```

16. Set an IP address for the iSCSI data LIF on the ENABLED CLC. This is used for NCs performing disk operations on the Vserver. If you want to configure multiple IPs, see see *Configure NetApp Multipathing*.

```
euca-modify-property -p <partition>.storage.ncpaths=<ip>
```

17. Set an IP address for the iSCSI data LIF on the ENABLED CLC. This is used by the SC for performing disk operations on the Vserver. The SC connects to the data LIFs on the Vserver in order to transfer snapshots to Walrus during snapshot operations. If you want to configure multiple IPs, see *Configure NetApp Multipathing*.

```
euca-modify-property -p <partition>.storage.scpaths=<ip>
```

Your NetApp Clustered Data ONTAP SAN is now ready to use with Eucalyptus.

Enable EMC VNX SANs

This adapter uses the newer VNX-Snapshot feature available on VNX devices running FLARE v5.32 or later that have a VNX-Snapshot license. This adapter also requires the Navisphere Secure CLI to be installed on the SCs. The Navisphere CLI must be version 7.32.0.5.54 or later.

Important: You must create a Clone Private LUN (CPL) of at least 1GB on each SP. For more information on creating private LUNs, see the VNX documentation.

1. We assume that the Navisphere CLI is installed in `/opt/Navisphere` on the SC.

Important: Eucalyptus currently supports version 7.32.0.5.54 or later of the Navisphere CLI.

2. Verify that the CLI is installed and can communicate with the VNX from the SCs.

On each SC that you are configuring, test the naviseccli command as follows:

```
/opt/Navisphere/bin/naviseccli -User <your SAN username> -Password <your SAN
password> -Scope 0 -Address <management port IP> connection -pingnode -address
 <a data port IP on your VNX>
```

Verify that the command runs successfully and the ping gets replies from the SAN.

3. On the CLC, run the following command to verify that the SC is listed; note that it may be in the `BROKEN` state:

```
euca_conf --list-scs
```

CC-BY-SA, Eucalyptus Systems, Inc.

4. Configure the SC to use the EMC VNX VNX-Snapshot-based manager for storage.

```
euca-modify-property -p <partition>.storage.blockstoragemanager=emc-vnx
```

   The output of the command should be similar to:

```
PROPERTY PARTI00.storage.blockstoragemanager emc-vnx was <unset>
```

5. Check the SC to be sure that it has transitioned out of the BROKEN state and is in either NOTREADY or DISABLED before configuring the rest of the properties for the SC. The following commands should be run on the ENABLED CLC to configure the SC.

   On the ENABLED CLC, run:

```
euca_conf --list-scs
```

6. On the CLC, enable SAN support in Eucalyptus by entering your SAN's hostname or IP address, the username, and password:

```
euca-modify-property -p [partition_name].storage.sanhost=[SAN_IP_address]
euca-modify-property -p [partition_name].storage.sanuser=[SAN_admin_user_name]
euca-modify-property -p
[partition_name].storage.sanpassword=[SAN_admin_password]
```

   If you have multiple management IP addresses for the SAN adapter, provide a comma-delimited list of IP addresses to the `[partition_name].storage.sanhost` property.

7. On the ENABLED CLC, set the login scope for the command line access. For most installs, the login scope will be `0`, which indicates a global login scope for the device. `1` indicates a local scope. `2` indicates LDAP authentication for the SAN device. Use login scope value of `2` only if your SAN is configured to use LDAP authentication and you have an admin user configured to use LDAP.

```
euca-modify-property -p <partition_name>.storage.loginscope=<login_scope>
```

8. On the ENABLED CLC, set the username for the Challenge Handshake Authentication Protocol (CHAP). This can be any value, however it should be unique when sharing VNX on multiple Eucalyptus clusters.

```
euca-modify-property -p <partition_name>.storage.chapuser=<chap_username>
```

9. On the ENABLED CLC, set the value for the unique storage pool that you have configured to use with the SC.

```
euca-modify-property -p <partition_name>.storage.storagepool=0
```

10. On the ENABLED CLC, set the iSCSI data port IP for NCs to use to perform disk operations on the SAN. If you want to configure multiple IPs, see *Configure EMC VNX Multipathing*.

```
euca-modify-property -p <partition_name>.storage.ncpaths=<ip>
```

11. On the ENABLED CLC, set the iSCSI data port IP for SCs to use to perform disk operations on the SAN. The SCs connect to the data ports on the SAN in order to transfer snapshots to Walrus during snapshot operations. If you want to configure multiple IPs, see the section on 'multipathing.

```
euca-modify-property -p <partition_name>.storage.scpaths=<ip>
```

12. On the ENABLED CLC, set the path to Navisphere CLI that you downloaded earlier to the SC. The following example shows the default path. This is that path on the SC, not on the CLC.

```
euca-modify-property -p
<partition_name>.storage.clipath=/opt/Navisphere/bin/naviseccli
```

Your EMC VNX SAN is now ready to use with Eucalyptus.

Tip: Note: The time it takes for a LUN migration to complete will depend on the exact VNX model, workload, and volume size, and the amount of data actually stored in the volume. The default timeout for LUN migrations is 12 hours. If your deployment uses volumes >50GB, or if you find that snapshots fail and a "migration timeout"

CC-BY-SA, Eucalyptus Systems, Inc.

message is seen in the SC logs, then you should increase the timeout to a larger value. It is recommended that if you plan on using volumes in the 100GB range that you set that timeout to 3600 or larger. You can set the timeout using `euca-modify-property` as follows:

```
euca-modify-property -p [partition].storage.lunmigrationtimeout=[time in
hours]
```

Enabling Ceph-RBD

This task assumes the following:

- You have an installed and working Ceph cluster.
- Ceph user credentials assigned to Eucalyptus SC and NCs are configured with appropriate permissions (different user credentials can be used for the SCs and NCs, below sections will explain how to configure them in Eucalyptus).

1. Node Controllers (NCs) are designed to communicate with Ceph cluster via libvirt. This interaction requires a hypervisor that supports Ceph-RBD. The following instructions will walk you through steps for verifying and or installing the hypervisor. Repeat this process for every NC in the Eucalyptus cluster.

a) Verify if `qemu-kvm` and `qemu-img` are already installed.

```
yum list qemu-kvm qemu-img
```

Proceed to the preparing the RHEV qemu packages step if they are not installed.

b) Verify qemu support for the `ceph-rbd` driver.

```
qemu-img --help
qemu-img version 0.12.1, Copyright (c) 2004-2008 Fabrice Bellard
...
Supported formats: raw cow qcow vdi vmdk cloop dmg bochs vpc vvfat qcow2
qed vhdx parallels nbd blkdebug host_cdrom
host_floppy host_device file gluster gluster gluster gluster rbd
```

Note: If 'rbd' is listed as one of the supported formats, no further action is required; otherwise proceed to the next step.

c) If the `eucalyptus-nc` service is running, terminate/stop all instances. After all instances are terminated, stop the eucalyptus-nc service.

```
service eucalyptus-nc stop
```

d) Uninstall the `qemu-kvm` and `qemu-img` packages.

Note: This will also uninstall `libvirt` if it was installed. Make a copy of the current `/etc/libvirt/libvirtd.conf` file. The yum uninstallation process will save the current copy to `/etc/libvirt/libvirtd.conf.rpmsave`.

```
yum remove qemu-kvm qemu-img
```

e) Prepare the RHEV qemu packages:

- If this NC is a RHEL system and the RHEV subscription to qemu packages is available, consult the RHEV package procedure to install the qemu-kvm-rhev and qemu-img-rhev packages. Blacklist the RHEV packages in the Eucalyptus repo to ensure that packages from the RHEV repo are installed. Skip to step **g** after installing the RHEV packages.
- If this NC is a RHEL system and RHEV subscription to qemu packages is unavailable, eucalyptus built and maintained qemu-rhev packages may be used. These packages are available in the same yum repository as other eucalyptus packages. Note that using eucalyptus built RHEV packages voids the original RHEL support for the qemu packages.
- If this NC is a non-RHEL (CentoOS) system, Eucalyptus-built and maintained qemu-rhev packages may be used. These packages are available in the same yum repository as other eucalyptus packages.

CC-BY-SA, Eucalyptus Systems, Inc.

f) Install Eucalyptus-built RHEV packages: `qemu-kvm-rhev` and `qemu-img-rhev`.

```
yum install qemu-kvm-rhev qemu-img-rhev
```

g) Install `libvirt`.

```
yum install libvirt
```

h) If libvirt was uninstalled in step **d**, copy the saved `libvirtd.conf` file to `/etc/libvirt/libvirtd.conf`. Otherwise modify the `/etc/libvirt/libvirtd.conf` file as per your requirements.

```
yum install libvirt
```

i) Start the `libvirtd` service.

```
service libvirtd start
```

j) Verify `qemu` support for the `ceph-rbd` driver.

```
qemu-img --help
qemu-img version 0.12.1, Copyright (c) 2004-2008 Fabrice Bellard
...
Supported formats: raw cow qcow vdi vmdk cloop dmg bochs vpc vvfat qcow2
qed vhdx parallels nbd blkdebug host_cdrom
host_floppy host_device file gluster gluster gluster gluster rbd
```

k) Make sure the eucalyptus-nc service is started.

```
service eucalyptus-nc start
```

2. Configure the SC to use Ceph-RBD for block storage.

```
euca-modify-property -p <partition>.storage.blockstoragemanager=ceph-rbd
```

The output of the command should be similar to:

```
PROPERTY PARTI00.storage.blockstoragemanager ceph-rbd was <unset>
```

3. Verify that the propery value is now `ceph-rbd`:

```
euca-describe-properties | grep blockstorage
```

4. Check the SC to be sure that it has transitioned out of the `BROKEN` state and is in the `NOTREADY`, `DISABLED` or `ENABLED` state before configuring the rest of the properties for the SC.

5. The ceph-rbd provider will assume defaults for the following properties

```
euca-describe-properties -v | grep 'PARTI00.storage.ceph'

PROPERTY        PARTI00.storage.cephconfigfile  /etc/ceph/ceph.conf
DESCRIPTION     PARTI00.storage.cephconfigfile  Absolute path to Ceph
configuration (ceph.conf) file. Default value is '/etc/ceph/ceph.conf'

PROPERTY        PARTI00.storage.cephkeyringfile
/etc/ceph/ceph.client.eucalyptus.keyring
DESCRIPTION     PARTI00.storage.cephkeyringfile Absolute path to Ceph keyring
 (ceph.client.eucalyptus.keyring) file. Default value is
'/etc/ceph/ceph.client.eucalyptus.keyring'

PROPERTY        PARTI00.storage.cephsnapshotpools       rbd
DESCRIPTION     PARTI00.storage.cephsnapshotpools       Ceph storage pool(s)
made available to Eucalyptus for EBS snapshots. Use a comma separated list for
 configuring multiple pools. Default value is 'rbd'

PROPERTY        PARTI00.storage.cephuser        eucalyptus
DESCRIPTION     PARTI00.storage.cephuser        Ceph username employed by
Eucalyptus operations. Default value is 'eucalyptus'
```

CC-BY-SA, Eucalyptus Systems, Inc.

```
PROPERTY        PARTI00.storage.cephvolumepools rbd
DESCRIPTION     PARTI00.storage.cephvolumepools Ceph storage pool(s) made
available to Eucalyptus for EBS volumes. Use a comma separated list for
configuring multiple pools. Default value is 'rbd'
```

6. The following steps are optional if the default values work for your cloud:
   a) To set the Ceph username (the default value for Eucalyptus is 'eucalyptus'):

   ```
   euca-modify-property -p <cluster-name>.storage.cephuser=qauser
   ```

   b) To set the absolute path to keyring file containing the key for the 'eucalyptus' user (the default value is '/etc/ceph/ceph.client.eucalyptus.keyring'):

   ```
   euca-modify-property -p
   <cluster-name>.storage.cephkeyringfile='/etc/ceph/ceph.client.qauser.keyring'
   ```

   Note: If cephuser was modified, ensure that cephkeyringfile is also updated with the location to the keyring for the specific cephuser:

   c) To set the absolute path to ceph.conf file (default value is '/etc/ceph/ceph.conf'):

   ```
   euca-modify-property -p
   <cluster-name>.storage.cephconfigfile=<path-to-ceph.conf>
   ```

   d) To change the comma-delimited list of Ceph pools assigned to Eucalyptus for managing EBS volumes (default value is 'rbd') :

   ```
   euca-modify-property -p <cluster-name>.storage.cephvolumepools=rbd,qavolumes
   ```

   e) To change the comma-delimited list of Ceph pools assigned to Eucalyptus for managing EBS snapshots (default value is 'rbd') :

   ```
   euca-modify-property -p <cluster-name>.storage.cephsnapshotpools=qasnapshots
   ```

7. Every NC will assume the following defaults:

   ```
   CEPH_USER_NAME="eucalyptus"
   CEPH_KEYRING_PATH="/etc/ceph/ceph.client.eucalyptus.keyring"
   CEPH_CONFIG_PATH="/etc/ceph/ceph.conf"
   ```

   a) To override the above defaults, add/edit the following properties in the `$EUCALYPTUS/etc/eucalyptus/eucalyptus.conf` on the specific NC. file:

   ```
   CEPH_USER_NAME="<ceph-username-for-use-by-this-NC>"
   CEPH_KEYRING_PATH="<path-to-keyring-file-for-ceph-username>"
   CEPH_CONFIG_PATH="<path-to-ceph.conf-file>"
   ```

   b) Repeat the previous two steps for every NC in the specific Eucalyptus cluster.

Configure Dell Equallogic Multipathing

Use multipathing to provide network-and-SP-redundancy for the iSCSI data path between the Dell Equallogic SAN and NCs.

Important: It is **strongly** recommended that you get your system up and running and able to create volumes and snapshots using the Dell Equallogic SAN prior to configuring multipathing. Multipathing can be configured after the cloud is fully functional and will apply to any volumes attached/snapshotted after multipathing is configured. Configuring multipathing on a non-multipathed system does not require a restart of the SC, NC, or CLC.

Important: The Dell Equallogic SAN has separate paths for data and management.

CC-BY-SA, Eucalyptus Systems, Inc.

The Dell Equallogic management interface is available for executing control operations only. If your Dell Equallogic SAN is configured to use the management port, please note the IP address of the management interface. The SC can be configured to use the management interface by specifying the IP address of the management interface using the `scpaths` property. For example:

```
euca-modify-property -p mypartition.storage.scpaths=192.168.3.1
```

The Dell Equallogic data interface is configured by specifying the IP address of the data interface using the `ncpaths` property. For example:

```
euca-modify-property -p mypartition.storage.ncpaths=192.168.3.1
```

To configure multipathing for a Dell Equallogic SAN:

1. Ensure that the mutipathd service is running on each NC:

   ```
   mpathconf --enable
   ```

2. Configure the /etc/multipath.conf file.

   Note: An example configuration for the Dell Equallogic SAN is installed with Eucalyptus. This file is located in `/usr/share/doc/eucalyptus-3.4.1/multipath.conf.example.equallogic` on each NC.

3. Start the mutipathd service:

   ```
   service multipathd start
   ```

4. If you modify the /etc/multipath.conf file, be sure to restart and reconfigure multipathd:

   a) Restart the multipathd service:

   ```
   service multipathd restart
   ```

   b) Run multipathd -k:

   ```
   multipathd -k
   ```

   c) Enter the following commands at the multipathd interactive prompt:

   ```
   reconfigure
   quit
   ```

5. Check that the multipath udev rules file is installed by verifying that the file `/etc/udev/rules.d/12-dm-permissions.rules` file exists.
6. Set the ISCSI paths:

   Tip: The path specification format is `iface0:ip0,iface1:ip1,...,ifaceN:ipN` where iface may be reused (i.e. multiple iface0 entries). Also, note that 'iface' is optional, you may just specify a comma-delimited list of IPs. Eucalyptus will detect which interfaces on the SC/NC can reach each specified IP and will use the first found. You must only specify the 'iface' value if you want precise control over which interfaces access which IPs. For using a single path only, just specify the IP of the iscsi data port to use on the VNX.

   Note: We recommend initially getting the system working with only one path. The path values can be modified at any time to enable multipathing, so it is possible to get everything working and confirm full VNX functionality before attempting multipathing. To use one path, simply specify a single IP for each the following steps.

   a) Set the NC paths. For example:

   ```
   euca-modify-property -p mypartition.storage.ncpaths=iface0:127.0.0.1
   ```

   b) Set the SC paths. For example:

   ```
   euca-modify-property -p mypartition.storage.scpaths=iface0:127.0.0.1
   ```

CC-BY-SA, Eucalyptus Systems, Inc.

Note: The NC and SC may each have different path lists, or you can optionally only enable multipathing on the NCs or SC if desire.

c) If you specified an `iface` when setting the SC paths, be sure to include a line in the eucalyptus.conf file of each NC in the cluster that defines each `iface`. For example:

```
STORAGE_INTERFACES="iface0=eth0"
```

7. Test and verify the configuration by creating (and attaching to) a volume and creating a snapshot on the partition.
8. If testing is successful, you can now configure multiple paths in your *.storage.ncpaths and *.storage.scpaths configuration. In the following example, the IP addresses for each interface correspond to the paths configured on the Dell Equallogic SAN:

```
euca-modify-property -p
mypartition.storage.ncpaths=iface0:192.168.1.1,iface1:192.168.1.2
```

9. Verify that multipathing is working on an NC by attaching a volume to an instance on that NC and running the following command:

```
multipath -ll
```

This command should return output similar to the following:

```
mpathb (36006016098b0300080722f971b2ee211) dm- 0 DGC,VRAID
size=1.0G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|- +-  policy='round- robin 0' prio=50 status=active
| `-  6:0:0:1 sdd 8:48 active ready running
`- +-  policy='round- robin 0' prio=10 status=enabled
`-  7:0:0:1 sdf 8:80 active ready running
```

You have now successfully configured multipathing for your Dell Equallogic SAN installation.

Configure EMC VNX Multipathing

Use multipathing to provide network-and-SP-redundancy for the iSCSI data path between the EMC VNX SAN and NCs.

Important: It is **strongly** recommended that you get your system up and running and able to create volumes and snapshots using the EMC VNX prior to configuring multipathing. Multipathing can be configured after the cloud is fully functional and will apply to any volumes attached/snapshotted after multipathing is configured. Configuring multipathing on a non-multipathed system does not require a restart of the SC, NC, or CLC.

To configure multipathing for a EMC VNX SAN:

1. Ensure that the mutipathd service is running on each NC:

```
mpathconf --enable
```

2. Configure the /etc/multipath.conf file.

Note: An example configuration for EMC VNX is installed with Eucalyptus. This file is located in `/usr/share/doc/eucalyptus-3.4.1/multipath.conf.example.vnx` on each NC.

3. Start the mutipathd service:

```
service multipathd start
```

4. If you modify the `/etc/multipath.conf` file, be sure to restart and reconfigure multipathd:
   a) Restart the multipathd service:

```
service multipathd restart
```

   b) Run multipathd -k:

CC-BY-SA, Eucalyptus Systems, Inc.

```
multipathd -k
```

c) Enter the following commands at the multipathd interactive prompt:

```
reconfigure
quit
```

5. Check that the multipath udev rules file is installed by verifying that the file `/etc/udev/rules.d/12-dm-permissions.rules` file exists.
6. Set the ISCSI paths:

Note: The path specification format is iface0:ip0,iface1:ip1,...,ifaceN:ipN where iface may be re- used (i.e. multiple iface0 entries). Also, note that 'iface' is optional, you may just specify a comma- delimited list of IPs. Eucalyptus will detect which interfaces on the SC/NC can reach each specified IP and will use the first found. You must only specify the 'iface' value if you want precise control over which interfaces access which IPs. For using a single path only, just specify the IP of the iscsi data port to use on the VNX.

Note: We recommend initially getting the system working with only one path. The path values can be modified at any time to enable multipathing, so it is possible to get everything working and confirm full VNX functionality before attempting multipathing. To use one path, simply specify a single IP for each the following steps.

a) Set the NC paths. For example:

```
euca-modify-property -p mypartition.storage.ncpaths=iface0:127.0.0.1
```

b) Set the SC paths. For example:

```
euca-modify-property -p mypartition.storage.scpaths=iface0:127.0.0.1
```

Note: The NC and SC may each have different path lists, or you can optionally only enable multipathing on the NCs or SC if desire.

c) If you specified an `iface` when setting the SC paths, be sure to include a line in the eucalyptus.conf file of each NC in the cluster that defines each `iface`. For example:

```
STORAGE_INTERFACES="iface0=eth0"
```

7. Test and verify the configuration by creating (and attaching to) a volume and creating a snapshot on the partition.
8. If testing is successful, you can now configure multiple paths in your *.storage.ncpaths and *.storage.scpaths configuration.
9. Verify that multipathing is working on an NC by attaching a volume to an instance on that NC and running the following command:

```
multipath -ll
```

This command should return output similar to the following:

```
mpathb (36006016098b0300080722f971b2ee211) dm- 0 DGC,VRAID
size=1.0G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|- +-  policy='round- robin 0' prio=50 status=active
| `-  6:0:0:1 sdd 8:48 active ready running
`- +-  policy='round- robin 0' prio=10 status=enabled
`-  7:0:0:1 sdf 8:80 active ready running
```

You have now successfully configured multipathing for your EMC VNX SAN installation.

Configure NetApp Multipathing

Use multipathing to provide network and controller redundancy for the iSCSI data path between the NetApp Cluster-mode SAN and NCs.

CC-BY-SA, Eucalyptus Systems, Inc.

Important: Eucalyptus supports multipathing for NetApp Clustered ONTAP only.

Important: It is **strongly** recommended that you get your system up and running and able to create volumes and snapshots using the NetApp SAN prior to configuring multipathing. Multipathing can be configured after the cloud is fully functional and will apply to any volumes attached/snapshotted after multipathing is configured. Configuring multipathing on a non-multipathed system does not require a restart of the SC, NC, or CLC.

To configure multipathing for a NetApp SAN:

1. Ensure that the mutipathd service is running on the SC and on each NC:
   ```
   mpathconf --enable
   ```
2. Configure the `/etc/multipath.conf` file.

   Note: An example configuration for NetApp is installed with Eucalyptus. This file is located in `/usr/share/doc/eucalyptus-3.4.1/multipath.conf.example.netapp` on each NC.

3. Start the mutipathd service:
   ```
   service multipathd start
   ```
4. If you modify the `/etc/multipath.conf` file, be sure to restart and reconfigure multipathd:
   a) Restart the multipathd service:
   ```
   service multipathd restart
   ```
   b) Run multipathd -k:
   ```
   multipathd -k
   ```
   c) Enter the following commands at the multipathd interactive prompt:
   ```
   reconfigure
   quit
   ```
5. Check that the multipath udev rules file is installed by verifying that the file `/etc/udev/rules.d/12-dm-permissions.rules` file exists.
6. Set the ISCSI paths:

   Note: The path specification format is iface0:ip0,iface1:ip1,...,ifaceN:ipN where iface may be re- used (i.e. multiple iface0 entries). Also, note that ‘iface’ is optional, you may just specify a comma- delimited list of IPs. Eucalyptus will detect which interfaces on the SC/NC can reach each specified IP and will use the first found. You must only specify the ‘iface’ value if you want precise control over which interfaces access which IPs. For using a single path only, just specify the IP of the iscsi data port to use on the NetApp Clustered ONTAP.

   Note: We recommend initially getting the system working with only one path. The path values can be modified at any time to enable multipathing, so it is possible to get everything working and confirm full NetApp functionality before attempting multipathing. To use one path, simply specify a single IP for each the following steps.

   a) Set the NC paths. For example:
   ```
   euca-modify-property -p mypartition.storage.ncpaths=iface0:127.0.0.1
   ```
   b) Set the SC paths. For example:
   ```
   euca-modify-property -p mypartition.storage.scpaths=iface0:127.0.0.1
   ```

   Note: The NC and SC may each have different path lists, or you can optionally only enable multipathing on the NCs or SC if desire.

CC-BY-SA, Eucalyptus Systems, Inc.

c) If you specified an `iface` when setting the SC paths, be sure to include a line in the eucalyptus.conf file of each NC in the cluster that defines each `iface`. For example:

```
STORAGE_INTERFACES="iface0=eth0"
```

7. Test and verify the configuration by creating (and attaching to) a volume and creating a snapshot on the partition.
8. If testing is successful, you can now configure multiple paths in your *.storage.ncpaths and *.storage.scpaths configuration.
9. Verify that multipathing is working on the SC and on an NC by attaching a volume to an instance on the SC and the NC and running the following command:

```
multipath -ll
```

This command should return output similar to the following:

```
mpathp (3600a098037542d69535d43514965354e) dm-2 NETAPP,LUN C-Mode
size=2.0G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1 alua'
 wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 18:0:0:0 sdd 8:48 active ready running
| `- 20:0:0:0 sdf 8:80 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 17:0:0:0 sdc 8:32 active ready running
  `- 19:0:0:0 sde 8:64 active ready running
```

You have now successfully configured multipathing for your NetApp Clustered ONTAP system.

Configure the Imaging Service

The Eucalyptus Imaging Service is a service, introduced in Eucalyptus 4.0, that makes it easier to deploy EBS images in your Eucalyptus cloud. This service automates many of the labor-intensive processes required for uploading data into EBS images. The Imaging Service is also involved in running paravirtual images: it converts the partitions of paravirtual images into disks and caches them in a system-owned bucket in Object Store (in Eucalyptus 4.0, raw disks are exclusively used within the system for deploying images to Node Controllers.)

Note: Imaging Service is **required** to run paravirtual images in your Eucalyptus cloud.

The Eucalyptus Imaging Service is implemented as a system-controlled "worker" virtual machine that is monitored and controlled via Auto Scaling. Once the Imaging Service is configured, the Imaging Service VM will be started automatically upon the first request that requires it: either an EBS volume ingress or a run of a paravirtual image that hasn't been converted. Specifically, in this release of Eucalyptus, there are three usage scenarios for the Eucalyptus Imaging Service:

- ***Importing a raw disk image as a volume:*** If you have a raw disk image (containing either a data partition or a full operating system with a boot record, e.g., an HVM image), you can use the Imaging Service to import this into your Eucalyptus cloud as a volume. This is accomplished with euca-import-volume command. If the volume was populated with a bootable disk, that volume can be snapshotted and registered as an EMI.
- ***Importing a raw disk image as an instance:*** If you have a raw disk image containing a bootable operating system, you can import this disk image into Eucalyptus as an instance: the Imaging Service automatically creates a volume, registers the EMI, and launches an instance from the EMI. This is accomplished with euca-import-instance command, which has options for specifying the instance type and the SSH key for the instance to use.
- ***Running a paravirtual image:*** If your system has partition-based paravirtual EMIs (which entered the system via the standard bundle/upload/register process), the Imaging Service will be engaged when you run such an EMI on a 4.0 system for the first time: Soon after run request is acknowledged, an Imaging Worker instance will start (unless one is already running) and will begin processing the paravirtual image. During this time, both the pending instance and the EMI will have tags associated with them showing the status of conversion. If conversion is successful, the instance will be deployed and the image will be usable again, without requiring conversion upon next run.

Install and Register the Imaging Worker Image

Eucalyptus provides a command-line tool for installing and registering the Imaging Worker image. Once you have run the tool, the Imaging Worker will be ready to use.

CC-BY-SA, Eucalyptus Systems, Inc.

1. Run the following command on the machine where you installed the `eucalyptus-service-image` RPM package (it will set the `imaging.imaging_worker_emi` property to the newly created EMI of the imaging worker):

```
euca-install-imaging-worker --install-default
```

2. Consider setting the `imaging.imaging_worker_instance_type` property to an Instance Type with enough ephemeral disk to convert any of your paravirtual images. The Imaging Worker root filesystem takes up about 2GB, so the maximum paravirtual image that the Imaging Worker will be able to convert is the disk allocation of the Instance Type minus 2GBs.

```
# euca-modify-property -p services.imaging.worker.instance_type=m3.xlarge
PROPERTY       services.imaging.worker.instance_type    m3.xlarge was m1.small
```

3. Consider setting the `imaging.imaging_worker_keyname` property to an SSH keyname (previously created with the `euca-create-keypair` command), so that you can perform troubleshooting inside the Imaging Worker instance, if necessary:

```
# euca-modify-property -p services.imaging.worker.keyname=mykey
```

Managing Imaging Worker Instance

Eucalyptus automatically starts Imaging Worker instances when there are tasks for workers to perform.

1. The cloud administrator can list running Imaging Worker instances, if any, by running the command:

```
# euca-describe-instances --filter tag-value=euca-internal-imaging-workers
```

2. Imaging Worker instances are not terminated automatically when they no longer have tasks to perform. Free up resources used by the Imaging Worker instances by reducing the Auto Scaling group size as follows:

   - `Min=0`
   - `Desired=0`
   - `Max=`leave at `1`

```
# euscale-update-auto-scaling-group asg-euca-internal-imaging-worker-01
--desired-capacity 0 --min-size 0
```

Note: When an Auto Scaling group is reduced in that manner, tasks that require the Imaging Worker (new paravirtual EMI deployment, import of volumes and instances) will not work.

3. To allow the Imaging Worker instances to be deployed again, increase the Auto Scaling group (ASG) capacity:

   - `Min=1`
   - `Desired=1`
   - `Max=`leave at `1`

```
# euscale-update-auto-scaling-group asg-euca-internal-imaging-worker-01
--desired-capacity 1 --min-size 1
```

Troubleshooting Imaging Worker

If the Imaging Worker is configured correctly, users will be able to import data into EBS volumes with `euca-import-*` commands, and paravirtual EMIs will run as instances. In some cases, though, paravirtual images may fail to convert (e.g., due to intermittent network failures or a network setup that doesn't allow the Imaging Worker to communicate with the CLC), leaving the images in a special state. To troubleshoot:

1. If the Imaging Worker Instance Type does not provide sufficient disk space for converting all paravirtual images, the administrator may have to change the Instance Type used by the Imaging Worker. After changing the instance type, the Imaging Worker instance should be restarted by terminating the old Imaging Worker instance:

```
euca-modify-property -p services.imaging.worker.instance_type=m2.2xlarge
euca-terminate-instances $(euca-describe-instances --filter
tag-value=euca-internal-imaging-workers | grep INSTANCE | cut -f 2)
```

CC-BY-SA, Eucalyptus Systems, Inc.

2. If the status of the conversion operation is 'Image conversion failed', but the image is marked as 'available' (in the output of euca-describe-images), the conversion can be retried by running the EMI again: :

```
euca-run-instance ...
```

Configure DNS

Eucalyptus provides a DNS service that you can configure to map instance IPs and Walrus bucket names to DNS host names. This section details how to configure the Eucalyptus DNS service.

The DNS service will automatically try to bind to port 53. If port 53 cannot be used, DNS will be disabled. Typically, other system services like dnsmasq are configured to run on port 53. To use the Eucalyptus DNS service, you must disable these services.

Configure the Domain and Subdomain

Before using the DNS service, configure the DNS sub domain name that you want Eucalyptus to handle using the steps that follow. Make sure that the Eucalyptus Cloud Controller (CLC) has been started.

1. Log in to the CLC and enter the following:

```
euca-modify-property -p
system.dns.dnsdomain=<eucadomain.yourdomain>
```

2. You can configure the load balancer DNS subdomain. To do so, log in to the primary CLC and enter the following:

```
euca-modify-property -p
loadbalancing.loadbalancer_dns_subdomain = <your-subdomain>
```

Turn on IP Mapping

To turn on mapping of instance IPs to DNS host names:

1. Enter the following command on the CLC:

```
euca-modify-property -p bootstrap.webservices.use_instance_dns=true
```

When this option is enabled, public and private DNS entries are set up for each instance that is launched in Eucalyptus. This also enables virtual hosting for Walrus. Buckets created in Walrus can be accessed as hosts. For example, the bucket `mybucket` is accessible as `mybucket.walrus.eucadomain.yourdomain`.

Instance IP addresses will be mapped as `euca-A.B.C.D.eucalyptus.<subdomain>`, where `A.B.C.D` is the IP address (or addresses) assigned to your instance.

2. If you want to modify the subdomain that is reported as part of the instance DNS name, enter the following command:

```
euca-modify-property -p cloud.vmstate.instance_subdomain=.<custom-dns-subdomain>
```

When this value is modified, the public and private DNS names reported for each instance will contain the specified custom DNS subdomain name, instead of the default value, which is `eucalyptus`. For example, if this value is set to `foobar`, the instance DNS names will appear as `euca-A.B.C.D.foobar.<subdomain>`.

Enable DNS Delegation

DNS delegation allows you to forward DNS traffic for the Eucalyptus subdomain to the Eucalyptus CLC host. This host acts as a name server. This allows interruption-free access to Eucalyptus cloud services in the event of a failure. The CLC host is capable of mapping cloud host names to IP addresses of the CLC and Walrus hosts.

For example, if the IP address of the CLC is `192.168.5.1`, and the IP address of Walrus is `192.168.6.1`, the host `eucalyptus.eucadomain.yourdomain` will resolve to `192.168.5.1` and `walrus.eucadomain.yourdomain` will resolve to `192.168.6.1`.

To enable DNS delegation:

1. Enter the following command on the CLC:

CC-BY-SA, Eucalyptus Systems, Inc.

```
euca-modify-property -p bootstrap.webservices.use_dns_delegation=true
```

2. Because the credentials are now slightly changed, you must generate the administrative credentials and source the `eucarc` file again. For more information, see *Generate Administrator Credentials*.

Configure the Master DNS Server

Set up your master DNS server to forward the Eucalyptus subdomain to the CLC server, which acts as a name server.

The following example shows how the Linux name server `bind` is set up to forward the Eucalyptus subdomain.

1. Open `/etc/named.conf` and set up the `eucadomain.yourdomain zone`. For example, your `/etc/named.conf` may look like the following:

```
zone "yourdomain" {
type master;
file "/etc/bind/db.yourdomain";
};

#Forward eucadomain.yourdomain
zone "eucadomain.yourdomain" {
type forward;
forward only;
forwarders { <CLC_IP>; };
};
```

   where `<CLC_IP>` is the IP address of your CLC.

2. Create `/etc/bind/db.yourdomain` if it does not exist. If your master DNS is already set up for `yourdomain`, you will need to add a name server entry for `<CLC_IP>`. For example:

```
$TTL 604800
@ IN SOA yourdomain. root.yourdomain. (
2 ; Serial
604800 ; Refresh
86400 ; Retry
2419200 ; Expire
604800 ) ; Negative Cache TTL
;
@ IN NS ns.yourdomain.
@ IN A <master_nameserver_IP>

ns.yourdomain. IN A <master_nameserver_IP>

;Add entry for CLC
eucadomain.yourdomain. IN NS clc.eucadomain.yourdomain.

clc.eucadomain.yourdomain. IN A <CLC_IP>
```

   where `clc.eucadomain.yourdomain` is the host name of your CLC server.

3. Restart the bind nameserver (`/etc/init.d/bind9 restart` or `/etc/init.d/named restart`, depending on your Linux distribution).
4. Test your setup by pointing `/etc/resolv.conf` on your client to your primary DNS server and attempt to resolve `eucalyptus.eucadomain.yourdomain` using ping or nslookup. It should return the IP address of the CLC server.

Advanced DNS Options

Recursive lookups and split-horizon DNS are available in Eucalyptus.

1. To enable any of the DNS resolvers, set `dns.enabled` to `true`.
2. To enable the recursive DNS resolver, set `dns.recursive.enabled` to `true`.

CC-BY-SA, Eucalyptus Systems, Inc.

3. To enable split-horizon DNS resolution for internal instance public DNS name queries, set `dns.split_horizon.enabled` to `true`.

Configure Node Controller

To prevent potential problems, we recommend that you perform the steps listed in this topic on each NC.

1. Log in to an NC server and open the `/etc/eucalyptus/eucalyptus.conf` file.
2. Change the `CONCURRENT_DISK_OPS` parameter to the number of disk-intensive operations you want the NC to perform at once. On some Linux installations, a sufficiently large amount of local disk activity can slow down process scheduling. This can cause other operations (e.g., network communication and instance provisioning) appear to stall. Examples of disk-intensive operations include preparing disk images for launch and creating ephemeral storage. Set this value to `1` to serialize all disk-intensive operations. Set to a higher number to increase the amount of disk-intensive operations the NC will perform in parallel.

Set Up Security Groups

In Managed and Managed (No VLAN) networking modes, you must configure the system with parameters that define how Eucalyptus will allocate and manage virtual machine networks. These virtual machine networks are known as security groups. The relevant parameters are set in the `eucalyptus.conf` on all machines running a CC.

These parameters are:

- `VNET_SUBNET`
- `VNET_NETMASK`
- `VNET_ADDRSPERNET`

The CC will read `VNET_SUBNET` and `VNET_NETMASK` to construct a range of IP addresses that are available to all security groups. This range will then be further divided into smaller networks based on the size specified in `VNET_ADDRSPERNET`. Note that Eucalyptus reserves eleven addresses per security group, so these networks will be smaller than the value specified in `VNET_ADDRSPERNET`.

The first time an instance runs in a given security group, Eucalyptus chooses an unused range of IPs of size specified in `VNET_ADDRSPERNET`. Eucalyptus then implements this network across all CCs. All instances that run within this given security group obtain a specific IP from this range.

Tip: Eleven of the IP addresses within each security group network are reserved for Eucalyptus to use as gateway addresses, broadcast address, etc. For example, if you set `VNET_ADDRSPERNET` to `32`, there will be 21 free IPs that are available for instances running in that security group.

In Managed mode, each security group network is assigned an additional parameter that is used as the VLAN tag. This parameter is added to all virtual machine traffic running within the security group. By default, Eucalyptus uses VLAN tags starting at 2, going to a maximum of 4094. The maximum is dependent on how many security group networks of the size specified in `VNET_ADDRSPERNET` fit in the network defined by `VNET_SUBNET` and `VNET_NETMASK`.

If your networking environment is already using VLANs for other reasons, Eucalyptus supports the definition of a smaller range of VLANs that are available to Eucalyptus. To configure Eucalyptus to use VLANs within a specified range:

1. Choose your range (a contiguous range of VLANs between 2 and 4095).
2. Configure your cluster controllers with a VNET_SUBNET/VNET_NETMASK/VNET_ADDRSPERNET that is large enough to encapsulate your desired range. For example, for a VLAN range of 1024-2048, you could set VNET_NETMASK to 255.254.0.0 to get a large enough network (131072 addresses), and VNET_ADDRSPERNET to 64, to give 2048 possible security groups.

Tip: The number of instances per security group can be calculated as follows:

subnets (SGs) = no. hosts / addrspernet
instances per subnet (SG) = addrspernet - 10

CC-BY-SA, Eucalyptus Systems, Inc.

3. Configure your cloud controller to work within that range. Use the following commands to verify that the range is now set to be 2-2048, a superset of the desired range.

```
euca-describe-properties | grep cluster.maxnetworktag
euca-describe-properties | grep cluster.minnetworktag
```

4. Constrict the range to be within the range that the CC can support as follows:

```
euca-modify-property -p cloud.network.global_max_network_tag=<max_vlan_tag>
euca-modify-property -p cloud.network.global_min_network_tag=<min_vlan_tag>
```

   This ensures that Eucalyptus will only use tags between 1024 and 2048, giving you a total of 1024 security groups, one VLAN per security group.

Tip: If VMs are already running in the system using a VLAN tag that is outside the range specified by global_min_network_tag-global_max_network_tag, that network will continue to run until all VMs within the network are terminated and the system removes reference to that network. Best practice is to configure these values in advance of running virtual machines.

Configure the Load Balancer

Eucalyptus provides optional support for Load Balancing. In order to use this support, you will need to register the Load Balancer image with the cloud.

Install and Register the Load Balancer Image

Eucalyptus provides a tool for installing and registering the Load Balancer image. Once you have run the tool, your Load Balancer will be ready to use.

Run the following command on the machine where you installed the eucalyptus-service-image package:

```
euca-install-load-balancer --install-default
```

Verify Load Balancer Configuration

If you would like to verify that Load Balancer support is enabled you can list installed Load Balancers. The currently active Load Balancer will be listed as enabled. If no Load Balancers are listed, or none are marked as enabled, then your Load Balancer support has not been configured properly.

1. Run the following command to list installed Load Balancer images:

```
euca-install-load-balancer --list
```

2. You can also check the enabled Load Balancer EMI with:

```
euca-describe-properties services.loadbalancing.worker.image
```

3. If you need to manually set the enabled Load Balancer EMI use:

```
euca-modify-property -p services.loadbalancing.worker.image=emi-12345678
```

CC-BY-SA, Eucalyptus Systems, Inc.

Eucalyptus Upgrade or Migration

This section details how upgrade a current installation and how to migrate a non-high availability deployment to high availability.

Eucalyptus Upgrade

This section details the tasks you need to perform in order to upgrade your current version of Eucalyptus.

You can upgrade to Eucalyptus 4.1.2 from 4.1.1 or 4.0.2; see the prescribed paths below. Follow the directions in that version's Installation Guide in the *documentation archive*, and then upgrade to 4.1.2 using the directions in this section.

You do not need to shutdown instances in order to upgrade. However, Auto Scaling instances will likely shut down and be replaced, based on each group's scaling policy and health check criteria.

Tip: If you want to upgrade from a release prior to 4.0.2, follow the prescribed upgrade path as shown:

- Upgrade from 3.1.2 -> 3.2.2
- Upgrade from 3.2.2 -> 3.3.2
- Upgrade from 3.3.2 -> 3.4.3
- Upgrade from 3.4.3 -> 4.0.2

Important: Eucalyptus does not support components that are at different releases, even at the sub-minor level. For example, you cannot have a CLC at 4.1.2 and a Walrus at 4.1.1. Please make sure that you update all Eucalyptus components when you upgrade.

Prepare the Configuration File

Complete the following steps to upgrade to Eucalyptus 4.1.2 on CentOS 6 or RHEL 6.

Note:

You should already have the repositories installed for euca2ools, EPEL, and ELRepo from your previous installation. If you do not have these installed, refer to the installation instructions to find out how to add these to your machines.

The steps in this section should be performed on all machines with Eucalyptus installed.

1. Remove any hand-written repository files for earlier versions of Eucalyptus and Euca2ools from `/etc/yum.repos.d`.
2. Install the new Eucalyptus release package on each host that will run a Eucalyptus component:

```
yum install
http://downloads.eucalyptus.com/software/eucalyptus/4.1/centos/6/x86_64/eucalyptus-release-4.1.el6.noarch.rpm
```

3.

   Tip: It's recommended that you install the new version of Euca2ools, although this is not required. If you don't install the new version of Euca2ools, you will not be able to use new features from the command line.

   Install the new Euca2ools release package on each host that will run a Eucalyptus component:

```
yum install
http://downloads.eucalyptus.com/software/euca2ools/3.2/centos/6/x86_64/euca2ools-release-3.2.el6.noarch.rpm
```

4. If you have a Eucalyptus subscription, run the following command on each machine that runs a Eucalyptus component:

```
yum install
http://subscription.eucalyptus.com/eucalyptus-enterprise-release-4.1-1.el6.noarch.rpm
```

CC-BY-SA, Eucalyptus Systems, Inc.

You are now ready to *Shutdown Components*.

Shutdown Components

To shut down Eucalyptus components:

1. Log in to the CLC host and shut down the CLC service.

```
service eucalyptus-cloud stop
```

2. Log in to an SC host and shut down the SC service.

```
service eucalyptus-cloud stop
```

   Repeat for any other machine hosting an SC.

3. Log in to the Walrus host and shut down the Walrus service.

```
service eucalyptus-cloud stop
```

4. Shut down the services on the CC host.

```
service eucalyptus-cloud stop
```

   Tip: This command shuts down any CLC, Walrus, and SC components co-located with the CC, and in the correct order.

5. Log in to a CC host and shut down the CC service.

```
service eucalyptus-cc stop
```

   Repeat for each machine hosting a CC.

6. Log in to each NC host and shut down eucanet.

```
service eucanetd stop
```

7. In each NC shut down the NC service.

```
service eucalyptus-nc stop
```

   Important: Running instances on the NC will continue running.

   Repeat the last two steps for each machine hosting an NC.

You are now ready to *Upgrade Eucalyptus Packages*.

Upgrade Euca2ools Packages

To use the new features available in Eucalyptus 4.1.2, you must upgrade to the latest version of the Euca2ools packages:

1. Enter the following command on each machine running a Eucalyptus component:

```
yum clean expire-cache
```

2. Enter the following command on each machine running a Eucalyptus component:

```
yum update euca2ools
```

You are now ready to *Start Eucalyptus*.

Upgrade Eucalyptus Packages

Before upgrading Eucalyptus packages, we suggest fully updating your systems using `yum update` where possible.

CC-BY-SA, Eucalyptus Systems, Inc.

Note: When using Walrus in a high availability (HA) configuration, *mount(8)* is now used for mounting and unmounting the DRBD device, instead of *mount(2)*. If Walrus HA is configured, add an entry to the `/etc/fstab` to mount the DRBD device to `/var/lib/eucalyptus/bukkits` on both primary and secondary Walrus. For example, if the DRBD device in the Eucalyptus DRBD resource file is defined as `/dev/drbd1` and the filesystem format is ext3, add the following line to `/etc/fstab`:

```
/dev/drbd1 /var/lib/eucalyptus/bukkits ext3 noauto,owner 0 0
```

.

To upgrade Eucalyptus packages:

1. Enter the following command on each machine running a Eucalyptus component:

   ```
   yum install
   http://downloads.eucalyptus.com/software/eucalyptus/4.1/rhel/6/x86_64/eucalyptus-release-4.1.el6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

2. If you are not a Eucalyptus subscriber, skip this step. If you are a Eucalyptus subscriber, you should have received an rpm package file containing your license for subscription-only components. Install that package, along with the Eucalyptus subscription package, on each host that will run a Eucalyptus component, as follows:

   ```
   yum install eucalyptus-enterprise-license*.noarch.rpm \
   http://subscription.eucalyptus.com/eucalyptus-enterprise-release-4.1-1.el6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

3. Enter the following command on each machine running a Eucalyptus component:

   ```
   yum clean expire-cache
   ```

4. Enter the following command on each machine running a Eucalyptus component:

   ```
   yum update 'eucalyptus*'
   ```

   If you have previously customized your configuration files, `yum` returns a warning, and installs the new configuration files with a different name. This preserves your customizations. Before you continue, customize and rename the new Configuration files.

   Tip: For larger deployments, use a script to upgrade the component host machines. For example:

   ```
   for host in 28 29 32 33 35 39 40; do echo 192.168.51.$host;
   ssh 192.168.51.$host 'yum -y update $( rpm -qa | grep euca )' ; done
   ```

5. Install the `qemu-kvm-rhev` package on each NC.

You are now ready to *Upgrade Euca2ools Packages*.

Switch to the Stock DHCP Package

This procedure is optional. However, Eucalyptus 4.1.2 uses the stock DHCP server package. Prior to starting Eucalyptus services after upgrading packages, edit the configurations on each Cluster Controller (CC) and Node Controller (NC).

To edit the values in the CCs and NCs:

1. Go to `/etc/eucalyptus/eucalyptus.conf` and change the value of `VNET_DHCPDAEMON` from `/usr/sbin/dhcpd41` to `/usr/sbin/dhcpd`.
2. Repeat on each CC and NC.
3. After you have started all Eucalyptus services as described in the next section, make sure that DHCP is functioning correctly. Make sure that instances are being assigned addresses from the DHCP server.

CC-BY-SA, Eucalyptus Systems, Inc.

Tip: Not applying this change does not cause the cloud to malfunction. But Eucalyptus no longer maintains the existing `VNET_DHCPDAEMON` package. From this release on, Eucalyptus will already have the stock package applied.

Start Eucalyptus

1. In the CLC, enter the following command.

```
service eucalyptus-cloud start
```

If you are upgrading from 4.1.1 you will see that the process starts the database upgrade. Eucalyptus returns output similar to the following example.

```
Starting Eucalyptus services: Attempting database upgrade from 4.1.1
at /var/lib/eucalyptus/upgrade/eucalyptus.backup.1326904600...
#                            UPGRADE INFORMATION
#=============================================================================
# Old Version:               4.1.1
# New Version:               4.1.2
# Upgrade keys:              false              using:
# Upgrade configuration:     false              using:
# Upgrade database:          true               using: upgrade_db
# Same version:              false              using:
# Start upgrading: db
Upgrading your database...
.
.
.
# Done upgrading:  db
done.
[debug:0387] redirecting stdout to //var/log/eucalyptus/startup.log
and stderr to //var/log/eucalyptus/startup.log
done.
```

Note: You might see some warnings in the output. These are a known issue.

2. Log in to the Walrus server and enter the following command:

```
service eucalyptus-cloud start
```

If you are upgrading from 4.1.1 you will see that Eucalyptus returns output similar to the following example.

```
Starting Eucalyptus services: Attempting database upgrade from 4.1.1
at /var/lib/eucalyptus/upgrade/eucalyptus.backup.1326905005...
#                            UPGRADE INFORMATION
#=============================================================================
# Old Version:               4.1.1
# New Version:               4.1.2
# Upgrade keys:              false              using:
# Upgrade configuration:     false              using:
# Upgrade database:          true               using: upgrade_db
# Same version:              false              using:
# Start upgrading: db
CLC is disabled.  Skipping DB upgrade.
# Done upgrading:  db
done.
[debug:0387] redirecting stdout to //var/log/eucalyptus/startup.log
and stderr to //var/log/eucalyptus/startup.log
done.
```

3. Log in to the CC server and enter the following:

```
service eucalyptus-cc start
```

CC-BY-SA, Eucalyptus Systems, Inc.

4. If you have a multi-cluster setup, repeat the previous step for each cluster.
5. Repeat for each CC server.
6. Log in to the SC server and enter the following command:

```
service eucalyptus-cloud start
```

If you are upgrading from 4.1.1 you will see that Eucalyptus returns output similar to the following example>

```
Starting Eucalyptus services: Attempting database upgrade from 4.1.1
at /var/lib/eucalyptus/upgrade/eucalyptus.backup.1326905005...
#                           UPGRADE INFORMATION
#=============================================================================
# Old Version:              4.1.1
# New Version:              4.1.2
# Upgrade keys:             false              using:
# Upgrade configuration:    false              using:
# Upgrade database:         true               using: upgrade_db
# Same version:             false              using:
# Start upgrading: db
CLC is disabled.  Skipping DB upgrade.
# Done upgrading:  db
done.
[debug:0387] redirecting stdout to //var/log/eucalyptus/startup.log
and stderr to //var/log/eucalyptus/startup.log
done.
```

7. Log in to an NC server and enter the following command:

```
service eucalyptus-nc start
```

8. Repeat for each NC server.

You are now ready to *Sync Keys on Components*.

Sync Keys on Components

After upgrading to 4.1.2, synchronize keys on remote java components. This registers user services and sets the storage backend.

The new Object Storage Gateway (OSG) and WalrusBackend components (WS) have new keys. Those keys are generated and need to be synced to all java components after upgrade, in order for everyone to have the same key set.

To sync keys on components after the CLC is started and the database is successfully upgraded:

1. Register User-Facing Services (UFS):

```
euca_conf --register-service -T user-api -H <UFS HOST IP> -N API_<last octet
of HOST IP>
```

For example:

```
# euca_conf --register-service -T user-api -H 10.111.1.135 -N API_135
```

Sample output:

```
Created new partition 'API_135'...
```

2. Verify that the `objectstorage.providerclient` property is set correctly to `walrus`:

```
euca-describe-property -p objectstorage.providerclient
```

This should return output similar to the following:

```
PROPERTY objectstorage.providerclient walrus
```

3. Synchronize the keys:

CC-BY-SA, Eucalyptus Systems, Inc.

```
euca_conf --sync-euca-p12
```

Sample output:

```
Copying euca.p12 to 10.111.1.174...
Copying euca.p12 to 10.111.5.15...
```

4. Install the imaging worker file:

```
yum install eucalyptus-service-image
```

5. Install the imaging worker image:

```
euca-install-imaging-worker --install-default
```

Sample output:

```
Installing default Imaging Service tarball...
```

As the commands process, the final output shows each service is enabled, indicating the commands ran successfully.

You are now ready to *Register Services*.

Register Services

To register the new user-facing services:

1. On the CLC server, enter the following command:

```
euca_conf --register-service -T user-api -H <host_url> -N <name>
```

For example:

```
euca_conf --register-service -T user-api -H 10.111.1.135 -N API_135
```

2. If you are upgrading an HA environment, repeat on the secondary CLC.

You are now ready to *Verify the Components*.

Verify the Components

Verify that all Eucalyptus components are running and properly connected to one another. Check to make sure that the status of each component is enabled.

To verify that all services are enabled:

1. Verify your Walruses:

```
euca_conf --list-walrusbackends
```

Eucalyptus returns a list, as in the following example.

```
WALRUS  walrus    walrus    192.168.51.28    ENABLED {}
```

2. Verify your CCs:

```
euca_conf --list-clusters
```

Eucalyptus returns a list, as in the following example.

```
CLUSTER test00    test00_cc    192.168.51.29    ENABLED {}
CLUSTER test01    test01_cc    192.168.51.35    ENABLED {}
```

3. Verify your SCs:

```
euca_conf --list-scs
```

Eucalyptus returns a list, as in the following example.

CC-BY-SA, Eucalyptus Systems, Inc.

```
STORAGECONTROLLER     test01     test01_sc     192.168.51.39      ENABLED {}
STORAGECONTROLLER     test00     test00_sc     192.168.51.32      ENABLED {}
```

4. Make sure that NCs are presenting available resources to the CC.

```
euca-describe-availability-zones verbose
```

The returned output should a non-zero number in the `free` and `max` columns, as in the following example.

```
AVAILABILITYZONE        test00  192.168.51.29
arn:euca:eucalyptus:test00:cluster:test00_cc/
AVAILABILITYZONE        |- vm types     free / max   cpu   ram  disk
AVAILABILITYZONE        |- m1.small     0004 / 0004   1    128     2
AVAILABILITYZONE        |- c1.medium    0004 / 0004   1    256     5
AVAILABILITYZONE        |- m1.large     0002 / 0002   2    512    10
AVAILABILITYZONE        |- m1.xlarge    0002 / 0002   2   1024    20
AVAILABILITYZONE        |- c1.xlarge    0001 / 0001   4   2048    20
AVAILABILITYZONE        test01  192.168.51.35
arn:euca:eucalyptus:test01:cluster:test01_cc/
AVAILABILITYZONE        |- vm types     free / max   cpu   ram  disk
AVAILABILITYZONE        |- m1.small     0004 / 0004   1    128     2
AVAILABILITYZONE        |- c1.medium    0004 / 0004   1    256     5
AVAILABILITYZONE        |- m1.large     0002 / 0002   2    512    10
AVAILABILITYZONE        |- m1.xlarge    0002 / 0002   2   1024    20
AVAILABILITYZONE        |- c1.xlarge    0001 / 0001   4   2048    20
```

You are now ready to *Update the Service Images*.

Update the Service Images

Eucalyptus uses images that serve as templates for load balancing and EBS deployment services. If you had these services in your previous installation, you must install these images again.

- Install the imaging worker image. Run the following command on the machine where you installed the Eucalyptus imaging worker image:

```
euca-install-imaging-worker --install-default
```

-

  Tip: If you have not configured load balancer support in your previous version of Eucalyptus, skip to *Upgrade Credentials*. If you want to configure load balancer support for the first time, see *Configure the Load Balancer*.

  Install the load balancer image. Run the following command on the machine where you installed the Eucalyptus load balancer package:

```
euca-install-load-balancer --install-default
```

You are now ready to *Upgrade Credentials*.

Upgrade Credentials

All users' credentials will still work after the upgrade. However the new Eucalyptus access control commands will not work until you upgrade your credentials. Other users must updates theirs as well.

To update your credentials, perform the following steps.

1. Download new credentials.

```
euca_conf --get-credentials <filename>
```

2. Unzip the credentials file.

```
unzip -o <filename>
```

3. Source the `eucarc` file.

CC-BY-SA, Eucalyptus Systems, Inc.

```
source eucarc
```

Your upgrade is now complete. If at any point your upgrade failed, see *Dealing with Failed Upgrades*.

Dealing with Failed Upgrades

The upgrade process creates a backup to `/var/lib/eucalyptus/upgrade/eucalyptus.backup.<timestamp>`. For example:

```
/var/lib/eucalyptus/upgrade/eucalyptus.backup.1326905212
```

If the upgrade fails and needs to be reverted to your earlier version, you can find your preserved data in this directory.

If the upgrade fails, all changes to the database and configuration files will be rolled back. You can retry the upgrade by following the upgrade instructions in the sections, *Shutdown Components* and *Upgrade Eucalyptus Packages*.

If you do not want to continue with the upgrade after a failure, you can downgrade your installation back to the previous version. Please note that downgrade instructions are different, depending on whether your Eucalyptus services are co-located or each run on their own machine. You will need to perform the downgrade for all services running on a single machine at the same time.

The `/var/lib/eucalyptus/db` and `/var/lib/eucalyptus/keys` directories should not be affected by the upgrade. If they have been removed subsequent to the upgrade, you must restore the contents of these directories from your backups before downgrading.

To downgrade from a failed upgrade, perform the tasks listed in the following sections.

Downgrade Eucalyptus

1. Downgrade to the Eucalyptus 4.1.1 release package on each host.

```
yum downgrade
http://downloads.eucalyptus.com/software/eucalyptus/4.1/centos/6/x86_64/eucalyptus-release-4.1.noarch.rpm
```

2. If you have a Eucalyptus subscription, downgrade your subscription release package on each host to the release package you obtained for Eucalyptus 4.1.1.

```
yum downgrade eucalyptus-enterprise-release-4.1*.rpm
```

3. Expire the cache for the yum repositories on each host.

```
yum clean expire-cache
```

4. Log in to each NC host and downgrade it. To downgrade to a specific version, append the version number to each package name. For example, to downgrade to 4.1.1, add `-4.1.1` to each of the packages in the following list:

```
yum downgrade eucalyptus eucalyptus-gl eucalyptus-admin-tools eucalyptus-nc
python-eucadmin
```

Important:

Use the `yum shell` command for the following instructions. This will allow you to perform more complex transactions that are required for the downgrade.

5. Log in to each machine running a Eucalyptus service and run the following command:

```
yum shell
```

6. Add the transaction commands listed below for each component installed on the machine. If more than one component asks you to use the same transactional command, you only need to specify that command once.

CLC Service Transaction Commands:

```
downgrade eucalyptus-cloud
downgrade eucalyptus
```

CC-BY-SA, Eucalyptus Systems, Inc.

```
downgrade eucalyptus-common-java
downgrade eucalyptus-common-java-libs
downgrade eucalyptus-admin-tools
downgrade python-eucadmin
```

Additional CLC Service Transaction Commands for Eucalyptus Subscription customers:

```
downgrade eucalyptus-enterprise-storage-san-common-libs
downgrade eucalyptus-enterprise-storage-san-libs
```

CC Service Transaction Commands:

```
downgrade eucalyptus-cc
downgrade eucalyptus
downgrade eucalyptus-gl
downgrade eucalyptus-admin-tools
downgrade python-eucadmin
```

SC Service Transaction Commands:

```
downgrade eucalyptus-sc
downgrade eucalyptus
downgrade eucalyptus-common-java
downgrade eucalyptus-common-java-libs
downgrade eucalyptus-admin-tools
downgrade python-eucadmin
```

Walrus Service Transaction Commands:

```
downgrade eucalyptus-walrus
downgrade eucalyptus
downgrade eucalyptus-common-java
downgrade eucalyptus-common-java-libs
downgrade eucalyptus-admin-tools
downgrade python-eucadmin
```

SAN EMC Transaction Commands:

```
remove eucalyptus-enterprise-storage-san-emc
downgrade eucalyptus-enterprise-storage-san-emc-libs
downgrade eucalyptus-enterprise-storage-san-common
downgrade eucalyptus-enterprise-storage-san-common-libs
```

SAN EqualLogic Transaction Commands:

```
downgrade eucalyptus-enterprise-storage-san-equallogic
remove eucalyptus-enterprise-storage-san-equallogic-libs
downgrade eucalyptus-enterprise-storage-san-common
downgrade remove eucalyptus-enterprise-storage-san-common-libs
```

SAN NetApp Transaction Commands:

```
downgrade eucalyptus-enterprise-storage-san-netapp
downgrade eucalyptus-enterprise-storage-san-netapp-libs
downgrade eucalyptus-enterprise-storage-san-common
downgrade eucalyptus-enterprise-storage-san-common-libs
```

7. When all transaction commands have been entered run the following command to verify that the transaction will be successful:

   ```
   ts solve
   ```

8. Perform the downgrade by running the following command in the transaction shell:

   ```
   run
   ```

9. Exit the transaction shell using the following command:

CC-BY-SA, Eucalyptus Systems, Inc.

```
exit
```

10. Remove the `/etc/eucalyptus/.upgrade` file from each machine:

```
rm /etc/eucalyptus/.upgrade
```

Enter `y` when prompted, to remove this file. It is important to remove this file from every Eucalyptus host.

Prepare System for Upgrade

1. Clear out the `/var/run/eucalyptus/` directory on all machines used for Eucalyptus.
2. Downgrade Euca2ools to 3.1.3.
3. Perform the upgrade tasks for your Eucalyptus version.
4. Start the cloud back up. Make sure all services show `ENABLED`.

Eucalyptus Migration to Edge Networking Mode

You can configure your existing cloud to use Edge networking mode. This topic provides instructions for configuring and installing additional Eucalyptus components in an existing environment for the purpose of moving to Edge.

Important: Migrating to Edge will require downtime of your cloud platform.

1. Terminate all running instances.
   a) Find out which instances are running:

```
euca-describe-instances
```

   b) List the instances to terminate:

```
euca-terminate-instances instance_id [instance_id ...]
```

2. Shut down all Eucalyptus components.

```
service eucalyptus-cloud stop
```

3. Edit all the config files on NC and CC for Edge networking mode. For more information, see *Configure for Edge Mode* or *Configure for Edge Mode* (for HA).
4. Install `eucanetd` on all NCs.

```
yum install eucanetd
```

5. Start `eucanetd` on all NCs

```
service eucanetd start
```

6. Start all components: CLC, CC, WS, SC, NCs. For more information, see *Start Eucalyptus* or *Start Eucalyptus* (for HA).
7. Set the Edge JSON property. For more information, see *#unique_121* or *#unique_122* (for HA).

Your Edge networking mode is now properly configured.

Eucalyptus Migration to High Availability

You can register additional components to bring high availability to your existing cloud. This topic provides instructions for registering additional Eucalyptus components (specifically Walrus) in an existing environment for the purpose of achieving high availability.

CC-BY-SA, Eucalyptus Systems, Inc.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Important: Adding an additional Walrus will require downtime of your cloud platform.

Registering redundant components to your existing Eucalyptus installation will take your platform into a highly available configuration. Registering the Cloud Controller, Cluster Controller and Storage Controller (with SAN adapter only) is fairly trivial but adding an additional Walrus requires that the user configure DRBD for bukkit storage replication.

Before you begin, ensure that you have completed the following:

- Installed the same operating system on any additional server you will add for High Availability of a Eucalyptus component.
- Installed and configured Eucalyptus with matching configuration on any additional system you will be adding.
- If registering an additional Walrus, ensure you are able to move the contents of /var/lib/eucalyptus/bukkits/ to a temporary storage area.
- A block device (disk or partition) is available for use as a DRBD device. Consider using LVM for future growth.

Important: The configuration of Highly Availiable components must match. For more information, see *Eucalyptus HA Planning*.

1. To register an additional Eucalyptus component please follow the instructions in *Register Eucalyptus*.
2. Additional steps are required for adding another Walrus component. You will need to configure a DRBD device for the bukkit store. Start by shutting down the eucalyptus-cloud service on both Walrus servers.

   ```
   service eucalyptus-cloud stop
   ```

3. On the original primary system, copy the current contents of /var/lib/eucalyptus/bukkits to a temporary location.

   ```
   cp -R --preserve /var/lib/eucalyptus/bukkits /newlocation/
   ```

4. Configure your new DRBD device. For more information, see *Configure DRBD*. After you are finished configuring your device, move to the next step.
5. On the primary Walrus, ensure the new DRBD resource is in a primary state before proceeding.

   ```
   drbd-overview
   ```

6. On the primary Walrus, copy the preserved contents of /var/lib/eucalyptus/bukkits from the temporary location to the new DRBD device.

   Tip: At this point you should run `service eucalyptus-cloud stop` to ensure cloud services are stopped during migration.

   ```
   cp -R --preserve /newlocation/* /var/lib/eucalyptus/bukkits/
   ```

   Tip: If you are migrating large amounts of data, consider skipping initial device synchronization as explained in *Skip Initial Device Synchronization*.

7. On the primary Walrus, monitor the state of the resource with `drbd-overview` to observe the sync. Data will not be replicated until the resource is marked UpToDate/UpToDate.

   ```
   drbd-overview
   ```

8. With the synchronization complete, start the eucalyptus-cloud service to bring up Walrus.

Your HA environment is now ready.

CC-BY-SA, Eucalyptus Systems, Inc.

Find More Information

This topic explains what to do once you have installed Eucalyptus, including further reading and other resources for understanding your cloud.

Read More

Eucalyptus has the following guides to help you with more information:

- The *Administration Guide* details ways to manage your Eucalyptus deployment. Refer to this guide to learn more about managing your Eucalyptus components, managing access to Eucalyptus, and managing Eucalyptus resources, like instances and images.
- The *User Guide* details ways to use Eucalyptus for your computing and storage needs. Refer to this guide to learn more about getting and using euca2ools, creating images, running instances, and using dynamic block storage devices.
- The *Image Management Guide* describes how to create and manage images for your cloud.
- The *Management Console Guide* describes how to create and manage cloud resources using the Eucalyptus Management Console.
- The *Euca2ools Reference Guide* describes the Euca2ools commands. Refer to this guide for more information about required and optional parameters for each command.

Get Involved

The following resources can help you to learn more, connect with other Eucalyptus users, or get actively involved with Eucalyptus development.

- The Eucalyptus IRC channel is #eucalyptus on Freenode. This channel is used for real-time communication among users and developers. Information on *how to use the network* is available from Freenode.
- Subscribe to one or more of the *Eucalyptus mailing lists*, which provide ways to ask questions and get assistance from the community.
- Check out the *Eucalyptus Support* pages for more ideas.

CC-BY-SA, Eucalyptus Systems, Inc.

Eucalyptus Installation from Local Package Repository

In certain situations, you might need to install Eucalyptus from a local repository. This section augments the standard installation instructions, and includes additional instructions for downloading and installing Eucalyptus from a local repository.

To install Eucalyptus from behind a firewall on CentOS 6 or RHEL 6:

1. Download the Eucalyptus repository to a local directory. For example:

```
wget -r --no-parent \
http://downloads.eucalyptus.com/software/eucalyptus/4.1/centos/6/x86_64/ \
-P /tmp/eucalyptus
```

2. Download euca2ools:

```
wget -r --no-parent \
http://downloads.eucalyptus.com/software/euca2ools/3.2/centos/6/x86_64/ \
-P /tmp/euca2ools
```

3. In step 1 of the *existing installation instructions*, modify the baseurl to point to your Eucalyptus local repository:

```
baseurl=file:///tmp/eucalyptus/downloads.eucalyptus.com/software/eucalyptus/4.1/centos/6/x86_64
```

4. In step 2 of the *existing installation instructions*, modify the baseurl to point to your local Euca2ools repository:

```
baseurl=file:///tmp/euca2ools/downloads.eucalyptus.com/software/euca2ools/3.2/centos/6/x86_64
```

5. Run `yum update`.

CC-BY-SA, Eucalyptus Systems, Inc.

Euca2ools Standalone Installation

Euca2ools is the Eucalyptus command line interface for interacting with Eucalyptus. This topic discusses how to perform a standalone installation of Euca2ools.

If you're running recent versions of Fedora, Debian, or Ubuntu, you can install Euca2ools using `yum` or `apt`.

If you're running RHEL/Centos, you can use the following instructions to install Euca2ools.

To perform a standalone installation of Euca2ools on RHEL/CentOS:

1. Configure the EPEL package repository:

```
yum install
http://downloads.eucalyptus.com/software/eucalyptus/4.1/centos/6/x86_64/epel-release-6.noarch.rpm
```

2. Configure the Euca2ools package repository:

```
yum install
http://downloads.eucalyptus.com/software/euca2ools/3.2/centos/6/x86_64/euca2ools-release-3.2.el6.noarch.rpm
```

3. Install Euca2ools:

```
yum install euca2ools
```

You've now performed a standalone installation of Euca2ools.

CC-BY-SA, Eucalyptus Systems, Inc.

Tech Preview: Eucalyptus HA Installation

This section details steps to install Eucalyptus.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

To install Eucalyptus HA, perform the tasks in the order presented in this section.

Plan Your Installation

In order to get the most out of a Eucalyptus deployment, we recommend that you create a plan that provides a complete set of features, performance, scaling, and resilience characteristics you want in your deployment.

Attention: If you are upgrading from an existing Eucalyptus release, see *Eucalyptus Upgrade*.

To successfully plan for your Eucalyptus installation, you must determine two things:

- **The infrastructure you plan to install Eucalyptus on:** Think about the application workload performance and resource utilization tuning. Think about how many machines you want on your system.
- **The amount of control you plan to give Eucalyptus on your network:** Use your existing architecture and policies to determine the Eucalyptus networking features you want to enable: elastic IPs, security groups, DHCP server, and Layer 2 VM isolation.

This section describes how to evaluate each tradeoff to determine the best choice to make, and how to verify that the resource environment can support the features that are enabled as a consequence of making a choice.

By the end of this section, you should be able to specify how you will deploy Eucalyptus in your environment, any tradeoffs between feature set and flexibility, and where your deployment will integrate with existing infrastructure systems.

Tip: For more help in planning your installation, see the *Eucalyptus Cloud Reference Architectures* page. This page includes use cases and reference architectures for various deployments.

Understanding the Eucalyptus HA Architecture

This topics describes the relationship of the components in a Eucalyptus HA installation.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

If you configure Eucalyptus for high availability (HA), you must have primary and secondary cloud and cluster components. In the event of a failure, the secondary component becomes the primary component.

CC-BY-SA, Eucalyptus Systems, Inc.

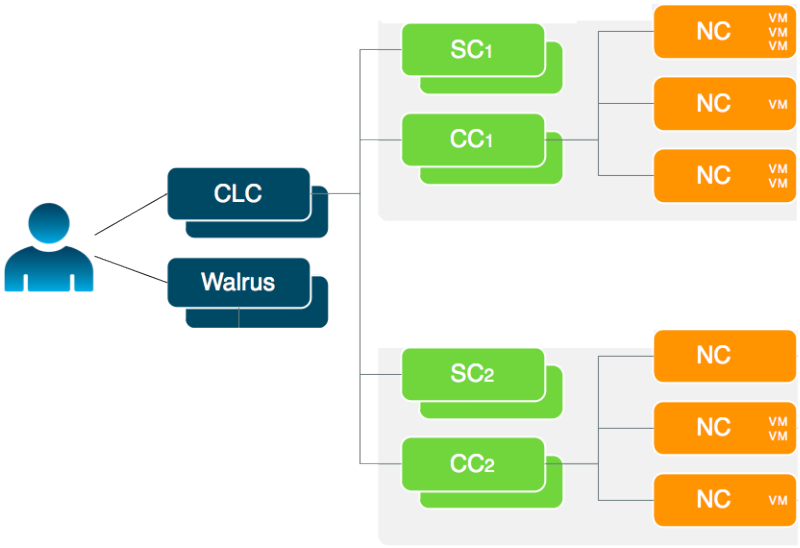

Eucalyptus HA uses a service called an Arbitrator that monitors connectivity between a user and a user-facing component (CLC, Walrus, and CC). An Arbitrator approximates reachability to a user. Each Arbitrator uses ICMP messages to periodically test reachability to an external entity (for example, a network gateway or border router) or to an external site (for example, google.com).

An Arbitrator is not required in HA. However, it is nice to have in order to test connectivity with a user.

If all Arbitrators fail to reach a monitored entity, Eucalyptus assumes there is a loss of connectivity between a user and the component. At that point a failover occurs. To allow for normal outages and maintenance, we recommend that you register more than one Arbitrator for each user-facing component.

Plan Component Placement

A Eucalyptus deployment is a set of cloud services (Cloud Controller and Walrus) and one or more clusters, each of which contains a Cluster Controller, a Storage Controller, and one or more Node Controllers.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

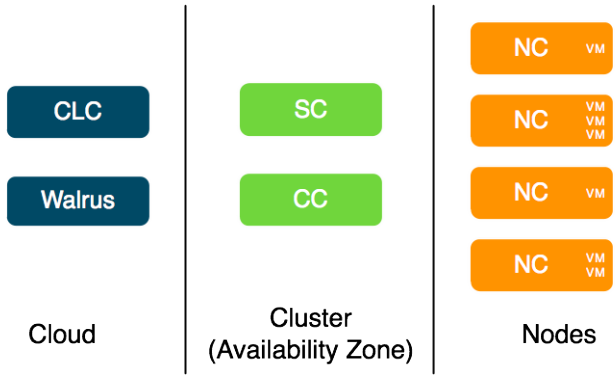

Cloud Components

The main decision for cloud components is whether to install the Cloud Controller (CLC) and Walrus on the same server. If they are on the same server, they operate as separate web services within a single Java environment, and they use a fast-path for inter-service communication. If they are not on the same server, they use SOAP and REST to work together.

CC-BY-SA, Eucalyptus Systems, Inc.

However, when installed on the same server, the CLC and Walrus must share a common memory footprint, both managed by the Java memory manager. Walrus self-tunes its performance based on the memory pressure it perceives and runs faster with more memory. So, while separating the CLC and Walrus decreases the efficiency of the messaging between the two, it often increases the responsiveness of the overall Eucalyptus system when Walrus is given a large memory footprint.

Sometimes the key factor for cloud components is not performance, but server cost and data center configuration. If you only have one server available for the cloud, then you have to install the components on the same server.

The CLC and Walrus components are not designed to be separated by wide-area, common carrier networks. They use aggressive time-outs to maintain system responsiveness so separating them over a long-latency, lossy network link will not work.

The CLC and Walrus communicate with Eucalyptus clients independently. End-users typically interact with Eucalyptus through a client interface. They can use either our provided euca2ools Linux command line client tools, or the Eucalyptus AWS-compatible API, or a third-party client that is compatible with Eucalyptus. In all cases, the end-user client must be able to send messages via TCP/IP to the machine on which the CLC is deployed.

In addition, the CLC must have TCP/IP connectivity to all other Eucalyptus components except for node controllers (NCs), which may reside on their own private networks. In addition, NC servers must be able to send messages to the Walrus server because images are downloaded by the NC using the Walrus URL. That is, the CLC does not need to be able to route network traffic directly to the NCs but Walrus does for the purposes of image delivery.

Cluster Components

The Eucalyptus components deployed in the cluster level of a Eucalyptus deployment are the Cluster Controller (CC) and Storage Controller (SC).

You can install all cluster components on a single machine, or you can distribute them on different machines. The choice of one or multiple machines is dictated by the demands of user workload in terms of external network utilization (CC) and EBS volume access (SC).

Things to consider for CC placement:

- For Managed and Managed (No VLAN) modes, the CC physical machine becomes a software IP gateway between VM instances and the public network. Because of this software routing function, the physical server on which the CC is deployed should have fast, dedicated network access to both the NC network, and the public network.
- For Edge mode, the CC physical machine will not act as a software gateway. Network traffic will be limited to small control messages.
- In all cases, place the CC on a machine that has TCP/IP connectivity to the Eucalyptus front end servers and the NC servers in its cluster.

Things to consider for SC placement:

- The machine on which the SC is deployed must always have TCP/IP connectivity to the CLC. If you are a subscriber and use one of Eucalyptus' provided SAN integration drivers, the SC must also have TCP/IP connectivity to the chosen SAN device. In this case, the SC only sends control messages to the SAN.
- If you do not configure a SAN, the SC requires only TCP/IP connectivity to the NCs in the cluster. The SC will use this TCP/IP connectivity to provide the NCs network access to the dynamic block volumes residing on the SC's storage. SC storage should consist of a fast, reliable disk pool (either local file-system or block-attached storage) so that the SC can create and maintain volumes for the NCs. The capacity of the disk pool should be sufficient to provide the NCs with enough space to accommodate all dynamic block volumes requests from end-users

Node Components

The Node Controllers are the components that comprise the Eucalyptus back-end. All NCs must have network connectivity to whatever hosts their EBS volumes. This host is either a SAN or the SC.

Plan Your Hardware

This topic describes ways you can install Eucalyptus components on your machines.

CC-BY-SA, Eucalyptus Systems, Inc.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

In Eucalyptus HA, we recommend that you install no more than one Eucalyptus component on each physical server you plan to use. Eucalyptus is designed to run in any combination on the various physical servers. However, in order to make the best use of HA, we recommend that you give each component maximal local resource usage by installing a maximum of one component on each server. This allows for high performance and high availability.

Verify Component Disk Space

Eucalyptus components need disk space for log files, databases, buckets, and instances. The following table details the needs of each component. Verify that the machines you plan to install the components on have adequate space.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

We recommend that you choose a disk for each Walrus that is large enough to hold all objects and buckets you ever expect to have, including all images that will ever be registered to your system, plus any Amazon S3 application data. For consistent performance, we recommend that you use identical disks for the primary and secondary Walrus.

Tip: We recommend that you use LVM (Logical Volume Manager). If you run out of disk space, LVM allows you to add disks and migrate the data.

Component	Directory	Minimum Size
Cluster Controller (CLC) CLC logging	`/var/lib/eucalyptus/db` `/var/log/eucalyptus`	20GB 2GB
Walrus Walrus logging	`/var/lib/eucalyptus/bukkits` `/var/log/eucalyptus`	250GB 2GB
Storage Controller (SC)	`/var/lib/eucalyptus/volumes` (EBS storage) **Important:** This disk space on the SC is only required if you are not using a SAN driver or if you are using Direct Attached Storage (DAS). For more information, see either *Configure the Storage Controller* or *Configure the Storage Controller* (HA).	250GB
Cluster Controller (CC) CC logging	`/var/lib/eucalyptus/CC` `/var/log/eucalyptus`	5GB 2GB
Node Controller (NC) NC logging	`/var/lib/eucalyptus/instances` `/var/log/eucalyptus`	250GB 2GB

If necessary, create symbolic links to larger filesystems from the above locations. Make sure that the eucalyptus user owns the directories.

Plan Networking Modes

Eucalyptus overlays a virtual network on top of your existing network. In order to do this, Eucalyptus supports three different networking modes: Edge, Managed, and Managed (No VLAN).

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

CC-BY-SA, Eucalyptus Systems, Inc.

Each mode is designed to allow you to choose an appropriate level of security and flexibility. The purpose of these modes is to direct Eucalyptus to use different network features to manage the virtual networks that connect VMs to each other and to clients external to Eucalyptus.

A Eucalyptus installation must be compatible with local site policies and configurations (e.g., firewall rules). Eucalyptus configuration and deployment interfaces allow a wide range of options for specifying how it should be deployed. However, choosing between these options implies tradeoffs.

Your choice of networking mode depends on the following considerations:

- Do you plan to support elastic IPs and security groups?
- Do you plan to provide your own network DHCP server?

These networking features are described in the following table:

Feature	Description	Mode
Elastic IPs	Eucalyptus instances typically have two IPs associated with them: a private one and a public one. Private IPs are intended for internal communications between instances and are usually only routable within a Eucalyptus cloud. Public IPs are used for external access and are usually routable outside of Eucalyptus cloud. How these addresses are allocated and assigned to instances is determined by a networking mode. The distinction between public and private addresses becomes important in Edge, Managed, and Managed (No VLAN) modes, which support elastic IPs. With elastic IPs the user gains control over a set of static IP addresses. Once allocated to the user, those same IPs can be dynamically associated to running instances, overriding pre-assigned public IPs. This allows users to run well-known services (for example, web sites) within the Eucalyptus cloud and to assign those services fixed IPs that do not change.	Edge Managed Managed (No VLAN)
Security groups	Security groups are sets of networking rules that define the access rules for all VM instances associated with a group. For example, you can specify ingress rules, such as allowing ping (ICMP) or SSH (TCP, port 22) traffic to reach VMs in a specific security group. When you create a VM instance, unless otherwise specified at instance run-time, it is assigned to a default security group that denies incoming network traffic from all sources. Thus, to allow login and usage of a new VM instance you must authorize network access to the default security group with the `euca-authorize` command.	Edge Managed Managed (No VLAN)
VM isolation	Although network traffic between VM instances belonging to a security group is always open, Eucalyptus can enforce isolation of network traffic between different security groups. This isolation is enforced using ebtables (Edge) or VLAN tags (Managed), thus, protecting VMs from possible eavesdropping by VM instances belonging to other security groups.	Edge Managed
DHCP server	Eucalyptus assigns IP addresses to VMs in all modes.	Edge Managed Managed (No VLAN)

If Eucalyptus can control and condition the networks its components use, your deployment will support the full set of API features. However, if Eucalyptus is confined to using an existing network, some of the API features might be disabled. So, understanding and choosing the right networking configuration is an important (and complex) step in deployment planning.

Each networking mode is detailed in the following sections.

CC-BY-SA, Eucalyptus Systems, Inc.

Managed Mode

Managed mode offers the most features of the networking modes, but also carries with it the most potential constraints on the setup of the network.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

In Managed mode, Eucalyptus manages the local network of VM instances and provides all networking features Eucalyptus currently supports, including VM network isolation, security groups, elastic IPs, and metadata service.

In Managed mode, you define a large network (usually private, unroutable) from which VM instances will draw their private IP addresses. Eucalyptus maintains a DHCP server with static mappings for each VM instance that is created. When you create a new VM instance, you can specify the name of the security group to which that VM will belong. Eucalyptus then selects a subset of the entire range of IPs, to hand out to other VMs in the same security group.

You can also define a number of security groups, and use those groups to apply network ingress rules to any VM that runs within that network. In this way, Eucalyptus provides functionality similar to Amazon's security groups. In addition, the administrator can specify a pool of public IP addresses that users may allocate, then assign to VMs either at boot or dynamically at run-time. This capability is similar to Amazon's 'elastic IPs'. Eucalyptus administrators that require security groups, elastic IPs, and VM network isolation must use this mode.

Managed mode uses a Virtual LAN (VLAN) to enforce network isolation between instances in different security groups. If your underlying physical network is also using a VLAN, there can be conflicts that prevent instances from being network accessible. So you have to determine if your network between the CC and NCs is VLAN clean (that is, if your VLANs are usable by Eucalyptus). To test if the network is VLAN clean, see VLAN Preparation.

Each VM receives two IP addresses: a public IP address and a private IP address. Eucalyptus maps public IP addresses to private IP addresses. Access control is managed through security groups.

Managed Mode Requirements

- There must be an available range of IP addresses for the virtual subnets. This range must not interfere with the physical network. Typically these IP addresses are selected from the private IP ranges: 192.168.x.x, 10.x.x.x, etc.
- The network between the CC and NCs must be VLAN clean, meaning that all switch ports that Eucalyptus components are connected to will allow and forward VLAN tagged packets.
- Any firewall running on the Cluster Controller must be compatible with the dynamic changes performed by Eucalyptus when working with security groups. (Note that Eucalyptus will flush the 'filter' and 'nat' tables upon boot).
- Any DHCP server on the subnet must be configured not to serve Eucalyptus instances.
- There must be a separate Layer 2 network for each cluster in a multi-cluster setup.

Managed Mode Limitations

- The machine that hosts the CC will be a router in the data path for any VM traffic that is not 'VM private IP to VM private IP, where both VMs are in the same security group'.
- Network switch must be properly configured. For more information, see *Prepare VLAN* or *Prepare VLAN* (for HA).
- In non-HA mode, the machine that hosts the CC is a single point of failure for most VM network communication.

Managed (No VLAN) Mode

In Managed (No VLAN) mode, Eucalyptus fully manages the local VM instance network and provides all of the networking features Eucalyptus currently supports, including security groups, elastic IPs, etc. However, it does not provide VM network isolation.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Without VLAN isolation at the bridge level, it is possible in Managed (No VLAN) mode for a root user on one VM to snoop and/or interfere with the ethernet traffic of other VMs running on the same layer 2 network.

CC-BY-SA, Eucalyptus Systems, Inc.

Tip: In Managed (No VLAN) mode, VM isolation is provided by having different security groups on different subnets—this translates into Layer-3 only VM isolation.

Managed (No VLAN) Mode Requirements

- There must be an available range of IP addresses for the virtual subnets. This range must not interfere with the physical network. Typically these IP addresses are selected from the private IP ranges: 192.168.x.x, 10.x.x.x, etc.
- Any firewall running on the Cluster Controller must be compatible with the dynamic changes performed by Eucalyptus when working with security groups. (Note that Eucalyptus will flush the 'filter' and 'nat' tables upon boot).
- A range of public IP addresses must be available for use by Eucalyptus.
- The CC must have a DHCP server daemon installed that is compatible with ISC DHCP Daemon version 3.0.X.
- If you plan to set up more than one cluster, you need to have a bridge for security groups to span the clusters.

Managed (No VLAN) Mode Limitations

- Limited (Layer-3) VM isolation.

Edge Mode

Edge mode offers the most features of the networking modes. It is designed to integrate into already extant (or straightforward to deploy) underlying network topologies. However, Edge mode can impose constraints in certain environments, in which case you can choose another mode.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

In Edge networking mode, the component responsible for implementing Eucalyptus VM networking artifacts is running at the edge of a Eucalyptus deployment: the Node Controller (NC). Eucalyptus provides a stand-alone component called `eucanetd` in each NC. This componenent dynamically receives changing Eucalyptus networking views and is responsible for configuring the Linux machine on which the NC is running to reflect the latest view.

Edge networking mode integrates with your existing network infrastructure, allowing you to 'tell' Eucalyptus (through configuration parameters for Edge mode) about the existing network, which Eucalyptus then will consume when implementing the networking view.

Edge networking mode integrates with two basic types of pre-existing network setups:

- One flat IP network used to service Eucalyptus component systems, Eucalyptus VM public IPs (elastic IPs), and Eucalyptus VM private IPs.
- Two networks, one for Eucalyptus components and Eucalyptus VM public IPs, and the other for Eucalyptus VM private IPs.

Important: Edge networking mode will not set up the network from scratch as Managed and Managed (No VLAN) modes do. Instead, it integrates with networks that already exist. If the network, netmask, and router don't already exist, you must create them outside of Eucalyptus before configuring Edge.

Edge Mode Requirements

- Each NC must have an interface configured with an IP on a VM public and a VM private network (which can be the same network).
  - There must be unused IP addresses on the VM public network for Eucalyptus to assign VM elastic IPs
  - There must be unused IP addresses on the VM private network for Eucalyptus to assign VM private IPs
- There must be IP connectivity from each NC machine (where `eucanetd` runs) and the CLC, for metadata re-directs for 169.254.169.254 to the active CLC to function.
- There must be a functioning router in place for the private network. This router will be the default gateway for VM instances.
- The private and public networks can be the same network, but they can also be separate networks.

CC-BY-SA, Eucalyptus Systems, Inc.

- The Node Controllers (NCs) need a bridge configured on the private network, with the bridge interface itself having been assigned an IP from the network.
- If you're using a public network, the NCs need an interface on the public network as well (if the public and private networks are the same network, then the bridge needs an IP assigned on the network).
- If you run in multi-cluster, each cluster can use the same network as its private network, or they can use separate networks as private networks. If you use separate networks, you need to have a router in place that is configured to route traffic between the networks.
- If you use private addressing only mode, the Cloud Controller machines must have a route back to the VM private network.

Edge Mode Limitations

- All NCs must have an interface on the VM public (Elastic IP) network.
- Global network updates (such as security group rule updates, security group VM membership updates, and elastic IP updates) are be applied through an "eventually consistent" mechanism, as opposed to an "atomic" mechanism. That is, there may be a brief period of time where one NC has the new state implemented but another NC has the previous state implemented.
- Mappings between VM MAC addresses and private IPs are strictly enforced.

Plan Eucalyptus Features

Before you install Eucalyptus, we recommend that you think about the features you plan to implement with Eucalyptus. These features are detailed in the following sections.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Windows Guest OS Support

This topic details what Eucalyptus needs in order to use Windows as a guest operating system.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

- A licensed installation copy (.iso image or CD/DVD disk) of a compatible Windows OS. Eucalyptus currently supports Windows virtual machines created from Windows Server 2003 R2 Enterprise (32/64 bit); Windows Server 2008 SP2, Datacenter (32/64 bit); Windows Server 2008 R2, Datacenter; and Windows 7 Professional.
- A VNC client such as RealVNC or Virtual Manager/Virtual Viewer for initial installation. Subsequent Eucalyptus-hosted Windows instances will use RDP, but the initial installation requires VNC.

For additional Windows-related licensing information, see the following links:

- *http://technet.microsoft.com/en-us/library/dd979803.aspx*
- *http://technet.microsoft.com/en-us/library/dd878528.aspx*
- *http://technet.microsoft.com/en-us/library/dd772269.aspx*

SAN Support

Eucalyptus includes optional, subscription only support for integrating enterprise-grade SAN (Storage Area Network) hardware devices into a Eucalyptus cloud.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

SAN support extends the functionality of the Eucalyptus Storage Controller (SC) to provide a high performance data conduit between VMs running in Eucalyptus and attached SAN devices. Eucalyptus dynamically manages SAN storage without the need for the administrator to manually allocate and de-allocate storage, manage snapshots or set up data connections.

CC-BY-SA, Eucalyptus Systems, Inc.

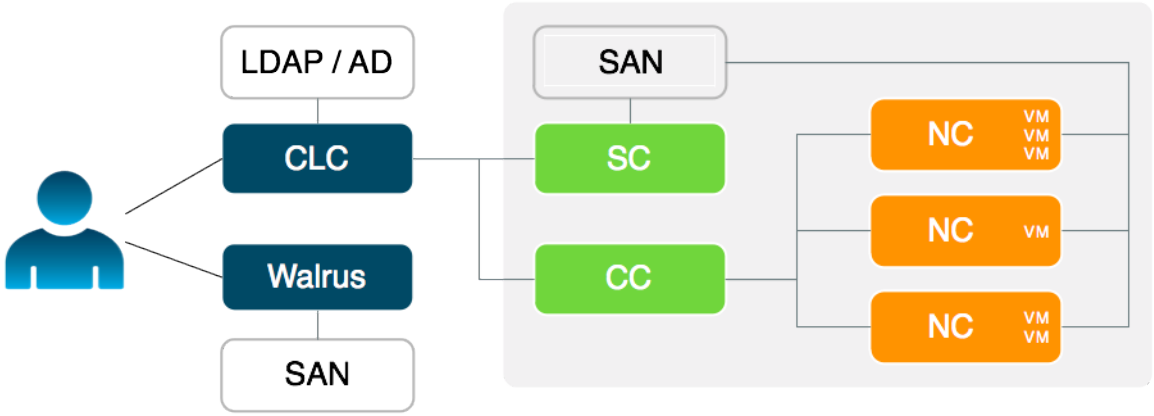

Eucalyptus with SAN support allows you to:

- Integrate Eucalyptus block storage functionality (dynamic block volumes, snapshots, creating volumes from snapshots, etc.) with existing SAN devices
- Link VMs in the Eucalyptus cloud directly to SAN devices, thereby removing I/O communication bottlenecks of the physical hardware host
- Incorporate enterprise-level SAN features (high-speed, large-capacity, reliability) to deliver a production-ready EBS (block storage) solution for the enterprise
- Attach SAN devices to Eucalyptus deployments on Xen and KVM

To use Eucalyptus with supported SAN storage, you must decide whether administrative access can be provided to Eucalyptus to control the SAN. If this is possible in your environment, Eucalyptus can automatically and dynamically manage SAN storage.

Currently, the Dell Equallogic series of SANs (PS 4000 and PS 6000), NetApp Filer FAS 2000 and FAS 6000 series and EMC VNX are supported. For Dell Equallogic, Eucalyptus requires SSH access to enable automatic provisioning. Eucalyptus will manage NetApp SANs via ONTAPI (version 7.3.3 and above). For EMC, Eucalyptus expects that the EMC NaviSecCLI software will be installed on the Storage Controller host.

SAN Support Prerequisites

Eucalyptus supports the following SAN devices:

- Dell EqualLogic, PS4000 series and PS6000 series (For more information about Dell EqualLogic SANs, go to *http://www.dell.com*)
- NetApp, FAS2000 series and FAS6000 series (For more information about NetApp SANs, go to *http://www.netapp.com*
- EMC VNX Series (For more information about EMC VNX, go to *VNX Family*

For additional information on SAN support for Eucalyptus, contact Eucalyptus Systems, Inc.

Availability Zone Support

Eucalyptus offers the ability to create multiple availability zones. In Eucalyptus, an availability zone is a partition in which there is at least one available cluster.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

CC-BY-SA, Eucalyptus Systems, Inc.

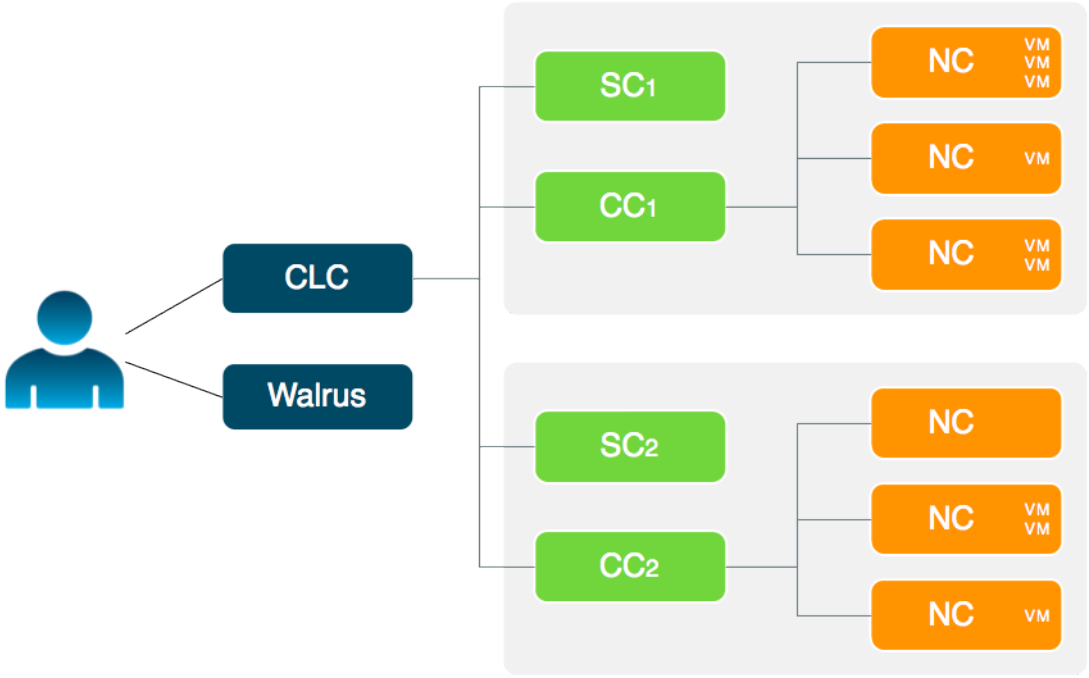

Object Storage

Eucalyptus supports Walrus and Riak CS as its object storage backend. There is no extra planning if you use Walrus. If you use Riak CS, you can use a single Riak CS cluster for several Eucalyptus clouds. Basho (the vendor of RiakCS) recommends five nodes for each Riak CS cluster. This also means that you have to set up and configure a load balancer between the Riak CS nodes and the object storage gateway (OSG).

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

High Availability Support

Eucalyptus includes the ability to run redundant, hot-swappable instances for the CLC, Walrus, CC, and SC. In a high availability (HA) configuration, a failure of any single component will not cause the system to halt.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

If your network configuration includes redundant networking hardware and routing paths, HA Eucalyptus can then tolerate a network component failure (for example, the loss of a networking switch) without halting.

The deployment choices for HA Eucalyptus are similar to a regular Eucalyptus deployment, with the following additional considerations:

- You must host redundant Eucalyptus software components on separate hardware components in order to be able to tolerate a hardware failure. If, for example, you install redundant CLCs on the same machine and the machine crashes, both CLCs will become inoperable.
- The redundant components occur in pairs, one primary, the other secondary. These components must be able to communicate with each other through the network to which they are both attached while they are running. For example, both CLC components in an HA installation must be able to exchange messages. If you use a firewall to separate them, one will not detect a failure of the other and a hot failover will not occur. This ability for pairs of components to communicate is required for the CLC, Walrus, CC, and SC for HA to operate properly.

The following image shows a single cluster deployment with the component pairs at the cloud and cluster level. The NCs are not redundant.

CC-BY-SA, Eucalyptus Systems, Inc.

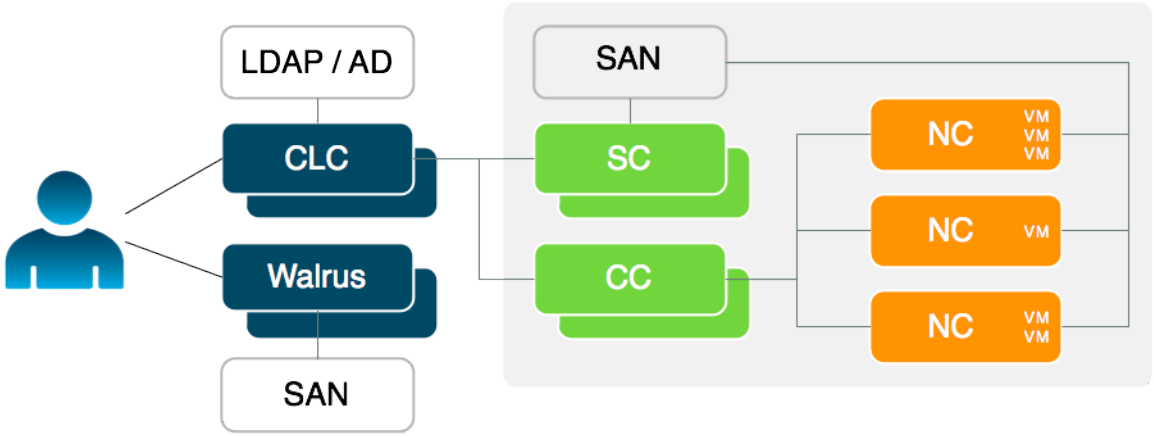

Note that the same considerations for a regular Eucalyptus deployment with respect to networking mode and components placement apply to HA Eucalyptus in addition to the need for redundant component pairs to be able to communicate. Note also that the NC components are deployed redundantly in an HA Eucalyptus deployment. If a machine running an NC fails, Eucalyptus will continue to be available for user requests. However, instances running on that specific NC will be lost.

For HA: The installation and configuration sections will note instructions specific to HA deployment by the HA icon.

Eucalyptus HA Requirements

Eucalyptus HA requires the same requirements as Eucalyptus. However, the infrastructure Eucalyptus HA will be deployed on must meet some additional requirements, listed in this topic.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Redundant Physical Servers for Eucalyptus Components

Each cloud component (CLC and OSG) and cluster component (CC and SC) in an HA deployment has a redundant hot backup. These redundant Eucalyptus components occur in pairs, and each member of a pair must be mapped to a separated physical server to ensure high availability.

Important: HA pairs must be able to connect to each other.

If the HA deployment is to be able to tolerate the failure of networking hardware, additional network interfaces are required for the physical servers that host Eucalyptus components. The physical servers hosting a CLC, Walrus, or CC must each have three network interface cards (NICs). Each remaining physical server (except the NC components) requires two NICs.

The following image shows a two-cluster deployment with the component pairs at the cloud and cluster level. The NCs are not redundant.

CC-BY-SA, Eucalyptus Systems, Inc.

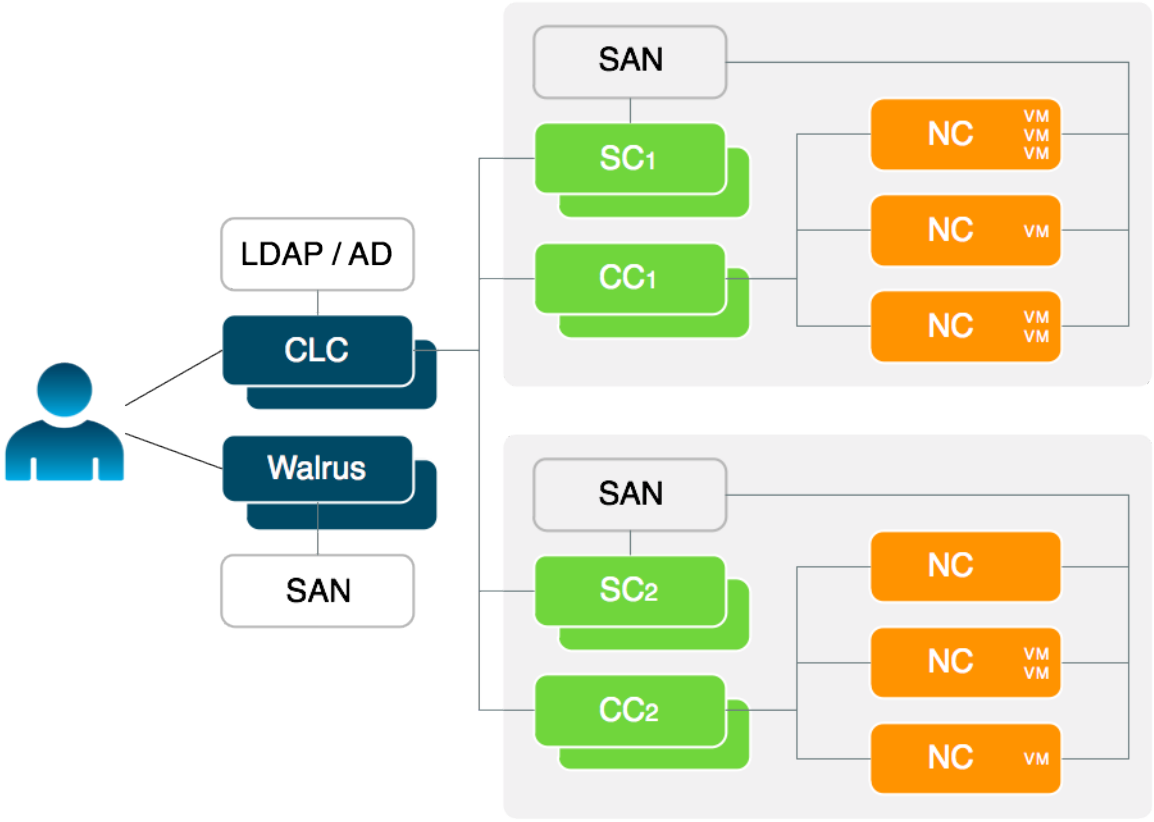

DNS Round-Robin Support

The DNS entries for the externally visible IP addresses of the physical servers hosting CLC or Walrus components must be configured to change round-robin style in an HA deployment.

Storage Mirroring

HA Eucalyptus uses a kernel-level storage technology called DRBD for storage integrity. DRBD must be configured to mirror data operations between physical servers that host Walrus components. For more information about DRBD, go to *What is DRBD*.

Storage Controllers

For HA Storage Controllers, you must be using a supported SAN. Only use HA SCs with NetApp or Equallogic drivers, not with the iSCSI or JBOD SC driver.

Eucalyptus HA Planning

High availability is the result of the combination of functionality provided by Eucalyptus and the environmental and operational support to maintain the systems proper operation.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Eucalyptus provides functionality aimed at enabling highly available deployments:

1. **Detection of hardware and network faults which impact system availability:** Availability of the system is determined by its ability to properly service a user request at a given time. The system is available when there is at least a set of functioning services to perform the operations which result from a user request (i.e., system is distributed and operations require orchestration involving some, possibly all, services in the system).
2. **Deployment of redundant services to accommodate host failure:** A failure is the observed consequence of an underlying fault which compromises the systems function in some way (possibly compromising availability).

CC-BY-SA, Eucalyptus Systems, Inc.

3. **Automated recovery from individual component failure:** Eucalyptus can take advantage of redundant host and network resources to accommodate singular failures while preserving the system's overall availability. As a result, the deployment of the system plays a large role in the level of availability that can be achieved.

To deliver services with high availability, Eucalyptus depends upon redundant hardware and network.

Considerations

A highly available deployment is able to mitigate the impact on system availability of faults from the following sources:

- **Machines hosting Eucalyptus services:** Hardware faults on machines hosting Eucalyptus services can result in component services being unavailable for use by the system or users. The state of the hosting machine is monitored by the system and determines whether it can contribute to work done. In support of high availability, you can configure redundant component services. With redundant component services, Eucalyptus can isolate and mask the a component's failure.
- **Inter-component networks:** Faults in the networks that connect the system's components to each other can prevent access to cloud resources and restrict the system's ability to process user requests. First, internal resources may become unavailable. For example, a single network outage could impact access to attached volumes or prevent access to running instances. Second, the coordination of services needed to process user requests may be impeded even if the service state is otherwise healthy.
- **User-facing network connections:** User-facing network faults can prevent access to an otherwise properly functioning system. The ability of a user to access the system is difficult to determine from the perspective of the system - can't look through the users eyes. Allowing for multiple inbound paths (for example, multiple disjoint routes) decreases the possibility of an availability-impacting outage occurring w/in the scope of the environment within which Eucalyptus is deployed. (See also: registering arbitrators)

Recommendations

To ensure availability in the face of any single failure, we recommend the following deployment strategy:

- **Host/Service Redundancy:** Each component which is registered should have a complementary service registered on a redundant host. For example, the cloud and walrus services should be installed and registered on two hosts. Additionally, for example, each partition should have two cluster controllers and storage controllers configured. Each such complementary pair of services can suffer a single outage before system availability is compromised.
- **Inter-component Network Redundancy:** Each host of a component service should have redundant and disjoint network connections to other internal component services and supporting systems (for example, SANs). The recommended approach is to have two ethernet devices (each connected to a disjoint layer-2 network) on each host and bonding the devices. Such a configuration is also suggested on node controllers. Then, the outage of a either layer-2 network or ethernet device on a host does not impact service availability or access to cloud resources.
- **User-facing Network Redundancy:** The wide area (where users are) network connection should be redundant and disjoint. Each such path should have an independent arbitrator host whose liveness (as determined by ICMP echo) is used to approximate the users' ability to access the system. Redundant network connections from the local area network to the wide area network and user reachability approximation (arbitrator)
- **System Reachability Approximation:** The wide area (where users are) network connection(s) path should have an independent host (arbitrator) whose liveness (as determined by ICMP echo) can serve as a reasonable approximation of users' ability to access the system. Ideally, the host “closest” to the user, but still within the domain of the deployment environment should be used (for example, the border gateway of the hosting AS network). With such an arbitrator host in the network path between the user and the system, a failure by the user to reach an otherwise working service and allow the system to enable the complementary service (which should have a separate network route) restoring user access.

SAN and Multipathing

Multipathing is a way to make the data path from the NC or SC to your SAN device highly available. Multipathing does this by giving the host two network paths that both lead to the same data volume. This allows the host to switch from one network path to the other, in the event that one path becomes unavailable. Essentially, multipathing decreases the likelihood that a volume will become unreachable from a host (NC). For information about configuring your SAN for multipathing, see *Configure the Storage Controller*.

CC-BY-SA, Eucalyptus Systems, Inc.

Prepare the Network

In order for Eucalyptus to function in your local environment, be sure to prepare your network. To prepare your network, perform the tasks listed in this section.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Reserve Ports

Eucalyptus components use a variety of ports to communicate. The following table lists the all of the important ports used by Eucalyptus.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Port	Description
TCP 5005	DEBUG ONLY: This port is used for debugging Eucalyptus (using the `--debug` flag).
TCP 8080	Port for the administrative web user interface. Forwards to 8443. Configurable with `euca-modify-property`.
TCP 8443	SSL port for the administrative web user interface. Configurable with `euca-modify-property`.
TCP 8772	DEBUG ONLY: JMX port. This is disabled by default, and can be enabled with the `--debug` or `--jmx` options for `CLOUD_OPTS`.
TCP 8773	Web services port for the CLC, Walrus, and SC; also used for external and internal communications by the CLC and Walrus. Configurable with `euca-modify-property`.
TCP 8774	Web services port on the CC. Configured in the `eucalyptus.conf` configuration file
TCP 8775	Web services port on the NC. Configured in the `eucalyptus.conf` configuration file.
TCP 8776	Used by the image cacher on the CC. Configured in the `eucalyptus.conf` configuration file.
TCP 8777	Database port on the CLC
TCP 8779 (or next available port, up to TCP 8849)	jGroups failure detection port on CLC, Walrus, and SC. If port 8779 is available, it will be used, otherwise, the next port in the range will be attempted until an unused port is found.
TCP 8888	The default port for the Eucalyptus Management Console. Configured in the `/etc/eucalyptus-console/console.init` file.
TCP 16514	TLS port on Node Controller, required for node migrations
UDP 7500	Port for diagnostic probing on CLC, Walrus, and SC
UDP 8773	HA membership port
UDP 8778	The bind port used to establish multicast communication
TCP/UDP 53	DNS port on the CLC

Verify Connectivity

Verify connectivity between the machines you'll be installing Eucalyptus on. Some Linux distributions provide default TCP/IP firewalling rules that limit network access to machines. Disable these default firewall settings before you install Eucalyptus components to ensure that the components can communicate with one another.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

CC-BY-SA, Eucalyptus Systems, Inc.

Note: Any firewall running on the CC must be compatible with the dynamic changes performed by Eucalyptus when working with security groups. Eucalyptus will flush the 'filter' and 'nat' tables upon boot.

Verify component connectivity by performing the following checks on the machines that will be running the listed Eucalyptus components.

1. Verify connection from an end-user to the CLC on TCP ports 8443 and 8773
2. Verify connection from an end-user to Walrus on TCP port 8773
3. Verify connection from the CLC, SC, and NC to SC on TCP port 8773
4. Verify connection from the CLC, SC, and NC to Walrus on TCP port 8773
5. Verify connection from Walrus and SC to CLC on TCP port 8777
6. Verify connection from CLC to CC on TCP port 8774
7. Verify connection from CC to NC on TCP port 8775
8. Verify connection from NC to Walrus on TCP port 8773. Or, you can verify the connection from the CC to Walrus on port TCP 8773, and from an NC to the CC on TCP port 8776
9. Verify connection from public IP addresses of Eucalyptus instances (metadata) and CC to CLC on TCP port 8773
10. Verify TCP connectivity between CLC, Walrus, and SC on TCP port 8779 (or the first available port in range 8779-8849)
11. Verify connection between CLC, Walrus, and SC on UDP port 7500
12. Verify multicast connectivity for IP address 228.7.7.3 between CLC, Walrus, and SC on UDP port 8773
13. If DNS is enabled, verify connection from an end-user and instance IPs to DNS ports
14. If you use tgt (iSCSI open source target) for EBS storage, verify connection from NC to SC on TCP port 3260
15. Test multicast connectivity between each CLC and Walrus, and SC.
    a) Clone the Eucalyptus `deveutils` repository

    ```
    git clone https://github.com/eucalyptus/deveutils
    ```

    b) Run the `network-tomography` tool on the Cloud Controller, Cluster Controller, Storage Controller, and any machines running Walrus, passing a list of IP addresses for each of these machines.

    ```
    cd deveutils/network-tomography
    ./network-tomography 192.168.51.174 192.168.51.196 192.168.51.86
    192.168.51.99
    ```

    This tool may take up to an hour to run. Check the output for reports of packet loss. If there is significant packet loss, ensure that your network is available and multicast enabled.

Prepare VLAN

Managed networking mode requires that switches and routers be “VLAN clean.” This means that switches and routers must allow and forward VLAN tagged packets. If you plan to use the Managed networking mode, you can verify that the network is VLAN clean between machines running Eucalyptus components by performing the following test.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Tip: You only need to read this section if you are using Managed mode. If you aren’t using Managed mode, skip this section.

1. Choose two IP addresses from the subnet you plan to use with Eucalyptus, one VLAN tag from the range of VLANs that you plan to use with Eucalyptus, and the network interface that will connect your planned CC and NC servers. The examples in this section use the IP addresses 192.168.1.1 and 192.168.1.2, VLAN tag 10, and network interface eth3, respectively.
2. On the planned CC server, choose the interface on the local Ethernet and run:

CC-BY-SA, Eucalyptus Systems, Inc.

```
vconfig add eth3 10
ifconfig eth3.10 192.168.1.1 up
```

3. On a planned NC server, choose the interface on the local network and run:

```
vconfig add eth3 10
ifconfig eth3.10 192.168.1.2 up
```

4. On the NC, ping the CC:

```
ping 192.168.1.1
```

5. On the CC, ping the NC:

```
ping 192.168.1.2
```

   - If this VLAN clean test fails, configure your switch to forward VLAN tagged packets. If it is a managed switch, see your switch's documentation to determine how to do this.
   - If the VLAN clean test passes, continue with the following steps to remove the test interfaces.

6. On the CC, remove the test interface by running:

```
vconfig rem eth3.10
```

7. On the planned NC, run:

```
vconfig rem eth3.10
```

Configure Dependencies

Before you install Eucalyptus HA, make sure you have the following dependencies installed and configured.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Before you install Eucalyptus, make sure you have the following dependencies installed and configured.

Configure Bridges

For Managed (No VLAN) and EDGE modes, you must configure a Linux ethernet bridge on all NC machines. This bridge connects your local ethernet adapter to the cluster network. Under normal operation, NCs will attach virtual machine instances to this bridge when the instances are booted.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

To configure a bridge in CentOS 6 or RHEL6, you need to create a file with bridge configuration (for example, ifcfg-brX) and modify the file for the physical interface (for example, ifcfg-ethX). The following steps describe how to set up a bridge on both CentOS 6 and RHEL 6. We show examples for configuring bridge devices that either obtain IP addresses using DHCP or statically.

1. Install the `bridge-utils` package.

```
yum install bridge-utils
```

2. Go to the `/etc/sysconfig/network-scripts` directory:

```
cd /etc/sysconfig/network-scripts
```

3. Open the network script for the device you are adding to the bridge and add your bridge device to it. The edited file should look similar to the following:

CC-BY-SA, Eucalyptus Systems, Inc.

```
DEVICE=eth0
# change the hardware address to match the hardware address your NIC uses
HWADDR=00:16:76:D6:C9:45
ONBOOT=yes
BRIDGE=br0
NM_CONTROLLED=no
```

4. Create a new network script in the `/etc/sysconfig/network-scripts` directory called `ifcfg-br0` or something similar. The br0 is the name of the bridge, but this can be anything as long as the name of the file is the same as the `DEVICE` parameter, and the name is specified correctly in the previously created physical interface configuration (ifcfg-ethX).

   - If you are using DHCP, the configuration will look similar to:

     ```
     DEVICE=br0
     TYPE=Bridge
     BOOTPROTO=dhcp
     ONBOOT=yes
     DELAY=0
     ```

   - If you are using a static IP address, the configuration will look similar to:

     ```
     DEVICE=br0
     TYPE=Bridge
     BOOTPROTO=static
     IPADDR=<static_IP_address>
     NETMASK=<netmask>
     GATEWAY=<gateway>
     ONBOOT=yes
     ```

5. Enter the following command:

   ```
   service network restart
   ```

Disable the Firewall

If you have existing firewall rules on your hosts, you should disable the firewall in order to install Eucalyptus. You should re-enable it after installation.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Tip: If you do not have a firewall enabled, skip this step.

1. To disable your firewall:
   a) Run the command `system-config-firewall-tui`
   b) Turn off the **Enabled** check box.
2. Repeat on each host that will run a Eucalyptus component: Cloud Controller, Walrus, Cluster Controller, Storage Controller, and Node Controllers.

Configure SELinux

Security-enabled Linux (SELinux) is security feature for Linux that allows you to set access control through policies. Eucalyptus is not compatible with SELinux.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

To configure SELinux to allow Eucalyptus access:

1. Open `/etc/selinux/config` and edit the line `SELINUX=enforcing` to `SELINUX=permissive`.

CC-BY-SA, Eucalyptus Systems, Inc.

2. Save the file.
3. Run the following command:

```
setenforce 0
```

Configure NTP

Eucalyptus requires that each machine have the Network Time Protocol (NTP) daemon started and configured to run automatically on reboot.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

To use NTP:

1. Install NTP on the machines that will host Eucalyptus components.

```
yum install ntp
```

2. Open the `/etc/ntp.conf` file and add NTP servers, as in the following example.

```
server 0.pool.ntp.org
server 1.pool.ntp.org
server 2.pool.ntp.org
```

3. Save and close the file.
4. Configure NTP to run at reboot.

```
chkconfig ntpd on
```

5. Start NTP.

```
service ntpd start
```

6. Synchronize your server.

```
ntpdate -u <your_ntp_server>
```

7. Synchronize your system clock, so that when your system is rebooted, it does not get out of sync.

```
hwclock --systohc
```

8. Repeat on each host that will run a Eucalyptus component.

Configure an MTA

All machines running the Cloud Controller must run a mail transport agent server (MTA) on port 25. Eucalyptus uses the MTA to deliver or relay email messages to cloud users' email addresses.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

You can use Sendmail, Exim, postfix, or something simpler. The MTA server does not have to be able to receive incoming mail.

Many Linux distributions satisfy this requirement with their default MTA. For details about configuring your MTA, go to the documentation for your specific product.

To test your mail relay for localhost, send email to yourself from the terminal using `mail`.

Enable Packet Routing

Edit the `sysctl.conf` on each machine you plan to install the Cluster Controller (CC) component and the Node Controller (NC) on.

CC-BY-SA, Eucalyptus Systems, Inc.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

In the `sysctl.conf` file, set the following parameters and values:

1. Enable IP forwarding.

```
net.ipv4.ip_forward = 1
```

2. Enable the bridge to forward traffic based on iptables rules.

```
net.bridge.bridge-nf-call-iptables = 1
```

Install Eucalyptus

Eucalyptus installation packages are available for CentOS 6 and RHEL 6. The following sections show installation steps on each supported Linux distribution.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Eucalyptus Subscription allows you access to additional software modules. If you are a subscriber, you will receive an entitlement certificate and a private key that allow you to download Eucalyptus subscription modules. You will also receive a GPG public key to be used to verify the Eucalyptus software's integrity. The files will come in the form of a platform specific package.

Important: For Eucalyptus HA, we recommend that you install each Eucalyptus component on a separate host. For example, if you are installing CLC, Walrus, CC, and SC, you will install each of these components on a separate host. You will also install each secondary component (the secondary CLC, Walrus, CC, and SC) on a separate host. In this case, you will need eight machines. Each additional cluster needs four more machines for its CCs and SCs. This does not account for NCs, which are not redundant.

Software Signing

We use a number of GPG keys to sign our software packages and package repositories. The necessary public keys are provided with the relevant products and can be used to automatically verify software updates. You can also verify the packages or package repositories manually using the keys on this page.

Use the `rpm --checksig` command on a download file to verify a RPM package for an HP Helion Eucalyptus product. For example:

```
rpm --checksig -v myfilename.rpm
```

Follow the procedure detailed on Debian's *SecureApt* web page to verify a deb package for an HP Helion Eucalyptus product.

Please do not use package signing keys to encrypt email messages.

The following keys are used for signing Eucalyptus software:

c1240596: Eucalyptus Systems, Inc. (release key) <security@eucalyptus.com>

This key is used for signing HP Helion Eucalyptus products released after July 2011 and their updates.

- *Download from Helion Eucalyptus*
- *Download from pgp.mit.edu*
- Fingerprint: `8639 B2D2 11BB 930D 16A2 D654 BE26 4D09 C124 0596`

CC-BY-SA, Eucalyptus Systems, Inc.

0260cf4e: Eucalyptus Systems, Inc. (pre-release key) <security@eucalyptus.com>

This key is used for signing HP Helion Eucalyptus pre-release products due for release after July 2011.

- *Download from Helion Eucalyptus*
- *Download from pgp.mit.edu*
- Fingerprint: `7363 5F5A 9531 308B E83D 3413 8B94 DFB5 0260 CF4E`

9d7b073c: Eucalyptus Systems, Inc. (nightly release key) <security@eucalyptus.com>

This key is used for signing nightly builds of HP Helion Eucalyptus products published after July 2011.

- *Download from Helion Eucalyptus*
- *Download from pgp.mit.edu*
- Fingerprint: `708C DB7C B08A 204E C4CB FA6A 4882 7EF9 9D7B 073C`

Install Eucalyptus from Release Packages

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

To install Eucalyptus on servers running CentOS 6 or RHEL 6:

1. Configure the Eucalyptus package repository on each host that will run a Eucalyptus component:

   ```
   yum install
   http://downloads.eucalyptus.com/software/eucalyptus/4.1/centos/6/x86_64/eucalyptus-release-4.1.el6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

2. If you are not a Eucalyptus subscriber, skip this step. If you are a Eucalyptus subscriber, you should have received an rpm package file containing your license for subscription-only components. Install that package, along with the Eucalyptus subscription package, on each host that will run a Eucalyptus component, as follows:

   ```
   yum install eucalyptus-enterprise-license*.noarch.rpm \
   http://subscription.eucalyptus.com/eucalyptus-enterprise-release-4.1-1.el6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

3. Configure the Euca2ools package repository on each host that will run a Eucalyptus component or Euca2ools:

   ```
   yum install
   http://downloads.eucalyptus.com/software/euca2ools/3.2/centos/6/x86_64/euca2ools-release-3.2.el6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

4. Configure the EPEL package repository on each host that will run a Eucalyptus component or Euca2ools:

   ```
   yum install
   http://downloads.eucalyptus.com/software/eucalyptus/4.1/centos/6/x86_64/epel-release-6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

5. If you are using Walrus as your object storage backend configure the ELRepo repository on each machine that will run Walrus. Otherwise, skip this step.

   ```
   yum install
   http://downloads.eucalyptus.com/software/eucalyptus/4.1/centos/6/x86_64/elrepo-release-6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

6. For RHEL 6 systems, you must enable the Optional repository in Red Hat Network for each NC, as follows:
   a) Go to *http://rhn.redhat.com* and navigate to the system that will run the NC.
   b) Click **Alter Channel Subscriptions**.

CC-BY-SA, Eucalyptus Systems, Inc.

c) Make sure the **RHEL Server Optional** checkbox is checked.
d) Click **Change Subscriptions**.

7. a) Install the Eucalyptus node controller software on each planned NC host:

```
yum install eucalyptus-nc
```

b) Check that the KVM device node has proper permissions.

Run the following command:

```
ls -l /dev/kvm
```

Verify the output shows that the device node is owned by user root and group kvm.

```
crw-rw-rw- 1 root kvm 10, 232 Nov 30 10:27 /dev/kvm
```

If your kvm device node does not have proper permissions, you need to reboot your NC host.

8. If you plan to run in *Edge networking mode*, install the package for Edge support on each planned NC host.

```
yum install eucanetd
```

9. On each planned CLC host, install the Eucalyptus cloud controller software.

```
yum install eucalyptus-cloud
```

10. Install the Imaging Worker image package on the machine hosting the primary CLC:

```
yum install eucalyptus-service-image
```

11. Install the software for the remaining Eucalyptus components. The following example shows most components being installed on the same host. We recommend that you use different hosts for each component:

```
yum install eucalyptus-cc eucalyptus-sc eucalyptus-walrus
```

12. If you would like Load Balancer support enabled in your Cloud, you will need to install the Load Balancer image package on the machine hosting the primary CLC:

```
yum install eucalyptus-service-image
```

13. If you are a subscriber and use SAN, run the appropriate command for your device on each machine hosting a CLC:

For EMC SAN:

```
yum install eucalyptus-enterprise-storage-san-emc-libs
```

For EqualLogic SAN:

```
yum install eucalyptus-enterprise-storage-san-equallogic-libs
```

For NetApp SAN:

```
yum install eucalyptus-enterprise-storage-san-netapp-libs
```

14. If you are a subscriber and use SAN, run the appropriate command for your device on each machine hosting a SC:

For EMC SAN:

```
yum install eucalyptus-enterprise-storage-san-emc
```

Important: To use Eucalyptus with EMC SAN support, you must have the `NaviCLI-Linux-64-latest.rpm` package installed on each SC. This package is not supplied with Eucalyptus, please see your SAN vendor if it is not already installed.

For EqualLogic SAN:

```
yum install eucalyptus-enterprise-storage-san-equallogic
```

CC-BY-SA, Eucalyptus Systems, Inc.

For NetApp SAN:

```
yum install eucalyptus-enterprise-storage-san-netapp
```

15. After you have installed Eucalyptus, test multicast connectivity between each CLC and OSG, and SC.
    a) Clone the Eucalyptus `deveutils` repository

    ```
    git clone https://github.com/eucalyptus/deveutils
    ```

    b) Run the `network-tomography` tool on the Cloud Controller, Cluster Controller, Storage Controller, and any machines running Walrus, passing a list of IP addresses for each of these machines.

    ```
    cd deveutils/network-tomography
    ./network-tomography 192.168.51.174 192.168.51.196 192.168.51.86
    192.168.51.99
    ```

    This tool may take up to an hour to run. Check the output for reports of packet loss. If there is significant packet loss, ensure that your network is available and multicast enabled.

    c) Repeat these tasks with the secondary controllers.

Your installation is complete.

You are now ready to *Configure Eucalyptus*.

Install Eucalyptus from Nightly Packages

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Important: Eucalyptus nightly packages are latest Eucalyptus builds. They should be considered unstable/"bleeding edge" software and should not be installed in production. In addition, upgrades from nightlies to released software are not supported.

To install Eucalyptus nightly builds on servers running CentOS 6 or RHEL 6:

1. On all servers, run the following commands:

   ```
   yum install http://downloads.eucalyptus.com/
   software/eucalyptus/nightly/4.1/centos/6/x86_64/eucalyptus-release-4.1.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

2. On all systems that will run either Eucalyptus or Euca2ools, run the following commands:

   ```
   yum install http://downloads.eucalyptus.com/software/
   euca2ools/3.2/centos/6/x86_64/euca2ools-release-3.2.el6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

3. If you are using Walrus as your object storage backend, install the ELRepo repository on each machine that will run Walrus. Otherwise, skip this step.

   ```
   yum install http://downloads.eucalyptus.com/
   software/eucalyptus/nightly/4.1/centos/6/x86_64/elrepo-release-6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

4. Configure the EPEL package repository:

   ```
   yum install
   http://downloads.eucalyptus.com/software/eucalyptus/nightly/4.1/centos/6/x86_64/epel-release-6.noarch.rpm
   ```

   Enter `y` when prompted to install this package.

5. On all servers, enter:

CC-BY-SA, Eucalyptus Systems, Inc.

```
yum update
```

6. Install Eucalyptus packages. The following example shows most components being installed all on the same server. You can use different servers for each component.

```
yum install eucalyptus-cloud
yum install eucalyptus-cc eucalyptus-sc eucalyptus-walrus
```

For HA: For Eucalyptus HA, you must install these packages on pairs of systems. For instance, "eucalyptus-cloud" is installed on the primary CLC and the secondary CLC.

7. Install the Imaging Worker image package on the machine hosting the primary CLC:

```
yum install eucalyptus-service-image
```

8. If you would like Load Balancer support in your cloud, you will need to install the Load Balancer image package on the machine hosting the primary CLC:

```
yum install eucalyptus-service-image
```

9. On each planned NC server, install the NC package:

```
yum install eucalyptus-nc
```

10. If you plan to run in *Edge networking mode*, install the package for Edge support on each planned NC host.

```
yum install eucanetd
```

Your installation is complete.

You are now ready to *Configure Eucalyptus*.

Configure Eucalyptus

This topic describes the parameters you need to set in order to launch Eucalyptus for the first time.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

The first launch of Eucalyptus is different than a restart of a previously running Eucalyptus deployment in that it sets up the security mechanisms that will be used by the installation to ensure system integrity.

Eucalyptus configuration is stored in a text file, `/etc/eucalyptus/eucalyptus.conf`, that contains key-value pairs specifying various configuration parameters. Eucalyptus reads this file when it launches and when various forms of reset commands are sent it the Eucalyptus components.

Important: Perform the following tasks after you install Eucalyptus software, but before you start the Eucalyptus services.

Configure Network Modes

This section provides detailed configuration instructions for each of the four Eucalyptus networking modes. Eucalyptus requires network connectivity between its clients (end-users) and the cloud components (CC, CLC, and Walrus).

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

- In Edge mode, most networking configuration is handled through settings in a global Cloud Controller (CLC) property file. For more information, see *Configure for Edge Mode* (or *Configure for Edge Mode* for HA).
- In Managed and Managed (No VLAN) modes, traffic to instances pass through the CC. In these two modes clients must be able to connect to the Cluster Controller (CC).

CC-BY-SA, Eucalyptus Systems, Inc.

The `/etc/eucalyptus/eucalyptus.conf` file contains all network-related options in the “Networking Configuration” section. These options use the prefix VNET_. The most commonly used VNET options are described in the following table. The set of networking settings that apply to a cloud varies based on its networking mode. Each setting in this section lists the modes in which it applies. Unless otherwise noted, all of these settings apply only to CCs.

The most commonly used VNET options are described in the following table.

Option	Description	Modes

`VNET_ADDRESSPERNET`	This option controls how many VM instances can simultaneously be part of an individual user's security group. This option is set to a power of 2 (8, 16, 32, 64, etc.) but it should never be less than 8 and it cannot be larger than: (the total number of available IP addresses - 2).
	
	This option is used with VNET_NETMASK to determine how the IP addresses that are available to VMs are distributed among security groups. VMs within a single security group can communicate directly. Communication between VMs within a security group and clients or VMs in other security groups is controlled by a set of firewall rules. For example, setting
	
	`VNET_NETMASK="255.255.0.0"`
	`VNET_ADDRESSPERNET="32"`
	
	defines a netmask of 255.255.0.0 that uses 16 bits of the IP address to specify a network number. The remaining 16 bits specify valid IP addresses for that network meaning that 2^16 = 65536 IP addresses are assignable on the network. Setting `VNET_ADDRESSPERNET="32"` tells Eucalyptus that each security group can have at most 32 VMs in it (each VM getting its own IP address). Further, it stipulates that at most 2046 security groups can be active at the same time since 65536 / 32 = 2048. Eucalyptus reserves two security groups for its own use.
	
	In addition to subnets at Layer 3, Eucalyptus uses VLANs at Layer 2 in the networking stack to ensure isolation (Managed mode only).
		Managed, Managed (No VLAN)

`VNET_BRIDGE`	On an NC, this is the name of the bridge interface to which instances' network interfaces should attach. A physical interface that can reach the CC must be attached to this bridge. Common setting for KVM is `br0`.	Edge (on NC) Managed (No VLAN)
`VNET_DHCPDAEMON`	The ISC DHCP executable to use. This is set to a distro-dependent value by packaging. The internal default is `/usr/sbin/dhcpd3`.	Edge (on NC) Managed Managed (No VLAN)
`VNET_DHCPUSER`	The user the DHCP daemon runs as on your distribution. For CentOS 6 and RHEL 6, this is typically `root`. Default: `dhcpd`	Managed Managed (No VLAN)

CC-BY-SA, Eucalyptus Systems, Inc.

Option	Description	Modes
`VNET_DNS`	The address of the DNS server to supply to instances in DHCP responses. Example: `VNET_DNS="173.205.188.129"`	Managed Managed (No VLAN)
`VNET_LOCALIP`	By default the CC automatically determines which IP address to use when setting up tunnels to other CCs. Set this to the IP address that other CCs can use to reach this CC if tunneling does not work.	Managed Managed (No VLAN)
`VNET_MACPREFIX`	This option is used to specify a prefix for MAC addresses generated by Eucalyptus for VM instances. The prefix has to be in the form `HH:HH` where H is a hexadecimal digit. Example: `VNET_MACPREFIX="D0:D0"`	Managed, Managed (No VLAN)
`VNET_MODE`	The networking mode in which to run. The same mode must be specified on all CCs and NCs in your cloud. Valid values: `EDGE` `MANAGED`, `MANAGED-NOVLAN`,	All
`VNET_PRIVINTERFACE`	The name of the network interface that is on the same network as the NCs. In Managed and Managed (No VLAN) modes this must be a bridge for instances in different clusters but in the same security group to be able to reach one another with their private addresses. Default: `eth0`	Edge (on NC) Managed
`VNET_PUBINTERFACE`	**On a CC**, this is the name of the network interface that is connected to the "public" network. **On an NC**, this is the name of the network interface that is connected to the same network as the CC. Depending on the hypervisor's configuration this may be a bridge or a physical interface that is attached to the bridge. Default: `eth0`	Edge (on NC) Managed Managed (No VLAN)
`VNET_PUBLICIPS`	A space-separated list of individual and/or hyphenated ranges of public IP addresses to assign to instances. If you do not set a value for this option, all instances will receive only private IP addresses. Example: `VNET_PUBLICIPS= "173.205.188.140-173.205.188.254"` **Tip:** To offer more public IPs, you can span subnets. However you must list each subnet range separately. For example: `"10.133.82.50-10.133.82.254 10.133.83.0-10.133.83.254"`	Managed Managed (No VLAN)

CC-BY-SA, Eucalyptus Systems, Inc.

Option	Description	Modes
VNET_SUBNET, VNET_NETMASK	These options control the internal private network used by instances within Eucalyptus. Eucalyptus assigns a distinct subnet of private IP addresses to each security group. This setting dictates how many addresses each of these subnets should contain. Specify a power of 2 between 16 and 2048. This is directly related, though not equal, to the number of instances that can reside in each security group. Eucalyptus reserves eleven addresses per security group.	Managed, Managed (No VLAN)

Configure for Edge Mode

To configure Eucalyptus for Edge mode, you must edit `eucalyptus.conf` on the Cluster Controller (CC) and Node Controller (NC) hosts. You must also create a JSON file and upload it the Cloud Controller (CLC).

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Configure for Managed Mode

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

In Managed mode, Eucalyptus manages the local network of VM instances and provides all networking features Eucalyptus currently supports, including VM network isolation, security groups, elastic IPs, and metadata service. Configure each CC to use an Ethernet device that lies within the same broadcast domain as all of its NCs.

Important: In Managed mode, each security group requires a separate subnet and a separate VLAN that Eucalyptus controls and maintains. So the underlying physical network must be "VLAN clean." For more information about VLAN clean, see *Prepare VLAN*.

To configure for Managed mode:

CLC Configuration

No network configuration required.

CC Configuration

Important:

We recommend allowing the CC to act as the gateway for NCs, in Managed mode. To do so, ensure that traffic from all NCs (on private network) is allowed to be masqueraded on the CC, and set the output interface to the public interface of the CC. You can do this using the following `iptables` command:

```
iptables -t nat -A POSTROUTING -s 10.101.104.0/16 -o em1 -j MASQUERADE
```

Where `10.101.104.0/16` is the private network containing all NCs, and `em1` is the public interface set on the CC.

1. Log in to the CC and open the `/etc/eucalyptus/eucalyptus.conf` file.
2. Go to the **Network Configuration** section, uncomment and set the following:

```
VNET_MODE="MANAGED"

VNET_SUBNET="<subnet for instances' private IPs. Example: 192.168.0.0>"
VNET_NETMASK="<your netmask for the vnet_subnet. Example: 255.255.0.0>"
VNET_DNS="<your DNS server's IP>"
VNET_ADDRSPERNET="<# of simultaneous instances per security group>"

VNET_PUBLICIPS="<your_free_public_ip1 your_free_public_ip2 …>"
```

CC-BY-SA, Eucalyptus Systems, Inc.

```
VNET_LOCALIP="<the IP of the local interface on the cc that is reachable from
 CLC>"

VNET_DHCPDAEMON="<path to DHCP daemon binary. Example: /usr/sbin/dhcpd3>"

VNET_DHCPUSER="<DHCP user name. Example: dhcpd>"
```

3. If your NCs are not reachable from end-users directly and the CC has two (or more) Ethernet devices of which one connects to the client/public network and one connects to the NC network, or the single Ethernet device that the CC uses to connect to both clients and NCs is NOT 'eth0', then you must also uncomment and set:

```
VNET_PRIVINTERFACE="<Ethernet device on same network as NCs. Example: eth1>"

VNET_PUBINTERFACE="<Ethernet device on 'public' network. Example: eth0>"
```

4. Save the file.
5. Repeat on each CC in your system.

Important: Each CC must have the same configuration with the exception of the `VNET_LOCALIP` value, which should be machine-specific. In a multi-cluster configuration, you must set `VNET_PUBLICIPS` identically on all CCs.

NC Configuration

Important:

We recommend allowing the CC to act as the gateway for NCs, in Managed mode. To do so, ensure that traffic from all NCs (on private network) is allowed to be masqueraded on the CC, and set the output interface to the public interface of the CC. You can do this using the following `iptables` command:

```
iptables -t nat -A POSTROUTING -s 10.101.104.0/16 -o em1 -j MASQUERADE
```

Where `10.101.104.0/16` is the private network containing all NCs, and `em1` is the public interface set on the CC.

1. Log into an NC machine and open the `/etc/eucalyptus/eucalyptus.conf` file.
2. Go to the **Network Configuration** section, uncomment and set the following:

```
VNET_MODE="MANAGED"
VNET_PUBINTERFACE="<Ethernet device/bridge reachable from cc machine. Example:
 eth0>"
```

3. Save the file.
4. Repeat on each NC.

Configure for Managed (No-VLAN) Mode

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

In Managed (No-VLAN) mode, Eucalyptus does not use VLANs to isolate the network bridges attached to VMs from each other. Configure each CC to use an Ethernet device that lies within the same broadcast domain as all of its NCs.

To configure for Managed (No VLAN) mode:

CLC Configuration

No network configuration required.

CC-BY-SA, Eucalyptus Systems, Inc.

CC Configuration

Important: You must set `VNET_PUBLICIPS` identically on all CCs in a multi-cluster configuration.

1. Log in to the CC and open the `/etc/eucalyptus/eucalyptus.conf` file.
2. Go to the **Network Configuration** section, uncomment and set the following:

```
VNET_MODE="MANAGED-NOVLAN"
VNET_SUBNET="[Subnet for VMs private IPs. Example: 192.168.0.0]"
VNET_NETMASK="[Netmask for the vnet_subnet. Example: 255.255.0.0]"
VNET_DNS="[DNS server IP]"
VNET_ADDRSPERNET="[Number of simultaneous instances per security group]"
VNET_PUBLICIPS="[Free public IP 1] [Free public IP 2] ..."
VNET_LOCALIP="[IP address that other CCs can use to reach this CC]"
VNET_DHCPDAEMON="[Path to DHCP daemon binary. Example: /usr/sbin/dhcpd3]"
VNET_DHCPUSER='[DHCP user. Example: dhcpd]"
```

3. If your NCs are not reachable from end-users directly and the CC has two (or more) Ethernet devices of which one connects to the client/public network and one connects to the NC network, or the single Ethernet device that the CC uses to connect to both clients and NCs is NOT ‘eth0’, then you must also uncomment and set:

```
VNET_PRIVINTERFACE="[Ethernet device on same network as NCs. Example: eth1]"

VNET_PUBINTERFACE="[Ethernet device on ‘public’ network. Example: eth0]"
```

4. Save the file.
5. Repeat on each CC in your system.

Important: Each CC must have the same configuration with the exception of the `VNET_LOCALIP` value, which should be machine-specific.

NC Configuration

1. Log into an NC machine and open the `/etc/eucalyptus/eucalyptus.conf` file.
2. Go to the **Network Configuration** section, uncomment and set the following:

```
VNET_MODE="MANAGED-NOVLAN"
VNET_BRIDGE="[bridge name. Example: br0]"
```

3. Save the file.
4. Repeat on each NC.

Configure Loop Devices

In order to start new instances, Eucalyptus needs a sufficient number of loop devices to use for SC and NC components. An SC with insufficient loop devices fails to create new EBS volumes. An NC with insufficient loop devices fails to start new instances.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Eucalyptus installs with a default loop device amount of 256. If you want to change this number, perform the following steps. Otherwise, skip this section.

Tip: We recommend that you err on the side of configuring too many loop devices. Too many loop devices result in a minor amount of memory tie-up and some clutter added to the system's `/dev` directory. Too few loop devices make Eucalyptus unable to use all of a system's resources. We recommend a minimum of 50 loop devices. If you have fewer than 50, the startup script will complain.

1. Log in to the SC server and open the `/etc/eucalyptus/eucalyptus.conf` file.

CC-BY-SA, Eucalyptus Systems, Inc.

2. Uncomment the following line:

```
# CREATE_SC_LOOP_DEVICES=256
```

3. Replace 256 with the number of loop devices.
4. Repeat for each SC on your system.
5. Log in to an NC server and open the `/etc/eucalyptus/eucalyptus.conf` file.
6. Uncomment the following line:

```
# CREATE_NC_LOOP_DEVICES=256
```

7. Replace 256 with the number of loop devices.
8. Repeat for each NC on your system.

Configure Multi-Cluster Networking

Eucalyptus supports multiple clusters within a single Eucalyptus cloud. This topic briefly describes how Eucalyptus manages the networking aspect of a multi-cluster setup.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

In Edge networking mode, Eucalyptus does not perform any special configuration for a multi-cluster setup. In Managed and Managed (No VLAN) modes, Eucalyptus sets up Layer 2 Tunneling Protocol (L2TP) between your clusters. This means that virtual machines in the same security group, but distributed across clusters (potentially each in their own broadcast domain), can communicate with one another. Eucalyptus uses the VTun package to handle all L2TP tunnels between clusters. If VTun is installed on each of your CCs, multi-cluster tunneling is automatically handled by each CC.

Depending on the networking mode and network topology, keep the following network configuration considerations in mind.

Managed Mode:	During normal operation, you will see many tunnel interfaces being created and destroyed as virtual networks are constructed and torn down.
Managed (No VLAN) Mode:	In order for VTun tunneling to work in this mode, you must configure each CC with a bridge as its primary, private interface (`VNET_PRIVINTERFACE`). All traffic from nodes in one cluster to nodes in another cluster is routed through the CCs. Each cluster requires that the interface that faces the nodes for the CC (the private interface) be a bridge device for the nodes themselves.
Managed Mode and Managed (No VLAN) Mode:	The CC attempts to auto-discover its list of local IP addresses upon startup, but if the IP that was used to register the CC is not locally available, you can override the CC's notion of 'self' by setting the `VNET_LOCALIP` variable in the `eucalyptus.conf` file.

Important: Note the following:

- You must set VNET_PUBLICIPS identically on all CCs in a multi-cluster configuration.
- To enable tunneling, set `DISABLE_TUNNELING=N` in `eucalyptus.conf` on both CC hosts.
- When L2TP tunneling is enabled in a multi-cluster setup, make sure that you are using different IP ranges for the nodes in each cluster.
- Do not run two CCs in the same broadcast domain with tunneling enabled, as this will potentially lead to a broadcast storm as tunnels start forwarding packets in a loop on your local network. Please disable tunneling by setting `DISABLE_TUNNELING=Y` in `eucalyptus.conf` on both CC hosts.

Configure the Firewall

This topic provides guidelines for restricting network access and managing iptables rules.

CC-BY-SA, Eucalyptus Systems, Inc.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Restricting Network Access

This section provides basic guidance on setting up a firewall around your Eucalyptus components. It is not intended to be exhaustive.

On CLC, Walrus, and SC, you should allow for the following jGroups traffic:

- TCP connections between CLC, Walrus, and SC on port 8779 (or the first available port in range 8779-8849)
- UDP connections between CLC, Walrus, and SC on port 7500
- Multicast connections between CLC, Walrus, and SC to IP 228.7.7.3 on UDP port 8773

On the CLC, you should additionally allow the following connections:

- TCP connections from end-users on ports 8773 and 8443
- TCP connections from CC and Eucalyptus instances (public IPs) on port 8773 (for metadata service)
- TCP connections from Walrus and SC on port 8777
- End-user and instance connections to DNS ports

On the CC, you should ensure that all firewall rules are compatible with the dynamic changes performed by Eucalyptus, described in the section below. You should also allow the following connections:

- TCP connections from CLC on port 8774
- TCP connections from NC on port 8776, if CC image proxying is enabled

On Walrus, you should also allow the following connections:

- TCP connections from end-users on port 8773
- TCP connections from SC and NC on port 8773
- TCP connections from CC on port 8773, if CC image proxying is enabled

On the SC, you should also allow the following connections:

- TCP connections from CLC and NC on TCP port 8773
- TCP connections from NC on TCP port 3260, if tgt (iSCSI open source target) is used for EBS storage

On the NC, you should allow the following connections:

- TCP connections from CC on port 8775
- TCP connections from other NCs on port 16514
- DHCP traffic forwarding to VMs
- Traffic forwarding to and from instances' private IP addresses

Managing iptables Rules for the CC

In Managed and Managed (No VLAN) modes, Eucalyptus flushes the CC's iptables rules for both `filter` and `nat`, then it sets the default policy for the `FORWARD` chain in `filter` to `DROP`. At run time, the CC adds and removes rules from `FORWARD` as users add and remove ingress rules from their active security groups. In addition, the `nat` table is configured to allow VMs access to the external network using IP masquerading, and dynamically adds/removes rules in the `nat` table as users assign and unassign public IPs to VMs at instance boot or run-time.

If you have rules you want to apply on the CC, make the following edit on the CC before you start Eucalyptus or while Eucalyptus is stopped:

```
iptables-save > /etc/eucalyptus/iptables-preload
```

Caution: Performing this operation to define special iptables rules that are loaded when Eucalyptus starts could cause Eucalyptus VM networking to fail. We recommend that you only do this if you are completely sure that it will not interfere with the operation of Eucalyptus.

CC-BY-SA, Eucalyptus Systems, Inc.

Start Eucalyptus

Start the Eucalyptus components in the order presented in this section.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Make sure that each host you installed a Eucalyptus component on resolves to an IP address. Edit the `/etc/hosts` file if necessary.

Note: Eucalyptus 4.1.2 requires version 7 of the Java Virtual Machine. Make sure that your CLOUD_OPTS settings in the /etc/eucalyptus/eucalyptus.conf file either do not set `--java-home`, or that `--java-home` points to a version 7 JVM. This needs to happen before services are started but after the upgraded packages are installed.

Start the CLC Pairs

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

1. Log in to the primary Cloud Controller (CLC).
2. Enter the following command to initialize the primary CLC:

   Note: Make sure that the `eucalyptus-cloud` process is not running prior to executing this command.

   ```
   /usr/sbin/euca_conf --initialize
   ```

   Note: This command might take a minute or more to finish.

3. Enter the following command to start the primary CLC:

   ```
   service eucalyptus-cloud start
   ```

4. Start the secondary CLC. Do not initialize the secondary CLC. Just start it.

Start the Walrus Pairs

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

To start Walrus:

1. Log in to the Walrus server and enter the following command:

   ```
   service eucalyptus-cloud start
   ```

2. Repeat this task on the secondary Walrus.

Start the CC Pairs

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

To start the CC:

1. Log in to the CC server and enter the following:

   ```
   service eucalyptus-cc start
   ```

CC-BY-SA, Eucalyptus Systems, Inc.

2. If you have a multi-cluster setup, repeat this step on the CC in each cluster.
3. Repeat this task on the secondary CC in each cluster.

Start the SC Pairs

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

If the SC component is not on the same machine as the CLC, do the following:

1. Log in to the SC server and enter the following command:

```
service eucalyptus-cloud start
```

Important: If you are re-installing the SC, please restart the tgt (iSCSI open source target) daemon.

2. If you have a multi-cluster setup, repeat this step on the SC in each cluster.
3. Repeat this task on the secondary SC in each cluster.

Start the NCs

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

1. Log in to an NC server and enter the following command:

```
service eucalyptus-nc start
```

2. If you are running in Edge networking mode, start the Edge component.

```
service eucanetd start
```

3. Repeat for each NC server.

Verify the Startup

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

At this point, all Eucalyptus components are enabled and starting up. Some of these services perform intensive initialization at start-up, particularly the first time they are started. You might have to wait a few minutes until they are fully operational.

One quick way to determine if the components are running is to run netstat on the various hosts and look to see when the service ports are allocated to a process. Specifically, the CLC, Walrus, and the SC allocate ports 8773. The CC listens to port 8774, and the NC uses port 8775.

Verify that everything has started without error. Expected outcomes include:

- The CLC is listening on ports 8443 and 8773
- Walrus is listening on port 8773
- The SC is listening on port 8773
- The CC is listening on port 8774
- The NCs are listening on port 8775
- Log files are being written to `/var/log/eucalyptus/`

Register Eucalyptus

After you start Eucalyptus for the first time, register the Eucalyptus components as described in this section.

CC-BY-SA, Eucalyptus Systems, Inc.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Eucalyptus implements a secure protocol for registering separate components so that the overall system can't be tricked into including a component run by an unauthorized administrator or user. You only need to register components the first time Eucalyptus is started after it was installed.

Most registration commands run on the CLC server. NCs, however, are registered on each CC. You must register each NC on every CC for the cluster on which the NC participates.

Note that each registration command will attempt an SSH as root to the remote physical host where the registering component is assumed to be running. The registration command also contacts the component so it must be running at the time of the command is issued. If a password is required to allow SSH access, the command will prompt the user for it.

Except for NCs, each registration command requires four pieces of information:

- The **component** (`--register-XYZ`) you are registering, because this affects where the commands must be executed.
- The **partition** (`--partition`) the component will belong to. The partition is the same thing as availability zone in AWS.
- The **name** (`--component`) you assign to the component, up to 256 characters. This is the name used to identify the component in a human-friendly way. This name is also used when reporting system state changes which require administrator attention. This name must be globally-unique with respect to other component registrations. To ensure this uniqueness, we recommend using a combination of the component type (CLC, SC, CC, etc) and system hostname or IP address when you choose your component names. For example: `clc-eucahost15` or `clc-192.168.0.15`.
- The **IP address** (`--host`) of the service being registered. The host must be specified by IP address to function correctly.

  Note: You must specify public IP addresses.

NCs only have two pieces of information: component name and IP address.

Note: We recommend that you use IP addresses rather than host names when registering Eucalyptus components. If you do use hostnames, the underlying IP address may not be a site-local, any-cast, loopback, link-local, or multicast address.

Note: Once you've registered a Eucalyptus component with a host name, to avoid connectivity issues, do not change the host name's underlying IP address.

Register the Secondary CLC

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Log in to the primary CLC and enter the following command to register the secondary CLC:

```
/usr/sbin/euca_conf --register-cloud --partition eucalyptus
--host [Secondary_CLC_IP] --component [CLC_Name]
```

The `partition` name for the CLC has to be `eucalyptus`. The `component` name is a unique name for this particular component: we recommend a format such as `CLC-[hostname]`.

Register Services

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

CC-BY-SA, Eucalyptus Systems, Inc.

To register the user-facing services:

1. On the primary CLC server, enter the following command:

```
euca_conf --register-service -T user-api -H <host_url> -N <name>
```

   For example:

```
euca_conf --register-service -T user-api -H 10.111.1.135 -N API_135
```

2. Repeat on the secondary CLC.

Register Walrus Pairs

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

To register Walrus:

1. On the CLC server, enter the following command:

```
/usr/sbin/euca_conf --register-walrusbackend --partition walrus --host
[walrus_IP_address] --component [walrus_name]
```

   The `partition` name for Walrus has to be `walrus`. Like the CLC, the component name is a unique name for this particular component: we recommend a format such as `walrus-[hostname]`.

2. Register the secondary Walrus the same way, using the secondary Walrus IP address and secondary Walrus name. Use the same partition name as the primary Walrus

Register the CC Pairs

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

To register the CC:

1. On the CLC, enter the following command:

```
/usr/sbin/euca_conf --register-cluster --partition [partition_name]
 --host [CC_IP_address] --component [cc_name]
```

   We recommend that you set the `partition` name to a descriptive name for the availability zone controlled by the CC. For example: `cluster01`.

   The `component` must be a unique name. We recommend that you use a short-hand name of the hostname or IP address of the machine, like `cc-[hostname]` or `cc-[IP address]`.

2. Repeat for each cluster, replacing the CC name, partition name, CC IP address, and CC name.
3. Register the secondary CC the same way, replacing the CC IP address and CC name, but using the same partition name as the primary CC.

Register the SC Pairs

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

To register the SC:

1. On the CLC, enter the following command:

Note: We recommend that you use IP addresses instead of DNS names when registering Eucalyptus components.

CC-BY-SA, Eucalyptus Systems, Inc.

```
/usr/sbin/euca_conf --register-sc --partition [partition_name] --host
[SC_IP_address]
 --component [SC_name]
```

An SC must have the same `partition` name as the CC in the same cluster. Like the other components, the `component` is a unique name for this particular component: we recommend a format such as `sc-[hostname]`.

Warning: Newly registered SCs will be in the BROKEN state until they are explicitly configured to use a backend storage provider. The output of the registration for the first SC registered in a partition will look like:

```
SERVICE storage         PARTI00         SC71            BROKEN     37
 http://192.168.51.71:8773/services/Storage
arn:euca:eucalyptus:PARTI00:storage:SC71/
Registered the first storage controller in partition 'PARTI00'.  You must
 choose a storage back end with ``euca-modify-property -p
PARTI00.storage.blockstoragemanager=$BACKEND''
```

This is completely normal and simply indicates that further action must be taken to configure the SC before it will become fully functional. For information about configuring the SC, see *Configure the Runtime Environment->Configure the Storage Controller*

2. Register the secondary SC the same way, using the secondary SC IP address and SC name, but using the same partition name as the primary SC.
3. Repeat for each cluster, replacing the SC name, partition name, SC IP address, and SC name.

Register the NCs

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Important: If you are using host names rather than IP addresses when registering your NCs, ensure that DNS is working properly, or populate /etc/hosts for all nodes in a cluster.

1. On a CC, register all NCs using the following command with the IP address of each NC server:

```
/usr/sbin/euca_conf --register-nodes "[node0_IP_address] ... [nodeN_IP_address]"
```

2. Repeat the previous step on the secondary CC.
3. Repeat the previous steps on each cluster in your cloud.

> The IP addresses of the NCs are space delimited, as in the following example:
>
> ```
> /usr/sbin/euca_conf --register-nodes "192.168.71.154 192.168.71.155
>  192.168.71.159"
> ```

Register Arbitrators

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Eucalyptus uses a periodic ICMP echo test to an Arbitrator. This test approximates an end user's ability to access the system. If Eucalyptus determines that it cannot reach the host associated with a registered Arbitrator, all Eucalyptus services operating on that host attempt to failover to the alternate hosts running those services.

For HA: In HA, you can register each Arbitrator service on the primary and secondary CLC and Walrus. If you are using either Managed or Managed (No VLAN) mode, you can also register Arbitrator services on both the primary CC and the secondary CC.

CC-BY-SA, Eucalyptus Systems, Inc.

We recommend that you register more than one Arbitrator for each Eucalyptus component. This will allow for normal outages and maintenance. There is no limit on the number of Arbitrators on a CLC and a Walrus. You can only register up to three on a CC.

Register an Arbitrator service on each host that has a cloud component (CLC or Walrus) installed. An Arbitrator is a host-wide component: when an Arbitrator is registered on a host, it is registered with all cloud components enabled on that host. A separate arbitrator has to be registered per each network entity that needs to be monitored from the host.

To register an Arbitrator:

1. Log in to the primary CLC.
2. Enter the following command to register an arbitrator:

   ```
   /usr/sbin/euca_conf --register-arbitrator --partition [ID]
    --component [ID] --host [target_IP]>
   ```

   where:

   - `[ID]` is a globally unique ID that identifies an Arbitrator. Note that you must use the same [ID] as both a partition and component ID.
   - `[target_IP]` is the IP of the machine running the Eucalyptus component that will run the Arbitrator.

   For example:

   ```
   euca_conf --register-arbitrator --partition EXAMPLE_ARB --component EXAMPLE_ARB
    --host 192.168.1.10
   ```

3. Repeat for the secondary CLC and for both Walrus servers.
4. Define the gateway for each Arbitrator:

   ```
   /usr/sbin/euca-modify-property -p <ID>.arbitrator.gatewayhost=<gateway>
   ```

   where:

   - `<ID>` is the globally unique ID of the registered Arbitrator.
   - `<gateway>` is an external hostname or IP address used to approximate connectivity to the end user.

   For example:

   ```
   euca-modify-property -p EXAMPLE_ARB.arbitrator.gatewayhost=192.168.1.1
   ```

5. Repeat for each registered Arbitrator.
6. To register on each CC, log in to the primary CC, and open the `/etc/eucalyptus/eucalyptus.conf` file.
7. Provide a list of Arbitrators (up to three) as values for the `CC_ARBITRATORS` property. For example:

   ```
   CC_ARBITRATORS="192.168.48.11 192.168.48.12"
   ```

8. Save the file and restart the CC.

   ```
   service eucalyptus-cc restart
   ```

9. Repeat on the secondary CC.

   In the following example, the primary CLC is on <CLC_host_p>, the secondary CLC is on <CLC_host_s>, the primary Walrus is on <Walrus_host_p>, and the secondary Walrus is on <Walrus_host_s>.

   ```
   /usr/sbin/euca_conf --register-arbitrator --host <CLC_host_p>
   --component ARB00 --partition ARB00
   /usr/sbin/euca_conf --register-arbitrator --host <CLC_host_p>
   --component ARB01 --partition ARB01
   /usr/sbin/euca_conf --register-arbitrator --host <Walrus_host_p>
   --component ARB02 --partition ARB02
   ```

CC-BY-SA, Eucalyptus Systems, Inc.

```
/usr/sbin/euca_conf --register-arbitrator --host <Walrus_host_p>
--component ARB03 --partition ARB03
/usr/sbin/euca_conf --register-arbitrator --host <CLC_host_s>
--component ARB04 --partition ARB04
/usr/sbin/euca_conf --register-arbitrator --host <CLC_host_s>
--component ARB05 --partition ARB05
/usr/sbin/euca_conf --register-arbitrator --host <Walrus_host_s>
--component ARB06 --partition ARB06
/usr/sbin/euca_conf --register-arbitrator --host <Walrus_host_s>
--component ARB07 --partition ARB07
```

Configure the Runtime Environment

After Eucalyptus is installed and registered, perform the tasks in this section to configure the runtime environment.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Generate Administrator Credentials

Now that you have installed and configured Eucalyptus, you're ready to start using it. To do so, you must generate credentials.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Important: When you run the `euca_conf --get-credentials` command, you are requesting the access and secret keys and an X.509 certificate and key. You cannot retrieve an existing X.509 certificate and key. You can only generate a new pair.

To generate a set of credentials:

1. Generate administrator credentials.

```
/usr/sbin/euca_conf --get-credentials admin.zip

unzip admin.zip
```

2. Source the `eucarc` file.

```
source eucarc
```

You are now able to run Eucalyptus commands.

Upload the JSON Configuration File (Edge Mode Networking)

If you've selected edge mode networking, you will need to upload the JSON configuration file you configured earlier in the installation process.

To upload the JSON file with the Edge mode networking configuration:

Important: This step can only be run after getting your credentials in *Generate Administrator Credentials*.

Run the following command to upload the configuration file to the CLC (with valid eucalyptus/admin credentials sourced):

```
euca-modify-property -f
cloud.network.network_configuration=</path/to/your/json_config_file>
```

CC-BY-SA, Eucalyptus Systems, Inc.

Configure Object Storage Gateway

This topic details how to configure OSG for either Walrus or Riak.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

You must configure an object storage gateway (OSG) before it will function properly. OSGs that have been registered but not properly configured will be listed in the `BROKEN` state when listed with the euca-describe-services command. For example:

```
SERVICE objectstorage objectstorage osg-192.168.1.16 BROKEN 23
http://192.168.1.16:8773/services/objectstorage
arn:euca:bootstrap:objectstorage:objectstorage:osg-192.168.1.16/
```

Configure OSG for Walrus

1. On the primary CLC, enter `walrus` as the storage provider using the `euca-modify-property` command.
   ```
   euca-modify-property -p objectstorage.providerclient=walrus
   ```
2. Repeat on the secondary CLC.
3. Check that the OSG is enabled.
   ```
   euca-describe-services
   ```
   If the state appears as `DISABLED` or `BROKEN`, check the cloud-*.log files in the `/var/log/eucalyptus` directory. A `DISABLED` state generally indicates that there is a problem with your network or credentials.

Configure OSG for Riak

1. On the primary CLC:
   a) Enter `riakcs` as the storage provider using the `euca-modify-property` command.
   ```
   euca-modify-property -p objectstorage.providerclient=riakcs
   ```
   b) Specify the RiakCS/S3 endpoint that you want to use with Eucalyptus. For example:
   ```
   euca-modify-property -p
   objectstorage.s3provider.s3endpoint=riakcs-01.riakcs-cluster.myorg.com
   ```
   c) Provide your RiakCS credentials to access your RiakCS installation:
   ```
   euca-modify-property -p objectstorage.s3provider.s3accesskey=<RiakCS access
    key>
   euca-modify-property -p objectstorage.s3provider.s3secretkey=<RiakCS secret
    key>
   ```
2. Repeat on the secondary CLC.
3. After successful configuration, check to ensure that the state of the OSG is `ENABLED` by running the `euca-describe-services` command. For example:
   ```
   SERVICE objectstorage objectstorage osg-192.168.1.16 ENABLED 23
   http://192.168.1.16:8773/services/objectstorage
   arn:euca:bootstrap:objectstorage:objectstorage:osg-192.168.1.16/
   ```
   If the state appears as `DISABLED` or `BROKEN`, check the cloud-*.log files in the `/var/log/eucalyptus` directory. A `DISABLED` state generally indicates that there is a problem with your network or credentials.

Configure the Storage Controller

Eucalyptus offers SAN support for Eucalyptus block storage (EBS). Eucalyptus directs the Storage Controller (SC) to manage any supported SAN devices.

CC-BY-SA, Eucalyptus Systems, Inc.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Eucalyptus automatically creates and tears down volumes, snapshots, and data connections from guest instances. The administrator does not need to pre-allocate volumes or LUNs for Eucalyptus.

Eucalyptus currently offers several backend providers for the SC:

- Overlay
- DAS
- Equallogic
- Netapp
- EMC-VNX

The SC must be configured explicitly upon registration. This is a change from previous versions (pre-3.2) of Eucalyptus, which would configure themselves to a default configuration using a tgtd-based filesystem-backed storage controller to provide volumes and snapshots directly from the SC. As of version 3.2, SCs automatically go to the BROKEN state after being registered with the CLC and will remain in that state until the administrator explicitly configures the SC by telling it which backend storage provider to use.

You can check the state of a storage controller by running

```
euca-describe-services -E
```

and note the state and status message of the SC(s). The output for an unconfigured SC will look like:

```
SERVICE storage         PARTI00         SC71            BROKEN     37
http://192.168.51.71:8773/services/Storage
arn:euca:eucalyptus:PARTI00:storage:SC71/
SERVICEEVENT 6c1f7a0a-21c9-496c-bb79-23ddd5749222
arn:euca:eucalyptus:PARTI00:storage:SC71/
SERVICEEVENT 6c1f7a0a-21c9-496c-bb79-23ddd5749222 ERROR
SERVICEEVENT 6c1f7a0a-21c9-496c-bb79-23ddd5749222 Sun Nov 18 22:11:13 PST 2012
SERVICEEVENT 6c1f7a0a-21c9-496c-bb79-23ddd5749222 SC blockstorageamanger not
configured. Found empty or unset manager(unset). Legal values are: das,overlay
```

Note the error above: `SC blockstoragemanager not configured. Found empty or unset manager(unset). Legal values are: das,overlay`.

This indicates that the SC is not yet configured. It can be configured by setting the `[partition].storage.blockstoragemanager` property to either 'das' or 'overlay'.

If you have installed the Eucalyptus Enterprise packages for your SAN, you will also see additional options in the output line above, and can set the block storage manager to 'netapp','emc-vnx-flare31','emc-vnx', or 'equallogic' as appropriate.

You can verify that the SC blockstoragemanager is unset using:

```
euca-describe-properties | grep blockstorage
```

To configure SAN support, follow the steps for your desired backend storage device: *Open-Source ISCSI Filesystem-backed*, *Dell Equallogic*, *JBOD*, *Netapp*, or *EMC VNX*.

Configuring the SC to use the local filesystem (Overlay)

This was the default configuration option for the SC in pre-3.2 Eucalyptus. In this configuration the SC itself hosts the volume and snapshots for EBS and stores them as files on the local filesystem. It uses standard linux iSCSI tools to serve the volumes to instances running on NCs.

1. Configure the SC to use the OverlayManager for storage.

   ```
   euca-modify-property -p <partition>.storage.blockstoragemanager=overlay
   ```

   The output of the command should be similar to:

   ```
   PROPERTY PARTI00.storage.blockstoragemanager overlay was <unset>
   ```

CC-BY-SA, Eucalyptus Systems, Inc.

2. Verify that the property value is now: 'overlay'

```
euca-describe-properties | grep blockstorage
```

Enable Dell Equallogic SANs

1. Configure the SC to use the EquallogicManager for storage.

```
euca-modify-property -p <partition>.storage.blockstoragemanager=equallogic
```

   The output of the command should be similar to:

```
PROPERTY PARTI00.storage.blockstoragemanager equallogic was <unset>
```

2. Verify that the property value is now: 'equallogic'

```
euca-describe-properties | grep blockstorage
```

3. On the CLC, run the following command to verify that the SC is listed; note that it may be in the `BROKEN` state:

```
euca_conf --list-scs
```

4. On the primary CLC, enable SAN support in Eucalyptus by entering your SAN's hostname or IP address, the username, and password:

```
euca-modify-property -p [partition_name].storage.sanhost=[SAN_IP_address]
euca-modify-property -p [partition_name].storage.sanuser=[SAN_admin_user_name]
euca-modify-property -p
[partition_name].storage.sanpassword=[SAN_admin_password]
```

   If you have multiple management IP addresses for the SAN adapter, provide a comma-delimited list of IP addresses to the `[partition_name].storage.sanhost` property.

Your Equallogic SAN is now ready to use with Eucalyptus.

Enable Direct Attached Storage (JBOD) SANs

Important: Direct Attached Storage still requires that `/var/lib/eucalyptus/volumes` has enough space for locally cached snapshots.

1. Configure the SC to use the (Direct Attached Storage) DASManager for storage.

```
euca-modify-property -p <partition>.storage.blockstoragemanager=das
```

   The output of the command should be similar to:

```
PROPERTY PARTI00.storage.blockstoragemanager das was <unset>
```

2. Verify that the property value is now: 'das'

```
euca-describe-properties | grep blockstorage
```

3. On the CLC, run the following command to verify that the SC is listed; note that it may be in the `BROKEN` state:

```
euca_conf --list-scs
```

4. On the primary CLC, set the DAS device name property. The device name can be either a raw device (/dev/sdX, for example), or the name of an existing Linux LVM volume group.

```
euca-modify-property -p <cluster name>.storage.dasdevice=<device name>
```

   For example:

```
euca-modify-property -p cluster0.storage.dasdevice=/dev/sdb
```

Your SAN is now ready to use with Eucalyptus.

CC-BY-SA, Eucalyptus Systems, Inc.

Enable NetApp SANs

Eucalyptus supports both NetApp Clustered ONTAP and traditional 7-mode SANs. NetApp Vservers and 7-mode Filers (FAS 2000 and FAS 600 series) are managed by Eucalyptus using NetApp Manageability Software Development Kit (NMSDK) and Data ONTAP APIs. This section covers enabling both NetApp Clustered ONTAP and traditional 7-mode SANs.

Enable NetApp 7-mode SANs

To configure NetApp 7-mode Filer and enable the SAN in Eucalyptus:

1. Verify Data ONTAP version for the 7-mode Filer is 7.3.3 or later.
2. Verify SSL access by typing `secureadmin status`
3. If SSL is marked inactive, enable with `secureadmin setup ssl` and generate a new certificate.
4. Turn on SSL access with `options httpd.admin.ssl.enable on`
5. Enable the iSCSI service on the NetApp device with `option iscsi.enable on` or `option licensed_feature.iscsi.enable on` if you have an embedded license on your array.
6. Turn on the iSCSI service with `iscsi start`
7. Enable the iSCSI service on the NetApp device with `enable iscsi service`
8. Verify that an aggregate with sufficient spare capacity exists.
   - If you have SSH access to the NetApp Filer, enter `aggr show_space`.
   - If an aggregate with spare capacity does not exist, create one using the `aggr create` command.
9. Verify that you have a license for FlexClone installed. At the shell prompt, enter `license` to see the list of all installed licenses.
10. Verify that administrator account credentials for NetApp Filer are available to be configured in Eucalyptus. If not, create a new administrator account for use by Eucalyptus
11. Configure the SC to use the NetappManager for storage.

    ```
    euca-modify-property -p <partition>.storage.blockstoragemanager=netapp
    ```

    The output of the command should be similar to:

    ```
    PROPERTY <partition>.storage.blockstoragemanager netapp was <unset>
    ```

12. Verify that the propery value is now: 'netapp'

    ```
    euca-describe-properties | grep blockstorage
    ```

13. On the CLC, run the following command to verify that the SC is listed; note that it may be in the `BROKEN` state:

    ```
    euca_conf --list-scs
    ```

14. Wait for the SC to transition to the NOTREADY or DISABLED state.
15. On the primary CLC, enable NetApp SAN support in Eucalyptus by entering the Filer's hostname or IP address, the username and password of the administrator account, and CHAP username.

    Note: Eucalyptus uses Challenge Handshake Authentication Protocol (CHAP) for disk operations. The CHAP username can be any value, however it should be unique when sharing a NetApp Filer across multiple Eucalyptus clusters.

    Note: CHAP support for NetApp has been added in Eucalyptus 3.3. An SC will not transition to ENABLED state until the CHAP username is configured.

    ```
    euca-modify-property -p <partition>.storage.sanhost=<Filer_IP_address>
    euca-modify-property -p <partition>.storage.sanuser=<Filer_admin_username>
    euca-modify-property -p <partition>.storage.sanpassword=<Filer_admin_password>
    euca-modify-property -p <partition>.storage.chapuser=<Chap_username>
    ```

16. Wait for the SC to transition to the ENABLED state.

CC-BY-SA, Eucalyptus Systems, Inc.

Note: The SC must be in the ENABLED state before configuring the following properties.

17. If no aggregate is set, Eucalyptus will query the NetApp Filer for all available aggregates and use the one that has the highest capacity (free space) by default. To make Eucalyptus use specific aggregate(s) configure the following property:

```
euca-modify-property -p
<partition>.storage.aggregate=<aggregate_1_name,aggregate_2_name,...>
```

If you want Eucalyptus to use the smallest aggregate first configure the following property:

```
euca-modify-property -p <partition>.storage.uselargestaggregate=false
```

18. Set the iSCSI data IP on the ENABLED CLC. This IP is used by NCs to perform disk operations on the Filer.

Note: Filer IP address can be used as the data port IP. If this is not set, Eucalyptus will automatically use the Filer IP address/hostname.

Note: Eucalyptus does not support Multipath I/O for NetApp 7-mode Filers.

```
euca-modify-property -p <partition>.storage.ncpaths=<ip>
```

19. Set the iSCSI data IP on the ENABLED CLC. This IP is used by the SC to perform disk operations on the Filer. The SC connects to the Filer in order to transfer snapshots to Walrus during snapshot operations.

Note: The Filer IP address can be used as the data port IP. If this is not set, Eucalyptus will automatically use the Filer IP address/hostname.

Note: Eucalyptus does not support Multipath I/O for NetApp 7-mode Filers.

```
euca-modify-property -p <partition>.storage.scpaths=<ip>
```

Your Netapp 7-mode SAN is now ready to use with Eucalyptus.

Enable NetApp Clustered Data ONTAP SAN

Eucalyptus integrates with NetApp Clustered ONTAP SAN by operating against a Vserver. SC must be configured to operate against Vserver contained in the NetApp Clustered ONTAP environment.

For more information on NetApp Clustered Data ONTAP, see *Clustered Data ONTAP 8.1 and 8.1.1: An Introduction*.

To configure NetApp Vserver and enable the SAN in Eucalyptus:

1. Verify Clustered Data ONTAP version for the SAN is 8.1.1 or later.
2. Verify that FlexClone and iSCSI licenses are installed on the SAN.
3. Verify that a Vserver with iSCSI data protocol is available for use by Eucalyptus.
4. Verify that Vserver administration is delegated to a user with administrative privileges for that Vserver. If not, create a new Vserver administrator account for use by Eucalyptus.
5. Verify that a management (only) Logical Interface (LIF) is configured for the Vserver and an IP address or hostname is assigned to it.
6. Verify that data LIFs are configured on the Vserver.
7. Verify that one or more aggregates with sufficient spare capacity exists.
8. Verify the network connectivity between Eucalyptus components and the Vserver. The SC must be able communicate with the Vserver over both management and data LIFs. The NC must be able to communicate with the Vserver using the data LIFs.
9. Configure the SC to use the NetApp SAN for storage:

```
euca-modify-property -p <partition>.storage.blockstoragemanager=netapp
```

CC-BY-SA, Eucalyptus Systems, Inc.

The output of the command should be similar to:

```
PROPERTY <partition>.storage.blockstoragemanager netapp was <unset>
```

10. Verify that the propery value is now: 'netapp'

```
euca-describe-properties | grep blockstorage
```

11. On the CLC, run the following command to verify that the SC is listed; note that it may be in the BROKEN state:

```
euca_conf --list-scs
```

12. Wait for the SC to transition to NOTREADY or DISABLED states.

13. On the primary CLC, enable NetApp SAN support in Eucalyptus by entering the Vserver's hostname or IP address, the username and password of the administrator account, CHAP username and Vserver name.

Note: Eucalyptus uses Challenge Handshake Authentication Protocol (CHAP) for disk operations. The CHAP username can be any value, however it should be unique when sharing a NetApp Vserver across multiple Eucalyptus clusters.

Note: CHAP support for NetApp has been added in Eucalyptus 3.3. The SC will not transition to ENABLED state until the CHAP username is configured.

```
euca-modify-property -p <partition>.storage.sanhost=<Vserver_IP_address>
euca-modify-property -p <partition>.storage.sanuser=<Vserver_admin_username>
euca-modify-property -p <partition>.storage.sanpassword=<Vserver_admin_password>
euca-modify-property -p <partition>.storage.chapuser=<Chap_username>
```

Note: The following command may fail if tried immediately after configuring the block storage manager. Retry the command a few times, pausing for a few seconds after each retry:

```
euca-modify-property -p <partition>.storage.vservername=<Vserver_name>
```

14. Wait for the SC to transition to ENABLED state.

Note: The SC must be in the ENABLED state before configuring the following properties.

15. If no aggregate is set, Eucalyptus will query the NetApp Vserver for all available aggregates and use the one that has the highest capacity (free space) by default. To make Eucalyptus use specific aggregate(s) configure the following property:

```
euca-modify-property -p <partition>.storage.aggregate=<aggregate_1_name,
       aggregate_2_name,...>
```

If you want Eucalyptus to use the smallest aggregate first configure the following property:

```
euca-modify-property -p <partition>.storage.uselargestaggregate=false
```

16. Set an IP address for the iSCSI data LIF on the ENABLED CLC. This is used for NCs performing disk operations on the Vserver. If you want to configure multiple IPs, see see *Configure NetApp Multipathing*.

```
euca-modify-property -p <partition>.storage.ncpaths=<ip>
```

17. Set an IP address for the iSCSI data LIF on the ENABLED CLC. This is used by the SC for performing disk operations on the Vserver. The SC connects to the data LIFs on the Vserver in order to transfer snapshots to Walrus during snapshot operations. If you want to configure multiple IPs, see *Configure NetApp Multipathing*.

```
euca-modify-property -p <partition>.storage.scpaths=<ip>
```

Your NetApp Clustered Data ONTAP SAN is now ready to use with Eucalyptus.

CC-BY-SA, Eucalyptus Systems, Inc.

Enable EMC VNX SANs

This adapter uses the newer VNX-Snapshot feature available on VNX devices running FLARE v5.32 or later that have a VNX-Snapshot license. This adapter also requires the Navisphere Secure CLI to be installed on the SCs. The Navisphere CLI must be version 7.32.0.5.54 or later.

Important: You must create a Clone Private LUN (CPL) of at least 1GB on each SP. For more information on creating private LUNs, see the VNX documentation.

1. We assume that the Navisphere CLI is installed in `/opt/Navisphere` on the SC.

   Important: Eucalyptus currently supports version 7.32.0.5.54 or later of the Navisphere CLI.

2. Verify that the CLI is installed and can communicate with the VNX from the SCs.

   On each SC that you are configuring, test the naviseccli command as follows:

   ```
   /opt/Navisphere/bin/naviseccli -User <your SAN username> -Password <your SAN
   password> -Scope 0 -Address <management port IP> connection -pingnode -address
    <a data port IP on your VNX>
   ```

   Verify that the command runs successfully and the ping gets replies from the SAN.

3. On the CLC, run the following command to verify that the SC is listed; note that it may be in the `BROKEN` state:

   ```
   euca_conf --list-scs
   ```

4. Configure the SC to use the EMC VNX VNX-Snapshot-based manager for storage.

   ```
   euca-modify-property -p <partition>.storage.blockstoragemanager=emc-vnx
   ```

   The output of the command should be similar to:

   ```
   PROPERTY PARTI00.storage.blockstoragemanager emc-vnx was <unset>
   ```

5. Check the SC to be sure that it has transitioned out of the BROKEN state and is in either NOTREADY or DISABLED before configuring the rest of the properties for the SC. The following commands should be run on the ENABLED CLC to configure the SC.

   On the ENABLED CLC, run:

   ```
   euca_conf --list-scs
   ```

6. On the primary CLC, enable SAN support in Eucalyptus by entering your SAN's hostname or IP address, the username, and password:

   ```
   euca-modify-property -p [partition_name].storage.sanhost=[SAN_IP_address]
   euca-modify-property -p [partition_name].storage.sanuser=[SAN_admin_user_name]
   euca-modify-property -p
   [partition_name].storage.sanpassword=[SAN_admin_password]
   ```

   If you have multiple management IP addresses for the SAN adapter, provide a comma-delimited list of IP addresses to the `[partition_name].storage.sanhost` property.

7. On the ENABLED CLC, set the login scope for the command line access. For most installs, the login scope will be `0`, which indicates a global login scope for the device. `1` indicates a local scope. `2` indicates LDAP authentication for the SAN device. Use login scope value of `2` only if your SAN is configured to use LDAP authentication and you have an admin user configured to use LDAP.

   ```
   euca-modify-property -p <partition_name>.storage.loginscope=<login_scope>
   ```

8. On the ENABLED CLC, set the username for the Challenge Handshake Authentication Protocol (CHAP). This can be any value, however it should be unique when sharing VNX on multiple Eucalyptus clusters.

   ```
   euca-modify-property -p <partition_name>.storage.chapuser=<chap_username>
   ```

CC-BY-SA, Eucalyptus Systems, Inc.

9. On the ENABLED CLC, set the value for the unique storage pool that you have configured to use with the SC.

```
euca-modify-property -p <partition_name>.storage.storagepool=0
```

10. On the ENABLED CLC, set the iSCSI data port IP for NCs to use to perform disk operations on the SAN. If you want to configure multiple IPs, see *Configure EMC VNX Multipathing*.

```
euca-modify-property -p <partition_name>.storage.ncpaths=<ip>
```

11. On the ENABLED CLC, set the iSCSI data port IP for SCs to use to perform disk operations on the SAN. The SCs connect to the data ports on the SAN in order to transfer snapshots to Walrus during snapshot operations. If you want to configure multiple IPs, see the section on 'multipathing.

```
euca-modify-property -p <partition_name>.storage.scpaths=<ip>
```

12. On the ENABLED CLC, set the path to Navisphere CLI that you downloaded earlier to the SC. The following example shows the default path. This is that path on the SC, not on the CLC.

```
euca-modify-property -p
<partition_name>.storage.clipath=/opt/Navisphere/bin/naviseccli
```

Your EMC VNX SAN is now ready to use with Eucalyptus.

Tip: Note: The time it takes for a LUN migration to complete will depend on the exact VNX model, workload, and volume size, and the amount of data actually stored in the volume. The default timeout for LUN migrations is 12 hours. If your deployment uses volumes >50GB, or if you find that snapshots fail and a "migration timeout" message is seen in the SC logs, then you should increase the timeout to a larger value. It is recommended that if you plan on using volumes in the 100GB range that you set that timeout to 3600 or larger. You can set the timeout using `euca-modify-property` as follows:

```
euca-modify-property -p [partition].storage.lunmigrationtimeout=[time in
hours]
```

Configure Dell Equallogic Multipathing

Use multipathing to provide network-and-SP-redundancy for the iSCSI data path between the Dell Equallogic SAN and NCs.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Important: It is **strongly** recommended that you get your system up and running and able to create volumes and snapshots using the Dell Equallogic SAN prior to configuring multipathing. Multipathing can be configured after the cloud is fully functional and will apply to any volumes attached/snapshotted after multipathing is configured. Configuring multipathing on a non-multipathed system does not require a restart of the SC, NC, or CLC.

Important: The Dell Equallogic SAN has separate paths for data and management.

The Dell Equallogic management interface is available for executing control operations only. If your Dell Equallogic SAN is configured to use the management port, please note the IP address of the management interface. The SC can be configured to use the management interface by specifying the IP address of the management interface using the `scpaths` property. For example:

```
euca-modify-property -p mypartition.storage.scpaths=192.168.3.1
```

The Dell Equallogic data interface is configured by specifying the IP address of the data interface using the `ncpaths` property. For example:

```
euca-modify-property -p mypartition.storage.ncpaths=192.168.3.1
```

To configure multipathing for a Dell Equallogic SAN:

CC-BY-SA, Eucalyptus Systems, Inc.

1. Ensure that the mutipathd service is running on each NC:

   ```
   mpathconf --enable
   ```

2. Configure the /etc/multipath.conf file.

   Note: An example configuration for the Dell Equallogic SAN is installed with Eucalyptus. This file is located in `/usr/share/doc/eucalyptus-3.4.1/multipath.conf.example.equallogic` on each NC.

3. Start the mutipathd service:

   ```
   service multipathd start
   ```

4. If you modify the /etc/multipath.conf file, be sure to restart and reconfigure multipathd:

   a) Restart the multipathd service:

   ```
   service multipathd restart
   ```

   b) Run multipathd -k:

   ```
   multipathd -k
   ```

   c) Enter the following commands at the multipathd interactive prompt:

   ```
   reconfigure
   quit
   ```

5. Check that the multipath udev rules file is installed by verifying that the file `/etc/udev/rules.d/12-dm-permissions.rules` file exists.

6. Set the ISCSI paths:

   Tip: The path specification format is `iface0:ip0,iface1:ip1,...,ifaceN:ipN` where iface may be reused (i.e. multiple iface0 entries). Also, note that 'iface' is optional, you may just specify a comma-delimited list of IPs. Eucalyptus will detect which interfaces on the SC/NC can reach each specified IP and will use the first found. You must only specify the 'iface' value if you want precise control over which interfaces access which IPs. For using a single path only, just specify the IP of the iscsi data port to use on the VNX.

   Note: We recommend initially getting the system working with only one path. The path values can be modified at any time to enable multipathing, so it is possible to get everything working and confirm full VNX functionality before attempting multipathing. To use one path, simply specify a single IP for each the following steps.

   a) Set the NC paths. For example:

   ```
   euca-modify-property -p mypartition.storage.ncpaths=iface0:127.0.0.1
   ```

   b) Set the SC paths. For example:

   ```
   euca-modify-property -p mypartition.storage.scpaths=iface0:127.0.0.1
   ```

   Note: The NC and SC may each have different path lists, or you can optionally only enable multipathing on the NCs or SC if desire.

   c) If you specified an `iface` when setting the SC paths, be sure to include a line in the eucalyptus.conf file of each NC in the cluster that defines each `iface`. For example:

   ```
   STORAGE_INTERFACES="iface0=eth0"
   ```

7. Test and verify the configuration by creating (and attaching to) a volume and creating a snapshot on the partition.

8. If testing is successful, you can now configure multiple paths in your *.storage.ncpaths and *.storage.scpaths configuration. In the following example, the IP addresses for each interface correspond to the paths configured on the Dell Equallogic SAN:

CC-BY-SA, Eucalyptus Systems, Inc.

```
euca-modify-property -p
mypartition.storage.ncpaths=iface0:192.168.1.1,iface1:192.168.1.2
```

9. Verify that multipathing is working on an NC by attaching a volume to an instance on that NC and running the following command:

   ```
   multipath -ll
   ```

   This command should return output similar to the following:

   ```
   mpathb (36006016098b0300080722f971b2ee211) dm- 0 DGC,VRAID
   size=1.0G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
   |- +-  policy='round- robin 0' prio=50 status=active
   | `-  6:0:0:1 sdd 8:48 active ready running
   `- +-  policy='round- robin 0' prio=10 status=enabled
   `-  7:0:0:1 sdf 8:80 active ready running
   ```

You have now successfully configured multipathing for your Dell Equallogic SAN installation.

Configure EMC VNX Multipathing

Use multipathing to provide network-and-SP-redundancy for the iSCSI data path between the EMC VNX SAN and NCs.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Important: It is **strongly** recommended that you get your system up and running and able to create volumes and snapshots using the EMC VNX prior to configuring multipathing. Multipathing can be configured after the cloud is fully functional and will apply to any volumes attached/snapshotted after multipathing is configured. Configuring multipathing on a non-multipathed system does not require a restart of the SC, NC, or CLC.

To configure multipathing for a EMC VNX SAN:

1. Ensure that the mutipathd service is running on each NC:

   ```
   mpathconf --enable
   ```

2. Configure the /etc/multipath.conf file.

   Note: An example configuration for EMC VNX is installed with Eucalyptus. This file is located in `/usr/share/doc/eucalyptus-3.4.1/multipath.conf.example.vnx` on each NC.

3. Start the mutipathd service:

   ```
   service multipathd start
   ```

4. If you modify the `/etc/multipath.conf` file, be sure to restart and reconfigure multipathd:

   a) Restart the multipathd service:

   ```
   service multipathd restart
   ```

   b) Run multipathd -k:

   ```
   multipathd -k
   ```

   c) Enter the following commands at the multipathd interactive prompt:

   ```
   reconfigure
   quit
   ```

5. Check that the multipath udev rules file is installed by verifying that the file `/etc/udev/rules.d/12-dm-permissions.rules` file exists.

CC-BY-SA, Eucalyptus Systems, Inc.

6. Set the ISCSI paths:

   Note: The path specification format is iface0:ip0,iface1:ip1,...,ifaceN:ipN where iface may be re- used (i.e. multiple iface0 entries). Also, note that 'iface' is optional, you may just specify a comma- delimited list of IPs. Eucalyptus will detect which interfaces on the SC/NC can reach each specified IP and will use the first found. You must only specify the 'iface' value if you want precise control over which interfaces access which IPs. For using a single path only, just specify the IP of the iscsi data port to use on the VNX.

   Note: We recommend initially getting the system working with only one path. The path values can be modified at any time to enable multipathing, so it is possible to get everything working and confirm full VNX functionality before attempting multipathing. To use one path, simply specify a single IP for each the following steps.

   a) Set the NC paths. For example:

   ```
   euca-modify-property -p mypartition.storage.ncpaths=iface0:127.0.0.1
   ```

   b) Set the SC paths. For example:

   ```
   euca-modify-property -p mypartition.storage.scpaths=iface0:127.0.0.1
   ```

   Note: The NC and SC may each have different path lists, or you can optionally only enable multipathing on the NCs or SC if desire.

   c) If you specified an `iface` when setting the SC paths, be sure to include a line in the eucalyptus.conf file of each NC in the cluster that defines each `iface`. For example:

   ```
   STORAGE_INTERFACES="iface0=eth0"
   ```

7. Test and verify the configuration by creating (and attaching to) a volume and creating a snapshot on the partition.
8. If testing is successful, you can now configure multiple paths in your *.storage.ncpaths and *.storage.scpaths configuration.
9. Verify that multipathing is working on an NC by attaching a volume to an instance on that NC and running the following command:

   ```
   multipath -ll
   ```

   This command should return output similar to the following:

   ```
   mpathb (36006016098b0300080722f971b2ee211) dm- 0 DGC,VRAID
   size=1.0G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
   |- +-  policy='round- robin 0' prio=50 status=active
   | `-  6:0:0:1 sdd 8:48 active ready running
   `- +-  policy='round- robin 0' prio=10 status=enabled
   `-  7:0:0:1 sdf 8:80 active ready running
   ```

You have now successfully configured multipathing for your EMC VNX SAN installation.

Configure NetApp Multipathing

Use multipathing to provide network and controller redundancy for the iSCSI data path between the NetApp Cluster-mode SAN and NCs.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Important: Eucalyptus supports multipathing for NetApp Clustered ONTAP only.

Important: It is **strongly** recommended that you get your system up and running and able to create volumes and snapshots using the NetApp SAN prior to configuring multipathing. Multipathing can be configured after

CC-BY-SA, Eucalyptus Systems, Inc.

the cloud is fully functional and will apply to any volumes attached/snapshotted after multipathing is configured. Configuring multipathing on a non-multipathed system does not require a restart of the SC, NC, or CLC.

To configure multipathing for a NetApp SAN:

1. Ensure that the multipathd service is running on the SC and on each NC:
   ```
   mpathconf --enable
   ```
2. Configure the `/etc/multipath.conf` file.

   Note: An example configuration for NetApp is installed with Eucalyptus. This file is located in `/usr/share/doc/eucalyptus-3.4.1/multipath.conf.example.netapp` on each NC.

3. Start the multipathd service:
   ```
   service multipathd start
   ```
4. If you modify the `/etc/multipath.conf` file, be sure to restart and reconfigure multipathd:
   a) Restart the multipathd service:
      ```
      service multipathd restart
      ```
   b) Run multipathd -k:
      ```
      multipathd -k
      ```
   c) Enter the following commands at the multipathd interactive prompt:
      ```
      reconfigure
      quit
      ```
5. Check that the multipath udev rules file is installed by verifying that the file `/etc/udev/rules.d/12-dm-permissions.rules` file exists.
6. Set the ISCSI paths:

   Note: The path specification format is iface0:ip0,iface1:ip1,...,ifaceN:ipN where iface may be re- used (i.e. multiple iface0 entries). Also, note that ‘iface’ is optional, you may just specify a comma- delimited list of IPs. Eucalyptus will detect which interfaces on the SC/NC can reach each specified IP and will use the first found. You must only specify the ‘iface’ value if you want precise control over which interfaces access which IPs. For using a single path only, just specify the IP of the iscsi data port to use on the NetApp Clustered ONTAP.

   Note: We recommend initially getting the system working with only one path. The path values can be modified at any time to enable multipathing, so it is possible to get everything working and confirm full NetApp functionality before attempting multipathing. To use one path, simply specify a single IP for each the following steps.

   a) Set the NC paths. For example:
      ```
      euca-modify-property -p mypartition.storage.ncpaths=iface0:127.0.0.1
      ```
   b) Set the SC paths. For example:
      ```
      euca-modify-property -p mypartition.storage.scpaths=iface0:127.0.0.1
      ```

      Note: The NC and SC may each have different path lists, or you can optionally only enable multipathing on the NCs or SC if desire.

   c) If you specified an `iface` when setting the SC paths, be sure to include a line in the eucalyptus.conf file of each NC in the cluster that defines each `iface`. For example:
      ```
      STORAGE_INTERFACES="iface0=eth0"
      ```

CC-BY-SA, Eucalyptus Systems, Inc.

7. Test and verify the configuration by creating (and attaching to) a volume and creating a snapshot on the partition.
8. If testing is successful, you can now configure multiple paths in your *.storage.ncpaths and *.storage.scpaths configuration.
9. Verify that multipathing is working on the SC and on an NC by attaching a volume to an instance on the SC and the NC and running the following command:

```
multipath -ll
```

This command should return output similar to the following:

```
mpathp (3600a098037542d69535d43514965354e) dm-2 NETAPP,LUN C-Mode
size=2.0G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1 alua'
 wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 18:0:0:0 sdd 8:48 active ready running
| `- 20:0:0:0 sdf 8:80 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 17:0:0:0 sdc 8:32 active ready running
  `- 19:0:0:0 sde 8:64 active ready running
```

You have now successfully configured multipathing for your NetApp Clustered ONTAP system.

Configure DNS

Eucalyptus provides a DNS service that you can configure to map instance IPs and Walrus bucket names to DNS host names and enable DNS delegation to support transparent failover in HA mode.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

The DNS service will automatically try to bind to port 53. If port 53 cannot be used, DNS will be disabled. Typically, other system services like dnsmasq are configured to run on port 53. To use the Eucalyptus DNS service, you must disable these services.

Configure the Domain and Subdomain

Before using the DNS service, configure the DNS domain name that you want Eucalyptus to handle using the steps that follow. Make sure that the Eucalyptus Cloud Controller (CLC) has been started.

1. Log in to the primary CLC and enter the following:

```
euca-modify-property -p
system.dns.dnsdomain=<eucadomain.yourdomain>
```

2. You can configure the load balancer DNS subdomain. To do so, log in to the primary CLC and enter the following:

```
euca-modify-property -p
loadbalancing.loadbalancer_dns_subdomain = <your-subdomain>
```

Turn on IP Mapping

To turn on mapping of instance IPs to DNS host names:

1. Enter the following command on the primary CLC:

```
euca-modify-property -p bootstrap.webservices.use_instance_dns=true
```

When this option is enabled, public and private DNS entries are set up for each instance that is launched in Eucalyptus. This also enables virtual hosting for Walrus. Buckets created in Walrus can be accessed as hosts. For example, the bucket `mybucket` is accessible as `mybucket.walrus.eucadomain.yourdomain`.

Instance IP addresses will be mapped as `euca-A.B.C.D.eucalyptus.<subdomain>`, where `A.B.C.D` is the IP address (or addresses) assigned to your instance.

CC-BY-SA, Eucalyptus Systems, Inc.

2. If you want to modify the subdomain that is reported as part of the instance DNS name, enter the following command:

```
euca-modify-property -p cloud.vmstate.instance_subdomain=.<custom-dns-subdomain>
```

When this value is modified, the public and private DNS names reported for each instance will contain the specified custom DNS subdomain name, instead of the default value, which is `eucalyptus`. For example, if this value is set to `foobar`, the instance DNS names will appear as `euca-A.B.C.D.foobar.<subdomain>`.

Enable DNS Delegation

For HA: If you are using HA and do not enable DNS delegation, you must manually update `EC2_URL`, `S3_URL` and `EUARE_URL` to point to the new primary hosts in case of failover.

DNS delegation allows you to forward DNS traffic for the Eucalyptus subdomain to the Eucalyptus CLC hosts. These hosts act as name servers. This allows interruption-free access to Eucalyptus cloud services in the event of a failure. Both primary and secondary CLC hosts are capable of mapping cloud host names to IP addresses of the primary CLC and Walrus hosts.

For example, if the IP address of the primary and secondary CLC are `192.168.5.1` and `192.168.5.2`, and the IP addresses of primary and secondary Walruses are `192.168.6.1` and `192.168.6.2`, the host `eucalyptus.eucadomain.yourdomain` will resolve to `192.168.6.1` and `walrus.eucadomain.yourdomain` will resolve to `192.168.6.1`.

If the primary CLC fails, the secondary CLC will become the primary and `eucalyptus.eucadomain.yourdomain` will resolve to `192.168.5.2`. If the primary Walrus fails, the secondary Walrus will be promoted and `walrus.eucadomain.yourdomain` will resolve to `192.168.6.2`.

To enable DNS delegation:

1. On the primary CLC, enter the following command:

```
euca-modify-property -p bootstrap.webservices.use_dns_delegation=true
```

2. Because the credentials are now slightly changed, you must generate the administrative credentials and source the `eucarc` file again. For more information, see *Generate Administrator Credentials*.

Configure the Master DNS Server

Set up your master DNS server to forward the Eucalyptus subdomain to the primary and secondary CLC servers, which act as name servers.

The following example shows how the Linux name server `bind` is set up to forward the Eucalyptus subdomain.

1. Open `/etc/named.conf` and set up the `eucadomain.yourdomain zone`. For example, your `/etc/named.conf` may look like the following:

```
zone "yourdomain" {
type master;
file "/etc/bind/db.yourdomain";
};

#Forward eucadomain.yourdomain
zone "eucadomain.yourdomain" {
type forward;
forward only;
forwarders { <CLC_0_IP>; <CLC_1_IP>; };
};
```

where `<CLC_0_IP>` is the IP address of your primary CLC and `<CLC_1_IP>` is the IP address of your secondary CLC.

2. Create `/etc/bind/db.yourdomain` if it does not exist. If your master DNS is already set up for `yourdomain`, you will need to add name server entries for `<CLC_0_IP>` and `<CLC_1_IP>`. For example:

CC-BY-SA, Eucalyptus Systems, Inc.

```
$TTL 604800
@ IN SOA yourdomain. root.yourdomain. (
2 ; Serial
604800 ; Refresh
86400 ; Retry
2419200 ; Expire
604800 ) ; Negative Cache TTL
;
@ IN NS ns.yourdomain.
@ IN A <master_nameserver_IP>

ns.yourdomain. IN A <master_nameserver_IP>

;Add entries for primary and secondary CLCs
eucadomain.yourdomain. IN NS clc0.eucadomain.yourdomain.
eucadomain.yourdomain. IN NS clc1.eucadomain.yourdomain.

clc0.eucadomain.yourdomain. IN A <CLC_0_IP>
clc1.eucadomain.yourdomain. IN A <CLC_1_IP>
```

where `clc0.eucadomain.yourdomain` and `clc1.eucadomain.yourdomain` are the host names of your primary and secondary CLC servers.

3. Restart the bind nameserver (`/etc/init.d/bind9 restart` or `/etc/init.d/named restart`, depending on your Linux distribution).
4. Test your setup by pointing `/etc/resolv.conf` on your client to your primary DNS server and attempt to resolve `eucalyptus.eucadomain.yourdomain` using ping or nslookup. It should return the IP address of the primary CLC server.

Advanced DNS options

Recursive lookups and split-horizon DNS are available in Eucalyptus.

1. To enable any of the DNS resolvers, set `dns.enabled` to `true`.
2. To enable the recursive DNS resolver, set `dns.recursive.enabled` to `true`.
3. To enable split-horizon DNS resolution for internal instance public DNS name queries, set `dns.split_horizon.enabled` to `true`.

Configure Node Controller

To prevent potential problems, we recommend that you perform the steps listed in this topic on each NC.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

1. Log in to an NC server and open the `/etc/eucalyptus/eucalyptus.conf` file.
2. Change the `CONCURRENT_DISK_OPS` parameter to the number of disk-intensive operations you want the NC to perform at once. On some Linux installations, a sufficiently large amount of local disk activity can slow down process scheduling. This can cause other operations (e.g., network communication and instance provisioning) appear to stall. Examples of disk-intensive operations include preparing disk images for launch and creating ephemeral storage. Set this value to `1` to serialize all disk-intensive operations. Set to a higher number to increase the amount of disk-intensive operations the NC will perform in parallel.

Configure DRBD

This topic details how to configure DRBD for Eucalyptus HA.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Before you begin, ensure that you have the following information:

CC-BY-SA, Eucalyptus Systems, Inc.

- The IP address and hostname of each Walrus
- The DRBD block device name of each Walrus. In the following examples, we assume that DRBD block device name is `/dev/drbd1`.
- The DRBD backing disk partition names on each Walrus. A partition (either on a new disk or an existing disk) should be dedicated to Walrus. The partition sizes should be identical.

Tip: Consider backing the DRBD resource with a logical volume using LVM, this will make growing the backing store easier in the future if you are running low on disk space.

Configuring DRBD requires that you edit the Eucalyptus DRBD file to include your Walrus information, and edit the master DRBD file to tell it to look for the Eucalyptus DRBD file.

Tip: We recommend that you use DRBD peer authentication. For more information, go to *http://www.drbd.org/users-guide/re-drbdconf.html*.

To configure DRBD:

1. Log in to the primary Walrus.
2. Load the DRBD module

   ```
   modprobe drbd
   ```

   There is no output from this command.
3. Copy the example Eucalyptus DRBD file (`/etc/eucalyptus/drbd.conf.example`) to `/etc/eucalyptus/drbd.conf`.
4. Open the `/etc/eucalyptus/drbd.conf` file and make the following edits:

   - Change the value of `<walrus-host-1>` to the hostname (output of `uname -n`) of the primary Walrus.
   - Change the value of `<drbd-block-dev, e.g., /dev/drbd1>` to `/dev/drbd1`
   - Change the value of `<drbd-backing-disk-dev, e.g. /dev/sdb1>` to `/dev/sdb1`
   - Change the value of `<walrus-host-1-ip>` to the IP address of the primary Walrus.
   - Change the value of `<walrus-host-2>` to the hostname (output of `uname -n`) of the secondary Walrus.
   - Change the value of `<drbd-block-dev, e.g., /dev/drbd1>` to `/dev/drbd1`
   - Change the value of `<drbd-backing-disk-dev, e.g. /dev/sdb1>` to `/dev/sdb1`
   - Change the value of `<walrus-host-2-ip>` to the IP address of the secondary Walrus.

   The file should look like the following example:

   ```
   common {
     protocol C;
   }

   resource r0 {

     on walrus00.eucalyptus.com {
       device    /dev/drbd1;
       disk      /dev/sdb1;
       address   192.168.58.1:7789;
       meta-disk internal;
     }

     on walrus01.eucalyptus.com {
       device    /dev/drbd1;
       disk      /dev/sdb1;
       address   192.168.58.2:7789;
       meta-disk internal;
     }

     syncer {
       rate 40M;
     }
   ```

CC-BY-SA, Eucalyptus Systems, Inc.

```
  net {
    after-sb-0pri discard-zero-changes;
    after-sb-1pri discard-secondary;
  }
}
```

Important: On RHEL 6 and Ubuntu, remove the `common` section (`common { protocol C; }`). The configuration in these distributions already include a common section.

5. Save and close the file.
6. Open the master DRBD file (`/etc/drbd.conf`) and append the following line:

```
include "/etc/eucalyptus/drbd.conf";
```

Important: For Ubuntu, you must also remove the `common` section (`common { protocol C; }`) and the line `include "drbd.d/*.res";`.

7. Save and close the file.
8. Log in to the secondary Walrus and load the DRBD module

```
modprobe drbd
```

There is no output from this command.

9. On both primary and secondary Walrus, open the file `/etc/fstab` and append the following line for the DRBD resource to allow the `eucalyptus` user to mount and unmount the device:

```
/dev/drbd1 /var/lib/eucalyptus/bukkits ext3 noauto,owner 0 0
```

Important: Make sure the `/etc/fstab` file has the right filesystem format for the device that DRBD will be using. If labeled incorrectly, status of DRBD will be `Secondary/Secondary`.

10. Save and close the file.
11. Open the `/etc/eucalyptus/eucalyptus.conf` file and make the following configuration:

```
CLOUD_OPTS="-Dwalrus.storage.manager=DRBDStorageManager"
```

12. Copy the `/etc/drbd.conf`, the `/etc/eucalyptus/drbd.conf`, and the `/etc/eucalyptus/eucalyptus.conf` files to the secondary Walrus server.
13. Restart Walrus, first on the primary, and then on the secondary. Restarting the primary Walrus will trigger an HA failover to the secondary, and restarting the secondary will fail back, preparing the entire system for the next steps.

```
service eucalyptus-cloud restart
```

14. Monitor the failover with euca-describe-services. If successful, stop the eucalyptus-cloud service again on both Walruses whilst you configure the DRBD device.

```
service eucalyptus-cloud stop
```

15. On the primary Walrus, associate the DRBD block device (`/dev/drbd1`) with the disk partition allocated for Walrus (`/dev/sdb1`).

```
drbdmeta --force /dev/drbd1 v08 /dev/sdb1 internal create-md
drbdadm up r0
```

Important: Repeat this step on the secondary Walrus.

16. Set up the DRBD block device on the primary Walrus:

Tip: With a large DRBD device, the initial synchronization can take a considerable amount of time. Consult *Skip Initial Device Synchronization* for instructions on how to skip the synchronization.

CC-BY-SA, Eucalyptus Systems, Inc.

```
drbdsetup /dev/drbd1 syncer -r 110M
drbdadm -- --overwrite-data-of-peer primary r0
```

17. On the primary Walrus only, run the following command to indicate whether the data on the DRBD primary and secondary is consistent:

```
drbdadm dstate r0
```

Wait for the output to display `UpToDate/UpToDate`, then continue to the next step.

Tip: To view the synchronization process in near-realtime, run `watch -n 2 cat /proc/drbd`.

18. On the primary Walrus, create a filesystem on `/dev/drbd1`. Eucalyptus supports ext3 or ext4. For example:

```
mkfs.ext3 /dev/drbd1
```

19. With the DRBD device now configured, start the eucalyptus-cloud service on both Walruses.

```
service eucalyptus-cloud start
```

20. On the primary CLC, tell Eucalyptus to use DRBD parameters configured in the DRBD config file so Walrus can write to the correct device:

```
euca-modify-property -p walrus.blockdevice=/dev/drbd1
euca-modify-property -p walrus.resource=r0
```

Skip Initial Device Synchronization

This topic details steps for skipping the initial device synchronization in Eucalyptus HA.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

For HA: Only perform these steps if you have no existing data on your DRBD devices. We assume you have completed the steps up to and including step *15* in *Configure DRBD*.

1. Before promoting the DRBD resource to primary, generate a new UUID and clear the bitmap on the primary Walrus ONLY.

```
drbdadm -- --clear-bitmap new-current-uuid r0
```

2. On the same Walrus, promote the resource to primary:

```
drbdadm primary r0
```

3. Resume the configuration at step *17* in *Configure DRBD*.

Synchronize Pairs Configuration

This topic details steps for keeping configuration files synchronized between Eucalyptus HA component pairs.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

The `csync2` utility can be used to achieve this, it's designed to keep configuration consistent between systems in large clusters. The following example uses two standalone Walrus systems.

For HA: Only perform these steps between HA pairs containing the same components to ensure the configuration matches. For assistance in implementing more advanced tasks the `csync2` man pages should be consulted directly.

CC-BY-SA, Eucalyptus Systems, Inc.

Note: Completion of this section is optional. You might also want to expand on this example and generate your own more complex configurations for other components, adding more groups and including additional files in the synchronization tasks.

1. Install the csync2 utility from the EPEL repository on both systems:

```
yum -y install csync2
```

2. On the primary system, generate a pre-shared key for the hosts you wish to sync:

```
csync2 -k /etc/csync2/eucalyptus.key
```

3. On the primary system edit the /etc/csync2/csync2.cfg configuration file to create the synchronization template. Use the one below as an example, creating a group for your two Walrus hosts and adding the eucalyptus and drbd configuration files. Replace the host lines below with the full hostnames of your Walrus systems:

```
   group eucalyptus {
   host euca-walrus-ha-1;
   host euca-walrus-ha-2;
   key /etc/csync2/eucalyptus.key;
   include /etc/csync2/csync2.cfg;
   include /etc/eucalyptus/eucalyptus.conf;
   include /etc/eucalyptus/drbd.conf;
   include /etc/drbd.conf;
   }
```

4. Copy the base configuration file and pre-shared key generated earlier to the secondary Walrus server:

```
scp /etc/csync2/csync2.cfg /etc/csync2/eucalyptus.key
euca-walrus-ha-2:/etc/csync2/
```

5. Next, ensure the csync2 service and xinetd are set to start at boot on both Walrus systems:

```
 chkconfig csync2 on
 chkconfig xinetd on
 service xinetd restart
```

6. Next, perform a dry-run sync to check for proposed changes. Run this on your primary Walrus.

```
csync2 -xvd
```

7. If you are happy with the proposed changes, go ahead and perform the sync.

```
csync2 -xv
```

8. With the initial synchronization complete, a periodic compare and sync could be run via a cronjob which runs the csync2 utility every 5 minutes. In the future, changes can be made to the synchronized files on your primary Walrus and these would be propagated to the other system by the cronjob. Add the following to /etc/cron.d/csync2 as root on the primary Walrus:

```
*/5 * * * * root csync2 -x
```

Set Up Security Groups

In Managed and Managed (No VLAN) networking modes, you must configure the system with parameters that define how Eucalyptus will allocate and manage virtual machine networks. These virtual machine networks are known as security groups. The relevant parameters are set in the `eucalyptus.conf` on all machines running a CC.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

CC-BY-SA, Eucalyptus Systems, Inc.

These parameters are:

- `VNET_SUBNET`
- `VNET_NETMASK`
- `VNET_ADDRSPERNET`

The CC will read `VNET_SUBNET` and `VNET_NETMASK` to construct a range of IP addresses that are available to all security groups. This range will then be further divided into smaller networks based on the size specified in `VNET_ADDRSPERNET`. Note that Eucalyptus reserves eleven addresses per security group, so these networks will be smaller than the value specified in `VNET_ADDRSPERNET`.

The first time an instance runs in a given security group, Eucalyptus chooses an unused range of IPs of size specified in `VNET_ADDRSPERNET`. Eucalyptus then implements this network across all CCs. All instances that run within this given security group obtain a specific IP from this range.

Tip: Eleven of the IP addresses within each security group network are reserved for Eucalyptus to use as gateway addresses, broadcast address, etc. For example, if you set `VNET_ADDRSPERNET` to `32`, there will be 21 free IPs that are available for instances running in that security group.

In Managed mode, each security group network is assigned an additional parameter that is used as the VLAN tag. This parameter is added to all virtual machine traffic running within the security group. By default, Eucalyptus uses VLAN tags starting at 2, going to a maximum of 4094. The maximum is dependent on how many security group networks of the size specified in `VNET_ADDRSPERNET` fit in the network defined by `VNET_SUBNET` and `VNET_NETMASK`.

If your networking environment is already using VLANs for other reasons, Eucalyptus supports the definition of a smaller range of VLANs that are available to Eucalyptus. To configure Eucalyptus to use VLANs within a specified range:

1. Choose your range (a contiguous range of VLANs between 2 and 4095).
2. Configure your cluster controllers with a VNET_SUBNET/VNET_NETMASK/VNET_ADDRSPERNET that is large enough to encapsulate your desired range. For example, for a VLAN range of 1024-2048, you could set VNET_NETMASK to 255.254.0.0 to get a large enough network (131072 addresses), and VNET_ADDRSPERNET to 64, to give 2048 possible security groups.

   Tip: The number of instances per security group can be calculated as follows:

   subnets (SGs) = no. hosts / addrspernet
   instances per subnet (SG) = addrspernet - 10

3. Configure your cloud controller to work within that range. Use the following commands to verify that the range is now set to be 2-2048, a superset of the desired range.

   ```
   euca-describe-properties | grep cluster.maxnetworktag
   euca-describe-properties | grep cluster.minnetworktag
   ```

4. Constrict the range to be within the range that the CC can support as follows:

   ```
   euca-modify-property -p cloud.network.global_max_network_tag=<max_vlan_tag>
   euca-modify-property -p cloud.network.global_min_network_tag=<min_vlan_tag>
   ```

   This ensures that Eucalyptus will only use tags between 1024 and 2048, giving you a total of 1024 security groups, one VLAN per security group.

Tip: If VMs are already running in the system using a VLAN tag that is outside the range specified by global_min_network_tag-global_max_network_tag, that network will continue to run until all VMs within the network are terminated and the system removes reference to that network. Best practice is to configure these values in advance of running virtual machines.

Configure the Load Balancer

Eucalyptus provides optional support for Load Balancing. In order to use this support, you will need to register the Load Balancer image with the cloud.

CC-BY-SA, Eucalyptus Systems, Inc.

Caution: Eucalyptus HA is currently in technical preview. To install Eucalyptus without high availability, see *Eucalyptus Installation*. For information about technical previews in Eucalyptus, see *Technology Preview*.

Install and Register the Load Balancer Image

Eucalyptus provides a tools for installing and registering the Load Balancer image. Once you have run the tool, your Load Balancer will be ready to use.

Run the following command on the machine where you installed the eucalyptus-service-image package:

```
euca-install-load-balancer --install-default
```

Verify Load Balancer Configuration

If you would like to verify that Load Balancer support is enabled you can list installed Load Balancers. The currently active Load Balancer will be listed as enabled. If no Load Balancers are listed, or none are marked as enabled, then your Load Balancer support has not been configured properly.

1. Run the following command to list installed Load Balancer images:

   ```
   euca-install-load-balancer --list
   ```

2. You can also check the enabled Load Balancer EMI with:

   ```
   euca-describe-properties loadbalancing.loadbalancer_emi
   ```

3. If you need to manually set the enabled Load Balancer EMI use:

   ```
   euca-modify-property -p loadbalancing.loadbalancer_emi=emi-12345678
   ```

CC-BY-SA, Eucalyptus Systems, Inc.

Index

A

arbitrator *37*, *108*
registering *37*, *108*
architecture *10*, *74*
components *74*
HA *74*

C

cloud controller (CLC) *33*, *104*, *106*
registering *106*
starting *33*, *104*
cluster controller (CC) *34*, *36*, *104*, *107*
registering *36*, *107*
starting *34*
starting (HA) *104*
components *7*, *10*, *12*, *35*, *75*, *77*, *105*
about *7*
cloud *10*, *75*
cluster *10*, *75*
disk space *12*, *77*
node *10*, *75*
registering *35*, *105*
configuration *26*, *31–32*, *96*, *101–102*
iptables *32*, *102*
loop devices *31*, *101*
configuring *39*, *56*, *58*, *110*, *123*, *125*, *129*
concurrency level *58*, *125*
credentials *39*, *110*
DNS *56*, *123*
DRBD *125*
security groups *58*, *129*
subdomains *56*, *123*
credentials *39*, *110*

D

DNS *56*, *123*
configuring *56*, *123*
delegation *56*, *123*
IP mapping *56*, *123*

E

euca2ools *73*

F

firewalls *21*, *90*
configuring *21*, *90*

H

high availability *85*, *125*
DRBD *125*
multipathing *85*

I

installation *22*
installing *25*, *72*, *92*, *95*
HA *92*
local repository *72*
CentOS *72*
nightly packages *25*, *95*

M

MTA *22*, *91*
configuring *22*, *91*
multi-cluster *32*, *102*
multipathing *85*

N

networking *20*, *26*, *29–30*, *32*, *58*, *89*, *96*, *99–100*, *102*, *129*
configuration *26*, *29–30*, *32*, *96*, *99–100*, *102*
managed *29*, *99*
managed (no VLAN) *30*, *100*
multi-cluster *32*, *102*
configuring bridges *20*, *89*
security groups *58*, *129*
networking modes *12*, *14–15*, *77*, *79–80*
edge *14*, *80*
managed *15*, *79*
managed (no VLAN) *15*, *79*
planning *12*, *77*
node controller *34*, *37*, *108*
HA *108*
registering *37*, *108*
starting *34*
NTP *21*, *91*
configuring *21*, *91*

O

object storage *17*
object storage (HA) *83*
object storage gateway (OSG) *17*
object storage gateway (OSG) for HA *83*

S

SELinux *21*, *90*
configuring *21*, *90*
startup *33*, *35*, *104–105*
verifying *35*, *105*
storage controller *36*, *107*
HA *107*
registering *36*, *107*
storage controller (SC) *34*
starting *34*
support *71*

system requirements *8*, *16–17*, *81*, *83*, *87*
high availability *83*
networking *17*, *87*
SAN *16*, *81*

U

upgrading *60*
user-facing services *36*
registering *36*
user-facing services (HA) *106*
registering *106*
user-facing services (upgrade) *65*
registering *65*

W

walrus *34*, *36*, *104*, *107*
HA *107*
registering *36*, *107*
starting *34*, *104*