

# **HP Continuous Delivery Automation**

For the Windows<sup>®</sup> and Linux operating systems

Software Version: 1.30

## **HP CDA 1.30 Cloud Setup Integration Guide**

Document Release Date: August 2013

Software Release Date: August 2013



## Legal Notices

### Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

### Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

### Copyright Notice

© Copyright 2012-2013 Hewlett-Packard Development Company, L.P.

### Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

## Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

**<http://h20230.www2.hp.com/selfsolve/manuals>**

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

**<http://h20229.www2.hp.com/passport-registration.html>**

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

## Support

Visit the HP Software Support Online web site at:

**<http://www.hp.com/go/hpsoftwaresupport>**

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

**<http://h20229.www2.hp.com/passport-registration.html>**

To find more information about access levels, go to:

**[http://h20230.www2.hp.com/new\\_access\\_levels.jsp](http://h20230.www2.hp.com/new_access_levels.jsp)**

# Contents

Contents .....	5
Overview .....	9
Terminology .....	10
Dashboard Look and Feel .....	12
Cloud Tab .....	12
Prerequisites View .....	12
Domain Controllers View .....	12
Compute Region Controllers View .....	13
Updates and Extensions View .....	14
Environment Tab .....	14
Setup View .....	14
Private Cloud .....	15
Hybrid-Only Cloud .....	15
Virtual Hybrid-Only Cloud .....	16
Server Types View .....	16
Connections View .....	17
Networks View .....	17
HP Cloud Connector Cloud Services Reference Architectures (CSRA) .....	18
Services .....	18
Controllers .....	19
Example: Cloud Infrastructure .....	19
HP Cloud Connector Physical Infrastructure Reference Architectures (PIRA) .....	21
Prerequisites .....	21
Infrastructure Sizing .....	22
Server Configuration .....	22
Switch Configuration .....	23
Network Configuration .....	23
Storage Configuration .....	24
Cloud Infrastructure Network Topology .....	24

Cloud Infrastructure Production Deployments .....	25
Bootstrap the Cloud Administration Node Using the ISO .....	27
Example of Server Deployment .....	28
Launch the Cloud Installation Dashboard .....	29
Set Up the Cloud Administration Node .....	30
Customize Admin and Public Network .....	30
Customizing Admin Network .....	30
Customizing Public Network .....	31
Complete the Setup Process .....	32
Configure the Cloud Administration Node .....	34
Download Third-Party Packages .....	34
Deploy and Allocate the Domain Controller Node .....	36
Deploy and Allocate the Domain Controller Node .....	36
Deployment States .....	38
Create an Alias for the Domain Controller .....	38
Complete Storage Configuration .....	39
Apply Domain Controller Barclamps .....	39
Deploy and Allocate Compute Region Controller Nodes .....	42
Deploy and Allocate Compute Region Controller Nodes .....	43
Create an Alias for a Compute Region Controller .....	44
Complete Storage Configuration .....	44
Apply the Nova Barclamp .....	45
Nova Barclamp Roles .....	46
Nova Barclamp Proposal Settings .....	47
Cloud Utilities Barclamp .....	48
Clean-up Scripts .....	48
Apply the HP Cloud Utils 120 Barclamp .....	48
Monitor RabbitMQ .....	49
Configure Network Infrastructure for Virtual Machines .....	50
Assumptions .....	50
Configure a Fixed Network .....	50

Configure a Floating Network .....	52
<b>Post Deployment Tasks .....</b>	<b>53</b>
Configure an Image Repository .....	53
Create a Keypair .....	54
Configure the Default Security Group .....	55
Launch an Instance .....	55
Validate Instance Accessibility .....	56
Create a Resource Pool .....	57
Create an Infrastructure Topology .....	58
Create an Infrastructure Design .....	59
Launch an Infrastructure Design Document .....	60
Validate an Infrastructure Design Document .....	61
<b>Troubleshooting .....</b>	<b>62</b>
Problem: Cloud Administration Node displays the "not ready" (grey) state. ....	62
Problem: Cloud Infrastructure - An error message displays when configuring Cloud Infrastructure prerequisites. ....	62
Problem: Cloud Infrastructure barclamp proposal fails. ....	63
Problem: Cloud Installation Dashboard MongoDB prerequisite installation failed. ....	64
Problem: Domain Controller Node or Compute Region Node displays the "not ready" (grey) state. ....	66
Problem: PXE - When creating a new PXE node, the PXE boot fails with a TFTP timeout error. ....	66
<b>Appendix A: Domain Controller Barclamps .....</b>	<b>67</b>
<b>Appendix B: Load Balancer Image .....</b>	<b>75</b>
Prerequisites .....	75
Ways To Use a Load Balancer Image .....	75
Static Load Balancer .....	75
Dynamically-Provisioned Load Balancer .....	76
Create a Standalone Load Balancer Instance .....	76
Convert a Load Balancer Instance to an Image Snapshot .....	79
<b>Appendix C: Chef Server Image .....</b>	<b>80</b>
Prerequisites .....	80

Ways To Use a Chef Server Image .....	80
Static Chef Server .....	81
Dynamically-Provisioned Chef Server .....	81
Creating a Chef Server Image .....	81
Adding a Chef Client .....	82

# Overview

In HP Continuous Delivery Automation (HP CDA), an HP Cloud Infrastructure requires planning and preparation by IT Administrators to set up, configure, and install a specific network, storage, and hardware infrastructure.

To help you plan and prepare for, and set up and install an HP Cloud Infrastructure, review the following sections:

- ["Terminology " on next page](#)
- ["Dashboard Look and Feel " on page 12](#)
- ["HP Cloud Connector Cloud Services Reference Architectures \(CSRA\) " on page 18](#)
- ["HP Cloud Connector Physical Infrastructure Reference Architectures \(PIRA\) " on page 21](#)
- ["Bootstrap the Cloud Administration Node Using the ISO " on page 27](#)
- ["Launch the Cloud Installation Dashboard " on page 29](#)
- ["Set Up the Cloud Administration Node " on page 30](#)
- ["Configure the Cloud Administration Node " on page 34](#)
- ["Deploy and Allocate the Domain Controller Node " on page 36](#)
- ["Deploy and Allocate Compute Region Controller Nodes " on page 43](#)
- ["Configure Network Infrastructure for Virtual Machines " on page 50](#)
- ["Post Deployment Tasks " on page 53](#)
- ["Troubleshooting " on page 62](#)
- ["Appendix A: Domain Controller Barclamps " on page 67](#)
- ["Appendix B: Load Balancer Image " on page 75](#)
- ["Appendix C: Chef Server Image " on page 80](#)

## Terminology

The following table defines key terms and concepts used in this guide:

Term	Description
Admin Network	Used for administrative functions, such as member node installation, TFTP booting, DHCP assignments, KVM, system logs, backups, and other monitoring tasks. There is only one VLAN set up for this function and it spans the entire network.
CDA Node	HP Continuous Delivery Automation (HP CDA) is a complete dev-ops solution that integrates with Cloud Infrastructure services.
CSA Node	HP Cloud Service Automation (HP CSA) is a portal to subscribe to applications, platforms, and infrastructures offered by HP CDA.
Cloud Administration Node	Ubuntu12.04 OS where Cloud Infrastructure media ISO is booted. Bootstraps the install to set up the Domain Controller Node and Compute Region Node. Hosts Cloud Installation Dashboard that helps deployment of private and hybrid-only Cloud Infrastructure environments by network booting the member nodes, which are typically Domain Controller and Compute Region nodes.
Cloud Administration Dashboard	Hosted on a Domain Controller node. Enables administrators to manage and administer a Cloud Infrastructure, such as launching instances, uploading images, creating resource pools, and creating network topologies. During deployment, this dashboard is typically used for creating network infrastructure for virtual machine instances.
Cloud Installation Dashboard	Hosted on the Cloud Administration Node as a bootstrap UI for setting up prerequisites, network, servers, and connections in a Cloud Infrastructure. Invokes Cloud Installation Dashboard after the setup wizard is completed. A typical access URL is <a href="http://192.168.124.10:9000">http://192.168.124.10:9000</a>
Compute Region Controller	A set of services, preferably across multiple nodes, to offer private cloud resources for compute, storage, and network to self-service users of your private cloud. In OpenStack, a Compute Region is commonly known as Nova. A Compute Region includes a Compute Region Controller, Storage, and a set of Compute and Network services. Set up a Compute Region only when you are deploying a private Cloud Infrastructure.
Compute Region Node	Also known as a KVM node because it hosts the cloud Virtual Machine instances. This node hosts OpenStack services called nova-compute and nova-network. Multiple Compute Region Nodes can be created to provide more cloud capacity.

Domain Controller	A set of services that are used to provide capabilities within your cloud, including services for identity management, image library, document repository, and so on. A domain controller manages multiple compute regions in a private cloud or hybrid cloud.
Domain Controller Node	Hosts Cloud Infrastructure services, such as Identity (Keystone), Image (Glance), Volume, OpenStack's Nova Controller, and HP's Domain Controller. This node can also host other HP software, such as HP Continuous Delivery Automation (HP CDA) and HP Cloud Services Automation (HP CSA). To increase cloud capacity, you can create multiple Compute Region Nodes.
Fixed Network	An IP address network that is associated with the same instance each time that instance boots. This type of network is used to manage the instance. In general, this network is not accessible to end users or the public internet.
Floating Network	An IP address network that is associated with one of OpenStack's virtual machines so that the instance has the same public IP address each time it boots. To maintain a consistent IP address for maintaining DNS assignment, create a pool of floating IP addresses and assign them to instances when they are launched.
Hybrid-Only Cloud	Requires a single bare-metal node to host Domain Controller and a Compute Region that relies on other OpenStack-based Cloud Infrastructures, such as HP Cloud Services ( <a href="http://www.hpcloud.com">www.hpcloud.com</a> ) .
Instance	Virtual machine that is provisioned in a Cloud Infrastructure.
Node Dashboard	Part of the Cloud Installation Dashboard for managing nodes, networks, barclamps, and utilities.
Private Cloud	Requires a bare-metal or virtual machine to host a Domain Controller, and one or more bare-metals to host Compute Region.
Public Network	Used for connections to devices that are external to the Cloud Infrastructure. These networks are externally visible for services, such as load balancers and web servers.
Virtual Hybrid Cloud	Similar to a hybrid-only cloud, except the Domain Controller is hosted on a virtual machine. HP recommends this as a preferred deployment for smaller environment or lab purposes.

## Dashboard Look and Feel

To understand what the Cloud Installation Dashboard looks like and what you use it for, review the following information about the Dashboard tabs and their views:

**Cloud Tab**—In the Cloud tab, you update and complete prerequisite information. After you finish the prerequisites, you can then begin creating the services required to stand up a cloud. You will need to create a Domain Controller and one or more Compute Region Controllers, depending on whether you are creating a private cloud or a hybrid-only cloud. See ["Cloud Tab " below](#)

**Environment Tab**—In the Environment tab, review the types of cloud environments you want to create for your environment, the setup requirements for each type of cloud environment, and how to set up the Cloud Administration Node to deploy a cloud environment. See ["Environment Tab " on page 14](#).

## Cloud Tab

To understand why and when you need to use different views in the Cloud tab, review the following topics:

- ["Prerequisites View" below](#)
- ["Domain Controllers View" below](#)
- ["Compute Region Controllers View" on the facing page](#)
- ["Updates and Extensions View" on page 14](#)

## Prerequisites View

This view provides a description, the last time the prerequisite was successfully updated, and, if necessary, allows you to edit certain details that are required to complete that prerequisite.

## Domain Controllers View

After you update and complete Prerequisites information, you must create a Domain Controller. A Domain Controller is a set of services for identity management, the image library, document repository, and so on.

- For a small cloud, a Domain Controller can be deployed to a single bare-metal server.
- For a virtual hybrid-only cloud, a Domain Controller can be deployed to a single virtual machine.

- For a large cloud, the services and components can be spread across multiple machines.

To create a Domain Controller:

1. Click the **Deploy Domain Controller** button.
2. In the Authentication Required dialog, enter the User Name: **crowbar** and Password: **crowbar**.
3. Apply new proposals to the following Cloud Infrastructure installation modules (barclamps), in the following order:
  - a. Postgresql 915
  - b. Hp Keystone 201211
  - c. Hp Glance 201211
  - d. Rabbitmq 271
  - e. Mongodb 284
  - f. Hp Eden 100
  - g. Hp Peer 100
  - h. Hp Eve 100
  - i. Hp Focus 100
  - j. Hp Skyline 120
  - k. HP Cloud Utils 120

**Note:** After you apply proposals to this set of barclamps, the cloud Domain Controller is installed. To update an existing Domain Controller, edit and re-apply existing proposals on this set of barclamps. For more information about these barclamps, see "[Appendix A: Domain Controller Barclamps](#)" on page 67.

## Compute Region Controllers View

After you create a Domain Controller and if you require private cloud resources for self-service users, you can create and deploy one or more Compute Region Controllers. A Compute Region Controller deploys on one or more servers, depending on the scale of resources that are needed to offer private cloud resources for compute, storage, and network.

To create a new Compute Region Controller:

1. Click the **Deploy Compute Region Controllers** button.
2. In the Authentication Required dialog, enter the User Name: **crowbar** and Password: **crowbar** to open the Node Dashboard.
3. In the Node Dashboard, apply a new proposal to the Cloud Infrastructure installation Nova module (barclamp).

**Note:** After you apply proposals to the Nova barclamp, the cloud Compute Region Controller is installed. Compute Region Controllers are elastic and can expand and contract at any time, as long as unallocated bare-metal servers are available and attached to the Cloud Administration Node's admin network and are powered on and discoverable. To update an existing Compute Region Controller, edit and re-apply existing proposals on the Nova barclamp.

## Updates and Extensions View

This view allows you to optionally import and install cloud service utility packages that extend the functionality of your environment. After a package is imported and installed, you can view product and package details. The Details Log tab tracks every install attempt by date and can be used to view installation output and troubleshoot.

## Environment Tab

To understand why and when you need to use different views in the Environment tab, review the following topics:

- ["Setup View" below](#)
- ["Server Types View" on page 16](#)
- ["Connections View" on page 17](#)
- ["Networks View" on page 17](#)

## Setup View

In the Cloud Installation Dashboard, set up the Cloud Administration Node to deploy these types of cloud environments:

- ["Private Cloud" below](#)
- ["Hybrid-Only Cloud" below](#)
- ["Virtual Hybrid-Only Cloud " on next page](#)

After you determine the type of cloud you need to set up, review the information in the **Cloud Installation Dashboard > Environment > Setup > Overview** tab.

## ***Private Cloud***

In this type of cloud environment, you deploy a private cloud on your own bare-metal servers, such as blades and rack-mounted servers, that you have and manage in your IT environment. A private cloud is the most advanced option for deploying a cloud environment.

To set up the Cloud Administration Node for a private cloud, you must gather information about the servers that will be nodes in your cloud, their networking configuration, and the overall network configuration of your IT environment. You will need to capture this information in the Cloud Installation Dashboard views: Server Types, Connections, and Networks. Consult your IT Administrator to make sure this information is accurate and complete.

In a common use case, deploying a private cloud requires a bare-metal server for a Domain Controller and at least one bare-metal server for a Compute Region Controller, which offers virtualized compute, network, and storage.

- After the Domain Controller and Compute Region nodes are installed, they are known as elastic. In the Cloud Installation Dashboard, this means they can be expanded or contracted to additional bare-metal servers to offer a variety of scales of virtualized resources to your cloud self-service users.
- A private cloud also supports hybrid Compute Region Controllers. This means you can use Compute Region Controllers from other OpenStack compatible clouds as if they were part of this cloud to expand the amount of resources available to your self-service users.

**Best Practice:** For long-term elastic cloud usage, choose a private cloud or a hybrid-only cloud environment.

**Note:** For instructions on how to install a private cloud, review the information in the **Cloud Installation Dashboard > Environment > Setup > Setup** tab.

## ***Hybrid-Only Cloud***

In this type of cloud environment, you deploy a hybrid-only cloud on your own bare-metal server.

To set up the Cloud Administration Node for a hybrid-only cloud, you must gather information about the servers that will be nodes in your cloud, their networking configuration, and the overall network configuration of your IT environment. You will need to capture this information in the Cloud Installation Dashboard views: Server Types, Connections, and Networks. Consult your IT Administrator to make sure this information is accurate and complete.

In a common use case, deploying a hybrid-only cloud requires a single bare-metal server for a Domain Controller. Because of the hybrid nature of this cloud, Compute Region Controllers (virtualized compute, network, and storage) are used by other OpenStack-compatible clouds.

**Best Practice:** For a long-term elastic cloud, choose a hybrid-only cloud environment or a private cloud.

HP Cloud ([www.hpcloud.com](http://www.hpcloud.com)) is an example of an OpenStack-compatible public cloud that can serve as the Compute Region Controller to a hybrid-only cloud.

**Tip:** After a hybrid-only cloud is set up, it can be converted to a private cloud at a later time. A separate HP Cloud account is required for using their resources and usage rates will apply. If you already have another OpenStack-compatible cloud in your environment, you can use that cloud's Compute Region Controller in this hybrid-only cloud to optimize the advanced functionality of the Cloud Infrastructure.

## ***Virtual Hybrid-Only Cloud***

In this type of cloud environment, you deploy a hybrid-only cloud in a virtual environment. This is the simplest cloud to deploy because the Cloud Administration Node does not need to be modified for your IT environment. The virtual machine used for the Domain Controller can be configured to match the default setup of the Cloud Administration Node.

A virtual hybrid-only cloud requires only one other virtual machine for the Domain Controller. This type of cloud cannot easily be converted to hybrid-only or a private cloud.

**Best Practice:** For short-term usage, such as for proof-of-concept, education about cloud environments, or development purposes, choose a virtual hybrid-only cloud.

## **Server Types View**

In the Server Types view, you can capture or use existing server definitions about the networking ports for various servers in your bare-metal cloud environment. If all servers in your cloud are the same machine type (homogeneous nodes), you are not required configure anything in this view. HP recommends that you keep and accept the default settings.

**Tip:** Based on your network and hardware infrastructure, you can accept defaults for or customize Server Types.

## Connections View

In the Connections view, you define a connection set. Review one of the Network Modes along with a list of logical connection interfaces with each one being bound to a port, defined by bandwidth and port number.

Examples of Logical Interface values are: `intf0`, `intf1`, and `intf2`.

For each Logical Interface, a Network Mode is defined, such as `dual`, `single`, `team`, and so on.

**Best Practice:** If your Connection Type is `dual` and the interfaces are not contiguous, use the Edit Connection view. For example, if `eth0` is `private 1` and `eth2` is `public`, the default settings assume that you have `eth0` of the member nodes connected to Private Network1 and `eth1` connected to a public network.

**Tip:** Based on your network and hardware infrastructure, you can accept defaults for or customize Connections.

## Networks View

In the Networks view, you will see the following types of networks: `admin`, `bmc`, `bmc_vlan`, and `public`.

These types of networks are used for different purposes in the cloud, such as:

- **admin** - networking between the nodes and the Cloud Administration Node
- **public** - the public or corporate network attached to your cloud infrastructure

For each Logical Interface, a Network Type such as `admin`, **`bmc`**, **`bmc_vlan`**, and **`public`**, is defined.

Examples of Logical Interface values are: `intf0`, `intf1`, and `intf2`.

**Note:** Based on your network and hardware infrastructure, you can accept defaults for or customize Network Types.

# HP Cloud Connector Cloud Services Reference Architectures (CSRA)

HP Cloud Infrastructure consists of individual services installed and integrated together across one or more machines to form a cloud environment. Each service is large grained with a REST API and can consist of many internal components that can also be distributed across one or more machines when installed. It is through plug-in points in each service integrated to the service API that services are composited together and integrated into a single cloud solution. Each service has a code name and is responsible for a separate set of functionality that is encapsulated behind its API.

To learn more about CSRA, review the following sections:

- ["Services " below](#)
- ["Controllers " on the facing page](#)
- ["Example: Cloud Infrastructure " on the facing page](#)

## Services

The list of possible services that can be deployed for HP Cloud Infrastructure are:

- **Keystone:** OpenStack service to provide identity management, access token, and service catalog functionality.
- **Glance:** OpenStack service to provide for discovering, registering, and retrieving virtual machine images.
- **Eden:** HP service to provide foundation support for other HP services in the area of security and management.
- **Peer:** HP service to provide for discovering, registering, and retrieving resource pool definitions for compute resources.
- **Eve:** HP service to provide for provisioning lifecycle of a TOSCA-based infrastructure topology template of compute resources in a defined resource pool.
- **Focus:** HP service to provide for discovering, registering, versioning, and retrieving of document types necessary to describe TOSCA-based infrastructure topologies.
- **Nova:** OpenStack service to provide a cloud computing fabric controller, the main part of an IaaS environment.

**Note:** HP Cloud Administration Dashboard is based on OpenStack Horizon for user and project management of resources.

## Controllers

While each service can be individually deployed, HP Cloud Infrastructure groups these services into two distinct groups for ease of architectural description:

- **Domain Controller:** Contains services that are considered singleton for a cloud environment, such as Keystone, Glance, Eden, Peer, Eve and Focus, and define the boundaries of the cloud environment from an identity standpoint. For small-scale cloud environments, the Domain Controller can be installed on a single node (machine) that is called a Domain Controller Node. Domain Controller Node sizing depends on the scale of the cloud solution.
- **Compute Region Controller:** A pool of resources, such as compute and storage, that can be consumed through a service API by consumers of the cloud, such as Nova.

A Compute Region Controller consists of the following internal components:

*Cloud Controller*—Can be installed on the same Domain Controller or separated out for large-scale cloud environments.

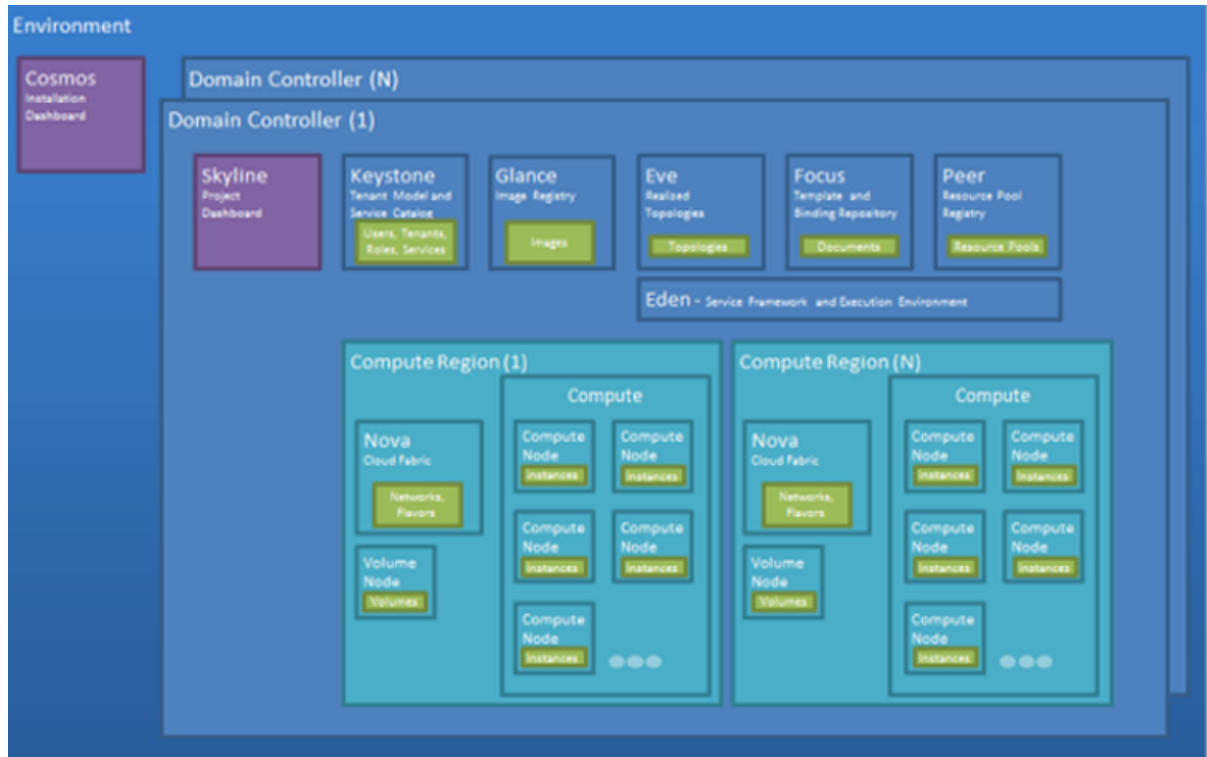
*Compute Worker*—Must be installed on all Compute Region nodes that make up the pool of compute resources to be consumed by cloud consumers. In a small cloud environment, a single node is used as a Cloud Administration Node. The Cloud Administration Node contains the cloud installation services, a controller node that contains the domain controller services, the compute region, cloud controller and volume services, and one or more compute nodes to make up the compute resource pool.

*Volume Controller*—Can be installed on the same Domain Controller Node or separated out for large-scale cloud environments.

**Note:** This can be scaled-out significantly, depending on the number of Compute Region Controllers required and the scale requirements. A cloud environment can include one or more Compute Region Controllers of a certain type or it can exclude Compute Region Controllers. For example, a cloud environment can have more than one Compute Region Controller that is divided by geography, availability, organization, hardware characteristics, and so on.

## Example: Cloud Infrastructure

The following diagram is an example of a cloud infrastructure that includes multiple Domain Controllers, where within on Domain Controller there are multiple Compute Region Controllers.



# HP Cloud Connector Physical Infrastructure Reference Architectures (PIRA)

To prepare your Cloud Environment, review the prerequisites, requirements, and examples that are described and illustrated in the following sections:

- ["Prerequisites" below](#)
- ["Infrastructure Sizing " on next page](#)
- ["Server Configuration " on next page](#)
- ["Switch Configuration " on page 23](#)
- ["Storage Configuration " on page 24](#)
- ["Cloud Infrastructure Network Topology " on page 24](#)
- ["Cloud Infrastructure Production Deployments " on page 25](#)

## Prerequisites

An HP Cloud Infrastructure requires the following nodes: Cloud Administration, Domain Controller, and Compute Region.

Node	Description
Cloud Administration Node (Required)	Ubuntu12.04 OS where Cloud Infrastructure media is booted. Bootstraps the install to set up the Domain Controller Node and Compute Region Node.
Domain Controller Node (Required)	Hosts Cloud Infrastructure services.
Compute Region Controller Node (Required)	Hosts provisioned virtual machines. Multiple Compute Region Controller Nodes can be created to provide more cloud capacity.
CDA Node (Optional)	HP Continuous Delivery Automation is a complete dev-ops solution that integrates with Cloud Infrastructure services.

CSA Node (Optional)	HP Cloud Service Automation is a portal to subscribe application/platform/infrastructure offered by CDA.
Windows Client (Optional)	This machine has a connection to the private network and will provide browser and SSH access to the Cloud Administration Node and other nodes.

- HP recommends that the Domain Controller and Compute Region Controller nodes are on separate bare-metals. When a Compute Region Controller Node reports availability of resources to host new virtual machines, it does not account for the memory and CPU consumed by Cloud Infrastructure services. By having the Compute Region Node on a separate node helps compute nodes avoid over-provisioning virtual machines.
- HP recommends that HP Continuous Service Automation (HP CSA) and HP Continuous Delivery Automation (HP CDA) are on the Domain Controller Node to make sure all resources assigned to the Domain Controller Node are optimized, based on HP CDA and HP CSA sizing guidelines.
- If you need a Windows Client in your cloud environment, HP recommends that you add it.

## Infrastructure Sizing

For an HP Cloud Infrastructure, HP recommends sizing requirements that are described in the following sections:

- ["Server Configuration " below](#)
- ["Switch Configuration " on the facing page](#)
- ["Network Configuration " on the facing page](#)
- ["Storage Configuration " on page 24](#)

## Server Configuration

Node Type	Virtual/Physical Node	CPU Cores	Memory	NICs	OS	Notes
-----------	-----------------------	-----------	--------	------	----	-------

Cloud Administration Node	Virtual Only	>= 2	>= 12GB	>= 2*	Ubuntu Server 12.04 LTS (64 bit)	* External Internet Connection Required
Domain Controller Node	Either	>= 2	>= 32GB	>= 1**	Ubuntu Server 12.04 LTS (64 bit)	** External Internet Connection Required for Hybrid Cloud/HP Cloud (Public) Usage
Compute Region Controller Node	Physical Only	>= 4****	>= 32GB	>= 1	Ubuntu Server 12.04 LTS (64 bit)	*** Intel or AMD Hardware Virtualization Support Required. The CPU Cores and Memory requirements must align with VM instances hosted by the compute node. Ideally, maintain 1 CPU core to 2 vCPU ratio.

## Switch Configuration

Your environment must support the 802.1Q specification (VLAN tagging/trunking) for tagged networks or it must provide physical switches and wiring for 3-4 networks, depending on the cloud configuration.

## Network Configuration

Network	Speed	Tagged*	Required	Notes
Admin Network	1gb	N	Y	

BMC Network*	1gb	Y (100)	N	* Bare-metal control network
Public Network*	1gb	Y (300)	Y	* Must provide a pool of IP addresses for floating IP assignments. Size of pool depends on usage requirements.
Private Network	1gb	Y (500)	Y	

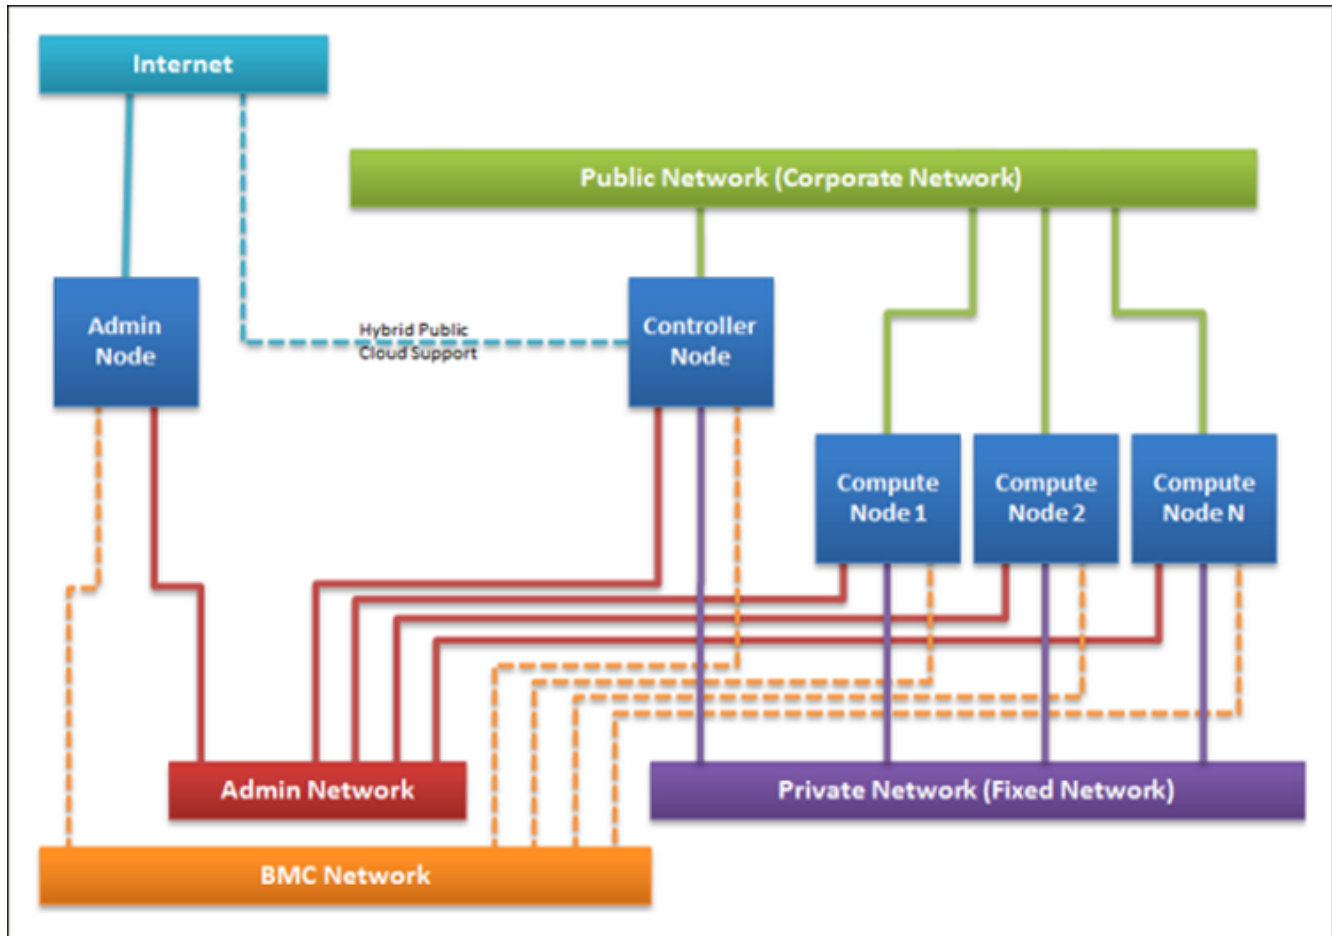
\* Default configuration of network tagging to combine logical networks on to a single physical network.

## Storage Configuration

Node Type	Internal Storage	Partition Scheme	File System Mounts	Notes
Cloud Administration Node	>= 20GB			
Domain Controller Node	>= 60GB	CSRA:Block Storage	/tmp: >= 200GB /var/lib/glance/images: >= 300GB* LVM Volume Group: >= 2TB	* Shared across controller and compute nodes
Compute Region Controller Node	>= 60GB		/tmp: >= 150GB /var/lib/glance/images: >= 300GB*	* Shared across controller and compute nodes

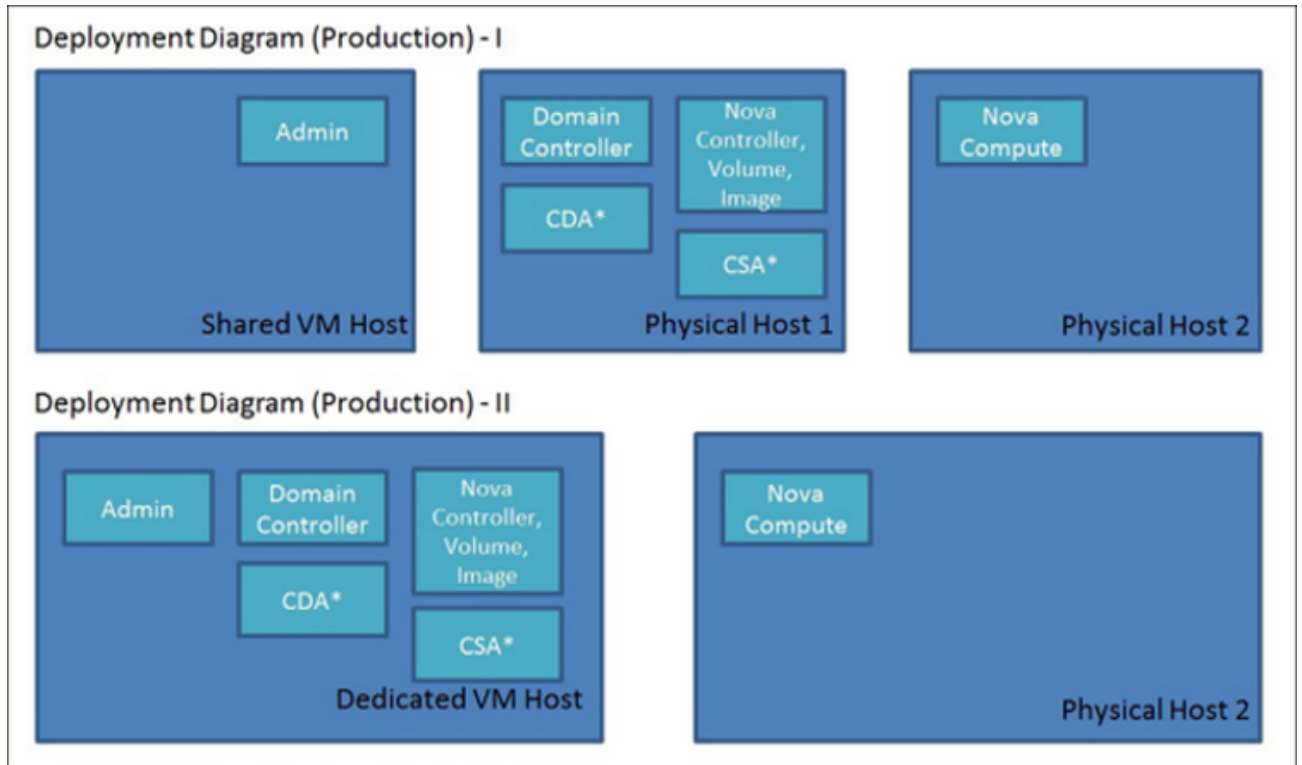
## Cloud Infrastructure Network Topology

The following diagram is an example of a logical network topology for an HP Cloud Infrastructure:



## Cloud Infrastructure Production Deployments

The following diagrams are examples of production deployment architectures:



# Bootstrap the Cloud Administration Node Using the ISO

**Note:** The ISO image must be located in the vCenter's datastore.

To bootstrap the Cloud Administration Node using the ISO so that it sets up the Cloud Administration Node as a guest virtual machine on vCenter, using ESX 4.1.0 or higher:

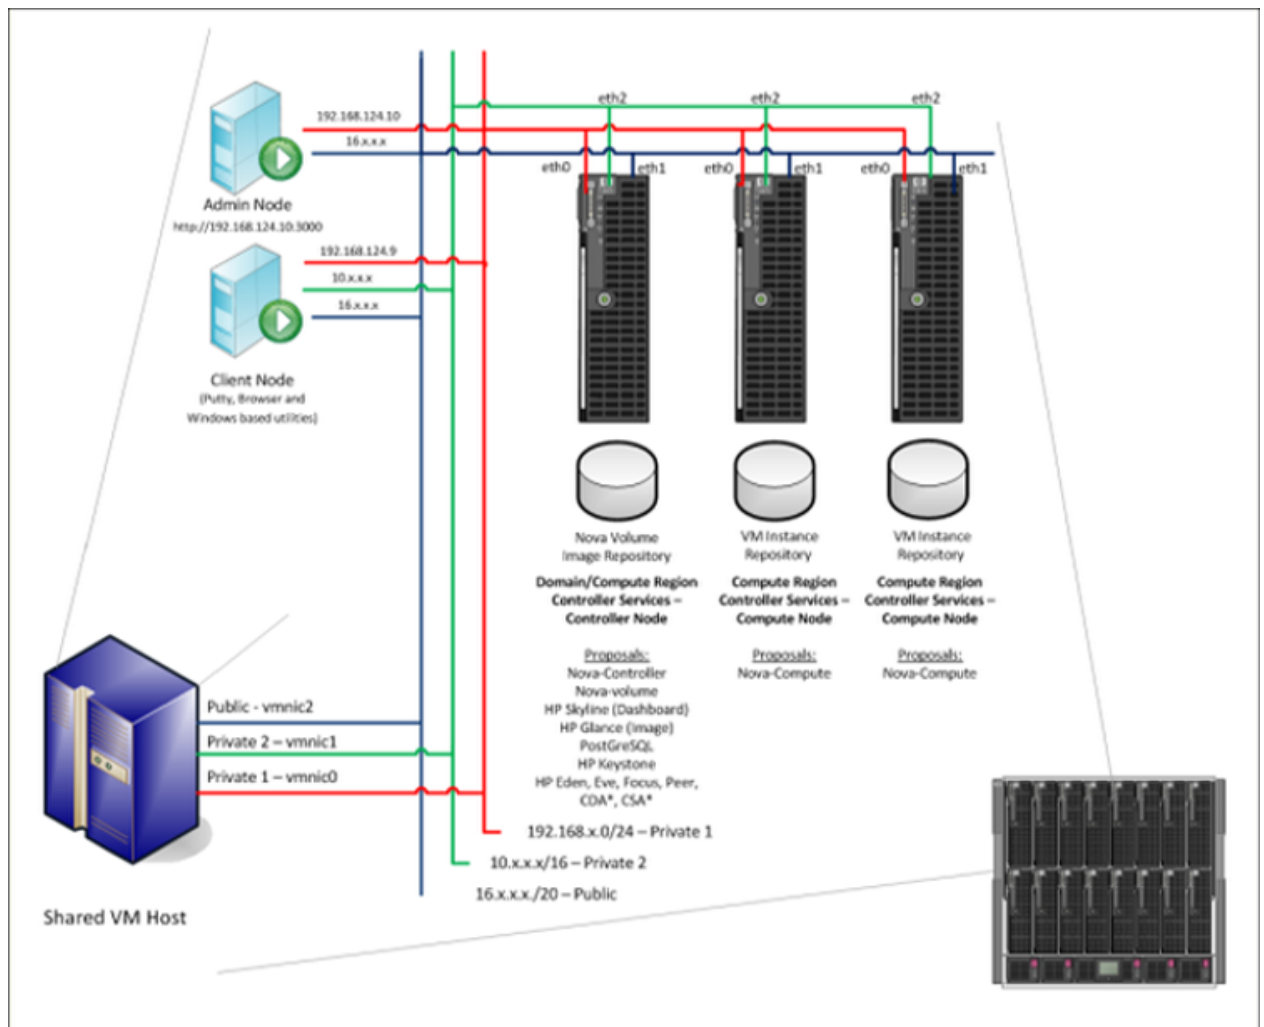
1. Make sure the Cloud Administration Node is powered OFF.
2. Edit the virtual machine settings.
3. Set the CD/DVD to connect at power ON and point it to the ISO file in the "Datastore ISO file" selection. The virtual machine will begin to boot.
4. In the "Configure the Network" dialog box, select eth0 and then enter required information:
  - a. Admin node's IP address. Select **<Continue>**.
    - i. Example: 192.168.124.10
  - Tip:** If you provide a different IP Address and related settings, make sure the Cloud Installation Dashboard's Networks page > Admin tab is also set up to reflect a non-default IP address range and its details.
  - b. Netmask. Select **<Continue>**.
    - i. Example: 255.255.255.0
  - c. Gateway. Select **<Continue>**.
    - i. Example: 192.168.124.1
  - d. Name server. Select **<Continue>**.
    - i. Example: 192.168.124.10
5. Wait until the system finishes detecting the link on eth0.
6. In the "Configure the Clock" dialog box, specify the appropriate time zone.
7. Wait until the ISO image completes the installation process. When this process is completed, the Cloud Administration Node displays its login screen.
8. Power DOWN the virtual machine.
9. Edit the virtual machine settings so that it will not boot from the ISO again. See ["Example of](#)

[Server Deployment " on next page](#)

10. Power UP the virtual machine. The Cloud Installation Dashboard has been successfully installed on the Cloud Administration Node.
11. Go to ["Launch the Cloud Installation Dashboard " on the facing page.](#)

## Example of Server Deployment

In the following diagram, the Cloud Administration Node is virtual and the other nodes are on bare-metal.



## Launch the Cloud Installation Dashboard

After you boot from the ISO, you are ready to launch the Cloud Installation Dashboard to set up the Cloud Administration Node.

To launch the Cloud Installation Dashboard:

1. From the Windows virtual machine client, open a browser that does not have a proxy set.
2. In the browser URL, enter **http://192.168.124.10:9000** to launch the Cloud Installation Dashboard user interface. **IMPORTANT:** Use Google Chrome or Mozilla Firefox to access the Cloud Administration Node.
3. In the Environment tab, review the types of servers, connections, and networks in your environment. ["Environment Tab " on page 14](#)
4. Go to ["Set Up the Cloud Administration Node " on next page.](#)

Related Topics: ["Bootstrap the Cloud Administration Node Using the ISO " on page 27.](#)

## Set Up the Cloud Administration Node

The following prerequisites assume that the Cloud Administration Node is using vCenter as a virtualized environment. Administrators can choose other virtualized environments that meet CSRA and PIRA requirements.

### **Before you begin setting up the Cloud Administration Node:**

- The Cloud Administration Node must be bootstrapped using the ISO. The release product ISO image file name is "Cloud\_Installer\_1\_2.iso". See ["Bootstrap the Cloud Administration Node Using the ISO " on page 27](#).
- Edit your network. See ["Customize Admin and Public Network" below](#).

### **After you have verified that these prerequisites are fulfilled:**

- Go to ["Complete the Setup Process " on page 32](#)

Related Topics: ["HP Cloud Connector Cloud Services Reference Architectures \(CSRA\) " on page 18](#) and ["HP Cloud Connector Physical Infrastructure Reference Architectures \(PIRA\) " on page 21](#).

## Customize Admin and Public Network

As a best practice, customize Admin and Public Networks for the following reasons:

- During bootstrapping the Cloud Administration Node using the ISO, you have provided IP address, DNS, Gateway configuration other than the one shown in the examples, such as 192.168.124.10. This indicates that you would like use your own set of private IP ranges and settings.
- You want to use an existing set of Public IP addresses to provide external access directly to the participating hosts/node.

## Customizing Admin Network

1. In the Environment tab, select **Networks**.
2. Click **Edit Network** for the Network Type **admin**.

3. Review each tab: Network, VLAN, and Router. Modify these, based on your required private network settings.
4. Click **Update Network** to save your changes.
5. Click **Edit Address Ranges** (in the drop-down list next to **Edit Network**) for the Network type **admin**.
6. Review the Node Type, IPV4 Start Addr, IPV4 End Addr configurations. The following Node Types are listed:
  - a. **admin**: IP range for the Cloud Administration Node; default value is 192.168.124.10 to 192.168.124.11
  - b. **dhcp**: Temporary IP address assignment for the PXE booted nodes; default value is 192.168.124.21 to 192.168.124.80
  - c. **host**: Permanent admin network IP address assignment for the participating nodes; default value is 192.168.124.81 to 192.168.124.160
  - d. **switch**: This is for test purposes only.
7. Modify the ranges to meet your requirements.
8. Click **Update Address Ranges** to save your changes.

**Note:** HP recommends that you keep the default values because this network type is considered an isolated, private network. Keeping the default values will overwrite any non-default IP address provided during bootstrapping of the Cloud Administration Node using the ISO. You still can access Cloud Installation Dashboard using the default URL: <http://192.168.124.10:9000>.

## Customizing Public Network

To provide external access directly to the participating nodes, make sure you have two sets of contiguous IP address ranges:

- **Range 1:** 16.x.x.2 to 16.x.x.14. This range is used for PXE booted nodes.
- **Range 2:** 16.x.x.122 to 16.x.x.224. This range is used for Floating IP assignment of provisioned instances within OpenStack. It will be used when you create a Floating IP Network in the Cloud Administration Dashboard. See ["Configure a Floating Network " on page 52](#)

1. In the Environment tab > select **Networks**.
2. Click **Edit Network** for the Network Type **public**.
3. Review each tab: Network, VLAN, and Router:

- a. In the Network tab, specify the subnet, netmask, and broadcast explicitly for your public IP address range and set Bridge Enabled to False.
  - b. In the VLAN tab, set **VLAN Enabled** to False. Do not change the **VLANID**.
  - c. In the Router tab, specify the Router value explicitly for your public IP address range. The Router Preference default is sufficient.
4. Click **Update Network** to save your changes.
5. Click **Edit Address Ranges** (drop-down list next to Edit Network) for Network Type **public**.
6. In the Edit Address Ranges window, **Network Type** and **Subnet** values are pre-populated by default, with the correct settings.
7. Click **Add Address Range**.
  - a. Set **Node Type** to **host**.
  - b. Set **IPv4 Start Addr** to the start of the IIP address range 1.
  - c. Set **IPv4 End Addr** to the end of the IP address range 1.
8. Click **Update Address Ranges** to save your changes.

## Complete the Setup Process

To finish setting up the Cloud Administration Node:

1. In the Cloud Installation Dashboard, go to the Environment tab.
2. In the Setup view, select the Setup tab.
3. Click **Complete Setup**.
4. In the "Confirm Complete Setup - Private Cloud" window, specify the following configuration information:
  - a. In the **Network Mode** drop-down list, select *dual*.
  - b. In the **IPMI/BMC Network Status** drop-down list, select *Disabled*.

**Tip:** The enable IPMI/BMC feature is for test purposes only.
5. Click **Complete Setup**. This action installs underlying components, such as TFTP Boot, DHCP, DNS, NTP servers, and monitoring tools, on the Cloud Administration Node. This may take several minutes.
6. A working status message displays in the Setup Complete column.

**Note:** Any failure during this part of the process requires rebuilding the Cloud Administration Node through the bootable ISO.

7. Click **Show Install Log** to view the install activity in the /var/log/install.log.
8. Click **Refresh** to monitor the install progress.
9. After the Cloud Installation Dashboard is installed, the screen will display the timestamp in the Setup Complete column.
10. If you did not already do this, click **Show Install Log** to view the install activity in the /var/log/install.log.
11. The Cloud Administration Node is now ready to be configured. Go to "[Configure the Cloud Administration Node](#)" on next page.

# Configure the Cloud Administration Node

## Before you begin configuring the Cloud Administration Node:

- The Cloud Administration Node must have Internet connectivity and you must download additional packages before you can deploy Cloud Infrastructure services.
- Make sure you have prepared the Cloud Administration Node as explained in "[Set Up the Cloud Administration Node](#)" on page 30.

## After you have verified these prerequisites are fulfilled:

- Go to "[Download Third-Party Packages](#)" below.

## Download Third-Party Packages

The Cloud Administration Node requires temporary access to the Internet to download and install required third-party packages, such as MongoDB.

**Note:** The settings for public access to the Cloud Administration Node will be disabled after the reboot.

To download the third-party packages:

1. Click **Deploy Cloud**. A Cloud tab displays in the Cloud Installation Dashboard navigation panel.
2. In the Prerequisites view, click **Edit Prerequisite**.
3. In the Edit Cloud Administration Node Internet Access Prerequisite window, select **eth1** from the **Network Interface** drop-down list.

**Note:** eth1 is connected to a network that provides internet access.

4. Keep the default for **Network Configuration** checked as **DHCP**. If your IT policies restrict using a DHCP-provided address, provide a static IP Address, Network Mask, Gateway and DNS Address.
5. Set **HTTP Proxy Information** applicable for your environment.
  - a. Enter the **Host**, such as web-proxy.corp.xx.com.
  - b. Enter the **Port**, such as 8080.

- c. (Optional) Enter the **Username** and **Password** for the HTTP Proxy Information.
6. Click **Update Prerequisite**.
7. Click **Complete Prerequisite for the Cloud Administration Node Internet Access** entry.
8. In the confirmation dialog box, click **Complete Prerequisite**.
9. Click **Complete Prerequisite for the Build MongoDB Installation Module**.
10. In the Prerequisite view, click **Complete Prerequisite**. It may take several minutes for the MongoDB install packages to finish downloading from the Internet to the Cloud Administration Node.
11. In the confirmation dialog box, click **Complete Prerequisite**.
  - a. Make sure that a timestamp (yyyy-mm-dd hh:mm) displays in the **Completed** column.
  - b. A success message displays when the MongoDB prerequisites are configured.
  - c. If the prerequisites fail to complete, an error message displays. See ["Troubleshooting " on page 62](#) for more information.
  - d. You are now ready to deploy the Domain Controller Node. Go to ["Deploy and Allocate the Domain Controller Node " on next page](#).

## Deploy and Allocate the Domain Controller Node

The Domain Controller contains infrastructure services that are considered singleton for a cloud environment, such as Keystone, Glance, Eden, Peer, Eve and Focus. From the identity standpoint, these services define the boundaries of the cloud environment. For small-scale cloud environments, the Domain Controller can be installed on a single node (machine) that is called a Domain Controller Node. Domain Controller Node sizing depends on the scale of the cloud solution. See ["HP Cloud Connector Physical Infrastructure Reference Architectures \(PIRA\) " on page 21](#).

### **Before you begin deploying and allocating:**

- Make sure the Cloud Administration Node prerequisites have been fulfilled. ["Configure the Cloud Administration Node " on page 34](#).
- Make sure you have access to the Cloud Installation Dashboard at <http://192.168.124.10:9000>. See ["Launch the Cloud Installation Dashboard " on page 29](#).







To help you plan for and then configure the Domain Controller, review the following sections:

- ["Deploy and Allocate the Domain Controller Node " below](#)
- ["Deployment States " on page 38](#)
- ["Create an Alias for the Domain Controller " on page 38](#)
- ["Complete Storage Configuration " on page 39](#)
- ["Apply Domain Controller Barclamps " on page 39](#)
- ["Appendix A: Domain Controller Barclamps " on page 67](#)

## Deploy and Allocate the Domain Controller Node

To deploy and allocate the Domain Controller Node:

1. In the **Cloud Installation Dashboard > Cloud** tab, select the Domain Controllers tab.
2. Click **Deploy Domain Controller**.
3. Enter the User Name: **crowbar** and Password: **crowbar**, and then click **Log In**.







- a. In the Node Dashboard view, the "admin" node displays a Ready state .
  - b. The member node for hosting the Domain Controller must be powered **ON**.
4. When the Domain Controller node boots up, the Cloud Administration Node automatically installs through a PXE (network) boot, the Ubuntu 12.04 LTS operating system.
5. When the operating system is being installed, the Node Dashboard view displays the Domain Controller node:
  - a. Initially, the deployment state indicator is Unknown .
  - b. The Domain Controller Node displays a generated, MAC Address. **Tip:** Watch the Domain Controller console to monitor the installation process.
6. Wait until the Node Dashboard view shows the Domain Controller node as Waiting .
7. To confirm that the MAC Address is associated with the Domain Controller node, review the Domain Controller console. The HOSTNAME in the Domain Controller console must match the MAC Address that is displayed in the Node Dashboard view.
8. In the Cloud Installation Dashboard, select **Nodes > Bulk Edit**.
9. In the Bulk Edit view, click the **Allocate?** checkbox and then click **Save**.
10. In the Cloud Installation Dashboard, select **Nodes>Dashboard**.
11. Wait for the Domain Controller node's deployment state to change from Pending  to Ready 
  - a. Watch the Domain Controller Node remote console to monitor the operating system installation. Some of your hardware requires non-free firmware files to operate. The firmware can be loaded from removable media, such as a USB stick or floppy.
  - b. If a dialog box displays, prompting you to load missing firmware, provide the appropriate driver.
12. When the Domain Controller node installation is complete, the Node Dashboard displays the Domain Controller node as Ready .

The Domain Controller login prompt displays the MAC Address that is displayed in the Node Dashboard view.

Related Topic: ["Deployment States " on next page](#)

## Deployment States

The following table explains the deployment state icons that display in the Cloud Installation Dashboard.

Icon	Status	Description	User Action
	Ready	Requested services provisioned.	Configure as needed.
	Waiting (blinking)	Waiting for user input.	Node waiting to be allocated.
	Pending (solid)	Hardware and operating system installation.	None. Crowbar is provisioning the node.
	In Process (spinning)	Crowbar and Chef actively deploying.	None. Crowbar is provisioning the node.
	Failed (blinking)	Failure detected operating on node.	Correct the problem.
	Unknown	In between states or not reporting for 20 minutes, such as powered OFF.	Restart server.

## Create an Alias for the Domain Controller

As a best practice, rename the generated Domain Controller node name to a meaningful name that represents its functionality.

To create an alias name for the Domain Controller node:

1. Select the Domain Controller Node's MAC Address. A details view for the Domain Controller Node displays.
2. Click the **Edit** link.
3. Change **Alias** to a meaningful name, such as **Controller**, and then click **Save**.

**Note:** The **Full Name** of the Domain Controller Node does not change. The MAC address remains the HOSTNAME on the deployed Domain Controller.

## Complete Storage Configuration

Make sure you understand the following HP Cloud Infrastructure concepts and architecture:

- ["HP Cloud Connector Cloud Services Reference Architectures \(CSRA\)" on page 18](#)
- ["HP Cloud Connector Physical Infrastructure Reference Architectures \(PIRA\)" on page 21](#)

A controller node that has been designated as the Domain Controller and Compute Region Controller(s) requires external or internal storage that is configured for:

- **Nova-Volume**—Provides persistent storage for cloud virtual machine instances.
- **Image Repository**—Provides storage for all launch-able images.
- **SQL Data/Log**—Hosts the Domain Controller and Compute Region Controller databases.


**Note:** The following instructions apply when you host the Domain Controller and Compute Region Controller services on the same node.

To configure storage:

1. As a best practice, add an external SAN storage (through iSCSI or Fiber Channel) as a block device.
2. Extend the root volume to accommodate the Image Repository and SQL Data/Log.
3. Add a raw device for Nova-Volume.

## Apply Domain Controller Barclamps

Cloud Infrastructure services are delivered as barclamps. Barclamps are a mechanism to install and configure a service on the Domain Controller Node.




**Caution:** These barclamps have dependencies and *must* be applied in the order they are listed in the Cloud Installation Dashboard. Each barclamp must be successfully applied, one at a time, and Ready  *before* you apply the next barclamp in the list. The following table lists the barclamps in the order they must be applied.

Barclamp	Description
----------	-------------

Postgres 915	Provides a PostGreSQL based database engine to the Domain Controller and Compute Region Controller(s).
HP Keystone 201211	Provides Identity Management services to the Cloud infrastructure.
HP Glance 201211	Provides Image Repository services to launch virtual machine instances.
Rabbitmq 271	Provides a RabbitMQ based message queuing mechanism to Cloud Infrastructure services.
Mongodb 284	Provides MongoDB based database engine for the Topology Provisioner service.
HP Eden 100	Provides a service framework for all Domain Controller services.
HP Peer 100	Provides a Resource Pool Registry service for the Domain Controller service.
HP Eve 100	Provides a Topology Provisioning service for the Cloud Infrastructure.
HP Focus 100	Provides a Topology Document Repository.
Hp Skyline 120	Provides the Cloud Administration Dashboard.

To apply barclamps:

1. In the Cloud Installation Dashboard, select **Barclamps > Cloud Infrastructure**. Barclamps that are specific to Cloud Infrastructure are displayed.
2. For each barclamp in the list:
  - a. Click the **<name of the barclamp>**.
  - b. In **Proposal**, change the default value ("proposal") to a custom name that identifies the barclamp name and the node to which this is being applied.
  - c. Click **Create**.
  - d. Make sure the default values for the proposal are correct. See ["Appendix A: Domain Controller Barclamps" on page 67](#).
  - e. At the bottom of the proposal, the items on the left are **Available Nodes** and the items on the right are the **Roles**.
    - i. In the Edit Proposal view, select the Domain Controller node name link in the **Available Nodes** list to drag it to the appropriate role.

- ii. Do not select the Domain Controller node link icon  and drag it to a role.
- iii. Select the Domain Controller node delete icon  to remove the association of the Domain Controller node from a role.
- f. Click **Apply** and then click **OK**.
- g. Wait for the proposal status to indicate Ready .
- h. If the proposal fails to apply, an error message displays. ["Troubleshooting " on page 62](#) for a solution.
- i. Select **Barclamps > Cloud Infrastructure** to review the barclamp list. The proposal Status is also displayed in the Cloud Infrastructure barclamp list.
- j. Apply the next barclamp in the list, from top to bottom, until all Cloud Infrastructure barclamps related to the Domain Controller Node have been applied.

## Deploy and Allocate Compute Region Controller Nodes

A Compute Region Controller Node is a pool of resources, such as compute and storage, that can be consumed through a service API by consumers of the cloud, such as Nova.

This pool of resources consists of the following internal components:

- **Cloud Controller**—Can be installed on the same Domain Controller or separated out for large-scale cloud environments. In the Nova barclamp, this is also referred as nova-multi-controller role.
- **Compute Worker**—Must be installed on all Compute Region Controller Nodes that make up the pool of compute resources to be consumed by cloud consumers. In the Nova barclamp, this is also referred as nova-multi-controller role.
- **Volume Controller**—Can be installed on the same Domain Controller Node or separated out for large-scale cloud environments. In the Nova barclamp, this is also referred as nova-multi-controller role.

**Tip:** This can be scaled-out significantly, depending on the number of Compute Region Controllers required and the scale requirements. A cloud environment can include one or more Compute Region Controllers of a certain type or it can exclude Compute Region Controllers. For example, a cloud environment can have more than one Compute Region Controller that is divided by geography, availability, organization, hardware characteristics, and so on. See "[HP Cloud Connector Cloud Services Reference Architectures \(CSRA\)](#)" on page 18.

### **Before you begin to deploy and allocate:**

- Make sure the Cloud Administration and Domain Controller nodes have been successfully deployed. See "[Deploy and Allocate the Domain Controller Node](#)" on page 36
- Make sure you have access to the Cloud Installation Dashboard at <http://192.168.124.10:3000>.

To prepare for and then deploy and allocate Compute Region Controller Nodes, review the following topics:


- "[Deploy and Allocate Compute Region Controller Nodes](#)" on the facing page
- "[Create an Alias for a Compute Region Controller](#)" on page 44
- "[Complete Storage Configuration](#)" on page 44



- ["Apply the Nova Barclamp " on page 45](#)
- ["Nova Barclamp Roles " on page 46](#)
- ["Nova Barclamp Proposal Settings " on page 47](#)




## Deploy and Allocate Compute Region Controller Nodes

Deploying and allocating Compute Region Controller Nodes involves bringing up Compute Region Controller Nodes, verifying that it is discovered and allocated, and applying a Nova proposal to both Cloud Controller and Compute Worker nodes.

To deploy and allocate Compute Region Controller Nodes:

1. In the **Cloud Installation Dashboard > Cloud** tab, select the Compute Region Controllers tab.
2. Click **Deploy Compute Region Controllers**.
3. Enter the User Name: **crowbar** and Password: **crowbar**, and then click **Log In**.
  - a. In the Node Dashboard view, the "admin" and "controller" nodes display a Ready state .
  - b. The Compute Region Controller Node must be powered ON.
4. When a Compute Region Controller node boots up, the Cloud Administration Node automatically installs through a PXE (network) boot, the Ubuntu 12.04 LTS operating system.

**Tip:** Watch the Compute Region Controller console to monitor the installation process.
5. The Node Dashboard displays the node as the OS is installed.
  - a. The status indicator is initially gray.
  - b. The node initially displays in the Node Dashboard view with a generated MAC address.
6. Wait until the Node Dashboard view shows the Compute Region Controller Node as Waiting .
7. Wait until the Node Dashboard shows the new node with a flashing yellow state .
8. To confirm that the MAC Address is associated with the Compute Region Controller Node, review the Compute Region Controller console. The HOSTNAME in the Compute Region Controller console must match the MAC Address that is displayed in the Node Dashboard view.

9. In the Cloud Installation Dashboard, select **Nodes > Bulk Edit**.
10. In the Bulk Edit view, click the **Allocate?** checkbox and then click **Save**.
11. In the Cloud Installation Dashboard, select **Nodes > Dashboard**.
12. Wait for the Compute Region Controller node's status to change from Pending  to Ready .
  - a. Watch the Compute Region Controller node remote console to monitor the installation process.
  - b. If a dialog box displays, prompting you to load missing firmware, click **No**.
13. When the Compute Region Controller Node installation is complete, the Node Dashboard displays the Compute Region Controller Node as Ready  and in the same group as the "admin" node.

**Tip:** The Domain Controller login prompt displays the MAC Address that is displayed in the Node Dashboard view.
14. Repeat steps 4 to 13 for all the nodes assigned to the Compute Worker role.

Related Topic: ["Deployment States " on page 38](#)

## Create an Alias for a Compute Region Controller

As a best practice, rename the generated Compute Region Controller Node name to a meaningful name that represents its functionality, such as CloudController, ComputeWorker1, ComputeWorker2, and so on.

To create an alias name for the Compute Region Controller Node:

1. Select the Compute Region Controller Node's MAC Address. A details view for the Compute Region Controller Node displays.
2. Click the **Edit** link.
3. Change **Alias** to a meaningful name and then click **Save**.

Repeat steps 1 to 3 for all nodes assigned to the Cloud Controller and Compute Worker roles.

**Note:** The Full Name of the Domain Controller Node does not change. The MAC address remains the HOSTNAME on the deployed Domain Controller.

## Complete Storage Configuration

Make sure you understand the following HP Cloud Infrastructure concepts and architecture:

- ["HP Cloud Connector Cloud Services Reference Architectures \(CSRA\)" on page 18](#)
- ["HP Cloud Connector Physical Infrastructure Reference Architectures \(PIRA\)" on page 21](#)

A member node that has been designated as the Compute Worker role requires external or internal storage to be configured as OpenStack's instances repository. This storage space would be used by the provisioned instances' non-persistent data, such as the operating system's boot volume or ephemeral disk. After the instance has served its purpose and is deleted, all state is reclaimed, excluding persistent volume.

At this point, configure storage to meet the above-mentioned requirement. HP recommends that you add an external SAN storage through iSCSI or fiber channel, and extend the root volume.

For more information about Cloud Controller role requirements, specifically for nova-volume and image repository, see ["Complete Storage Configuration" on page 39](#) in ["Deploy and Allocate the Domain Controller Node" on page 36](#).

**Note:** The assumption here is that both the Domain Controller role and the Cloud Controller role are applied to the same node.

## Apply the Nova Barclamp

Cloud Infrastructure services are delivered as barclamps. Barclamps are a mechanism to install and configure a service on the Compute Region Control Node.

To apply the Nova barclamp:

1. In the Cloud Administration Dashboard, select **Barclamps > Cloud Infrastructure**.  
Barclamps that are specific to Cloud Infrastructure are displayed. The majority of Cloud Infrastructure barclamps have already been deployed, as indicated by a Ready state ●.
2. In the Cloud Infrastructure view, select **Nova** in the **Name** column.
3. Specify the **Compute Region name** as the proposal.
  - a. Click **Create**.
  - b. Verify that default values are correct for the proposal. Depending on your cloud environment, some Nova attributes must be modified. See ["Nova Barclamp Proposal Settings" on page 47](#).
  - c. Make sure you drag and drop the designated nodes to their appropriate Role. See ["Nova Barclamp Roles" on next page](#)
4. Go to **Barclamps > Cloud Infrastructure** to review the proposal status in the barclamp list.
  - d. Click **Apply** and then click **OK**.

e. Wait for the proposal status to indicate Ready. If the proposal fails to apply, an error message displays. See ["Troubleshooting " on page 62](#) for a solution.

**Tip:** If you want to add another Compute Worker Node later, complete the steps to discover and allocate the new node, and then add the node to an *existing* Nova proposal. Do *not* create a separate proposal for the newly added node.

**Note:** All Cloud Infrastructure barclamps are now applied to the Domain Controller and the Compute Region nodes. The next step is configuring networking (Fixed and Floating Network) for virtual machine instances launched by the Cloud Infrastructure. See ["Configure Network Infrastructure for Virtual Machines " on page 50](#).

Related Topics: ["Nova Barclamp Roles " below](#) and ["Nova Barclamp Proposal Settings " on the facing page](#)

## Nova Barclamp Roles

The Nova barclamp supports the following roles:

- **Nova-multi-controller**—Determines which nodes perform the infrastructure management and API functions. The default node allocation for the controller role is to use the same node as the PostGres barclamp.
- **Nova-multi-compute**—Identifies nodes that act as virtualization hosts. The majority of the nodes in the Compute Region deployment will perform this role.
- **Nova-multi-volume**—Identifies a single node on which a Nova volume will be created. The default node allocation for the volume role is to use the same node as the controller role.

**Tip:** As a best practice, the node used earlier for the Domain Controller is assigned *both* the **Nova-multi-controller** role and the **Nova-multi-volume** role. The remaining member nodes that are designated as Compute Worker nodes can be assigned to the **Nova-multi-compute** role.

Related Topics: ["Apply the Nova Barclamp " on previous page](#) and ["Nova Barclamp Proposal Settings " on the facing page](#)

## Nova Barclamp Proposal Settings

The following table explains how the Nova barclamp is structured and describes its relationship with a proposal.

In the Cloud Installation Dashboard:

- You can edit barclamp attributes and node deployment information.
- You can set attribute values that are based on roles associated with the barclamp.
- Each barclamp may have specific logic that requires minimums or maximums for node assignments.

Attribute	Value	Description
Compute Region	RegionOne	The name of the Compute Region. If you are adding an additional compute region, change this value to Region Two or any custom name.
Postgres	(generated)	The Postgres proposal to use.
Keystone	(generated)	The Keystone proposal to use.
Keystone Service User	Nova	The type of user that Nova uses when authenticating with Keystone.
Keystone Service Password	(generated)	The password for the Nova Keystone authentication user.
Glance	(generated)	The Glance proposal to use.
Verbose	false	Indicates if Nova will run in verbose mode.
Use NoVNC (otherwise VPN-VNC)	true	Indicates the VNC package to use.
Hypervisor	kvm	Indicates what hypervisor Nova should use when spinning up virtual machines. Select qemu if your are running Nova on virtual machines. The default is kvm, but will be switched to qemu if virtual machines are detected.
Name of Volume	nova-volumes	The name of the volume-group created on the nova-volume node.

Type of Volume	Raw	Indicates the type of volume to create. If raw is specified, the system attempts to use the remaining unused disks to create a volume group. If the system does not have additional free drives, the system will switch to local. Local uses a local file in the existing file system, based on other parameters. Volume File Name /var/lib/nova/volume.raw is used when local type is chosen or fallen back to. This field is the name of the file in the file system to use.
Maximum File Size	2000	Specified in gigabytes. When local type is chosen or fallen back to, this defines the maximum size of that file. If the file is too big for the file system, the size of the file will be capped at 90% of the free space in that file system (at the time of creation).
Disk Selection Method	all	When raw type is chosen, this indicates how to select the disks to use for volume construction. "all" means use all available. "first" means use the first one detected. "selected" means use the disks selected in the list below this option.

Related Topics: ["Apply the Nova Barclamp " on page 45](#) and ["Nova Barclamp Roles " on page 46](#)

## Cloud Utilities Barclamp

The Cloud Utilities Barclamp uses the Nagios web interface for scheduling database clean-up tasks and monitoring RabbitMQ on the Cloud Controller Node.

## Clean-up Scripts

The cloud utilities scripts keep the database meaningful by removing all references to deleted resources and invalid data. The following three new service checks are scheduled through Nagios to run once daily, between 23:00 – 00:00 hours local time.

- **Cloud clean db objects**—Runs the empty\_trash script to delete users, projects, and flavors.
- **Cloud clean deleted objects**—Runs the clean\_databases script.
- **Cloud clean outdated objects**—Runs the clean\_historical\_data script.

## Apply the HP Cloud Utils 120 Barclamp

1. In the Cloud Administration Dashboard, select **Barclamps > Cloud Infrastructure**.
2. Click the last barclamp in the list: **HP Cloud Utils 120**.

3. Click **Create** and enter a proposal name.
4. Click the **Services** link in the left panel of the Nagios web interface to see the cloud service checks, including the Multi Nova RabbitMQ service.
  - a. Review the Attribute and Node Deployment sections of the proposal:

Attribute	Value	Description
Postgres	(generated)	The Postgres proposal to use.
Keystone	(generated)	The Keystone proposal to use.

## Monitor RabbitMQ

Nagios is also used to monitor RabbitMQ on the Nova Controller Node. If the status changes to CRITICAL, Nagios automatically restarts RabbitMQ.

1. Access the Nagios dashboard (<http://192.168.124.10> or <Cloud Administration Node IP>/nagios3/) using the credentials nagiosadmin/password.
2. Use the following login credentials: nagiosadmin/password.
3. Click the **Services** link in the left panel of the Nagios web interface to see the cloud service checks, including the Multi Nova RabbitMQ service.
4. In the Node Deployment section, make sure you drag and drop the node that has the Cloud Controller role applied in the nova proposal.

# Configure Network Infrastructure for Virtual Machines

After the Cloud Infrastructure is deployed, you must configure the network infrastructure for virtual machines in the Cloud environment.

To prepare for and configure the network infrastructure for virtual instances, review the following topics:

- ["Assumptions " below](#)
- ["Configure a Fixed Network " below](#)
- ["Configure a Floating Network " on page 52](#)

## Assumptions

Your network environment consists of:

- eth0 - private 1 as administration network
- eth1 - public/corporate network attached to your cloud infrastructure
- eth2 - private 2 network available for configuring Fixed IP networks

Related Topics: ["HP Cloud Connector Cloud Services Reference Architectures \(CSRA\) " on page 18](#) and ["HP Cloud Connector Physical Infrastructure Reference Architectures \(PIRA\) " on page 21](#)

## Configure a Fixed Network

1. Go to the Cloud Administration Dashboard (<http://192.168.124.81>).
2. Enter the User Name: **Admin** and Password: **secretword**.
3. Select the Region tab and then Networks.
4. Click **Create Fixed IP**.
5. Specify a **Network Label**.
6. Leave **Project** unselected to enable the fixed network to be used across the Cloud. This

setting indicates that the network is a flat DHCP fixed IP network. Or, set **Project** to a specific project if the fixed network needs to be separate for the project. This requires VLAN configuration on the eth2 port, including the connecting switch.

7. Set **IPv4 CIDR** to **192.168.123.0/24**. A Class B CIDR can be provided to accommodate more virtual machines.
8. Set **Gateway** to **192.168.123.1**.
9. Set **DNS1** to **192.168.124.10**. By default, the Cloud Administration Node acts as a DNS server. However, you can specify your own DNS servers that can be reachable through Public network.
10. Set **Multi-Host** to **True**.
11. Set **Number of Networks** to **1**.
12. Set **Network Size** to **256**.
13. Set **Bridge ID** to **br100** or any custom name.
14. Set **Bridge Interface** to **eth2**. This setting must be unique if you are adding multiple fixed network.

## Example

If the private Network 2 is shared for other purposes, set the VLAN correctly on the switch, with the appropriate tag. Also, create the VLAN interface on both the Domain Controller and Compute Region nodes.

The following commands create a VLAN interface on eth2 with tag 500:

- `sudo vconfig add eth2 500`
- `ifconfig eth2.500 up`

In this example, the following fields are configured as:

- Set **VlanID** to **500**.
- Set **BridgeID** to **br500**.
- Set **Bridge** to **eth2.500**.

## Configure a Floating Network

1. Go to the Cloud Administration Dashboard (<http://192.168.124.81>).
2. Enter the User Name: **Admin** and Password: **secretword**.
3. Select the Region tab and then Networks.
4. Click **Create Floating IP**.
5. Set **IPv4 CIDR** to the static IP range applicable to the corporate/public network. See ["Customize Admin and Public Network" on page 30](#)
6. Set **Interface** to the interface that is configured for public/corporate access, such as eth1.
7. Click **Create Network**.

## Post Deployment Tasks

To validate the compute and topology implementation of the Cloud Infrastructure, the following post deployment tasks are required:

- "Configure an Image Repository " below
- "Create a Keypair " on next page
- "Configure the Default Security Group " on page 55
- "Launch an Instance " on page 55
- "Validate Instance Accessibility " on page 56
- "Create a Resource Pool " on page 57
- "Create an Infrastructure Topology " on page 58
- "Create an Infrastructure Design " on page 59
- "Launch an Infrastructure Design Document " on page 60
- "Validate an Infrastructure Design Document " on page 61

## Configure an Image Repository

An image repository contains required images that must be provisioned in the Cloud Infrastructure, as part of a topology design. The initial, required images are base operating system images, such as:

- **Cirros Image**—[https://launchpad.net/cirros/trunk/0.3.0/+download/cirros-0.3.0-x86\\_64-disk.img](https://launchpad.net/cirros/trunk/0.3.0/+download/cirros-0.3.0-x86_64-disk.img)
- **Ubuntu Image** —<http://uec-images.ubuntu.com/precise/current/precise-server-cloudimg-amd64-disk1.img>

To upload these images to the Cloud Infrastructure:

1. Download the required images to your client node. Using winscp or pscp tools, copy them to Cloud Administration Node's /tftpboot folder. The SSH credentials to access the Cloud Administration Node are **crowbar/crowbar**.

**Note:** The /tftpboot folder requires read/write permission.

2. Log in to the Cloud Administration Dashboard, using User Name: **Admin** and Password: **secretword**.
3. Navigate to **Domain > Images** and then click **Create Image**.
4. In the Create an Image window, enter a **Name** for the image.
5. In **Image Location**, enter the HTTP URL of the Cloud Administration node's **/tftpboot** folder, including the image name, such as <http://192.168.124.10:8091/precise-server-cloudimg-amd64-disk1.img>.
6. In the **Format** drop-down list, select **QCOW2 - QEMU Disk Image**.
7. In **Minimum Disk**, enter **0** (zero).
8. In **Minimum RAM**, enter **0** (zero).
9. Check the **Public** checkbox.
10. Click **Create Image**.
11. In the Images view, click the **Image Name** to validate the image's details and size after it is loaded.

**Note:** To enable your cloud environment with Load Balancer and Chef Server services, see ["Appendix B: Load Balancer Image " on page 75](#) and ["Appendix C: Chef Server Image " on page 80](#)

## Create a Keypair

For a secure environment, every virtual machine instance to be launched must have a keypair injected. This keypair is used to SSH into the launched virtual machine instance.

To create a keypair:

1. In the Cloud Administration Dashboard, go to **Project > Access & Security**.
2. In the Keypairs section, click **Create Keypair**.
3. In the Create Keypair window, enter a Keypair Name, such as **testonlykeypair** or any custom name.
4. Click **Create Keypair**.
5. Click the **Download keypair <name of keypair>**. Save this file to access launched instances using SSH.

## Configure the Default Security Group

To validate an instance's accessibility, the instance must be on a network segment (security group) that has both the PING and SSH ports open.

To do this, edit the default security group:

1. In the Cloud Administration Dashboard, go to **Project > Access & Security**.
2. In the Security Groups section, click **Edit Rules** for the default security group.
3. In the Edit Security Group Rules window, in the Add Rule section, specify the following information:
  - a. For SSH - **IP Protocol**: TCP, **From Port**: 22, **To Port**: 22, **Source Group**: CIDR, and **CIDR**: 0.0.0.0/0.
  - b. For PING **IP Protocol**: ICMP, **From Port**: -1, **To Port**: -1; **Source Group**: CIDR, and **CIDR**: 0.0.0.0/0.
4. Click **Update Security Group Rules** and then add the next rule.

## Launch an Instance

Validate whether you can launch an instance using the newly uploaded image:

1. In the Cloud Installation Dashboard, go to **Project > Images**.
2. Click **Launch** for the newly uploaded image.
3. Provide the following information:
  - a. **Instance Source**: Image
  - b. **Image**: precise-server-cloudimg
  - c. **Server Name**: TestPurposeOnly1 or any custom name
  - d. **Flavor**: m1.tiny
  - e. **Instance Count**: 1
4. Go to **Project > Access & Security** in the Launch Instance window and then select the newly created keypair. Make sure that **default** has a checkmark in the Security Groups section.
5. Click **Launch**.

## Validate Instance Accessibility

To validate whether the virtual machine was successfully launched:

1. In the Cloud Administration Dashboard, go to **Project > Instances**.
2. Locate the recently launched virtual machine instance and then click the **Instance Name**.
3. In the Overview tab, record the **IP Address** assigned to this instance.
4. Go to the Log tab and make sure the following information has been logged:

```
cloud-init boot finished at Mon, 17 Dec 2012 19:04:33 +0000. Up 209.73 seconds
```

5. Validate whether the network infrastructure has been correctly set up to access the recently launched virtual machine. Using the IP address of the instance from the previous section, try to PING and SSH into it.

a. For PING:

- i. Make sure you are on the same network as your virtual machine instance or that the proper routing is enabled.
- ii. Open the command prompt and try pinging the IP Address.

b. For SSH:

- i. Make sure you are on the same network as your virtual machine instance or that the proper routing is enabled.
- ii. Convert the file that you downloaded during keypair creation from a .pem format to a .ppk format. **Tip:** Use PuTTYgen to convert the file format from .pem to .ppk.
- iii. If you are using a Putty Client to do SSH, provide the .ppk file under **Connection > SSH > Auth > Authentication Parameters > Private key file** for authentication.

Or

- iv. If you are already in an SSH window for one of the Compute Region nodes, enter the following command to validate that SSH is working:

```
root@de4-11-5b-b7-b3-6e:/home/crowbar# ssh ubuntu@192.168.123.161 -i
testonlykeypair.pem
Warning: Permanently added '192.168.11.161' (ECDSA) to the list of known
hosts.
Welcome to Ubuntu 12.04.1 LTS (GNU/Linux 3.2.0-34-virtual x86_64)

 * Documentation:  https://help.ubuntu.com/

System information as of Mon Dec 17 19:32:51 UTC 2012

System load:  0.0                       Processes:            59
Usage of /:   33.3% of 1.96GB           Users logged in:     0
Memory usage: 8%                       IP address for eth0: 192.168.123.161
Swap usage:   0%

Graph this data and manage this system at
https://landscape.canonical.com/

0 packages can be updated.
0 updates are security updates.

Get cloud support with Ubuntu Advantage Cloud Guest
http://www.ubuntu.com/business/services/cloud

The programs included with the Ubuntu system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Ubuntu comes with ABSOLUTELY NO WARRANTY, to the extent permitted by
applicable law.

To run a command as administrator (user "root"), use "sudo <command>".
See "man sudo_root" for details.

ubuntu@testpurposeonly1:~$
```

## Create a Resource Pool

To enable and validate domain controller services that launch Infrastructure Design documents, complete the following steps to create a Resource Pool:

1. In the Cloud Administration Dashboard, go to **Project > Resource Pools**.
2. Click **Create Resource Pool**.
3. In the Create Resource Pool window, provide the following information:
  - a. From the **Cloud Type** drop-down list, select HP Cloud Infrastructure - OpenStack (Essex).
  - b. From the **Compute Region** drop-down list, select Domain/RegionOne.
  - c. In the **Name** field, select <ProvideCustomName> or AdminProjectRP.
  - d. From the **Type** drop-down list, select OpenStack.
  - e. From the **Version** drop-down list, select Essex 2012.1.X.
  - f. In the **Region ID** field, enter RegionOne.
  - g. In the **Domain URL** field, enter <http://192.168.124.81:5000/v2.0/tokens> or <URLtoKeystoneHost>.
4. Click **Create Resource Pool**.
5. Validate that the newly created resource pool is listed under **Resource Pools**.

## Create an Infrastructure Topology

**Caution:** Every Infrastructure Topology document must bind with an Infrastructure Design document.

To create an Infrastructure Topology:

1. In the Cloud Administration Dashboard, go to **Project > Documents**.
2. In the Documents view, click **Create Infrastructure Topology**.
3. In the Create Infrastructure Topology window, click the **Help** link in the upper right corner.
4. In the left navigation pane, select **Project > Documents > Infrastructure Topology> Create an Infrastructure Topology** for step-by-step instructions.
5. After you have created and saved a topology document, validate that it is listed in the Documents view.

## Create an Infrastructure Design

**Caution:** Every Infrastructure Topology document must bind with an Infrastructure Design document.

To create an Infrastructure Design document:

1. In the Cloud Administration Dashboard, go to **Project > Documents**.
2. In the Documents view, click the **Help** link in the upper right corner.
3. In the left navigation pane, select **Welcome to Cloud Administration Dashboard Help > Project > Documents > Infrastructure Design > Create an Infrastructure Design** for step-by-step instructions.
  - a. In the Create Infrastructure Design wizard, in the Resource Binding view, enter the following information:
    - i. **Name:** A customer name or TestPurposeDesign1.
    - ii. **Infrastructure Topology:** Select a value from the drop-down list.
    - iii. **ResourcePool:** Select a value from the drop-down list.
4. Click **Next** to advance to the next wizard step called Binding Details.
  - a. In the **Server Group** section, enter the following information:
    - i. **Minimum Instances:** 1
    - ii. **Instance Name Prefix:** test
    - iii. **Machine Flavor:** Select a flavor from the drop-down list.
    - iv. **Machine Image:** precise-server-clouding
    - v. **Key Pair Name:** Select a keypair from the drop-down list.
  - b. In the **Volume Group** section, enter the following information:
    - i. **Minimum Volumes:** (Optional) The minimum number of volumes in the group. An associated server group's Minimum Instances count must be the same as this count to avoid a provisioning failure.
    - ii. **Volume Name Prefix:** Provide the first characters of a volume name.
    - iii. **Volume Size (GB):** The size of one volume in the group.

- c. In the **Network Segment** section, enter the following information:
  - i. **Name:** TestNetwork1
  - ii. **Type:** Select a segment type from the drop-down list, such as Cloud Edge Gateway or Security Group. Cloud Edge Gateway provides both Fixed IP and Floating IP addresses for the virtual machine instance.
  - iii. **Description:** Test Purpose Only
  - iv. **Open Port List:** 22,80
- d. In the **Configuration Manager Service** section, enter the following information:
  - i. **Instance Name Prefix:** Provide the first characters of a host name.
  - ii. **Resource Pool Service:** Select a service from the drop-down list.
  - iii. **Key Pair Name:** Select a key pair from the drop-down list.
  - iv. **Knife Key Pair Path:** Tells Eve where to store the knife private key on the Configuration Management server to be provisioned. The full path and filename must be specified, such as /etc/knife/knife.pem.
- e. In the **Load Balance Service** section, enter the following information:
  - i. **Instance Name Prefix:** Provide the first characters of a name for this entity.
  - ii. **Resource Pool Service:** Select a service from the drop-down list.
  - iii. **Key Pair Name:** Select a key pair from the drop-down list.
5. Click **Create Infrastructure Design**.
6. After you have created and saved a topology document, validate that it is listed in the Documents view.

## Launch an Infrastructure Design Document

After you have created and saved an Infrastructure Topology document and an Infrastructure Design document, you can launch the Design document:

1. In the Cloud Administration Dashboard, go to **Project > Documents**.
2. In the Documents view, click the **Help** link in the upper right corner.
3. In the left navigation pane, select **Welcome to Cloud Administration Dashboard Help >**

**Project > Documents > Infrastructure Design > Launch an Infrastructure Design** for step-by-step instructions.

## Validate an Infrastructure Design Document

To validate an Infrastructure document:

1. Validate the status of a launched infrastructure in **Topologies**:
  - a. In the Cloud Administration Dashboard, go to **Project > Topologies**.
  - b. Check the **State** of the newly launched Infrastructure Document.
  - c. Click the Topology name to open the Topology Detail view.
    - i. Select the Job tab to verify the **State** of a successfully provisioned document is **Succeeded**.
    - ii. Select the Content tab to identify the IP Address allocated to the newly provisioned virtual machine. **Tip**: Search for **ip\_address\_type**.
2. Validate accessibility by using SSH.
  - a. You must have both Fixed IP and Floating IP addresses associated with the newly provisioned virtual machine. Use SSH through both Fixed IP and Floating IP addresses.
  - b. You can also locate this instance of the virtual machine. Go to **Project > Instances** and then click the instance name to see more details.

## Troubleshooting

This section describes Cloud Installation Dashboard known problems and solutions.

### **Problem: Cloud Administration Node displays the "not ready" (grey) state.**

Symptoms	You may not be able to PXE boot any new node.
Primary software component	Cloud Administration Node
Failure message	
Probable cause	You changed the date on the Cloud Administration Node or you brought up a previously allocated node to the Cloud Administration Node.

#### **Solution:**

1. Reboot the Cloud Administration Node.
2. When a new Cloud Administration Node is installed, all previous participating nodes must be kept off of PXE booted through the Cloud Administration Node. If you are in this situation, reinstall of the Cloud Administration Node and its participating node.

### **Problem: Cloud Infrastructure - An error message displays when configuring Cloud Infrastructure prerequisites.**

Symptoms	Completing the prerequisites fails.
Primary software component	Cloud Installation Dashboard
Failure message	Error: There was an error submitting the form. Please try again.

Probable cause	<ul style="list-style-type: none"> <li>• Verify that the Cloud Administration node has both a Public Network and a Private Network configured.</li> <li>• Make sure the proxy information is correct and that the settings are for a working proxy server and port.</li> </ul>
----------------	--

### Solution:

- SSH into the Cloud Administration Node using the credentials crowbar/crowbar and review the following log files:

```
/var/log/apache2/error.log
```

```
/tftpbboot/ubuntu_dvd/extra/prereq/prereq.log
```

- Verify that a Public Network is configured in the Cloud Administration Node.
- Repeat the steps in the **Configure Cloud Administrator Node** section for **Enable Cloud Administration node for Internet Access**.
- Correct the proxy entries and/or specify a different proxy host and/or port.

## Problem: Cloud Infrastructure barclamp proposal fails.

Symptoms	An error message displays and the status turns red if a proposal fails.
Primary software component	Cloud Installation Dashboard
Failure message	Failed to apply the proposal to: <Domain Controller Node name>.
Probable cause	Any number of factors.

### Solution:

- Try applying the proposal again. If that fails:
  - Deactivate the proposal.
  - Delete the proposal and create it again.

- SSH in to the Cloud Administration Node using the credentials: **crowbar/crowbar**.
  - Review the log files /opt/dell/crowbar\_framework/log/production.log and MacAddressHostname>.chef.log.
  - Review the log file /var/log/apache2/error.log.
  - Review the folder /var/tmp/cosmos and verify the installer settings.

## Problem: Cloud Installation Dashboard MongoDB prerequisite installation failed.

This problem has been seen in the Cloud Installation Dashboard while installing MongoDB where the adapter (default is eth 0) or proxy settings in Internet prerequisites are incorrect.

Symptoms	The failure message below is returned when the system attempts to install MongoDB.
Primary software component	Cloud Installation Dashboard
Failure message	"Failed completing the Prerequisite" in the Cloud Installation Dashboard user interface.
Probable cause	Incorrect adapter setting.

### Solution:

- Log in to the Cloud Administration node and review the log file /tftpboot/ubuntu\_dvd/extra/prereq/prereq.log.
- The following shows an example log where the specified interface is eth0 (the primary management interface) or where incorrect proxy settings were specified:

```
Spider mode enabled. Check if remote file exists.

--2013-02-23 07:46:05-- http://www.hp.com/

Resolving www.hp.com (www.hp.com)... 15.192.45.28, 15.240.238.53

Connecting to www.hp.com (www.hp.com)|15.192.45.28|:80... failed:
Connection timed out.
```

```
Connecting to www.hp.com (www.hp.com)|15.240.238.53|:80... failed:
Connection timed out.
```

Giving up.

- Review the interface settings and proxy settings in the “Edit Cloud Admin Node Internet Access Prerequisite” window in the Cloud Installation Dashboard. The interface must point to eth1 and the proxy address and port must be set as suggested by your administrator.

**Note:** The following shows an example log where the specified interface is eth1 (the Internet backbone) and the correct proxy settings are specified.

```
Spider mode enabled. Check if remote file exists.
```

```
--2013-02-23 07:49:44-- http://www.hp.com/
```

```
Resolving web-proxy.ind.hp.com (web-proxy.ind.hp.com)...
16.150.210.10
```

```
Connecting to web-proxy.ind.hp.com (web-proxy.ind.hp.com)
|16.150.210.10|:8080... connected.
```

```
Proxy request sent, awaiting response... 200 OK
```

```
Length: unspecified [text/html]
```

```
Remote file exists and could contain further links, but recursion is
disabled -- not retrieving.
```

```
--2013-02-23 07:49:49--
```

```
http://us.archive.ubuntu.com/ubuntu/pool/main/p/pcre3/libpcrecpp0_
8.12-4_amd64.deb
```

```
Resolving web-proxy.ind.hp.com (web-proxy.ind.hp.com)...
16.150.210.10
```

```
Connecting to web-proxy.ind.hp.com (web-proxy.ind.hp.com)
|16.150.210.10|:8080... connected.
```

```
Proxy request sent, awaiting response... 200 OK
```

```
Length: 16156 (16K) [application/x-debian-package]
```

```
Saving to: `libpcrecpp0_8.12-4_amd64.deb'
```

```
OK ..... 100% 1.34M=0.01s
```

## Problem: Domain Controller Node or Compute Region Node displays the "not ready" (grey) state.

Symptoms	Cloud Installation Dashboard is unable to monitor the Domain Controller Node or the Compute Region Node.
Primary software component	Cloud Administration Node
Failure message	
Probable cause	The Domain Controller Node or the Compute Region Node has not updated the live status to the Cloud Administration Node.

### Solution:

- Log in to the suspect node through SSH and run `sudo chef-client`. This will force the node to update its status with Chef.

## Problem: PXE - When creating a new PXE node, the PXE boot fails with a TFTP timeout error.

Symptoms	Cloud Administration Node takes a long time to reboot.
Primary software component	Cloud Administration Node.
Failure message	
Probable cause	This occurs in a few cases after the Cloud Administration Node is rebooted.

### Solution:

- Log in to the Cloud Administration Node and run the following commands:

```
sudo bluepill tftpd stop
```

```
sudo bluepill tftpd start
```

- After these commands have run, PXE boot the nodes.

## Appendix A: Domain Controller Barclamps

This appendix explains how barclamps are structured and describes the relationship between a proposal and a barclamp.

In the Cloud Installation Dashboard:

- You can edit barclamp attributes and node deployment information.
- You can set attribute values that are based on roles associated with the barclamp.
- Each barclamp might have specific logic that requires minimums or maximums for node assignments.

The following table contains descriptions and node deployment information for the following domain controller barclamps and proposals:

- ["Hp Eden 100" below](#)
- ["Hp Eve 100" on next page](#)
- ["Hp Focus 100" on page 69](#)
- ["Hp Glance 201211" on page 69](#)
- ["Hp Keystone 201211" on page 72](#)
- ["Hp Peer 100" on page 72](#)
- ["Hp Skyline 120" on page 73](#)
- ["Mongodb 284" on page 73](#)
- ["Postgresql 915" on page 73](#)
- ["Rabbitmsgq 271" on page 74](#)

Barclamp	Hp Eden 100	
Description	Configures value added eden service	
Node Deployment	Domain Controller Node	
Attributes	Values	Description

Proxy Host (Optional)	<IPAddress>	Web proxy server's IP address to access internet
Proxy Port (Optional)	<PortNo>	Port used by web proxy server
Non Proxy Host(s) (Optional)	192.168.*  10.1.* localhost	Do not use proxy servers for addresses beginning with
Proxy User (Optional)	<username>	Username required for authenticating to proxy server
Proxy Password (Optional)	<password>	Password of the user required for authenticating to proxy server
Keystone Instance	[generated]	The Keystone proposal to use
PostgreSQL Instance	[generated]	The PostgreSQL proposal to use
RabbitMQ Instance	[generated]	The RabbitMQ proposal to use
Service User	eden	The user that Eden uses when authenticating with Keystone
Service Password	[generated]	The password for the Eden Keystone authentication user

<b>Barclamp</b>	Hp Eve 100	
<b>Description</b>	Configures value added service peer	
<b>Node Deployment</b>	Domain Controller Node	
<b>Attributes</b>	<b>Values</b>	<b>Description</b>
Proxy Host (Optional)	<IPAddress>	Web proxy server's IP address to access internet
Proxy Port (Optional)	<PortNo>	Port used by web proxy server
Non Proxy Host(s) (Optional)	192.168.*  10.1.* localhost	Do not use proxy servers for addresses beginning with
Proxy User (Optional)		Username required for authenticating to proxy server
Proxy Password (Optional)		Password of the user required for authenticating to proxy server
PostgreSQL Instance		The PostgreSQL proposal to use
Mongodb Instance		The MongoDB proposal to use

Rabbitmq Instance		The RabbitMQ proposal to use
Peer Instance		The Peer proposal to use
Eden Instance		The Eden proposal to use
Keystone Instance		The Keystone proposal to use

<b>Barclamp</b>	Hp Focus 100	
<b>Description</b>	Configures value added service focus	
<b>Node Deployment</b>		
<b>Attributes</b>	<b>Values</b>	<b>Description</b>
Eden Instance		The Eden proposal to use
Keystone Instance		The Keystone proposal to use

<b>Barclamp</b>	Hp Glance 201211	
<b>Description</b>	Hp glance 201211 service (image registry and delivery service) for the cloud	
<b>Node Deployment</b>	Domain Controller Node	
<b>Attributes</b>	<b>Values</b>	<b>Description</b>
Working Directory	/var/lib/glance	Glance working directory
PID Directory	/var/run/glance	Location of Glance's PID files
Notifier Strategy	Noop	The only option is "No Operation"
Image Store Directory	/var/lib/glance/images	Location of images
Location of images Scrubber:		
Log File	/var/log/glance/scrubber.log	The location where the scrubber will log
Config File	/etc/glance/glance-scrubber.conf	The configuration file for the scrubber

Debug	FALSE	Indicates if the scrubber will run in debug mode Verbose true Indicates if the scrubber will run in verbose mode
<b>Reaper:</b>		
Log File	/var/log/glance/reaper.log	The location where the reaper will log
Config File	/etc/glance/glance-reaper.conf	The configuration file for the reaper
Debug	FALSE	Indicates if the reaper will run in debug mode
Verbose	TRUE	Indicates if the reaper will run in verbose mode
<b>Pruner:</b>		
Log File	/var/log/glance/pruner.log	The location where the pruner will log
Config File	/etc/glance/glance-pruner.conf	The configuration file for the pruner
Debug	FALSE	
<b>Prefetcher:</b>		
Log File	/var/log/glance/prefetcher.log	The location where the prefetcher will log
Config File	/etc/glance/glance-prefetcher.conf	The configuration file for the prefetcher Debug false Indicates if the prefetcher will run in debug mode
Verbose		
<b>Indicates whether the prefetcher will run in verbose mode API:</b>		
Log File	/var/log/glance/api.log	The location where the API will log
Config File	/etc/glance/glance-api.conf	The configuration file for the API
Paste INI File	/etc/glance/glance-api-paste.ini	Paste Deploy configuration file for the API
Debug	FALSE	Indicates if the API will run in debug mode
Verbose	TRUE	Indicates if the API will run in verbose mode

Bind to All Addresses	TRUE	Controls if the API will bind to all addresses or the public address only
Access Port	9292	The port the API service will run on
<b>Registry:</b>		
Log File	/var/log/glance/registry.log	The location where the registry will log
Config File	/etc/glance/glance-registry.conf	The configuration file for the registry
Paste INI File	/etc/glance/glance-registry-paste.ini	Paste Deploy configuration file for the registry
Debug	FALSE	Indicates if the registry will run in debug mode
Verbose	TRUE	Indicates if the registry will run in verbose mode
Bind to All Addresses	TRUE	Controls if the registry will bind to all addresses or the public address only
Access Port	9191	The port the registry service will run on
<b>Caching:</b>		
Enable Caching	FALSE	Indicates if caching should be on
Turn On Cache Management	FALSE	Enables the use of glance-cache-manage CLI & the corresponding API
Directory	/var/lib/glance/image-cache	The location where images are cached
Grace Period	3600	The timeout for accessing the image
Stall Timeout	86400	The timeout to wait for a stalled GET request
<b>Database:</b>		
SQL Idle Timeout	3600	PostgreSQL idle time check
PostgreSQL Instance	[generated]	The PostgreSQL proposal to use
Use Keystone	TRUE	Indicates to Crowbar if Keystone is to be used for authentication
Keystone Instance	[generated]	The Crowbar Keystone proposal to use

Service User	glance	The user that Glance uses when authenticating with Keystone
Service Password	[generated]	The password for the Glance Keystone authentication user
Use Syslog	FALSE	Indicates to Glance to not log to syslog

<b>Barclamp</b>	Hp Keystone 201211	
<b>Description</b>	Centralized authentication and authorization service for openstack	
<b>Node Deployment</b>		
<b>Attributes</b>	<b>Values</b>	<b>Description</b>
PostgreSQL Instance	<Generated>	The Postgresql proposal to use
Domain Name	Default: Domain	Name of the domain hosted by domain controller
Domain Admin Password	Default: secretword	Admin user's password; Granted domain admin role on AdminProject
Domain Arch Password	Default: secretword	Arch user's password; Granted domain architect role on AdminProject project
Domain Trash Password	Default: secretword	Trash user's password; Granted user role on trash project

<b>Barclamp</b>	Hp Peer 100	
<b>Description</b>	Configures value added service peer	
<b>Node Deployment</b>	Domain Controller Node	
<b>Attributes</b>	<b>Values</b>	<b>Description</b>
PostgreSQL Instance	[generated]	The PostgreSQL proposal to use
Eden Instance	[generated]	The Eden proposal to use
Keystone Instance	[generated]	The Keystone proposal to use

<b>Barclamp</b>	Hp Skyline 120	
<b>Description</b>	Installs hp skyline for release 1.2.	
<b>Node Deployment</b>		
<b>Attributes</b>	<b>Values</b>	<b>Description</b>
Eden Instance		The Eden proposal to use
Keystone Instance		The Keystone proposal to use
Proxy Host (Optional)	<IPAddress>	Web proxy server's IP address to access internet
Proxy Port (Optional)	<PortNo>	Port used by web proxy server. If your environment does not require specifying Proxy Host, do not clear the default value, such as 8080.
Non Proxy Host(s) (Optional)	192.168.*  10.1.* localhost	Do not use proxy servers for addresses beginning with

<b>Barclamp</b>	Mongodb 284	
<b>Description</b>	Configure mongodb	
<b>Node Deployment</b>	Domain Controller Node	
<b>Attributes</b>	<b>Values</b>	<b>Description</b>
LogPath	/var/log/mongodb	The location where mongodb will log

<b>Barclamp</b>	Postgresql 915	
<b>Description</b>	Configures a postgresql server	
<b>Node Deployment</b>	Domain Controller Node	
<b>Attributes</b>	<b>Values</b>	<b>Description</b>
Datadir	/var/lib/postgresql	Location where database files will be stored.

<b>Barclamp</b>	Rabbitmq 271	
<b>Description</b>	Configures a rabbitmq server	
<b>Node Deployment</b>	Domain Controller Node	
<b>Attributes</b>	<b>Values</b>	<b>Description</b>
LogPath	/var/log/rabbitmq	The location where RabbitMQ will log

## Appendix B: Load Balancer Image

To help you create and implement a load balancer image in your Cloud Infrastructure, review the following sections:

- ["Prerequisites " below](#)
- ["Ways To Use a Load Balancer Image " below](#)
- ["Create a Standalone Load Balancer Instance " on next page](#)
- ["Convert a Load Balancer Instance to an Image Snapshot" on page 79](#)

### Prerequisites

See ["Configure an Image Repository " on page 53](#)

### Ways To Use a Load Balancer Image

You can use a load balancer in the following ways:

- ["Static Load Balancer " below](#)
- ["Dynamically-Provisioned Load Balancer " on next page](#)

**Tip:** To determine the which load balance approach you need to use, analyze the application that is responsible for provisioning.

### Static Load Balancer

When a load balancer needs to be shared in your Cloud Infrastructure, a static load balancer (referenced instance) is created. This one instance is used to provide load balancing features for all provisioned topologies.

In an Infrastructure Design document, the static load balancer image is referenced. The static load balancer service is referenced using "service\_reference", whose value should be the same as the "service\_id" in the realized topology template of the static Load balancer service.

## Dynamically-Provisioned Load Balancer

A load balancer instance can be provisioned dynamically when a topology is provisioned. This approach results in a unique load balancer instance for *each* topology. The provisioning application must be able to discover the load balancer's IP address.

In a topology binding document, a resource pool's load balancer service must be specified.

**Note:** The resource pool must be created *before* the topology binding document is created, using the applicable load Balancer image in the specified Cloud Infrastructure.

In the Cloud Administration Dashboard, you can specify a load balancer image as a service to a resource pool. When you use the Create Infrastructure Design wizard to define the topology binding document, the load balancer service is listed in the drop-down list.

## Create a Standalone Load Balancer Instance

1. In the Cloud Administration Dashboard, upload an Ubuntu operating system image, if it does not already exist. **Best Practice:** HP recommends the Ubuntu 64 bits 12.04 LTS server image.
  - a. HP recommends that you use this Ubuntu Image: <http://uec-images.ubuntu.com/precise/current/precise-server-cloudimg-amd64-disk1.img>.
  - b. For more information, see <http://docs.openstack.org/trunk/openstack-compute/admin/content/starting-images.html>.
2. In the Cloud Administration Dashboard, create resources required for this activity:
  - a. Select **Project > Access & Security**.
    - i. In the Access & Security view, create a new key pair or use an existing one. Click **Create Keypair**.
    - ii. Save the associated private key (.pem) file to access the provisioned load balancer instance with SSH.
  - b. Select **Project > Access & Security**.
    - i. Create a new security group or use an existing one.
    - ii. Click **Create Security Group**.
    - iii. For the new security group, add the following rules to allow access to the provisioned load balancer instance:

IP Protocol	From Port	To Port	Source	Access
TCP	22	22	0.0.0.0/0 (CIDR)	SSH
TCP	21212	21212	0.0.0.0/0 (CIDR)	Load Balancer REST service

3. In the Cloud Administration Dashboard, provision a new virtual machine, using the uploaded Ubuntu image.
  - a. Select **Project > Images**.
  - b. Click **Launch** for the uploaded Ubuntu image.
    - i. In the Details tab, specify the server (load balancer instance) name, and a flavor size, such as m1.tiny.
    - ii. In the Access & Security view:
      - i. Select a key pair.
      - ii. Select a security.
  - c. In the Instances view, find the provisioned load balancer instance and make sure it is in a "Running" power state.
4. In the Cloud Administration Dashboard, associate a floating (public) IP address to the provisioned load balancer instance.
  - a. Select **Project > Instances**.
  - b. In the Instances view, select the Associate Floating IP action for the new Ubuntu instance. Both the private and floating (public) IPs will display for the provisioned load balancer instance.
5. SSH to the provisioned load balancer instance, using its floating (public) IP address, using the private key (.pem) file.
6. If you are using PuTTY to access the provisioned load balancer instance, you must modify the private key (.pem) file.
  - a. The saved private key (.pem) file must be converted from PEM format to PPK format. For conversion instructions, see instructions. In these instructions, use PuTTY to access the launched instance's Fixed IP and Floating IP addresses, using the private key for authentication.

**Note:** The login account for the provisioned load balancer instance is ubuntu.

7. Two debian packages, **haproxy-vm** and **loadbalancer-api** must be downloaded to the provisioned load balancer instance.
  - a. To work around install issues, complete this step *before* you configure a proxy server. This requires the public IP address for the Cloud Administration Node. Use ifconfig to find the IP address. This requires the full filename to be downloaded. In the Cloud Administration Node, the filenames are located in tftpboot/ubuntu\_dvd/extra/files/loadbalancer.
    - `sudo wget http://<Admin Node's public IP address>:8091/ubuntu_dvd/extra/files/loadbalancer/haproxy-vm_1.20.0+SNAPSHOT-all.deb`
    - `sudo wget http://<Admin Node's public IP address>:8091/ubuntu_dvd/extra/files/loadbalancer/loadbalancer-api_1.20.0+SNAPSHOT-all.deb`
8. If a proxy server is required to access the Internet, make sure the proxy settings are configured on the provisioned load balancer instance.
  - a. Modify `/etc/bash.bashrc`, adding the following two lines. HP recommends that you use IP addresses because DNS is not always available. The following example uses **web-proxy.corp.hp.com**.
    - i. `*export http_proxy=http://16.85.175.150:8080/*`
    - ii. `*export https_proxy=http://16.85.175.150:8080/*`
  - b. Modify the **sudoers** file to make sure the proxy information is available to all programs:  
`sudo visudo`
    - i. Add this entry: "Defaults env\_keep += "http\_proxy,https\_proxy"
    - ii. Make sure there is a character space before the "+=" and another character space after it.
    - iii. Save your changes.
  - c. Log out of the provisioned load balancer instance and then re-connect via SSH to set the proxy settings in the environment.
9. On the provisioned load balancer instance, run the following commands, in this order, to install the prerequisite software:
  - a. `sudo apt-get update`
  - b. `sudo apt-get install openjdk-7-jre-headless`
  - c. `sudo apt-get install haproxy`
10. Install and start the two debian packages, **haproxy-vm** and **loadbalancer-api**, on the provisioned load balancer instance.

```
sudo dpkg -i haproxy-vm_1.20.0+SNAPSHOT-all.deb
```

```
sudo dpkg -i loadbalancer-api_1.20.0+SNAPSHOT-all.deb
```

```
sudo start loadbalancer-api
```

11. In a browser, go to *http://<Load Balancer instance's public IP address>:21212/loadbalancer* to verify that the load balancer REST service is running.

## Convert a Load Balancer Instance to an Image Snapshot

To create an image snapshot of the provisioned load balancer instance:

1. In the Cloud Administration Dashboard, select **Project > Instances**.
2. In the Instances view, find the provisioned load balancer by its **Instance Name**.
  - a. Click **Create Snapshot**.
  - b. Enter a snapshot name.
3. In the Cloud Administration Dashboard, select **Project > Snapshots**.
  - a. In the Snapshots view, find the snapshot you just created.
  - b. Select the new snapshot to view its details.
  - c. View the snapshot ID. This is the load balancer image ID that must be referenced in a resource pool as a load balancer service definition.

## Appendix C: Chef Server Image

To help you create and implement a chef server image in your Cloud Infrastructure, review the following sections:

- ["Prerequisites " below](#)
- ["Creating a Chef Server Image " on the facing page](#)
- ["Adding a Chef Client " on page 82](#)

### Prerequisites

In the Cloud Administration Dashboard, upload an Ubuntu operating system image, if it does not already exist. See ["Configure an Image Repository " on page 53](#).

**Best Practice:** HP recommends the Ubuntu 64 bits 12.04 LTS server image.

- HP recommends that you use this Ubuntu Image: [http://uecimages.](http://uecimages.ubuntu.com/precise/current/precise-server-cloudimg-amd64-disk1.img)

[ubuntu.com/precise/current/precise-server-cloudimg-amd64-disk1.img](http://uecimages.ubuntu.com/precise/current/precise-server-cloudimg-amd64-disk1.img).

- For more information, see <http://docs.openstack.org/trunk/openstackcompute/admin/content/starting-images.html>.

### Ways To Use a Chef Server Image

You can use a Chef server in the following ways:

- ["Static Chef Server" on the facing page](#)
- ["Dynamically-Provisioned Chef Server" on the facing page](#)

**Tip:** To determine which Chef Server approach you need to use, analyze the application that is responsible for provisioning.

## Static Chef Server

When a Chef Server needs to be shared in your Cloud Infrastructure, a static Chef Server (referenced instance) is created. This one instance is used to provide Chef Server features for all provisioned topologies.

In an Infrastructure Design document, the static Chef Server image is referenced. The static Chef Server service is referenced using "service\_reference", whose value should be the same as the "service\_id" in the realized topology template of the static Chef Server service.

## Dynamically-Provisioned Chef Server

A Chef Server instance can be provisioned dynamically when a topology is provisioned. This approach results in a unique Chef Server instance for each topology. The provisioning application must be able to discover the Chef Server's IP address. In a topology binding document, a resource pool's Chef Server service must be specified.

**Caution:** The resource pool must be created *before* the topology binding document is created, using the applicable Chef Server service in the specified Cloud Infrastructure.

In the Cloud Administration Dashboard, you can specify a Chef Server image as a service to a resource pool. When you use the Create Infrastructure Design wizard to define the topology binding document, the Chef Server service is listed in the drop-down list.

## Creating a Chef Server Image

In the Cloud Installation Dashboard, you must create a new virtual machine from an Ubuntu virtual machine image that:

- Uses a security group that allows port 22 for ssh. Optionally, add port 4040 to the security group so you can access the chef web-ui service.
- Uses a key name that you have the private key (.pem) for. **Note:** You can create a new keypair in the Cloud Administration Dashboard.
- Associates a floating (public) IP address so you can access it.

To create a new virtual machine from an Ubuntu virtual machine image:

1. **Update /etc/hosts.allow** to include any hosts that you want to access, using ssh. If you are using Chef deployment from within CDA, this must include the CDA system.

2. ssh to the new virtual machine, using its public IP address and the private ssh key.
  - a. If a proxy server is used to access the internet, modify the proxy settings accordingly for the aptitude program. There are two general approaches:
    - i. Modify the .conf file for aptitude.
    - ii. Set the environment variable http\_proxy in the user profile or in the /etc/environment file. Add http\_proxy to the env\_keep list in the sudoer file, such as Defaults env\_keep="http\_proxy."
3. Install the chef-server package. For more information, see [http://wiki-opscode.com/display/chef/Installing+Chef+Server+on+Debian+or+Ubuntu+using+Packages](http://wiki.opscode.com/display/chef/Installing+Chef+Server+on+Debian+or+Ubuntu+using+Packages).
4. Configure a chef repository. For more information, see [http://docs.opscode.com/essentials\\_repository\\_create.html](http://docs.opscode.com/essentials_repository_create.html).
5. In the Cloud Installation Dashboard, create an image snapshot of the virtual machine.

## Adding a Chef Client

For instructions on how to add a chef client on a new system, see <http://wiki.opscode.com/display/chef/Installing+Chef+Client+on+Ubuntu+or+Debian>.

(Optional) You can use an image that already has the client installed.

