HP Continuous Delivery Automation

Software Version: 1.10

Cloud Installation Dashboard Guide

Document Release Date: January 2013

Software Release Date: December 2012

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 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.

Oracle and Java are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

RED HAT READY™ Logo and RED HAT CERTIFIED PARTNER™ Logo are trademarks of Red Hat, Inc.

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.

Support

Visit the HP Software Support Online web site at: 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 user 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

Cloud Installation Dashboard Guide

Prerequisites

Basic Configuration

The Cloud environment consists of Admin, Controller, and Compute nodes. An optional Windows client can be added to the configuration.

Node	Role
Admin Node	Ubuntu12.04 OS where Cloud Infrastructure media is booted. Bootstraps the install to set up the Controller and Compute nodes.
Controller Node	Hosts the Cloud Infrastructure services.
Compute Node	Hosts the provisioned virtual machines. Multiple compute nodes can be created to provide more cloud capacity.
CDA Node	(Optional) Continuous Delivery Automation - complete dev-ops solution that integrates with Cloud Infrastructure services.
CSA Node	(Optional) Cloud Service Automation - a portal to subscribe application/platform/infrastructure offered by CDA.
Windows client	This machine has a connection to the private network and will provide browser and SSH access to the Admin and other nodes.

In small cloud environments, the Controller and Compute features are co-located on the same node. In larger cloud environments, the Compute node may be separate from the Controller node, with additional Compute nodes available to provide larger provisioned cloud capacity.

CDA and CSA are installed in their own respective nodes.

Cloud Infrastructure Sizing

Refer to the HP Cloud Connector Physical Infrastructure Reference Architectures (PIRA) document.

Cloud Infrastructure Suggested Deployment Diagrams

Refer to the HP Cloud Connector PIRA document.

CDA & CSA Node Sizing

Refer to the CDA and CSA Installation Guides.

Set Up the Cloud Infrastructure Environment

To set up the Cloud Infrastructure environment, go to the Prepare Cloud Environment section.

Prepare Cloud Environment

Overview

Preparing the Cloud Environment requires IT administrators to set up the appropriate network, storage, and hardware infrastructure as described in the HP Cloud Connector PIRA document. Suggested deployments:

Deployment Diagram (Production) - I

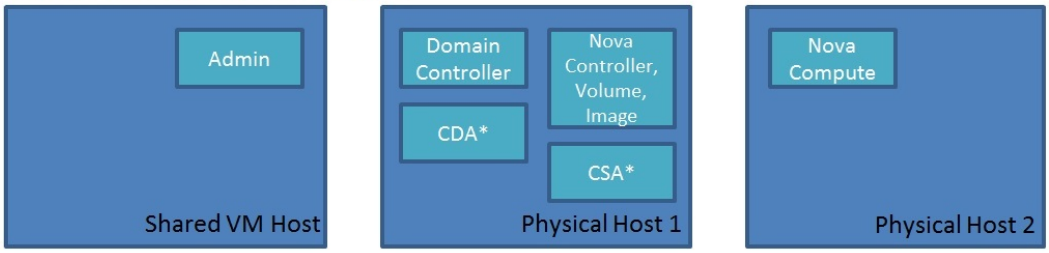

Deployment Diagram (Production) - II

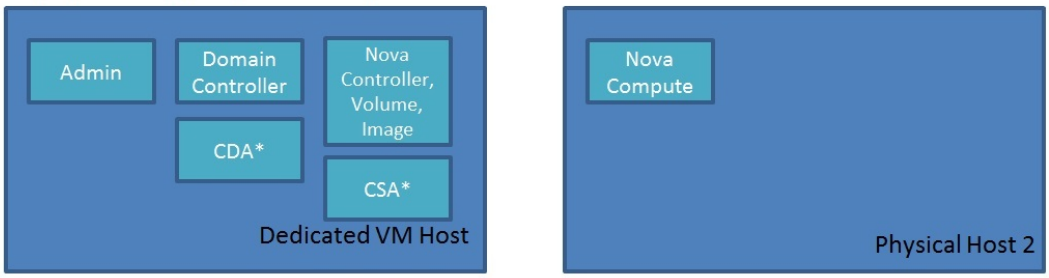

Prepare the Servers (Deployment Diagram I)

The steps below outline Deployment Diagram I where the Admin node is virtual and the other nodes are on physical machines.

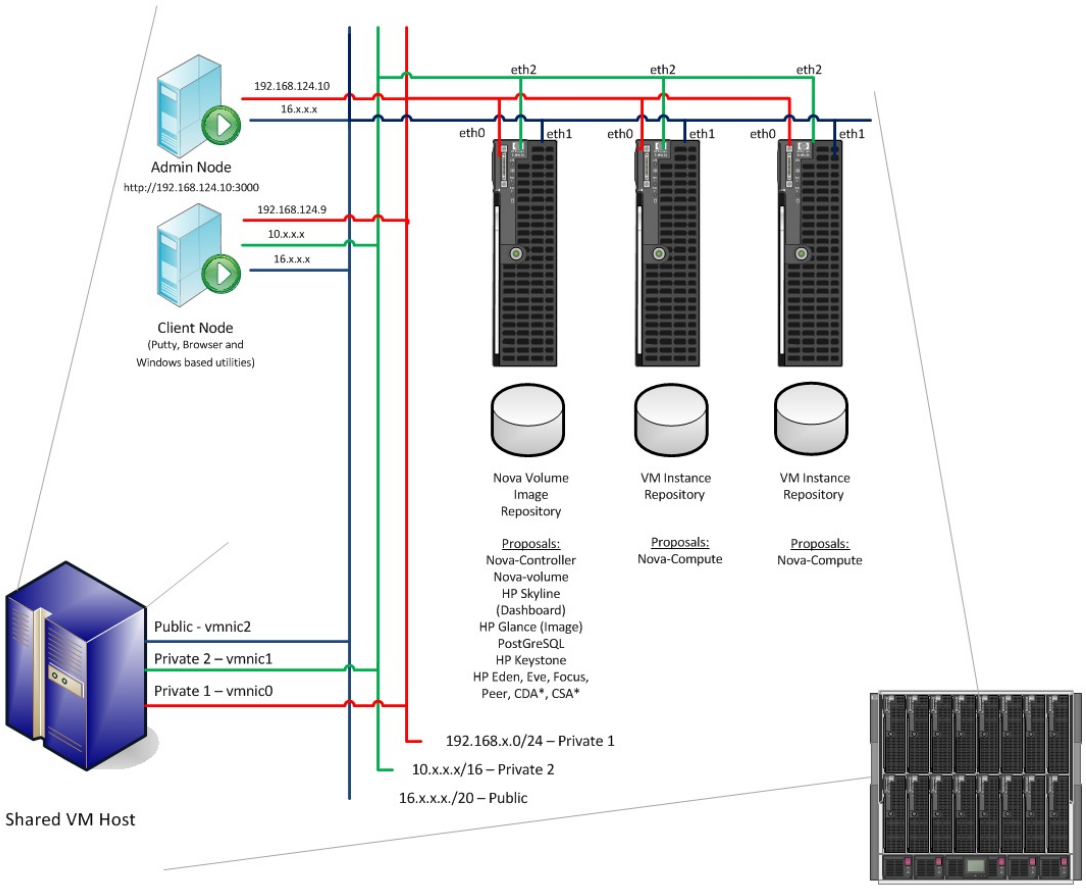

Controller and Compute Node

The Controller and Compute node must be powered OFF before installing Cloud Infrastructure. If you are using HP Blade servers, these are the typical steps:

1. Using the Integrated Lights-Out 3 (iLO) web interface, access the Controller and Compute node.
2. Log into the iLO session.
3. In the right pane of the iLO screen, the server power displays **ON** or **OFF**.
4. In the left pane, select **Power Management > Server Power**.
5. In the right pane, select the **Press and Hold** button to do a forced power OFF. Confirm the action.
6. The right pane displays that system power is **OFF**.
7. Log out of the iLO session.

Admin Node

The Admin node must be booted using either a CD/DVD or an ISO image. The release product ISO image file name is **Cloud_Installer_1_0.iso**.

Booting the Admin Node from the ISO Image

The following steps outline how to boot the Admin node from the ISO image setting up the Admin Node as a guest VM on vCenter (using ESX 4.1.0 or above).

Note: the ISO image should be located in the vCenter's datastore.

1. Ensure the Admin node is powered OFF.
2. Edit the VM settings.
3. Set the CD/DVD to connect at power ON and point it to the ISO file in the “Datastore ISO file” selection.
4. Select **Options > Advanced > Boot Options**. In the Force BIOS Setup, Check the box for **The next time the virtual machine boots, force entry into the BIOS screen**.
5. Click **OK.**
6. Power ON the VM.
7. View the Admin node's console. A series of dialog boxes display. Use the <Tab>, <Enter>, and <Arrow> keys to navigate.
8. The VM console should display the BIOS Setup Utility screen.
9. Use the arrow keys to select the **Boot** section.
10. Using the keyboard, move the CD/DVD drive to the first item in the list so it will boot from the ISO.
11. Press F10 or use the arrow keys to select **Exit**. Press <Enter>.
12. Select **Yes.**
13. The VM will begin to boot.
14. A "Configure the Network" dialog box displays. Select **eth0**.
15. Enter the following information in the "Configure the Network" dialog box:
    a. Admin node's IP address. Select **<Continue>**.
        i. Example: 192.168.124.10
        ii. **NOTE**: If you provide a different IP Address and related settings, ensure that the Cosmos Networks UI is also set up this way.
    b. Netmask. Select **<Continue>**.
        i. Example: 155.255.255.0
    c. Gateway. Select **<Continue>**.
        i. Example: 192.168.124.1
    d. Name server. Select **<Continue>**.
        i. Example: 192.168.124.10
16. Wait while the system detects the link on eth0.
17. Specify the appropriate time zone in the "Configure the Clock" dialog box.
18. Wait for the ISO image to complete installation.
19. The Admin node will display a login screen when finished.
20. Power the VM down.
21. Edit the VM settings so that it won't boot from the ISO again (refer to the picture).
22. Power the VM up.

The Cosmos installer has now been successfully installed on the Admin node.

From the Windows VM client, use a browser to access the URL, http://192.168.124.10:9000 to access the Admin node. **NOTE**: The browser must not have a proxy set.

Next Step

Refer to the Set Up Cloud Administration Node section to begin installation and configuration of the Cloud Infrastructure.

Set Up Cloud Administration Node

Overview

Using a browser that has access to the Cloud Administration node and cloud network, invoke the URL http://192.168.124.10:9000 to access the Cloud Installation Dashboard.

NOTE: If you are installing on a private network, a proxy must be set in the browser.

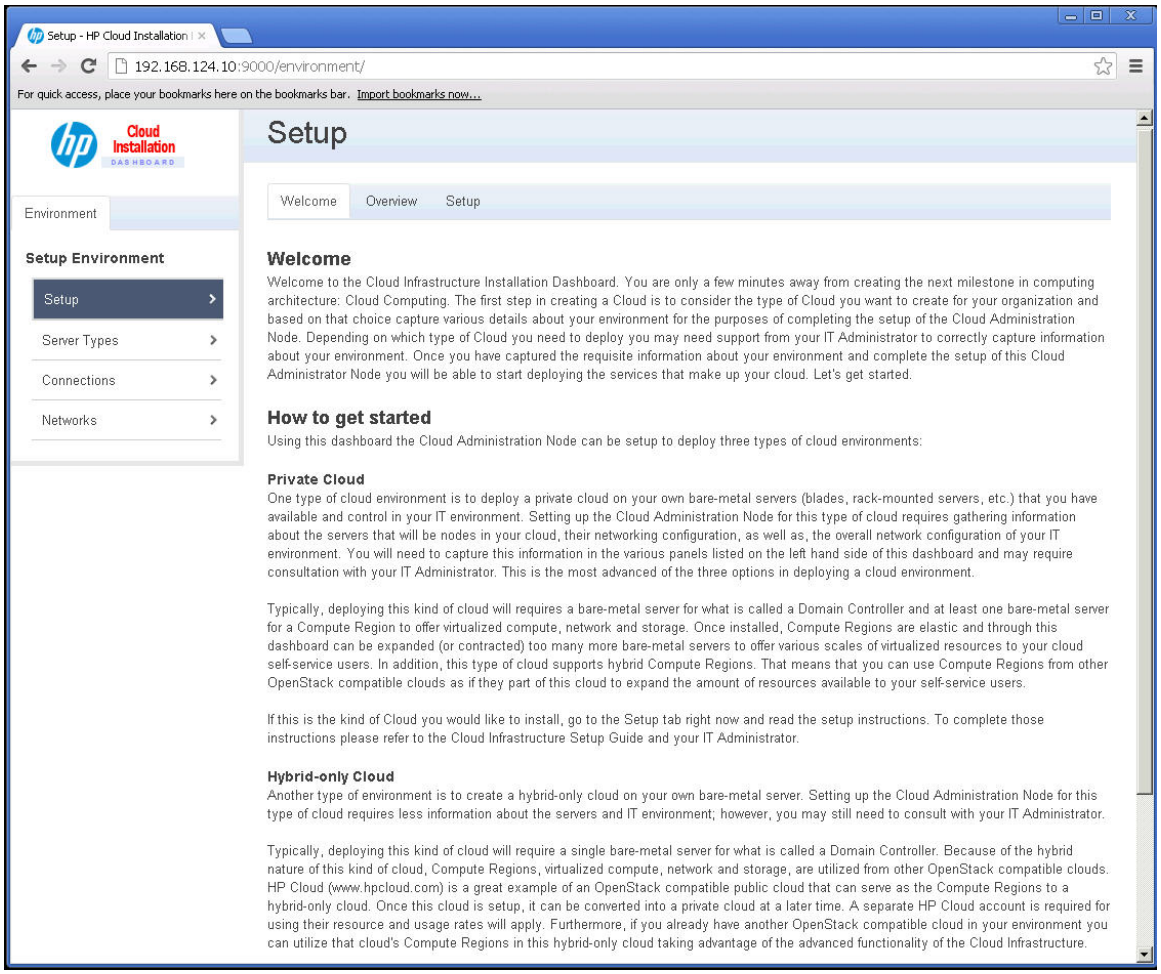

Set Up the Cloud Administration Node

Review Configuration Prior to Setup

From the HP Cloud Installation Dashboard, review the **Welcome** and **Overview** tabs to understand the concepts of Cloud environments and setup configurations. Customize Server Types, Connections, and Networks based on your network and hardware infrastructure.

Server Types: This panel lets you capture (or use existing) server definitions about the networking ports for the various servers in your bare-metal cloud environment. If all of the servers in your cloud are the same machine type (homogeneous nodes), you will not need to configure anything in this section. HP recommends keeping the defaults.

Connections: To define a connection set, review one of the network modes (single, dual, team, etc.) along with a list of logical connection interfaces (intf0, intf1, etc.) with each one being bound to a port, defined by bandwidth and port number. You may want to use the **Edit Connection** option if your Connection Type is Dual and the interfaces are not contiguous, For example, eth0 - private 1 and eth2 is public. The default settings assume that you have **eth0** of the participating hosts/nodes connected to Private Network1 and **eth1** connected to a public network.

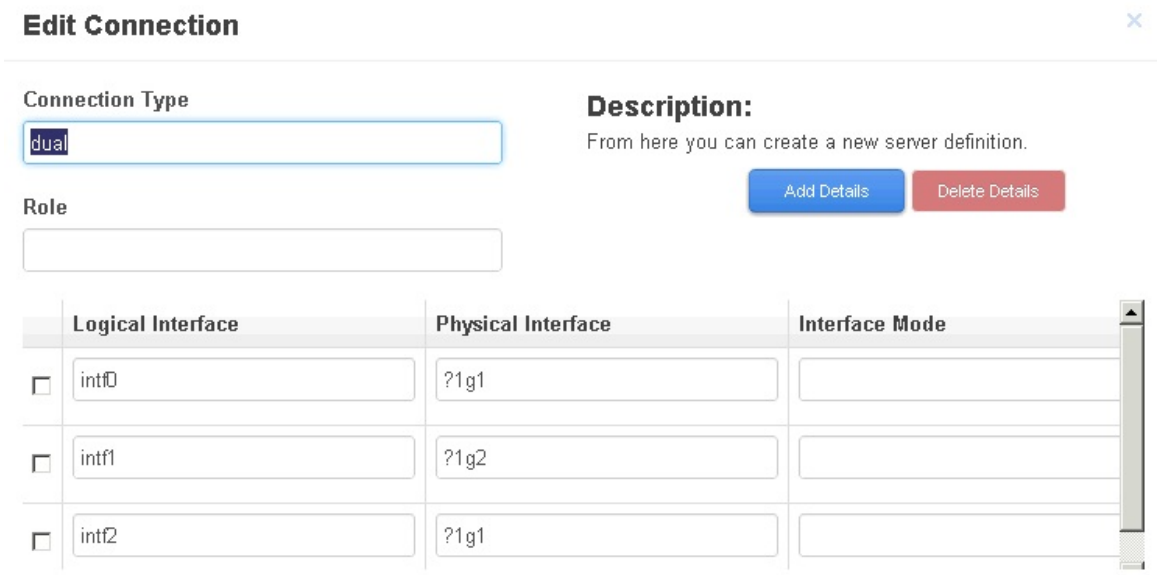

NOTE: If you are not sure, keep the default settings to avoid installation failure.

Networks: For each logical connection (i.e., intf0, intf1, etc.) a network is defined (i.e., admin network, public network, BMC network, etc.). The different networks are used for varying purposes in the cloud, such as networking between the nodes and this Cloud Administration Node (i.e. admin network), and the public or corporate network (i.e., public network) attached to your cloud. The default settings configure VLAN (tagged as 300) based on a public network on **eth1** if installing in dual mode, and on **eth0** if installing in single mode. In such cases, you will need an external router for traffic to internet/public.

If you plan to provide external access directly to the participating hosts/nodes, ensure you have two sets of contiguous IP ranges:

- Range 1: 16.x.x.2 to 16.x.x.10 (This range will get assigned to the bootstrapped participating nodes.)
- Range 2: 16.x.x.121 to 16.x.x.254 (This range is for Floating IP Configuration in OpenStack, used later in the Installation Guide.)

The customization is done through a two-step process:

1. Edit Network: Specify Subnet, Netmask, Broadcast, VLAN Enabled and Router details specific to your public/corporate network.

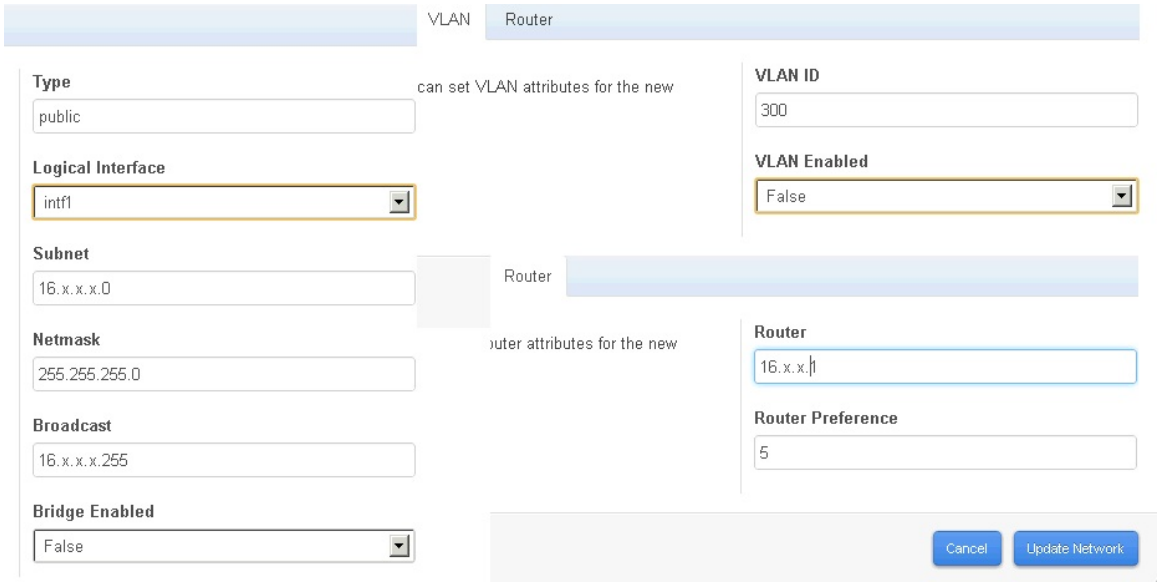

2. Edit Ranges: Specify a range for Node Type host for the participating nodes.

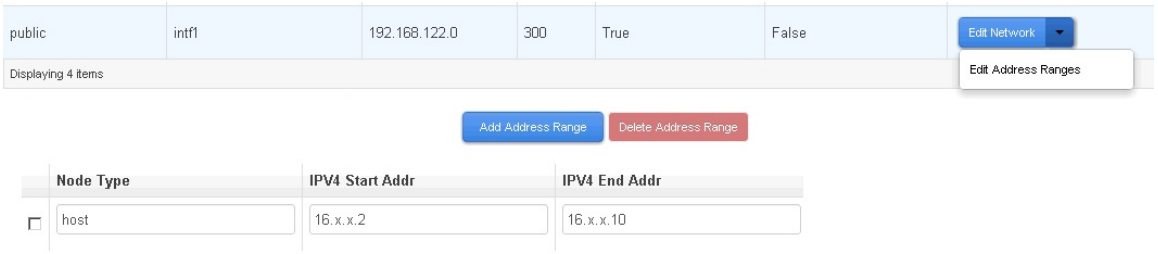

NOTE: If you are not sure, keep the default settings to avoid installation failure.

During the boot process, if you have provided ip address other than 192.168.124.10, you must update the **Admin** network.

Begin the Setup Process

To start the setup process, use the following steps:

1. Select the **Setup** tab.

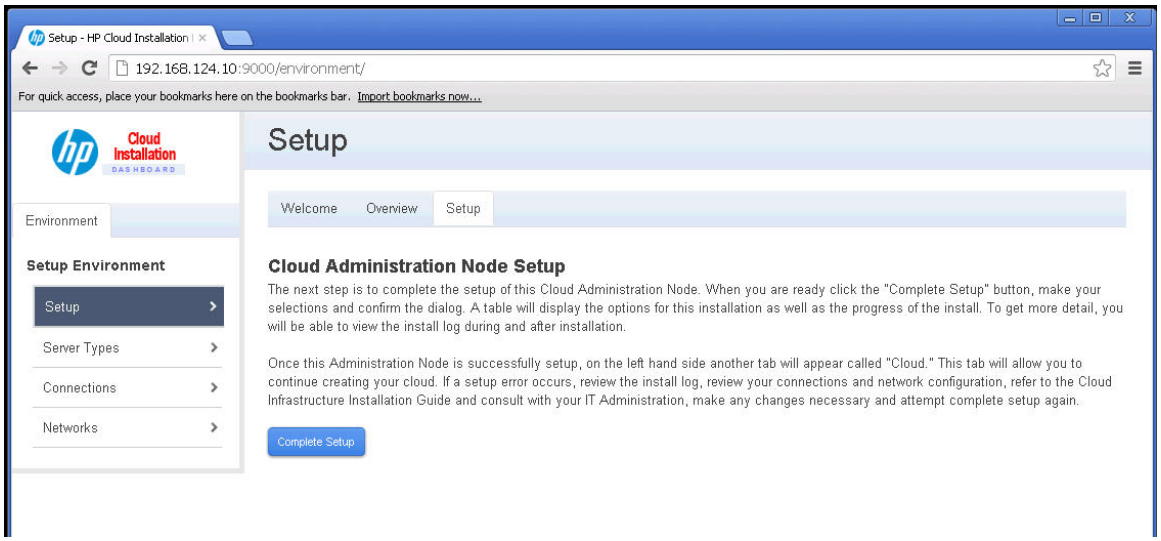

2. Click **Complete Setup**.
3. Specify the necessary configuration:
    a. Set **Network Mode** to dual.
    b. Set **IPMI/BMC Network Status** to Disabled.

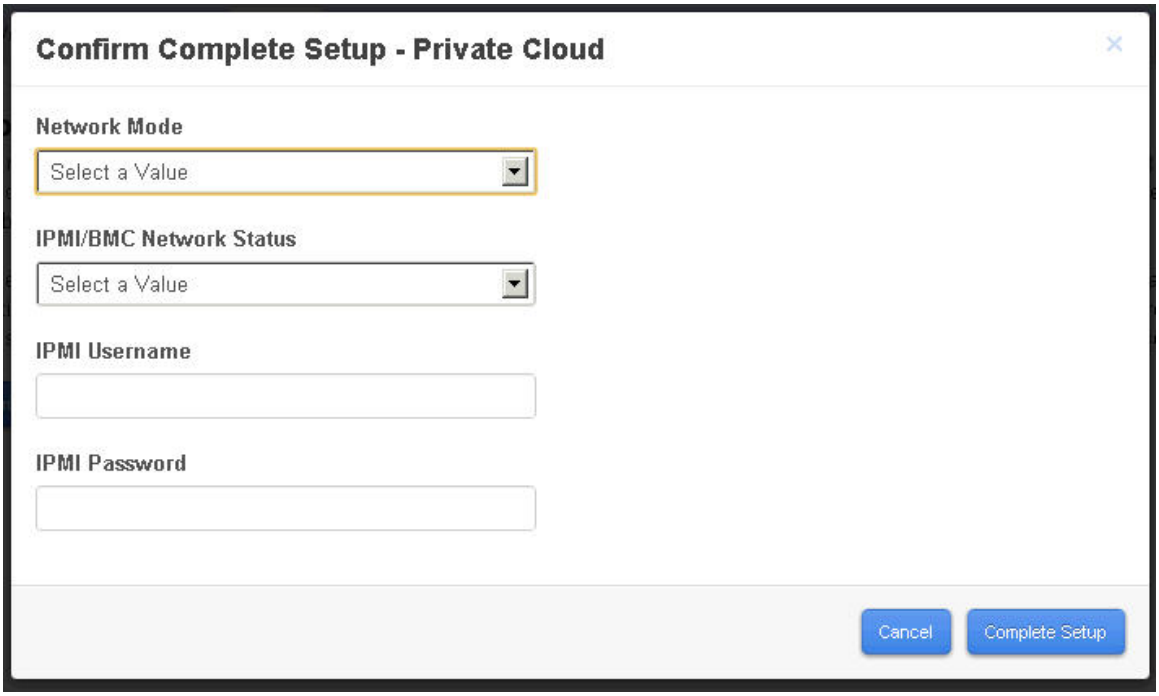

NOTE: Enable IPMI/BMC feature is for test purposes only.

4. Click **Complete Setup**. This action installs crowbar on the Cloud Administration node. This may take several minutes.
5. A working status message displays in the **Setup Complete** column. **NOTE**: Any failure at this stage requires rebuilding the Admin node through the bootable ISO.

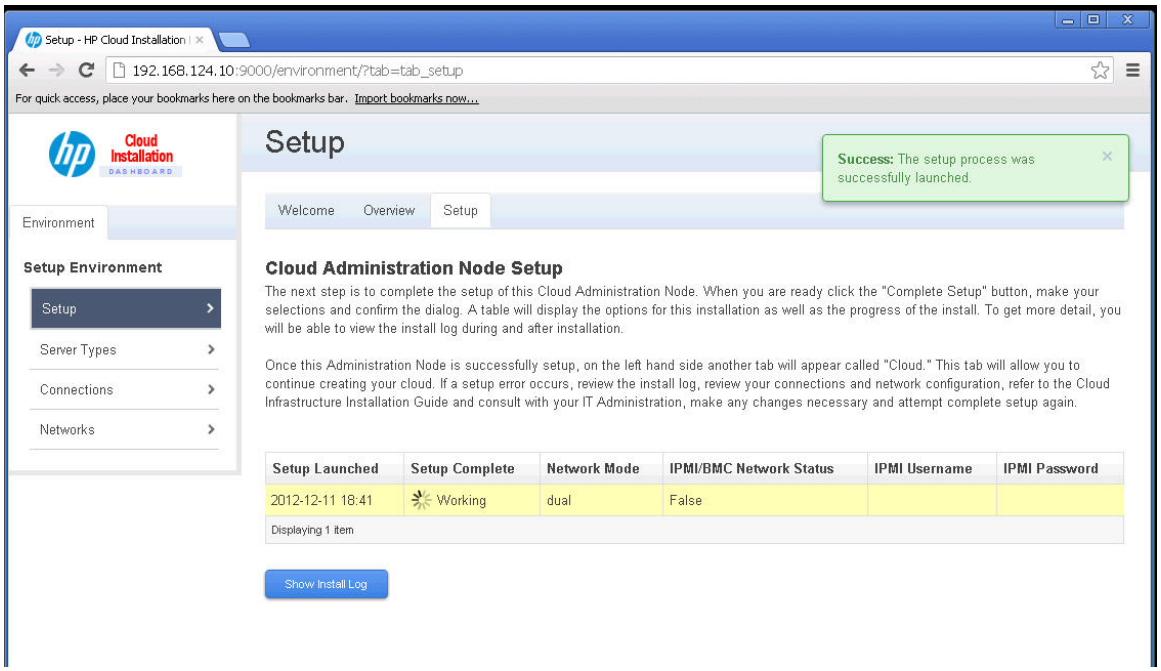

6. Click **Show Install Log** to view the install activity. Click **Refresh** to monitor the install progress.

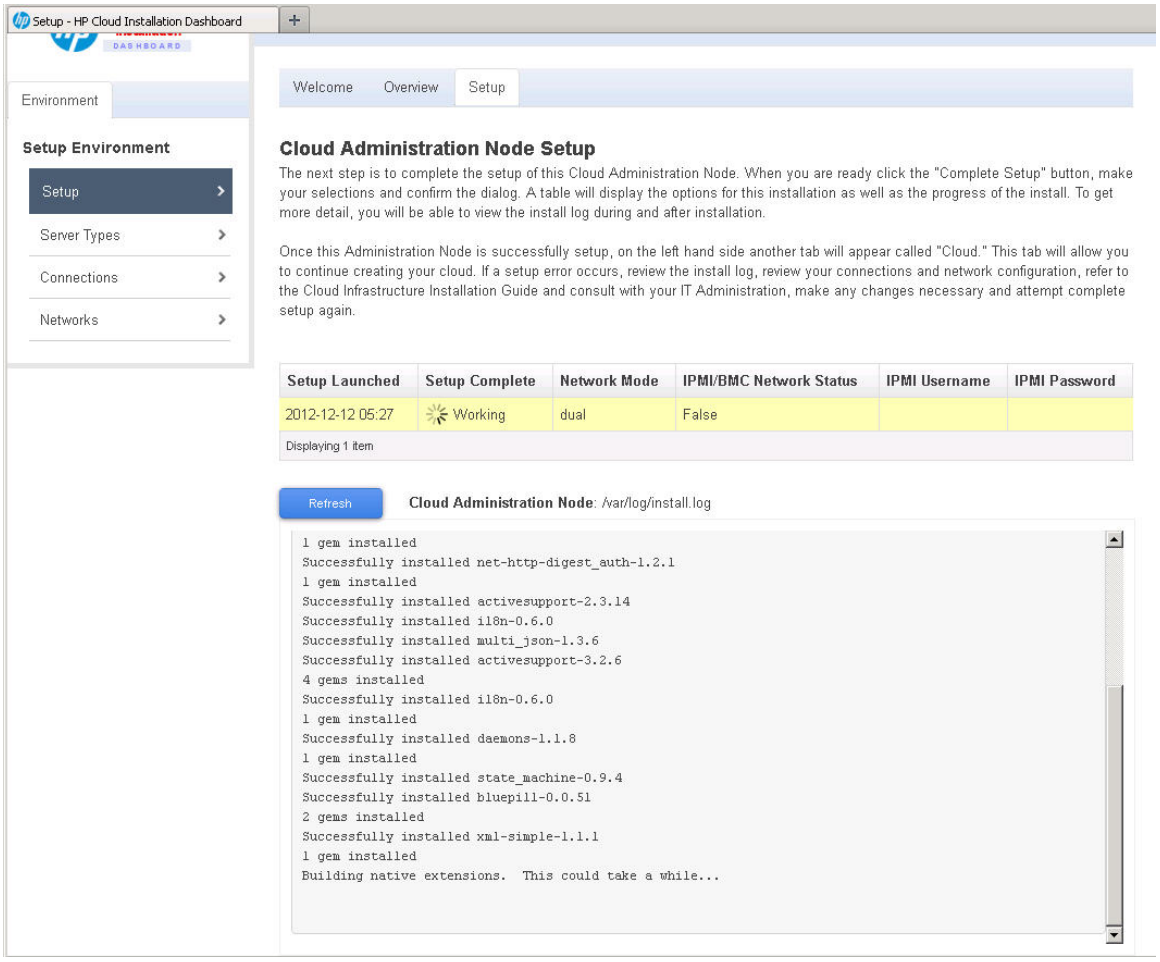

7. Once crowbar is installed, the screen will display the timestamp in the **Setup Complete** column.

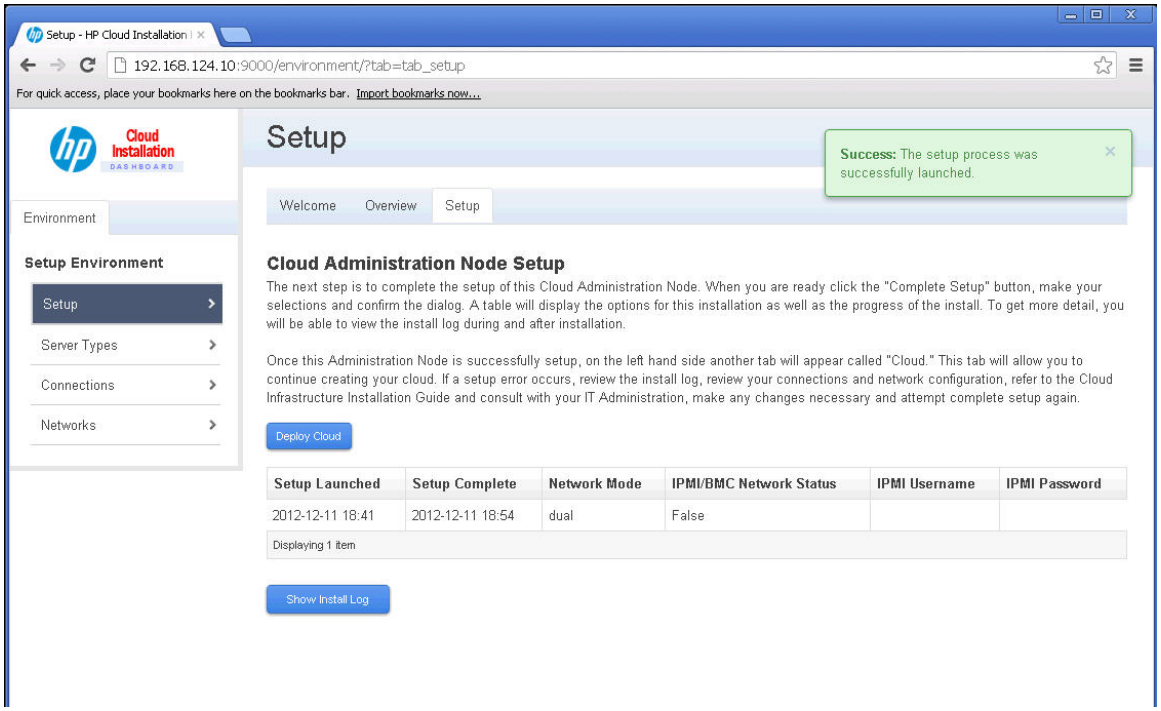

8. Click **Show Install Log** (if you did not do this previously) to view install activity.

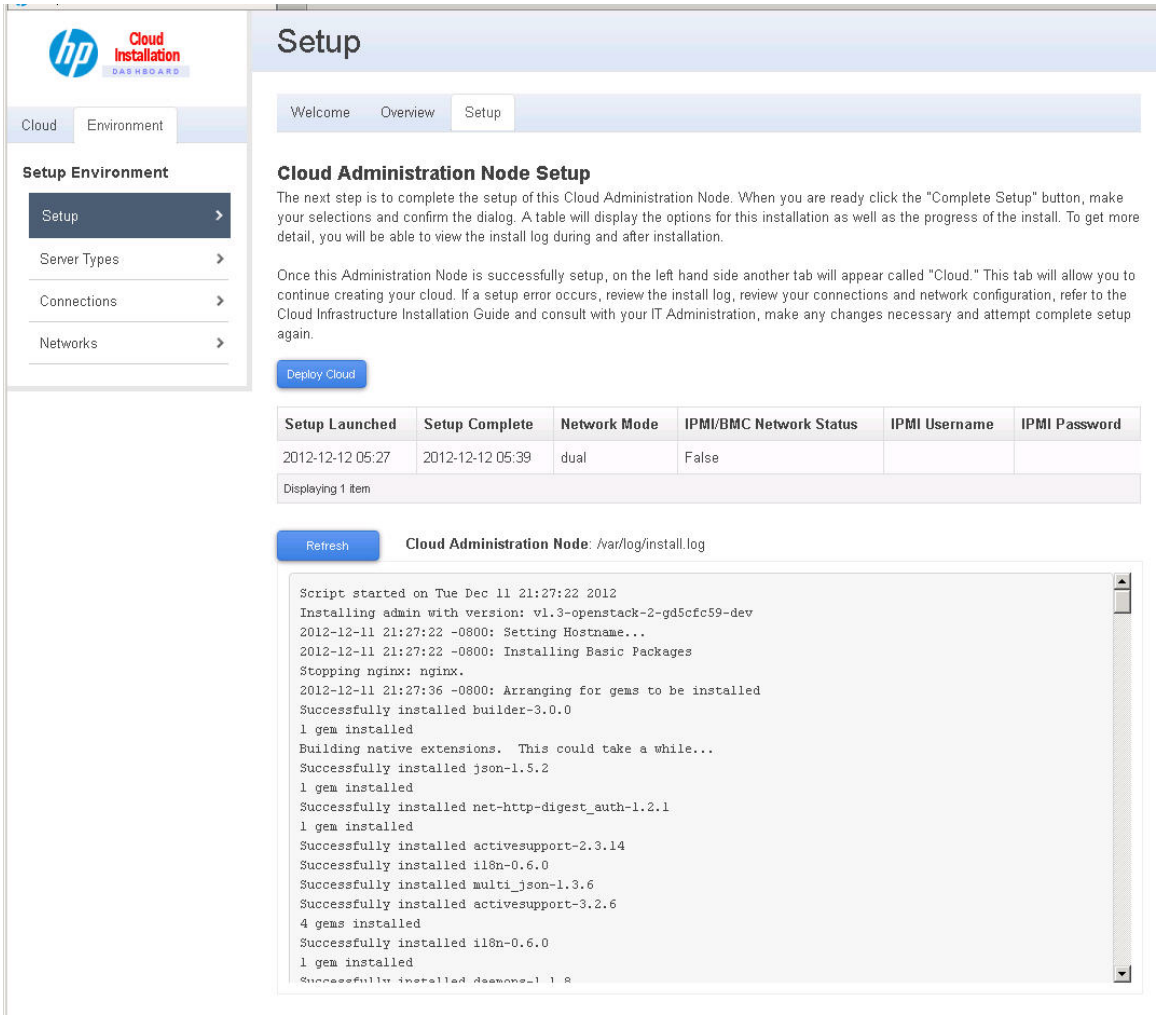

Next Step

The Cloud Administration node is ready to be configured. Refer to the Configure Cloud Administrator Node section.

Configure Cloud Administrator Node

Overview

Ensure the Cloud Administration node has been prepared as outlined in the previous section.

The following steps must be executed to begin configuring the Cloud Administration node.

1. Click **Deploy Cloud**.

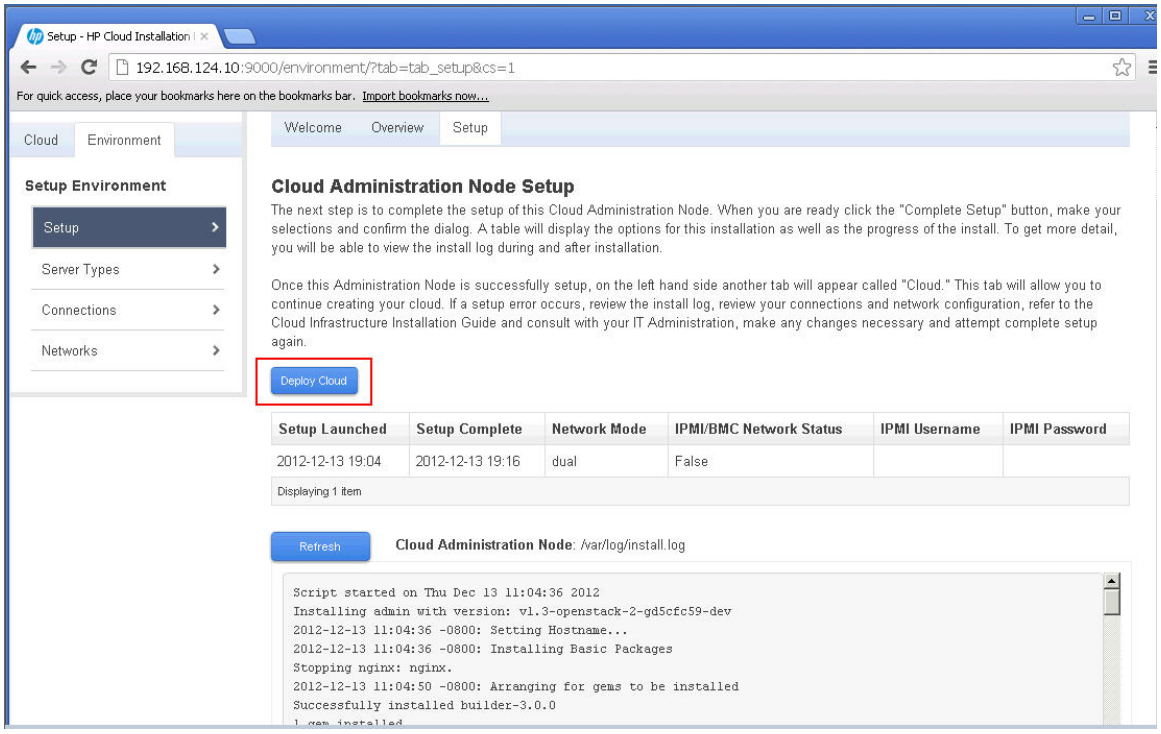

2. A **Cloud** tab becomes visible in the left navigation panel.

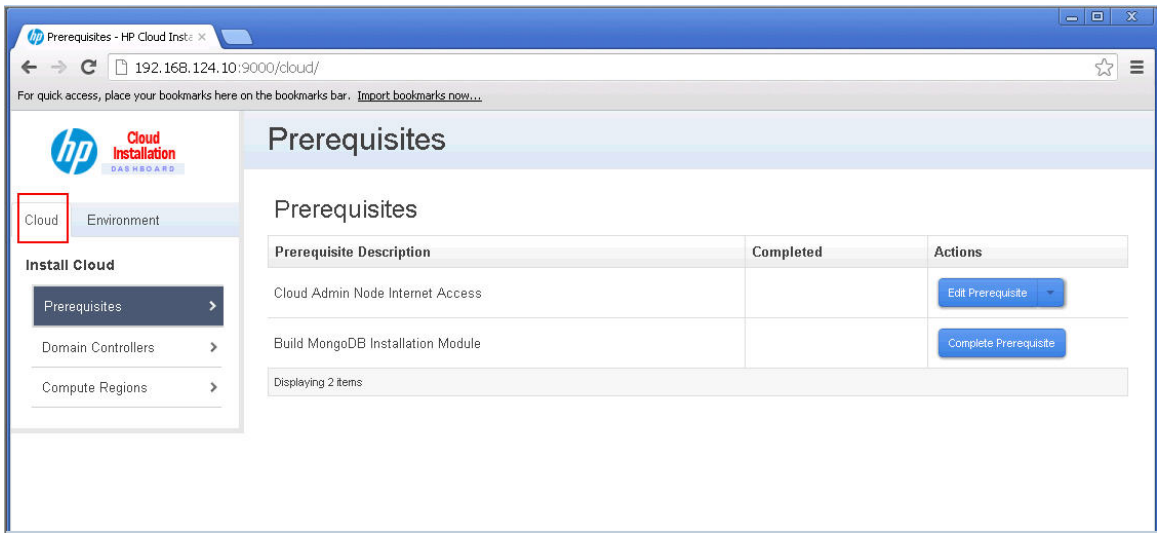

Enable Cloud Administration Node for Internet Access

Cloud Administration node temporarily needs access to the Internet to download and install necessary third-party packages such as MongoDB. The settings for public access of Admin node will be disabled post reboot.

Follow these steps to enable networking:

1. Click **Edit Prerequisite** in the **Cloud Admin Node Internet Access** entry.

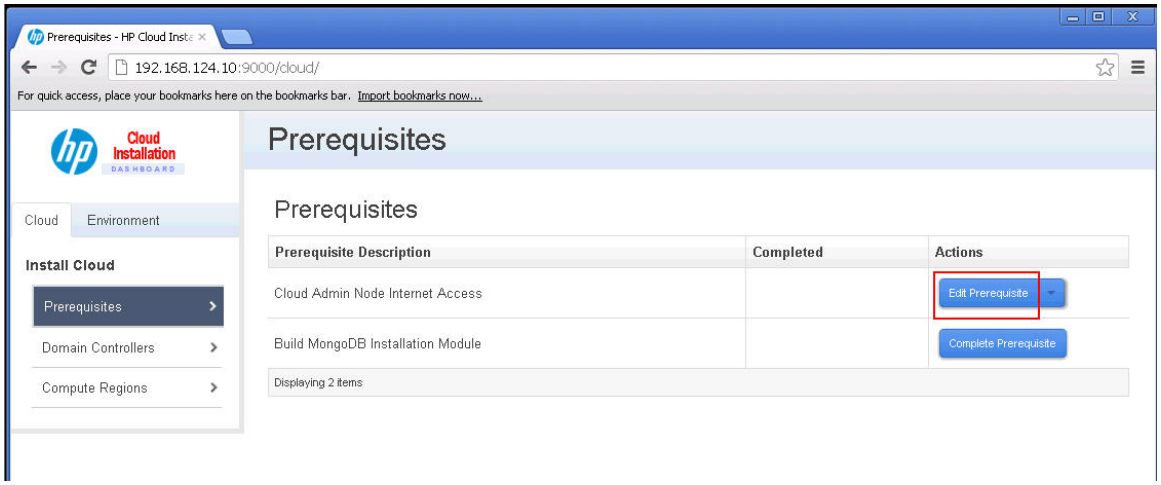

2. The Edit Cloud Admin Node Internet Access Prerequisite screen displays.

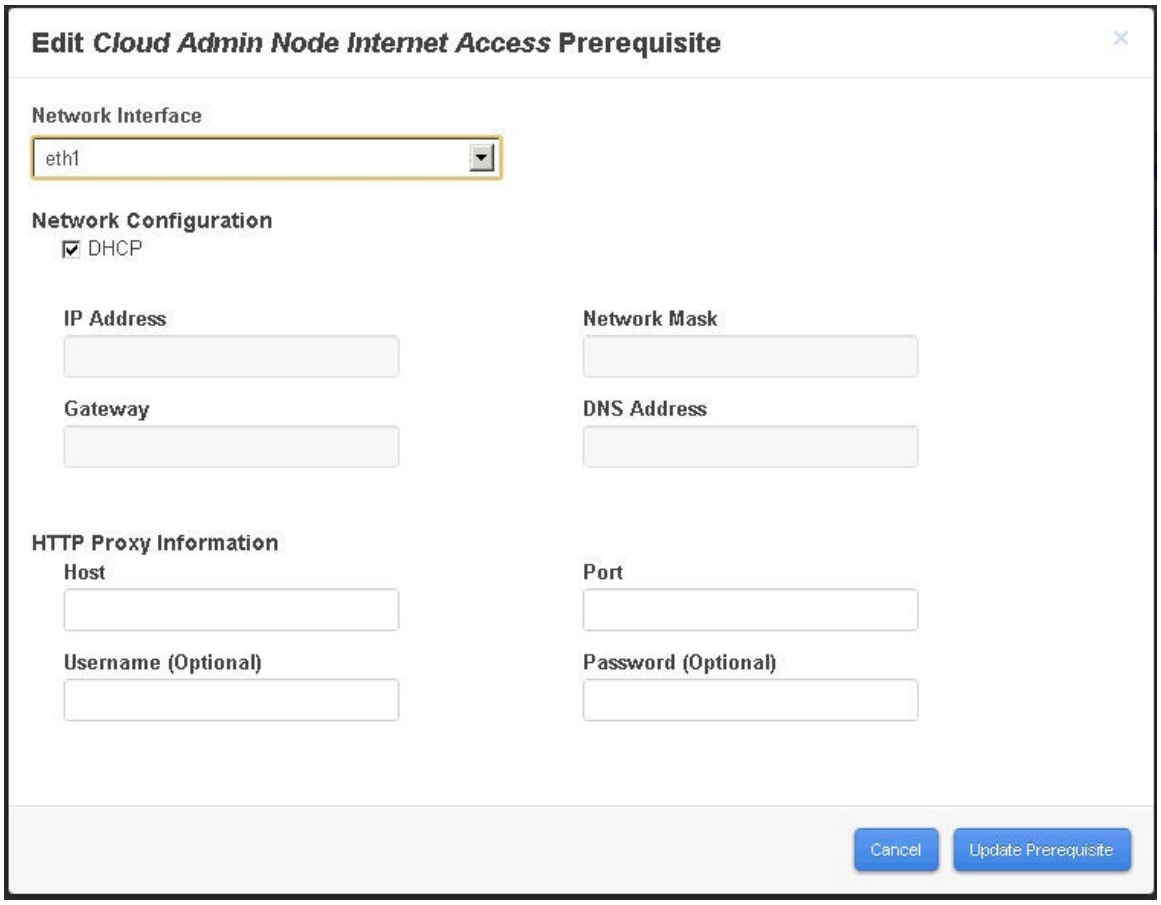

3. Set **Network Interface** to **eth1** (**NOTE**: eth1 is connected to a network providing internet access.)
4. Use the default **Network Configuration** checked as **DHCP**
5. Set **HTTP Proxy Information** applicable for your environment.
    a. Specify the **Host**
    b. Specify the **Port**
    c. If necessary, specify the optional **Username** and **Password**.
6. Click **Update Prerequisite**.
7. Click **Complete Prerequisite** for the **Cloud Admin Node Internet Access** entry. The button is located in the drop-down menu to the right of the **Edit Prerequisite** button.

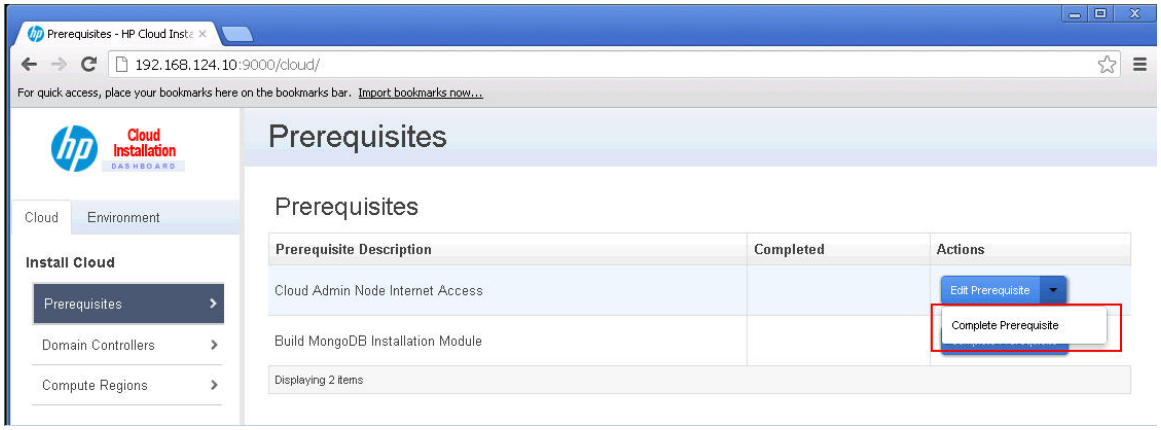

8. Click **Complete Prerequisite** in the confirmation dialog box.

9. Click **Complete Prerequisite** for the **Build MongoDB Installation Module** entry.

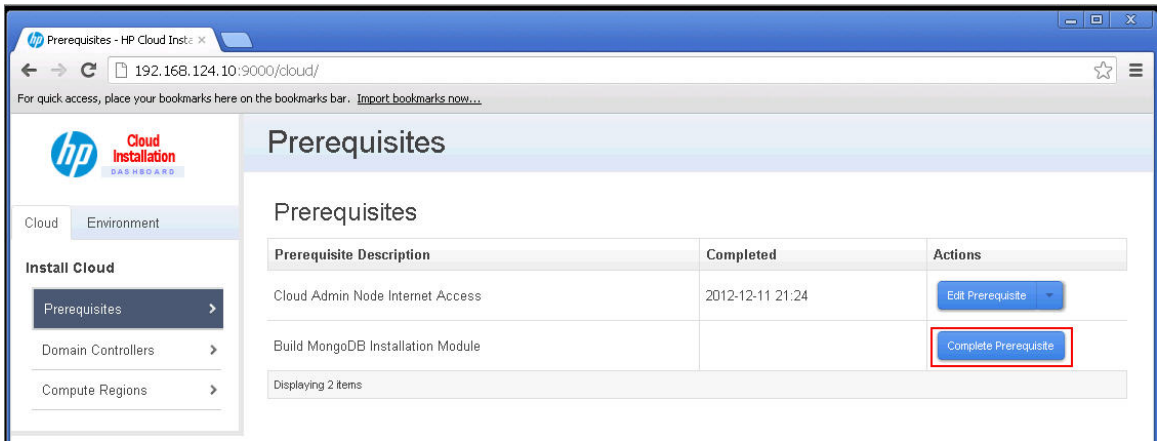

10. Click **Complete Prerequisite** in the confirmation dialog box.

11. It may take several minutes as the MongoDB install packages are being downloaded from the Internet to the Cloud Administration node.
12. Ensure that a time stamp is displayed in the **Completed** column. A success message displays when the MongoDB prerequisites are configured.

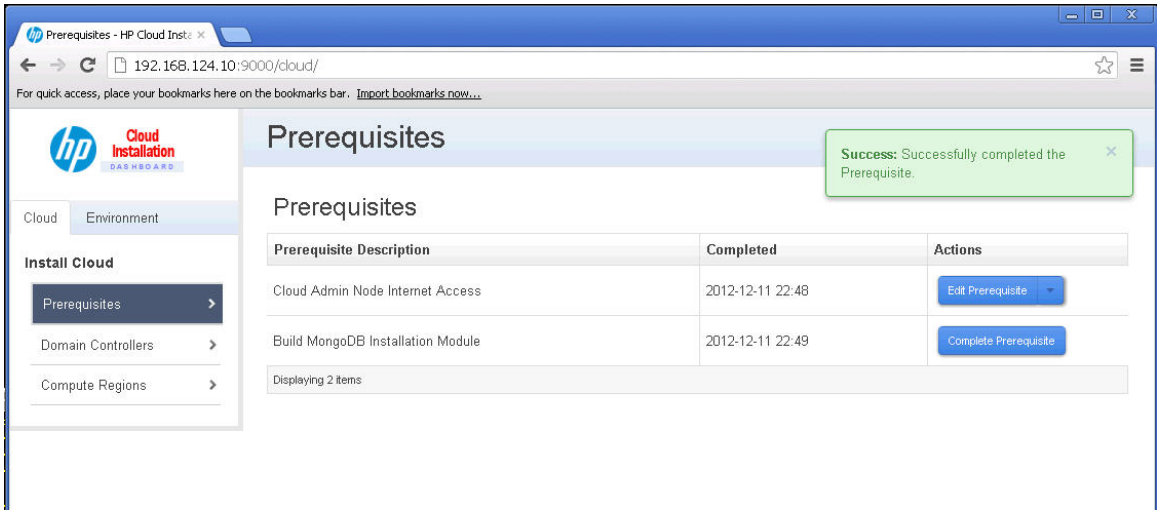

13. If the prerequisites fail to complete, the following error message displays. Consult the Troubleshooting section for a solution

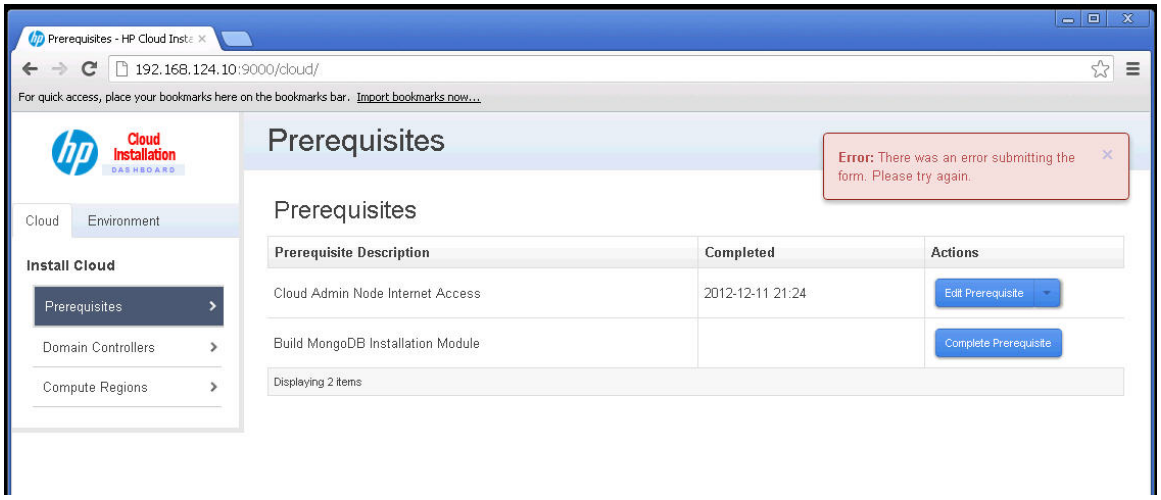

Next Step

The Domain Controller node is ready to be deployed. Refer to the Deploy Domain Controller section.

Deploy Domain Controller

Overview

Ensure that the Cloud Administration node prerequisites have been configured.

Use the following steps to deploy Cloud Infrastructure services to the Domain Controller node from the Cloud Administration (Admin) node. The Admin node is also accessible directly through the web URL: http://192.168.124.10:3000

NOTE: Use Google Chrome or Mozilla Firefox to access the Admin node.

1. Select the **Domain Controllers** tab.

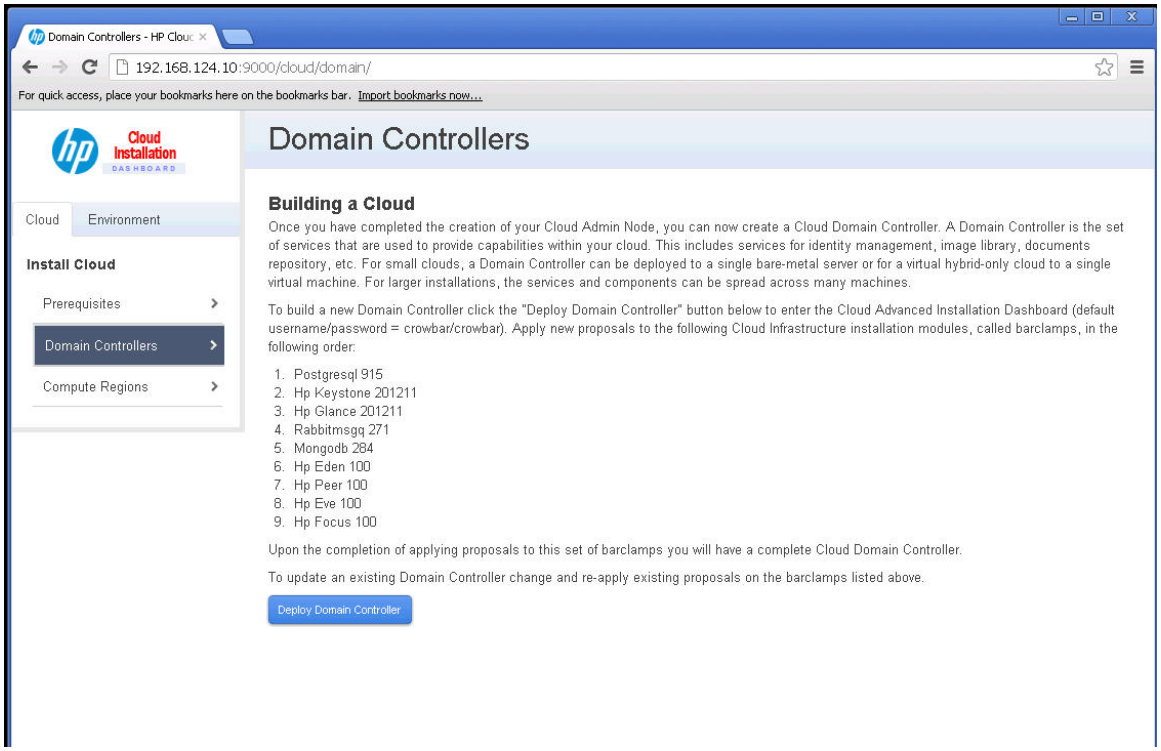

2. Click **Deploy Domain Controller**.
3. Enter the credentials:
    a. User Name: **crowbar**
    b. Password: **crowbar**
4. The **Node Dashboard** displays with a green Admin node indicating it is online.

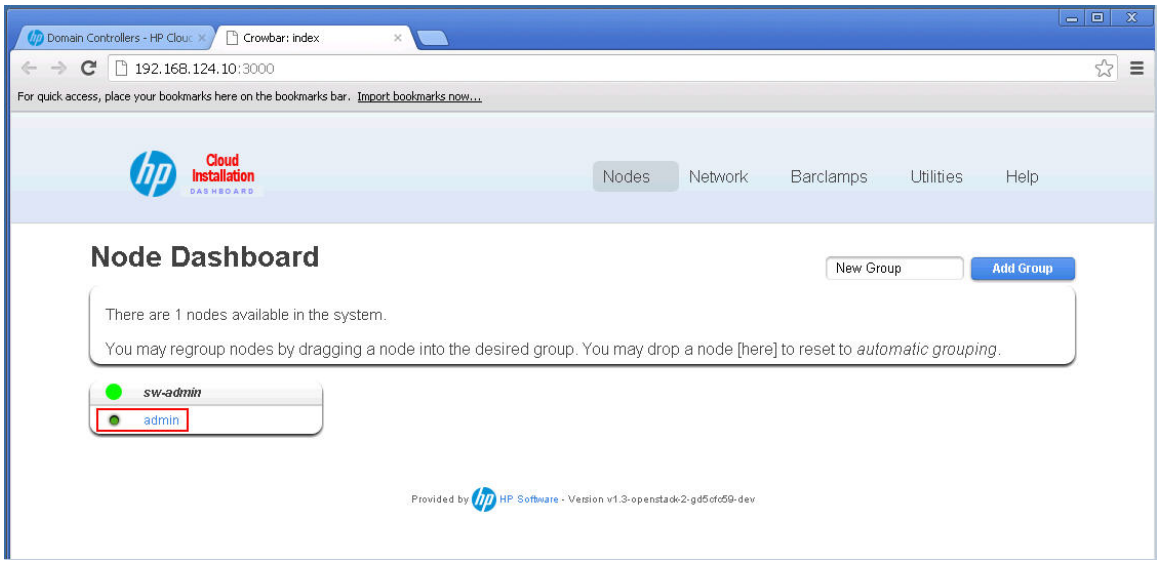

Deploy the Domain Controller

Discover and Allocate Controller node

1. The Domain Controller node needs to be powered ON.
2. When the Domain Controller node boots up, the Cloud Administration node automatically installs through a PXE boot, the Ubuntu operating system.
3. The **Node Dashboard** displays the Domain Controller node as the operating system is installed:
    a. The status indicator is initially grey.
    b. The Domain Controller node displays in the **Node Dashboard** with a generated, MAC Address name.

c. **Tip:** Note the Domain Controller node's console to monitor the system as it is being installed.

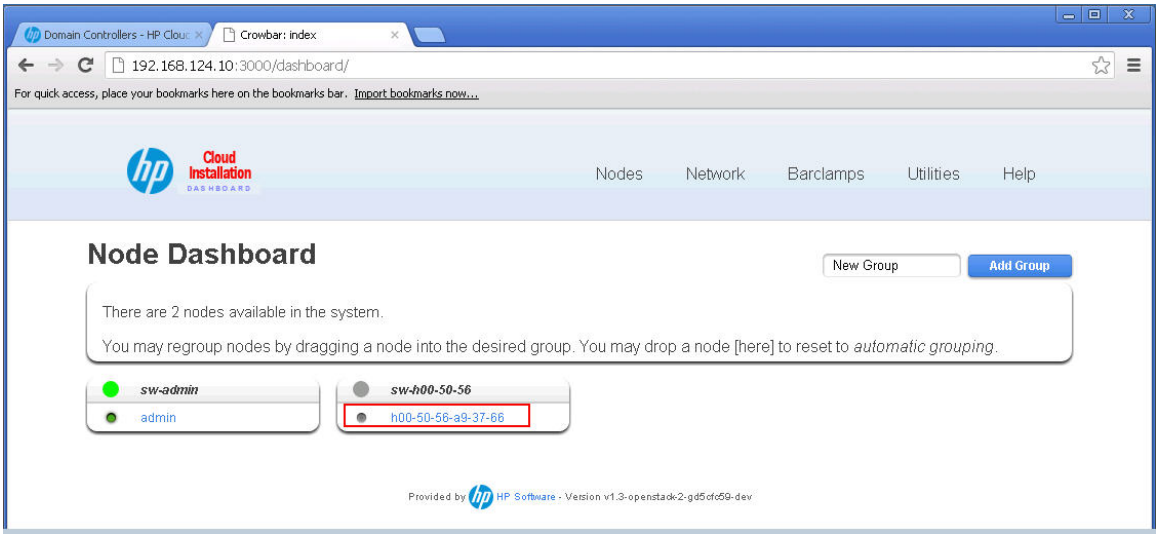

4. Wait until the **Node Dashboard** shows the Domain Controller node with a flashing yellow status.

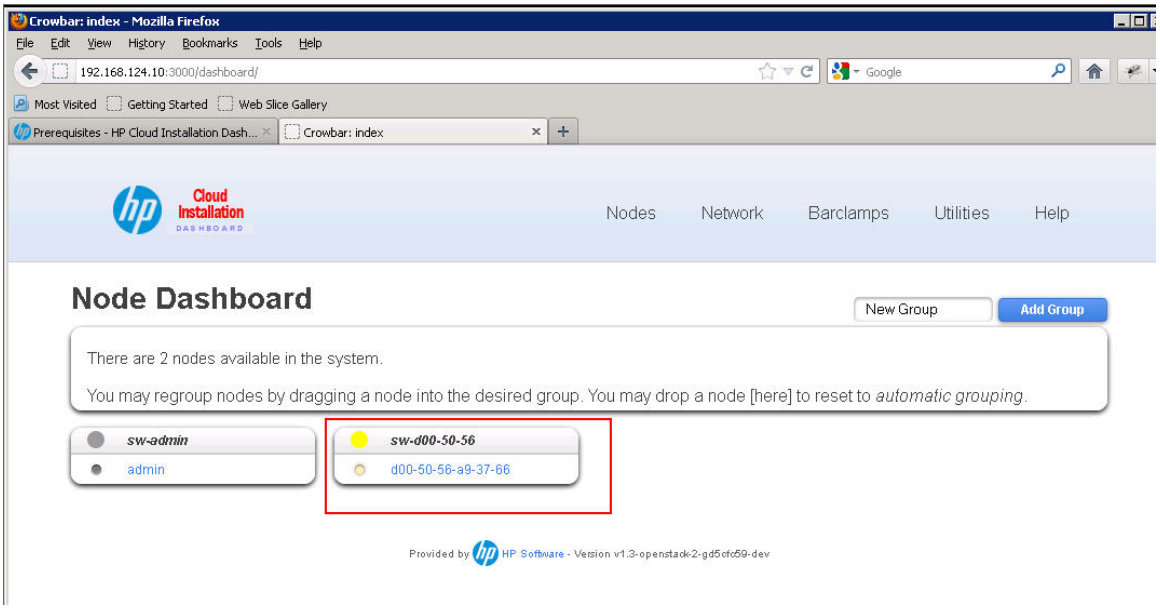

5. To confirm that the MAC Address name is associated to the Domain Controller node, view the Domain Controller node's console.The hostname displayed in the Domain Controller node's console will match the MAC Address name displayed in the **Node Dashboard**.

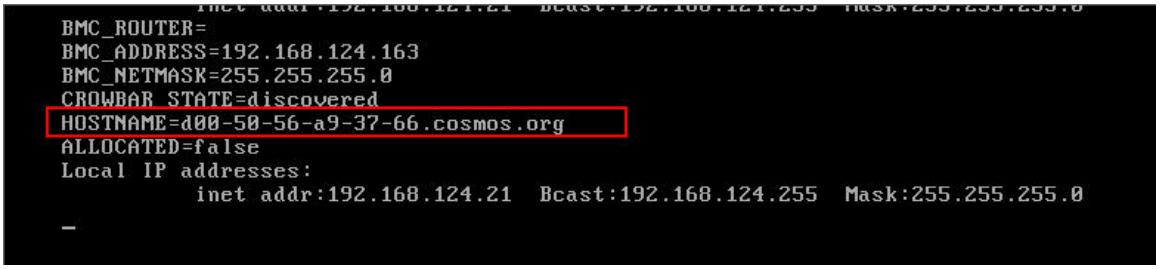

6. On the Cloud Installation Dashboard, select **Nodes > Bulk Edit**.

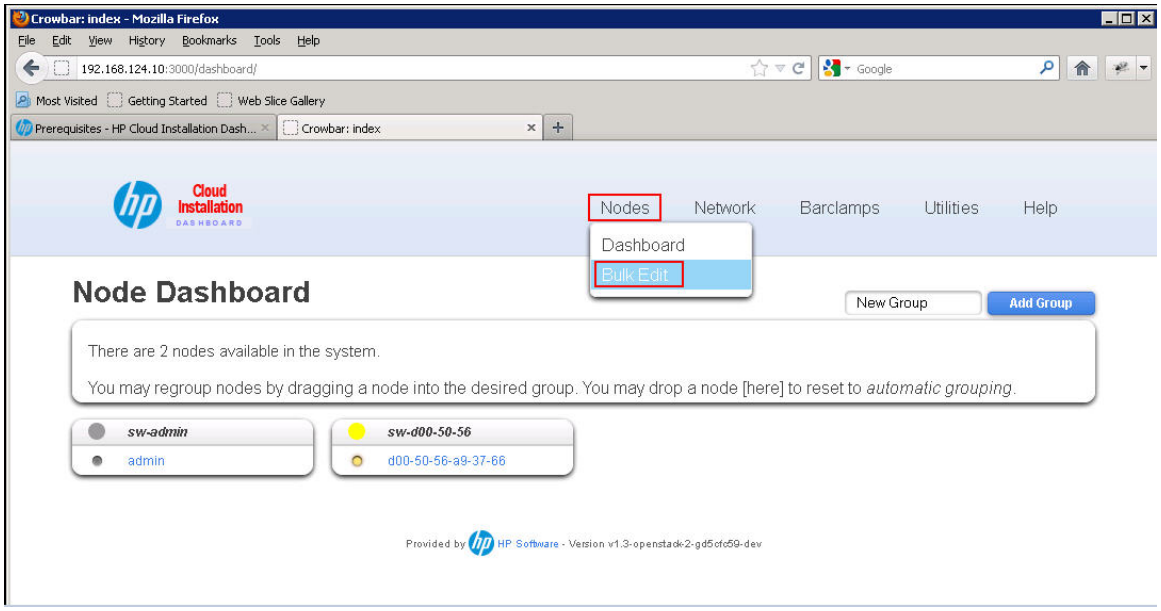

7. Click the **Allocate?** checkbox.

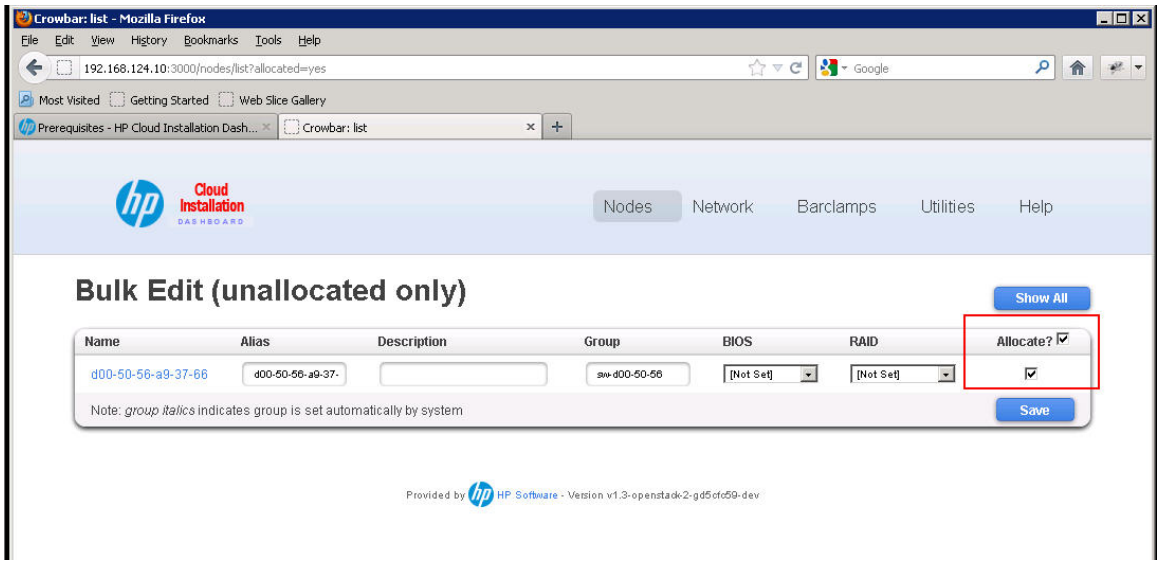

8. Select **Save**.
9. On the Cloud Installation Dashboard, select **Nodes** > **Dashboard**.

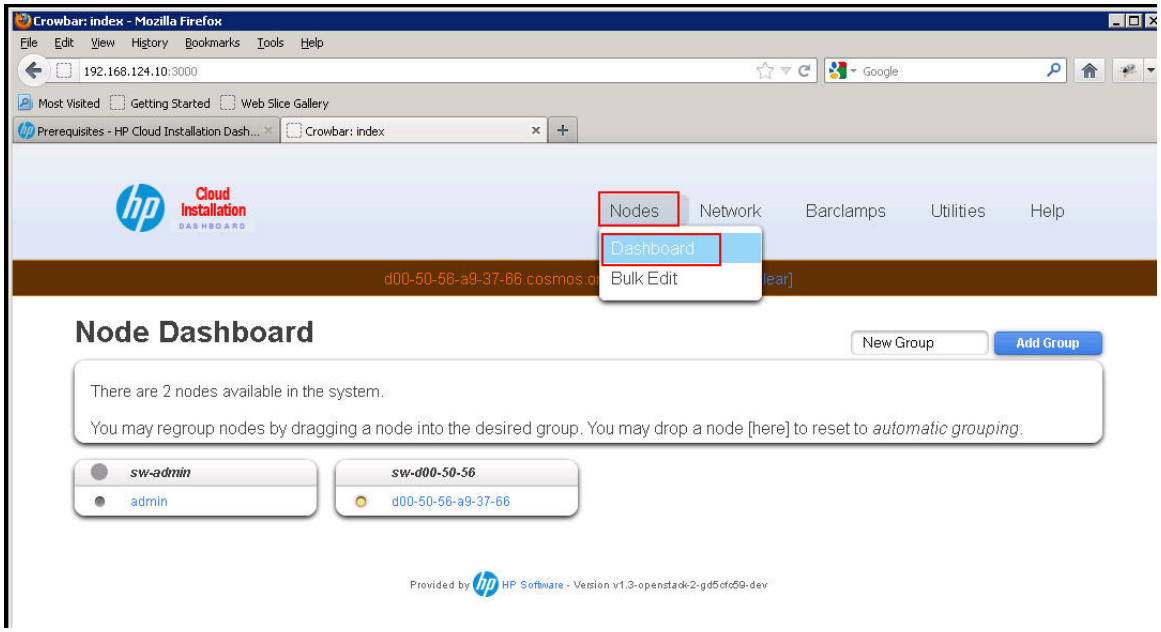

10. Wait for the Domain Controller node's status to change from yellow to green.
    a. View the Domain Controller node's remote console. If you are using HP hardware, monitor installation progress through iLO.
    b. If a dialog box opens with a message to load missing firmware, provide the appropriate driver.

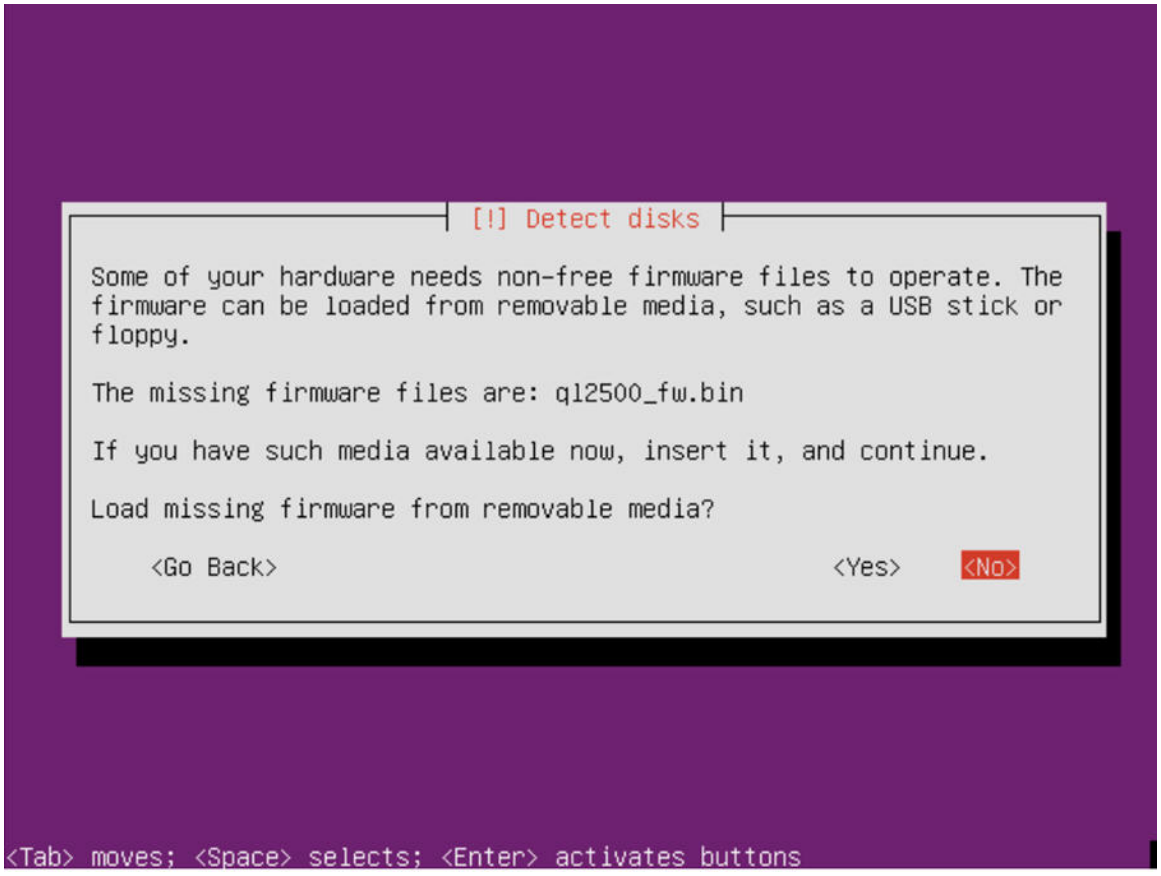

11. When the Domain Controller node installation is complete, the **Node Dashboard** displays the Domain Controller node as Ready.

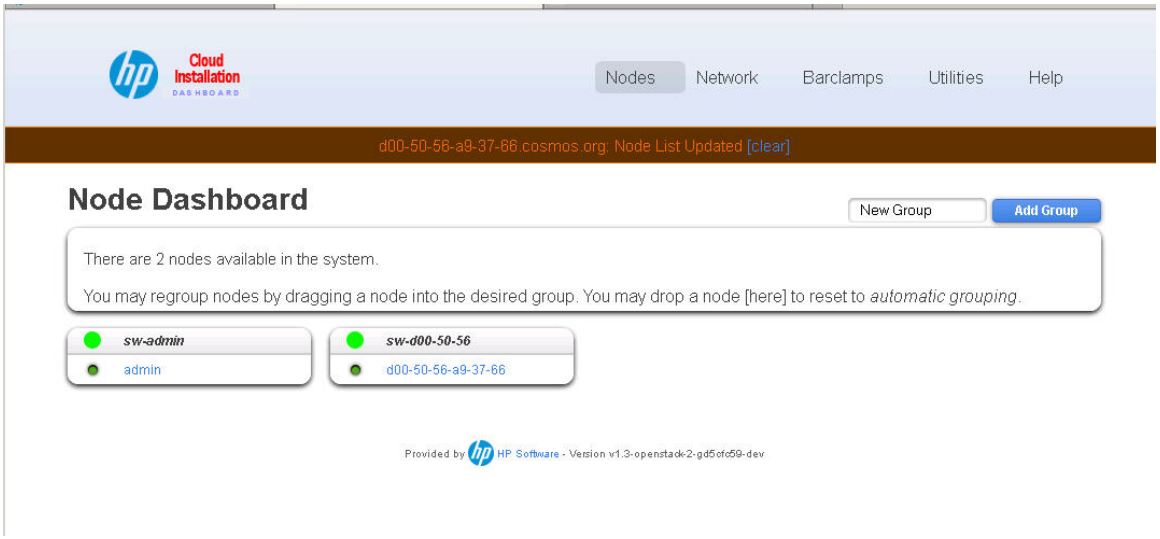

12. **Tip:** The Domain Controller node's login prompt displays the MAC Address name displayed in the **Node Dashboard**.

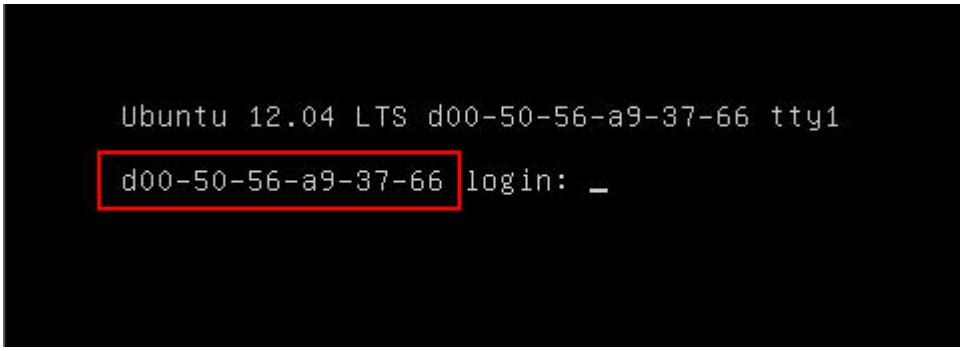

Rename the Domain Controller

It is helpful to rename the generated Domain Controller node name to one reflecting its functionality.

Follow these steps to rename the Domain Controller node:

1. Select the Domain Controller node's **MAC Address name**.

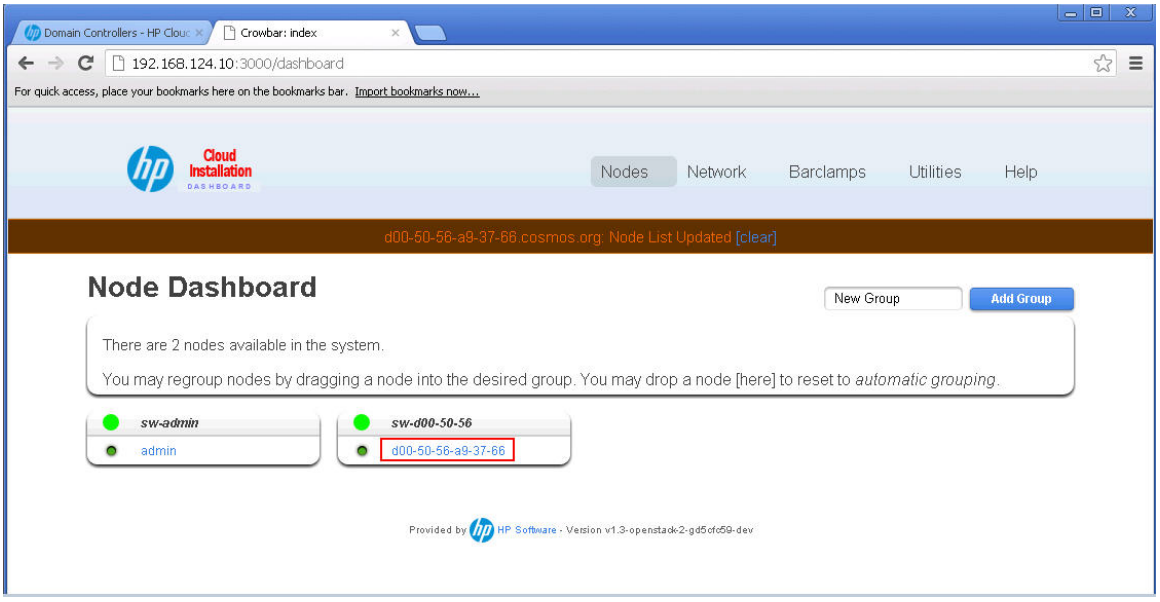

2. Details for the Domain Controller node display.

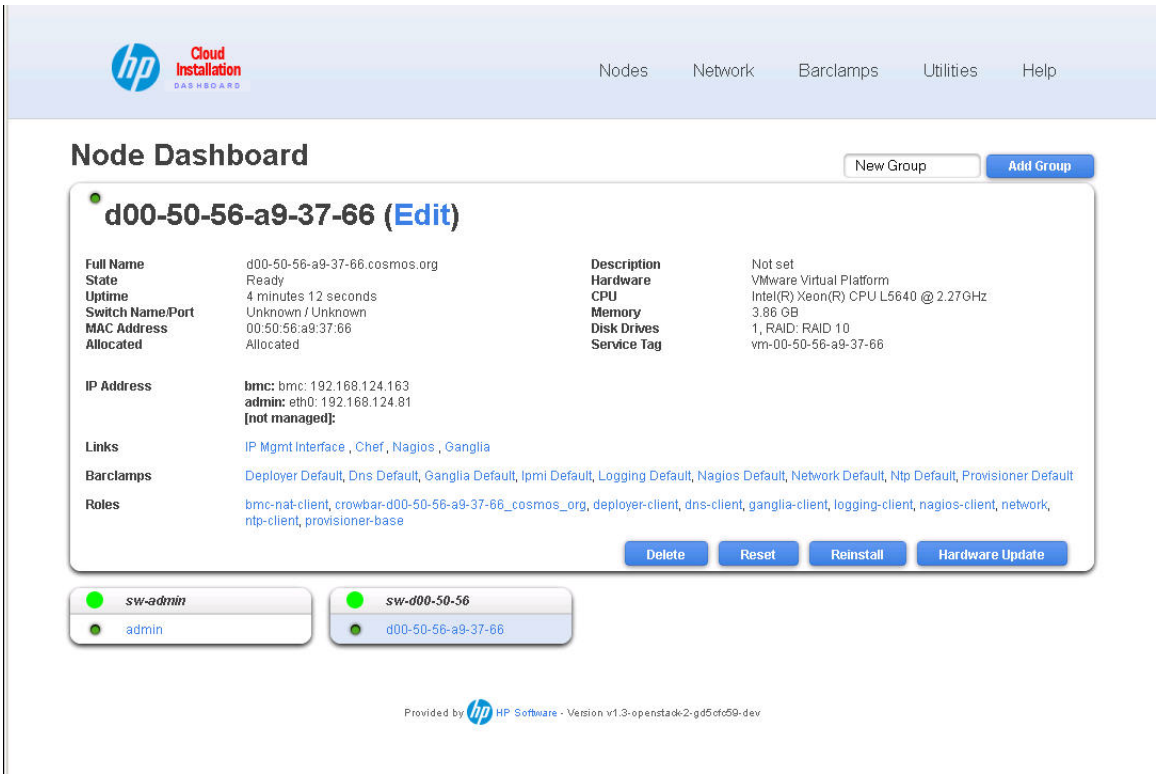

3. Click **Edit**.

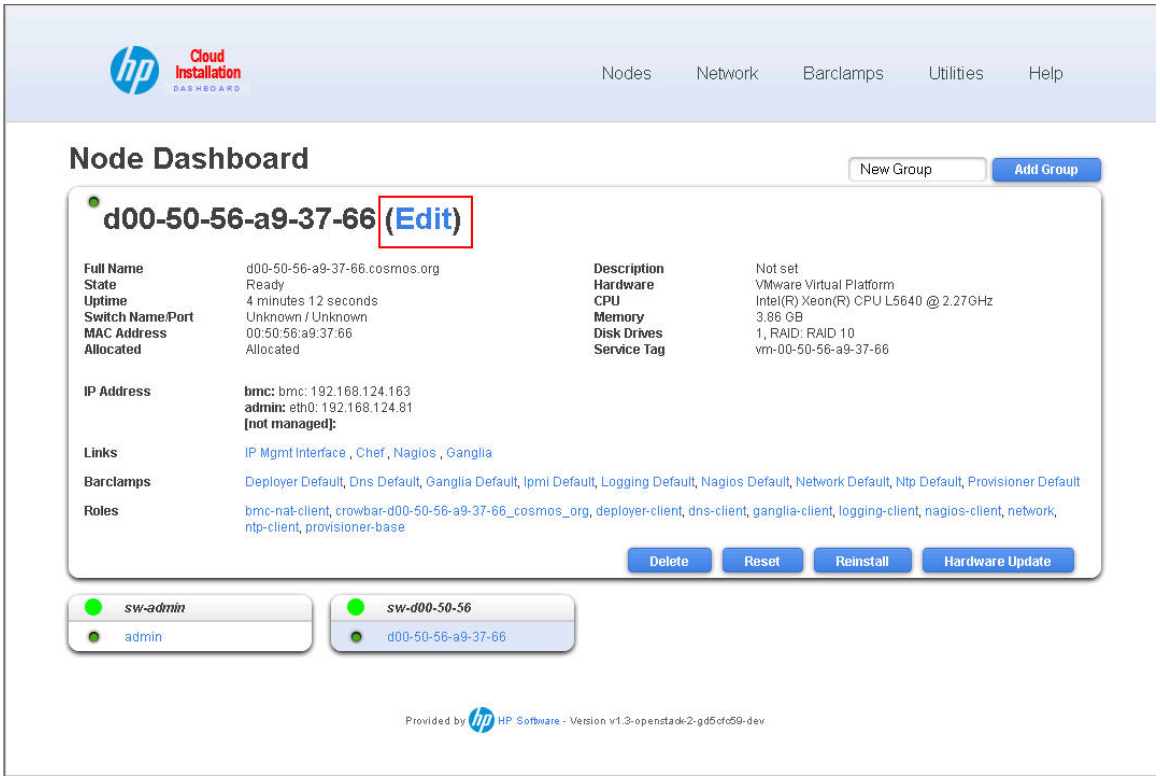

4. Change the alias field from the MAC Address name to a meaningful name.
    a. From:

b. To:

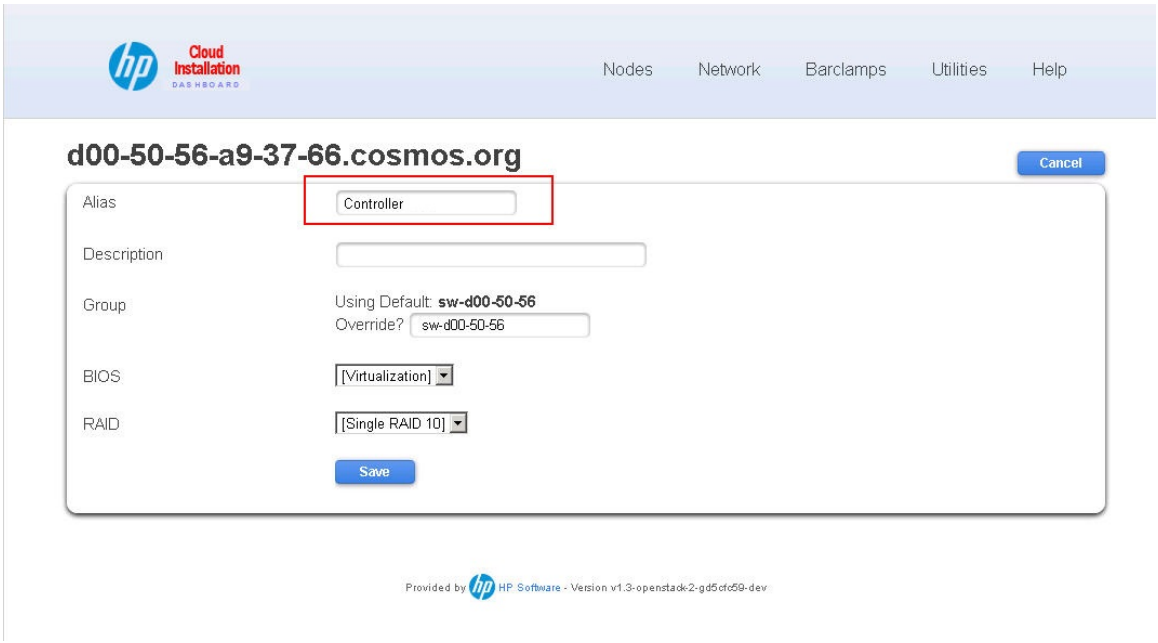

5. Click **Save**. The Domain Controller node's new name is displayed.

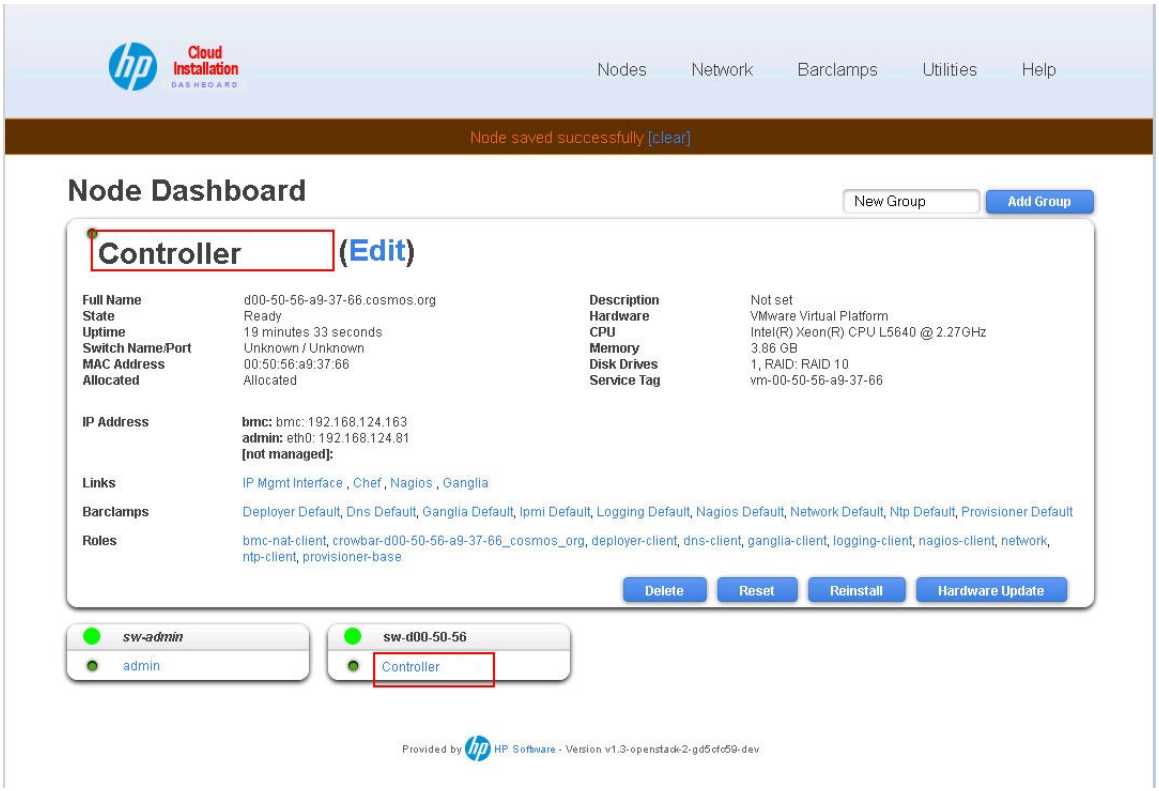

6. Notice the **Full Name** of the Domain Controller node remains unchanged.

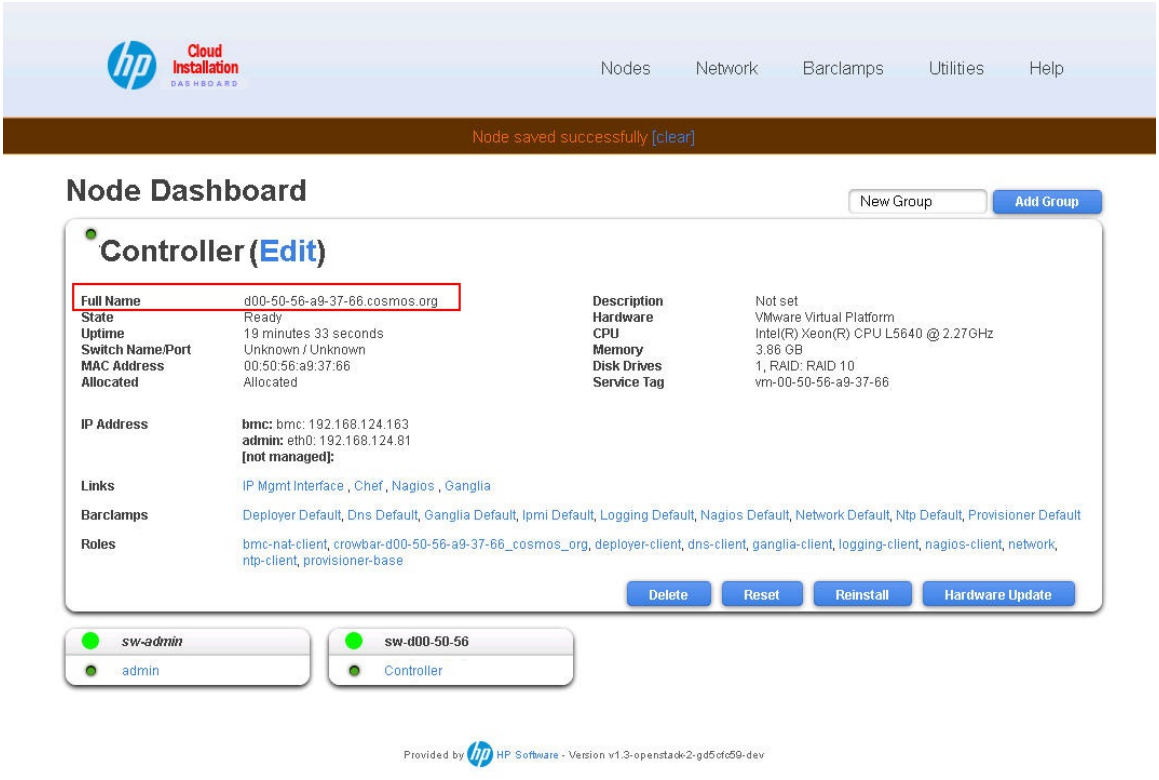

Complete Storage Configuration

With OS up, the next step is to extend the root volume to accommodate nova-volume and image repository. If you plan to configure Nova-Volume to use a raw device, ensure it is correctly discovered by the OS. In some environments, you may have to configure Multi-Path support at this stage.

Apply Barclamps to the Domain Controller

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

Use the following steps to apply barclamps:

1. From the Domain Controller node detail page, select **Barclamps** > **Cloud Infrastructure**.

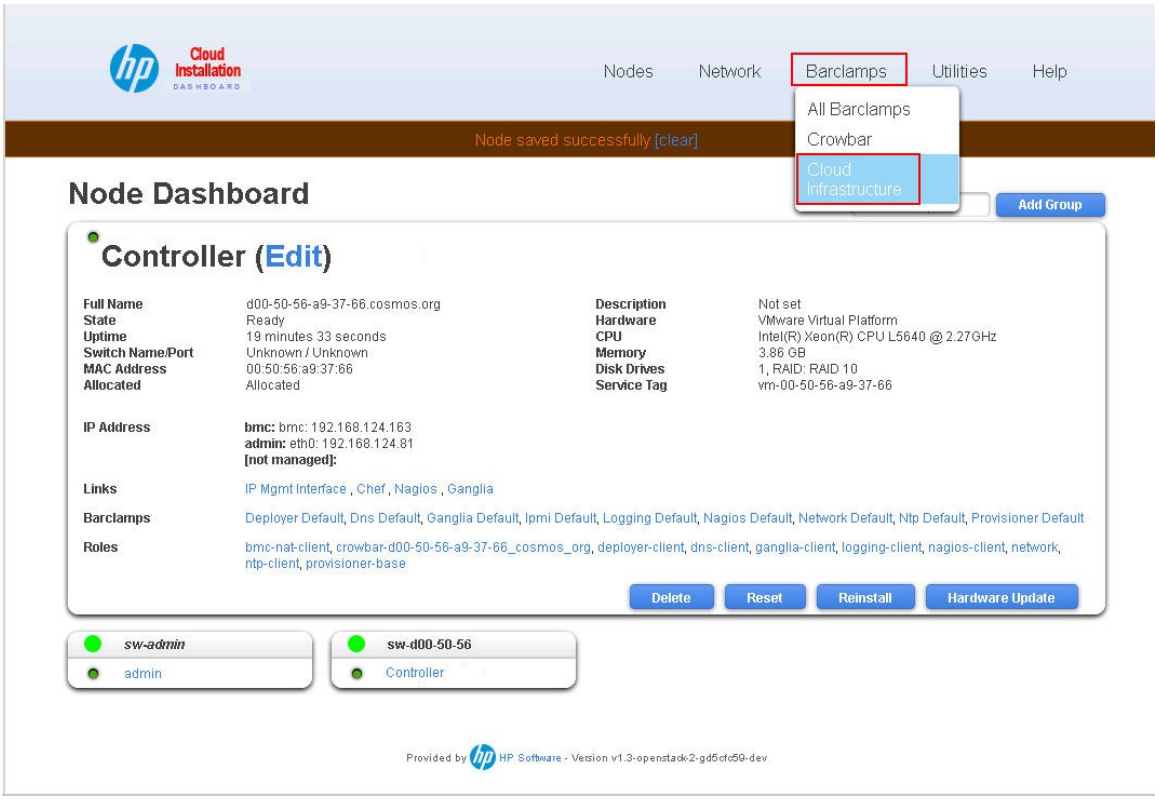

2. The barclamps specific to Cloud Infrastructure display.
    - **NOTE**: It is important that these barclamps be applied in sequential order, each one successfully completing before applying the next barclamp.

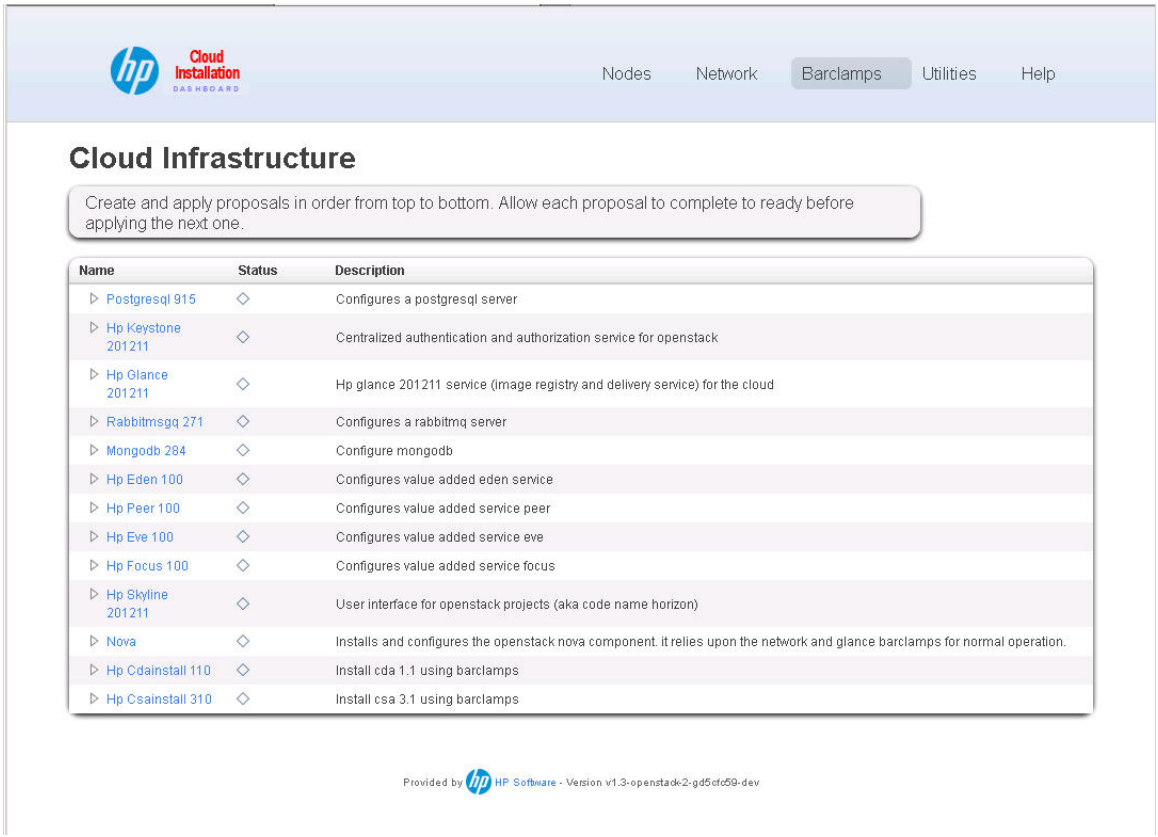

1. For each barclamp in the list, perform the following steps:
   a. Select the **barclamp's name**.

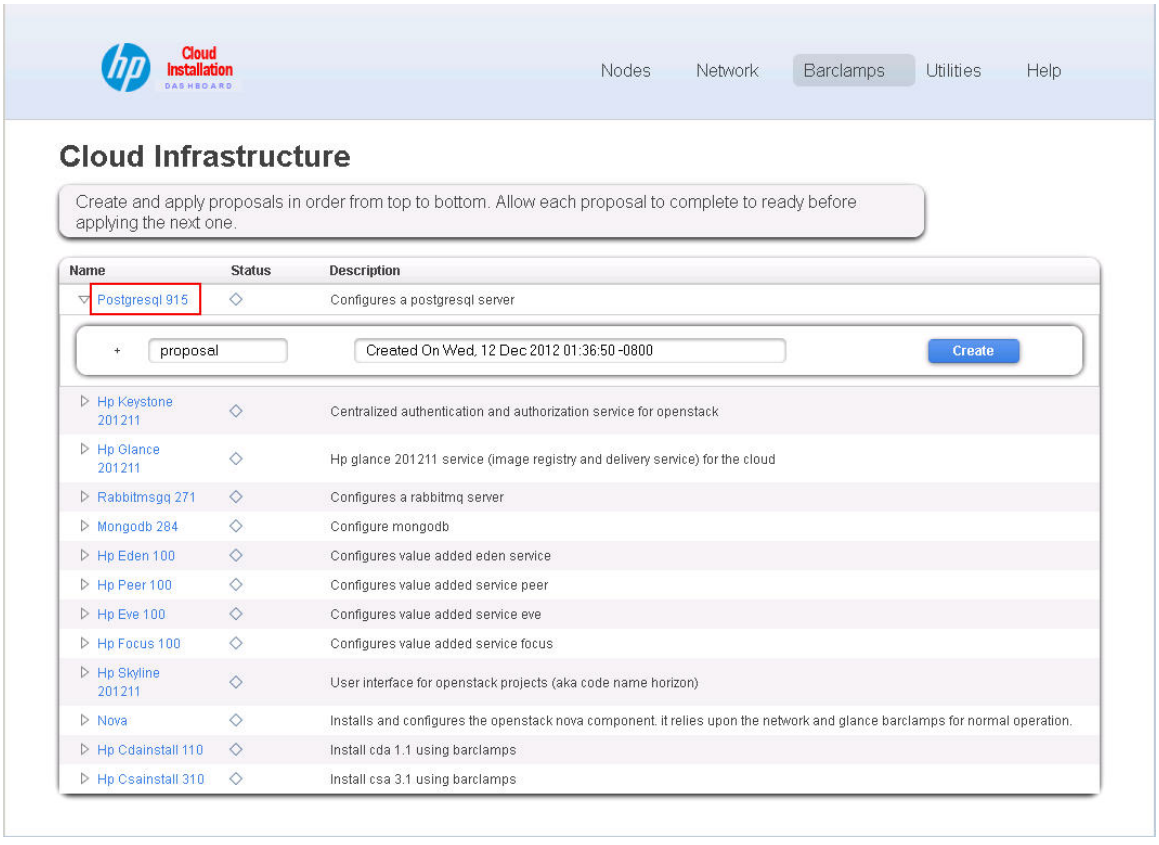

   b. Specify the **Domain Controller's name** as the proposal.

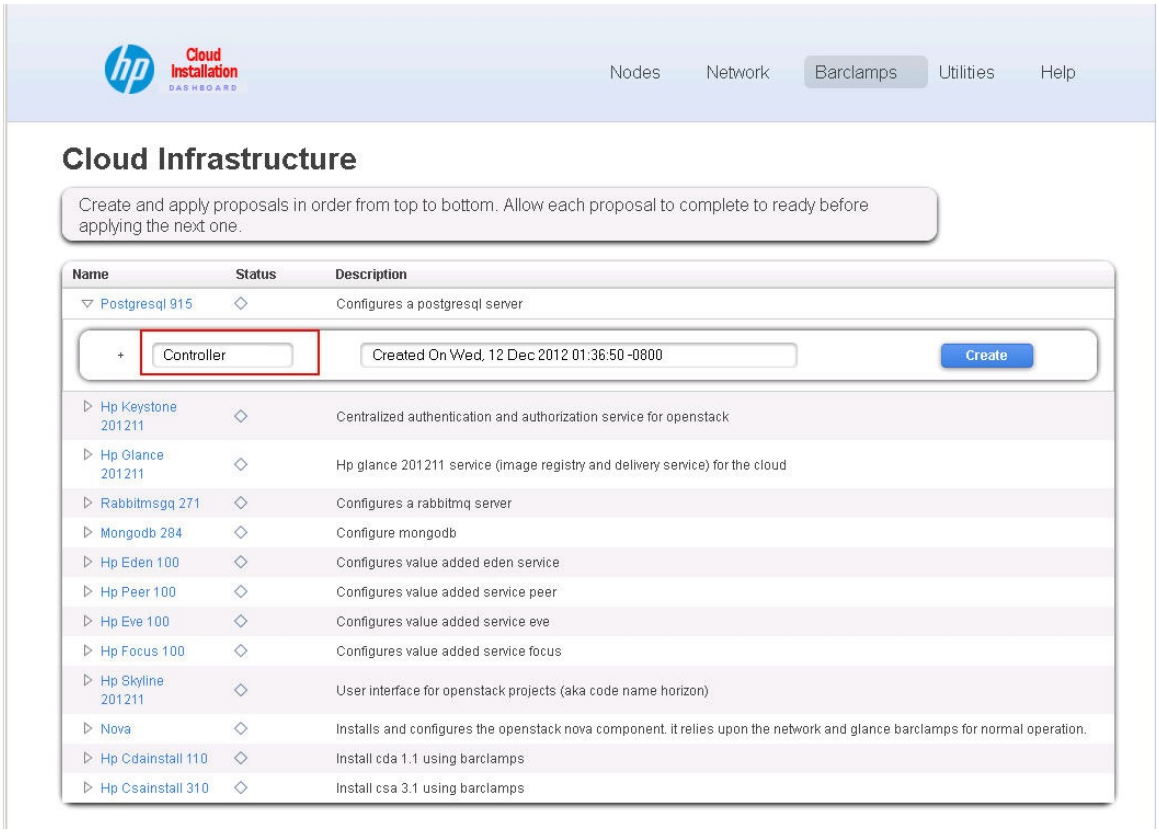

c. Click **Create**.
d. Verify that the default values are correct for the proposal. Some optional values should be set.
    i. Postgres proposal
    ii. Keystone proposal
    iii. Glance proposal
    iv. Rabbit MQ proposal
    v. Mongo DB proposal
    vi. Eden proposal. It is recommended you specify a **Proxy Host** and **Proxy Port** for the proposal. **NOTE**: If specifying a Proxy Host, ensure that the Non Proxy Host(s) field specifies an admin or public network, for example: 192.*|10.*
    vii. Peer proposal
    viii. Eve proposal. It is recommended you specify a **Proxy Host** and **Proxy Port** for the proposal. **NOTE**: If specifying a Proxy Host, ensure that the Non Proxy Host(s) field specifies an admin or public network, for example: 192.*|10.*
    ix. Focus proposal
    x. Skyline proposal. It is recommended you specify a **Proxy Host** and **Proxy Port** for the proposal.
e. Verify that the Domain Controller node's name is listed for each applicable role. The defaults should be sufficient.
    i. At the bottom of the proposal, the items on the left are **Available Nodes**; the items on the right are the **Roles**.
    ii. Select the Domain Controller node name link in the **Available Nodes** list to drag it to the appropriate role (if not done previously).

1. Selecting the Domain Controller node link icon and dragging it to a role will not work.

2. Selecting the Domain Controller node delete icon removes the association of the Domain Controller node from a role.

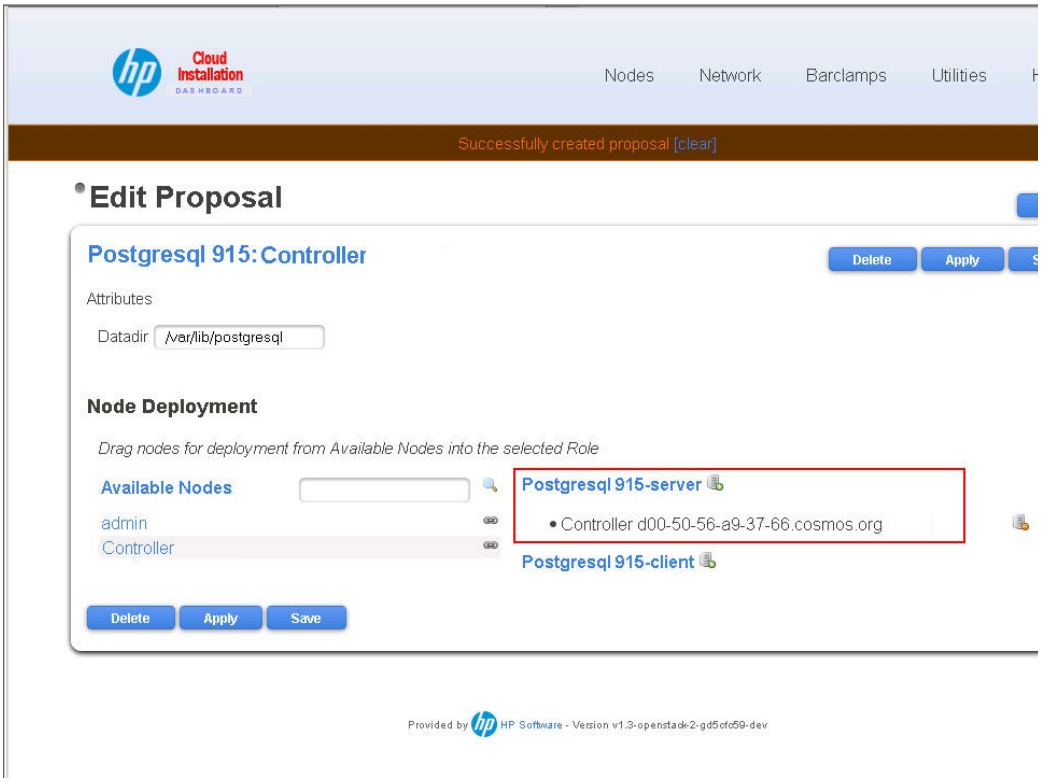

f. Click **Apply**. Click **OK**.
g. Wait for the proposal status (displayed at the top left of the dialog box) to turn green.

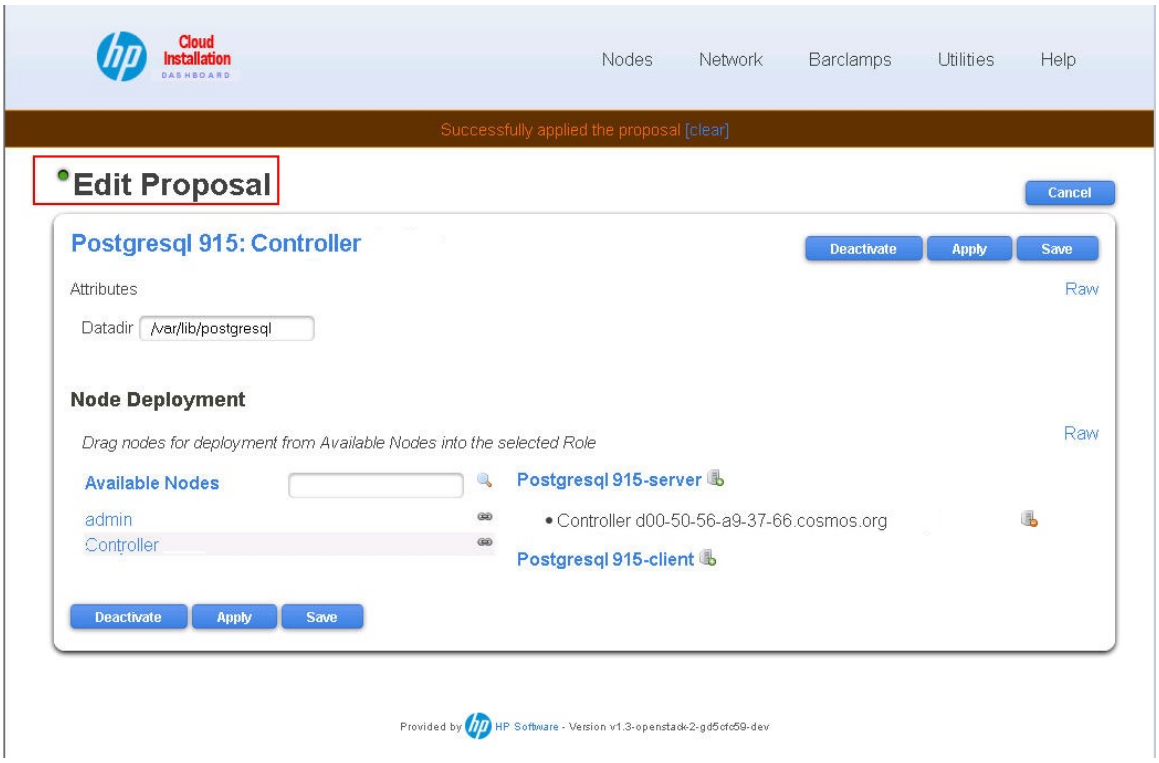

h. If the proposal fails to apply, the following error message displays. Consult the Troubleshooting section for a solution.

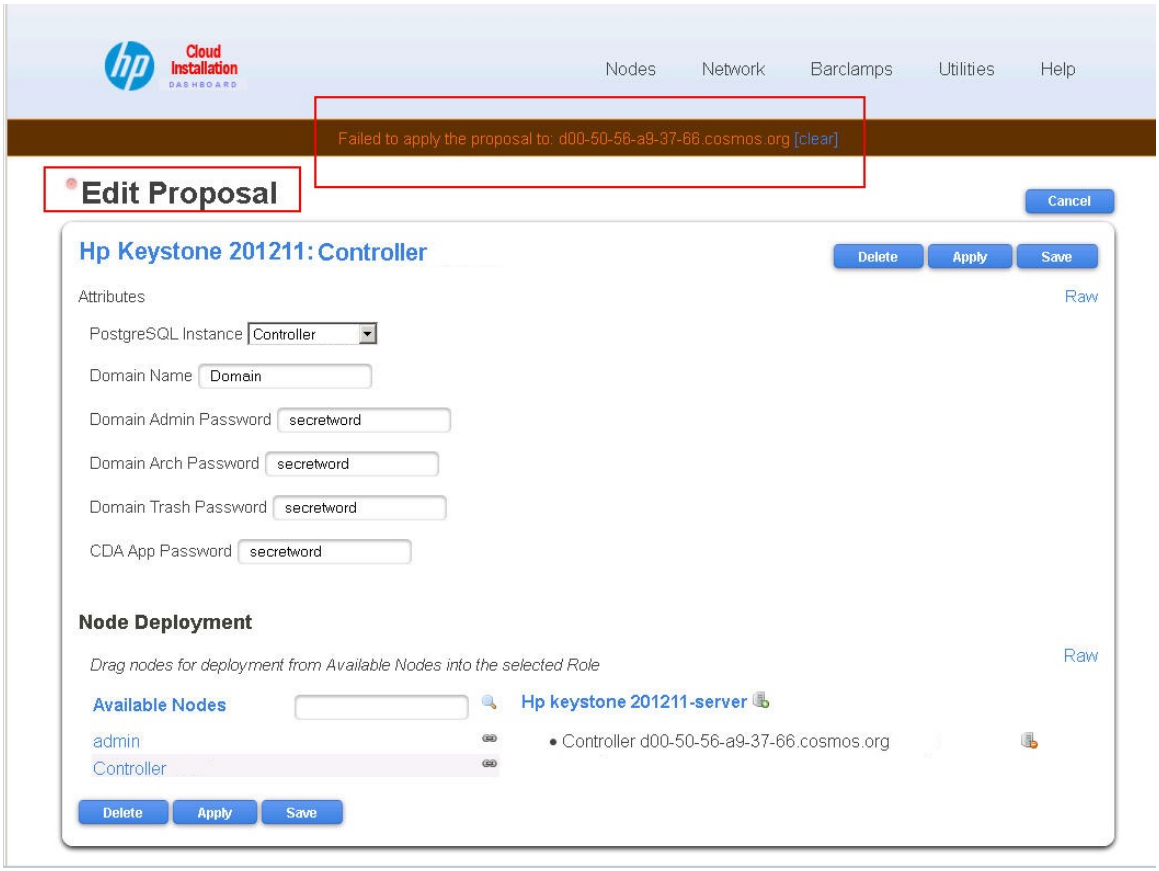

i. Re-select **Barclamps** > **Cloud Infrastructure** to view the barclamp list.

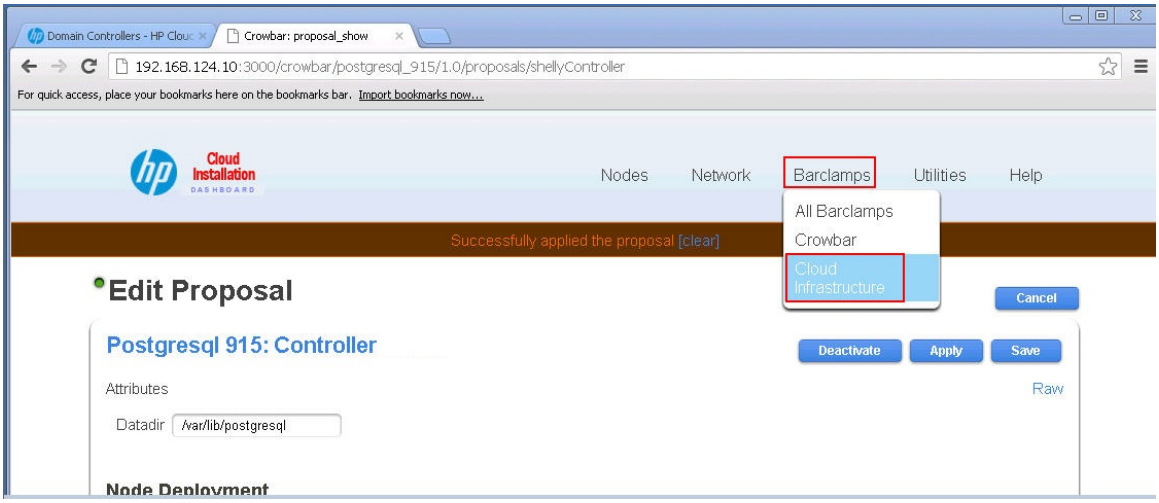

j. The proposal status is also displayed within the Cloud Infrastructure barclamp list.

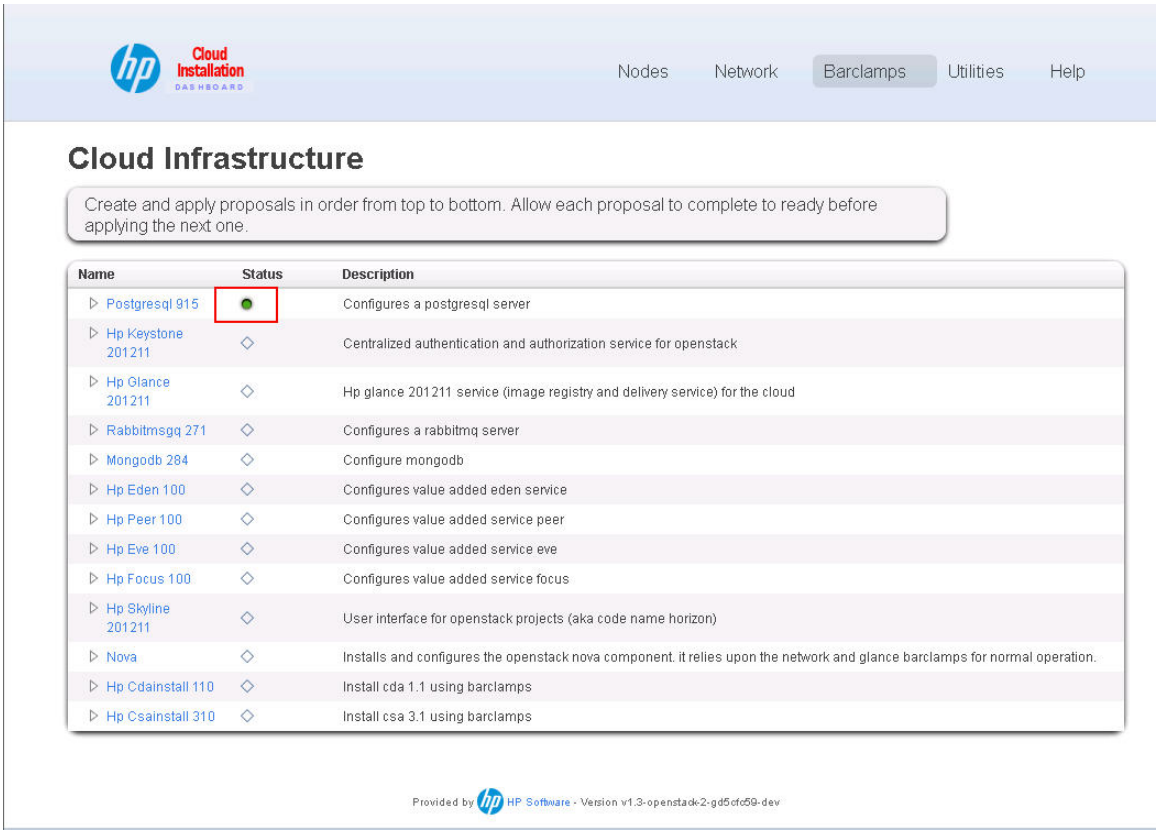

k. Apply the next barclamp in the list, from top to bottom until all Cloud Infrastructure barclamps have been applied.
l. At this stage, you may optionally apply CDA and CSA barclamps. Ensure you have completed the pre-requisite "Coexisting of CSA and CDA on same host" before applying CSA barclamp. Also, ensure you have completed the post-installation steps of CSA barclamp. Refer to Deploying CSA and CDA sections.

Next Step

The Compute Region node is ready to be deployed. Refer to the Deploy Compute Region section.

Deploy Compute Region

Overview

Ensure that the Cloud Administration and Domain Controller nodes have been deployed successfully. Use the following steps to deploy Cloud Infrastructure services to the Compute node from the Cloud Administration node.

1. Select the **Compute Regions** tab.

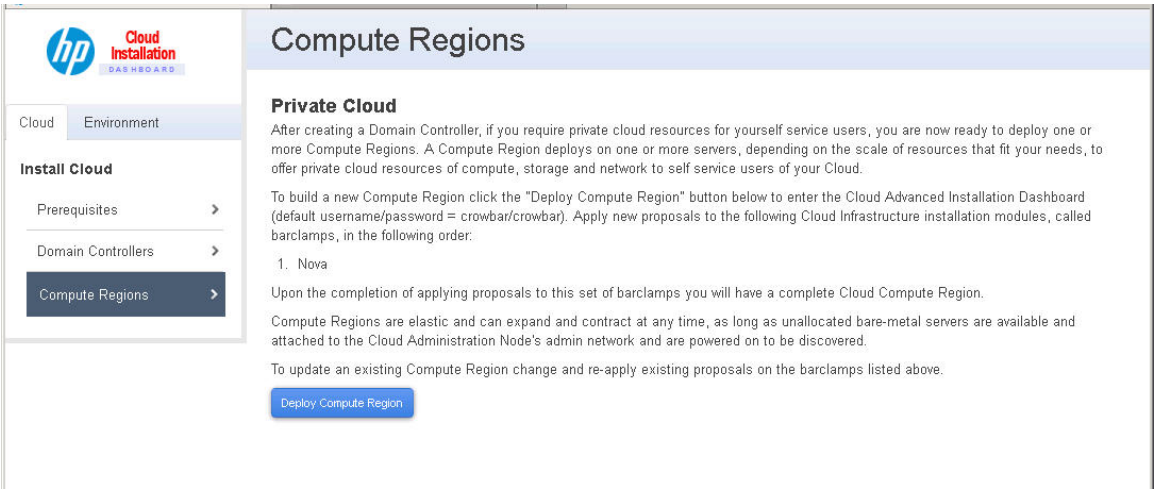

2. Select the **Deploy Compute Region** button.
3. Enter crowbar/crowbar for the credentials.

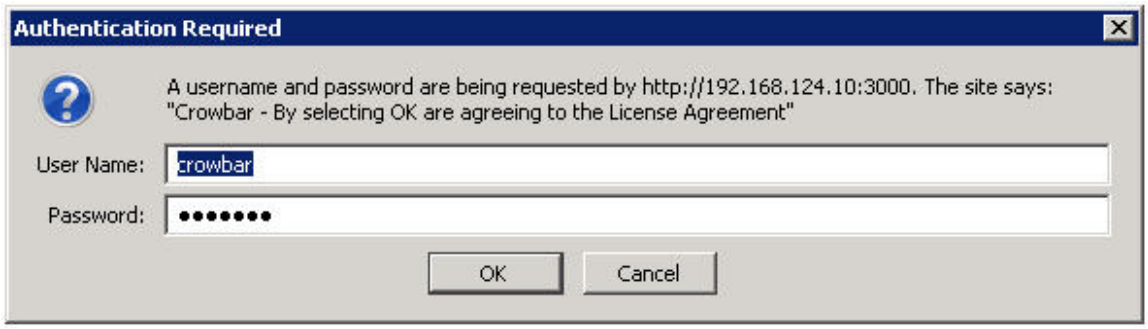

4. The **Node Dashboard** displays with the Admin and Controller nodes displaying a green online status.

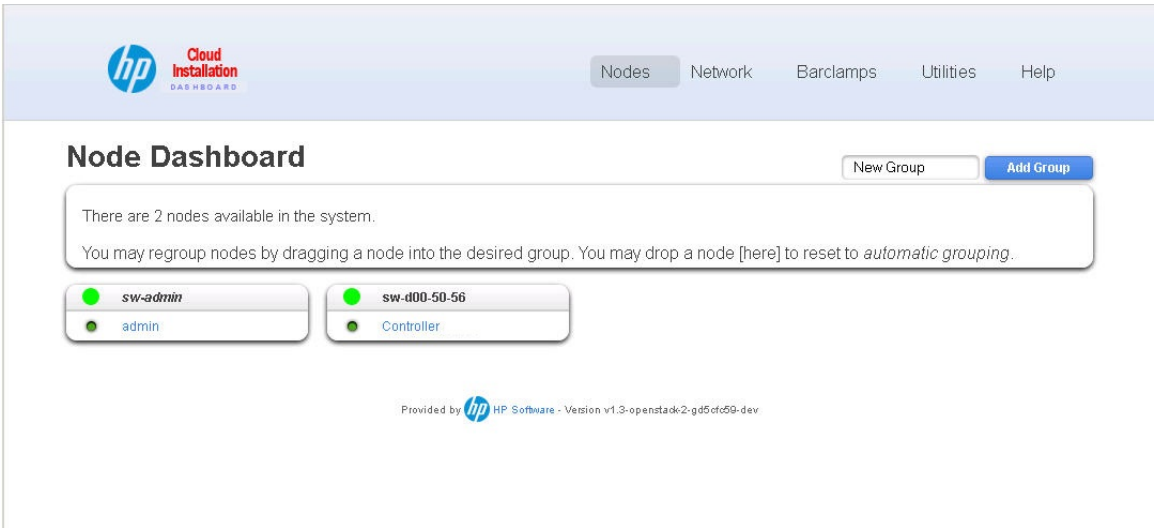

Deploy the Compute Region

Deploying the Compute Region involves bringing up Compute node, ensuring it is discovered and allocated, and applying Nova proposal to both Controller and Compute nodes. Follow the steps below:

Discover and Allocate the Compute Node

1. The Compute node needs to be powered ON.
2. When the Compute node boots up, the Admin node automatically installs the OS using a PXE boot. The **Node Dashboard** displays the

node as the OS is installed.

a. The status indicator is initially grey.
b. The node initially displays in the **Node Dashboard** with a generated, MAC address name.
c. Tip: view the node's console to monitor the system as it is being installed.

3. Wait until the **Node Dashboard** shows the new node with a flashing yellow status.

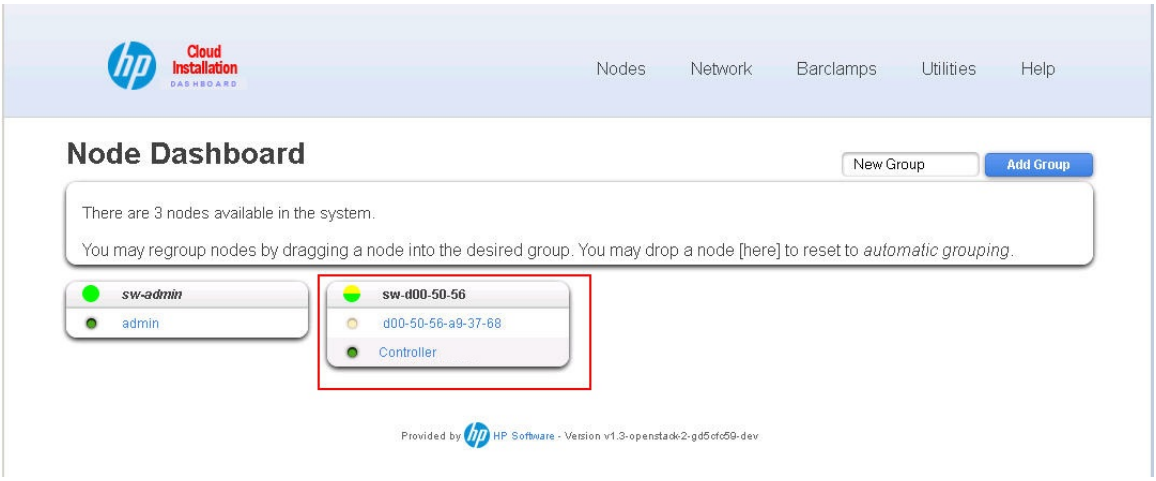

4. Tip: To confirm that the MAC Address name is associated to the node, view the node's console. The hostname displayed in the node's console will match the MAC Address name displayed in the Node Dashboard.

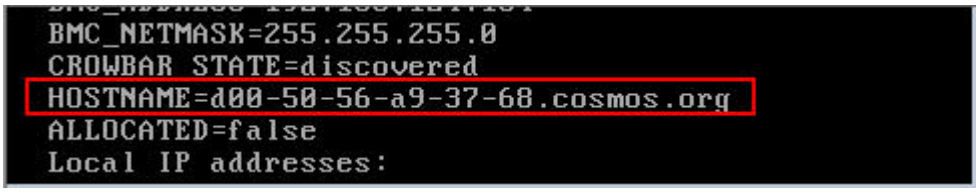

5. On the Cloud Installation Dashboard, select Nodes > Bulk Edit.

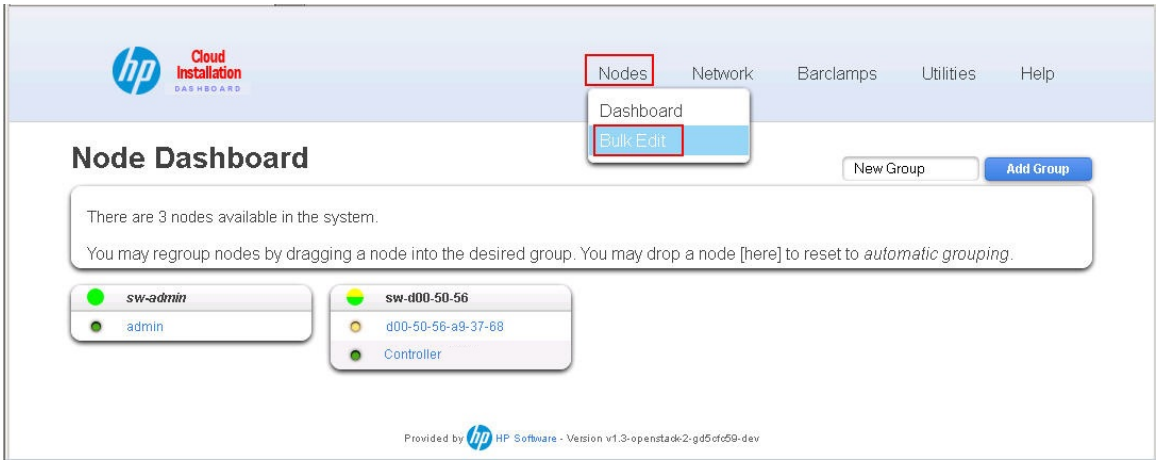

6. Click the **Allocate?** checkbox.

7. Click **Save**.
8. On the Cloud Installation Dashboard, select **Nodes** > **Dashboard**.

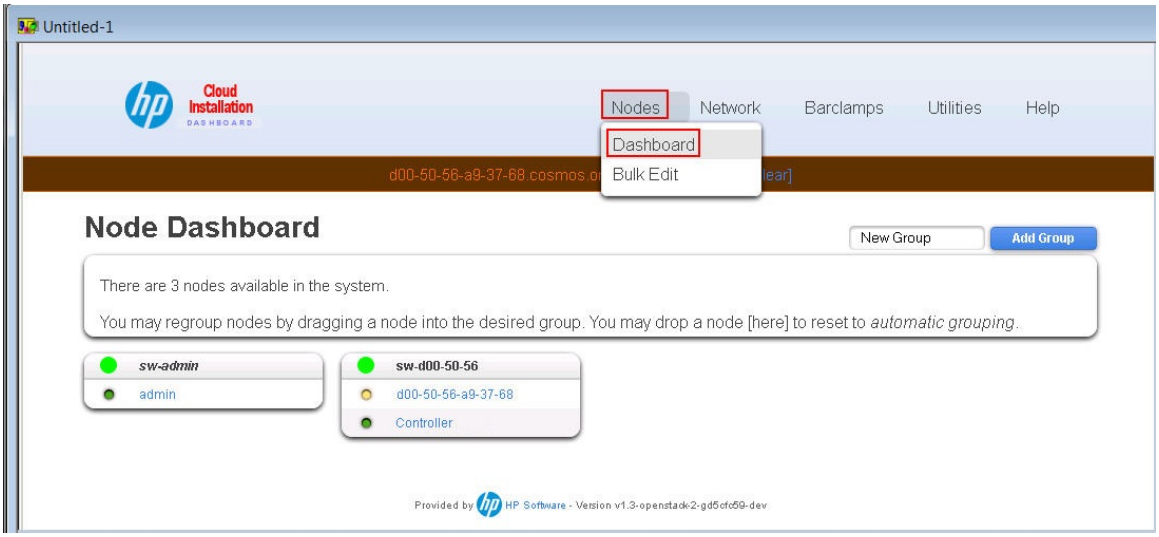

9. Wait for the node's status to change from yellow to green. View the node's console for installation progress.
10. If a dialog asking to load missing firmware displays, select NO.
11. When the node's install has completed, the *Node Dashboard *displays the node in the same group as the Admin node and displays a green status.

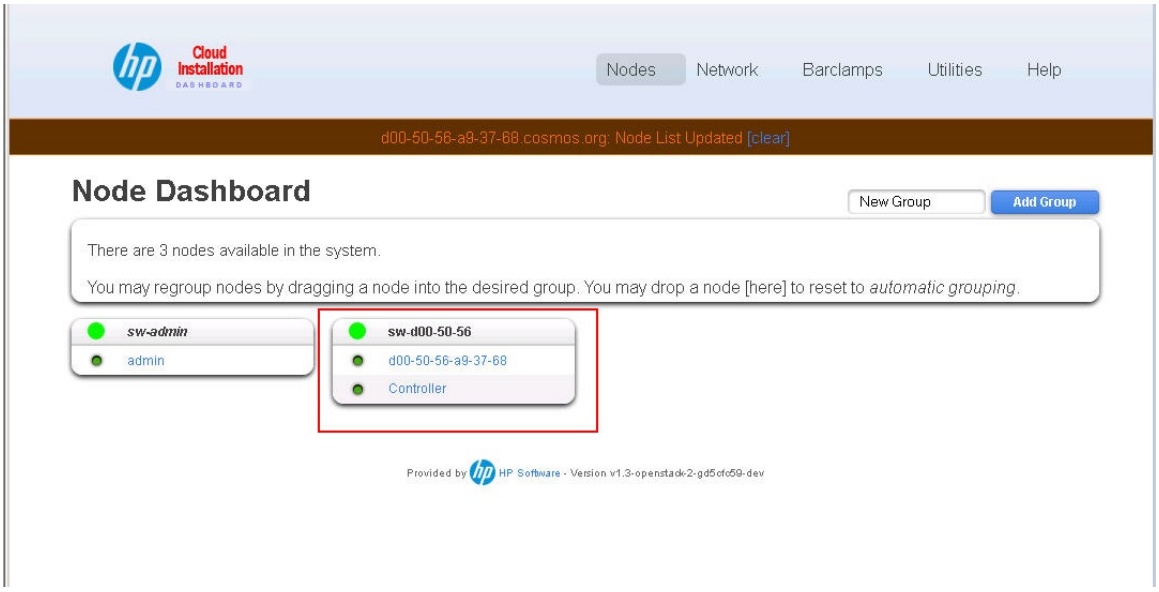

12. Tip: the node's login prompt displays the MAC Address as hostname displayed in the}**Node Dashboard**.

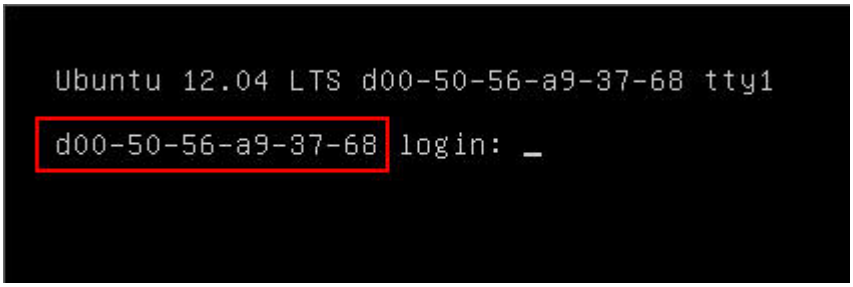

Rename the Compute Node

In general, it is helpful to name the nodes installed by Cloud Installation Dashboard something meaningful to their function.

Follow these steps to rename the Compute node:

1. Select the **node's MAC Address name**.

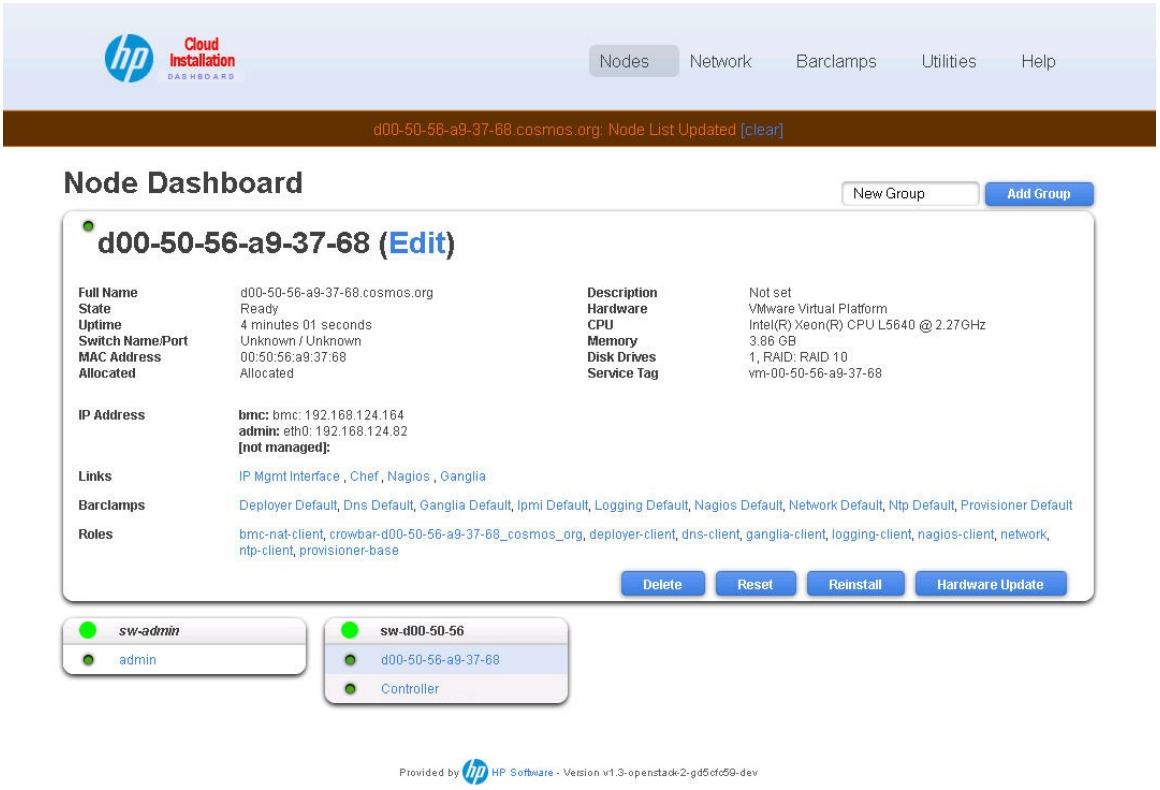

2. Click **Edit**.

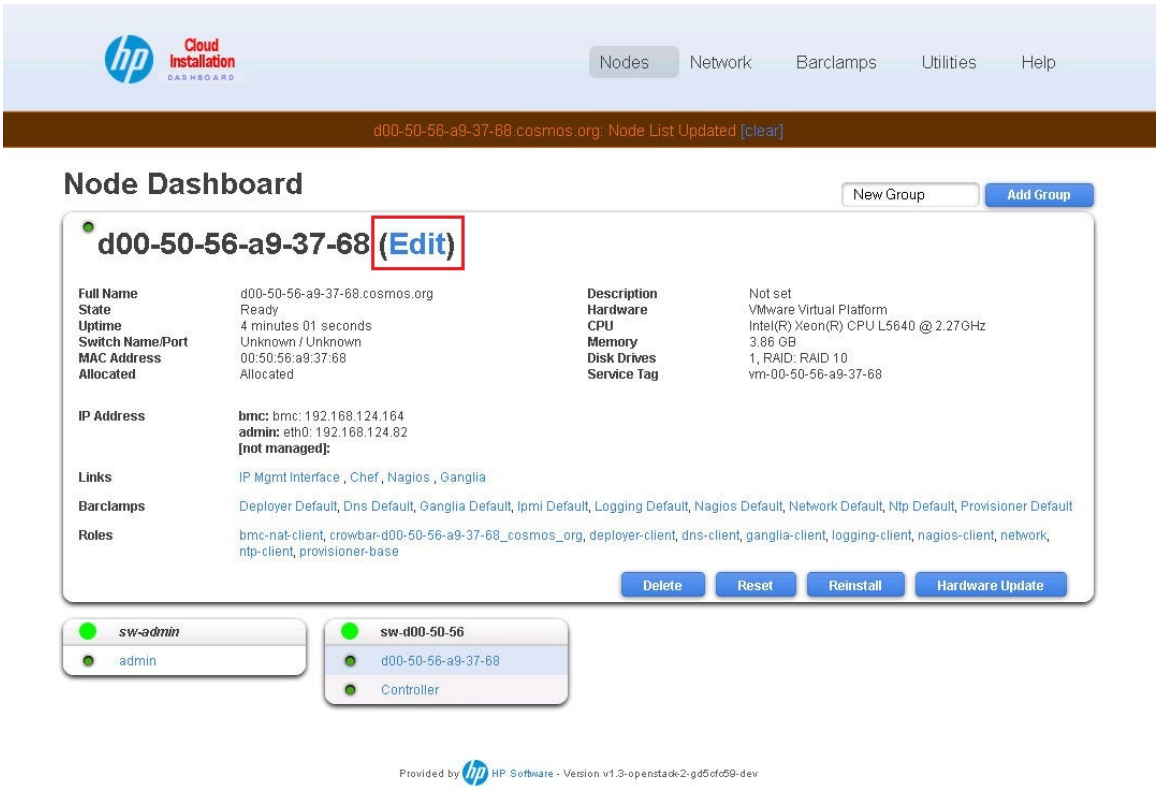

3. Change the alias field from the MAC Address name to a meaningful name.
4. Click **Save**.
5. Notice the node's new name is reflected in the display.

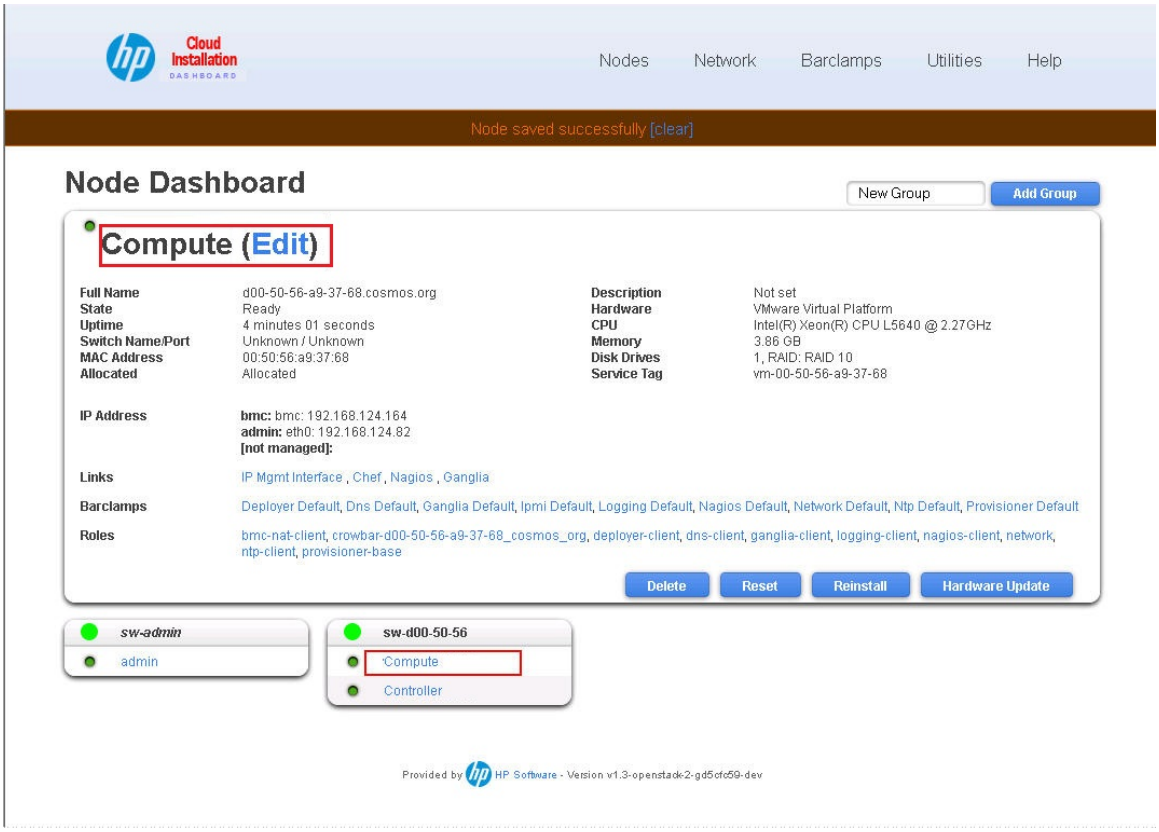

6. Notice the **Full Name** of the node remains unchanged.

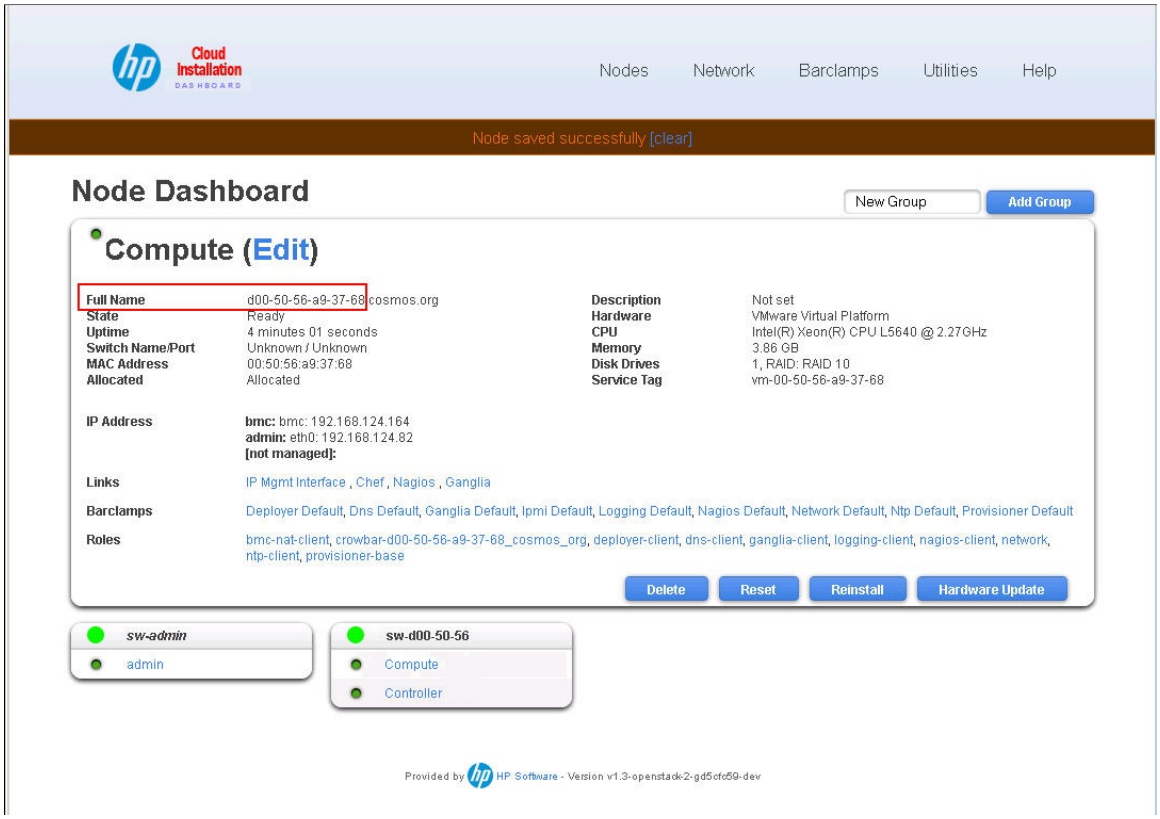

Apply Nova proposal to Compute Region

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

Use the following steps to apply barclamps:

1. From the node's detail page, select **Barclamps > Cloud Infrastructure**.

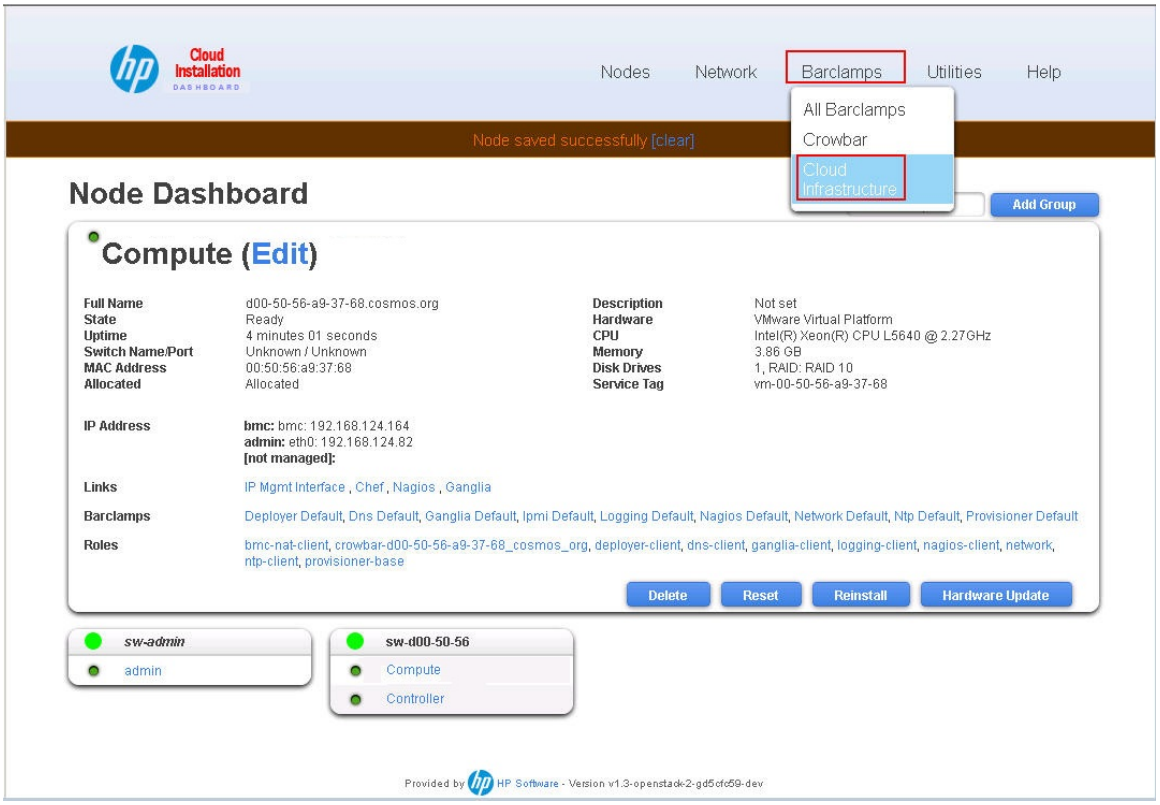

2. The barclamps specific to Cloud Infrastructure display.

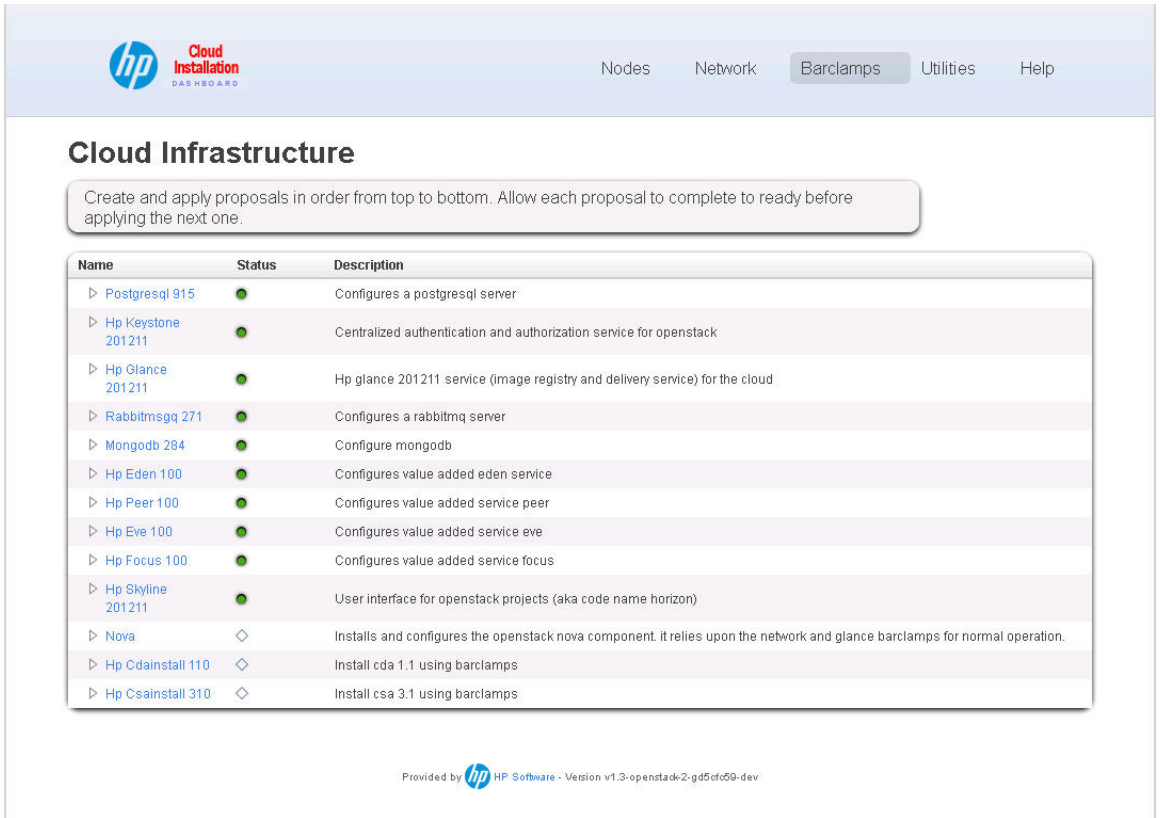

3. Take special note that the majority of the barclamps have been deployed, as indicated by the green status. Nova is the next barclamp to deploy. It must be deployed on both the Controller and Compute nodes.
4. For the Nova barclamp, perform the following steps:
    a. Select the **barclamp's name**.

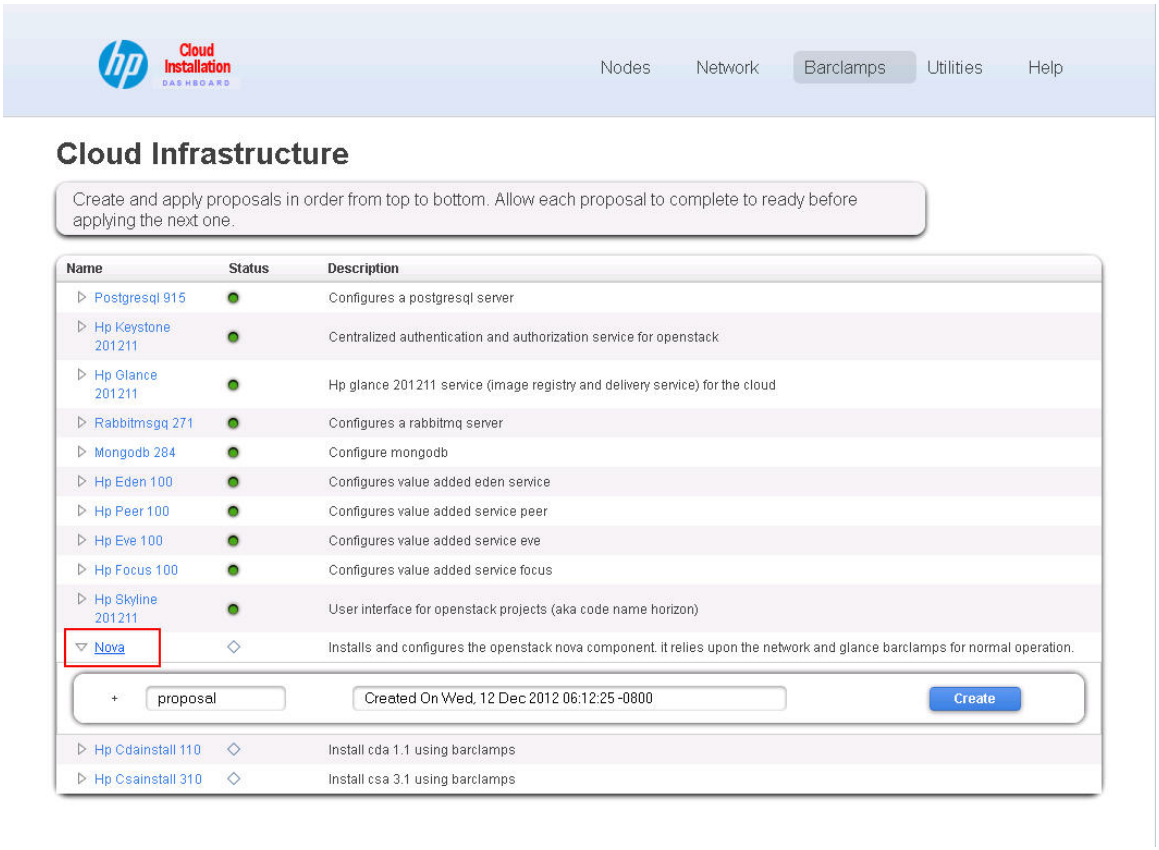

b. Specify the **Compute Region's name** as the proposal.

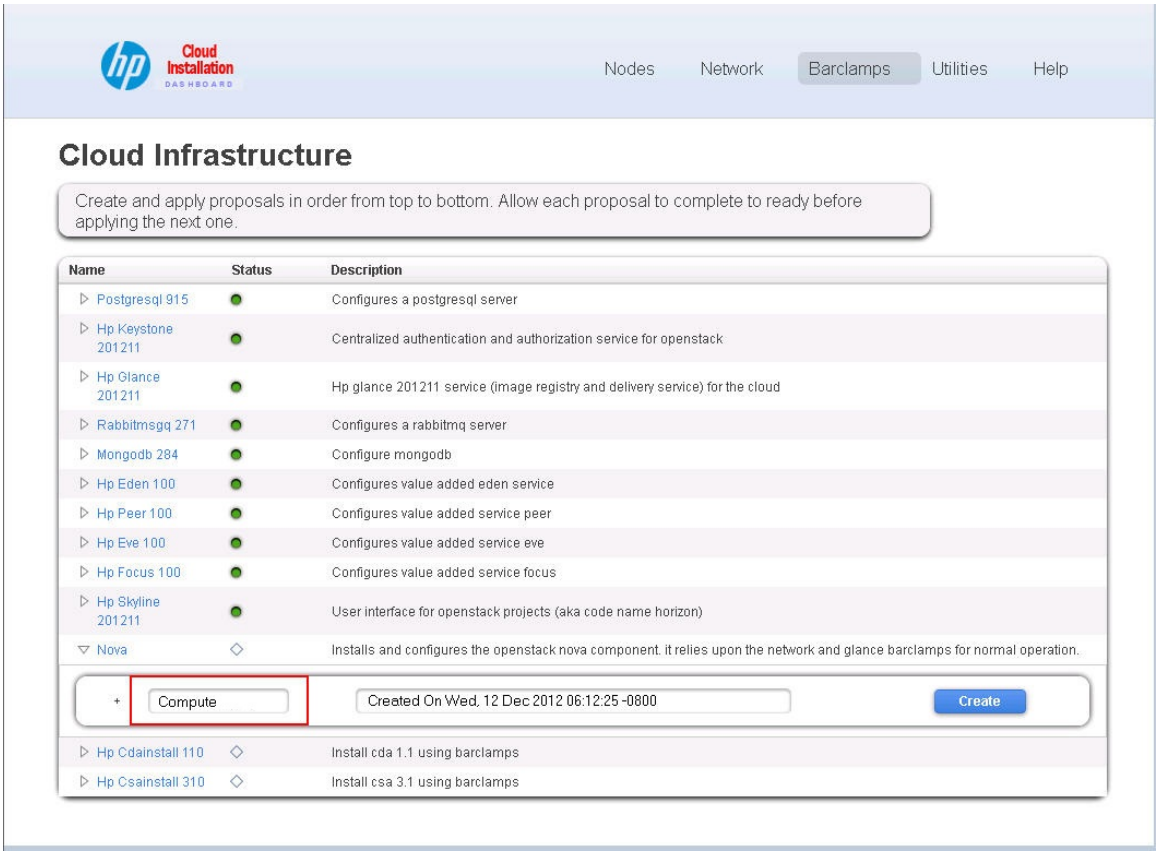

c. Click **Create**.

d. Verify that the default values are correct for the proposal. Some Nova values need adjusting.
    i. Sample nova proposal part 1. Use the default values
    ii. Nova-Volume role by default is hosted by the Controller. If it has an additional block device designated for Nova-Volume, set the **Type of Volume** value to **raw**.
    iii. If not, and it has sufficient space on the /root volume, set the **Type of Volume** value to **local**.
    iv. In case of local, set the **Maximum File Size** value to a size reasonable for the environment. Dedicate sufficient space for image repository. Typically, dedicate 80% of the space to nova-volume. A nova-volume host persistent virtual disk of VM instances launched by OpenStack.
    v. Sample nova proposal part 2. The following roles must be set:
        1. At the bottom of the proposal, the items on the left are **Available Nodes**; the items on the right are the **Roles**.
        2. Ensure the Compute node is set for only the **Nova-multi-compute** role.
        3. Ensure the Controller node is set for Both *Nova-multi-controller and Nova-multi-volume roles.
        4. Select the node name's link in the **Available Nodes** list to drag it to the appropriate role (if not already done).

        5. Selecting the node's link icon and dragging it to a role will not work.

        6. Selecting the node's delete icon removes the association of the node to a role.
e. Click **Apply**. Click **OK**.
f. Wait for the proposal status (displayed at the top left of the dialog box) to turn green.

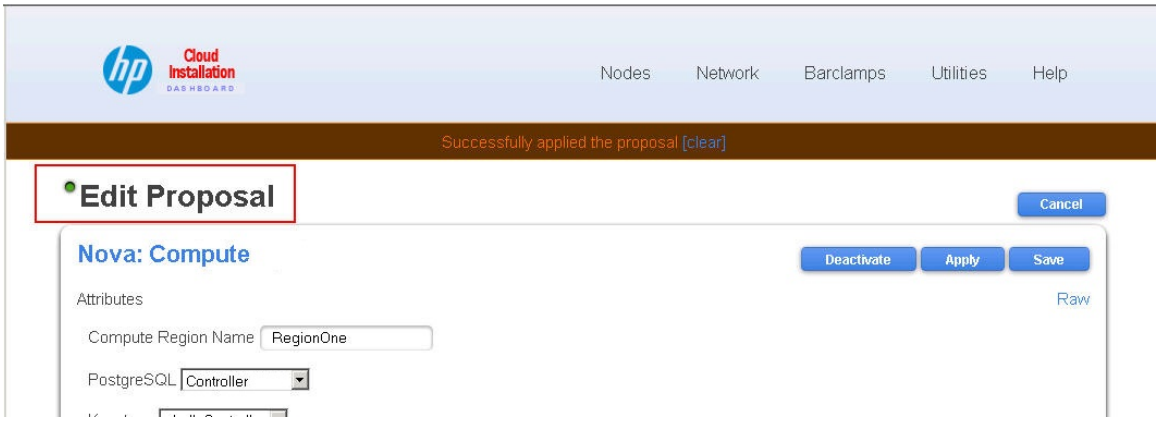

The proposal status can also be seen within the barclamp list.

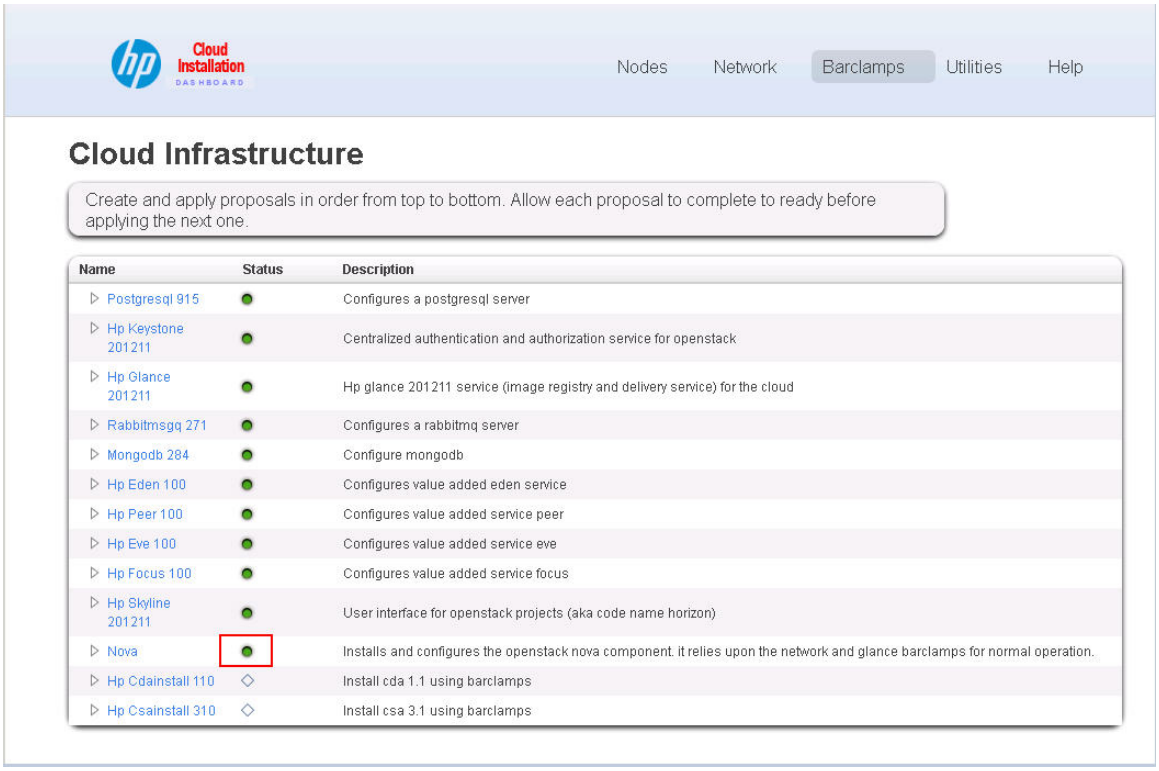

5. Re-select **Barclamps > Cloud Infrastructure** to view the barclamp list again.

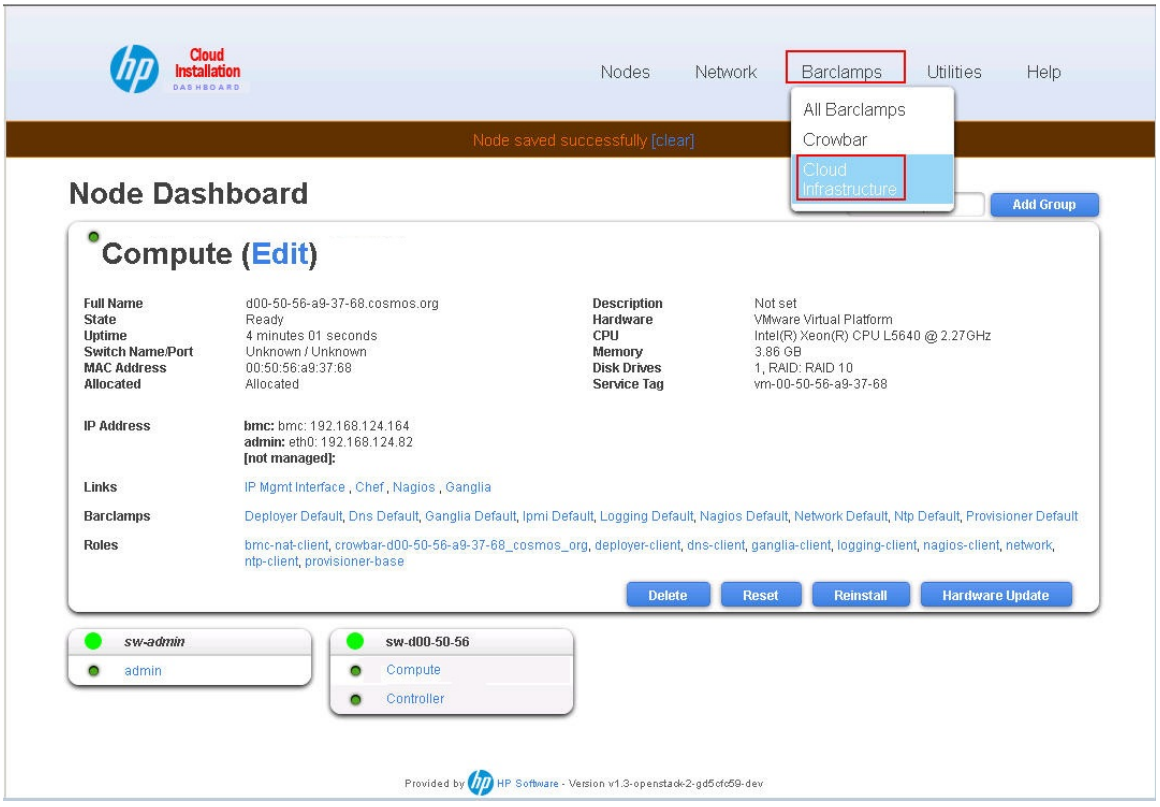

NOTE: If you wish to add another Compute Node later, follow the steps of discover and allocate the new node and add the node to an existing nova proposal. Do not create a separate proposal for the newly added node.

Next Step

All Cloud Infrastructure barclamps are now applied to the Domain Controller and the Compute node(s). The next step is configuring Networking (Nova Fixed and Floating Network) for VM instances launched by the Cloud Infrastructure. Refer to the Configure Networking Infrastructure for Virtual Machines in the Cloud section.

Configure Networking Infrastructure for Virtual Machines in the Cloud

Overview

Once the Cloud Infrastructure is deployed, the network infrastructure for the Cloud needs to be configured.

Referring back to the Deployment Diagram I, your networking environment should consist of:

- eth0 - private 1 as administration network
- eth1 - public/corporate network
- eth2 - private 2 network available for configuring Fixed IP networks

Setting a Fixed Network

1. Access the Cloud Administration Dashboard URL http://<MACAddressHostnameOfDomainController>; provide Credentials **Admin/secretword**
2. Go to Region tab > Networks.
3. Click **Create Fixed IP**.
4. Set **Network Label** to the desired name
5. Leave **Project** not selected, enabling the fixed network to be used across the Cloud. This setting indicates that the network is a flat DHCP fixed IP network.
    a. Set **Project** to a specific project the fixed network needs to be separate for the project. This requires VLAN configuration on the eth2 port as well as the connecting switch.
6. Set **IPv4 CIDR** to **192.168.123.0/24**
    a. A Class B CIDR may be provided to accommodate more VMs.
7. Set **Gateway** to **192.168.123.1**
8. Set **DNS1** to **192.168.124.10**
9. **Set Multi-Host** to **True.**
10. Set **Number of Networks** to **1**
11. Set **Network Size** to **256**
12. **Set Bridge ID** to **br100** (or any custom name).
13. Set **Bridge Interface** to **eth2** (needs to be unique if you are adding multiple fixed networks).

NOTE: 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 Controller and Compute nodes.

For example, the following command creates a VLAN interface on eth2 with tag 500:

- sudo vconfig add eth2 500

In the above example, the following fields will have additional information:

1. Set **VlanID** to **500**
2. Set **BridgeID** to **br500**
3. Set **Bridge** to **eth2.500**

Setting a Floating Network

1. Access the Cloud Administration Dashboard URL http://<MACAddressHostnameOfDomainController>; provide Credentials **Admin/secretword**
2. Go to Region tab > Networks.
3. Click **Create Floating IP**.
4. Set **IPv4 CIDR** to the static IP range applicable to the corporate/public network.
5. Set **Interface** to the interface that is configured for public/corporate access. In this example, it is **eth1**.

Post Deployment and Validation

Overview

The Cloud Infrastructure installation and configuration is complete. The post-deployment tasks outlined below validate the Compute and Topology implementation of the Cloud Infrastructure.

Setting an Image Repository

An image repository contains necessary images to be provisioned into the Cloud as part of a Topology design. Initial images needed are base operating system images.
For example, download the following operating system images:

- CirrOS Image - 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

Follow these steps to upload these images to Cloud Infrastructure:

1. Copy these images to the Admin node's **/tftpboot** folder. This folder requires read/write permission.
2. Log into the Cloud Administration Dashboard, using default credentials **Admin/secretword**.
3. Navigate to Domain > Images and click **Create Image**.

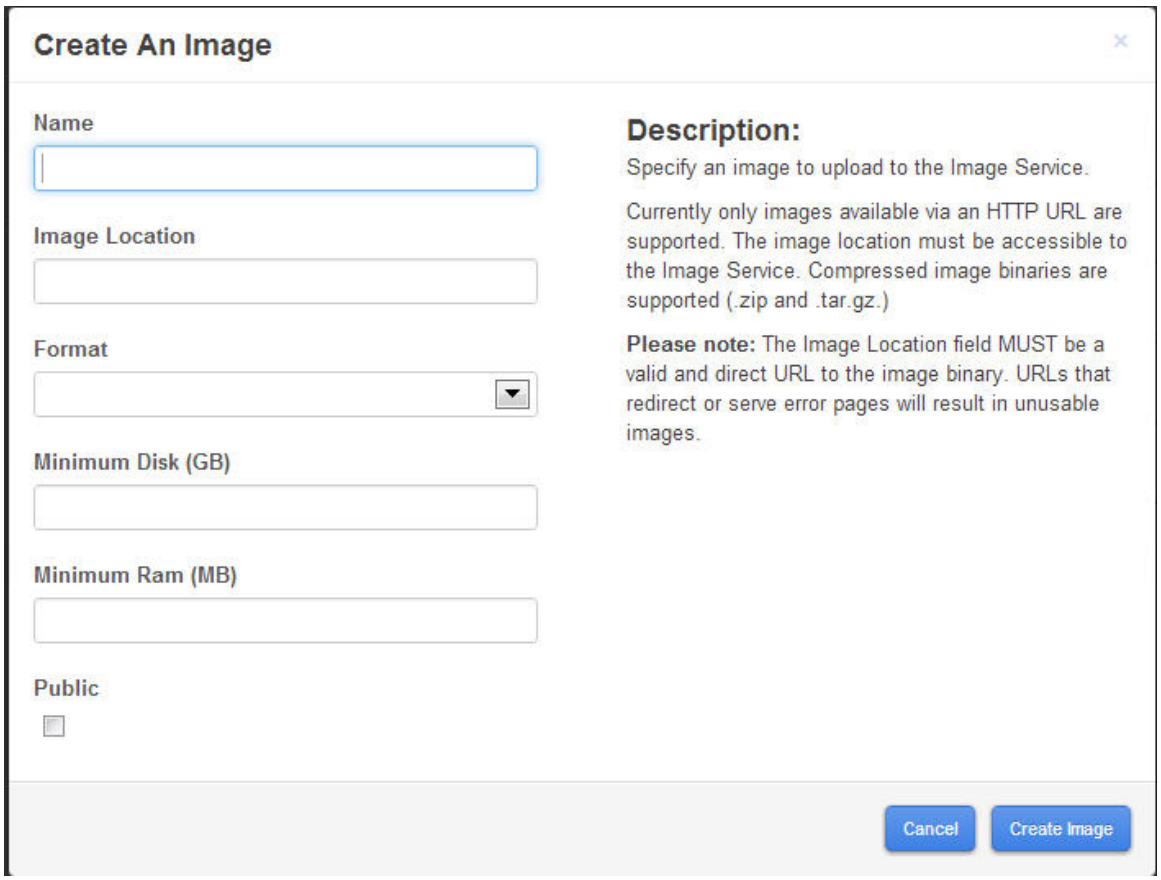

4. Set **Name** to the desired image name.
5. Set **Image Location** to the HTTP URL of the Admin node's **/tftpboot** folder, including the image name.
    a. For example: http://192.168.124.10:8091/precise-server-cloudimg-amd64-disk1.img
6. Set **Format** to **QCOW2**
7. Set **Minimum Disk** to **0**
8. Set **Minimum RAM** to **0**
9. Set **Public** to *Yes
10. Click **Create Image**.
11. From the Images table, click the image's name to validate the image's details and size once loaded.

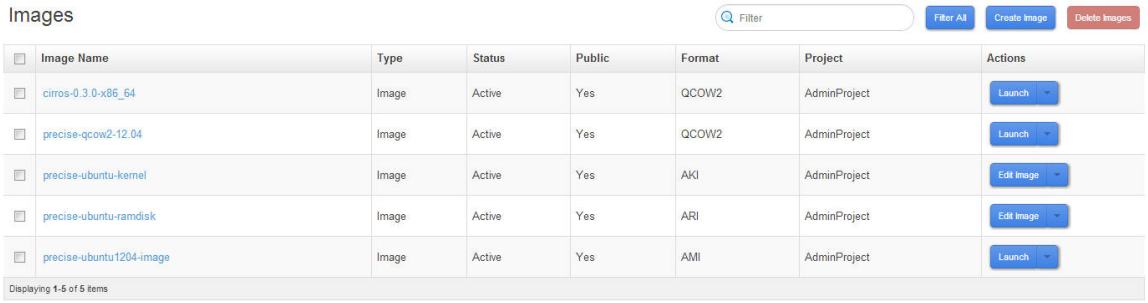

Creating a Keypair

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

To create:

1. Navigate through Project > Manage Compute > Access & Security and click **Create Keypair**.
2. Provide the Keypair Name (use 'testonlykeypair' or any custom name).
3. Download the keypair to a safe location.

Configuring the Default Security Group

To validate the VM accessibility, the VM needs to be on a network segment (security group) that has both "ping" and SSH ports open. To to this, edit the default security group:

1. Navigate through Project > Manage Compute > Access & Security
2. Under Security Groups, click **Edit Rules** on the default security group.
3. Under Add Rule, specify the following values:
    a. For SSH, IP Protocol: TCP; From Port: 22; To Port:22; Source Group: CIDR; CIDR: 0.0.0.0/0
    b. For PING, IP Protocol: ICMP; From Port: -1; To Port:-1; Source Group: CIDR; CIDR: 0.0.0.0/0
4. Click **Update Security Group Rules** and add the next rule.

Launching a VM Instance

Next, validate whether you can launch a VM instance using the newly uploaded image:

1. Navigate through Project > Manage Compute > Images
2. Click **Launch** on the newly uploaded image name
3. Provide following values:
    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 Access and Security tab within the Launch Instance window and select the newly created keypair and ensure that 'default' has a check mark under Security Groups.
5. Click **Launch**.

Validating VM Accessibility

Validate whether the VM was launched successfully:

1. Navigate through Project > Manage Compute > Instances.
2. Locate the recently launched VM instance and click on the Instance Name.
3. From the Overview tab, record the IP Address assigned to this instance.
4. Switch to the Log tab and ensure the following has been logged:

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

5. Validate whether network infrastructure has been set up correctly to access the recently launched VM. Using the IP address of the

instance from the previous section, try to ping and ssh into it.

6. For Ping, do the following:
    a. Ensure you are on the same network as your VM Instance or have the proper routing enabled.
    b. Open up the command prompt, and try pinging the IP Address.
7. For SSH, do the following:
    a. Ensure you are on the same network as your VM Instance or have the proper routing enabled.
    b. Convert the file downloaded during keypair creation from .pem to .ppk format. You can use a tool like PuTTYgen to convert.
    c. 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
    d. If you are already in an SSH window of one of the Compute nodes, enter the following command to validate that the 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:~$
```

Creating Resource Pool

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

1. Navigate through Project > Manage Topology > Resource Pools.
2. Click **Create Resource Pool**.
3. Provide the following values:
    a. Cloud Type: HP Cloud Infrastructure - OpenStack (Essex) (select from the drop-down)
    b. Compute Region: Domain/RegionOne (select from the drop-down)
    c. Name: <ProvideCustomName> or AdminProjectRP
    d. Type: OpenStack
    e. Version: Essex 2012.1.x
    f. Region ID: RegionOne

    g. Domain URL: 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**.

Creating Infrastructure Topology

To create an Infrastructure Topology, follow these steps:

1. Navigate through Project > Manage Topology > Documents
2. Click **Create Infrastructure Topology**.
3. Drag a Group and Network and connect them.
4. Click **Save** and provide a custom name 'TestPurposeTopo1.'
5. Validate whether the newly created Infrastructure Topology document is listed under **Documents**.

Creating Infrastructure Design

Every Infrastructure Topology document needs to bind with an Infrastructure Design document. To create anInfrastructure Design document, follow these steps:

1. Navigate through Project > Manage Topology > Documents
2. Click **Create Infrastructure Design**.
3. In the Resource Binding section, provide the following values:
    a. Name: a customer name or 'TestPurposeDesign1'
    b. Infrastructure Topology: <Select from the drop-down> or 'TestPurposeTopo1'
    c. ResourcePool: <Select from the drop-down> or 'AdminProjectRP'
4. Click **Next**, The Binding Details wizard displays.
5. Provide the following values under the Server Group section:
    a. Minimum Instances: 1
    b. Instance Name Prefix: test
    c. Machine Flavor: m1.tiny or <Select a Flavor from the drop-down>
    d. Machine Image: precise-server-cloudimg
    e. Key Pair Name: testonlykeypair
6. Provide the following values under Network Segment section:
   Name: TestNetwork1
   Type: Cloud Edge Gateway <This will provide both Fixed and Floating IP to the VM instance>
   Description: Test Purpose Only
   Open Port List: 22,80
7. Click **Create Infrastructure Design**.
8. Validate whether the newly created Infrastructure Design document is listed under **Documents**.

Launching an Infrastructure Design Document

With both Design and Topology documents created, you will be able launch a Design document:

1. Navigate through Project > Manage Topology > Documents
2. Locate the newly created Infrastructure Design document and click **Launch**.
3. Provide a custom name 'TestInfrastrctureName1' and click **Launch**.

Validating the Infrastructure Document

Validation of the Infrastructure document is a multi-step process.

1. Validate the status of launched infrastructure in **Topologies**.
    a. Navigate through Project > Manage Topology > Topologies.
    b. Locate and check the **State** of the newly launched Infrastructure Document.
    c. Click the Topology Name to see details (specifically the Job tab). The State of a successfully provisioned document is **Succeeded**.
    d. Go to the Content tab within Topology Detail to identify the IP Address allocated to the newly provisioned VM (search for 'ip_address_type').
2. Validate the accessibility through SSH:
    a. You must have both Fixed and Floating IPs associated to the newly provisioned VM from the previous section. Try SSH through both of the IPs (details above).
    b. You can also locate this instance of the VM by navigating through Project > Manage Compute > Instances. Click the instance name to see more details (details above).

Deploy CDA (Optional)

Overview

CDA can be deployed on the Domain Controller if it meets the sizing requirements. Refer to the CDA Installation Guide for more details. However, if it needs to be on separate node, follow the Deploy CDA Node section below. Or, follow these instructions for applying the CDA proposal:

1. Select the **Compute Regions** tab.
2. Click **Deploy Compute Region**.
3. Enter crowbar/crowbar for the credentials.
4. The **Node Dashboard** displays. The Admin and Controller nodes display a green online status.

Deploy CDA Node

Discover and Allocate CDA Node

1. The CDA node needs to be powered ON.
2. When the CDA node boots up, the Admin node automatically installs the OS using a PXE boot. The **Node Dashboard** displays the node as the OS is installed. The status indicator is initially grey.
    a. The CDA node initially displays in the **Node Dashboard** with a generated, hexadecimal name.
    b. Tip: view the node's console to monitor the system as it is being installed.
3. Wait until the **Node Dashboard** shows the CDA node with a flashing yellow status.
    a. Tip: To confirm that the hexadecimal name is associated to the CDA node, view the CDA node's console. The hostname displayed in the CDA node's console will match the hexadecimal name displayed in the **Node Dashboard**.
4. On the Cloud Installation Dashboard, select **Nodes** > **Bulk Edit**.
5. Click the **Allocate?** checkbox.
6. Click **Save**.
7. On the Cloud Installation Dashboard, select **Nodes** > **Dashboard**.
8. Wait for the CDA node's status to change from yellow to green. View the CDA node's console for installation progress.
    a. If a dialog box displays a message asking you to load missing firmware, select NO.
9. When the CDA node's install has completed, the **Node Dashboard** displays the CDA node.
    a. Tip: the CDA node's login prompt displays the hexadecimal name displayed in the **Node Dashboard**.

Rename CDA Node

Follow these steps to rename the CDA node:

1. Select the CDA node's **hexadecimal name**.
2. Click **Edit**.
3. Change the alias field from the hexadecimal name to a meaningful name.
4. Click **Save**. Notice the CDA node's new name is reflected in the display and the **Full Name** of the CDA node remains unchanged.

Apply CDA Proposal

Use the following steps to apply barclamps:

1. From the CDA node's detail page, select **Barclamps > Cloud Infrastructure**.
2. The barclamps specific to Cloud Infrastructure display.
3. For the CDA barclamp, perform the following steps:
    a. Select the **barclamp's name**.
    b. Specify the **CDA node's name** as the proposal.
    c. Click **Create**.
    d. Ensure the default values are correct for the proposal (see screenshot below). The values must have:
    PostgreSQL Instance: <ProposalThroughWhichPostgresApplied>
    Admin Email Address: <an email address>
    FQDN for Hostname: <MACAddressHostName>.cosmos.org
    Domain Name: cosmos.org
    Passphrase for LWSSO:<AnyParaPhraseKey>
    e. Click **Apply**. Click **OK**.
    f. Wait for the proposal status (displayed at the top left of the dialog box) to turn green.

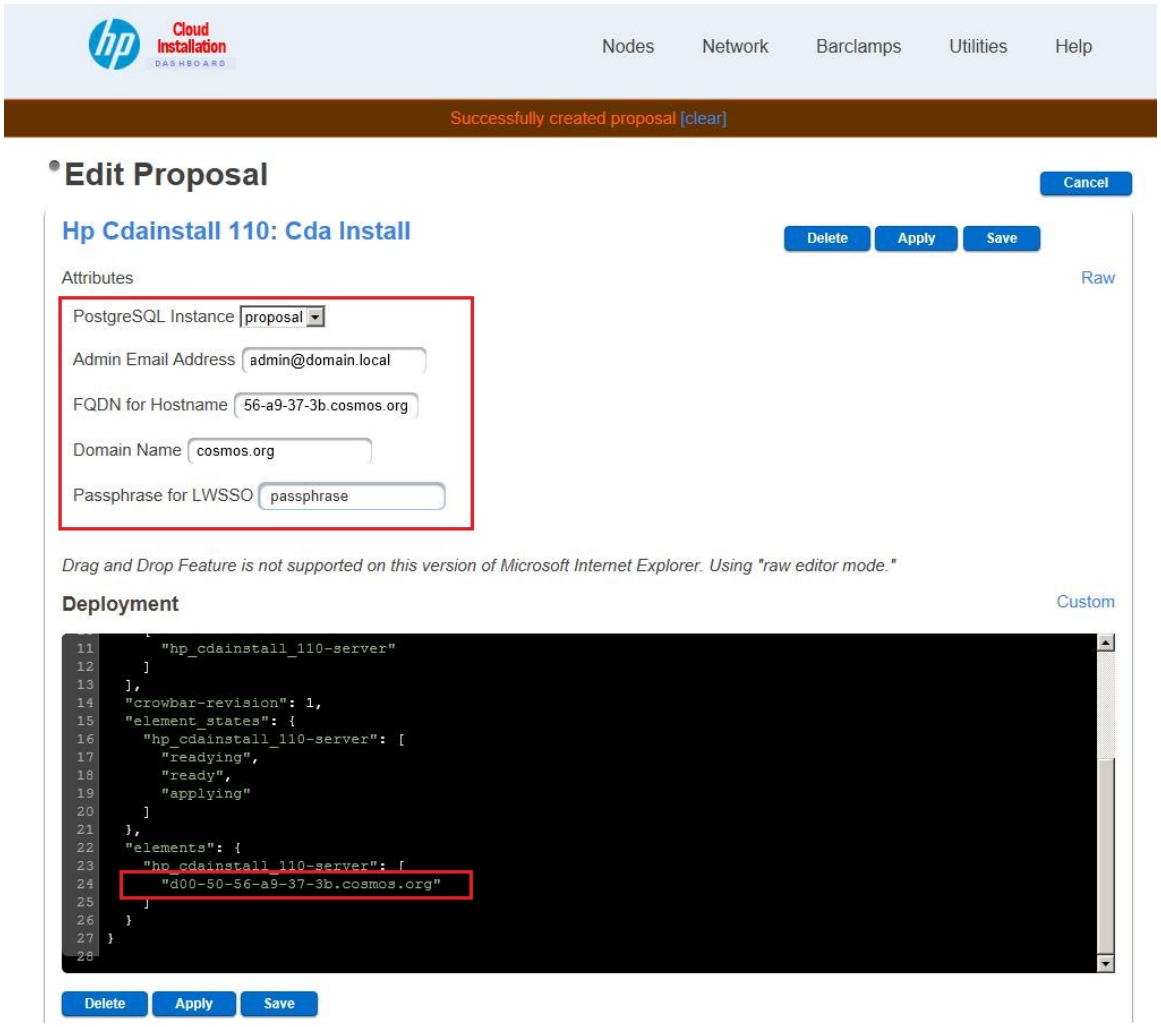

Accessing CDA Application

Now that the CDA application is configured, you can access it by launching a browser and connecting to https://<node name>:8443/cda.

For example: https://d00-50-56-a9-37-3b.cosmos.org:8443/cda

Default username/password = admin/changeit

You can refer to the HP CDA 1.1 User Guide for details on how to use CDA with this local cloud installation.

Next Step

If you plan to deploy CSA on the same node as CDA, refer CDA 1.1 Installation and Configuration Guide, Appendix C: Configuring HP CDA and HP CSA to Run on the Same Server.

Go to Deploy CSA (Optional).

Deploy CSA (Optional)

Overview

CSA can be deployed on the Domain Controller if it meets the sizing requirements. Refer to the CSA guide for more sizing details. However, if it needs to be on separate node, follow Deploy CSA Node section. Or, follow these instructions for Applying CSA Proposal:

1. Select the **Compute Regions** tab.
2. Click the **Deploy Compute Region** button.
3. Enter crowbar/crowbar for the credentials.
4. The **Node Dashboard** displays and the Admin and Controller nodes display a green online status.

Deploy CSA Node

Discover and Allocate CSA Node

1. The CSA node needs to be powered ON.
2. When the CSA node boots up, the Admin node automatically installs the OS using a PXE boot. The **Node Dashboard** displays the node as the OS is installed. The status indicator is initially grey.
    a. The CSA node initially displays in the **Node Dashboard** with a generated, hexadecimal name.
    b. Tip: view the node's console to monitor the system as it is being installed.
3. Wait until the **Node Dashboard** shows the CSA node with a flashing yellow status.
    a. Tip: To confirm that the hexadecimal name is associated to the CSA node, view the CSA node's console. The hostname displayed in the CSA node's console will match the hexadecimal name displayed in the **Node Dashboard**.
4. On the Cloud Installation Dashboard, select **Nodes** > **Bulk Edit**.
5. Click the **Allocate?** checkbox.
6. Click **Save**.
7. On the Cloud Installation Dashboard, select **Nodes** > **Dashboard**.
8. Wait for the CSA node's status to change from yellow to green. View the CSA node's console for installation progress.
    a. If a dialog box displays a message asking you to load missing firmware, select NO.
9. When the CSA node's install has completed, the **Node Dashboard** displays the CSA node.
    a. Tip: the CSA node's login prompt displays the hexadecimal name displayed in the **Node Dashboard**.

Rename CSA Node

Follow these steps to rename the CSA node:

1. Select the CSA node's **hexadecimal name**.
2. Click **Edit**.
3. Change the alias field from the hexadecimal name to a meaningful name.
4. Click **Save**. The CSA node's new name is reflected in the display, and the **Full Name** of the CSA node remains unchanged.

Apply Barclamps to the CSA Node

Use the following steps to apply barclamps:

1. From the CSA node's detail page, select **Barclamps > Cloud Infrastructure**.
2. The barclamps specific to Cloud Infrastructure display.
3. For the CSA barclamp, do the following steps:
    a. Select the **barclamp's name**.
    b. Specify the **CSA node's name** as the proposal.
    c. Click **Create**.
    d. Verify that the default values are correct for the proposal (see screenshot below). You must set the FQDN information correctly
    e. hostname: <MACAddressName>.cosmos.org
    f. Click **Apply**. Click **OK**.
    g. Wait for the proposal status (displayed at the top left of the dialog) to turn green.

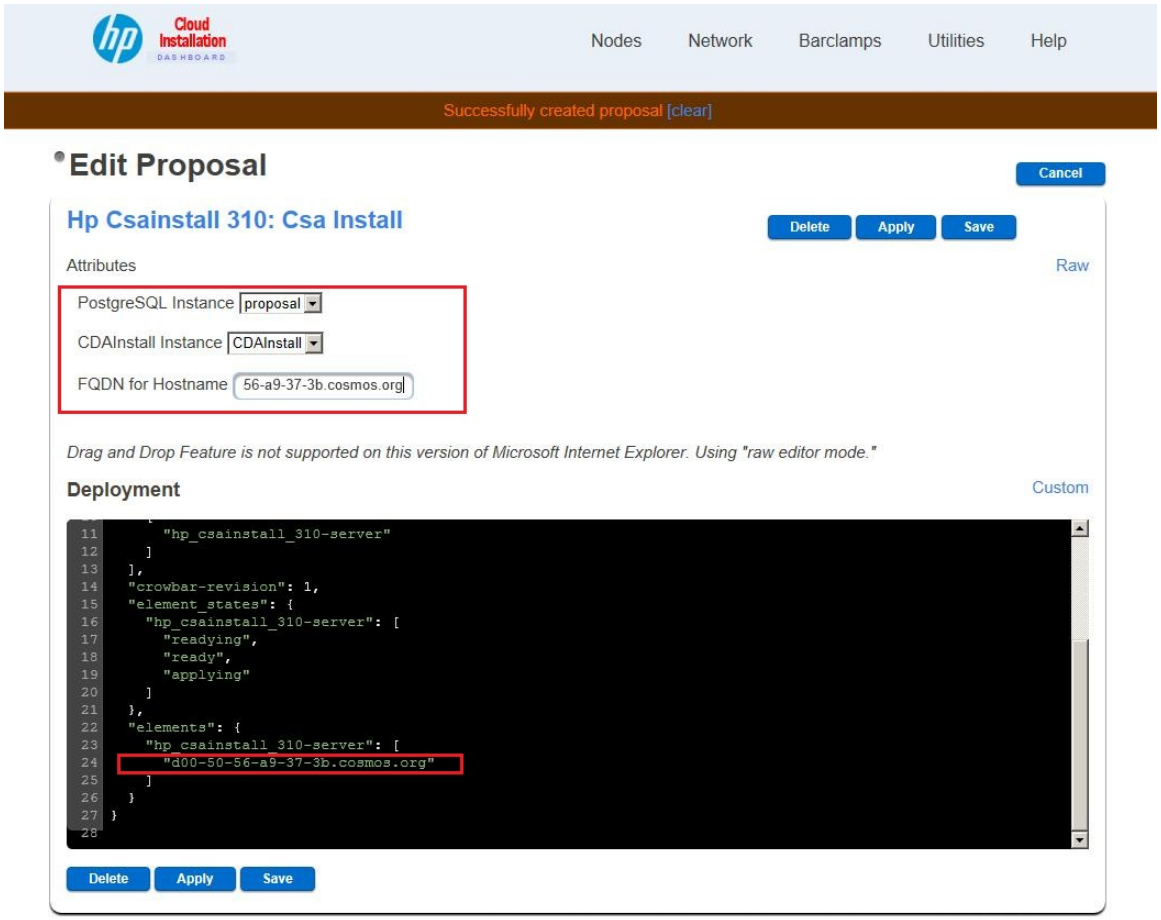

Configure CDA and CSA Integration

To synchronize CDA and CSA:

- Make sure the machines running CDA and CSA are able to connect to each other on the network and that name resolution is working.
- On the node running CSA, log in using the default user name/password.
- Switch to root user by running "sudo -i".
- Enter ./tmp/postinstall.sh and then press <Enter>.
- The postinstall.sh script contains:

  #!/bin/sh
  export JAVA_HOME=/usr
  export PATH=$PATH:$JAVA_HOME/bin
  java -jar /opt/hp/csa/3.10/Tools/ProcessDefinitionTool/process-defn-tool.jar -d /var/hp/csa/conf/postgresdb.properties -p /var/hp/csa/conf/cda.xml
  ---
- This script will synchronize CSA and CDA and configure them for use.

Accessing CSA Application

Now that the CSA application is configured, you can access it by launching a browser and connecting to https://<node name>:8444/csa for admin portal and https://<node name>:8444/csp for consumer portal.

For example: https://d00-50-56-a9-37-3b.cosmos.org:8444/csa
For example: https://d00-50-56-a9-37-3b.cosmos.org:8444/csp

Default Admin Portal username/password = admin/cloud
Default Consumer Portal username/password = consumer/cloud

You can refer to HP CSA 3.10 User Guide for details on how to use CSA with CDA.

Troubleshooting

Overview

Problem: An error message displays when configuring the Cloud prerequisites

Symptoms	Completing the prerequisites fails.
Primary software component	Cloud Infrastructure Dashboard
Failure message	**Error:** There was an error submitting the form. Please try again.
Probable cause	• Ensure the Cloud Administration node has both a public and private network configured. • Ensure the proxy information is correct and references a working proxy server and port.

Solution

- On the Cloud Administration node, view the log file, /var/log/apache2/error.log, for details.
- Verify a public network is configured in the Cloud Administration node.
- Redo 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: A Cloud Infrastructure barclamp proposal fails

Symptoms	An error message displays and the status turns red if a proposal fails.
Primary software component	Cloud Infrastructure 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
- On the Cloud Administration node:
  - View the log file, /opt/dell/crowbar_framework/log/production.log and <MacAddressHostname>.chef.log, for details.
  - View the log file, /var/log/apache2/error.log, for details.
  - View the folder, /var/tmp/cosmos, for installer settings.

Problem: The Admin node displays the "not ready" (grey) state

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

Solution

1. Reboot the Admin node.
2. Whenever a new Admin node is installed, all the previous participating nodes must be kept off of PXE booted through the new admin node. If you end up with this issue, a reinstall of the Admin node and its participating node is suggested.

Problem: The Controller or Compute node displays the "not ready" (grey) state

Symptoms	Dashboard is unable to monitor the Controller or Compute nodes.
Primary software component	Cloud Installation Administration Node
Failure message	
Probable cause	The Controller or Compute nodes have not updated the live status to Admin node.

Solution

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

Problem: When creating a new pxe node, the pxe boot fails with a TFTP timeout error

Symptoms	Admin node takes a lot of time to reboot.
Primary software component	Cloud Installation Administration Node
Failure message	
Probable cause	This has been seen in a few cases after Admin node is rebooted.

Solution

1. Log into the Admin node and run the following commands and then PXE boot the nodes:
    a. sudo bluepill tftpd stop
    b. sudo bludpill tftpd start