



## **OpenStack 1.1 for Cisco DFA Install Guide for Using Pre-Built OpenStack for Cisco DFA Images**

**First Published:** April 10, 2014

### **Americas Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 527-0883

Text Part Number: OL-32000-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2014 Cisco Systems, Inc. All rights reserved.



## CONTENTS

---

### Preface

#### Preface v

Audience v

Document Conventions v

Related Documentation for Cisco DFA vi

Obtaining Documentation and Submitting a Service Request vii

---

### CHAPTER 1

#### Hardware and Software Requirements 1

Requirements for Hardware and Software 1

---

### CHAPTER 2

#### Prerequisite 3

Prerequisite for OpenStack for DFA 3

---

### CHAPTER 3

#### OpenStack for DFA Overview 5

Overview of OpenStack for DFA 6

---

### CHAPTER 4

#### Before Starting 9

Before Starting 9

---

### CHAPTER 5

#### Servers/UCS Topology and OpenStack for DFA Installation 11

UCS Resources Required 11

Servers/UCS Topology 12

Installation 12

Restore Target UCS Host Names (for all nodes) 14

Reboot the UCS OpenStack Nodes (for all nodes) 15

Generate Site.pp 15

Invoke Scripts to Restore Nodes (All Nodes) 16

Invoke restore\_openstack on Build Server Node 16

Invoke restore_openstack on Control and Compute Nodes	16
Confirm Keystone Started	17
Reboot the Nodes	17
Images for Launching a VM	17
Configuration in the Leaf or DCNM	18
Leaf Interface Configuration and Global Configuration	18
DHCP Configuration in DCNM	18

---

**CHAPTER 6****Create Project and Launch VM 21**

Create Project and Launch VM	21
Known Caveats	21
Steps to Create a Project	21
Steps to Create a User for the Project	22
Steps to Create the Network	22
Steps to Launch the VM	23

---

**CHAPTER 7****Scalability 25**

Scalability	25
-------------	----

---

**CHAPTER 8****Known Limitations and Caveats 27**

Known Limitations and Caveats	27
-------------------------------	----

---

**CHAPTER 9****Technical Support Model 29**

Technical Support Model	29
-------------------------	----



## Preface

---

The Preface contains the following sections:

- [Audience, page v](#)
- [Document Conventions, page v](#)
- [Related Documentation for Cisco DFA, page vi](#)
- [Obtaining Documentation and Submitting a Service Request, page vii](#)

## Audience

This publication is for experienced network administrators who configure and maintain Cisco Dynamic Fabric Automation.

## Document Conventions

Command descriptions use the following conventions:

Convention	Description
<b>bold</b>	Bold text indicates the commands and keywords that you enter literally as shown.
<i>Italic</i>	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x   y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x   y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.

Convention	Description
[x {y   z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
<i>variable</i>	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use the following conventions:

Convention	Description
<code>screen font</code>	Terminal sessions and information the switch displays are in screen font.
<b>boldface screen font</b>	Information you must enter is in boldface screen font.
<i>italic screen font</i>	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:



#### Note

Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.



#### Caution

Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

## Related Documentation for Cisco DFA

The Cisco Dynamic Fabric Automation documentation is at the following URL: [http://www.cisco.com/c/en/us/solutions/data-center-virtualization/unified-fabric/dynamic\\_fabric\\_automation.html#~Products](http://www.cisco.com/c/en/us/solutions/data-center-virtualization/unified-fabric/dynamic_fabric_automation.html#~Products).

The Cisco Nexus 6000 Series documentation is at the following URL: <http://www.cisco.com/c/en/us/support/switches/nexus-6000-series-switches/tsd-products-support-series-home.html>.

The Cisco Nexus 7000 Series documentation is at the following URL: <http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/tsd-products-support-series-home.html>.

The Cisco Nexus 5500 Series documentation is at the following URL: <http://www.cisco.com/c/en/us/support/switches/nexus-5000-series-switches/tsd-products-support-series-home.html>.

The Cisco Nexus 1000V switch for VMware vSphere documentation is at the following URL: [http://www.cisco.com/en/US/products/ps9902/tsd\\_products\\_support\\_series\\_home.html](http://www.cisco.com/en/US/products/ps9902/tsd_products_support_series_home.html). The documentation therein includes the following guides for Cisco DFA. Additional information pertaining to troubleshooting can be located in the Cisco Nexus 1000V documentation for Cisco NX-OS Release 4.2(1)SV2(2.2).

- *Cisco Nexus 1000V DFA Configuration Guide, Release 4.2(1)SV2(2.2)*
- *Cisco Nexus 1000V VDP Configuration Guide, Release 4.2(1)SV2(2.2)*

The Cisco Prime Data Center Network Manager (DCNM) documentation is at the following URL: [http://www.cisco.com/en/US/products/ps9369/tsd\\_products\\_support\\_series\\_home.html](http://www.cisco.com/en/US/products/ps9369/tsd_products_support_series_home.html). The Cisco Prime DCNM documentation for Cisco DFA includes but is not limited to the following guides:

- *Cisco DCNM 7.0 OVA Installation Guide.*
- *Cisco DCNM 7.0 Fundamentals Guide*
- *Cisco DCNM DFA REST 7.0 API Guide*

The Cisco Prime Network Services Controller (NSC) documentation is at the following URL: [http://www.cisco.com/en/US/products/ps13213/tsd\\_products\\_support\\_series\\_home.html](http://www.cisco.com/en/US/products/ps13213/tsd_products_support_series_home.html).

The OpenStack for Cisco DFA install documentation includes the following guide and documents:

- *Open Source Used In OpenStack for Cisco DFA 1.0* at the following URL: [http://www.cisco.com/en/US/docs/switches/datacenter/dfa/openstack/opensource/OpenStack\\_for\\_Cisco\\_DFA\\_1.0\\_Open\\_Source\\_Documentation.pdf](http://www.cisco.com/en/US/docs/switches/datacenter/dfa/openstack/opensource/OpenStack_for_Cisco_DFA_1.0_Open_Source_Documentation.pdf)
- *OpenStack for Cisco DFA Install Guide Using Cisco OpenStack Installer* at the following URL: <http://www.cisco.com/en/US/docs/switches/datacenter/dfa/openstack/install/guide/os-dfa-coi.pdf>
- *OpenStack for Cisco DFA Install Guide for Using Pre-built OpenStack for Cisco DFA Images* at the following URL: <http://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/dfa/openstack/install/guide/prebld-image.pdf>
- *Quick Guide to Clonezilla* at the following URL: <http://www.cisco.com/en/US/docs/switches/datacenter/dfa/openstack/install/guide/clonezilla-image-restore.pdf>

## Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation*, at: <http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>.

Subscribe to *What's New in Cisco Product Documentation*, which lists all new and revised Cisco technical documentation, as an RSS feed and deliver content directly to your desktop using a reader application. The RSS feeds are a free service.





# Hardware and Software Requirements

---

- [Requirements for Hardware and Software, page 1](#)

## Requirements for Hardware and Software

OpenStack for DFA software is currently packaged with COI (Cisco OpenStack Installer). The Cisco OpenStack Installer is qualified on:

- Ubuntu 12.04 LTS serves as a base operating system.
- KVM serves as the hypervisor.
- Cisco UCS C-Series serve as physical compute/storage hardware. The required BIOS version is C220M3.1.5h.0 (Build Date: 07/10/2013) .
- OpenStack for DFA software is based on open source software Grizzly release with Cisco DFA plugin developed at Cisco.

We have developed an alternative mechanism that allows a user to install OpenStack for DFA software from pre-built/cloned images used for build server, OpenStack controller and compute nodes in the COI installation and deployment model. The target UCS needs to have a minimum of 500.1G hard disk space.





## Prerequisite

---

- [Prerequisite for OpenStack for DFA, page 3](#)

## Prerequisite for OpenStack for DFA

As a prerequisite to run OpenStack for DFA, please follow *Dynamic Fabric Automation Deployment Guideline* document to bring up DFA system first, including Nexus switches and DCNM software.





## OpenStack for DFA Overview

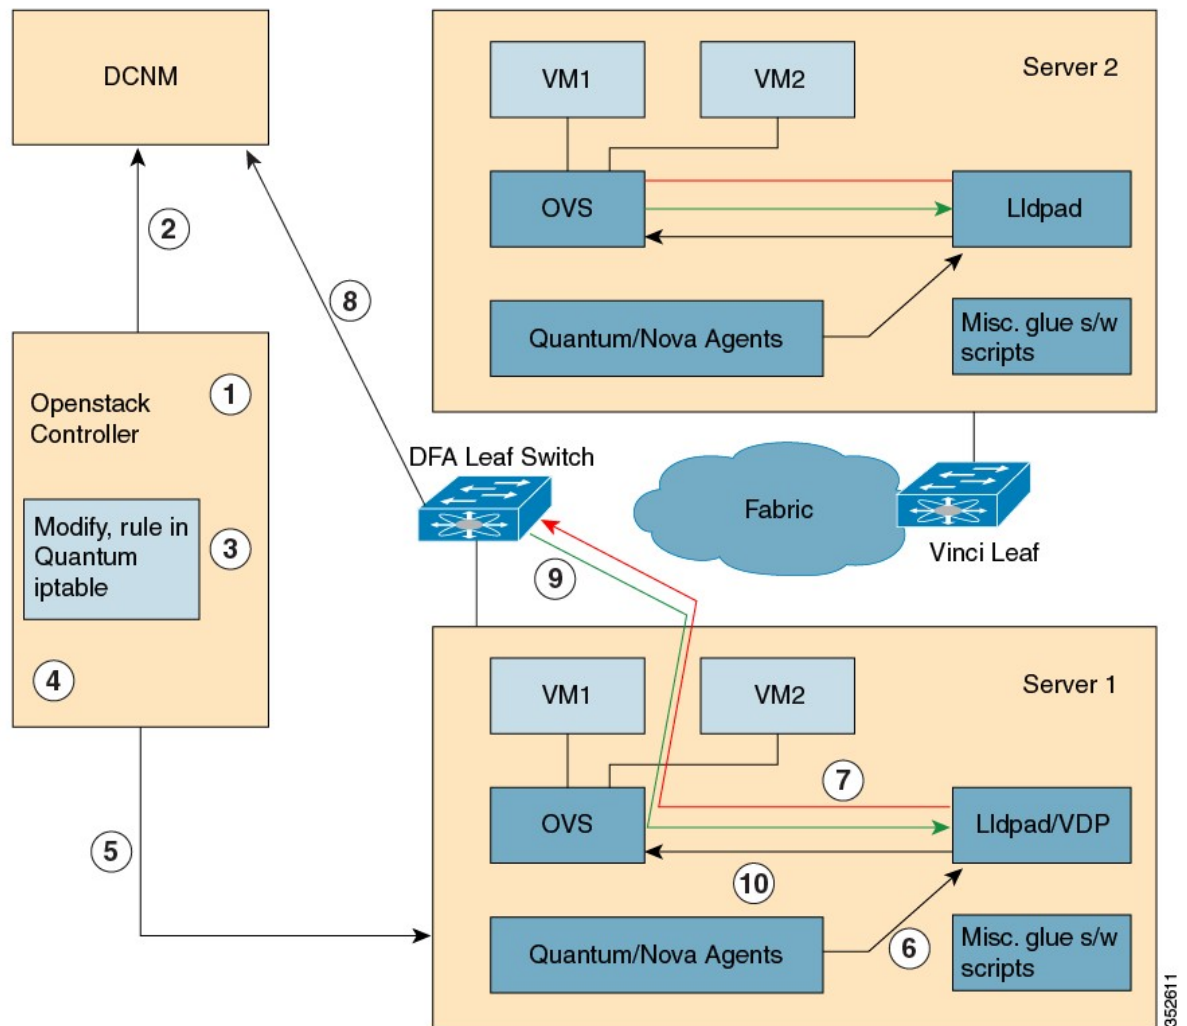
---

- [Overview of OpenStack for DFA, page 6](#)

# Overview of OpenStack for DFA

OpenStack serves as one of orchestrators of the cloud enabled through DFA. For this release, all the orchestration is done using OpenStack's dashboard Horizon graphic user interface.

**Figure 1: OpenStack for DFA System Flow**



## Note

Steps 2, 4, and 6 can be disabled, along with other external DFA modules, so OpenStack will continue to function natively, without DFA.

The above diagram provides an overview of the system in the DFA context. The control flow can be summarized as follows:

- 1 A network of type "tunnel" with a tunnel ID range that is greater than 4000 is created. In OpenStack, the tunnel ID is a segment ID.
- 2 Network Information (Subnet/mask, tenant name) is sent to the Cisco Prime Data Center Network Manager (DCNM) through the Cisco Prime DCNM representational state transfer (REST) application programming interfaces (APIs).
- 3 A virtual machine (VM) instance is launched and the network for this instance is specified.
- 4 The host- specific IP table rules in the quantum agent are modified to unblock incoming Dynamic Host Configuration Protocol (DHCP) frames from the outside world and to add rules that allow only incoming frames with OpenStack-assigned source IP addresses.
- 5 Information about the network (subnet/mask, tenant name) and the VM is sent to the compute node.
- 6 Vision station interface discovery and provisioning (VDP) is notified about the VM and the segment ID for the VM.
- 7 VDP passes the information and segment ID for the VM to the Cisco DFA leaf switch.
- 8 The leaf switch contacts the Cisco Prime DCNM with the VM segment ID for retrieving the network attributes.
- 9 The leaf switch responds with the VLAN ID for tagging the traffic from the VM.
- 10 The Open vSwitch (OVS) is configured for tagging the packets from the VM to the network with the VLAN ID provided by the leaf switch. The virtual interface card (VNIC) is now operational.





## Before Starting

---

- We assume you have some general understanding of DFA (Dynamic Fabric Automation) system architecture.
- Please refer the Server/UCS Topology and OpenStack for DFA Installation section for prerequisites.
- Ensure CIMC connectivity is there.
- Ensure that DCNM is reachable via the Openstack Controller as well as from the Vinci cluster.
- [Before Starting, page 9](#)

## Before Starting

- We assume you have some general understanding of DFA (Dynamic Fabric Automation) system architecture.
- Please refer the Server/UCS Topology and OpenStack for DFA Installation section for prerequisites.
- Ensure CIMC connectivity is there.
- Ensure that DCNM is reachable via the Openstack Controller as well as from the Vinci cluster.





## Servers/UCS Topology and OpenStack for DFA Installation

---

- [UCS Resources Required, page 11](#)
- [Servers/UCS Topology, page 12](#)
- [Installation, page 12](#)
- [Images for Launching a VM, page 17](#)
- [Configuration in the Leaf or DCNM, page 18](#)

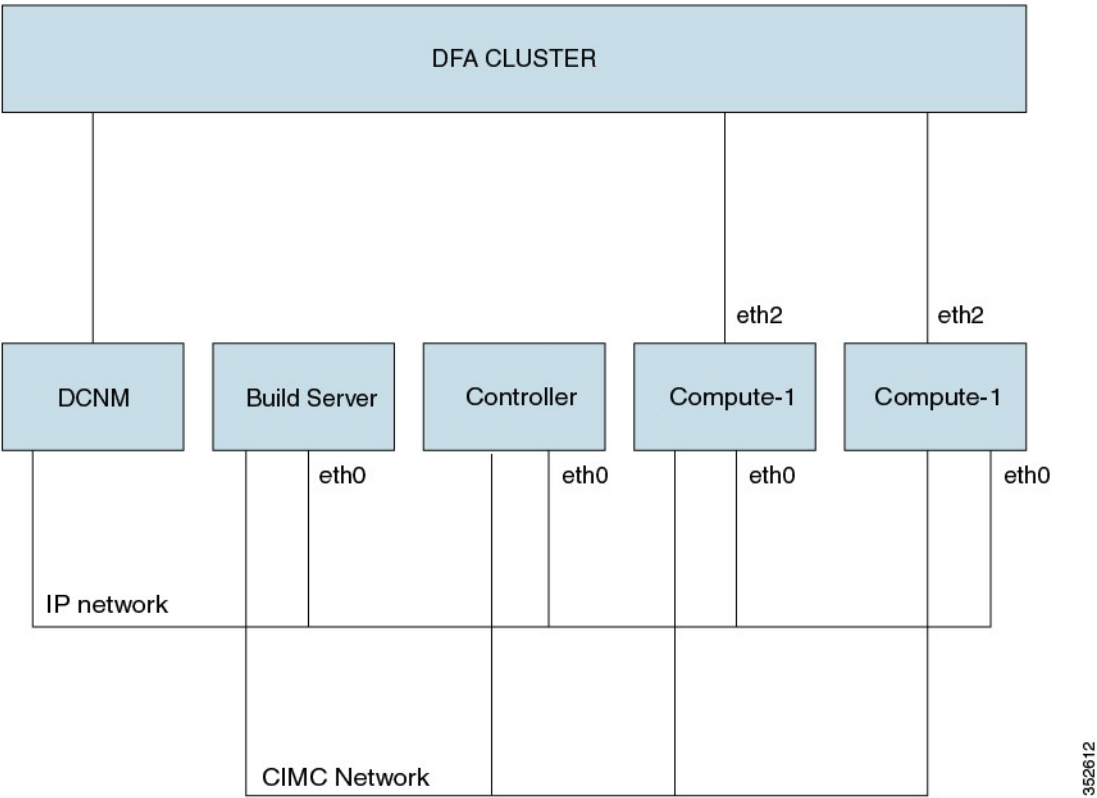
### UCS Resources Required

- One UCS is dedicated for the puppet build server.
- One UCS is dedicated for the OpenStack Controller where the dashboard Horizon is running for doing orchestration.
- As many UCS as the number of Computes are needed. VM's run in the Computes.
- Ensure that all the nodes are configured with CIMC and they are reachable.

# Servers/UCS Topology

A sample topology is shown below. This is a critical step as it lays down your foundation so it is highly recommended that users follow the same wiring scheme.

Figure 2: Sample Topology



Ensure that all the UCS’ are connected to each other. The sample topology has the build server, the controller and the computes connected through eth0. This is referred as the COI network.

The computes must have their uplink interface(s) connected to the DFA cluster.

All the CIMC ports of the OpenStack nodes including the build server must be connected to the CIMC network.

The DCNM also must have IP connectivity to the OpenStack control node. This is needed when verifying the OpenStack for DFA functionality after the installation is completed.

## Installation

The installation procedure described here includes the following major steps. But more precisely, you don’t specifically install COI-DFA. It is the pre-built/cloned images that already have a COI OpenStack for DFA installation and the procedure described here takes you to bring them up on your target UCS servers for OpenStack.

- 1 Copy the 3 cloned images and the Clonezilla utility (iso image) to a repository server (any Linux machine with about 80G free disk space and SSH daemon running) that can be reached by the target UCSs. The prebuilt/cloned images are at: [http://software.cisco.com/download/release.html?mdfid=281722751&flowid=36334&softwareid=282088134&release=7.0\(2\)&reind=AVAILABLE&rellifecycle=&retype=latest](http://software.cisco.com/download/release.html?mdfid=281722751&flowid=36334&softwareid=282088134&release=7.0(2)&reind=AVAILABLE&rellifecycle=&retype=latest)

The open source Clonezilla s/w for restoring the cloned/pre-built images is at:

<http://www.clonezilla.org/downloads.php> (we used clonezilla-live-20131125-saucy-amd64.iso)

Then extract the build\_node, control and compute node .tgz files using **tar -xvzf file name**.

To make the same files owner as the account owner, please use the *no-same-owner* option as shown below:

```
tar zxvf openstack_dfa_1_1_build_node.tgz --no-same-owner
```

"ucs-xx" are the original hostnames in the OpenStack setup where the images were cloned/prebuilt.

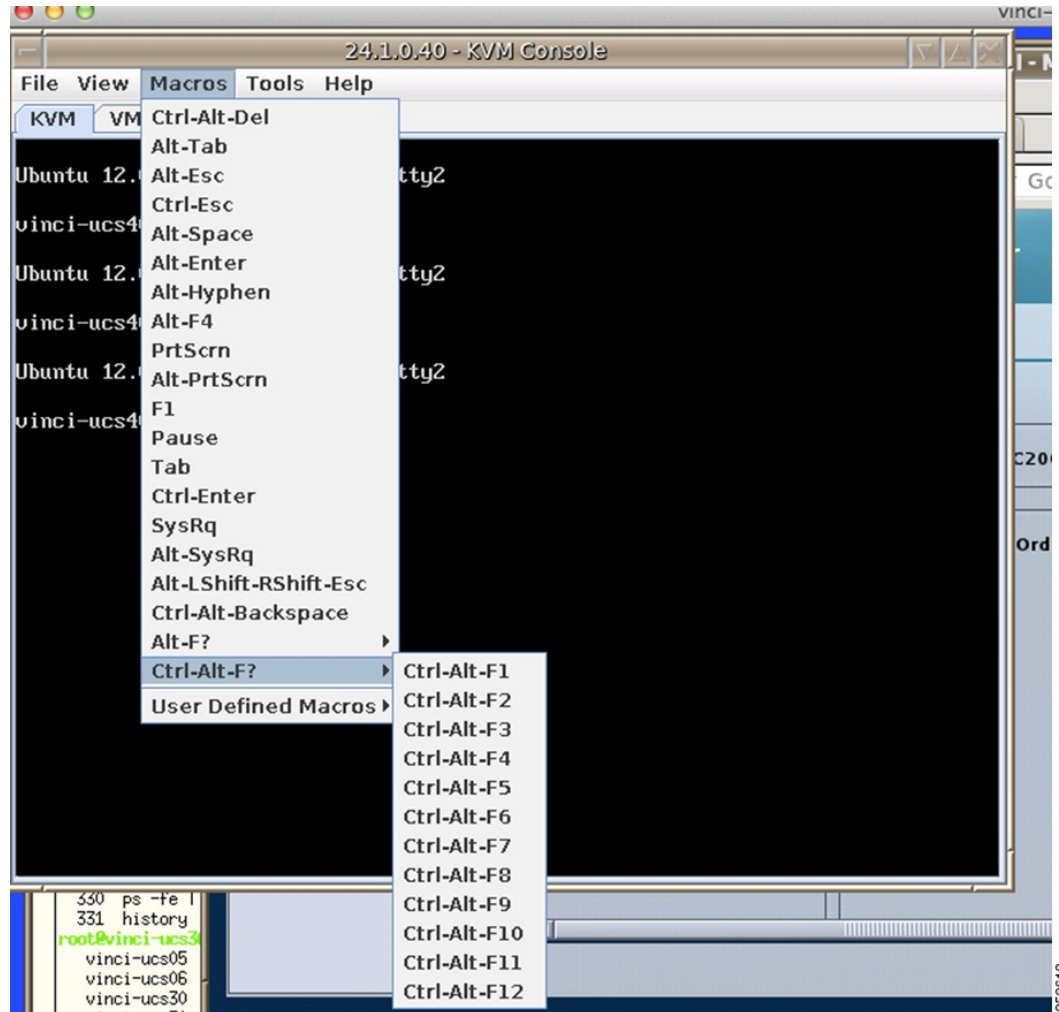
- 2 Please follow *Quick Guide to Clonezilla* for how to restore a cloned UCS image to a target UCS system. From the repository server above, launch the browser and remotely access to the CIMC port IP address (<http://<your UCS's CIMC port IP address>>) of your target UCS, then use Cisco Integrated Management Controller for UCS to proceed.

Assuming the UCS's are all wired correctly, their connectivity is good and the steps described below are followed strictly, restoring a cloned UCS system each will take about 15-20 minutes. So if you have two compute nodes to install, then the total number of target systems will be four and that can take about 1 hour if the restoring execution is done here flawlessly. But this is not really related to OpenStack installation and it is simply a cloning restoring process.

The following steps are done after the clonezilla restoration and after rebooting the nodes:

- 1 On the KVM console , if it is a graphical interface, hit CTRL-ALT-F2 to get the text login window:

**Figure 3: Text Login Window**



```
build/Control/compute login/password : localadmin/cisco123
```

- 2 Post-cloning COI/OpenStack, complete the instructions detailed in the following section.

## Restore Target UCS Host Names (for all nodes)

Edit /etc/hosts: On all nodes, add one entry per node in the following setup:

```
<ip> <hostname.domain.name> <hostname> <build ip address> <build hostname.domain.name> <build
hostname> <controller ip address> <controller hostname.domain.name> <controller hostname> <compute
ip address> <compute hostname.domain.name> <compute hostname>
```

Edit /etc/hostname: On all nodes, modify the hostname to new one.

Edit /etc/network/interfaces: On all nodes, add network IP address for public interface.

Here is an example:

```
auto eth2
iface eths inet static
    address 172.28.12.88
    netmask 255.255.255.0
    network 172.28.12.0
    broadcast 172.28.12.255
    gateway 172.28.12.1
# dns -* options are impletmented by the resolvconf package, if installed
    dns - nameservers 172.28.12.131
    dns - search cisco.com
```



**Note**

For control node and compute nodes, you can use build node IP address as the DNS nameservers because dnsmasq is running on the build node.

## Reboot the UCS OpenStack Nodes (for all nodes)

After rebooting, you should be able to ping outside from your UCS node.

## Generate Site.pp

Go to /etc/puppet/manifests, edit site.config:

Following is a sample file site.config:

```
vinci-ucs07:~/openstackscripts> cat site.config
# general config, input domain name, ntp server, and general password for keystone, mysql,
nova....
# the user can modify site.pp if they want password to be different in each component, seach
$mypasswd...
[general]
domain_name = cisco.com
ntp_servers = ntp.esl.cisco.com
password = cisco123

##### input dcnm ip address, account name and password and tunnel base address, this
is the segment id starting value
# gateway_mac is the mac you configured in the dcnm poap
[dcnm]
node_ip = 172.28.11.155
user_account = root
user_passwd = cisco123
dhcp_server=10.10.10.152 // specify the external dhcp_server ip address
// so you won't need input it into the DCNM software (start 7.0.2 release)
gateway_mac = 20:20:00:00:00:AA //This should be the same as what you configured in dcnm
dfa_tunnel_base = 90000 //This is the start of the segment ID base. Recommended min value
is 8k. Ensure that this doesn't conflict
// with other orchestrators segment ID or segment_id_range in DCNM
(under Admin/DFA/Settings)

# fill in hostname and ip of build node
[build]
node_name = vinci-ucs36
node_ip = 172.28.12.99

# fill in hostname and ip/network of control node
[control]
node_name = vinci-ucs34
node_ip = 172.28.12.97
node_network = 172.28.12.0

# There can be more than one compute nodes, list their corresponding node name, ip, and
```

```

uplink interfaces.
# Uplink interface is the interface which connects to leaf
# for example, for 3 compute nodes:
# node_name = name1,name2,name3
# node_ip = ip1,ip2,ip3
# node_uplink = uplink1, uplink2, uplink3
[compute]
node_name = vinci-ucs05,vincci-ucs06
node_ip = 172.28.12.68, 172.28.12.69
node_uplink = eth2,eth2

```

Invoke `sudo ./generate_site` at `dir /etc/puppet/manifests`.

## Invoke Scripts to Restore Nodes (All Nodes)

### Invoke `restore_openstack` on Build Server Node

Invoke the `/home/localadmin/restore_openstack`, and when prompted, give the full qualified domain name of build server node and decide if you want to apply `site.pp`.

An example follows:

```

sudo -H bash
root@vinci-ucs29:/home/localadmin# ./restore_openstack
please input FQDN of build node: vinci-ucs29.cisco.com
.....

```

Would you like to apply `site.pp` (y/n)? y

Type y if your `site.pp` is ready. Type n if your `site.pp` is not ready and you want to manually edit it.

After your `site.pp` is ready (or the previous apply `site.pp` is errored out from the script), enter the following:

```
puppet apply -v /etc/puppet/manifests/site.pp
```

This step takes a short while to complete so be patient. Verify that `apache2` is running by entering the following:

```
- service apache2 status
```

Restart `dnsmasq`, by running the following on the build node :

```
service dnsmasq restart
```

### Invoke `restore_openstack` on Control and Compute Nodes

Invoke the `/home/localadmin/restore_openstack` as `sudo`, and when prompted, give the full qualified domain name of build server node:

```

sudo -H bash
root@vinci-ucs138:/home/localadmin# ./restore_openstack
please input FQDN of build node: vinci-ucs29.cisco.com

```

If you plan to deploy a service router VM (CSR), please apply the `iptables` firewall patch on each compute node, as follows:

```

1. cp /usr/share/pyshared/quantum/agent/linux/iptables_firewall.py_new
   /usr/share/pyshared/quantum/agent/linux/iptables_firewall.py
2. Edit /etc/quantum/quantum.conf (as sudo)
sudo -H bash
vi /etc/quantum/quantum.conf
[keystone_authtoken]
auth_host = (control node ip)
admin_tenant_name = services
admin_user = quantum
admin_password = cisco123

```

## Confirm Keystone Started

Next, confirm that keystone has started correctly.

Ensure openrc has the following line:

```
export OS_AUTH_URL=http://<your controller ip address>:5000/v2.0/)
```

Enter the following:

```
source /root/openrc
```

```
keystone service-list - has output
tkeystone endpoing-list - has output
```

•

## Reboot the Nodes

Reboot the nodes, using the following sequence:

- 1 Power down the compute nodes.
- 2 Reboot the control node.
- 3 Check that you can login from your Chrome browser using admin/cisco123 as the username and password.
- 4 Power on the compute nodes.

After the compute node reboots and is up, wait for the puppet to run and then ensure that the daemons LLDAP and Packet Capture are running as follows (also wait for the puppet run to complete as given above):

```
pidof lldpad
pidof pktcpt
```

## Images for Launching a VM

One of the following images can be used to launch an instance from the image list:

- cirros image
  - Login: cirros
  - Password: cubswin:)
- Ubuntu-mini image
  - Login: ubuntu
  - Password: ubuntu

# Configuration in the Leaf or DCNM

## Leaf Interface Configuration and Global Configuration

You must connect the leaf interface to the UCS server and accomplish the global configuration by the following steps:

- 1 You connect the leaf interface to the UCS server, by entering the “fabric forwarding port-tracking” configuration:
 

```
int e1/35
switchport mode trunk
```
- 2 Set the speed of the interface connected to compute node correctly (currently DCNM PoAP screen does not allow you to configure this), by entering the following:
 

```
fabric forwarding port-tracking
speed 1000
no shut
```
- 3 Enter the following global command:
 

```
evb batch-response disable
```
- 4 The leaf node connected to the OpenStack compute node must have the same anycast gateway MAC as in the site.pp file of the build server. Verify this by entering the following command, which is shown with a sample display below:
 

```
N6k-23# sh run | inc anycast
fabric forwarding anycast-gateway-mac 0000.DEAD.BEEF
```

You might see some modules on which the interface connecting the leaf and UCS does not come up. If this should occur, take the following steps:

- 1 Load the debug plugin on the leaf.
- 2 Enter the following command:
 

```
cd /mnt/pss
```
- 3 Enter the following command:
 

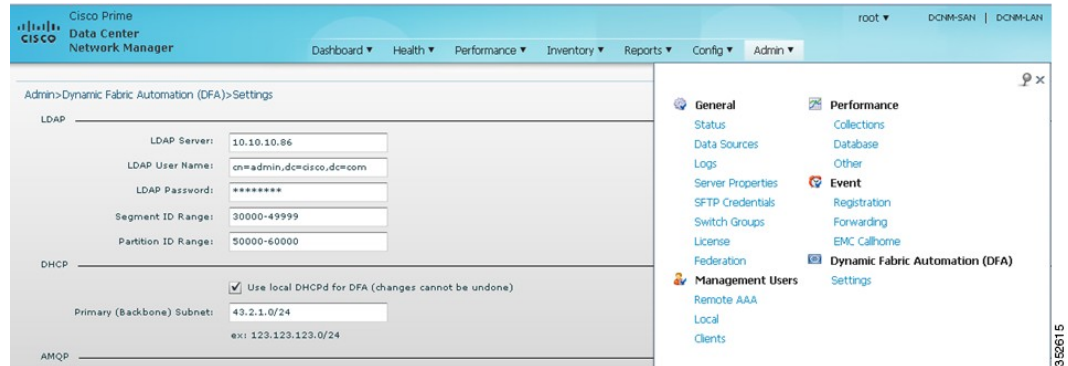
```
rm norcal_disable_qsfp
```
- 4 Reload the leaf.

## DHCP Configuration in DCNM

For DHCP configuration in DCNM:

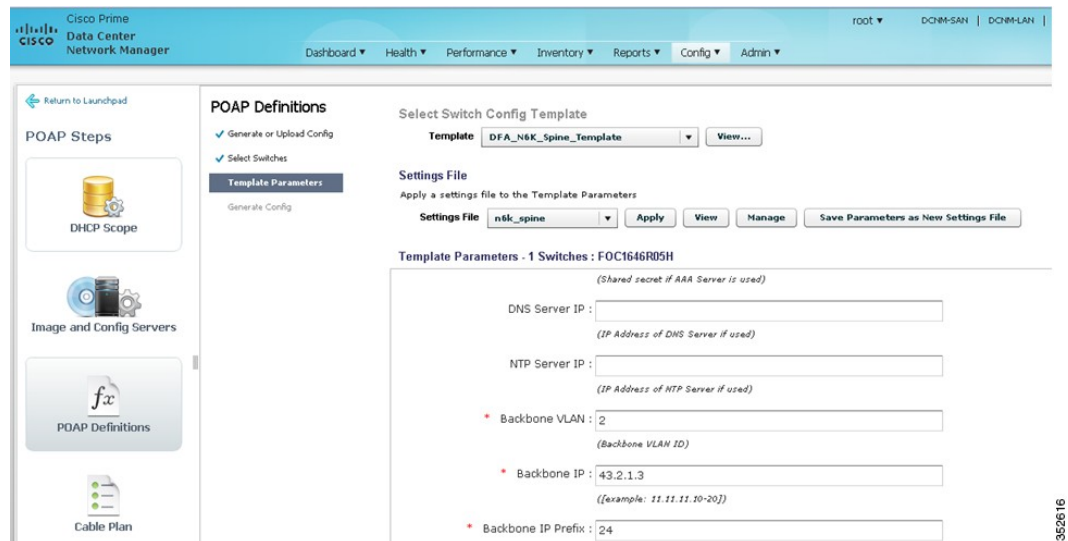
- 1 Fill the DFA backbone VLAN subnet value in the Settings of the Dynamic Fabric Automation (DFA) section under the Admin tab in the DCNM console.

**Figure 4: Admin Tab**



- 2 Ensure that Primary (Backbone) subnet with Use local DHCPd for DFA checked is consistent with the backbone VLAN subnet for the leaf and spine nodes in the DFA fabric. If the fabric is brought up from PoAP in DCNM, the configuration screen will look like the following:

**Figure 5: PoAP Definitions**







## Create Project and Launch VM

---

- [Create Project and Launch VM, page 21](#)
- [Known Caveats, page 21](#)
- [Steps to Create a Project, page 21](#)
- [Steps to Create a User for the Project, page 22](#)
- [Steps to Create the Network, page 22](#)
- [Steps to Launch the VM, page 23](#)

## Create Project and Launch VM

The information provided in this section is generic to OpenStack and it is provided here for your convenience with the exception of ConfigProfile, which is DFA specific.

## Known Caveats

The following are known caveats when you are creating a project and launching VM:

- Do not put hyphens in the project name.
- The Bulk Create and Delete functionality of VMs is not completely supported. (Refer the Known Limitations and Caveats section.)
- Do not give space for a network name or project name or a VM name.

## Steps to Create a Project

Take the following steps to create a project:

- 1 Login to Horizon using the following username and password:  
username: admin  
password: cisco 123

- 2 Click Projects and then Create Project.
- 3 Enter a name for the project.

**Note**


---

Do not use hyphens in the project name.

---

- 4 Click Create Project.

## Steps to Create a User for the Project

Take the following steps to create a user for the project:

- 1 Click Users.
- 2 Click Create User.
- 3 Fill in all the fields, select the project you just created and select the role as admin

**Note**


---

The network information will not be populated to DCNM correctly, if you fail to specify the role as admin.

---

## Steps to Create the Network

Take the following steps to create the network:

- 1 Login as the user you created in the "Steps to Create a User for the Project" section.
- 2 Click the Project tab.
- 3 Click Networks and click Create Network.
- 4 Specify a name for the network and click the Subnet tab.

**Note**


---

Steps 3 and 4 are mandatory.

---

- 5 Specify a network address for the subnet.
- 6 Click the Subnet Detail tab and leave the Enable External DHCP option checked if External DHCP instead of static IP address, for VM to be used.
- 7 Click the CiscoDFA tab, It gives you a new window, and then select the config profile from the pull-down menu.

The other parameters specified in this window is for Cisco's Prime Network Service Controller (PNSC) integration.

Without service, make sure you select the basic non service related profiles:

- defaultNetworkIpv4EFProfile
- defaultNetworkIpv4TfProfile

- defaultNetworkdefaultPartitionIpv4TfProfile
- defaultNetworkdefaultPartitionIpv4EfProfile
- defaultNetworkL2Profile

Then select Network Role "Host Network," as shown below.

**Figure 6: Create Network Screen**

**Create Network**

Network Subnet Subnet Detail **Cisco DFA**

**Config Profile**  
 defaultNetworkIpv4EfProfile

**Network Role**  
 Host Network

**Secondary Gateway IP (optional)**

**Static IP Start**

**Static IP End**

**Partition Name**

You can configure Cisco DFA and PNSC Parameters for the network.

Cancel Create

## Steps to Launch the VM

Take the following steps to launch the VM:

- 1 Click Instances and then click Launch Instance.

- 2 Click the Image drop-down menu and select the image.  
There will be cirros image by default.
- 3 Specify a name for the Instance.
- 4 Click the Networking tab and select the network from the Available network list.
- 5 Click Launch.



# Scalability

---

- [Scalability, page 25](#)

## Scalability

OpenStack is supposed to support many compute nodes from the control node running the Horizon dashboard. This should not be related to the added DFA functionality.

By default without specific configuration, OpenStack supports ten VMs per compute node. This is the recommended number for a user.

Batching VM creation and deletion with OpenStack should be used with great caution. We support ten or fewer VMs per batch creation and deletion on two compute nodes. Batching capability beyond that is best-effort.





## Known Limitations and Caveats

- [Known Limitations and Caveats, page 27](#)

### Known Limitations and Caveats

The following are known limitations and caveats for the project:

- Currently live-migration/vmotion is not supported because it is not supported in COI current release.
- OpenStack does not support IPv6 yet.
- All added DFA functionality in OpenStack is done through the Horizon dashboard for this release.
- If a compute node with live VMs gets reloaded, OpenStack will lose these VMs.
- Currently a single port of a server running OpenStack is allowed to connect to a DFA leaf port.
- The following are required:
  - The compute nodes must be connected to the build server always, at least during reboot.s
  - Puppet cannot be disabled.
- The VMs' VNIC interface up/down is not detected by vdp. (Use VM deletion/creation as a potential workaround.)
- Observed LLDPad does not come up after COI installation, so you must manually reload LLDPad.
- All the subnet creation for network should be done through the Horizon project/network menu. Other methods of creating network subnet might not work.
- All orchestration tasks are done through the Horizon dashboard.
- VMs provisioned through Horizon always get the IP address via DHCP server in DCNM.
- The Cirros and mini Ubuntu images are currently tested.
- If a VLAN is instantiated on the spine then that corresponding 4k segment range cannot be used as the segment ID . For example if VLAN 1 is instantiated in the spine then 4K-to-8K segment cannot be used. Similarly, if VLAN 2 is instantiated in the spine then 8K-to-12K cannot be used and so on. VLAN 1 is always enabled on the spine, so the tunnel base in OpenStack should be at least greater than 8K. Moreover,

depending on the backbone VLAN configured on the spine, the tunnel base in OpenStack needs to be appropriately configured such that it doesn't overlap with the corresponding segment range in the spine.

- Modifying network parameters in the Horizon is not supported. You need to delete the network and create a new network.



## Technical Support Model

---

- [Technical Support Model](#), page 29

## Technical Support Model

The following is the technical support model for using prebuilt OpenStack for Cisco DFA images:

- OpenStack is based on open source and is generally supported through its community using a best-effort approach.
- All general COI installation support will be from the Cisco COI team. See the "Before You Start" section.
- COI-DFA part support will be from the Cisco DFA team.
- The DFA part of the OpenStack DFA support will come from Cisco DFA team.

