Cisco Dynamic Fabric Automation Troubleshooting Guide

First Published: March 12, 2014

Last Modified: May 20, 2015

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-31750-02

CONTENTS

Preface

CHAPTER 1

CHAPTER 2

Preface

The Preface contains the following sections:

Audience

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

Document Organization

This document is organized into the following chapters:

Chapter	Description
"Information About Cisco DFA"	Provides an overview of Cisco Dynamic Fabric Automation (DFA) and descriptions of the Cisco DFA building blocks.
"Migration to Cisco DFA"	Provides information about how to prepare for migration to Cisco DFA, including migration steps and migration configuration.
"Troubleshooting Migration"	Provides the verification steps on how to troubleshoot Cisco DFA migration.

Document Conventions

Note As part of our constant endeavor to remodel our documents to meet our customers' requirements, we have modified the manner in which we document configuration tasks. As a result of this, you may find a deviation in the style used to describe these tasks, with the newly included sections of the document following the new format.

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	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.
[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.
`variable`	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
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	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.

Convention	Description
!, #	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 Dynamic Fabric Automation

The Cisco Dynamic Fabric Automation documentation is at the following URL: http://www.cisco.com/c/en/us/support/cloud-systems-management/dynamic-fabric-automation/tsd-products-support-series-home.html.

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 and 5600 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. 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. 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.

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
- *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/preblt-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

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to: ciscodfa-docfeedback@cisco.com.

We appreciate your feedback.

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.

CHAPTER 1

Information About Cisco DFA

This chapter includes the following sections:

Terminology

The following figure shows the terms that are used for a Cisco Dynamic Fabric Automation (DFA) deployment. You should understand these terms and definitions before you deploy Cisco DFA.

Figure 1: Terms Used in a Cisco Dynamic Fabric Automation Deployment

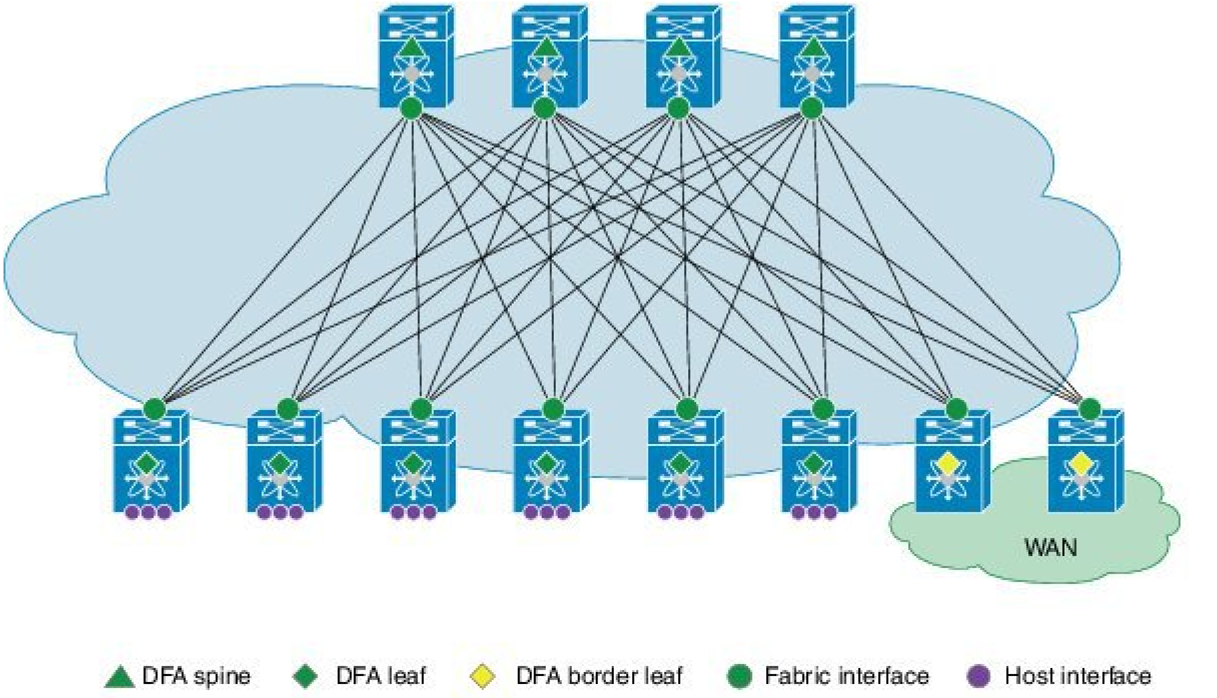

- Cisco DFA fabric—A multistage, switching network in which every connected device is reachable through the same number of hops. The Cisco DFA fabric enables the use of a scale-out model for optimized growth.
- Cisco DFA switch—A leaf, border leaf, or spine device.
- Leaf—Switches with ports that are connected to ethernet devices, such as servers (host interfaces) and ports (fabric interfaces), that are connected to the Cisco DFA fabric. Leaf switches forward traffic based on the enhanced control plane functionality of Cisco DFA optimized networking, which requires segment ID-based forwarding.
- Border leaf—Switches that connect external network devices or services, such as firewalls and router ports, to a Cisco DFA fabric. Border leaf switches are similar to leaf switches and can perform segment ID-based forwarding.
- Spine—Switches through which all leaf and border leaf switches are connected to each other and to which no end nodes are connected. Spine switches forward traffic based on Cisco DFA optimized networking with enhanced or traditional forwarding.

- Host interface—Leaf to server interfaces that receive traffic for connected VLANs to be extended across the Cisco DFA fabric.
- Fabric interface—Ports through which Cisco DFA switches are connected to one another.

Cisco Dynamic Fabric Automation Overview

Cisco Dynamic Fabric Automation optimizes data centers through integration. This architecture eliminates the need for overlay networks that can hinder traffic visibility and optimization and reduce scalability when physical server and virtual machine environments are integrated. The architecture enables zero-touch provisioning and greater orchestration, while delivering more predictable performance and latency for large cloud networks. The following building blocks are the foundation of Cisco DFA:

- Fabric Management—Simplifies workload visibility, optimizes troubleshooting, and automates fabric component configuration.
- Workload Automation—Integrates with automation and orchestration tools through northbound application programming interfaces (APIs) and also provides control for provisioning fabric components by automatically applying templates that leverage southbound APIs and standard-based protocols. These automation mechanisms are also extensible to network services.
- Optimized Networking—Uses a simple distributed gateway mechanism to support any subnet, anywhere, concurrently. Existing redundancy models are also used to provide N+ redundancy across the entire fabric.
- Virtual Fabrics—Extends the boundaries of segmented environments to different routing and switching instances by using logical fabric isolation and segmentation within the fabric. All of these technologies can be combined to support hosting, cloud, and multi-tenancy environments.
- DCI Automation—Automate the configuration of connecting tenants within the unified fabric to the external world, be it the Internet or other unified fabric networks. These features works in tandem with DCNM (7.1.1 onwards) to enable auto configuration of such requirement.

Note Global VLAN mutually exclude segment ID, (at least for Layer-2 Traffic). A segment ID is a global identifier, there cannot be two global identifier = VLAN + segment ID, you have to decide one or the other. Global VLANs and segment ID can co-exist in the same fabric, if the outer header is not overlapping.

Fabric Management

The fabric management network in Cisco Dynamic Fabric Automation represents a dedicated out-of-band network that is responsible for bootstrapping and managing the individual networking devices, such as spines, leafs, and border leaf switches that are controlled by fabric management. The fabric management network is responsible for transporting the protocols that are required for the different fabric management functions. The following table lists the functions and protocols across the fabric management network.

Table 1: Functions and Protocols Across the Fabric Management Network

Function	Protocol
Power On Auto provisioning (POAP) for automatically configuring network devices	• Dynamic Host Configuration Protocol (DHCP) • Trivial File Transfer Protocol (TFTP) • Secure Copy Protocol (SCP)
Fabric discovery	Simple Network Management Protocol (SNMP)
User-to-machine and machine-to-machine communication	Extensible Messaging and Presence Protocol (XMPP)
Automated network provisioning	Lightweight Directory Access Protocol (LDAP)
DCI Automation	Auto Provisioning of Data Center Interconnect on a border leaf.

The management network, also known as the management access, is the network administrator-facing interface for accessing fabric management. The management network represents the portion of your network from which you, as the network administrator, can connect to an element manager or a network management station (NMS) and to switches and routers.

The Cisco Prime Data Center Network Manager (DCNM) is a turn-key management system for fabric management, visibility, and an extensible set of functions to more efficiently control the data center fabric. Cisco Prime DCNM uses standards-based control protocol components to provide you with an extensive level of customization and integration with an operations support system (OSS) network.

Cisco Prime Data Center Network Manager

An Open Virtual Appliance (OVA) is a prebuilt software solution that comprises one or more virtual machines (VMs) that are packaged, maintained, updated, and managed as a single unit. The Cisco DCNM OVA includes an application functionality that is necessary for Cisco DFA. Cisco Prime DCNM as an OVA can be deployed on a VMware vSphere infrastructure.

Cisco Prime DCNM provides the following functionalities:

- Device auto configuration is the process of bringing up the Cisco DFA fabric by applying preset configuration templates to any device that joins the fabric. Auto configuration installs an image or applies the basic configuration.
- Cable-plan consistency checks the physical connectivity of the fabric against a documented cable-plan for compliance. The lack of compliance prevents specific links from being active and protects the fabric from unwanted errors.
- Common point of fabric access allows you, as a network administrator, to interact with the fabric as a single entity (system) to simplify queries and to eliminate switch by switch troubleshooting efforts.

- Automated network provisioning provides a new layer of automation integration in which the data center fabric-switching infrastructure is automatically provisioned for the physical or virtual workload that is being instantiated.
- Automated profile refresh allows keeping the fabric and the network information in sync in a non-disruptive manner.
- DCI Automation provides a touchless provisioning of datacenter interconnections for the tenants.
- Network, virtual fabric, and host visibility is provided by the management GUI and displays a single set of active network elements that belong to an organization in the fabric.

The Cisco DFA DCNM access network is the network administrator facing interface for accessing fabric management and for connecting northbound application program interfaces (APIs) to orchestrators.

Automated Network Provisioning

Cisco DFA fabric automatically provisions tenant networks using a database of network information. Network information database can be looked up using either tenant's traffic information or by VSI Discovery Protocol (VDP) running on the connected Vswitches. The network information database can be stored and managed using Cisco Prime DCNM. This makes it possible for a complete tenant VM orchestration with automated network provisioning to be absolutely touchless from the fabric perspective. For more information on tenant provisioning, see http://www.cisco.com/c/en/us/td/docs/switches/datacenter/dfa/configuration/b-dfa-configuration.html.

Mobility Domain

In a fabric, when auto-configuration is done using tenant's traffic, the dot1q from the traffic is used to locate the network information. Dot1Q is always used with a notion of mobility domain. A mobility domain represents a set of network ports in the fabric where dot1q is treated symmetrically.

From 7.1.x release, each network interface of a leaf can be configured with a mobility domain in addition to global leaf mobility domain configuration. By translating tenant’s dot1Q values to internal leaf dynamic VLANs, true multi-tenancy is achieved with touchless orchestration. A tenant can orchestrate its own range of server VLANs without the need for coordinating the VLAN usage in the fabric. However, with Cisco Nexus 55XX Series Switches as a leaf, mobility domain can only be specified global to the leaf and no translation is possible. For more information on configuration details, see http://www.cisco.com/c/en/us/td/docs/switches/datacenter/dfa/configuration/b-dfa-configuration.html.

VDP-Based Configuration

When a Vswitch connected to the network port is VDP (Virtual station interface Discovery Protocol) capable, VDP can be used to learn segment information of the connected virtual machines in a reliable out of band manner. The segment information being global to the fabric is alone to look up to the network information. In this method, the leaf communicates a dynamically allocated VLAN to the Vswitch through the VDP messages. VDP protocol implementation is based on IEEE standard 802.1QBG. Nexus 1000V and an open source LLDPAD application (for OpenStack) have this VDP implementation.

From release 7.1 onwards, VDP can be used for virtual machines that are provisioned in a VLAN network without using the segment. VDP can also be enabled on Cisco Nexus 55XX Series Switches.

Simplified Profile Management

Network information is stored as a set of parameters in the database; these parameters are then applied to the desired profile to achieve a configuration set for a particular tenant network. Each network can be mapped to its own profile; for example, a network may need only IPv4 parameters and hence it can use a default NetworkIpv4EfProfile and a certain network may use both, where it will use its own profile. Since the 7.1 release, fabric supports universal profiles, where certain parameters can be left empty. If a particular network does not need IPv6 parameters, they can be left unfilled while the profile still contains configuration related to IPv6. This hugely simplifies profile management as only a few profiles will accomplish multiple needs. Also, profile refresh with universal profiles fabric and the network information will be in synchronization in a non-disruptive manner.

CLI-Based Auto-Configuration

Cisco DFA supports a command-line interface (CLI) based auto-configuration for pre-provisioning network devices. The auto-configuration is the same as any configuration that is based on network triggers such as data packet and Virtual Discovery Protocol (VDP). After an auto-configuration is created on a switch, you can use existing Cisco DFA commands, such as the **clear fabric database host** command, to manage the switch configuration.

Automation of Border Leaf L3 External Connectivity

This feature works in conjunction with DCNM (7.1.1 release) to enable auto-configuration of fabric external connectivity on a per-tenant basis. Enhancements have been made to UCSD 5.2 , OpenStack, border leaf POAP template, LDAP Schema, DCNM GUI, and on the switch-side software. These enhancements are done to automate the extension of the tenant towards the DC Edge router and optionally beyond to connect to other fabrics using a BGP MPLS VPN. The DFA 2.0 release completely automates the border leaf auto-configuration for the most common topologies that customers use to connect to the DC Edge box. The creation of the topology is enabled by enhancement to POAP templates for border leaf and a new POAP template is created for a Cisco Nexus 7000-based DC Edge box running a Cisco NX-OS 6.2(10) image. After these devices are booted up, they are imported into Cisco Prime DCNM. At the Cisco Prime DCNM, the imported devices are paired as per network design and assigned attributes such as maximum number of tenants to be deployed on them, the configuration profile associated with the extension. After the topology is complete at Cisco Prime DCNM, the auto-configuration can be globally enabled at Cisco Prime DCNM. At this point, the border leaf auto-configuration is ready for deployment of tenants. This extension can be initiated from the orchestrator (UCSD 5.2 or OpenStack 2). It can also be initiated from Cisco Prime DCNM itself. In Cisco NX-OS 6.2(10) release for Cisco Nexus 7000 platform, the configuration can be generated on Cisco Prime DCNM and copied and pasted manually on the N7000 DC edge device. Similar support is available for ASR9K. The N7000 border leaf (the HUB PE model) will also be supported with auto-configuration in the future releases of Cisco Prime DCNM and N7000. This feature is driven by Cisco Prime DCNM. You can refer to the *Cisco DCNM Fundamentals Guide, Release 7.x*.

After the network is ready for orchestration, the extension can be done by either UCSD or OpenStack. Similarly, the L3 extension can be removed from the orchestrator. For more details, refer to the *Cisco UCS Director Dynamic Fabric Automation Management Guide* and the *Openstack 2.0 User Guide*.

Optimized Networking

Optimized networking in Cisco DFA uses a simple distributed gateway mechanism to support any subnet, anywhere, concurrently.

Frame Encapsulation

Optimized networking in a Cisco DFA deployment uses Cisco FabricPath Frame Encapsulation (FE) for efficient forwarding based on a Shortest Path First (SPF) algorithm for unicast and multicast IP traffic. Host route distribution across the fabric is accomplished using a scalable multi-protocol Border Gateway Protocol (MP-BGP) control plane.

The Cisco DFA enhanced forwarding improves Cisco FabricPath FE by optimizing the conversational learning from Layer-2 to Layer-3. In addition to the enhanced control and data plane for unicast and multicast forwarding, Cisco DFA reduces the Layer-2 failure domain by having the Layer-2/Layer-3 demarcation on the host-connected leaf switch, which terminates the host-originated discovery protocols at this layer.

A distributed anycast gateway on all of the Cisco DFA leaf switches for a VLAN improves resilience and enables the fabric to scale to more hosts by keeping a shorter path for intra and inter-VLAN forwarding. Cisco DFA leaf switches that operate as border leaf switches interconnect the Cisco DFA fabric to external networks. Cisco DFA border leaf switches peer with external standard unicast and multicast routing protocols.

Dynamic VLAN Management

Managing VLANs that are used to interact with the servers is always complicated due to the need for more than 4K tenants. Fabric dynamic VLAN allocations can solve this problem. With a VDP-capable Vswitch, leafs can communicate with Vswitch using VDP and discover the presence of VMs. VDP can communicate segment information of a network to the leaf. The leaf then maps the segment to the next available VLAN. These allocated VLANs are communicated back to the Vswitch for use with the traffic that the VM sends out. A tenant or VM Orchestrator is completely unaware of the VLAN space that needs to be managed across all of the fabric. For a Vswitch that cannot communicate using VDP, a mobility domain can be specified for each network interface where a Vswitch is connected. Each mobility domain in a leaf can be mapped to a VLAN pool. When a tenant network is orchestrated for a particular dot1Q, the dot1Q is normalized to the next available VLAN in the leaf's VLAN pool for forwarding. The VLAN that is mapped can also be configured to carry tenant's traffic over the fabric using a segment. The number of tenant VMs that can be orchestrated under a leaf is drastically increased by enabling tenant VLANs only on the ports where the tenant is detected. When an auto-configuration of a tenant network is done for a network using either VDP or tenant's traffic, the leaf provisions the VLAN that is required for the tenant. The provisioned VLAN is brought up only on the port where the network was provisioned. Refer to the DFA Configuration guide for more details as described in the sections *Multiple Mobility Domain* and *Dynamic Virtual Port*.

Cisco Dynamic Fabric Automation Services Support

Services such as a firewall, load balancer, and virtual private networks (VPNs) are deployed at the aggregation layer in the traditional data center. In a Cisco DFA deployment, services nodes are deployed at regular leaf switches for both east-west and north-south traffic. Services can be physical or virtual services nodes.

The following figure shows the interaction between the Cisco Prime Network Services Controller (NSC) and the Cisco DFA deployment through Cisco Prime Data Center Network Manager (DCNM).

Figure 2: Cisco DFA with Services

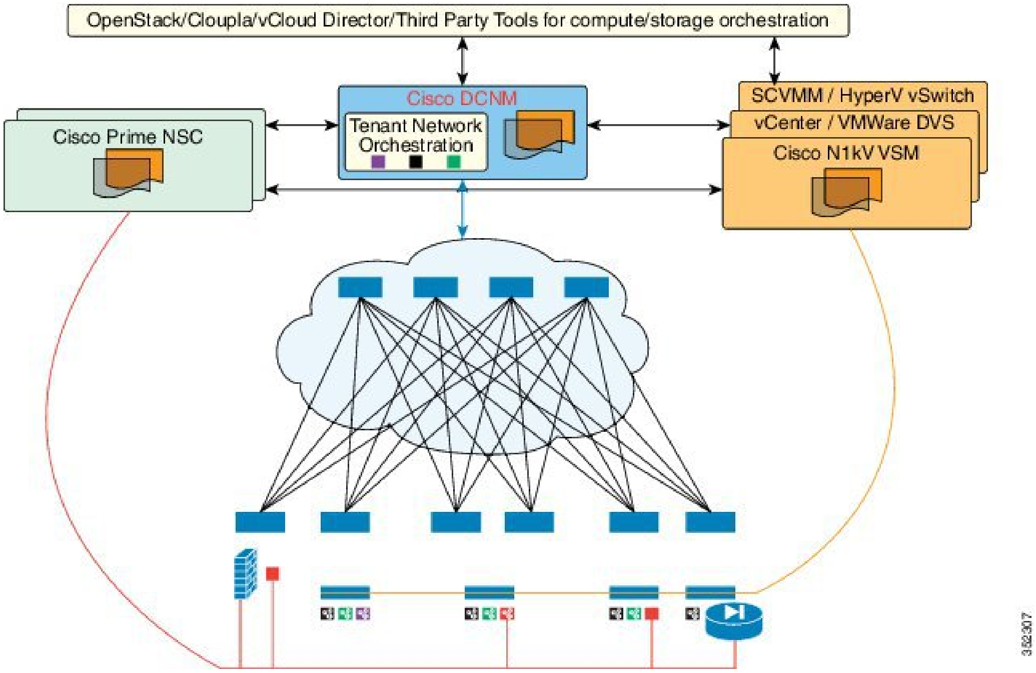
OpenStack/Cloupla/vCloud Director/Third Party Tools for compute/storage orchestration
Cisco Prime NSC
Cisco DCNM
Tenant Network Orchestration
SCVMM / HyperV vSwitch
vCenter / VMWare DVS
Cisco N1kV VSM
362307

The Cisco Prime NSC is the services orchestrator for Cisco DFA. The NSC Adapter in the Cisco Prime DCNM Open Virtual Appliance (OVA) performs the following functions:

- Provides connectivity between Cisco Prime DCNM and the Cisco Prime NSC services orchestrator
- Automatically populates the Cisco Prime NSC with the organizations, partitions, and networks that are created in Cisco Prime DCNM
- Populates Cisco Prime DCNM with the services that are stitched through Cisco Prime NSC
- Allows the use of multiple Cisco Prime NSC instances to match the Cisco Prime DCNM scale

Fabric can be provisioned for services using Cisco UCSD as well without using PNSC for certain scenarios. Containers can be used to orchestrate policies for tenant edge firewall using Physical ASA or ASAv. Containers are integrated with Cisco Prime DCNM to use DFA VLANs to create networks for a firewall’s inside and outside interfaces. VSG service networks can also be orchestrated using UCSD; however, in this scenario, PNSC is required for provisioning the VSG. UCSD deploys all the virtual form factor service nodes (ASAv, VSG) using the port groups with DFA VLANs. These networks are also pushed to Cisco Prime DCNM through the Rest APIs. Note that interaction between PNSC and Cisco Prime DCNM is not needed for this approach; UCSD implements this functionality for services.

In Cisco DFA, configuration profile templates and instantiating the profiles on a leaf switch provide network automation. The templates are extended to support services in Cisco DFA. The profile templates are packaged in Cisco Prime DCNM for the services orchestrator. The table below includes a list of profile templates that are available for Cisco DFA services. It is important that you select the correct profile to orchestrate and automate services in the Cisco DFA fabric.

Table 2: Cisco Templates for Services Support

Service	Network	Routing	Service Profile
Edge Firewall	Host Network	N/A	defaultUniversalTfProfile
	Edge Firewall	Static	serviceNetworkUniversalTfStaticRoutingProfile
		Dynamic	serviceNetworkUniversalDynamicRoutingESProfile
	Tenant External Service Network	Static	externalNetworkUniversalTfStaticRoutingESProfile
		Dynamic	externalNetworkUniversalDynamicRoutingESProfile
Service Node as Router/Default Gateway	Host Network	N/A	defaultNetworkL2Profile

For NSC Adapter installation information, see the *Cisco DCNM 7.1 OVA Installation Guide*.

OpenStack for Cisco DFA

OpenStack creates a human and machine-accessible service for managing the entire life cycle of the infrastructure and applications within OpenStack clouds. The technology consists of a series of inter-related projects that control pools of processing, storage, and networking resources throughout a data center that can be managed or provisioned through a web-based dashboard, command line tools, a RESTful application programming interface (API), or Python scripts based on OpenStack Python SDK.

The OpenStack for Cisco DFA software is an application-level enabler that works with the latest Juno release. OpenStack for Cisco DFA provides orchestration of the cloud that is enabled by Cisco DFA. OpenStack for Cisco DFA provides orchestration of the cloud that is enabled by Cisco DFA.

Users can choose to install OpenStack using their preferred mechanism on their chosen target servers. After the OpenStack installation, the lightweight DFA enabler installation will make the OpenStack DFA ready. The enabler will work with the Juno OpenStack release and will be qualified for prior releases (such as Icehouse) as well.

In the diagram below, OpenStack control and compute nodes are connected together after the generic OpenStack installation is finished. The compute nodes (DC servers of user choice) are connected to the leaf switches. DCNM and OpenStack control node needs to be connected using an IP network.

Figure 3: Sample Topology

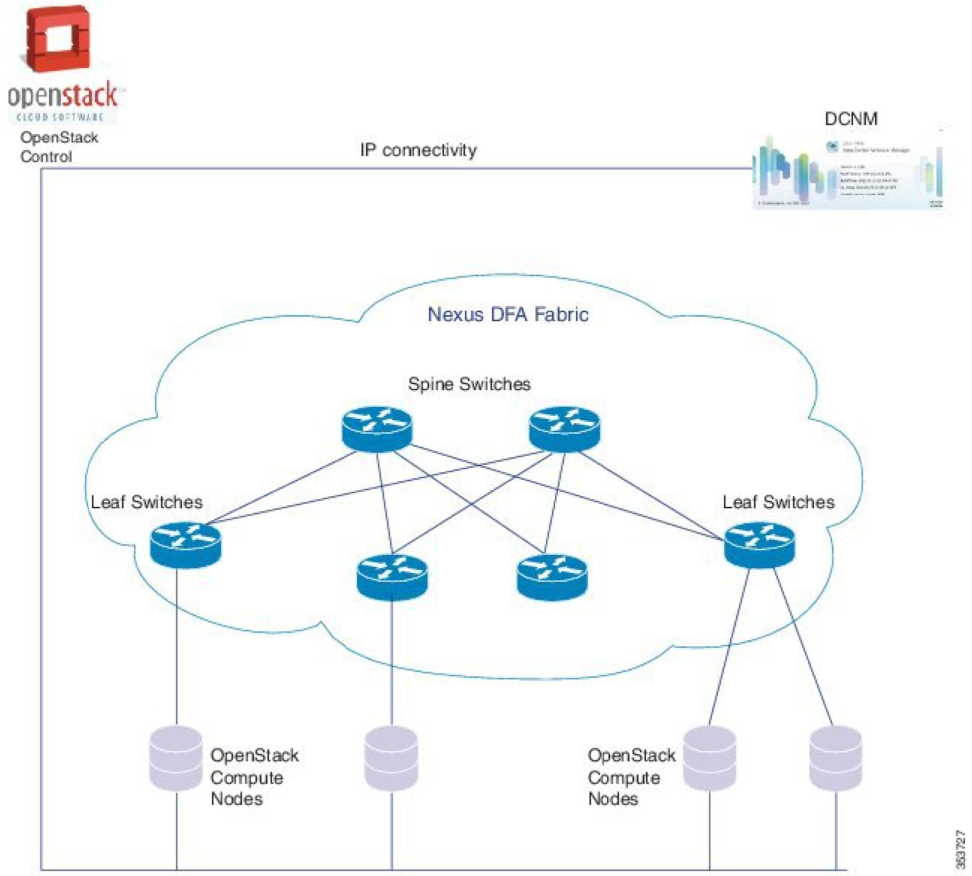

For information about Open Source used in OpenStack for Cisco DFA 2.0, see the Open Source used in *OpenStack for Cisco DFA 2.0* document.

CHAPTER 2

Troubleshooting Platform and Software Requirements

This chapter includes troubleshooting for platform and software requirements that are associated with Cisco Dynamic Fabric Automation (DFA).

This chapter contains the following sections:

Troubleshooting Platform and Software Requirements for the Cisco Nexus 6000 Series Switches

This section describes troubleshooting procedures for various issues that involve platform and software requirements for Cisco Nexus 6000 Series switches that are associated with a Cisco Dynamic Fabric Automation deployment.

Symptom	Cause	Resolution
Layer 3 forwarding is not functioning.	On a Cisco Nexus 6000 Series switch, Layer 3 forwarding does not work without the LAN base license.	Ensure that you have each of the following licenses installed: • ENTERPRISE_PKG • ENHANCED_LAYER2_PKG • LAN_BASE_SERVICES_PKG • LAN_ENTERPRISE_SERVICES_PKG

Symptom	Cause	Resolution
The Border Gateway Protocol (BGP) is not functioning.	On the Cisco Nexus 6000 Series switch, an Enterprise License is required to enable BGP.	Ensure that you have each of the following licenses installed: • ENTERPRISE_PKG • ENHANCED_LAYER2_PKG • LAN_BASE_SERVICES_PKG • LAN_ENTERPRISE_SERVICES_PKG
Performance Monitoring and Configuration Archive is not functioning on a device.	A Cisco Prime Data Center Network Manager (DCNM) license is required on each device for Performance Monitoring and Configuration Archive to function. Demo licenses will not work. **Note** No special licensing is required for other Cisco Dynamic Fabric Automation (DFA) functions within Cisco Prime DCNM.	Access a license at http://www.cisco.com/go/license . **Note** If you installed Cisco DCNM using the DCNM OVA installation process, these licenses are included as part of that installation. You can obtain a demo license using the DCNM web interface under licensing.

Troubleshooting Platform and Software Requirements for Cisco Nexus 7000 Series Switches

This section describes troubleshooting procedures for various issues that involve platform and software requirements for the Cisco Nexus 7000 Series switches that are associated with a Cisco Dynamic Fabric Automation deployment.

Symptom	Cause	Resolution
A Virtual Private Network (VPN) address family configuration is rejected. **Note** VPN address family configuration is only required if the Cisco Nexus 7000 Series switch is configured as a route reflector.	A Multiprotocol Label Switching (MPLS) license is not installed.	Upgrade the Cisco NX-OS software to Release 6.2(6a). This version eliminates the requirement for an MPLS license.

Symptom	Cause	Resolution
Performance Monitoring and Configuration Archive is not functioning on a device.	A Cisco Prime Data Center Network Manager (DCNM) license is required on each device for Performance Monitoring and Configuration Archive to function. **Note** No special licensing is required for other Cisco Dynamic Fabric Automation (DFA) functions within Cisco Prime DCNM.	Access a license at http://www.cisco.com/go/license . **Note** If you installed Cisco Prime DCNM using the DCNM OVA installation process, these licenses are included as part of the DCNM installation.

Troubleshooting Platform and Software Requirements of the Cisco Nexus 1000V Series Switch

See the following URLs for configuration information about the Cisco Nexus 1000V Series switch for Cisco Dynamic Fabric Automation:

- http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_2_2_2/DFA/configuration/b_Cisco_Nexus_1000V_DFA_Configuration_Guide_421_SV2_2_2.html
- http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_2_2_2/VDP/configuration/b_Cisco_Nexus_1000V_VDP_Configuration_Guide_421_SV2_2_2.html

See the following URL for information about troubleshooting the Cisco Nexus 1000V Series switch, including information about the VSI Discovery and Configuration Protocol (VDP):

- http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_2_2_2/troubleshooting/configuration/guide/n1000v_trouble.html

Troubleshooting OpenStack for Cisco Dynamic Fabric Automation

See the following URL for information about Open Source software used in OpenStack for Cisco DFA:

- http://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/dfa/openstack/opensource/OpenStack_for_Cisco_DFA_1-0_Open_Source_Documentation.pdf

See the following URL for information about the Pre-Built OpenStack for Cisco DFA Images:

- http://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/dfa/openstack/install/guide/preblt-image.pdf

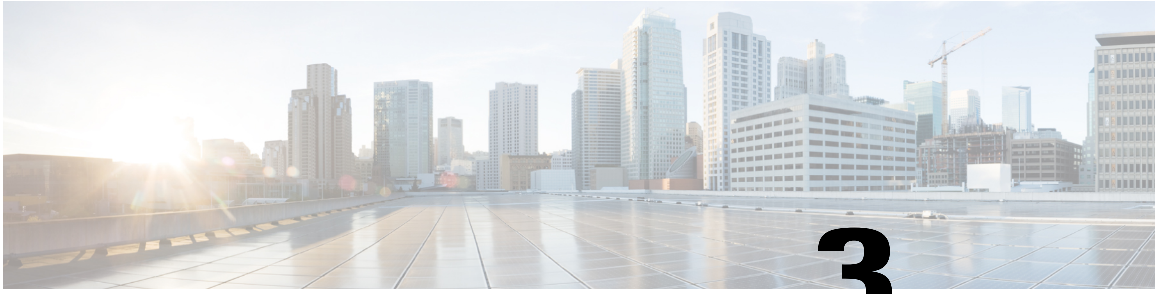

CHAPTER 3

Troubleshooting Your Cable Plan

This chapter contains the following sections:

Troubleshooting Your Cable Plan

This section describes troubleshooting procedures for various scenarios that involve the cabling or cabling plan for switches within a Cisco Dynamic Fabric Automation (DFA) deployment.

Symptom	Cause	Resolution
Unable to import a cable plan from Cisco Prime DCNM to a device.	• The switch does not have the **feature cable-management** command enabled. • Cisco Prime DCNM is unable to make a Secure Shell (SSH) connection to the device. • The cable plan is not correctly formatted for the switch.	**1** Check log files: `/usr/local/cisco/dcm/fm/logs/fmserver*.log*` for indication of the error condition. **2** Use the **show accounting log** command on the switch to review the following: • Errors from the switch • Whether the cable plan was successfully executed on the switch • Reported errors
	The switch already has an XML file with the same name.	**1** Manually delete the existing XML file. **2** Import the file again.

Symptom	Cause	Resolution
The following error is displayed when you attempt to import a cable plan: `Failed to import cable plan or Remove copy failure`	The switch was previously registered with a different Cisco Prime DCNM instance.	**1** Open an SSH terminal session. **2** Clear the SSH key on the switch on which you want to deploy the cable plan. For example: `n6k-leaf-2017#` **`clear ssh hosts`**
A cable or tier mismatch error is expected but did not occur.	The DFA switch poll is not updated.	**1** Open an SSH terminal session. **2** Use the **show fabric conn neighbor** command to identify the state of cable or tier links. **3** Review the output of the **show fabric conn neighbor** command for cable/tier mismatches. **4** If the **show fabric conn neighbor** command output displayed the proper connections, DFA should show those same connections. However, the updated state might take up to 5 minutes. **Note** DFA polls the switches every 5 minutes to learn the error state information.

Symptom	Cause	Resolution

The following system message error is displayed when you attempt to import the cable plan:
`errDisabled`
	In the cable plan XML file, sourceChassis and destChassis were not specified in a hostname.domainname fully qualified domain name (FQDN) format.	Specify the sourceChassis and destChassis in a hostname.domainname format.
		Note When you enter the **ip domain-name** command, the Link Layer Discovery Protocol (LLDP) sends the chassis type-length-value (TLV) in an FQDN format.
		For example:
		`switch # config terminal`
		`switch (config) #` **`ip domain-name cisco.com`**
		`<<<<cableplan.xml snip>>>>`
		`<CHASSIS_INFO sourceChassis="spine1.cisco.com"`
		`type-"n6k">`
		`<LINK_INFO sourcePort="Eth2/1" destChassis="leaf1.cisco.com" destPort="Eth2/1"/>`

No neighbors are detected.	LLDP is not running on the interface and cannot transport information.	Check to ensure that the LLDP is enabled on the interface. Use the **feature lldp** command to enable LLDP on the switch and the **lldp** {**receive** | **transmit**} command to enable the reception or transmission of LLDP packets on an interface.
Problems occur with neighbor detection.	Neighbor connections are stale.	• Use the **feature cable-management** command to enable the cable management feature. • Use the **clear fabric connectivity neighbors stale** command to clear neighbor cache information for stale or purged neighbors.

Symptom	Cause	Resolution
When cable-plan is not enforced, the following status is displayed for links when you issue **show fabric connectivity cable-plan —** "ErrC, S".	Actual link errors or faulty cable plan is causing the errors. After the errors are seen, disabling the cable-plan enforcement.	**1** Open an SSH terminal session. **2** Actual link or cable errors could have caused the ErrC status displayed via the **show fabric connectivity cable-plan** command. Since cable-plan enforcement is disabled after the errors are seen via the **show fabric connectivity cable-plan** command, there is no automated way to clear those entries. You can either ignore the show command output or use the **clear fabric connectivity cable-plan** command to clear the entries. To fix the actual link errors and enable cable-plan, perform the following steps. **1** Use the **clear fabric connectivity cable-plan** command to clear the error/stale entries. **2** Verify/re-generate cable-plan to be imported. **3** Import and enforce cable-plan. **4** Use the **show fabric connectivity cable-plan** to identify the status of the links. **5** If the connections are OK, the output of **show fabric connectivity neighbors** command should list the cable plan information. If not, check the connections, cable-plan, and cabling.

Symptom	Cause	Resolution
The entry for cable plan in the output of **show fabric connectivity neighbors** command is 'unknown'.	Cable-plan is not enforced.	**1** Open an SSH terminal session. **2** To verify if the cable plan enforcement is disabled via the **show fabric connectivity neighbors** command output, check if the output reads the following message "Cable-Plan Enforce: Disabled" and the cable-plan column displays “unknown". **3** In order to generate/import/enforce cable-plan, use the **fabric connectivity cable-plan..** command. **4** Once a cable-plan is successfully enforced, the output of **show fabric connectivity neighbors** will list the appropriate cable-plan entry.

CHAPTER 4

Troubleshooting the DCNM OVA Installation

This chapter contains the following sections:

Troubleshooting the OVA Installation

This section describes troubleshooting procedures for various scenarios that involve the Cisco Prime DCNM OVA installation that is associated with Cisco Dynamic Fabric Automation deployment.

Symptom	Cause	Resolution
The following error is displayed in the vSphere client while you are deploying the Cisco DCNM OVA: `The OVA package requires support for OVR Properties. Details: Unsupported element "Property".`	The vSphere client is connected directly to the ESXi host.	Connect the vSphere client to a vCenter and not directly to the ESXi host.
The VMware vCloud Director (vCD) or Cisco Prime Network Registrar (CPNR) scripts run with connection exceptions.	The vCD, Virtual Supervisor Module (VSM) or CPNR are not reachable.	Make sure that the external entities (vCD, vSM, or CPNR) are reachable and that the vCD web interface is up and running.

Symptom	Cause	Resolution
Unable to login with default credentials, after the Cisco DCNM OVA is deployed.	Administrative password entered at Management Properties section of OVF deploy contains the '@' character.	Do not use the character with the '@' symbol for password. If used, reinstall to access WebUI with the recommended password criteria.

Troubleshooting Accessibility and Connectivity Issues

This section describes troubleshooting procedures for various scenarios when you are unable to access Cisco Prime DCNM after a successful OVA deployment.

Symptom	Cause	Resolution
Cisco Prime DCNM and other applications in the OVA have connectivity issues; a noticeable drop in ping packets are directed to the virtual appliances.	The datastore used by Cisco Prime DCNM virtual appliance may be almost full or full. The VM performance is compromised and unpredictable.	1 Free up datastore space and then reboot the Cisco Prime DCNM OVA. 2 Make sure that the IP address for the Cisco Prime DCNM management access and Enhanced Fabric Management networks are in different subnets. a If they are not in different subnets, redeploy the OVA. b Enter different subnets for the Cisco Prime DCNM management access and Enhanced Fabric Management networks.
The OVA is successfully deployed, but either the Cisco Prime DCNM management access or Enhanced Fabric Management interface is down.	The IP addresses for Cisco Prime DCNM and Enhanced Fabric Management networks are in the same subnet.	1 Make sure that the IP address for the DCNM management access and Enhanced Fabric Management networks are in different subnets. a If they are not in different subnets, redeploy the OVA. b Enter different subnets for the Cisco Prime DCNM management access and Enhanced Fabric Management networks.

Symptom	Cause	Resolution

Cisco Prime DCNM is not accessible from the Web UI. The following message is displayed:
`System Message: DCNM(pid 22094) process running, but may not be accessible from Web UI.`
	Cisco Prime DCNM is starting due to a scheduled restart or due to an appliance reboot.	1 After waiting a few minutes, use the **appmgr status dcnm** command to determine if Cisco DCNM is still not accessible.
		2 Use the **appmgr start dcnm** command to start Cisco Prime DCNM.
		3 Review the server logs at the following location:
		`/usr/local/cisco/dcnm/`
		`/jboss-4.2.2GA/server/fm/log/server.log`

Users are unable to log in to the Cisco Prime DCNM Web UI.	The password does not meet security requirements.	Make sure that the administrative password created during the OVA deployments meets password security criteria. For version 7.0.1, redeploy the OVA with a new password. The password must meet the following requirements: • Eight characters in length • A combination of alpha and numeric characters • Limited to include only the following special characters ◦ . (dot) ◦ + (plus) ◦ _ (underscore) ◦ - (hyphen)
The Cisco DFA view in the Cisco Prime DCNM Web UI is taking too long to load.	A large number of devices are managed with a local/packaged Postgres database; significant delays can occur in load times.	1 Verify the number of devices managed by Cisco DFA. 2 If the number of devices managed by Cisco DFA exceeds 50 devices, make sure that you are using an external Oracle Database. See http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/7_x/dcnm/installation/master_files/OVA_Installation_Guide.html for instructions on configuring the Oracle database for Cisco Prime DCNM.

Symptom	Cause	Resolution
The **IP address** field does not display a fourth octet.	Your monitor resolution is not properly set.	Change your monitor display settings to a higher resolution.
DHCPD crashes if a /8 subnet is entered.	An /8 subnet is in use. We recommend that you do not use a /8 subnet; this is a problem with the DHCPD that is packaged as part of the DCNM bundle.	If DHCPD crashes, you must manually restart it by using an SSH terminal session in DCNM.
The creation of a network fails.	DHCPD may already be employed for that network. DHCPD does not allow overlapping subnets, even if they are created across different virtual routing and forwarding (VRF) instances.	Ensure that the new network does not have a subnet that conflicts with an existing network.

Configuration-profile instantiation or application errors	Conflicting Profiles Exists	Use the following commands to debug the log file:
		• Use the **show fabric database host** to check what profile have been applied successfully.
		• Use the **show fabric database host statistics** if data is missing. Determine if there are any the statistics
		`sh system internal config-profile history`
		`debug logfile fabric-autoconfig`
		`debug fabric forwarding auto-config all`
		`copy log:fabric-autoconfig bootflash:`
		`debug logfile port-profile`
		`debug port-profile all`
		`copy log:port-profile bootflash:`

Troubleshooting Database Issues

This section describes troubleshooting procedures for various scenarios when you encounter problems with Cisco Prime DCNM and database connections in Cisco Dynamic Fabric Automation (DFA).

Symptom	Cause	Resolution
The Cisco Prime DCNM logs report TNSListener exceptions from the Oracle database.	Sessions, process, and open cursors are set incorrectly.	1 Review the Oracle database configuration. 2 Make sure that the sessions, processes, and open cursors are set appropriately.
The Cisco DFA view in Cisco Prime DCNM is taking too long to load.	A large number of devices are managed with a local/packaged Postgres database.	1 Verify the number of devices managed by Cisco DFA. 2 If the number of devices managed by Cisco DFA exceeds 50, make sure that you are using an external Oracle database. See http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/7_x/dcnm/installation/master_files/OVA_Installation_Guide.html for instructions on configuring the Oracle database for Cisco DCNM.

Symptom	Cause	Resolution
Devices are not able to query the network database.	The cause can be one of the following: • The Lightweight Directory Access Protocol (LDAP) is not up and running. • There is a problem with the connection to the profile database.	1 Log in to the DCNM web UI. 2 On the menu bar, choose **Admin** > **DFA Settings.** 3 Determine if the LDAP used as the network/profile database is local or external to the Cisco Prime DCNM virtual appliance. • If LDAP is local, use the **appmgr status ldap** command to verify that LDAP is up and running. • If LDAP is external, verify that LDAP is up and running. 4 Log in to the device and turn on debugging using the **debug adbm trace** command and the **debug fabric forwarding auto-config evt** command. 5 Test the connection to the network database using the **test fabric database network segment 12300** command. 6 Test the connection to the profile database using the **test fabric database profile MyTestProfile** command. 7 Use the **show fabric database statistics** command on the device to review the primary databases that the switch is connected to and the success and failure counts. 8 Use the **show logging logfile** command to display messages in the log file.

Troubleshooting XMPP Issues

This section describes troubleshooting procedures for various scenarios that involve the Extensible Messaging and Presence Protocol (XMPP) application within the Cisco Prime DCNM OVA deployment for Cisco Dynamic Fabric Automation (DFA) .

Symptom	Cause	Resolution
The XMPP application is down.	Either XMPP has not come up shortly after the OVA deployment (it takes a few minutes) or a fully qualified domain name was not used during OVA deployment.	**1** Check the logs at `/root/postInstallApps.log`. **2** Make sure that a fully qualified name was entered for the **hostname** attribute during the OVA deployment.
Devices are having problems connecting to XMPP.	The device user and groups have not been created.	**1** Review the log files for XMPP and Jabber. • Log file location for XMPP: `/usr/local/cisco/dcm/fm/logs/fms_xmpp.log` • Log file location for Jabber: `/opt/jabber/xcp/var/log/jabberd.log` **2** If devices cannot connect, make sure that device users and groups have been created.

Symptom	Cause	Resolution
Devices are not able to join XMPP groups.	XMPP may not be up and running.	**1** Log in to the DCNM web UI. **2** Determine if XMPP is local or external to the Cisco Prime DCNM virtual appliance. • If XMPP is local, use the **appmgr status xmpp** command to verify that the XMPP is up and running. • If XMPP is external, verify that XMPP is up and running: **1** Use the **show fabric access group** command to display the groups that a device is subscribed to, or to display a list of members existing in a particular group. **2** Use the **show fabric access group device** command to list the groups that are currently logged in to the device. **3** Use the **show fabric access connection** command to display the connection status of a device or a user in the fabric access network.

Symptom	Cause	Resolution
A device is not re-pulling from LDAP when the autoconfiguration network is updated.	The XMPP server is down.	Use the **appmgr status xmpp** command to make sure the XMPP server is up and running.
	The device is pointed to an incorrect XMPP server.	**1** Log in to the DCNM web UI. **2** Click the **View Details** link to view more information. **3** Make sure that the device points to the correct XMPP server. **4** Choose **Admin** > **Settings.** **5** Make sure the Cisco Prime DCNM uses the XMPP user specified on the **DFA Settings** page to join the group.
	The device is not accepting SVI updates.	Use the **clear fabric database host vni segmentID re-apply** command to verify that the device accepts the SVI update. This update is sent by Cisco Prime DCNM for device notification.

Symptom	Cause	Resolution
Cisco Prime DCNM displays fewer devices that have matched virtual machines (VMs) or virtual routing and forwarding (VRF) instances that are associated with them.	The VM or VRF is not associated with the device.	**1** Use the **show evb host** command to verify that the VM is associated with the device. **2** Use the **show vrf** command to verify that the VRF is associated with the device.
	The device is pointing to the incorrect XMPP server.	**1** Log in to the DCNM web UI. **2** Click the **View Details** link to view more information. **3** Make sure that the device points to the correct XMPP server. **4** In the Cisco DCNM web UI, choose **Admin** > **Settings.** **5** Make sure that Cisco Prime DCNM uses the XMPP users specified on the **DFA Settings** page. **6** Make sure that Cisco Prime DCNM and the device joined the same XMPP group.
	The device could be timed out.	**1** Log in to the DCNM web UI. **2** On the menu bar, choose **Admin** > **Settings.** **3** Verify if the XMPP response timeout specified on the **DFA Settings** page allows enough time.
	The XMPP server is down.	Use the **appmgr status xmpp** command to ensure that the XMPP server is up and running.

Troubleshooting DHCP Issues

This section describes troubleshooting procedures for various scenarios that involve DHCP in the Cisco Prime DCNM OVA installation for Cisco Dynamic Fabric Automation (DFA).

Symptom	Cause	Resolution
DHCP does not come up after you use the **appmgr setup ha** command.	No free IP address ranges were entered for the default scope.	**1** Log in to the DCNM web UI. **2** On the menu bar, choose **Configuration** > **POAP** > **DHCP Scope.** **3** Enter the free IP address ranges for the default scope: **enhanced_fabric_mgmt_scope.**
Devices do not see a DHCP server or receive a DHCP response.	DHCP is not running; no lease is available.	**1** Use the **appmgr status dhcp** command to determine if DHCP is running. **2** Check the /var/log messages for any error message from DHCP or to determine if no lease is available.
	The IP address is not active.	Review the following file to determine whether the IP address is active (allocated), aborted (no more address is available), or whether the device is free or has released the allocated address after Power-on Auto Provisioning (POAP): `file/var/lib/dhcpd/dhcpd.leases`
	The DHCP packet is sent to the incorrect interface.	Use TCPDUMP on port 67 and 68 to make sure that the DHCP packet is sent to the correct interface (eth1): tcpdump-ieth1-vv"port67 or port68"
The DHCP client is not getting an IP address when using Cisco Prime DCNM as the DHCP server.	The backbone VLAN subnet is missing.	**1** Log in to the DCNM Web UI. **2** On the menu bar, choose **Admin** >**DFA** > **Settings** > **DHCP.** **3** Click the **Edit scope** icon to edit the scope and add a VLAN subnet.

Troubleshooting AMQP Issues

This section describes troubleshooting procedures for various scenarios that involve the AMQP application within the Cisco Prime DCNM OVA deployment for Cisco Dynamic Fabric Automation.

Symptom	Cause	Resolution
Cisco Prime DCNM is not sending or is missing AMQP notifications.	The AMQP server is not up or the exchange does not exist on the AMQP server.	1 Use the **appmgr status amqp** command to make sure that the AMQP server is up and running. 2 Use the **rabbitmqctl list exchanges** command to make sure the exchange specified on the **DFA Settings** page in DCNM exists on the AMQP server. 3 Please ensure that the **Use local DHCPd for DFA** option is selected. 4 Enter the subnet address that corresponds to the iBGP or backbone VLAN employed in the DFA cluster managed by this DCNM. 5 Make sure that a fully qualified domain name was entered for the **hostname** attribute during the OVA deployment.

Troubleshooting LDAP Issues

This section describes troubleshooting procedures for various scenarios that involve Lightweight Directory Access Protocol (LDAP) within the Cisco Prime DCNM OVA deployment for Cisco Dynamic Fabric Automation.

Symptom	Cause	Resolution
The following system error message is displayed:`500 Internal Server Error. LDAP server communication failure. Failed to add new scope, because of IP Range values already in use.`	The IP address range is already in use or has overlapped with another network.	1 Log in to the Cisco Prime DCNM web UI. 2 On the menu bar, choose **Configuration** > **POAP** > **DHCP Scope**. 3 Update the free IP address ranges for the default scope: **enhanced_fabric_mgmt_scope.**

Symptom	Cause	Resolution
A device is not re-pulling from LDAP when the autoconfiguration network is updated.	The device is pointed to an incorrect XMPP server.	1 Use the **appmgr status xmpp** command to make sure that the XMPP server is up and running. 2 Make sure that the device points to the correct XMPP server. 3 In the Cisco Prime DCNM web UI, choose **Admin** > **Settings.** 4 Make sure that Cisco Prime DCNM uses the XMPP user specified on the DFA Settings page to join the group.
	The device is not accepting Switch Virtual Interface (SVI) updates.	Use the **clear fabric database host vni segmentID re-apply** command to verify that the device accepts the SVI update. This update is sent by Cisco Prime DCNM for device notification.
A device does not autoconfigure SVIs.	The device is pointed to an incorrect LDAP server.	1 Ensure that the device points to the correct LDAP server for both the network and profile type data. 2 Enable debugging on the device.
	The device is not properly configured in LDAP.	Ensure that the network is properly configured in LDAP. All fields must be filled in.
	The device does not have proper licenses installed.	Ensure the device has the proper licenses installed: • ENTERPRISE_PKG • ENHANCED_LAYER2_PKG • LAN_BASE_SERVICES_PKG • LAN_ENTERPRISE_SERVICES_PKG

Symptom	Cause	Resolution
The Org/Partition drop-down list is empty in DCNM.	LDAP is unreachable or the Org/Partition definition is unavailable in the LDAP.	**1** Log in to the DCNM web UI. **2** Review the **DFA Health** tab to review the health of the LDAP and its contents.

Troubleshooting High Availability Issues

This section describes troubleshooting procedures for various scenarios that involve a high availability (HA) environment after a Cisco Prime DCNM OVA installation for Cisco Dynamic Fabric Automation (DFA).

Symptom	Cause	Resolution
DHCP does not come up after you use the **appmgr setup ha** command	No free IP address ranges were entered for the default scope.	**1** Log in to the Cisco Prime DCNM web UI. **2** On the menu bar, choose **Configuration** > **POAP** > **DHCP Scope**. **3** Enter the free IP address ranges for the default scope: **enhanced_fabric_mgmt_scope**.

Symptom	Cause	Resolution

It is difficult to access or bring up applications or virtual IP addresses after you set up an HA environment.	• The HA setup has issues.
	• Listening inconsistencies have occurred.
		1 Review the following log for HA setup issues: `/root/cluster/ha.log`
		2 Determine any listening inconsistencies.
		The following example shows that the Cisco Prime DCNM virtual IP address is listening on the HTTP port on the server on which the **ipvsadm command** is entered (indicated by 'Local') and that AMQP is listening on the peer server (indicated by 'Route').
		`[root@dcnm139 ~]# ipvsadm`
		`IP Virtual Server version 1.2.1 (size=4096)`
		`IP Virtual Server version 1.2.1 (size=4096)`
		`Prot LocalAddress:Port Scheduler Flags`
		`-> RemoteAddress:Port Forward Weight ActiveConn InActConn`
		`TCP 192.168.57.157:http wlc`
		`-> 192.168.57.139:http Local 1 0 0`
		`TCP 10.77.247.157:amqp wlc`
		`-> dcnm155:amqp Route 1 0 0`

CHAPTER 5

Troubleshooting Power-on Auto Provisioning Issues

This chapter contains the following sections:

Troubleshooting POAP Issues

This section describes troubleshooting procedures for various scenarios that involve Power-on Auto Provisioning (POAP) and autoconfiguration associated with Cisco Dynamic Fabric Automation.

Symptom	Cause	Resolution
POAP has failed.	POAP failed during a specific phase.	Review the following POAP files in the switch bootflash to understand in what phase the POAP failed: `poap.log.*` and `.#poap.*init.log`
The device is not booted up with the updated configuration in Cisco Prime DCNM.	The configuration was saved but not published. Clicking Save does not publish the configuration to the file server; it is invisible to the device.	1 Log in to the Cisco Prime DCNM web UI. 2 On the menu bar, choose **Config** > **Power-on Auto Provisioning (POAP)**. 3 Choose POAP switch definitions from the list and click the **Publish** button. 4 Reload the switch for the updated configuration to be applied.

Symptom	Cause	Resolution
A configuration error occurs during provisioning.	The switch's own IP address is configured as the route-reflector peer.	Make sure that the switch IP address and the route reflector IP address are correct (and different).
A user-imported template does not appear in the Template selection on the **POAP Creation** page.	The template has either not been marked as a POAP template or has not been published.	**1** Log in to the Cisco Prime DCNM web UI. **2** On the menu bar, choose **Config** > **Power-on auto Provisioning Open** > **POAP Definitions**. **3** In the Configuration Steps, click the **template** hyperlink in the **POAP Definitions** section. **4** Select the template and click the **Modify/View template** icon. **5** Make sure that the template is marked as a POAP template type and that it is marked as Published. **6** If the template is not published, click the **Publish** button. **7** Reload the switch for the updated configuration to be applied.
The device cannot download the Python boot script (poap_dcnm.py) from the DCNM server.	The serial number entered in the POAP definition does not match the device chassis serial number.	**1** View the output from the **show license host-id** command. **2** Make sure that the serial number you entered in the POAP definition matches the device chassis serial number, which is displayed in the output
	The Trival File Transfer Protocol (TFTP) server IP address is not accessible from the switch.	Make sure that the TFTP server IP address is accessible from the switch.

Symptom	Cause	Resolution
The device does not autoconfigure switch virtual interfaces (SVIs).	The device is pointed to the incorrect Lightweight Directory Access Protocol (LDAP) server.	**1** Ensure that the device points to the correct LDAP server for both network and profile type data. **2** Enable debugging on the device.
	The device is not properly configured in LDAP.	Ensure that the network is properly configured in LDAP and that all fields are filled in.
	The device does not have proper licenses installed.	Ensure that the device has proper licenses installed: • ENTERPRISE_PKG • ENHANCED_LAYER2_PKG • LAN_BASE_SERVICES_PKG • LAN_ENTERPRISE_SERVICES_PKG
Autoconfigurations are not restored after using the **appmgr backup dcnm**and **appmgr restore dcnm** commands.	The autoconfiguration tenant information is stored in the database/LDAP/DHCP.	Use the **appmgr restore all** command to restore all applications.
DHCP does not come up after using the **appmgr setup ha** command.	No available IP address ranges were entered for the default scope.	**1** Log in to the Cisco Prime DCNM web UI. **2** On the menu bar, choose **Configuration** > **POAP** > **DHCP Scope** **3** Enter the free IP address ranges for the default scope: **enhanced fabric mgmt scope**.
The following message is displayed: `500 Internal Server Error. LDAP server communication failure. Failed to add new scope, because of IP Range values already in use.`	The IP address range is already in use or overlaps with another network.	**1** Log in to the Cisco Prime DCNM web UI. **2** On the menu bar, choose **Configuration** > **POAP** > **DHCP Scope** **3** Enter the free IP address ranges for the default scope: **enhanced fabric mgmt scope**.

Symptom	Cause	Resolution
You cannot edit a published POAP template.	You cannot modify a published POAP template.	**1** Before editing the POAP definition, save the settings in your Settings file. **2** On an existing and published POAP template, make desired changes to the template. **3** Click **Save As** to store the modified template with a different name. **4** Open the template. **5** Choose the **POAP** check box. **6** Apply the previously stored settings file. **7** Reboot the switch with no configuration or allowance to go into the POAP process. **8** Check that the template appears during the switch configuration. **9** Publish the new template.
A device remains in discovery mode.	The POAP script is not finished.	**1** Log in to the Cisco Prime DCNM web UI. **2** On the menu bar, choose **Config** > **Power-on Auto Provisioning (POAP)** > **DHCP Scope.** **3** Review the Bootscript Status column for error messages. **4** If there are no error messages, the Bootscript Status column should indicate that the POAP script is finished and the Bootscript Last Updated Time should be current.
	Incorrect access credentials were used in the POAP creation or editing process.	Make sure that the access credential provided in the UI during POAP creation or POAP editing is correct.
	The Management IP address provided in the template has not been learned from the uploaded configurations.	Make sure that the Management IP address provided in the POAP template is correct.

Symptom	Cause	Resolution
POAP process ends in a boot loop -> kick start and system images are interchanged in the definition file.	The kick start or system images are corrupt or bad.	Obtain new images and restart the POAP process.

Troubleshooting Inband Management and Inband POAP Issues

This section describes troubleshooting procedures for Inband Management and Inband POAP issues that are associated with a Cisco Dynamic Fabric Automation deployment.

Symptom	Cause	Resolution
Download failure due to wrong vrf.	Inband management not configured for vfs.	Within DCNM the POAP script must be run in order to change from management to default vfs.
POAP failure on Management Leaf no DHCP response.	Leaf not configured correctly upstream of DCNM.	The Leaf switch immediately upstream of DCNM (Management Leaf) must be connected to DCNM through an access port at native (10 Gig or 40 Gig) port speed.
POAP failure on Leaves and Spines with no DHCP response.	Using breakout cables between leafs and spines.	Breakout cables may not be used between leaves and spines, or between the Management Leaf and DCNM.
POAP failure on Management Leaf no DHCP response.	Cable plan is enforcement throughout the fabric.	Cable plan enforcement throughout the fabric must be disabled or cable plan files removed prior to running POAP.
POAP failures as well as VPC keepalive, LDAP query, visualization failures if POAP is bypassed.	DCNM is not connected to the Management Leaf. using a classic access port configured for VLAN management. Management Leaf not set to mode fabricpath.	DCNM must be connected to the Management Leaf using a classic access port configured for the Management VLAN, which must be set in mode **fabricpath** throughout the fabric. The Management IP must be configured on each switch via an SVI on the management VLAN.

CHAPTER 6

Troubleshooting Post-POAP, Autoconfiguration, and Device Discovery Issues

This chapter includes the following sections:

Troubleshooting Post-POAP Autoconfiguration and Device Discovery Issues

This section describes troubleshooting procedures for various issues that involve post-Power-on Auto Provisioning autoconfiguration and device discovery issues associated with Cisco Dynamic Fabric Automation (DFA).

Symptom	Cause	Resolution
The spine and leaf are not displayed correctly.	The device is not configured with the **fabric forwarding switch-role** command to identify the switch role, or the device does not support this feature. **Note** If the device does not have the role configured or the image does not support this feature, all Cisco Nexus 5000 Series and Cisco Nexus 6000 Series switches are categorized as leafs and Cisco Nexus 7000 Series switches are categorized as spines.	Use one of the following methods to configure the switch role: • Use the **fabric forwarding switch-role** [**border**] {**leaf** | **spine**} command to identify the switch role. • Log in to the DCNM Web UI. ◦ Choose **Dashboard** > **Dynamic Fabric Automation** > **Inter Switch Links View** ◦ Click **Override Switch Role** to change the role of a switch from leaf to spine or vice-versa.

Symptom	Cause	Resolution
The spine is not in FabricPath mode transit	Not all segments are instantiated on the spine.	**1** Instantiate all segments on the spine: • Layer 3 segment ID for routed traffic • Layer 2 segment ID for ARP, RARP, bridged-traffic **2** Use the **vni** command on the spine.

Symptom	Cause	Resolution
Cisco Prime DCNM displays fewer devices that have matched virtual machines (VMs) or virtual routing and forwarding (VRF) instances associated with them.	The VM or VRF is not associated with the device.	1 Use the **show evb host** command to verify that the VM is associated with the device. 2 Use the **show vrf** command to verify that the VRF is associated with the device.
	The device is pointing to an incorrect XMPP server.	1 Log in to the Cisco Prime DCNM web UI. 2 Click the **View Details** link to view more information. 3 Make sure that the device points to the correct XMPP server. 4 In the Cisco Prime DCNM web UI, on the menu bar choose **Admin** > **Settings.** 5 Make sure that DCNM uses the XMPP users specified on the DFA Settings page. 6 Make sure that DCNM and the device joined the same XMPP group.
	The device might be timed out.	1 Log in to the Cisco Prime DCNM web UI. 2 On the menu bar, choose **Admin** > **Settings.** 3 Verify if the XMPP response timeout specified on the DFA Settings page allows enough time.
	The XMPP server may be down.	Use the **appmgr status xmpp** command to ensure that the XMPP server is up and running.

Symptom	Cause	Resolution
Switches are not appearing in the Cisco DFA fabric.	The Cisco Discovery Protocol (CDP) is disabled on switches. Cisco DFA is unable to detect neighbor links.	**1** Review the switch configuration. **2** If CDP is not enabled, use the **cdp enable** command to enable CDP.
	Cisco DCNM has not completed discovery for the switches you expect to be in the fabric.	**1** Log in go the Cisco Prime DCNM web UI. **2** On the menu bar, choose **Inventory** > **Switches**. **3** Locate the switch and review the Status column. • If the Status column is OK, the switch is discoverable and reachable. • If the Status column is not OK, check relevant discovery logs at `/usr/local /cisco/dcm/fm/log/fmserver*.log*` to identify the root cause.
The Edge port view is not showing in virtual port channel (vPC) peers.	Cisco DCNM is not recognizing the vPC pairing.	**1** Log in to the Cisco Prime DCNM web UI. **2** Choose **Health** > **VPC**. **3** Look for missing vPC peers. **4** If DCNM does not recognize the vPC pairing, review the configuration on the vPC peers. **5** Use the **show vpc peer** command to view the output and determine the state of the switches.

Symptom	Cause	Resolution
A device remains in discovery mode.	The POAP script is not finished.	1 Log in to the Cisco Prime DCNM web UI. 2 On the menu bar, choose **Config** > **Power-on Auto Provisioning (POAP)** > **DHCP Scope.** 3 Review the Bootscript Status column for error messages. 4 Review the POAP log files in the switch bootflash to understand in what phase it failed. The files are `poap.log.*` and `.*poap.*init.log)` 5 If there are no error messages, the Bootscript Status column should indicate that the POAP script is finished and the Bootscript Last Updated Time should be current.
	Incorrect access credentials were used in the POAP creation or editing process.	Make sure that the access credentials that were provided in the UI during POAP Creation or POAP editing is correct.
	The Management IP address provided in the template has not been learned from the uploaded configurations.	Make sure that the Management IP address provided in the POAP template is correct.
Autoconfigurations are not restored after you used the **appmgr backup dcnm** and **appmgr restore dcnm** commands.	The autoconfiguration tenant information is stored in the database/LDAP/DHCP.	Use the **appmgr restore all** command to restore all applications.

Symptom	Cause	Resolution

The VLAN information is not available when you use the **show run output** command.	Configuration that is instantiated through autoconfiguration is not shown as part of **show run** output.	Use the **show running-config expand-port-profile** command to view autoconfigured VLAN information.
		For example:
		`>!Command: show running-config`
		`interface Vlan11`
		`expand-port-profile`
		`>!Time: Thu Aug 25 23:02:37 2011`
		`>`
		`>version 7.0(0)N1(1)`
		`>`
		`>interface Vlan11`
		`>  no shutdown`
		`>  vrf member Dept:Marketing`
		`>  no ip redirects`
		`>  ip address 11.1.1.1/24`
		`>  fabric forwarding mode`
		`proxy-gateway`
		`>  ip dhcp relay address`
		`10.1.1.100 use-vrf`
		`management`

POAP definition changes are not being reflected in configurations on a device.	The configuration is applied only during the next reload and POAP for the device.	There are two options: **1** Apply a configuration at run time using XMPP. **a** Ensure that all of the switches have joined a fabric access XMPP group. **b** Send the command to all devices in the group using the XMPP messenger. **2** Push a configuration to a switch that is up and running. **a** Log in to the Cisco Prime DCNM Web UI. **b** On the menu bar, choose **Config** > **Templates.** **c** Select or create a config template. **Note** For example, if you want to add a single configuration, only include the single configuration in the template. **d** Launch the job creation wizard. **e** Choose the switch. **f** Submit or schedule the job.

Symptom	Cause	Resolution
A device cannot execute a downloaded network profile successfully.	Basic network profiles are not on the device.	• Make sure that VRF of Org: Partition configuration profiles (such as vrf-tenant-profile, vrf-common, vrf-common-services) are already on the device. • In most cases, the basic network profiles must be included in user-defined network profiles.
	Incorrect values have been entered for the Profile parameters, or there is a mismatch between the values entered for some primary fields such as segmentId, vlanId, etc.,	Possible errors: • The Mobility-Domain was entered instead of the segmentID • An incorrect netMaskLength was used. For example, 255.255.255.0 or /24 instead of 24.
	Profile parameter values are missing.	• From LDAP, verify that all the profile parameters (configArgs) are filled in, including "$include_vrfSegmentId" and "dhcpServerAddr". • If the service IP address is specified for the partition, "$include_serviceIpAddress" should also be part of configArgs. • Verify that the segment ID in configArgs, for example $segmentId, has the same value as is specified for the network.
	The VRF name is inconsistent in profile parameters for all networks.	

Symptom	Cause	Resolution
		Make sure that the VRF name is entered consistently in the profile parameters for all networks in that VRF. The format should be "organizationName:partitionName". **Note** Inconsistencies in the name can result in inconsistent network behavior due to a mismatch in VRF to segment-ID mapping.
Address Resolution Protocol (ARP) address is not hitting leaf.	No ARP is detected.	Use the **debug ip arp xxx** command to review why the ARP is not hitting the leaf.
ARP is hitting, but profile does not get instantiated.	No connectivity to the database server exists	Use the **show fabric database statistics** command to display fabric database statistics.

CHAPTER 7

Troubleshooting Prime Network Services Controller Issues

This section contains the following sections:

Troubleshooting Integration Issues

If you encounter issues with the and DCNM integration, you can look for information in the following locations:

- On the DCNM server, review the log files in /opt/nscadapter/var/log for information.
- In the GUI:
    - Review faults for services by choosing **Resource Management** > **Managed Resources** > **root** > *tenant* > **Network Services** > *network-service* > **Edit** > **Faults** tab.
    - Review audit logs and faults by choosing **Resource Management** > **Diagnostics** > **Audit Logs** or **Faults**.

For either option, double-click a fault to view more information.

The following table describes specific issues that you might encounter and how to address them:

Symptom	Cause	Resolution
Organizations and partitions are created in DCNM but no tenants or virtual device contexts (VDCs) are displayed in	The configurations in DCNM and are incomplete.	1 Confirm that the Service Configuration parameters are complete for networks created in DCNM. 2 Confirm that is registered with the VM Manager IP parameter.

Symptom	Cause	Resolution
Networks are created in DCNM but no tenants, VDCs, or subnetworks are displayed in .	The Network Services Controller (NSC) Adapter does not have an active connection to .	Use the **nsc-adapter-mgr adapter connections** command to ensure there is an active connection to .
	The NSC Adapter is not active on DCNM.	Use the **nsc-adapter-mgr adapter connections** command to ensure there is an active connection to DCNM.
	does not have the VM Manager IP.	Confirm that is registered with the correct VM Manager and provide the VM Manager IP address in the VM Manager IP parameter.
	Networks were added to DCNM while or the NSC Adapter was down.	**1** Enter the command **nsc-adapter-mgr adapter connections** and verify that the connections are correct. **2** In the DCNM GUI, choose the auto-config interface, choose the network, click **Edit**, and then click **OK** without making changes.
Service networks were deleted in DCNM but the tenants, VDCs, and subnetworks are still shown in .	Networks were deleted from DCNM while or the NSC Adapter was down.	**1** Enter the **nsc-adapter-mgr adapter connections** command and verify that the connections are correct. **2** In the DCNM GUI, choose the auto-config interface, choose the network, click **Edit**, and then click **OK** without making changes.
An edge service was removed from but the Service Node IP Address is still shown in DCNM.	The service was deleted from while DCNM or the NSC Adapter was down.	Manually delete the Service Node IP Address in DCNM for the affected partition.
An edge service was deployed in but the Service Node IP Address is not shown in DCNM.	The service was deployed in while DCNM or the NSC Adapter was down.	Manually update the Service Node IP Address in DCNM auto-config for the affected partition.

Symptom	Cause	Resolution
Host traffic does not reach the service node.	• The wrong profile is specified in DCNM for host networks. • The service is not attached to the leaf.	• Make sure that the correct profile is specified in DCNM for the host network. • Make sure that the auto-config profile and parameters are correct with particular attention to the Service Node IP address.

CHAPTER 8

Decoding Leaf-Spine Traffic

This chapter contains the following sections:

Decoding Leaf-Spine Traffic

You can decode a leaf-spine packet capture using Wireshark.

Before You Begin

You must have version 1.8 of Wireshark installed. Download Wireshark at www.wireshark.org.

Step 1 In Wireshark, on the **Analyze** menu, choose **Enabled Protocols**.
The Enabled Protocols dialog box opens.

Step 2 Enable **Cisco FabricPath protocol (CFP)**.

Step 3 Click on the unknown data following ethertype 893b.

Step 4 On the **Analyze** menu, choose **Decode As** and choose **VLAN**.
This forces Wireshark to decode ethertype 893b as VLAN (dot1q).

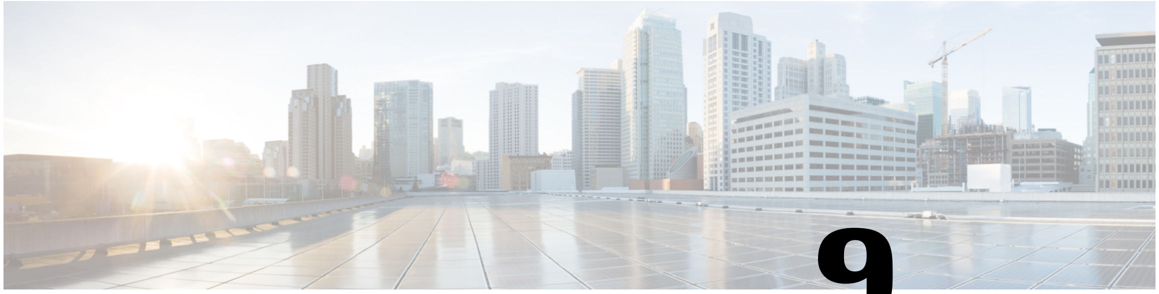

CHAPTER 9

Debugging Tenant Traffic

This chapter contains the following sections:

Debugging Tenant Traffic

To verify the VLAN to VN-Segment mapping use the **show platform fwm info vdc all verbose | begin fwm_avl_vlan_tree_by_vni** command.

To verify the if programming works fine, use the **show platform fwm info qinq-xlate-table <asic-num>** command as this will show the mapping from the internal VLAN to VN-Segment mapping and vice-versa.

The following example shows how to verify that both the VLAN and QinQ values are pointing to the same internal VLAN:

```
switch# show platform fwm info  xlate-vlan-table  1 | grep " 200 "
Dir  Xlate-idx Key-vlan Res-vlan Ref-count Masked Location     is_l2_if
Ig   17        200      199      1         no     1.784.0     1
Eg   17        199      200      1         no     1.3262.0    1


switch# show platform fwm info qinq-xlate-table  1  | grep “ 200 "
Number of xlate containers pending PSS: 0
Dir  Xlate-idx Key-vlan Res-vlan Ref-count Masked Location     is_l2_if
Eg   17        199      20000    1         no     1.3024.0     1
Ig   17        20000    199      1         no     1.3189.0     1
```

For certain VLAN and VN-Segment mapping, the VLAN is seen from the server-side and hence it is xlate-vlan-table and QinQ is seen from the fabric-side and hence it is qinq-xlate-table.

In the above example of xlate-vlan-table, 'Res-vlan' in the Ig direction is the internal context that ASIC uses to forward server traffic of a tenant. In the CLI output of qinq-xlate-table, 'Res-vlan' in the Ig direction is the forwarding context that ASIC uses to forward FabricPath traffic of a tenant. It is required that, for a VLAN with VN-Segment, both the VLAN in the Ig direction and VN-Segment in the Ig direction should point to the same 'Res-vlan'.

The following example shows how to determine the ASIC number:

```
switch# show platform fwm info  pif  ethernet 1/1 | grep asic
Eth1/1 pd: slot 0  logical port num 0  slot_asic_num 1  global_asic_num 1 fw_ins t 0
phy_fw_inst 0 fc 0
```

The global_asic_num value is 1 for ethernet 1/1 in the above example.

Note The VLAN and VN-Segment programming will be global and will be programmed symmetrically in all the ASICs.