Cisco Dynamic Fabric Automation Solution Guide

Last Modified: August 08, 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:

CONTENTS

Preface

The Preface contains the following sections:

Audience

This publication is for network administrators who configure and maintain Cisco Nexus devices.

Document Conventions

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.

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.
`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.
!, #	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/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.

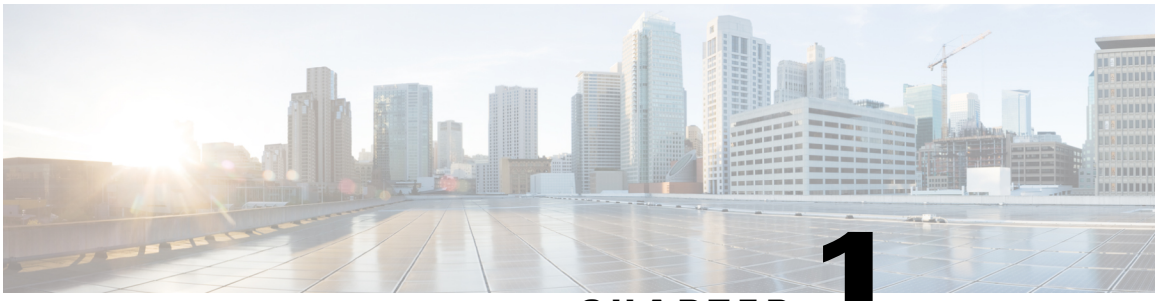

CHAPTER 1

New and Changed Information

This chapter includes the following sections:

New and Changed Information

The table below summarizes the new and changed features for this document and shows the releases in which each feature is supported. Your software release might not support all the features in this document. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release.

Table 1: New and Changed Information

Date	Description	Where Documented
December 18, 2014	Added a new section on the Border Leaf and DC Edge router auto-configuration for Layer 3 Interconnect. This section describes the capabilities that come with Cisco Prime DCNM 7.1.1 and the Cisco NX-OS 7.1(0)N1(1) release in the context of Border Leaf Auto Configuration. Added explanation for new features in Cisco NX-OS 7.1(0)N1(1) (multi-mobility domain, Dynamic Virtual Port, and Universal Profiles).	"Information About Cisco DFA" chapter

Date	Description	Where Documented
July 7, 2012	Updated licensing information for the Dynamic FCoE Using DFA feature on the Cisco Nexus 5600 Series and 6000 Series switches in the "Licensing Requirements for Cisco DFA" section. Updated for Cisco NX-OS 6.2(6b) in the “Platform Requirements” section.	“Deploying Cisco DFA” chapter
May 19, 2014	Updated licensing information for the Cisco Nexus 5600 Series and 6000 Series switches as a spine or leaf switch in the "Licensing Requirements for Cisco DFA" section.	“Deploying Cisco DFA” chapter
April 25, 2014	• Updated for Cisco NX-OS Release 7.0(2)N1.(1). • Added support Cisco Nexus 5600 Series switches as a spine, leaf, border leaf, and route reflector to the "Platform Requirements" and "Licensing Requirements for Cisco DFA" sections. • Added information about the CLI-based auto configuration for Cisco Nexus 55xx Series switches.	“Deploying Cisco DFA” chapter “CLI-Based Auto Configuration” section in the “Information About Cisco DFA” chapter
April 1, 2014	Added Cisco Nexus 5500 Series switches to the "Licensing Requirements for Cisco DFA" section.	“Deploying Cisco DFA” chapter
March 27, 2014	Added support for Cisco Nexus 5600 Series switches as a spine, leaf, border leaf, and route reflector to the "Platform Requirements" and "Licensing Requirements for Cisco DFA" sections.	“Deploying Cisco DFA” chapter

Date	Description	Where Documented
February 26, 2014	Added support for Cisco Nexus 7000 Series switches as a route reflector to the "Platform Requirements" and "Licensing Requirements for Cisco DFA" sections.	“Deploying Cisco DFA” chapter
February 13, 2014	Added Cisco Nexus 5500 Series switches to the "Platform Requirements" section.	“”Deploying Cisco DFA” chapter
January 31, 2013	This book was created for Cisco Dynamic Fabric Automation (DFA) 1.0.	—

New and Changed Information

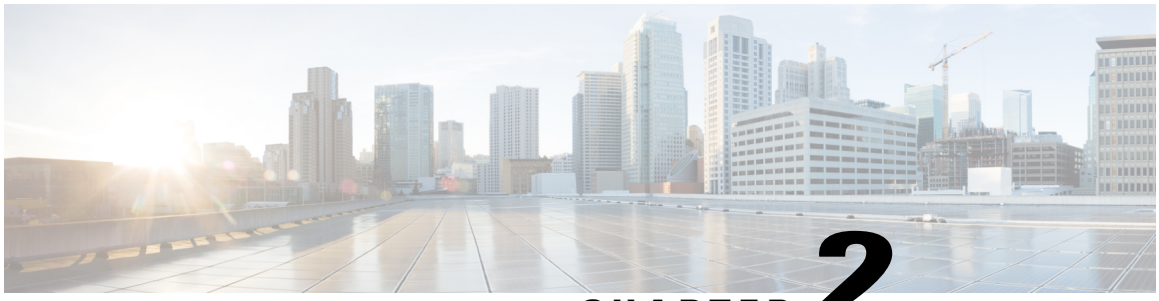

CHAPTER 2

Information About Cisco DFA

This chapter includes the following sections:

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter.

Terminology

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

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

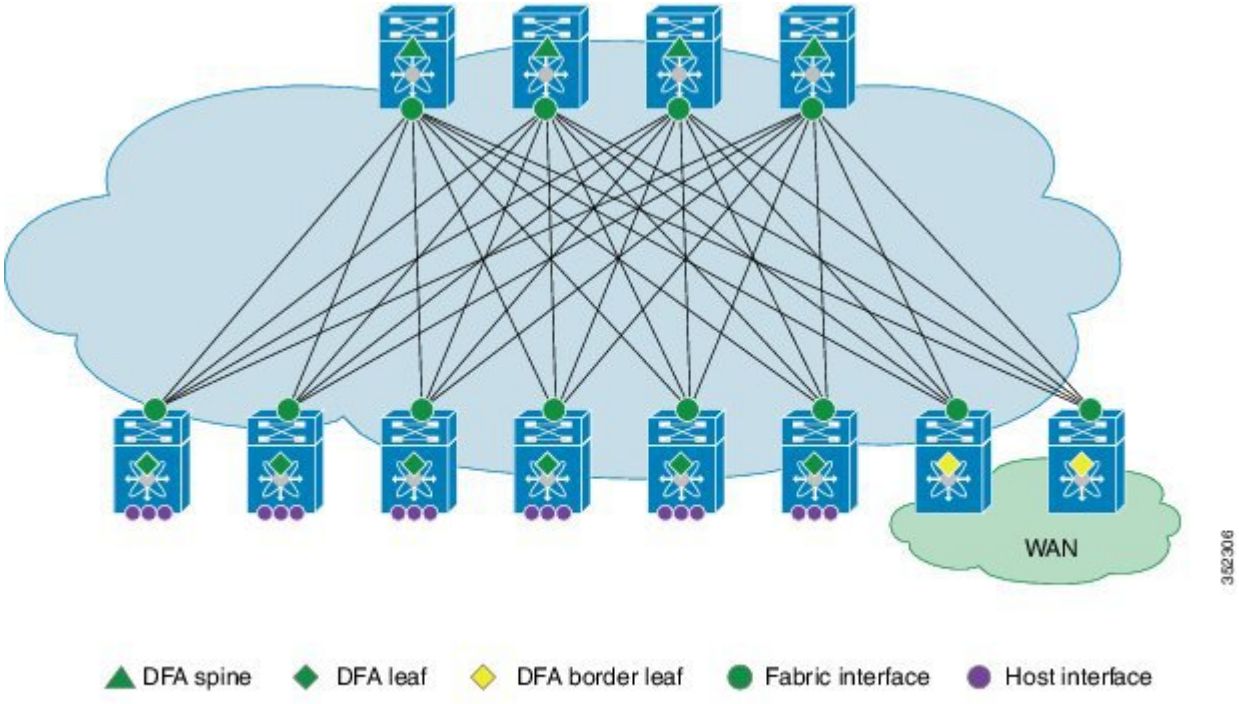

- Cisco DFA fabric—A multistage, switching network in which every connected device is reachable through the same number of hops. The Cisco DFA fabric enables the use of a scale-out model for optimized growth.
- Cisco DFA switch—A leaf, border leaf, or spine device.
- Leaf—Switches with ports that are 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 Unified Fabric Automation Overview

Cisco Unified 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 Unified Fabric Automation:

- 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/or 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/or multitenancy 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. This 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 on 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 Unified 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 2: 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 Unified Fabric Automation. Cisco Prime Data Center Network Manager (DCNM) as an OVA can be deployed on a VMware vSphere infrastructure.

Cisco Prime DCNM provides the following functionality:

- Device auto configuration is the process of bringing up the Cisco Unified Fabric Automation 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 touch less 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

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 DCNM. This makes it possible for a complete tenant VM orchestration with automated network provisioning to be absolutely touch-less from the fabric perspective. Refer to the DFA Configuration Guide for details on tenant provisioning.

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 touch-less 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. Refer to the DFA Configuration Guide for configuration details.

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 Dynamic Fabric Automation (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 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 DCNM. At the 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 DCNM, the auto-configuration can be globally enabled at 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 DCNM itself. In Cisco NX-OS 6.2(10) release for Nexus 7000 platform, the configuration can be generated on 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 DCNM and N7000. This feature is driven by 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 Dynamic Fabric Automation (DFA) uses a simple distributed gateway mechanism to support any subnet, anywhere, concurrently.

Frame Encapsulation

Optimized networking in a Cisco Dynamic Fabric Automation (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 Layer2/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 short 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 Unified 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 Unified Fabric Automation 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 Unified Fabric Automation deployment through Cisco Data Center Network Manager (DCNM).

Figure 2: Cisco Unified Fabric Automation with Services

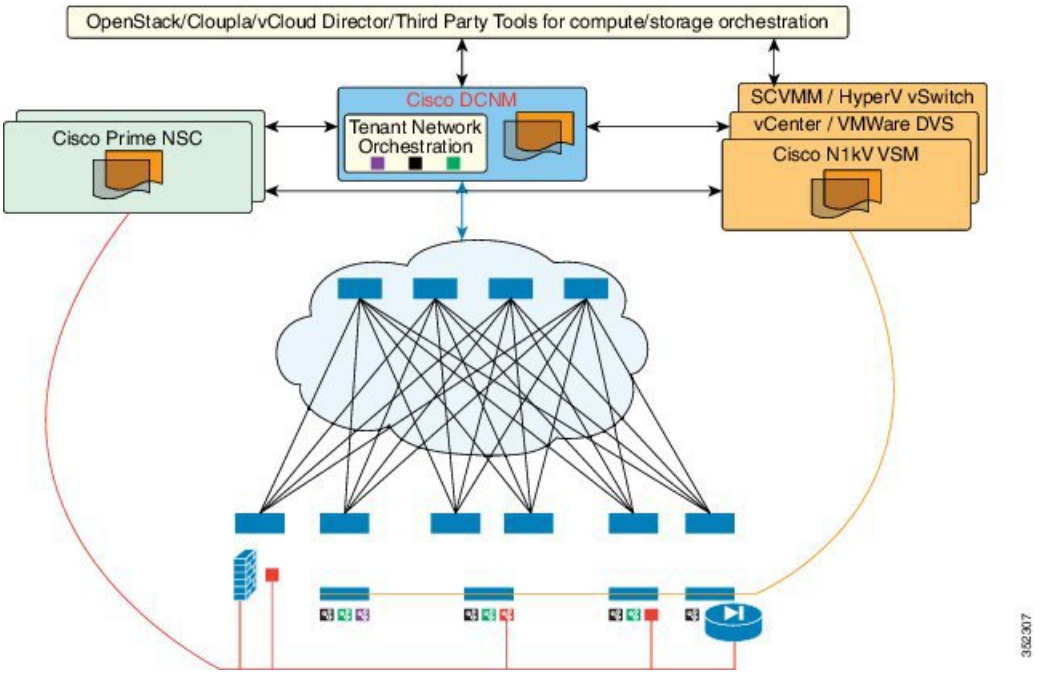

The Cisco Prime NSC is the services orchestrator for Cisco Unified Fabric Automation. The NSC Adapter in the Cisco Prime DCNM Open Virtual Appliance (OVA) performs the following functions:

- 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 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 DCNM through the Rest APIs. Note that interaction between PNSC and DCNM is not needed for this approach; UCSD implements this functionality for services.

In Cisco Unified Fabric Automation, configuration profile templates and instantiating the profiles on a leaf switch provide network automation. The templates are extended to support services in Cisco Unified Fabric Automation. 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 Unified Fabric Automation services. It is important that you select the correct profile to orchestrate and automate services in the Cisco Unified Fabric Automation fabric.

Table 3: 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 interrelated 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

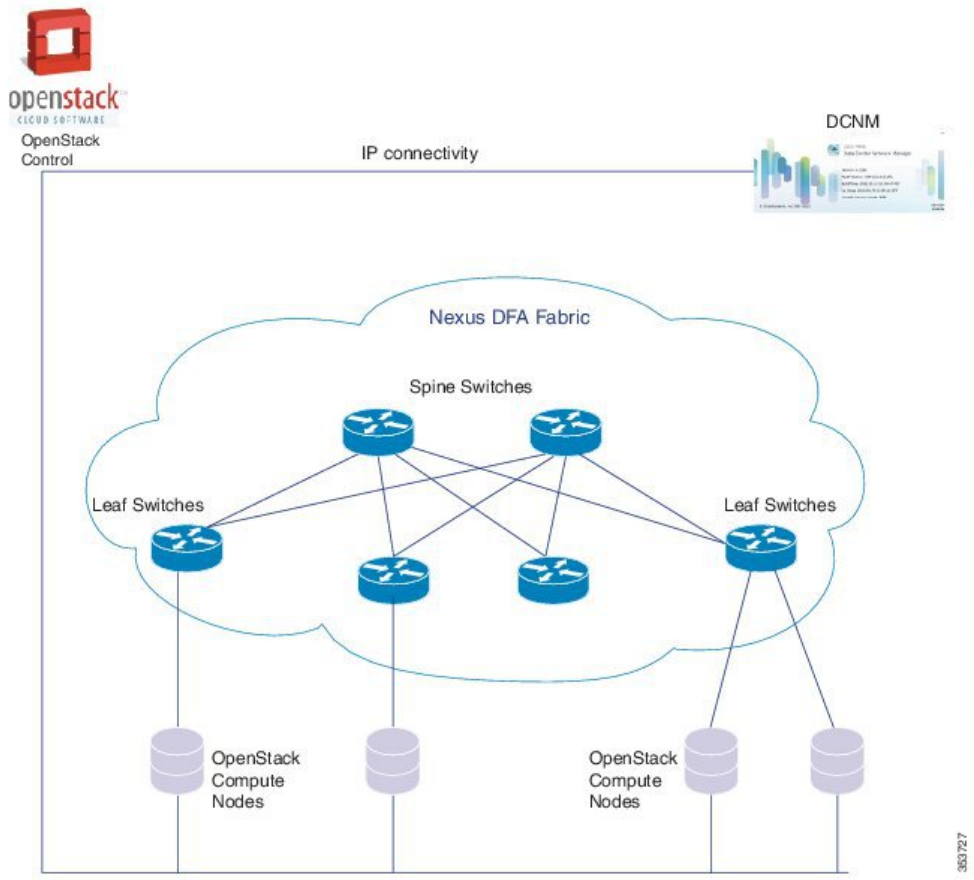

For information about Open Source used in OpenStack for Cisco DFA 2.0, see the Open Source used in *OpenStack for Cisco DFA 2.0* document.

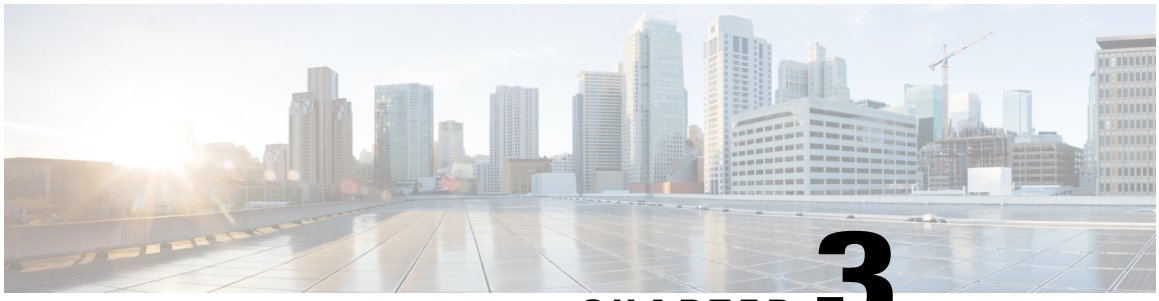

CHAPTER 3

Deploying Cisco DFA

This section describes how to deploy Cisco Dynamic Fabric Automation (DFA).

This section includes the following topics:

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter.

Platform Requirements

Table 4: Cisco Dynamic Fabric Automation Platform Support

Product	Function					Software Release (and later releases)
	Spine	Leaf	Border-Leaf	RR[3]	Other	
Cisco Nexus 6001 Series switch	Yes	Yes	Yes	Yes		Cisco NX-OS Release 7.0(0)N1(1)

Product	Function					Software Release (and later releases)
	Spine	Leaf	Border-Leaf	RR[3]	Other	
Cisco Nexus 6004 Series switch	Yes	Yes	Yes	Yes		Cisco NX-OS Release 7.0(0)N1(1)
Cisco Nexus 7000 Series switch	Yes[1]	—	—	Yes		Cisco NX-OS Release 6.2(6b)
Cisco Nexus 5672UP switch	Yes	Yes	Yes	Yes		Cisco NX-OS Release 7.0(1)N1(1)
Cisco Nexus 56128 switch	Yes	Yes	Yes	Yes		Cisco NX-OS Release 7.0(2)N1(1)
Cisco Nexus 5596UP switch	—	—	—	—	Layer 2-only leaf	Cisco NX-OS Release 7.0(0)N1(1)
						Cisco NX-OS Release 7.0(2)N1(1) supports the CLI-based auto configuration option
Cisco Nexus 5548P and 5548UP switches	—	—	—	—	Layer 2-only leaf	Cisco NX-OS Release 7.0(0)N1(1)
						Cisco NX-OS Release 7.0(2)N1(1) supports CLI-based auto configuration
Cisco Nexus 1000V switch for VMware vSphere 5.1 and 5.5	—	—	—	—	Virtual switch with VDP signaling	Cisco NX-OS Release 4.2(1)SV2(2.2)
Cisco Prime Data Center Network Manager (DCNM)	—	—	—	—	Fabric management	Cisco Prime Data Center Network Manager Release 7.0
Cisco Prime Network Services Controller (NSC)	—	—	—	—	Services support	Release 3.2
OpenStack for Cisco DFA	—	—	—	—	(Optional) - Controller	OpenStack for Cisco DFA 1.0

Note

1 With Cisco Nexus 7000 F2, F2e, and F3 Series modules.

2 With Cisco Nexus 7000 F3 Series module.

3 Cisco DFA requires a minimum of one multiprotocol BGP route-reflector (RR). As an integrated function of Cisco DFA, the following platforms can support this function:

- Nexus 6000 Series switches with Cisco NX-OS Release 7.0(0)N1(1) and later releases
- Nexus 5600 Series switches with Cisco NX-OS Release 7.0(1)N1(1) and later releases
- Nexus 7000 Series switches with Cisco NX-OS Release 6.2(6b) and an MPLS feature, grace period, or evaluation license

Licensing Requirements for Cisco DFA

Review the other hardware and software components of your existing fabric with respect to the Cisco Dynamic Fabric Automation (DFA) release requirements and compatibility constraints. Because Cisco DFA implements an architectural solution with a switch topology that is different from what you have previously used, devices might be required to perform different roles when used in a Cisco DFA implementation, and might be subject to new licensing requirements. For more information, see the "Platform Requirements" section of this guide.

Product	License Requirement
Cisco Nexus 5500 Series switches	Cisco DFA requires the FabricPath Services package (ENHANCED_LAYER2_PKG) license. For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply licenses, see the *Cisco NX-OS Licensing Guide*.

Product	License Requirement
Cisco Nexus 5600 Series switches, including the 5672UP and 56128 switches	• Cisco DFA requires the FabricPath Services package (ENHANCED_LAYER2_PKG) license. • For a Cisco Nexus 5672UP or 56128 switch as a Cisco DFA spine or leaf switch, the Enterprise Services Package (LAN_ENTERPRISE_SERVICES_PKG) is required. • For a Cisco Nexus 5672UP or 56128 switch as a Cisco DFA spine or leaf switch, the Layer 3 Base Services Package (LAN_BASE_SERVICES_PKG) is required. • For Dynamic Fibre Channel over Ethernet (FCoE) using DFA on a Cisco Nexus 56xx switch, the Storage Protocols Services Package (FC_FEATURES_PKG, ENTERPRISE_PKG) is required. For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply licenses, see the *Cisco NX-OS Licensing Guide*.
Cisco Nexus 6000 Series switches	• Cisco DFA requires the FabricPath Services package (ENHANCED_LAYER2_PKG) license. • For a Cisco Nexus 6000 Series switch as a Cisco DFA spine or leaf switch, the Enterprise Services Package (LAN_ENTERPRISE_SERVICES_PKG) is required. • For a Cisco Nexus 6000 Series switch as a Cisco DFA spine or leaf switch, the Layer 3 Base Services Package (LAN_BASE_SERVICES_PKG) is required. • For Dynamic Fibre Channel over Ethernet (FCoE) using DFA on a Cisco Nexus 6xxx switch, the Storage Protocols Services Package (FC_FEATURES_PKG, ENTERPRISE_PKG) is required. For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply licenses, see the *Cisco NX-OS Licensing Guide*.

Product	License Requirement
Cisco Nexus 7000 Series switches	• Cisco DFA requires the FabricPath Services package (ENHANCED_LAYER2_PKG) license. • For a Cisco Nexus 7000 Series switch as a Cisco DFA spine switch, the Enterprise Services Package (LAN_ENTERPRISE_SERVICES_PKG) is required. • For a Cisco Nexus 7000 Series switch as a Cisco DFA route reflector, the MPLS Services Package (MPLS_PKG) license is required. Starting with NX-OS 6.2(10) the requirement for MPLS license and feature-set is removed. The Multi-Protocol BGP VPNv4/VPNv6 AFI are being available with a new feature called "fabricpath-vpn". For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply licenses, see the *Cisco NX-OS Licensing Guide*.
Cisco Prime Data Center Network Manager (DCNM)	**Note** The switch feature licenses must be installed before you install the Cisco Prime DCNM license. Cisco DFA features and capabilities are covered by the Cisco DCNM Base license. The basic unlicensed version of Cisco DCNM-SAN Server is included in the software download. To get licensed features, such as Performance Manager, remote client support, and continuously monitored fabrics, you must buy and install the Cisco DCNM-SAN Server package. For information, see the *Cisco DCNM Installation and Licensing Guide, Release 7.x*.

Guidelines and Limitations for Cisco Unified Fabric Automation

Cisco Unified Fabric Automation has the following guidelines and limitations:

- The fabric management network can support only one Dynamic Host Configuration Protocol (DHCP) server. You can use either the DHCP server in Cisco Prime Data Center Network Manager (DCNM) or another designated DHCP server, but not both.
- To ensure that Cisco Unified Fabric Automation device auto configuration does not interfere with other DHCP servers on your network, we recommend that you use a dedicated VLAN and subnet for the fabric management network. Cisco Prime DCNM and the Ethernet out-of-band ports of the Cisco Unified

Fabric Automation switches (mgmt0) reside in the fabric management network. You have the option to interconnect the fabric management network with your existing out-of-band management network.

- The management connectivity for Cisco Unified Fabric Automation must come through the Cisco NX-OS device management interface (mgmt0).
- The management port on any Cisco Unified Fabric Automation switch must be connected to the same management subnet that includes the Cisco Prime DCNM user interface.
- Every Cisco Unified Fabric Automation switch to be managed by fabric management must be connected to the fabric management network through the Ethernet out-of-band network.
- A console connection for fabric management is recommended but not required for Cisco Unified Fabric Automation.
- If Cisco Prime DCNM is your repository server, you must upload the Cisco NX-OS kickstart and system images to Cisco Prime DCNM using the Serial Copy Protocol (SCP) or Secure File Transfer Protocol (SFTP).

How to Cable the Network Fabric and Servers for Cisco DFA

Fabric Management Network and Console

Every Cisco DFA switch that is to be managed by Cisco Dynamic Fabric Automation (DFA) fabric management must connect to the fabric management network through the Ethernet out-of-band port (mgmt0).

Figure 4: Cabling the Fabric Management Network

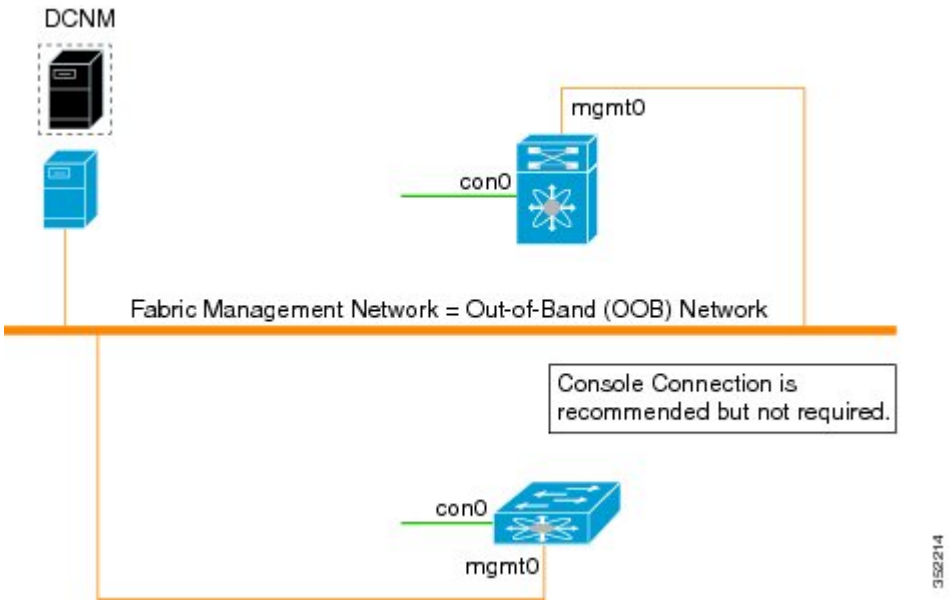

Fabric Connectivity

The fabric interfaces of the Cisco Dynamic fabric Automation (DFA) fabric connect the Cisco DFA switches to one another. Fabric interfaces are configured with Cisco FabricPath Frame Encapsulation (FE) for efficient forwarding based on a Shortest Path First (SPF) algorithm. You do not configure VLAN trunking or pruning for the transported VLANs on Cisco DFA fabric interfaces .

Figure 5: Cabling the Cisco DFA Network Fabric and Servers

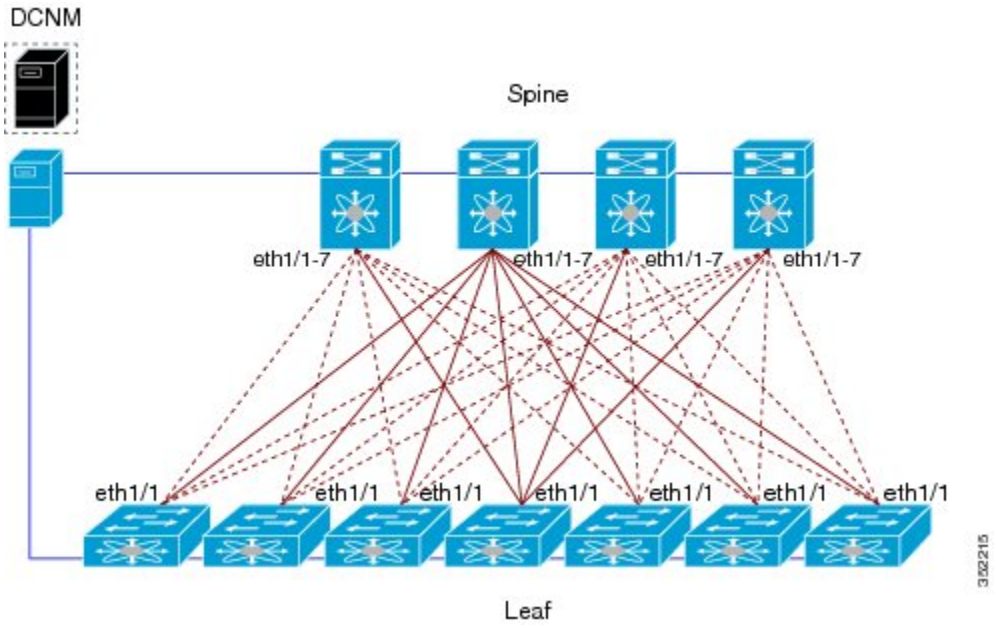

Server Connectivity

To transport data traffic across the Cisco Dynamic Fabric Automation (DFA) Fabric, the leaf switch must receive the traffic for connected VLANs that are to be extended across the fabric. The leaf-to-server interfaces are called host interfaces.

Note Always connect servers to Cisco DFA leaf or border leaf switches. You must not connect servers to Cisco DFA spine switches.

Figure 6: How to Cable Server Connectivity

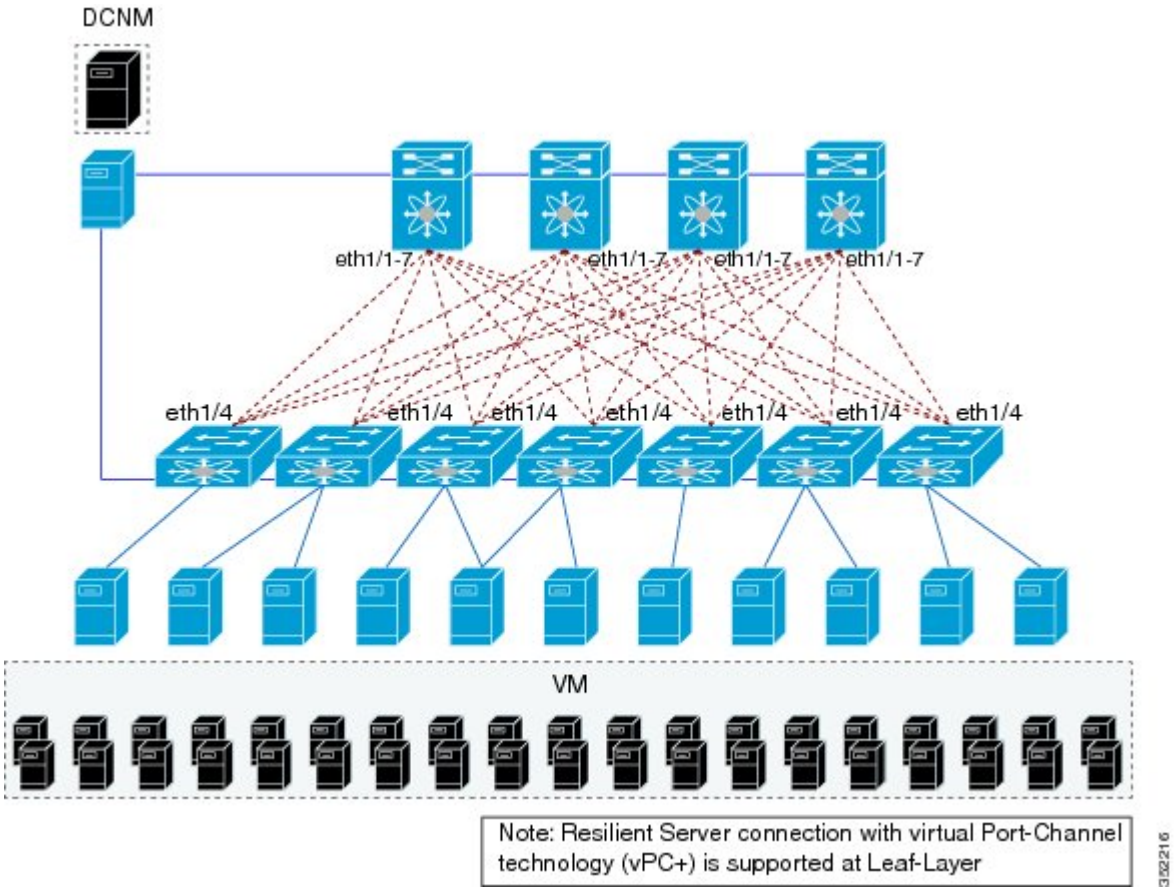

Fabric Management

The Cisco Prime Data Center Network Manager (DCNM) is the central point of management for Cisco DFA.

Figure 7: Preparing to Deploy Cisco Prime DCNM

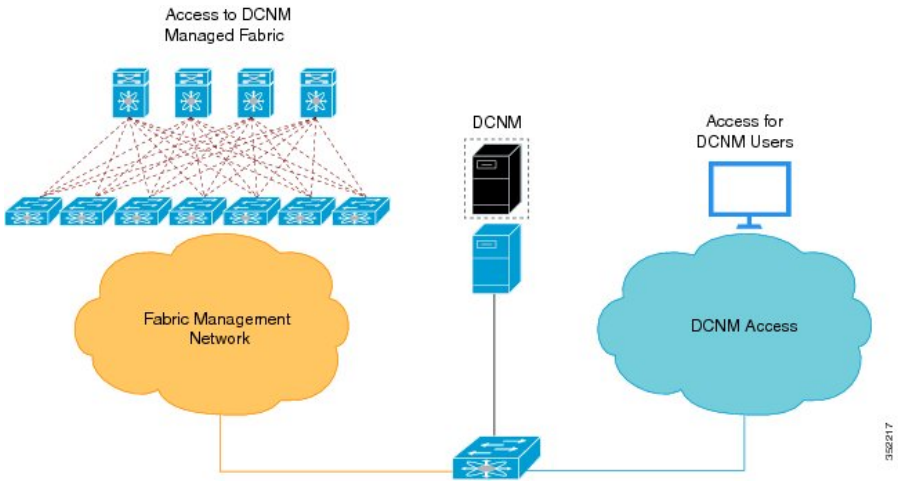

Deploying Cisco DFA

Note If this is not a new Cisco Dynamic Fabric Automation (DFA) deployment, see the Cisco Dynamic Fabric Automation Migration Guide for migrating your existing fabric to a Cisco DFA deployment.

1. Ensure that you have the appropriate Cisco Nexus devices with the minimum required Cisco NX-OS software releases to support Cisco Dynamic Fabric Automation DFA). See the "Platform requirements" section of this guide.
2. Install the Data Center devices. For information, see the appropriate install guides for your Cisco DFA switches.
3. Install and configure the Cisco Nexus 1000V switch for VMware vSphere for Cisco DFA. For information, see the Cisco Nexus 1000V Installation and Upgrade Guide and theCisco Nexus 1000V DFA Configuration Guide.

Note To deploy Cisco Prime DCNM, two port groups or port profiles are required on the virtual switch.

4. Create a cabling plan and cable your Cisco Nexus devices for Cisco DFA. For information, see the "How to Cable the Network Fabric and Servers for Cisco DFA" section of this guide.
5. Install the Cisco Prime Data Center Network Manager (DCNM) Open Virtual Appliance (OVA) to manage all the applications for the central point of management. For information, see the Cisco DCNM 7.0 OVA Installation Guide.

6 Start the Prime NSC adapter in the Cisco Prime DCNM OVA and configure Services support for Cisco DFA. For information, see the "Network Services" section of the Cisco DCNM 7.0 OVA Installation Guide.

7 (Optional) Use one of the following options to install OpenStack for Cisco DFA:

a Install the Cisco OpenStack Installer to install the OpenStack for Cisco DFA orchestrator. For information, see the OpenStack for Cisco DFA Install Guide Using Cisco OpenStack Installer.

Note

- Before installing the Cisco OpenStack installer, the Cisco DFA fabric, switches, and Cisco Prime DCNM OVA must be already installed.
- To support OpenStack for Cisco DFA, Cisco Prime DCNM must be accessible via the OpenStack controller and the Cisco DFA fabric.

b Use the pre-built OpenStack for Cisco DFA images to install the OpenStack for Cisco DFA orchestrator. For information, see the following guides:

- OpenStack for Cisco DFA Install Guide for Using Pre-built OpenStack for Cisco DFA Images
- Quick Guide to Clonezilla