



Information About Cisco DFA


This chapter includes the following sections:

- [Terminology, page 2](#)
- [Cisco Dynamic Fabric Automation Overview, page 3](#)
- [Fabric Management, page 3](#)
- [Automated Network Provisioning, page 5](#)
- [Optimized Networking, page 6](#)
- [Dynamic VLAN Management, page 7](#)
- [Cisco Dynamic Fabric Automation Services Support, page 7](#)
- [OpenStack for Cisco DFA, page 9](#)

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	<ul style="list-style-type: none"> • 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

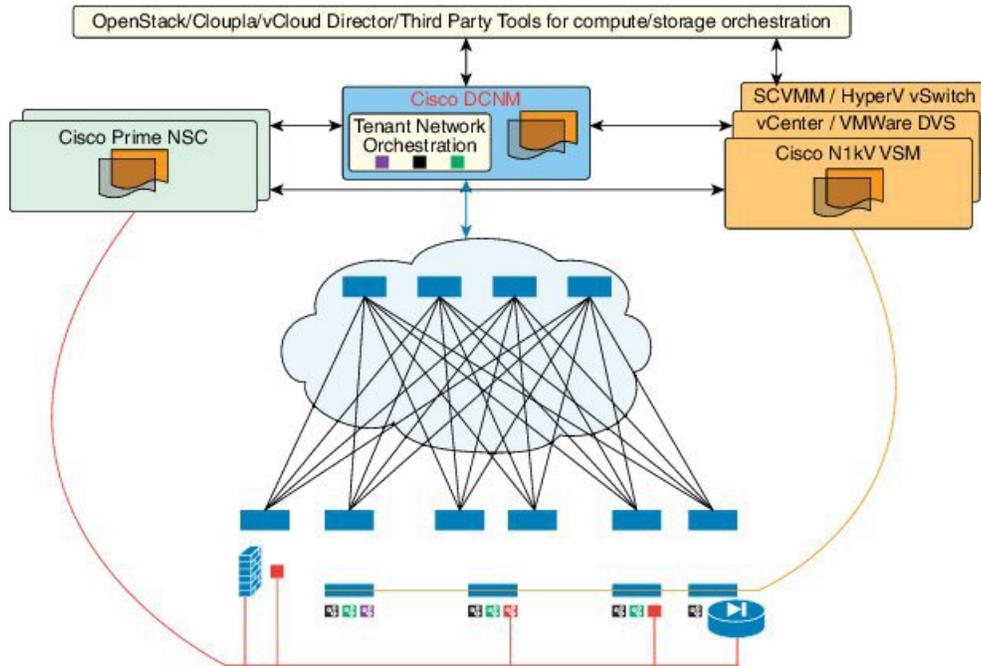
Managing VLANs that are used to interact with the servers is always complicated due to the need for more than 4K tenants. Fabric dynamic VLAN allocations can solve this problem. With a VDP-capable Vswitch, leafs can communicate with Vswitch using VDP and discover the presence of VMs. VDP can communicate 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



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

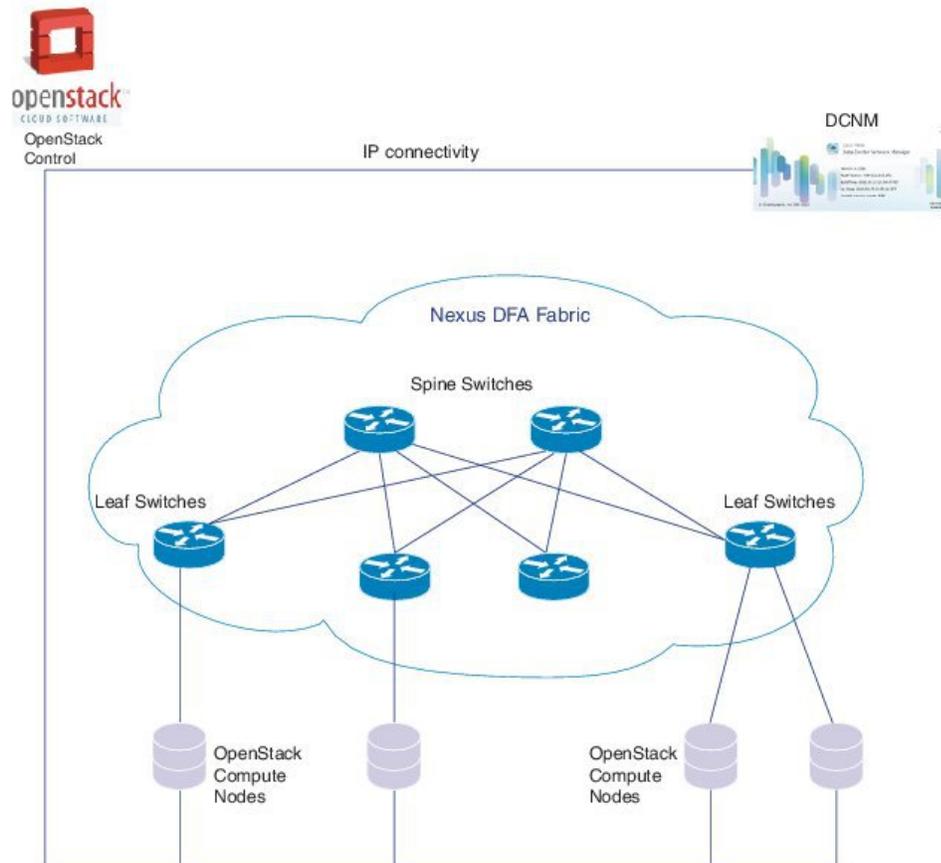
OpenStack creates a human and machine-accessible service for managing the entire life cycle of the infrastructure and applications within OpenStack clouds. The technology consists of a series of inter-related projects that control pools of processing, storage, and networking resources throughout a data 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.