Cisco Unified Fabric Innovations

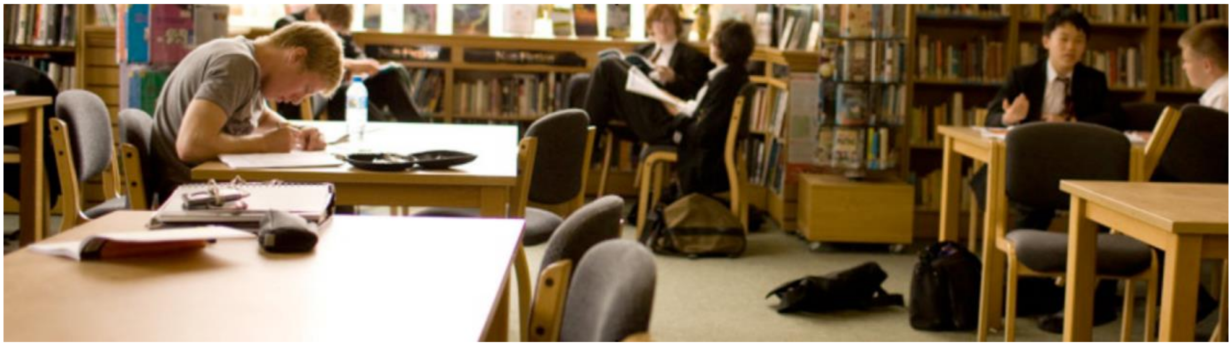

This Q&A document consists of five sections, summarized in Table 1. However, a given question and its answer may overlap across several sections. Please use the Find or Search option (Ctrl+F in Microsoft Windows or ⌘+F on a Mac) of your document viewer for quickest results.

For acronyms and abbreviations, please also refer to the document located at http://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/dynamic-fabric-automation/white-paper-c11-732230.html.

Table 1. Cisco Unified Fabric Innovations Q&A Topics

Section	Description
Fabric Management	Questions related to general and specific aspects of setup, maintenance, and operations of the fabric components
Optimized Networking	Questions related to data forwarding and data path creation within the fabric
Virtual Fabrics	Questions related to features that promote multitenancy capabilities of the fabric
Workload Automation	Questions related to features that enable automated workload instantiation, network provisioning, and the purging of obsolete configurations from the fabric
Miscellaneous	Questions related to management tools installation, licensing, etc.

Fabric Management

Q. What is Cisco® Prime™ Data Center Network Manager (DCNM) and what functions does it serve in Cisco Unified Fabric?

A. Cisco Prime DCNM serves as a single pane from which you can monitor, provision, and control various aspects of your Cisco Unified Fabric network. For more information, see http://www.cisco.com/c/en/us/products/cloud-systems-management/prime-data-center-network-manager/index.html.

Q. Is Cisco Prime DCNM required to run Cisco Unified Fabric innovations? Can I just use the command-line interface (CLI)?

A. Cisco Prime DCNM is an important tool that simplifies multiple aspects of fabric management. Most tasks performed by the Cisco Prime DCNM prepackaged tools can be accomplished manually, either using the CLI or with open source applications (Lightweight Directory Access Protocol [LDAP], Extensible Messaging and Presence Protocol [XMPP], etc.). However, for the most simplified management, you should use Cisco Prime DCNM.

Q. What is the Cisco Prime Network Services Controller (NSC), and what is its relationship to Cisco Unified Fabric?

A. Cisco Prime NSC is software for centralized management of multiple devices and policies for network virtual services. For more information, see http://www.cisco.com/c/en/us/products/cloud-systems-management/prime-network-services-controller/index.html. Cisco Prime NSC can integrate with Cisco Unified Fabric to automatically provision virtual services in response to workload deployment in the network.

Q. What is a fabric management network?

A. In Cisco Unified Fabric, the fabric management network is an out-of-band (OOB) network for all switching nodes: leaf nodes, spines, border leaf nodes, etc. Features such as Power On Auto Provisioning (POAP) are facilitated over this network.

Q. What is an out-of-band network?

A. An OOB network is typically used to interconnect management interfaces into a single networking domain. An OOB network is specifically kept separate from the in-band data paths, to avoid overlap with the customer's data traffic and to restrict access to network management functions to authorized personnel.

Q. What are the components of Cisco Prime DCNM?

A. Cisco Prime DCNM consists of the following components:

- The LDAP server stores Virtual Routing and Forwarding (VRF) information (organization and partition) and network profile information (VLAN, segment ID, and profile type).
- The Trivial FTP (TFTP) server facilitates the POAP process.
- The XMPP server facilitates instant messaging between XMPP clients and fabric nodes (leaf and spine).
- The Simple Network Management Protocol (SNMP) subsystem provides shallow SNMP discovery of fabric switches (leaf and spine nodes).
- The SQL server stores information about discovered fabric switches.
- The Secure Copy (SCP) server facilitates secure file transfer between network nodes and the repository.

Q. What is a network profile? What kind of network profiles exist?

A. A network profile is a snippet of configuration settings that characterizes various aspects of bridge domain (VLAN) configuration. It includes the VLAN ID, segment ID, forwarding mode (traditional forwarding, enhanced forwarding, or Layer 2), IP addresses for switch virtual interfaces (SVIs), VRF membership, and various VRF parameters, among other information. The network profile can be defined manually through Cisco Prime DCNM, or through the orchestrator using Representational State Transfer (REST) APIs (that is, the orchestrator can use the REST API to define the details of the network profile). This network profile then can be used by the leaf nodes to dynamically configure relevant runtime information.

Q. How do the orchestrator and hypervisors communicate?

A. The orchestrator uses REST APIs to communicate with hypervisors.

Q. What is the Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP)?

A. VDP is part of the IEEE 802.1Qbg Edge Virtual Bridging standard. VDP can detect and signal the presence of end-host and exchange capabilities with an adjacent VDP-capable bridge (leaf switch). For more information, click here.

Optimized Networking

Q. What is the host mobility manager (HMM)?

A. HMM is a component of Cisco NX-OS Software and a fundamental innovation of Cisco Unified Fabric. To support both virtual and physical devices anywhere in a data center, Cisco Unified Fabric needs the capability to track and detect end hosts. Tracking end-host movement mainly involves discovering the end host and propagating the end-host reachability information to the other switches (leaf switches) in the fabric. End-host detection can occur through Address Resolution Protocol (ARP), Neighbor Discovery Protocol (NDP), or VDP. When a connected end host starts and sends an ARP, NDP, or VDP packet, the supervisor intercepts these packets and builds a local end-host entry point in the adjacency manager and updates the HMM software process. It subsequently populates the HMM database (if the SVI on which the host is learned is configured in fabric forwarding mode). HMM provides the main enhancement for optimized networking through the distributed gateways on the leaf switches. HMM along with the fabric feature set is enabled through the feature fabric forwarding command.

Q. What is a virtual LAN (VLAN)?

A. A VLAN today is synonymous with a bridge domain. A VLAN is uniquely identified by a VLAN ID. The VLAN ID specifications are defined in the IEEE 802.1q standard. A VLAN is limited to 4096 unique identifiers (12-bit IDs). A VLAN is considered to be a broadcast domain, so in traditional Ethernet environments, every end host within the same VLAN receives a broadcast message from every other end host in that same VLAN. In Cisco Unified Fabric with both enhanced and traditional forwarding, a VLAN ID configured on the access or trunk switch port is translated to a segment ID and is then carried over the fabric links.

Q. What is a bridge domain?

A. See question "What is a virtual LAN (VLAN)?"

Q. What is a segment ID?

A. The segment ID is a hardware-based innovation offered by the leaf nodes and is part of the Cisco FabricPath enhancements. The segment ID uses double IEEE 802.1Q tags, for a total of 24 bits, and can uniquely identify up to 16.7 million segments, a huge advancement over the traditional IEEE 802.1Q VLAN tag. In the context of Cisco Unified Fabric, segment IDs have fabricwide significance, whereas VLANs have switch or port significance. A segment ID can be used to identify a unique bridge domain or a tenant in Cisco Unified Fabric.

Q. What does the frame with the segment ID look like?

A. Figure 1 shows the segment ID in an Ethernet frame.

Figure 1. Ethernet Frame with Segment ID and Cisco FabricPath Encapsulation

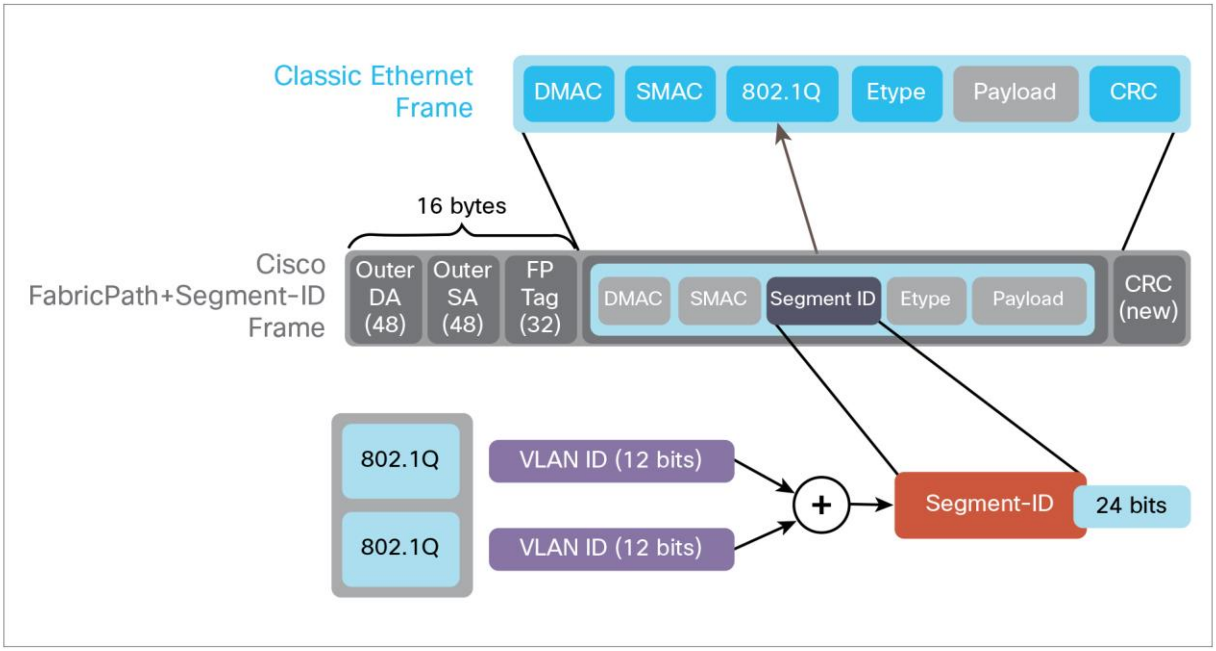

Q. What is the Virtual Extensible LAN (VXLAN) network identified (VNI)?

A. The VNI, also known as the VXLAN ID, is a 24-bit ID that allows over 16 million unique identifiers.

Q. What is Virtual Extensible LAN?

A. VXLAN is an encapsulation mechanism that can address more than 16 million IDs, or bridge domains, because of the 24-bit field in the header (VNI). It uses a generic routing encapsulation (GRE) like technique to encapsulate MAC addressbased Layer 2 frames within Layer 3 User Datagram Protocol (UDP) packets.

Q. What does the frame with VXLAN look like?

A. Figure 2 shows VXLAN in the Ethernet frame.

Figure 2. Ethernet Frame with Segment ID and Cisco FabricPath Encapsulation

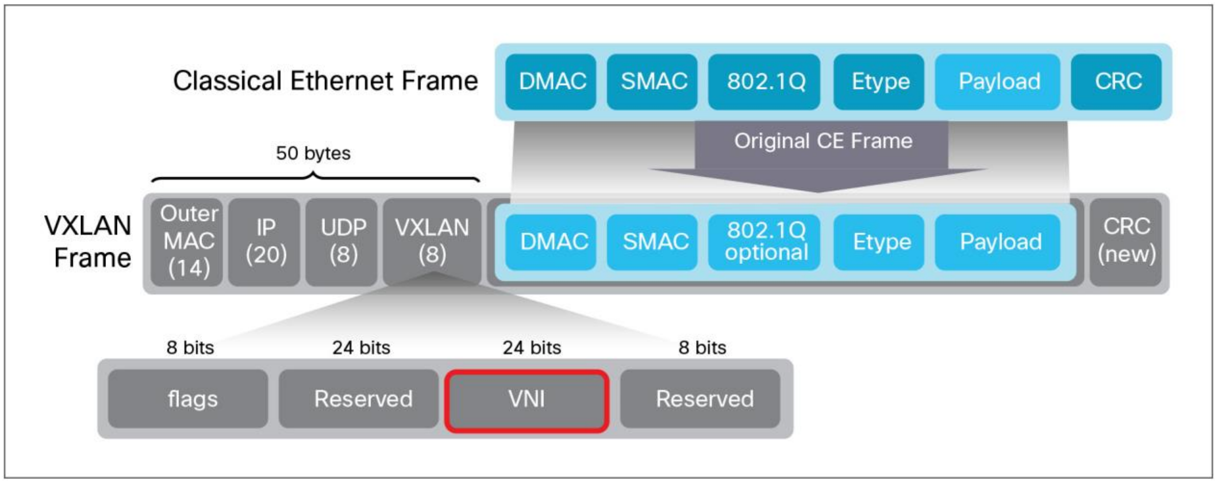

Q. What is a tenant?

A. In Cisco Unified Fabric, a tenant is uniquely represented by a VRF instance coupled with one or more bridge domains and is assigned a Layer 3 segment ID. From a networking perspective, a tenant can be thought of as a set of networks (VLANs or bridge domains) that belong to a given customer.

Q. What is a switch virtual interface?

A. An SVI is a Layer 3 termination point for a particular bridge domain (VLAN). It serves as a default gateway for that particular Layer 2 domain. A given SVI by default belongs to the default VRF instance and can be configured to be a member of other VRF instances.

Q. What is distributed virtual switch (DVS)?

A. A DVS is a set of centrally controlled virtual switching elements distributed across multiple hypervisors. Typically, a DVS provides a virtual access or trunk switch port to virtual machines and virtual appliances. The DVS also interfaces with the leaf node through an Ethernet uplink.

Q. What is a silent host?

A. Any host that attaches to the network implements a certain kind of standard TCP/IP stack, which includes ARP. Whenever a host needs to discover its default gateway, it initiates an ARP request. During this event, the default gateway creates a cache entry for this end host: an ARP entry. This ARP entry times out after the time specified in the preconfigured timer elapses. Occasionally, the end host resends the ARP request, which refreshes the ARP entry on the default gateway and thus, prevents the ARP entry from timing out. If for some reason the TCP/IP stack for a given IP and MAC address pair does not regularly send ARP packets to the default gateway and does not respond to ARP requests from the default gateway, the cached ARP entry times out and is never reinstated again. Hosts with such behavior are called silent hosts. Silent hosts in the network are very rare; however, some clustering applications with emulated virtual IP addresses can result in silent hosts.

Q. What is the data plane? What are typical examples of a data plane?

A. The data plane is a logical set of methods and mechanisms for establishing a switching data path for user data. Typically, switching application-specific integrated circuits (ASICs) and line cards containing those ASICs are considered to make up the data plane.

Q. What is the control plane? What protocols work on the control plane?

A. The control plane is a logical concept that encompasses methods and mechanisms for distribution, monitoring, and maintenance of the information needed to establish and maintain the data path. Typically, various protocols (Spanning Tree Protocol, Border Gateway Protocol [BGP], Open Shortest Path First [OSPF], Label Distribution Protocol [LDP], Bidirectional Forwarding Detection [BFD], etc.) are considered to constitute the control plane, and supervisors and line-card CPUs are considered to be control-plane hardware.

Q. What is a distributed gateway and how is it different from First-Hop Redundancy Protocol (FHRP)?

A. A distributed gateway is a set of features that can be enabled on a per-VLAN basis. A distributed gateway is configured with **fabric forwarding mode proxy-gateway** for enhanced forwarding mode, and with **fabric forwarding mode anycast-gateway** for traditional forwarding mode. With these settings along with others, any leaf can automatically instantiate an SVI and tenant for a given VLAN (bridge domain), but only if workloads are connected and active in that VLAN. This feature set allows the leaf nodes to terminate end-host-originated control-plane traffic, such as ARP and NDP, Gratuitous ARP (GARP), Domain Host Configuration Protocol (DHCP), and Interior Gateway Management Protocol (IGMP). The result is better scalability and much smaller flood and fault domains in comparison to FHRP.

Q. What modes can a VLAN use in Cisco Unified Fabric with optimized networking? How do I choose the best mode to use?

A. A given VLAN can be configured in three forwarding modes: enhanced forwarding, traditional forwarding, and Layer 2. Enhanced forwarding and traditional forwarding are similar in that they automatically instantiate a distributed gateway on the relevant leaf nodes. Refer to the question "What is a distributed gateway and how is it different from First-Hop Redundancy Protocol (FHRP)?" for more details.

Enhanced forwarding does not forward ARP, NDP, GARP, DHCP, and IGMP traffic that originates at the end host. These protocols are always terminated on the local leaf SVIs; hence, the flooding domain is limited to the local leaf. Enhanced forwarding replies with a predefined MAC address (it is identical on all leaf nodes and can be configured using POAP) to any ARP request, regardless of whether it is requesting a MAC address for the default gateway or any end host within the same subnet. In addition, any network communication within the enhanced forwarding mode VLAN always is routed through the network, even for data traffic within the same VLAN.

The SVI for a traditional forwarding mode VLAN still replies with a predefined MAC address to ARP requests toward the default gateway; however, ARP requests for the end host within a subnet are allowed to be flooded in the fabric. These broadcast frames are forwarded along the Cisco FabricPath multidestination tree (MDT).

A VLAN configured in Layer 2 mode is identical to a regular VLAN in Cisco FabricPath mode. The only difference is that with the segment ID, the network can have more than 16 million bridge domains.

Virtual Fabrics

Q. What is Virtual Routing and Forwarding?

A. VRF is a software and hardware feature set that allows multiple separate and independent routing tables to coexist on the same device. These routing tables may belong to different customers (tenants) and may contain overlapping IP addresses. VRF separation allows Cisco Unified Fabric to host hundreds of customers on the same fabric while providing comprehensive services and network management capabilities.

Q. What is the mobility domain?

A. Cisco Unified Fabric can take advantage of VDP with the Cisco Nexus® 1000V or Open Virtual Switch (OVS) to automatically resolve segment ID to - VLAN mapping. However, nonvirtualized physical appliances, bare-metal servers, and any other networking devices can directly connect to leaf nodes. Such devices use only standard IEEE 802.1q VLAN tags, so upon receiving these frames, the leaf switch can't uniquely identify the segment ID and VRF to which that particular VLAN maps. The mobility domain is the mechanism that enables this association. The VLAN ID coupled with the mobility domain allows the switch to uniquely identify the segment ID of that bridge domain as well as the VRF membership of this bridge domain.

Q. What function does the Border Gateway Protocol have in the Cisco Unified Fabric network?

A. Multiprotocol BGP (MP-BGP) is the version of BGP used in Cisco Unified Fabric. It provides IPv4, IPv6, Virtual Private Network Version 4 (VPNv4, VPNv6, Multicast VPNv4 (mVPNv4), and mVPNv6 prefix exchange within the fabric. MP-BGP is responsible for reliable distribution of end-host reachability information within the fabric. Essentially all/32 and/128 prefixes learned by HMM are carried to the MP-BGP route reflectors, and later to the rest of the leaf and border leaf nodes. Note that not all unicast routing information base (URIB) prefixes received from the MP-BGP route reflector are installed in hardware, because only the relevant prefixes are downloaded to hardware with the optional conversational learning.

Q. What function does Cisco FabricPath Intermediate SystemtoIntermediate System (IS-IS) Protocol have in the Cisco Unified Fabric network?

A. Cisco FabricPath IS-IS is responsible for building a logical underlay using all available paths. The best paths between all devices in the network are calculated using the Shortest-Path First (SPF) algorithm. This logical underlay network is used to transport user data traffic using the shortest available path. Cisco FabricPath IS-IS also calculates several MDTs to facilitate loop-free forwarding of multidestination traffic. Cisco FabricPath IS-IS in Cisco Unified Fabric with optimized networking also carries mappings between the switch ID and IP and MAC addresses of the SVI for the backbone VLAN. This information is exchanged with all fabric switches in the network to facilitate distributed gateway functions.

Q. What function does Cisco FabricPath have in Cisco Unified Fabric with optimized networking?

A. Cisco FabricPath provides underlay transport for user data traffic in Cisco Unified Fabric.

Q. What is an interautonomous system (InterAS) and why is it needed?

A. InterAS, a Multiprotocol Label Switching (MPLS) term used in Cisco Unified Fabric with optimized networking, signifies a way of connecting Cisco Unified Fabric to an external network using BGP. Such technology is needed to distribute tenant-specific (or VRF-specific) prefixes to remote networks.

Workload Automation

Q. What is an application programming interface?

A. An API consists of methods and protocols that describe how software and hardware components interact with each other. For more information, see http://en.wikipedia.org/wiki/Application_programming_interface. Cisco Unified Fabric uses multiple APIs, including REST APIs.

Q. What are northbound and southbound APIs?

A. Northbound and southbound APIs allow a particular network component to communicate with higher-level or lower-level components, respectively. In other words, northbound and southbound APIs enable various applications with control and feedback mechanisms. Orchestrators are one type of such applications. That is, through northbound APIs, the orchestrator can communicate with Cisco Prime DCNM, with virtual switches (for example, the Cisco Nexus 1000V or OVS), and with the virtual machine manager (for example, VMware vCenter) and push configurations that enable network connectivity for particular workloads.

Q. What is a Representational State Transfer API?

A. A REST API is a standard programming interface used to exchange information mainly through HTTP calls. Fabric management uses the REST API to support information exchange between Cisco Prime DCNM and NSC, orchestrators, and other northbound systems.

Miscellaneous

Q. Where can I find installation guidance for Cisco Prime DCNM?

A. Go to http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/7_x/dcnm/installation/master_files/OVA_Installation_Guide.html.

Q. Where can I find a quick-start guide for Cisco Unified Fabric innovations?

A. Go to http://www.cisco.com/c/en/us/products/cloud-systems-management/dynamic-fabric-automation/solution-overview-listing.html.

Q. Where can I find a Cisco Unified Fabric command reference guide?

A. Go to http://www.cisco.com/c/en/us/products/cloud-systems-management/dynamic-fabric-automation/literature.html.

Q. What are the scalability limits for Cisco Unified Fabric innovations?

A. Please refer to the verified scalability document at http://www.cisco.com/c/en/us/td/docs/switches/datacenter/dfa/verified-scalability/guide/b-dfa-verified-scalability-guide/b-dfa-verified-scalability_chapter_01.html.

Q. What is the Open Virtual Appliance (OVA)?

A. OVA is an open standard for packaging and distributing virtual appliances that are deployed as virtual machines. A virtual appliance is a prebuilt software solution maintained, updated, and managed as a single piece of software. Cisco Prime DCNM Release 7.0 or later is available in OVA format.

Q. What is a hypervisor?

A. A hypervisor is software or an operating system installed on a computing server that provides an abstraction layer for hosted virtual machines. Common supported hypervisors include Red Hat Kernel-based Virtual Machine (KVM), VMware ESXi, Citrix XenServer, and Microsoft Hyper-V.

Q. What is the Open Virtual Switch?

A. OVS is an open virtual switch and is usually installed on the Red Hat KVM hypervisor.

Q. What protocols supplement Ethernet?

A. Ethernet protocol needs additional support from several control protocols to facilitate its work in IP networks. ARP, GARP, Reverse ARP (RARP), NDP, and IGMP are some of these protocols. End hosts typically communicate with their respective default gateways as well as other end hosts using such protocols. Typically, these protocols are terminated at the Layer 2 and Layer3 boundary and do not propagate past their respective Layer 2 domains. A large number of hosts in a single Layer 2 domain may translate into a very large number of end-host control-protocol conversations, which creates high levels of CPU utilization on all Layer 2-adjacent end hosts, appliances, and gateways.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C67-732812-00 09/14