Campus Fabric Design Guide

October 2016

CISCO

REFERENCE
NETWORK
ARCHITECTURE
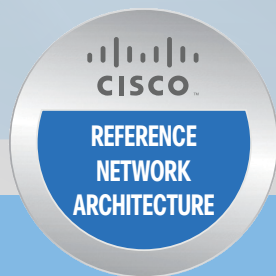
Table of Contents

Campus Fabric Introduction

Cisco Digital Network Architecture (DNA) provides a roadmap to digitization and a path to realize immediate benefits of network automation, assurance, and security. The *campus fabric* architecture is the Cisco DNA evolution from existing campus LAN designs. Campus fabric introduces programmable overlays enabling easy-to-deploy network virtualization across the campus. In addition to network virtualization, campus fabric allows for software-defined segmentation and policy enforcement based on user identity and group membership, seamlessly integrated with Cisco TrustSec technology. The capability to easily instantiate virtual networks with integrated security provides significant reduction in operational expenses and increased ability to drive business innovations quickly with minimum risk.

This guide provides an overview of the requirements driving the evolution of campus network designs, followed by a discussion about the latest technologies and designs that are available for building a campus fabric to address those requirements. This guide is a companion to the associated deployment guides for campus fabric, which provide configurations required to implement the designs that are described in this guide. The intended audience is a technical decision maker who wants to understand Cisco's campus offerings and to learn about the technology options available and design best practices.

For related design guides, deployment guides, and white papers, see the following page:

http://www.cisco.com/go/designzone

NETWORK REQUIREMENTS FOR THE DIGITAL ORGANIZATION

With digitization, software applications are evolving from simply supporting business processes to becoming, in some cases, the primary source of business revenue and competitive differentiation. Organizations are now constantly challenged by the need to scale their network capacity in order to quickly react to, and support the growth of, application demands. Because the campus LAN is the network through which users and devices within a location access applications, campus LAN capabilities should be enhanced to support those changing needs.

The following are key requirements that are driving the evolution of existing campus networks.

Flexible Ethernet Foundation for Growth and Scale

- **Increased capacity of wireless access points**—The bandwidth demands on wireless access points (APs) with the latest 802.11ac Wave 2 technology now exceed 1 Gbps, and the IEEE has now ratified the 802.3bz standard that defines 2.5 Gbps and 5 Gbps Ethernet. Cisco Catalyst Multigigabit technology supports that bandwidth demand without requiring an upgrade of the existing copper Ethernet wiring plant.

- **Increased power requirements from Ethernet devices**—New devices may require higher power to operate, such as lighting, remote access switches, and APs. Your design should have the ability to support power over Ethernet with 60W per port. Cisco Universal Power Over Ethernet in the access layer achieves this goal.

- **Increased bandwidth needs**—Bandwidth needs are doubling potentially multiple times over the lifetime of a network, resulting in new networks needing to be prepared to aggregate using 10 Gbps Ethernet to 40 Gbps to 100 Gbps capacities over time.

Network Integrated Security

- **Identity services**—Identifying users and devices connecting to the network provides the contextual information required to implement security policies for access control, network segmentation, and mapping of devices into virtual networks (VN).

- **Group-based policies**—Creating security policies based on user group information provides a much easier and scalable way to deploy and manage security policies. Traditional access control lists (ACLs) can be difficult to implement, manage, and scale because they rely on network constructs such as IP addresses and subnets.

- **Network segmentation**—Security group tags (SGTs) assigned from group-based policies can be used to segment a network in order to achieve data plane isolation over physical and virtual networks.

- **Network virtualization**—The capability to share a common infrastructure while supporting multiple VNs with isolated data and control planes provides multi-tenancy and security.

Campus Fabric Architecture

The campus fabric architecture enables the use of virtual networks (overlay networks) running on a physical network (underlay network) in order to create alternative topologies to connect devices. Overlay networks are commonly used to provide Layer 2 and Layer 3 logical networks with virtual machine mobility in data center fabrics (examples: ACI, VXLAN, and FabricPath) and also in wide-area networks to provide secure tunneling from remote sites (examples: MPLS, DMVPN, and GRE). This section provides information about the architecture elements that define the campus fabric. Design recommendations for the campus fabric are covered in the Design Considerations section.

UNDERLAY NETWORK

The underlay network is defined by the physical switches and routers that are part of the campus fabric. All network elements of the underlay must establish IP connectivity via the use of a routing protocol. Theoretically, any topology and routing protocol can be used, but the implementation of a well-designed Layer 3 foundation to the campus edge is highly recommended to ensure performance, scalability, and high availability of the network. In the campus fabric architecture, end-user subnets are not part of the underlay network.

> **Tech Tip**
>
> Initial versions support only IPv4 underlay networks. To verify if IPv6 underlay networks are supported, check the release notes for your software version.

OVERLAY NETWORK

An overlay network runs over the underlay in order to create a virtualized network. Virtual networks isolate both data plane traffic and control plane behavior among the virtualized networks from the underlay network. Virtualization is achieved inside the campus fabric by encapsulating user traffic over IP tunnels that are sourced and terminated at the boundaries of the campus fabric. Network virtualization extending outside of the fabric is preserved using traditional virtualization technologies such as VRF-Lite and MPLS VPN. Overlay networks can run across all or a subset of the underlay network devices. Multiple overlay networks can run across the same underlay network to support multi-tenancy through virtualization.

Layer 2 Overlays

Layer 2 overlays emulate a LAN segment and can be used to transport IP and non-IP frames. Layer 2 overlays carry a single subnet over the Layer 3 underlay. Layer 2 overlays are useful in emulating physical topologies and are subject to L2 flooding.

Figure 1   *Layer 2 overlay–connectivity logically switched*

Layer 2 Overlay

Border Nodes

Edge Nodes

Physical Fabric Topology
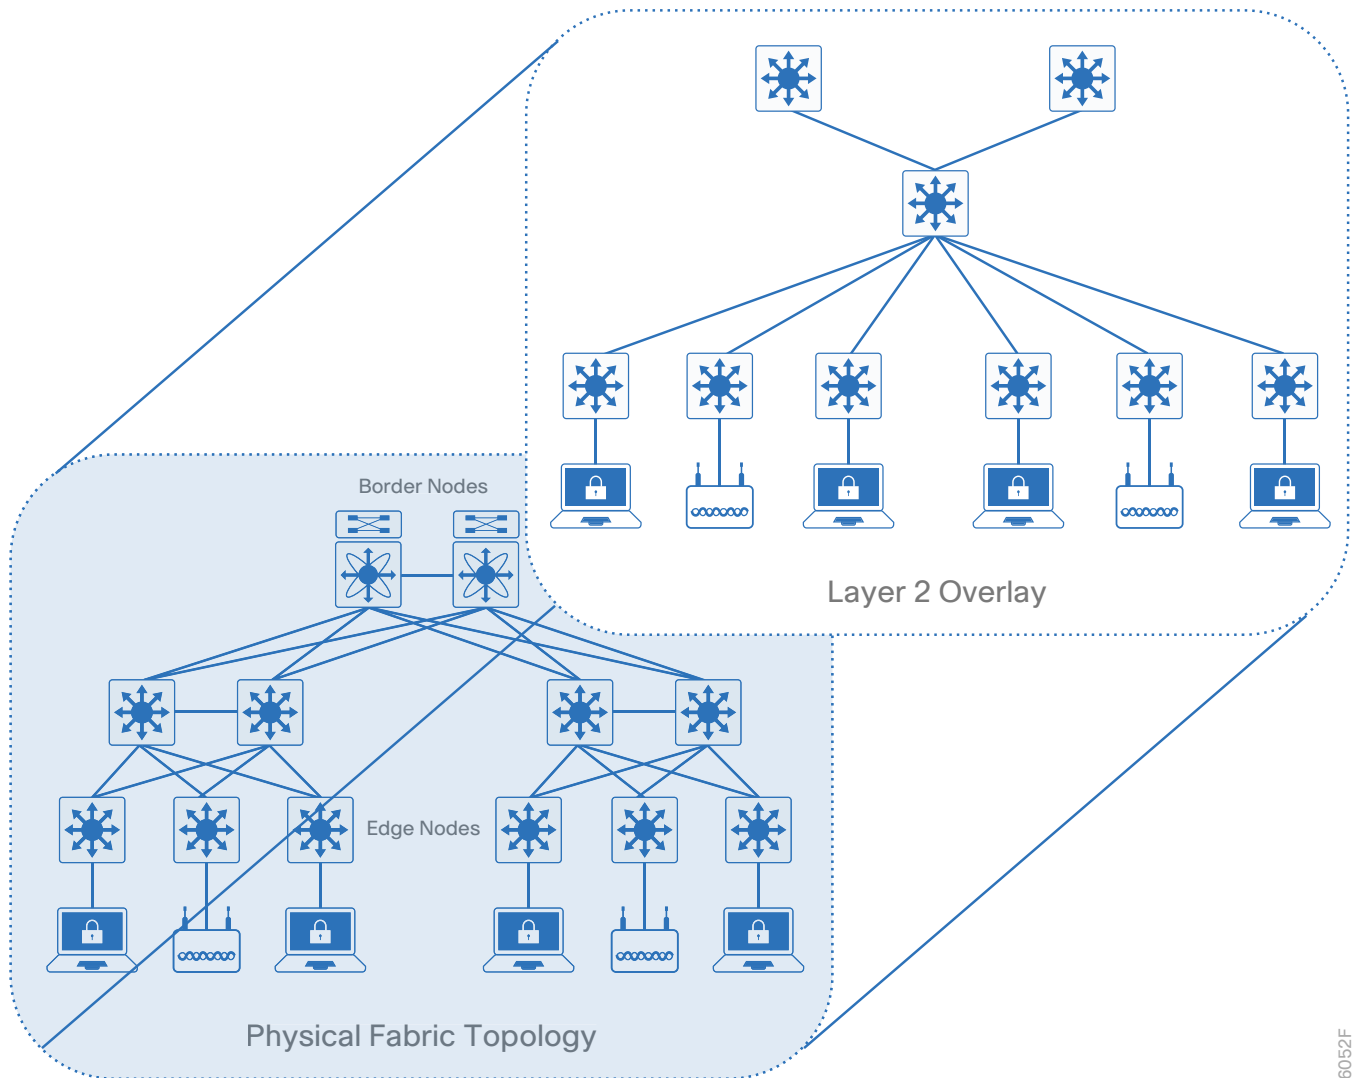
6052F

Tech Tip

Initial versions only support transport of IP frames. Check the release notes for your software version to verify if transport of non-IP frames is supported.

Layer 3 Overlays

Layer 3 overlays abstract IP based connectivity from physical connectivity and allow multiple IP networks as part of each virtual network. Overlapping IP address space is supported across different Layer 3 overlays as long as the network virtualization is preserved outside of the fabric, using existing network virtualization functions such as VRF–Lite and MPLS L3VPN.

**Figure 2** *Layer 3 overlay–connectivity logically routed*
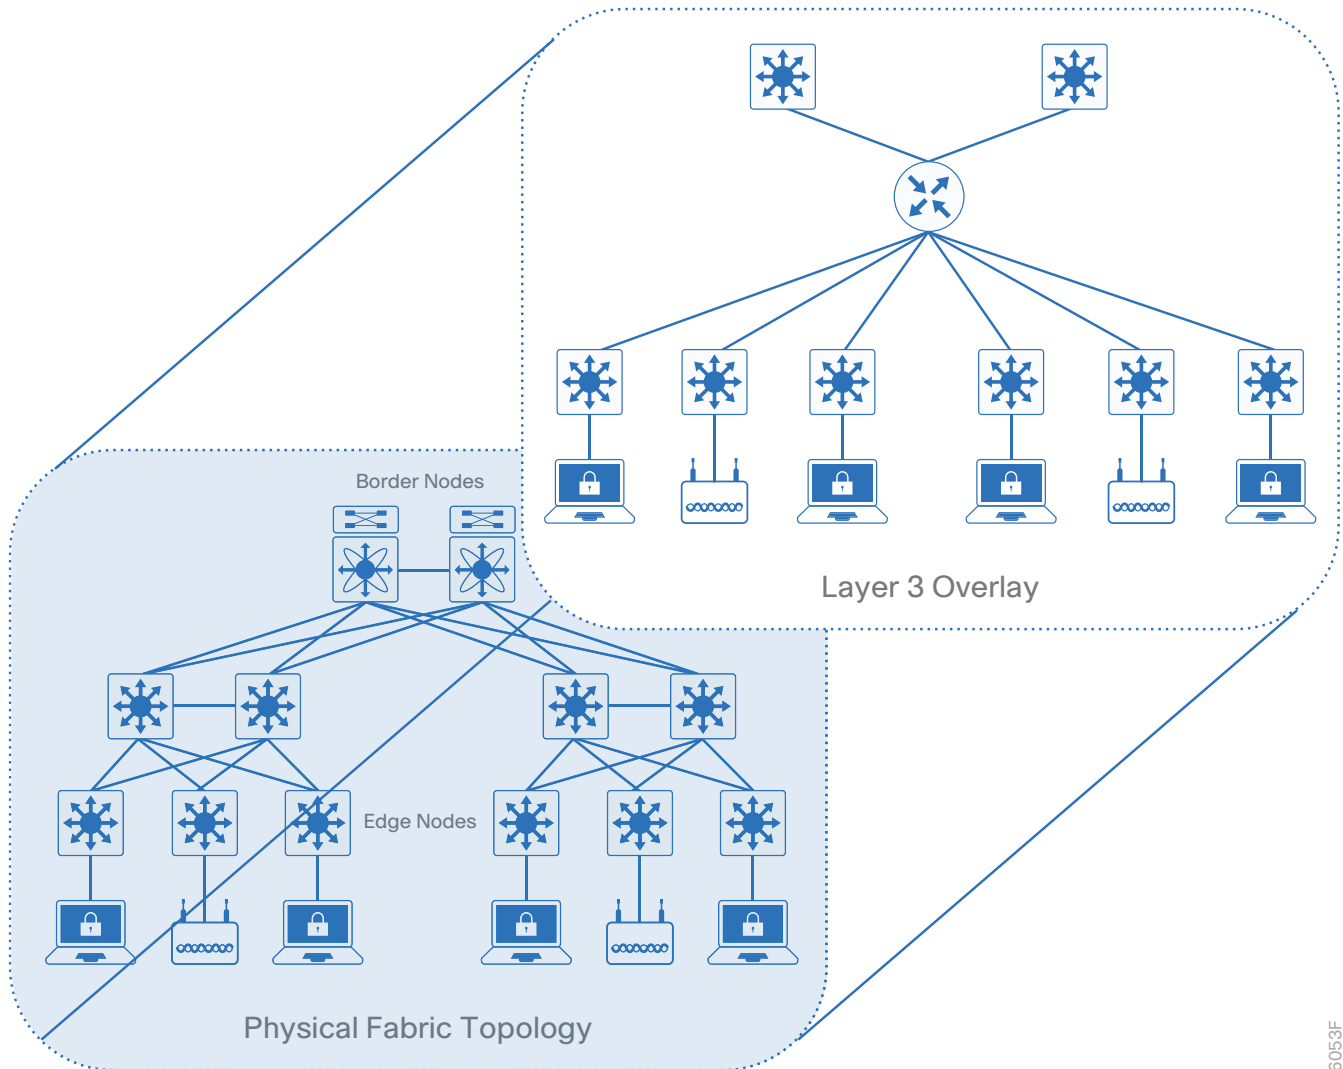
CAMPUS FABRIC DATA PLANE

RFC 7348 defines the use of virtual extensible LAN (VXLAN) as a way to overlay a Layer 2 network on top of a Layer 3 network. Using VXLAN, you tunnel the original Layer 2 frame using UDP/IP over the Layer 3 network. The tunnel interface at each node is called a *VXLAN tunnel endpoint* (VTEP). VTEPs rely on data-plane learning or a control plane in order to determine the remote endpoint to VTEP mapping for traffic encapsulation. Each overlay network is called a *VXLAN segment* and is identified using a 24-bit VXLAN network identifier (VNI), which supports up to 16 million VXLAN segments.

Figure 3 *RFC 7348 VXLAN header*
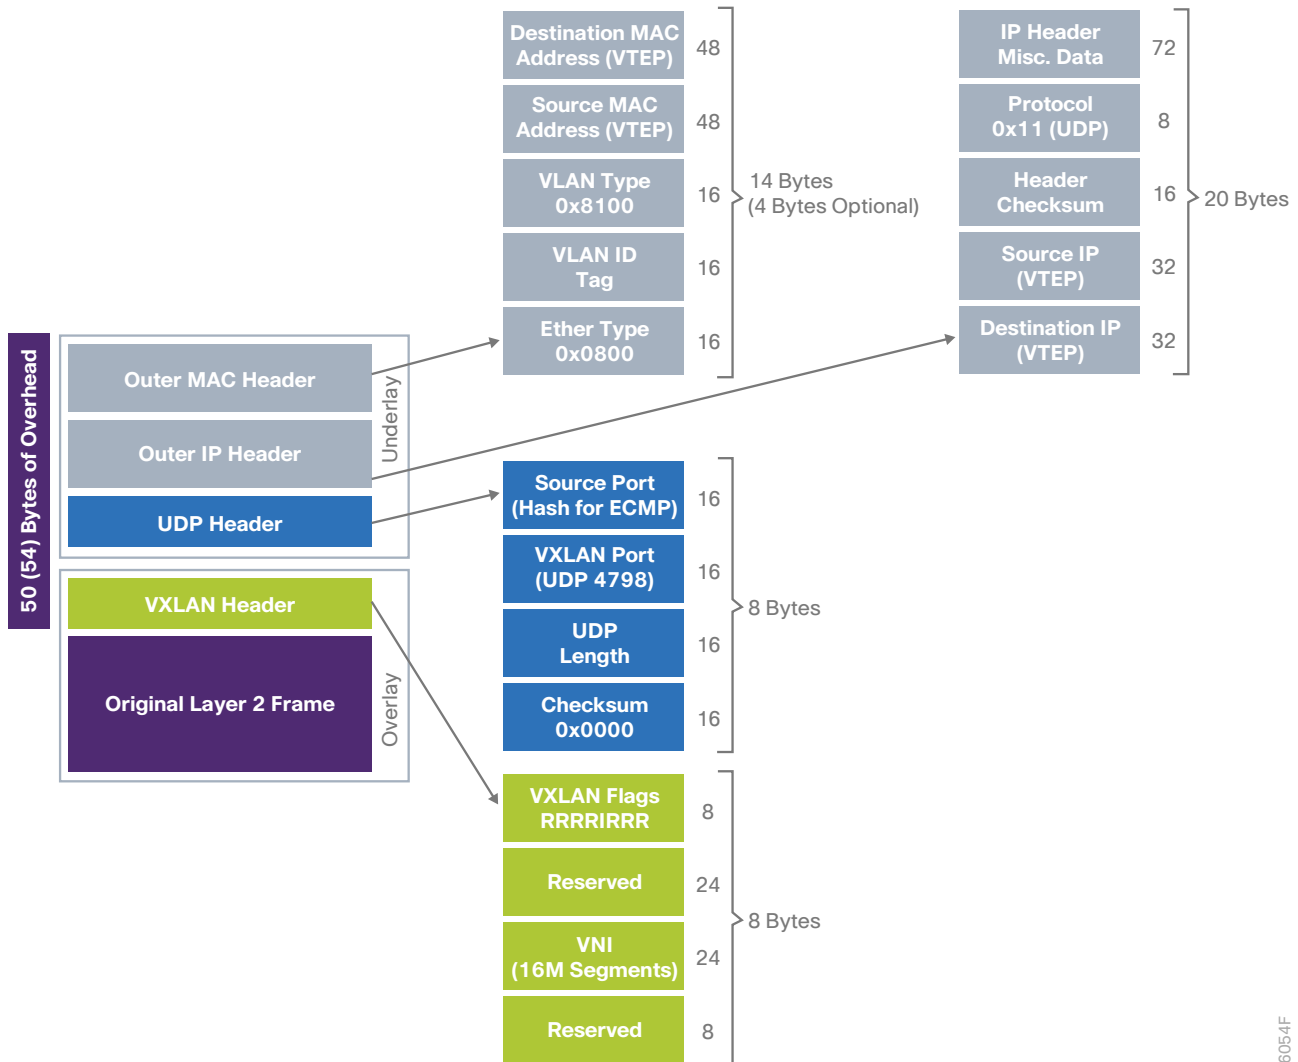
The campus fabric uses the VXLAN data plane in order to provide transport of the full original Layer 2 frame and additionally uses Locator/ID Separation Protocol (LISP) as the control-plane in order to resolve endpoint-to-VTEP mappings. The campus fabric replaces 16 of the reserved bits in the VXLAN header in order to transport up to 64,000 SGTs.

The VNI maps to virtual routing and forwarding (VRF) and provides the mechanism to isolate data and control plane across different virtual networks. The SGT carries user group membership information and is used to provide data-plane segmentation inside the virtualized network.

CAMPUS FABRIC CONTROL PLANE

RFC 6830 and other RFCs define LISP as a network architecture and set of protocols that implement a new semantic for IP addressing and forwarding. In traditional IP networks, the IP address is used to identify both an endpoint and its physical location as part of a subnet assignment on a router. In a LISP-enabled network, an IP address is used as the endpoint identifier (EID) for a device, and an additional IP address is used as a routing locator (RLOC) to represent the physical location of that device (typically a loopback address of the router to which the EID is attached). The EID and RLOC combination provides the necessary information for traffic forwarding. The RLOC address is part of the routing domain, and the EID can be assigned independently of the location.

The LISP architecture requires a mapping system that stores and resolves EIDs to RLOCs. This is analogous to using DNS to resolve IP addresses for host names and also similar to the previously mentioned VTEP mapping in the VXLAN data plane. EID prefixes (IPv4 addresses with /32 "host" masks) are registered into the map server along with their associated RLOCs. When sending traffic to an EID, a source RLOC queries the mapping system in order to identify the destination RLOC for traffic encapsulation.

Although a full understanding of LISP and VXLAN isn't required to deploy a campus fabric, it is helpful to understand how these technologies support the deployment goals of the campus fabric. Some of the benefits that the LISP architecture provides for the campus fabric include:

- **Network virtualization**—A LISP Instance ID is used to maintain independent VRF topologies. From a data-plane perspective, the LISP Instance ID maps to the VNI.

- **Subnet stretching**—A single subnet can be extended to exist at multiple RLOCs. The separation of EID from RLOC enables the capability to extend subnets across different RLOCs. The RLOC in the LISP architecture represents the VTEP functionality in VXLAN as it is the ingress and egress tunnel used to encapsulate EID traffic over a Layer 3 network.

- **Smaller routing tables**—Only RLOCs need to be reachable in the global routing table. Local EIDs are cached at the local node while remote EIDs are learned through conversational learning. *Conversational learning* is the process of populating forwarding tables with only endpoints that are communicating through the node. This allows for efficient use of forwarding tables.

The following diagram shows an example of two subnets that are part of the overlay network and are stretched across routers that are physically separated. The RLOC interface is the only routable address that is required to establish connectivity between endpoints of the same or different subnet.
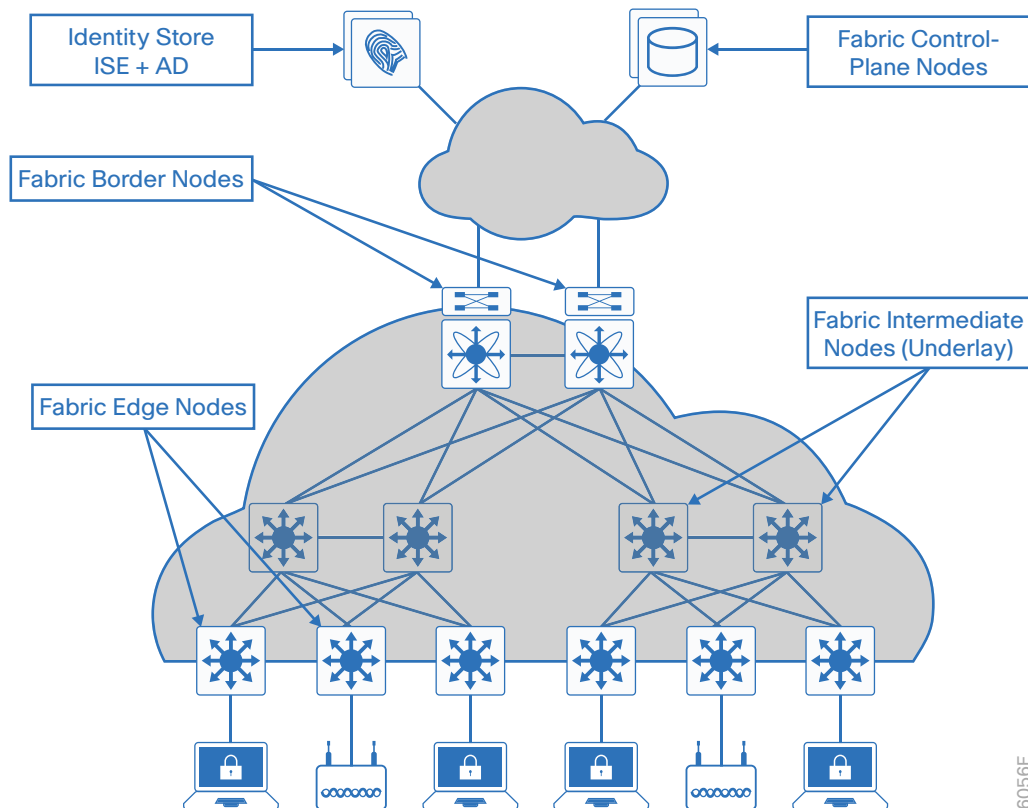
Figure 4   *LISP sample topology*
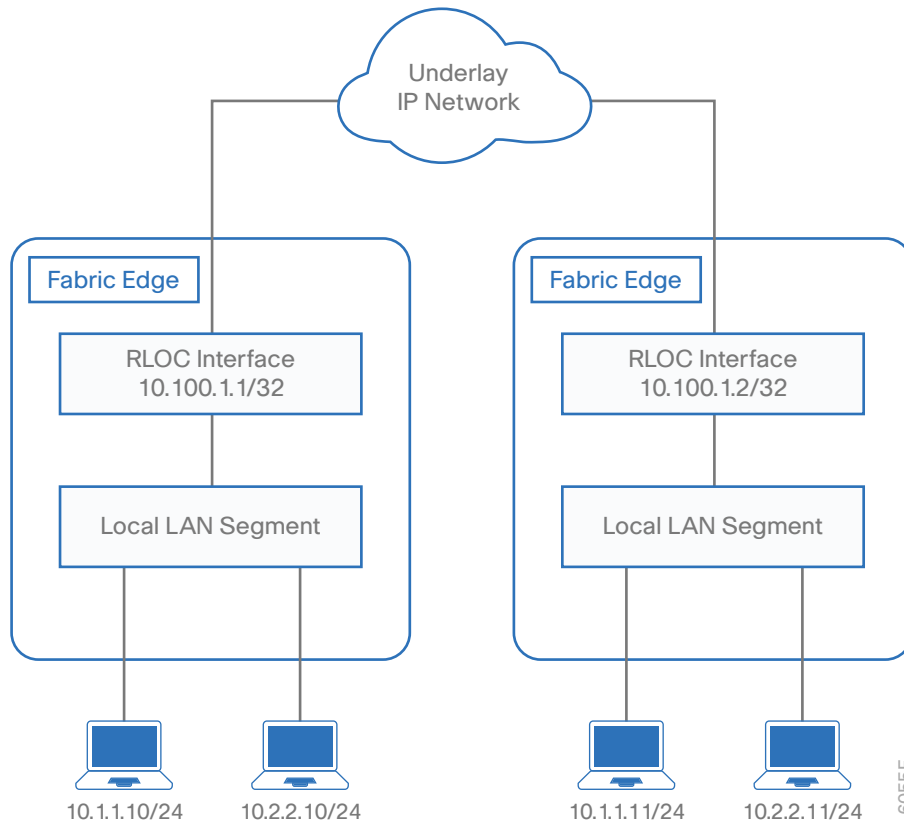
Solution Components

The campus fabric is composed of fabric control plane nodes, edge nodes, intermediate nodes, and border nodes. This section describes the functionality for each role and how the roles map to the physical campus topology.

Figure 5 *Campus fabric components*

FABRIC CONTROL-PLANE NODE

The fabric control plane node is based on the LISP Map-Server (MS) and Map-Resolver (MR) functionality combined on the same node. The fabric control plane node functionality is instantiated on the fabric border node or a dedicated node. The fabric control plane node enables the following functions:

- **Host Tracking Database**—The host tracking database (HTDB) is a central repository of EID-to-fabric-edge node bindings.

- **Map-Server**—The LISP MS is used to populate the HTDB from registration messages from fabric edge devices.

- **Map-Resolver**—The LISP MR is used to respond to map queries from fabric edge devices looking to determine RLOC mapping information for destination EIDs.

FABRIC EDGE NODE

The fabric edge nodes are the equivalent of an access layer switch in a traditional campus design. The fabric edge nodes implement a Layer 3 access design with the addition of the following fabric functions:

- **Endpoint registration**—After an endpoint is detected by the fabric edge, it is added to a local host tracking database called the EID-table. The edge device also issues a LISP map-register message in order to inform the control plane node of the endpoint detected so that it can populate the HTDB.

- **Mapping of user to virtual network**—Endpoints get placed into virtual networks by assigning the endpoint to a VLAN, which is mapped to a LISP instance. The mapping of endpoints into VLANs can done statically or dynamically using 802.1X. An SGT can also be assigned to provide segmentation and policy enforcement at the fabric edge.

- **Anycast Layer 3 gateway**—A common gateway (IP and MAC addresses) can be used at every node that shares a common EID subnet in order to provide for optimal forwarding and mobility across different RLOCs.

- **LISP forwarding**—Instead of a typical routing-based decision, the fabric edge nodes query the map server in order to determine the RLOC associated with the destination IP and use that information to encapsulate the traffic in VXLAN. In case of a failure to resolve the destination RLOC, the traffic is sent to the fabric border in which the global routing table is used for forwarding. The response received from the map server is stored in the LISP map-cache, which is merged to the CEF table and installed in hardware. If VXLAN-encapsulated traffic is received at the fabric edge for an endpoint not locally connected, a LISP solicit map request is sent to the sending fabric edge in order to trigger a new map request; this addresses the case where the end-point may be present on a different fabric edge switch.

FABRIC INTERMEDIATE NODE

The fabric intermediate nodes are part of the Layer 3 network that interconnects the edge nodes to border nodes. In case of a three-tier campus design using a core, distribution, and access, the fabric intermediate nodes are the equivalent of a distribution switches. Fabric intermediate nodes only route IP traffic inside the fabric. No VXLAN encapsulation/de-encapsulation or LISP control plane messages are required from the fabric intermediate node.

FABRIC BORDER NODE

The fabric border nodes serve as the gateway between the fabric domain and the network outside of the fabric. The fabric border node is responsible for network virtualization inter-working and SGT propagation from the campus fabric to the rest of the network. The fabric border nodes implement the following functions:

- **Advertisement of EID subnets**—The fabric border runs either an interior gateway protocol (IGP) or border gateway protocol (BGP) as a routing protocol in order to advertise the EID prefixes outside of the fabric and traffic destined to EID subnets from outside the campus fabric goes through the border nodes. These EID prefixes appear only on the routing tables at the border—throughout the rest of the fabric, the EID information is accessed using the fabric control plane node.

- **Fabric domain exit point**—The fabric border is the gateway of last resort for the fabric edge nodes. This is implemented using LISP Proxy Tunnel Router functionality.

- **Mapping of LISP instance to VRF**—The fabric border can extend network virtualization from inside the campus fabric to outside the campus fabric by using external VRF instances in order to preserve the virtualization.

- **Policy mapping**—The fabric border node also maps SGT information from within the fabric to be appropriately maintained when exiting that fabric. Tags from the VXLAN header are mapped to Cisco Meta Data (CMD) when inline tagging capabilities are used, or alternatively the tags are transported by SGT exchange protocol (SXP), allowing for seamless integration with the Cisco TrustSec solution.

Design Considerations

Designing for a campus fabric is not a one-design-fits-all proposition. The scale of a campus fabric can be as small as an access-distribution block or as big as a three-tier campus deployment. In a single network, multiple fabrics can be deployed as long as the fabric elements are assigned to a single fabric only.

PLATFORM SUPPORT

Platform choices for a campus fabric deployment are usually driven by the capacity required from the network.

Table 1 *Platform and capabilities comparison*

|	Catalyst 3850 and 3650 Series	Catalyst 6807XL	Nexus 7700 with Supervisor 2E	CSR 1Kv
Supervisor	–	Supervisor 6T and Supervisor 2T	Supervisor 2E	–
Supported Linecards	–	6900 Series	M3 Series	–
Edge Node Functionality	Yes	No	No	–
Border Node Functionality	Yes	Yes	Yes	–
Control Plane Node Functionality	Yes	Yes	No	Yes

Tech Tip

To achieve the functionality shown, you must meet minimum software release requirements. For more information, see the software release notes for your platform.
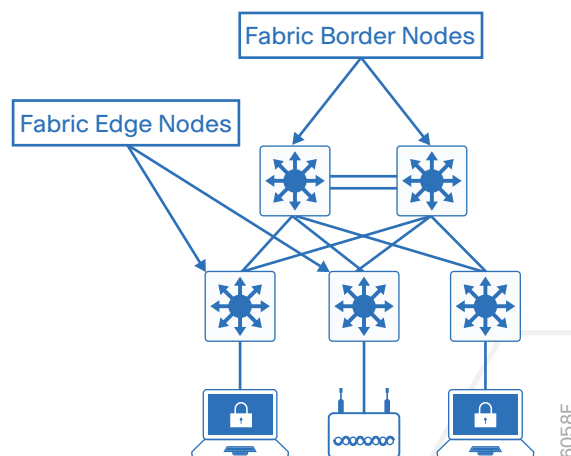
PHYSICAL TOPOLOGIES

Campus fabric topologies should follow the same design principles and best practices associated with a hier-archical design by splitting the network into modular groups. You create design elements that can be replicated throughout the network by using modular designs. The following example shows the physical topology of a three-tier campus design in which all nodes are dual homed with equal-cost links that will provide for load-balancing, redundancy, and fast convergence. The cross links at each aggregation layer is only used for optimal routing in case of an uplink failure.

Figure 6   *Three-tier campus fabric*
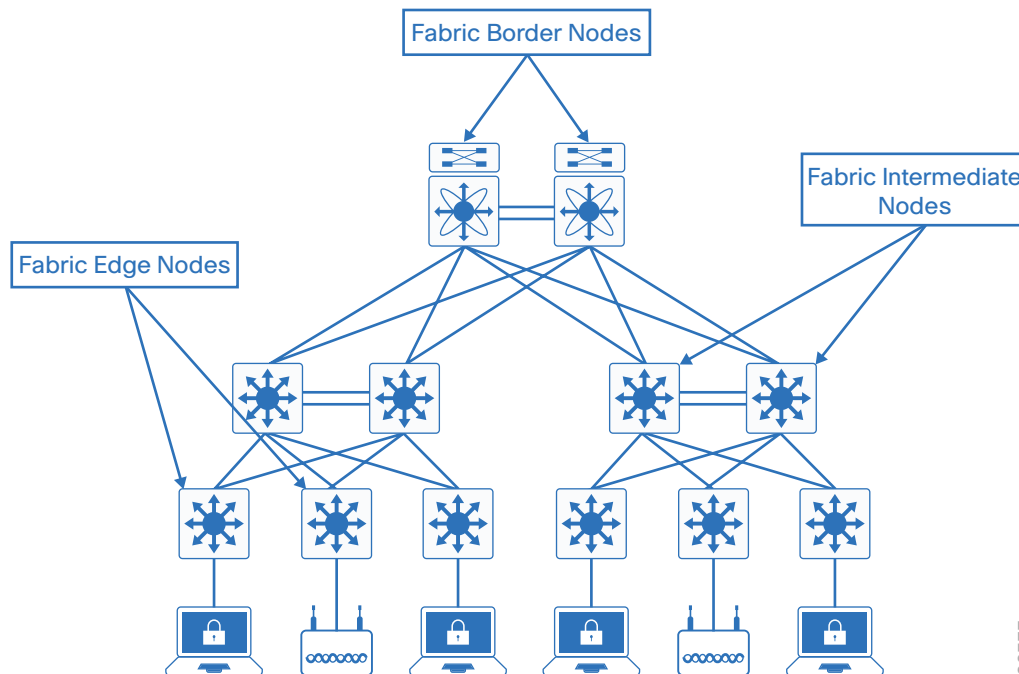
For smaller deployments, a campus fabric can be implemented using a two-tier design. The same design prin-ciples should be applied but without the need for an aggregation layer implemented by intermediate nodes.

Figure 7   *Two-tier campus fabric*

Campus fabric topologies should be deployed as spoke networks with the fabric border node at the exit point hub for the spokes, although other physical topologies can be used. Topologies in which the fabric is a transit network should be planned carefully in order to ensure optimal forwarding. If the border node is implemented at a node that is not the aggregation point for exiting traffic, sub-optimal routing results when traffic exits the fabric at the border and then doubles back to the actual aggregation point.

UNDERLAY DESIGN

Having a well-designed underlay network will ensure the stability, performance, and efficient utilization of the campus fabric network. The following underlay design best practices are highly recommended:

- **Increase default MTU**—The VXLAN header adds 50 and optional 54 bytes of encapsulation overhead. Some Ethernet switches support a maximum transmission unit (MTU) of 9216 while others may have an MTU of 9196 or smaller. Given that server MTUs typically go up to 9,000 bytes, enabling a network wide MTU of 9100 ensures that Ethernet jumbo frames can be transported without any fragmentation inside and outside of the fabric.

- **Layer 3 to the access design**—The use of a Layer 3 routed network for the fabric provides the highest level of availability without the need to use loop avoidance protocols or interface bundling techniques.

- **Use point-to-point links**—Point-to-point links provide the quickest convergence times because they eliminate the need to wait for the upper layer protocol timeouts typical of more complex topologies. Combining point-to-point links with the recommended physical topology design provides fast convergence after a link failure. The fast convergence is a benefit of quick link failure detection triggering immediate use of alternate topology entries preexisting in the routing and forwarding table. Implement the point-to-point links using optical technology and not copper, because optical interfaces offer the fastest failure detection times to improve convergence.

- **Dedicated IGP process for the fabric**—The underlay network of the fabric only requires IP reachability from the fabric edge to the border node. In a fabric deployment a single area IGP design can be implemented for which a dedicated IGP process can be implemented at the campus fabric. Addressing space used for links inside the fabric does not need to be advertised outside of the fabric and can be reused across multiple fabrics.

- **Loopback propagation**—The loopback addresses assigned to the underlay devices need to propagate outside of the fabric in order to establish connectivity to infrastructure services such as fabric control plane nodes, DNS, DHCP, and AAA. As a best practice, use /32 host masks. To propagate the loopback host routes, use route tags in order to enable an easy mechanism for redistributing only the loopbacks, to avoid maintaining prefix lists.

OVERLAY DESIGN

In the campus fabric, the overlay networks are used for transporting user traffic within the fabric. The fabric encapsulation also carries security group information that can be used for traffic segmentation inside the overlay. The following design considerations should be taken into account when deploying virtual networks:

- **Virtualize only when needed**—Segmentation using SGTs allows for simple-to-manage group-based policies and enables granular data plane isolation between groups of users and applications. Virtualized networks are isolated at both the data plane and control plane. If communication is required between different virtual networks, you use an external device or firewall to enable inter-VN communication.

- **Reduce subnets and simplify DHCP management**—In the overlay, IP subnets can be stretched across the fabric without flooding issues that can happen on large Layer 2 networks. Use fewer subnets and DHCP scopes for simpler IP addressing and DHCP scope management.

- **Avoid overlapping IP subnets**—Different overlay networks can support overlapping address space, but be aware that most deployments require shared services across all VNs and other inter-VN communication. Avoid overlapping address space so that the additional operational complexity of adding a network address translation device is not required for shared services and inter-VN communication.

CONTROL PLANE DESIGN

The fabric control plane node contains the database used to identify endpoint location for the fabric elements. This is a central and critical function for the fabric to operate. A control plane that is overloaded and slow to respond results in application traffic loss on initial packets. If the fabric control plane is down, endpoints inside the campus fabric will fail to establish communication to remote endpoints that are not in the local database.

You should enable control plane functionality in a dedicated control plane node residing in a services block or data center. For redundancy, you should deploy two control plane nodes to ensure high availability of the fabric. Follow best practices for data center design, including deploying the redundant control plane node virtual machines in separate compute nodes with top-of-rack switch redundancy. For the best application performance, the latency from the fabric edge node to the fabric control plane node should be less than 100 milliseconds. For small fabrics, you can deploy the control plane function at the border node, but monitoring of the HTDB, CPU, and memory is a critical operational factor to consider in a design consolidating the control plane node and border node functionality.

FABRIC BORDER DESIGN

The fabric border design is dependent on how the fabric is connected to the outside network. VNs inside the fabric should map to VRF-Lite instances outside the fabric. Depending on where shared services are placed in the network the border design will have to be adapted. For more information, see "End-to-End Virtualization Considerations," later in this guide.

INFRASTRUCTURE SERVICES

The campus fabric does not require any changes to existing infrastructure services, with the exception of the DHCP server. In a typical DHCP relay design, the unique gateway IP address determines the subnet address assignment for an endpoint, in addition to the location where the DHCP server should direct the offered address. In a campus fabric overlay network, the gateway IP address is not unique at each fabric edge device. Instead, the same IP address is used as an Anycast IP across all fabric edge devices within an overlay. This situation results in an ambiguous DHCP relay gateway destination when returning the DHCP offer from the DHCP server. To remove

the ambiguity, you should use DHCP option 82 with the information option for circuit ID insertion on the relay agent at the fabric edge. Adding the information provides additional sub-options to identify the specific source relay agent. You will have to make appropriate adjustments to the DHCP scope selection criteria in order to ensure predictable DHCP operation inside a campus fabric. For information about implementing the DHCP server scopes, DHCP options, and selection criteria, refer to campus fabric deployment guides.

CENTRALIZED WIRELESS INTEGRATION

A centralized design model, also known as a *local-mode model*, is recommended for large site deployments. The benefits of a centralized design include IP address management, simplified configuration and troubleshooting, and roaming at scale. In a centralized model, the WLAN controller and APs are both located within the same site. You can connect the WLAN controller to a data center services block or a dedicated block off of the campus core. Wireless traffic between WLAN clients and the LAN is tunneled by using the control and provisioning of wireless access points (CAPWAP) protocol between the controller and the AP. APs can reside inside or outside the fabric without any change to the recommended centralized WLAN design.

Figure 8   *Centralized wireless over campus fabric*
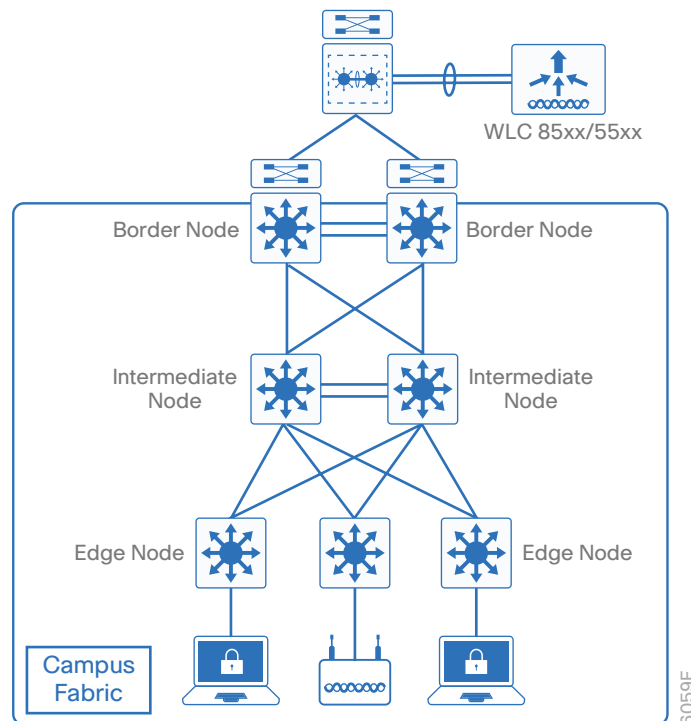
For additional information about campus wireless design, see the Campus LAN and Wireless LAN Design Summary.

> ### Tech Tip
>
> Converged Access and FlexConnect are not supported inside the campus fabric.

SECURITY/POLICY DESIGN

Security policies vary by organization—it is not possible to define one-size-fits-all security design. Security designs are driven by information security policies and legal compliance. The planning phase for a security design is key to ensuring the right balance of security and user experience. You should consider the following aspects designing your security policy for the campus fabric:

- **Openness of the network**—Some organizations allow only organization-issued devices in the network, and some support a "Bring Your Own Device" approach. Alternatively, you can balance user choice and allow easier-to-manage endpoint security by deploying a "Choose Your Own Device" model in which a list of IT-approved endpoints is offered to the users for business use. And an identity-based approach is also possible in which the network security policies can be deployed depending of the device ownership. For example, organization-issued devices may get group-based access, while personal devices may get Internet access only.

- **Identity management**—In the simplest form, identity management can be a username and password used for authenticating users. Adding embedded security functions and application visibility in the network devices provides telemetry for advanced policy definitions that can include additional context such as physical location, device used, type of access network, application used, and time of day.

- **Authentication, Authorization, and Accounting policies**—*Authentication* is the process of establishing and confirming the identity of a client requesting access to the network. *Authorization* is the process of authorizing the endpoint to some set of network resources. Segmentation policies do not necessarily have to be enforced at the access layer, and can be deployed in multiple locations. Policies are enforced with the use of security group ACLs (SGACLs) for segmentation and dynamic VLAN assignment for mapping endpoints into VNs at the fabric edge node.

- **Endpoint security**—Endpoints can be infected with malware, compromising data and create network disruptions. Malware detection, endpoint management, and data exports from the network devices provide insight into endpoint behavior. Tight integration of the network with security appliances and analytics platforms enable the network with the necessary intelligence to quarantine and help remediate compromised devices.

- **Data integrity and confidentiality**—Network segmentation can be used to control access to applications; encryption of the data path in the switching environment using IEEE 802.1AE is used to provide encryption at Layer 2 to prevent eavesdropping and to ensure that the data cannot be modified.

- **Network device security**—Hardening the security of the network devices is essential because they are common targets for security attacks. The use of the most secure device management options, such as enabling device authentication using TACACS+ and disabling unnecessary services, are best practices to ensure the network devices are secured.

The use of SGTs provides the capability to tag endpoint traffic based on group membership policies in Cisco Identity Services Engine. On most deployments, Active Directory is used as the identity store for user accounts, credentials, and group membership information. Group assignments can be created based on job role, which can be used to create segmentation policies and virtual network assignment rules. SGT information is carried across the network in several forms:

- **Inside the campus fabric—**The campus fabric header transports SGT information. Fabric edge nodes and border nodes can enforce SGACLs to enforce the security policy.

- **Outside of the fabric on a TrustSec-capable device—**Inline TrustSec-capable devices carry the SGT information in a CMD header on the Layer 2 frame. This is the recommended mode of transport outside of the campus fabric.

- **Outside of the fabric over devices without TrustSec capability—**SXP allows the transport of SGTs over a TCP connection. This can be used to bypass network devices that do not support SGT inline.

For additional information about Cisco TrustSec, see www.cisco.com/go/trustsec.

End-to-End Virtualization Considerations

In a virtualized network there is full isolation of data and control planes over a shared networking infrastructure. In the case of the campus fabric, a user on one VN is completely isolated and will not be able to communicate with a user on a different VN. The fabric border node is responsible for extending network virtualization beyond the campus fabric. Organizations may have business requirements that call for this type of isolation. Some example of vertical specific use cases where network virtualization maybe useful include:

- **Education**—College campus divided into administrative and student residence networks.

- **Retail**—Isolation for point-of-sale machines supporting payment card industry compliance.

- **Manufacturing**—Isolation for machine-to-machine traffic in manufacturing floors.

- **Healthcare**—Dedicated networks for medical equipment, patient guest access and HIPAA compliance.

- **Enterprise**—Integration of networks during mergers, where overlapping address spaces may exist. Separation of building control systems and video surveillance devices.

Designing for end-to-end network virtualization requires detailed planning in order to ensure the integrity of the virtual networks. In most cases, there is a need to have some form of shared services that can be reused across multiple virtual networks. It is important that those shared services are deployed correctly in order to preserve the isolation between different virtual networks sharing those services. Although route leaking provides a quick way to extend prefixes of shared services across multiple networks, the use of firewalls provides an additional layer of security and monitoring of traffic between virtual networks. Examples of shared services include:

- **Wireless infrastructure**—Radio frequency performance and cost efficiency is increased using common wireless LANs (single SSID). Traffic isolation is achieved by assigning dedicated VLANs at the WLC and using dynamic VLAN assignment using 802.1X authentication to map wireless endpoints into their corresponding VNs.

- **DHCP, DNS, and IP address management**—The same set of infrastructure services can be reused as long as they have support for virtualized networks. Special capabilities such as advanced DHCP scope selection criteria, multiple domains, and support for overlapping address space are some of the capabilities required to extend the services beyond a single network.

- **Internet access**—The same set of Internet firewalls can be used for multiple virtual networks. If firewall policies need to be unique for each virtual network, the use of a multi-context firewall is recommended.

- **IP communications**—When IP phones and other unified communications devices are connected in multiple virtual networks, the call control signaling to the communications manager and the IP traffic between those devices needs to be able to traverse multiple VNs in the infrastructure.

NETWORK VIRTUALIZATION TECHNOLOGIES

Network virtualization technologies vary based on the place in the network. For example, network virtualization in the data center typically uses different technologies compared to network virtualization in the WAN because of the unique requirements of the type of clients attached to the virtual network. An overlay network interconnecting servers in a data center will have a different set of requirements from an overlay network interconnecting remote site routers over a public network. Guidance for virtualizing your end-to-end network is beyond the scope of this guide. However, this section provides a brief introduction to the most commonly used technologies that you can investigate when virtualizing your network.

Device Level Virtualization

Within the same device physical device, logical separation capabilities at Layer 2 and Layer 3 can be used to extend virtual networks:

Virtual LANs

The most basic form of device-level virtualization is isolating network traffic using different virtual LANs (VLANs). This form of virtualization applies to Layer 2 devices and can be extended across switched domains. VLANs are also used to virtualize point-to-point links between routers and security appliances that require connectivity to multiple virtual networks via the same physical interface.

Virtual Routing and Forwarding

VRF is a device-level virtualization technology for creating multiple Layer 3 routing tables on the same device. VRFs can be tied to existing Layer 2 domains in order to provide Layer 3 edge functionality to multiple VLANs and also between Layer 3 routed interfaces in order to extend a multiple virtualized control plane over the same set of interfaces.

Path Isolation

To maintain isolation on the paths of links interconnecting devices, there are many technology options that provide network virtualization among devices. For the campus fabric, the recommended path-isolation technologies are VRF-Lite and MPLS VPN. The number of virtualized networks required typically dictates the design. If you forecast a need for more than a few VRFs, deploying MPLS VPNs simplifies configuration and management.

VRF-Lite End-to-End

VRF-Lite is deployed on a hop-by-hop basis in a campus network, making use of 802.1Q trunks between devices in order to isolate data and control plane for each virtual network. For ease of operation, you should use the same set of VLANs across every hop and use BGP with per-VN address families providing attributes that can be leveraged for easy route-leaking for shared services.

MPLS

Although often considered a service-provider technology, MPLS is common on larger enterprises needing a large number of virtualized networks, most commonly in the WAN but also extended to the campus network. While VRF-Lite is common to most routing platforms, MPLS is not supported across all platforms. A combination of VRF-Lite at the edge with MPLS VPN is another design that should be considered.

> ### Tech Tip
>
> Initial versions support only VRF-Lite handoff at the fabric border. For more information, see the software release notes for your platform.

Appendix–Glossary

AAA authentication, authorization, and accounting

ACL access control list

AP access point

BGP border gateway protocol

CMD Cisco Meta Data

DNA Cisco Digital Network Architecture

EID endpoint identifier

HTDB host tracking database

IGP interior gateway protocol

LISP Locator/ID Separation Protocol

MR Map-Resolver

MS Map-Server

MTU maximum transmission unit

RLOC routing locator

SGACL security group access control list

SGT security group tag

SXP security group tag exchange protocol

VLAN virtual local area network

VN virtual network

VNI virtual extensible LAN network identifier

VRF virtual routing and forwarding

VTEP tunnel endpoint

VXLAN virtual extensible LAN

Please use the [feedback form](#) to send comments and suggestions about this guide.