Boson®

ENCOR
Curriculum

350-401

Labs powered by

NetSim®
NETWORK SIMULATOR®

Boson®

Implementing Cisco Enterprise Network Core Technologies (ENCOR)

350-401 Curriculum

LM20201031/BV1.0

Boson® NetSim® NETWORK SIMULATOR®

25 Century Blvd., Ste. 500, Nashville, TN 37214 | Boson.com

The labs referenced in this book have been printed in the Boson Lab Guide, which is included with the purchase of the curriculum. These labs can be performed with real Cisco hardware or in the Boson NetSim Network Simulator version 11 or later. To learn more about the benefits of using NetSim or to purchase the software, please visit www.boson.com/netsim.

Module 1

Architecture

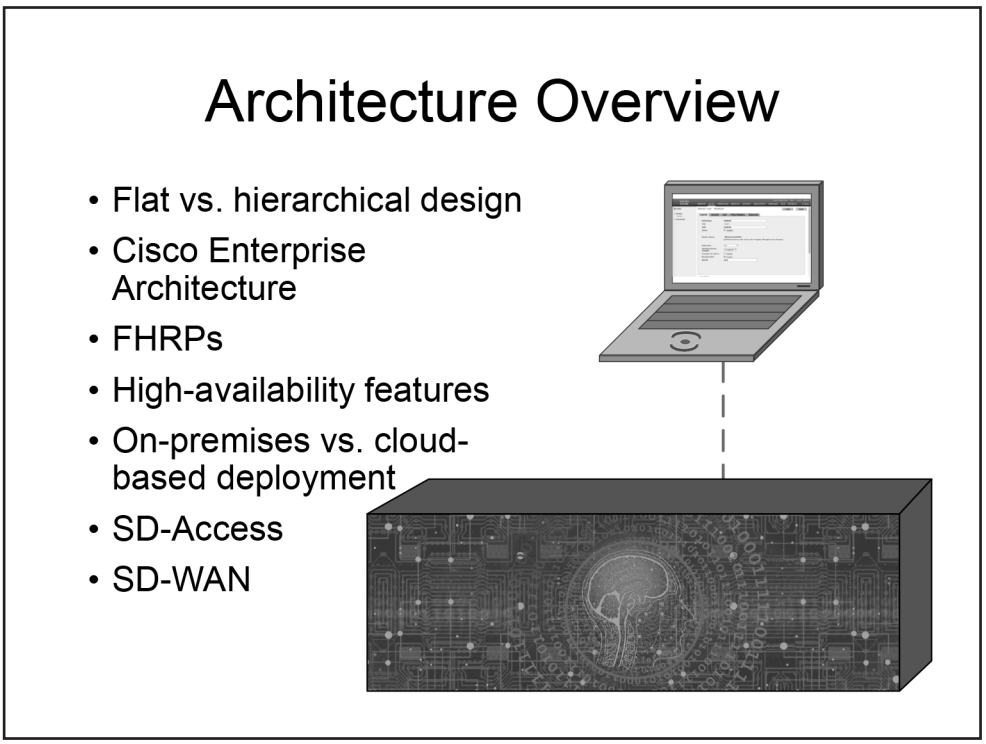

Overview

Current best practice eschews flat network designs in favor of leveraging the resilience and simplicity of hierarchical design patterns. A hierarchical design, such as Cisco's Enterprise Campus Architecture model, relies on modular, structured components to deliver efficient, flexible, and manageable network solutions. This module explores several approaches to LAN and WAN architecture design.

Objectives

After completing this module, you should have the basic knowledge required to complete all of the following tasks:

- Understand flat and hierarchical network design.
- Understand the Cisco Enterprise Architecture model.
- Understand First-Hop Redundancy Protocols (FHRPs).
- Understand high-availability features.
- Understand the advantages of on-premises versus cloud-based deployments.
- Understand Software-Defined Access (SD-Access).
- Understand Software-Defined WAN (SD-WAN).

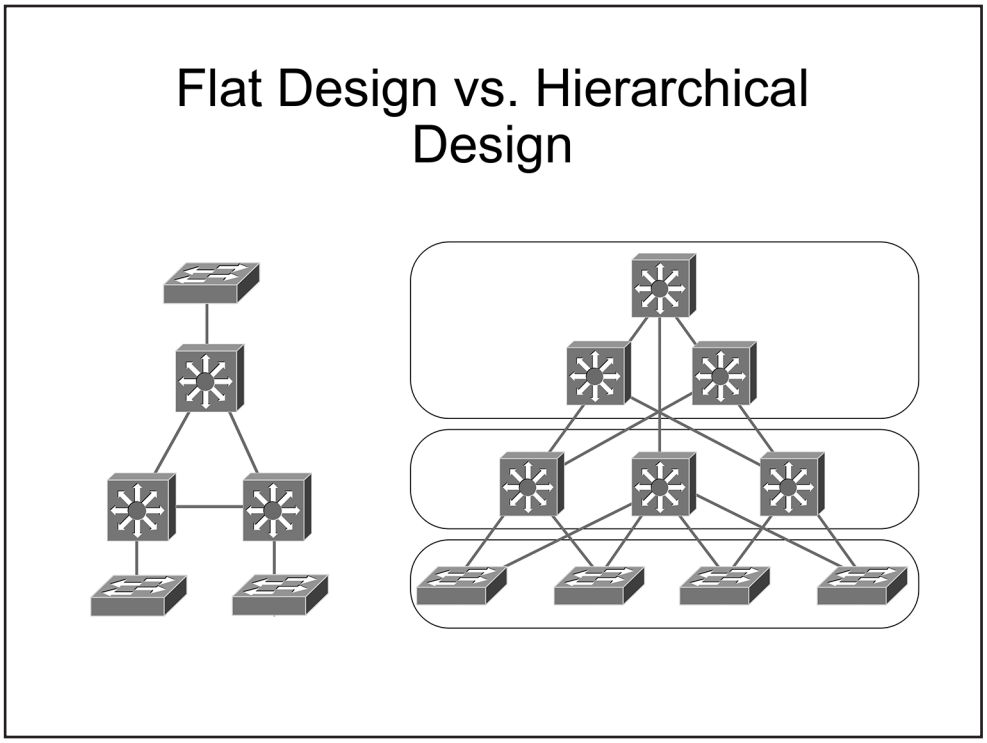

Flat Design vs. Hierarchical Design

Many early network design models produced flat topologies in which network devices often shared a single broadcast or collision domain. These design models were constrained by the shared-medium access technologies of their day and by the high cost of devices, such as bridges, switches, and routers, that might have provided a necessary level of network segmentation. Without adequate network segmentation, the scalability of a flat topology is limited. The bandwidth and CPU resources that are required to process broadcast traffic increase exponentially even as the network undergoes linear growth. Adding and maintaining resilient and efficient paths for network traffic as the network scales is difficult, if not impossible, in a flat topology because of the lack of both physical and logical structure. In addition, the lack of logical structure reduces the overall manageability of the network; managing the impact of adding, modifying, or removing network devices becomes increasingly more complex as the network grows and can result in large-scale design revisions.

The advent of inexpensive, hardware-based routing and switching platforms have enabled network designers to produce hierarchical network design models to address the scalability, efficiency, and manageability limitations of older design models. In a hierarchical network design model, network devices are logically and physically grouped into a layered, fault-tolerant structure based on their functionality or role within the topology.

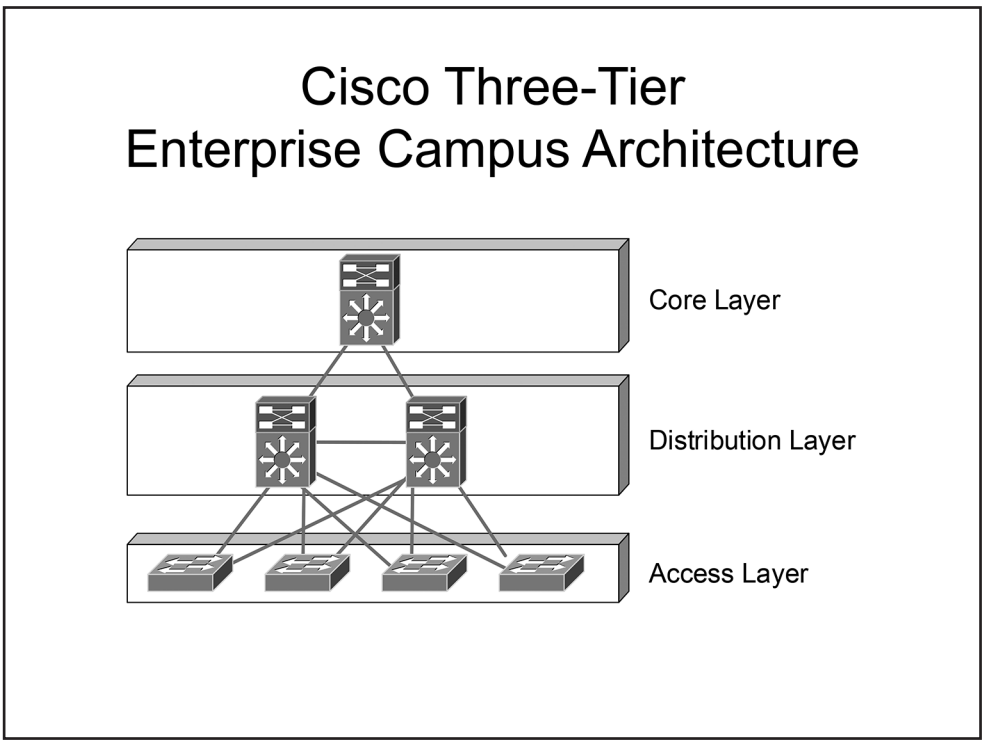

Cisco Three-Tier Enterprise Campus Architecture

The Cisco Enterprise Campus Architecture is an example of a hierarchical design model that is commonly in use today. Switches and routers are used to segment network devices into a layered, modular topology wherein collision and broadcast domains can be optimized to increase network efficiency and scalability.

Cisco recommends using a three-tier Enterprise Campus Architecture model with the following layers:

- Core layer
- Distribution layer
- Access layer

Dividing a network design into layers simplifies the scalability and troubleshooting of a network. Modularity ensures that network components can be added, modified, or removed without a high probability of drastic revisions to design or implementation details. In addition, a layered, hierarchical structure reduces management complexity and facilitates the implementation of redundancy and optimized data paths.

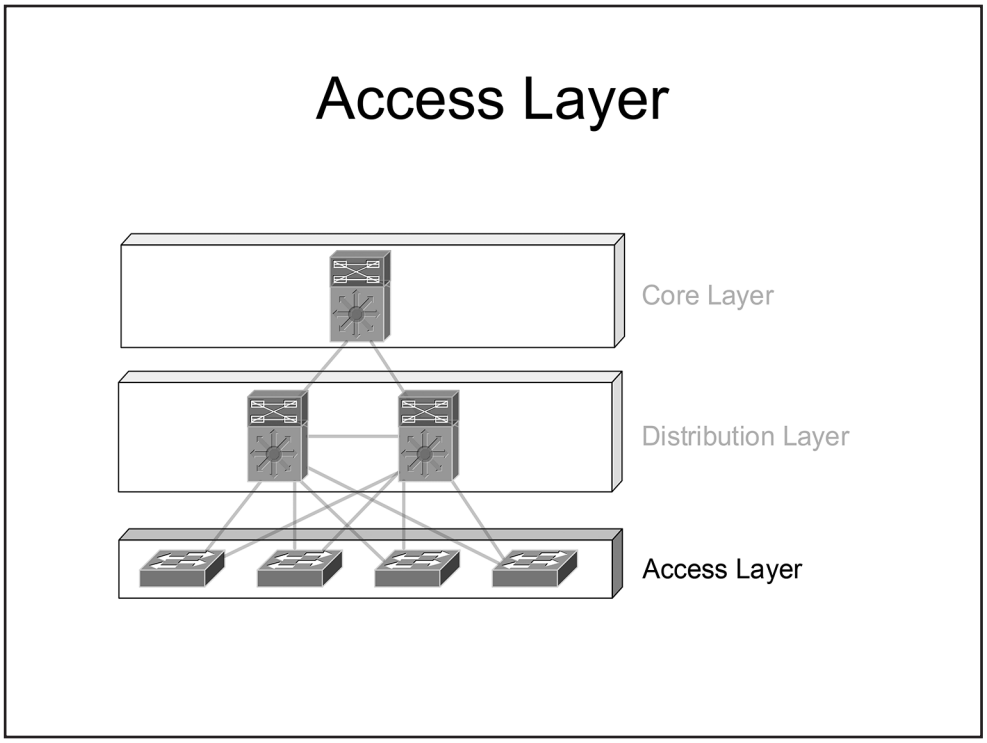

Access Layer

The Access layer uses switches and wireless access points (WAPs) to provide network connectivity for end-user devices such as computers, printers, and IP telephones. Because Access layer devices provide users with access to the network, the Access layer is the ideal place to perform user authentication, port security, Network Admission Control (NAC), Quality of Service (QoS) classification and marking, and Power over Ethernet (PoE). NAC is used to prevent access to the network from unauthorized end-user devices. QoS classification and marking ensure that different categories of traffic, such as voice, video, and data, receive enough bandwidth to ensure acceptable levels of service throughout the network. PoE is used to deliver power to end-user devices, such as IP phones, over the same physical cable that is used for network traffic.

Virtual LANs (VLANs) are implemented in the Access layer to provide network segmentation and logical organization. For example, VLANs can be used to separate traffic from different departments within an organization. In addition, VLANs are commonly used in converged networks to separate voice, data, and wireless traffic.

Traditionally, the Access layer consists of Open Systems Interconnection (OSI) Layer 2 switches only; when a packet must be routed to a separate network, the packet must first be sent to a Layer 3 device in the Distribution layer before it can be routed to the correct destination. However, some designs employ Layer 3 switches in the Access layer, which in effect moves the demarcation between Layer 2 and Layer 3 switching to the Access layer.

Cisco recommends using redundant network links at the Access layer to ensure high availability. Servers and other critical end-user devices should connect to Access layer switches over redundant network links. Likewise, Access layer switches should connect to the Distribution layer over redundant network links to ensure that the failure of a single uplink or network device will not result in Access layer devices being isolated from the rest of the network.

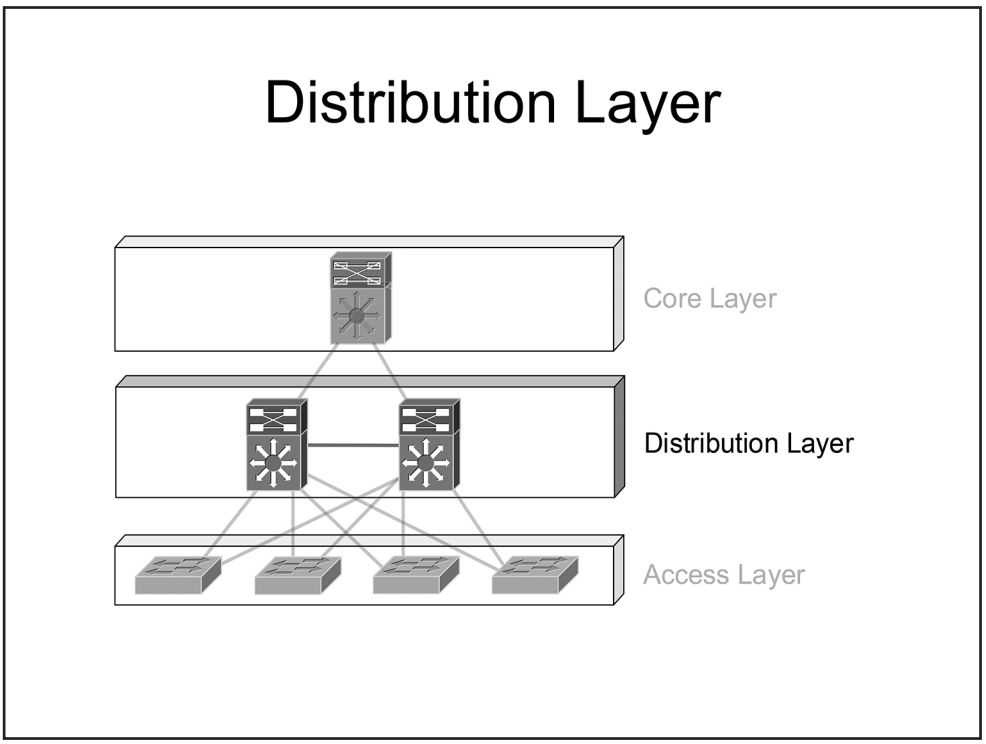

Distribution Layer

The Distribution layer, which is sometimes referred to as the aggregation layer, is positioned between the Access and Core layers. Traditionally implemented with multilayer switches, the Distribution layer provides Layer 3 uplinks to the Core layer and aggregates Layer 2 links from the Access layer. Interconnecting Access layer devices at the Distribution layer instead of through a mesh of direct connections reduces network complexity and increases the scalability and manageability of the network. Troubleshooting is simplified because failed links or devices are more easily isolated and affect less of the overall network than in a flat network topology.

The Distribution layer is responsible for performing QoS resource reservation, interVLAN routing, packet manipulation, route filtering, and summarization. In addition, access control lists (ACLs) and Intrusion Prevention System (IPS) filtering can be implemented in the Distribution layer to enforce organizational security policies and to filter the flow of traffic between the Access and Core layers.

In a three-tier network design, the default gateway for Access layer devices resides in the Distribution layer. However, because Access layer devices typically have redundant Layer 2 paths to the Distribution layer, an FHRP is recommended in order to prevent a link or device failure from disconnecting Access layer devices from their configured default gateways. Alternatively, the Distribution layer can aggregate Layer 3 links from the Access layer, in which case an FHRP would not be required.

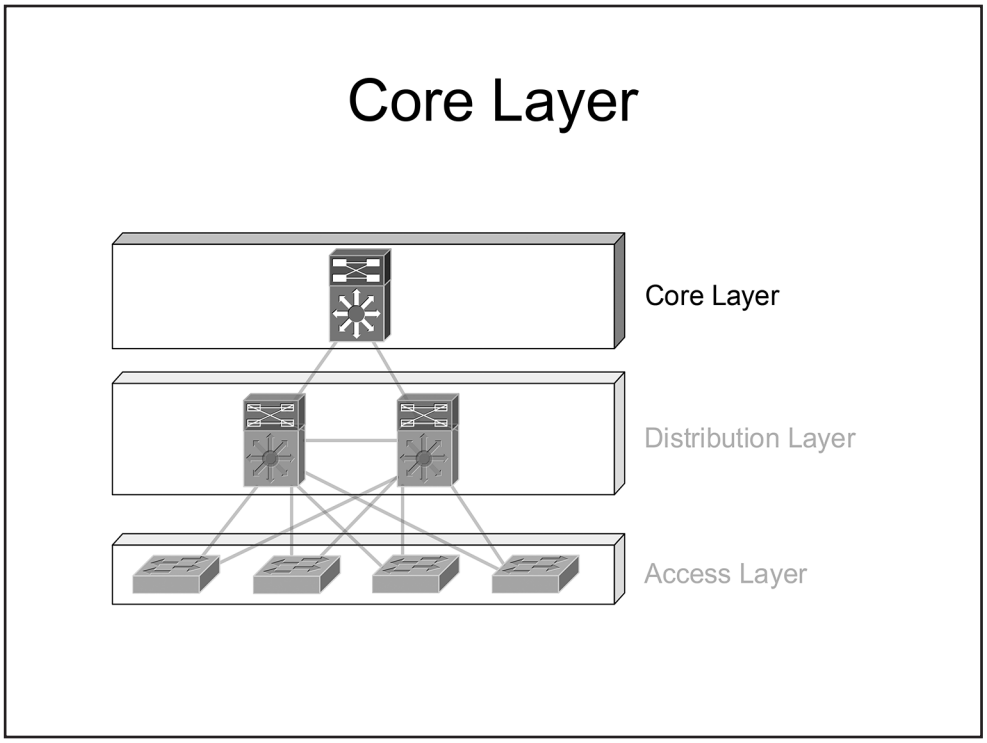

Core Layer

The Core layer typically provides the fastest switching path in the network. As the network backbone, the Core layer is primarily associated with low latency and high reliability.

The Core layer provides a high-speed backbone that interconnects Distribution layer devices. The Core layer should be optimized to minimize the latency through the backbone; thus very little packet manipulation and route processing should occur in the Core layer. The Core layer is dependent on the Distribution layer to enforce security and QoS policies and to perform packet manipulation and route optimization.

Typically implemented as a mesh of multilayer switches, the Core layer is highly scalable because it requires only minimal connectivity to each Distribution layer switch. Without a Core layer, the number of interconnections between Distribution layer switches to ensure high availability would quickly grow unmanageable and would drastically limit network scalability. However, with a Core layer, Distribution layer switches do not need to be directly interconnected; they can rely on the redundancy of the Core layer to ensure connectivity.

The Core layer consists entirely of Layer 3 connections between Core layer devices and Distribution layer devices. This configuration results in an inherent lack of Layer 2 loops and relies on the deterministic convergence of Layer 3 routing protocols to guarantee high availability.

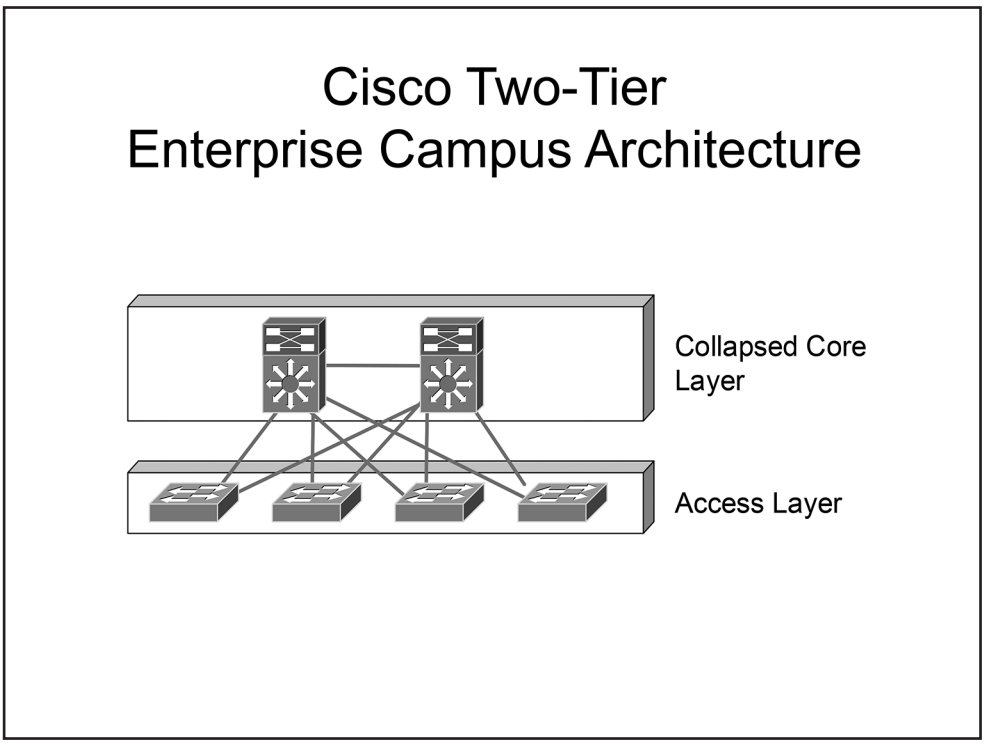

Cisco Two-Tier Enterprise Campus Architecture

The Cisco two-tier Enterprise Campus Architecture model is sometimes called the collapsed-core network design model. In a two-tier network design model, the functionality of the Core layer is collapsed into the Distribution layer. The functionality of the Core layer is provided by the Distribution layer, and a distinct Core layer does not exist.

Maintaining a full mesh of connections between Distribution layer switches grows exponentially more difficult as the number of Distribution layer switches is increased. Therefore, a collapsed core topology does not scale beyond a small number of switches. In addition, the large number of routing peers or neighbors in a full mesh increases the complexity of the routing configuration within the collapsed core.

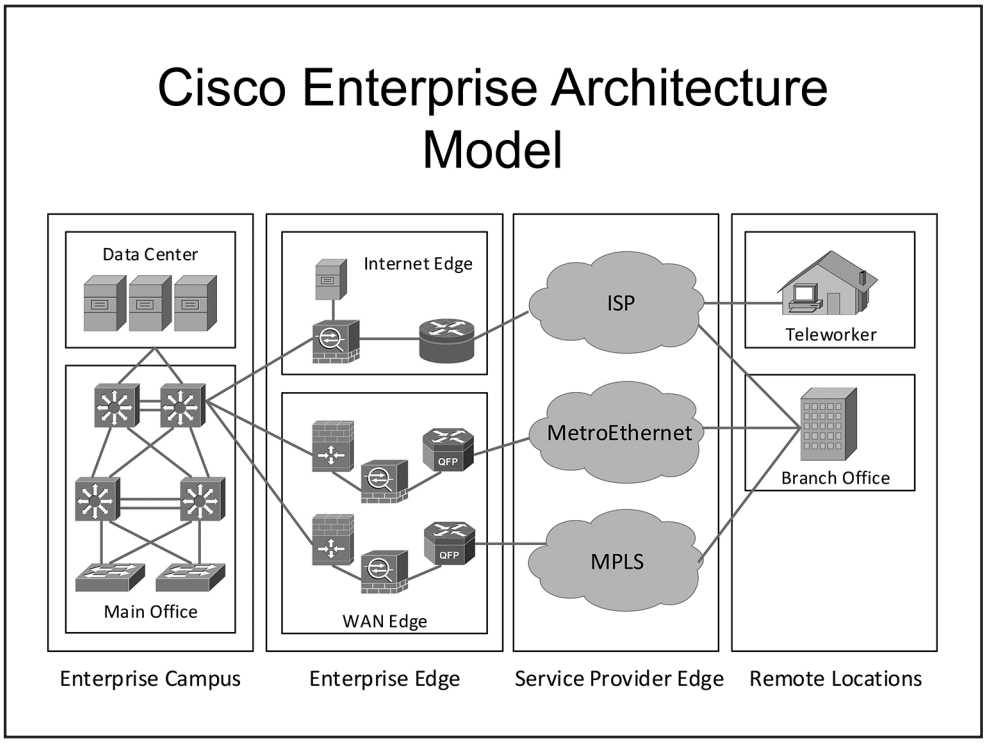

Cisco Enterprise Architecture Model

The Cisco Enterprise Architecture Model consists of the following four modules:

- Enterprise Campus
- Enterprise Edge
- Service Provider Edge
- Remote Locations

The Enterprise Campus module follows the three-tier hierarchical model with Core, Distribution, and Access layers. In addition, the data center submodule provides network services, such as application, email, Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), and file services.

The Enterprise Edge module consists of the Internet Edge and WAN Edge submodules. These submodules provide connectivity between the service provider infrastructure and the internal network. Multiple service providers and multiple paths to these service providers can be implemented for redundancy.

The Service Provider Edge module provides connectivity between the enterprise location and remote locations. Connectivity is provided by an Internet, Multiprotocol Label Switching (MPLS), Metro Ethernet, or other WAN service provider.

The Remote Locations module consists primarily of teleworkers and branch offices. However, any part of the enterprise network that is geographically distant from the main headquarters can be included within this module.

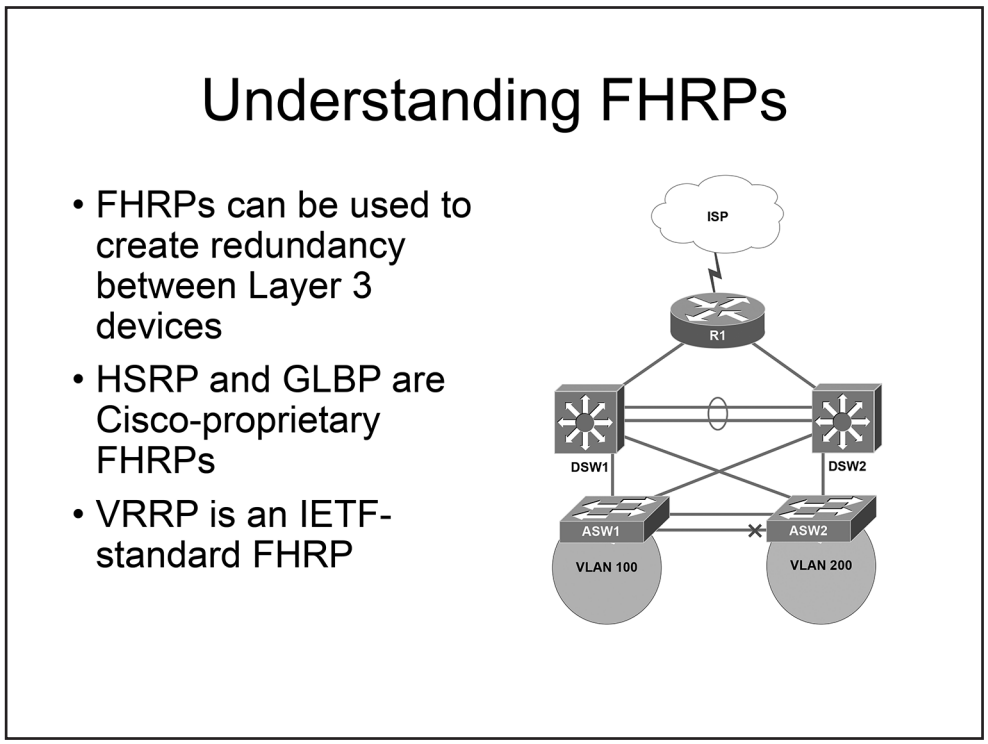

Understanding FHRPs

In a routed Access layer design, Layer 3 switches are implemented in the Access layer. Because the Layer 3 boundary exists at the Access layer, Access layer devices can use a directly connected Layer 3 switch as a default gateway and the Layer 3 switch can use dynamic routing protocols, such as Enhanced Interior Gateway Routing Protocol (EIGRP), to provide redundancy and load sharing. However, in a switched Access layer design, the Layer 3 boundary exists at the Distribution layer. Therefore, Access layer devices must use a default gateway that is no longer directly connected. If access to a default gateway is disrupted in a switched access design, any Access layer devices using that gateway will lose connectivity to devices outside of their own VLAN or IP subnet unless they are manually reconfigured with an alternate default gateway or until access to the original gateway is restored.

FHRPs can mitigate the need to manually reconfigure gateway information on Access layer devices in a switched Access layer design. FHRPs provide a mechanism for Layer 3 gateway redundancy and load sharing that is transparent to Access layer devices. Hot Standby Router Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP), and Gateway Load Balancing Protocol (GLBP) are all FHRPs. Both HSRP and GLBP are Cisco-proprietary FHRPs. VRRP, on the other hand, is an Internet Engineering Task Force (IETF)-standard FHRP that is supported by both Cisco and non-Cisco devices. However, it is important to note that Cisco-specific enhancements to VRRP might not be available when a Cisco device that is using VRRP is connected to a non-Cisco device that is using VRRP. If only Cisco devices are used in the topology and a choice between HSRP and VRRP is available, Cisco recommends using HSRP.

Because FHRPs are commonly configured on Layer 3 switches and routers, both devices will be referred to as routers throughout this module.

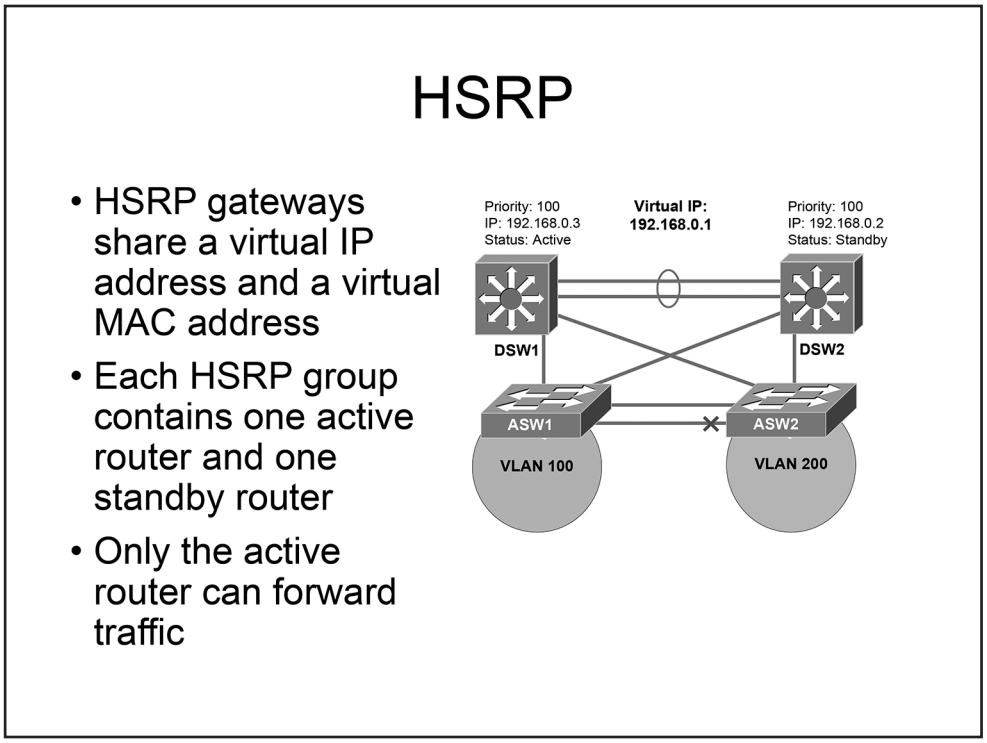

HSRP

HSRP enables multiple Layer 3 devices, such as routers or Layer 3 switches, to act as a single gateway for the network by sharing a single virtual IP address and a single virtual Media Access Control (MAC) address between redundant gateways in an HSRP group. Each HSRP group contains a single active router and a single standby router. Only the active router responds to Address Resolution Protocol (ARP) requests and forwards traffic. If the active router fails, the standby router assumes the role of the active router and a new standby router is elected among the remaining routers in the HSRP group.

HSRP will be discussed in detail in Module 6: IP Services.

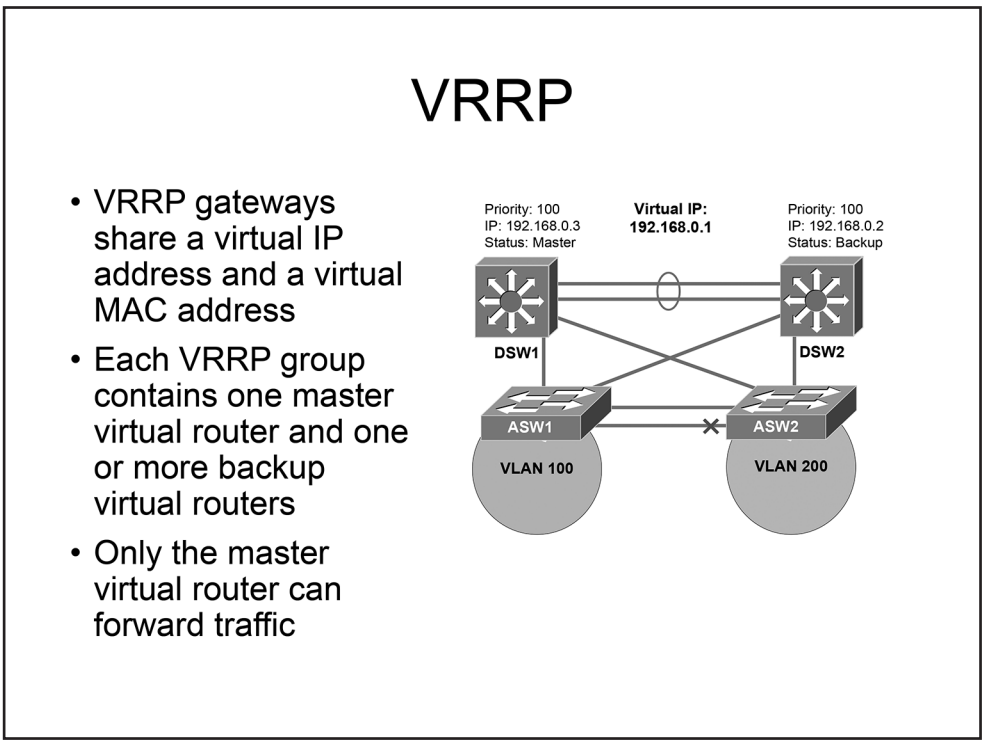

VRRP

VRRP is an open-standard FHRP that provides much of the functionality offered by Cisco's proprietary HSRP protocol. Each VRRP group contains a single master virtual router and one or more backup virtual routers to provide first-hop redundancy for client devices. Only the master virtual router responds to ARP requests and forwards traffic. If the master virtual router fails, one of the backup virtual routers will assume the role of master virtual router and will begin forwarding traffic.

VRRP will be discussed in detail in Module 6: IP Services.

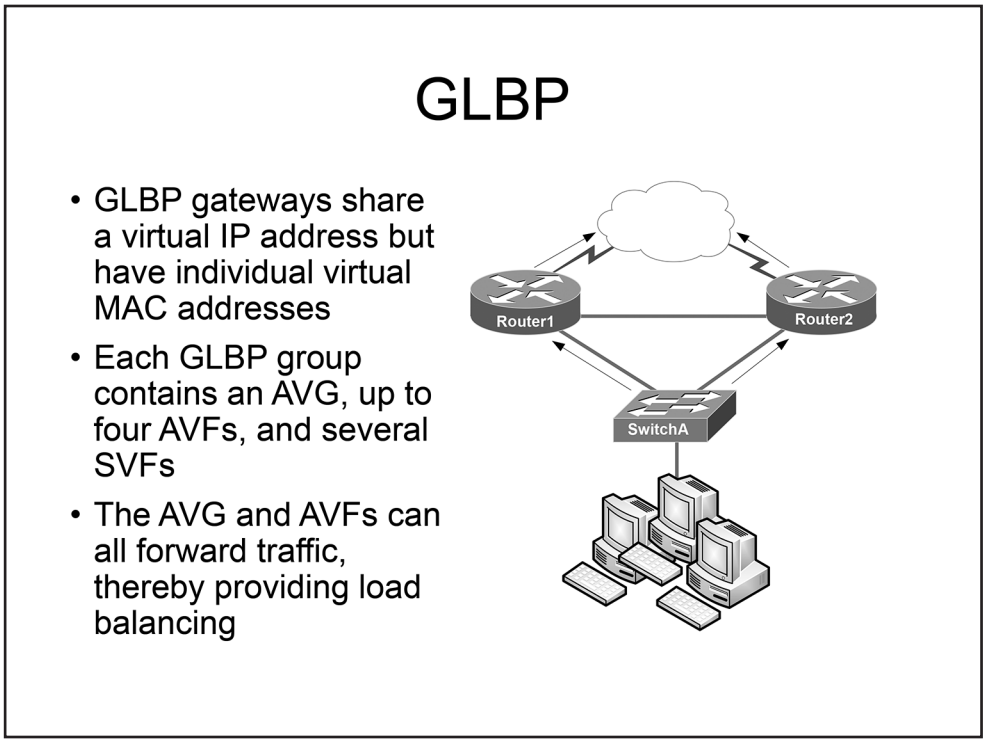

GLBP

GLBP is a Cisco-proprietary protocol that provides router redundancy and load balancing. The routers in a GLBP group receive traffic sent to a virtual IP address that is configured for the group. However, unlike HSRP and VRRP, GLBP is capable of load balancing traffic among every router in the group.

Each GLBP group contains an active virtual gateway (AVG) and up to four primary active virtual forwarders (AVFs); the remaining routers become secondary virtual forwarders (SVFs). The AVG assigns a virtual MAC address to the primary AVFs. When the AVG receives ARP requests that are sent to the virtual IP address for the GLBP group, the AVG responds with a different virtual MAC address, including its own. This provides load balancing, because each of the primary AVFs can participate by forwarding a portion of the traffic sent to the virtual IP address.

If the AVG fails, the AVF with the next highest priority, which is referred to as the standby virtual gateway (SVG), will take over for the AVG. If one of the AVFs fails, the AVG will assign another router the role of that AVF.

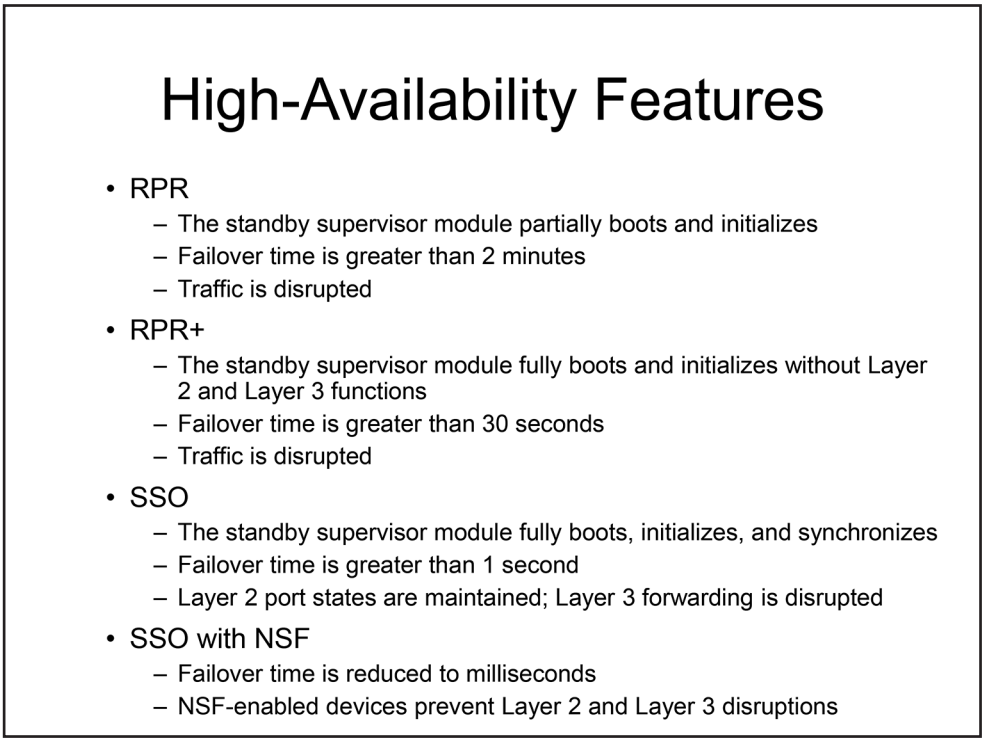

High-Availability Features

When a network device fails, network instability can occur. For example, a router failure can cause routing protocols to recalculate paths throughout the network. Cisco devices that include redundant supervisor modules can be configured to use one of the following redundancy modes:

- Route-processor redundancy (RPR)
- RPR Plus (RPR+)
- Stateful switchover (SSO)
- SSO with Nonstop Forwarding (NSF)

When RPR is configured, the standby supervisor module is partially booted, and the startup configuration, boot variables, configuration register, and VLAN database are synchronized with the active supervisor module. When the active supervisor module fails, the standby module will load the other modules and will initialize the supervisor functions. Failover time for RPR mode is greater than two minutes; during this time, traffic is disrupted as physical ports restart.

With RPR+, the standby supervisor module boots and initializes supervisor functions and the route engine. However, Layer 2 and Layer 3 functions will not start until the active module fails. Failover time for RPR+ mode is greater than 30 seconds.

With SSO, the standby supervisor module fully boots and initializes. The startup and running configurations are synchronized between the active and standby modules. Failover time for SSO mode is greater than one second; during this time, the Layer 2 port states will be maintained, but Layer 3 packet forwarding will be disrupted.

To prevent Layer 3 disruptions, NSF can be enabled with SSO. When a supervisor module failover occurs on an NSF-enabled device, NSF continues to use the Cisco Express Forwarding (CEF) database to route traffic for a short time until the Routing Information Base (RIB) and Forwarding Information Base (FIB) are rebuilt. Failover time for SSO with NSF can be less than 150 milliseconds.

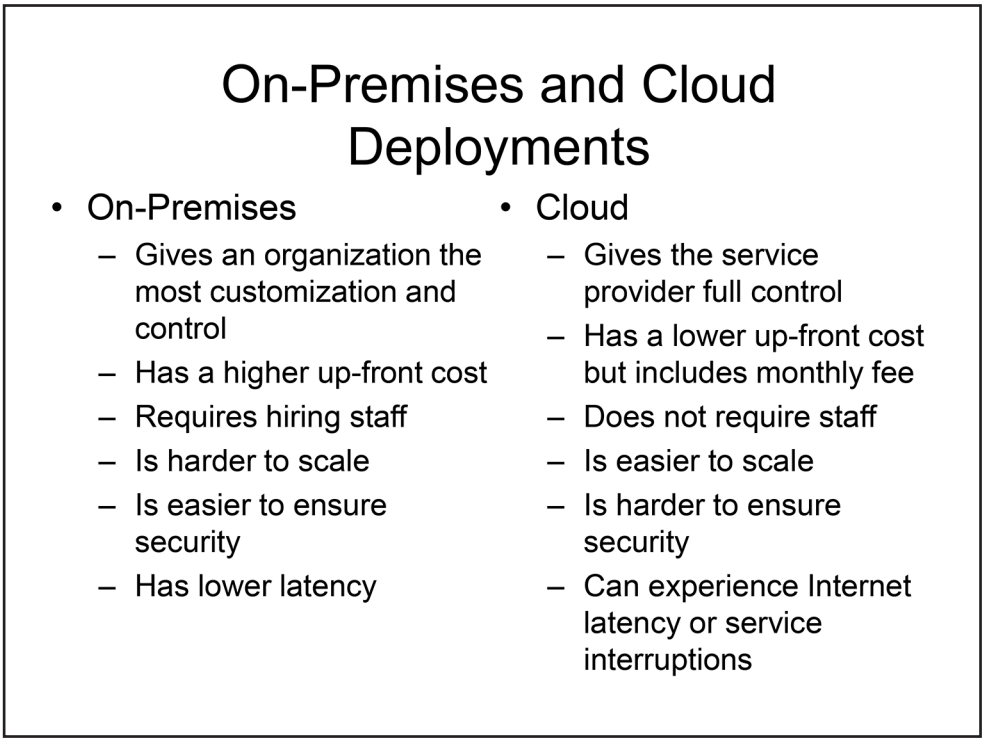

On-Premises and Cloud Deployments

On-premises deployments, which are also known as traditional deployments, involve the purchase, configuration, and maintenance of the deployment at the local level. This means that an organization has full control over the hardware or software solution. For example, a company that deploys an on-premises badge access solution will have complete control of the system and its data. However, the company must hire staff to maintain the deployment. In addition, if the hardware or software is overutilized or underutilized, it might be difficult or costly to right-size the hardware or software resources. However, a company might be required to choose an on-premises deployment in order to ensure full compliance with regulations or to maintain full security control.

A cloud deployment is owned and maintained by the cloud-hosting provider, meaning the provider has control over both the hardware and software. Therefore, a cloud deployment is less likely to offer an organization more customization and control than an on-premises deployment. However, cloud deployments are quite scalable, as hardware or software resources can be added or removed quickly. A cloud deployment does not require an investment in new hardware or software, so cloud deployments often have a lower up-front cost than on-premises deployments. In addition, the customer does not have to hire, train, and employ technical staff to maintain the deployment. However, cloud hosting providers charge the customer a monthly service fee.

Although cloud deployments can decrease operational costs, they can increase risks. Because cloud-based resources exist on the Internet, latency can occur, causing slower access to cloud resources. If Internet service is lost, access to cloud resources will also be lost. Also, a company's confidential data might be stored on a third-party server for which the company does not have full administrative control; if security is not adequate, data breaches can occur.

SD-Access

- Provides the following
  - Network automation
  - Network segmentation
  - Analytics
  - Identity services
  - Policy enforcement
- Consists of the following layers
  - Management layer
  - Controller layer
  - Network layer
  - Physical layer

SD-Access

Traditional network configuration and maintenance typically involves an administrator either physically connecting a console to a single device or remotely connecting a management application to the device in order to issue commands. The output of commands must then be interpreted by the administrator to either verify or troubleshoot a configuration.

As networks become more complex, they become harder and harder for administrators to manage and troubleshoot. Configurations can change over time as devices and applications are added to the network, and these changes and additions can sometimes cause network instability and downtime.

SD-Access helps to simplify, standardize, and secure networks by providing network automation, network segmentation, analytics, identity services, and policy enforcement. SD-Access consists of a Cisco campus fabric solution that is managed by a Cisco Digital Network Architecture (DNA) Center controller. This type of controller-based network is also known as Software-Defined Networking (SDN).

Controller-based networks are different from traditional networks in that a central component, the controller, is responsible for all network decision-making. Aside from the general benefit of easing administrative overhead, controller-based networking and network automation provide the specific benefit of using data models that are formalized and defined by a central controller. This means that configurations can be more reliably deployed than they are in traditional networks.

Conceptually, the SD-Access architecture can be segmented into the following four layers:

- Management layer
- Controller layer
- Network layer
- Physical layer

These layers will be discussed in detail in the following sections.

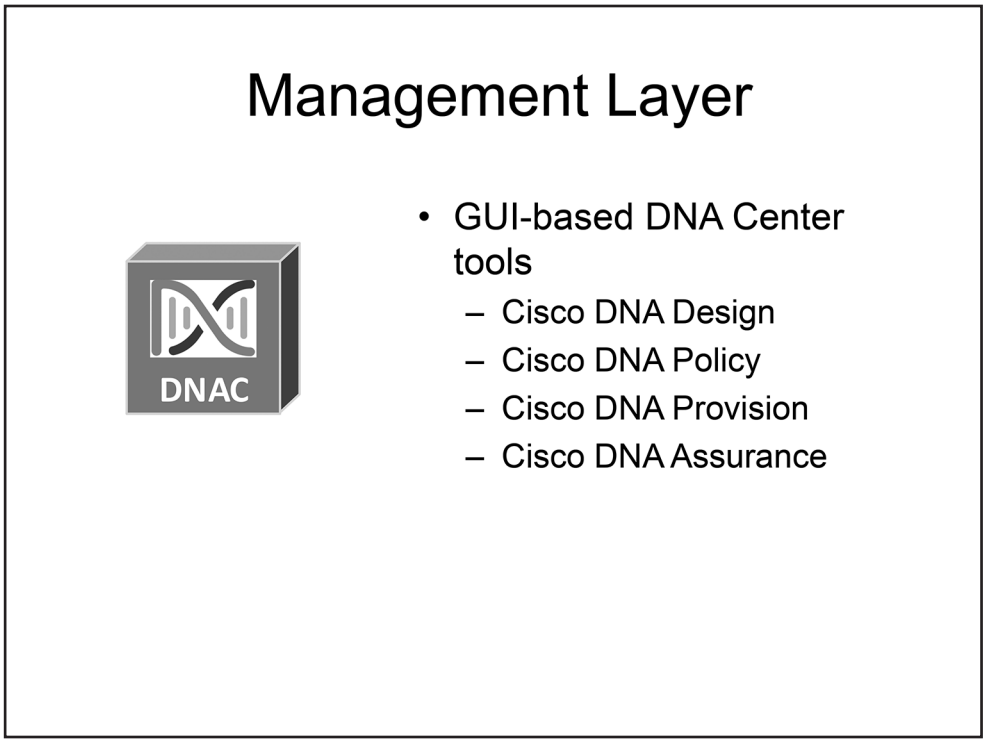

Management Layer

The management layer consists of the Cisco DNA Center tools that can be used by administrators to manage the network. Unlike traditional Cisco IOS-based implementations, the Cisco DNA Center tools are provided in a GUI. These management tools help to abstract and simplify the complex network interactions that exist on the other layers.

Cisco DNA Center applications are divided into four primary workflows:

- Cisco DNA Design
- Cisco DNA Policy
- Cisco DNA Provision
- Cisco DNA Assurance

These workflows will be covered in detail in Module 7: Network Assurance.

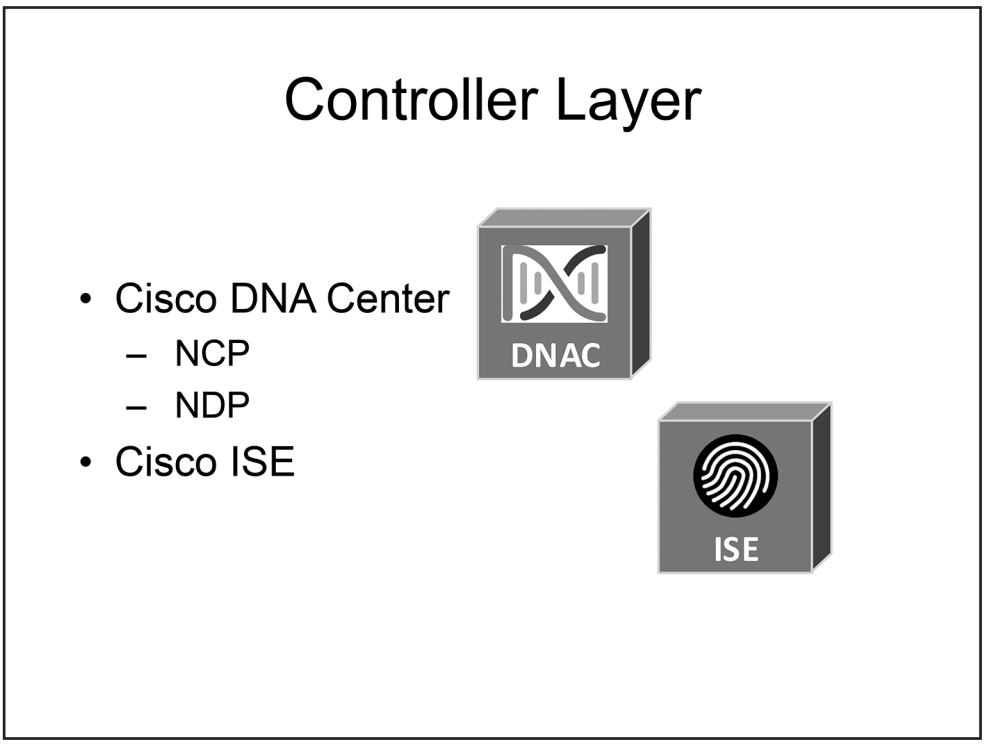

Controller Layer

The controller layer provides the two systems that are used by the management layer: Cisco DNA Center and the Cisco Identity Services Engine (ISE). Cisco DNA Center contains two subsystems that operate at the controller layer:

- Cisco Network Control Platform (NCP)
- Cisco Network Data Platform (NDP)

NCP provides the underlay and fabric automation for the network layer and physical layer. Application Programming Interfaces (APIs) facilitate the automation of management tasks by enabling the controller to communicate with applications rather than relying on an administrator to manually intervene. NCP configures network devices by using a southbound API, such as NETCONF. NCP communicates status information to the management plane by using a northbound API, such as Representational State Transfer (REST).

NDP collects data from multiple sources, such as NetFlow and Switched Port Analyzer (SPAN). NDP then analyzes the data and presents it in a contextualized format to NCP and ISE. NDP also provides status messages to the management layer.

Cisco ISE provides NAC and identity services by using 802.1x, MAC Authentication Bypass (MAB), and Web Authentication (WebAuth). ISE also provides policy services; administrators create group-based policies at the management layer, and ISE will translate those policies into configuration changes that will be automatically applied to network devices within the fabric.

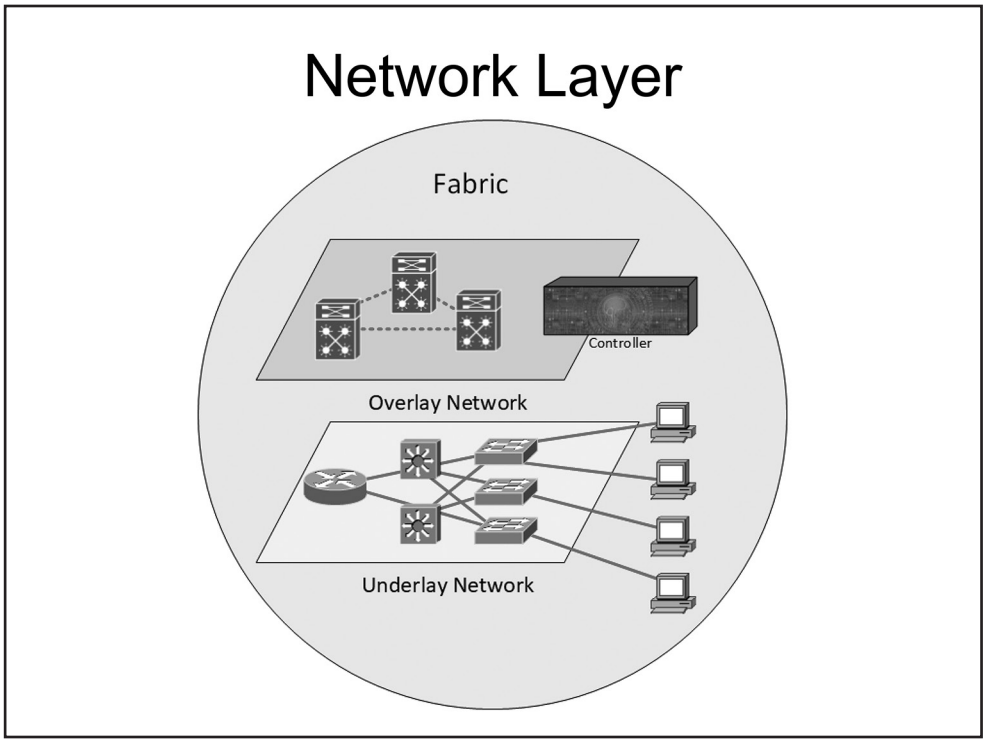

Network Layer

The network layer consists of the underlay network and the overlay network. Together, the underlay network and the overlay network are considered the SDN fabric. When the SDN architecture is layered in a diagram, the lower layer in the diagram is the underlay network and the higher layer in the diagram is the overlay network.

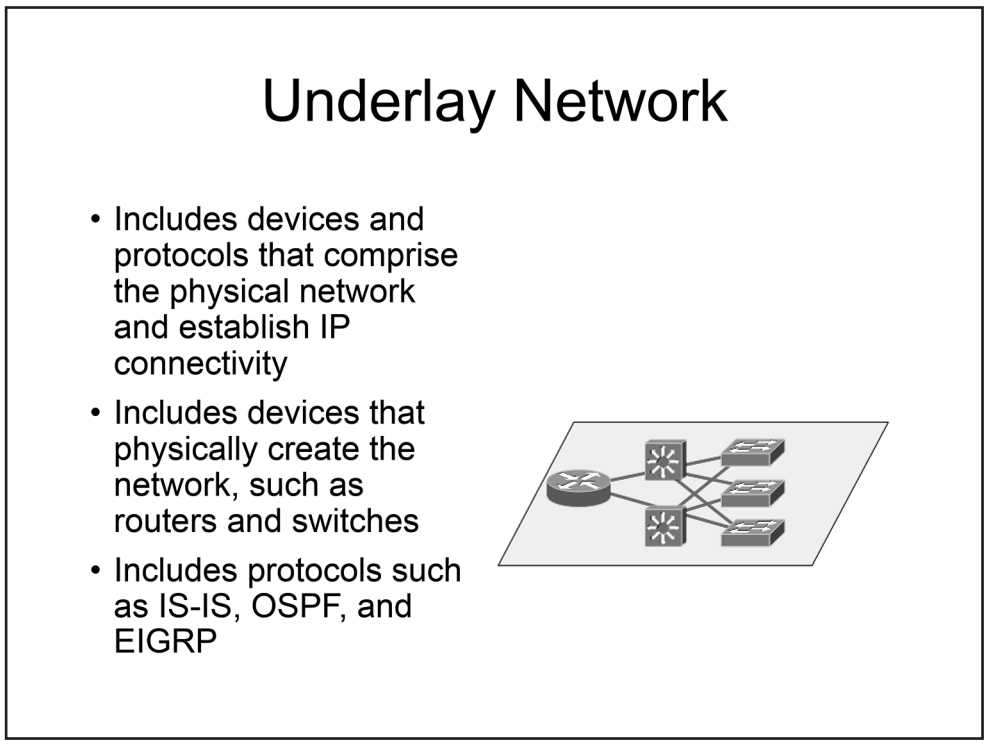

Underlay Network

The underlay network is the traditional physical composition of the network. It is a collection of devices, interfaces, and media that comprises the IP network that connects each fabric node.

Typically, routing protocols such as Intermediate System-to-Intermediate System (IS-IS), OSPF, and EIGRP are used to support SDN underlay networks. However, Cisco recommends using IS-IS instead of OSPF or EIGRP for SDN networks. This is because IS-IS performs better than OSPF or EIGRP and forms neighbor relationships without a dependence on IP protocol configuration or operation.

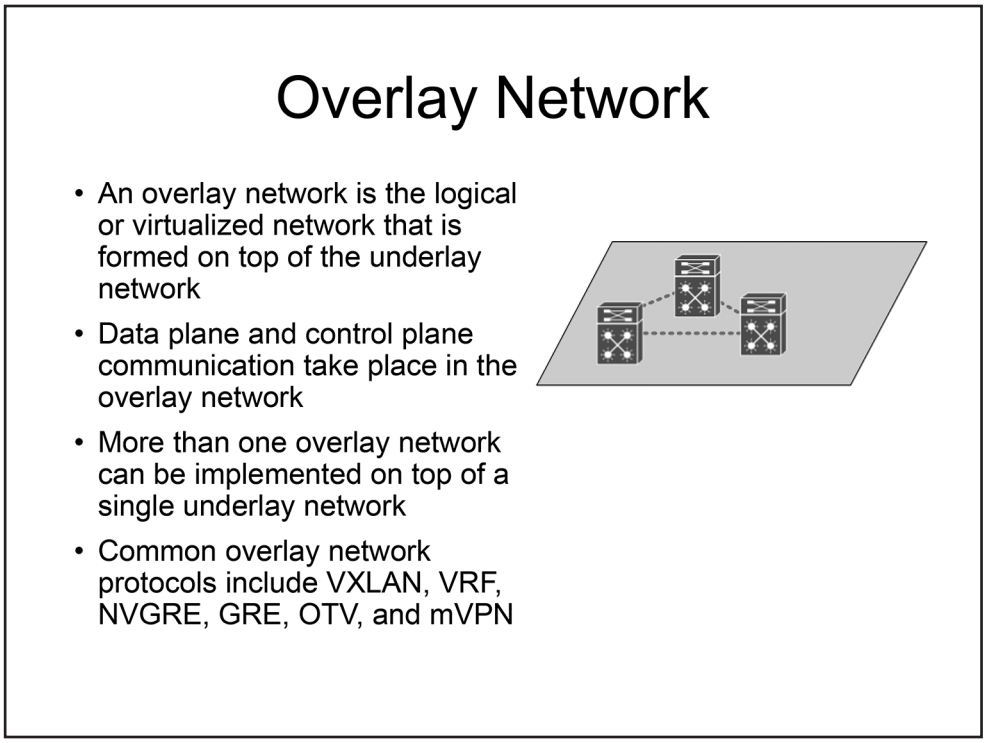

Overlay Network

An SDN overlay network is a logical or virtualized network that is formed on top of the underlay network. Both data plane and control plane communication typically take place in the overlay network. Multiple overlay networks can be constructed over a single underlay network. The overlay network creates Virtual Extensible LAN (VXLAN) tunnels between SDN switches. The tunnels send and receive traffic between fabric endpoints.

The underlay network is part of a dynamic discovery process that is involved in creating the overlay network's VXLAN tunnels. When an endpoint in a Cisco SDA network sends traffic to another endpoint, the traffic flows from the endpoint through the overlay network's VXLAN tunnels.

Other protocols involved in the construction and flow of an overlay network include Virtual Routing and Forwarding (VRF), Network Virtualization using Generic Routing Encapsulation (NVGRE), Overlay Transport Virtualization (OTV), and Multicast Virtual Private Networking (mVPN).

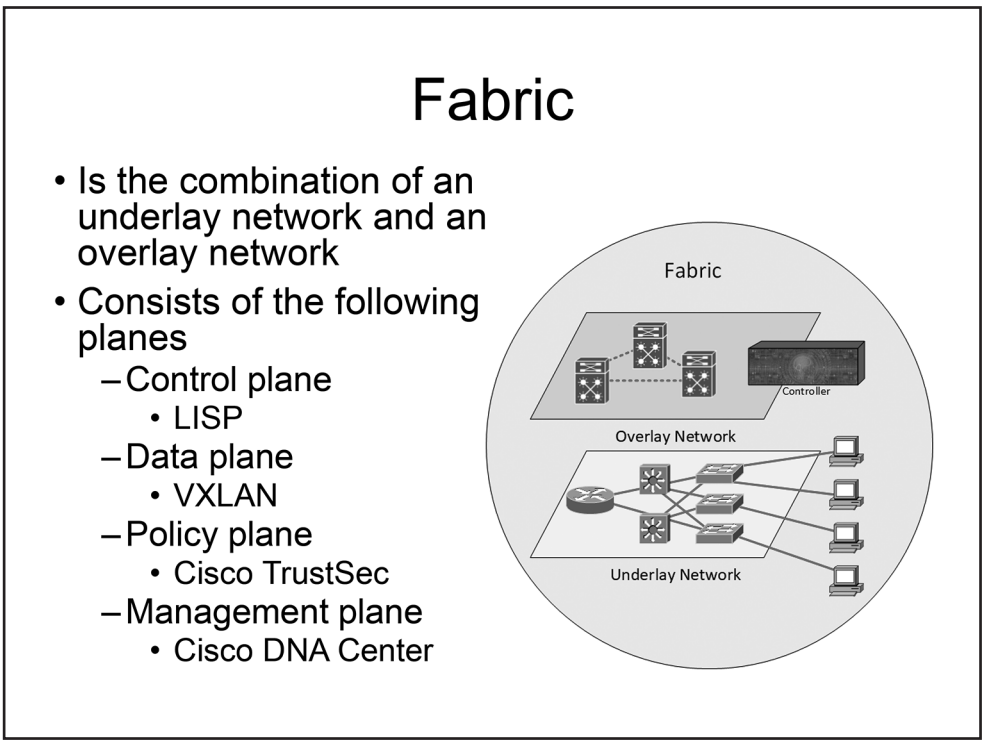

Fabric

The combination of an underlay network and an overlay network and all the components that are used to communicate between them and from them is known as the SDN fabric. In an SD-Access fabric, the fabric is managed through Cisco DNA Center.

An SD-Access fabric consists of four planes of operation:

- Control plane
- Data plane
- Policy plane
- Management plane

Each of these planes has a corresponding component in SD-Access. These components will be discussed in the following sections.

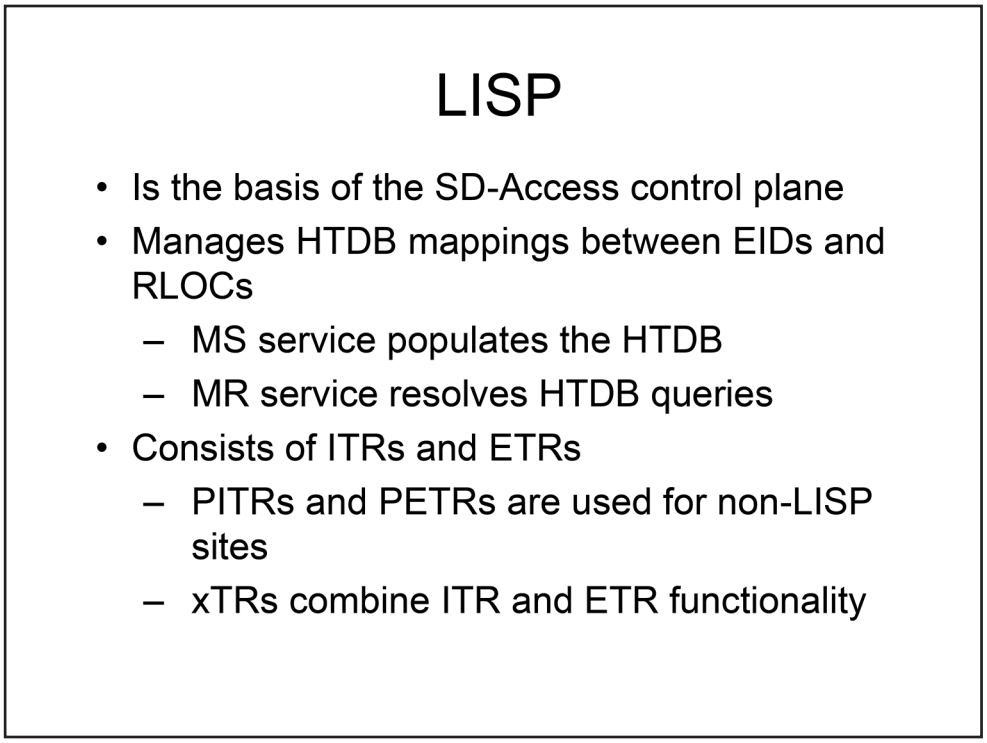

LISP

Locator/ID Separation Protocol (LISP) is the basis of the SD-Access control plane and is used to manage the mappings between endpoint identifiers (EIDs), which are assigned to hosts, and routing locators (RLOCs), which are assigned to routers. This mapping associates each endpoint with a fabric node rather than using the traditional coupling of an endpoint MAC address or an endpoint IP address with the closest available gateway. LISP manages these mappings in the host tracking database (HTDB). The HTDB is populated by the LISP Map-Server (MS) service, and queries to the HTDB are resolved by the LISP Map-Resolver (MR) service.

An ingress tunnel router (ITR) is a LISP device that LISP-encapsulates IP packets from EIDs that have destinations that lie outside the LISP site. For example, if a device with the IP address of 192.168.1.1 in LISP Site 1 wants to send a packet to a host with the IP address of 192.168.2.1 in LISP Site 2, the ITR that connects LISP Site 1 to the Internet or network core will LISP-encapsulate that packet prior to routing it over the Internet or network core. Non-LISP sites use ITR-equivalent routers known as proxy ITRs (PITRs) to send traffic to EIDs at LISP sites.

An egress tunnel router (ETR) is a LISP device that de-encapsulates LISP packets from EIDs that reside outside the LISP site but have destinations that lie inside the LISP site. For example, an ETR that connects LISP Site 2 to the Internet or network core might receive a LISP packet that is destined for a host inside LISP Site 2 that has an IP address of 192.168.2.1. The ETR will de-encapsulate the LISP packet and send the remaining IP packet to the host within LISP Site 2. EIDs in LISP sites can use ETR-equivalent routers known as proxy ETRs (PETRs) to send traffic to non-LISP sites.

The functions of an ITR and an ETR can be combined into a single device that is known as a tunnel router (xTR). Like an ITR and an ETR, an xTR connects a LISP site to the Internet or network core. It both sends

LISP packets to the Internet or network core and receives LISP packets from the Internet or network core. An xTR can encapsulate IP packets from the connected LISP site for routing over the Internet or network core. It can also de-encapsulate LISP packets that it receives from other LISP sites over the Internet or network core. Most routers that connect LISP sites to the Internet or network core are xTR routers.

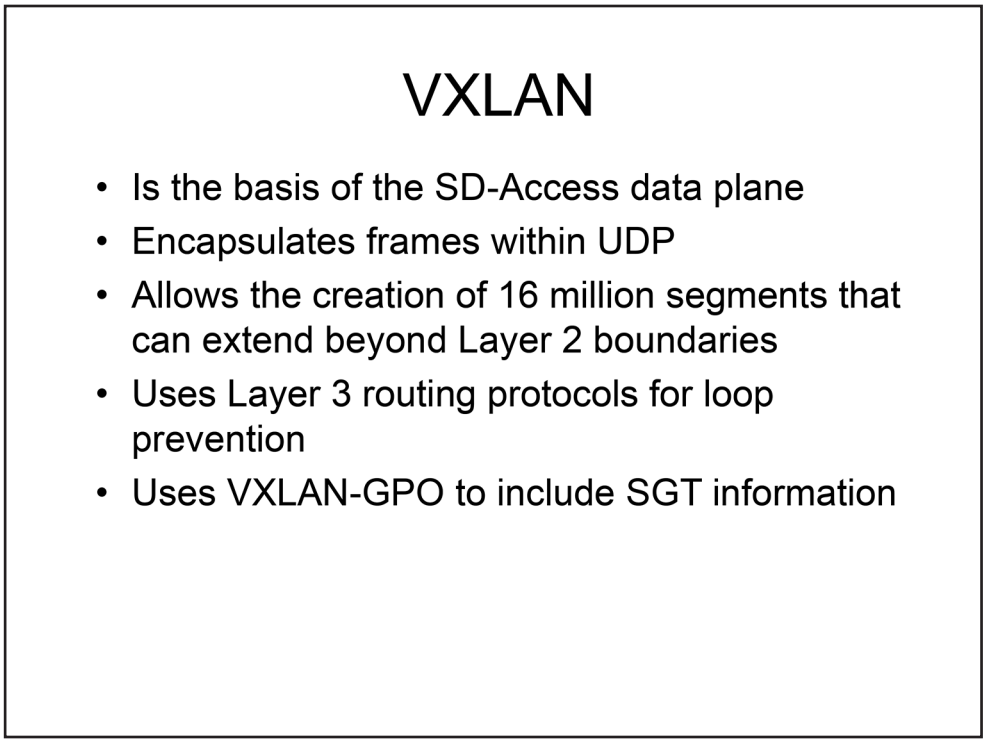

VXLAN

VXLAN is the basis of the SD-Access data plane and is used to extend and enhance VLAN capabilities, overlaying a Layer 2 network on top of a Layer 3 network. With VXLAN,
Layer 2 data plane frames are encapsulated within User Datagram Protocol (UDP), thereby enabling those packets to be transported over any IP-based underlay infrastructure.

Cisco devices are limited to 4,094 VLANs, and these VLANs are restricted to Layer 2 boundaries. VXLANs allow the creation of 16 million segments, and these segments can extend beyond the Layer 2 boundaries.

VLANs use Spanning Tree Protocol (STP) for loop prevention. By contrast, VXLAN uses Layer 3 routing protocols for loop prevention.

The Cisco SD-Access fabric uses an IETF draft standard enhancement to VXLAN, referred to as VXLAN Group Policy Option (VXLAN-GPO), that redefines a reserved portion of the standard VXLAN header to include Security Group Tag (SGT) information. SGT information is used to identify user group security memberships and is used by Cisco TrustSec.

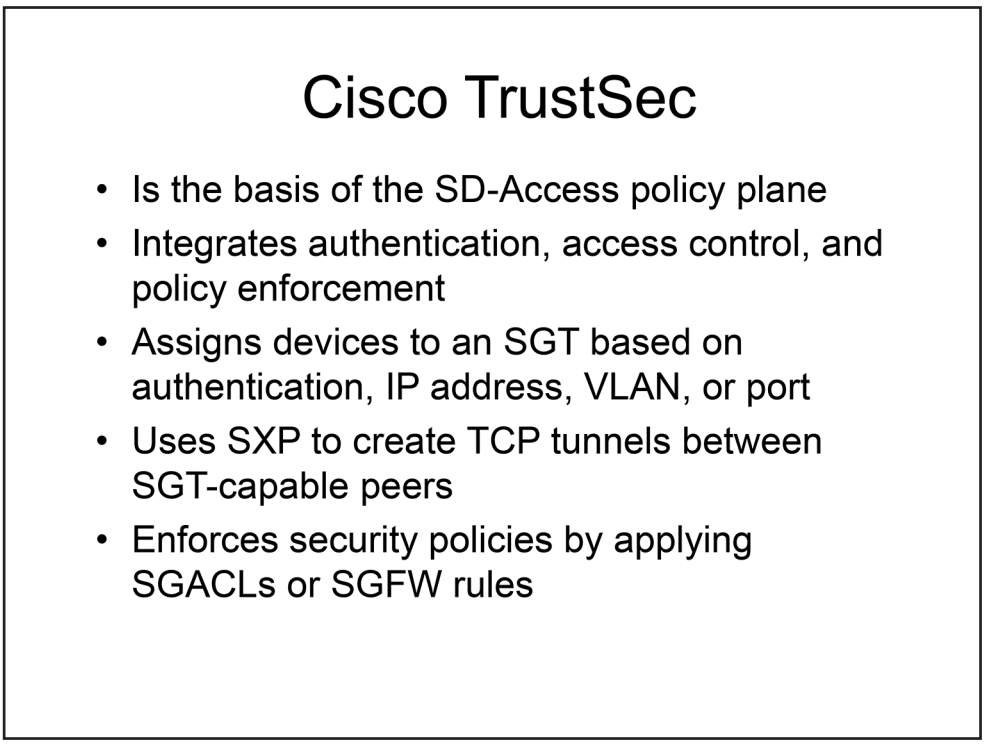

Cisco TrustSec

Cisco TrustSec is the basis of the SD-Access policy plane and integrates authentication, access control, and policy enforcement. With TrustSec, when a device connects to the network, it is assigned to an associated security group, which is represented by an SGT. In TrustSec terms, this assignment is referred to as classification and can occur either dynamically as part of the authentication process or statically by associating the device with a supported identifier, such as an IP address, VLAN, or port. SGT-capable devices can embed SGT information directly into Ethernet frames through a process called inline-tagging. However, if there are devices in the data path that do not support inline-tagging, SGT Exchange Protocol (SXP) must be used to create Transmission Control Protocol (TCP) tunnels between SGT-capable peers to ensure that tags are preserved during transit. TrustSec devices use the SGT information to enforce security policies by applying Security Group Access Control Lists (SGACLs) or Security Group Firewall (SGFW) rules.

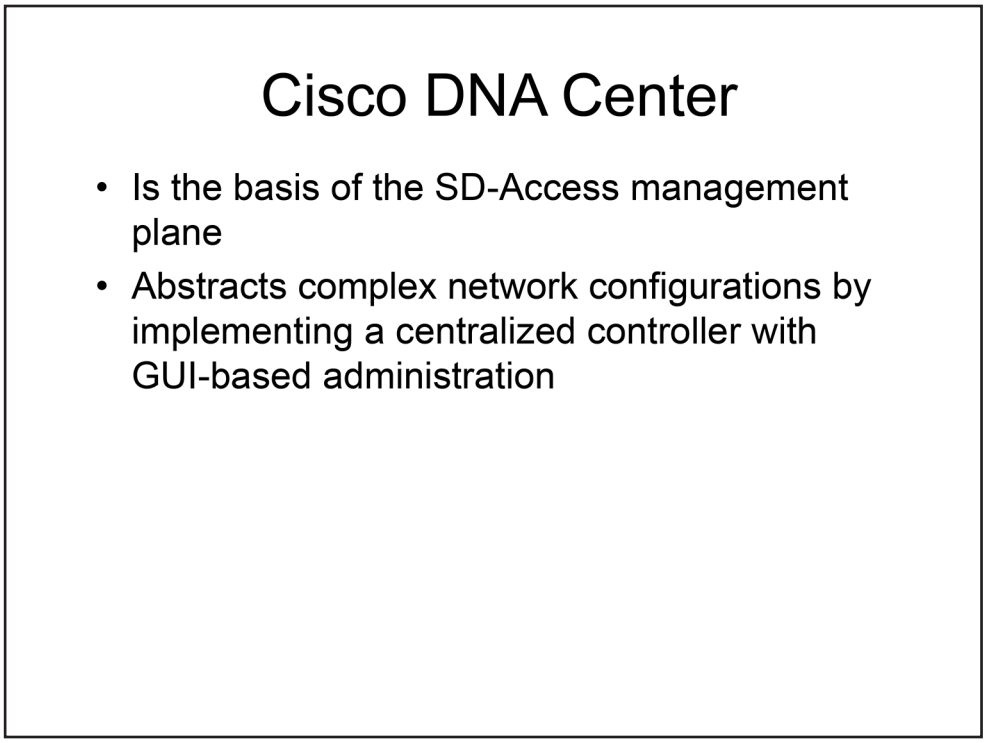

Cisco DNA Center

Cisco DNA Center is the basis of the SD-Access management plane. As previously discussed, Cisco DNA Center abstracts the complexity of network configuration by implementing a centralized controller with GUI-based administration.

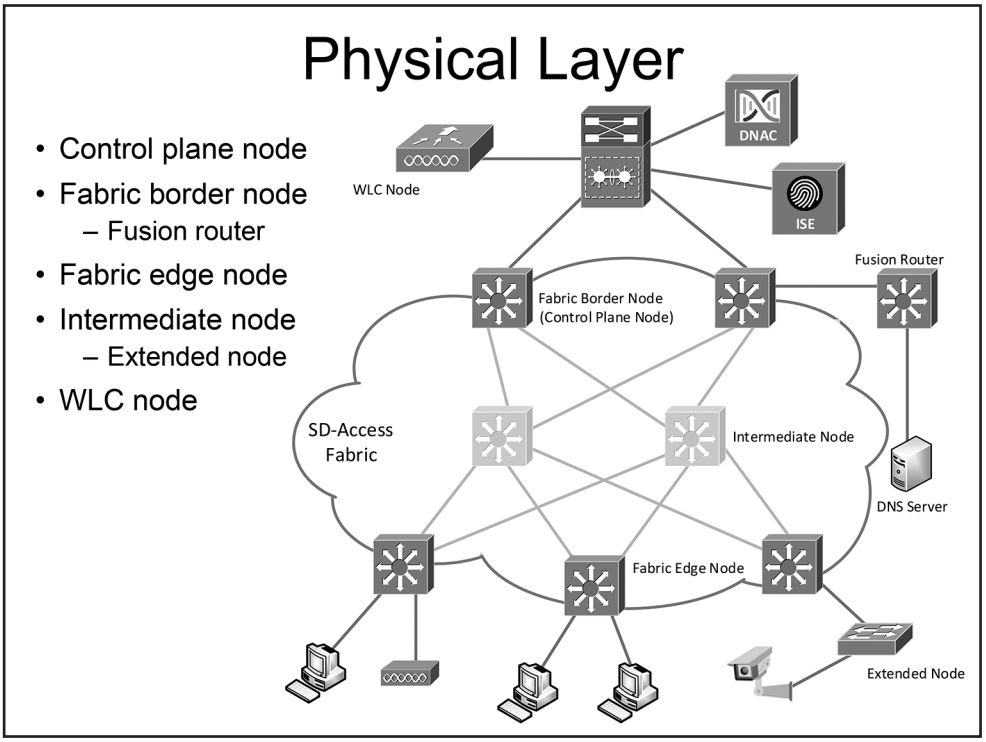

Physical Layer

The physical layer consists of Cisco routers, switches, wireless devices, and controllers. Each device within the fabric overlay is associated with a device role:

- Control plane node
- Fabric border node
- Fabric edge node
- Intermediate node
- WLC node

A control plane node manages the HTDB, which is used to map EIDs to RLOCs. This mapping associates each endpoint with a fabric node rather than using the traditional coupling of an endpoint MAC address or an endpoint IP address with the closest available gateway. Every edge node and border node must register with any available control nodes and will use the control plane nodes to register and resolve endpoints. Control plane node functionality can be provided by one or more dedicated devices or by one or more border nodes configured with control plane node functionality.

Fabric border nodes connect the SD-Access fabric to external Layer 3 networks. There are three general categories of fabric border node: internal border, default border, and internal+default border. Internal border nodes connect to networks that are external to the fabric but within the administrative domain of an organization. By contrast, a default border node connects to a network that is outside the administrative domain of the organization, such as the Internet. Internal+default border nodes combine the capabilities of internal border nodes and default border nodes. Fabric border nodes use BGP to advertise EID prefixes to

external networks. Fusion routers connect to fabric border nodes and perform route leaking between VRFs and shared services, such as DNS and DHCP.

Fabric edge nodes connect endpoint devices, such as computers and APs, to the SD-Access fabric and function as anycast Layer 3 gateways. All fabric edge nodes maintain LISP control plane connections to available control plane nodes to register their local endpoints with the HTDB. In addition, each edge node maintains a local host-tracking database referred to as the EID table. Because edge nodes are functionally similar to Layer 3 access switches in a traditional network design, they are ultimately responsible for mapping endpoints to their appropriate virtual networks (VNs). Fabric edge nodes query the LISP MS to determine the RLOC that is associated with each destination EID; after the RLOC that is associated with an EID is resolved, the edge node can use that information to encapsulate traffic destined to that EID with the appropriate VXLAN header information.

An intermediate node is a Layer 3 transport device within the underlay network; it is capable of supporting VXLAN traffic but does not otherwise participate in fabric operations. Intermediate nodes need not be aware of the LISP control plane or of SGT information; the primary requirement for an intermediate node is support for the increased maximum transmission unit (MTU) size in order to accommodate VXLAN encapsulated packets. Most intermediate nodes reside between border nodes and edge nodes in the underlay network. However, an extended node is a special type of intermediate node that resides outside the SD-Access fabric and connects to an edge node with an 802.1Q trunk to extend fabric capabilities to Layer 2 devices such as industrial endpoints and Internet of Things (IoT) devices.

A wireless LAN controller (WLC) node resides outside the SD-Access fabric and provides wireless connectivity to endpoints connected to lightweight wireless access points (LWAPs) that are attached to edge nodes. In a traditional wireless network, control plane and data plane traffic from wireless endpoints passes through Control and Provisioning of Wireless Access Points (CAPWAP) tunnels from the LWAPs to the WLC. However, with a fabric WLC node and the associated fabric LWAPs, control plane and data plane traffic is encapsulated in VXLAN tunnels. The control plane traffic continues to flow to the WLC; however, the data plane traffic can be forwarded directly through the fabric to the appropriate edge node without needing to pass through the WLC. Although a fabric WLC and its associated fabric APs reside outside the fabric, they are considered a part of the fabric overlay network because their reachability is not dependent on the details of the underlay network implementation.

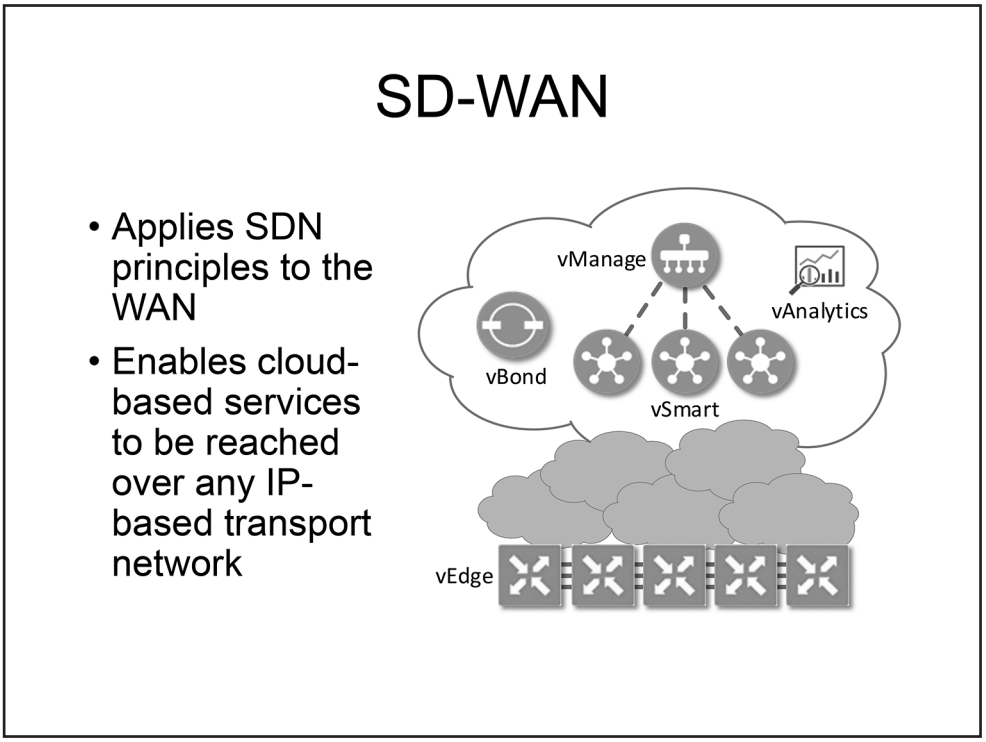

SD-WAN

Many organizations are taking greater advantage of the benefits of migrating resources and applications to the cloud. However, this escalating demand for cloud-based connectivity comes with an increase in WAN complexity.

SD-WAN simplifies and centralizes WAN automation, operations, provisioning, and monitoring. By applying SDN principles to the WAN, SD-WAN separates the control plane from the data plane, thereby enabling cloud-based services to be reached over any IP-based transport network.

There are two types of SD-WAN solutions currently offered by Cisco:

- Cisco SD-WAN
- Meraki SD-WAN

Cisco SD-WAN, which is based on Viptela technology, is generally recommended for businesses that require an SD-WAN solution. However, businesses that require unified threat management (UTM) solutions with SD-WAN functionality or that already have Cisco Meraki equipment should consider the Meraki SD-WAN solution.

The Cisco SD-WAN solution includes a feature called Cloud OnRamp that extends the SD-WAN fabric to the cloud, providing optimized Software as a Service (SaaS) application access and Infrastructure as a Service (IaaS) connectivity. Cloud OnRamp monitors fabric performance statistics and selects the best path for cloud-based application traffic based on these statistics.

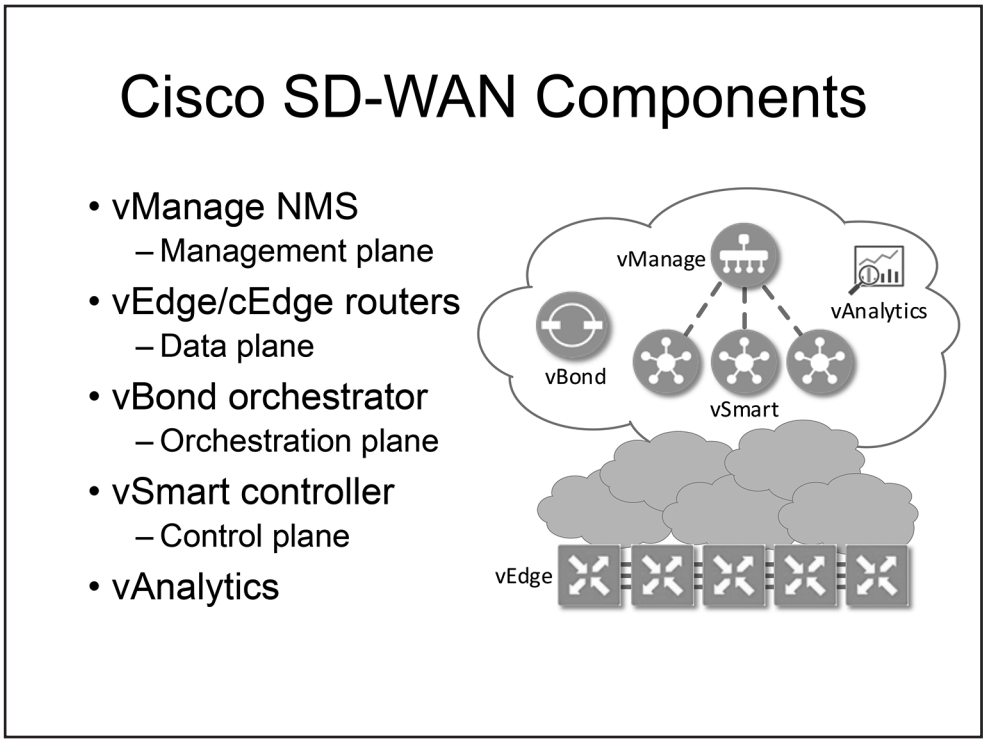

Cisco SD-WAN Components

The Cisco SD-WAN solution consists of the following primary components:

- vManage Network Management System (NMS)
- vEdge and cEdge routers
- vBond orchestrator
- vSmart controller

These primary components correspond to the four planes of the SD-WAN overlay network architecture:

- Management plane
- Data plane
- Orchestration plane
- Control plane

An optional component, vAnalytics, provides detailed information about the WAN infrastructure and WAN-based applications. If an unusual network disruption is detected by vAnalytics, it will alert administrators of the problem. vAnalytics can also be used to predict bandwidth requirements for each location.

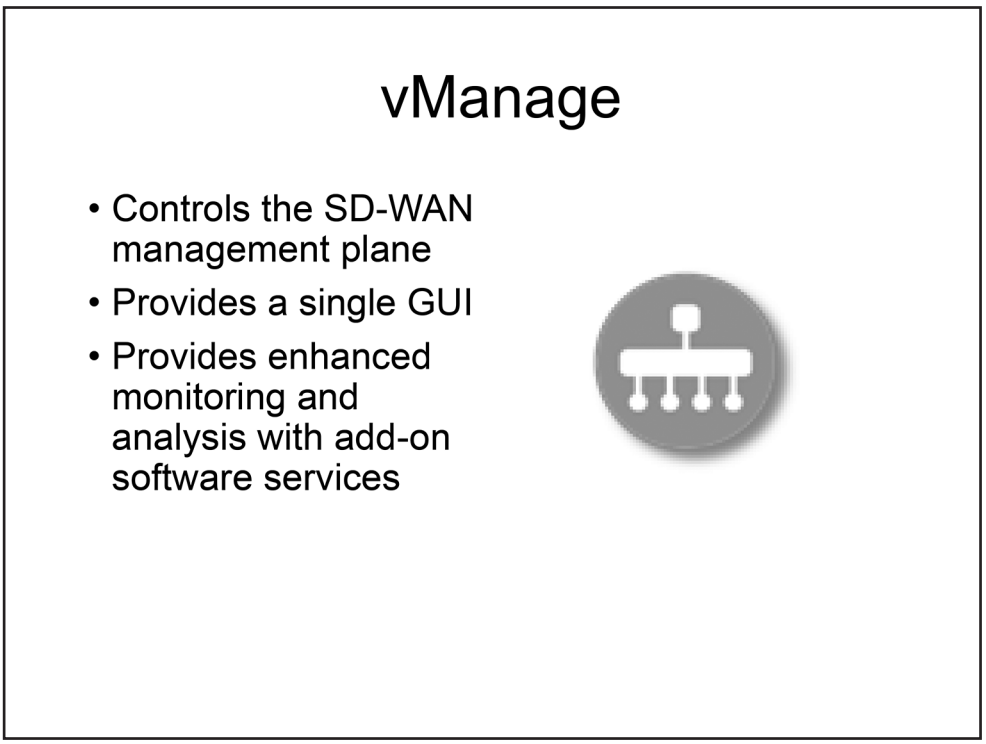

vManage

The vManage NMS controls the management plane of the SD-WAN overlay network architecture. vManage is a software solution that provides a single GUI to configure and manage the components of the SD-WAN architecture. Additional software services can be added to vManage NMS to provide enhanced monitoring and analysis capabilities.

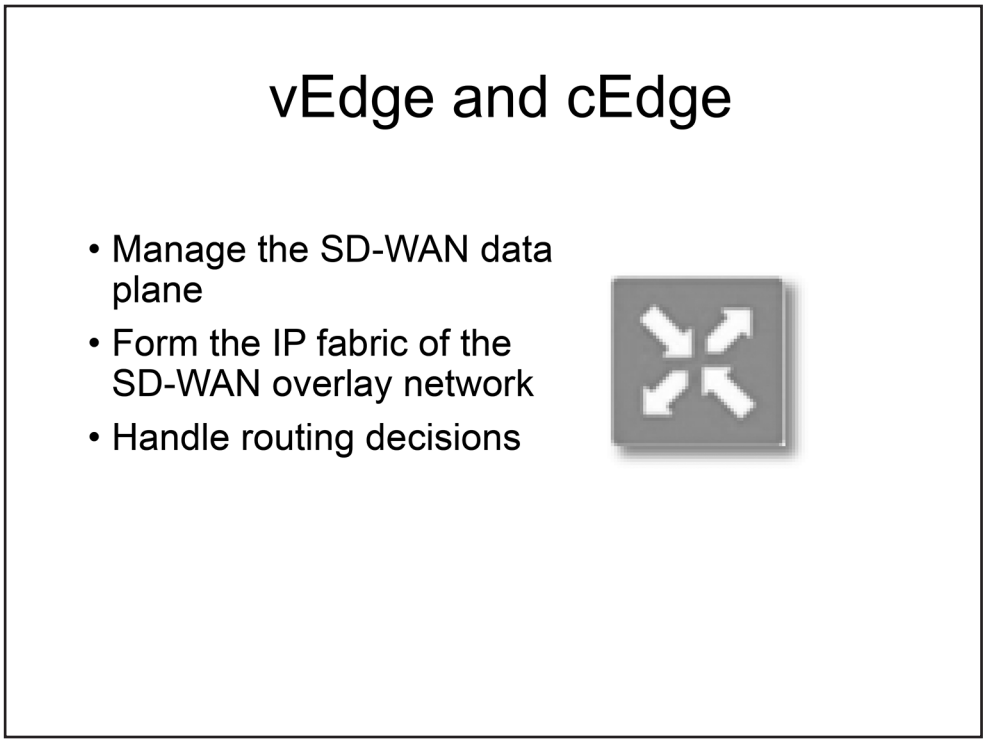

vEdge and cEdge

The vEdge and cEdge routers manage the data plane of the SD-WAN overlay network. vEdge routers consist of the original Viptela platforms running Viptela software; cEdge routers integrate Viptela software with Cisco IOS-XE. The IP fabric of the SD-WAN overlay network consists of Edge routers interconnected by IP Security (IPSec) tunnels. Routing information for the IP fabric is reflected by using Overlay Management Protocol (OMP) from Edge routers to vSmart routers, which will reflect those updates to all edge routers. Routes reflected from the vSmart controller are redistributed into the routing table at each associated Edge router so that all routing decisions can be handled locally by the Edge routers.

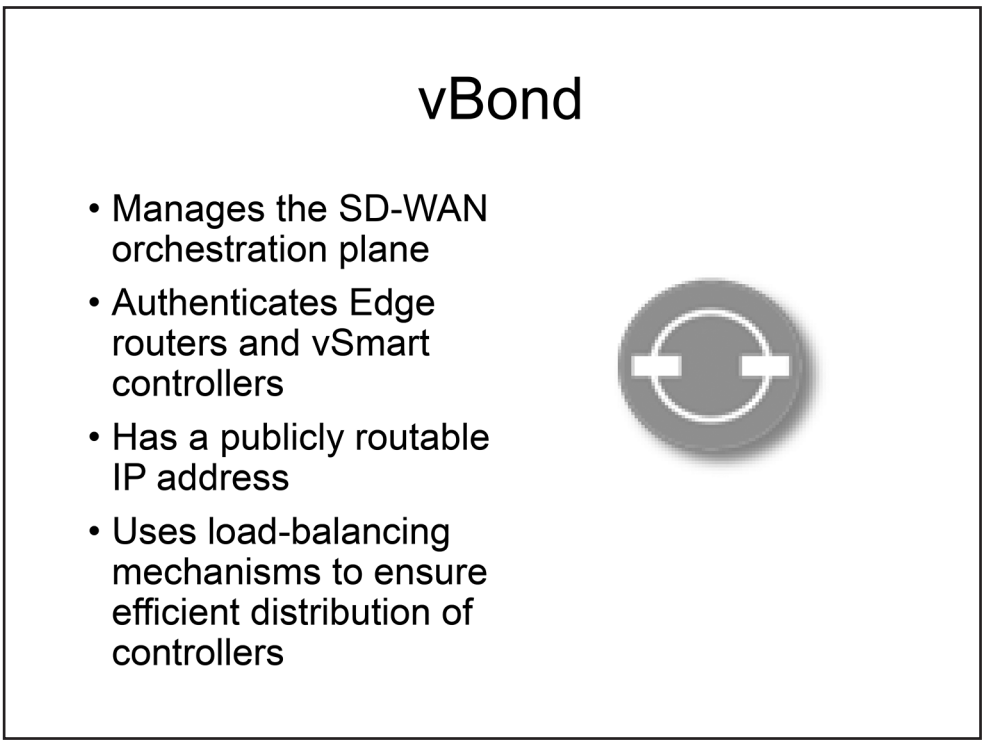

vBond

The vBond orchestrator manages the orchestration plane of the SD-WAN overlay network. vBond authenticates Edge routers and vSmart controllers over a Datagram Transport Layer Security (DTLS) tunnel connection as they perform their initial startup sequences. The vBond orchestrator is the only SD-WAN component that is required to have a publicly routable IP address. Because the vBond orchestrator has a public IP address, it can be accessed by other SD-WAN components even if they reside behind Network Address Translation (NAT) devices such as firewalls or routers. This accessibility enables the vBond orchestrator to facilitate connections among all other SD-WAN components. In addition, the vBond orchestrator uses load balancing mechanisms to ensure that Edge routers are efficiently distributed between available vSmart controllers when the Edge routers are initially configured.

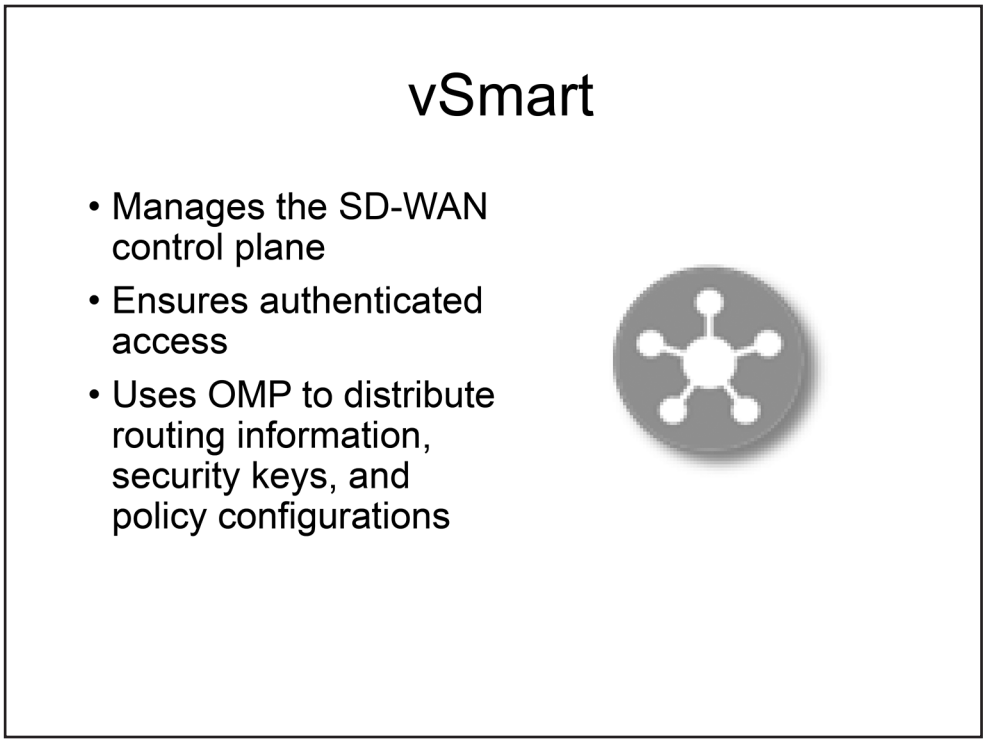

vSmart

The vSmart controller manages the control plane of the SD-WAN overlay network architecture and ensures that only authenticated devices can access the SD-WAN. The vBond orchestrator and Edge routers must maintain DTLS connections to at least one vSmart controller. Control plane traffic passes through DTLS tunnels between the vSmart controller and the other SD-WAN components. For example, the vSmart controller uses OMP to distribute routing information, security keys, and policy configurations through DTLS tunnels to Edge routers. The Edge routers can then use this information to determine the appropriate next hop for data plane traffic, to create IPSec tunnels to other Edge routers for data plane traffic, and to ensure that Service Level Agreements (SLAs) are met and that traffic policies are enforced.

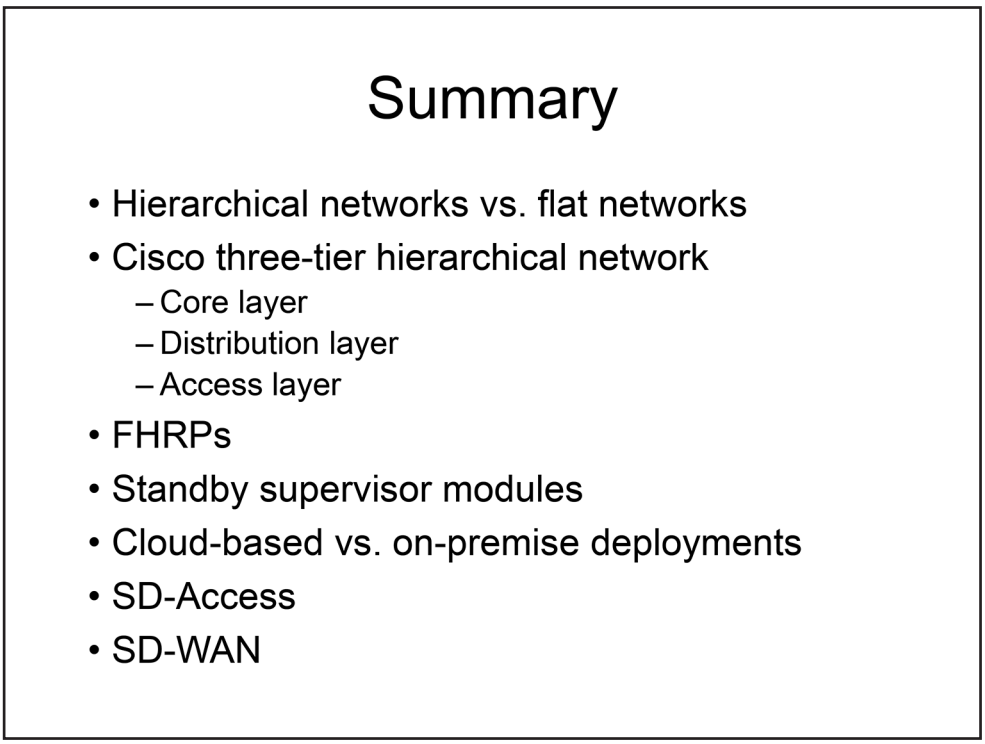

Summary

Hierarchical network designs are more scalable and manageable than flat network designs. Cisco's hierarchical network design includes the Core layer, the Distribution layer, and the Access layer. FHRPs can be used to create redundancy between Layer 3 devices. Standby supervisor modules are available on some Cisco models to increase device availability and route stability.

Cloud-based deployments are becoming more common and are replacing traditional on-premise deployments. SD-Access and SD-WAN can be implemented to handle the additional complexity of larger, geographically disparate networks.

Review Question 1

Which of the following layers are combined in the Cisco two-tier Enterprise Campus Architecture model?

A. access and core

B. core and distribution

C. distribution and access

D. core, distribution, and access

Review Question 1

Which of the following layers are combined in the Cisco two-tier Enterprise Campus Architecture model?

A. access and core

B. core and distribution

C. distribution and access

D. core, distribution, and access

The Core and Distribution layers are combined in the Cisco two-tier Enterprise Campus Architecture Model. The Cisco two-tier Enterprise Campus Architecture model, which is sometimes called the collapsed-core network design model, is recommended only for smaller networks. In a two-tier network design model, the functionality of the Core layer is collapsed into the Distribution layer. The functionality of the Core layer is provided by the Distribution layer, and a distinct Core layer does not exist.

The Access layer is not combined with either the Core layer or the Distribution layer in the Cisco two-tier Enterprise Campus Architecture Model. Access layer devices are typically high-density switches that are focused on providing direct network access to end-user devices. Combining the Access layer with another layer would reduce the functionality and effectiveness of each layer.

Review Question 2

Which of the following statements is true regarding VRRP?

A. VRRP is Cisco-proprietary.

B. VRRP provides load balancing between gateways.

C. VRRP groups contain one active router and one standby router.

D. VRRP gateways share a virtual IP address and a virtual MAC address.

Review Question 2

Which of the following statements is true regarding VRRP?

A. VRRP is Cisco-proprietary.

B. VRRP provides load balancing between gateways.

C. VRRP groups contain one active router and one standby router.

D. VRRP gateways share a virtual IP address and a virtual MAC address.

Virtual Router Redundancy Protocol (VRRP) gateways share a virtual IP address and a virtual Media Access Control (MAC) address. VRRP is an open-standard First-Hop Redundancy Protocol (FHRP) that provides much of the functionality offered by Cisco's proprietary Hot Standby Router Protocol (HSRP). Each VRRP group contains a single master virtual router and one or more backup virtual routers to provide first-hop redundancy for client devices. Only the master virtual router responds to Address Resolution Protocol (ARP) requests and forwards traffic. If the master virtual router fails, one of the backup virtual routers will assume the role of master virtual router and will begin forwarding traffic.

HSRP enables multiple Layer 3 devices to act as a single gateway for the network by sharing a single virtual IP address and a single virtual MAC address. Each HSRP group contains a single active router and a single standby router. Only the active router responds to ARP requests and forwards traffic. If the active router fails, the standby router assumes the role of the active router and a new standby router is elected among the remaining routers in the HSRP group.

Gateway Load Balancing Protocol (GLBP) is a Cisco-proprietary protocol that provides router redundancy and load balancing. The routers in a GLBP group receive traffic sent to a virtual IP address that is configured for the group. However, unlike HSRP and VRRP, GLBP is capable of load balancing traffic among every router in the group.

Each GLBP group contains an active virtual gateway (AVG) and up to four primary active virtual forwarders (AVFs); the remaining routers become secondary virtual forwarders (SVFs). The AVG assigns a virtual MAC address to the primary AVFs. When the AVG receives ARP requests that are sent to the virtual IP address for the GLBP group, the AVG responds with a different virtual MAC address, including its own. This provides load balancing, because each of the primary AVFs can participate by forwarding a portion of the traffic sent to the virtual IP address. If the AVG fails, the AVF with the next highest priority, which is referred to as the standby virtual gateway (SVG), will take over for the AVG. If one of the AVFs fails, the AVG will assign another router the role of that AVF.

Review Question 3

Which of the following technologies is the basis of the Cisco SD-Access control plane?

A. LISP

B. VXLAN

C. Cisco TrustSec

D. Cisco DNA Center

Review Question 3

Which of the following technologies is the basis of the Cisco SD-Access control plane?

A. LISP

B. VXLAN

C. Cisco TrustSec

D. Cisco DNA Center

Locator/ID Separation Protocol (LISP) is the basis of the Cisco Software-Defined Access (SD-Access) control plane. An SD-Access fabric consists of four planes of operation:

- Control plane
- Data plane
- Policy plane
- Management plane

LISP is used to manage the mappings between endpoint identifiers (EIDs) and routing locators (RLOCs). This mapping associates each endpoint with a fabric node rather than using the traditional coupling of an endpoint Media Access Control (MAC) address or an endpoint IP address with the closest available gateway. LISP manages these mappings in the host tracking database (HTDB). The HTDB is populated by the LISP Map-Server (MS) service, and queries to the HTDB are resolved by the LISP Map-Resolver (MR) service.

Virtual Extensible LAN (VXLAN) is the basis of the SD-Access data plane, not the SD-Access control plane. VXLAN encapsulates entire Layer 2 data plane frames for User Datagram Protocol (UDP) transport through the overlay network. The Cisco SD-Access fabric uses an Internet Engineering Task Force (IETF) draft standard enhancement to VXLAN, referred to as VXLAN Group Policy Option (VXLAN-GPO), that redefines a reserved portion of the standard VXLAN header to include Security Group Tag (SGT) information. Because VXLAN encapsulates an entire Layer 2 frame into a UDP datagram, it can be used to create Layer 2 and Layer 3 overlay networks over any IP-based underlay infrastructure.

Cisco TrustSec is the basis of the SD-Access policy plane, not the SD-Access control plane. TrustSec integrates authentication, access control, and policy enforcement. With TrustSec, when a device connects to the network, it is assigned to an associated security group, which is represented by an SGT. In TrustSec terms, this assignment is referred to as classification and can occur either dynamically as part of the

authentication process or statically by associating the device with a supported identifier, such as an IP address, virtual LAN (VLAN), or port. SGT-capable devices can embed SGT information directly into Ethernet frames through a process called inline-tagging. However, if there are devices in the data path that do not support inline-tagging, SGT Exchange Protocol (SXP) must be used to create Transmission Control Protocol (TCP) tunnels between SGT-capable peers to ensure that tags are preserved during transit. TrustSec devices use the SGT information to enforce security policies by applying Security Group Access Control Lists (SGACLs) or Security Group Firewall (SGFW) rules.

Cisco DNA Center is the basis of the SD-Access management plane, not the SD-Access control plane. Unlike traditional Cisco IOS-based implementations, the Cisco DNA Center tools are provided in a GUI. These management tools help to abstract and simplify the complex network interactions that exist on the other layers.

Certification Candidates

Boson Software's ExSim-Max practice exams are designed to simulate the complete exam experience. These practice exams have been written by in-house authors who have over 30 years combined experience writing practice exams. ExSim-Max is designed to simulate the live exam, including topics covered, question types, question difficulty, and time allowed, so you know what to expect. To learn more about ExSim-Max practice exams, please visit www.boson.com/exsim-max-practice-exams or contact Boson Software.

Organizational and Volume Customers

Boson Software's outstanding IT training tools serve the skill development needs of organizations such as colleges, technical training educators, corporations, and governmental agencies. If your organization would like to inquire about volume opportunities and discounts, please contact Boson Software at orgsales@boson.com.

Contact Information

E-Mail: support@boson.com
Phone: 877-333-EXAM (3926)
615-889-0121
Fax: 615-889-0122
Address: 25 Century Blvd., Ste. 500
Nashville, TN 37214

Boson.com

877.333.3926 support@boson.com