Boson®

SWITCH Curriculum

300-115

Labs powered by

NetSim®

NETWORK SIMULATOR®

SWITCH

300-115 Curriculum

LM20170901/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.

Content in these modules is available in the full version of the curriculum. Please visit www.boson.com for more information.

Module 1

Network Design Basics

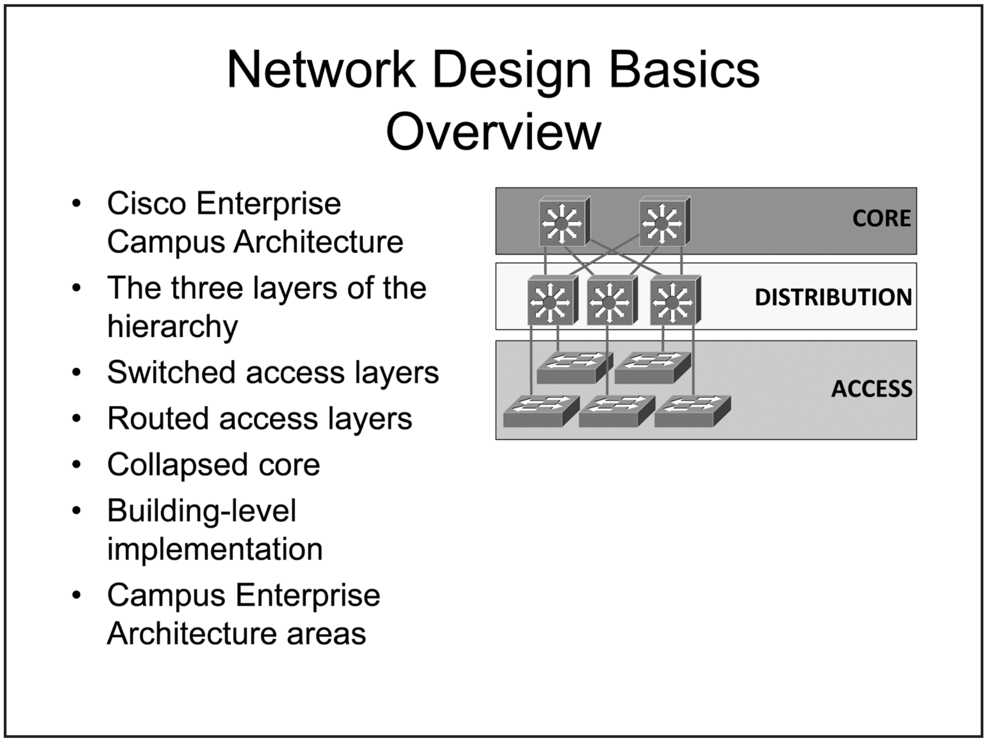

Network Design Basics Overview

- Cisco Enterprise Campus Architecture
- The three layers of the hierarchy
- Switched access layers
- Routed access layers
- Collapsed core
- Building-level implementation
- Campus Enterprise Architecture areas

Overview

Current best practice eschews flat network designs and leverages the resilience and simplicity of hierarchical design patterns. A hierarchical design pattern, such as Cisco's Enterprise Campus Architecture, relies on modular, structured components to deliver efficient, flexible, and manageable network solutions. This module explores the Cisco Enterprise Campus Architecture and the advantages of hierarchical design.

Objectives

After completing this module, you should have the basic knowledge required to understand the following concepts:

- Cisco Enterprise Campus Architecture
- Layers of the hierarchical network
- Switched versus routed access layer
- Traditional versus collapsed core layer
- Building-level implementation
- Cisco Enterprise Campus Architecture areas

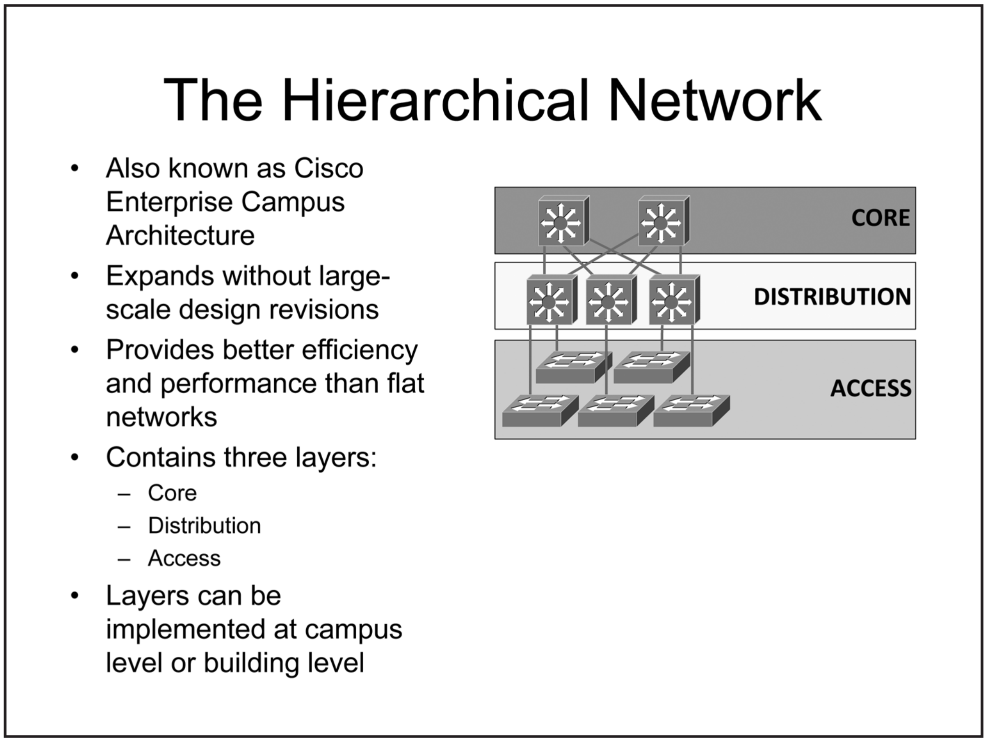

The Hierarchical Network

- Also known as Cisco Enterprise Campus Architecture
- Expands without large-scale design revisions
- Provides better efficiency and performance than flat networks
- Contains three layers:
  - Core
  - Distribution
  - Access
- Layers can be implemented at campus level or building level

The Hierarchical Network

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 that might have provided a necessary level of network segmentation, such as bridges, switches, and routers. Without adequate network segmentation, the scalability of a flat topology is limited. The bandwidth and CPU resources 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 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 structure based on their functionality or role within the topology. 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. 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, modularity and a layered, hierarchical structure reduce management complexity and thus facilitate the implementation of redundancy and optimized data paths.

The Cisco Enterprise Campus Architecture model contains three layers:

- Access
- Distribution
- Core

Each of these layers is a conceptual entity that corresponds to a group of network devices with shared functionality. The network devices residing in these layers may be deployed within a single building or spread out across a campus. In some implementations, two of the layers may be combined within a single group of physical devices. For example, the distribution and core layers may be combined for small, building-level implementations that do not require a distinct core layer.

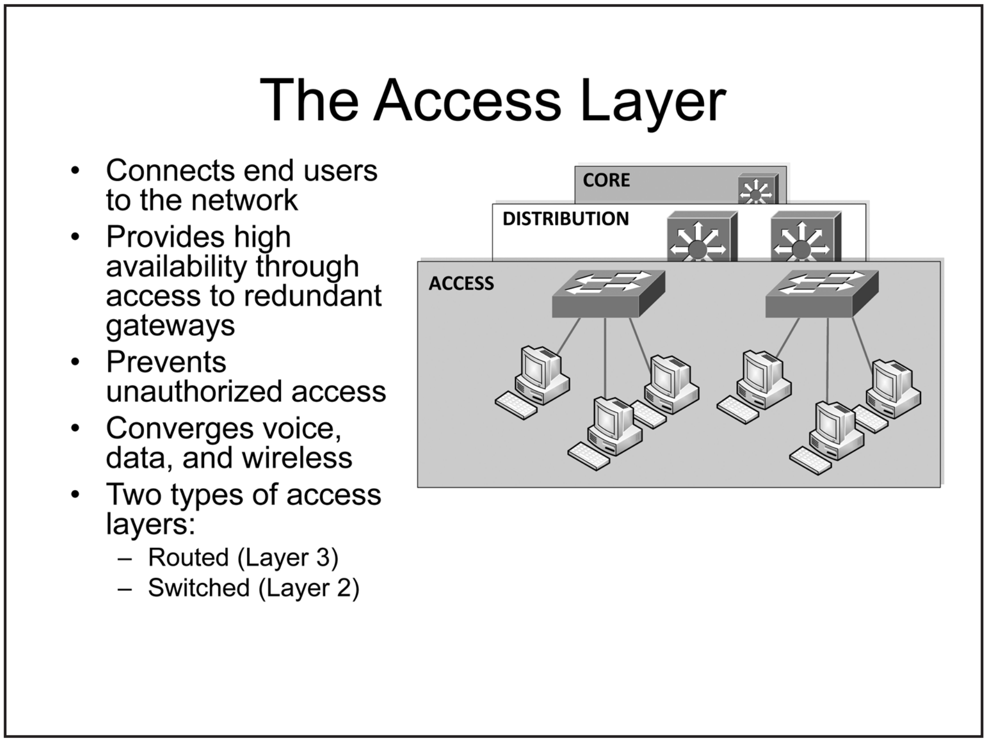

The Access Layer

- Connects end users to the network
- Provides high availability through access to redundant gateways
- Prevents unauthorized access
- Converges voice, data, and wireless
- Two types of access layers:
  - Routed (Layer 3)
  - Switched (Layer 2)

The 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. In the Cisco Enterprise Campus Architecture design model, the access layer provides services such as 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 sufficient amounts of 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.

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.

In the Cisco Enterprise Campus Architecture design model, the access layer can be configured as a switched access layer or as a routed access layer. In a switched access layer, connectivity to the distribution layer is over Layer 2 uplinks, such as Institute of Electrical and Electronics Engineers (IEEE) 802.1Q trunks; the access layer switches do not require IP routing capability. By contrast, in a routed access layer, connectivity to the distribution layer is over Layer 3 uplinks, such as PortChannel virtual interfaces.

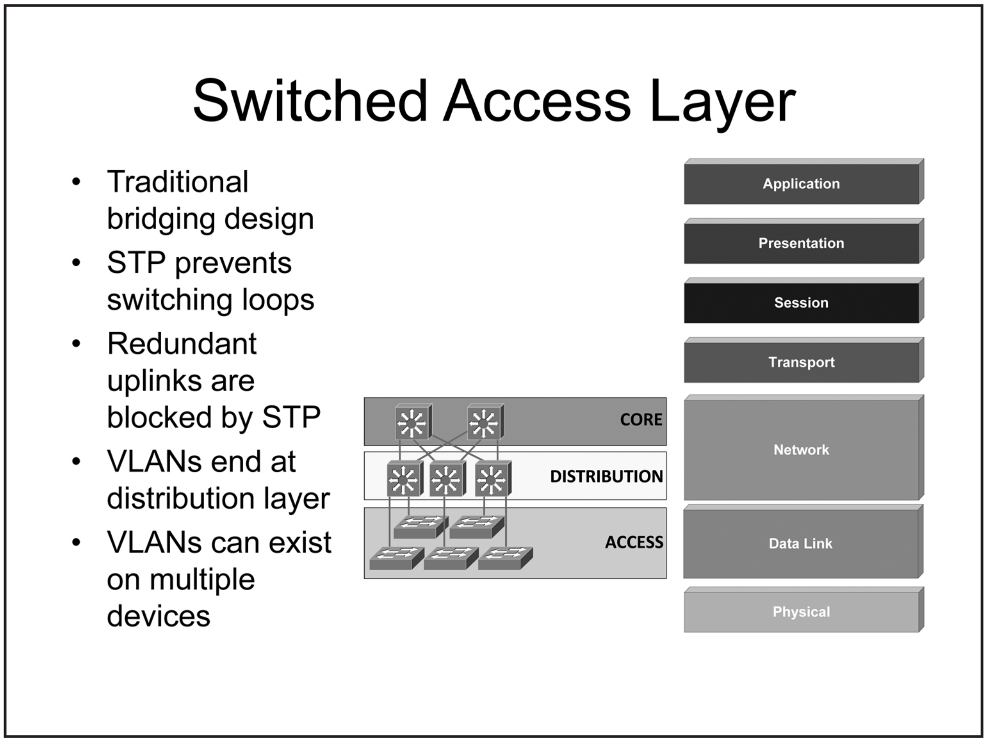

Switched Access Layer

- Traditional bridging design
- STP prevents switching loops
- Redundant uplinks are blocked by STP
- VLANs end at distribution layer
- VLANs can exist on multiple devices

Switched Access Layer

A switched access layer design uses traditional Layer 2 uplinks to connect access layer switches to the distribution layer. Because the entirety of the switched access layer operates at the Data Link layer of the Open Systems Interconnection (OSI) model, switching loops can occur if multiple paths exist between devices. Therefore, Spanning Tree Protocol (STP) is used by access layer switches to build a common tree view of the network and to mitigate switching loops. STP blocks links within the tree that could potentially cause switching loops, such as redundant uplinks.

In a switched access layer design, VLANs are extended into the distribution layer by tagging VLAN traffic on uplink trunk interfaces. Because VLAN information terminates in the distribution layer, VLANs can span multiple access layer devices, even if they are not directly interconnected. Distribution layer switches function as gateways for access layer VLANs and can route traffic between them; the access layer switches do participate in routing VLAN traffic.

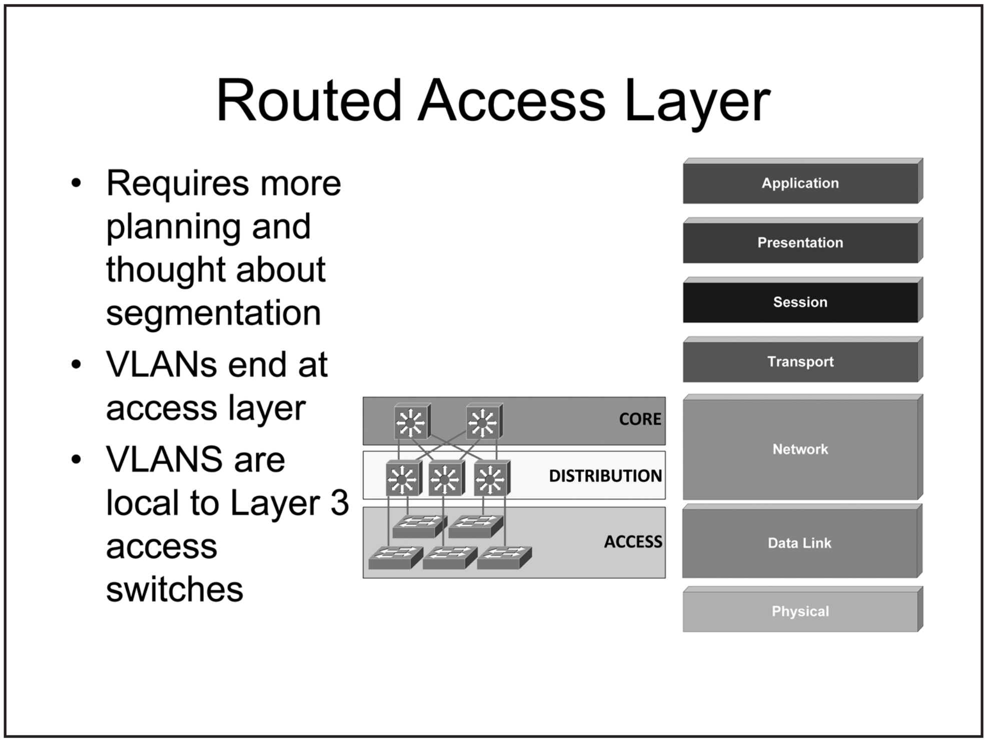

Routed Access Layer

- Requires more planning and thought about segmentation
- VLANs end at access layer
- VLANS are local to Layer 3 access switches

Routed Access Layer

In a routed access layer design, the uplinks between the access layer and distribution layer switches are Layer 3 connections. Because the Layer 2 topology in this design is reduced to a potential trunk link between access layer switches, Layer 2 loops are eliminated and all uplinks to the distribution layer are in a forwarding state. Although STP is no longer necessary in this design, Cisco recommends configuring STP on ports that connect to end-user devices, such as servers and workstations, to prevent user-side loops from entering the network.

Routed access layer designs require more planning than switches access layer designs because uplinks must be configured with Network layer addresses. Additional thought must be put into the appropriate network addressing, summarization, and general segmentation to ensure that scalability and flexibility are maintained. Although they are not required to run a dynamic routing protocol, Layer 3 access switches participate in routing traffic throughout the network, including the routing of traffic between VLANs local to the switch, configured gateways, and statically configured routes.

Because access layer switches are configured with routed interfaces in this design model, configured VLANs remain local to each switch and do not extend into the distribution layer. In addition, VLANs can only be extended between access layer switches if they share a Layer 2 link over which to communicate VLAN data. However, Cisco recommends limiting VLANs to individual access layer switches in the routed access layer design model.

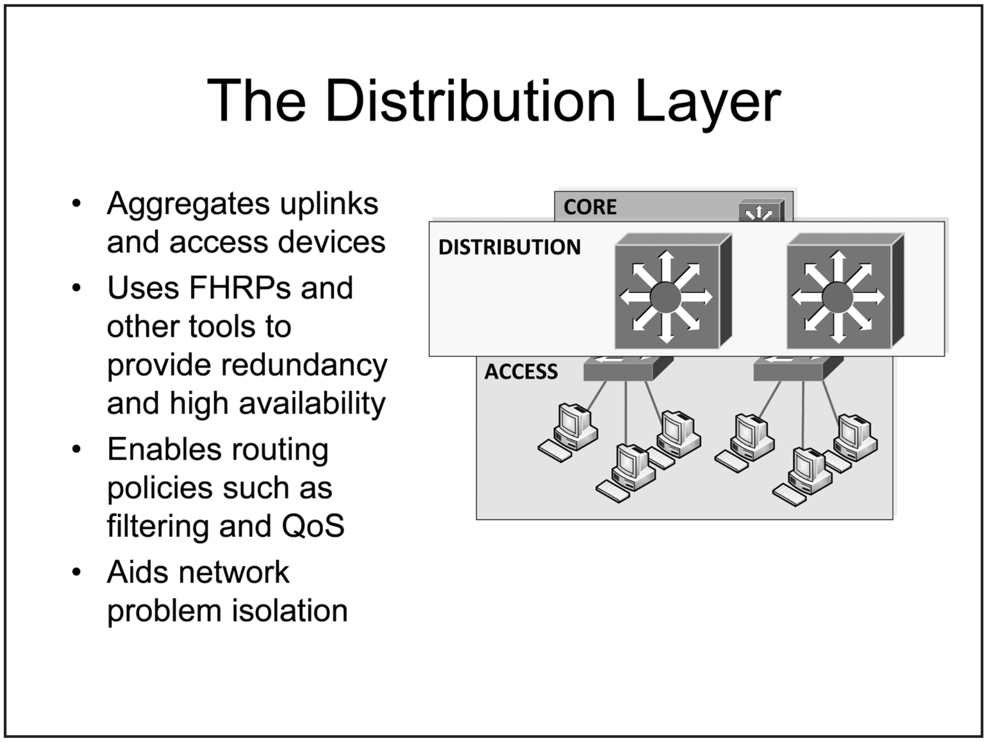

The Distribution Layer

- Aggregates uplinks and access devices
- Uses FHRPs and other tools to provide redundancy and high availability
- Enables routing policies such as filtering and QoS
- Aids network problem isolation

CORE
DISTRIBUTION
ACCESS

The Distribution Layer

The distribution layer aggregates uplinks from access layer devices and provides services to both the access layer and the core layer. Positioned between the access and core layers, the distribution layer is the ideal place to perform functions such as QoS resource reservation, interVLAN routing, packet manipulation, and route summarization. In addition, access control lists (ACLs) can be used in the distribution layer to enforce organizational security policies and to filter the flow of traffic 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. In this type of 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, a First-Hop Redundancy Protocol (FHRP) is required in this design 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.

Some Cisco distribution layer switch platforms support a technology known as virtual switching system (VSS). With VSS, two physical switches operate as a single, virtual distribution layer switch. VSS removes the potential for Layer 2 loops, enables VLANs to span multiple access layer devices, and mitigates the need for an FHRP in a switched access layer design.

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.

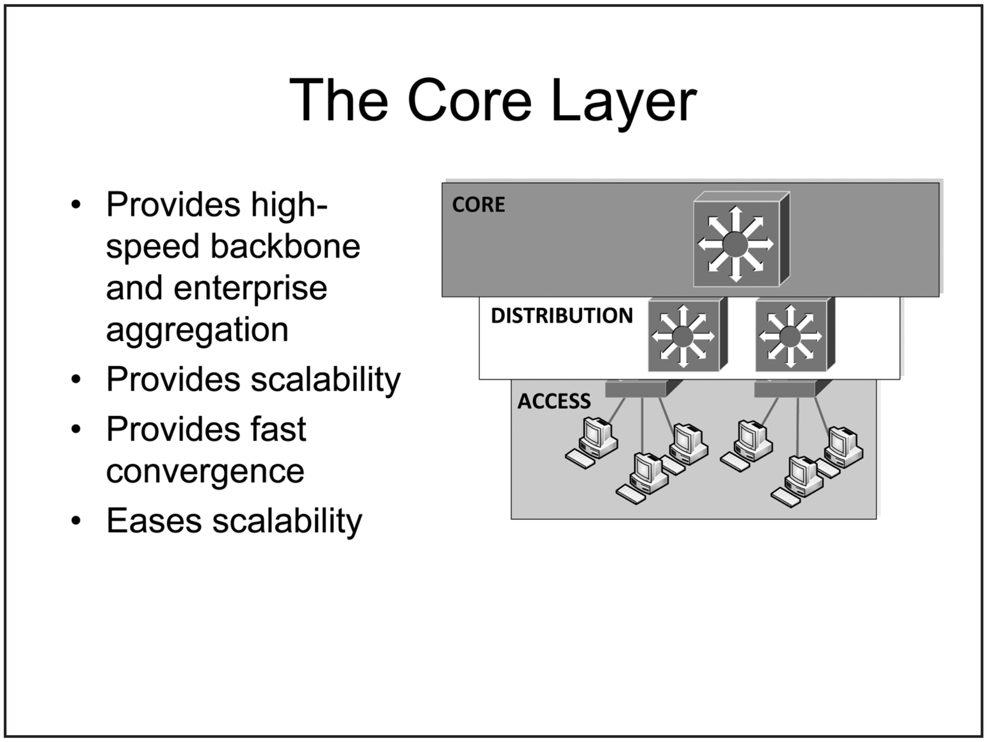

The Core Layer

- Provides high-speed backbone and enterprise aggregation
- Provides scalability
- Provides fast convergence
- Eases scalability

The Core Layer

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 and 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.

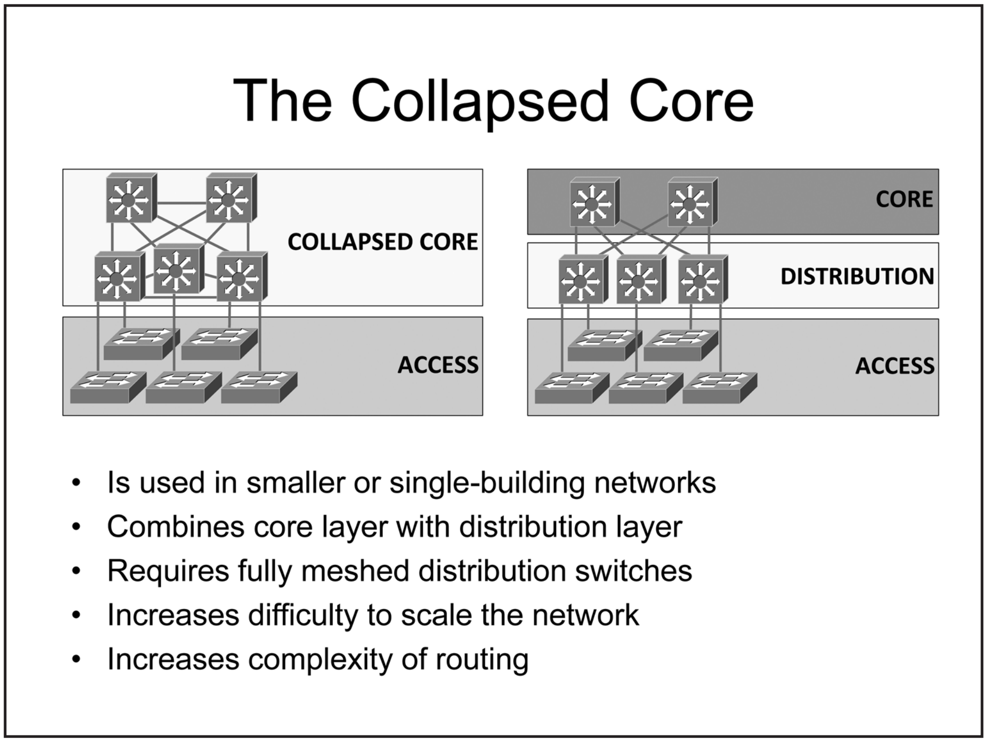

The Collapsed Core

COLLAPSED CORE

ACCESS

CORE

DISTRIBUTION

ACCESS

- Is used in smaller or single-building networks
- Combines core layer with distribution layer
- Requires fully meshed distribution switches
- Increases difficulty to scale the network
- Increases complexity of routing

The Collapsed Core

When the distribution layer contains a small number of switches or resides in a single building, a physically distinct core layer may not be necessary. In this case, the functionality of the core and distribution layers can be combined. However, it becomes necessary for the distribution layer switches to maintain a full-mesh topology to ensure that a link or device failure will not isolate part of the network.

Maintaining a full mesh of connections between distribution layer switches grows exponentially more difficult as the number of distribution layer switches is increases. 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.

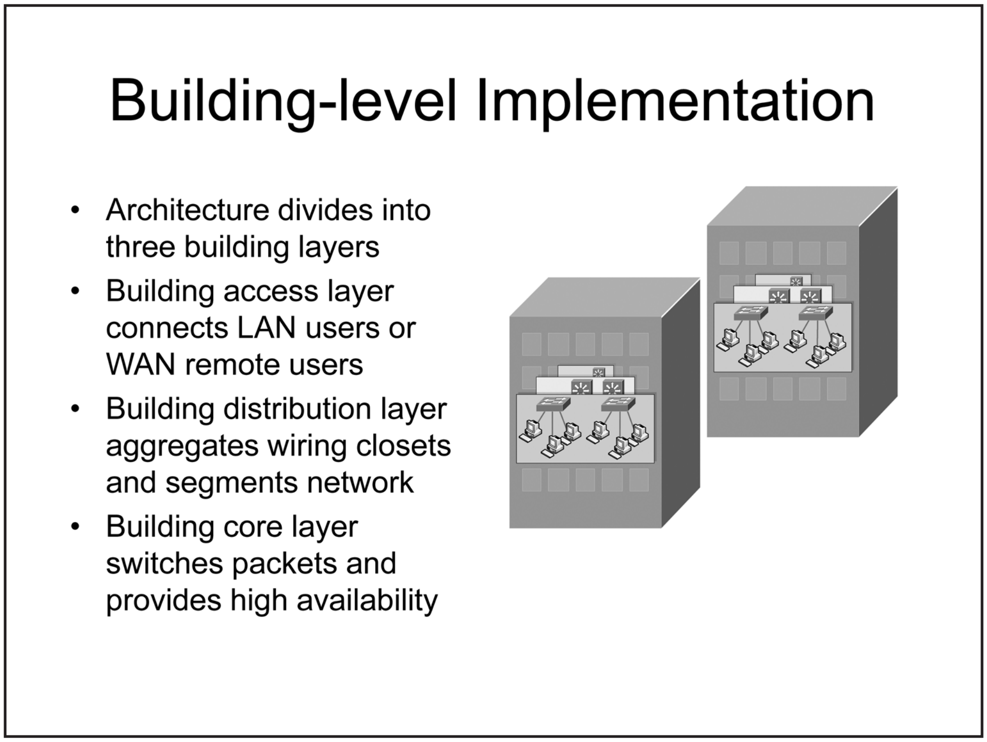

Building-level Implementation

- Architecture divides into three building layers
- Building access layer connects LAN users or WAN remote users
- Building distribution layer aggregates wiring closets and segments network
- Building core layer switches packets and provides high availability

Building-level Implementation

The hierarchical components of the Cisco Enterprise Campus Architecture can be implemented at the building level or across a campus or regional area. When implemented at the building level, the names of the hierarchical components are sometimes changed and their function somewhat refined. At the building level, the three component layers of the Cisco Enterprise Campus Architecture become the following:

- The building access layer
- The building distribution layer
- The building core layer

The building access layer provides LAN and possibly WAN connectivity to end-user devices, including workstations and servers. The building distribution layer aggregates uplinks from the building access layer and provides interVLAN routing and other network services for access layer devices. In building-level implementations, the distribution layer switches typically reside in a wiring closet known as the main distribution frame (MDF) whereas access layer switches might reside in smaller wiring closets known as intermediate distribution frames (IDFs) throughout the same floor as the MDF. The building core layer provides high-speed interconnectivity between distribution layer switches. In practice, core layer switches reside in the MDF and provide floor-to-floor connectivity between MDFs.

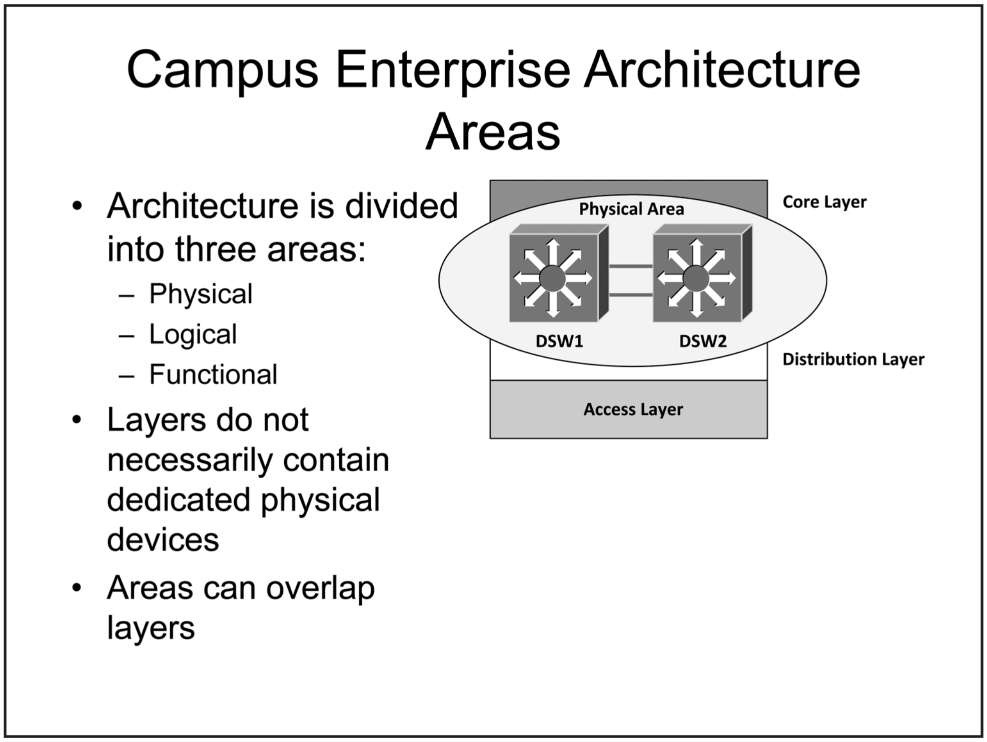

Campus Enterprise Architecture Areas

- Architecture is divided into three areas:
  - Physical
  - Logical
  - Functional
- Layers do not necessarily contain dedicated physical devices
- Areas can overlap layers

Campus Enterprise Architecture Areas

Hierarchical network designs in general, and the Cisco Enterprise Campus Architecture design model in particular, divide the network into modular components. These components represent physical, logical, and functional areas of separation. In some implementations, these three areas are identical; however, that may not always be the case. Within the Cisco Enterprise Campus Architecture design model, these areas can overlap and physical devices can reside within multiple logical or functional areas. The collapsed core is an example of this overlap because the physical area of the collapsed core spans both the distribution and core layers of the logical and functional areas. In addition, within a logical or functional area, there may be several tiers of physical devices instead of the single tier that is logically or functionally represented.

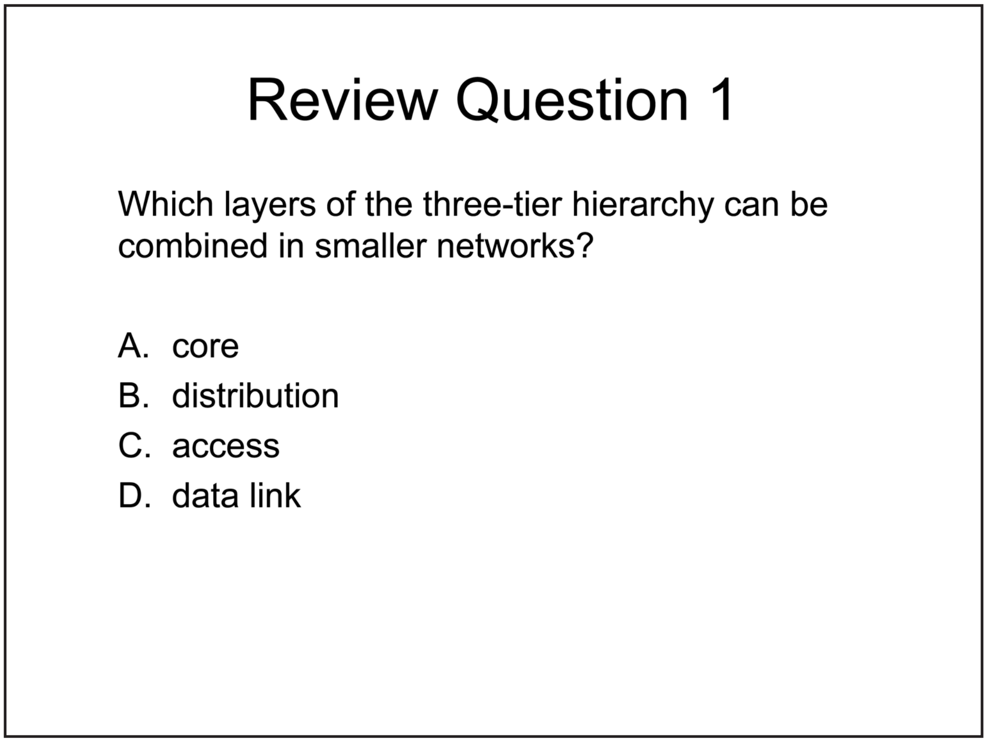

Review Question 1

Which layers of the three-tier hierarchy can be combined in smaller networks?

A. core
B. distribution
C. access
D. data link

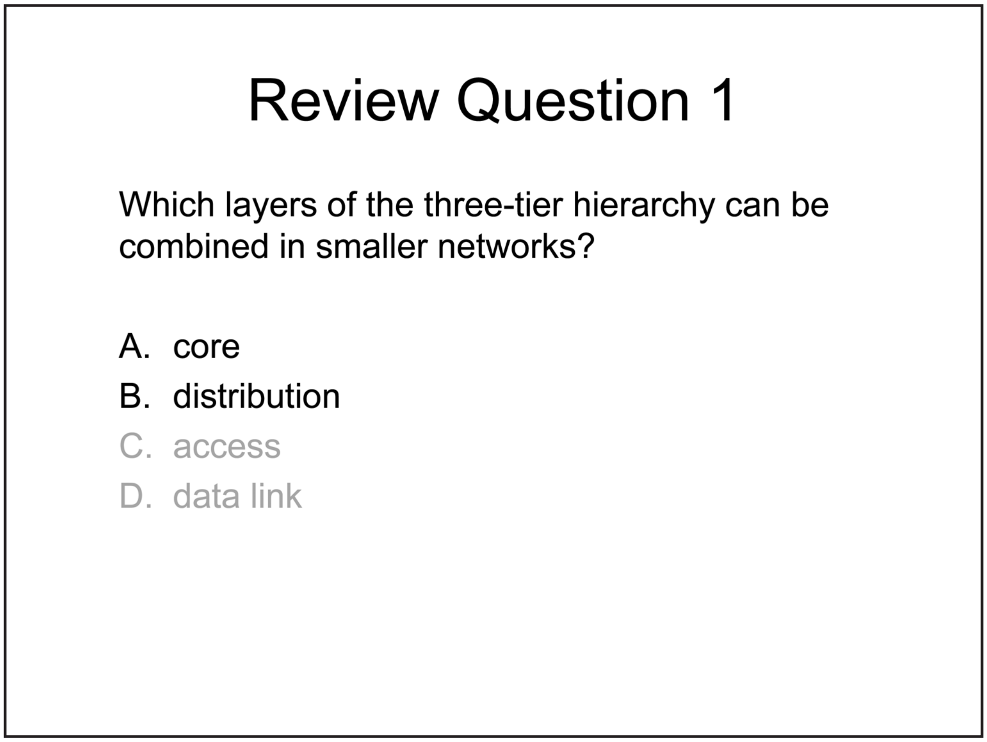

Review Question 1

Which layers of the three-tier hierarchy can be combined in smaller networks?

A. core
B. distribution
C. access
D. data link

The core and distribution layers of the three-tier hierarchy can be combined in smaller networks. When the distribution layer contains a small number of switches or resides in a single building, a physically distinct core layer may not be necessary. In this case, the functionality of the core and distribution layers can be combined.

The access layer is typically not combined with other layers in the three-tier hierarchy. Access layer devices are typically high-density switches focused on providing direct network access to end-user devices. Combining the access and distribution layers would reduce the effectiveness of each layer and undermine the purpose of the modular design of the three-tier hierarchy.

The Data Link layer is a layer of the Open Systems Interconnection (OSI) model and not a component of the three-tier hierarchy.

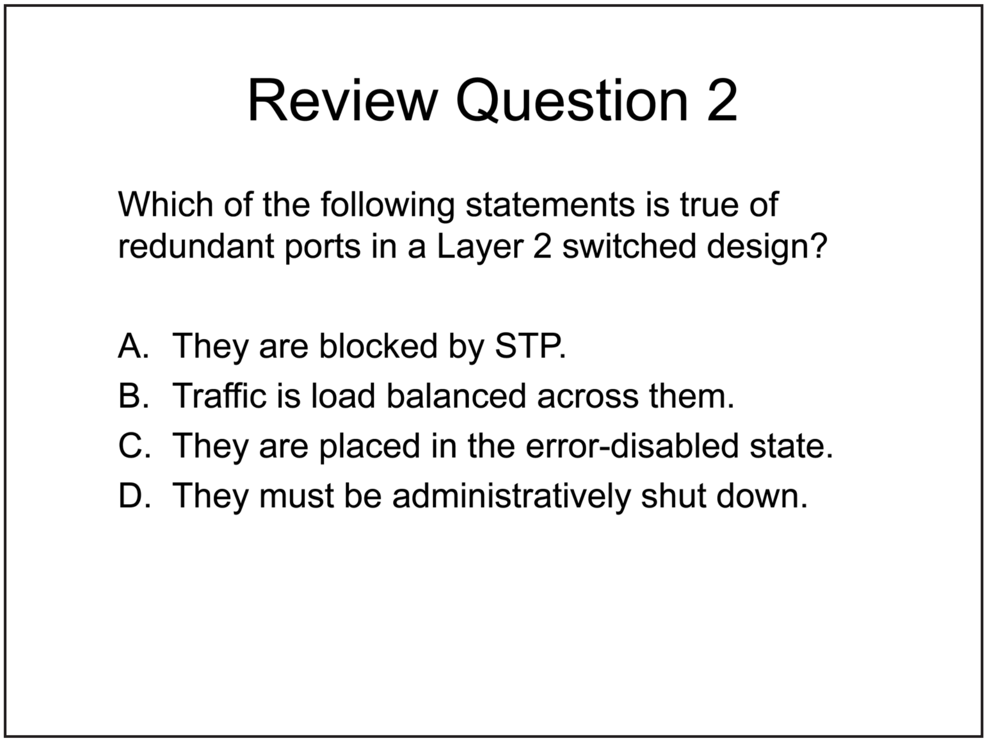

Review Question 2

Which of the following statements is true of redundant ports in a Layer 2 switched design?

A. They are blocked by STP.
B. Traffic is load balanced across them.
C. They are placed in the error-disabled state.
D. They must be administratively shut down.

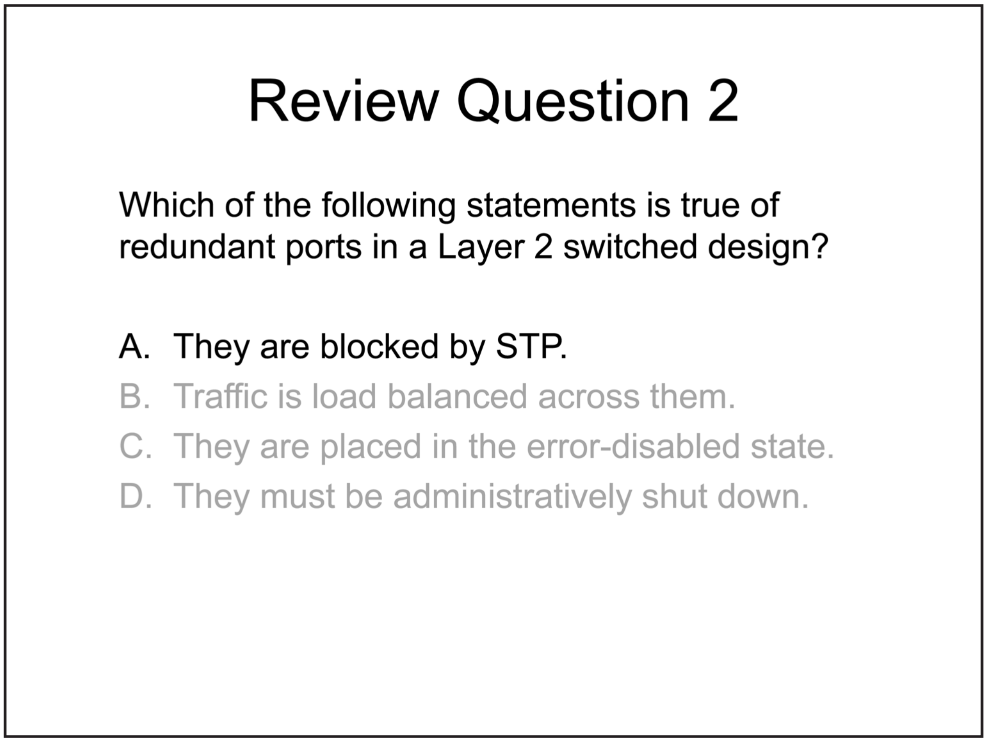

Review Question 2

Which of the following statements is true of redundant ports in a Layer 2 switched design?

A. They are blocked by STP.
B. Traffic is load balanced across them.
C. They are placed in the error-disabled state.
D. They must be administratively shut down.

In a Layer 2 switched design, Spanning Tree Protocol (STP) blocks redundant ports. STP is a Layer 2 protocol that is used to prevent switching loops from occurring on a network with redundant links. This process prevents the flooding of multiple copies of the same packet across a network by placing switch ports into either a forwarding or a blocking state. Ports in the forwarding state will forward received packets to the rest of the network. Ports in the blocking state will not forward packets.

STP does not enable load-balancing across redundant ports. Load balancing is supported in configurations where the potential for Layer 2 loops is mitigated, such as EtherChannel bundles or redundant Layer 3 ports.

Redundant ports are not placed in an error-disabled state nor must they be administratively shut down. STP places redundant ports into a blocking state and effectively disables the ports with respect to data transmission. Although the ports are in a blocking state, they still participate in STP processes and can become active if the topology changes because of a link or device failure.

Module 2

Switch Basics

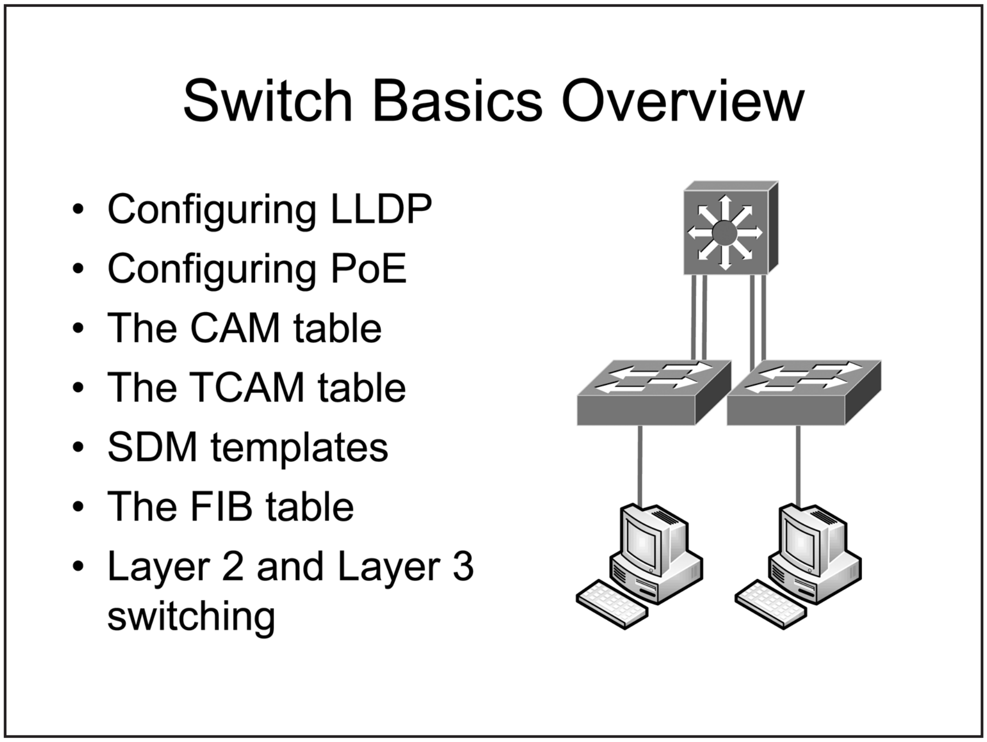

Switch Basics Overview

- Configuring LLDP
- Configuring PoE
- The CAM table
- The TCAM table
- SDM templates
- The FIB table
- Layer 2 and Layer 3 switching

Overview

This module explores the basic mechanisms that Cisco switches use to discover neighboring devices, negotiate power distribution, and forward network frames in a heterogeneous network environment. Switch memory pools, resource allocation, and resource management are also discussed.

Objectives

After completing this module, you should have the basic knowledge required to understand the following concepts:

- Neighbor discovery using Link-Layer Discovery Protocol (LLDP)
- Power negotiation and distribution using Power over Ethernet (PoE)
- Content-addressable memory (CAM) and ternary CAM (TCAM) tables
- Switch memory allocation and resource management using Switch Database Management (SDM) templates
- Layer 2 and Layer 3 network forwarding mechanisms

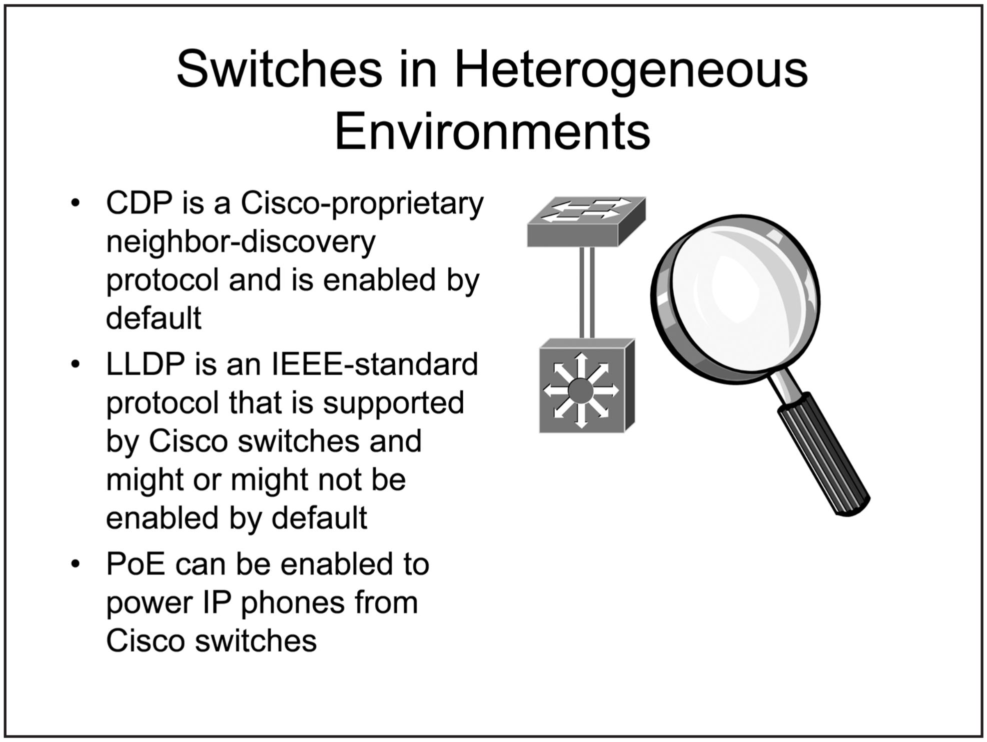

Switches in Heterogeneous Environments

- CDP is a Cisco-proprietary neighbor-discovery protocol and is enabled by default
- LLDP is an IEEE-standard protocol that is supported by Cisco switches and might or might not be enabled by default
- PoE can be enabled to power IP phones from Cisco switches

Switches in Heterogeneous Environments

In an ideal world, every device in a network design can be manufactured by a single vendor that offers hardware and software for every feature and service required by the network. In reality, such a homogenous network design is rarely possible. Inevitably, a network design will require devices from several vendors to interact and, in such a heterogeneous environment, standards-based support for features such as neighbor discovery and power distribution becomes vital.

For example, Cisco Discovery Protocol (CDP) is a proprietary, neighbor discovery protocol natively supported by nearly every Cisco-branded hardware platform; however, almost no third-party vendors offer support for CDP. In a heterogeneous network, a standards-based neighbor discovery protocol, such as LLDP, must be used if Cisco devices must interact with devices from other vendors. Likewise, if PoE power distribution is required in a heterogeneous network, a standards-based solution, such as Institute of Electrical and Electronics Engineers (IEEE) 802.3af, should be used instead of Cisco's native PoE implementation.

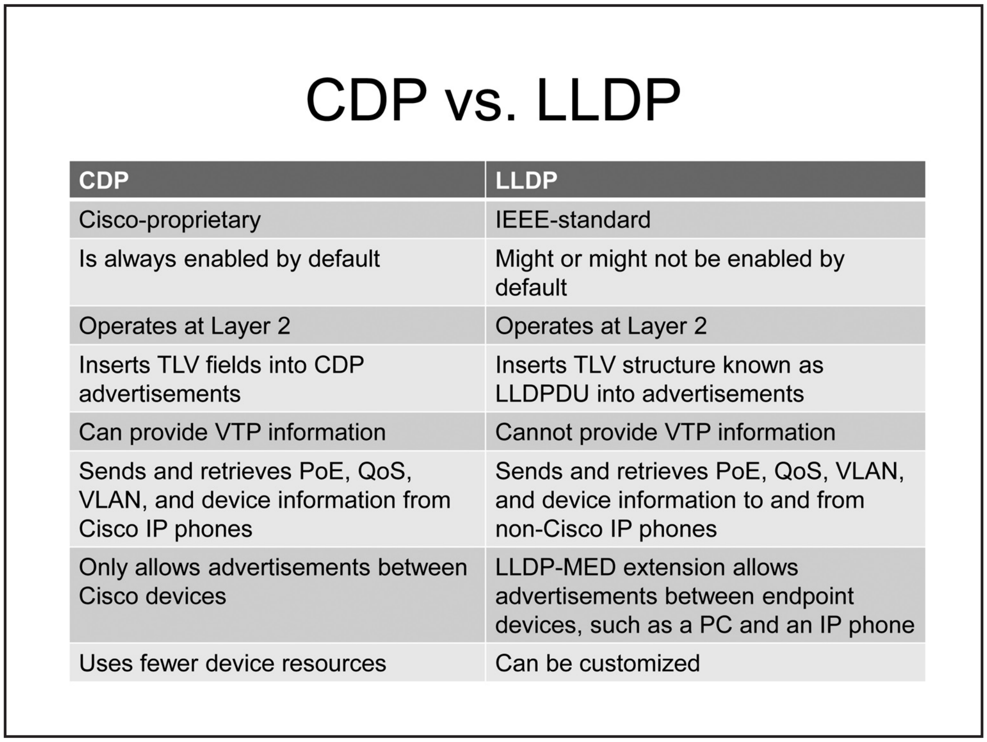

CDP vs. LLDP

CDP	LLDP
Cisco-proprietary	IEEE-standard
Is always enabled by default	Might or might not be enabled by default
Operates at Layer 2	Operates at Layer 2
Inserts TLV fields into CDP advertisements	Inserts TLV structure known as LLDPDU into advertisements
Can provide VTP information	Cannot provide VTP information
Sends and retrieves PoE, QoS, VLAN, and device information from Cisco IP phones	Sends and retrieves PoE, QoS, VLAN, and device information to and from non-Cisco IP phones
Only allows advertisements between Cisco devices	LLDP-MED extension allows advertisements between endpoint devices, such as a PC and an IP phone
Uses fewer device resources	Can be customized

CDP vs. LLDP

CDP is a Cisco-proprietary protocol and is enabled by default on Cisco hardware platforms. CDP can be used to convey information from other Cisco-proprietary protocols, such as VLAN Trunking Protocol (VTP). In addition, because CDP is limited to Cisco hardware platforms and the features they provide, the protocol is optimized to use as few system resources as possible.

By contrast, LLDP is an open-standard protocol officially known as IEEE 802.1AB. Although LLDP is an open standard, the standard does not specify whether LLDP should be enabled or disabled by default. Therefore, LLDP can be found either enabled by default or disabled by default in various implementations. For example, LLDP is disabled by default on most Cisco platforms. Because it is an open-standard protocol, LLDP is supported on a wide variety of hardware platforms and can be customized to exchange information from a large number of protocols and network services. For example, the LLDP Media Endpoint Discover (LLDP-MED) extension enables LLDP to exchange information with endpoints, such as IP phones, and to discover and negotiate PoE power distribution. LLDP cannot source information from Cisco-proprietary protocols, such as VTP, but it can source information from standards-based protocols.

CDP and LLDP share many attributes essential to the neighbor discovery process. Both are Link layer protocols; therefore, they are able to operate without the necessity of Network layer addressing. They both use type, length, and value (TLV) fields to advertise device capabilities and configuration information. However, CDP and LLDP TLVs are not compatible; CDP TLVs are contained in CDP advertisements, whereas LLDP TLVs are contained in LLDP Data Unit (LLDPDU) advertisements. Both protocols are capable of exchanging PoE, Quality of Service (QoS), virtual LAN (VLAN), and device configuration information.

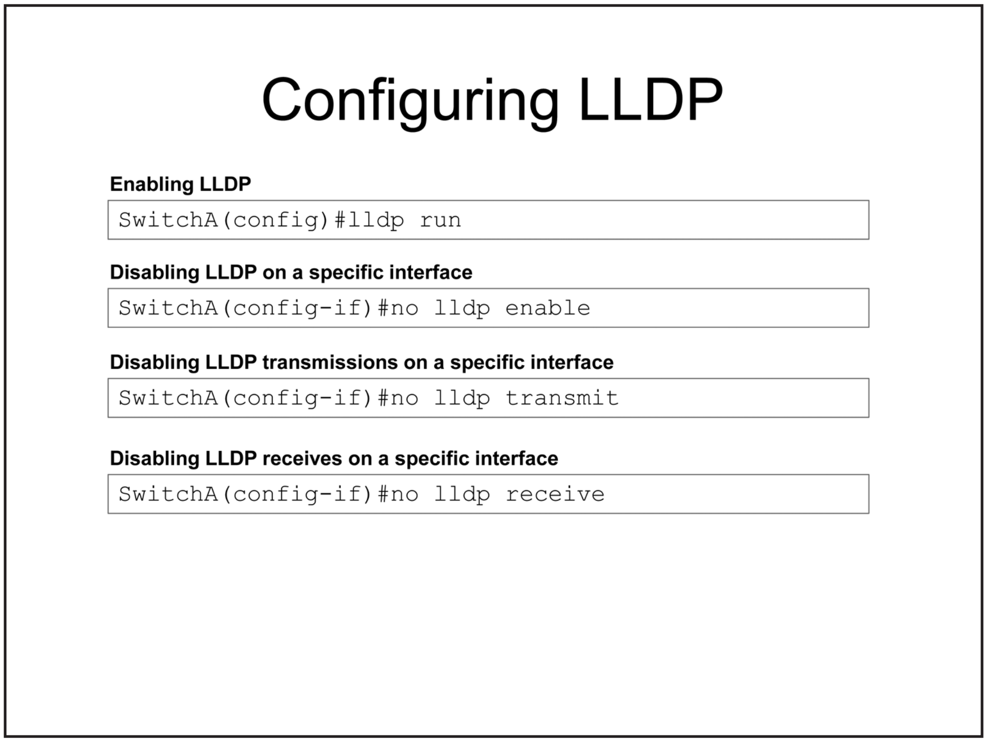

Configuring LLDP

Enabling LLDP

```
SwitchA(config)#lldp run
```

Disabling LLDP on a specific interface

```
SwitchA(config-if)#no lldp enable
```

Disabling LLDP transmissions on a specific interface

```
SwitchA(config-if)#no lldp transmit
```

Disabling LLDP receives on a specific interface

```
SwitchA(config-if)#no lldp receive
```

Configuring LLDP

On supported Cisco hardware platforms, you can issue the **lldp run** command from global configuration mode to enable LLDP globally. Once LLDP has been globally enabled, its operation can be limited on individual interfaces. You can issue the **no lldp enable** command from interface configuration mode to completely disable LLDP on a specific interface. Alternatively, LLDP can be partially disabled at the interface level. You can issue the **no lldp transmit** command from interface configuration mode to disable LLDP advertisements from a specific interface. Likewise, you can issue the **no lldp receive** command from interface configuration mode to prevent an interface from processing received LLDP advertisements.

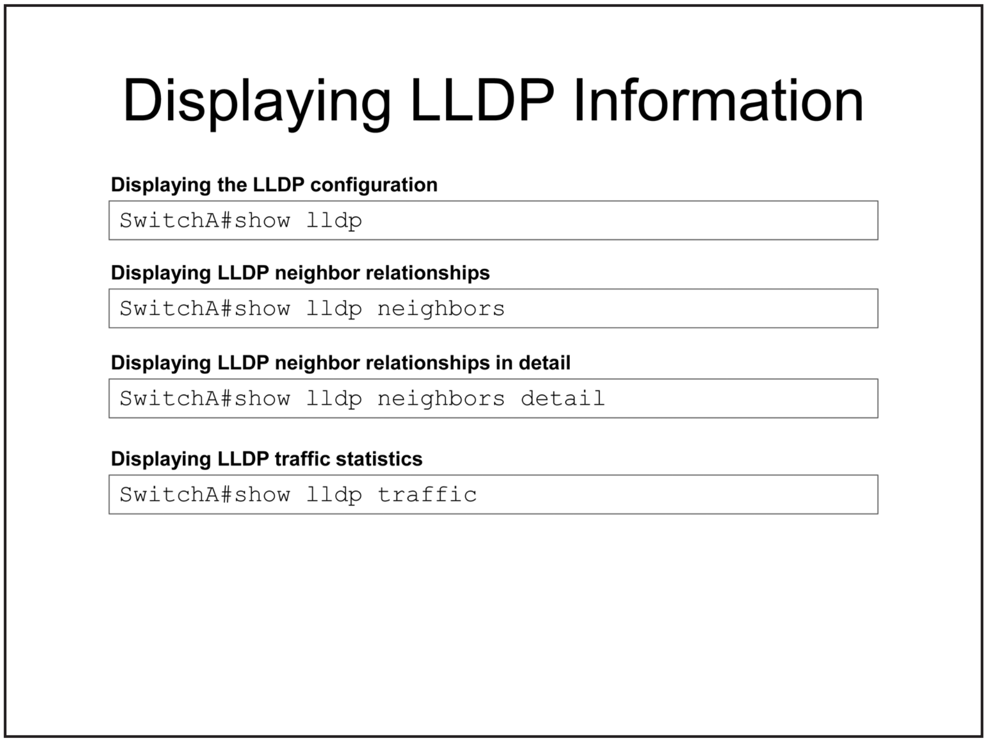

Displaying LLDP Information

Displaying the LLDP configuration

```
SwitchA#show lldp
```

Displaying LLDP neighbor relationships

```
SwitchA#show lldp neighbors
```

Displaying LLDP neighbor relationships in detail

```
SwitchA#show lldp neighbors detail
```

Displaying LLDP traffic statistics

```
SwitchA#show lldp traffic
```

Displaying LLDP Information

You can issue the **show lldp** command from privileged EXEC mode to verify global LLDP configuration parameters, such as the frequency of advertisement transmissions, the length of time a device should store information about a neighboring device before discarding it, and the length of time a device should wait before initializing an LLDP interface when needed. In addition, the **show lldp** command reports to the state of LLDP on the device. The LLDP state should be listed as `ACTIVE` if LLDP is globally enabled.

You can issue the **show lldp neighbors** [*interface-id*] [**detail**] command to view information about neighboring devices. When used without optional keywords, the **show lldp neighbors** command displays a tabular list of neighboring devices and their general capabilities. The example below indicates that SwitchB is a repeater and a network bridge, both of which are capabilities that define a modern switch:

```
SwitchA#show lldp neighbors

 Capability codes:
     (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
     (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other

 Device ID           Local Intf     Hold-time  Capability      Port ID
 SwitchB             Fa0/1          120        P,B             Fa0/11
```

You can issue the **show lldp neighbors detail** command to display more detailed information about neighboring devices, such as their management IP addresses, firmware revisions, and MED details. In addition, you can use an interface ID as an optional parameter to limit the displayed information to neighbors discovered on a specific interface.

The **show lldp traffic** command displays traffic statistics for various types of LLDP frames. The following example indicates that SwitchA has transmitted 412 LLDP advertisements and received 304 TLV advertisements from neighboring devices:

```
SwitchA#show lldp traffic

LLDP traffic statistics:

   Total frames out: 412

   Total entries aged: 0

   Total frames in: 304

   Total frames received in error: 0

   Total frames discarded: 0

   Total TLVs discarded:  63

   Total TLVs unrecognized:  63
```

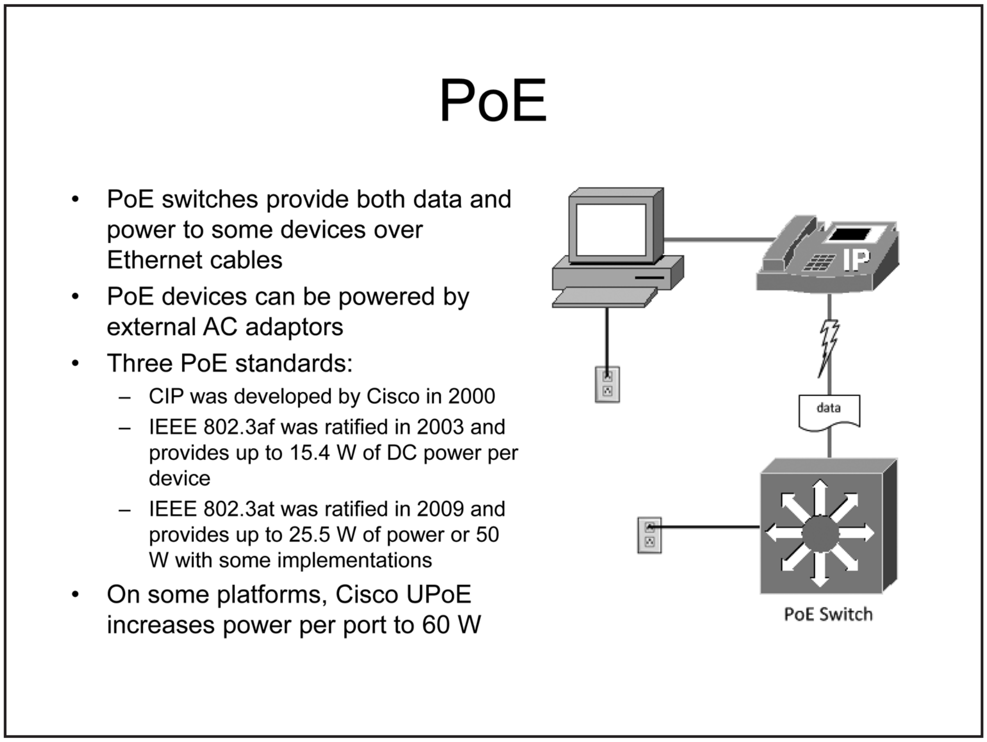

PoE

- PoE switches provide both data and power to some devices over Ethernet cables
- PoE devices can be powered by external AC adaptors
- Three PoE standards:
  - CIP was developed by Cisco in 2000
  - IEEE 802.3af was ratified in 2003 and provides up to 15.4 W of DC power per device
  - IEEE 802.3at was ratified in 2009 and provides up to 25.5 W of power or 50 W with some implementations
- On some platforms, Cisco UPoE increases power per port to 60 W

PoE

PoE enables a device, referred to as a powered device (PD), to draw power from the same physical cable that delivers network data if there is a compatible power sourcing equipment (PSE), such as a PoE switch, at the other end of the link. Some PoE-capable devices can be powered through alternate means, such as through a separate power adapter or through an inline, power injector.

Cisco pioneered the development of inline power (ILP) delivery starting in 2000 with a proprietary solution known as Cisco Inline Power (CIP). In 2003, the IEEE ratified the 802.3af standard to provide a uniform method of providing up to 15.4 watts (W) of power over standard Category 5 cabling. Although an 802.3af-compliant PSE can provide up to 15.4 W of power, a PD is only guaranteed to receive up to 12.95 W as some power is lost during transmission. In 2009, the IEEE ratified the 802.3at standard, which is also known as Power over Ethernet Plus (PoE+). The 802.3at standard increased the maximum power distribution to 30 W, with a maximum of 25.5 W of guaranteed power, for 802.3at-compliant PDs. With 802.3at, a PD can receive up to 50 W of power if all four copper pairs are used for power transmission. As of 2013, both 802.3af and 802.3at have been integrated into the 802.3 Ethernet standard.

The IEEE PoE committee continues to work on improving standardized PoE delivery. A proposed PoE standard, which will likely emerge as 802.3bt, will ultimately be able to provide up to 50 W of power for mid-powered PDs and 100 W of power for high-powered PDs. In the interim, Cisco has developed its own Universal PoE (UPoE) standard to provide up to 60 W of power to Cisco PDs.

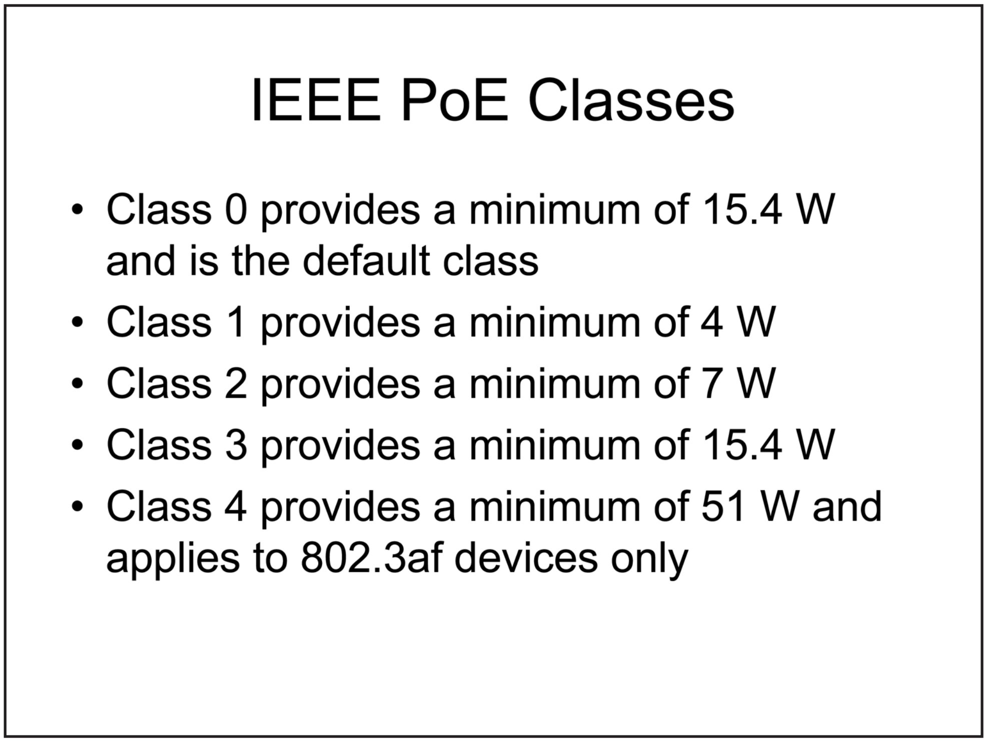

IEEE PoE Classes

- Class 0 provides a minimum of 15.4 W and is the default class
- Class 1 provides a minimum of 4 W
- Class 2 provides a minimum of 7 W
- Class 3 provides a minimum of 15.4 W
- Class 4 provides a minimum of 51 W and applies to 802.3af devices only

IEEE PoE Classes

The 802.3af standard categorizes PDs into classes based on the amount of power they require for operation. Class 0 is the default class and is used if a PSE cannot determine the appropriate class of a PD through its available means of discovery. Classes 1, 2, and 3 refer to Very Low Power, Low Power, and Mid Power PDs, respectively. These four classes are also supported within the 802.3at standard.

However, the 802.3at standard adds a fifth class, Class 4, that refers to High Power devices. High Power devices are also referred to as Type 2 devices and are not supported by the 802.3af standard.

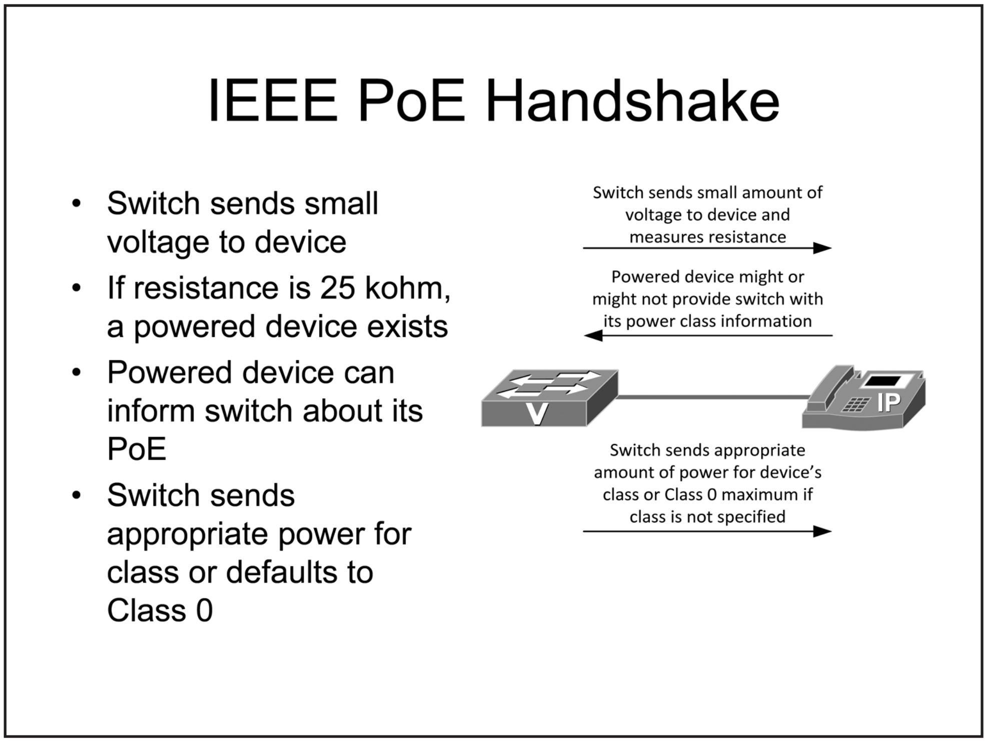

IEEE PoE Handshake

- Switch sends small voltage to device
- If resistance is 25 kohm, a powered device exists
- Powered device can inform switch about its PoE
- Switch sends appropriate power for class or defaults to Class 0

Switch sends small amount of voltage to device and measures resistance

Powered device might or might not provide switch with its power class information

Switch sends appropriate amount of power for device's class or Class 0 maximum if class is not specified

IEEE PoE Handshake

Because sending an electrical current to a device that does not support PoE could potentially damage the receiving device, a PSE will first apply a small voltage to a PoE-capable port to determine whether a PD is attached to the port. The IEEE PoE standards require a PD to provide a measurable resistance of approximately 25 kilo Ohms (kohms) when probed by a PSE. If the PSE detects a PD, the PSE can then send a signal with a higher voltage to determine the class of the PD. When an 802.3af-capable or 802.3at-capable PD receives this higher-voltage signal from a PSE, its response will inform the PSE about the PD's power requirements. The PSE will categorize the PD into an appropriate class, if possible, and will then guarantee a minimum amount of power relative to the class of the PD. If the PSE cannot identify the appropriate class for a PD, the PD will be categorized into the default class and will receive the default amount of power.

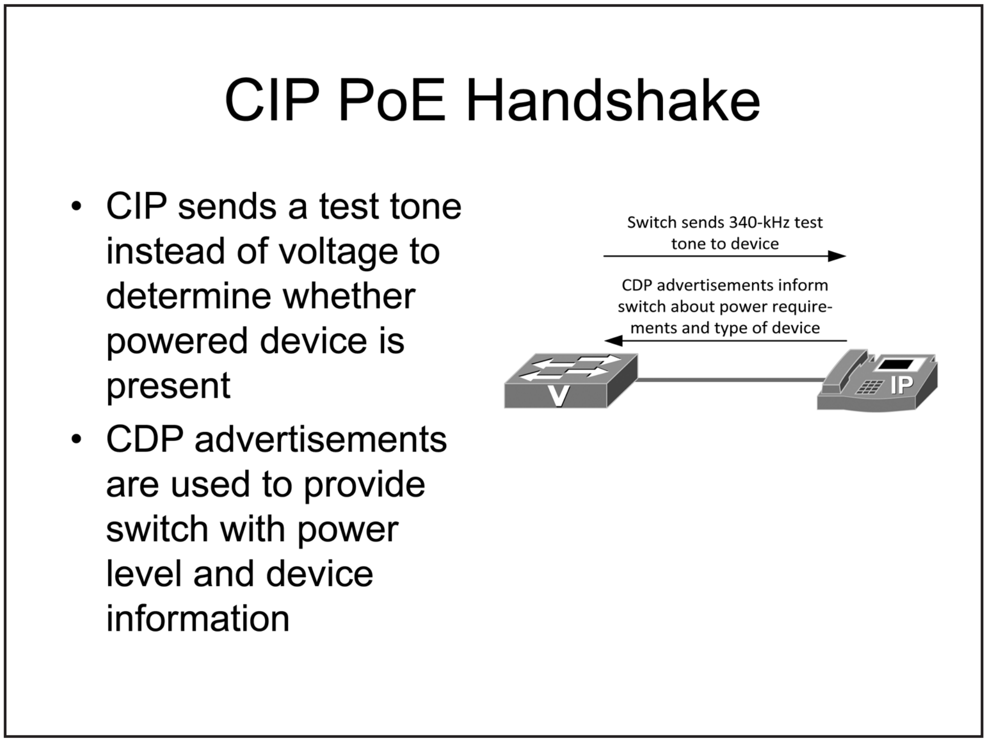

CIP PoE Handshake

- CIP sends a test tone instead of voltage to determine whether powered device is present
- CDP advertisements are used to provide switch with power level and device information

Switch sends 340-kHz test tone to device

CDP advertisements inform switch about power requirements and type of device

CIP PoE Handshake

CIP uses a power-initialization method that differs from the method standardized by the IEEE. With CIP, a PSE transmits a 340-kilohertz (kHz) analog test tone of a frequency on the transmit pins of a port. Because a CIP-capable device initially loops its transmit and receive pairs, a received test tone would simply be transmitted back to the PSE. If the PSE detects the test tone on the receive pins of a port, it has determined that a CIP-capable PD is attached to the port and it provides 48 volts (V) voltage to the port. Otherwise, it does not provide power to the port.

Once a CIP-capable PD has received its initial voltage from the PSE, it can use CDP to exchange information with the PSE and to request a specific amount of power for normal operations.

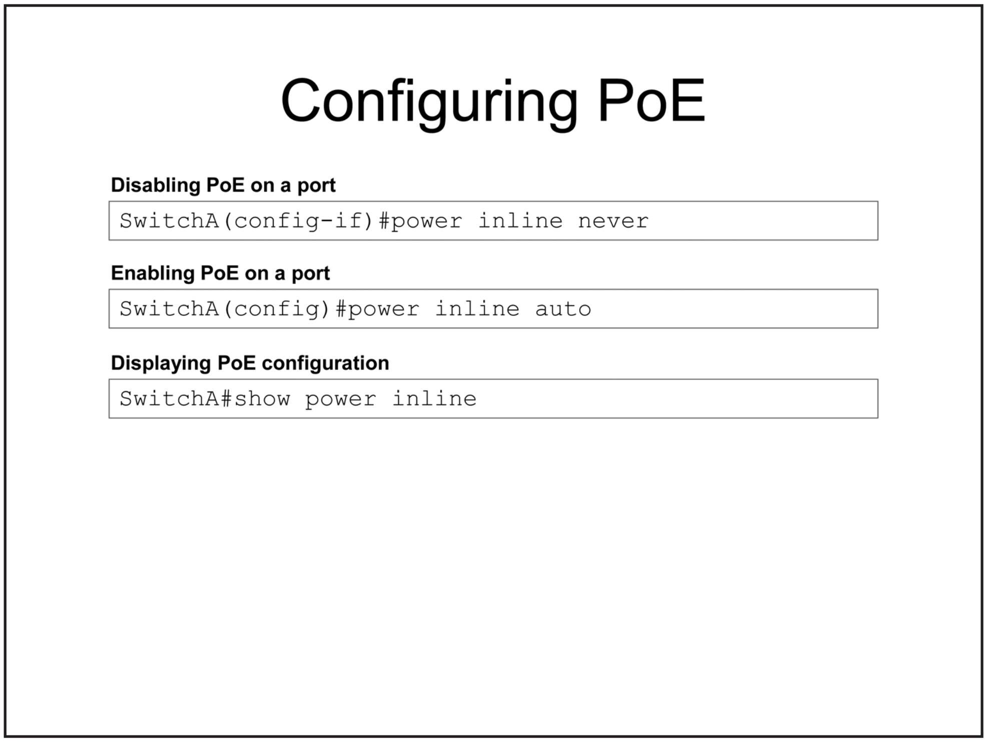

Configuring PoE

Disabling PoE on a port

```
SwitchA(config-if)#power inline never
```

Enabling PoE on a port

```
SwitchA(config)#power inline auto
```

Displaying PoE configuration

```
SwitchA#show power inline
```

Configuring PoE

By default, power is automatically provided on port if a PSE detects a PD; however, power is never supplied to ports that are disabled or shut down. There may be times when a port must be configured so that it can be enabled without the risk of inadvertently providing power to attached devices. You can issue the **power inline never** command from interface configuration mode on a Cisco switch to disable PoE for an interface. Conversely, you can issue the **power inline auto** command to restore the PoE capabilities of a port on which PoE has been disabled.

You can issue the **show power inline** command from privileged EXEC mode on a Cisco switch to examine a summary of the inline power usage on the switch. The output from the **show power inline** command displays the maximum total power that the switch can provide, the total amount of power currently being provided, and a summary of the PoE status of every individual port. In the example output below, the switch has detected a Cisco IP Phone on interface FastEthernet 0/2 and is providing the IP phone with 15.4 W of power:

```
Switch>show power inline
Module   Available     Used     Remaining
          (Watts)     (Watts)    (Watts)
------   ---------   --------   ---------
1             720        184.8         535.2
Interface Admin  Oper       Power   Device             Class Max
                            (Watts)
--------- ------ ---------- ------- ------------------- ----- ----
Fa0/1     auto   off        0.0     n/a                 n/a   30.0
Fa0/2     auto   off        15.4    Cisco IP Phone      n/a   15.4
Fa0/3     auto   off        0.0     n/a                 n/a   30.0
```

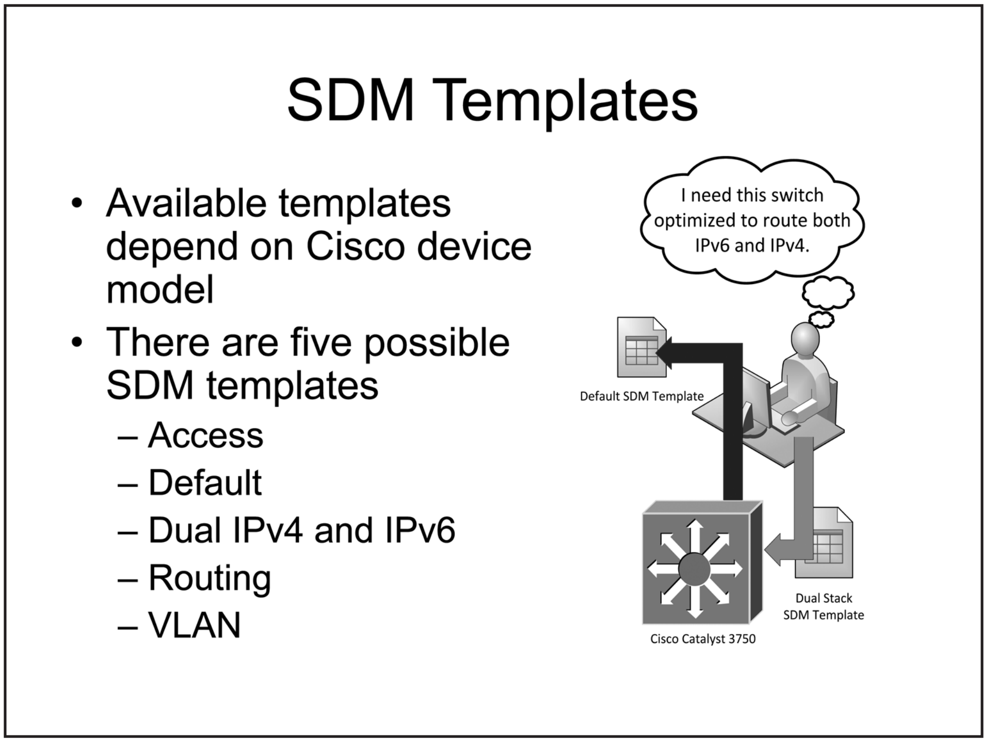

SDM Templates

- Available templates depend on Cisco device model
- There are five possible SDM templates
  - – Access
  - – Default
  - – Dual IPv4 and IPv6
  - – Routing
  - – VLAN

I need this switch optimized to route both IPv6 and IPv4.
Default SDM Template
Dual Stack SDM Template
Cisco Catalyst 3750

SDM Templates

Some Cisco switch platforms, such as the Catalyst 3750 series, support the customization of system resources to optimize performance for a specific type of network deployment. SDM templates provide a convenient way to configure the system resources on supported switches. The number and type of SDM templates available on a hardware platform depend on the hardware platform and the revision of the system software. In general, there are several SDM templates specific to IP version 4 (IPv4) networks and a comparable set of SDM templates for dual-stack networks.

The following four SDM templates are specific to IPv4 network implementations:

- Access
- Default
- Routing
- VLAN

The following four templates are specific to dual-stack network implementations:

- Dual IPv4 and IPv6 Default
- Dual IPv4 and IPv6 Routing
- Dual IPv4 and IPv6 VLAN
- Indirect Dual IPv4 and IPv6 Routing

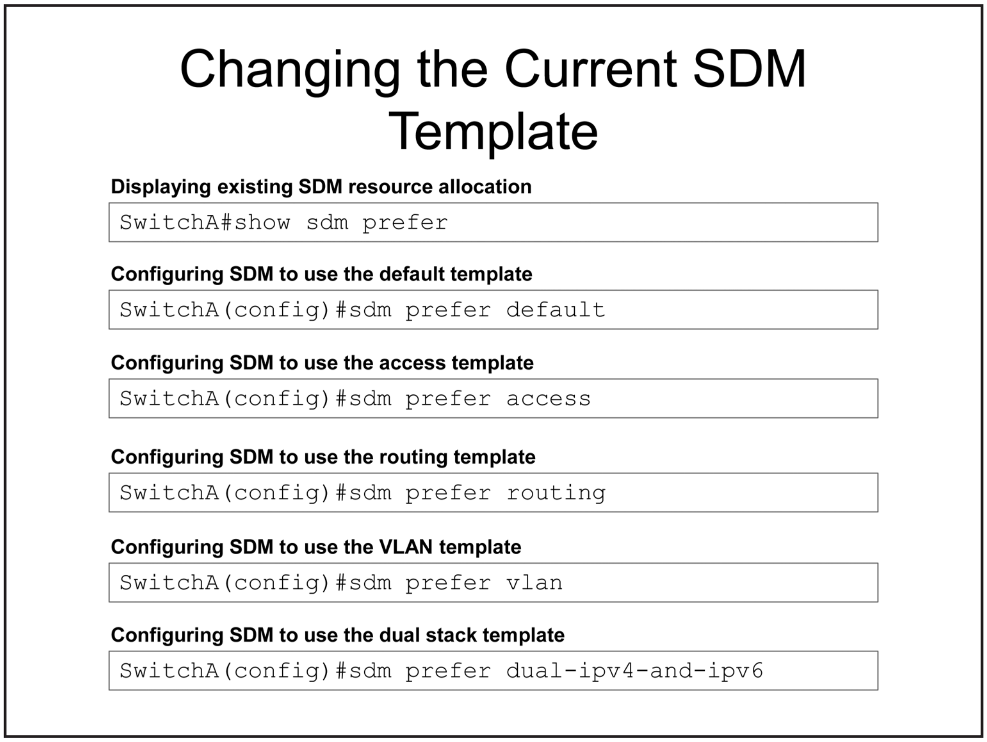

Changing the Current SDM Template

Displaying existing SDM resource allocation

```
SwitchA#show sdm prefer
```

Configuring SDM to use the default template

```
SwitchA(config)#sdm prefer default
```

Configuring SDM to use the access template

```
SwitchA(config)#sdm prefer access
```

Configuring SDM to use the routing template

```
SwitchA(config)#sdm prefer routing
```

Configuring SDM to use the VLAN template

```
SwitchA(config)#sdm prefer vlan
```

Configuring SDM to use the dual stack template

```
SwitchA(config)#sdm prefer dual-ipv4-and-ipv6
```

Changing the Current SDM Template

You can issue the **show sdm prefer** command from privileged EXEC mode to determine the SDM template that is currently applied to a switch. If no other SDM template has been configured on a switch, the switch will use the **default** SDM template, which provides a balanced distribution of system resources. The **default** SDM template provides sufficient resources for most general applications, except for policy-based routing (PBR). You can issue the **sdm prefer default** command from global configuration mode to restore the **default** SDM template on a switch.

In an environment requiring a large number of access control lists (ACLs), Cisco recommends using the **access** SDM template. Compared to the **default** SDM template, the access SDM template reduces the memory allocated for unicast Media Access Control (MAC) addresses and unicast routes. However, it increases the memory allocated for most types of Access Control Entries (ACEs), such as PBR and Security ACEs. You can issue the **sdm prefer access** command from global configuration mode to enable the **access** SDM template.

The **routing** SDM template optimizes switch resources for an environment wherein the switch is responsible for a large number of indirect unicast routes. Compared to the **default** SDM template, the **routing** SDM template reduces the memory allocated for unicast MAC addresses and directly connected hosts. However, it increases the memory allocated for indirect unicast routes and PBR ACEs. You can issue the **sdm prefer routing** command from global configuration mode to enable the **routing** SDM template.

The **vlan** SDM template optimizes switch resources for a Layer 2 switched environment by disabling routing and reallocating resources that had previously been allocated to handling Layer 3 information, such as unicast routes and PBR ACEs. Because the **vlan** SDM template disables routing, the switch must rely on an external router for interVLAN routing. Compared to the **default** SDM template, the **vlan** SDM template allocates nearly twice as much system resources for handling unicast MAC address information. You can issue the **sdm prefer vlan** command from global configuration mode to enable the **vlan** SDM template.

For dual-stack network implementations, Cisco provides three SDM templates that are analogous to their IPv4 counterparts. The **dual IPv4-and-IPv6** default SDM template provides a balanced allocation of system resources for IPv4 and IPv6 traffic. The **dual IPv4-and-IPv6 vlan** SDM template disables routing and maximizes system resource allocation for unicast MAC addresses and IPv6 unicast routes. The **dual IPv4-and-IPv6 routing** SDM template reduces the resources allocated to unicast MAC addresses directly connected IPv4 routes and maximizes the resources allocated to indirect routing information, such as indirect IPv4 and IPv6 unicast routes. You can issue the **sdm prefer ipv4-and-ipv6** {**default** | **vlan** | **routing**} command from global configuration mode to configure a dual-stack SDM template.

On some systems, a fourth dual-stack SDM template is available. The **indirect IPv4-and-IPv6 routing** SDM template is similar to the **dual IPv4-and-IPv6 routing** SDM template; however, it reduces the resources allocated to ACEs and increases the resource allocation for indirect IPv4 and IPv6 routing information.

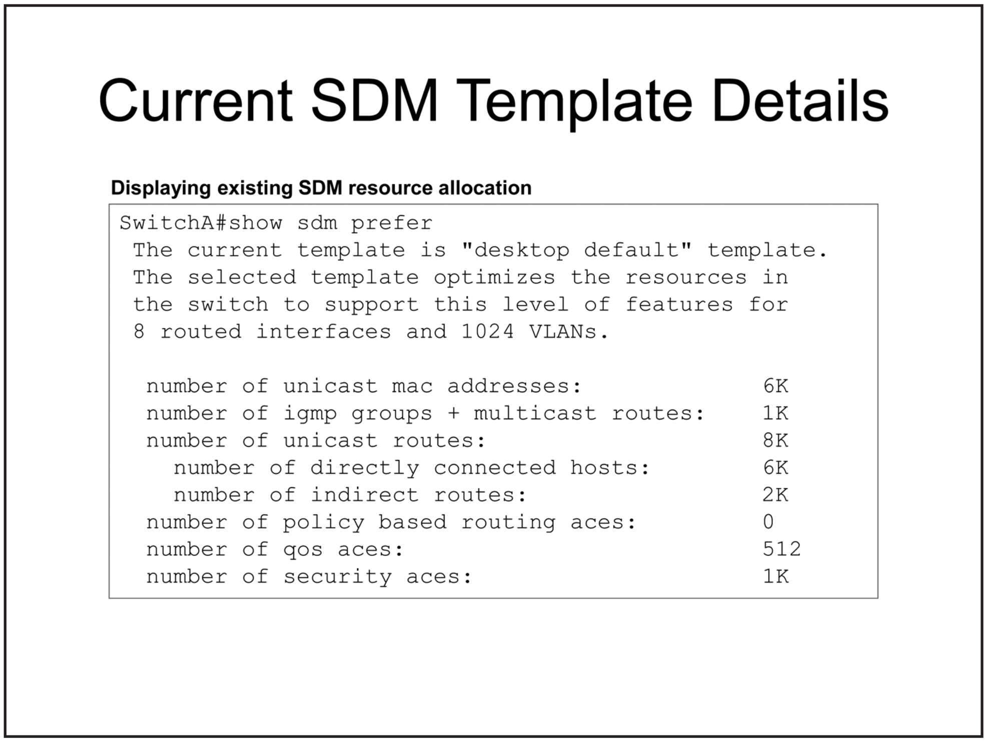

Current SDM Template Details

Displaying existing SDM resource allocation

```
SwitchA#show sdm prefer
 The current template is "desktop default" template.
 The selected template optimizes the resources in
 the switch to support this level of features for
 8 routed interfaces and 1024 VLANs.

  number of unicast mac addresses:             6K
  number of igmp groups + multicast routes:    1K
  number of unicast routes:                    8K
    number of directly connected hosts:        6K
    number of indirect routes:                 2K
  number of policy based routing aces:         0
  number of qos aces:                          512
  number of security aces:                     1K
```

Current SDM Template Details

You can issue the **show sdm prefer** command from privileged EXEC mode to display information about system resource allocation under the current SDM template. The command output can be useful for determining the limits for a category of information within the switch database under the current SDM template. For example, you could issue the **show sdm prefer** command to determine the maximum number of unicast MAC addresses or indirect IPv4 routes supported by the switch in its current configuration.

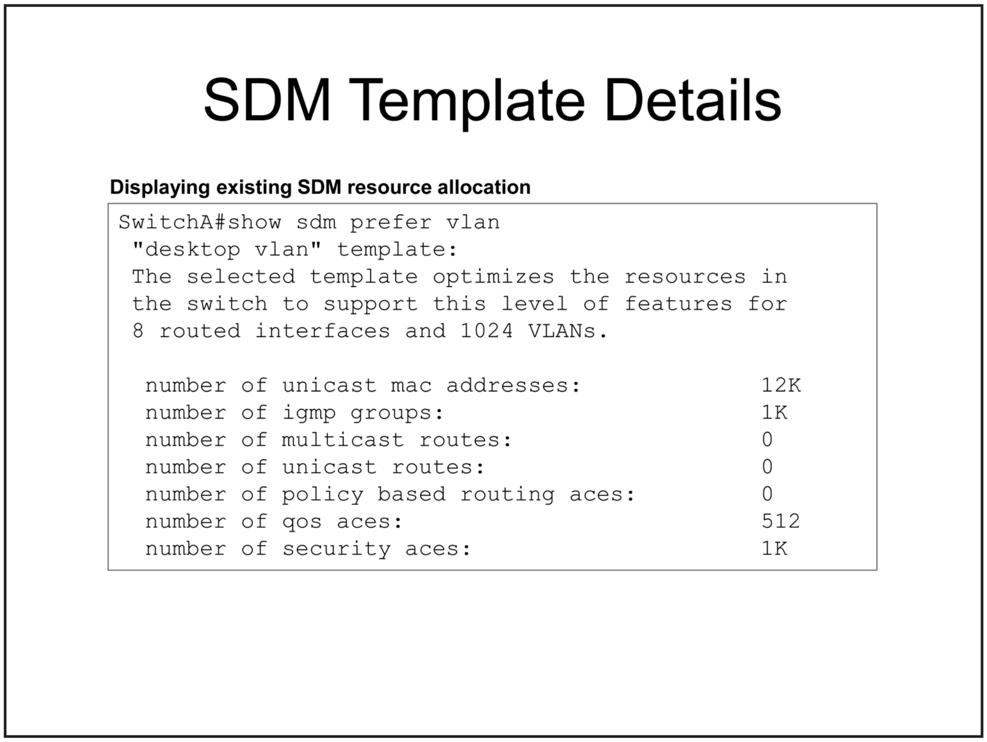

SDM Template Details

Displaying existing SDM resource allocation

```
SwitchA#show sdm prefer vlan
 "desktop vlan" template:
 The selected template optimizes the resources in
 the switch to support this level of features for
 8 routed interfaces and 1024 VLANs.

  number of unicast mac addresses:             12K
  number of igmp groups:                       1K
  number of multicast routes:                  0
  number of unicast routes:                    0
  number of policy based routing aces:         0
  number of qos aces:                          512
  number of security aces:                     1K
```

SDM Template Details

The **show sdm prefer** command can be used to examine the resource allocation for a specific SDM template supported by a switch, even if that SDM templates is not currently the active SDM template. You can specify the name of SDM template to examine how it would allocate system resources if it were configured as the active SDM template. For example, you could issue the **show sdm prefer vlan** command to examine the resource allocation by the **vlan** SDM template.

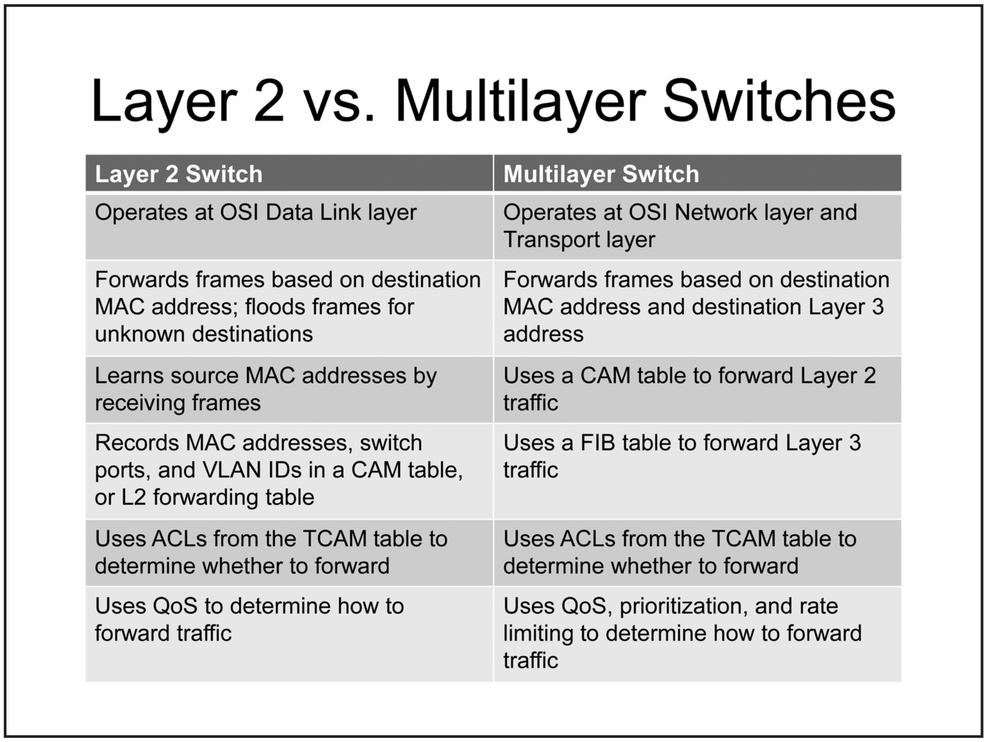

Layer 2 vs. Multilayer Switches

Layer 2 Switch	Multilayer Switch
Operates at OSI Data Link layer	Operates at OSI Network layer and Transport layer
Forwards frames based on destination MAC address; floods frames for unknown destinations	Forwards frames based on destination MAC address and destination Layer 3 address
Learns source MAC addresses by receiving frames	Uses a CAM table to forward Layer 2 traffic
Records MAC addresses, switch ports, and VLAN IDs in a CAM table, or L2 forwarding table	Uses a FIB table to forward Layer 3 traffic
Uses ACLs from the TCAM table to determine whether to forward	Uses ACLs from the TCAM table to determine whether to forward
Uses QoS to determine how to forward traffic	Uses QoS, prioritization, and rate limiting to determine how to forward traffic

Layer 2 vs. Multilayer Switches

A network device is traditionally considered to operate at the highest layer of the Open Systems Interconnection (OSI) model from which it can extract data to make forwarding decisions. An Ethernet switch uses Data Link layer information, such as MAC addresses, to make forwarding decisions; thus, it is said to operate at Layer 2. Likewise, a multilayer switch can use Network layer information, such as IP addresses, to make forwarding decisions; thus, a multilayer switch is considered to operate at Layer 3 and is often referred to as a Layer 3 switch. Although most Cisco multilayer switches can make forwarding decisions based on information from higher levels of the OSI model, such as the Transport layer, they are still commonly referred to as Layer 3 switches.

Layer 2 and Layer 3 switches both learn the MAC addresses of directly connected hosts. These addresses are stored in a Layer 2 forwarding table referred to as the CAM table, which is also known as the MAC address table. Switches populate the CAM table by extracting MAC addresses from received frames and associating those addresses with the received switch port and configured VLAN ID. The information stored in the CAM table is referenced when a switch needs to forward a frame destined for a particular MAC address. If the MAC address is found in the CAM table, the switch can immediately forward the frame to the appropriate port. However, if the destination MAC address is not found, the switch floods the frame to every port with the same VLAN configuration as the port from which the frame was received.

Multilayer switches maintain an additional forwarding table that is used to make forwarding decisions for Layer 3 traffic. The Layer 3 forwarding table, referred to as the Forwarding Information Base (FIB), merges information from the CAM table and the routing table. The FIB enables a multilayer switch to determine the appropriate forwarding information, such as next-hop MAC address, VLAN ID, and egress port for an IP packet with a single table lookup.

Layer 2 and Layer 3 switches can use Quality of Service (QoS) policies and ACLs to manage forwarding decisions. QoS policy parameters and ACL ACEs are stored in a special type of table, known as a TCAM table. TCAM tables enable a switch to rapidly process information queries without negatively impacting frame-forwarding rates. A switch can use separate TCAM tables for different data types, such as QoS parameters and ACL ACEs, so that table lookups can be performed in parallel to increase performance. In addition, Layer 3 switches can use IP priority values and rate limiting to manage traffic flow.

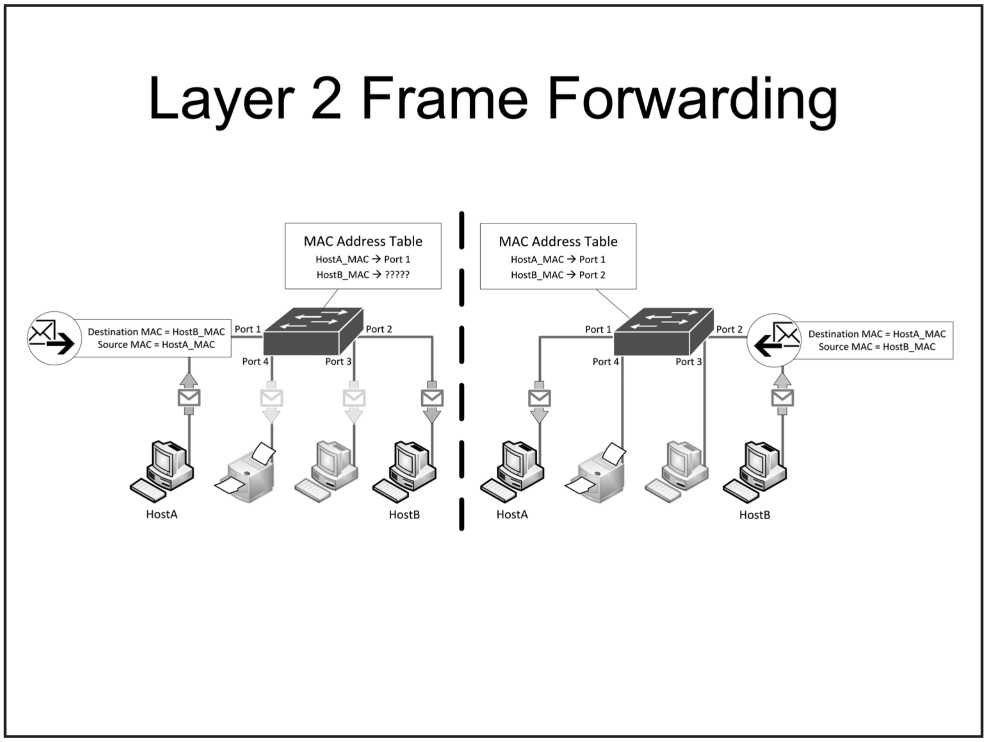

Layer 2 Frame Forwarding

Layer 2 Frame Forwarding

When a Layer 2 switch forwards a network frame, the switch effectively copies the frame from a memory location on the ingress port to a memory location on the appropriate egress port. These memory locations are referred to as queues. Every port is allocated a portion of system memory that is used for transmit and receive queues. Although frames are typically described as passing through a switch, in reality the data from the frame is copied from one memory location to another and then converted into the appropriate signal type for the transmission media involved.

For example, HostA constructs a network frame with its own MAC address as the source and the MAC address of HostB as the destination. HostA then transmits the frame to a Layer 2 switch. When the switch receives the frame from HostA, the switch examines the source MAC address to ensure that HostA is associated with the port on which the frame received the frame. The switch then examines the destination MAC address of the frame to make a forwarding decision. If the destination MAC address is unknown, the switch decides to forward the frame to every port other than the port on which the frame was received. This effectively creates a copy of the frame in the transmit queue of every other egress port. When HostB responds with a frame for HostA, the switch associates the MAC address of HostB with the port on which it received the response and then copies the frame directly to the transmit queue of the port attached to HostA because HostA was already associated with the appropriate egress interface.

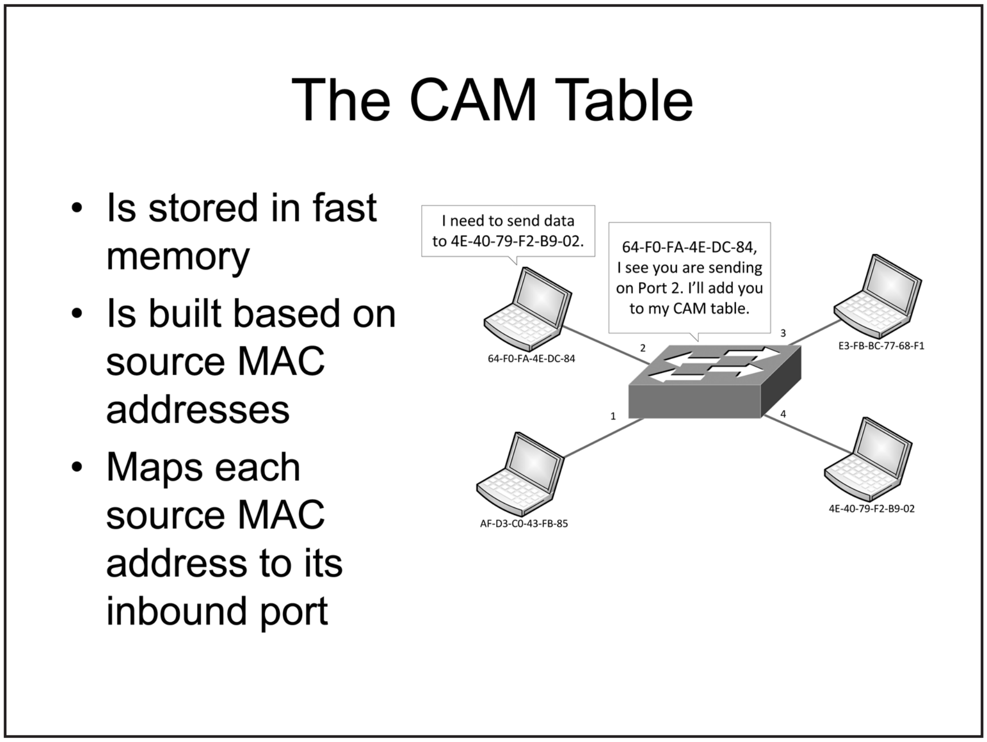

The CAM Table

- Is stored in fast memory
- Is built based on source MAC addresses
- Maps each source MAC address to its inbound port

I need to send data to 4E-40-79-F2-B9-02.
64-F0-FA-4E-DC-84, I see you are sending on Port 2. I'll add you to my CAM table.
64-F0-FA-4E-DC-84
2
3
E3-FB-BC-77-68-F1
1
4
AF-D3-C0-43-FB-85
4E-40-79-F2-B9-02

The CAM Table

The CAM table is allocated from a designated portion of system memory in a Layer 2 switch. System memory is faster than other types of memory, such as non-volatile random-access memory (NVRAM), and facilitates rapid searches and data manipulation. However, because the CAM table is allocated a fixed amount of system memory, attempting to store more data in the table than it can hold can cause adverse behavior on the switch.

The CAM table is constructed dynamically as the switch receives frames from the network. The switch associates source MAC addresses from network frames with the ports on which the frames are received. An entry is created for each source MAC address, its associated ingress port, and the VLAN ID associated with the port.

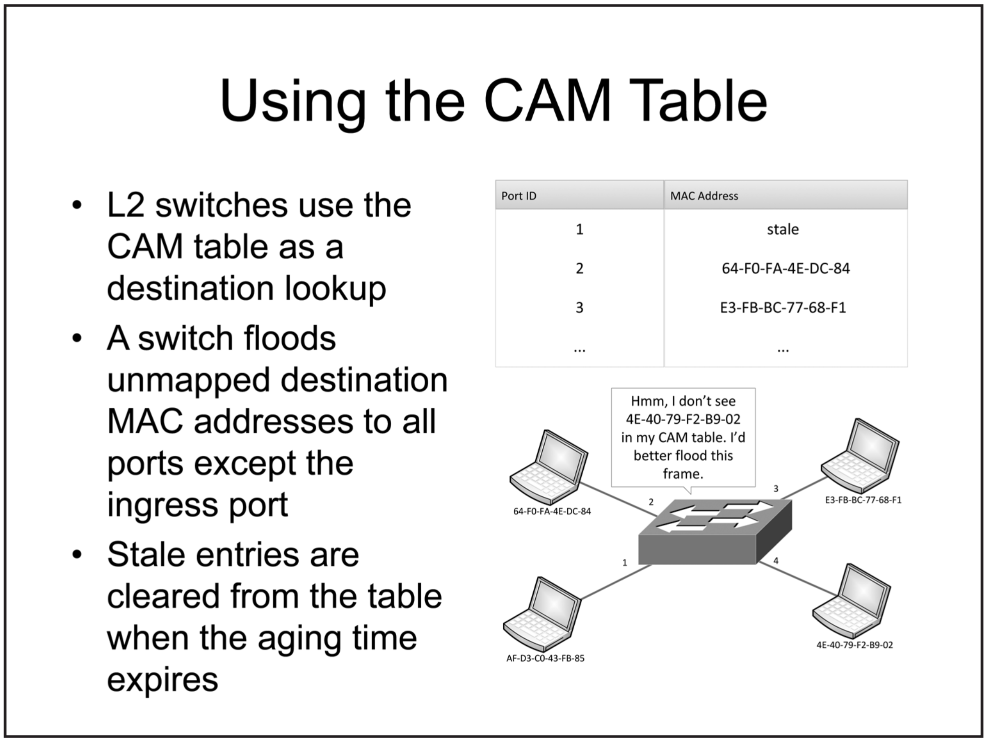

Using the CAM Table

- L2 switches use the CAM table as a destination lookup
- A switch floods unmapped destination MAC addresses to all ports except the ingress port
- Stale entries are cleared from the table when the aging time expires

Port ID	MAC Address
1	stale
2	64-F0-FA-4E-DC-84
3	E3-FB-BC-77-68-F1
…	…

Hmm, I don't see 4E-40-79-F2-B9-02 in my CAM table. I'd better flood this frame.
64-F0-FA-4E-DC-84
E3-FB-BC-77-68-F1
AF-D3-C0-43-FB-85
4E-40-79-F2-B9-02
1
2
3
4

Using the CAM Table

When a Layer 2 switch receives a frame on a port, the switch examines the source MAC address of the frame and creates an entry in its CAM table for that MAC address if one does not already exist. An entry is retained in the CAM table for a period known as the aging time. When the amount of time an entry has resided in the CAM table exceeds the aging time, the entry is marked as stale. Stale entries are ignored and ultimately removed from the CAM table if the switch does not receive another frame with the same source MAC address, port ID, and VLAN ID.

A Layer 2 switches uses its CAM table to determine the appropriate egress interface for a network frame. The CAM table is indexed by the destination MAC address. A switch examines the destination MAC address of a received frame and searches the CAM table for a corresponding entry. If a matching entry for the destination MAC address is found in the CAM table, the switch forwards the frame to the corresponding port in the CAM table. The VLAN ID of the source and destination ports must match, otherwise the switch cannot directly forward the frame from the source to the destination. If an entry for the destination MAC address is not found in the CAM table, the switch floods the frame to all ports within the VLAN identified by the source port.

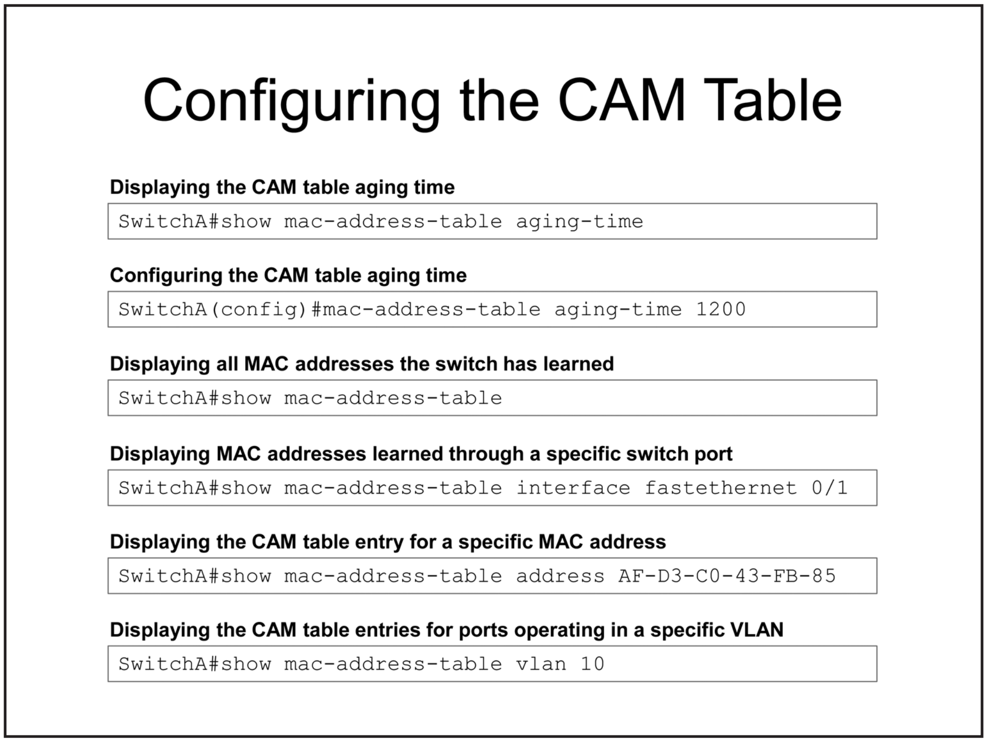

Configuring the CAM Table

Displaying the CAM table aging time

```
SwitchA#show mac-address-table aging-time
```

Configuring the CAM table aging time

```
SwitchA(config)#mac-address-table aging-time 1200
```

Displaying all MAC addresses the switch has learned

```
SwitchA#show mac-address-table
```

Displaying MAC addresses learned through a specific switch port

```
SwitchA#show mac-address-table interface fastethernet 0/1
```

Displaying the CAM table entry for a specific MAC address

```
SwitchA#show mac-address-table address AF-D3-C0-43-FB-85
```

Displaying the CAM table entries for ports operating in a specific VLAN

```
SwitchA#show mac-address-table vlan 10
```

Configuring the CAM Table

You can issue the **show mac-address-table aging-time** command from privileged EXEC mode to examine the value of the CAM aging time. The default aging time depends on the hardware platform. For example, the following sample output from a Catalyst 3750 switch indicates a default aging time of 300 seconds:

```
Switch#show mac-address-table aging-time
Vlan Aging Time
---- ----------
1 300
```

You can issue the **mac-address-table aging-time** command from global configuration mode to modify the CAM table aging time for every VLAN on a switch. An aging time of **0** disables the aging timer. When the aging timer is disabled, CAM table entries become static and are not removed from the table. Alternatively, you can use the **vlan** keyword to modify the aging time for a specific VLAN.

You can use the **show mac-address-table** command from privileged EXEC mode to examine the contents of the CAM table.

```
Switch#show mac-address-table
          Mac Address Table
------------------------------------------
Vlan    Mac Address       Type       Ports
----    -----------       ----       -----
   1    000C.3962.6232    DYNAMIC    Fa0/1
   1    000C.5988.4462    DYNAMIC    Fa0/2

Total Mac Addresses for this criterion: 2
```

Alternatively, you can display specific information from the CAM table. For example, you can issue the **show mac-address-table interface** command to display MAC addresses learned on a specified interface or the **show mac-address-table vlan** command to display MAC addresses learned on ports within a specified VLAN. If you are looking for information relating to a specific MAC address, you can issue the **show mac-address-table address** command to query the CAM table for an entry matching a specific MAC address.

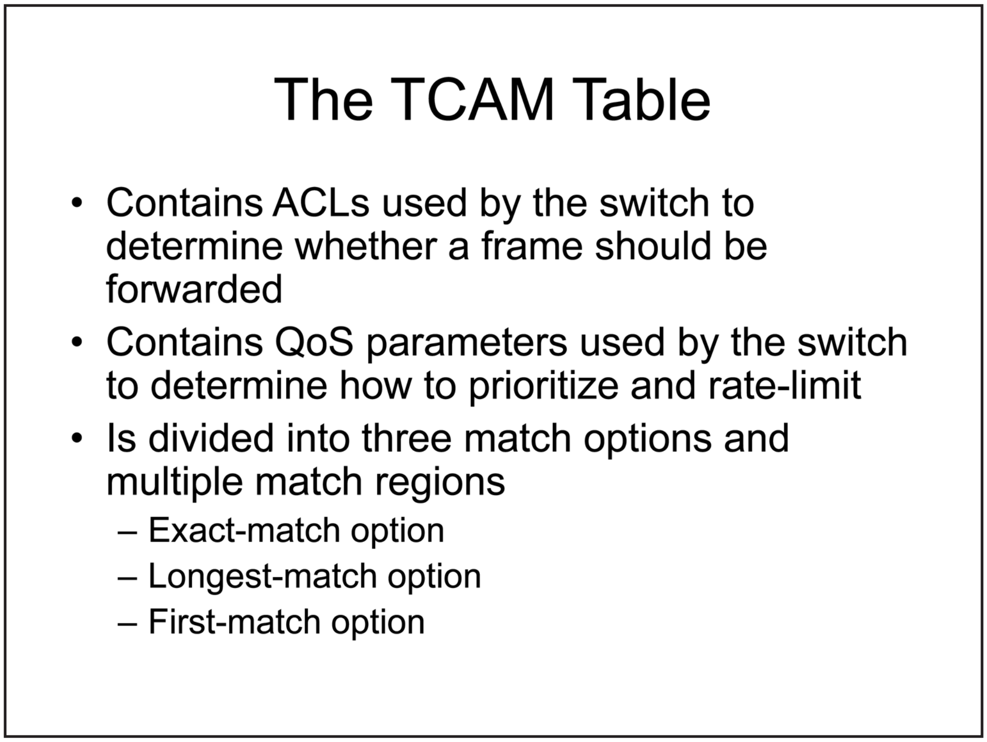

The TCAM Table

- Contains ACLs used by the switch to determine whether a frame should be forwarded
- Contains QoS parameters used by the switch to determine how to prioritize and rate-limit
- Is divided into three match options and multiple match regions
  - Exact-match option
  - Longest-match option
  - First-match option

The TCAM Table

Cisco switches use TCAM tables to facilitate the processing of ACLs and QoS parameters. Like the CAM table, the TCAM table uses fast, system memory to mitigate delay caused by table queries. In addition, because Cisco switches support multiple TCAM tables, queries can be executed in parallel to further mitigate query-related delays to frame forwarding.

Unlike the CAM table, which has fixed-length fields and uses exact-match queries, the TCAM table space is divided into application and protocol-centric regions. These regions enable the TCAM table space to store variable-length fields that are optimized for the data being stored. TCAM table regions also facilitate additional match options, such as longest-match and first-match queries.

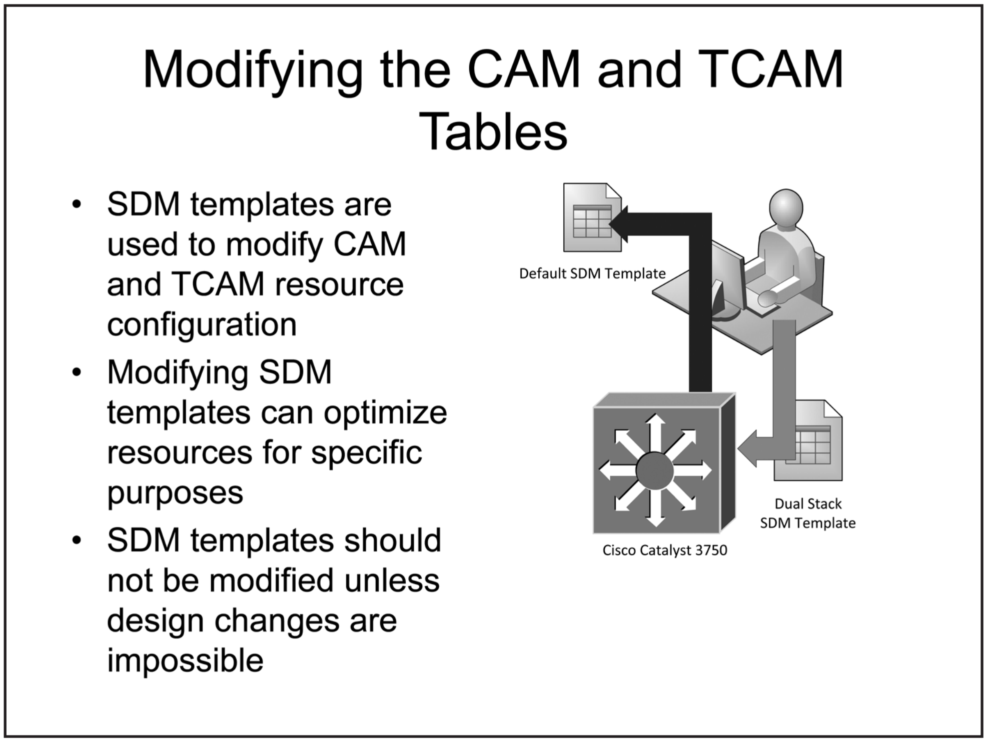

Modifying the CAM and TCAM Tables

- SDM templates are used to modify CAM and TCAM resource configuration
- Modifying SDM templates can optimize resources for specific purposes
- SDM templates should not be modified unless design changes are impossible

Default SDM Template
Dual Stack
SDM Template
Cisco Catalyst 3750

Modifying the CAM and TCAM Tables

On workgroup class switches, like the Catalyst 3700 series, CAM and TCAM table parameters cannot be directly modified. Instead, these switches provide SDM templates to reallocate system resources in predefined configurations. SDM templates are designed to optimize the allocation of system resources for common workgroup switch use cases. Cisco recommends exploring design changes before considering the use of SDM templates to reconfigure CAM and TCAM resources beyond their defaults.

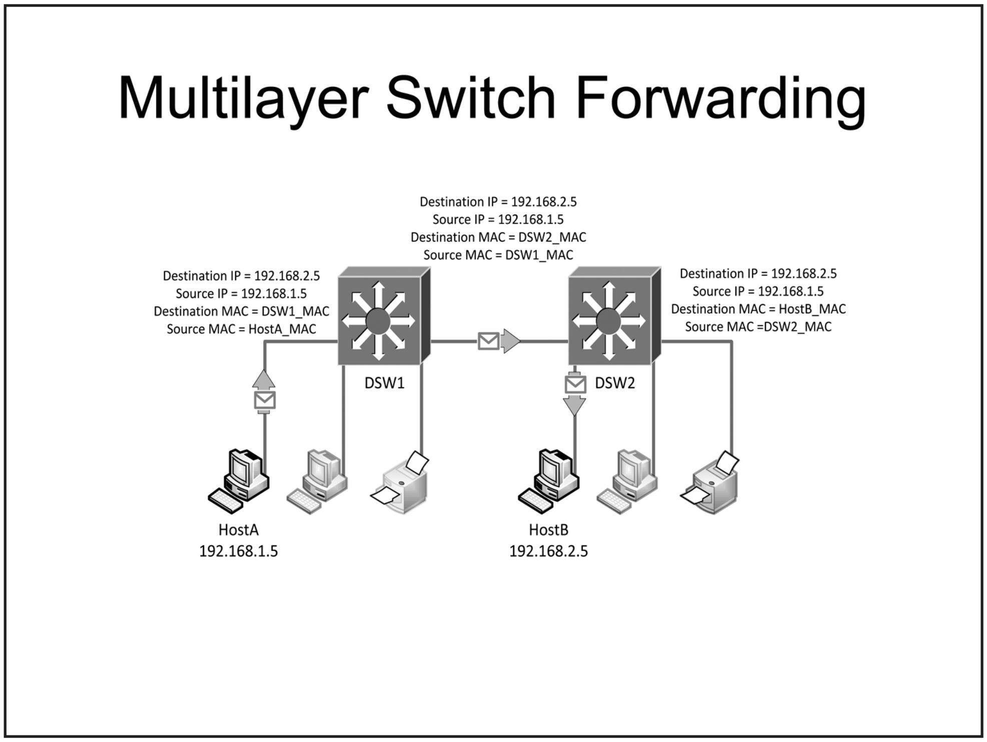

Multilayer Switch Forwarding

Destination IP = 192.168.2.5
Source IP = 192.168.1.5
Destination MAC = DSW2_MAC
Source MAC = DSW1_MAC

Destination IP = 192.168.2.5
Source IP = 192.168.1.5
Destination MAC = DSW1_MAC
Source MAC = HostA_MAC

Destination IP = 192.168.2.5
Source IP = 192.168.1.5
Destination MAC = HostB_MAC
Source MAC =DSW2_MAC

Multilayer Switch Forwarding

A multilayer switch forwards network frames in a manner similar to the way a Layer 2 switch processes frames if those frames are destined to MAC addresses that do not belong to the switch itself. However, when a frame is received with a MAC address that belongs to the switch, then the switch will use its Routing Information Base (RIB) to make a Layer 3 forwarding decision based on the header fields of the packet encapsulated by the received frame.

For example, if HostA needs to send a packet to HostB, HostA will construct an IP packet with its own IP address as the source and the IP address of HostB as the destination. However, because HostB is not in the same IP subnet, HostA must send that packet to its default gateway. HostA constructs a network frame to encapsulate the IP packet. The network frame will use the MAC address of HostA as the source and the MAC address of the default gateway, DSW1, as the destination. HostA then transmits the frame to DSW1, which is a Layer 3 switch. Because the frame is addressed to DSW1, the switch will de-encapsulate the IP packet and then examine its destination IP address. DSW1 will consult its RIB to make a forwarding decision. In this scenario, DSW1 will forward the packet to DSW2. DSW1 will construct a network frame with its own MAC address as the source and the MAC address of DSW2 as the destination and will encapsulate the IP packet within that frame before forwarding the frame to DSW2.

When DSW2 receives the frame, it will follow the same procedure that DSW1 followed. Because the frame is addressed to DSW2, it will de-encapsulate the IP packet and then examine its destination IP address. DSW2 will consult its RIB to make a forwarding decision. HostB is directly connected to DSW2. Therefore, DSW1 can encapsulate the packet in a new frame with its own MAC address as the source and the MAC address of HostB as the destination. DSW2 can then forward the frame directly to HostB.

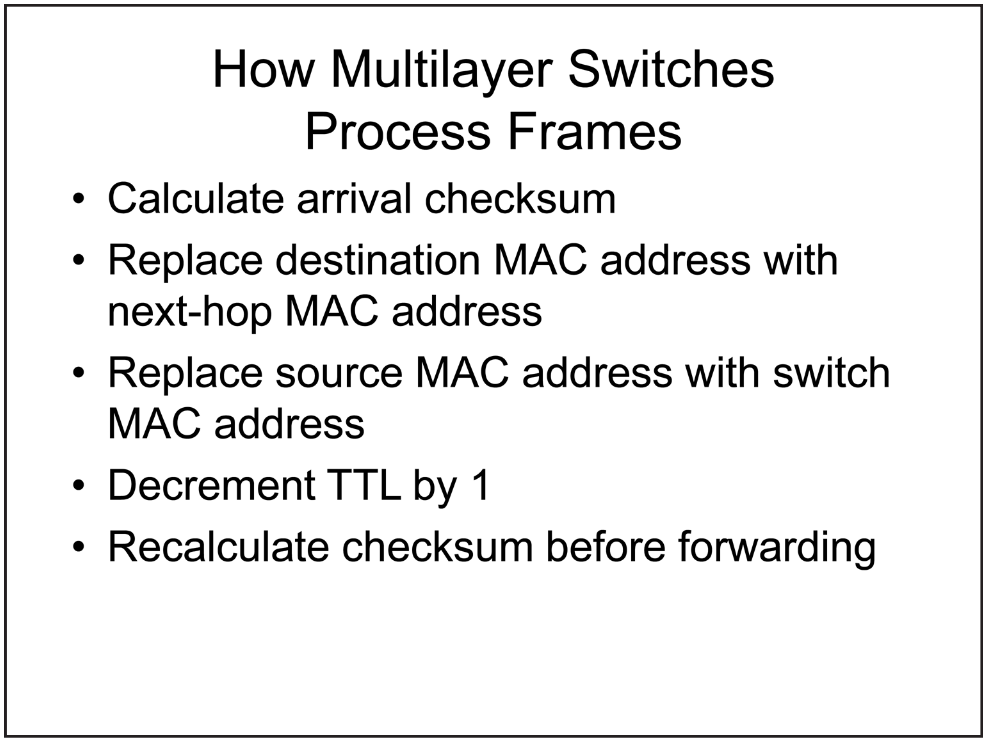

How Multilayer Switches Process Frames

- Calculate arrival checksum
- Replace destination MAC address with next-hop MAC address
- Replace source MAC address with switch MAC address
- Decrement TTL by 1
- Recalculate checksum before forwarding

How Multilayer Switches Process Frames

When a multilayer switch processes frames, whether at Layer 2 or at Layer 3, it must examine the integrity of the frame before making a forwarding decision. Network frames typically have some form of integrity check to ensure that they arrive uncorrupted after transmission. Ethernet frames use a checksum value, which is stored in the frame trailer. The checksum value is calculated using several input values, including the source MAC address and destination MAC address. If either the source MAC address or destination MAC address must be changed in the frame header before the frame is forwarded, the checksum must be recalculated. Changing frame header or trailer fields is commonly referred to as frame rewriting and is necessary when a multilayer switch processes frames that require Layer 3 forwarding.

On a multilayer switch, the Layer 3 forwarding process follows a series of simple steps. These steps are typically implemented in Application-Specific Integrated Circuits (ASICs) to minimize the processing delay. First, the Layer 2 checksum is calculated for the received frame and the value is compared to the checksum field in the frame trailer. If these values match, the frame is considered intact and the Layer 3 checksum can be calculated. Like the Layer 2 checksum, the Layer 3 checksum is calculated using the source and destination addresses in the header. In the case of IP packets, the Layer 3 checksum calculation includes the source and destination IP addresses in the packet header. If the calculated checksum value is equal to the checksum value in the packet header, the switch considers the packet intact and uses the packet header information to make a forwarding decision.

Once a forwarding decision is made, the multilayer switch performs a frame rewrite. The switch decrements the packet time-to-live (TTL) value by one and then encapsulates the packet in a new network frame. The new frame header uses the MAC address of the switch as the source MAC address and the MAC address of the next hop as the destination MAC address. The switch then recalculates the frame checksum based on the new addressing and places this value in the checksum field in the frame trailer. The frame is then placed in the transmit queue of the appropriate egress port and sent to the next hop.

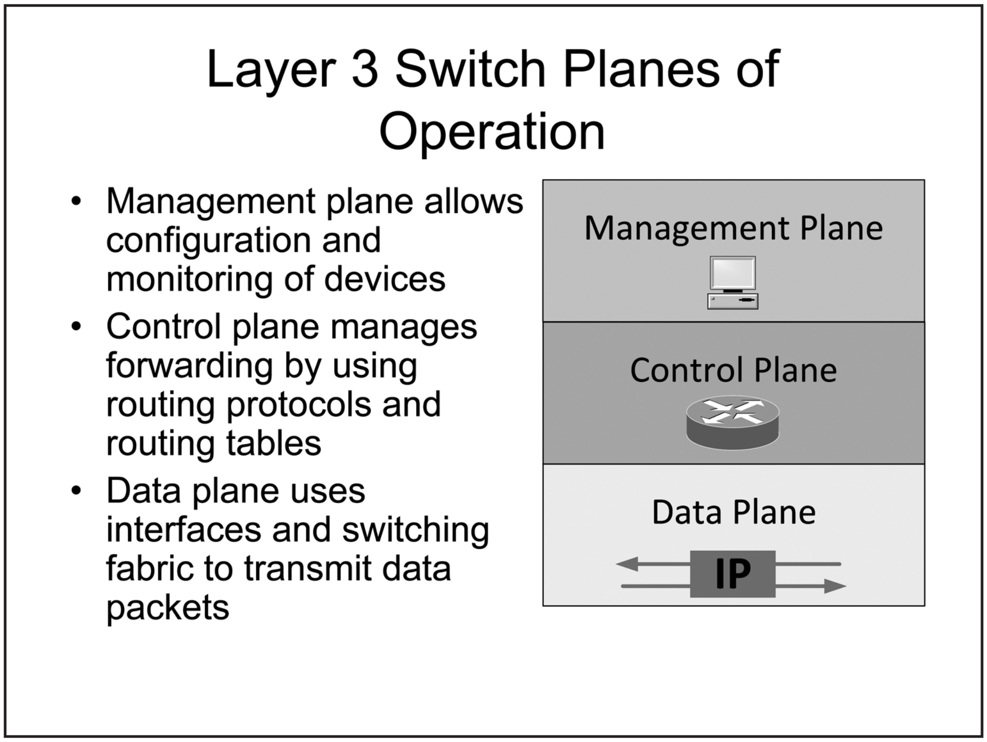

Layer 3 Switch Planes of Operation

- Management plane allows configuration and monitoring of devices
- Control plane manages forwarding by using routing protocols and routing tables
- Data plane uses interfaces and switching fabric to transmit data packets

Layer 3 Switch Planes of Operation

Cisco routers and switches are divided into three logical planes of operation:

- The management plane
- The control plane
- The data plane

The management plane is responsible for processing packets that are destined to the Layer 3 address of the device, such as Secure Shell (SSH), Simple Network Management Protocol version 3 (SNMPv3), and Syslog traffic.

The control plane is responsible for the creation and maintenance of structures related to routing and forwarding. These functions are heavily dependent on the CPU and memory availability. Control plane traffic is generated by dynamic routing protocols, VTP, and Spanning Tree Protocol (STP).

The data plane, which is also referred to as the forwarding plane, is responsible for traffic passing through the device. Traffic from the data plane consists primarily of user-generated traffic that is forwarded from one interface to another; this type of traffic is also referred to as transit traffic.

In most network designs, data plane traffic makes up the majority of network traffic during normal operation; therefore, router and switch hardware, such as the switching fabric, route processor modules, and ASICs, is optimized for the processing of data plane traffic.

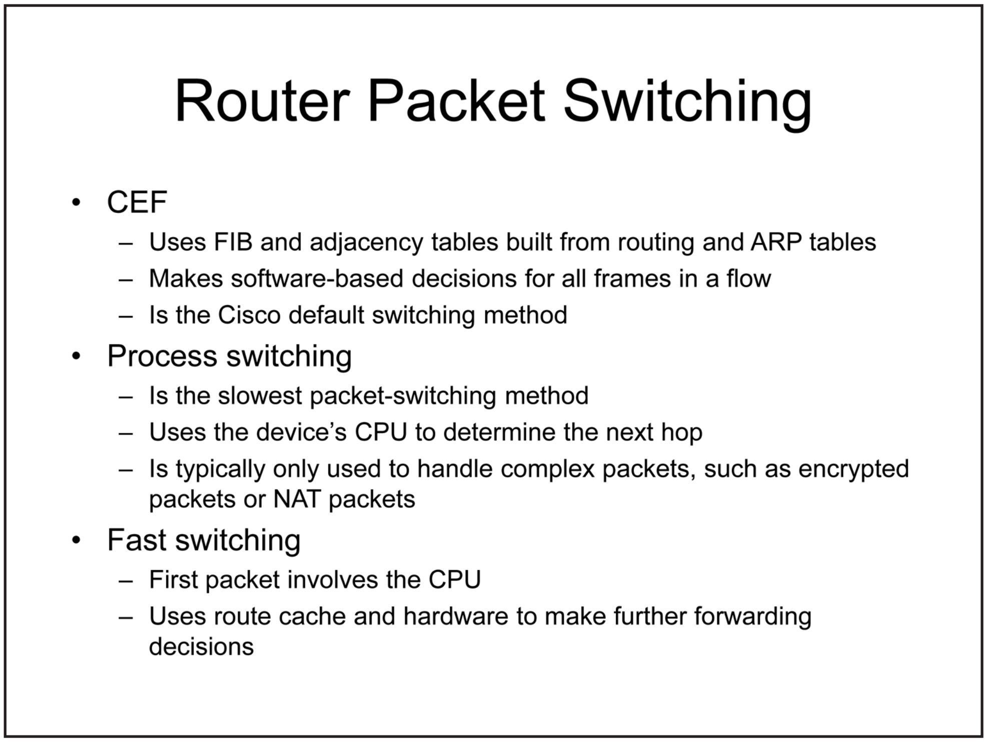

Router Packet Switching

- CEF
  - Uses FIB and adjacency tables built from routing and ARP tables
  - Makes software-based decisions for all frames in a flow
  - Is the Cisco default switching method
- Process switching
  - Is the slowest packet-switching method
  - Uses the device's CPU to determine the next hop
  - Is typically only used to handle complex packets, such as encrypted packets or NAT packets
- Fast switching
  - First packet involves the CPU
  - Uses route cache and hardware to make further forwarding decisions

Router Packet Switching

Cisco Express Forwarding (CEF) switching is the default switching method on supported Cisco routers. CEF switching is faster than other packet-switching methodologies, such as fast switching or process switching. CEF relies on a FIB table to make fast routing decisions. The FIB is built from information in the IP routing table and the Address Resolution Protocol (ARP) table, which is comparable to the CAM table on a switch. With CEF, forwarding data exists even before frames arrive for a particular data flow. Therefore, routers can perform the entire forwarding process in hardware without requiring the involvement of the CPU. When a router receives a CEF-switched packet, it consults the FIB and forwards the packet based on the destination address. However, if the packet's destination address is not in the FIB, the packet is dropped.

Process switching is the slowest of the three switching methods. Routers that perform process switching on a packet must rely on the router CPU to consult the routing table, create a new Layer 2 frame header, and forward the packet. Even on routers supporting faster mechanisms, process switching might be used to process a select number of packets. For example, encrypted packets or packets requiring address translation are typically handled by process switching.

Fast switching uses a switching cache that contains the most recently used destination lookups. If a packet's destination is not in the cache, the packet is process switched, the Layer 2 frame header is stored in the fast switching cache, and subsequent packets to that destination are fast switched. Only the first packet of an unknown packet flow involves the CPU.

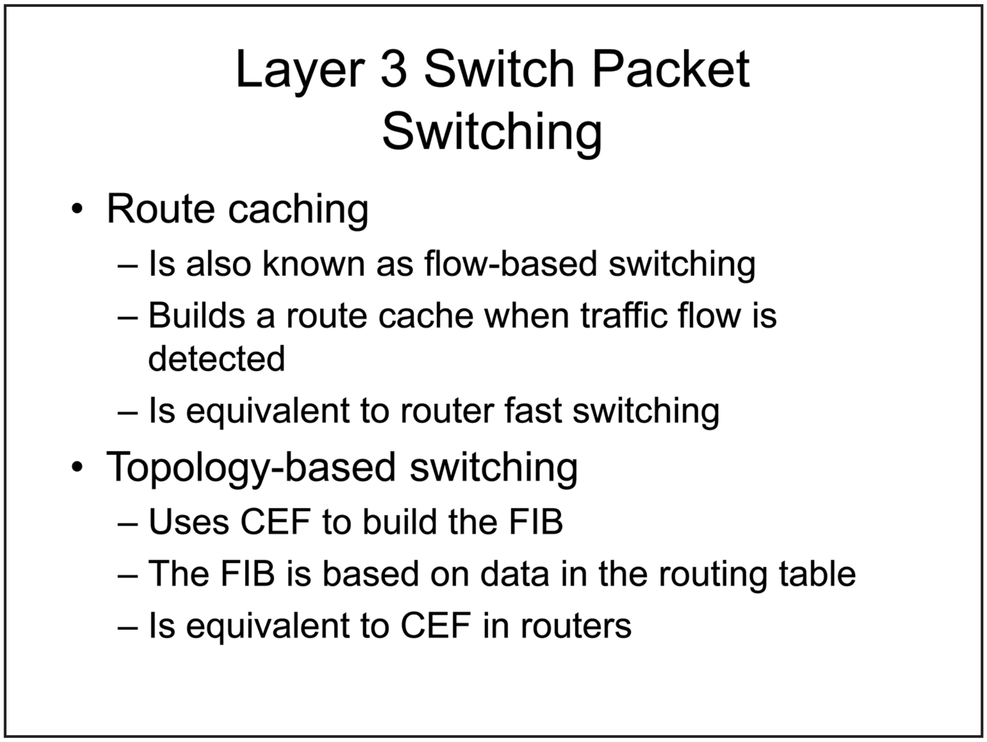

Layer 3 Switch Packet Switching

- Route caching
  - Is also known as flow-based switching
  - Builds a route cache when traffic flow is detected
  - Is equivalent to router fast switching
- Topology-based switching
  - Uses CEF to build the FIB
  - The FIB is based on data in the routing table
  - Is equivalent to CEF in routers

Layer 3 Switch Packet Switching

Like Cisco routers, Cisco multilayer switches support advanced forwarding methodologies to facilitate high forwarding rates in common network deployments. Because switches typically contain a limited number of network interface types and support a limited number of higher-level protocols, packet-switching processes can be more optimized than they are on routers, which support a greater variety of interfaces and protocols. For this reason, multilayer switches do not use packet switching to switch packets. Instead, CPU involvement is minimized and the majority of processing is handled in hardware by ASICs.

Cisco multilayer switches support the following packet-switching methodologies:

- Route caching
- Topology-based switching

Route caching, which is also referred to as flow-based switching, is analogous to fast switching on a Cisco router. With route caching, the switch builds a route cache for each traffic flow. Only the initial packet for each traffic flow requires CPU involvement. Once a route cache entry exists for a traffic flow, all subsequent packets are processed in hardware.

Likewise, topology-based switching is analogous to CEF on a Cisco router. With topology-based switching, the switch builds a FIB using information from the routing table and the CAM table. Because the switch can populate the FIB without the involvement of the CPU or the presence of any data flows, all packet switching can be performed in hardware.

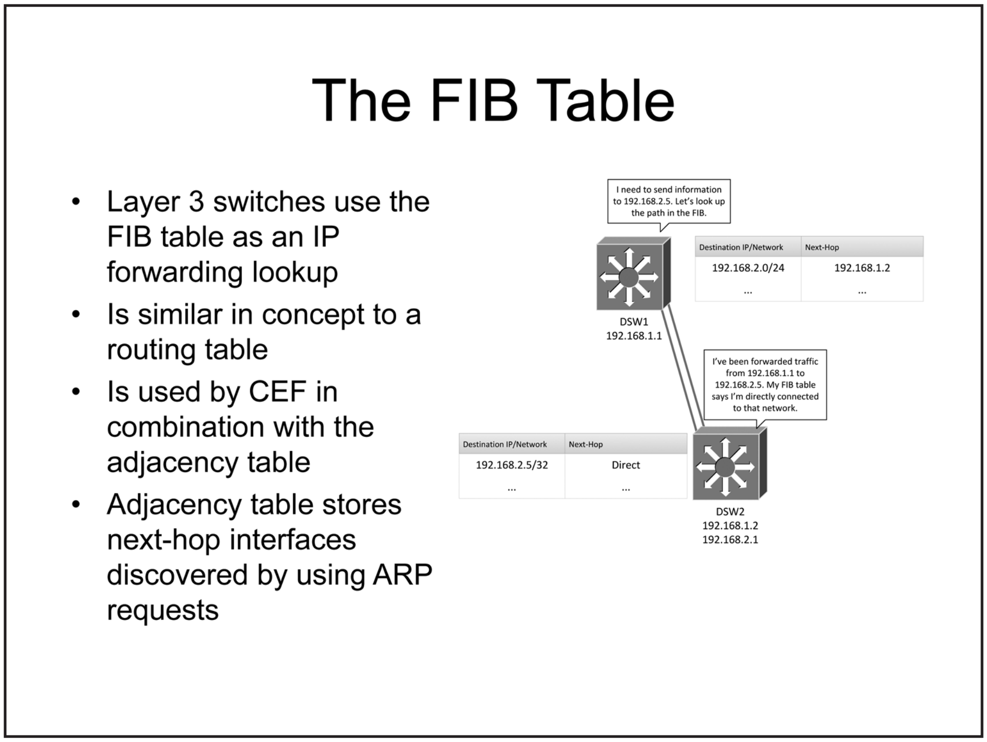

The FIB Table

- Layer 3 switches use the FIB table as an IP forwarding lookup
- Is similar in concept to a routing table
- Is used by CEF in combination with the adjacency table
- Adjacency table stores next-hop interfaces discovered by using ARP requests

I need to send information to 192.168.2.5. Let's look up the path in the FIB.
DSW1
192.168.1.1
I've been forwarded traffic from 192.168.1.1 to 192.168.2.5. My FIB table says I'm directly connected to that network.
DSW2
192.168.1.2
192.168.2.1

Destination IP/Network	Next-Hop
192.168.2.0/24	192.168.1.2
...	...

Destination IP/Network	Next-Hop
192.168.2.5/32	Direct
...	...

The FIB Table

Cisco Layer 3 switches use the FIB table by default when making forwarding decisions. The FIB table is similar to the IP routing table in that it contains all the IP prefixes of which the switch is aware; these prefixes are known to the switch either through static configuration or from dynamic routing protocols. The FIB remains synchronized with the IP routing table during topology changes. However, unlike the routing table, the FIB contains additional information that has been gleaned from ARP traffic or from the CAM table. The adjacency table contains the Layer 2 addressing information for each next-hop address listed in the FIB. An adjacency is a device that is accessible via a single Layer 2 hop.

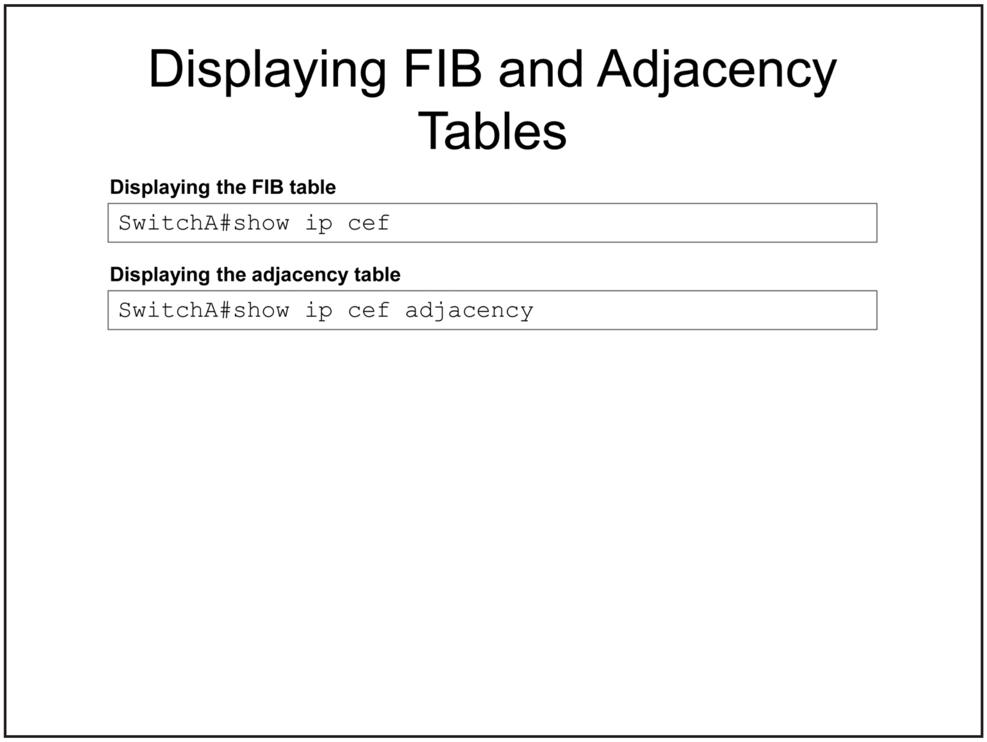

Displaying FIB and Adjacency Tables

Displaying the FIB table

```
SwitchA#show ip cef
```

Displaying the adjacency table

```
SwitchA#show ip cef adjacency
```

Displaying FIB and Adjacency Tables

A FIB entry comprises an IP network prefix, a next-hop address, and the outgoing interface used to reach the next-hop address. You can issue the **show ip cef** command from privileged EXEC mode to examine the contents of the FIB. The following example output from the **show ip cef 10.1.2.0/24** command indicates that an entry for the 10.1.2.0/24 network exists in the FIB:

```
Prefix               Next Hop             Interface
10.1.2.0/24           192.168.2.1           FastEthernet0/22
```

It is important to note that CEF discards packets that do not match an entry in the FIB. With Layer 2 processing, a switch will forward a frame out every interface if no match is found in the CAM table. However, with Layer 3 processing, a switch must have a matching entry in the FIB to forward a packet.

You can issue the **show ip cef adjacency** command to examine details about a specific adjacency or about a specific type of adjacency. For example, you could issue the **show ip cef adjacency punt** command to view adjacencies that cannot be processed directly in hardware and require CPU involvement for processing.

Module Notes

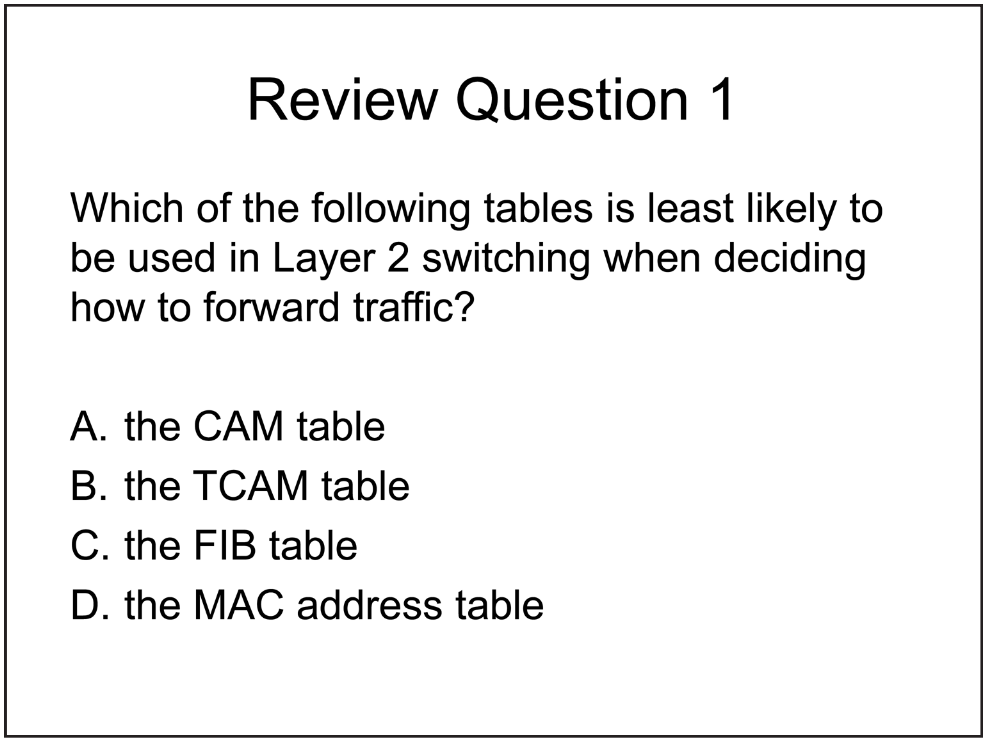

Review Question 1

Which of the following tables is least likely to be used in Layer 2 switching when deciding how to forward traffic?

A. the CAM table
B. the TCAM table
C. the FIB table
D. the MAC address table

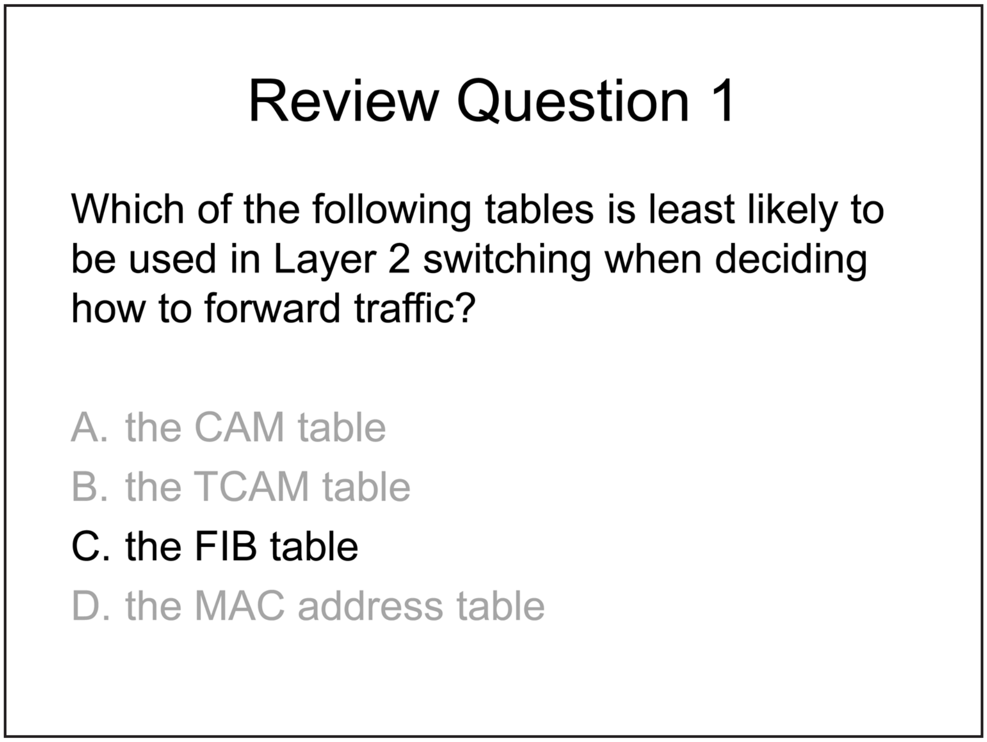

Review Question 1

Which of the following tables is least likely to be used in Layer 2 switching when deciding how to forward traffic?

A. the CAM table
B. the TCAM table
C. the FIB table
D. the MAC address table

The Forwarding Information Base (FIB) table is least likely to be used in Layer 2 switching when deciding how to forward traffic. In Layer 2 switching, a switch examines the destination Media Access Control (MAC) address in the header of a network frame and consults its content-addressable memory (CAM) table to make a forwarding decision. The CAM table is also referred to as the MAC address table.

Layer 2 switches also contain ternary CAM (TCAM) tables. TCAM tables are used to efficiently perform lookups for access control lists (ACLs), Quality of Service (QoS) parameters, and other parameters that may affect Layer 2 forwarding decisions.

Cisco Layer 3 switches, which are also referred to as multilayer switches, use the FIB to make Layer 3 forwarding decisions. The FIB combines information from the routing table with Layer 2 information gleaned from Address Resolution Protocol (ARP) traffic and the CAM table.

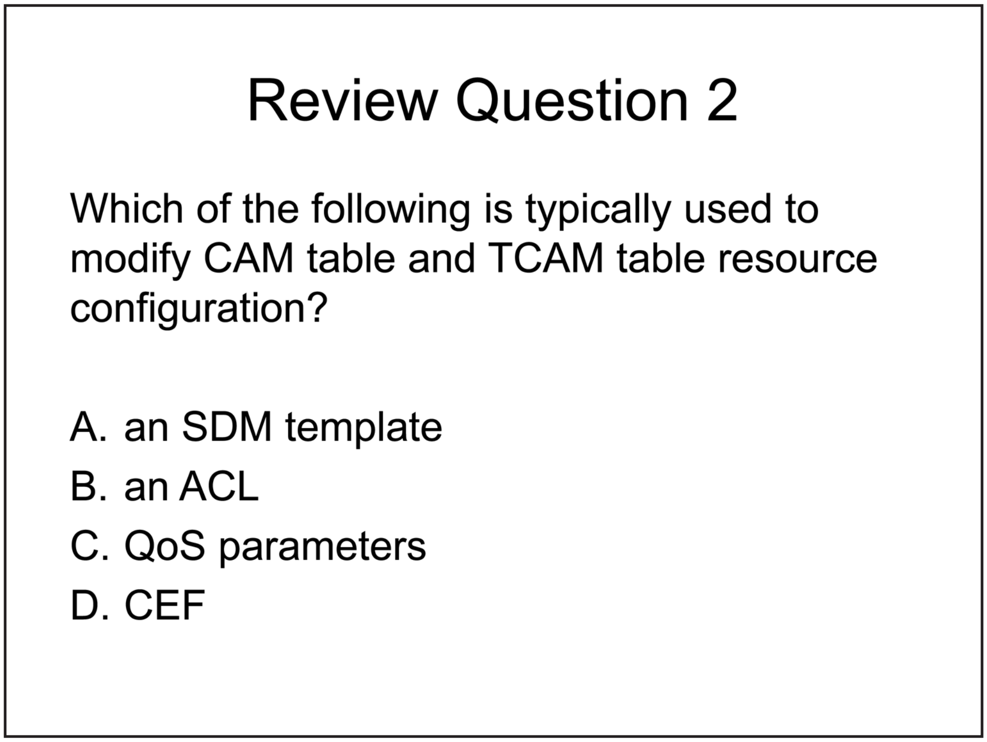

Review Question 2

Which of the following is typically used to modify CAM table and TCAM table resource configuration?

A. an SDM template
B. an ACL
C. QoS parameters
D. CEF

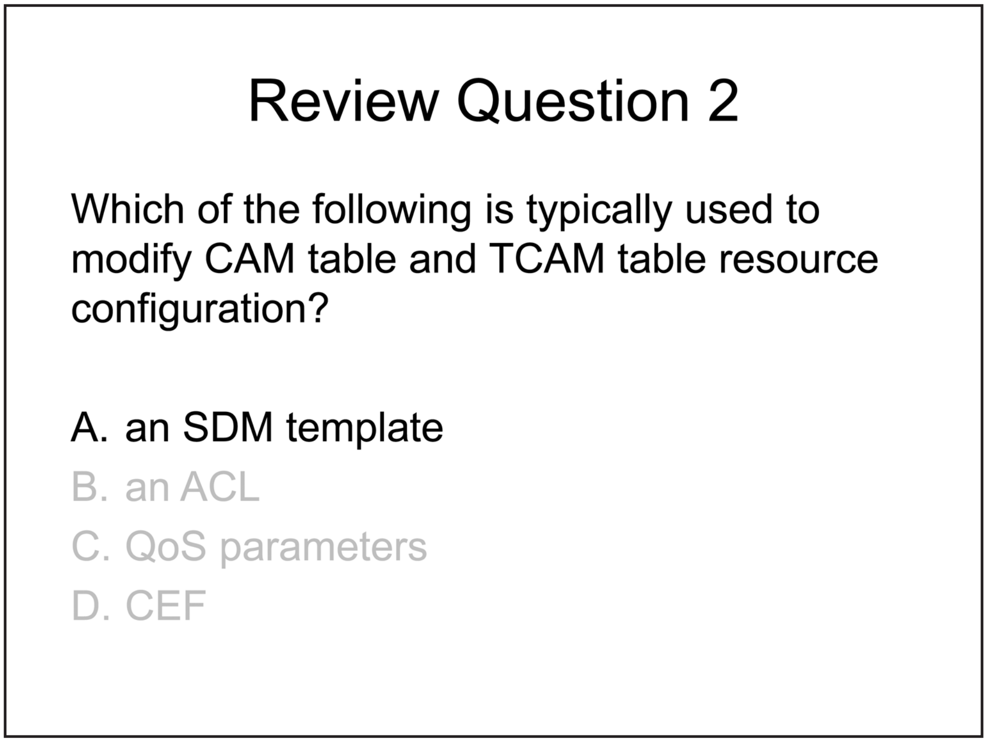

Review Question 2

Which of the following is typically used to modify CAM table and TCAM table resource configuration?

A. an SDM template
B. an ACL
C. QoS parameters
D. CEF

A Switch Database Management (SDM) template is typically used to modify content-addressable memory (CAM) table and ternary CAM (TCAM) table resource configuration. The CAM and TCAM tables are allocated a portion of system memory to store information such as Media Access Control (MAC) addresses, access control list (ACL) parameters, and Quality of Service (QoS) parameters. On Cisco workgroup switches, such as the Catalyst 3750, the amount of memory allocated to the CAM and TCAM tables cannot be directly modified. Instead, Cisco workgroup switches provide preconfigured resource allocation profiles called SDM templates to facilitate the reconfiguration of system memory for a specific type network design.

Cisco Express Forwarding (CEF) is not used to modify CAM table and TCAM table resource configuration. CEF is a packet-switching mechanism that is used by default on Cisco multilayer switches.

Neither QoS parameters nor an ACL are used to modify CAM and TCAM resource configuration. QoS parameters and ACL parameters are stored in TCAM tables to expedite processing and mitigate delays from table queries.

Lab Exercises

Module 2: Switch Basics

Lab 2.1 – IOS Switching Initial Configuration

Labs powered by

NetSim® NETWORK SIMULATOR®

The labs referenced in this book have been printed in the Boson Lab Guide, which is available for purchase. To learn more about the Boson NetSim or to purchase and download the software, please visit www.boson.com/netsim-cisco-network-simulator.

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