



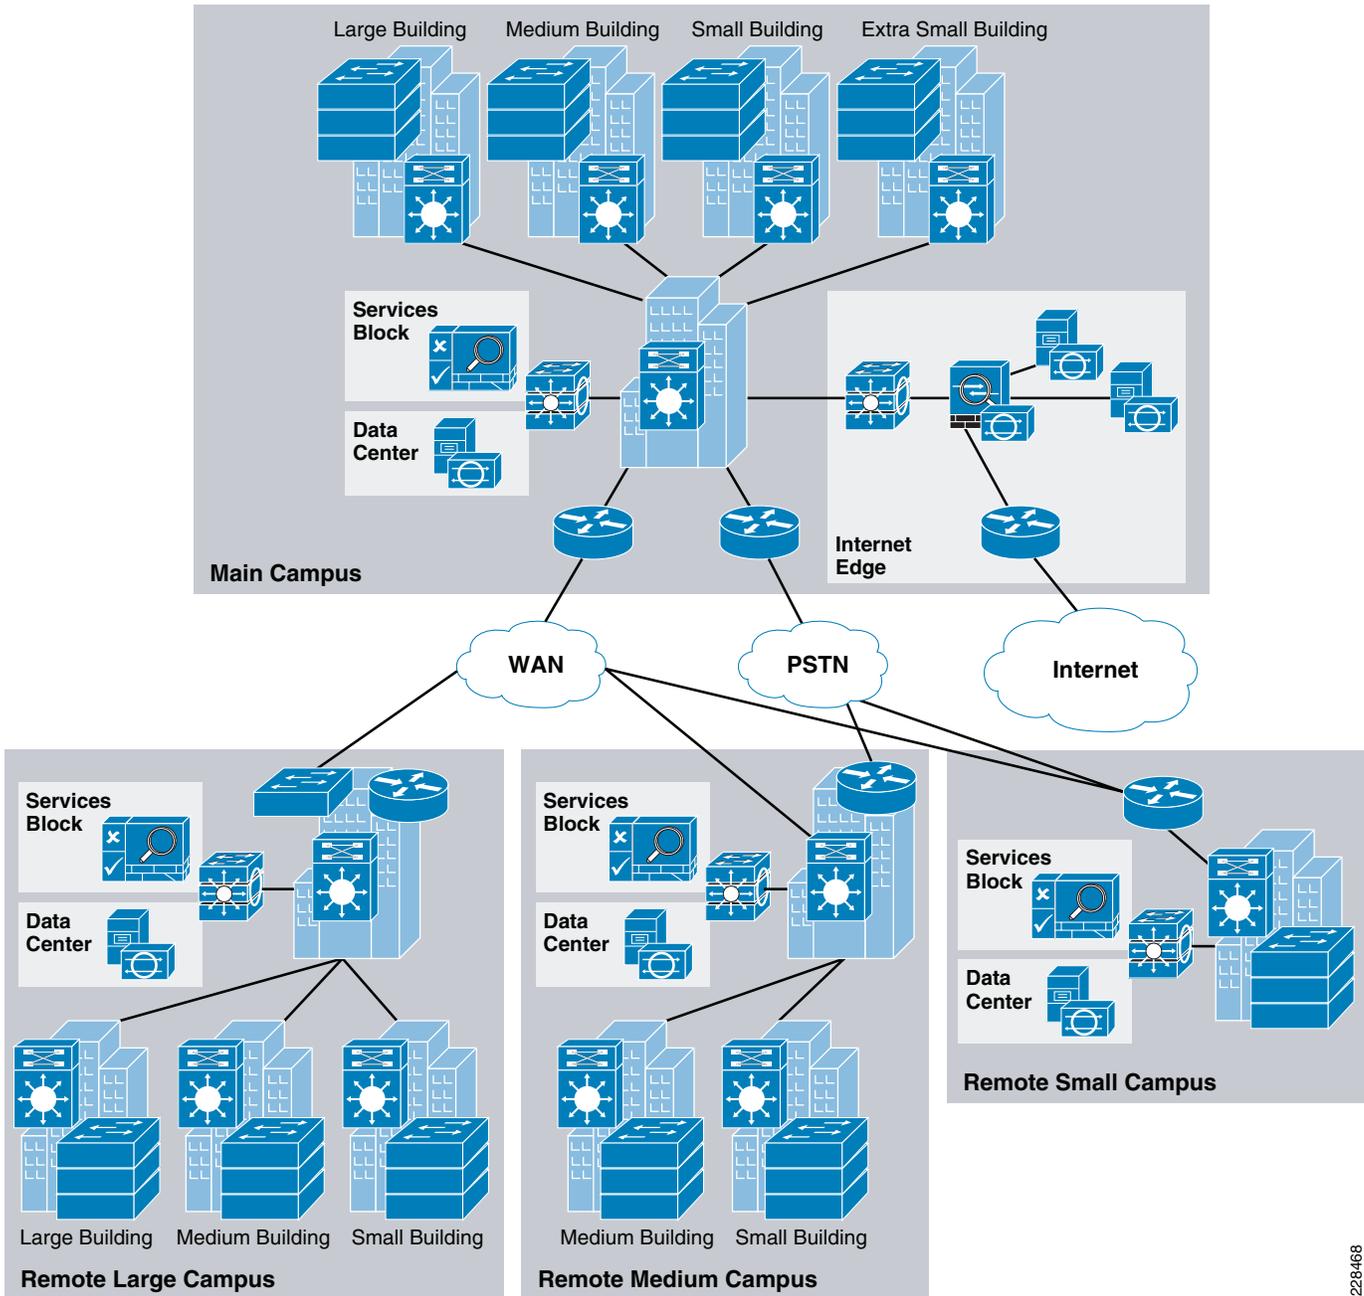
CHAPTER 3

Community College LAN Design Considerations

LAN Design

The community college LAN design is a multi-campus design, where a campus consists of multiple buildings and services at each location, as shown in [Figure 3-1](#).

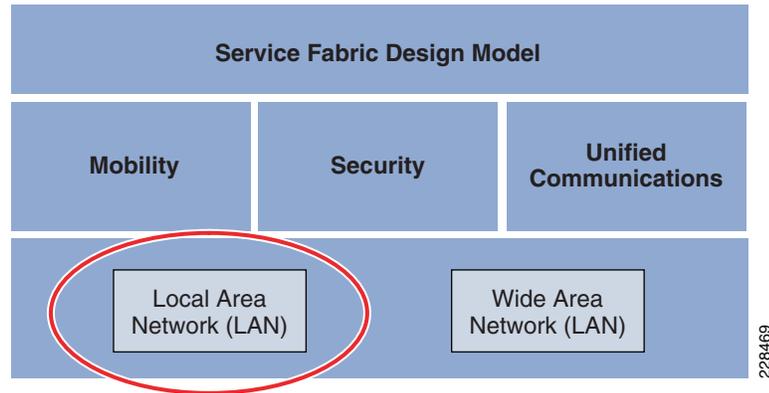
Figure 3-1 Community College LAN Design



228468

Figure 3-2 shows the service fabric design model used in the community college LAN design.

Figure 3-2 Community College LAN Design



This chapter focuses on the LAN component of the overall design. The LAN component consists of the LAN framework and network foundation technologies that provide baseline routing and switching guidelines. The LAN design interconnects several other components, such as endpoints, data center, WAN, and so on, to provide a foundation on which mobility, security, and unified communications (UC) can be integrated into the overall design.

This LAN design provides guidance on building the next-generation community college network, which becomes a common framework along with critical network technologies to deliver the foundation for the service fabric design. This chapter is divided into following sections:

- *LAN design principles*—Provides proven design choices to build various types of LANs.
- *LAN design model for the community college*—Leverages the design principles of the tiered network design to facilitate a geographically dispersed college campus network made up of various elements, including networking role, size, capacity, and infrastructure demands.
- *Considerations of a multi-tier LAN design model for community colleges*—Provides guidance for the college campus LAN network as a platform with a wide range of next-generation products and technologies to integrate applications and solutions seamlessly.
- *Designing network foundation services for LAN designs in community colleges*—Provides guidance on deploying various types of Cisco IOS technologies to build a simplified and highly available network design to provide continuous network operation. This section also provides guidance on designing network-differentiated services that can be used to customize the allocation of network resources to improve user experience and application performance, and to protect the network against unmanaged devices and applications.

LAN Design Principles

Any successful design or system is based on a foundation of solid design theory and principles. Designing the LAN component of the overall community college LAN service fabric design model is no different than designing any large networking system. The use of a guiding set of fundamental engineering design principles serves to ensure that the LAN design provides for the balance of availability, security, flexibility, and manageability required to meet current and future college and technology needs. This chapter provides design guidelines that are built upon the following principles to allow a community college network architect to build college campuses that are located in different geographical locations:

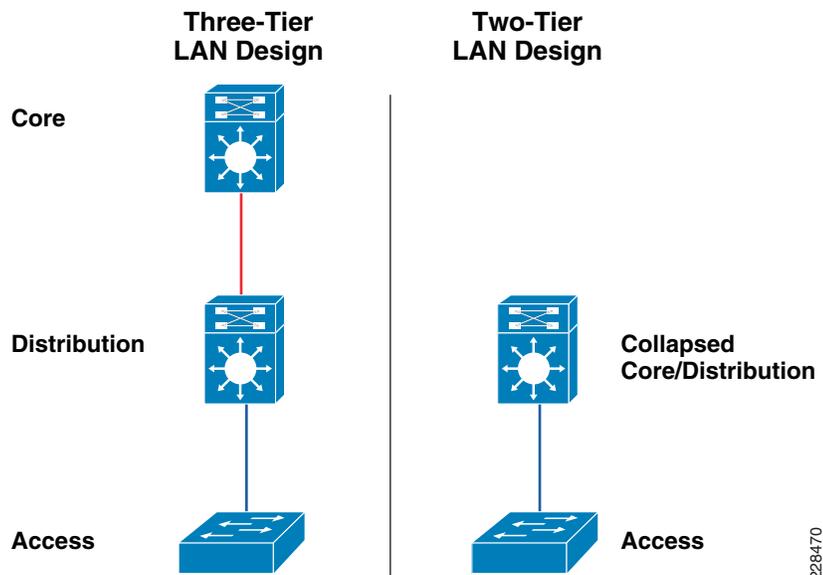
- *Hierarchical*

- Facilitates understanding the role of each device at every tier
- Simplifies deployment, operation, and management
- Reduces fault domains at every tier
- *Modularity*—Allows the network to grow on an on-demand basis
- *Resiliency*—Satisfies user expectations for keeping network always on
- *Flexibility*—Allows intelligent traffic load sharing by using all network resources

These are not independent principles. The successful design and implementation of a college campus network requires an understanding of how each of these principles applies to the overall design. In addition, understanding how each principle fits in the context of the others is critical in delivering a hierarchical, modular, resilient, and flexible network required by community colleges today.

Designing the community college LAN building blocks in a hierarchical fashion creates a flexible and resilient network foundation that allows network architects to overlay the security, mobility, and UC features essential to the service fabric design model, as well as providing an interconnect point for the WAN aspect of the network. The two proven, time-tested hierarchical design frameworks for LAN networks are the three-tier layer and the two-tier layer models, as shown in [Figure 3-3](#).

Figure 3-3 Three-Tier and Two-Tier LAN Design Models



The key layers are access, distribution and core. Each layer can be seen as a well-defined structured module with specific roles and functions in the LAN network. Introducing modularity in the LAN hierarchical design further ensures that the LAN network remains resilient and flexible to provide critical network services as well as to allow for growth and changes that may occur in a community college.

- *Access layer*

The access layer represents the network edge, where traffic enters or exits the campus network. Traditionally, the primary function of an access layer switch is to provide network access to the user. Access layer switches connect to the distribution layer switches to perform network foundation technologies such as routing, quality of service (QoS), and security.

To meet network application and end-user demands, the next-generation Cisco Catalyst switching platforms no longer simply switch packets, but now provide intelligent services to various types of endpoints at the network edge. Building intelligence into access layer switches allows them to operate more efficiently, optimally, and securely.

- *Distribution layer*

The distribution layer interfaces between the access layer and the core layer to provide many key functions, such as the following:

- Aggregating and terminating Layer 2 broadcast domains
- Aggregating Layer 3 routing boundaries
- Providing intelligent switching, routing, and network access policy functions to access the rest of the network
- Providing high availability through redundant distribution layer switches to the end-user and equal cost paths to the core, as well as providing differentiated services to various classes of service applications at the edge of network

- *Core layer*

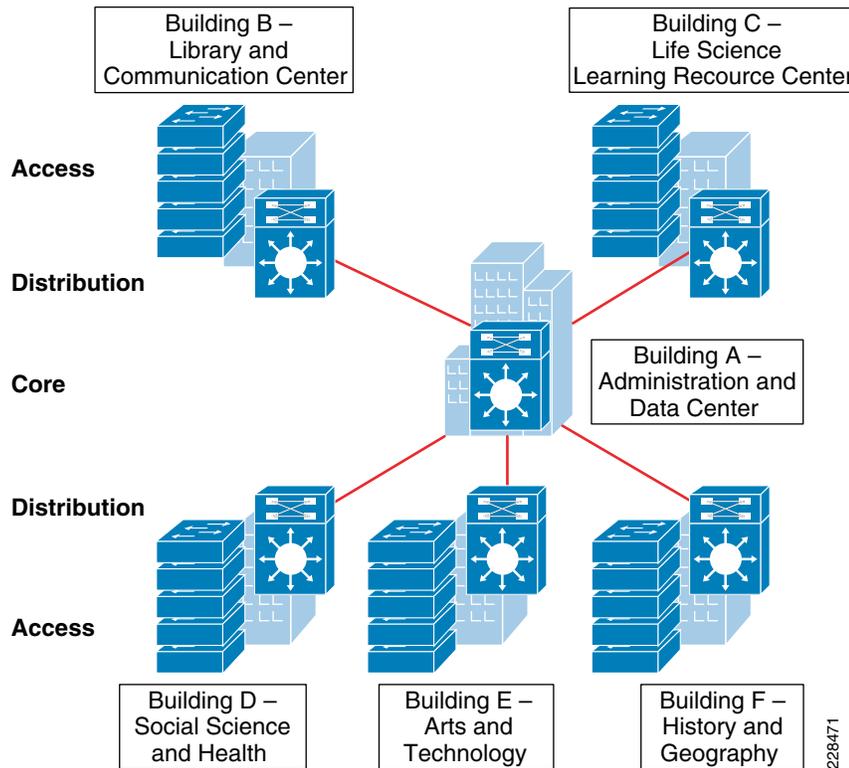
The core layer is the network backbone that connects all the layers of the LAN design, providing for connectivity between end devices, computing and data storage services located within the data center and other areas, and services within the network. The core layer serves as the aggregator for all the other campus blocks, and ties the campus together with the rest of the network.

**Note**

For more information on each of these layers, see the enterprise class network framework at the following URL: <http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/campover.html>.

Figure 3-4 shows a sample three-tier LAN network design for community colleges where the access, distribution, and core are all separate layers. To build a simplified, cost-effective, and efficient physical cable layout design, Cisco recommends building an extended-star physical network topology from a centralized building location to all other buildings on the same campus.

Figure 3-4 Three-Tier LAN Network Design Example



228471

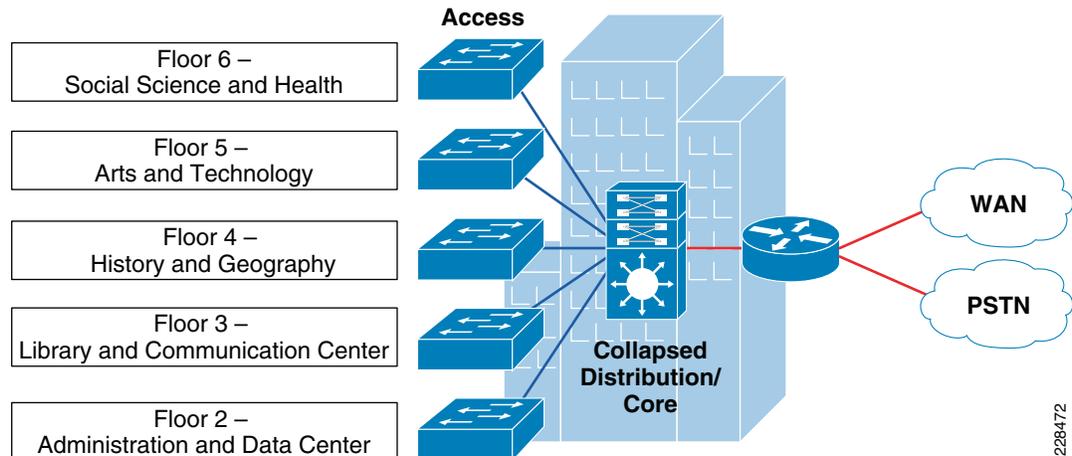
Collapsed Core Campus Network Design

The primary purpose of the core layer is to provide fault isolation and backbone connectivity. Isolating the distribution and core into separate layers creates a clean delineation for change control between activities affecting end stations (laptops, phones, and printers) and those that affect the data center, WAN, or other parts of the network. A core layer also provides for flexibility in adapting the campus design to meet physical cabling and geographical challenges. If necessary, a separate core layer can use a different transport technology, routing protocols, or switching hardware than the rest of the campus, providing for more flexible design options when needed.

In some cases, because of either physical or network scalability, having separate distribution and core layers is not required. In smaller locations where there are less users accessing the network or in college campus sites consisting of a single building, separate core and distribution layers are not needed. In this scenario, Cisco recommends the two-tier LAN network design, also known as the collapsed core network design.

Figure 3-5 shows a two-tier LAN network design example for a community college LAN where the distribution and core layers are collapsed into a single layer.

Figure 3-5 Two-Tier Network Design Example



If using the small-scale collapsed campus core design, the college network architect must understand the network and application demands so that this design ensures a hierarchical, modular, resilient, and flexible LAN network.

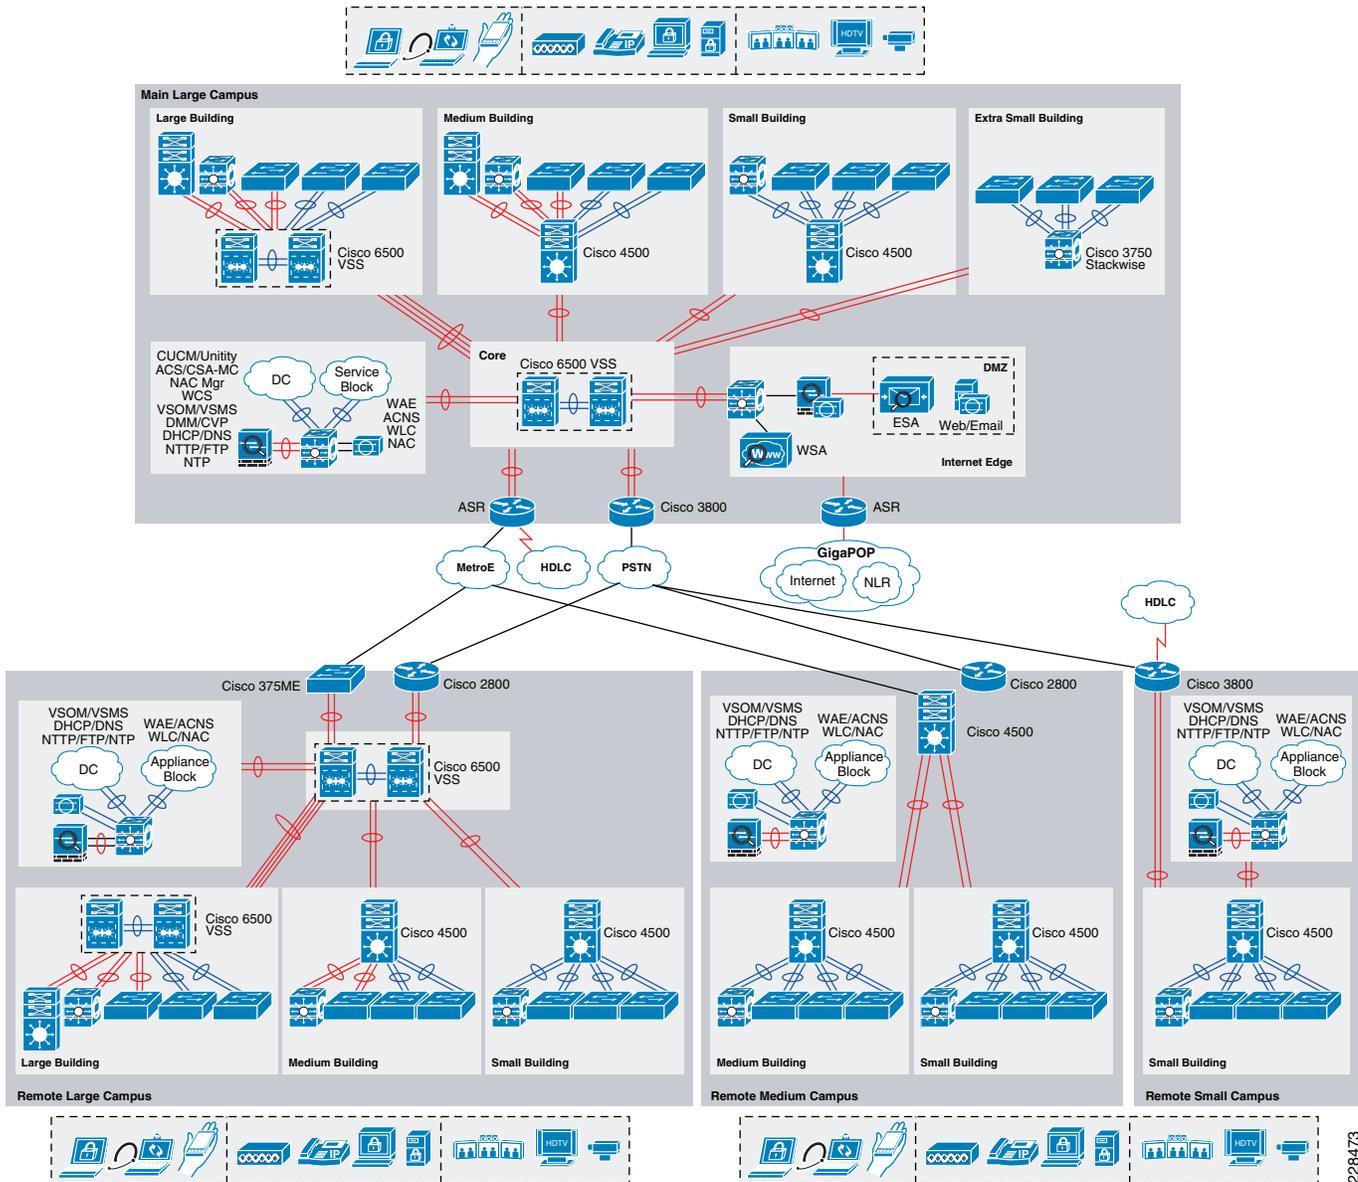
Community College LAN Design Models

Both LAN design models (three-tier and two-tier) have been developed with the following considerations:

- *Scalability*—Based on Cisco enterprise-class high-speed 10G core switching platforms for seamless integration of next-generation applications required for community colleges. Platforms chosen are cost-effective and provide investment protection to upgrade network as demand increases.
- *Simplicity*—Reduced operational and troubleshooting cost via the use of network-wide configuration, operation, and management.
- *Resilient*—Sub-second network recovery during abnormal network failures or even network upgrades.
- *Cost-effectiveness*—Integrated specific network components that fit budgets without compromising performance.

As shown in [Figure 3-6](#), multiple campuses can co-exist within a single community college system that offers various academic programs.

Figure 3-6 Community College LAN Design Model



Depending on the number of available academic programs in a remote campus, the student, faculty, and staff population in remote campuses may be equal to or less than the main college campus site. Campus network designs for the remote campus may require adjusting based on overall college campus capacity.

Using high-speed WAN technology, all the remote community college campuses interconnect to a centralized main college campus that provides shared services to all the students, faculty, and staff, independent of their physical location. The WAN design is discussed in greater detail in the next chapter, but it is worth mentioning in the LAN section because some remote sites may integrate LAN and WAN functionality into a single platform. Collapsing the LAN and WAN functionality into a single Cisco platform can provide all the needed requirements for a particular remote site as well as provide reduced cost to the overall design, as discussed in more detail in the following section.

Table 3-1 shows a summary of the LAN design models as they are applied in the overall community college network design.

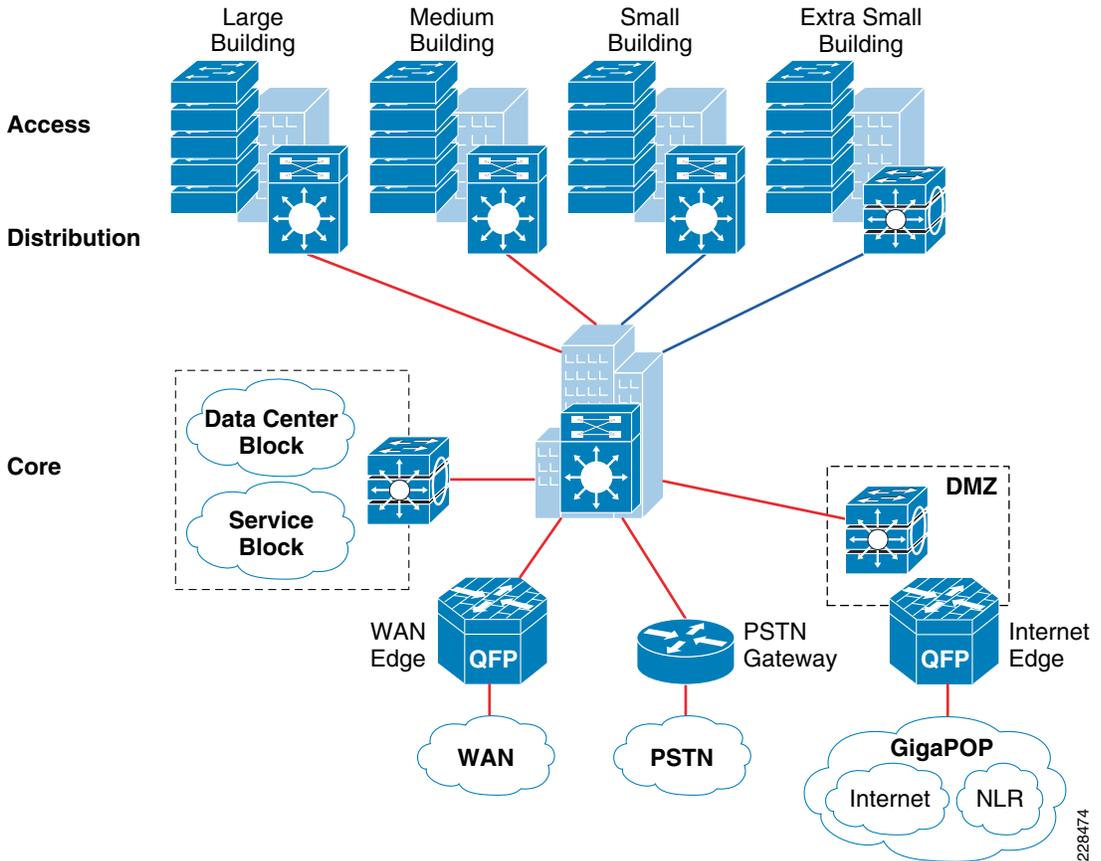
Table 3-1 Community College Recommended LAN Design Model

Community College Location	Recommended LAN Design Model
Main campus	Three-tier
Remote large campus	Three-tier
Remote medium campus	Three-tier with collapsed WAN edge
Remote small campus	Two-tier

Main College Campus Network Design Overview

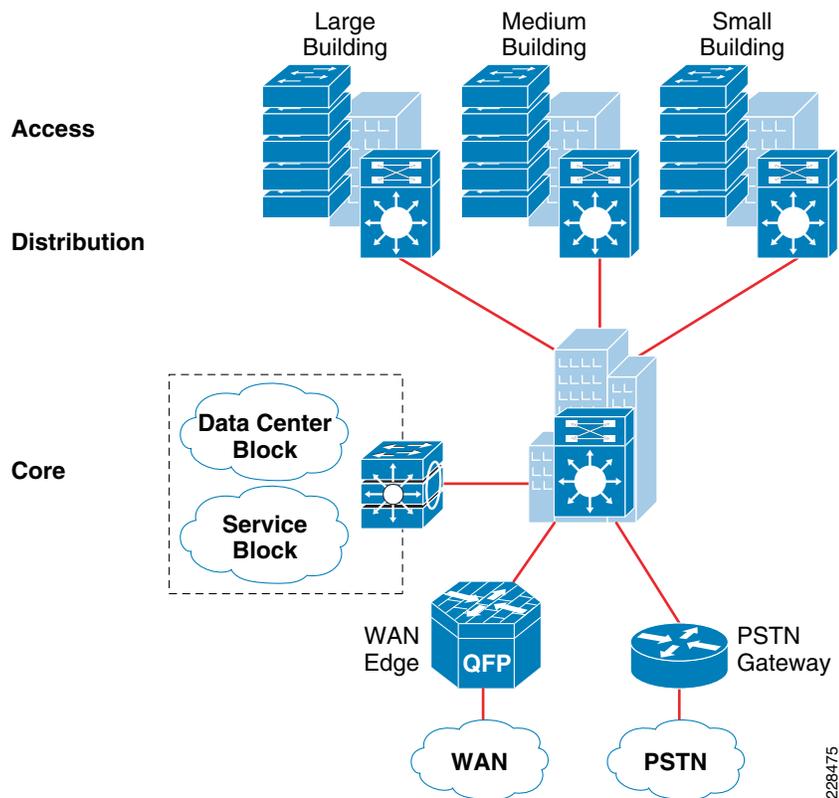
The main college campus in the community college design consists of a centralized hub campus location that interconnects several sizes of remote campuses to provide end-to-end shared network access and services, as shown in Figure 3-7.

Figure 3-7 Main College Campus Site Reference Design



The main college campus typically consists of various sizes of building facilities and various education department groups. The network scale factor in the main college campus site is higher than the remote college campus site, and includes end users, IP-enabled endpoints, servers, and security and network edge devices. Multiple buildings of various sizes exist in one location, as shown in Figure 3-8.

Figure 3-8 Main College Campus Site Reference Design



The three-tier LAN design model for the main college campus meets all key technical aspects to provide a well-structured and strong network foundation. The modularity and flexibility in a three-tier LAN design model allows easier expansion and integration in the main college network, and keeps all network elements protected and available.

To enforce external network access policy for each end user, the three-tier model also provides external gateway services to the students and staff for accessing the Internet as well as private education and research networks.



Note

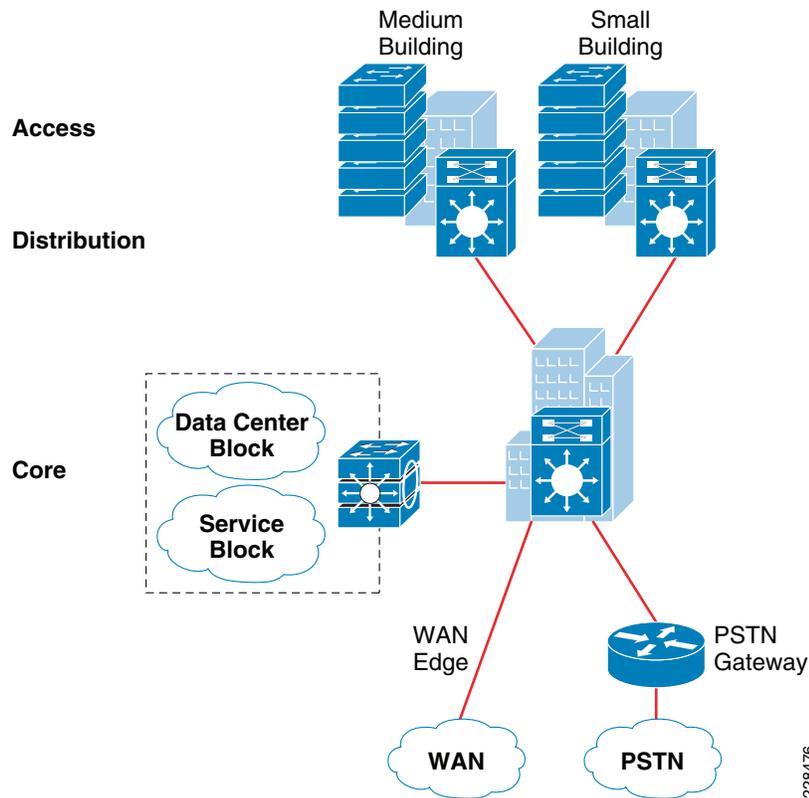
The WAN design is a separate element in this location, because it requires a separate WAN device that connects to the three-tier LAN model. WAN design is discussed in more detail in [Chapter 4, “Community College WAN Design Considerations.”](#)

Remote Large College Campus Site Design Overview

From the location size and network scale perspective, the remote large college is not much different from the main college campus site. Geographically, it can be distant from the main campus site and requires a high-speed WAN circuit to interconnect both campuses. The remote large college can also be considered as an alternate college campus to the main campus site, with the same common types of applications, endpoints, users, and network services. Similar to the main college campus, separate WAN devices are recommended to provide application delivery and access to the main college campus, given the size and number of students at this location.

Similar to the main college campus, Cisco recommends the three-tier LAN design model for the remote large college campus, as shown in Figure 3-9.

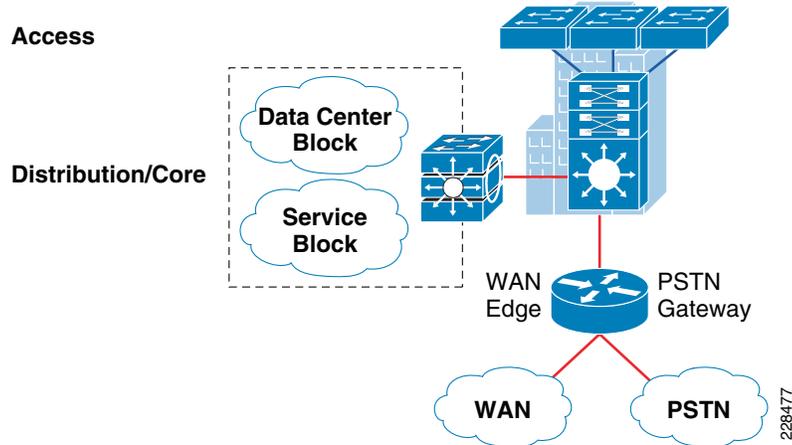
Figure 3-9 Remote Large College Campus Site Reference Design



Remote Medium College Campus Site Design Overview

Remote medium college campus locations differ from a main or remote large campus in that there are less buildings with distributed education departments. A remote medium college campus may have a fewer number of network users and endpoints, thereby reducing the need to build a similar campus network to that recommended for main and large college campuses. Because there are fewer students, faculty, and end users at this site as compared to the main or remote large campus sites, the need for a separate WAN device may not be necessary. A remote medium college campus network is designed similarly to a three-tier large campus LAN design. All the LAN benefits are achieved in a three-tier design model as in the main and remote large campus, and in addition, the platform chosen in the core layer also serves as the WAN edge, thus collapsing the WAN and core LAN functionality into a single platform. Figure 3-10 shows the remote medium campus in more detail.

Figure 3-10 Remote Medium College Campus Site Reference Design



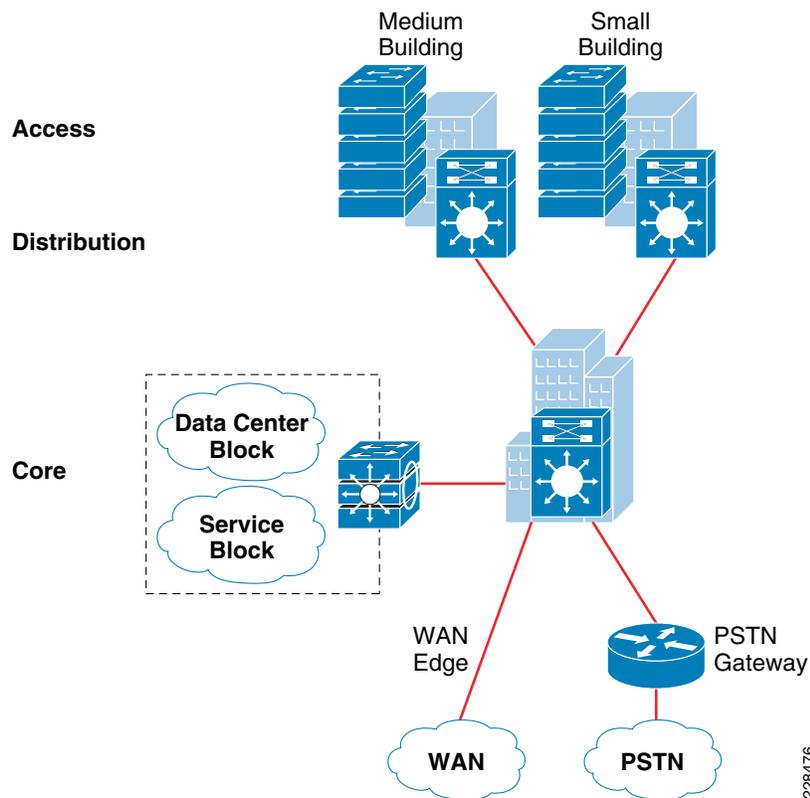
Remote Small College Campus Network Design Overview

The remote small college campus is typically confined to a single building that spans across multiple floors with different academic departments. The network scale factor in this design is reduced compared to other large college campuses. However, the application and services demands are still consistent across the community college locations.

In such smaller scale campus network deployments, the distribution and core layer functions can collapse into the two-tier LAN model without compromising basic network demands. Before deploying a collapsed core and distribution layer in the remote small campus network, considering all the scale and expansion factors prevents physical network re-design, and improves overall network efficiency and manageability.

WAN bandwidth requirements must be assessed appropriately for this remote small campus network design. Although the network scale factor is reduced compared to other larger college campus locations, sufficient WAN link capacity is needed to deliver consistent network services to student, faculty, and staff. Similar to the remote medium campus location, the WAN functionality is also collapsed into the LAN functionality. A single Cisco platform can provide collapsed core and distribution LAN layers. This design model is recommended only in smaller locations, and WAN traffic and application needs must be considered. [Figure 3-11](#) shows the remote small campus in more detail.

Figure 3-11 Remote Small College Campus Site Reference Design



Considering Multi-Tier LAN Design Models for Community Colleges

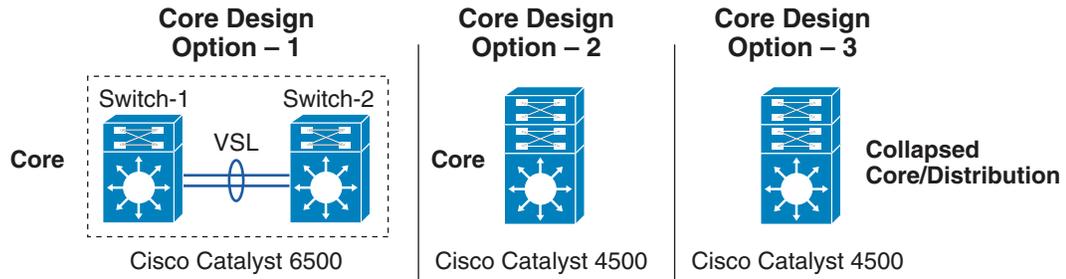
The previous section discussed the recommended LAN design model for each community college location. This section provides more detailed design guidance for each tier in the LAN design model. Each design recommendation is optimized to keep the network simplified and cost-effective without compromising network scalability, security, and resiliency. Each LAN design model for a community college location is based on the key LAN layers of core, distribution, and access.

Campus Core Layer Network Design

As discussed in the previous section, the core layer becomes a high-speed intermediate transit point between distribution blocks in different premises and other devices that interconnect to the data center, WAN, and Internet edge.

Similarly to choosing a LAN design model based on a location within the community college design, choosing a core layer design also depends on the size and location within the design. Three core layer design models are available, each of which is based on either the Cisco Catalyst 6500 Series or the Cisco Catalyst 4500 Series Switches. [Figure 3-12](#) shows the three core layer design models.

Figure 3-12 Core Layer Design Models for Community Colleges



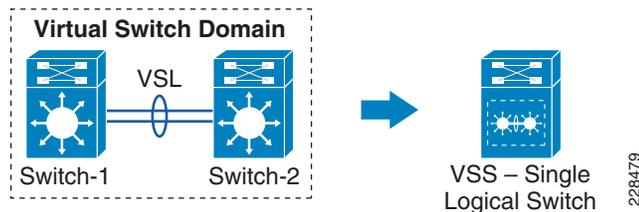
Each design model offers consistent network services, high availability, expansion flexibility, and network scalability. The following sections provide detailed design and deployment guidance for each model as well as where they fit within the various locations of the community college design.

Core Layer Design Option 1—Cisco Catalyst 6500-Based Core Network

Core layer design option 1 is specifically intended for the main and remote large campus locations. It is assumed that the number of network users, high-speed and low-latency applications (such as Cisco TelePresence), and the overall network scale capacity is common in both sites and thus, similar core design principles are required.

Core layer design option 1 is based on Cisco Catalyst 6500 Series switches using the Cisco Virtual Switching System (VSS), which is a software technology that builds a single logical core system by clustering two redundant core systems in the same tier. Building a VSS-based network changes network design, operation, cost, and management dramatically. Figure 3-13 shows the physical and operational view of VSS.

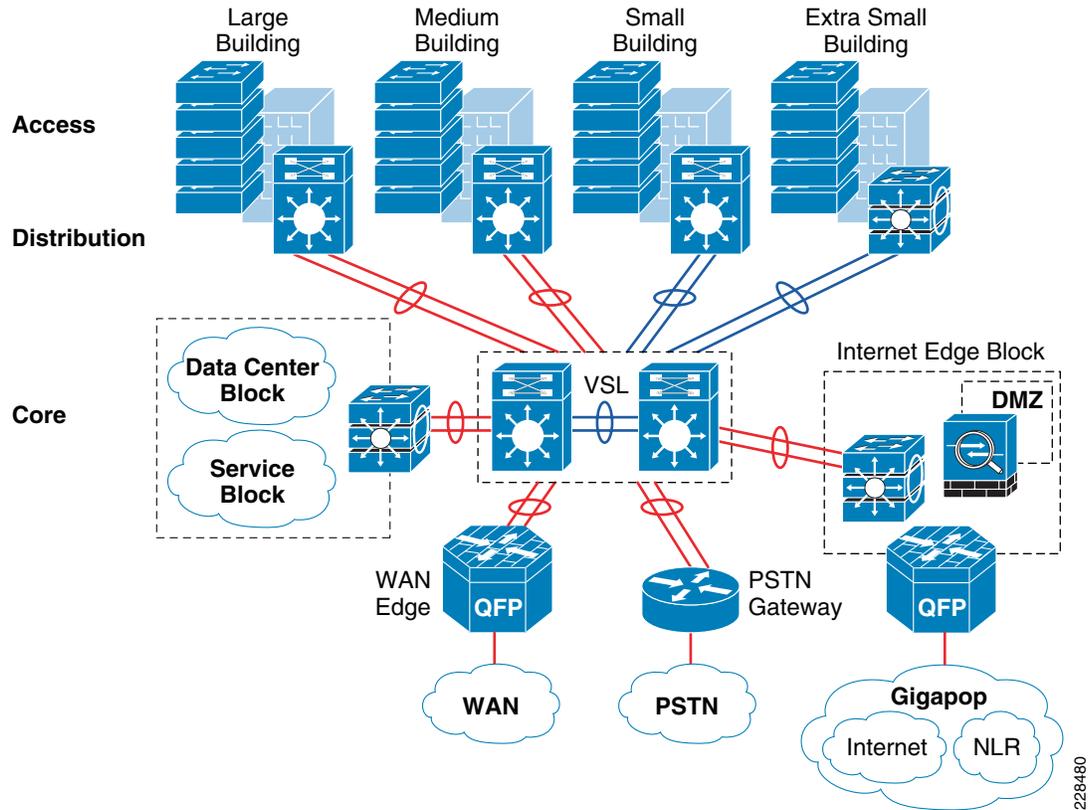
Figure 3-13 VSS Physical and Operational View



To provide end-to-end network access, the core layer interconnects several other network systems that are implemented in different roles and service blocks. Using VSS to virtualize the core layer into a single logical system remains transparent to each network device that interconnects to the VSS-enabled core. The single logical connection between core and the peer network devices builds a reliable, point-to-point connection that develops a simplified network topology and builds distributed forwarding tables to fully use all resources. Figure 3-14 shows a reference VSS-enabled core network design for the main campus site.

Figure 3-14

VSS-Enabled Core Network Design

**Note**

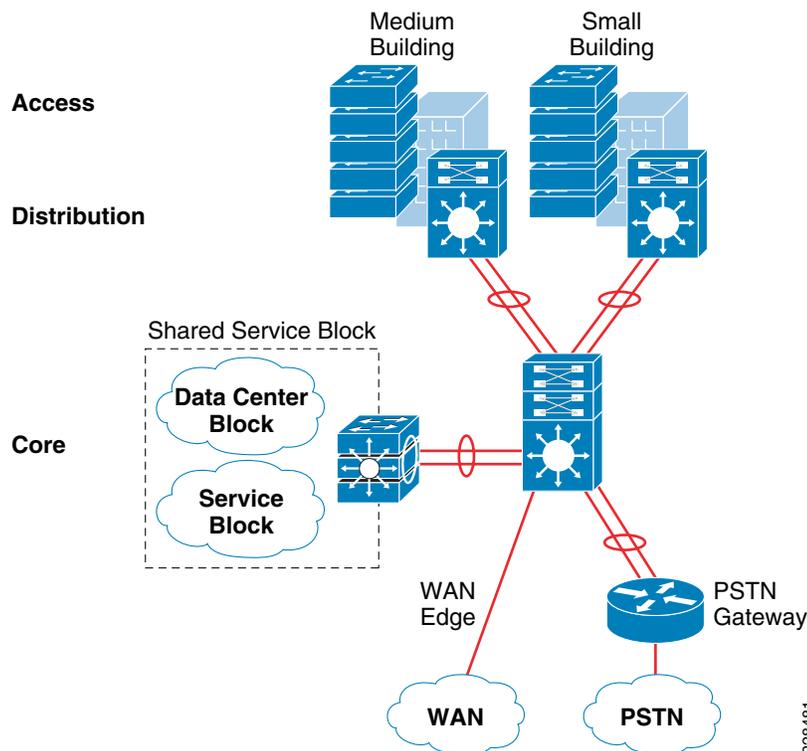
For more detailed VSS design guidance, see the *Campus 3.0 Virtual Switching System Design Guide* at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/VSS30dg/campusVSS_DG.html

Core Layer Design Option 2—Cisco Catalyst 4500-Based Campus Core Network

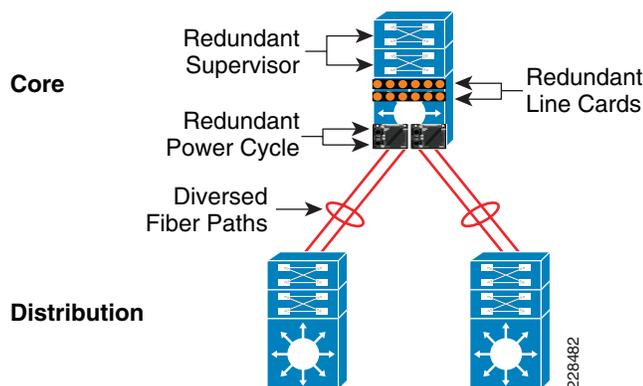
Core layer design option 2 is intended for a remote medium-sized college campus and is built on the same principles as for the main and remote large campus locations. The size of this remote site may not be large, and it is assumed that this location contains distributed building premises within the remote medium campus design. Because this site is smaller in comparison to the main and remote large campus locations, a fully redundant, VSS-based core layer design may not be necessary. Therefore, core layer design option 2 was developed to provide a cost-effective alternative while providing the same functionality as core layer design option 1. Figure 3-15 shows the remote medium campus core design option in more detail.

Figure 3-15 Remote Medium Campus Core Network Design



The cost of implementing and managing redundant systems in each tier may introduce complications in selecting the three-tier model, especially when network scale factor is not too high. This cost-effective core network design provides protection against various types of hardware and software failure and offers sub-second network recovery. Instead of a redundant node in the same tier, a single Cisco Catalyst 4500-E Series Switch can be deployed in the core role and bundled with 1+1 redundant in-chassis network components. The Cisco Catalyst 4500-E Series modular platform is a one-size platform that helps enable the high-speed core backbone to provide uninterrupted network access within a single chassis. Although a fully redundant, two-chassis design using VSS as described in core layer option 1 provides the greatest redundancy for large-scale locations, the redundant supervisors and line cards of the Cisco Catalyst 4500-E provide adequate redundancy for smaller locations within a single platform. Figure 3-16 shows the redundancy of the Cisco Catalyst 4500-E Series in more detail.

Figure 3-16 Highly Redundant Single Core Design Using the Cisco Catalyst 4500-E Platform



This core network design builds a network topology that has similar common design principles to the VSS-based campus core in core layer design option 1. The future expansion from a single core to a dual VSS-based core system becomes easier to deploy, and helps retain the original network topology and the management operation. This cost-effective single resilient core system for a medium-size college network meets the following four key goals:

- *Scalability*—The modular Cisco Catalyst 4500 chassis enables flexibility for core network expansion with high throughput modules and port scalability without compromising network performance.
- *Resiliency*—Because hardware or software failure conditions may create catastrophic results in the network, the single core system must be equipped with redundant system components such as supervisor, line card, and power supplies. Implementing redundant components increases the core network resiliency during various types of failure conditions using Non-Stop Forwarding/Stateful Switch Over (NSF/SSO) and EtherChannel technology.
- *Simplicity*—The core network can be simplified with redundant network modules and diverse fiber connections between the core and other network devices. The Layer 3 network ports must be bundled into a single point-to-point logical EtherChannel to simplify the network, such as the VSS-enabled campus design. An EtherChannel-based campus network offers similar benefits to an Multi-chassis EtherChannel (MEC)- based network.
- *Cost-effectiveness*—A single core system in the core layer helps reduce capital, operational, and management cost for the medium-sized campus network design.

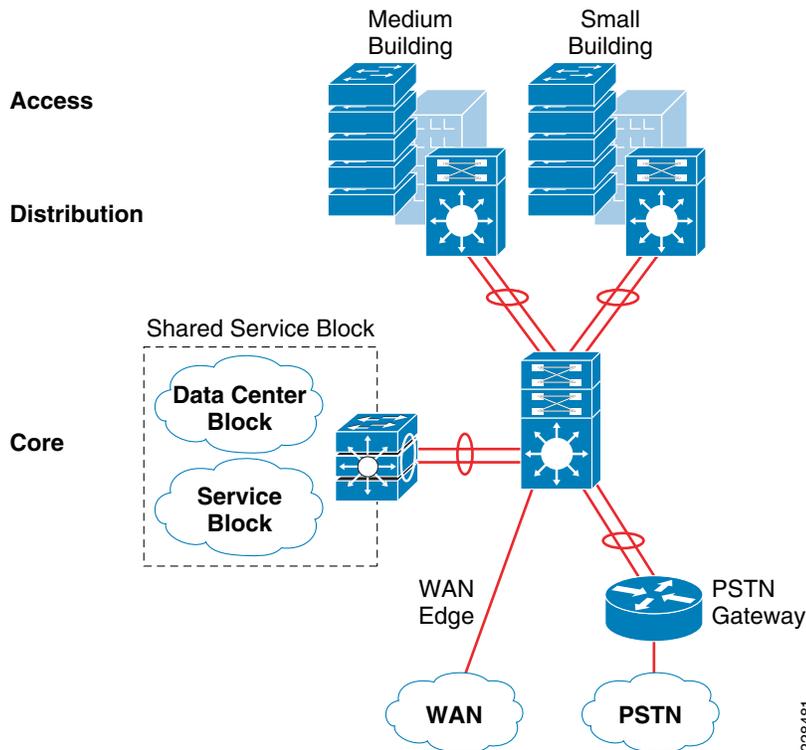
Core Layer Design Option 3—Cisco Catalyst 4500-Based Collapsed Core Campus Network

Core layer design option 3 is intended for the remote small campus network that has consistent network services and applications service-level requirements but at reduced network scale. The remote small campus is considered to be confined within a single multi-story building that may span academic departments across different floors. To provide consistent services and optimal network performance, scalability, resiliency, simplification, and cost-effectiveness in the small campus network design must not be compromised.

As discussed in the previous section, the remote small campus has a two-tier LAN design model, so the role of the core system is merged with the distribution layer. Remote small campus locations have consistent design guidance and best practices defined for main, remote large, and remote medium-sized campus cores. However, for platform selection, the remote medium campus core layer design must be leveraged to build this two-tier campus core.

Single highly resilient Cisco Catalyst 4500 switches with a Cisco Sup6L-E supervisor must be deployed in a centralized collapsed core and distribution role that interconnects to wiring closet switches, a shared service block, and a WAN edge router. The cost-effective supervisor version supports key technologies such as robust QoS, high availability, security, and much more at a lower scale, making it an ideal solution for small-scale network designs. [Figure 3-17](#) shows the remote small campus core design in more detail.

Figure 3-17 Core Layer Option 3 Collapsed Core/Distribution Network Design in Remote Small Campus Location

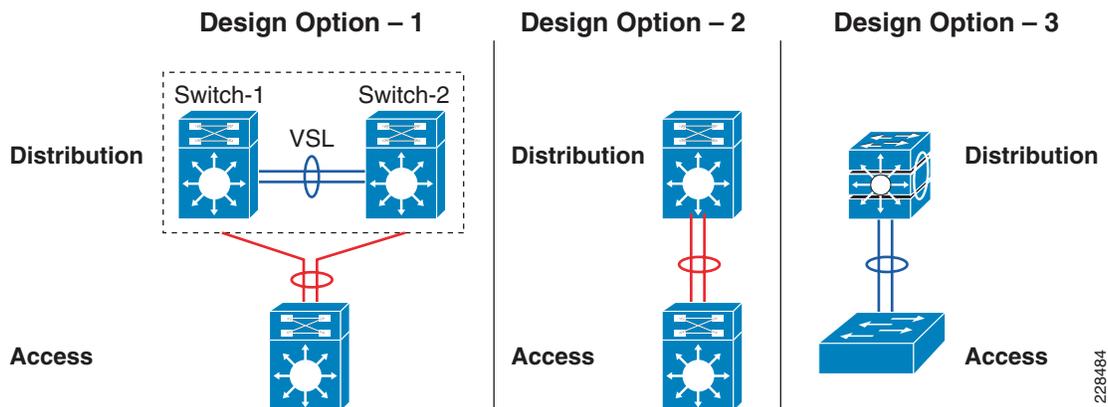


Campus Distribution Layer Network Design

The distribution or aggregation layer is the network demarcation boundary between wiring-closet switches and the campus core network. The framework of the distribution layer system in the community college design is based on best practices that reduce network complexities and accelerate reliability and performance. To build a strong campus network foundation with the three-tier model, the distribution layer has a vital role in consolidating networks and enforcing network edge policies.

Following the core layer design options in different campus locations, the distribution layer design provides consistent network operation and configuration tools to enable various network services. Three simplified distribution layer design options can be deployed in main or remote college campus locations, depending on network scale, application demands, and cost, as shown in [Figure 3-18](#). Each design model offers consistent network services, high availability, expansion flexibility, and network scalability.

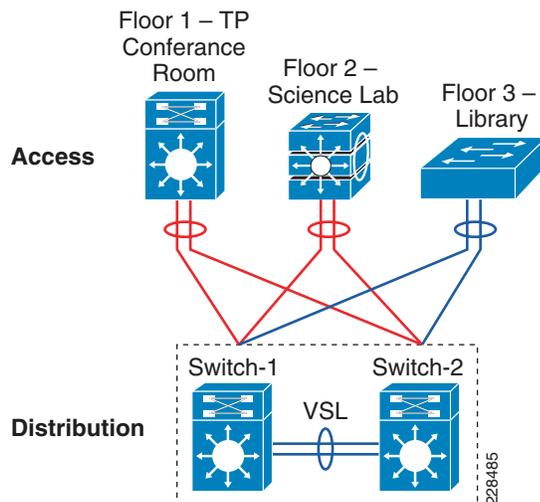
Figure 3-18 Distribution Layer Design Model Options



Distribution Layer Design Option 1—Cisco Catalyst 6500-E Based Distribution Network

Distribution layer design option 1 is intended for main campus and remote large campus locations, and is based on Cisco Catalyst 6500 Series switches using the Cisco VSS, as shown in Figure 3-19.

Figure 3-19 VSS-Enabled Distribution Layer Network Design



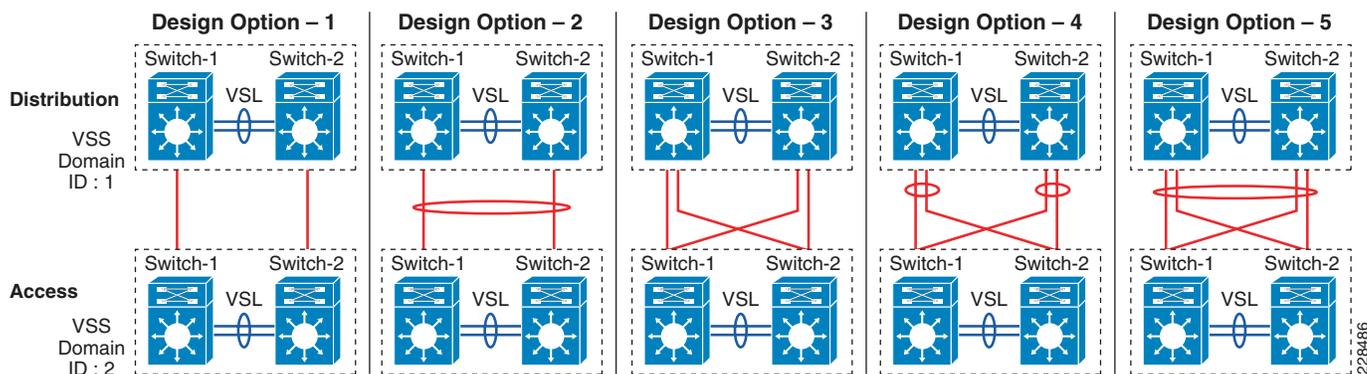
The distribution block and core network operation changes significantly when redundant Cisco Catalyst 6500-E Series switches are deployed in VSS mode in both the distribution and core layers. Clustering redundant distribution switches into a single logical system with VSS introduces the following technical benefits:

- A single logical system reduces operational, maintenance, and ownership cost.
- A single logical IP gateway develops a unified point-to-point network topology in the distribution block, which eliminates traditional protocol limitations and enables the network to operate at full capacity.
- Implementing the distribution layer in VSS mode eliminates or reduces several deployment barriers, such as spanning-tree loop, Hot Standby Routing Protocol (HSRP)/Gateway Load Balancing Protocol (GLBP)/Virtual Router Redundancy Protocol (VRRP), and control plane overhead.

- Cisco VSS introduces unique inter-chassis traffic engineering to develop a fully-distributed forwarding design that helps in increased bandwidth, load balancing, predictable network recovery, and network stability.

Deploying VSS mode in both the distribution layer switch and core layer switch provides numerous technology deployment options that are not available when not using VSS. Designing a common core and distribution layer option using VSS provides greater redundancy and is able to handle the amount of traffic typically present in the main and remote large campus locations. Figure 3-20 shows five unique VSS domain interconnect options. Each variation builds a unique network topology that has a direct impact on steering traffic and network recovery.

Figure 3-20 Core/Distribution Layer Interconnection Design Considerations



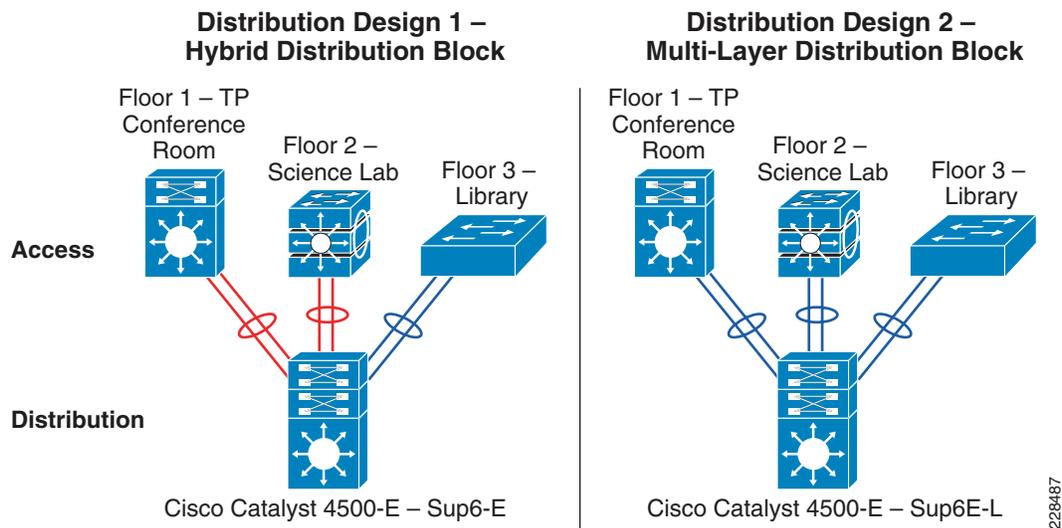
The various core/distribution layer interconnects offer the following:

- *Core/distribution layer interconnection option 1*—A single physical link between each core switch with the corresponding distribution switch.
- *Core/distribution layer interconnection option 2*—A single physical link between each core switch with the corresponding distribution switch, but each link is logically grouped to appear as one single link between the core and distribution layers.
- *Core/distribution layer interconnection option 3*—Two physical links between each core switch with the corresponding distribution switch. This design creates four equal cost multi-path (ECMP) with multiple control plane adjacency and redundant path information. Multiple links provide greater redundancy in case of link failover.
- *Core/distribution layer interconnection option 4*—Two physical links between each core switch with the corresponding distribution switch. There is one link direction between each switch as well as one link connecting to the other distribution switch. The additional link provides greater redundancy in case of link failover. Also these links are logically grouped to appear like option 1 but with greater redundancy.
- *Core/distribution layer interconnection option 5*—This provides the most redundancy between the VSS-enabled core and distribution switches as well as the most simplified configuration, because it appears as if there is only one logical link between the core and the distribution. Cisco recommends deploying this option because it provides higher redundancy and simplicity compared to any other deployment option.

Distribution Layer Design Option 2—Cisco Catalyst 4500-E-Based Distribution Network

Two cost-effective distribution layer models have been designed for the medium-sized and small-sized buildings within each campus location that interconnect to the centralized core layer design option and distributed wiring closet access layer switches. Both models are based on a common physical LAN network infrastructure and can be chosen based on overall network capacity and distribution block design. Both distribution layer design options use a cost-effective single and highly resilient Cisco Catalyst 4500 as an aggregation layer system that offers consistent network operation like a VSS-enabled distribution layer switch. The Cisco Catalyst 4500 Series provides the same technical benefits of VSS for a smaller network capacity within a single Cisco platform. The two Cisco Catalyst 4500-E-based distribution layer options are shown in [Figure 3-21](#).

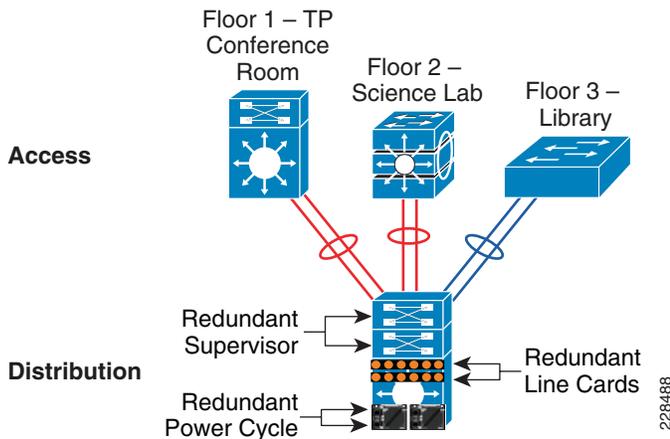
Figure 3-21 Two Cisco Catalyst 4500-E-Based Distribution Layer Options



The hybrid distribution block must be deployed with the next-generation supervisor Sup6-E module. Implementing redundant Sup6-Es in the distribution layer can interconnect access layer switches and core layer switches using a single point-to-point logical connection. This cost-effective and resilient distribution design option leverages core layer design option 2 to take advantage of all the operational consistency and architectural benefits.

Alternatively, the multilayer distribution block option requires the Cisco Catalyst 4500-E Series Switch with next-generation supervisor Sup6E-L deployed. The Sup6E-L supervisor is a cost-effective distribution layer solution that meets all network foundation requirements and can operate at moderate capacity, which can handle a medium-sized college distribution block.

This distribution layer network design provides protection against various types of hardware and software failure, and can deliver consistent sub-second network recovery. A single Catalyst 4500-E with multiple redundant system components can be deployed to offer 1+1 in-chassis redundancy, as shown in [Figure 3-22](#).

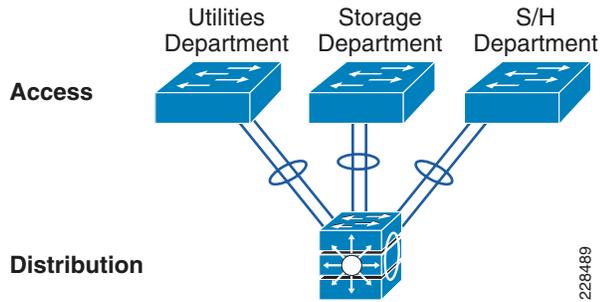
Figure 3-22 Highly Redundant Single Distribution Design

Distribution layer design option 2 is intended for the remote medium-sized campus locations, and is based on the Cisco Catalyst 4500 Series Switches. Although the remote medium and the main and remote large campus locations share similar design principles, the remote medium campus location is smaller and may not need a VSS-based redundant design. Fortunately, network upgrades and expansion become easier to deploy using distribution layer option 2, which helps retain the original network topology and the management operation. Distribution layer design option 2 meets the following goals:

- *Scalability*—The modular Cisco Catalyst 4500 chassis provides the flexibility for distribution block expansion with high throughput modules and port scalability without compromising network performance.
- *Resiliency*—The single distribution system must be equipped with redundant system components, such as supervisor, line card, and power supplies. Implementing redundant components increases network resiliency during various types of failure conditions using NSF/SSO and EtherChannel technology.
- *Simplicity*—This cost-effective design simplifies the distribution block similarly to a VSS-enabled distribution system. The single IP gateway design develops a unified point-to-point network topology in the distribution block to eliminate traditional protocol limitations, enabling the network to operate at full capacity.
- *Cost-effectiveness*—The single distribution system in the core layer helps reduce capital, operational, and ownership cost for the medium-sized campus network design.

Distribution Layer Design Option 3—Cisco Catalyst 3750-E StackWise-Based Distribution Network

Distribution layer design option 3 is intended for a very small building with a limited number of wiring closet switches in the access layer that connects remote classrooms or and office network with a centralized core, as shown in [Figure 3-23](#).

Figure 3-23 Cisco StackWise Plus-enabled Distribution Layer Network Design

While providing consistent network services throughout the campus, a number of network users and IT-managed remote endpoints can be limited in this building. This distribution layer design option recommends using the Cisco Catalyst 3750-E StackWise Plus Series platform for the distribution layer switch.

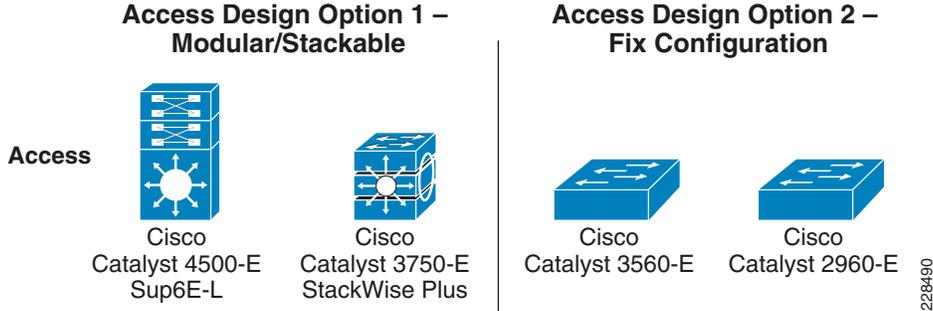
The fixed-configuration Cisco Catalyst 3750-E Series Switch is a multilayer platform that supports Cisco StackWise Plus technology to simplify the network and offers flexibility to expand the network as it grows. With Cisco StackWise Plus technology, the Catalyst 3750-E can be clustered into a high-speed backplane stack ring to logically build as a single large distribution system. Cisco StackWise Plus supports up to nine switches into single stack ring for incremental network upgrades, and increases effective throughput capacity up to 64 Gbps. The chassis redundancy is achieved via stacking, in which member chassis replicate the control functions with each member providing distributed packet forwarding. This is achieved by stacked group members acting as a single virtual Catalyst 3750-E switch. The logical switch is represented as one switch by having one stack member act as the master switch. Thus, when failover occurs, any member of the stack can take over as a master and continue the same services. It is a 1:N form of redundancy where any member can become the master. This distribution layer design option is ideal for the remote small campus location.

Campus Access Layer Network Design

The access layer is the first tier or edge of the campus, where end devices such as PCs, printers, cameras, Cisco TelePresence, and so on attach to the wired portion of the campus network. It is also the place where devices that extend the network out one more level, such as IP phones and wireless access points (APs), are attached. The wide variety of possible types of devices that can connect and the various services and dynamic configuration mechanisms that are necessary, make the access layer one of the most feature-rich parts of the campus network. Not only does the access layer switch allow users to access the network, the access layer switch must provide network protection so that unauthorized users or applications do not enter the network. The challenge for the network architect is determining how to implement a design that meets this wide variety of requirements, the need for various levels of mobility, the need for a cost-effective and flexible operations environment, while being able to provide the appropriate balance of security and availability expected in more traditional, fixed-configuration environments. The next-generation Cisco Catalyst switching portfolio includes a wide range of fixed and modular switching platforms, each designed with unique hardware and software capability to function in a specific role.

Community college campuses may deploy a wide range of network endpoints. The campus network infrastructure resources operate in shared service mode, and include IT-managed devices such as Cisco TelePresence and non-IT-managed devices such as student laptops. Based on several endpoint factors such as function and network demands and capabilities, two access layer design options can be deployed with college campus network edge platforms, as shown in [Figure 3-24](#).

Figure 3-24 Access Layer Design Models



Access Layer Design Option 1—Modular/StackWise Plus Access Layer Network

Access layer design option 1 is intended to address the network scalability and availability for the IT-managed critical voice and video communication network edge devices. To accelerate user experience and college campus physical security protection, these devices require low latency, high performance, and a constant network availability switching infrastructure. Implementing a modular and Cisco StackWise Plus-capable platform provides flexibility to increase network scale in the densely populated campus network edge.

The Cisco Catalyst 4500-E with supervisor Sup6E-L can be deployed to protect devices against access layer network failure. Cisco Catalyst 4500-E Series platforms offer consistent and predictable sub-second network recovery using NSF/SSO technology to minimize the impact of outages on college business and IT operation.

The Cisco Catalyst 3750-E Series is the alternate Cisco switching platform in this design option. Cisco StackWise Plus technology provides flexibility and availability by clustering multiple Cisco Catalyst 3750-E Series Switches into a single high-speed stack ring that simplifies operation and allows incremental access layer network expansion. The Cisco Catalyst 3750-E Series leverages EtherChannel technology for protection during member link or stack member switch failure.

Access Layer Design Option 2—Fixed Configuration Access Layer Network

This entry-level access layer design option is widely chosen for educational environments. The fixed configuration Cisco Catalyst switching portfolio supports a wide range of access layer technologies that allow seamless service integration and enable intelligent network management at the edge.

The fixed configuration Cisco Catalyst 3560-E Series is a commonly deployed platform for wired network access that can be in a mixed configuration with critical devices such as Cisco IP Phones and non-mission critical endpoints such as library PCs, printers, and so on. For non-stop network operation during power outages, the Catalyst 3560-E must be deployed with an internal or external redundant power supply solution using the Cisco RPS 2300. Increasing aggregated power capacity allows flexibility to scale power over Ethernet (PoE) on a per-port basis. With its wire-speed 10G uplink forwarding capacity, this design reduces network congestion and latency to significantly improve application performance.

For a college campus network, the Cisco Catalyst 3560-E is an alternate switching solution for the multilayer distribution block design option discussed in the previous section. The Cisco Catalyst 3560-E Series Switches offer limited software feature support that can function only in a traditional Layer 2 network design. To provide a consistent end-to-end enhanced user experience, the Cisco Catalyst 2960-E supports critical network control services to secure the network edge,

intelligently provide differentiated services to various class-of-service traffic, as well as simplified management. The Cisco Catalyst must leverage the 1G dual uplink port to interconnect the distribution system for increased bandwidth capacity and network availability.

Both design options offer consistent network services at the campus edge to provide differentiated, intelligent, and secured network access to trusted and untrusted endpoints. The distribution options recommended in the previous section can accommodate both access layer design options.

Community College Network Foundation Services Design

After each tier in the model has been designed, the next step for the community college design is to establish key network foundation services. Regardless of the application function and requirements that community colleges demand, the network must be designed to provide a consistent user experience independent of the geographical location of the application. The following network foundation design principles or services must be deployed in each campus location to provide resiliency and availability for all users to obtain and use the applications the community college offers:

- Network addressing hierarchy
- Network foundation technologies for LAN designs
- Multicast for applications delivery
- QoS for application performance optimization
- High availability to ensure user experience even with a network failure

Design guidance for each of these five network foundation services are discussed in the following sections, including where they are deployed in each tier of the LAN design model, the campus location, and capacity.

Network Addressing Hierarchy

Developing a structured and hierarchical IP address plan is as important as any other design aspect of the community college network to create an efficient, scalable, and stable network design. Identifying an IP addressing strategy for the network for the entire community college network design is essential.



Note

This section does not explain the fundamentals of TCP/IP addressing; for more details, see the many Cisco Press publications that cover this topic.

The following are key benefits of using hierarchical IP addressing:

- *Efficient address allocation*
 - Hierarchical addressing provides the advantage of grouping all possible addresses contiguously.
 - In non-contiguous addressing, a network can create addressing conflicts and overlapping problems, which may not allow the network administrator to use the complete address block.
- *Improved routing efficiencies*
 - Building centralized main and remote college campus site networks with contiguous IP addresses provides an efficient way to advertise summarized routes to neighbors.
 - Route summarization simplifies the routing database and computation during topology change events.

- Reduces network bandwidth utilization used by routing protocols.
- Improves overall routing protocol performance by flooding less messages and improves network convergence time.
- *Improved system performance*
 - Reduces the memory needed to hold large-scale discontinuous and non-summarized route entries.
 - Reduce higher CPU power to re-compute large-scale routing databases during topology change events.
 - Becomes easier to manage and troubleshoot.
 - Helps in overall network and system stability.

Network Foundational Technologies for LAN Design

In addition to a hierarchical IP addressing scheme, it is also essential to determine which areas of the community college design are Layer 2 or Layer 3 to determine whether routing or switching fundamentals need to be applied. The following applies to the three layers in a LAN design model:

- *Core layer*—Because this is a Layer 3 network that interconnects several remote locations and shared devices across the network, choosing a routing protocol is essential at this layer.
- *Distribution layer*—The distribution block uses a combination of Layer 2 and Layer 3 switching to provide for the appropriate balance of policy and access controls, availability, and flexibility in subnet allocation and VLAN usage. Both routing and switching fundamentals need to be applied.
- *Access layer*—This layer is the demarcation point between network infrastructure and computing devices. This is designed for critical network edge functions to provide intelligent application and device-aware services, to set the trust boundary to distinguish applications, provide identity-based network access to protected data and resources, provide physical infrastructure services to reduce greenhouse emission, and more. This subsection provides design guidance to enable various types of Layer 1 to 3 intelligent services, and to optimize and secure network edge ports.

The recommended routing or switching scheme of each layer is discussed in the following sections.

Designing the Core Layer Network

Because the core layer is a Layer 3 network, routing principles must be applied. Choosing a routing protocol is essential, and routing design principles and routing protocol selection criteria are discussed in the following subsections.

Routing Design Principles

Although enabling routing functions in the core is a simple task, the routing blueprint must be well understood and designed before implementation, because it provides the end-to-end reachability path of the college network. For an optimized routing design, the following three routing components must be identified and designed to allow more network growth and provide a stable network, independent of scale:

- *Hierarchical network addressing*—Structured IP network addressing in the community college LAN and/or WAN design is required to make the network scalable, optimal, and resilient.

- *Routing protocol*—Cisco IOS supports a wide range of Interior Gateway Protocols (IGPs). Cisco recommends deploying a single routing protocol across the community college network infrastructure.
- *Hierarchical routing domain*—Routing protocols must be designed in a hierarchical model that allows the network to scale and operate with greater stability. Building a routing boundary and summarizing the network minimizes the topology size and synchronization procedure, which improves overall network resource use and re-convergence.

Routing Protocol Selection Criteria

The criteria for choosing the right protocol vary based on the end-to-end network infrastructure. Although all the routing protocols that Cisco IOS currently supports can provide a viable solution, network architects must consider all the following critical design factors when selecting the right routing protocol to be implemented throughout the internal network:

- *Network design*—Requires a proven protocol that can scale in full-mesh campus network designs and can optimally function in hub-and-spoke WAN network topologies.
- *Scalability*—The routing protocol function must be network- and system-efficient and operate with a minimal number of updates and re-computation, independent of the number of routes in the network.
- *Rapid convergence*—Link-state versus DUAL re-computation and synchronization. Network re-convergence also varies based on network design, configuration, and a multitude of other factors that may be more than a specific routing protocol can handle. The best convergence time can be achieved from a routing protocol if the network is designed to the strengths of the protocol.
- *Operational*—A simplified routing protocol that can provide ease of configuration, management, and troubleshooting.

Cisco IOS supports a wide range of routing protocols, such as Routing Information Protocol (RIP) v1/2, Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), and Intermediate System-to-Intermediate System (IS-IS). However, Cisco recommends using EIGRP or OSPF for this network design. EIGRP is a popular version of an Interior Gateway Protocol (IGP) because it has all the capabilities needed for small to large-scale networks, offers rapid network convergence, and above all is simple to operate and manage. OSPF is popular link-state protocol for large-scale enterprise and service provider networks. OSPF enforces hierarchical routing domains in two tiers by implementing backbone and non-backbone areas. The OSPF area function depends on the network connectivity model and the role of each OSPF router in the domain. OSPF can scale higher but the operation, configuration, and management might become too complex for the community college LAN network infrastructure.

Other technical factors must be considered when implementing OSPF in the network, such as OSPF router type, link type, maximum transmission unit (MTU) considerations, designated router (DR)/backup designated router (BDR) priority, and so on. This document provides design guidance for using simplified EIGRP in the community college campus and WAN network infrastructure.

**Note**

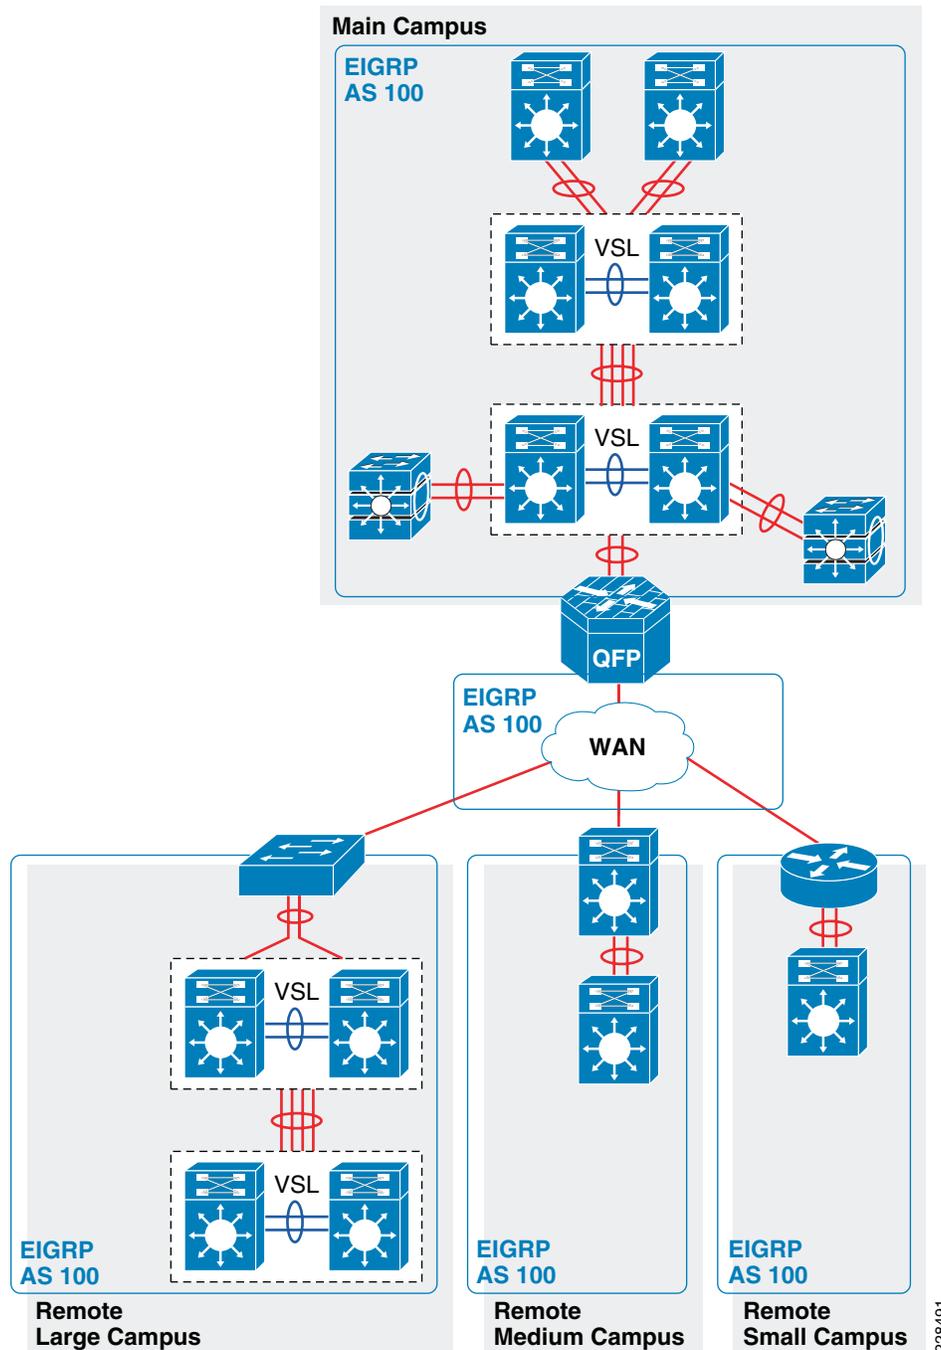
For detailed information on EIGRP and OSPF, see the following URL:
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/routed-ex.html>

Designing an End-to-End EIGRP Routing Network

EIGRP is a balanced hybrid routing protocol that builds neighbor adjacency and flat routing topology on a per autonomous system (AS) basis. Cisco recommends considering the following three critical design tasks before implementing EIGRP in the community college LAN core layer network:

- *EIGRP autonomous system*—The Layer 3 LAN and WAN infrastructure of the community college design must be deployed in a single EIGRP AS, as shown in Figure 3-25. A single EIGRP AS reduces operational tasks and prevents route redistribution, loops, and other problems that may occur because of misconfiguration.

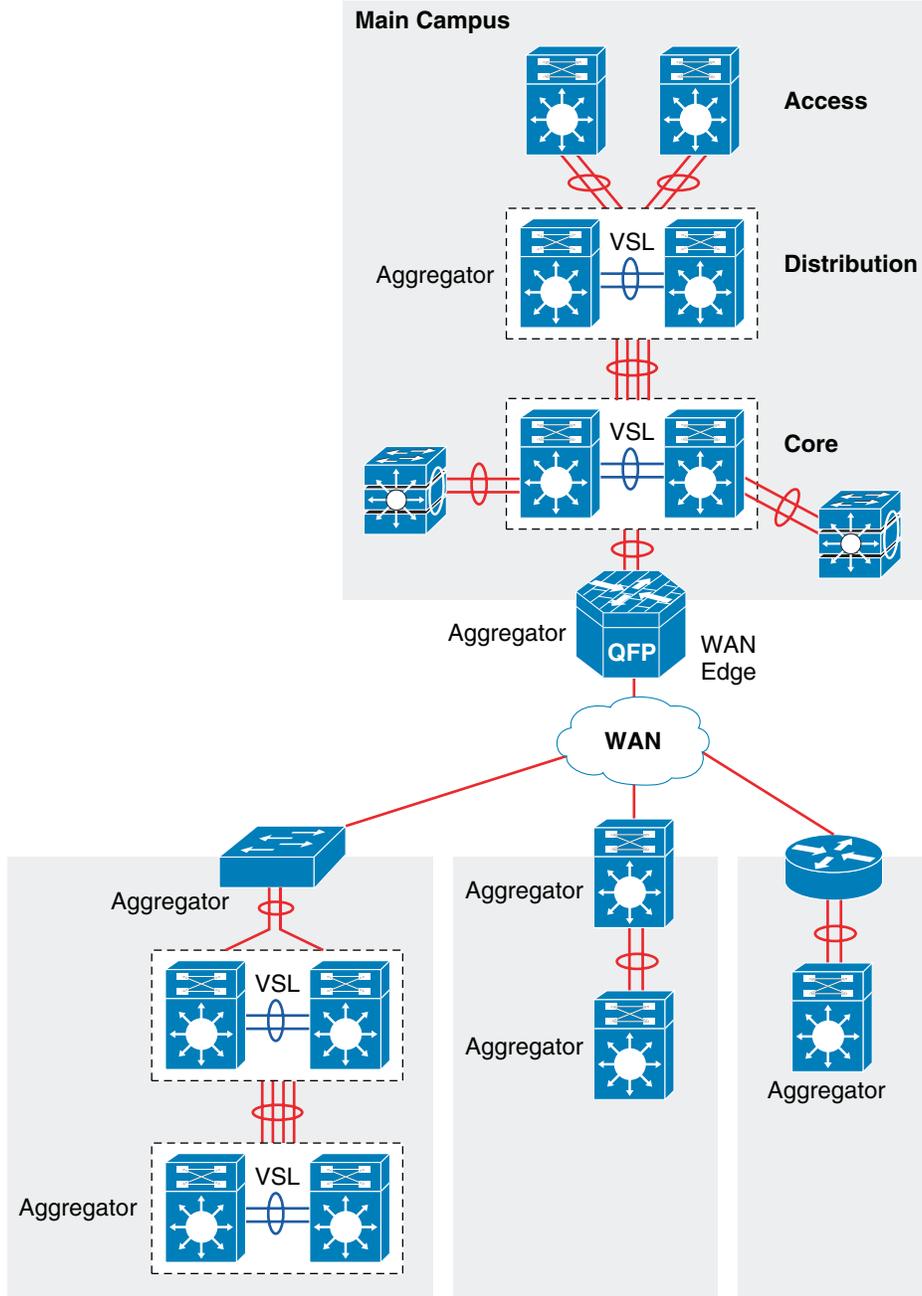
Figure 3-25 Sample End-to-End EIGRP Routing Design in Community College LAN Network



In the example in Figure 3-25, AS100 is the single EIGRP AS for the entire design.

- *EIGRP adjacency protection*—This increases network infrastructure efficiency and protection by securing the EIGRP adjacencies with internal systems. This task involves two subset implementation tasks on each EIGRP-enabled network devices:
 - Increases system efficiency—Blocks EIGRP processing with passive-mode configuration on physical or logical interfaces connected to non- EIGRP devices in the network, such as PCs. The best practice helps reduce CPU utilization and secures the network with unprotected EIGRP adjacencies with untrusted devices.
 - Network security—Each EIGRP neighbor in the LAN/WAN network must be trusted by implementing and validating the Message-Digest algorithm 5 (MD5) authentication method on each EIGRP-enabled system in the network.
- *Optimizing EIGRP topology*—EIGRP allows network administrators to summarize multiple individual and contiguous networks into a single summary network before advertising to the neighbor. Route summarization helps improve network performance, stability, and convergence by hiding the fault of an individual network that requires each router in the network to synchronize the routing topology. Each aggregating device must summarize a large number of networks into a single summary route. [Figure 3-26](#) shows an example of the EIGRP topology for the community college LAN design.

Figure 3-26 EIGRP Route Aggregator Design



By default, EIGRP speakers transmit Hello packets every 5 seconds, and terminates EIGRP adjacency if the neighbor fails to receive it within 15 seconds of hold-down time. In this network design, Cisco recommends retaining default EIGRP Hello and Hold timers on all EIGRP-enabled platforms.

Designing the Campus Distribution Layer Network

This section provides design guidelines for deploying various types of Layer 2 and Layer 3 technology in the distribution layer. Independent of which implemented distribution layer design model is deployed, the deployment guidelines remain consistent in all designs.

Because the distribution layer can be deployed with both Layer 2 and Layer 3 technologies, the following two network designs are recommended:

- Multilayer
- Routed access

Designing the Multilayer Network

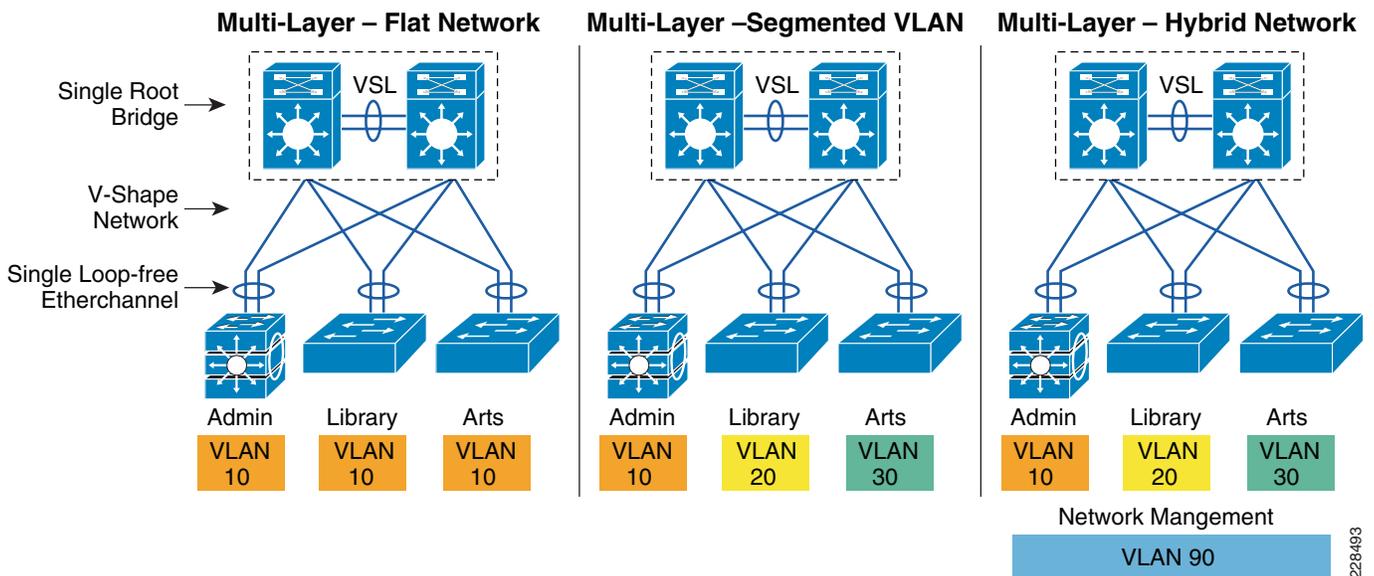
A multilayer network is a traditional, simple, and widely deployed scenario, regardless of network scale. The access layer switches in the campus network edge interface with various types of endpoints and provide intelligent Layer 1/2 services. The access layer switches interconnect to distribution switches with the Layer 2 trunk, and rely on the distribution layer aggregation switch to perform intelligent Layer 3 forwarding and to set policies and access control.

There are the following three design variations to build a multilayer network; all variations must be deployed in a V-shape physical network design and must be built to provide a loop-free topology:

- *Flat*—Certain applications and user access requires that the broadcast domain design span more than a single wiring closet switch. The multilayer network design provides the flexibility to build a single large broadcast domain with an extended star topology. Such flexibility introduces scalability, performance, and security challenges, and may require extra attention to protect the network against misconfiguration and miswiring that can create spanning-tree loops and de-stabilize the network.
- *Segmented*—Provides a unique VLAN for different education divisions and college business function segments to build a per-department logical network. All network communication between education and administrative groups passes through the routing and forwarding policies defined at the distribution layer.
- *Hybrid*—A hybrid logical network design segments VLAN workgroups that do not span different access layer switches, and allows certain VLANs (for example, that net management VLAN) to span across the access-distribution block. The hybrid network design enables flat Layer 2 communication without impacting the network, and also helps reduce the number of subnets used.

Figure 3-27 shows the three design variations for the multilayer network.

Figure 3-27 Multilayer Design Variations



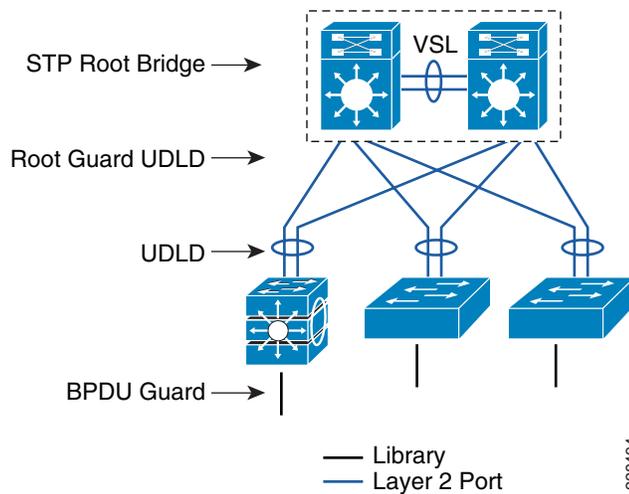
228493

Cisco recommends that the hybrid multilayer access-distribution block design use a loop-free network topology, and span a few VLANs that require such flexibility, such as the management VLAN.

Ensuring a loop-free topology is critical in a multilayer network design. Spanning-Tree Protocol (STP) dynamically develops a loop-free multilayer network topology that can compute the best forwarding path and provide redundancy. Although STP behavior is deterministic, it is not optimally designed to mitigate network instability caused by hardware miswiring or software misconfiguration. Cisco has developed several STP extensions to protect against network malfunctions, and to increase stability and availability. All Cisco Catalyst LAN switching platforms support the complete STP toolkit suite that must be enabled globally on individual logical and physical ports of the distribution and access layer switches.

Figure 3-28 shows an example of enabling various STP extensions on distribution and access layer switches in all campus sites.

Figure 3-28 Protecting Multilayer Network with Cisco STP Toolkit



Note

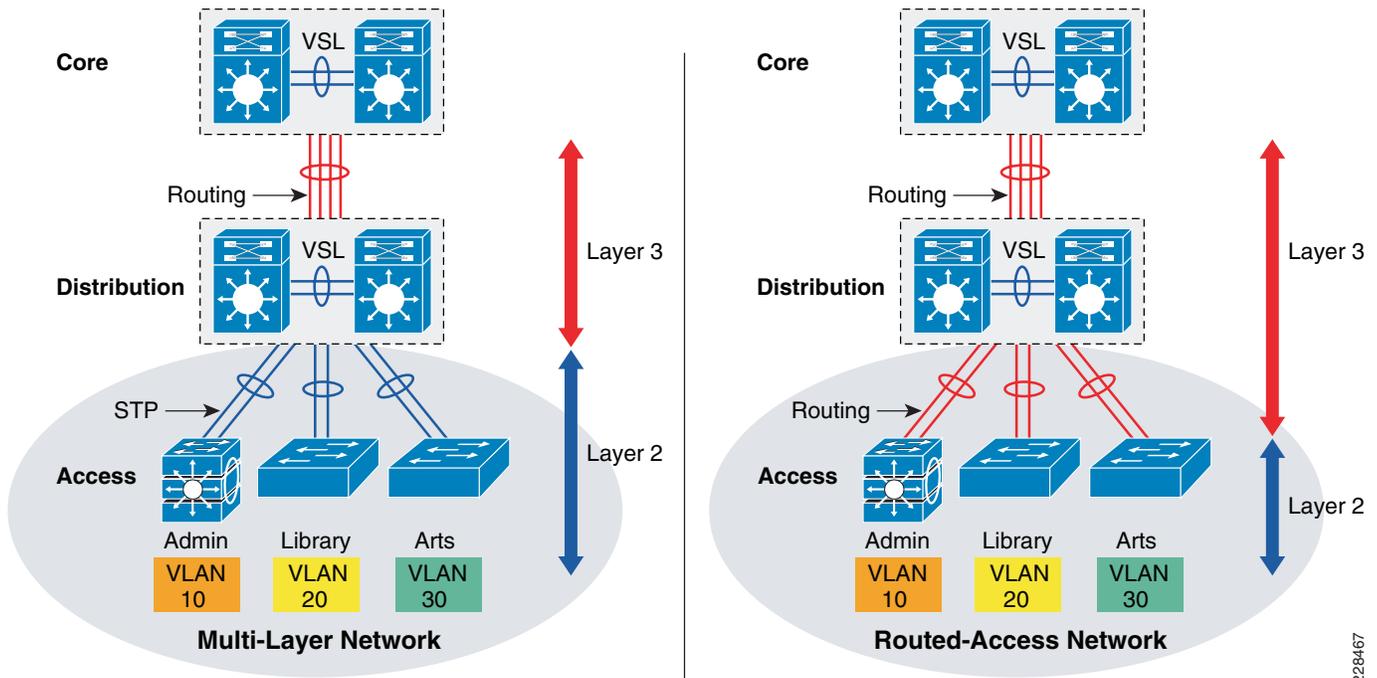
For additional STP information, see the following URL:

http://www.cisco.com/en/US/tech/tk389/tk621/tsd_technology_support_troubleshooting_technotes_list.html.

Designing the Routed Access Network

Routing functions in the access layer network simplify configuration, optimize distribution performances, and provide end-to-end troubleshooting tools. Implementing Layer 3 functions in the access layer replaces Layer 2 trunk configuration to a single point-to-point Layer 3 interface with a collapsed core system in the aggregation layer. Pushing Layer 3 functions one tier down on Layer 3 access switches changes the traditional multilayer network topology and forwarding development path. Implementing Layer 3 functions in the access switch does not require any physical or logical link reconfiguration; the access-distribution block can be used, and is as resilient as in the multilayer network design. Figure 3-29 shows the differences between the multilayer and routed access network designs, as well as where the Layer 2 and Layer 3 boundaries exist in each network design.

Figure 3-29 Layer 2 and Layer 3 Boundaries for Multilayer and Routed Access Network Design



228467

Routed-access network design enables Layer 3 access switches to perform Layer 2 demarcation point and provide Inter-VLAN routing and gateway function to the endpoints. The Layer 3 access switches makes more intelligent, multi-function and policy-based routing and switching decision like distribution-layer switches.

Although Cisco VSS and a single redundant distribution design are simplified with a single point-to-point EtherChannel, the benefits in implementing the routed access design in community colleges are as follows:

- Eliminates the need for implementing STP and the STP toolkit on the distribution system. As a best practice, the STP toolkit must be hardened at the access layer.
- Shrinks the Layer 2 fault domain, thus minimizing the number of denial-of-service (DoS)/distributed denial-of-service (DDoS) attacks.
- Bandwidth efficiency—Improves Layer 3 uplink network bandwidth efficiency by suppressing Layer 2 broadcasts at the edge port.
- Improves overall collapsed core and distribution resource utilization.

Enabling Layer 3 functions in the access-distribution block must follow the same core network designs as mentioned in previous sections to provide network security as well as optimize the network topology and system resource utilization:

- *EIGRP autonomous system*—Layer 3 access switches must be deployed in the same EIGRP AS as the distribution and core layer systems.
- *EIGRP adjacency protection*—EIGRP processing must be enabled on uplink Layer 3 EtherChannels, and must block remaining Layer 3 ports by default in passive mode. Access switches must establish secured EIGRP adjacency using the MD5 hash algorithm with the aggregation system.

- *EIGRP network boundary*—All EIGRP neighbors must be in a single AS to build a common network topology. The Layer 3 access switches must be deployed in EIGRP stub mode for a concise network view.

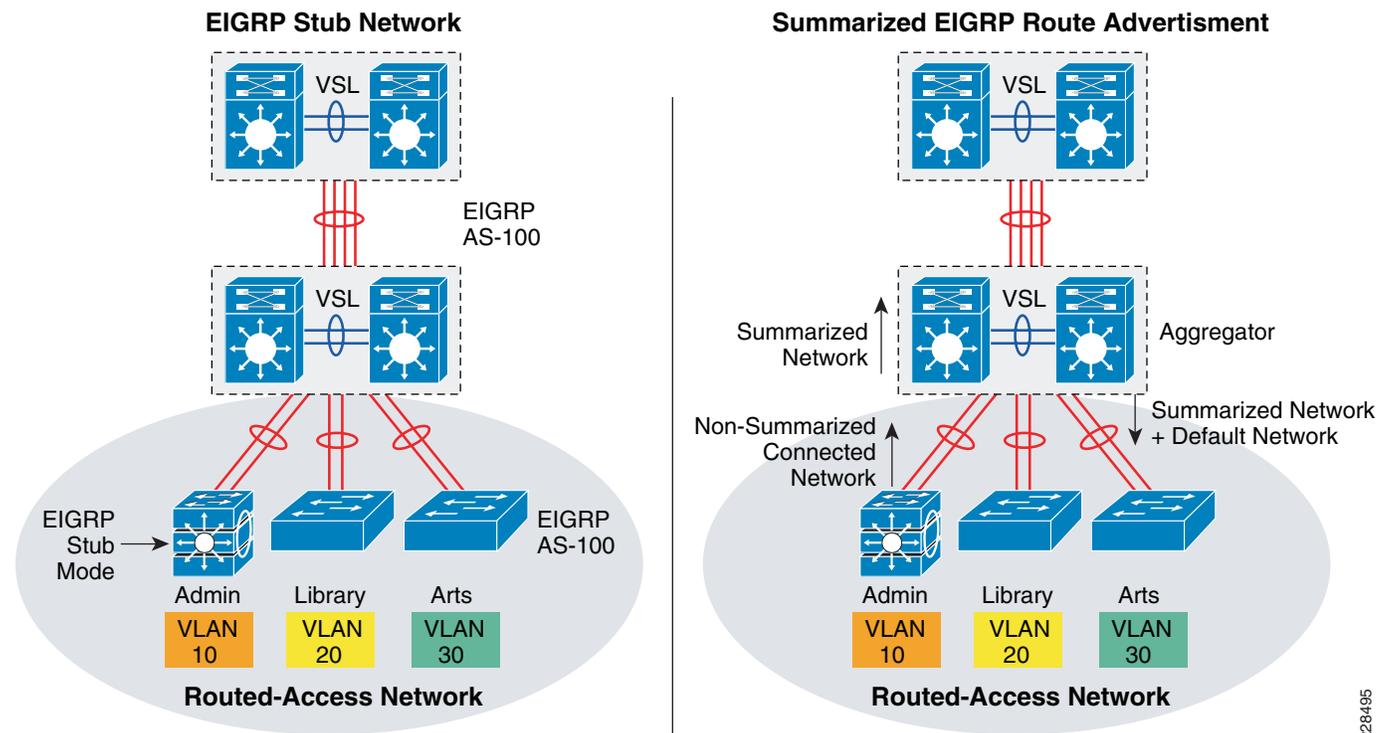
Designing the Layer 3 Access Layer

EIGRP creates and maintains a single flat routing topology network between EIGRP peers. Building a single routing domain in a large-scale campus core design allows for complete network visibility and reachability that may interconnect multiple campus components, such as distribution blocks, services blocks, the data center, the WAN edge, and so on.

In the three- or two-tier deployment models, the Layer 3 access switch must always have single physical or logical forwarding to a distribution switch. The Layer 3 access switch dynamically develops the forwarding topology pointing to a single distribution switch as a single Layer 3 next hop. Because the distribution switch provides a gateway function to rest of the network, the routing design on the Layer 3 access switch can be optimized with the following two techniques to improve performance and network reconvergence in the access-distribution block, as shown in [Figure 3-30](#):

- Deploying the Layer 3 access switch in EIGRP stub mode
- Summarizing the network view with a default route to the Layer 3 access switch for intelligent routing functions

Figure 3-30 Designing and Optimizing EIGRP Network Boundary for the Access Layer



228495

Multicast for Application Delivery

Because unicast communication is based on the one-to-one forwarding model, it becomes easier in routing and switching decisions to perform destination address lookup, determine the egress path by scanning forwarding tables, and to switch traffic. In the unicast routing and switching technologies discussed in the previous section, the network may need to be made more efficient by allowing certain applications where the same content or application must be replicated to multiple users. IP multicast delivers source traffic to multiple receivers using the least amount of network resources as possible without placing an additional burden on the source or the receivers. Multicast packet replication in the network is done by Cisco routers and switches enabled with Protocol Independent Multicast (PIM) as well as other multicast routing protocols.

Similar to the unicast methods, multicast requires the following design guidelines:

- Choosing a multicast addressing design
- Choosing a multicast routing protocol
- Providing multicast security regardless of the location within the community college design

Multicast Addressing Design

The Internet Assigned Numbers Authority (IANA) controls the assignment of IP multicast addresses. A range of class D address space is assigned to be used for IP multicast applications. All multicast group addresses fall in the range from 224.0.0.0 through 239.255.255.255. Layer 3 addresses in multicast communications operate differently; while the destination address of IP multicast traffic is in the multicast group range, the source IP address is always in the unicast address range. Multicast addresses are assigned in various pools for well-known multicast-based network protocols or inter-domain multicast communications, as listed in [Table 3-2](#).

Table 3-2 **Multicast Address Range Assignments**

Application	Address Range
Reserved—Link local network protocols.	224.0.0.0/24
Global scope—Group communication between an organization and the Internet.	224.0.1.0 – 238.255.255.255
Source Specific Multicast (SSM)—PIM extension for one-to-many unidirectional multicast communication.	232.0.0.0/8
GLOP—Inter-domain multicast group assignment with reserved global AS.	233.0.0.0/8
Limited scope—Administratively scoped address that remains constrained within a local organization or AS. Commonly deployed in enterprise, education, and other organizations.	239.0.0.0/8

During the multicast network design phase, community college network architects must select a range of multicast sources from the limited scope pool (239/8).

Multicast Routing Design

To enable end-to-end dynamic multicast operation in the network, each intermediate system between the multicast receiver and source must support the multicast feature. Multicast develops the forwarding table differently than the unicast routing and switching model. To enable communication, multicast requires specific multicast routing protocols and dynamic group membership.

Multicast Routing Protocol Design

The community college LAN design must be able to build packet distribution trees that specify a unique forwarding path between the subnet of the source and each subnet containing members of the multicast group. A primary goal in distribution trees construction is to ensure that no more than one copy of each packet is forwarded on each branch of the tree. The two basic types of multicast distribution trees are as follows:

- *Source trees*—The simplest form of a multicast distribution tree is a source tree, with its root at the source and branches forming a tree through the network to the receivers. Because this tree uses the shortest path through the network, it is also referred to as a shortest path tree (SPT).
- *Shared trees*—Unlike source trees that have their root at the source, shared trees use a single common root placed at a selected point in the network. This shared root is called a rendezvous point (RP).

The PIM protocol is divided into the following two PIM modes to support both types of multicast distribution trees:

- *Dense mode (DM)*—Assumes that almost all routers in the network need to distribute multicast traffic for each multicast group (for example, almost all hosts on the network belong to each multicast group). PIM in DM mode builds distribution trees by initially flooding the entire network and then pruning back the small number of paths without receivers.
- *Sparse mode (SM)*—Assumes that relatively few routers in the network are involved in each multicast. The hosts belonging to the group are widely dispersed, as might be the case for most multicasts over the WAN. Therefore, PIM-SM begins with an empty distribution tree and adds branches only as the result of explicit Internet Group Management Protocol (IGMP) requests to join the distribution. PIM-SM mode is ideal for a network without dense receivers and multicast transport over WAN environments, and it adjusts its behavior to match the characteristics of each receiver group.

Selecting the PIM mode depends on the multicast applications that use various mechanisms to build multicast distribution trees. Based on the multicast scale factor and centralized source deployment design for one-to-many multicast communication in community college LAN infrastructures, Cisco recommends deploying PIM-SM because it is efficient and intelligent in building multicast distribution tree. All the recommended platforms in this design support PIM-SM mode on physical or logical (switched virtual interface [SVI] and EtherChannel) interfaces.

Designing PIM Rendezvous Point

The following sections discuss best practices in designing the PIM RP.

PIM-SM RP Placement Best Practices

It is assumed that each community college site has a wide range of local multicast sources in the data center for distributed community college IT-managed media and student research and development applications. In such a distributed multicast network design, Cisco recommends deploying PIM RP on

each site for wired or wireless multicast receivers and sources to join and register at the closest RP. The community college reference design recommends PIM-SM RP placement on a VSS-enabled and single resilient core system in the three-tier campus design, or on the collapsed core/distribution system in the two-tier campus design model.

PIM-SM RP Mode Best Practices

PIM-SM supports RP deployment in the following three modes in the network:

- *Static*—In this mode, RP must be statically identified and configured on each PIM router in the network. RP load balancing and redundancy can be achieved using anycast RP.
- *Auto-RP*—This mode is a dynamic method for discovering and announcing the RP in the network. Auto-RP implementation is beneficial when there are multiple RPs and groups that often change in the network. To prevent network reconfiguration during a change, the RP mapping agent router must be designated in the network to receive RP group announcements and to arbitrate conflicts, as part of the PIM version 1 specification.
- *Bootstrap Router (BSR)*—This mode performs the same tasks as Auto-RP but in a different way, and is part of the PIM version 2 specification. Auto-RP and BSR cannot co-exist or interoperate in the same network.

In a small to mid-sized multicast network, static RP configuration is recommended over the other modes. Static RP implementation offers RP redundancy and load sharing, and an additional simple access control list (ACL) can be applied to deploy RP without compromising multicast network security. Cisco recommends designing the community college LAN multicast network using the static PIM-SM mode configuration.

PIM-SM RP Redundancy Best Practices

PIM-SM RP redundancy and load sharing becomes imperative in the community college LAN design, because each recommended core layer design model provides resiliency and simplicity. In the Cisco Catalyst 6500 VSS-enabled core layer, the dynamically discovered group-to-RP entries are fully synchronized to the standby switch. Combining NSF/SSO capabilities with IPv4 multicast reduces the network recovery time and retains the user and application performance at an optimal level. In the non-VSS-enabled network design, PIM-SM uses Anycast RP and Multicast Source Discovery Protocol (MSDP) for node failure protection. PIM-SM redundancy and load sharing is simplified with the Cisco VSS-enabled core. Because VSS is logically a single system and provides node protection, there is no need to implement Anycast RP and MSDP on a VSS-enabled PIM-SM RP.

Inter-Site PIM RP Best Practices

MSDP allows PIM RPs to share information about the active sources. PIM-SM RPs discover local receivers through PIM join messages, while the multicast source can be in a local or remote network domain. MSDP allows each multicast domain to maintain an independent RP that does not rely on other multicast domains, but does enable RPs to forward traffic between domains. PIM-SM is used to forward the traffic between the multicast domains.

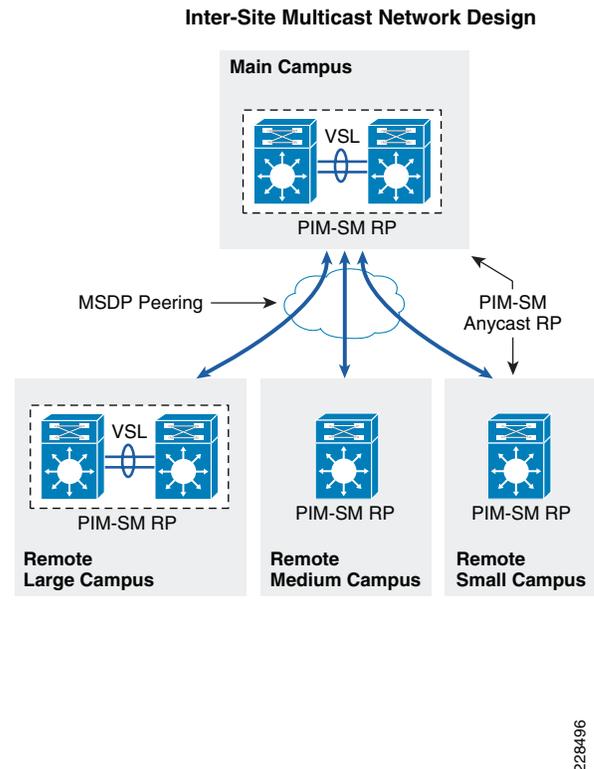
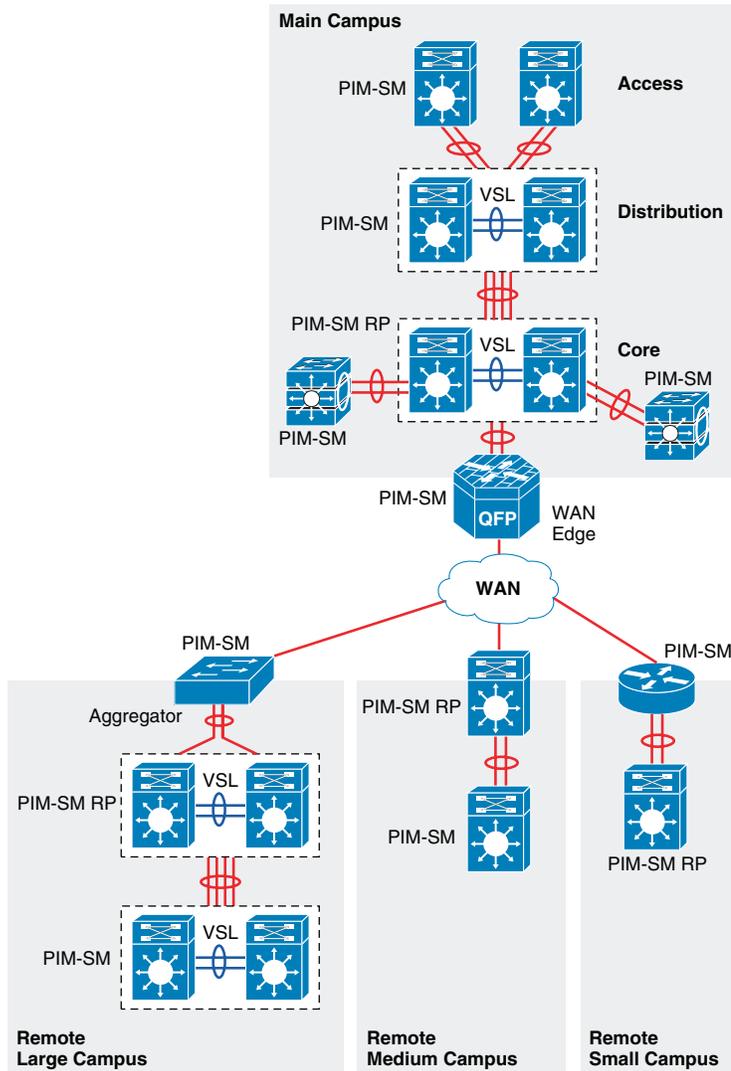
Anycast RP is a useful application of MSDP. Originally developed for interdomain multicast applications, MSDP used with Anycast RP is an intradomain feature that provides redundancy and load sharing capabilities. Large networks typically use Anycast RP for configuring a PIM-SM network to meet fault tolerance requirements within a single multicast domain.

The community college LAN multicast network must be designed with Anycast RP. PIM-SM RP at the main or the centralized core must establish an MSDP session with RP on each remote site to exchange distributed multicast source information and allow RPs to join SPT to active sources as needed.

[Figure 3-31](#) shows an example of a community college LAN multicast network design.

Figure 3-31 Community College LAN Multicast Network Design

CCVE PIM-SM Network Design



228496

Dynamic Group Membership Design

Multicast receiver registration is done via IGMP protocol signaling. IGMP is an integrated component of an IP multicast framework that allows the receiver hosts and transmitting sources to be dynamically added to and removed from the network. Without IGMP, the network is forced to flood rather than multicast the transmissions for each group. IGMP operates between a multicast receiver host in the access layer and the Layer 3 router at the distribution layer.

The multicast system role changes when the access layer is deployed in the multilayer and routed access models. Because multilayer access switches do not run PIM, it becomes complex to make forwarding decisions out of the receiver port. In such a situation, Layer 2 access switches flood the traffic on all ports. This multilayer limitation in access switches is solved by using the IGMP snooping feature, which is enabled by default and is recommended to not be disabled.

IGMP is still required when a Layer 3 access layer switch is deployed in the routed access network design. Because the Layer 3 boundary is pushed down to the access layer, IGMP communication is limited between a receiver host and the Layer 3 access switch. In addition to the unicast routing protocol, PIM-SM must be enabled at the Layer 3 access switch to communicate with RPs in the network.

Designing Multicast Security

When designing multicast security in the community college LAN design, two key concerns are preventing a rogue source and preventing a rogue PIM-RP.

Preventing Rogue Source

In a PIM-SM network, an unwanted traffic source can be controlled with the **pim accept-register** command. When the source traffic hits the first-hop router, the first-hop router (DR) creates the (S,G) state and sends a PIM source register message to the RP. If the source is not listed in the accept-register filter list (configured on the RP), the RP rejects the register and sends back an immediate Register-Stop message to the DR. The drawback with this method of source filtering is that with the **pim accept-register** command on the RP, the PIM-SM (S,G) state is still created on the first-hop router of the source. This can result in traffic reaching receivers local to the source and located between the source and the RP. Furthermore, because the **pim accept-register** command works on the control plane of the RP, this can be used to overload the RP with fake register messages and possibly cause a DoS condition.

Preventing Rogue PIM-RP

Like the multicast source, any router can be misconfigured or can maliciously advertise itself as a multicast RP in the network with the valid multicast group address. With a static RP configuration, each PIM-enabled router in the network can be configured to use static RP for the multicast source and override any other Auto-RP or BSR multicast router announcement from the network.

QoS for Application Performance Optimization

The function and guaranteed low latency bandwidth expectation of network users and endpoints has evolved significantly over the past few years. Application and device awareness has become a key tool in providing differentiated service treatment at the campus LAN edge. Media applications, and particularly video-oriented media applications, are evolving as the education community enters the digital era of delivering education, as well as the increased campus network and asset security requirements. Integrating video applications in the community college LAN network exponentially increases bandwidth utilization and fundamentally shifts traffic patterns. Business drivers behind this media application growth include remote learning, as well as leveraging the network as a platform to build an energy-efficient network to minimize cost and go “green”. High-definition media is transitioning from the desktop to conference rooms, and social networking phenomena are crossing over into educational settings. Besides internal and college research applications, media applications are fueling a new wave of IP convergence, requiring the ongoing development of converged network designs.

Converging media applications onto an IP network is much more complex than converging voice over IP (VoIP) alone. Media applications are generally bandwidth-intensive and bursty (as compared to VoIP), and many different types of media applications exist; in addition to IP telephony, applications can include live and on-demand streaming media applications, digital signage applications, high-definition room-based conferencing applications, as well as an infinite array of data-oriented applications. By

embracing media applications as the next cycle of convergence, community college IT departments can think holistically about their network design and its readiness to support the coming tidal wave of media applications, and develop a network-wide strategy to ensure high quality end-user experiences.

The community college LAN infrastructure must set the administrative policies to provide differentiated forwarding services to the network applications, users and endpoints to prevent contention. The characteristic of network services and applications must be well understood, so that policies can be defined that allow network resources to be used for internal applications, to provide best-effort services for external traffic, and to keep the network protected from threats.

The policy for providing network resources to an internal application is further complicated when interactive video and real-time VoIP applications are converged over the same network that is switching mid-to-low priority data traffic. Deploying QoS technologies in the campus allows different types of traffic to contend inequitably for network resources. Real-time applications such as voice, interactive, and physical security video can be given priority or preferential services over generic data applications, but not to the point that data applications are starving for bandwidth.

Community College LAN QoS Framework

Each group of managed and un-managed applications with unique traffic patterns and service level requirements requires a dedicated QoS class to provision and guarantee these service level requirements. The community college LAN network architect may need to determine the number of classes for various applications, as well as how should these individual classes should be implemented to deliver differentiated services consistently in main and remote college campus sites. Cisco recommends following relevant industry standards and guidelines whenever possible, to extend the effectiveness of your QoS policies beyond your direct administrative control.

With minor changes, the community college LAN QoS framework is developed based on RFC4594 that follows industry standard and guidelines to function consistently in heterogeneous network environment. These guidelines are to be viewed as industry best-practice recommendations. Community college and service providers are encouraged to adopt these marking and provisioning recommendations, with the aim of improving QoS consistency, compatibility, and interoperability. However, because these guidelines are not standards, modifications can be made to these recommendations as specific needs or constraints require. To this end, to meet specific business requirements, Cisco has made a minor modification to its adoption of RFC 4594, namely the switching of call-signaling and broadcast video markings (to CS3 and CS5, respectively).

RFC 4594 outlines twelve classes of media applications that have unique service level requirements, as shown in [Figure 3-32](#).

Figure 3-32 Community College LAN Campus 12-Class QoS Policy Recommendation

Application Class	Media Application Examples	PHB	Admission Control	Queuing and Dropping
VoIP Telephony	Cisco IP Phone	EF	Required	Priority Queue (PQ)
Broadcast Video	Cisco IPVS, Enterprise TV	CS5	Required	(Optional) PQ
Real-Time Interactive	Cisco TelePresence	CS4	Required	(Optional) PQ
Multimedia Conferencing	Cisco CUPC, WebEx	AF4	Required	BW Queue + DSCP WRED
Multimedia Streaming	Cisco DMS, IP/TV	AF3	Recommended	BW Queue + DSCP WRED
Network Control	EIGRP, OSPF, HSRP, IKE	CS6		BW Queue
Call-Signaling	SCCP, SIP, H.323	CS3		BW Queue
Ops/Admin/Mgmt (OAM)	SNMP, SSH, Syslog	CS2		BW Queue
Transactional Data	ERP Apps, CRM Apps	AF2		BW Queue + DSCP WRED
Bulk Data	E-mail, FTP, Backup	AF1		BW Queue + DSCP WRED
Best Effort	Default Class	DF		Default Queue + RED
Scavenger	YouTube, Gaming, P2P	CS1		Min BW Queue

226467

The twelve classes are as follows:

- *VoIP telephony*—This service class is intended for VoIP telephony (bearer-only) traffic (VoIP signaling traffic is assigned to the call-signaling class). Traffic assigned to this class should be marked EF. This class is provisioned with expedited forwarding (EF) per-hop behavior (PHB). The EF PHB-defined in RFC 3246 is a strict-priority queuing service and, as such, admission to this class should be controlled (admission control is discussed in the following section). Examples of this type of traffic include G.711 and G.729a.
- *Broadcast video*—This service class is intended for broadcast TV, live events, video surveillance flows, and similar *inelastic* streaming video flows, which are highly drop sensitive and have no retransmission and/or flow control capabilities. Traffic in this class should be marked class selector 5 (CS5) and may be provisioned with an EF PHB; as such, admission to this class should be controlled. Examples of this traffic include live Cisco Digital Media System (DMS) streams to desktops or to Cisco Digital Media Players (DMPs), live Cisco Enterprise TV (ETV) streams, and Cisco IP Video Surveillance.
- *Real-time interactive*—This service class is intended for (inelastic) room-based, high-definition interactive video applications and is intended primarily for voice and video components of these applications. Whenever technically possible and administratively feasible, data sub-components of this class can be separated out and assigned to the transactional data traffic class. Traffic in this class should be marked CS4 and may be provisioned with an EF PHB; as such, admission to this class should be controlled. A sample application is Cisco TelePresence.
- *Multimedia conferencing*—This service class is intended for desktop software multimedia collaboration applications and is intended primarily for voice and video components of these applications. Whenever technically possible and administratively feasible, data sub-components of this class can be separated out and assigned to the transactional data traffic class. Traffic in this class should be marked assured forwarding (AF) Class 4 (AF41) and should be provisioned with a guaranteed bandwidth queue with Differentiated Services Code Point (DSCP)-based Weighted Random Early Detection (WRED) enabled. Admission to this class should be controlled;

additionally, traffic in this class may be subject to policing and re-marking. Sample applications include Cisco Unified Personal Communicator, Cisco Unified Video Advantage, and the Cisco Unified IP Phone 7985G.

- *Multimedia streaming*—This service class is intended for video-on-demand (VoD) streaming video flows, which, in general, are more elastic than broadcast/live streaming flows. Traffic in this class should be marked AF Class 3 (AF31) and should be provisioned with a guaranteed bandwidth queue with DSCP-based WRED enabled. Admission control is recommended on this traffic class (though not strictly required) and this class may be subject to policing and re-marking. Sample applications include Cisco Digital Media System VoD streams.
- *Network control*—This service class is intended for network control plane traffic, which is required for reliable operation of the enterprise network. Traffic in this class should be marked CS6 and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, because network control traffic should not be dropped (if this class is experiencing drops, the bandwidth allocated to it should be re-provisioned). Sample traffic includes EIGRP, OSPF, Border Gateway Protocol (BGP), HSRP, Internet Key Exchange (IKE), and so on.
- *Call-signaling*—This service class is intended for signaling traffic that supports IP voice and video telephony. Traffic in this class should be marked CS3 and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, because call-signaling traffic should not be dropped (if this class is experiencing drops, the bandwidth allocated to it should be re-provisioned). Sample traffic includes Skinny Call Control Protocol (SCCP), Session Initiation Protocol (SIP), H.323, and so on.
- *Operations/administration/management (OAM)*—This service class is intended for network operations, administration, and management traffic. This class is critical to the ongoing maintenance and support of the network. Traffic in this class should be marked CS2 and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, because OAM traffic should not be dropped (if this class is experiencing drops, the bandwidth allocated to it should be re-provisioned). Sample traffic includes Secure Shell (SSH), Simple Network Management Protocol (SNMP), Syslog, and so on.
- *Transactional data (or low-latency data)*—This service class is intended for interactive, “foreground” data applications (foreground refers to applications from which users are expecting a response via the network to continue with their tasks; excessive latency directly impacts user productivity). Traffic in this class should be marked AF Class 2 (AF21) and should be provisioned with a dedicated bandwidth queue with DSCP-WRED enabled. This traffic class may be subject to policing and re-marking. Sample applications include data components of multimedia collaboration applications, Enterprise Resource Planning (ERP) applications, Customer Relationship Management (CRM) applications, database applications, and so on.
- *Bulk data (or high-throughput data)*—This service class is intended for non-interactive “background” data applications (background refers to applications from which users are not awaiting a response via the network to continue with their tasks; excessive latency in response times of background applications does not directly impact user productivity). Traffic in this class should be marked AF Class 1 (AF11) and should be provisioned with a dedicated bandwidth queue with DSCP-WRED enabled. This traffic class may be subject to policing and re-marking. Sample applications include E-mail, backup operations, FTP/SFTP transfers, video and content distribution, and so on.
- *Best effort (or default class)*—This service class is the default class. The vast majority of applications will continue to default to this best-effort service class; as such, this default class should be adequately provisioned. Traffic in this class is marked default forwarding (DF or DSCP 0) and should be provisioned with a dedicated queue. WRED is recommended to be enabled on this class.

- *Scavenger (or low-priority data)*—This service class is intended for non-business-related traffic flows, such as data or video applications that are entertainment and/or gaming-oriented. The approach of a less-than Best-Effort service class for non-business applications (as opposed to shutting these down entirely) has proven to be a popular, political compromise. These applications are permitted on enterprise networks, as long as resources are always available for business-critical voice, video, and data applications. However, as soon as the network experiences congestion, this class is the first to be penalized and aggressively dropped. Traffic in this class should be marked CS1 and should be provisioned with a minimal bandwidth queue that is the first to starve should network congestion occur. Sample traffic includes YouTube, Xbox Live/360 movies, iTunes, BitTorrent, and so on.

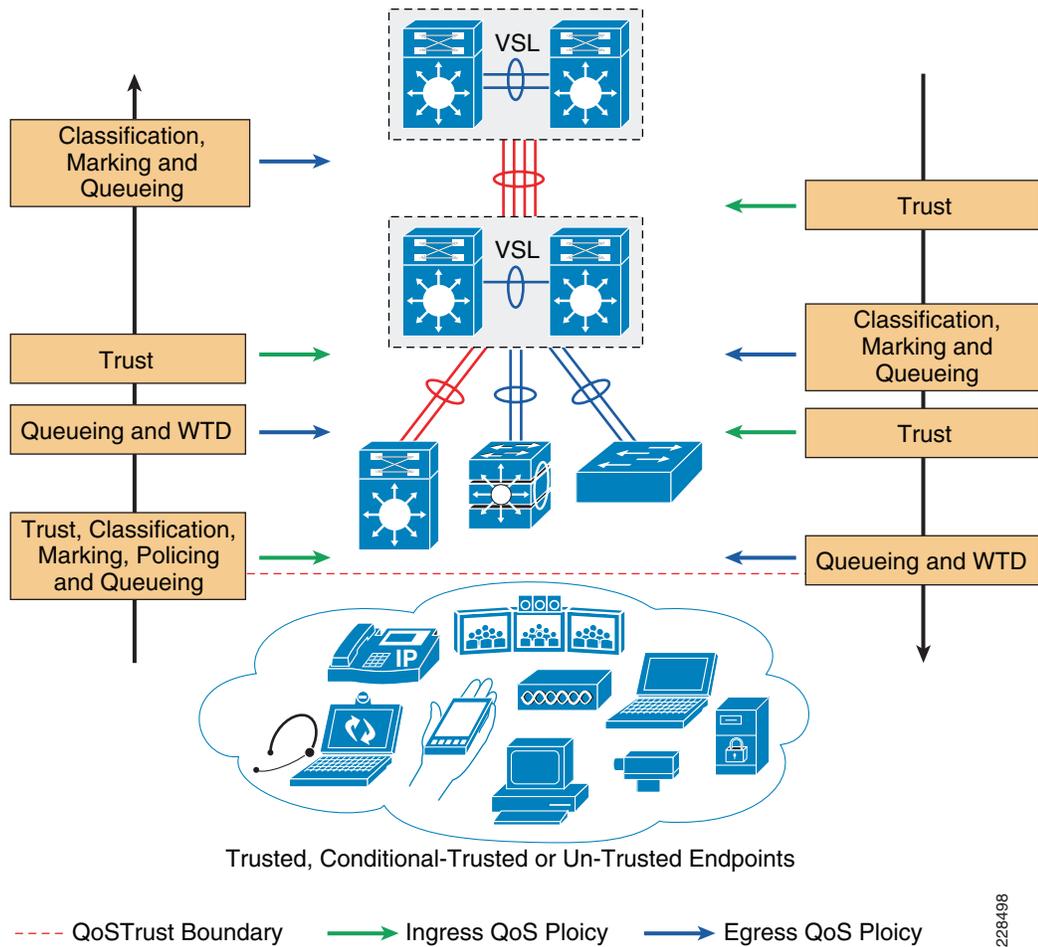
Designing Community College LAN QoS Trust Boundary and Policies

To build an end-to-end QoS framework that offers transparent and consistent QoS service without compromising performance, it is important to create a blueprint of the network, classifying a set of trusted applications, devices, and forwarding paths; and then define common QoS policy settings independent of how QoS is implemented within the system.

QoS settings applied at the LAN network edge sets the ingress rule based on deep packet classification and marks the traffic before it is forwarded inside the campus core. To retain the marking set by access layer switches, it is important that other LAN network devices in the college campus trust the marking and apply the same policy to retain the QoS settings and offer symmetric treatment. Bi-directional network communication between applications, endpoints, or other network devices requires the same treatment when traffic enters or leaves the network, and must be taken into account when designing the trust model between network endpoints and core and edge campus devices.

The trust or un-trust model simplifies the rules for defining bi-directional QoS policy settings. [Figure 3-33](#) shows the QoS trust model setting that sets the QoS implementation guidelines in community college campus networks.

Figure 3-33 Campus QoS Trust and Policies



Community College LAN QoS Best Practices

With an overall application strategy in place, end-to-end QoS policies can be designed for each device and interface, as determined by their roles in the network infrastructure. However, because the Cisco QoS toolset provides many QoS design and deployment options, a few succinct design principles can help simplify strategic QoS deployments, as discussed in the following sections.

Hardware versus Software QoS

A fundamental QoS design principle is to always enable QoS policies in hardware rather than software whenever possible. Cisco IOS routers perform QoS in software, which places incremental loads on the CPU, depending on the complexity and functionality of the policy. Cisco Catalyst switches, on the other hand, perform QoS in dedicated hardware application-specific integrated circuits (ASICs) on Ethernet-based ports, and as such do not tax their main CPUs to administer QoS policies. This allows complex policies to be applied at line rates even up to Gigabit or 10-Gigabit speeds.

Classification and Marking Best Practices

When classifying and marking traffic, a recommended design principle is to classify and mark applications as close to their sources as technically and administratively feasible. This principle promotes end-to-end differentiated services and PHBs.

In general, it is not recommended to trust markings that can be set by users on their PCs or other similar devices, because users can easily abuse provisioned QoS policies if permitted to mark their own traffic. For example, if an EF PHB has been provisioned over the network, a PC user can easily configure all their traffic to be marked to EF, thus hijacking network priority queues to service non-realtime traffic. Such abuse can easily ruin the service quality of realtime applications throughout the college campus. On the other hand, if community college network administrator controls are in place that centrally administer PC QoS markings, it may be possible and advantageous to trust these.

Following this rule, it is recommended to use DSCP markings whenever possible, because these are end-to-end, more granular, and more extensible than Layer 2 markings. Layer 2 markings are lost when the media changes (such as a LAN-to-WAN/VPN edge). There is also less marking granularity at Layer 2. For example, 802.1P supports only three bits (values 0–7), as does Multiprotocol Label Switching Experimental (MPLS EXP). Therefore, only up to eight classes of traffic can be supported at Layer 2, and inter-class relative priority (such as RFC 2597 Assured Forwarding Drop Preference markdown) is not supported. Layer 3-based DSCP markings allow for up to 64 classes of traffic, which provides more flexibility and is adequate in large-scale deployments and for future requirements.

As the network border blurs between enterprise and education community network and service providers, the need for interoperability and complementary QoS markings is critical. Cisco recommends following the IETF standards-based DSCP PHB markings to ensure interoperability and future expansion. Because the community college voice, video, and data applications marking recommendations are standards-based, as previously discussed, community colleges can easily adopt these markings to interface with service provider classes of service.

Policing and Markdown Best Practices

There is little reason to forward unwanted traffic that gets policed and drop by a subsequent tier node, especially when unwanted traffic is the result of DoS or worm attacks in the college network. Excessive volume attack traffic can destabilize network systems, which can result in outages. Cisco recommends policing traffic flows as close to their sources as possible. This principle applies also to legitimate flows, because worm-generated traffic can masquerade under legitimate, well-known TCP/UDP ports and cause extreme amounts of traffic to be poured into the network infrastructure. Such excesses should be monitored at the source and marked down appropriately.

Whenever supported, markdown should be done according to standards-based rules, such as RFC 2597 (AF PHB). For example, excess traffic marked to AFx1 should be marked down to AFx2 (or AFx3 whenever dual-rate policing such as defined in RFC 2698 is supported). Following such markdowns, congestion management policies, such as DSCP-based WRED, should be configured to drop AFx3 more aggressively than AFx2, which in turn should be dropped more aggressively than AFx1.

Queuing and Dropping Best Practices

Critical media applications require uncompromised performance and service guarantees regardless of network conditions. Enabling outbound queuing in each network tier provides end-to-end service guarantees during potential network congestion. This common principle applies to campus-to-WAN/Internet edges, where speed mismatches are most pronounced; and campus interswitch links, where oversubscription ratios create the greater potential for network congestion.

Because each application class has unique service level requirements, each should be assigned optimally a dedicated queue. A wide range of platforms in varying roles exist in community college networks, so each must be bounded by a limited number of hardware or service provider queues. No fewer than four queues are required to support QoS policies for various types of applications, specifically as follows:

- Realtime queue (to support a RFC 3246 EF PHB service)
- Guaranteed-bandwidth queue (to support RFC 2597 AF PHB services)
- Default queue (to support a RFC 2474 DF service)
- Bandwidth-constrained queue (to support a RFC 3662 scavenger service)

Additional queuing recommendations for these classes are discussed next.

Strict-Priority Queuing Recommendations

The realtime or strict priority class corresponds to the RFC 3246 EF PHB. The amount of bandwidth assigned to the realtime queuing class is variable. However, if the majority of bandwidth is provisioned with strict priority queuing (which is effectively a FIFO queue), the overall effect is a dampening of QoS functionality, both for latency- and jitter-sensitive realtime applications (contending with each other within the FIFO priority queue), and also for non-realtime applications (because these may periodically receive significant bandwidth allocation fluctuations, depending on the instantaneous amount of traffic being serviced by the priority queue). Remember that the goal of convergence is to enable voice, video, and data applications to transparently co-exist on a single community college network infrastructure. When realtime applications dominate a link, non-realtime applications fluctuate significantly in their response times, destroying the transparency of the converged network.

For example, consider a 45 Mbps DS3 link configured to support two Cisco TelePresence CTS-3000 calls with an EF PHB service. Assuming that both systems are configured to support full high definition, each such call requires 15 Mbps of strict-priority queuing. Before the TelePresence calls are placed, non-realtime applications have access to 100 percent of the bandwidth on the link; to simplify the example, assume there are no other realtime applications on this link. However, after these TelePresence calls are established, all non-realtime applications are suddenly contending for less than 33 percent of the link. TCP windowing takes effect and many applications hang, timeout, or become stuck in a non-responsive state, which usually translates into users calling the IT help desk to complain about the network (which happens to be functioning properly, albeit in a poorly-configured manner).



Note

As previously discussed, Cisco IOS software allows the abstraction (and thus configuration) of multiple strict priority LLQs. In such a multiple LLQ context, this design principle applies to the sum of all LLQs to be within one-third of link capacity.

It is vitally important to understand that this strict priority queuing rule is simply a best practice design recommendation and is not a mandate. There may be cases where specific business objectives cannot be met while holding to this recommendation. In such cases, the community college network administrator must provision according to their detailed requirements and constraints. However, it is important to recognize the tradeoffs involved with over-provisioning strict priority traffic and its negative performance impact, both on other realtime flows and also on non-realtime-application response times.

And finally, any traffic assigned to a strict-priority queue should be governed by an admission control mechanism.

Best Effort Queuing Recommendation

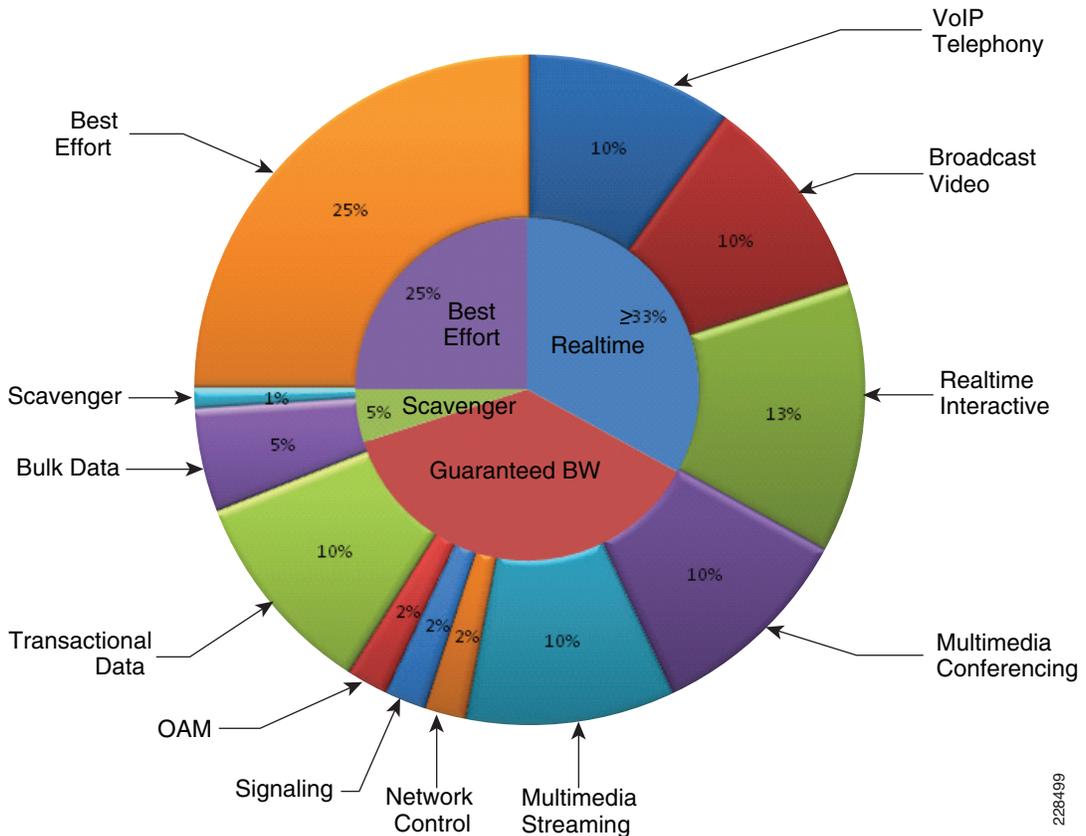
The best effort class is the default class for all traffic that has not been explicitly assigned to another application-class queue. Only if an application has been selected for preferential/deferential treatment is it removed from the default class. Because most community colleges may have several types of applications running in networks, adequate bandwidth must be provisioned for this class as a whole to handle the number and volume of applications that default to it. Therefore, Cisco recommends reserving at least 25 percent of link bandwidth for the default best effort class.

Scavenger Class Queuing Recommendations

Whenever the scavenger queuing class is enabled, it should be assigned a minimal amount of link bandwidth capacity, such as 1 percent, or whatever the minimal bandwidth allocation that the platform supports. On some platforms, queuing distinctions between bulk data and scavenger traffic flows cannot be made, either because queuing assignments are determined by class of service (CoS) values (and both of these application classes share the same CoS value of 1), or because only a limited amount of hardware queues exist, precluding the use of separate dedicated queues for each of these two classes. In such cases, the scavenger/bulk queue can be assigned a moderate amount of bandwidth, such as 5 percent.

These queuing rules are summarized in Figure 3-34, where the inner pie chart represents a hardware or service provider queuing model that is limited to four queues and the outer pie chart represents a corresponding, more granular queuing model that is not bound by such constraints.

Figure 3-34 Compatible 4-Class and 12-Class Queuing Models



228499

High-Availability in LAN Network Design

Network reliability and availability is not a new demand, but is well planned during the early network design phase. To prevent a catastrophic network failure during an unplanned network outage event, it is important to identify network fault domains and define rapid recovery plans to minimize the application impact during minor and major network outage conditions.

Because every tier of the LAN network design can be classified as a fault domain, deploying redundant systems can be effective. However, this introduces a new set of challenges, such as higher cost and the added complexity of managing more systems. Network reliability and availability can be simplified using several Cisco high availability technologies that offer complete failure transparency to the end users and applications during planned or unplanned network outages.

Cisco high availability technologies can be deployed based on critical versus non-critical platform roles in the network. Some of the high availability techniques can be achieved with the LAN network design inherent within the community college network design, without making major network changes. However, the critical network systems that are deployed in the main campus that provide global connectivity may require additional hardware and software components to provide non-stop communications. The following three major resiliency requirements encompass most of the common types of failure conditions; depending on the LAN design tier, the resiliency option appropriate to the role and network service type must be deployed:

- *Network resiliency*—Provides redundancy during physical link failures, such as fiber cut, bad transceivers, incorrect cabling, and so on.
- *Device resiliency*—Protects the network during abnormal node failure triggered by hardware or software, such as software crashes, a non-responsive supervisor, and so on.
- *Operational resiliency*—Enables resiliency capabilities to the next level, providing complete network availability even during planned network outage conditions, using In Service Software Upgrade (ISSU) features.

Community College LAN Design High-Availability Framework

This high availability framework is based on the three major resiliency strategies described in the previous section. Several high availability technologies must be deployed at each layer to provide higher network availability and rapid recovery during failure conditions, to prevent communication failure or degraded network-wide application performance. (See [Figure 3-35](#).)

Figure 3-35 Community College LAN Design High-Availability Goals, Strategy, and Technologies

Resilient Goal	Network Service Availability		
Resilient Strategies	Network Resiliency	Device Resiliency	Operational Resiliency
Resilient Technologies	EtherChannel/MEC UDLD IP Event Dampening	NSF/SSO Stack Wise	ISSU eFSU

228500

Network Resiliency Best Practices

The most common network fault occurrence in the LAN network is a link failure between two systems. Link failures can be caused by issues such as a fiber cut, miswiring, and so on. Redundant parallel physical links between two systems can increase availability, but also change how overall higher layer protocols construct the adjacency and loop-free forwarding topology to the parallel physical paths.

Deploying redundant parallel paths in the recommended community college LAN design by default develops a non-optimal topology that keeps the network under-utilized and requires protocol-based network recovery. In the same network design, the routed access model eliminates such limitations and enables the full load balancing capabilities to increase bandwidth capacity and minimize the application impact during a single path failure. To develop a consistent network resiliency service in the centralized main and remote college campus sites, the following basic principles apply:

- Deploying redundant parallel paths are the basic requirement to employ network resiliency at any tier. It is critical to simplify the control plane and forwarding plane operation by bundling all physical paths into a single logical bundled interface (EtherChannel). Implement a defense-in-depth approach to failure detection and recovery mechanisms. An example of this is configuring the UniDirectional Link Detection (UDLD) protocol, which uses a Layer 2 keep-alive to test that the switch-to-switch links are connected and operating correctly, and acts as a backup to the native Layer 1 unidirectional link detection capabilities provided by 802.3z and 802.3ae standards. UDLD is not an EtherChannel function; it operates independently over each individual physical port at Layer 2 and remains transparent to the rest of the port configuration. Therefore, UDLD can be deployed on ports implemented in Layer 2 or Layer 3 modes.
- Ensure that the design is self-stabilizing. Hardware or software errors may cause ports to flap, which creates false alarms and destabilizes the network topology. Implementing route summarization advertises a concise topology view to the network, which prevents core network instability. However, within the summarized boundary, the flood may not be protected. Deploy IP event dampening as a tool to prevent the control and forwarding plane impact caused by physical topology instability.

These principles are intended to be a complementary part of the overall structured modular design approach to the campus design, and serve primarily to reinforce good resilient design practices.

Device Resiliency Best Practices

Another major component of an overall campus high availability framework is providing device or node level protection that can be triggered during any type of abnormal internal hardware or software process within the system. Some of the common internal failures are a software-triggered crash, power outages, line card failures, and so on. LAN network devices can be considered as a single-point-of-failure and are considered to be major failure condition because the recovery type may require a network administrator to mitigate the failure and recover the system. The network recovery time can remain undeterministic, causing complete or partial network outage, depending on the network design.

Redundant hardware components for device resiliency vary between fixed configuration and modular Cisco Catalyst switches. To protect against common network faults or resets, all critical community college campus network devices must be deployed with a similar device resiliency configuration. This subsection provides basic redundant hardware deployment guidelines at the access layer and collapsed core switching platforms in the campus network.

Redundant Power System

Redundant power supplies for network systems protect against power outages, power supply failures, and so on. It is important not only to protect the internal network system but also the endpoints that rely on power delivery over the Ethernet network. Redundant power systems can be deployed in the two following configuration modes:

- *Modular switch*—Dual power supplies can be deployed in modular switching platforms such as the Cisco Catalyst 6500 and 4500-E Series platforms. By default, the power supply operates in redundant mode, offering the 1+1 redundant option. Overall power capacity planning must be done to dynamically allow for network growth. Lower power supplies can be combined to allocate power to all internal and external resources, but may not be able to offer power redundancy.
- *Fixed configuration switch*—The power supply in fixed configuration switches can be internal or use Cisco RPS 2300 external power supplies. A single Cisco RPS 2300 power supply uses a modular power supply and fan for flexibility, and can deliver power to multiple switches. Deploying an internal and external power supply solution protects critical access layer switches during power outages, and provides complete fault transparency and constant network availability.

Redundant Control Plane

Device or node resiliency in modular Cisco Catalyst 6500/4500 platforms and Cisco StackWise provides a 1+1 redundancy option with enterprise-class high availability and deterministic network recovery time. The following sub-sections provide high availability design details, as well as graceful network recovery techniques that do not impact the control plane and provide constant forwarding capabilities during failure events.

Stateful Switchover

The stateful switchover (SSO) capability in modular switching platforms such as the Cisco Catalyst 4500 and 6500 provides complete carrier-class high availability in the campus network. Cisco recommends distribution and core layer design model be the center point of the entire college communication network. Deploying redundant supervisors in the mission-critical distribution and core system provides non-stop communication throughout the network. To provide 99.999 percent service availability in the access layer, the Catalyst 4500 must be equipped with redundant supervisors to critical endpoints, such as Cisco TelePresence.

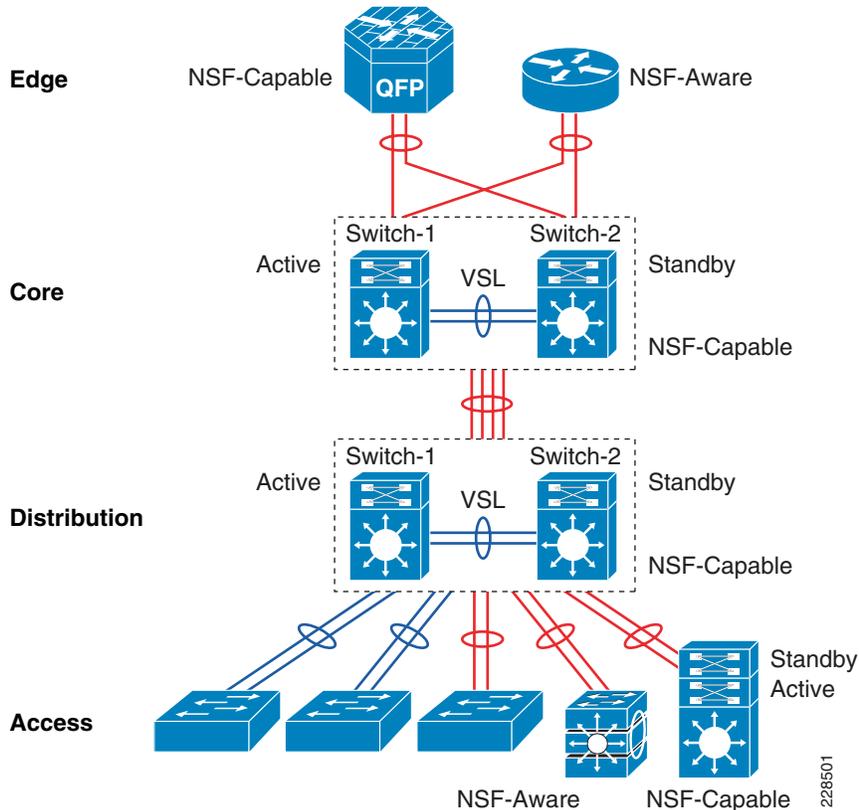
Cisco StackWise is a low-cost solution to provide device-level high availability. Cisco StackWise is designed with unique hardware and software capabilities that distribute, synchronize, and protect common forwarding information across all member switches in a stack ring. During master switch failure, the new master switch re-election remains transparent to the network devices and endpoints. Deploying Cisco StackWise according to the recommended guidelines protects against network interruption, and recovers the network in sub-seconds during master switch re-election.

Bundling SSO with NSF capability and the awareness function allows the network to operate without errors during a primary supervisor module failure. Users of realtime applications such as VoIP do not hang up the phone, and IP video surveillance cameras do not freeze.

Non-Stop Forwarding

Cisco VSS and the single highly resilient-based campus system provides uninterrupted network availability using non-stop forwarding (NSF) without impacting end-to-end application performance. The Cisco VSS and redundant supervisor system is an NSF-capable platform; thus, every network device that connects to VSS or the redundant supervisor system must be NSF-aware to provide optimal resiliency. By default, most Cisco Layer 3 network devices are NSF-aware systems that operate in NSF helper mode for graceful network recovery. (See Figure 3-36.)

Figure 3-36 Community College LAN Design NSF/SSO Capable and Aware Systems



Operational Resiliency Best Practices

Designing the network to recover from failure events is only one aspect of the overall campus non-stop design. Converged network environments are continuing to move toward requiring true 7x24x365 availability. The community college LAN network is part of the backbone of the college network and must be designed to enable standard operational processes, configuration changes, and software and hardware upgrades without disrupting network services.

The ability to make changes, upgrade software, and replace or upgrade hardware becomes challenging without a redundant system in the campus core. The ability to upgrade individual devices without taking them out of service is similarly based on having internal component redundancy (such as with power supplies and supervisors), complemented with the system software capabilities. The Cisco Catalyst 4500 and 6500 support realtime upgrade software in the campus.

Catalyst 4500—ISSU

Full-image ISSU on the Cisco Catalyst 4500 leverages dual supervisors to allow for a full, in-place Cisco IOS upgrade, such as moving from 12.2(50)SG to 12.2(53)SG for example. This leverages the NSF/SSO capabilities of the switch and provides for less than 200 msec of traffic loss during a full Cisco IOS upgrade.

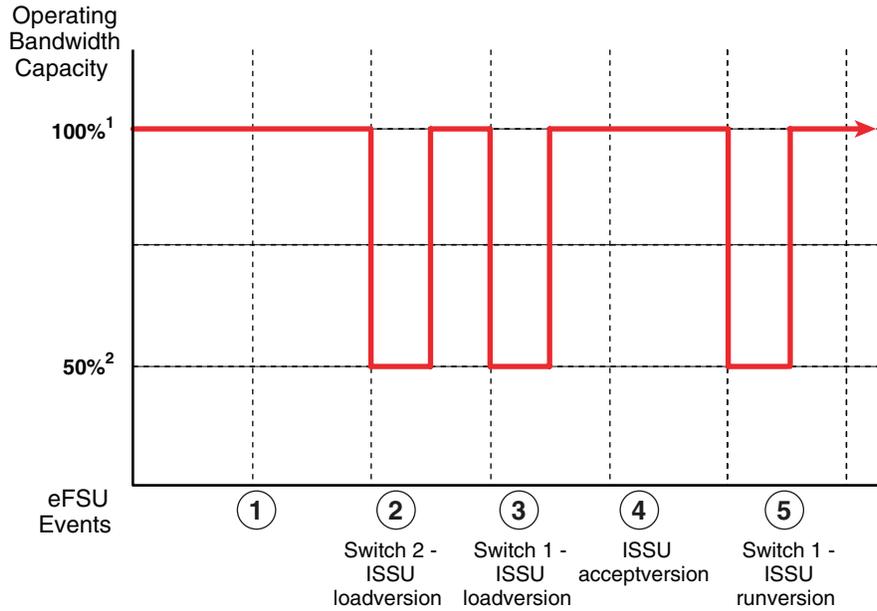
Having the ability to operate the campus as a non-stop system depends on the appropriate capabilities being designed-in from the start. Network and device level redundancy, along with the necessary software control mechanisms, guarantee controlled and fast recovery of all data flows following any network failure, while concurrently providing the ability to proactively manage the non-stop infrastructure.

Catalyst 6500 VSS—eFSU

A network upgrade requires planned network and system downtime. VSS offers unmatched network availability to the core. With the Enhanced Fast Software Upgrade (eFSU) feature, the VSS can continue to provide network services during the upgrade. With the eFSU feature, the VSS network upgrade remains transparent and hitless to the applications and end users (see [Figure 3-37](#)). Because eFSU works in conjunction with NSF/SSO technology, the network devices can gracefully restore control and forwarding information during the upgrade process, while the bandwidth capacity operates at 50 percent and the data plane can converge within sub-seconds.

For a hitless software update, the ISSU process requires three sequential upgrade events for error-free software install on both virtual switch systems. Each upgrade event causes traffic to be re-routed to a redundant MEC path, causing sub-second traffic loss that does not impact realtime network applications, such as VoIP.

Figure 3-37 Real-time Network Upgrade Events and Availability



¹Both virtual-system nodes forwarding traffic
²Single virtual-system node forwarding traffic

228502

Summary

Designing the LAN network aspects for the community college network design establishes the foundation for all other aspects within the service fabric (WAN, security, mobility, and UC) as well as laying the foundation to provide safety and security, operational efficiencies, virtual learning environments, and secure classrooms.

This chapter reviews the two LAN design models recommended by Cisco, as well as where to apply these models within the various locations of a community college network. Each of the layers is discussed and design guidance is provided on where to place and how to deploy these layers. Finally, key network foundation services such as routing, switching, QoS, multicast, and high availability best practices are given for the entire community college design.

