



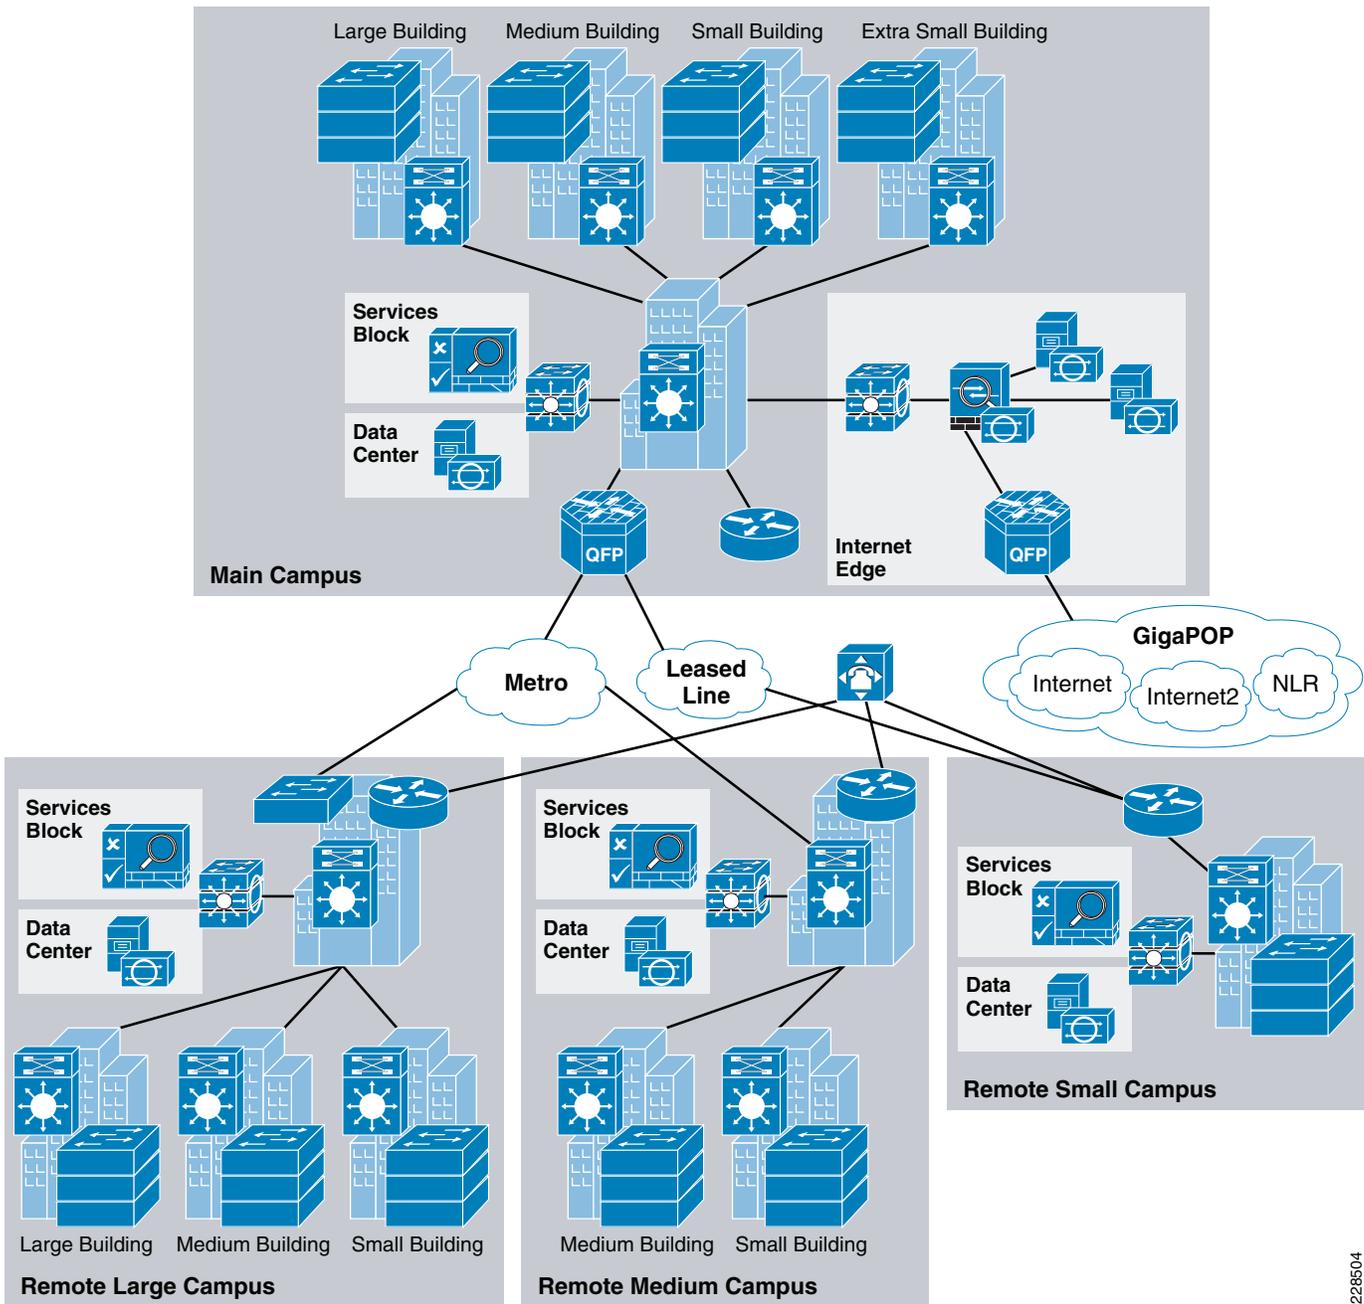
CHAPTER 4

Community College WAN Design Considerations

WAN Design

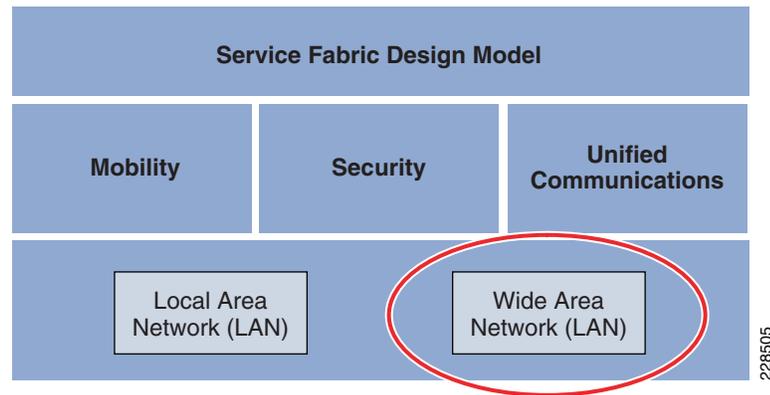
The Cisco Community College reference design is a multi-campus design where a campus consists of multiple buildings and services. The campuses are interconnected through various WAN transports as shown in [Figure 4-1](#).

Figure 4-1 Community College WAN Design Diagram



228504

Within the Community College reference design, the service fabric network provides the foundation on which all the solutions and services are built upon to solve the business challenges facing community colleges. These challenges include virtual learning, secure connected classrooms, and safety and security. This service fabric consists of four distinct components as shown in [Figure 4-2](#).

Figure 4-2 The Service Fabric Design Model

This chapter discusses the WAN design component of the community college service fabric design. This section discusses how the WAN design is planned for community colleges, the assumptions made, the platforms chosen, and the justification for choosing a platform. The WAN design is highly critical to provide network access for remote campus locations to the main campus site, as well as connectivity between community colleges, and general Internet access for the entire college. The WAN design should not be viewed merely for providing access, but mainly to see how the business requirements can be met. In today's collaborative learning environment, it is important for communication to exist between students and teachers. This communication could be with voice, video, or data applications. Moreover, the video applications, may possess, flavors ranging from desktop video to real-time video. To provide this collaborative environment, highly resilient and, highly performing WAN designs are required.

The main components of Community College WAN design are as follows:

- WAN transport
- WAN devices
- Network Foundation services—Routing, QoS, and multicast

WAN Transport

This section discusses the different WAN transports present in the community college.

Private WAN Service

One of the main requirements for community colleges is the ability to collaborate with other colleges within North America and globally. To achieve the inter connectivity between the colleges, the network should be connected to certain providers, such as Lambda rail, Internet2. The community colleges need to connect to Gigapops—regional networks, which provide access to these private WAN networks. The following sections provide a brief description on these two network types:

Internet2 is a not-for-profit advanced networking consortium comprising more than 200 U.S. universities in cooperation with 70 leading corporations, 45 government agencies, laboratories, and other institutions of higher learning as well over 50 international partner organizations. Internet2 provides its members both leading-edge network capabilities and unique partnership opportunities that together facilitate the development, deployment and use of revolutionary Internet technologies. The Internet2 network's physical implementation is comprised of an advanced IP network, virtual circuit network and core optical network. It provides the necessary scalability for member institutions to efficiently provision

resources to address bandwidth-intensive requirements of their campuses such as, collaborative applications, distributed research experiments, grid-based data analysis and social networking. For more information on the Internet2 network, refer to the following URL:

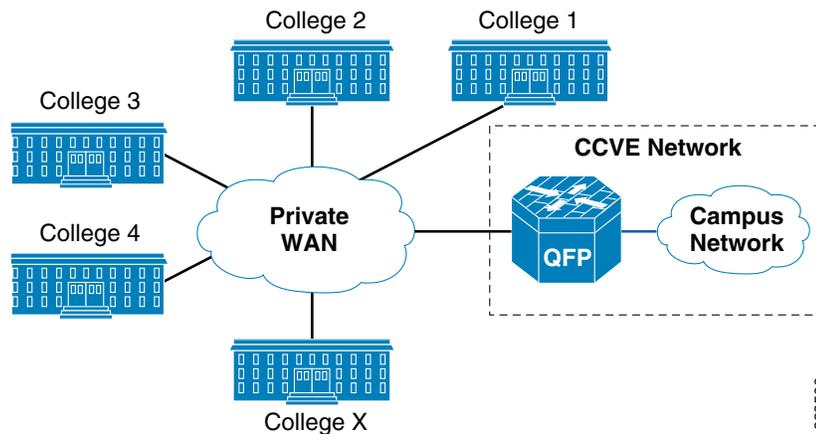
<http://www.internet2.edu/network/>

National LambdaRail (NLR) is a high-speed, fiber-optic network infrastructure linking over 30 cities in 21 states. It is owned by the U.S. research and education community and is dedicated to serving the needs of researchers and research groups. NLR's high-performance network backbone offers unrestricted usage and bandwidth, a choice of cutting-edge network services and applications, and customized service for individual researchers and projects. NLR services include high-capacity 10Gb Ethernet LAN-PHY or OC-192 lambdas, point-to-point or multipoint Ethernet-based transport, routed IP-based services, and TelePresence video-conference services. For more information on the NLR network and its services, refer to the following URL:

<http://www.nlr.net/>

This design assumes that community colleges are connected to one of these networks using either Layer 2 or Layer 3 networks for WAN connectivity, using WAN speeds of 100Mbps. The physical connection is assumed to be one connection to the service provider, but there will be two logical connections—one for accessing private networks, and the second one for Internet access. Figure 4-3 depicts how community college would connect to different colleges, universities, and research networks using either NLR or Internet2 service.

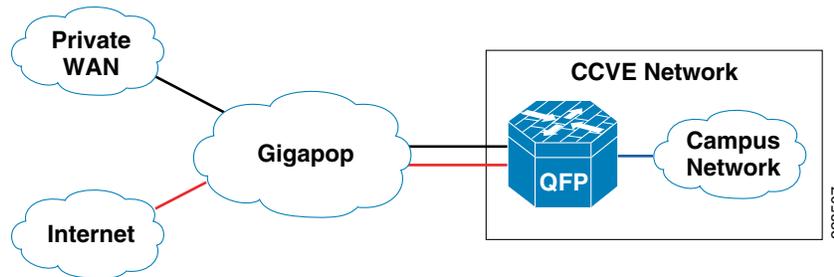
Figure 4-3 Community College Connection to Other Colleges Using Private WAN



Internet Service

The physical connection for reaching the Internet and the private WAN network is same; however, both circuits are logically separated using different sub-interfaces. Therefore, it is similar to a situation where a customer is connected to different service providers. See Figure 4-4.

Figure 4-4 Community College Internet Service



Metro Service

Metro Ethernet is one of the fastest growing WAN transport technologies in the telecommunications industry. The advantages of using this WAN transport are as follows:

- Scalability and reachability
 - The services offered would scale from 1Mbps to 10Gbps, and beyond in granular increments, which makes this transport highly scalable.
 - Service providers worldwide are migrating their networks to provide metro services; thereby, it is available at large number of places.
- Performance, QoS, and suitability for convergence
 - Inherently Ethernet networks require less processing to operate and manage and operate at higher bandwidth than other technologies.
 - The granular options in bandwidth, ability to provide different SLA based on voice, video, and data applications that provide QoS service to customers.
 - Low latency and delay variation make it the best solution for video, voice, and data.
- Cost savings
 - Metro Ethernet brings the cost model of Ethernet to the WAN.
- Expediting and enabling new applications
 - Accelerates implementations with reduced resources for overburdened IT departments.
 - Enables new applications requiring high bandwidth, and low latency that were previously not possible or prohibited by high cost.

There are two popular methods of service for Metro Ethernet:

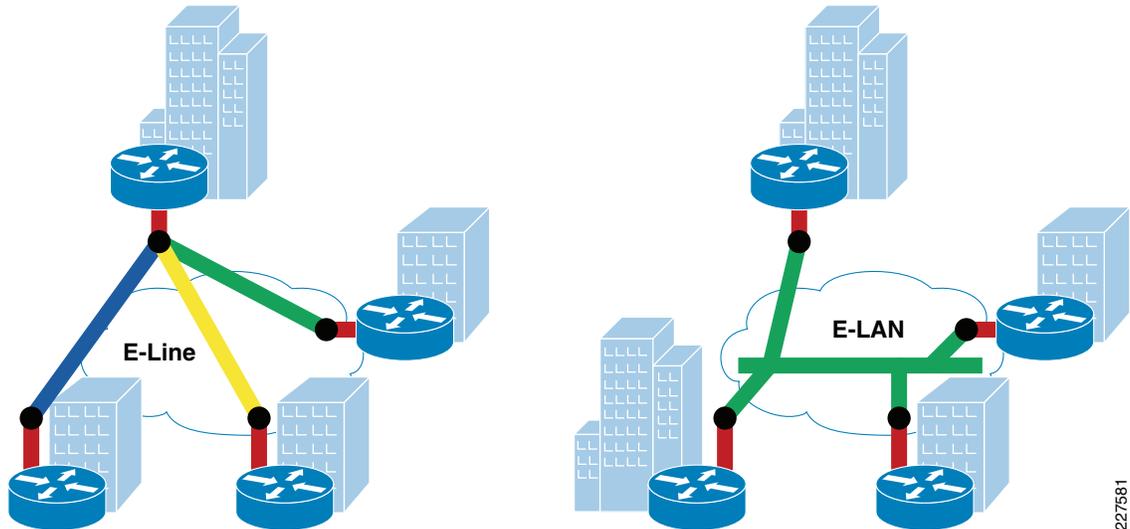
1. E-line, which is also known as Ethernet Virtual Private Line (EVPL) provides a point-to-point service.
2. E-LAN which provides multipoint or any-to-any connectivity.

EVPL, like Frame Relay, provides for multiplexing multiple point-to-point connections over a single physical link. In the case of Frame Relay, the access link is a serial interface to a Frame Relay switch with individual data-link connection identifiers (DLCIs), identifying the multiple virtual circuits or connections. In the case of EVPL, the physical link is Ethernet, typically FastEthernet or Gigabit Ethernet, and the multiple circuits are identified as VLANs by way of an 802.1q trunk.

E-LAN, also known as Virtual Private LAN Services (VPLS), provides any-to-any connectivity within the Metro area, which allows flexibility. It passes 802.q trunks across the SP network known as Q-in-Q.

Figure 4-5 shows the difference between these services.

Figure 4-5 *Different Services Available*

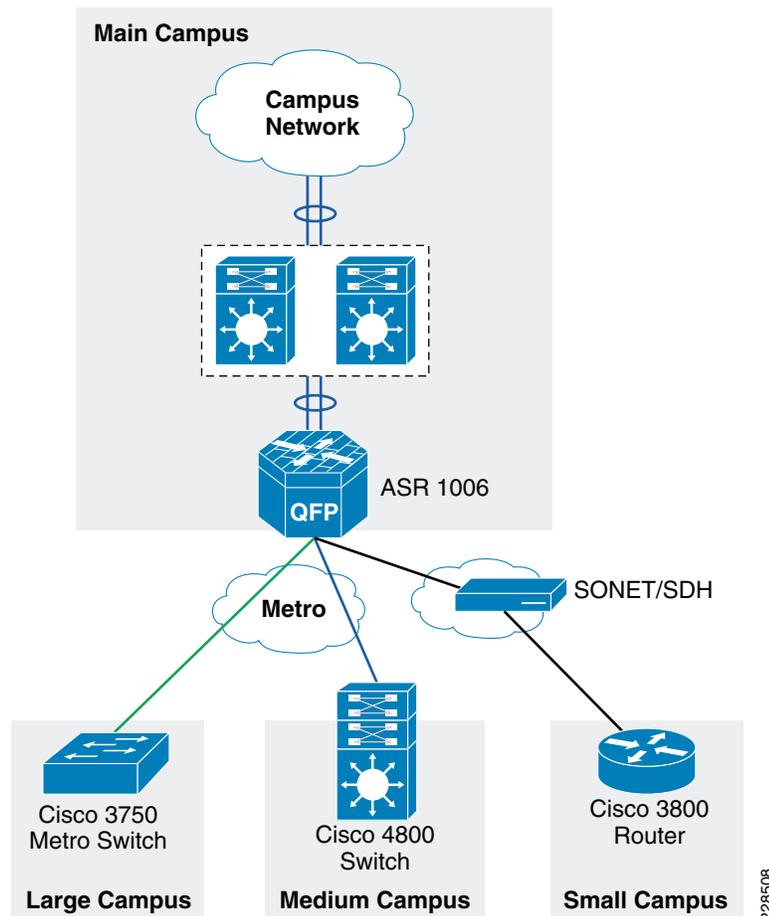


This section discusses how the Metro service is designed in the Community College reference design. The Metro service is used to provide connectivity between the remote campuses to the main campus site. The key reasons for recommending Metro service for community college are as follows:

- *Centralized administration and management*—E-line service provides point-to-point connectivity, where as, E-LAN provides point-to-multipoint connectivity. Having a point-to-point connectivity mandates that all the remote campus sites need to traverse the main campus site to reach the other, making the centralized administration applicable.
- *Performance*—Since all the application services are centrally located at main campus site, the WAN bandwidth required for remote campus sites to main campus site should be at least 100 Mbps. The Metro transport can provide 100Mbps, and more if needed in the future.

Therefore, in this design, it is recommended that the remote large and remote medium campus locations use E-line service to connect to the main campus site. [Figure 4-6](#) shows how the remote campus locations are connected to main campus site using Metro service.

Figure 4-6 The Metro Transport Deployment in Community College WAN Design



Leased-Line Service

The WAN bandwidth requirement for a small remote site is assumed to be 20Mbps. Cisco recommends that the remote small campus site connect to the main campus site using a private leased-line service. The leased-line service is more readily available for these type of locations and the bandwidth is sufficient for the remote small campus application requirements.

WAN Aggregation Platform Selection in the Community College Reference Design

In addition to selecting the WAN service for connectivity between college campus locations and access to the Internet, choosing the appropriate WAN aggregation router is essential. For each location in the Community College reference design, various WAN aggregation platforms are selected based on the requirements.

Main Campus WAN Aggregation Platform Selection

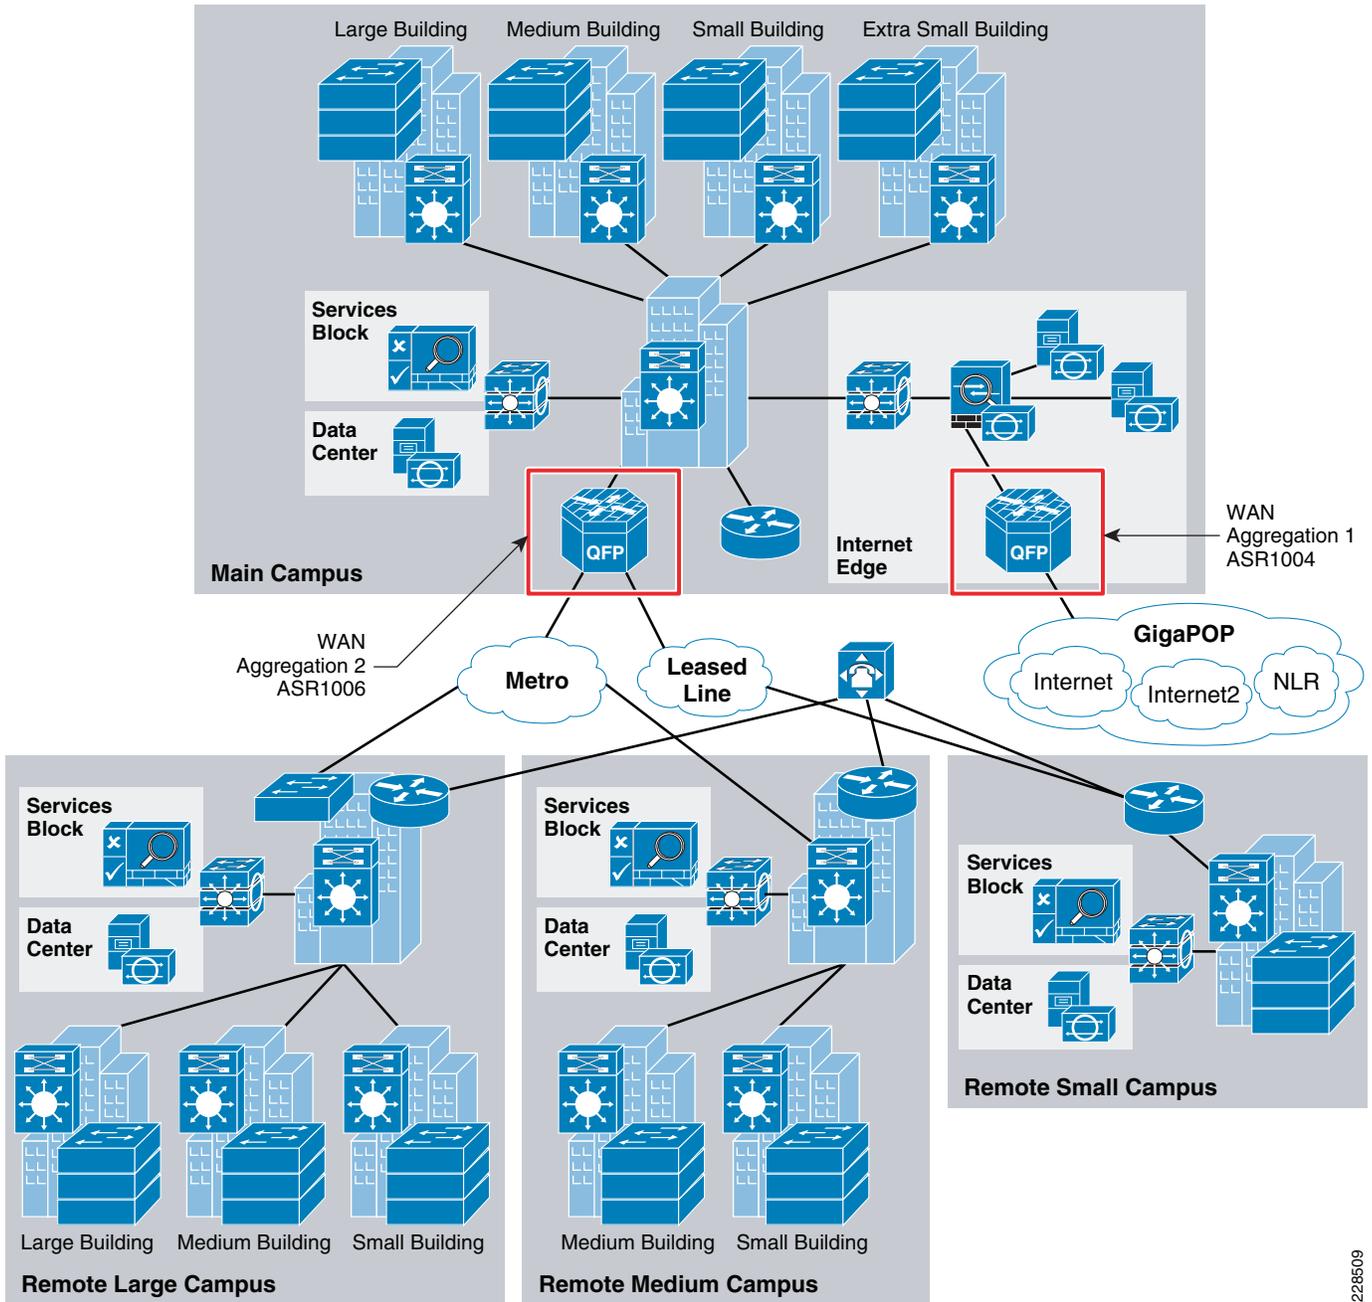
A WAN aggregation router aggregates all the incoming WAN circuits from various locations in the network as well as the Internet and also provides the proper QoS required for application delivery. Cisco recommends the Cisco ASR family of routers as the WAN aggregation platform for the main campus location.

The Cisco ASR 1000 Series Router family consists of three different models:

- The Cisco ASR 1002 Router is a 3-SPA, 2-rack-unit (RU) chassis with one Embedded Services Processor (ESP) slot that comes with an integrated Router Processor (RP), integrated Cisco ASR 1000 Series Shared Port Adapter Interface Processor (SIP), and integrated four Gigabit Ethernet ports.
- The Cisco ASR 1004 Router is an 8-SPA, 4-RU chassis with one ESP slot, one RP slot, and two SIP slots.
- The Cisco ASR 1006 Router is a 12-SPA, 6-RU, hardware redundant chassis with two ESP slots, two RP slots and three SIP slots.

In community college WAN design, there are two places where the WAN aggregation occurs in the main campus location. The first place is where the main campus location connects to outside world using private WAN and Internet networks. The second place is where all the remote campus locations connect to the main campus sites. [Figure 4-7](#) shows the two different WAN aggregation devices.

Figure 4-7 The WAN Aggregation Points in Community College



228509

WAN Aggregation 1

Cisco ASR 1004 Series router is recommended as WAN aggregation platform for private WAN/Internet connectivity. This choice was made considering the cost and required features—performance, QoS, routing, and resiliency, which are essential requirements for WAN aggregation router. Moreover, this platform contains built-in resiliency capabilities such as ISSU and IOS-based redundancy.

WAN Aggregation 2

The second WAN aggregation device provides connectivity to the large and medium remote community college campuses. To perform this aggregation, the Cisco ASR 1006 router with redundant route processors and redundant ESP's has been recommended for the following reasons:

- *Performance*—Up to 20 Gbps throughput
- *Port density*—Up to 12 shared port adapters (SPAs), the highest port density solution of the three Cisco ASR 1000 routers
- *Resiliency*—Cisco ASR 1006 router supports hardware redundancy and in-service software upgrades (ISSU). This chassis would support dual route processors, and dual ESP modules to support the hardware redundancy. Moreover, this router would also support EtherChannel load balancing feature.

Remote Large Campus WAN Aggregation Platform Selection

The WAN connectivity between the large remote campus sites to the main campus site is fairly simpler because of the lack of requirements of advanced encryption technologies. Therefore, the main idea is to reduce the cost and try to consolidate the WAN functionality into the distribution device at the large campus site. However, at the large campus site, as per the campus LAN design document VSS has been chosen as distribution switch, and it does not support WAN functionality. Therefore, a dedicated WAN aggregation device needed to perform that functionality, and the choice can be an ASR, 7200, or 3750ME switches. Out of these choices, considering the cost/performance criteria, the Cisco 3750ME switch was selected to perform the WAN aggregation. The Cisco 3750 Metro switch has the following features/capabilities to adequately meet the requirements:

- Hierarchical QoS
- Routing support: OSPF, EIGRP, BGP
- Multicast support: PIM
- Redundant power supply

Remote Medium Campus WAN Aggregation Platform Selection

As discussed in [Chapter 3, “Community College LAN Design Considerations,”](#) the remote medium campus collapses the WAN edge and core-layer LAN functionality into a single switch to provide cost effectiveness to meet the budget needs for this size location. The remote medium campus location is connected to the main campus location through Metro service. At the remote medium campus location, the WAN and LAN aggregation platform is the Cisco Catalyst 4507 switch. This switch has necessary features to perform as WAN router. However, if there is the need for advanced WAN features such as MPLS, the Cisco Catalyst 3750 ME or Cisco ISR Series router or upgrading to the Cisco Catalyst 6500 series could be explored as an option. For this design, the Cisco Catalyst 4500 Series switches has been chosen to perform the dual functionality as WAN router, in addition to its role as core-layer LAN switch.

Remote Small Campus WAN Aggregation Platform Selection

The remote small campus is connected to main campus using a private leased-line service. The WAN speed between the remote small campus and the main campus location is assumed to be around 20Mbps, and this service is provided by a traditional leased line. Since it is a leased-line circuit, WAN devices such as Cisco 3750 Metro or 4507 switch can not be used. Therefore, an integrated services router is needed to meet the requirement. For this reason, the Cisco 3845 Series router is chosen as WAN platform for remote small campus. The main advantages of using the Cisco 3845 Series router are as follows:

- Enhanced Network Module Slot
- Support for over 90 existing and new modules
- Voice Features: Analog and digital voice call support and optional voice mail support
- Support for majority of existing AIMS, NMs, WICs, VWICs, and VICs
- Integrated GE ports with copper and fiber support

Network Foundation Services

The key network foundation services are as follows:

- Routing
- QoS
- Resiliency
- Multicast

Routing Design

This section discusses how routing is designed in Community College reference WAN design. As indicated in the WAN transport design, the Community College reference design has multiple transports—NLR or I2 networks, Internet, Metro Service, and leased-line services. The NLR or I2 networks would provide access to reach other community colleges, universities, and research networks globally. Internet service would help the Community College to reach Internet. Metro/leased-line service would help to connect remote campus locations to the main campus. To provide connectivity using these transport services we have designed two distinct routing domains – external and internal. The external routing domain is where the Community College would connect with external autonomous system, and the internal routing domain is where the entire routing domain is within single autonomous system. The following section would discuss about the external routing domain design, and the internal routing domain design.

External Routing Domain

As indicated above, the external routing domain would connect with different service providers, NLR or I2, and the Internet service. This is applicable only to the WAN aggregation router 1, which interfaces with both NLR or I2, and the Internet service, because it the only router which interfaces with the external domain.

The main design considerations for routing for the Internet/private WAN edge router are as follows:

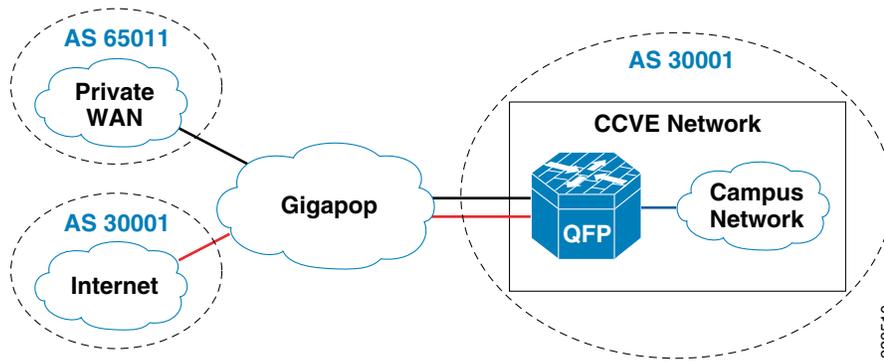
- Scale up to large number of routes
- Support for multi-homing—connection to different service providers
- Ability to implement complex polices—Have separate policies for incoming and outgoing traffic

To meet the above requirements, BGP has been chosen as the routing protocol because of the following reasons:

- *Scalability*—BGP is far superior when routing table entries is quite large.

- *Complex policies*—IGP protocol is better in environments where the neighbors are trusted, whereas when dealing with different service providers' complex policies are needed to deal with incoming entries, and outgoing entries. BGP supports having different policies for incoming and outgoing prefixes. Figure 4-8 shows the BGP design.

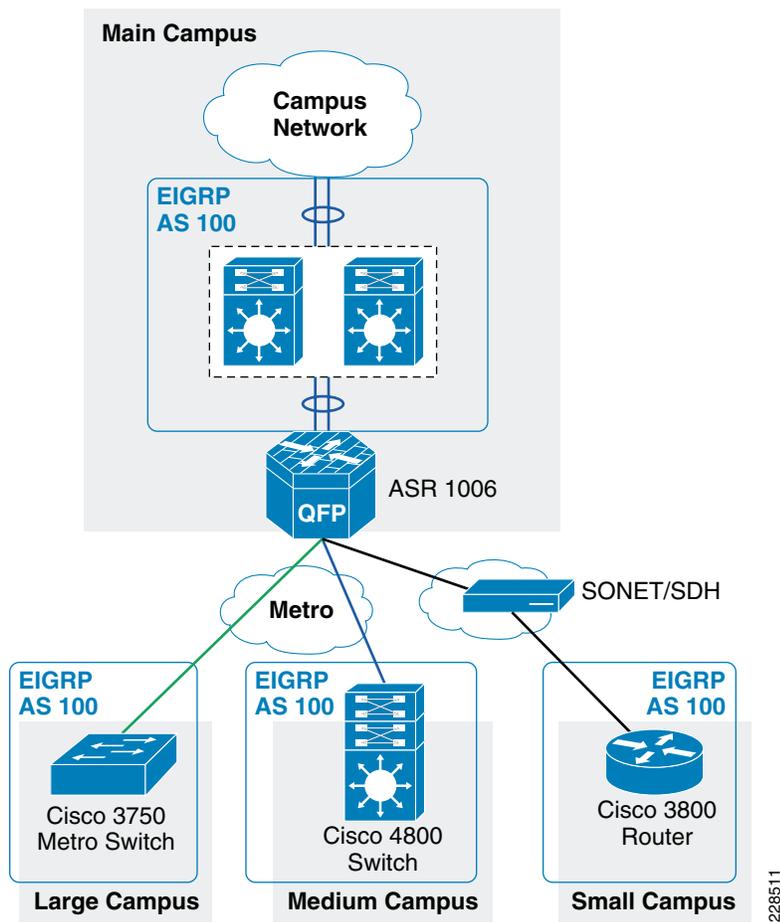
Figure 4-8 BGP Design in Community College



Internal Routing Domain

EIGRP is chosen as the routing protocol for designing the internal routing domain, which is basically connecting all the devices in the campus network. EIGRP is a balanced hybrid routing protocol that builds neighbor adjacency and flat routing topology on per autonomous-system (AS)-basis. It is important to design EIGRP routing domain in college infrastructure with all the design principles defined earlier in this section. CCVE SRA network infrastructure must be deployed in recommended EIGRP protocol design to secure, simplify, and optimize the network performance. Figure 4-9 depicts the design of EIGRP for internal network.

Figure 4-9 EIGRP Design Diagram



To obtain more information about EIGRP design, refer to the “[Community College Network Foundation Services Design](#)” section on page 3-25.

QoS

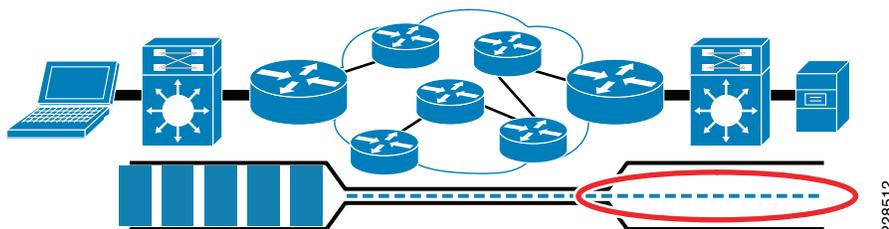
QoS is a part of foundation services, which is very critical to the application performance. Today’s networks are rapidly converging into IP network. The traditional applications, which used the networks, were voice, video, and data. However, broadcast video, real-time video, video surveillance, and many other applications have all converged into IP networks. Moreover, each of these applications require different performance characteristics on the network. For example, data applications may need only high throughput, but are tolerant to delay and loss. Similarly, voice applications need constant low bandwidth and low delay performance. To cater to these performance characteristics, Cisco IOS has several rich QoS tools such as classification and marking, queuing, WRED, policing, shaping, and many other tools to effect the traffic characteristics. Before discussing the QoS design, the following subsection provides a brief introduction on these characteristics.

Traffic Characteristics

The main traffic characteristics are bandwidth, delay, loss, and jitter.

- *Bandwidth*—Lack of proper bandwidth can cause applications from performing poorly. This problem would be exacerbated if there were more centralized applications. The bandwidth constraint occurs because of the difference between the bandwidth available at LAN and the WAN. As shown in Figure 4-10, the bandwidth of the WAN transport dictates the amount of traffic received at each remote site. Applications are constrained by the amount of WAN bandwidth.

Figure 4-10 Bandwidth Constraint Due to Difference in Speeds



- *Jitter*—Occurs when there are bandwidth mismatches between the sender, and receiver which could result in poor performance of delay sensitive applications like voice, and video.
- *Loss*—occurs when the queues become full, and there is not enough bandwidth to send the packets.
- *Delay*—Is an important characteristic, which plays a large role in determining the performance of the applications. For a properly designed voice network the one-way delay must be less then 150 msec.

QoS Design for WAN Devices

For any application regardless of whether it is video, voice, or data the traffic characteristics just mentioned need to be fully understood before making any decisions on WAN transport or the platforms needed to deploy these services. Cisco QoS tools help to optimize these characteristics so that voice, video, and data applications performance is optimized. The voice and video applications are highly delay-and drop-sensitive, but the difference lies in the bandwidth requirement. The voice applications have a constant and low bandwidth requirement, but the video applications have variable bandwidth requirements. Therefore, it is important to have a good QoS policy to accommodate these applications.

Regardless of the WAN transport chosen, QoS design is the most significant factor in determining the success of network deployment. There are number of benefits in deploying a consistent, coherent QoS scheme across all network layers. It helps not only in optimizing the network performance, it helps to mitigate network attacks, and also manage the control plane traffic. Therefore, when the platforms are selected at each network layer, QoS must always be considered in the design choice.

In the WAN links the congestion can occur when there are speed mismatches. This may occur because there is significant difference between LAN speeds and WAN speeds. To prevent that from occurring, the following two major tools can be used:

- Low-Latency Queuing (LLQ), which is used for highest-priority traffic (voice/ video).
- Class-based Weighted-Fair Queuing (CBWFQ), which can be used for guaranteeing bandwidth to data applications.

The general guidelines for deploying the WAN edge device considerations are as follows:

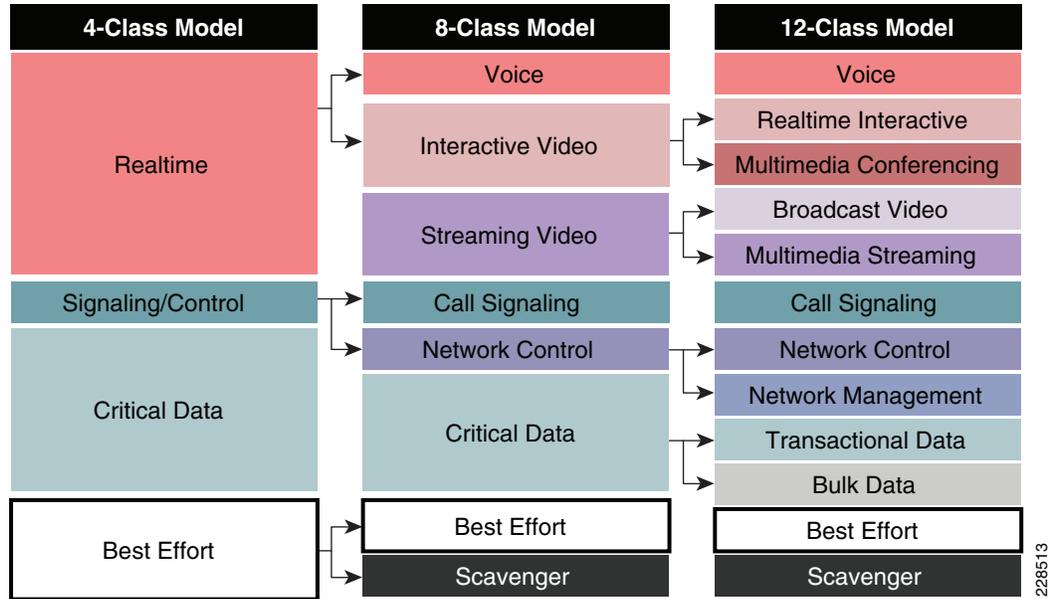
- For WAN speeds between 1Mbps to 100Mbps, use hierarchical policies for sub-line-rate Ethernet connections to provide shaping and CBWFQ/LLQ.
- For WAN speeds between 100Mbps to 10Gbps, use ASR1000 with QFP or hardware queuing via Cisco Catalyst 3750-Metro and 6500/7600 WAN modules.

When designing the QoS for WAN architecture, there are two main considerations to start with:

- Whether the service provider will provide four classes of traffic.
- The service provider will only provide one class of traffic.

This document assumes that the service provider will support at least 4 classes of traffic such as REAL_TIME, GOLD, SILVER, and DEFAULT. The community college campus LAN supports 12 classes of traffic, which will be mapped to 4 classes of traffic on the WAN side. Figure 4-11 illustrates the recommended markings for different application traffic.

Figure 4-11 Mapping of 12-Class Model to 4-classes



Once the QoS policy is designed the next pertinent question is the appropriate allocation of bandwidth for the 4 classes of traffic. Table 4-1 describes the different classes, and the percentage, and actual bandwidth allocated for each class of traffic.

Table 4-1 Classes of Traffic

Class of Traffic	4-class SP Model	Bandwidth Allocated	Actual Bandwidth
Voice, Broadcast Video, Real Time Interactive	SP- Real-Time	30%	30 Mbps
Network Control Signaling Transactional Data	SP-Critical 1	20%	20 Mbps
Multi-media Conferencing Multimedia streaming OAM	SP-Critical 2	20%	20 Mbps
Bulk data Scavenger Best Effort	SP-Best Effort	30%	30 Mbps

Redundancy

Redundancy must be factored into the WAN design for a number of reasons. Since the WAN may span across several service provider networks, it is likely that network will be subjected to different kinds of failures occurring all the time. One of the following failures can occur over a period of time: route flaps, brownouts, fibers being cut, and device failures. The probability of these occurring over a short period of time is low, but the occurrence is highly likely over a long period of time. To meet these challenges, different kind of redundancy should be planned. The following are the some of the ways to support redundancy:

- NSF/SSO—For networks to obtain 99.9999% of availability, technologies such as NSF/SSO are needed. The NSF would route packets until route convergence is complete, where as SSO allows standby RP to take immediate control and maintain connectivity protocols.
- Service Software Upgrade (ISSU) allows software to be updated or modified while packet forwarding continues with minimal interruption.
- Ether channel load balancing—Enabling this feature provides link resiliency and load balancing of traffic. This feature is enabled on the WAN aggregation 2 device. [Figure 4-12](#) shows where this feature is enabled.

Figure 4-12 Link Resiliency

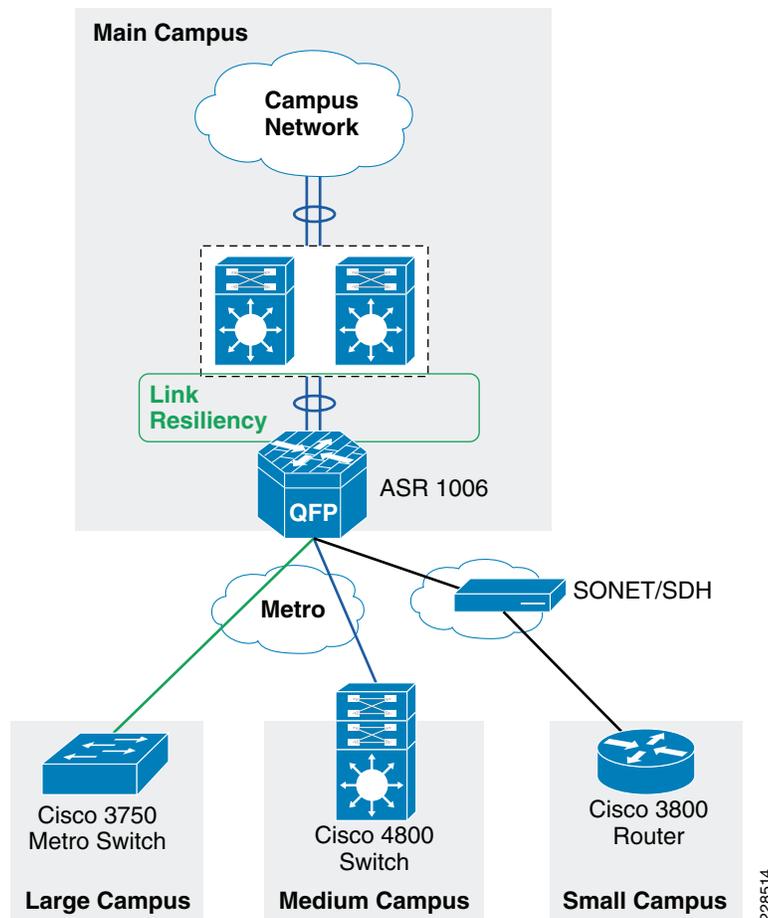


Table 4-2 shows the various WAN devices that are designed for resiliency.

Table 4-2 WAN Devices

Device	WAN transport	Resiliency feature
WAN aggregation 1	Private WAN/Internet	ISSU, IOS based redundancy
WAN aggregation 2	Metro	Redundant ESP, RP'

This section discusses how to incorporate the resiliency principle in Cisco Community College reference design for the WAN design. To enable resiliency adds cost and complexity to the design. Therefore, resiliency has been added at certain places where it is absolutely critical to the network architecture rather than designing redundancy at every place of the network.

In the Cisco Community College reference design the redundancy is planned at both WAN aggregation router1, and WAN aggregation router 2 in the main campus location. As explained in the WAN aggregation platform selection for the main campus location discussion ASR routers have been selected at both WAN aggregation locations places. However, we have different models, at both WAN aggregation places. When the ASR router interfaces with the private WAN, Internet networks the ASR 1004 with IOS-based redundancy has been chosen. Similarly, for the ASR router that interfaces with Metro connections, the ASR 1006 with dual RP, and dual ESP to provide for hardware-based redundancy has been chosen. Both of these models support In Service Software Upgrade (ISSU) capabilities to allow a user to upgrade Cisco IOS XE Software while the system remains in service. To obtain more information on ASR resiliency capabilities, see the ASR page at following URL: <http://www.cisco.com/go/asr1000>

Multicast

The main design considerations for multicast are as follows:

- The number of groups supported by the WAN edge device. This is scalability factor of the WAN edge device. The platform chosen must support the number of required groups.
- The placement of the RP—There are couple of options available with RP placement, which include Anycast with Static, Anycast with Auto-RP, or Anycast with BSR.
- Multicast protocols—PIM-Sparse mode, IGMP
- QoS policy must be configured for multicast traffic, so that this traffic does not affect the unicast traffic.

In the Community College reference design, we are assuming that multicast traffic would be present only within the campus, and not between the community colleges. Therefore, to obtain more information about multicast design for campus, refer to “[Community College Network Foundation Services Design](#)” section on page 3-25.

Summary

Designing the WAN network aspects for the Cisco Community College reference design interconnects the various LAN locations as well as lays the foundation to provide safety and security, operational efficiencies, virtual learning environments, and secure classrooms.

This chapter reviewed the WAN design models recommended by Cisco and where to apply these models within the various locations within a community college network. Key WAN design principles such as WAN aggregation platform selection, QoS, multicast and redundancy best practices were discussed for the entire community college design. Designing the WAN network of a community college using these

recommendations and best practices will establish a network that is resilient in case of failure, scalable for future growth, simplified to deploy and manage and cost efficient to meet the budget needs of a community college.