Small Enterprise Design Profile (SEDP)—WAN Design

This chapter discusses how to design and deploy WAN architecture for Small Enterprise Design Profile. The primary components of the WAN architecture are as follows:

- WAN technology
- Bandwidth capacity planning
- WAN IP addressing structure
- Routing
- QoS

WAN Design

The success of Small Enterprise Design depends on having an effective and better collaborative environment between employees located at different geographic locations. Some of the technologies that enhance this environment are interactive-video, on-demand video, voice and web collaboration, video to mobile devices, and TelePresence. To enable these technologies that foster the collaborative environment, WAN design is a critical consideration. There are several WAN technologies available today to provide WAN services, such as MPLS/VPN, Internet, and Metro Ethernet. See Figure 1.

Figure 1 WAN Technology Representation

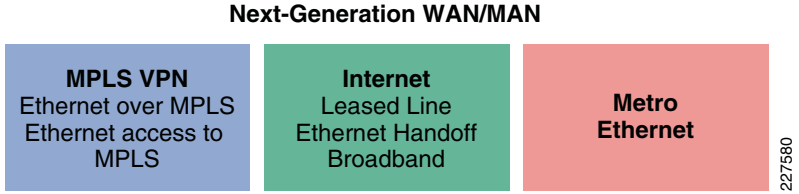

MPLS/VPN

The MPLS/VPN provides Layer-2 or Layer-3 VPN services. It provides the capability to an IP network infrastructure that delivers private network services over a shared network infrastructure. For more information about deploying MPLS/VPN, refer to the following URL: www.cisco.com/go/srnd

Internet

Out of the three WAN technologies, Internet is the cheapest and easiest to deploy. However, to deploy a VPN service over Internet requires an overlay VPN network such as DMVPN to provide secure VPN service.

Metro Ethernet

Metro Ethernet is one of the fastest growing transport technologies in the telecommunications industry. In the current WAN design, it is recommended to use Metro Ethernet Service as WAN transport between remote sites and main office. The following subsection describes the advantages of Metro service.

Why Metro Ethernet Service Is Needed

The deployment of carrier Ethernet services had raced to an astounding $7 billion by 2007* with predictions of more than $30 billion by 2012. This is driven by the following benefits to end users (IT, network, and applications departments) and service providers alike:

- *Scalability, ubiquity, and reachability*
  - Global availability of standardized services independent of physical access type dramatically reduce complexity and cost
- *Performance, QoS, and suitability for convergence*
  - Inherently, Ethernet networks require less processing to operate and manage at higher bandwidth than other technologies
  - Low latency and delay variation make it the best solution for video, voice, and data
- *Cost savings*
  - Carrier Ethernet brings the cost model of Ethernet to the wide area network WAN)
- *Control, simplicity, familiarity*
  - IT departments manage Ethernet networks every day and now are in control of their IP routed networks, worldwide
- *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

Types of Services Available in Metro Ethernet

The following are two popular methods of Metro service:

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

EVPL, like Frame Relay, provides 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.

One of the major advantage E-LAN provides is 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 2 shows the difference between E-line and E-LAN.

Figure 2 Different Services Available

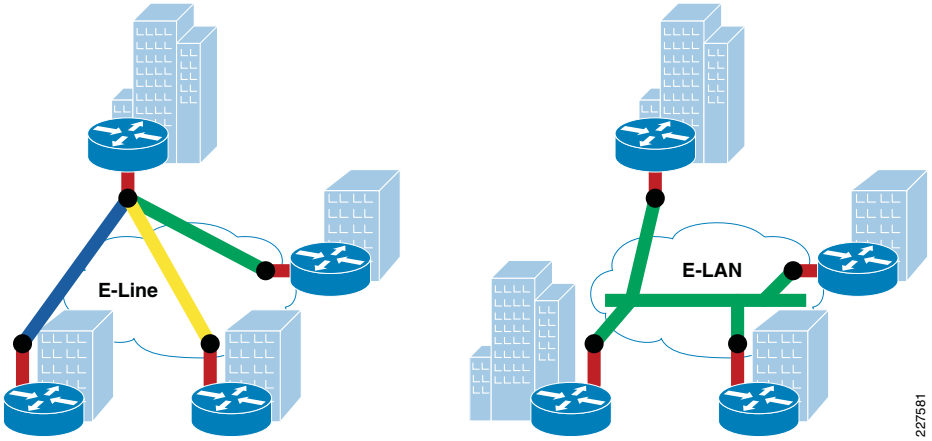

WAN Service Deployed in the Small Enterprise Design

In this version of the guide, E-line with point-to-point service is chosen as the model for connecting remote sites to main site. All the remote sites have a point-to-point connection to the main site. Each remote site is assumed to have 100Mpbs of Metro service to the service provider. As mentioned in the previous section, each circuit is represented by a VLAN using dot1q trunk. Figure 3 shows how this is implemented in this design.

Figure 3 EVPN Service Used in Remote WAN Architecture

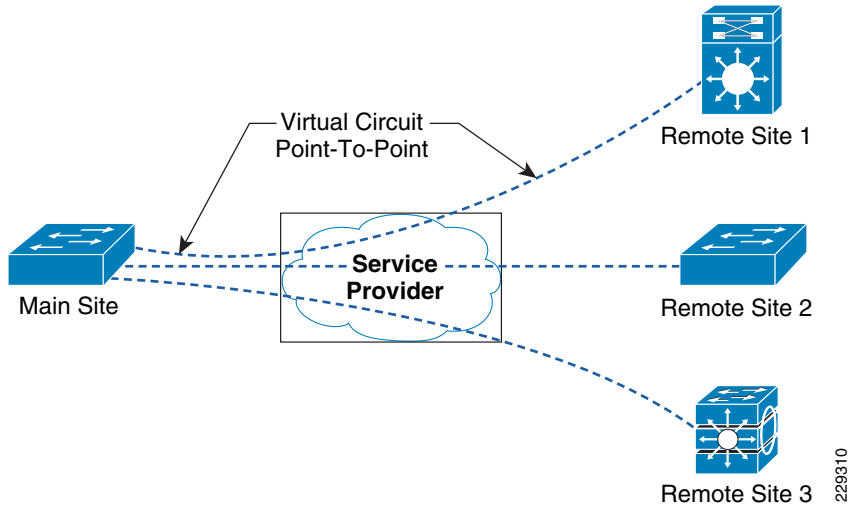

The following is the configuration of the WAN interface at the main site:

```
interface GigabitEthernet1/1/1
 description Connected to SP-MPLS-Core-cr24-6500-1
 switchport trunk native vlan 801
 switchport trunk allowed vlan 501-550
 switchport mode trunk
 logging event trunk-status
 load-interval 30
 carrier-delay msec 0
 priority-queue out
 mls qos trust dscp
 spanning-tree portfast trunk
 spanning-tree bpdufilter enable
 spanning-tree guard root
 service-policy output Remote-1to50-Parent-Policy-Map
 hold-queue 2000 in
 hold-queue 2000 out
```

As shown in the above configuration, the link is carrying 50 VLANs, which are connected to 50 small office sites. The solution shown above is one of the ways to deploy WAN service; there are other ways to deploy the WAN service. The next section discusses when to look at alternative to Metro service.

When to Consider WAN Technologies

In this design, all the remote sites are connected to the main site using point-to-point links that provide high bandwidth, symmetric traffic flows, and rich media services to all the sites. Since it is a point-to-point circuit, when the number of circuits increases to 1K and beyond, then we need not only to have 1K point-to-point circuits at the main site but also have 1K routing adjacency relationships, which will increase the CPU utilization. Therefore, when the scalability numbers increase more than 1K, an alternative Layer-3 WAN technology such as MPLS/VPN should be considered. MPLS/VPN Layer-3 technology deployments will not require different hardware, but only changing the WAN service. Figure 4 shows the routing behavior difference between Layer 2 (Metro service) and Layer 3 (MPLS/VPN service) from the routing protocol perspective.

Figure 4 Difference in Routing Adjacency Between Layer 2 and Layer 3 Service

Layer 2 (L2) Service

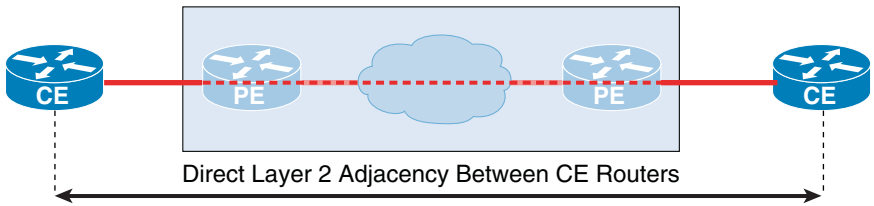

Layer 3 (L3) Service

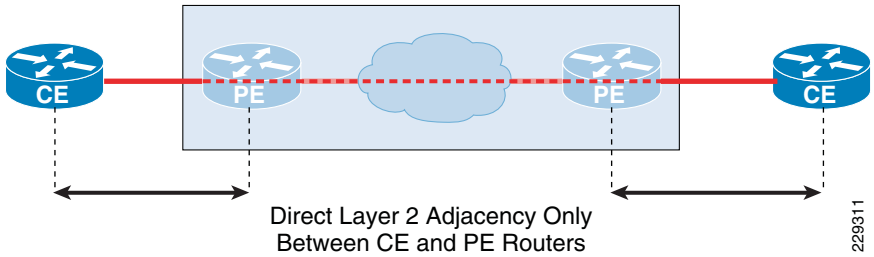

Bandwidth Capacity

This is critical component of the overall WAN design architecture because the application performance largely depends upon guaranteed level of bandwidth at remote sites and the main site. This section primarily discusses the general WAN bandwidth capacity planning and implementation steps.

Planning

To start with, each site must purchase the Metro Ethernet service from local service provider (SP) with bandwidth capacity that can meet the current minimum network load and provides enough flexibility to scale higher in the future, when additional applications need to be added. This bandwidth must be shared based on a 4 class QoS model for optimal delivery., as described in the section Deploying QoS in Network in the document *Small Enterprise Design Profile (SEDP)—Network Foundation Design* (http://www.cisco.com/en/US/docs/solutions/Enterprise/Small_Enterprise_Design_Profile/chap2sba.pdf). Logical bandwidth assignment to each circuit must be symmetric at the remote small office sites and at the main site. A mismatch in bandwidth capacity will force the traffic drop in the SP core due to inconsistent bandwidth service-level agreement.

This design requirement remains consistent for all the remote sites. Depending on large, medium, or small size, the WAN bandwidth capacity may also vary; however, the bandwidth sharing, routing, QoS, multicast etc. principles remains consistent.

Note Main site WAN router is an aggregator that connects multiple small office sites logically over a single media. This media type *must* be GigE, because the platform is 3750ME and the WAN egress port is a non-negotiated fiber port.

Calculating the optimum guaranteed bandwidth required at each remote site and the main site must be done by considering the following factors:

- Number of remote sites
- Bandwidth required at each location
- Platforms scalability limits, mainly, at main office

To explain how much bandwidth is required at each remote site, two terminologies (CIR and interface bandwidth) are used:

- *CIR*—The committed bandwidth that is guaranteed from the service provider, based on the logical connection.
- *Interface speed*—The interface speed is actual Ethernet handoff, which is 1Gbps for both the remote sites and the main site.

For example, consider a scenario where there are three remote sites and one main site. The CIR required at each remote site is 100Mbps and, because there are three remote sites, the main office needs three virtual circuits each having a CIR of 100Mbs. This also means that the aggregate bandwidth at main office is 300Mbps. Figure 5 illustrates this example.

Figure 5 Bandwidth Capacity Planning for Three Remote Sites

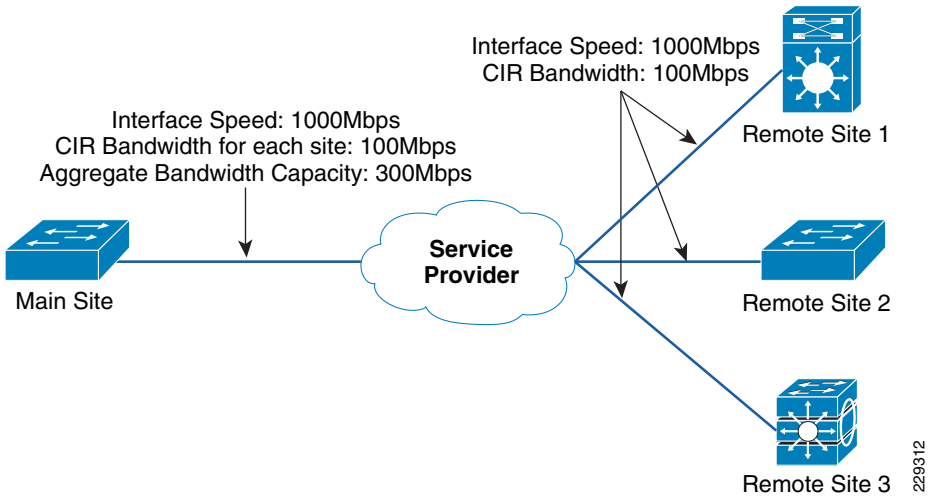

Similarly, if the CIR required at each remote site is 100Mbps and if there are 100 remote sites, then the bandwidth required at main office is 10Gbps. To support an aggregate CIR bandwidth of 10Gbps at main office, we need to think about scalability limits on the WAN aggregation switch. If the aggregated logical connection speed to the remote sites exceed the media capacity on 3750-ME, which is WAN aggregation switch for this design, then there are two approaches to mitigate the problem:

- *Migrate to a Modular switching platform*—Migrating to a modular switching platform for the WAN aggregation tier will enable wan capability to scale higher, reduce operational and management complexities and may offer better performance and resiliency depending on modular platform capabilities. When planning to migrate to a modular switching platform, please make sure that it supports the same QoS features of 3750ME, along with redundancy.

- *Deploy another 3750-ME in same tier*—Alternatively, it is possible to deploy secondary 3750-ME system to connect another set of remote networks. Deploying another set of 3750-ME does not change any WAN principles, In fact, except VLAN and IP address, all the configurations can be replicated to the secondary system.

Implementation

This section discusses how WAN bandwidth was implemented and validated at remote sites and main office of this design. The following are the considerations used to validate this model:

- The main office has dual WAN connections, and each connection is 1Gbps.
- On each WAN connection at the main office, there are 50 virtual circuits, each with CIR of 20Mbps.
- The remote sites have 1Gbps connection, and CIR value of 20Mbs on each circuit.

Figure 6 shows how this design is validated with 100 remote sites.

Figure 6 Bandwidth Capacity Planning for 100 Remote Sites

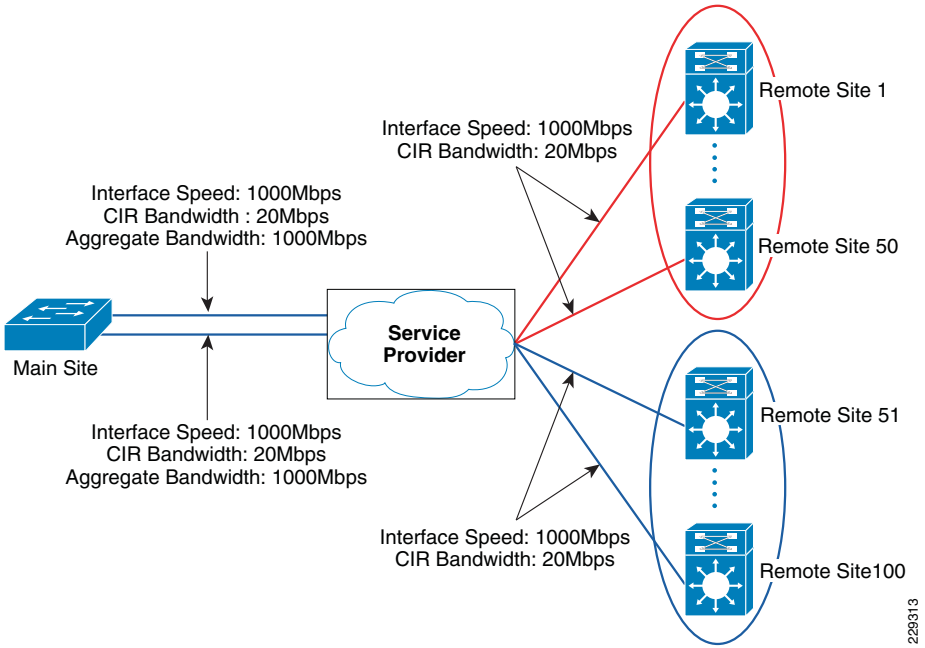

The next section discusses how to implement an IP addressing structure for the WAN interfaces.

Aggregation Structure

This section describes how to aggregate IP address space at remote sites and the main site on the WAN interface.

For designing IP address spaces for the LAN, refer to the section Multicast IP Addressing in the document *Small Enterprise Design Profile (SEDP)—Network Foundation Design* (http://www.cisco.com/en/US/docs/solutions/Enterprise/Small_Enterprise_Design_Profile/chap2sba.pdf)..

As explained in the previous section, all the remote sites are connected to the main site using point-to-point links, which means that every remote site should be on a different subnet. One common method that some customers deploy is using /30 subnets at both ends. However, this would mean that every subnet uses up to four addresses. The recommendation is to use /31 subnets, this way only two addresses are used in every link. The following is a sample configuration for deploying this link at the remote site:

Remote Site	Main Office
interface Vlan501 description Connected to cr24-3750ME-DO dampening **ip address 10.127.0.1 255.255.255.254** no ip redirects no ip unreachables ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp-key ip pim sparse-mode ip summary-address eigrp 100 10.127.0.0 255.255.248.0 5 load-interval 30	interface Vlan501 description Connected to cr35-4507-SS1 dampening **ip address 10.127.0.0 255.255.255.254** no ip redirects no ip unreachables ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp-key ip pim sparse-mode ip summary-address eigrp 100 10.124.0.0 255.252.0.0 5 load-interval 30 hold-queue 2000 in hold-queue 2000 out

After planning an IP addressing scheme, the next part is to complete a routing design for the WAN connections, which is discussed in the following section.

Routing for WAN Connections

This section discusses how to implement routing on WAN interfaces. The key consideration while designing routing scheme is as follows:

- Summarization on network boundaries—This is very important aspect of the design as it prevents unnecessary routing updates to flow across the WAN interface when there is a link state change in the network. For example, let us consider the main site where there are subnets in the following range:

```
10.127.0.0/26
10.127.0.64/26
10.127.0.128/26
        .
        .
        .
10.127.7.64/26
```

Since the above subnets belong to a particular remote site (the main site pair, they could be summarized as 10.127.0.0/21. The following configuration shows how to perform summarization on the WAN interface, which is Vlan501 in this example:

Remote Site	Main Office
interface Vlan501 description Connected to cr24-3750ME-DO dampening ip address 10.1270.1 255.255.255.254 no ip redirects no ip unreachables ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp-key ip pim sparse-mode ip summary-address eigrp 100 10.127.0.0 255.255.248.0 5 load-interval 30	interface Vlan501 description Connected to cr35-4507-SS1 dampening ip address 10.1270.0 255.255.255.254 no ip redirects no ip unreachables ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp-key ip pim sparse-mode ip summary-address eigrp 100 10.124.0.0 255.252.0.0 5 load-interval 30 hold-queue 2000 in hold-queue 2000 out

The next most important component is QoS design, which the following section discusses.

QoS Design

This section discusses how QoS is implemented at the main office and remote sites. With Ethernet defined as the WAN medium, QoS design becomes most important because the router or switch on the WAN edge may believe that it has the complete line rate available for it to transmit. If so, the service provider would start dropping packets due to policers defined at its end. Figure 7 shows what may happen without a proper QoS design.

Figure 7 Policing at Service Provider Due to Lack of Proper QoS at Main Office and Remote Sites

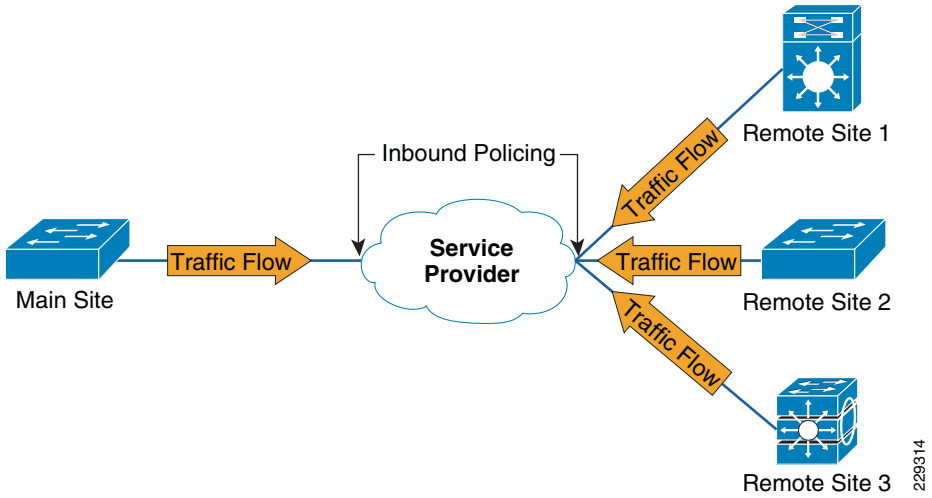

To prevent packets getting dropped at the service provider network, it is very important to have proper and consistent QoS policies at the remote sites and main site. Having a proper QoS policy would ensure that all the traffic is queued and serviced as per the bandwidth defined. This in turn ensures that packets will not be dropped by the service provider. The following subsections discuss how QoS policies are implemented at the remote and main sites.

QoS Policy at Main Site

As mentioned in the When to Consider WAN Technologies, the main office has several point-to-point circuits to the remote sites, and the number (point-to-point circuits) depends upon the number of remote sites. Therefore, the QoS policy at the main site has the following objectives:

- The aggregate traffic going out to the remote site does not exceed 20Mbs.
- All the traffic going out is put in four classes.

To accomplish the above objectives, an hierarchical CBWFQ (HCBWFQ) is implemented. For more information about HCBWFQ, refer to "Chapter 4, Medianet QoS Design Considerations" of the *Medianet Reference Guide* at following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/Medianet_Ref_Gd/medianet_ref_gd.html

To implement HCWFQ, the following two policies are needed:

1. Parent policy that defines the aggregate shape rate.
2. Child policy that enables queuing within the shaped rate.

Figure 8 shows the representation of hierarchical policy.

Figure 8 Hierarchical Policy Implementation at Main Office

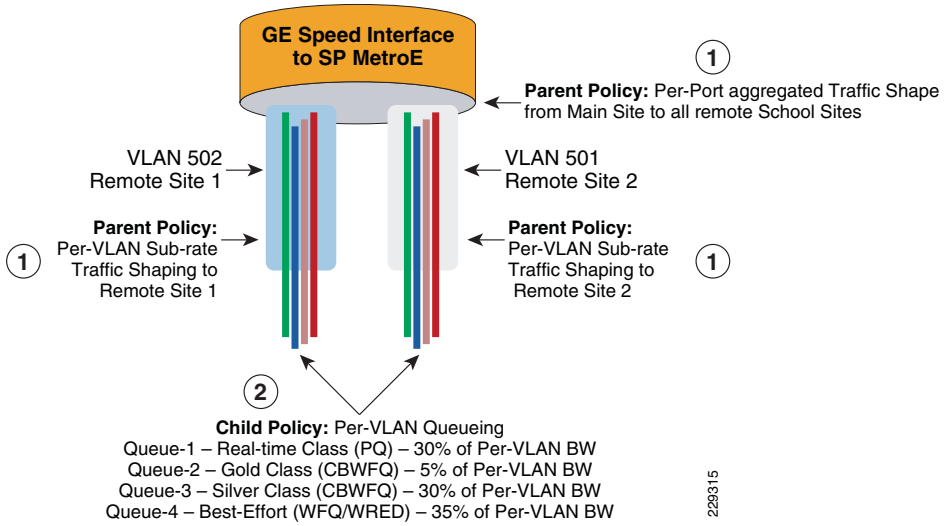

Table 1 shows how these classes are defined.

Table 1 QoS Policies

Class	Queuing type	Bandwidth Allocation
REAL_TIME	LLQ	30%
GOLD	CBWFQ	5%
SILVER	CBWFQ	30%
DEFAULT	CBWFQ	35%

The following is the configuration of the QoS policy at the main office. For brevity, only configuration for one VLAN is shown.

```
class-map match-all Remote_Site1 ← This class map would match the
traffic going to a Remote site
  description cr2-4507-SS1
 match vlan  501

policy-map Remote-Child-Policy-Map  ← This is child policy
 class REAL_TIME
    priority
   police cir percent 30 conform-action set-cos-transmit 5 exceed-action
drop violate-action drop
 class GOLD
    bandwidth percent 5
  set cos 3
```

```
 class SILVER
    bandwidth percent 30
  set cos 2
 class class-default
    bandwidth percent 35
  set cos 0
!
```

The following configuration is of the parent policy, which in this design shapes to 20Mbps for each remote site:

```
Policy Map Remote-1to50-Parent-Policy-Map ← This is parent policy map
    Class Remote_Site1
      shape average 20000000 (bits/sec)
      service-policy Remote-Child-Policy-Map ← This is child policy map
```

After defining the policies, they are applied to WAN interfaces. The following example shows the configuration of the Metro switch on its WAN interface:

```
interface GigabitEthernet1/1/1
 description Connected to SP-MPLS-Core-cr24-6500-1
 switchport trunk native vlan 801
 switchport trunk allowed vlan 501-550
 switchport mode trunk
 logging event trunk-status
 load-interval 30
 carrier-delay msec 0
 priority-queue out
 mls qos trust dscp
 spanning-tree portfast trunk
 spanning-tree bpdufilter enable
 spanning-tree guard root
 max-reserved-bandwidth 100
 service-policy output Remote-1to50-Parent-Policy-Map ← The policy-map
 hold-queue 2000 in
 hold-queue 2000 out
```

After completing the QoS policy at main office, we need to define the QoS policy at remote sites. The following subsection describes how to define QoS policy at the remote site.

QoS Policy at Remote Site

At the remote sites, the objectives are similar to the one at main site, which is to ensure that 20Mbps is the aggregate traffic leaving out the switch and the ingress traffic is queued in 4 classes. However, the implementation is different from the main office due to lack of hierarchical QoS support on the particular switch deployed at the remote site. It is however possible to have a line card with the capability to do HCBWFQ, if it is needed. The following configuration shows how to do QoS policy when the device does not support HCBWFQ.

The first step is to queue the ingress traffic in the four queues. Table 2 shows the queues and the bandwidth allocated.

Table 2 QoS Policies for Remote Site

Class	Queuing type	Bandwidth Allocation
REAL_TIME	Per-class	6mbps
GOLD	Per-class	1mbps
SILVER	Per-class	7mpbs
DEFAULT	Per-class	6mbps

The following is configuration of egress interface on the remote site:

```
interface GigabitEthernet1/1
 description Connected to MetroE-Core-cr25-6500-1
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 801
 switchport trunk allowed vlan 501
 switchport mode trunk
 logging event link-status
 load-interval 30
 carrier-delay msec 0
 qos trust dscp
 udld port disable
 tx-queue 1
   bandwidth 1 mbps
 tx-queue 2
   bandwidth 7 mbps
 tx-queue 3
   bandwidth 6 mbps
   priority high
 tx-queue 4
   bandwidth 6 mbps
 no cdp enable
 spanning-tree portfast trunk
 spanning-tree bpdufilter enable
 spanning-tree guard root
 service-policy output WAN-EGRESS-PARENT
```

The above configuration ensures that each class of the traffic is queued as shown in Table 2. However, the bandwidth only ensures the minimum amount of bandwidth available. It does not control the upper threshold, which needs to be 20Mbs in our solution. Therefore, to ensure that the traffic on the egress interface does not exceed 20Mps, WAN-EGRESS-PARENT policy is used to police the traffic to 20Mbps. The following is configuration of the WAN egress policy:

```
cr35-4507-SS1#show policy-map WAN-EGRESS-PARENT
  Policy Map WAN-EGRESS-PARENT
    Class class-default
      police 20 mbps 1000 byte conform-action transmit exceed-action drop

      service-policy WAN-EGRESS-CHILD
cr35-4507-SS1#
```