Schools Network Foundation Deployment Guide

Last Updated: October 18, 2009

Schools Design Principles and Foundational Technologies

This document describes the Schools Service Ready Architecture network design, which is a well designed and tested network architecture that is flexible, and cost effective to support a wide range of educational services. Key features of the Schools SRA include:

- High Availability
- Single Fabric—Multi Services
- Differentiated Services
- Layer 2 and Layer 3 Access

This document provides design guidance to build a highly resilient, manageable and cost-effective school network which provides a solid in foundation for seamless integration and operation of applications and network services. The network has been specifically designed to meet the challenges of the education environment.

This document consist of three major section:

- Schools Network Infrastructure
- Core Distribution Integration
- District Office Integration

The first section focuses upon the networking technology that is the foundation of the SRA, such as Layer 2 Networking, IP Routing, and QoS. The following two section specifically address the integration of core components into the the foundation, for example the connection of components such as ISR Routers, Wireless LAN Controllers, and Services Blocks to the network.

Building Unified Schools Network Infrastructure

Cisco has many years of experience developing high performance, highly available, multi service networks. The key to developing a robust design is applying a proven methodology. The following design principles were applied to develop the School SRA network architecture:

- Hierarchy
- Modularity
- Resiliency
- Flexibility

The Unified Schools Network is designed to be highly available, and cost effective, while delivering capabilities necessary to enable advanced services, such as IP telephony, video, security, wireless LANs. The network design includes the following key features;

- Heirarchichal design with collapsed Core
- Quality of Service to enure real time data (telephony, video) are given higher priority
- Application of Resilient design principles
- Multicast
- Routed Access
- Redundancy

Hierarchical Network Design with Collapsed Core

A three-tier heirarchical design maximizes performance, network availability, and the ability to scale the network design. Most school campus' do not grow significantly larger over time, and most school campus' are small enough to be well served by a two-tier hierarchical design, where the Core and Distribution layers are collapsed into one layer. The primary motivation for the collapsed core design is reducing network cost, while maintaining most of the benefits of the three-tier hierarchical model.

Deploying a collapsed core network results in the Distribution layer and Core layer functions being implemented in a single device. The collapsed Core/Distribution device must provide:

- High speed physical and logical paths connecting to the network
- Layer-2 aggregation and demarcation point
- Define routing and network access policies
- Intelligent network services—QoS, Network virtualization, etc.

Note If the District Office or a School Campus has multiple buildings, and is expected to grow over time, then implementing the three-tier hierarchical model is a better choice.

Figure 1 illustrates a sample network diagram for a single building school district office.

Figure 1 Collapsed Core School District Office Network Design

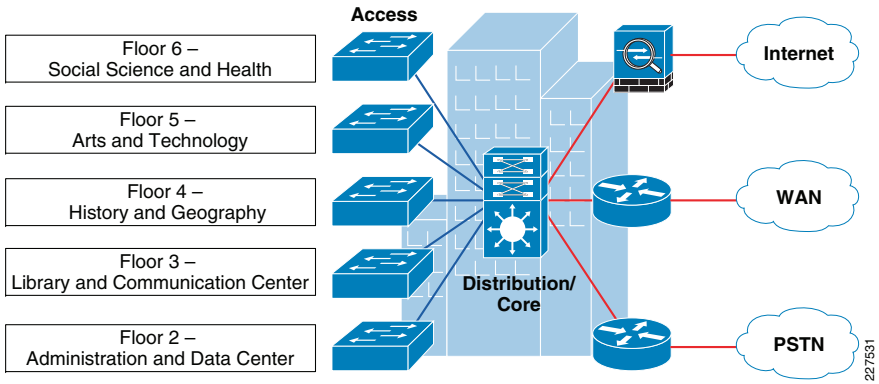

District Office Network Design

If the District Office has multiple buildings, or the District Office site is expected to grow significantly over time; then implementing the 3 tier hierarchical model is a good choice. For smaller District Office sites which are unlikely to grow significantly, the collapsed core model is more cost effective. The School Service Ready Architecture utilizes the collapsed core network design in the District Office.

The Collapsed Core network may be deployed with redundant core/distibution router, or consolidated core/distribution router. See Figure 2.

Figure 2 Collapsed core district office school network models

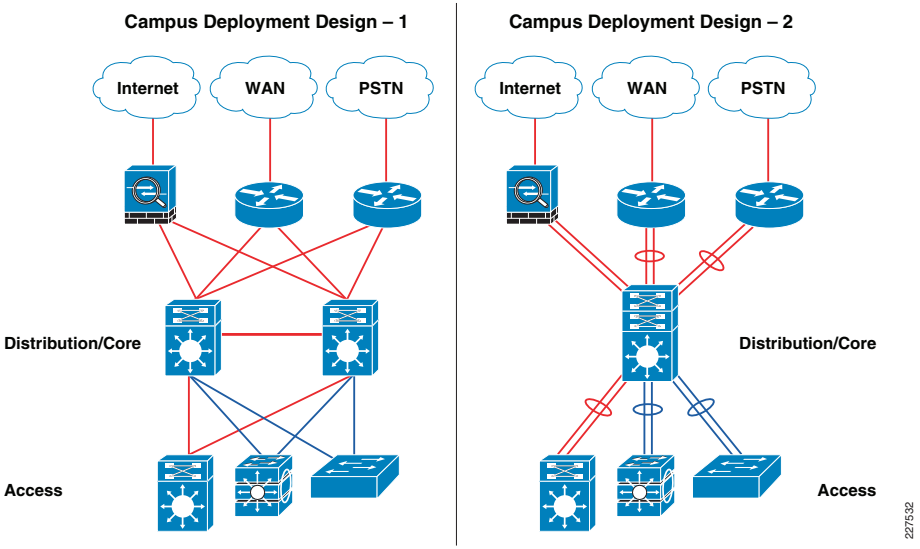

The redundant design is more complex, since all of the core/distribution functions must be implemented on two routers in a complimentary fashion. To learn more about the redundant designs, refer to *High Availability School Recovery Analysis Design Guide*

http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/HA_recovery_DG/campusRecovery.html

The School SRA District Office is designed with a consolidated core/distribution router to maximize performance, while keeping costs affordable. With this design, the default behavior of layer-2 and layer-3 nework control protocols is to create a redundant view between two systems. The core router builds a ECMP routing topology which results in symetric forwarding paths beyond the District Office.

Default layer-2 configuration eliminates the need for FHRP,. This simplifies the network operation, since there is no need to configure or tune FHRP protocols. Since this design employs point-to-point links between the collapsed core and peer devices, the solution is to tune the network to enable a single control plane, to improve forwarding efficienty and resource utilization. The recommendation is to aggregate all physical ports into a single logical channel-group. This logical aggregated Ethernet bundle interface is known as *EtherChannel*.

EtherChannel

EtherChannel provides inverse-multiplexing of multiple ports into a single logical port to a single neighbor. This technique increases bandwidth, link efficiency, and resiliency. EtherChannel technology operates on the MAC layer. Upper layer protocols require a single instance to operate over the logical interface. EtherChannel provides efficient network operation and graceful recovery to higher layer protocols during bundle port failure and restoration.

EtherChannel helps improve the overall network stability and availability. Failure of individual physical link will cause network topology recomputation, restoration, and may be rerouted. Such process requires CPU interruption that could impact the overall application performance. EtherChannel significantly simplifies the network response to a individual link failure. If an individual link in EtherChannel fails, the interface will not trigger any network topology changes. All underlying hardware changes remain transparent to higher-layer protocols, thus minimizing impact to network and application performance, and improving network convergence. Figure 3 illustrates how enabling EtherChannel in Layer-2 and Layer-3 network simplifies control-plane and forwarding-plane.

Figure 3 Optimized control and forwarding paths with EtherChannel

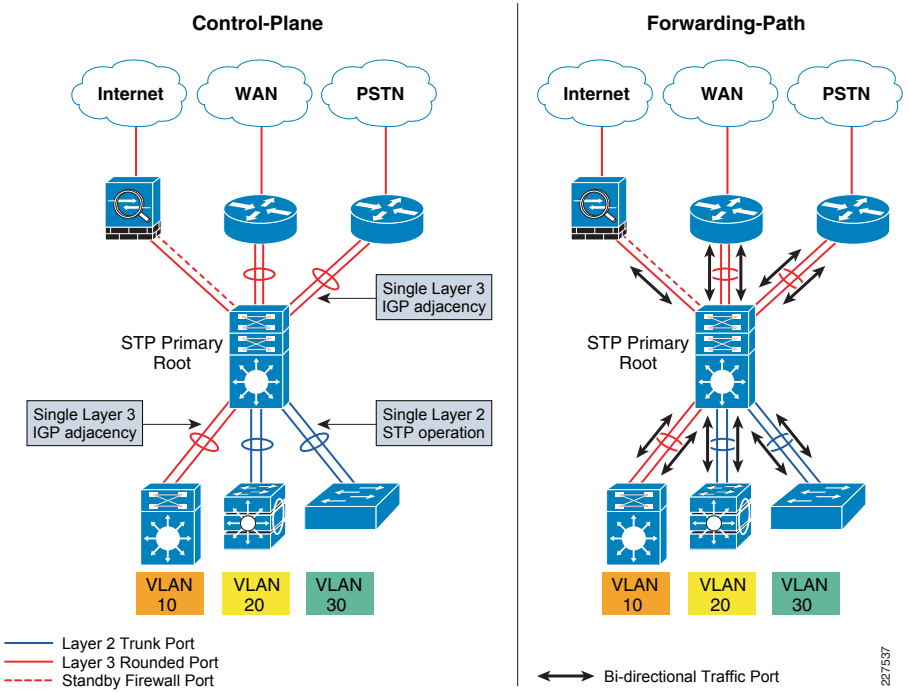

Resilient Distributed System

The consolidated core/distribution layer may become a single-point-of-failure (SPOF) in the network. A software upgrade or a route-processor failure may cause network outage for minutes.

The school SRA uses the Cisco Catalyst 4500 with next-generation Supervisor-6E in the consolidated core/distribution layer. It is choosen for it's price performance, and the high availability features within the device. The Cisco Catalyst 4500 switch supports redundant supervisor engines and provides Stateful Switchover (SSO) and Non-Stop Forwarding (NSF) capabilities. SSO ensures the L2 and L3 protocol state-machines and network forwarding entries on the standby supervisor engine are maintained, and can quickly assume control-plane responsibilities and gracefully restore the control-plane in the event of a primary supervisor failure. While the control-plane is gracefully recovering, the NSF function continues to switch traffic in hardware.

Access-Layer Edge Services

The access layer is the first tier or edge of the network. It is the layer where end devices (PCs, printers, cameras, etc.) attach to the school network. It is also the layer where devices that extend the network out one more level are attached; IP phones and wireless access points (APs) are examples of devices that extend the connectivity out from the access switch. The wide variety of devices that can connect and the various services and dynamic configuration mechanisms required, make the access layer the most feature-rich layer of the school network. Figure 4 illustrates a district office network deployment with various types of trusted and untrusted endpoints. .

Figure 4 Access-Layer Trust Boundary and Network Control Services

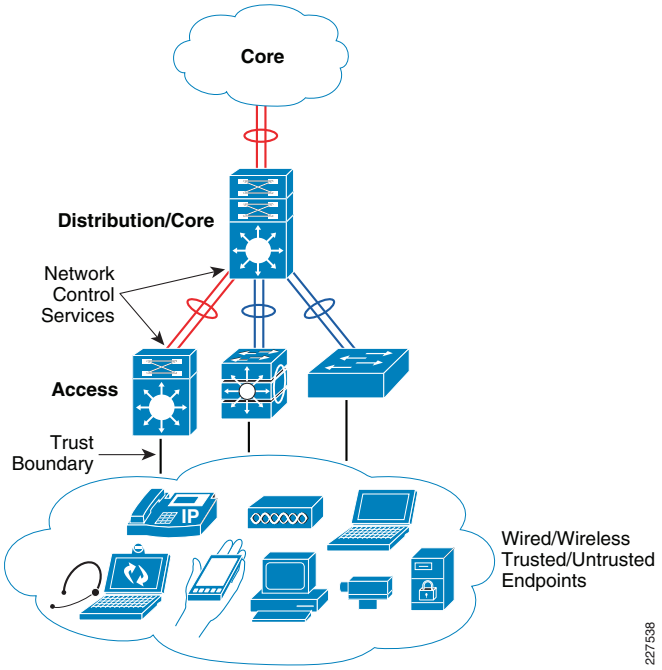

Table 3 examples of the types of services and capabilities that need to be defined and supported in the access layer of the network.

Table 1 Access-layer Services and Capabilities

Service Requirements	Service Features
Discovery and Configuration Services	802.1AF, CDP, LLDP, LLDP-MED
Integrated Security Services	IBNS (802.1X), CISF – Port-Security, DHCP Snooping, DAI and IPSG
Network Identity and Access	802.1X, MAB, Web-Auth
Application Recognition Services	QoS marking, policing, queueing, deep packet inspection NBAR
Intelligent Network Control Services	PVST+, Rapid PVST+, EIGRP, OSPF, DTP, PAgP/LACP, UDLD, FlexLink, Portfast, UplinkFast, BackboneFast, LoopGuard, BPDUGuard, Port Security, RootGuard

Table 1 Access-layer Services and Capabilities

Service Requirements	Service Features
Energy Efficient Services	Power over Ethernet, EnergyWise, Energy efficient systems
Management Services	Auto-SmartPort Macro, Cisco Network Assistant

The access layer provides the intelligent demarcation between the network infrastructure and the computing devices that use the infrastructure. It provides network edge security, QoS, and policy trust boundary. It is the first point of negotiation between the network infrastructure and the end devices seeking access to the network.

Design details, explaining how to select the features needed for a given deployment, and how to implement the features is provided in *Schools Access Layer Features* chapter,

Access-Layer Design

School SRA is designed with 2 to 4 uplink ports for each access switch, providing link-failure protection, the Schools SRA Acces Network supports a Hybrid Access Layer of either a flat, segmented and/or Layer 3 as required by the school. (these Access layer options are discussed in more detail later in this document) .

For mission critical endpoints, School SRA employs the Cisco StackWise or StackWise Plus solution in the access. It is designed to physically stack and interconnect multiple Layer-2 or Layer-3 switches using special cables. Stacking multiple switches into a logical ring creates a single unified and resilient access-layer network topology. The Cisco Catalyst 2975 StackWise can be deployed in Layer-2 network domain and the Cisco Catalyst 3750-E StackWise or StackWise Plus is deployed for routed access implementations. See Figure 5.

Figure 5 Resilient, Scalable and Efficient Access-Layer Network Design

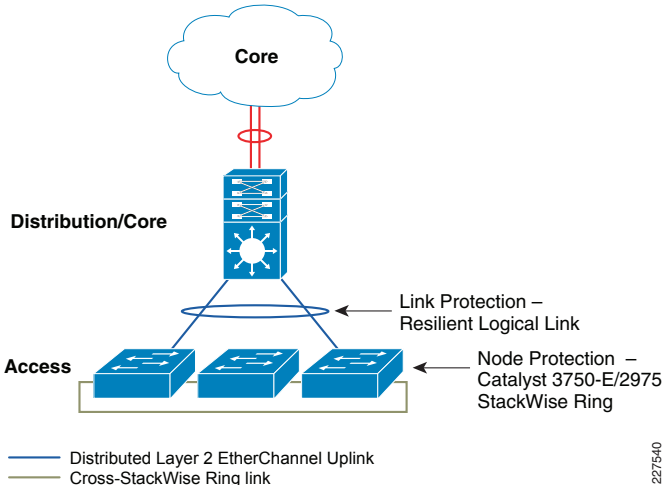

School Site Design

The School Service Ready Architecture includes two school site designs. One for larger schools, and another for medium to smaller schools. The typical school site is a single building with a limited popluation which makes the collapsed core network design a suitable choice.

School Collapsed Core Design

The key criteria to consider when designing a school network are the network size, bandwidth capacity and high-availability requirements. The School Service Ready Architecture includes two models; one for smaller schools, and another for larger schools. Both designs offer high capacity, performance and availability. illustrates the two school network design models. See Figure 6.

Figure 6 Collapsed core school network models

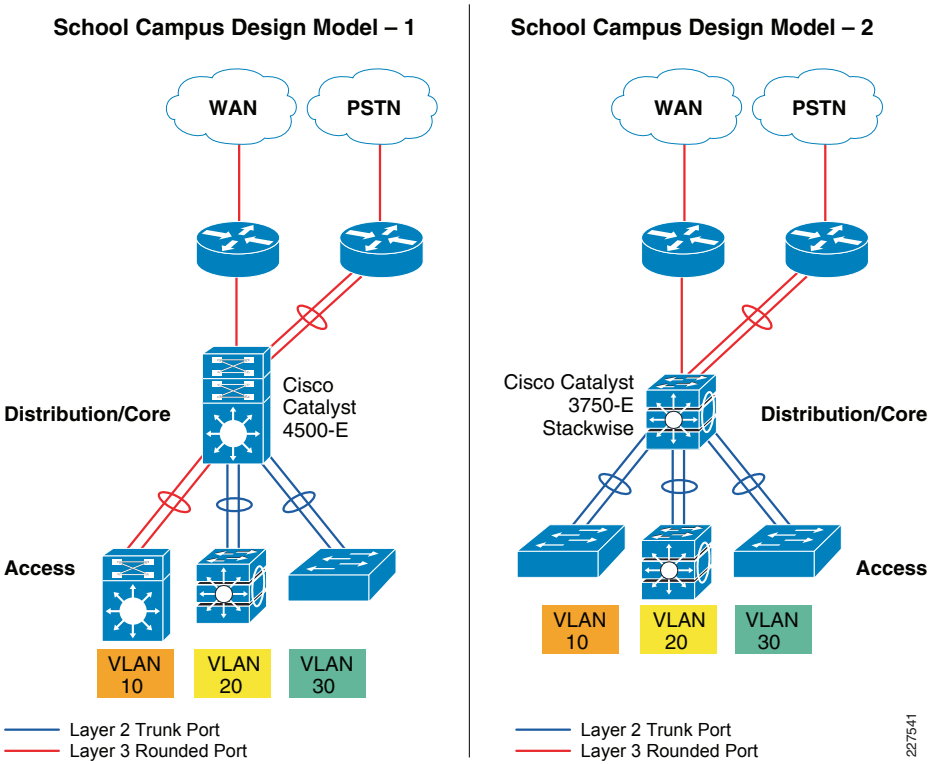

Design Model-1 is for a larger school site. The network design is the same as the District Office network design, with the same performance capabilities, scalability options, and high availability features.

Design Model-2 is for a medium to small school site. The primary difference is the use of the Cisco Catalyst 3750-E Stack Wise Plus switch in the collapsed core/distribution layer. The 3750-E Stack Wise Plus deploys up to nine 3750-E switches in a ring topology as a single virtual switch. Each chassis replicates the control functions, and provides packet forwarding. If the master switch fails, another switch will quickly assume the control plane 'master' function. This results in a cost effective, high performance, scalable solution, with built in resliiency.

- *Performance* —Provides wire-rate network connection to access switches
- *Scalable*—May deploy up to 9 switches in a stack to aggregate a reasonble number of access switches
- *High Availability*—Stack provides a virtual switch, with distributed control plane, delivering subsecond system failure recovery times

The Cisco 3750-E StackWise Plus delivers high performance routing and switching capability and robust IOS feature support. The control-plane and forwarding paths functions for the Cisco 3750-E StackWise Plus in the collapsed core network design remain the same. For more information about the Cisco 3750-E StackWise architecture, refer to the following URL:

http://www.cisco.com/en/US/partner/prod/collateral/switches/ps5718/ps5023/prod_white_paper09186a00801b096a_ps7077_Products_White_Paper.html

School Access-Layer Design

The Access-layer network design at the school site is the same as at the district office. The same devices are available, and the same design choices may be deployed to achieve a high performance, secure and resilient access layer. To simplify the overall system design, and network operations, it is recommended to use consistent design and platform selections in the access-layer role, at the district office and school sites. This will allow a common configuration template and simplify operations and troubleshooting procedures.

Deploying Schools Foundation Services

The two-tier hierarchical design delivers a *reliable* and *resilient*, *scalable*, and *manageable* foundation network design. This subsection provides design and deployment guidelines for the school core layer, and access-distribution block.

The Access - Distribution block, as described in the "School Network Design" section on page 20, uses a combination of Layer-2 and Layer-3 switching to provide a balance of policy and access controls, availability, and flexibility in subnet allocation and VLAN usage. Deployment guidelines are provided to implement multi-layer, and routed access designs in the access-distribution block.

Implementing EtherChannel in School Network

Etherchannel is used throughout the network design, and the implementation guidelines are the same for multi-layer, and routed-access models, and in the WAN edge design. As recommended in the "EtherChannel Fundamentals" section on page 14, there should be single logical point-to-point EtherChannel deployed between collapsed core and access-layer. The EtherChannel configuration on each end of the link in the access-distribution block must be consistent to prevent a link bundling problem. EtherChannels use link bundling protocols to dynamically bundle physical interfaces into a logical interface.

The following are the benefits of building EtherChannel in dynamic mode:

- Ensure link aggregation parameters consistency and compatibility between switches.

- Ensure compliance with aggregation requirements.
- Dynamically react to runtime changes and failures on local and remote Etherchannel systems
- Detect and remove unidirectional links and multi-drop connections from the Etherchannel bundle.

EtherChannels can be deployed in dynamic or static modes. Both EtherChannel modes can coexist in a single system; however, the protocols (PagP, LACP) can not interoperate with each other.

- Cisco proprietary link aggregation-Cisco's implementation of Port Aggregation group Protocol (PAgP) is supported on all the Cisco Catalyst platforms. The PAgP protocol is not supported when the Cisco Catalyst 2975 or 3750 Series switches are deployed in StackWise mode.
- EEE 802.3ad link aggregation-Link Aggregation Control Protocol (LACP) is based on IEEE 802.3ad specification to operate in vendor-independent network environment. LACP link bundling protocol is developed with same goal as Cisco's PAgP. Cisco Catalyst switches in StackWise mode must use LACP to dynamically bundle.

See Figure 7.

Figure 7 Implementing EtherChannel in District Office School Network

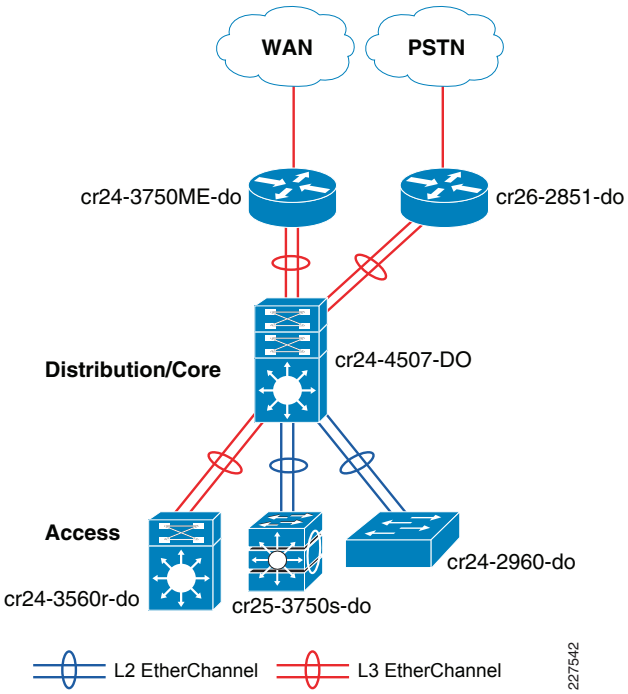

The following sample configuration shows how to build Layer-2 and Layer-3 EtherChannel configuration and bundling physical ports into appropriate logical EtherChannel-group:

```
cr24-4507-DO#config t
Enter configuration commands, one per line.  End with CNTL/Z.
cr24-4507-DO(config)#interface Port-channel1
cr24-4507-DO(config-if)# description Connected to cr24-3750ME-DO
cr24-4507-DO(config-if)#
cr24-4507-DO(config-if)#interface Port-channel11
cr24-4507-DO(config-if)# description Connected to cr24-2960-DO
cr24-4507-DO(config-if)# switchport
cr24-4507-DO(config-if)#
cr24-4507-DO(config-if)#interface Port-channel16
cr24-4507-DO(config-if)# description Connected to cr25-3750s-DO
cr24-4507-DO(config-if)# switchport
cr24-4507-DO(config-if)#
cr24-4507-DO(config-if)#interface range Gig 3/3 , Gig 4/3
cr24-4507-DO(config-if-range)# description Connected to cr24-3750ME-DO
cr24-4507-DO(config-if-range)# channel-protocol pagp
cr24-4507-DO(config-if-range)# channel-group 1 mode desirable
cr24-4507-DO(config-if-range)#
cr24-4507-DO(config-if-range)#interface range Gig 1/1 , Gig 2/1
cr24-4507-DO(config-if-range)# description Connected to cr24-2960-DO
cr24-4507-DO(config-if-range)# channel-protocol pagp
cr24-4507-DO(config-if-range)# channel-group 11 mode desirable
cr24-4507-DO(config-if-range)#
cr24-4507-DO(config-if-range)#interface range Gig 1/6 , Gig 2/6
cr24-4507-DO(config-if-range)#description Connected to cr26-3750s-DO
cr24-4507-DO(config-if-range)# channel-protocol lacp
cr24-4507-DO(config-if-range)# channel-group 16 mode active
```

Enabling EtherChannel on each switch endpoint will automatically form a logical connection and can be verified using following CLI command:

```
cr24-4507-DO#show etherchannel summary | inc Po
Group  Port-channel  Protocol    Ports
1      Po1(RU)         PAgP      Gi3/3(P)    Gi4/3(P)
11     Po11(SU)        PAgP      Gi1/1(P)    Gi2/1(P)
16     Po16(SU)        LACP      Gi1/6(P)    Gi2/6(P)
```

EtherChannel Load Balancing

EtherChannel load-sharing is based on a polymorphic algorithm. On per protocol basis, load sharing is done based on source XOR destination address or port from Layer 2 to 4 header and ports. For higher granularity and optimal utilization of each member-link port, an EtherChannel can intelligently load-share egress traffic using different algorithms.

EtherChannel load-balancing mechanisms function on a per-system basis. By default, EtherChannel will use the hash computation algorithm. The network administrator can globally configure the load balancing mechanism. In Cisco Catalyst platforms, EtherChannel load balancing is performed in hardware and it cannot perform per-packet-based load balancing among different member links within EtherChannel. Bandwidth utilization of each member-link may not be equal in default load balancing mode. The Ether Channel load balancing method should be changed to source and destination IP address based throughout the district office and school network Tuning the load-balancing to source-and-destination IP address allows for statistically-superior load-distribution. When loads are balanced in this manner, packets belonging to a single flow will retain their packet order. See Figure 8.

Figure 8 EtherChannel Load-Balance Method

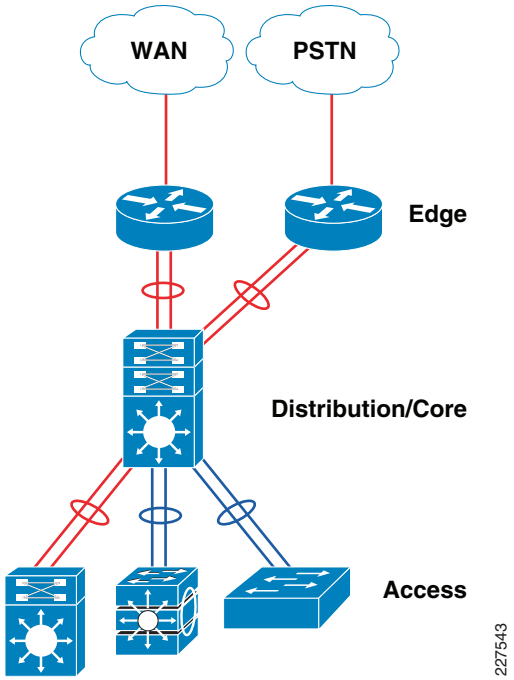

The following output provides sample configuration guideline for changing the default port-channel load-balance setting to source-destination-ip based. Aside from Layer-2 or Layer-3 EtherChannel mode, similar configuration must be applied on each system in the access-distribution block and WAN edge.

```
cr24-4507-DO(config)#port-channel load-balance src-dst-ip
cr24-4507-DO#show etherchannel load-balance
EtherChannel Load-Balancing Configuration:
        src-dst-ip
EtherChannel Load-Balancing Addresses Used Per-Protocol:
Non-IP: Source XOR Destination MAC address
  IPv4: Source XOR Destination IP address
  IPv6: Source XOR Destination IP address
```

Implementing Dynamic Routing

This section provides implementation and best practice guidelines for deploying the core-layer in both the district office and school site. Proper design of the core network layer ensures reachability, transparency and availability. This section focuses on building a unicast routing topology.

Enabling routing in the school network is a simple task. However, the network physical layout must be carefully planned and designed to ensure flexible, stable and efficient routing. Developing a hierarchical network addressing scheme enables a stable, efficient and scalable design.

Efficient address allocation - hierarchical addressing enables efficient use of address space, since groups are contiguous.

- Improves routing efficienty—Using contiguous ip addresses enables efficient route summarization. Route summarization simplifies the routing database, and computations during topology changes. This reduces the nework bandwidth used by the routing protocol, and improves routing protocol performance by reducing network convergence time.
- *Improves system performance*—Hierarchical, contiguous ip addressing reduces router memory usage by eliminating dis-contiguous and non-summarized route entries. It saves on CPU cycles needed to compute the routing database during topology changes. This contributes to a more stable routing network, and simplifies the task of network operations and management.

Cisco IOS supports many Interior Gateway Protocols (IGP), including EIGRP and OSPF, either of which are suitable for large network deployments. While OSPF is capable of greater scale, it is also more complex, and hence more difficult to configure, operate and manage. The Schools Service Ready Architecture is designed and validated using EIGRP, since it is a stable, high performance, efficient protocol, which is simple to implement and manage. The same design principles apply whether using EIGRP or OSPF.

Designing End-to-End EIGRP Routing Domain

EIGRP is a balanced hybrid routing protocol that builds neighbor adjacency and a flat routing topology on a per-autonomous-system (AS) basis. The LAN/WAN infrastructure of School Service Ready Architecture should be deployed in a single EIGRP AS to prevent route redistribution, loops, and other problems that may occur due to misconfiguration. See Figure 9.

Figure 9 End-to-End EIGRP Routing Design in School Architecture

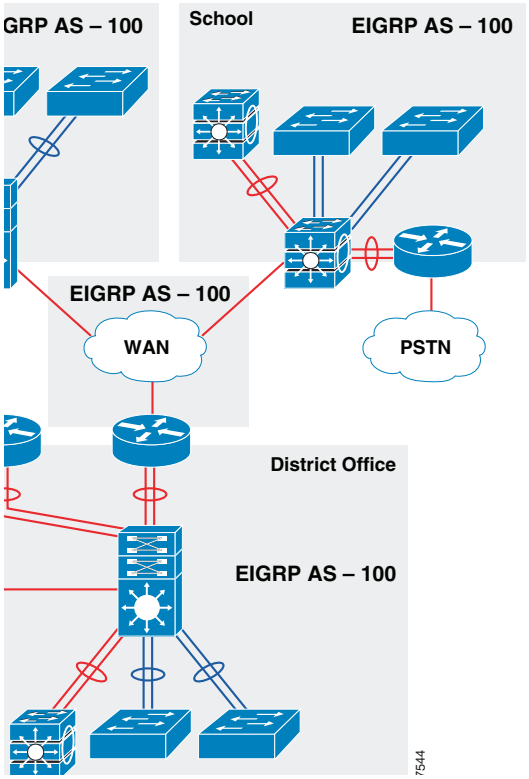

Implementing EIGRP Routing

The district office is the central hub in the network. Each school site is connected to the district office over the WAN infrastructure. The district office network includes the Internet gateway and provides access to the central data-center. Since both the school sites, and district office networks use the collapsed core design, the routing configuration of the core routers is the same..

The following is a sample configuration to enable EIGRP routing process at the edge of the district office collapsed core network. EIGRP is enabled in the school site network with the same configuration:

```
cr24-4507-DO(config)#interface Loopback0
cr24-4507-DO(config-if)# ip address 10.125.100.1 255.255.255.255

cr24-4507-DO(config-if)#interface Port-channel1
cr24-4507-DO(config-if)# description Connected to cr24-3750ME-DO
cr24-4507-DO(config-if)#no switchport
cr24-4507-DO(config-if)# ip address 10.125.32.4 255.255.255.254

cr24-4507-DO(config-if)#interface Port-channel2
cr24-4507-DO(config-if)# description Connected to cr24-2851-DO
cr24-4507-DO(config-if)#no switchport
cr24-4507-DO(config-if)# ip address 10.125.32.6 255.255.255.254

cr24-4507-DO(config)#interface Vlan200
cr24-4507-DO(config-if)# description Connected to cr24_ASA_Inside_Port
cr24-4507-DO(config-if)# ip address 10.125.33.9 255.255.255.0


cr24-4507-DO(config)#router eigrp 100
cr24-4507-DO(config-router)# no auto-summary
cr24-4507-DO(config-router)# eigrp router-id 10.125.100.1
cr24-4507-DO(config-router)# network 10.125.0.0 0.0.255.255

cr24-4507-DO#show ip eigrp neighbor port-channel 13
EIGRP-IPv4:(100) neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.125.33.10Vl200111d00h       1   200  0  171
0   10.125.32.7Po2161d02h       1   200  0  304
2   10.125.32.5Po1141d02h       2   200  0  25038
```

EIGRP Adjacency Protection

Implementing summarization in the EIGRP routing process automatically enables EIGRP routing process on each interface that is summarized. By default, the router transmits and accept EIGRP hello messages from remote device to form an adjacency on all EIGRP enabled interfaces. This behavior needs to be modified to ensure a secure, efficient and stable routing design.

- *System efficiency*—There is no need to send EIGRP hellos on an interface where there is no trusted EIGRP neighbor. In a large network, sending EIGRP hello messages periodically to such interfaces consumes unnecessary CPU resource. EIGRP route processing should only be enabled on interfaces where trusted network devices are connected. All other interfaces can be suppressed in passive mode. The following configuration shows how to automatically disable EIGRP processing on all the Layer 3 interfaces and only enable on the trusted interface. This design principle must be applied on each EIGRP router, including distribution and core routers:

```
cr24-4507-DO(config)#router eigrp 100
cr24-4507-DO(config-router)# network 10.125.0.0 0.0.255.255
cr24-4507-DO(config-router)# passive-interface default
cr24-4507-DO(config-router)# no passive-interface Port-channel1
cr24-4507-DO(config-router)# no passive-interface Port-channel2
cr24-4507-DO(config-router)# no passive-interface Vlan200

cr24-3560r-DO#show ip eigrp interface
EIGRP-IPv4:(100) interfaces for process 100

                    Xmit           Queue   Mean   Pacing Time   Multicast Pending
Interface    Peers      Un/Reliable  SRTT   Un/Reliable   Flow Timer  Routes
Vl2001        0/0 1    0/1      50           0
Po1  1        0/0 2    0/1      50           0
Po2  1        0/0 4    0/1      50           0


cr24-4507-DO#show ip protocols | inc Passive|Vlan
  Passive Interface(s):
    Vlan1
    Vlan101
    Vlan102
    Vlan103
    Vlan104
```

- *Network Security*—Sending unnecessary EIGRP Hello messages opens a security vulnerability in two ways. An attacker can detect EIGRP operation and send flood of EIGRP hello messages to destabilize the network. Or an attacker could establish a "fake" EIGRP adjacency and advertise a best metric default-route into the network to black hole and compromise all critical traffic. Each EIGRP system should implement MD5 authentication, and each EIGRP neighbor should validate MD5 authentication is enabled on adjacent systems . This provides a secure method of transmitting and receiving routing information between devices in the network. Following is a sample configuration to enable EIGRP neighbor authentication using MD5:
  - -Distribution

```
cr24-4507-DO(config-keychain)# key 1
cr24-4507-DO(config-keychain-key)#   key-string <password>

cr24-4507-DO(config)#interface Port-channel1
cr24-4507-DO(config-if)# description Connected to cr24-3750ME-DO
cr24-4507-DO(config-if)# ip authentication mode eigrp 100 md5
cr24-4507-DO(config-if)# ip authentication key-chain eigrp 100
eigrp-key
```

  - WAN Aggregation

```
cr24-3750ME-DO(config)#key chain eigrp-key
cr24-3750ME -DO(config-keychain)# key 1
cr24-3750ME -DO(config-keychain-key)#   key-string <password>

cr24-3750ME -DO(config)#interface Port-channel1
cr24-3750ME -DO(config-if)# description Connected to cr24-4507-DO
cr24-3750ME -DO(config-if)# ip authentication mode eigrp 100 md5
cr24-3750ME -DO(config-if)# ip authentication key-chain eigrp 100
eigrp-key
```

- *System Stability*—As mentioned in Table 8, EIGRP allows network adminstrator to summarize multiple individual and contiguous networks into a single summarized network before advertising to neighbors. Route summarization improves performance, stability, and convergence times, and it makes the network easier to manage operate and troubleshoot. EIGRP provides the flexibility to summarize at any point in the network. Proper design requires determining which routers will serve as Aggregators,and advertise summarized network information to peers. Routers which connect multiple access devices, or connect to the WAN edge should be made Aggregators. Figure 17 provides an example Schools SRA network with route aggregator devices identified with the direction of route summarization illustrated.

Figure 10 Route Aggregator and Summary Route Advertisement Direction

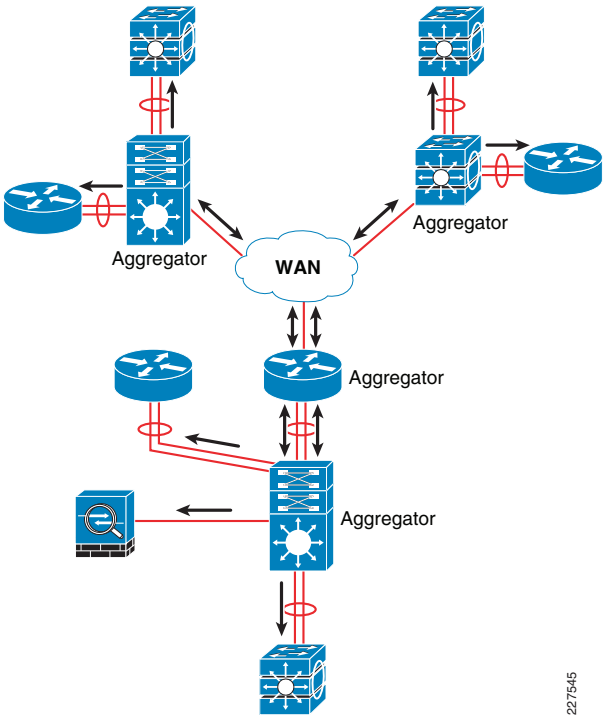

The following sample configuration shows EIGRP route summarization. In this example, the entire access-layer network is summarized into a single classless network and advertised to the WAN edge, the ASA firewall and the PSTN gateway.

- *Distribution*

```
cr24-4507-DO(config)#interface Port-channel1
cr24-4507-DO(config-if)# description Connected to cr24-3750ME-DO
cr24-4507-DO(config-if)# ip summary-address eigrp 100 10.125.0.0 255.255.0.0

cr24-4507-DO(config-if)#interface Port-channel2
cr24-4507-DO(config-if)# description Connected to cr24-2851-DO
cr24-4507-DO(config-if)# ip summary-address eigrp 100 10.125.0.0 255.255.0.0

cr24-4507-DO(config-if)#interface Vlan200
cr24-4507-DO(config-if)# description Connected to cr24_ASA_Inside_Port
cr24-4507-DO(config-if)# ip summary-address eigrp 100 10.125.0.0 255.255.0.0

cr24-4507-DO#show ip protocols | inc Address|10.125.0.0
Address Family Protocol EIGRP-IPv4:(100)
  Address Summarization:
    10.125.0.0/16 for Port-channel1, Vlan200, Port-channel2
```

- WAN Aggregation

  Verifying district office EIGRP summarized route status at WAN aggregation layer as follows:

```
cr24-3750ME-DO#show ip route 10.125.0.0 255.255.0.0
Routing entry for 10.125.0.0/16
```

```
Known via "eigrp 100", distance 90, metric 1792, type internal
Redistributing via eigrp 100
Last update from 10.125.32.4 on Port-channel1, 1d04h ago
Routing Descriptor Blocks:
* 10.125.32.4, from 10.125.32.4, 1d04h ago, via Port-channel1
    Route metric is 1792, traffic share count is 1
    Total delay is 20 microseconds, minimum bandwidth is 2000000 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 1
```

Tuning EIGRP Protocol Timers

EIGRP uses Hello messages to form adjacencies and determine if neighbors are alive. EIGRP adjacency is declared down if it fails to receive Hello messages within the Hold down timer interval. All the prefixes discovered from a dead neighbor are removed from the routing table. By default, EIGRP transmits a Hello message every 5 seconds to notify neighbors that it is still alive. The EIGRP hold-down timer gets reset each time the router receives a EIGRP Hello message. Default EIGRP adjacency Hold-down timer is 15 seconds.

Lowering EIGRP hello and hold-down timer intervals improves network convergence times (ie time to detect and respond to an outage). For Schools SRA design it is recommended to use the default EIGRP Hello and Hold timer values for the following reasons:

- *EtherChannel Benefits*—EIGRP operates over the Layer-3 EtherChannel. In the event of a single member-link failure condition, layer 2 will respond more quickly than the routing protocol, and switchover traffic from the impacted link to an alternate member link. EIGRP routing is not impacted by individual link member and no change in the routing table is required. Thus reducing the EIGRP timers will not result in quicker convergence, and may adversely impact system stability.
- *High-Availability*—The Cisco Catalyst 4500, 37xx (non-Stack Wise) and 35xx series Layer 3 switches support Stateful-Switch Over (SSO) which enables a backup supervisor to gracefully assume the active role while maintaining adjacency with neighbors, during a supervisor failure condition. The backup supervisor requires sufficient time to detect a failure and initiate graceful recovery with neighbors. Implementing aggressive timers may abruptly terminate adjacency and cause network outage before a stateful switch over is accomplished. Thus, default EIGRP Hello and Hold timers are recommended on Cisco Catalyst 4500, 37xx (non-Stackwise) and 35xx Series Layer-3 platforms.

Deploying A Multi-Layer Access Network

Multilayer design is one of the two access-distribution block designs included in the Schools Service Ready Architecture. This section provides implementation and best practices guidelines the multi-layer design. The deployment and configuration guidelines for the multi-layer access-distribution block are the same for both district office and school site networks.

Spanning-Tree in Multilayer Network

Spanning Tree (STP) is a Layer-2 protocol that prevents logical loops in switched networks with redundant links. The Schochool SRA design uses Etherchannel (point-to-point logical Layer-2 bundle) connection between access-layer and distribution switch which inherently simplifies the STP topology and operation. In this design, the STP operation is done on a logical port, therefore, it will be assigned automatically in forwarding state.

Over the years, the STP protocols have evolved into the following versions:

- Per-VLAN Spanning Tree Plus (PVST+)—Provides a separate 802.1D STP for each active VLAN in the network.
- EEE 802.1w-Rapid PVST+—Provides an instance of RSTP (802.1w) per VLAN. It is easy to implement, proven in large scale networks that support up to 3000 logical ports and greatly improves network restoration time.
- IEEE 802.1s MST—Provides up to 16 instances of RSTP (802.1w) and combines many VLANS with the same physical and logical topology into a common RSTP instance.

Following is the example configuration to enable STP protocol in multi-layer network:

Distribution

```
cr24-4507-DO(config)#spanning-tree mode rapid-pvst

cr24-4507-DO#show spanning-tree summary | inc mode
Switch is in rapid-pvst mode
```

Access-Layer Switch

```
cr24-2960-DO(config)#spanning-tree mode rapid-pvst
```

Default STP parameters optimize the network for packet forwarding. Best practice design includes hardening STP parameters in the access and distribution switch to protect against STP misconfiguration, or malicious user by deploying spanning-tree toolkit in the access-distribution block. See Figure 18.

Figure 11 Figure 18Hardening Spanning-Tree Toolkit in Multi-Layer Network

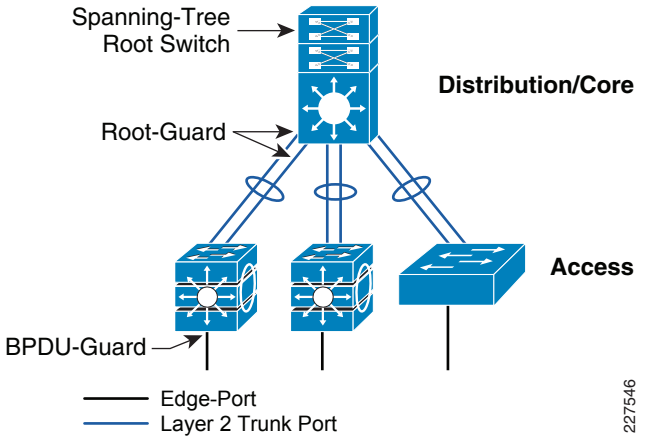

The following is the configuration deploys spanning-tree toolkit in the access-distribution block:

Distribution

```
cr24-4507-DO(config)#spanning-tree vlan 1-4094 root primary
cr24-4507-DO(config)#interface range Gig 1/1 - 2 , Gig 2/1 - 2
cr24-4507-DO(config)#spanning-tree guard root
```

Access

```
cr26-2975-DO(config)#interface GigabitEthernet1/0/1
cr26-2975-DO(config-if)#description CONNECTED TO UNTRUSTED-PC
cr26-2975-DO(config-if)#spanning-tree bpduguard enable
```

Other STP Toolkit Consideration

When the access-distribution block multi-layer design is deployed using the recommended best practices, it automatically minimizes the need for deploying the following additional spanning-tree toolkit technologies:

- *UplinkFast*—UplinkFast feature improves the network convergence time by providing direct access to the root switch link failure. UplinkFast is not necessary in this design, because there is no alternate STP path and RSTP protocol natively includes rapid recovery mechanism.
- *BackBone Fast*—BackboneFast provides rapid convergence from indirect Layer-2 link failures in a redundant distribution switch configuration. This is feature is not necessary for the same reason as stated for UplinkFast.
- *LoopGuard*—LoopGuard protects Layer-2 networks from loops that occur due to any malfunction that prevents normal BPDU forwarding. A STP loop is created when a blocking port in a redundant topology erroneously transitions to the forwarding state. This usually happens because one of the ports in a physically redundant topology (not necessarily the blocking port) stopped receiving BPDUs. Because there is single point-to-point STP forwarding port in this design, enabling Loopguard does not provide any additional benefit. UDLD protocol must be implemented to prevent STP loop that may occur in the network due to network malfunction, mis-wiring, etc.

Access Network Design Options

VLAN assignment can have a significant impact on network performance and stability. There are three basic ways to assign VLANs within the access-distribution block.

Flat Logical Network Design

Spanning a single VLAN across multiple access-layer switches is much simpler with a single collapsed core-distribution device versus a design with redundant distribution devices. The flat multi-layer design has a single VLAN across multiple access devices, as shown in Figure 12.

Figure 12 Multi-Layer Flat Network Design

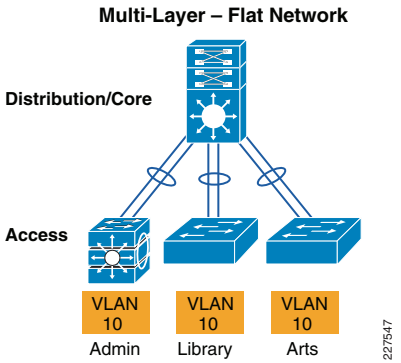

A flat multi-layer network deployment introduces the following challenges:

- *Scalability*—Spanning the same VLAN in different access-layer switches will create a large Layer-2 broadcast domain that dynamically discovers and populates MAC address entries for endpoints that may not need to communicate. In a large network, this may become a scalability issue (ie memory required to hold large CAM table)
- *Performance*—In a large network, spanning a large number of broadcast domains will impact the performance of all network devices in the access-distribution block, because the switch will have to process many more broadcast packets such as ARP
- *Security*—The flat muli layer design widens the fault domain which increases possible attacks to a larger number of users.The number of users is not necessarily due to the number switches spanned and applications during DoS or viruses attack.

Segmented Logical Network Design

Best practice design includes identifying meaningful groups within the user community, and assigning a unique VLAN to each group. These groups may be departments, user groups, or any other logical grouping of users. Enabling a unique VLAN for each group will segment the network and build a logical network structure. .All network communication between groups will pass through the routing and forwarding policies defined at the distribution layer. See Figure 13.

Figure 13 Multi-Layer Segmented Network Design

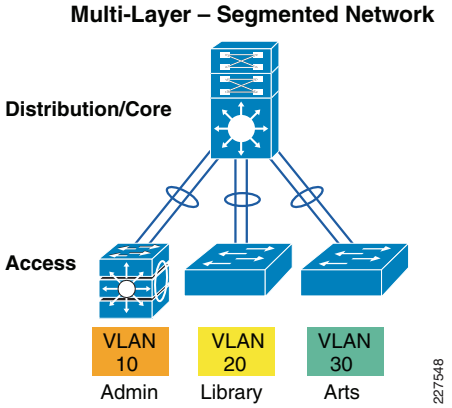

A Segmented VLAN design is the solution to the challenges described in the flat network design. VLAN segmentation improves the scalability, performance, and security of the network.

Hybrid Logical Network Design

The Segmented logical network design improves scalability, performance and security, and addresses the challenges of a flat network design. In real world deployments, there is usually a need for some users or applications to communicate with all users (eg system administrator). The Hybrid network design is the segmented design, with the addition of a exceptional VLAN which spans the entire access-distribution block. See Figure 14.

Figure 14 Multi-Layer Hybrid Network Design

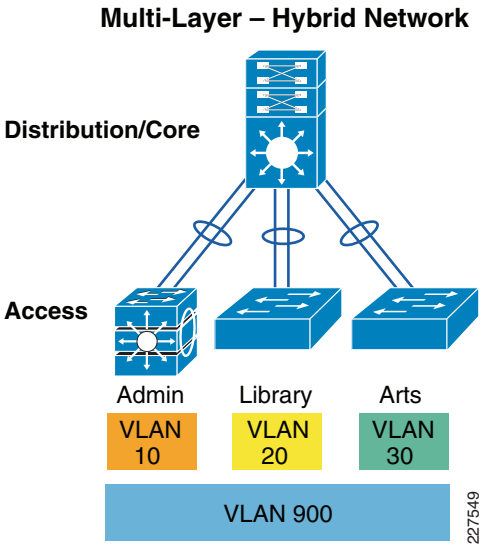

Cisco recommends the segmented VLAN network design and optionally hybrid network for centralized users or applications that requires distributed function across the access-layer network.

Following are the sample VLAN configuration steps in the access and the distribution layer switches:

- Distribution

  VLAN Trunking Protocol (VTP) is a Cisco proprietary Layer -messaging protocol that manages the addition, deletion, and renaming of VLANs on a network-wide basis. Cisco's VTP simplifies administration in a switched network. VTP can be configured in three modes: server, client, and transparent. Set the VTP domain name and change the mode to the transparent mode as follows:

```
cr24-4507-DO(config)#vtp domain District-Office
cr24-4507-DO(config)#vtp mode transparent

cr24-4507-DO(config)#vlan 10
cr24-4507-DO(config-vlan)#name cr24-3750-Admin-Dept
cr24-4507-DO(config-vlan)#vlan 20
cr24-4507-DO(config-vlan)#name cr24-3560-Library-Dept
cr24-4507-DO(config-vlan)#vlan 30
cr24-4507-DO(config-vlan)#name cr24-2960-Arts-Dept
```

- Access

  Set VTP domain name and change the mode to the transparent mode as follows:

```
cr24-3750-DO(config)#vtp domain District-Office
cr24-3750-DO(config)#vtp mode transparent

cr24-3750-DO(config)#vlan 10
cr24-3750-DO(config-vlan)#name cr24-3750-Admin-Dept
```

Implementing Layer 2 Trunk

In a typical network design, a single access switch will have more than one VLAN, for example a Data VLAN and a Voice VLAN. The network connection between Distribution and Access device is a trunk. VLAN's tag their traffic to maintain separation between VLANS across the trunk. By default on Cisco Catalyst switches, the native VLAN on each layer 2 trunk port is VLAN 1, and cannot be disabled or removed from VLAN database. The native VLAN remains active on all access switches layer 2 ports.

There are two choices for encapsulating the tagged VLAN traffic on the trunk: IEEE 802.1Q or Cisco ISL. It is recommended to implement trunk encapsulation in static mode instead of negotiating mode, to improve the rapid link bring-up performance. Not all Cisco Catalyst platforms support ISL encapsulation; therefore IEEE 802.1Q is recommended, and validated in the access and distribution switches.

Enabling the layer-2 trunk on a port-channel, automatically enables communication for all of the active VLANs between the access and distribution. This means an access-switch which has implemented, for example, VLANs 10 to 15, will receive flood traffic destined for VLANs 20 to 25, which are implemented on another access switch.. RPVST +, using logical ports, operates on a per-VLAN basis to load balance traffic. In a large network, it is important to limit traffic on layer 2 trunk ports to only the assigned VLANS, to ensure efficient and secure network performance. Allowing only assigned VLANs on a trunk port automatically filters rest.

The default native VLAN must be properly configured to avoid several security risks - Attack, worm and virus or data theft.. Any malicious traffic originated in VLAN 1 will span across the access-layer network. With a VLAN-hopping attack it is possible to attack a system which does not reside in VLAN 1. Best practice to mitigate this security risk is to implement a unused and unique VLAN ID as a native VLAN on the Layer-2 trunk between the access and distribution switch. For example, configure VLAN 802 in the access-switch and in the distribution switch. Then change the default native VLAN setting in both the switches. Thereafter, VLAN 802 must not be used anywhere for any purpose in the same access-distribution block.

Following is the configuration example to implement Layer-2 trunk, filter VLAN list and configure the native-VLAN to prevent attacks on port channel interface. When the following configurations are applied on port-channel interface (i.e., Port-Channel 11), they are automatically inherited on each bundled member-link (i.e., Gig1/1 and Gig2/1):

Distribution

```
cr24-4507-DO(config)#vlan 802
cr24-4507-DO(config-vlan)#name Admin-Hopping-VLAN

cr24-4507-DO(config)#interface Port-channel 11
cr24-4507-DO(config-if)# description Connected to cr24-3750-DO
cr24-4507-DO(config-if)# switchport
cr24-4507-DO(config-if)# switchport mode trunk

cr24-4507-DO(config-if)# switchport trunk allowed vlan 101-110,900
cr24-4507-DO(config-if)# switchport trunk native vlan 802



cr24-4507-DO#show interface port-channel 11 trunk

Port        Mode             Encapsulation  Status        Native vlan
Po11        on               802.1q         trunking      802

Port        Vlans allowed on trunk
Po11        101-110,900

Port        Vlans allowed and active in management domain
Po11        101-110,900

Port        Vlans in spanning tree forwarding state and not pruned
Po11        101-110,900
```

Access-switch

```
cr24-3750-DO(config)#vlan 802
cr24-3750-DO(config-vlan)#name Admin-Hopping-VLAN

cr24-3750-DO(config)#interface Port-channel 1
cr24-3750-DO(config-if)# description Connected to cr24-4507-DO
cr24-3750-DO(config-if)# switchport
cr24-3750-DO(config-if)# switchport mode trunk
cr24-3750-DO(config-if)# switchport trunk allowed vlan 101-110,900
cr24-3750-DO(config-if)# switchport trunk native vlan 802
```

Unidirectional Link Detection

UDLD is a Layer 2 protocol that works with the Layer 1 features to determine the physical status of a link. At Layer 1, auto-negotiation takes care of physical signaling and fault detection. UDLD performs tasks that auto-negotiation cannot perform, such as detecting the identity of neighbors and shutting down misconnected ports. When both auto-negotiation and UDLD are enabled, Layer 1 and Layer 2 detection works together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

Copper media ports use Ethernet link pulse as a link monitoring tool and are not susceptible to unidirectional link problems. Because one-way communication is possible in fiber-optic environments, mismatched transmit/receive pairs can cause a link up/up condition even though bidirectional upper-layer protocol communication has not been established. When such physical connection errors occur, it can cause loops or traffic black holes. UDLD functions transparently on Layer 2 or Layer 3 physical ports. UDLD operates in one of two modes:

- *Normal mode*—If bidirectional UDLD protocol state information times out; it is assumed there is no-fault in the network, and no further action is taken. The port state for UDLD is marked as undetermined. The port behaves according to its STP state.
- *Aggressive mode*—If bidirectionalUDLD protocol state information times out,, UDLD will attempt to reestablish the state of the port, if it detects the link on the port is operational. Failure to reestablish communication with UDLD neighbor will force the port into the err-disable state. That must be manually recovered by user or the switch can be configured for auto recovery within specified interval of time.

Following is the configuration example to implement UDLD protocol:

Distribution

```
cr24-4507-DO(config)#interface range Gig 1/2 , Gig 2/2
cr24-4507-DO(config-int)#udld port

cr24-4507-DO#show udld neighbor
PortDevice Name   Device ID     Port IDNeighbor State
------     ----------------    ------------ -------- --------------
Gi1/2    FOC1318Y06V    1            Gi1/0/49Bidirectional
Gi2/2    FOC1318Y06J     1            Gi3/0/49Bidirectional
```

Access

```
cr26-2975-DO(config)#interface Gig 1/0/49 , Gig 3/0/49
cr26-2975-DO(config-if)#description Connected to cr24-4507-DO
cr26-2975-DO(config-if)#udld port

cr26-2975-DO#show udld neighbor
PortDevice Name   Device ID     Port IDNeighbor State
------     ----------------    ------------ -------- --------------
Gi1/0/49 FOX1216G8LT     1            Gi1/2      Bidirectional
Gi3/0/49 FOX1216G8LT     1            Gi2/2      Bidirectional
```

Deploying a Routed-Access Network

This section provides implementation and best practices guidelines to deploy routed-access in the access-distribution block. The routed access design moves the boundary between Layer 2 and Layer 3 from the distribution layer to the access layer as shown in Figure 15.

Figure 15 Control Function in Multi-Layer and Routed-Access Network Design

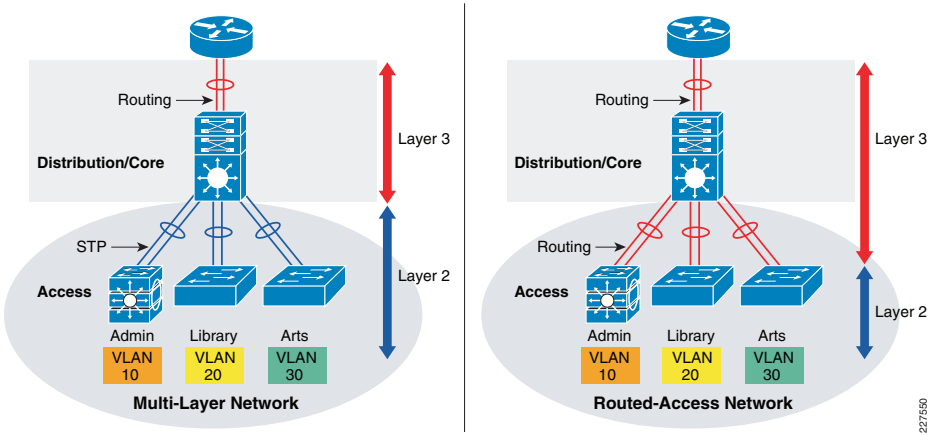

Routing in the access-layer simplifies configuration, optimizes distribution performance, and improves end-to-end troubleshooting tools. Implementing routing in the access-layer replaces Layer-2 trunk configuration with single point-to-point Layer-3 interface in distributon layer. Placing Layer-3 function one tier down on access-switches, changes the multilayer network topology and forwarding path. Implementing Layer-3 function in the access-switch does not require a physical or logical link reconfiguration; the same EtherChannel in access-distribution block can be used.

At the network edge, Layer-3 access-switches provides an IP gateway and become the Layer-2 demarcation point to locally connected endpoints that could be logically segmented into multiple VLANs. Following are the benefits of implementing routed-access in the access-distribution block:

- Eliminates the need to implement STP and the STP toolkit in the distribution layer. As a best practice, STP toolkit must be hardened at the access-layer.
- Shrinks the Layer-2 fault domain, which minimizes the number of endpoints affected by a DoS/DDoS attack.
- Improves Layer-3 uplink bandwidth efficiency by suppressing Layer-2 broadcasts at the access edge port
- Improves performance by reducing resource utilization in collapsed core-distribution layer. In a large multilayer network, the aggregation layer may consume more CPU cycles due to the large number of MAC and ARP discovery and processing and storing required for each end-station. Routed-access reduces the load of this layer 2 processing and storage in the distribution layer, by moving the load to Layer-3 access-switches. Figure 16 illustrates where Layer-2 and Layer-3 forwarding entry processing and storage takes plance when access-distribution block is implemented as multi-layer versus routed-access network.

Figure 16 Forwarding entry development in multi-tier network

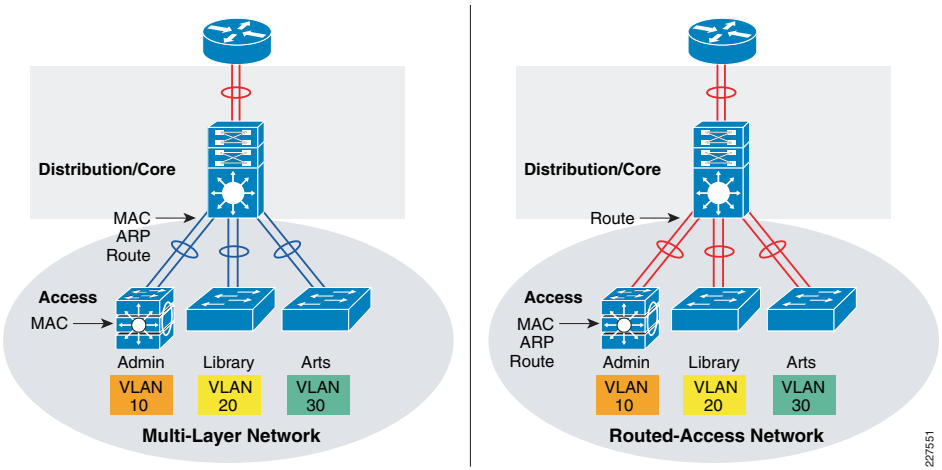

While the routed access design is appropriate for many school networks it is not suitable for all environments. Routed access does not allow a VLAN to span multiple access switches. Refer to following URL for detailed design guidance for the routed access distribution block design:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/routed-ex.html

Implementing EIGRP Routing in Access-Distribution Block

The School Service Ready Architecture uses EIGRP routing protocol, and all the devices in the LAN and WAN sub-networks are deployed in a single AS. This subsection focuses on implementing EIGRP in the access-distribution block. All the deployment and configuration guidelines in this section are the same for deploying in the district office or school site network.

Following is the example configuration to enable basic EIGRP routing in the distribution layer and in the access layer:

Distribution

```
cr24-4507-DO(config)#interface Port-channel13
cr24-4507-DO(config-if)# description Connected to cr24-3560r-DO
cr24-4507-DO(config-if)#no switchport
cr24-4507-DO(config-if)# ip address 10.125.32.0 255.255.255.254

cr24-4507-DO(config)#router eigrp 100
cr24-4507-DO(config-router)# no auto-summary
cr24-4507-DO(config-router)# eigrp router-id 10.125.100.1
cr24-4507-DO(config-router)# network 10.125.0.0 0.0.255.255

cr24-4507-DO#show ip eigrp neighbor port-channel 13
EIGRP-IPv4:(100) neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO   Q  Seq
                                            (sec)         (ms)       Cnt Num
3   10.125.32.1             Po13              14 00:02:14    2         200   0  385
```

Access

```
cr24-3560r-DO(config)#interface Loopback0
cr24-3560r-DO(config-if)# ip address 10.125.100.4 255.255.255.255
cr24-3560r-DO(config-if)#
cr24-3560r-DO(config-if)#interface Port-channel1
cr24-3560r-DO(config-if)# description Connected to cr24-4507-DO
cr24-3560r-DO(config-if)# no switchport
cr24-3560r-DO(config-if)# ip address 10.125.32.1 255.255.255.254


cr24-3560r-DO(config)#ip routing

cr24-3560r-DO(config)#router eigrp 100
cr24-3560r-DO(config-router)# no auto-summary
cr24-3560r-DO(config-router)# eigrp router-id 10.125.100.4
cr24-3560r-DO(config-router)# network 10.125.0.0 0.0.255.255

cr24-3560r-DO#show ip eigrp neighbor port-channel 1
EIGRP-IPv4:(100) neighbors for process 100
H   Address                 Interface       Hold  Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.125.32.0             Po1               13  00:10:00    1     200  0
176
```

Building EIGRP Network Boundary

EIGRP creates and maintains a single flat routing network topology between EIGRP peers. Building a single routing domain enables complete network visibility and reach ability between all of the elements within the network.(access, distribution, core, data center, WAN, etc)

In a tiered design, the access layer always has a single physical or logical forwarding path to the distribution layer. The access switch will build a forwarding topology pointing to same distribution switch as a single Layer-3 next-hop. Since the distribution switch provides a gateway function to the access switch, the routing design can be optimized with the following two techniques to improve performance and network convergence in the access-distribution block:

- Deploy Layer 3 access-switch in EIGRP stub mode
- Summarize network view to Layer-3 access-switch for intelligent routing function

Deploy Layer 3 Access-Switch in EIGRP Stub Mode

The Layer-3 access switch can be deployed to announce itself as a stub router that acts as a non-transit router and does not connect any other Layer-3 stub or non-stub routers. Announcing itself as a non-transit stub Layer-3 router is one way to notify the distribution router that it should not include the Layer-3 access switch in the EIGRP topology recomputation process. This optimized recomputation process will prevent unnecessary EIGRP network queries, which reduces network traffic, and simplifies the route computation. As illustrated in Figure 17, implementing EIGRP stub function in the access switches, greatly reduces the number of EIGRP network queries.

Figure 17 EIGRP Query Path with and without Stub Implementation

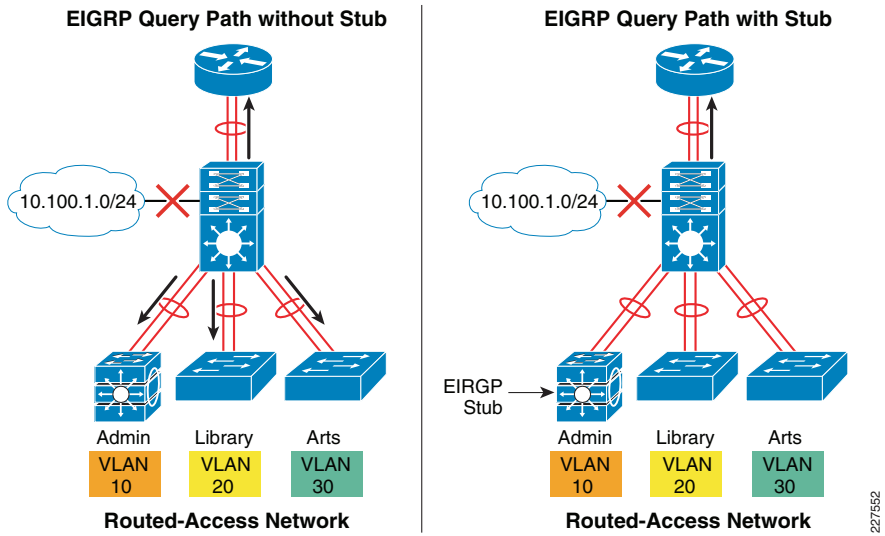

EIGRP stub router in Layer-3 access-switch can announce routes to a distribution-layer router with great flexibility.

EIGRP stub router can be deployed to announce routes dynamically discovered or statically configured. Best practice design is to deploy EIGRP stub router to announce locally learned routes to aggregation layer.

Following is the example configuration to enable EIGRP stub routing in the Layer-3 access-switch, no configuration changes are required in distribution system:

Access

```
cr24-3560r-DO(config)#router eigrp 100
cr24-3560r-DO(config-router)#eigrp stub connected
cr24-3560r-DO#show eigrp protocols detailed
Address Family Protocol EIGRP-IPv4:(100)
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  EIGRP NSF-aware route hold timer is 240
  EIGRP stub, connected
  Topologies : 0(base)
```

Distribution

```
cr24-4507-DO#show ip eigrp neighbors detail port-channel 13
EIGRP-IPv4:(100) neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.125.32.1             Po13              13 00:19:19   16   200  0  410
   Version 12.2/3.0, Retrans: 0, Retries: 0, Prefixes: 11
   Topology-ids from peer - 0
   Stub Peer Advertising ( CONNECTED ) Routes
   Suppressing queries
```

Summarizing the Routed Networks

Enabling the EIGRP stub function on the access switch does not change the distribution router behavior of forwarding the full EIGRP topology table. The Distribution router must be configured to advertise summarized routes that do not compromise end-to-end reach ability, and help access switches maintain minimal routing information. In a network with a well designed IP addressing scheme, the aggregation system can advertise summarized routes in a classless address configuration, that reduce individual network advertisements, improve network scalability and network convergence. The,distribution router must have full network topology information to ensure efficient reachability paths. Hence it is recommended to summarize at the distribution router, and not summarize at the access-layer.

Route summarization must be implemented on the distribution layer of district office and each school site network. This includes devices such as the WAN aggregation in the district office. The distribution router must advertise the following summarized network information to Layer 3 access-switch:

- *Local Network*—Distribution router can be implemented in hybrid access-distribution configuration that interconnects several multi-layer or routed-access enabled access-layer switches. Independent of route origination source (connected or dynamic route) and network size within the access-distribution block, the distribution router in district office and school site network must advertise a single, concise and summarized Layer 3 network to each Layer 3 access-switch and to core devices.
- *Remote Network*—Summarized network will be propagated dynamically across the network. Single summarization of all remote networks may be advertised to local Layer 3 access-switchs, since it improves bandwidth efficiency. During a network outage, Layer 3 access-switch may drop traffic at the network edge instead of transmitting it to the distribution router to black hole traffic.
- *WAN Network*—Announcing a single summarized WAN network provides flexibility to troubleshoot and verify network availability.
- *Default Network*—When Layer 3 access-switch receives unknown destination traffic from the edge that does not match any of the above mentioned summarized networks, then it is sent to the distribution router to make a forwarding decision. The distribution router performs a forwarding table lookup and may forward to appropriate path or black hole the traffic. In a typical school environment, a default route is announced by an Internet edge system, to forward all internet traffic. Distribution router must propagate this default route to the Layer 3 access-switch.

Figure 18 illustrates a summarized EIGRP network advertisement, by route aggregation system, that provides end-to-end internal and external network reachability.

Figure 18 End-to-End Routed-Access Network

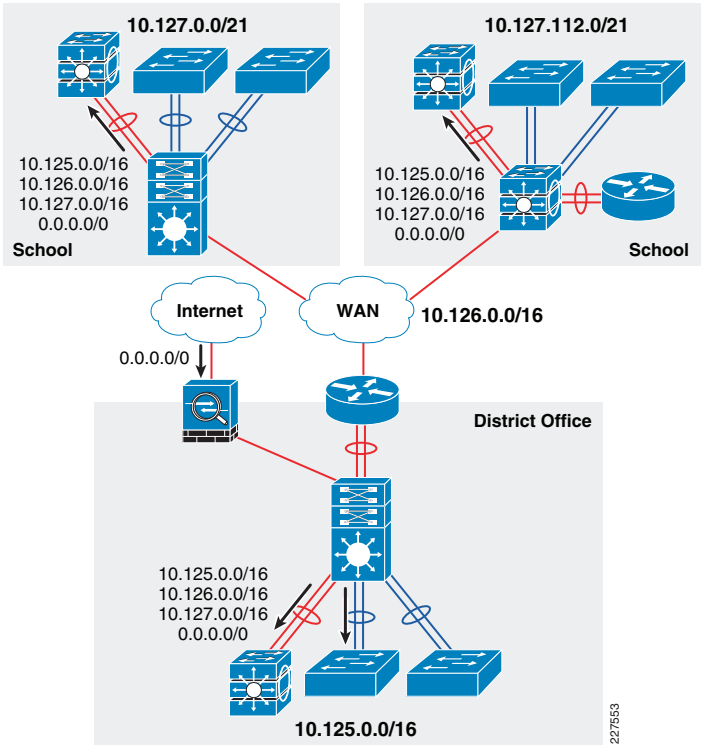

Following is configuration example to deploy summarized and filtered Layer-3 network information to Layer-3 access-switch.

Distribution

```
interface Port-channel13
 description Connected to cr24-3560r-DO
 dampening
 ip address 10.125.32.0 255.255.255.254
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 eigrp-key
 ip summary-address eigrp 100 10.125.0.0 255.255.0.0 5
 load-interval 30
 carrier-delay msec 0
!
!configure ACL and route-map to allow summarized route advertisement to Layer 3
access-
switch
!
access-list 1 permit 0.0.0.0
access-list 1 permit 10.126.0.0
access-list 1 permit 10.127.0.0
access-list 1 permit 10.125.0.0
!
route-map EIGRP_STUB_ROUTES permit 10
 match ip address 1
!
```

```
router eigrp 100
 distribute-list route-map EIGRP_STUB_ROUTES out Port-channel13


cr24-4507-DO#show ip protocols | inc Outgoing|filtered
  Outgoing update filter list for all interfaces is not set
    Port-channel13 filtered by
```

Access

```
cr24-3560r-DO#show ip route eigrp
     10.0.0.0/8 is variably subnetted, 15 subnets, 4 masks
D       10.126.0.0/16 [90/3328] via 10.125.32.0, 01:37:21, Port-channel1
D       10.127.0.0/16 [90/3584] via 10.125.32.0, 01:37:21, Port-channel1
D       10.125.0.0/16 [90/1792] via 10.125.32.0, 01:34:29, Port-channel1
D*EX 0.0.0.0/0 [170/515072] via 10.125.32.0, 00:03:15, Port-channel1
cr24-3560r-DO#
```

EIGRP Adjacency Protection

EIGRP adjacency protection guidelines discussed earlier for the core network, apply equally to routed access in the access-distribution block. ,The two challenges, system efficiency, and network security also apply equally to the routed access design, and the same solution is applied.

- *System efficiency*—EIGRP hello transmission must be blocked on an interface where there are no trusted EIGRP neighbors, to reduce CPU utilization and prevent network attacks. EIGRP routing process should only be enabled on interfaces where trusted school devices are connected. All other interfaces can be suppressed in passive mode.

  Following is the example configuration on Layer-3 access-switch that advertises networks enabled on SVI interfaces; however, keeps them in passive mode and explicitly allows EIGRP function on uplink port-channel to distribution router. Same configuration principle must be applied on each EIGRP router including distribution and core routers:

```
cr24-3560r-DO(config)#router eigrp 100
cr24-3560r-DO(config-router)# network 10.125.0.0 0.0.255.255
cr24-3560r-DO(config-router)# passive-interface default
cr24-3560r-DO(config-router)# no passive-interface Port-channel1

cr24-3560r-DO#show ip eigrp interface
EIGRP-IPv4:(100) interfaces for process 100

                      Xmit           Queue   Mean   Pacing Time   Multicast
Pending
Interface        Peers      Un/Reliable  SRTT   Un/Reliable   Flow Timer  Routes
Po1                  1        0/0         1       0/1           50           0

cr24-3560r-DO#show ip protocols | inc Passive|Vlan
  Passive Interface(s):
    Vlan1
    Vlan11
    Vlan12
    Vlan13
    Vlan14
```

- *Network Security*—EIGRP adjacency between distribution and Layer-3 access-switch must be secured. Following is the example configuration to enable EIGRP neighbor authentication using MD5:

Distribution

```
cr24-4507-DO(config)#key chain eigrp-key
cr24-4507-DO(config-keychain)# key 1
cr24-4507-DO(config-keychain-key)#   key-string <password>

cr24-4507-DO(config)#interface Port-channel13
cr24-4507-DO(config-if)# description Connected to cr24-3560r-DO
cr24-4507-DO(config-if)# ip authentication mode eigrp 100 md5
cr24-4507-DO(config-if)# ip authentication key-chain eigrp 100 eigrp-key
```

Access

```
cr24-3560r-DO(config)#key chain eigrp-key
cr24-3560r-DO(config-keychain)# key 1
cr24-3560r-DO(config-keychain-key)#   key-string <password>

cr24-3560r-DO(config)#interface Port-channel1
cr24-3560r-DO(config-if)# description Connected to cr24-4507-DO
cr24-3560r-DO(config-if)# ip authentication mode eigrp 100 md5
cr24-3560r-DO(config-if)# ip authentication key-chain eigrp 100 eigrp-key
```

Tuning EIGRP Protocol Timers

EIGRP protocol functions the same in routed-access as it does in the core network. It is highly recommended to retain default EIGRP hello and hold timers on distribution and Layer 3 access-switch and rely on EtherChannel and SSO-based recovery mechanisms, that offers sub-second network convergence, during individual link or supervisor failure scenarios.

Building a Resilient Network

The Schools Service Ready Architecture is a high performance, resilient and scalable network design. A network outage may be caused by the system, human error, or natural disaster. The Schools SRA is designed to minimize the impact of a failure regardless of the cause. Network outages may be either Planned or Unplanned.

- *Planned Outage*—Planned network outage occurs when a portion of the network is taken out of service as part of a scheduled event (e.g., a software upgrade).
- *Unplanned Outage*—Any unscheduled network outage is considered an unplanned outage. Such outages may be caused by internal faults in the network, or devices due to hardware or software malfunctions.

The network is designed to recover from most un planned outages in less than a second (milliseconds). In many situations, the user will not even notice the outage occurred. If the outage lasts longer (several seconds), then the user will notice the lack of application responsiveness. The network is designed to minimize the overall impact of a unplanned network outage, and gracefully adjust and recover from many outage conditions. Figure 19 shows an example of a real-time VoIP application and user impact depending on duration of outage event.

Figure 19 VoIP User Impact for minor and major network outage

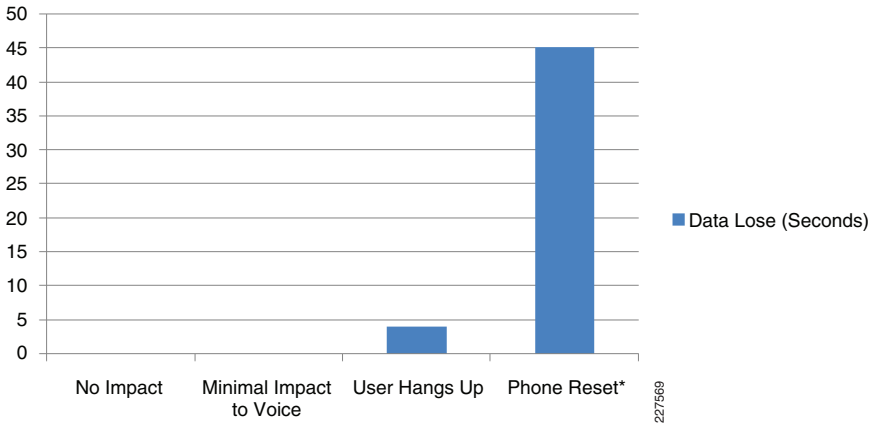

Several techniques are used to make the network design more resilient. Deploying redundant devices and redundant connections between devices, enables the network to recover from fault conditions. Identifying critical versus non-critical applications, and network resources optimizes cost performance, by focusing on the most important elements of the network design. The resiliency of a system design (see Figure 20) is often catagorized as follows:

- *Network Resiliency*—Provides redundancy during physical link outages (e.g., fiber cut, bad transceivers, incorrect cabling, etc).
- *Device Resiliency*—Protects network during device outage triggered by hardware or software (e.g. software crash, non-responsive supervisor, etc).
- *Operational Resiliency*—Capabilities which provide network availability even during planned network outage conditions (e.g., ISSU features which enable software upgrades while device is operating).

Figure 20 Resiliency Deployment Strategy

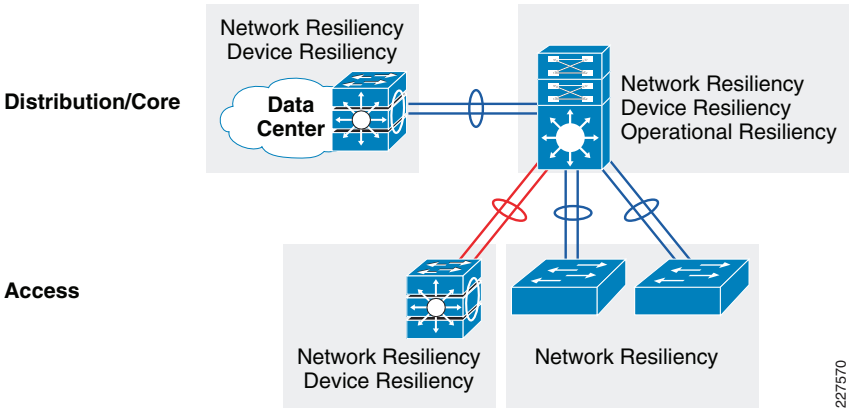

The high availability framework is based upon the three resiliency categories described in the previous section. Figure 21 shows which technologies are implemented to achieve each category of resiliency.

Figure 21 High-Availability Categories and Technologies

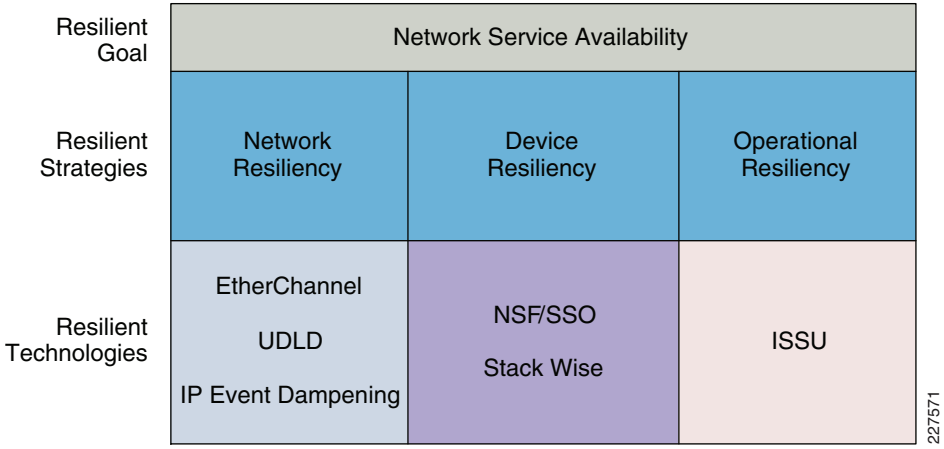

Redundant Hardware Components

Redundant hardware implementations vary between fixed configuration and modular Cisco Catalyst switches. Selective deployment of redundant hardware is an important element of the Schools SRA design which delivers device resiliency.

Redundant Power System

Redundant power supplies protect the device from power outage or power supply failure. Protecting the power is not only important for the network device, but also the endpoints that rely on power delivery over the Ethernet network. Redundant power supplies are deployed differently depending on the switch type:

- *Modular Switch*—Dual power supplies can be deployed in the modular switching platforms like the Cisco Catalyst 4500-E. By default, the Cisco Catalyst 4500 power supply operates in 1+1 redundant mode (both power supplies are active).
- *Fixed configuration switch*—Fixed configuration switches are deployed with internal power supplies and they may also use Cisco RPS 2300 external power supply. A single Cisco RPS 2300 power supply has modular power supplies and fans to deliver power to multiple switches. Deploying internal and external power-supplies provides a redundant power solution for fixed configuration switches.

Redundant Network Connectivity

Redundant network connections protect the system from failure due to cable or transceiver faults. Redundant network connections attached to a single fixed configuration switch or network module in the Cisco Catalyst 4500 switch do not protect against internal device hardware or software fault.

Best practice design is to deploy redundant network modules within the Catalyst 4500 switch and the Cisco 3750-E StackWise Plus solution in the small school site collapsed core network. Deploying the 3750-E StackWise Plus in critical access-layer switches in the data center network and in the district office is also best practice. Connecting redundant paths to different hardware elements provides both network and device resiliency. See Figure 22.

Figure 22 Redundant Network Connectivity

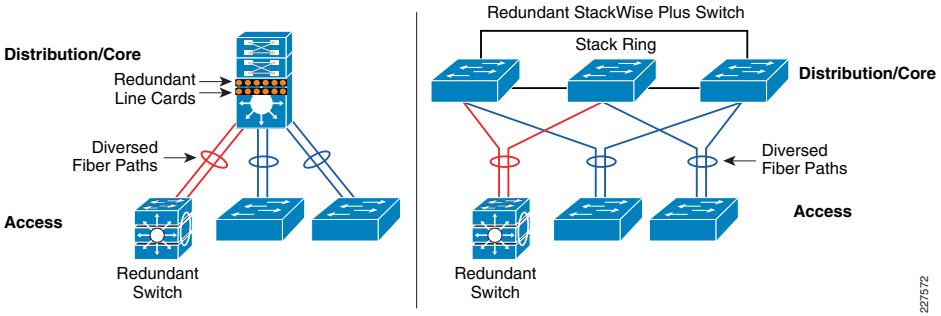

Redundant Control-Plane

The processing software operation is different in standalone or StackWise fixed configuration switches, and on a supervisor module of a modular switch. Network communication and forwarding operations can be disrupted when the processing unit fails, causing a network outage. Network recovery techniques vary based on the different platforms.

The standalone and non-stackable fixed configuration switches like the Cisco Catalyst 2960 or 3560-E feature power redundancy and network resiliency support; however they do not protect against a processing unit failure. During a processing unit failure event, all endpoints attached to the switch are impacted and network recovery time is undeterministic.

Device resiliency in Cisco StackWise and modular switching platforms provides 1+1 redundancy with enterprise-class high availability and deterministic network recovery time.

Cisco StackWise

Cisco Catalyst 2975 and Catalyst 3750-E switches can be deployed in StackWise mode using a special stack cable. Up to nine switches can be integrated into a single stack that delivers distributed forwarding architecture and unified single control and management plane. Device level redundancy in StackWise mode is achieved via stacking multiple switches using the Cisco StackWise technology. One switch from the stack is selected automatically to serve as the master, which manages the centralized control-plane process. Cisco StackWise solution provides 1:N redundancy. In the event of a active master-switch outage, a new master is selected automatically. See Figure 23.

Figure 23 Cisco Stack Wise Switching Architecture

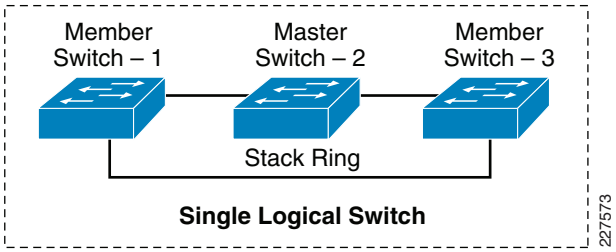

Since Cisco StackWise enables up to 9 switches to appear as one logical switch, it has centralized management and control functions. Most layer 2 and layer 3 functions are centrally performed, however layer 2 topology development is distributed, ie each switch performs the function independently. Table 2 lists network protocol functions and identifies which are centralized and which are distributed.

Table 2 Cisco StackWise Centralized and Distributed Control-Plane

Protocols		Function
Layer 2 Protocols	MAC Table	Distributed
	Spanning-Tree Protocol	Distributed
	CDP	Centralized
	VLAN Database	Centralized
	EtherChannel - LACP	Centralized
Layer 3 Protocols	Layer 3 Management	Centralized
	Layer 3 Routing	Centralized

Cisco StackWise solution offers network and device resiliency with distributed forwarding. In the event of a master switch outage, Non Stop Forwarding (NSF) enables packet forwarding to continue based on current state information, while a new master switch is selected. New master switch selection is accomplished in the range of 700 to 1000 milliseconds; the amount of time to reestablish the control-plane and develop distributed forwarding will vary depending on the size and complexity of the network.

Following is a best practice to reduce Layer-3 disruption in the event of a master switch outage: Determine the master switch with the higher switch priority, and isolate the uplink Layer-3 EtherChannel bundle path by using physical ports from member switches (i.e. don't use the master switches ports for Etherchannel uplinks). With NSF capabilities enabled, this design decreases network downtime during a master-switch outage.

An understanding of SSO and StackWise components and failover events associated with NSF provides significant insight in designing a network that enables supervisor redundancy. The following subsection uses the above concepts and principles to identify the design parameters, and applies them to develop a best-practice hierarchical network with the highest availability.

Cisco Modular Switch

The Cisco Catalyst 4500 modular switch supports redundant supervisors, and Stateful Switch Over (SSO). When deployed along with NSF, the 4500 provides a enterprise-class highly available system with network and device resiliency.

SSO is a Cisco IOS service used to synchronize critical forwarding and protocol state information between redundant supervisors configured in a single chassis. With SSO enabled, one supervisor in the system assumes the role of *active* and the other supervisor becomes the *hot-standby.* Each is ready to backup the other, thus providing 1:1 hot redundancy to protect from a control-plane outage. Since both supervisors are active, the system benefits by using the physical ports from both supervisors during normal operation.. SSO synchronizes system services such as DHCP snooping, Switched Port Analyzer (SPAN), security access control lists (ACLs), and QoS policies so ensure the switch provides the same level of protection and service after a supervisor failover event.

NSF enables packets to continue to be forwarded using existing routing table information, during switchover. NSF also provides graceful restart to the routing protocol such that during the failover, the routing protocol remains aware of the change and does not react by resetting its adjacency. If the routing protocol were to react to the failure event, and alter routing path information, the effectiveness of stateful switch over would be diminished.

Operational Resiliency Strategy

Designing the network to recover from unplanned outages is important. It is also important to consider how to minimize the disruption caused by planned outages. These planned outages can be due to standard operational processes, configuration changes, software and hardware upgrades, etc.

The same redundant components which mitigate the impact of unplanned outages can also be used to minimize the disruption caused by planned outages.. The ability to upgrade individual devices without taking them out of service is enabled by having internal component redundancy (such as with power supplies, and supervisors) complemented with the system software capabilities. Two primary mechanisms exist to upgrade software in a live network:

- Full-image In-Service Software Upgrade (ISSU) on the Cisco Catalyst 4500 leverages dual supervisors to allow for a full, in-place Cisco IOS upgrade. 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.
- Network and device level redundancy, along with the necessary software control mechanisms, guarantee controlled and fast recovery of all data flows following a fault condition, and provide the ability to manage the fault tolerant infrastructure during planned outage events. .

Deploying Resiliency in the Schools Network

Many of the design features of the Schools Service Ready Architecture which were described in 'Deploying School Foundation Services' section earlier, contribute to the network high availability capabilities. This section focuses on how to implement additional features which complete the Schools SRA high availablity design.

Network Resiliency

Etherchannel and UDLD are two design features which are included in the network foundation services, which contribute to network resiliency.

Implementing IP Event Dampening

Poor signaling or a loose connection may cause continuous port-flap (port alternates between active state and inactive state).. A single interface flapping can impact the stability and availability of the network. Route summarization is one technique which mitigates the impact of a flapping port. Summarization isolates the fault domain with a new metric announcement by the aggregator and thereby hides the local networks fault within the domain.

A best practice to mitigate local network domain instability due to port-flap, is implementing IP Event Dampening on all layer 3 interfaces. Each time the Layer-3 interface flaps the IP dampening tracks and records the flap event. Upon multiple flaps, a logical penalty is assigned to the port and suppresses link status notification to IP routing until the port becomes stable. IP event dampening is a local function and does not have a signaling mechanism to communicate with a remote system. It can be implemented on each individual physical or logical Layer-3 interface: physical ports, SVI or port-channels. Following is a example configuration to implement IP Event Damenping

Distribution/Core

```
cr24-4507-DO(config)#int range Port1 , Gig5/6 , Gig6/6 , Vlan 101 - 110
cr24-4507-DO(config-if-range)#dampening
```

```
    cr24-4507-DO#show interface dampening | be Port
    Port-channel1 Connected to cr24-3750ME-DO
    Flaps Penalty    Supp ReuseTm   HalfL  ReuseV   SuppV  MaxSTm    MaxP Restart
      0       0   FALSE       0           5       1000       2000      20
16000       0
```

The following output illustrates how the IP event dampening keeps track of port flaps and makes a decision to notify IP routing process based on interface suppression status:

```
cr24-4507-DO#debug dampening interface
cr24-4507-DO#show logging | inc EvD|IF-EvD

12:32:03.274: EvD(GigabitEthernet5/6): charge penalty 1000, new accum. penalty
1000, flap count 2
12:32:03.274: EvD(GigabitEthernet5/6): accum. penalty 1000, not suppressed
12:32:03.274: IF-EvD(GigabitEthernet5/6): update IP Routing state to DOWN,
interface is not suppressed

cr24-4507-DO#show interface dampening | be 5/6
Flaps Penalty    Supp ReuseTm   HalfL  ReuseV   SuppV  MaxSTm    MaxP Restart
```

```
2       0    FALSE       0            5       1000       2000       20
16000       0
```

In a multilayer access-distribution design, the Layer-2 and Layer-3 demarcation is at the collapsed core-distribution device. IP event dampening is enabled on per-logical VLAN (SVI) interface basis on the collapsed core device. IP event dampening becomes more effective when each access-layer switch is deployed with a unique set of Layer-2 VLANs.

Assigning unique VLANs on each access-layer switch also helps IP event dampening to isolate the problem and prevent network faults triggered in a multilayer network. The following output illustrates how IP event dampening keeps track of individual logical VLAN networks associated to same Layer-2 physical trunk ports. When a Layer-2 trunk port flaps, the state of SVI also flaps, and forces dampening to track and penalize unstable interfaces:

```
12:58:41.332: EvD(Vlan101): charge penalty 1000, new accum. penalty 2627, flap
count 3
12:58:41.332: EvD(Vlan101): accum. penalty 2627, now suppressed with a reuse
intervals of 7
12:58:41.332: IF-EvD(Vlan101): update IP Routing state to DOWN, interface is
suppressed

    cr24-4507-DO#show interface dampening
    Vlan101 Connected to cr24_2960_Dept_1_VLAN
    Flaps Penalty    Supp ReuseTm   HalfL  ReuseV   SuppV  MaxSTm    MaxP Restart
    3       71    FALSE       0            5       1000       2000      20
16000       0
```

Device Resiliency

As described earlier, redundant hardware is an important technique for achieveing device resiliency. The Schools SRA network design applies hardware redundancy considering the cost / performance tradeoffs.

Implementing Redundant Power Supply

Redundant power supplies can prevent a system outage due to power outage, power supply or fan hardware failure. All Cisco Catalyst switching platforms supports robust 1+1 redundant power capabilities that can be deployed with internal or external power source management. See Figure 24.

Figure 24 Cisco Catalyst Internal and External Power Redundancy Option

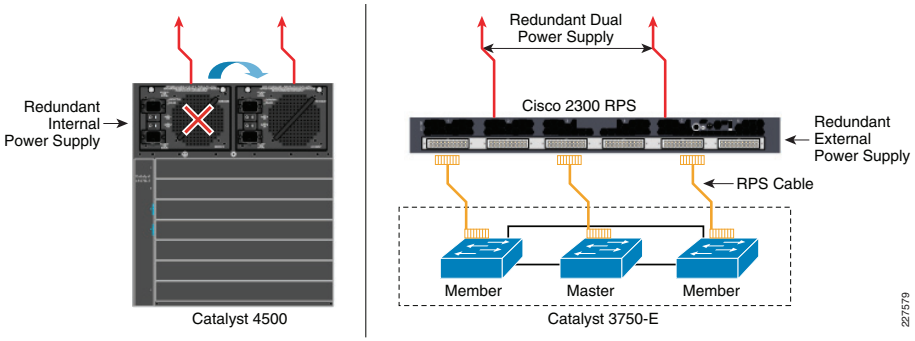

Catalyst 4500-Redundant Internal Power Supply

The Cisco Catalyst 4500 provides power to internal hardware components and external devices like IP phones. All the power is provided by the internal power supply. Dual-power supplies in the Catalyst 4500 can operate in one of two different modes:

- *Redundant Mode*—By default, Catalyst 4500 power supply operates in redundant mode offering 1+1 redundant option. The system determines power capacity and number of power supplies required based on the power required for all internal and external power components. Both power supplies must have sufficient power to support all the installed modules and operate in 1+1 redundant mode.

```
cr24-4507-DO(config)#power redundancy-mode redundant

cr24-4507-DO#show power supplies
Power supplies needed by system    :1
Power supplies currently available :2
```

- *Combined Mode*—If the system power requirement exceeds the capacity of a single power supply, then both power supplies can be combined to increase the capacity. In this mode, the power system does not provide 1+1 power redundancy. The following global configuration will enable power supplies to operate in combined mode:

```
cr24-4507-DO(config)#power redundancy-mode combined

cr24-4507-DO#show power supplies
Power supplies needed by system:2
Power supplies currently available:2
```

Catalyst 29xx and 3xxx - Redundant External Power Supply with RPS

Cisco Redundant Power Supply (RPS) 2300 provides up to 6 RPS ports to provide backup power to critical access-layer switches in the school network. Additional power resiliency can be added by deploying dual-power supply to backup to two devices simultaneously. The Cisco RPS 2300 can be provisioned through the Cisco 3750-E or 3560-E Series switches using the enable mode CLI:

```
cr36-3750s-ss100#power rps <switch id> name CiscoRPS
cr36-3750s-ss100#power rps <switch id> port <rps port id> active


cr36-3750s-ss100#show env rps

SW   StatusRPS NameRPS Serial# RPS Port#HN'
----  -------- ------------  --------------------‡ ?
1    ActiveCiscoRPSFDO1246SG3L1`
2    ActiveCiscoRPSFDO1246SG3L3
3    ActiveCiscoRPSFDO1246SG3L5

RPS Name: CiscoRPS
State: Activexs
PID: PWR-RPS2300
Serial#: FDO1246SG3L
Fan: Good
```

```
Temperature: Green

 RPS Power Supply A: Present
 PID              : C3K-PWR-1150WAC
 Serial#          : DTM124000XX
 System Power  : Good
 PoE Power: Good
 Watts            : 300/800 (System/PoE)
                          Redundant RPS
 RPS Power Supply B: PresentPower Supply
 PID              : C3K-PWR-1150WAC
 Serial#          : DTM124000WW
 System Power: Good
 PoE Power        : Good
 Watts            : 300/800 (System/PoE)

DCOut  State   Connected  Priority  BackingUp  WillBackup  Portname      SW#
-------- -------    --------------  ---------  --------------  ---------------
---------------      ---
   1  Active  Yes               6       NoYes         cr36-3750s-SS100    1
   2  Active  Yes               6       NoYes         <>       <>
   3  Active  Yes               6       NoYes         cr36-3750s_SS100   2
   4  Active  Yes               6       NoYes         <>       <>
   5  Active  Yes               6       NoYes         cr36-3750s_SS100   3
   6  Active  Yes               6       NoYes         <>       <>
```

Implementing Redundant Control Plane System

The collapsed core device in the district office and school sites (Catalyst 4500 or 3750-E StackWise) is deployed with redundant supervisor, or StackWise Plus to enable graceful recovery from switch hardware outage. Any access-switch which is deemed critical may be deployed as StackWise Plus to improve device resiliency. The implementation for each switch is different, and is discussed separately in the sections which follows

Resilient Cisco StackWise

Cisco Catalyst 2975 supports StackWise, and is used when a resilient layer-2 access switch is required. Cisco Catalyst 3750-E supports StackWise, and is used when a resilient layer 2 or layer 3 access switch is required. Cisco 3750-E StackWise Plus is deployed for the collapsed core in the small school site network.

StackWise switch provisioning is done dynamically by the StackWise protocol. Cisco IOS automatically adjusts the interface addressing and its associated configuration based on the number of provisioned switches in the stack.

```
cr26-2975-DO#show run | inc provision

switch 1 provision ws-c2975gs-48ps-l
switch 2 provision ws-c2975gs-48ps-l
switch 3 provision ws-c2975gs-48ps-l
```

Master Switch Election

The centralized control-plane and management plane is managed by the master switch in the stack. By default, the master switch selection within the ring is performed dynamically by negotiating several parameters and capabilities between each switch within the stack. Each StackWise-capable switch is by default configured with priority 1.

```
cr26-3750r-DO#show switch
Switch/Stack Mac Address : 0023.eb7b.e580
                                                 H/W   Current
Switch# Role   Mac AddressPriority     Version       State
-----------------------------------------------------------------------------
----------------
* 1 Member0023.eb7b.e580     1              0          Ready
   2Master  0026.5284.ec80      1                 0        Ready
   3Member0025.eb7b.e680      1              0          Ready
```

As described in previous section, the Cisco StackWise architecture is not SSO-capable. This means all the centralized Layer-3 functions must be reestablished with the neighbor switch during a master-switch outage. To minimize the control-plane impact and improve network convergence the Layer 3 up links should be diverse, originating from member switches, instead of the master switch. The default switch priority must be increased manually after identifying the master switch and switch number. The new switch priority becomes effective after switch reset.

```
cr26-3750r-DO(config)#switch 2 priority 15
Changing the Switch Priority of Switch Number 2 to 15

cr26-3750r-DO#show switch
Switch/Stack Mac Address : 0026.5284.ec80
                                                 H/W   Current
Switch# Role   Mac AddressPriority     Version       State
-----------------------------------------------------------------------------
----------------
    1Member0023.eb7b.e580     1              0          Ready
* 2 Master  0026.5284.ec80      15              0         Ready
   3Member0025.eb7b.e680      1              0          Ready
```

StackWise Layer-3 MAC Management

To provide a single unified logical network view in the network, the MAC addresses of Layer-3 interfaces on the StackWise (physical, logicl, SVI's, port channel) are derived from the Ethernet MAC address pool of the master switch in the stack. All the Layer-3 communication from the StackWise switch to the endpoints (like IP phone, PC, servers and core network system) is based on the MAC address pool of the master switch.

```
cr26-3750r-DO#show switch
Switch/Stack Mac Address : 0026.5284.ec80
                                                 H/W   Current
Switch# Role   Mac AddressPriority     Version       State
```

```
------------------------------------------------------------------------
---------------
    1Member0023.eb7b.e580     1               0          Ready
* 2 Master  0026.5284.ec80      15              0          Ready
   3Member0025.eb7b.e680      1               0          Ready

cr26-3750r-DO#show version
.  .  .
Base ethernet MAC Address       : 00:26:52:84:EC:80
.  .  .
```

After a master-switch outage, the new master switch in the stack assigns new MAC addresses to all Layer-3 interfaces, from the local MAC address pool. Once the new MAC address is assigned, it will force the switch to generate a gratuitous ARP in the network to make sure no other system is using the same MAC address. The default timer to retain the MAC address from the failed master switch is four minutes. While the new MAC address is not assigned on Layer-3 interface and not being propagated and updated in the network, the traffic will blackhole in the network.

```
cr26-3750r-DO#reload slot 2
Proceed with reload? [confirm]
Switch 2 reloading...

cr26-3750r-DO#show switch
Switch/Stack Mac Address : 0023.eb7b.e580
                                          H/W   Current
Switch# Role   Mac AddressPriority     Version       State
------------------------------------------------------------------------
---------------
* 1 Master  0023.eb7b.e580     1              1          Ready
   2Member000.0000.0000      0              1         Removed
   3Member0025.eb7b.e680      1              1          Ready
```

To prevent this network instability, the old MAC address assignments on Layer-3 interfaces can be retained even after the master switch fails. The new active master switch can continue to use the MAC addresses assigned by the old master switch, which prevents ARP and routing outages in the network. The default stack-mac timer settings must be changed in Cisco Catalyst 2975 and 3750-E StackWise switch mode using the global configuration CLI mode as shown below:

```
cr26-3750r-DO(config)#stack-mac persistent timer 0

cr26-3750r-DO#show switch
Switch/Stack Mac Address : 0026.5284.ec80
Mac persistency wait time: Indefinite

                                          H/W   Current
Switch# Role   Mac AddressPriority     Version       State
------------------------------------------------------------------------
---------------
    1Member0023.eb7b.e580     1               0          Ready
* 2 Master  0026.5284.ec80      15              0          Ready
   3Member0025.eb7b.e680      1               0          Ready
```

Non-Stop Forwarding (NSF)

The Cisco Catalyst 3750-E switch in StackWise mode is not SSO-capable. When the master switch fails, the new master switch is required to reform the Layer-3 adjacencies with the neighbors in the network. The forwarding architecture in StackWise switch is designed to provide non-stop forwarding during the master switch outage using NSF technology. Each 3750-E switch in the stack maintains distributed Layer-3 FIB from the old master switch and continues to forward upstream traffic, until they are updated by the new master in the stack ring

To enable NSF capability, explicit configuration must be enabled under the routing process. NSF-aware feature is enabled by default on all Layer-3 Ethernet switches to function in helper mode to perform graceful recovery during NSF-capable Cisco 3750-E master switch outage. NSF-capable system can also operate in NSF aware role:

NSF Capable Layer-3 Switch

```
cr36-3750s-SS100(config)#router eigrp 100
cr36-3750s-SS100(config-router)#nsf


cr36-3750s-SS100#show ip protocols | inc NSF
*** IP Routing is NSF aware ***
  EIGRP NSF-aware route hold timer is 240
  EIGRP NSF enabled
     NSF signal timer is 20s
     NSF converge timer is 120s

NSF-Aware Layer-3 Switch

cr24-3560r-DO#show ip protocols | inc NSF
*** IP Routing is NSF aware ***
  EIGRP NSF-aware route hold timer is 240
```

NSF Timers

As depicted in the above show commands, the default NSF-aware system hold timer is 240 seconds. Lowering the timer value may abruptly terminate graceful recovery, causing network instability.. Best practice is to use the default NSF hold timer, unless it is observed that NSF recovery takes longer than 240 seconds.

600 seconds after a graceful-recovery starts on a NSF-aware system , NSF clears the route stale marking and resumes using the synchronized routing database.

! NSF Aware received graceful-restart message from new master switch

```
11:56:15.365: %DUAL-5-NBRCHANGE: EIGRP-IPv4:(100) 100: Neighbor 10.125.32.3
(Port-channel15) is resync: peer graceful-restart
11:56:15.365: EIGRP: NSF: AS100, NSF or GR initiated by 10.125.32.3 at 00:00:00, flags 0x4

! NSF route hold timer expires and searches and removes all stale route entries
received graceful-restart message from new master switch
12:00:15.392: EIGRP: NSF: AS100. route hold timer expiry
12:00:15.392: DUAL: Search for outdated routes from 10.125.32.3
```

Resilient Cisco Catalyst 4500

A modular switching platform like the Cisco Catalyst 4500 is fully NSF/SSO-capable, providing 1+1 control plane redundancy. In the Catalyst 4500, all the intelligent Layer-2 and Layer-3 functions are performed centrally on the supervisor module. Deploying redundant supervisor in SSO mode in same system will allow the primary supervisor to fully synchronize the adjacencies, forwarding, configuration, counters, and more information on redundant hot-standby supervisor.

The Cisco Catalyst 4500 ports are independent of the supervisor state. Because of this hardware design, during a supervisor switchover, the ports connected to the failed supervisor do not go down. Because paths and ports are not down, hardware keeps forwarding the packet to a valid next-hop while supervisor switchover is occurring.

The configuration and implementation guidelines for implementing NSF/SSO on the Cisco Catalyst 4500 are the same for district office and school site network designs.

Increasing Supervisor Uplink Port Availability

There are restrictions on which supervisor uplink ports can be actively configured. Multiple ports can be simultaneously active on the supervisor. However Cisco IOS Release 12.2(25)SG or later is required for concurrent use of both 10G and 1G. The Schools SRA uses the 1G interface to connect to the Cisco 3750-MetroE WAN aggregation switch. To use 10G port in 1G mode with redundancy, the following configuration must be applied on collapsed core Catalyst 4500 switch:

```
cr24-4507-DO(config)#hw-module uplink mode shared-backplane
cr24-4507-DO(config)#hw-module module 3 port-group 1 select gigabitethernet
cr24-4507-DO(config)#hw-module module 4 port-group 1 select gigabitethernet

cr24-4507-DO#show hw-module uplink
 Active uplink mode configuration is Shared-backplane
```

Stateful Switchover (SSO)

SSO redundancy mode in the Cisco Catalyst 4500 supervisor is turned on by default starting with Cisco IOS Release 12.2(20)EWA. To provide 1+1 redundancy, all the technical specifications between active and standby supervisor must be identical. Also note that the Cisco Catalyst 4507R and 4510R are the only models that support supervisor redundancy. SSO is supported on all supervisors running IOS except Sup II-Plus-TS. The NSF-awareness feature is supported by all the supervisors supporting EIGRP, OSPF, IS-IS, and BGP routing protocols, while the NSF-capable feature is supported only on supervisors IV, V, and V-10G. NSF/SSO on Catalyst 4500 requires a minimum boot ROM version and must be the same on both supervisors.

For additional details on hardware requirements, refer to the Release Notes at the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/53SG/configuration/NSFwSSO.html#wp1135767

```
cr24-4507-DO(config)#redundancy
cr24-4507-DO(config-red)# mode sso
cr24-4507-DO(config-red)# main-cpu
cr24-4507-DO(config-r-mc)# auto-sync standard

cr24-4507-DO#show module | inc Chassis | 6-E | SSO
Chassis Type : WS-C4507R-E

 3     6  Sup 6-E 10GE (X2), 1000BaseX (SFP)     WS-X45-SUP6-E
JAE1132SXQ3
 4     6  Sup 6-E 10GE (X2), 1000BaseX (SFP)     WS-X45-SUP6-E
JAE1132SXRQ

 3   Active Supervisor   SSOActive
 4   Standby Supervisor  SSOStandby hot
```

The active supervisor dynamically detects the secondary supervisor installed in the same chassis and initiates several SSO dependency configuration checks. If the SSO dependency check fails, then the standby supervisor falls back into RPR mode. For example, IOS release mismatch between two supervisors may not allow SSO to synchronize.

If the SSO dependency configuration checks successfully pass, then SSO communication between both supervisors goes through several synchronization states before it transitions to hot-standby state as illustrated in the following output:

```
cr24-4507-DO#show redundancy states
my state = 13 -ACTIVE
peer state = 8  -STANDBY HOT
. . .

Redundancy Mode (Operational) = Stateful Switchover
Redundancy Mode (Configured)  = Stateful Switchover
Redundancy State              = Stateful Switchover
    Maintenance Mode = Disabled
    Manual Swact = enabled
    Communications = Up
    . . .
```

All the state-machines and dynamic information of SSO-capable protocols are automatically synchronized to the standby supervisor module. The hot-standby supervisor takes over the ownership of control-plane process if the active supervisor suffers an outage or is removed from the chassis.

Non-Stop Forwarding (NSF)

The IOS based NSF implementation in Cisco Catalyst 4500 is identical to Cisco 3750-E StackWise-based architecture. If the active supervisor fails or is removed, the standby supervisor detects the outage and transitions into the active supervisor role to take over control-plane ownership, within milliseconds. The new supervisor immediately initializes graceful routing recovery with NSF-aware Layer-3 routers, while it continues to forward traffic based on pre-switchover hardware FIB information.

NSF-Capable Layer 3 Switch

```
cr24-4507-DO(config)#router eigrp 100
cr24-4507-DO (config-router)#nsf

cr24-4507-DO#show ip protocols | inc NSF
*** IP Routing is NSF aware ***
  EIGRP NSF-aware route hold timer is 240
  EIGRP NSF enabled
     NSF signal timer is 20s
     NSF converge timer is 120s
```

NSF Aware Layer 3 Switch

```
cr24-3560r-DO#show ip protocols | inc NSF
*** IP Routing is NSF aware ***
  EIGRP NSF-aware route hold timer is 240
```

WAN Design

In the Schools Service Ready Architecture, the school sites are connected to the district office over Wide Area Network (WAN) links. This section discusses how to design and deploy the WAN for Schools SRA. The primary components of WAN architecture are as follows:

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

The Schools SRA design uses Metro Ethernet service as the WAN transport between school sites, and the district office.

Service Deployed in the Design

Schools SRA design uses an E-line Metro Ethernet with point-to-point service to connect school sites to the district office. Each school site has a 100 Mbps Metro point-to-point connection to the service provider Network. As mentioned in the previous section, each circuit is represented by a VLAN using dot1q trunk. Figure 25 illustrates how this is implemented.

Figure 25 EVPN Service Used in School WAN Architecture

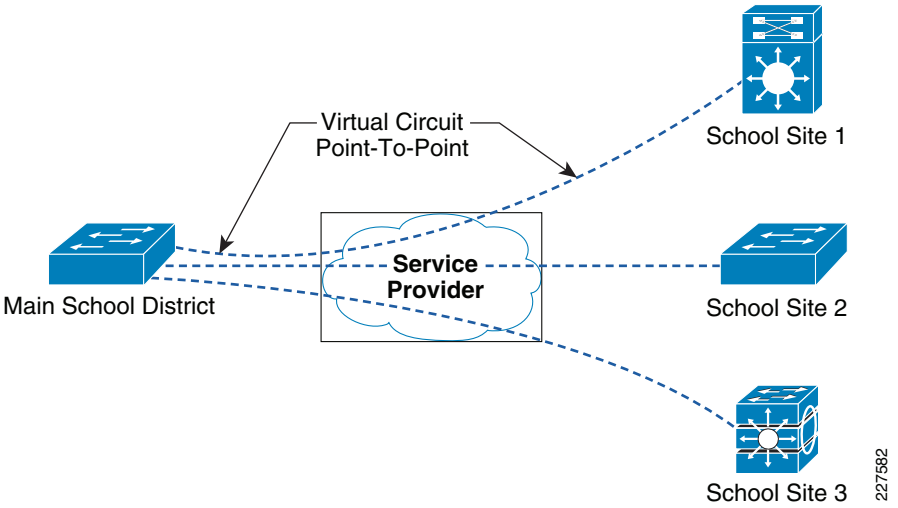

Following is a sample configuration of the WAN interface at the district office:

```
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 School-1to50-Parent-Policy-Map
 hold-queue 2000 in
 hold-queue 2000 out
```

In the above configuration, the link is carrying 50 VLANs, which are connected to 50 school sites.

Bandwidth Capacity Planning

Planning sufficient bandwidth capacity is a critical component of the overall WAN design. Application performance depends largely on guaranteed level of bandwidth at school sites and the district office. This section discusses the general WAN bandwidth capacity planning steps, and how this has been implemented in the School SRA.

The School district muct purchase the MetroE service from the local Service Provider. The amount of bandwidth capacity at each school, and at the district office must be sufficient to meet the anticipated network load, and some margin for peak usage, and to allow future growth. The bandwidth is shared based on the four-class QoS model, for optimal service delivery. Logical bandwidth assignment to each circuit must be symmetric at the school and at the district office. A mismatch in bandwidth capacity will force traffic drop in the SP core due to in-consistent bandwidth SLAs.

The WAN bandwidth capacity may vary between school sites, but the design principles and implementation (bandwidth sharing, routing, QoS, multicast) are the same.

Note The district office WAN router is an aggregator that logically connects to multiple schools over a single media. The School SRA district office WAN device is the Cisco 3750ME. The media type connecting to the Metro WAN MUST be GigE, since the 3750ME does not negotiate to lower speeds.

Calculating the optimum guaranteed bandwidth required at the district office is done by considering the following factors:

- Number of school sites.
- Bandwidth required at each location.
- Platform scalability limits, mainly, at district office.

To describe how much bandwidth is required at each school site, we use two terms, CIR, Interface speed.

- *CIR*—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 school sites and district office.

For example, let us consider a scenario where there are three school sites, and one district office. The CIR required at each school site is 100Mbps, and since there are three school sites, the district office needs three virtual circuits each having a CIR of 100Mbs, which also means that the aggregate bandwidth at district office is 300Mbps. Figure 26 illustrates this point.

Figure 26 Bandwidth Capacity Panning for Three School Sites

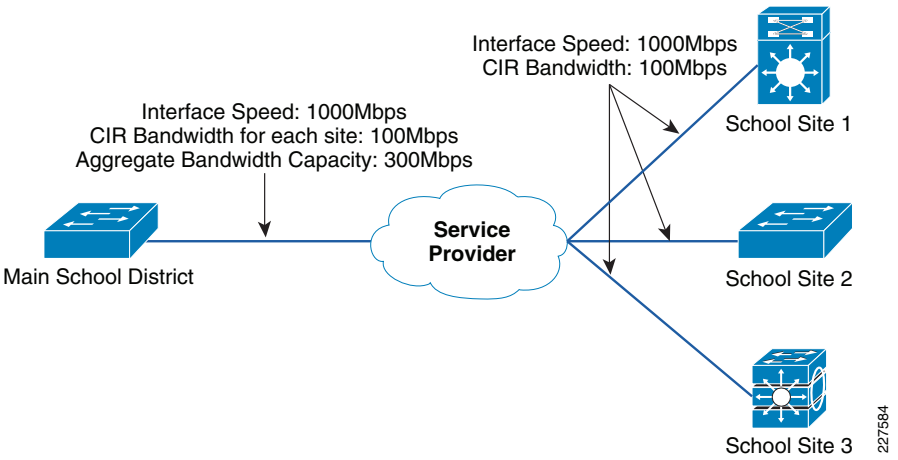

Similarly, if the CIR required at each school site is 100Mbps, and if there are 100 school sites, then the bandwidth required at district office is 10Gbps. To support an aggregate CIR bandwidth of 10Gbps at the district office, we need to think about scalability limits on the WAN aggregation box.

If the aggregated logical connection speed to the schools exceed the media capacity on 3750-ME, which is the WAN aggregator for our design, then there are two design options:

1. .Migrate to Modular a switching platform (such as Catalyst 6500) - Migrating to a modular switching platform in the WAN aggregation tier enables higher bandwidth capacityy,and may reduce operational and management complexities.
2. Deploy another 3750-ME -Deploying another 3750-ME is the simplest way to scale the WAN capacity. Deploying another set of 3750-ME does not change any WAN design principles, and except for VLAN and IP address, all the configurations can be replicated to the secondary system.

Implementation

This section describes how the WAN is implemented in the Schools SRA reference design, which was validated in the lab..

- The district office has dual WAN connections, and each connection is 1Gbps.
- On each WAN connection there are 50 Virtual circuits, each with CIR of 20Mbps.
- The school sites have 1Gbps connection, and CIR value of 20Mbs on each circuit.

Figure 27 shows how this design is validated with 100 school sites.

Figure 27 Bandwidth Capacity Planning for 100 School Sites

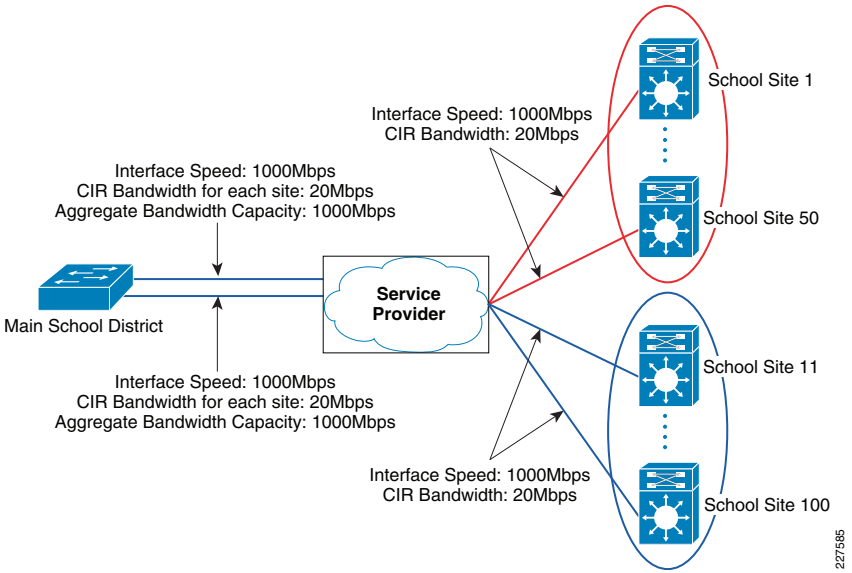

IP Address Aggregation

This section describes how to aggregate IP address space at the school sites, and the district office on the WAN interface.

As explained in the previous section, all the school sites are connected to the district office using point-to-point links, which means that every school site should be on a different IP subnet. It is common for customers to deploy using /30 subnet for each school site. This uses up to 4 addresses for each subnet. Best practice recommendation is to use /31 subnets. This approach only uses two addresses in each link. The following configurations show how to deploy this at a school site or district office:

School Site:

```
interface Vlan501
 description Connected to cr24-3750ME-DO
 dampening
 ip address 10.126.0.1 255.255.255.254
 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
```

District Office:

```
interface Vlan501
 description Connected to cr35-4507-SS1
 dampening
 ip address 10.126.0.0 255.255.255.254
 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
```

Routing for WAN Connections

This section discusses how to implement routing on the WAN interfaces. The key consideration when designing the routing protocol is summariztion.

Summarization on network boundaries is very important to 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 a school district 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 school district, 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:

School Site:

```
interface Vlan501
 description Connected to cr24-3750ME-DO
 dampening
 ip address 10.126.0.1 255.255.255.254
 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
```

District Site:

```
interface Vlan501
 description Connected to cr35-4507-SS1
 dampening
 ip address 10.126.0.0 255.255.255.254
 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
```

WAN QoS Design

QoS design is particularly important when using Ethernet as the WAN, since the router or switch on the WAN edge might believe they have complete line rate available to transmit. If the WAN device transmits at full line rate into the WAN, the Service Provider network will drop packets exceeding the CIR (ie the committed rate agreed to with Service Provider). Figure 28 shows what may happen without a proper QoS design.

Figure 28 Policing at Service provider due to lack of proper QoS at district office, and school sites

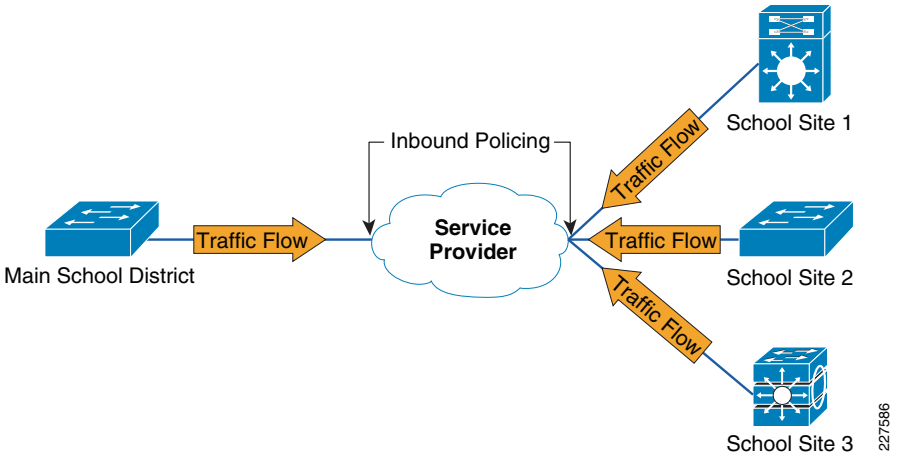

Proper QoS policies implemented at the district office and school site will prevent packets from being dropped at service provider network.

WAN QoS Policy at District Office

The district office has several point to point circuits; one connecting to each school site. The QoS policy at the district office has the following objectives:

- The aggregate traffic going out to the school site does not exceed the school site CIR (20Mbps for the lab testbed).
- All the traffic going out is put into four classes.

To accomplish the above objectives, Hierarchical Class Based Weighted Fair Queuing (HCBWFQ) is implemented. To learn more about HCBWFQ, refer to Ethernet Access for Next Generation Metro and Wide Area Networks Design Guide at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/Ethernet_Access_for_NG_MAN_WAN_V3.1_external.html

Implementing HCWFQ, requires two policies:

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

Figure 29 shows the representation of hierarchical policy.

Figure 29 Hierarchical policy implementation at District Ofice

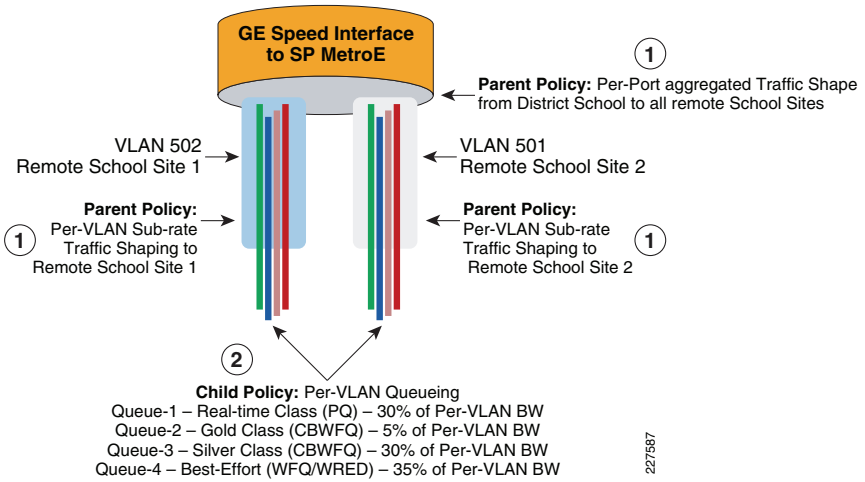

Table 3 shows how these classes are defined.

Table 3 Class Definitions

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

Following is the configuration of the QoS policy at the district office. This configuration is for one VLAN:

```
class-map match-all School_Site1   This class map would match the traffic going to a
school site
  description cr2-4507-SS1
 match vlan  501

policy-map School-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
!
```

Following is the configuration of the parent policy:, which shapes to 20Mbps for each school site in this example:

```
Policy Map School-1to50-Parent-Policy-Map   This is parent policy map
```

```
    Class School_Site1
      shape average 20000000 (bits/sec)
      service-policy School-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 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 School-1to50-Parent-Policy-Map   The policy-map
 hold-queue 2000 in
 hold-queue 2000 out
```

After completing the QoS policy at district office, we need to define the QoS policy at school sites.

WAN QoS Policy at School Site

The objectives at the school sites are similar to the one at the district office, which is:

- Ensure that 20Mbps is the maximum aggregate traffic leaving the WAN device,
- Ingress traffic is queued in four classes.

The school site implementation is different from the district office (due to lack of HCBWFQ support). The following configuration shows how to implement the Qos policy without HCBWFQ.

The first step is to queue the ingress traffic in the four queues. Table 4 shows the queues, and the bandwidth allocated for each.

Table 4 Queuing Ingress Traffic

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

Following is the configuration of egress interface on the school-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 traffic is queued as per the table shown above. However, the "bandwidth" would only ensure the minimum amount of bandwidth available. It does not control the upper threshold, which needs to be 20Mbps in our example. Therefore, to make sure that the traffic on the egress interface does not exceed 20Mps, we have a WAN-EGRESS-PARENT policy that polices the traffic to 20Mbps. Following is the 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#
```

Core Distribution Integration

The core/distribution component of the schools SRA is a key element in delivering a resilient network, while providing a network configuration that is easy to manage and to deploy. This chapter discusses both core/distribution models for the school sites, the Cisco Catalyst 3750 stack model and the Cisco Catalyst 4500 modular switch model. This chapter summarizes different connection types to the core/distribution models, and the key features of those connections.

Large School Design

The basic modular switch School topology is shown in Figure 30. This design is based upon a collapsed core/distribution mode, and a Layer-2 access distribution model. In this type of design all of the IP subnets are defined on the Catalyst 4500 modular switch, and access to these subnets is controlled by the VLANs that are trunked to the switches.

If desired a Layer-3 access switch model may be implemented, the physical topology does not change, and centrality of EtherChannel to the design does not change. The simplicity of the network design not only allows Layer-2 or Layer-3 access layers, it also allows a hybrid deployment. This allows the majority of clients on the switch use Layer-3 access features, but a group of legacy client are able to continue to use a Layer-2 network. This can be useful when migrating from clients that do not use IP, or rely heavily upon locally broadcast information to learn about services or devices on the network.

The Catalyst 4500 modular core provides resilient connections to the access LAN switches, local server switch, WLC, and SRST Router through EtherChannel. The NAC CAS appliance does not support EtherChannel, and the two connections are used to connect to the trusted and untrusted interfaces of the appliance.

The WAN connection to the Catalyst 4500 modular is a single connection

Figure 30 Stacked Switch School Schematic

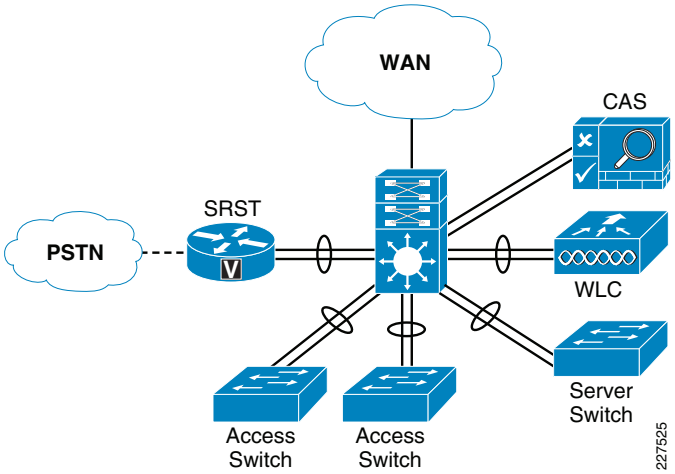

Core/Distribution Virtual Interfaces

The following is an example configuration of the switch virtual interfaces (SVIs) configured on the core/distribution Catalyst 4500modular switch. This SVIs are trunked to the access switches as required, and access to the VLANs are controlled by the switchport trunk allowed vlan command applied on the port channels. The same basic configuration is used for the server switch

```
.
interface Vlan101
description Connected to cr35_2960_Dept_1_VLAN
dampening
ip address 10.127.0.1 255.255.255.192
ip helper-address 10.125.31.2
no ip redirects
no ip unreachables
ip pim sparse-mode
load-interval 30
…
!
interface Vlan110
description Connected to cr35_2960_Dept_10_VLAN
dampening
ip address 10.127.2.65 255.255.255.192
ip helper-address 10.125.31.2
no ip redirects
no ip unreachables
ip pim sparse-mode
load-interval 30
```

Example Port Channel Configuration

The following are examples of the port channel configuration on core/distribution Catalyst 4500 modular switch and an example access switch. A similar configuration would be applied to each access switch connection with the same or different VLANs as required. From an IP routing or services level there is no requirement to span the same VLAN to multiple switches, but if there is a requirement to support legacy protocols such as AppleTalk at the school these AppleTalk VLANs can be easily spanned to different access switches as required

Example Catalyst 4500Modular Switch Port Channel Configuration

```
interface Port-channel11
description Connected to cr35-2960-SS1
switchport
switchport trunk encapsulation dot1q
switchport trunk native vlan 802
switchport trunk allowed vlan 101-110
switchport mode trunk
logging event link-status
load-interval 30
carrier-delay msec 0
qos trust dscp
```

Example 2960 Port Channel Configuration

```
interface Port-channel1
description Connected to cr35-4507-SS1
switchport trunk native vlan 802
switchport trunk allowed vlan 101-110,201
switchport mode trunk
ip arp inspection trust
load-interval 30
```

```
carrier-delay msec 0
hold-queue 2000 in
hold-queue 2000 out
ip dhcp snooping trust
```

WLC Connection

The WLC Connection to the core/distribution stack is fundamentally the same as an access switch connection, with different VLANs, and the exception of using a different QoS trust mode, where the CoS values from the WLC, are trusted. The following is an example Catalyst 4500modular switch Port Channel configuration:

```
Interface Port-channel12
description Connected to WLC-SS2
switchport trunk encapsulation dot1q
switchport trunk native vlan 802
switchport trunk allowed vlan 111-120
switchport mode trunk
load-interval 30
carrier-delay msec 0
ip dhcp snooping trust
```

NAC CAS Connection

The NAC CAS connection to the core/distribution switch. This is not an EtherChannel connection, but two switch ports are consumed. The two ports consist of a untrusted port for connecting client VLANs to the CAS prior to them completing the NAC process, and a trusted port that connects the NAS to the client VLANs used once clients have successfully completed the NAC process. The two, trusted and untrusted, ports are required even if OOB NAC is used, as the CAS requires access to the trusted VLANs during the NAC process. The following is an example of the configuration.

Core/Distribution NAC CAS Configuration

```
interface GigabitEthernet1/0/4
 description NAC Trusted Eth0
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 48,57,62
 switchport mode trunk
 spanning-tree portfast trunk
!
…!
interface GigabitEthernet1/0/8
 description NAC Untrusted Eth1
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 61,248,257
 switchport mode trunk
 spanning-tree portfast trunk
```

SRST Connection

Sample Configuration

The SRST connection to the core/distribution is another EtherChannel connection. The differences between the SRST connection and the access switch connections are that a trunk connection is not required, and that the SRST interfaces are router interfaces, requiring a slightly different connection. The following is an example of the configuration.

```
interface Port-channel3
 description to isr for simulated PSTN
GW for school2
 switchport access vlan 303
 switchport mode access

interface GigabitEthernet2/0/20
 switchport access vlan 303
 switchport mode access
 mls qos trust dscp
 channel-group 3 mode on
end

interface GigabitEthernet3/0/20
 switchport access vlan 303
 switchport mode access
 mls qos trust dscp
 channel-group 3 mode on
end
```

```
interface Port-channel3
 description port-channel to 4500
…
!
interface GigabitEthernet0/0
 description
$ETH-LAN$$ETH-SW-LAUNCH$$INTF-INFO-GE 0/0$
 no ip address
 duplex auto
 speed auto
 media-type rj45
 no keepalive
 channel-group 3
!
interface GigabitEthernet0/1
 no ip address
 duplex auto
 speed auto
 media-type rj45
 no keepalive
 channel-group 3
```

WAN Connection

The WAN connection is a single port connection from the core/distribution switch, and therefore there is no EtherChannel. The key component in the WAN connection configuration is the QoS implementation, that provides traffic shaping and limiting on this interface to ensure that the voice and video are given appropriate priority, but do not starve other applications of the throughput. An example of the WAN connection configuration is shown below.

WAN Port Sample Configuration-Core/Distribution

```
interface GigabitEthernet3/0/52
description Connected to MetroE-Core
switchport trunk encapsulation dot1q
switchport trunk native vlan 801
switchport trunk allowed vlan 650
switchport mode trunk
load-interval 30
carrier-delay msec 0
srr-queue bandwidth shape 35 15 25 25
srr-queue bandwidth limit 10
priority-queue out
mls qos trust dscp
no cdp enable
spanning-tree portfast trunk
spanning-tree bpdufilter enable
hold-queue 2000 in
hold-queue 2000 out
```

```
interface Vlan650
dampening
ip address 10.126.1.99 255.255.255.254
no ip redirects
no ip unreachables
ip pim sparse-mode
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 eigrp-key
ip summary-address eigrp 100 10.127.112.0 255.255.248.0 5
load-interval 30
hold-queue 2000 in
hold-queue 2000 out
```

Small School Design

The basic stacked switch school topology is shown in Figure 2.This design is based upon a collapsed core/distribution mode, and a Layer-2 access distribution model. In this type of design, all of the IP subnets are defined on the Catalyst 3750 stacked switch, and access to these subnets is controlled by the VLANs that are trunked to the switches.

The Catalyst 3750 stackwise core provides resilient connections to the access LAN switches, local server switch, WLC, and SRST router through EtherChannel. The NAC Appliance does not support EtherChannel, and the two connections are used to connect to the trusted and untrusted interfaces of the CAS.

The WAN connection to the Catalyst 3750 stack is a single Ethernet connection

Figure 31 Stacked Switch School Schematic

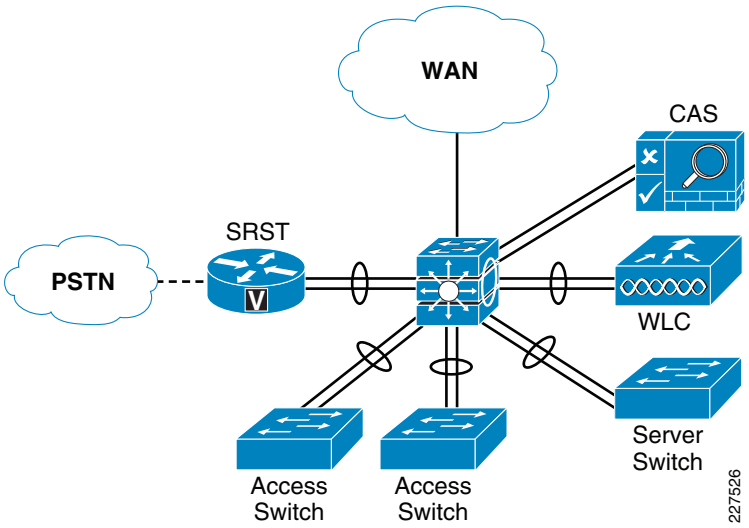

Below is an example configuration of the SVIs configured on the core/distribution Catalyst 3750 stack. These SVIs are trunked to the access switches as required, and access to the VLANs are controlled by the switchport trunk allowed vlan command applied on the port channels. The same basic configuration is used for the server switch.

Core/Distribution Virtual Interfaces

Example Port Channel Configuration

The following example shows an example of the port channel configuration on core/distribution Catalyst 3750 stack and an example access switch. A similar configuration would be applied to each access switch connection with the same or different VLANs as required. From an IP routing or services level there is no requirement to span the same VLAN to multiple switches, but if there is a requirement to support legacy protocols such as AppleTalk at the school these AppleTalk VLANs can be easily spanned to different access switches as required.

Example Catalyst 3750 Stack Port Channel Configuration	Example 2960 Port Channel Configuration
Interface Port-channel11 description Connected to 2960-SS2 switchport trunk encapsulation dot1q switchport trunk native vlan 802 switchport trunk allowed vlan 101-110,900 switchport mode trunk load-interval 30 carrier-delay msec 0 ip dhcp snooping trust	Interface Port-channel1 description Connected to 3750-Core-SS2 switchport trunk native vlan 802 switchport trunk allowed vlan 101-110 switchport mode trunk ip arp inspection trust load-interval 30 ip dhcp snooping trust
interface GigabitEthernet1/0/49 description Connected to 2960-SS2 switchport trunk encapsulation dot1q switchport trunk native vlan 802 switchport trunk allowed vlan 101-110,900 switchport mode trunk load-interval 30 carrier-delay msec 0 srr-queue bandwidth share 1 30 35 5 priority-queue out udld port channel-group 11 mode active spanning-tree guard root ip dhcp snooping trust ! interface GigabitEthernet3/0/49 description Connected to 2960-SS2 switchport trunk encapsulation dot1q switchport trunk native vlan 802 switchport trunk allowed vlan 101-110,900 switchport mode trunk load-interval 30 carrier-delay msec 0 srr-queue bandwidth share 1 30 35 5 priority-queue out udld port channel-group 11 mode active spanning-tree guard root ip dhcp snooping trust	interface GigabitEthernet0/1 description Connected to 3750-Core-SS2 switchport trunk native vlan 802 switchport trunk allowed vlan 101-110 switchport mode trunk ip arp inspection trust load-interval 30 srr-queue bandwidth share 1 30 35 5 priority-queue out udld port mls qos trust dscp channel-protocol lacp channel-group 1 mode active ip dhcp snooping trust ! interface GigabitEthernet0/2 description Connected to 3750-Core-SS2 switchport trunk native vlan 802 switchport trunk allowed vlan 101-110 switchport mode trunk ip arp inspection trust load-interval 30 srr-queue bandwidth share 1 30 35 5 priority-queue out udld port mls qos trust dscp channel-protocol lacp channel-group 1 mode active ip dhcp snooping trust

WLC Connection

The WLC connection to the core/distribution stack is fundamentally the same as an access switch connection, with different VLANs, and the exception of using a different QoS trust mode, where the CoS values from the WLC, are trusted. The following is an example of the configuration.

Example Catalyst 3750 Stack Port Channel Configuration

```
Interface Port-channel12
description Connected to 2960-SS2
switchport trunk encapsulation dot1q
switchport trunk native vlan 802
switchport trunk allowed vlan 111-120
switchport mode trunk
load-interval 30
carrier-delay msec 0
ip dhcp snooping trust

interface GigabitEthernet1/0/48
description Connected to WLC-SS2
switchport trunk encapsulation dot1q
switchport trunk native vlan 802
switchport trunk allowed vlan 110-120
switchport mode trunk
load-interval 30
carrier-delay msec 0
srr-queue bandwidth share 1 30 35 5
priority-queue out
udld port

mls qos trust coschannel-group 11 mode active
spanning-tree guard root
!
interface GigabitEthernet3/0/48
description Connected to WLC-SS2
switchport trunk encapsulation dot1q
switchport trunk native vlan 802
switchport trunk allowed vlan 110-110,
switchport mode trunk
load-interval 30
carrier-delay msec 0
srr-queue bandwidth share 1 30 35 5
priority-queue out
udld port
mls qos trust coschannel-group 11 mode active
spanning-tree guard root
```

NAC CAS Connection

The NAC CAS connection to the core/distribution switch. This is not an EtherChannel connection, but two switch ports are consumed. The two ports consist of a untrusted port for connecting client VLANs to the CAS prior to them completing the NAC process, and a trusted port that connects the NAS to the client VLANs used once clients have successfully completed the NAC process. The two, trusted and untrusted, ports are required even if OOB NAC is used, as the CAS requires access to the trusted VLANs during the NAC process. The following is an example of the configuration.

Core/Distribution NAC CAS Configuration

```
interface GigabitEthernet1/0/4
 description NAC Trusted Eth0
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 48,57,62
 switchport mode trunk
 spanning-tree portfast trunk
!
!
interface GigabitEthernet1/0/8
 description NAC Untrusted Eth1
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 61,248,257
 switchport mode trunk
 spanning-tree portfast trunk
```

SRST Connection

The SRST connection to the core/distribution is another EtherChannel connection. The differences between the SRST connection and the access switch connections are that a trunk connection is not required, and that the SRST interfaces are router interfaces, requiring a slightly different connection. The following is an example of the configuration.

```
interface Port-channel3
 description to isr for simulated PSTN GW
for school1
 switchport access vlan 303
 switchport mode access
 interface GigabitEthernet2/0/20
 switchport access vlan 303
 switchport mode access
 mls qos trust dscp
 channel-group 3 mode on
end

interface GigabitEthernet3/0/20
 switchport access vlan 303
 switchport mode access
 mls qos trust dscp
 channel-group 3 mode on
end
```

```
interface Port-channel3
 description port-channel to core stack
 ip address 10.40.63.9 255.255.255.252
 hold-queue 150 in
!
interface GigabitEthernet0/0
 description
$ETH-LAN$$ETH-SW-LAUNCH$$INTF-INFO-GE 0/0$
 no ip address
 duplex auto
 speed auto
 media-type rj45
 no keepalive
 channel-group 3
!
interface GigabitEthernet0/1
 no ip address
 duplex auto
 speed auto
 media-type rj45
 no keepalive
 channel-group 3
```

WAN Connection

The WAN connection is a single port connection from the core/distribution switch, and therefore there is no EtherChannel. The key component in the WAN connection configuration is the QoS implementation, that provides traffic shaping and limiting on this interface to ensure that the Voice and Video are given appropriate priority, but do not starve other applications of the throughput. An example of the WAN connection configuration is shown below.

```
interface GigabitEthernet3/0/52
description Connected to MetroE-Core
switchport trunk encapsulation dot1q
switchport trunk native vlan 801
switchport trunk allowed vlan 650
switchport mode trunk
load-interval 30
carrier-delay msec 0
srr-queue bandwidth shape 35 15 25 25
srr-queue bandwidth limit 10
priority-queue out
mls qos trust dscp
no cdp enable
spanning-tree portfast trunk
spanning-tree bpdufilter enable
hold-queue 2000 in
hold-queue 2000 out

interface Vlan650
dampening
ip address 10.126.1.99 255.255.255.254
no ip redirects
no ip unreachables
ip pim sparse-mode
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 eigrp-key
ip summary-address eigrp 100 10.127.112.0 255.255.248.0 5
load-interval 30
hold-queue 2000 in
hold-queue 2000 out
```

District Office Design

The school district office network provides the infrastructure that both severs the users local to the school headquarters and that supports the centralized services used across all the school sites. This chapter discusses the key aspects of the network design of the school district office, including core/distribution models, Internet connectivity and services block.

Compared to the previously discussed school site design, there are four main differences in the district office design:

- The use of the Supervisor 6—The Supervisor 6 supports hierarchical QoS.
- The Metro Ethernet Switch Connection—The aggregation and QoS policy enforcement point for the Metro Ethernet WAN connection to the schools
- The Services Block Switch Connection—The district office "mini-Data Center" for the management and services servers for the district and the schools
- The ASA Firewall Connection-The firewall connection to the Internet

District Office Partial Schematic shows a schematic of the district office network. Aside from providing core/distribution services to the access switches, the Cisco Catalyst 4500 modular switch in the district office connects the school WAN to the district office, the services such as Internet access and the Services Block of the SRA.

Figure 32 District Office Partial Schematic

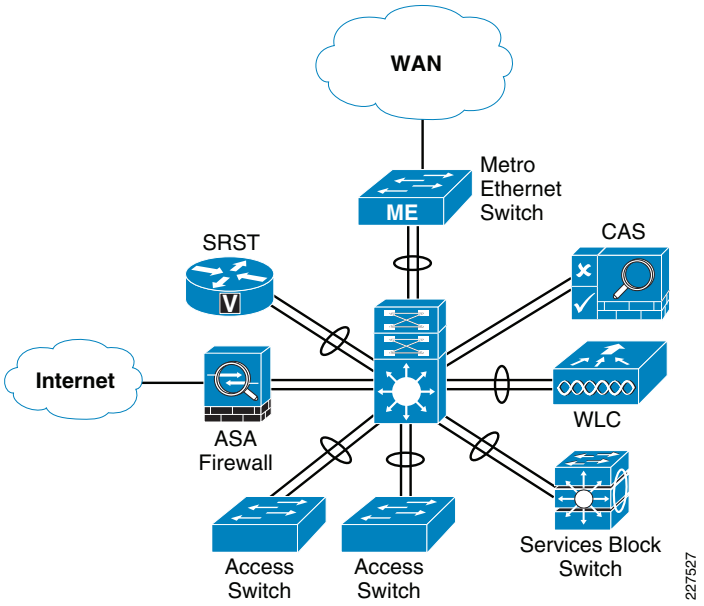

Metro Ethernet Connection Configuration

Port-Channel Configuration on the Core/Distribution Catalyst 4500 Modular Switch and the Catalyst 3750Metro Ethernet Sw shows an example of the port-channel configuration on the core/distribution Catalyst 4500 modular switch and the Catalyst 3750 Metro Ethernet switch. This is a Layer-3 connection where both the core/distribution switch and the Metro Ethernet switch are part of the same EIGRP AS. The most significant difference

in this configuration from the School design using the Catalyst 4500 modular switch is the QoS configuration on the Catalyst 4500 interface; this is primarily due to the district office using hierarchical QoS features of a Supervisor 6 module, rather than the Supervisor 5 used in the School SRA.

Table 5 Port-Channel Configuration on the Core/Distribution Catalyst 4500 Modular Switch and the Catalyst 3750Metro Ethernet Switch

Example Catalyst 4500 Modular Switch Configuration	Example Catalyst 3750ME Switch Configuraiton
interface Port-channel1 description Connected to 3750ME-DO dampening ip address 10.125.32.4 255.255.255.254 ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp-key ip pim dr-priority 100 ip pim sparse-mode ip summary-address eigrp 100 10.125.0.0 255.255.0.0 5 logging event link-status load-interval 30 carrier-delay msec 0 service-policy output PQ-POLICER	interface Port-channel1 description Connected to 4507-DO no switchport dampening ip address 10.125.32.5 255.255.255.254 ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp-key ip pim sparse-mode ip summary-address eigrp 100 10.127.0.0 255.255.0.0 5 ip summary-address eigrp 100 10.126.0.0 255.255.0.0 5 logging event bundle-status load-interval 30 carrier-delay msec 0 hold-queue 2000 in hold-queue 2000 out

Table 5 Port-Channel Configuration on the Core/Distribution Catalyst 4500 Modular Switch and the Catalyst 3750Metro Ethernet Switch (continued)

Example Catalyst 4500 Modular Switch Configuration	Example Catalyst 3750ME Switch Configuraiton
interface GigabitEthernet3/3 no switchport no ip address load-interval 30 carrier-delay msec 0 udld port channel-protocol pagp channel-group 1 mode desirable service-policy output EGRESS-POLICY ! interface GigabitEthernet4/3 no switchport no ip address load-interval 30 carrier-delay msec 0 udld port channel-protocol pagp channel-group 1 mode desirable service-policy output EGRESS-POLICY	interface GigabitEthernet1/0/1 description Connected to cr24-4507-DO no switchport no ip address logging event bundle-status load-interval 30 carrier-delay msec 0 srr-queue bandwidth share 1 30 35 5 priority-queue out udld port mls qos trust dscp channel-protocol pagp channel-group 1 mode desirable ! interface GigabitEthernet1/0/2 description Connected to cr24-4507-DO no switchport no ip address logging event bundle-status load-interval 30 carrier-delay msec 0 srr-queue bandwidth share 1 30 35 5 priority-queue out udld port mls qos trust dscp channel-protocol pagp channel-group 1 mode desirable
policy-map PQ-POLICER class PRIORITY-QUEUE police cir 300000000 conform-action transmit exceed-action drop policy-map EGRESS-POLICY class PRIORITY-QUEUE priority class CONTROL-MGMT-QUEUE bandwidth remaining percent 10 class MULTIMEDIA-CONFERENCING-QUEUE bandwidth remaining percent 10 class MULTIMEDIA-STREAMING-QUEUE bandwidth remaining percent 10 class TRANSACTIONAL-DATA-QUEUE bandwidth remaining percent 10 dbl class BULK-DATA-QUEUE bandwidth remaining percent 4 dbl class SCAVENGER-QUEUE bandwidth remaining percent 1 class class-default bandwidth remaining percent 25 dbl	

ASA Connection

The ASA firewall connection to the Catalyst 4500 modular switch is fundamentally different from the other network device connections to this switch-it uses the redundant interface feature of the ASA. The ASA redundant interface is a logical interface that pairs two physical interfaces, called active and standby interfaces. Under normal operation, the active interface is the only one passing traffic. The active interface uses the IP address defined at the redundant interface, and the MAC address of the first physical interface associated with the redundant interface. When the active interface fails, the standby interface becomes active and starts passing traffic. The same IP address and MAC address are maintained so that traffic is not disrupted. See Table 6.

Table 6 ASA Connection Configuration

Example Catalyst 4500 Modular Switch Configuration	Example ASA Interface Configuration
interface GigabitEthernet4/4 description backup link to cr26-asa5520-DO switchport access vlan 200 switchport mode access switchport block unicast load-interval 30 spanning-tree portfast spanning-tree bpduguard enable service-policy output EGRESS-POLICY ! interface GigabitEthernet5/3 description Connected to cr26-asa5520-DO switchport access vlan 200 switchport mode access switchport block unicast load-interval 30 media-type rj45 spanning-tree portfast spanning-tree bpduguard enable service-policy output EGRESS-POLICY !	interface GigabitEthernet0/0 description Connected to cr24-4507-DO no nameif no security-level no ip address ! interface GigabitEthernet0/1 description backup to cr24-4507-DO no nameif no security-level no ip address ! ! Defines logical redundant interface associated with physical interfaces. Configures IP and logical interface parameters. interface Redundant1 description Connected to cr24-4507-DO member-interface GigabitEthernet0/0 member-interface GigabitEthernet0/1 nameif inside security-level 100 ip address 10.125.33.10 255.255.255.0 authentication key eigrp 100 <removed> key-id 1 authentication mode eigrp 100 md5
interface Vlan200 description cr24_4507_FW_Inside ip address 10.125.33.9 255.255.255.0 ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp-key ip pim sparse-mode ip summary-address eigrp 100 10.125.0.0 255.255.0.0 5 logging event link-status load-interval 30 carrier-delay msec 0	

Services Block Connection

The Services Block supports the centralized servers and services for the district. The Cisco Catalyst 4500 modular switch connection to the Services Block switch uses EtherChannel, but in this case the connection between the switches is a Layer-3 connection, allowing the services block switch to keep its VLANs from stack is fundamentally the same as an access switch connection, with different VLANs.

Table 7 provides sample configurations for the Cisco 4500 modular switch and the Services Block switch.

Table 7 Service Block Configuration

Example 4500 Modular Switch Configuration	Example Services Block Switch Configuration
interface Port-channel17 description Connected to cr26-3750DC-DO switchport switchport trunk native vlan 806 switchport trunk allowed vlan 141-150,900 switchport mode trunk logging event link-status load-interval 30 carrier-delay msec 0 service-policy output PQ-POLICER	interface Port-channel1 description Connected to cr24-4507-DO switchport trunk encapsulation dot1q switchport trunk native vlan 806 switchport trunk allowed vlan 141-150,900 switchport mode trunk logging event bundle-status load-interval 30 carrier-delay msec 0 hold-queue 2000 in hold-queue 2000 out
interface GigabitEthernet1/1 description Connected to cr24-2960-DO switchport trunk native vlan 802 switchport trunk allowed vlan 101-110,900 switchport mode trunk logging event link-status load-interval 30 carrier-delay msec 0 udld port channel-protocol pagp channel-group 11 mode desirable spanning-tree guard root service-policy output EGRESS-POLICY ! interface GigabitEthernet2/1 description Connected to cr24-2960-DO switchport trunk native vlan 802 switchport trunk allowed vlan 101-110,900 switchport mode trunk logging event link-status load-interval 30 carrier-delay msec 0 udld port channel-protocol pagp channel-group 11 mode desirable spanning-tree guard root service-policy output EGRESS-POLICY	interface GigabitEthernet0/1 description Connected to cr24-4507-DO switchport trunk native vlan 802 switchport trunk allowed vlan 101-110,201,900 switchport mode trunk ip arp inspection trust load-interval 30 srr-queue bandwidth share 1 30 35 5 priority-queue out udld port mls qos trust dscp channel-protocol pagp channel-group 1 mode desirable hold-queue 2000 in hold-queue 2000 out ip dhcp snooping trust ! interface GigabitEthernet0/2 description Connected to cr24-4507-DO switchport trunk native vlan 802 switchport trunk allowed vlan 101-110,201,900 switchport mode trunk ip arp inspection trust load-interval 30 srr-queue bandwidth share 1 30 35 5 priority-queue out udld port mls qos trust dscp channel-protocol pagp channel-group 1 mode desirable hold-queue 2000 in hold-queue 2000 out ip dhcp snooping trust

Core/Distribution Virtual Interfaces

The following is an example configuration of the Switch Virtual Interfaces configured on the core/distribution Catalyst 4500 modular switch. This SVIs are trunked to the access switches as required, and access to the VLANs are controlled by the switchport trunk allowed vlan command applied on the port channels. The same basic configuration is used for the Server Switch

```
interface Vlan101
description Connected to cr24_2960_Dept_1_VLAN
dampening
ip address 10.125.1.1 255.255.255.128
ip helper-address 10.125.31.2
no ip redirects
no ip unreachables
ip pim sparse-mode
load-interval 30
!
...
interface Vlan110
description Connected to cr24_2960_Dept_10_VLAN
dampening
ip address 10.125.5.129 255.255.255.128
ip helper-address 10.125.31.2
no ip redirects
no ip unreachables
ip pim sparse-mode
load-interval 30
```

Table 8 provides examples of the port-channel configuration on core/distribution Catalyst 4500 modular switch and an access switch. A similar configuration would be applied to each access switch connection with the same or different VLANs as required. From an IP routing or services level there is no requirement to span the same VLAN to multiple switches, but if there is a requirement to support legacy protocols such as AppleTalk at the school these AppleTalk VLANs can be easily spanned to different access switches as required.

Table 8 Core/Distribution Virtual Interfaces

Example 4500 Modular switch Port Channel Configuration	Example 2960 Port Channel Configuration
Example 4500 Modular switch Port Channel Configuration	**Example 2960 Port Channel Configuration**
interface Port-channel11 description Connected to cr24-2960-DO switchport switchport trunk native vlan 802 switchport trunk allowed vlan 101-110,900 switchport mode trunk logging event link-status load-interval 30 carrier-delay msec 0 service-policy output PQ-POLICER	interface Port-channel1 description Connected to cr24-4507-DO switchport trunk native vlan 802 switchport trunk allowed vlan 101-110,201,900 switchport mode trunk ip arp inspection trust load-interval 30 carrier-delay msec 0 hold-queue 2000 in hold-queue 2000 out ip dhcp snooping trust
interface GigabitEthernet1/1 description Connected to cr24-2960-DO switchport trunk native vlan 802 switchport trunk allowed vlan 101-110,900 switchport mode trunk logging event link-status load-interval 30 carrier-delay msec 0 udld port channel-protocol pagp channel-group 11 mode desirable spanning-tree guard root service-policy output EGRESS-POLICY interface GigabitEthernet2/1 description Connected to cr24-2960-DO switchport trunk native vlan 802 switchport trunk allowed vlan 101-110,900 switchport mode trunk logging event link-status load-interval 30 carrier-delay msec 0 udld port channel-protocol pagp channel-group 11 mode desirable spanning-tree guard root service-policy output EGRESS-POLICY	interface GigabitEthernet0/1 description Connected to cr24-4507-DO switchport trunk native vlan 802 switchport trunk allowed vlan 101-110,201,900 switchport mode trunk ip arp inspection trust load-interval 30 srr-queue bandwidth share 1 30 35 5 priority-queue out udld port mls qos trust dscp channel-protocol pagp channel-group 1 mode desirable hold-queue 2000 in hold-queue 2000 out ip dhcp snooping trust ! interface GigabitEthernet0/2 description Connected to cr24-4507-DO switchport trunk native vlan 802 switchport trunk allowed vlan 101-110,201,900 switchport mode trunk ip arp inspection trust load-interval 30 srr-queue bandwidth share 1 30 35 5 priority-queue out udld port mls qos trust dscp channel-protocol pagp channel-group 1 mode desirable hold-queue 2000 in hold-queue 2000 out ip dhcp snooping trust

WLC Connection

The WLC Connection to the Core/Distribution Stack is fundamentally the same as an Access Switch connection, with different VLANs, and the exception of using a different QoS trust mode, where the CoS values from the WLC, are trusted. The following is an example of the configuration.

```
Interface Port-channel12
description Connected to WLC-SS2
switchport trunk encapsulation dot1q
switchport trunk native vlan 802
switchport trunk allowed vlan 111-120
switchport mode trunk
load-interval 30
carrier-delay msec 0
ip dhcp snooping trust

interface GigabitEthernet1/0/48
description Connected to WLC-SS2
switchport trunk encapsulation dot1q
switchport trunk native vlan 802
switchport trunk allowed vlan 110-120
switchport mode trunk
load-interval 30
carrier-delay msec 0
srr-queue bandwidth share 1 30 35 5
priority-queue out
udld port
mls qos trust coschannel-group 11 mode active
spanning-tree guard root
!
interface GigabitEthernet3/0/48
description Connected to WLC-SS2
switchport trunk encapsulation dot1q
switchport trunk native vlan 802
switchport trunk allowed vlan 110-110,
switchport mode trunk
load-interval 30
carrier-delay msec 0
srr-queue bandwidth share 1 30 35 5
priority-queue out
udld port
mls qos trust coschannel-group 11 mode active
spanning-tree guard root
```

NAC CAS Connection

The NAC CAS connection to the core/distribution switch. This is not an EtherChannel connection, but two switch ports are consumed. The two ports consist of a untrusted port for connecting client VLANs to the CAS prior to them completing the NAC process, and a trusted port that connects the NAS to the client VLANs used once clients have successfully completed the NAC process. The two, trusted and untrusted, ports are required even if OOB NAC is used, as the CAS requires access to the trusted VLANs during the NAC process. The following is an example of the configuration

```
.
interface GigabitEthernet 3/9
 description NAC Trusted Eth0
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 48,57,62
 switchport mode trunk
 spanning-tree portfast trunk
!
interface GigabitEthernet 4/9
 description NAC Untrusted Eth1
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 61,248,257
 switchport mode trunk
 spanning-tree portfast trunk
```

SRST Connection

The SRST connection to the core/distribution is another EtherChannel connection. The differences between the SRST connection and the access switch connections are that a trunk connection is not required, and that the SRST interfaces are router interfaces, requiring a slightly different connection. Table 9 provides an example of the configuration.

Table 9 SRST Connection

Core/Distribution	ISR Routers
interface Port-channel1 description Connected to ISR …! interface GigabitEthernet3/10 no switchport no ip address load-interval 30 carrier-delay msec 0 udld port channel-protocol pagp channel-group 1 mode desirable service-policy output EGRESS-POLICY ! interface GigabitEthernet4/10 no switchport no ip address load-interval 30 carrier-delay msec 0 udld port channel-protocol pagp channel-group 1 mode desirable service-policy output EGRESS-POLICY	interface Port-channel3 description port-channel to 4500 … ! interface GigabitEthernet0/0 description $ETH-LAN$$ETH-SW-LAUNCH$$INTF-INFO-GE 0/0$ no ip address duplex auto speed auto media-type rj45 no keepalive channel-group 3 ! interface GigabitEthernet0/1 no ip address duplex auto speed auto media-type rj45 no keepalive channel-group 3

NTP

The use of Network Time Protocol (NTP) to synchronize the clocks of network devices is a well established best practice, is fundamental for the analysis of logs/events and security, but might not warrant a mention in a design guide that is focused upon introducing new designs and practices to support new services in the network.

Given that a number of key components (for example, CUWN and Cisco NAC) of the Schools SRA rely upon or benefit from time synchronization, it was decided to include a short discussion on Time Synchronization as part of the Schools SRA.

The preferred mechanism for time synchronization in the network is NTP (other systems may use their own time synchronization mechanism) and this network NTP discussion is not proposed as a the design to synchronize all devices (hosts) in the network, its goal is synchronization of the network components of the Schools SRA. At the same time, whatever alternative times synchronization systems used in other parts of the network need to have agreement on the time, and should have a common time source at the beginning of their timing hierarchy. This will allow sufficient synchronization between hosts and network devices for the solutions deployed in the SRA.

The Schools SRA network has a hierarchy based upon the district office, as hub, and the schools as spokes. The NTP hierarchy should be the same, with the highest stratum NTP server located at the district office serving as the time reference for the district network. In order to spread the load, a hierarchy of NTP servers is used. The district office NTP server acting as the server for the district office network devices, and for the NTP server at each school, and the NTP Server for each school will act as the NTP server for network devices in that school.

The preferred network device to act as the NTP server in the schools SRA is the ISR router at each site. The ISR is the preferred device as routers have a greater CPU capacity than switches used in the Schools SRA due to many of the general purpose task that a router may be required it perform in CPU, compared to switches that have been more optimized to perform their more limited number of tasks in ASIC.

Figure 33 shows a schematic of the NTP hierarchy in the school district.

For more information about NTP, refer to the N*etwork Time Protocol: Best Practices White Paper* at the following URL:

http://www.cisco.com/en/US/tech/tk869/tk769/technologies_white_paper09186a0080117070.shtml

Figure 33 NTP School District Hierarchy

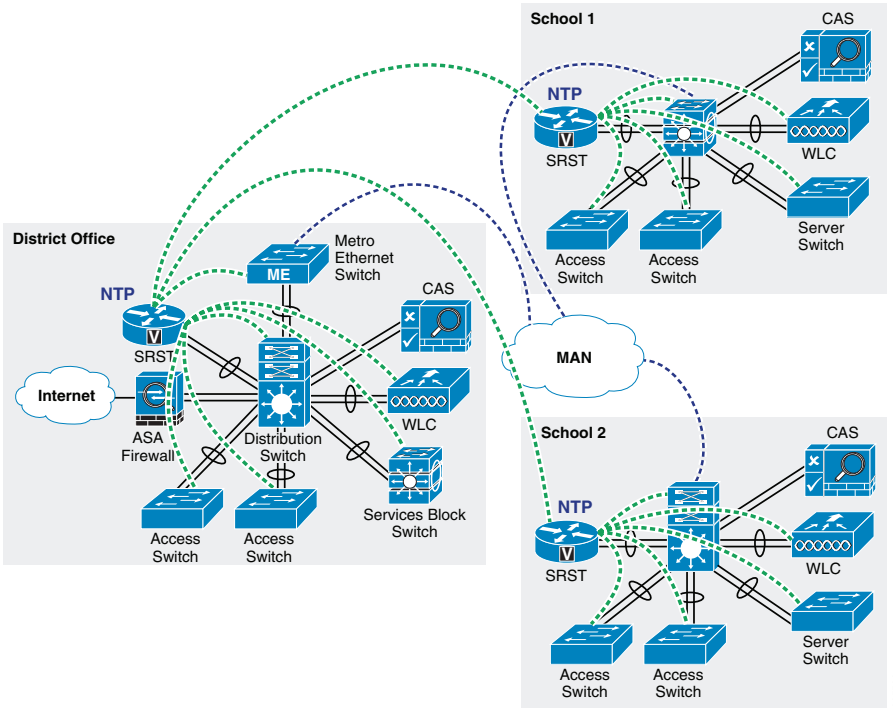

When creating the NTP configuration care should be taken to protect the NTP system.

The NTP associations should be limited, and controlled by an access list, to protect against DoS attacks. The NTP system should also use NTP authentication, where possible, to protected against spoofing attacks.

DO-ISR NTP	School1-ISR
access-list 99 permit x.x.x.x 0.0.0.255 access-list 99 permit y.y.y.y 0.0.0.255 ntp authentication-key 2 md5 Riewoldt ntp authenticate ntp source Port-channel3 ntp max-associations 150 ntp server a.a.a.a ntp access-group serve-only *99*	access-list 98 permit *z.z.z.z* 0.0.0.255 ntp authentication-key 2 md5 Riewoldt ntp trusted-key 2 ntp clock-period 17179685 ntp max-associations 150 ntp server *<DO-ISR>* key 2 ntp access-group serve-only *98*