Campus Wired LAN

Technology Design Guide

August 2014 Series

Table of Contents

Preface

Cisco Validated Designs (CVDs) present systems that are based on common use cases or engineering priorities. CVDs incorporate a broad set of technologies, features, and applications that address customer needs. Cisco engineers have comprehensively tested and documented each design in order to ensure faster, more reliable, and fully predictable deployment.

CVDs include two guide types that provide tested design details:

- **Technology design guides** provide deployment details, information about validated products and software, and best practices for specific types of technology.

- **Solution design guides** integrate existing CVDs but also include product features and functionality across Cisco products and sometimes include information about third-party integration.

Both CVD types provide a tested starting point for Cisco partners or customers to begin designing and deploying systems.

CVD Foundation Series

This CVD Foundation guide is a part of the *August 2014 Series*. As Cisco develops a CVD Foundation series, the guides themselves are tested together, in the same network lab. This approach assures that the guides in a series are fully compatible with one another. Each series describes a lab-validated, complete system.

The CVD Foundation series incorporates wired and wireless LAN, WAN, data center, security, and network management technologies. Using the CVD Foundation simplifies system integration, allowing you to select solutions that solve an organization's problems—without worrying about the technical complexity.

To ensure the compatibility of designs in the CVD Foundation, you should use guides that belong to the same release. For the most recent CVD Foundation guides, please visit the CVD Foundation web site.

Comments and Questions

If you would like to comment on a guide or ask questions, please use the feedback form.

CVD Navigator

The CVD Navigator helps you determine the applicability of this guide by summarizing its key elements: the use cases, the scope or breadth of the technology covered, the proficiency or experience recommended, and CVDs related to this guide. This section is a quick reference only. For more details, see the Introduction.

Use Cases

This guide addresses the following technology use cases:

- **Connecting Wired Devices to an Organization's Network**—Wired devices use Ethernet for providing or accessing services and communication at the workspaces and meeting places in an organization's remote sites and headquarters. Deployed with efficiency and consistency on LANs, the connectivity provides security, reliability, and manageability.

- **LAN and Services Interconnection to Scale within a Site**—At a larger site with increasing numbers of devices, a highly available, hierarchical network interconnects an organization's devices and services, for scale and growth. This network aids manageability, operational efficiency, and resiliency, while minimizing complexity.

- **Enhancing LAN Capacity and Functionality**—As the needs of an organization change, LAN capacity and functionality must be able to be refreshed to accommodate new requirements. Design modularity and software flexibility enhance an organization's efficiency to easily adapt to and accommodate updated network requirements.

For more information, see the "Use Cases" section in this guide.

Scope

This guide covers the following areas of technology and products:

- Ethernet wired access and device interconnection using Cisco Catalyst switches

- Hierarchical local area network design model, including access, distribution, and core layers, with simplified design options using Virtual Switching System (VSS)

- Advanced technology support for voice and video, including quality of service (QoS) marking and treatment

- Security, including management authentication, Catalyst Infrastructure Security Features (CISF), and IPv6 First Hop Security.

- Unicast routing, using Enhanced Interior Gateway Routing Protocol (EIGRP) or Open Shortest Path First (OSPF), and multicast routing using Protocol Independent Multicast (PIM) sparse mode

For more information, see the "Design Overview" section in this guide.

Related CVD Guides

Campus Wireless LAN Technology Design Guide

Device Management Using ACS Technology Design Guide

To view the related CVD guides, click the titles or visit the CVD Foundation web site.

Proficiency

This guide is for people with the following technical proficiencies–or equivalent experience:

· **CCNA Routing and Switching**–1 to 3 years installing, configuring, and maintaining routed and switched networks

Introduction

The *Campus Wired LAN Technology Design Guide* describes how to design a wired network access with ubiquitous capabilities that scale from small environments (for instance, those environments with one to just a few LAN switches) to a large, campus-size LAN. Resiliency, security, and scalability are included to provide a robust communications environment. Quality of Service (QoS) is integrated to ensure the base architecture can support a multitude of applications including low latency, drop-sensitive multimedia applications, that coexist with data applications on a single network.

The campus LAN architecture is designed to meet the needs of organizations with wired LAN connectivity requirements that range from a small, remote-site LAN to a large, multi-building location. The purpose of a campus network is to support arbitrary device connectivity for workers and users in the office and business spaces or meeting places of a building, such as for laptops, telephones, printers, and video conferencing systems. This is in contrast to the highly controlled connectivity for servers in a data center or machine and device connectivity in an industrial network or a WAN.

Many organizations have campus LAN requirements that include both wired and wireless access. The *Campus Wired LAN Technology Design Guide* offers guidance designed, deployed, and tested in conjunction with wireless guidance covered in the Campus Wireless LAN Technology Design Guide. Separation of the guides allows more concise coverage of each design. Depending on the needs of the organization this provides flexibility to use a single guide or multiple guides together as a set.

Technology Use Cases

This guide addresses the requirements of organizations when designing Local Area Networks (LANs) for their data communications needs. The guidance offered is useful for greenfield designs, for optimizing existing networks, and as a reference design offering operational consistency for an organization as its LAN grows. The scope of coverage applies to small, remote-site LANs with a single router up to large multi-building campuses with a routed core supporting connectivity to multiple-routed distribution modules.

This guide addresses four primary wired LAN requirements shared by organizations, including the need to:

- Offer reliable access to organization resources
- Minimize time required to absorb technology investments
- Provide a productive and consistent user experience
- Reduce operation costs

Use Case: Connecting Wired Devices to an Organization's Network

Organizations of all sizes have a need to connect data devices used by their employees such as desktop computers, laptops, and IP phones enabling communications with resources such as printers, business applications systems, voice and video endpoints, and conference bridges, along with Internet accesses, for interaction with partners and customers. Ethernet is the ubiquitous wired technology to make these communication connections. Using this guide, a LAN design of a few Ethernet interconnected devices can scale up to many thousands of devices in a multi-building campus over time.

This design guide enables the following network capabilities when connecting wired devices to an organization's network:

- **Consistent end user and network administrator experience**—Uses consistent design methodology in order to allow small remote sites with just a few Ethernet connections to be able to use the same access switch configurations as large campus Ethernet designs

- **Network security**—Protects the network and users from malicious attacks by applying security using Catalyst Infrastructure Security Features (CISF) and secure communication to devices, and integrating external authentication, authorization, and accounting (AAA) services

- **Protection of multimedia and critical applications traffic**—Enables critical applications and rich media communications, such as streaming and interactive voice and video media, through the use of end-to-end quality of service (QoS) enforcement, marking, and transmission policies—ensures appropriate network treatment of all types of business communications and deprioritization of background and non-business entertainment traffic

- **Rapid deployment**—Offers a choice of platforms with a range of power over Ethernet (PoE) support for deployment of media endpoints, such as phones and cameras, aided by in-line power technology

- **Manageability**—Allows the ability for network components to be managed from a central management network

- **Reliable connectivity**—Uses a Layer 2 LAN access design with resilient components and links for loop-free connections in order to ensure communications remain dependable, without wasted resources, such as unused links caused by spanning tree port blocking

Use Case: LAN and Services Interconnection to Scale within a Physical Site

As an organization grows, the network must grow to accommodate the increased number of devices connecting to the network, as well as offer connectivity to additional services components of increased size.

This design guide enables the following network capabilities supporting LAN and services interconnection within a physical site:

- **Reduced design complexity**—Uses replicable LAN access building blocks for Ethernet connectivity, network modularity concepts, and network hierarchy in order to allow network design to be assembled in a consistent approach to the scale that is dictated by organization growth.

- **Connectivity to IP services**—Uses resilient connectivity to a Layer 3 campus distribution or site router.

- **Ability to scale to large topologies**—Includes a design option of a resilient routed core, using a single pair of core devices, based on Virtual Switching System (VSS) technology.

- **High availability**—Offers resilient platform options and use of resilient connectivity configurations, allowing for maintenance of components without disruption of network services and mitigating single link failures from disrupting business communication.

- **Operational efficiency**—Uses consistent configurations across all areas of the network, increasing speed to deployment and reducing risk of configuration mistakes.

Use Case: Enhancing LAN Capacity and Functionality

As the needs of an organization change, the network should be able to be refreshed easily to adapt and support the new requirements for LAN capacity and functionality delivered.

This design guide enables the following network capabilities that support enhancing LAN capacity and functionality:

- **High design modularity**—Uses network modularity and hierarchy in order to easily introduce network components along with component options that support alternative functionality and new connectivity methods as requirements change.
- **Software flexibility**—Offers resilient platform software upgrade options and feature set licensing to minimize disruption of business communication while introduce new features to support an organization.
- **Operational efficiency**—Allows for bandwidth and capacity refresh as needed by an organization, in a consistent way that is not a burden to network administrators.

Design Overview

The LAN is the networking infrastructure that provides access to network communication services and resources for end users and devices spread over a single floor or building. A campus network is created by interconnecting a group of LANs that are spread over a small geographic area. Campus network design concepts are inclusive small networks that use a single LAN switch up to very large networks with thousands of connections.

This guide provides a design that enables communications between devices in a building or group of buildings, as well as interconnection to the WAN and Internet edge modules at the network core.

Specifically, this document shows you how to design the network foundation and services in order to enable:

- Tiered LAN connectivity
- Wired network access for employees
- IP Multicast for efficient data distribution
- Wired infrastructure ready for multimedia services

Hierarchical Design Model

This architecture uses a hierarchical design model to divide the design into modular groups or layers. Breaking up the design into layers allows each layer to implement specific functions. This simplifies the network design and therefore the deployment and management of the network.

Modularity in network design allows you to create design elements that can be replicated throughout the network. Replication provides an easy way to scale the network as well as a consistent deployment method.

In flat or meshed network architectures, changes tend to affect a large number of systems. Hierarchical design helps constrain operational changes to a subset of the network, which makes it easy to manage as well as improve resiliency. Modular structuring of the network into small, easy-to-understand elements also facilitates resiliency via improved fault isolation.

A hierarchical LAN design includes the following three layers:

- **Access layer**—Provides endpoints and users direct access to the network.
- **Distribution layer**—Aggregates access layers and provides connectivity to services.
- **Core layer**—Provides connections between distribution layers for large environments.

Figure 1 - LAN hierarchical design
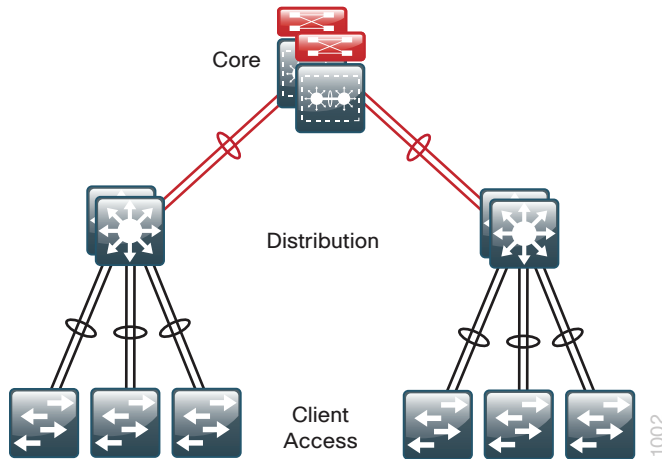
Each layer—access, distribution, and core—provides different functionality and capability to the network. Depending on the characteristics of the network deployment site, you might need one, two, or all three of the layers. For example, a site that occupies a single building might only require the access and distribution layers, while a campus of multiple buildings will most likely require all three layers.

Regardless of how many layers are implemented at a location, the modularity of this design ensures that each layer will provide the same services, and in this architecture, will use the same design methods.

Figure 2 - Scalability by using a modular design
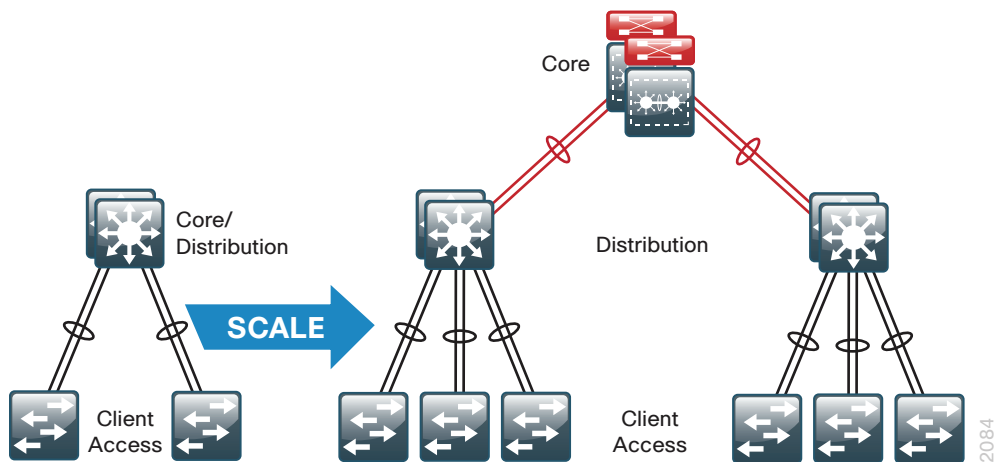
Access Layer

The access layer is where user-controlled devices, user-accessible devices, and other end-point devices are connected to the network. The access layer provides both wired and wireless connectivity and contains features and services that ensure security and resiliency for the entire network.

Device Connectivity

The access layer provides high-bandwidth device connectivity. Once expensive options, high-bandwidth access technologies like Gigabit Ethernet and 802.11n and 802.11ac wireless are now standard configurations on end-user devices. While an end-user device in most cases will not use the full capacity of these connections for long periods of time, the ability to burst up to these high bandwid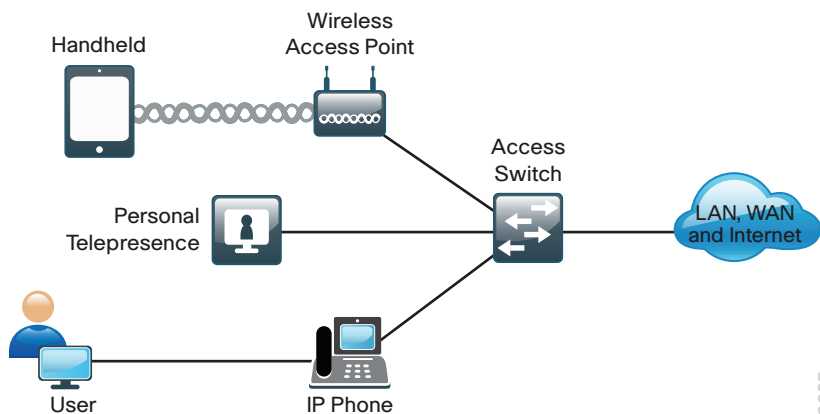ths when performing routine tasks does help make the network a transparent part of an end-users day-to-day job. The longer someone has to wait to back up their machine, send an email, or open a file off an internal web page, the harder it is for the network to be transparent.

Figure 3 - Access layer connectivity

It is common for many different types of devices to connect at the access layer. Personal computers, IP phones, wireless access points, and IP video surveillance cameras all might connect to the same access layer switch. Since it can be beneficial for performance, management, and security reasons to segment these different devices, the access layer provides the capability to support many logical networks on one physical infrastructure.

Resiliency and Security Services

In general, the goal of the resiliency and security services in the infrastructure is to ensure that the network is available for use without impairment for everyone that needs it. Because the access layer is the connection point between the network and client devices, it plays a role in ensuring the network is protected from human error and from malicious attacks. This protection includes making sure the devices connecting to the network do not attempt to provide services to any end users that they are not authorized for, that they do not attempt to take over the role of any other device on the network, and, when possible, that they verify the device is allowed on the network.

Enabling these services in the access layer contributes not only to the overall security of the network, but also to the resiliency and availability of the network.

Advanced Technology Capabilities

Finally, the access layer provides a set of network services that support advanced technologies. Voice and video are commonplace in today's organizations and the network must provide services that enable these technologies. This includes providing specialized access for these devices, ensuring others do not impair the traffic from these devices, and providing efficient delivery of traffic that is needed by many devices in the network.

Distribution Layer

The distribution layer supports many important services for the LAN. The primary function is to serve as an aggregation point for multiple access layer switches in a given location or campus, and serve as the demarcation between the layer-2 switching and layer-3 routing functions in this design. In a network where connectivity needs to traverse the campus network end-to-end, whether between different access layer devices or from an access layer device to the WAN, the distribution layer facilitates this connectivity.

Scalability

In any network where multiple access layer devices exist at a location to serve end-user connectivity, it becomes impractical to completely interconnect all access switches as the access layer grows beyond two or three switches.

The distribution layer provides a logical point to summarize addressing and to create a boundary for protocols and features necessary for the access layer operation. Another benefit of the distribution layer boundary is that it creates fault domains that serve to contain failures or network changes to those parts of the network directly affected.

The end result to the organization is that the distribution layer can lower the cost of operating the network by making it more efficient, by requiring less memory, and by processing resources for devices elsewhere in the network. The distribution layer also increases network availability by containing failures to smaller domains.

Reduce Complexity and Increase Resiliency

This design uses a simplified distribution layer. Organizations benefit from the consistency and reduced complexity features of the simplified distribution layer design by lower operational costs of configuring and maintaining the network.

The simplified distribution layer design consists of a single logical entity that can be implemented using a pair of physically separate switches operating as one device, or a physical stack of switches operating as one device. Using a single logical entity reduces complexity of configuring and operating the distribution layer, as fewer protocols are required and little or no tuning is needed to provide near-second or sub-second convergence around failures or disruptions.

The design resiliency is provided using physically redundant components such as power supplies, supervisors, and modules, as well as implementing Stateful Switchover with redundant logical control planes. There are other variations not validated as part of this design, which may meet the needs of an organization with less stringent redundancy requirements for their distribution layer. For example, a single physical device with redundant components could be suitable for a high-density space-constrained environment.

Flexible Design

The distribution layer provides connectivity to network-based services, to the WAN, and to the Internet edge, either with directly connected or connected through a core layer. Network-based services can include and are not limited to Wide Area Application Services (WAAS) and wireless LAN controllers. Depending on the size of the campus network, these services and the inter-connection to the WAN and Internet edge may reside on a distribution layer switch that also aggregates the LAN access layer connectivity. This is also referred to as a *collapsed core* design because the distribution serves as the Layer 3 aggregation layer for all devices.

Figure 4 - Two tier design: Distribution layer functioning as a collapsed Core
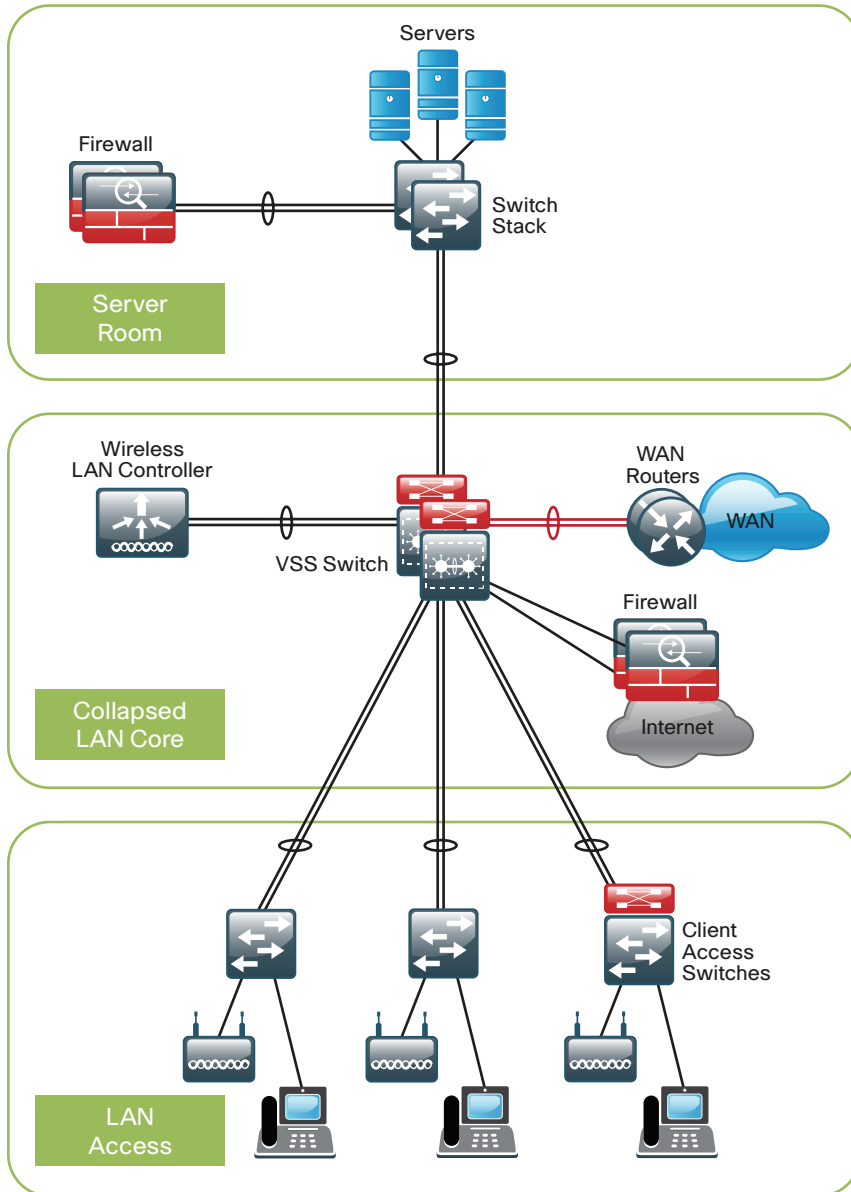
Larger LAN designs require a dedicated distribution layer for network-based services versus sharing connectivity with access layer devices. As the density of WAN routers, WAAS controllers, Internet edge devices, and wireless LAN controllers grows, the ability to connect to a single distribution layer switch becomes hard to manage. There are a number of factors that drive LAN design with multiple distribution layer modules:

- The number of ports and port bandwidth that the distribution layer platform can provide affects network performance and throughput.

- Network resilience is a factor when all LAN and network-based services rely on a single platform, regardless of that platform's design, it can present a single point of failure or an unacceptably large failure domain.

- Change control and frequency affects resilience. When all LAN, WAN, and other network services are consolidated on a single distribution layer, operational or configuration errors can affect all network operation.

- Geographic dispersion of the LAN access switches across many buildings in a larger campus facility would require more fiber optic interconnects back to a single collapsed core.

Figure 5 - Network-services distribution layer
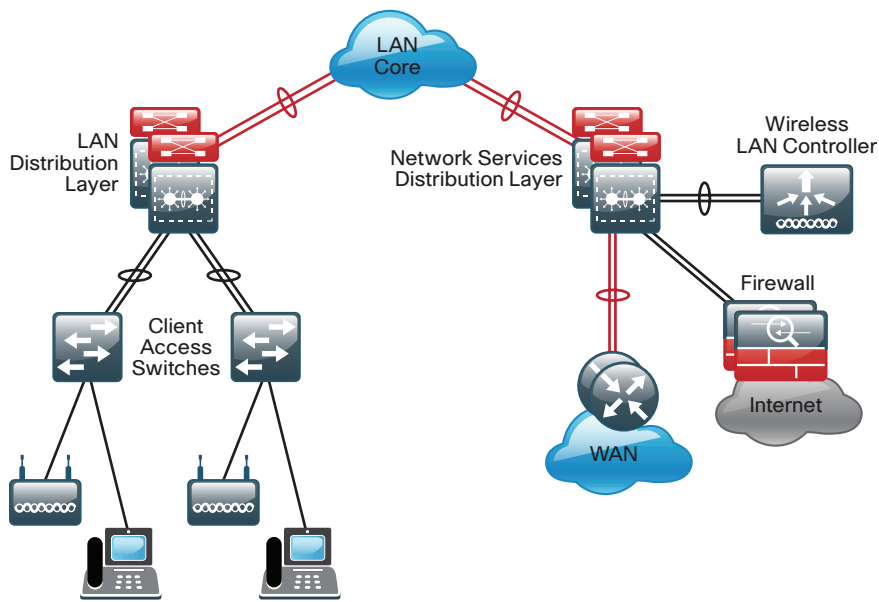
Like the access layer, the distribution layer also provides QoS for application flows to guarantee critical applications and multimedia applications perform as designed.

Core Layer

In a large LAN environment there often arises a need to have multiple distribution layer switches. One reason for this is that when access layer switches are located in multiple geographically dispersed buildings, you can save potential costly fiber-optic runs between buildings by locating a distribution layer switch in each of those buildings. As networks grow beyond three distribution layers in a single location, organizations should use a core layer to optimize the design.

Another reason to use multiple distribution layer switches is when the number of access layer switches connecting to a single distribution layer exceeds the performance goals of the network designer. In a modular and scalable design, you can colocate distribution layers for data center, WAN connectivity, or Internet edge services.

In environments where multiple distribution layer switches exist in close proximity and where fiber optics provide the ability for high-bandwidth interconnect, a core layer reduces the network complexity, as shown in the following two figures.

Figure 6 - LAN topology with a core layer
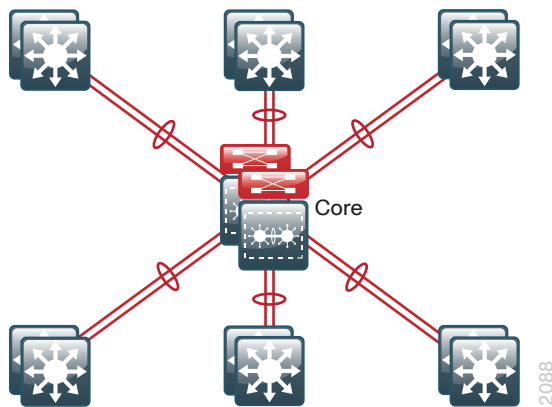
Figure 7 - LAN topology without a core layer
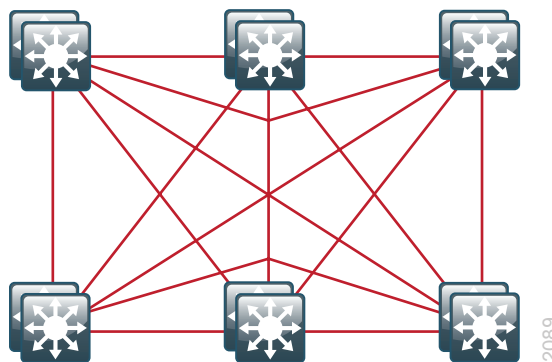
The core layer of the LAN is a critical part of the scalable network, and yet it is one of the simplest by design. The distribution layer provides the fault and control domains, and the core represents the 24x7x365 nonstop connectivity between them, which organizations must have in the present business environment where connectivity to resources to conduct business is critical.

In the campus core-layer design where Cisco Catalyst 6800 or 6500 Series Switches are used, a Catalyst VSS Layer-3 core design is chosen as the preferred alternative to traditional designs, which often use two independently configured and managed platforms. Connectivity to and from the core is Layer 3 only, which drives increased resiliency and stability.

Reader Tip

For an in-depth VSS configuration guide and configuration options, see the *Campus 3.0 Virtual Switching System Design Guide*, available on the Cisco Validated Designs tab at www.cisco.com/go/cvd/campus.

Quality of Service (QoS)

Real-time communication traffic is very sensitive to delay and drop. The network must ensure that this type of traffic is handled with priority so that the stream of audio or video is not interrupted. QoS is the technology that answers this need.

The primary role of QoS in rich-media campus networks is to manage packet loss, where high-bandwidth links with instantaneous congestion on the order of milliseconds can cause buffer overruns and a poor user experience. Another goal of campus QoS is to apply policies at the edge to allow consistent treatment of traffic for a predictable user experience across the entire network.

QoS allows an organization to define different traffic types and to create more deterministic handling for real-time traffic. QoS is especially useful in congestion handling where a full communications channel might prevent voice or video streams from being intelligible at the receiving side. Congestion is common when links are oversubscribed while aggregating traffic from a number of devices and when the link bandwidth to a device is steps down from a higher bandwidth uplink causing a delay in delivery of some of the traffic in transit. Rather than creating bandwidth, QoS takes bandwidth from one class (that is, generally the default traffic class) and gives it to another class during periods of congestion.

Within this design, the QoS profiles are as simple as possible while ensuring support for applications that need special delivery. This approach establishes a solid, scalable, and modular framework to implement QoS across the entire network.

The primary goals of implementing QoS within the network are:

- Expedited delivery service of communications for supported, real-time applications.
- Business continuance for business-critical applications.
- Fairness among all other applications when congestion occurs.
- Deprioritized background applications and non-business entertainment-oriented applications so that these do not delay interactive or business-critical applications.
- A trusted edge around the network to guarantee that users cannot inject their own arbitrary priority values and to allow the organization to trust marked traffic throughout the network.

To accomplish these goals, the design implements QoS across the network as follows:

- Establish a limited number of traffic classes (that is, typically one to eight classes) within the network that need special handling (for example, real-time voice, real-time video, high-priority data, interactive traffic, batch traffic, and default classes).
- Classify applications into the traffic classes.
- Apply special handling to the traffic classes to achieve intended network behavior.

In this design, QoS configurations are as simple as possible, and are applied only to those applications that require special handling.

This approach establishes a solid, scalable, and modular framework to implement QoS across the entire network.

Access Layer

Design Overview

The access layer is the point at which user-controlled and user-accessible devices are connected to the network and it is the one architecture component that is found in every LAN.

Infrastructure Security Features

Because the access layer is the connection point between network-based services and client devices, it plays an important role in protecting other users, the application resources, and the network itself from human error and malicious attacks. Network resiliency and security in the access layer is achieved through the use of Cisco Catalyst Infrastructure Security Features (CISF) including Dynamic Host Configuration Protocol (DHCP) snooping, IP Source Guard, port security, and Dynamic Address Resolution Protocol (ARP) Inspection.

MAC flooding attacks are used to force a LAN switch to flood all switch traffic out to all the switch interfaces. Port security limits the number of MAC addresses that can be active on a single port to protect against such attacks.

Port security lets you to configure Layer 2 interfaces to allow inbound traffic from only a restricted set of MAC addresses. The MAC addresses in the restricted set are called secure MAC addresses. In addition, the device does not allow traffic from these MAC addresses on another interface within the same VLAN.

The number of MAC addresses that the device secures on each interface is configurable. For ease of management, the device can learn the addresses dynamically. Using the dynamic learning method, the device secures MAC addresses while ingress traffic passes through the interface. If the address is not yet secured and the device has not reached any applicable maximum, it secures the address and allows the traffic. The device ages dynamic addresses and drops them when the age limit is reached.

DHCP snooping is a security feature for DHCP that filters and rate-limits DHCP traffic from untrusted sources. An untrusted source is any interface on the switch not specifically configured as a known DHCP server or path towards a known DHCP server, including all client-facing interfaces, allowing DHCP replies to be blocked from those interfaces.

The DHCP snooping feature helps simplify management and troubleshooting by tracking MAC address, IP address, lease time, binding type, VLAN number, and interface information that correspond to the local untrusted interfaces on the switch. DHCP snooping stores this information in the DHCP binding table, which is then used as a reference for comparison against observed traffic.

Dynamic ARP inspection (DAI) mitigates ARP poisoning attacks. An ARP poisoning attack is a method by which an attacker sends false ARP information to a local segment. This information is designed to poison the ARP cache of devices on the LAN, allowing the attacker to execute man-in-the-middle attacks.

Figure 8 - DHCP snooping and Dynamic ARP inspection
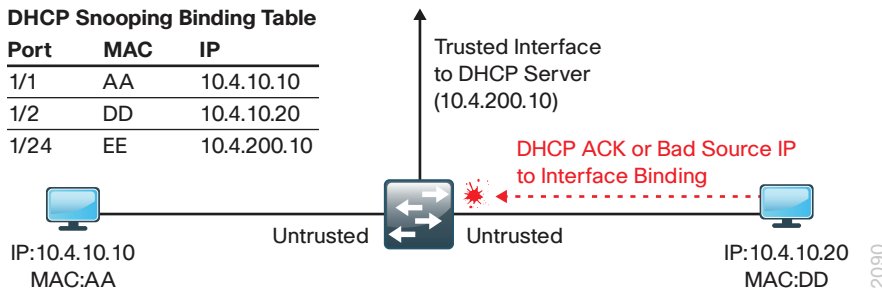
DAI uses the data generated by the DHCP snooping feature and intercepts and validates the IP-to-MAC address relationship of all ARP packets on untrusted interfaces. ARP packets that are received on trusted interfaces are not validated and invalid packets on untrusted interfaces are discarded.

IP Source Guard is a means of preventing a packet from using an incorrect source IP address to obscure its true source, also known as *IP spoofing*. IP Source Guard uses information from DHCP snooping to dynamically configure a port access control list (PACL) on the interface that denies any traffic from IP addresses that are not in the DHCP binding table.

Common Design Method to Simplify Installation and Operation

To provide consistent access capabilities and simplify network deployment and operation, the design uses a common deployment method for all access layer devices, whether they are located in the headquarters or at a remote site. To reduce complexity, the access layer is designed so that you can use a single interface configuration to accommodate a variety of device connectivity, such as for a standalone computer, an IP phone, an IP phone with an attached computer, or a wireless access point.

The LAN access layer provides high-bandwidth connections to devices via 10/100/1000 Ethernet with both Gigabit and 10-Gigabit uplink connectivity options. The 10 Gigabit uplinks also support Gigabit connectivity to provide flexibility and help business continuity during a transition to 10 Gigabit Ethernet. The LAN access layer is configured as a Layer 2 switch, with all Layer 3 services being provided either by the directly-connected distribution layer or router.

Figure 9 - Access layer overview
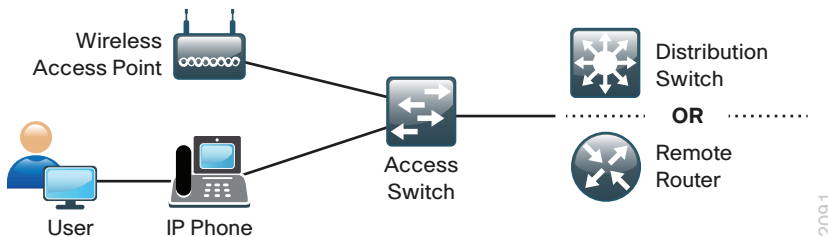
Features to Support Voice and Video Deployment

Voice and video are enabled in the access layer via network services such as Power over Ethernet (PoE), QoS, multicast support, and Cisco Discovery Protocol (CDP) with the voice VLAN.

PoE enables devices such as IP phones, wireless access points, virtual desktops, and security cameras to be powered by the access layer device. This removes the expense of installing or modifying building power to support devices in difficult to reach locations and allows for the consolidation of back-up power supplies and Uninterruptable Power Supplies (UPSs) to the access closet.

To support the increasing requirements of devices powered by the network, all of the access layer devices support the IEEE 802.3at standard, also known as PoE+. The devices, and or line cards support all the previous implementations of PoE up to 15 watts per port as well as the new IEEE 802.3at implementation of up to 30 watts per port. For the most demanding PoE environments, like virtual desktops, the Cisco Catalyst 4500, Catalyst 3850, and Catalyst 3750-X in the access layer have options to provide up to 60 watts of power per port with Cisco Universal Power Over Ethernet (UPOE) over the same cable plant as you use for PoE+.

Cisco Discovery Protocol supports voice and video device integration into the access layer. Cisco IP Phones that are plugged into the access layer communicate bidirectionally with the access layer switch via Cisco Discovery Protocol. Cisco Discovery Protocol provides the IP Phone with configuration information and provides the access layer switch with the IP Phones power requirements and the ability to selectively prioritize traffic from the IP Phone.

Access Layer Platforms

Wiring Closets Requiring up to 48 Ports

Cisco Catalyst 2960-X and 3650 Series are economical 10/100/1000 Ethernet fixed-port switches that provide flexibility and common features required for wiring closets that can be supported by a single fixed port switch. Cisco Catalyst 2960-X and 3650 are available in both PoE+ and non-power-supplying versions and have optional support for stacking.

Cisco Catalyst 3650 additional capabilities include support for dual replaceable redundant power supplies, dual redundant fans, and enhanced enterprise options such as Cisco TrustSec and NetFlow. The Cisco Catalyst 3650 has an integrated wireless controller which optionally can be enabled for converged wired and wireless access.

Wiring Closets Requiring Greater than 48 Ports

When a wiring closet requires greater interface density than can be provided by a single switch, an intelligent stack of fixed configuration switches or a modular switch is recommended.

Intelligent stacks or modular Ethernet switches provide the following major benefits:

- **Single point of management**—All switches in the stack are managed as one.
- **Built-in redundancy and high availability**—The high-bandwidth dedicated stack connections provide redundant communication for each stack member.
- **Scalable to fit network needs**—As the need for additional access interfaces grows, adding a new switch to a stack or a module to a modular switch is easy.

The following series of Cisco Catalyst switches are used in this design when intelligent stacking or a modular deployment is required in a single access layer in a wiring closet: Cisco Catalyst 2960-X, 3750-X, 3650, 3850, and 4500E Series.

Cisco Catalyst 2960-X Series are fixed-configuration, stackable, 10/10/1000 Ethernet switches, with PoE+ and non-power-supplying versions designed for entry-level enterprise, midmarket, and remote site networks.

- Cisco FlexStack+ is implemented by adding a stacking module to the Cisco Catalyst 2960-X Series Switch. This enables up to eight Catalyst 2960-X series switches to be stacked together. Cisco FlexStack+ links are full duplex 20-Gbps links with typical recovery time between 1-2 seconds.

Cisco Catalyst 3750-X Series are fixed-port, stackable, 10/100/1000 Ethernet switches, with PoE+, Cisco UPOE, and non-power-supplying versions, which provide enhanced resiliency through StackWise Plus and StackPower technologies.

- Cisco StackWise Plus enables up to nine Cisco Catalyst 3750-X Series Switches to be stacked together using a 64-Gbps stack interconnect with rapid failure recovery.

- Cisco StackPower shares power across the Cisco Catalyst 3750-X Series Switch stack. This allows the flexible arrangement of power supplies in the stack, and enables a zero-footprint redundant power supply deployment and intelligent load shedding.

- Cisco 3750-X Series Switches have modular uplinks and support upgrading the Cisco IOS feature set and enhanced enterprise capabilities like TrustSec, Flexible NetFlow, and Cisco IOS Sensor to ensure that the switch functionality grows as the organization grows.

Cisco Catalyst 3650 Series and Catalyst 3850 Series Switches are fixed-port, stackable, 10/100/1000 Ethernet switches, with PoE+ and non-power-supplying versions, which provide enhanced switching performance and resiliency through StackWise-160 (Cisco Catalyst 3650) or StackWise-480 and StackPower technologies (Cisco Catalyst 3850), with Flexible NetFlow capabilities on all ports.

- Cisco Catalyst 3650 stacking is implemented with an optional stacking module. Switches stack together using StackWise-160 mode with up to nine switches in single stack-ring.

- Cisco Catalyst 3850 Series Switches have built-in stacking capability, and stack together using StackWise-480 mode with up to nine switches in single stack-ring.

- Cisco StackPower technology increases system-level resiliency during catastrophic power failure on a stack-member switch. Cisco StackPower enables power redundancy across a group of four Cisco Catalyst 3850 Series Switches within same stack. This allows the flexible arrangement of power supplies in the stack, and enables a zero-footprint redundant power supply deployment and intelligent load shedding.

- Cisco 3650 Series Switches have fixed uplinks that can be configured as Gigabit Ethernet or 10-Gigabit Ethernet.

- Cisco 3850 Series Switches have modular uplinks that can be configured as Gigabit Ethernet or 10-Gigabit Ethernet.

- Cisco Catalyst 3650 and Cisco Catalyst 3850 Series supports Stateful Switchover, which allows a switch in the active role in a stack to rapidly switchover to a switch in the standby role with minimum disruption to the network.

- With appropriate licenses, the Cisco Catalyst 3650 and Cisco Catalyst 3850 Series hardware supports wireless LAN controller functionality in order to support a unified access policy for converged wired and wireless designs.

Cisco Catalyst 4500E Series are modular switches that support multiple Ethernet connectivity options, including 10/100/1000 Ethernet, 100-Megabit fiber, Gigabit fiber, and 10-Gigabit fiber. The Catalyst 4500E Series Switches also have an upgradable supervisor module that enables future functionality to be added with a supervisor module upgrade while maintaining the initial investment in the chassis and the modules.

- All key switching and forwarding components are located on the supervisor module; upgrading the supervisor upgrades the line cards.

- The Cisco Catalyst 4500E Series Supervisor Engine 7-E and Supervisor Engine 7L-E have uplink interfaces that can be configured as Gigabit Ethernet or 10-Gigabit interfaces, allowing organizations to easily increase bandwidth in the future. The Supervisor Engine 8-E includes eight 10-Gigabit interfaces on board, along with integrated wireless LAN controller hardware which can be enabled in the future, and extends the range of performance up to 48-Gbps per slot and 928 Gbps system switching capacity.

- The Cisco Catalyst 4500E Series provides maximum PoE flexibility with support of IEEE 802.3af, 802.3at, and now Cisco UPOE, and supplies up to 60 watts per port of power over Ethernet. Cisco UPOE line cards are backward compatible to earlier PoE and PoE+ connected end points as well.

- The Cisco Catalyst 4507R+E chassis supports redundant supervisor modules and power supplies, which increases system availability by providing 1:1 redundancy for all critical systems. When configured with dual Supervisor modules, Stateful Switchover, which allows a supervisor switchover to occur with minimum disruption to the network.

- With a dual Supervisor Engine system, the entire software upgrade process is simplified by using In-Service Software Upgrade (ISSU). Not only does ISSU help eliminate errors in the software upgrade process, but additional checks are incorporated that allow the new software version to be tested and verified before completing the upgrade.

Deployment Details

How to Read Commands

This guide uses the following conventions for commands that you enter at the command-line interface (CLI).

Commands to enter at a CLI prompt:
```
configure terminal
```

Commands that specify a value for a variable:
```
ntp server 10.10.48.17
```

Commands with variables that you must define:
```
class-map [highest class name]
```

Commands at a CLI or script prompt:
```
Router# enable
```

Long commands that line wrap are underlined. Enter them as one command:
```
police rate 10000 pps burst 10000
packets conform-action
```

Noteworthy parts of system output (or of device configuration files) are highlighted:
```
interface Vlan64
   ip address 10.5.204.5 255.255.255.0
```

As you review this guide, you may find it useful to understand the IP addressing and VLAN assignments used. Although your design requirements may differ, by addressing the various distribution layers at a location with contiguous IP address space, you can summarize the IP address range to the rest of the network. This design uses VLAN assignments that reflect the third octet of the IP address range for a given access layer switch for ease of reference. Alternatively, many organizations may choose to reuse the same VLAN IDs in each distribution—use and document the method that makes the most sense for your organization. The LAN Core IP addressing is a combination of 30-bit subnets for point-to-point Layer 3 links, and 32-bit host addresses for loopback addresses.

Table 1 - IP addressing for Campus Wired LAN Technology Design Guide

Address block	Access VLAN	IP addressing	Usage
Distribution #1 10.4.0.0/20	100	10.4.0.0/24	Data-Access Switch 1
	101	10.4.1.0/24	Voice-Access Switch 1
	102	10.4.2.0/24	Data-Access Switch 2
	103	10.4.3.0/24	Voice-Access Switch 2
	Continue through 113	10.4.4.0/24–10.4.13.0/24	alternate Data and Voice
	115	10.4.15.0/25	Management
	None	10.4.15.128/32–10.4.15.255/32	Loopback Interfaces
Distribution #2 10.4.64.0/20	164	10.4.64.0/24	Data-Access Switch 1
	165	10.4.65.0/24	Voice-Access Switch 1
	166	10.4.66.0/24	Data-Access Switch 2
	167	10.4.67.0/24	Voice-Access Switch 2
	Continue through 177	10.4.68.0/24–10.4.77.0/24	alternate Data and Voice
	179	10.4.79.0/25	Management
	None	10.4.79.128/32–10.4.79.255/32	Loopback Interfaces
Distribution #3 10.4.80.0/20	180	10.4.80.0/24	Data-Access Switch 1
	181	10.4.81.0/24	Voice-Access Switch 1
	182	10.4.82.0/24	Data-Access Switch 2
	183	10.4.83.0/24	Voice-Access Switch 2
	Continue through 193	10.4.84.0/24–10.4.93.0/24	alternate Data and Voice
	195	10.4.95.0/25	Management
	None	10.4.95.128/32–10.4.95.255/32	Loopback Interfaces
Distribution #4 10.4.96.0/20	196	10.4.96.0/24	Data-Access Switch 1
	197	10.4.97.0/24	Voice-Access Switch 1
	198	10.4.98.0/24	Data-Access Switch 2
	199	10.4.99.0/24	Voice-Access Switch 2
	Continue through 209	10.4.100.0/24–10.4.109.0/24	alternate Data and Voice
	211	10.4.111.0/25	Management
	None	10.4.111.128/32–10.4.111.255/32	Loopback Interfaces
Distribution #5 10.4.112.0/20	212	10.4.112.0/24	Data-Access Switch 1
	213	10.4.113.0/24	Voice-Access Switch 1
	214	10.4.114.0/24	Data-Access Switch 2
	215	10.4.115.0/24	Voice-Access Switch 2
	Continue through 225	10.4.116.0/24–10.4.125.0/24	alternate Data and Voice
	227	10.4.127.0/25	Management
	None	10.4.127.128/32–10.4.127.255/32	Loopback Interfaces
Core 10.4.40.0/24	None	10.4.40.0/30–10.4.40.124/30	Core to Distribution Links
	None	10.4.40.128/32–10.4.40.255/32	Loopback Interfaces

Configuring the Access Layer

1. Configure the platform
2. Configure LAN switch universal settings
3. Configure access switch global settings
4. Configure client connectivity
5. Connect to distribution or WAN router

Procedure 1 Configure the platform

Some platforms require a one-time initial configuration prior to configuring the features and services of the switch. If you do not have a platform listed for an option, skip that option and its steps.

Option 1: Configure the Cisco Catalyst 2960-X and 3750-X

Step 1: If you are configuring a stack of switches, set the stack master switch.

```
switch [switch number] priority 15
```

When there are multiple Cisco Catalyst 2960-X or 3750-X Series Switches configured in a stack, one of the switches controls the operation of the stack and is called the stack master. When three or more switches are configured as a stack, configure the stack master switch functionality on a switch that does not have uplinks configured.

Figure 10 - Stack master placement in a switch stack
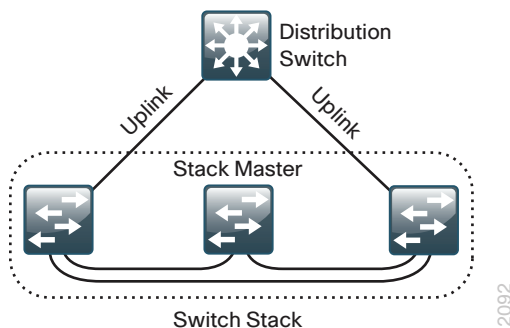
If you configure stack master switch priority on a Cisco Catalyst 2960-X or Cisco 3750-X switch stack, a single reload is required to force the stack master to operate on the switch that you configured with the highest priority. Reload the switch stack after all of your configuration is complete for this entire "Configuring the Access Layer" process.

Step 2: If you are configuring a stack, run the **stack-mac persistent timer 0** command. This ensures that the original stack master MAC address is used by any switch in the stack that takes the stack master role after a switchover.

```
Switch(config)#stack-mac persistent timer 0
```

The default behavior when the stack master switch fails is for the newly active stack master switch to assign a new stack MAC address. This new MAC address assignment can cause the network to reconverge because the link aggregation control protocol (LACP) and many other protocols rely on the stack MAC address and must restart.

Step 3: For each platform, define two macros that you will use in later procedures to apply the platform specific QoS configuration. This makes consistent deployment of QoS easier.

```
macro name AccessEdgeQoS
 auto qos voip cisco-phone
@
!
macro name EgressQoS
 mls qos trust dscp
 queue-set 1
 srr-queue bandwidth share 1 30 35 5
 priority-queue out
@
```

Option 2: Configure the Cisco Catalyst 3650 and 3850 platform

Step 1: To configure a Cisco Catalyst 3650 or 3850 stack, use the CLI global exec mode (not configuration mode) to set the preferred active switch.

```
switch [switch number] priority 15
```

When there are multiple Cisco Catalyst 3650 or 3850 Series Switches configured in a stack, one of the switches takes the ACTIVE switch role, and another member takes the HOT-STANDBY role. Upon reload, the switch configured with the highest priority assumes the active role. If this is a new configuration, only the active switch console is active during the initial configuration. When two or more switches are configured as a stack, configure the active switch functionality on a switch of your preference.

Step 2: For each platform, define two macros that you will use in later procedures to apply the platform-specific QoS configuration. This makes consistent deployment of QoS easier.

---

i | **Tech Tip**

The match statements in the class-maps can be combined into one line. However, listing them separately provides additional per-DSCP counters available from SNMP and from the following verification command.

```
show policy-map interface
```

```
table-map policed-dscp
 map from 0 to 8
 map from 10 to 8
 map from 18 to 8
 map from 24 to 8
 map from 46 to 8
 default copy
!
class-map match-any CISCO-PHONE-VOICE
```

```
    match cos 5
class-map match-any CISCO-PHONE-SIGNALING
    match cos 3
!
class-map match-any PRIORITY-QUEUE
    match  dscp ef
class-map match-any VIDEO-PRIORITY-QUEUE
    match  dscp cs5
    match  dscp cs4
class-map match-any CONTROL-MGMT-QUEUE
    match  dscp cs7
    match  dscp cs6
    match  dscp cs3
    match  dscp cs2
class-map match-any MULTIMEDIA-CONFERENCING-QUEUE
    match  dscp af41
    match  dscp af42
    match  dscp af43
class-map match-any MULTIMEDIA-STREAMING-QUEUE
    match  dscp af31
    match  dscp af32
    match  dscp af33
class-map match-any TRANSACTIONAL-DATA-QUEUE
    match  dscp af21
    match  dscp af22
    match  dscp af23
class-map match-any BULK-SCAVENGER-DATA-QUEUE
    match  dscp af11
    match  dscp af12
    match  dscp af13
    match  dscp cs1
!
policy-map CISCO-PHONE
 class CISCO-PHONE-VOICE
    set dscp ef
   police cir 128000 bc 8000
     conform-action transmit
     exceed-action set-dscp-transmit dscp table policed-dscp
 class CISCO-PHONE-SIGNALING
    set dscp cs3
   police cir 32000 bc 8000
     conform-action transmit
     exceed-action set-dscp-transmit dscp table policed-dscp
 class class-default
    set dscp default
!
policy-map 2P6Q3T
```

```
     class PRIORITY-QUEUE
        priority level 1
        police rate percent 10
     class VIDEO-PRIORITY-QUEUE
        priority level 2
        police rate percent 20
     class CONTROL-MGMT-QUEUE
        bandwidth remaining percent 10
        queue-buffers ratio 10
     class MULTIMEDIA-CONFERENCING-QUEUE
        bandwidth remaining percent 10
        queue-buffers ratio 10
        queue-limit dscp af43 percent 80
        queue-limit dscp af42 percent 90
        queue-limit dscp af41 percent 100
     class MULTIMEDIA-STREAMING-QUEUE
        bandwidth remaining percent 10
        queue-buffers ratio 10
        queue-limit dscp af33 percent 80
        queue-limit dscp af32 percent 90
        queue-limit dscp af31 percent 100
     class TRANSACTIONAL-DATA-QUEUE
        bandwidth remaining percent 10
        queue-buffers ratio 10
        queue-limit dscp af23 percent 80
        queue-limit dscp af22 percent 90
        queue-limit dscp af21 percent 100
     class BULK-SCAVENGER-DATA-QUEUE
        bandwidth remaining percent 5
        queue-buffers ratio 10
        queue-limit dscp values af13 cs1 percent 80
        queue-limit dscp values af12 percent 90
        queue-limit dscp values af11 percent 100
     class class-default
        bandwidth remaining percent 25
        queue-buffers ratio 25
 !
 macro name AccessEdgeQoS
  trust device cisco-phone
  service-policy input CISCO-PHONE
  service-policy output 2P6Q3T
 @
 !
 macro name EgressQoS
  service-policy output 2P6Q3T
 @
```

Option 3: Configure the Cisco Catalyst 4507R+E platform

Step 1: For each platform, define two macros that you will use in later procedures to apply the platform-specific QoS configuration. This makes consistent deployment of QoS easier.

```
class-map match-any PRIORITY-QUEUE
  match  dscp ef
  match  dscp cs5
  match  dscp cs4
class-map match-any CONTROL-MGMT-QUEUE
  match  dscp cs7
  match  dscp cs6
  match  dscp cs3
  match  dscp cs2
class-map match-any MULTIMEDIA-CONFERENCING-QUEUE
  match  dscp af41
  match  dscp af42
  match  dscp af43
class-map match-any MULTIMEDIA-STREAMING-QUEUE
  match  dscp af31
  match  dscp af32
  match  dscp af33
class-map match-any TRANSACTIONAL-DATA-QUEUE
  match  dscp af21
  match  dscp af22
  match  dscp af23
class-map match-any BULK-DATA-QUEUE
  match  dscp af11
  match  dscp af12
  match  dscp af13
class-map match-any SCAVENGER-QUEUE
  match  dscp cs1
!
policy-map 1P7Q1T
 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
    !
    macro name AccessEdgeQoS
     auto qos voip cisco-phone
    @
    !
    macro name EgressQoS
      service-policy output 1P7Q1T
    @
```

Step 2: When a Cisco Catalyst 4507R+E is configured with two Supervisor Engine 7L-E, 7-E, or 8-E modules, configure the switch to use Stateful Switchover (SSO) when moving the primary supervisor functionality between modules. SSO synchronizes active process information as well as configuration information between supervisor modules, which enables a fast transparent data plane failover.

```
    redundancy
      mode sso
```

> **i** | **Tech Tip**
>
> To enable SSO mode you must have a license level of ipbase or entservices operating on the switch supervisors. You can check the current license level of operation with a **show version** command.

Procedure 2 Configure LAN switch universal settings

Within this design, there are features and services that are common across all LAN switches, regardless of the type of platform or role in the network. These are system settings that simplify and secure the management of the solution.

This procedure provides examples for some of those settings. The actual settings and values will depend on your current network configuration.

Table 2 - Common network services used in the deployment examples

Setting	Value
Domain Name	cisco.local
Active Directory, DNS, DHCP Server	10.4.48.10
Authentication Control System	10.4.48.15
Network Time Protocol Server	10.4.48.17

Step 1: Configure the device hostname to make it easy to identify the device.

```
hostname [hostname]
```

Step 2: If the switch VTP mode has been changed from default, configure VTP transparent mode. This design uses VTP transparent mode because the benefits of dynamic propagation of VLAN information across the network are not worth the potential for unexpected behavior resulting from operational error.

VLAN Trunking Protocol (VTP) allows network managers to configure a VLAN in one location of the network and have that configuration dynamically propagate out to other network devices. However, in most cases, VLANs are defined once during switch setup with few, if any, additional modifications.

```
vtp mode transparent
```

Step 3: Enable Rapid Per-VLAN Spanning-Tree (PVST+). Rapid PVST+ provides an instance of RSTP (802.1w) per VLAN. Rapid PVST+ greatly improves the detection of indirect failures or linkup restoration events over classic spanning tree (802.1D).

Although this architecture is built without any Layer 2 loops, you should still enable spanning tree with the most up-to-date network safeguards. By enabling spanning tree, you ensure that if any physical or logical loops are accidentally configured, no actual layer 2 loops occur.

```
spanning-tree mode rapid-pvst
```

Step 4: Enable Unidirectional Link Detection (UDLD) as the default for fiber ports.

UDLD is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. When UDLD detects a unidirectional link, it disables the affected interface and alerts you. Unidirectional links can cause a variety of problems, including spanning-tree loops, black holes, and non-deterministic forwarding. In addition, UDLD enables faster link failure detection and quick reconvergence of interface trunks, especially with fiber, which can be susceptible to unidirectional failures.

```
udld enable
```

Step 5: Enable the recovery mechanism to allow ports disabled as a result of errors to automatically clear the err-disable status and attempt a recovery to operational behavior and connected status. Enabling the recovery mechanism avoids having to intervene manually by using the CLI to shutdown and enable the port after the cause of the error is cleared. By default, the recovery mechanism waits five minutes to attempt clearing of the interface err-disable status.

```
errdisable recovery cause all
```

Step 6: Set EtherChannels to use the traffic source and destination IP address when calculating which link to send the traffic across. This normalizes the method in which traffic is load-shared across the member links of the EtherChannel. EtherChannels are used extensively in this design because of their resiliency capabilities.

```
port-channel load-balance src-dst-ip
```

Step 7: Configure DNS for host lookup.

At the command line of a Cisco IOS device, it is helpful to be able to type a domain name instead of the IP address for a destination.

```
ip name-server 10.4.48.10
```

Step 8: Configure device management protocols.

Secure HTTP (HTTPS) and Secure Shell (SSH) are more secure replacements for the HTTP and Telnet protocols. They use Secure Sockets Layer (SSL) and Transport Layer Security (TLS) to provide device authentication and data encryption.

The SSH and HTTPS protocols enable secure management of the LAN device. Both protocols are encrypted for privacy, and the unencrypted protocols, Telnet and HTTP, are turned off. Enabling HTTPS automatically generates a cryptographic key to use the service. When SSH is configured after HTTPS, you do not have to explicitly generate the cryptographic key that SSH requires, unless you wish to change the default key size.

Secure Copy (SCP) provides a secure and authenticated method for copying configuration and image files by making use of SSH as a secure transport. Enable SCP to allow secure file management with the device to avoid the use of less secure protocols such as TFTP and FTP.

Specify the transport preferred none on vty lines in order to prevent errant connection attempts from the CLI prompt. Without this command, if the ip name server is unreachable, long timeout delays may occur for mistyped commands.

```
no ip http server
ip http secure-server
ip domain-name cisco.local
ip ssh version 2
ip scp server enable
!
line vty 0 15
  transport input ssh
  transport preferred none
```

Step 9: Enable Simple Network Management Protocol (SNMP) in order to allow the network infrastructure devices to be managed by a Network Management System (NMS), and then configure SNMPv2c both for a read-only and a read-write community string.

```
snmp-server community [SNMP RO name] RO
snmp-server community [SNMP RW name] RW
```

Step 10: If your network operational support is centralized, you can increase network security by using an access list to limit the networks that can access your device. In this example, only devices on the 10.4.48.0/24 network will be able to access the device via SSH or SNMP.

```
access-list 55 permit 10.4.48.0 0.0.0.255
line vty 0 15
  access-class 55 in
!
snmp-server community [SNMP RO name] RO 55
snmp-server community [SNMP RW name] RW 55
```

The Cisco Catalyst 3650 and 3850 Series Switches have an additional keyword to be added to the access-class, which allows console access from other switch members to not be affected.

```
line vty 0 15
  access-class 55 in vrf-also
```

> **! Caution**
>
> If you configure an access-list on the vty interface, you may lose the ability to use SSH to log in from one device to the next for hop-by-hop troubleshooting.

Step 11: Configure local login and password.

The local login account and password provide basic device access authentication to view platform operation. The enable password secures access to the device configuration mode. By enabling password encryption, you prevent the use of plain text passwords when viewing configuration files. The **aaa new-model** command enables new access control commands and functions, and causes the local username and password on the router to be used in the absence of other AAA statements.

```
username admin secret [user password]
enable secret [enable password]
service password-encryption
aaa new-model
```

By default, https access to the switch uses the enable password for authentication.

Step 12: If you want to reduce operational tasks per device, configure centralized user authentication by using the TACACS+ protocol to authenticate management logins on the infrastructure devices to the authentication, authorization and accounting (AAA) server.

As networks scale in the number of devices to maintain, there is an operational burden to maintain local user accounts on every device. A centralized AAA service reduces operational tasks per device and provides an audit log of user access for security compliance and root cause analysis. When AAA is enabled for access control, all management access to the network infrastructure devices (SSH and HTTPS) is controlled by AAA.

TACACS+ is the primary protocol used to authenticate management logins on the infrastructure devices to the AAA server. A local AAA user database is also defined on each network infrastructure device to provide a fallback authentication source in case the centralized TACACS+ server is unavailable.

```
tacacs server TACACS-SERVER-1
 address ipv4 10.4.48.15
 key [tacacs key]
!
aaa group server tacacs+ TACACS-SERVERS
  server name TACACS-SERVER-1
!
aaa authentication login default group TACACS-SERVERS local
aaa authorization exec default group TACACS-SERVERS local
aaa authorization console
ip http authentication aaa
```

Step 13: Configure a synchronized clock by programming network devices to synchronize to a local NTP server in the network. The local NTP server typically references a more accurate clock feed from an outside source. Configure console messages, logs, and debug output to provide time stamps on output, which allows cross-referencing of events in a network.

```
ntp server 10.4.48.17
ntp update-calendar
!
clock timezone PST -8
clock summer-time PDT recurring
!
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
```

The **ntp update-calendar** command configures the switch to update the hardware clock from the ntp time source periodically. Since not all switches have a hardware clock, this command is not supported by all devices.

Procedure 3	Configure access switch global settings

The access layer devices use VLANs to separate traffic from different devices into the following logical networks:

- The data VLAN provides access to the network for all attached devices other than IP phones.
- The voice VLAN provides access to the network for IP phones.

    Both the data and the voice VLAN are configured on all user-facing interfaces.

- The management VLAN provides in-band access to the network for the switches management interface. The management VLAN is not configured on any user-facing interface and the VLAN interface of the switch is the only member.

Step 1: Configure VLANs on the switch.

Configure the data, voice, and management VLANs on the switch so that connectivity to clients, IP phones, and the in-band management interfaces can be configured. These are the most common examples, and organizations can reduce or increase VLANs for access segmentation as needed to support security systems, IP cameras, Wireless LANs, etc.

```
vlan [data vlan]
  name Data
exit
vlan [voice vlan]
  name Voice
exit
vlan [management vlan]
  name Management
exit
```

> **i | Tech Tip**
>
> If the switch is the only switch at the site and is directly connected to a router or firewall, do not configure a management VLAN. Instead, use the data VLAN for both data and switch management. When you use this configuration, the next step requires that you also configure the in-band management interface VLAN using the data VLAN ID. For example:
>
> ```
> interface vlan [data vlan]
>   description In-band switch management using data VLAN
>   ip address [ip address] [mask]
>   no shutdown
> ```

Step 2: Configure the switch with an IP address so that it can be managed via in-band connectivity.

```
interface vlan [management vlan]
  ip address [ip address] [mask]
  no shutdown
ip default-gateway [default router]
```

Do not use the **ip default-gateway** command on Cisco Catalyst 4500 because it has IP routing enabled by default and this command will not have any effect. Instead use the following command on the Cisco Catalyst 4500.

```
ip route 0.0.0.0 0.0.0.0 [default router]
```

Step 3: Configure DHCP snooping and enable it on the data and voice VLANs. The switch intercepts and safeguards DHCP messages within the VLAN. This ensures that an unauthorized DHCP server cannot serve up addresses to end-user devices.

```
ip dhcp snooping vlan [data vlan]-[voice vlan]
no ip dhcp snooping information option
ip dhcp snooping
```

Step 4: Configure ARP inspection on the data and voice VLANs.

```
ip arp inspection vlan [data vlan],[voice vlan]
```

Step 5: Configure the Bridge Protocol Data Unit (BPDU) Guard global setting to protect PortFast-enabled interfaces.

```
spanning-tree portfast bpduguard default
```

This automatically disables any PortFast-enabled interface if it receives BPDUs protecting against an accidental topology loop which could cause data packet looping and disrupt switch and network operation. You configure the PortFast feature for interfaces in a later step.

If a PortFast-configured interface receives a BPDU, an invalid configuration exists, such as the connection of an unauthorized device. The BPDU guard feature prevents loops by moving a nontrunking interface into an errdisable state when a BPDU is received on an interface when PortFast is enabled.

Figure 11 - Scenario that BPDU Guard protects against

Spanning tree doesn't detect the
loop because PortFast is enabled

User-Installed
Low-End Switch

Access-Layer
Switch

Loop caused by mis-cabling the switch
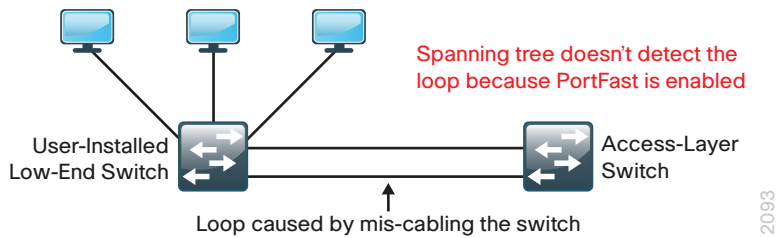
Step 6: Configure the IPv6 First Hop Security global policy for host ports. This global policy is referenced by the access-layer port configuration to intercept and drop IPv6 router advertisements from connected devices. Blocking the advertisements mitigates intentional and unintentional denial-of-service attacks and man-in-the-middle attacks among devices connected to the access layer, and is beneficial regardless of the use of IPv4 or IPv6 routing configuration in the rest of the network.

```
ipv6 nd raguard policy HOST_POLICY
   device-role host
```

ℹ️ **Tech Tip**

Default configurations and default options may not be displayed when you use the **show running-configuration** command. To see the configuration including default options, you can use the **all** option and filter the results as you desire. For example, to see the previous configuration with the default device-role that is applied, enter the following command:

```
show running-config all | section raguard
```

Procedure 4 ▸ Configure client connectivity

To make configuration easier when the same configuration is applied to multiple interfaces on the switch, use the **interface range** command. This command allows you to issue a command once and have it apply to many interfaces at the same time. Since most of the interfaces in the access layer are configured identically, it can save a lot of time. For example, the following command allows you to enter commands on all 24 interfaces (Gig 0/1 to Gig 0/24) simultaneously.

```
interface range GigabitEthernet 0/1-24
```

Step 1: Configure switch interfaces to support clients and IP phones.

The host interface configurations support PCs, phones, or wireless access points. Inline power is available on switches that support 802.3AF/AT for capable devices.

```
interface range [interface type] [port number]-[port number]
   switchport access vlan [data vlan]
   switchport voice vlan [voice vlan]
```

Step 2: Because only end-device connectivity is provided at the access layer, optimize the interface for device connectivity by applying the switchport host command.

```
switchport host
```

This command does three things: it applies switchport access mode, which disables negotiation of trunking, and enables participation as an access port in a VLAN; it enables PortFast, which moves the interface directly into spanning-tree forwarding state, reducing the time it takes for the interface to start forwarding packets; it also disables any channel-group configuration, which is incompatible with an access configuration.

Step 3: Reduce the length of time used for computing interface load average statistics. Reducing the time allows you to display more current and instantaneous data, which is useful for quickly observing interface load value changes during traffic bursts.

```
load-interval 30
```

Step 4: Configure 11 MAC addresses to be active on the interface at one time; additional MAC addresses are considered to be in violation, and their traffic will be dropped.

```
switchport port-security maximum 11
switchport port-security
```

The number of MAC addresses allowed on each interface is specific to the organization. However, the popularity of virtualization applications, IP phones, and passive hubs on the desktop drives the need for the number to be larger than one might guess at first glance. This design uses a number that allows flexibility in the organization while still protecting the network infrastructure.

Step 5: Set an aging time to remove learned MAC addresses from the secured list after 2 minutes of inactivity.

```
switchport port-security aging time 2
switchport port-security aging type inactivity
```

The timeout you choose is an arbitrary time. You may tune the time to fit your environment. Using aggressive timers can impact the switch CPU, so use caution when lowering this from the default value defined on your switch.

Step 6: Configure the restrict option to drop traffic from MAC addresses that are in violation, but do not shut down the port. This configuration ensures that an IP phone can still function on this interface when there is a port security violation.

```
switchport port-security violation restrict
```

Step 7: Configure DHCP snooping and ARP inspection on the interface to process 100 packets per second of traffic on the port.

```
ip arp inspection limit rate 100
ip dhcp snooping limit rate 100
```

The packets per second rate that you choose is an arbitrary rate. You may tune this value to fit your environment.

Step 8: Configure IP Source Guard on the interface. IP Source Guard is a means of preventing IP spoofing.

```
ip verify source
```

If you have a Cisco Catalyst 4500, use the following command instead because Catalyst 4500 requires an additional keyword for the **ip verify source** command.

```
ip verify source vlan dhcp-snooping
```

Step 9: Attach the IPv6 First Hop Security policy to the interface. This configuration uses the IPv6 FHS global policy you previously created for devices connecting to the access layer.

```
ipv6 nd raguard attach-policy HOST_POLICY
```

Step 10: Enable QoS by applying the access edge QoS macro that was defined in the platform configuration procedure. This macro generates a QoS configuration appropriate for the platform.

```
macro apply AccessEdgeQoS
```

All client-facing interfaces allow for an untrusted PC and/or a trusted Cisco IP phone to be connected to the switch and automatically set QoS parameters. When a Cisco IP Phone is connected, trust is extended to the phone and any device that connects to the phone will be considered untrusted and all traffic from that device will be remarked to best-effort or class of service (CoS) 0.

ⓘ Tech Tip

When you apply this macro, device-specific QoS using is applied and a service policy is imposed on the interface. An example policy application to the interface may look like:

```
service-policy input AutoQos-4.0-Cisco-Phone-Input-Policy
service-policy output AutoQos-4.0-Output-Policy
```

In this case, the policy-map called by the service-policy comes preconfigured in the software running on the platform. Detailed examples of the final configurations can be found in the Campus Wired LAN Configuration Files Guide.

Step 11: If the access switch is a Cisco Catalyst 3750-X or 2960-X, increase the buffers for the default queue. This modification of the global QoS settings improves the ability to handle high bandwidth bursty traffic in the dafault queue, by overriding one of the settings previously applied using the AccessEdgeQoS macro. In global configuration mode, add the following command:

```
mls qos queue-set output 1 threshold 3 100 100 100 3200
```

Example: Connected to Distribution Switch

VLAN 100
Data VLAN

VLAN 101
Voice VLAN

IP: 10.4.15.5/25
VLAN 115
Management VLAN
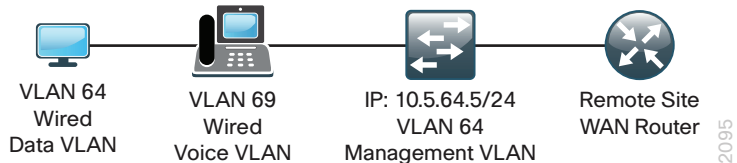
LAN
Distribution
Switch
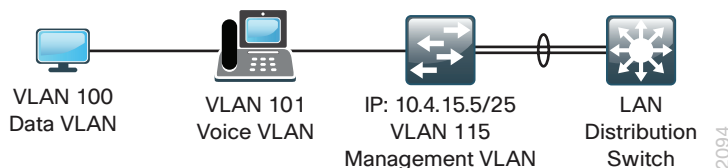
```
vlan 100
  name Data
vlan 101
  name Voice
vlan 115
  name Management
!
interface vlan 115
  description In-band Management
  ip address 10.4.15.5 255.255.255.0
  no shutdown
!
ip default-gateway 10.4.15.1
!
ip dhcp snooping vlan 100,101
```

```
    no ip dhcp snooping information option
    ip dhcp snooping
    ip arp inspection vlan 100,101
    !
    spanning-tree portfast bpduguard default
    !
    ipv6 nd raguard policy HOST_POLICY
      device-role host
    !
    interface range GigabitEthernet 1/0/1-24
      switchport access vlan 100
      switchport voice vlan 101
      switchport host
      switchport port-security maximum 11
      switchport port-security
      switchport port-security aging time 2
      switchport port-security aging type inactivity
      switchport port-security violation restrict
      ip arp inspection limit rate 100
      ip dhcp snooping limit rate 100
      ip verify source
      ipv6 nd raguard attach-policy HOST_POLICY
      macro apply AccessEdgeQoS
    !
    ! Next QoS Command for Cisco Catalyst 3750-X or 2960-X
    mls qos queue-set output 1 threshold 3 100 100 100 3200
```

Example: Connected to WAN Router at a small site

VLAN 64
Wired
Data VLAN

VLAN 69
Wired
Voice VLAN

IP: 10.5.64.5/24
VLAN 64
Management VLAN

Remote Site
WAN Router

2095

```
    vlan 64
      name WiredData
    vlan 69
      name WiredVoice
    !
    interface vlan 64
      description In-band Management to WAN Router
      ip address 10.5.64.5 255.255.255.0
      no shutdown
    !
    ip default-gateway 10.5.64.1
    !
    ip dhcp snooping vlan 64,69
    no ip dhcp snooping information option
```

```
      ip dhcp snooping
      ip arp inspection vlan 64,69
      !
      spanning-tree portfast bpduguard default
      !
      ipv6 nd raguard policy HOST_POLICY
        device-role host
      !
      interface range GigabitEthernet 1/0/1-24
        switchport access vlan 64
        switchport voice vlan 69
        switchport host
        switchport port-security maximum 11
        switchport port-security
        switchport port-security aging time 2
        switchport port-security aging type inactivity
        switchport port-security violation restrict
        ip arp inspection limit rate 100
        ip dhcp snooping limit rate 100
        ip verify source
        ipv6 nd raguard attach-policy HOST_POLICY
        macro apply AccessEdgeQoS
      !
      ! Next QoS Command for Cisco Catalyst 3750-X or 2960-X
      mls qos queue-set output 1 threshold 3 100 100 100 3200
```

Procedure 5	Connect to distribution or WAN router

Access layer devices can be one component of a larger LAN and connect to a distribution switch, or, in the case of a small remote site, might be the only LAN device and connect directly to a WAN device. Unless the access layer device is a single fixed configuration switch connecting to a WAN router, Layer 2 EtherChannels are used to interconnect the devices in the most resilient method possible.

When using EtherChannel, the member interfaces should be on different switches in the stack or different modules in the modular switch for the highest resiliency.

The physical interfaces that are members of a Layer 2 EtherChannel are configured prior to configuring the logical port-channel interface. This allows for minimal configuration because most of the commands entered to a port-channel interface are copied to its members' interfaces and do not require manual replication.

Figure 12 - EtherChannel example
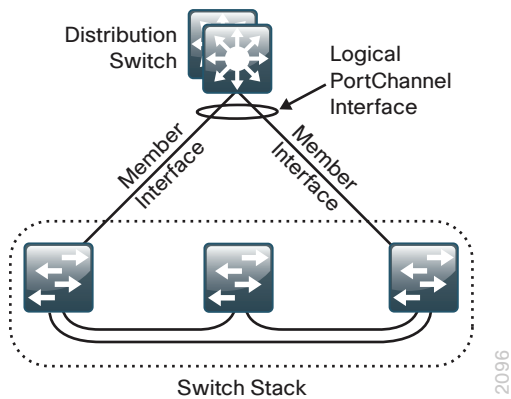
Configure two or more physical interfaces to be members of the EtherChannel. It is recommended that they are added in multiples of two.

This procedure details how to connect any access layer switch (Cisco Catalyst 4500, 3850, 3650, 3750-X, or 2960-X) or to a distribution switch or WAN router. Where there are differences for configuring a specific switch, the differences are called out in the step. If the upstream device is not a distribution switch (such as a remote site connection to a router) use Option 2 with an interface type that is appropriate for the deployment.

Option 1: Configure EtherChannel to distribution switch

Step 1: Configure EtherChannel member interfaces.

This design uses Layer 2 EtherChannels to connect all access layer switches to the distribution layer. When connecting to another switch, use two links or a multiple of two links distributed for maximum resiliency. A configuration with four links is shown in the example. Set Link Aggregation Control Protocol negotiation to active on both sides to ensure a proper EtherChannel is formed. Also, apply the egress QoS macro that was defined in the platform configuration procedure in order to ensure traffic is prioritized appropriately.

Cisco Catalyst 2960-X Series Switches do not require the **switchport** command, and the Cisco Catalyst 4500 does not use the **logging event bundle-status** command.

```
interface [interface type] [port 1]
  description Link to Distribution Layer Port 1
interface [interface type] [port 2]
  description Link to Distribution Layer Port 2
interface [interface type] [port 3]
  description Link to Distribution Layer Port 3
interface [interface type] [port 4]
  description Link to Distribution Layer Port 4
!
interface range [interface type] [port 1], [interface type] [port 2], [interface
type] [port 3], [interface type] [port 4]
  switchport
  channel-protocol lacp
  channel-group [number] mode active
  logging event link-status
  logging event bundle-status
  load-interval 30
  macro apply EgressQoS
```

Step 2: Configure the VLAN trunk interface to the upstream device.

An 802.1Q trunk is used for the connection to this upstream device, which allows the uplink to provide Layer 3 services to all the VLANs defined on the access layer switch. Using a trunk even for a single access VLAN allows for easier VLAN additions in the future. Prune the VLANs allowed on the trunk to only the VLANs that are active on the access switch. Set DHCP Snooping and ARP Inspection to trust.

Because the upstream device is a distribution switch, you use an EtherChannel—the interface type is port-channel and the number must match channel-group configured in Step 1.

The Cisco Catalyst 3750 Series Switch requires the **switchport trunk encapsulation dot1q** command.

```
interface [interface type] [number]
   description EtherChannel Link to Distribution Layer
   switchport trunk allowed vlan [data vlan],[voice vlan],
      [management vlan]
   switchport mode trunk
   ip arp inspection trust
   ip dhcp snooping trust
   logging event link-status
   logging event trunk-status
   no shutdown
   load-interval 30
exit
```

If the interface type is not a port-channel, you must configure an additional command **macro apply EgressQoS** on the interface.

In the next step, you mitigate VLAN hopping on the trunk for switch-to-switch connections.

There is a remote possibility that an attacker can create a double 802.1Q encapsulated packet. If the attacker has specific knowledge of the 802.1Q native VLAN, a packet could be crafted that when processed, the first or outermost tag is removed when the packet is switched onto the untagged native VLAN. When the packet reaches the target switch, the inner or second tag is then processed and the potentially malicious packet is switched to the target VLAN.

Figure 13 - VLAN hopping attack
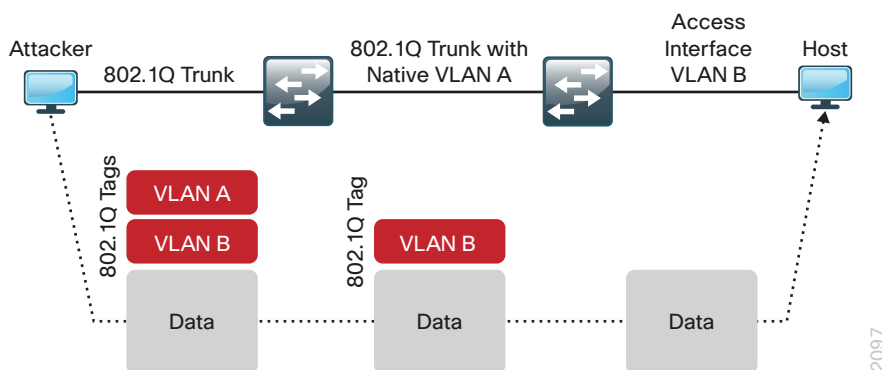
At first glance, this appears to be a serious risk. However, the traffic in this attack scenario is in a single direction and no return traffic can be switched by this mechanism. Additionally, this attack cannot work unless the attacker knows the native VLAN ID.

Step 3: Configure an unused VLAN on all switch-to-switch 802.1Q trunk links from access layer to distribution layer. This configuration mitigates the remote risk of a VLAN hopping attack. Choosing an arbitrary, non-default, unused VLAN assignment for the native VLAN reduces the possibility that a double 802.1Q-tagged packet can hop VLANs. If you are running the recommended EtherChannel uplink to the LAN access layer switch, configure the **switchport trunk native vlan** on the port-channel interface.

```
vlan 999
  name AntiVLANhopping
exit
!
interface [port-channel] [number]
  switchport trunk native vlan 999
```

Step 4: After leaving configuration mode, save the running configuration that you have entered so it will be used as the startup configuration file when your switch is reloaded or power-cycled.

```
copy running-config startup-config
```

Step 5: If you have configured your access-layer Cisco Catalyst 3750-X switch stack for an EtherChannel link to the distribution layer switch, reload your switch stack now to ensure proper operation of EtherChannel. A single reload of a newly configured switch may be necessary to ensure that EtherChannel operates with other features configured on the switch stack.

```
reload
```

Option 2: Configure Connection to WAN router

If your access layer switch is a single fixed configuration switch connecting to a single remote-site router without using EtherChannel, you can skip Step 1.

Step 1: Configure EtherChannel member interfaces.

When connecting to a network infrastructure device that does not support LACP, like a router, set the **channel-group mode** to be forced on.

Also, apply the egress QoS macro that was defined in the platform configuration procedure to ensure traffic is prioritized appropriately.

Cisco Catalyst 2960S and 2960-X do not require the **switchport** command, and the Cisco Catalyst 4500 does not use the **logging event bundle-status** command.

```
interface [interface type] [port 1]
  description Link to Router Port 1
interface [interface type] [port 2]
  description Link to Router Port 2
!
interface range [interface type] [port 1], [interface type] [port 2]
  switchport
  channel-group [number] mode on
  logging event link-status
  logging event trunk-status
  logging event bundle-status
  load-interval 30
  macro apply EgressQoS
```

Step 2: Configure the VLAN trunk interface to the upstream device.

An 802.1Q trunk is used for the connection to this upstream device, which allows the router to provide Layer 3 services to all the VLANs defined on the access layer switch. Prune the VLANs allowed on the trunk to only the VLANs that are active on the access switch. Set DHCP snooping and ARP Inspection to trust.

When using EtherChannel, the interface type is port-channel, and the number must match channel-group configured in Step 1 in Option 2 of this procedure. For deployments other than EtherChannel, interface type is one appropriate for your deployment.

The Cisco Catalyst Series Switch 3750 requires the **switchport trunk encapsulation dot1q** command.

```
interface [interface type] [number]
  description Trunk Interface to Router
  switchport trunk allowed vlan [data vlan],[voice vlan],[optional transit vlan]
  switchport mode trunk
  ip arp inspection trust
  ip dhcp snooping trust
  spanning-tree portfast trunk
  logging event link-status
  logging event trunk-status
  load-interval 30
  no shutdown
```

If the interface type is not a port-channel, you must configure additional commands **switchport** and **macro apply EgressQoS** on the interface.

Step 3: Save the running configuration that you have entered so it will be used as the startup configuration file when your switch is reloaded or power-cycled.

```
copy running-config startup-config
```

Step 4: If you have configured your access layer Cisco Catalyst 3750-X switch stack for EtherChannel to the WAN router, reload your switch stack now to ensure proper operation of EtherChannel. A single reload of a newly configured switch is necessary to ensure that EtherChannel operates with other features configured on the switch stack.

```
reload
```

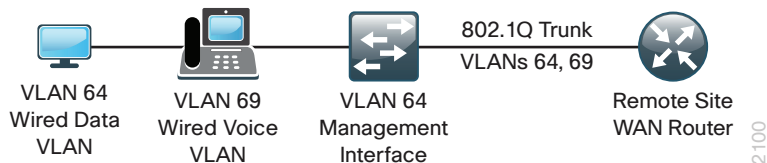**Example: Procedure 5, Option 1**
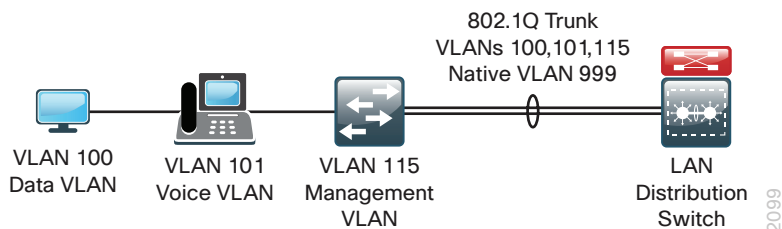

```
vlan 999
  name AntiVLANhopping
!
interface GigabitEthernet 1/0/25
  description Link to Distribution Layer port 1
interface GigabitEthernet 3/0/25
```

```
      description Link to Distribution Layer port 2
interface GigabitEthernet 1/0/26
      description Link to Distribution Layer port 3
interface GigabitEthernet 3/0/26
      description Link to Distribution Layer port 4
!
interface range GigabitEthernet 1/0/25, GigabitEthernet 3/0/25, GigabitEthernet
1/0/26, GigabitEthernet 3/0/26
      logging event link-status
      logging event bundle-status
      channel-protocol lacp
      channel-group 1 mode active
      load-interval 30
      macro apply EgressQoS
!
interface Port-channel 1
      description Etherchannel to Distribution Layer
      switchport trunk encapsulation dot1q
      switchport trunk native vlan 999
      switchport trunk allowed vlan 100,101,115
      switchport mode trunk
      logging event trunk-status
      ip arp inspection trust
      ip dhcp snooping trust
      load-interval 30
      no shutdown
```

Example: Procedure 5, Option 2 without EtherChannel

VLAN 64
Wired Data
VLAN

VLAN 69
Wired Voice
VLAN

VLAN 64
Management
Interface

802.1Q Trunk
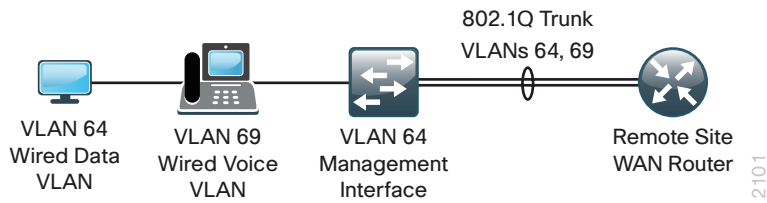VLANs 64, 69

Remote Site
WAN Router

```
interface GigabitEthernet 1/0/24
      description Link to WAN Router
      switchport trunk encapsulation dot1q
      switchport trunk allowed vlan 64,69,99
      switchport mode trunk
      ip arp inspection trust
      ip dhcp snooping trust
      spanning-tree portfast trunk
      macro apply EgressQoS
      no shutdown
```

Example: Procedure 5, Option 2 with EtherChannel


```
interface GigabitEthernet 1/0/25
  description Link to WAN Router Port 1
interface GigabitEthernet 3/0/25
  description Link to WAN Router Port 2
!
interface range GigabitEthernet 1/0/25, GigabitEthernet 3/0/25
  logging event link-status
  logging event trunk-status
  logging event bundle-status
  load-interval 30
  channel-group 1 mode on
  macro apply EgressQoS
!
interface Port-channel 1
  description EtherChannel to WAN Router
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 64,69
  switchport mode trunk
  ip arp inspection trust
  ip dhcp snooping trust
  spanning-tree portfast trunk
  load-interval 30
  no shutdown
```
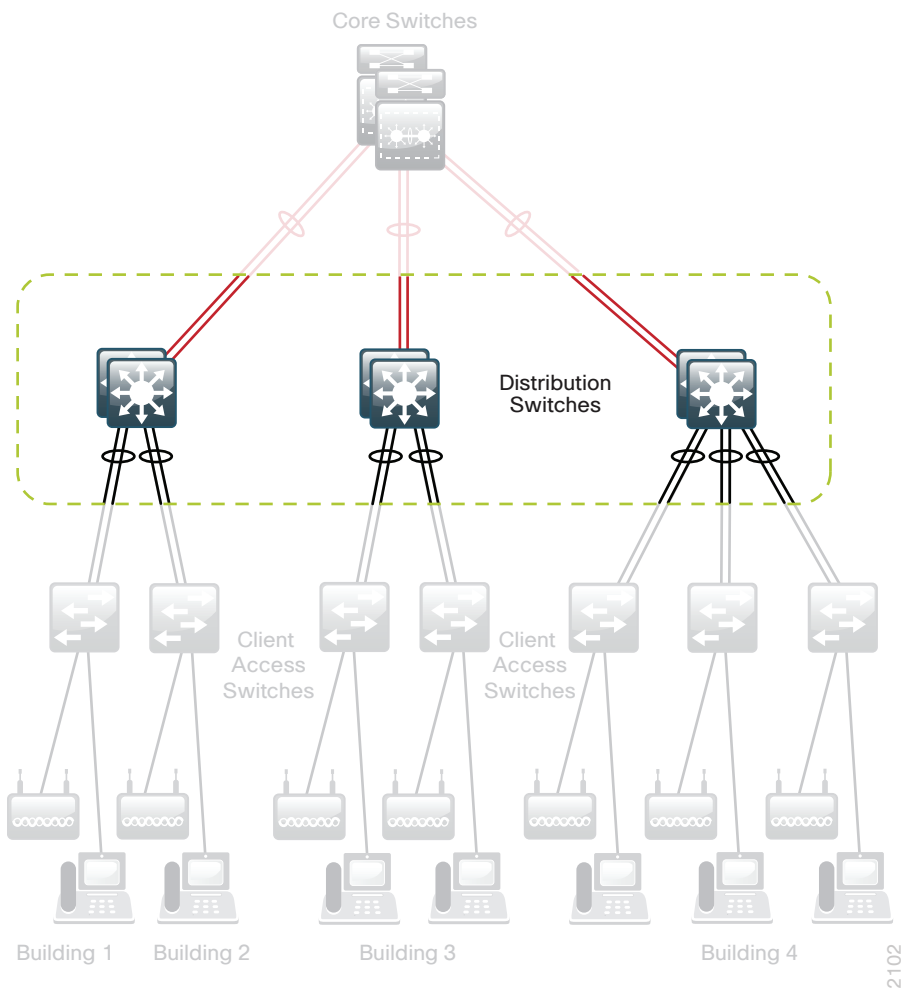
Distribution Layer

Design Overview

The primary function of the distribution layer is to aggregate access layer switches in a given building or campus. The distribution layer provides a boundary between the Layer 2 domain of the access layer and the Layer 3 domain that provides a path to the rest of the network. This boundary provides two key functions for the LAN. On the Layer 2 side the distribution layer creates a boundary for Spanning Tree Protocol limiting propagation of Layer 2 faults. On the Layer 3 side the distribution layer provides a logical point to summarize IP routing information before it enters the network and reduce IP route tables for easier troubleshooting and faster recovery from failures.

The LAN distribution layer uses a simplified distribution layer design that is easier to operate and troubleshoot than the traditional and routed access designs.

Figure 14 – Distribution layer overview

Traditional Distribution Layer Design

Traditional LAN designs use a multitier approach with Layer 2 from the access layer to the distribution layer, where the Layer 3 boundary exists. The connectivity from the access layer to the distribution layer can result in either a loop-free or looped design.

In the traditional network design, the distribution layer has two standalone switches for resiliency. It is recommended that you restrict a Layer 2 VLAN to a single wiring closet or access uplink pair to reduce or eliminate topology loops that Spanning Tree Protocol must block and that are a common point of failure in LANs. Restricting a VLAN to a single switch provides a loop-free design, but it does limit network flexibility.

To create a resilient IP gateway for VLANs in the traditional design, you must use first-hop redundancy protocols, which provide hosts with a consistent MAC address and gateway IP for a VLAN. Hot Standby Routing Protocol (HSRP) and Virtual Router Redundancy Protocol (VRRP) are the most common gateway redundancy protocols, but they only allow hosts to send data out one of the access uplinks to the distribution layer, and require additional configuration for each aggregation switch to allow you to distribute VLANs across uplinks. Gateway Load Balancing Protocol (GLBP) does provide greater uplink utilization for traffic exiting the access layer by balancing load from hosts across multiple uplinks, but you can only use it in a non-looped topology.

All of these redundancy protocols require that you fine-tune the default timer settings to allow for subsecond network convergence, which can impact switch CPU resources.

Some organizations require the same Layer 2 VLAN be extended to multiple access layer closets to accommodate an application or service. The looped design causes spanning tree to block links, which reduces the bandwidth from the rest of the network and can cause slower network convergence.

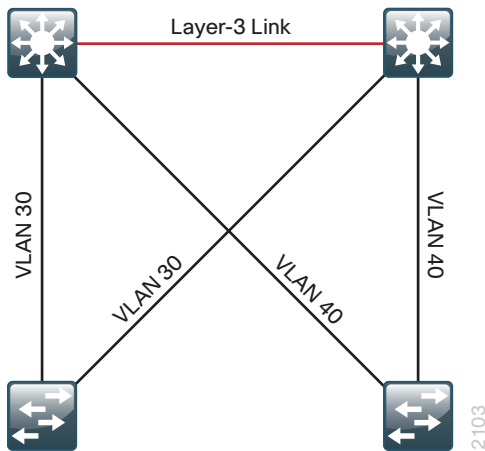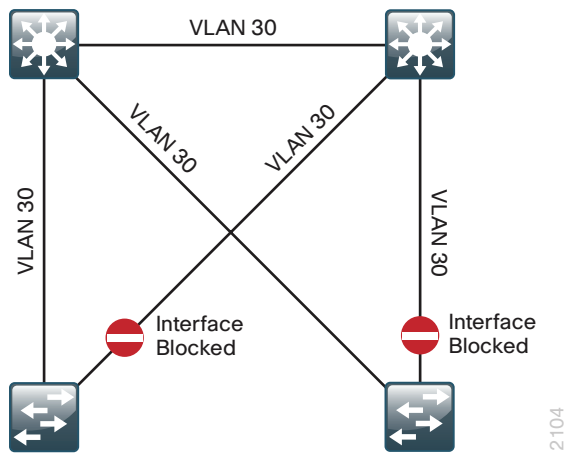*Figure 15 - Traditional loop-free design with a VLAN per access switch*

Figure 16 - Traditional looped design with VLANs spanning access switches

## Routed Access Distribution Layer Design

In another approach to access and distribution layer design, you can use Layer 3 all the way to the access layer. The benefits of this design are that you eliminate spanning tree loops and reduce protocols because the IP gateway is now the access switch. Because there are no spanning-tree blocking links, you can use both uplinks to the access layer and increase effective bandwidth available to the users.

The challenge with the routed access layer design is that the Layer 2 domains are confined to a single access closet, which limits flexibility for applications that require Layer 2 connectivity that extends across multiple access closets.

Simplified Distribution Layer Design

The distribution layer design chosen for the campus design uses multiple physical switches that act as a single logical switch or a single, highly-redundant physical switch. One advantage of this design is that spanning tree dependence is minimized, and all uplinks from the access layer to the distribution are active and passing traffic. Even in the distributed VLAN design, spanning tree blocked links due to looped topologies are eliminated. You reduce dependence on spanning tree by using EtherChannel to the access layer with dual-homed uplinks. This is a key characteristic of this design and you can load balance up to eight links if needed for additional bandwidth.

Figure 17 - Simplified design with a VLAN per access switch
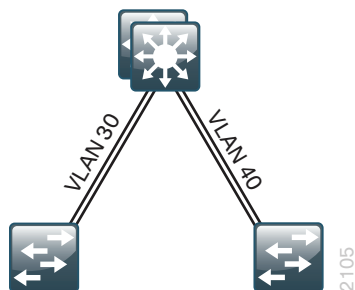
Figure 18 - Simplified design with VLANs spanning access switches
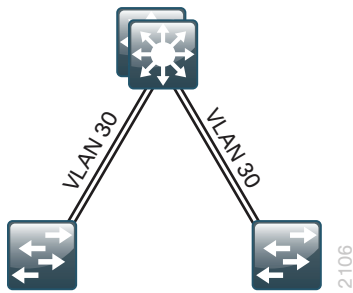
EtherChannel is a logical interface that can use a control plane protocol to manage the physical members of the bundle. It is better to run a channel protocol instead of using forced-on mode because a channel protocol performs consistency checks for interfaces programmed to be in the channel and provides protection to the system from inconsistent configurations.

Cisco Catalyst switches provide both Port Aggregation Protocol (PAgP), which is a widely deployed Cisco designed protocol, and Link Aggregation Protocol (LACP) based on IEEE 802.3ad. This design uses LACP for EtherChannel because it is the only protocol supported in a Cisco Catalyst 3750 cross-stack configuration and can be used in all configurations in this design.

There are several other advantages to the simplified distribution layer design. You no longer need IP gateway redundancy protocols like HSRP, VRRP, and GLBP because the default IP gateway is now on a single logical interface and resiliency is provided by the distribution layer switch or switches. Also, the network will converge faster now that it is not depending on spanning tree to unblock links when a failure occurs because EtherChannel provides fast subsecond failover between links in an uplink bundle.

The topology of the network from the distribution layer to the access layer is logically a hub-and-spoke topology, which reduces complexity of design and troubleshooting. The hub-and-spoke topology design provides a more efficient operation for IP Multicast in the distribution layer because there is now a single logical designated router to forward IP Multicast packets to a given VLAN in the access layer.

Finally, by using the single logical distribution layer design, there are fewer boxes to manage, which reduces the amount of time spent on ongoing provisioning and maintenance.

Distribution Layer Roles

Much emphasis has been placed on the distribution layer as the access layer aggregation point because this is the most common role. The distribution layer serves other roles in LAN designs.

In many smaller locations, the WAN head end and Internet edge terminate at the headquarters location, along with a server farm or small data center and the LAN access for user connectivity. In these situations a single distribution layer or *collapsed core* design may be appropriate to allow the network to stay within budget limits while serving a smaller LAN access environment. Although the port density and configuration complexity may not be an issue, operational complexity of supporting many functions on one device must be monitored as the organization grows.

Figure 19 - Two-tier collapsed LAN core design
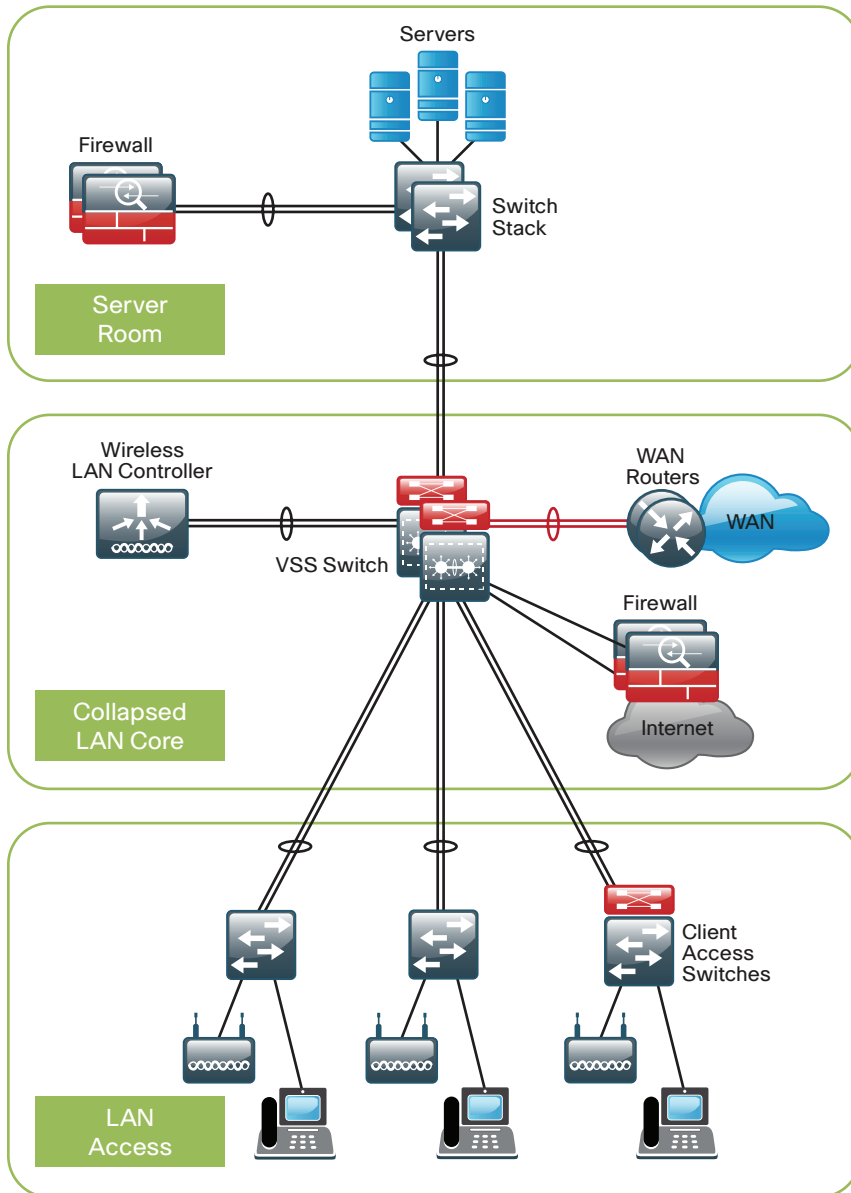
In larger LAN locations where the access layer density along with the number of network-service devices and WAN routers exceeds platform density or operational complexity additional distribution layer modules can break up the design.

The addition of a separate *services* distribution layer provides:

- Modular growth for high densities of WAN headend routers and WAN services like WAAS appliance.
- Wireless LAN controller termination in a central location for larger campus populations.
- Fault domains separate from the LAN access for a more resilient overall network.
- IP address summarization from WAN or Internet edge toward the core of the network.

Figure 20 - Network services distribution layer
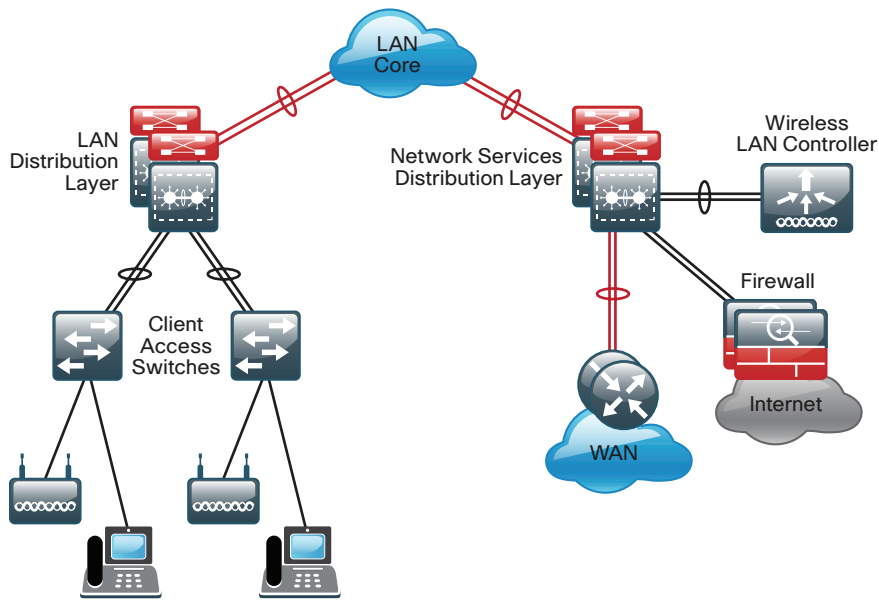
Whether the distribution layer role in your network design is serving as purely LAN access aggregation, a collapsed core, or network-services aggregation, the distribution layer configuration provides the processes and procedures to prepare this layer of the LAN for your application.

Distribution Layer Platforms

You can use multiple platforms to deploy the simplified distribution layer design. Physically, the distribution layer can be a Cisco Catalyst 6500 Virtual Switching System (VSS), a Cisco Catalyst 4500-X VSS, a highly available Cisco Catalyst 4507R+E switch pair in VSS mode, or a stack of Cisco Catalyst 3850 or 3750-X switches. It is important to note that although each switch has different physical characteristics, each appears to the rest of the network as a single node and provides a fully resilient design.

Cisco Catalyst 6807-XL and 6500-E VSS

The Cisco Catalyst 6500-E and 6807-XL chassis with the Supervisor Engine 2T are the premier distribution layer platforms.

- Cisco Catalyst 6807-XL VSS uses two physical chassis with Cisco Catalyst 6500 Supervisor Engine 2T, which offers a per slot switching capacity of 220 Gbps in the Cisco Catalyst 6807-XL chassis, and delivers hardware-enabled scalability and features. The performance is increased over the same Supervisor Engine in the 6500-E chassis enabling the system to provide 40-Gigabit Ethernet connectivity and growth capability for future 100-gigabit Ethernet connectivity.

- Cisco Catalyst 6500 VSS uses Cisco Catalyst 6500 Supervisor Engine 2T, which offers per slot switching capacity of 80 Gbps in the Cisco Catalyst 6500-E Series chassis and delivers hardware-enabled scalability and features. This level of performance enables the system to provide 40-gigabit Ethernet uplinks for core layer connectivity.

- Adding an additional Cisco Catalyst 6500 Supervisor Engine 2T to each chassis in the VSS pair for a total of four supervisors creates a Quad-Supervisor SSO (VS4O) configuration, offering the ability to have an in-chassis standby supervisor capability. The in-chassis standby enables Enhanced Fast Software Upgrades (eFSU) for minimal downtime during software upgrades, along with the ability to recover from a degraded state of performance upon loss of a supervisor, without human intervention.

- Cisco 6500 Supervisor Engine 2T supports the line cards with Distributed Forwarding Card 4-E (DFC4-E), including the WS-X6816-10G, WS-X6908-10G, and WS-X6904-40G-2T, which provide enhanced hardware capabilities. The WS-X6908-10G provides eight 10-Gbps Ethernet ports with 1:1 oversubscription. The WS-X6904-40G-2T provides up to four 40-Gbps Ethernet ports or up to sixteen 10-Gbps Ethernet ports using modular adapters and can be programmed to run in 2:1 or 1:1 oversubscription mode.

- The Supervisor Engine 2T supports DFC4-A based line cards, including the WS-X6824 and WS-X6848, to provide gigabit Ethernet ports. The WS-X6724 and WS-X6748 gigabit Ethernet cards are also supported when installed with CFC or DFC4-A modules.

- The Supervisor Engine 2T-based switch enhances support for Cisco TrustSec (CTS) by providing MacSec encryption and role-based access control lists (RBACL), and delivers improved control plane policing to address denial-of-service attacks.

- VSS effectively allows the merging of two physical chassis into a logical entity that can be operated as a single device. This configuration provides redundant chassis, supervisors, line cards, and power supplies and can provide the highest density of the product options for Gigabit Ethernet, 10 Gigabit Ethernet, and 40-Gigabit EtherChannel uplinks using Cisco Multi-chassis EtherChannel (MEC).

- Provides Stateful Switch-Over (SSO) to synchronize infrastructure and forwarding state between chassis, along with Non-Stop Forwarding (NSF) for graceful-restart of L3 routing protocols, in the event of a chassis failure. Also allows Enhanced Fast Software Upgrades (EFSU) with In-Service Software Upgrades (ISSU) for minimizing downtime for system upgrades.

- The Cisco Catalyst 6500-E and 6807-XL chassis with the Supervisor Engine 2T are the premier distribution layer platforms. They allow for high density aggregation of wiring closets connected with Gigabit Ethernet and 10-Gigabit Ethernet, while providing an advanced feature set and the highest resiliency of the available platforms.

Cisco Catalyst 6880-X VSS

- Cisco Catalyst 6880-X VSS uses Cisco Catalyst 6880-X Series extensible fixed aggregation switch, with the Cisco Catalyst 6500 feature set in a small form factor.

- The Cisco Catalyst 6800-X Series is a resilient chassis offering N+1 redundant fans and 1+1 resilient power supplies, along with the capability for two switches to be paired into a resilient single logical Virtual Switching System.

- The base chassis comes with 16 SFP+ ports supporting 1-Gigabit Ethernet and 10-Gigabit Ethernet services, and an additional four slots available for extensible port cards to allow for future growth. For example, adding C6880-X-16P10G 16-port 10-Gigabit Ethernet cards enables up to a total of 80-ports of 10-Gigabit Ethernet in the chassis. The chassis is also designed with the capability to support 40-Gigabit Ethernet and 100-Gigabit Ethernet, with future port cards, and a backplane capable of delivering 220-Gbps per slot.

- Provides Stateful Switch-Over (SSO) to synchronize infrastructure and forwarding state between chassis, along with Non-Stop Forwarding (NSF) for graceful-restart of L3 routing protocols, in the event of a chassis failure. Also allows Enhanced Fast Software Upgrades (EFSU) with In-Service Software Upgrades (ISSU) for minimizing downtime for system upgrades.

- Cisco Catalyst 6880-X is the premier fixed distribution layer platform in this design. It allows for medium density aggregation of Gigabit Ethernet and 10 Gigabit Ethernet connected wiring closets, while providing an advanced feature set and the highest resiliency of the available platforms.

Cisco Catalyst 4500-X VSS

- Cisco Catalyst 4500-X Series switch family includes 32-port 10-Gigabit Ethernet and 16-port 10-Gigabit Ethernet switches both with a slot for adding an optional 8-port 10-Gigabit Ethernet module.

- Cisco Catalyst 4500-X Series has resiliency capabilities including redundant hot swappable fans and power supplies, in addition to the capability for two switches to be paired into a resilient single logical Virtual Switching System.

- Provides Stateful Switchover to synchronize infrastructure and forwarding state between chassis, along with Nonstop Forwarding for graceful-restart of L3 routing protocols in the event of a chassis failure, which also allows In-Service Software Upgrade (ISSU) functionality for the system.

- Cisco Catalyst 4500-X Series can be used at locations where there is a smaller number of Gigabit Ethernet or 10 Gigabit Ethernet connected wiring closets that need to be aggregated.

Cisco Catalyst 4507R+E VSS

- The Cisco Catalyst 4507R+E switch supports redundant supervisors, line cards, and power supplies. In this design, two 4507R+E chassis are paired into a resilient single logical Virtual Switching System distribution layer platform. The Cisco Catalyst 4500 Supervisor Engine 7-E has the ability to provide a medium density of Gigabit Ethernet and even 10-Gigabit Ethernet EtherChannel links to the access layer.

- Provides Stateful Switchover between supervisors in the VSS pair, to synchronize infrastructure and forwarding state between chassis, along with Nonstop Forwarding for graceful-restart of L3 routing protocols in the event of a chassis failure, which also allows In-Service Software Upgrade (ISSU) functionality for the system.

Cisco Catalyst 3850 Stack

- Cisco Catalyst 3850 Series Switches are fixed-port Ethernet switches, with SFP-based versions for fiber connectivity recommended in the distribution layer. For resiliency, you configure the Cisco Catalyst 3850 Series Switches as a stack, and then configure and manage the stack as a single unit.

- Cisco Catalyst 3850 Series Switches have built-in stacking capability and stack together using StackWise-480 mode with up to nine switches in single stack-ring.

- Cisco StackPower technology increases system-level resiliency during catastrophic power failure on a stack-member switch. Cisco StackPower enables power redundancy across a group of four Cisco Catalyst 3850 Series Switches within same stack. This allows the flexible arrangement of power supplies in the stack and enables a zero-footprint redundant power supply deployment and intelligent load shedding.

- Cisco 3850 Series Switches have modular uplinks that can be configured as Gigabit Ethernet or 10-Gigabit Ethernet.

- Cisco Catalyst 3850 Series supports Stateful Switchover, which allows a switch in the active role in a stack to rapidly switchover to a switch in the standby role with minimum disruption to the network.

- With appropriate licenses, the Cisco Catalyst 3850 Series hardware supports wireless LAN controller functionality in order to support a unified access policy for converged wired and wireless designs with up to 40-Gbps of wireless capacity per switch.

Cisco Catalyst 3750-X Stack

- Cisco Catalyst 3750-X is configured as a single unit, but has independent load-sharing power supplies and processor for each switch in the StackWise Plus stack. The LAN architecture uses a pair of stacked 3750X-12S-E switches that provide Layer 2 and Layer 3 switching. The switches use Small Form-Factor Pluggable (SFP) transceivers for a port-by-port option of copper or fiber optic Gigabit Ethernet EtherChannel uplinks to access closets.

- Cisco StackWise Plus enables up to nine Catalyst 3750-X switches to be stacked together using a 64-Gbps stack interconnect with rapid failure recovery.

- Cisco StackPower shares power across the Cisco Catalyst 3750-X switch stack. This allows the flexible arrangement of power supplies in the stack, and enables a zero-footprint redundant power supply deployment and intelligent load shedding.

- Cisco 3750-X Series have modular uplinks for connectivity to the core layer at Gigabit or 10-Gigabit Ethernet rates, and support upgrading the Cisco IOS feature set and enhanced enterprise capabilities like TrustSec and Flexible NetFlow in order to ensure that the switch functionality grows as the organization grows.

- The Cisco Catalyst 3750-X switch stack can be used at locations where there is only a small number of gigabit connected wiring closets that need to be aggregated.

Deployment Details

The single, logical, resilient, distribution layer deployed with this validated design simplifies the distribution switch configuration over traditional dual system designs.

Configuring the Distribution Layer

1. Configure the platform
2. Configure LAN switch universal settings
3. Configure distribution global settings
4. Configure IP unicast routing
5. Configure IP Multicast routing
6. Configure IP Multicast RP
7. Connect to access layer
8. Connect to LAN core or WAN router

Procedure 1 Configure the platform

Some platforms require a one-time initial configuration prior to configuring the features and services of the switch. If you do not have a platform listed for an option, skip that option and its steps.

Option 1: Configure Cisco Catalyst 6807-XL, 6500-E, or 6880-X Virtual Switching System

Cisco Catalyst 6807-XL and 6500-E Virtual Switching System merges two physical 6807-XL or 6500-E switches together as a single logical switch using a single or optionally dual Cisco Supervisor Engine 2T modules in each physical switch. Cisco Catalyst 6880-X Virtual Switching System merges two physical 6880-X switches together as a single logical switch. A Supervisor module in the 6500-E or built-in Supervisor hardware in the 6880-X acts as the active control plane for both chassis by controlling protocols such as Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF) Spanning Tree, CDP, and so forth, while supervisor hardware in each chassis actively switches packets.

The following configuration example shows you how to convert two standalone Cisco Catalyst 6807-XL, 6500-E, or 6880-X switches to a Virtual Switching System (VSS). If you are migrating your switches from an existing in-service dual chassis role to a VSS system, go to www.cisco.com and search on "Migrate Standalone Cisco Catalyst 6500 Switch to Cisco Catalyst 6500 Virtual Switching System" for information that describes how to do this migration. For an in-depth VSS configuration guide and configuration options, see the *Campus 3.0 Virtual Switching System Design Guide,* available on the Cisco Validated Designs tab at www.cisco.com/go/cvd/campus.

When you set up the Cisco Catalyst 6807-XL, 6500-E, or 6880-X Virtual Switching System, connect two 10-Gigabit or 40-Gigabit Ethernet links between the chassis to provide the Virtual Switch Link (VSL). Always use at least two links.

You can use up to eight links as members of the VSL, and the links should be distributed across Supervisor Engines or line cards for resiliency. This design uses the two 10-Gigabit Ethernet interfaces on each Supervisor Engine for the Cisco Catalyst 6807-XL or 6500-E, and two 10 Gigabit Ethernet ports on the Catalyst 6880-X. You connect the VSL interfaces together before you configure the VSS.

To aid in understanding the connections between switches supporting the VSS configuration, the following tables can be used. The optional Fast-Hello connection is used in this design, but can be replaced or augmented by functionality available when implementing Ehanced PAgP, which is not covered in this release. The PortChannel interface numbers are arbitrary and should be adapted to best suit your deployment. The chosen PortChannel numbers reflect the highest values supported across all VSS platforms validated in this release.

Table 3 - Example VSS connections for Cisco Catalyst 6807-XL chassis pair with two Cisco Supervisor Engine 2T

VSS connection	VSS Switch 1 Port (PortChannel)	VSS Switch 2 Port (PortChannel)
10-Gbps, VSL 1	Ten3/4 (Po63)	Ten3/4 (Po64)
10-Gbps, VSL 2	Ten3/5 (Po63)	Ten3/5 (Po64)
1-Gbps, Fast-Hello	Gig7/24	Gig7/24

Table 4 - Example VSS connections for Cisco Catalyst 6509-E chassis pair with two Cisco Supervisor Engine 2T

VSS connection	VSS Switch 1 Port (PortChannel)	VSS Switch 2 Port (PortChannel)
10-Gbps, VSL 1	Ten5/4 (Po63)	Ten5/4 (Po64)
10-Gbps, VSL 2	Ten5/5 (Po63)	Ten5/5 (Po64)
1-Gbps, Fast-Hello	Gig9/24	Gig9/24

For additional resiliency and to allow support for an Enhanced Fast Software Upgrade (eFSU), a VSS quad-supervisor Stateful Switchover (VS4O) setup can be created. The four Supervisors must have fully meshed connectivity to support minimal traffic disruption during software upgrades as shown in the following example.

Table 5 - Example VSS connections for the Cisco 6807-XL chassis pair with four Supervisor Engine 2T, meshed VSL connections

VSS connection	VSS Switch 1 Port (PortChannel)	VSS Switch 2 Port (PortChannel)
10-Gbps, VSL 1	Ten3/4 (Po63)	Ten3/4 (Po64)
10-Gbps, VSL 2	Ten3/5 (Po63)	Ten4/4 (Po64)
10-Gbps, VSL 3	Ten4/4 (Po63)	Ten3/5 (Po64)
10-Gbps, VSL 4	Ten4/5 (Po63)	Ten4/5 (Po64)
1-Gbps, Fast-Hello	Gig7/48	Gig7/48

Table 6 - Example VSS connections for the Cisco 6509-E chassis pair with four Supervisor Engine 2T, meshed VSL connections

VSS connection	VSS Switch 1 Port (PortChannel)	VSS Switch 2 Port (PortChannel)
10-Gbps, VSL 1	Ten5/4 (Po63)	Ten5/4 (Po64)
10-Gbps, VSL 2	Ten5/5 (Po63)	Ten6/4 (Po64)
10-Gbps, VSL 3	Ten6/4 (Po63)	Ten5/5 (Po64)
10-Gbps, VSL 4	Ten6/5 (Po63)	Ten6/5 (Po64)
1-Gbps, Fast-Hello	Gig9/48	Gig9/48

Table 7 - Example VSS connections, connecting Cisco Catalyst 6880-X chassis pair without Extensible Port Cards installed

VSS connection	VSS Switch 1 Port (PortChannel)	VSS Switch 2 Port (PortChannel)
10-Gbps, VSL 1	Ten5/5 (Po63)	Ten5/5 (Po64)
10-Gbps, VSL 2	Ten5/13 (Po63)	Ten5/13 (Po64)
1-Gbps, Fast-Hello	Ten5/14 (1000-Mbps)	Ten5/14 (1000-Mbps)

Table 8 - Example VSS connections, connecting Cisco Catalyst 6880-X chassis pair with Extensible Port Cards installed

VSS connection	VSS Switch 1 Port (PortChannel)	VSS Switch 2 Port (PortChannel)
10-Gbps, VSL 1	Ten5/5 (Po63)	Ten5/5 (Po64)
10-Gbps, VSL 2	Ten1/5 (Po63)	Ten1/5 (Po64)
1-Gbps, Fast-Hello	Ten5/14 (1000-Mbps)	Ten5/14 (1000-Mbps)

---

i | **Tech Tip**

When you use a Cisco Catalyst 6880-X chassis pair with no Extensible Port Cards installed, select interfaces so that each VSL connection is in a different port group. The ports selected in the table include a port from the set of ports which is active for that port group regardless of the choice of performance mode or oversubscription mode. If an Extensible Port Card is added, the VSL connections should be distributed across available cards for additional resiliency, with one port of the VSL connection using a base port for the fastest VSS booting time. Similarly, when selecting ports for the VSL on a 6500-E chassis, at least one Supervisor Engine 2T 10-Gigabit Ethernet port should be used for the fastest VSS booting time.

Step 1: If you are configuring a VSS quad-supervisor VS4O system, ensure that all four Supervisor Engines are available with the same version of code installed and set to boot for each of them. To make configuration easier by limiting the possible active console connections, leave only one Supervisor Engine fully inserted in each chassis.

Step 2: If you are using a Cisco Catalyst 6807-XL or 6500-E chassis with the Cisco Catalyst 6900 Series 40-Gigabit Ethernet Interface Module with FourX adapters to convert CFP ports into four 10-Gigabit Ethernet ports, configure the switch to enable the line card to use 10-Gigabit Ethernet functionality for the associated port-group. If you configure this after converting to VSS, the **switch** parameter is required. The port-group can be **1**, **2**, or **all**.

```
D6500-VSS(config)# hw-module switch [switch] slot [slot] operation-mode port-
group [port-group] TenGigabitEthernet
```

Example—command used before VSS conversion

```
D6500-VSS(config)# hw-module slot 1 operation-mode port-group 2
TenGigabitEthernet
```

Example—command used after VSS conversion

```
D6500-VSS(config)# hw-module switch 1 slot 1 operation-mode port-group 2
TenGigabitEthernet
```

At the confirmation prompt, press enter to confirm the change. The module resets and the port configuration is cleared.

Step 3: Convert standalone Cisco switches to VSS.

Configure a temporary hostname on each switch so you can keep track of your configuration steps. In a later step after the conversion is complete, you apply a replacement hostname to the merged VSS configuration.

On the standalone switch #1:

```
Router#config t
Router#(config)#hostname VSS-Sw1
```

On the standalone switch #2:

```
Router#config t
Router#(config)#hostname VSS-Sw2
```

To form a VSS pair, each switch in the pair must have a matching domain ID assigned. To support the interconnection of multiple VSS pairs, the domain ID selected for the pair should be unique. In this example, the domain number is 100. Each switch is also given a unique identifier within the domain, switch 1 or switch 2.

Figure 21 - VSS domain
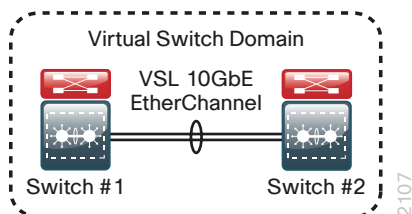
On the standalone switch #1:

```
VSS-Sw1(config)#switch virtual domain 100
VSS-Sw1(config-vs-domain)# switch 1
```

On the standalone switch #2:

```
VSS-Sw2(config)#switch virtual domain 100
VSS-Sw2(config-vs-domain)# switch 2
```

Step 4: Configure the Virtual Switch Link (VSL).

The VSL is a critical component of the Virtual Switching System. For each physical switch you must select a unique port-channel number identifying the same VSL. This allows the switch to maintain a separate identity for the interfaces used when making traffic forwarding decisions. This example uses port-channel number 63 on switch 1 and port-channel number 64 on switch 2. The PortChannel interface numbers are arbitrary and should be adapted to best suit your deployment. The chosen PortChannel numbers reflect the highest values supported across all VSS platforms validated in this release. You must configure **channel-group mode on** for the VSL port channel because it is an infrastructure link actively managed withing the VSS using Virtual Switch Link Protocol (VSLP). This example uses the 10-Gigabit Ethernet interfaces on the supervisor of a Cisco Catalyst 6500 Series switch for the EtherChannel member ports of the VSL.

On standalone switch #1:

```
VSS-Sw1(config)#interface port-channel 63
VSS-Sw1(config-if)#switch virtual link 1
VSS-Sw1(config-if)#no shutdown
VSS-Sw1(config)#interface range tengigabit 5/4-5
VSS-Sw1(config-if)#channel-group 63 mode on
VSS-Sw1(config-if)#no shutdown
```

On standalone switch #2:

```
VSS-Sw2(config)#interface port-channel 64
VSS-Sw2(config-if)#switch virtual link 2
VSS-Sw2(config-if)#no shutdown
VSS-Sw2(config)#interface range tengigabit 5/4-5
VSS-Sw2(config-if)#channel-group 64 mode on
VSS-Sw2(config-if)#no shutdown
```

At this point you should be able to see that port-channel 63 and 64 are up, and both links are active on standalone switch #1 and standalone switch #2 respectively. The switches are not in VSS mode yet.

```
VSS-Sw1# show etherchannel 63 port
VSS-Sw2# show etherchannel 64 port
```

The previous two commands show the same output below.

```
        Ports in the group:
        -------------------
Port: Te5/4
------------
Port state  = Up Mstr In-Bndl
…
Port: Te5/5
------------
Port state  = Up Mstr In-Bndl
…
```

Step 5: Enable virtual switch mode operation.

Now that a port-channel has been established between the switches, convert each switch to virtual switch mode operation. At the enable prompt (that is, not in configuration mode) on each switch, enter the following commands for each switch.

On standalone switch #1:

```
VSS-Sw1# switch convert mode virtual
```

On standalone switch #2:

```
VSS-Sw2# switch convert mode virtual
```

When asked if you want to proceed, answer yes.

```
This command will convert all interface names to naming convention "interface-
type switch-number/slot/port", save the running config to startup-config and
reload the switch.

NOTE: Make sure to configure one or more dual-active detection methods once the
conversion is complete and the switches have come up in VSS mode.

Do you want to proceed? [yes/no]: yes
```

Each switch now renumbers its interfaces from interface y/z (where y is the slot number and z is the interface number) to interface x/y/z (where x is the switch number, y is the module number in that switch, and z is the interface on that module). This numbering scheme allows the two chassis to be addressed and configured as a single system from a single supervisor, which is the supervisor with the active control plane.

Once the configuration changes, it prompts you to save the configuration to bootflash. Press Return <CR> or Enter to accept the destination filename and location on each switch.

```
Converting interface names
Building configuration...
[OK]
Saving converted configuration to bootdisk: ...
Destination filename [startup-config.converted_vs-20140429-034524]?
9789 bytes copied in 0.680 secs (14396 bytes/sec)
*Apr 29 03:45:44.253: %SYS-5-RELOAD: Reload requested by Exec. Reload Reason:
Conversion.
Resetting .......
```

Both switches reload. The switch pair negotiates using VSLP over the VSL and becomes a VSS, with one of the switches resolved as the ACTIVE supervisor for the merged VSS switch. All configuration commands now must be entered on the single active switch console. The other physical chassis in the VSS pair contains the STANDBY HOT supervisor with a console port that displays the Standby prompt.

Use the following command to verify that both switches can see each other, that they are in SSO mode, and that the second supervisor is in STANDBY HOT status.

```
VSS-Sw1#show switch virtual redundancy
```

Confirm that the two Cisco Catalyst 6500 or Catalyst 6880-X switches are now operating as a single VSS system by using configuration mode to rename the switch hostname.

```
VSS-Sw1(config)#hostname D6500-VSS
D6500-VSS(config)#
```

A critical aspect of the Cisco Catalyst VSS is the control plane and data plane operating models. From a control plane standpoint the VSS uses an active-standby operating model. This means that supervisor hardware on one chassis becomes the active control plane for the entire VSS while the other supervisor hardware on the paired chassis becomes the standby. The control plane handles protocol operations like IP routing, peering, route table updates, and spanning tree BPDUs. The dataplane handles the hardware forwarding of packets, and both switches are actively forwarding traffic in an active-active operating model.

The VSL allows the switches to communicate and stay in synchronization. The VSS uses the Stateful Switchover (SSO) redundancy facility to keep the control plane synchronized between the two switches. As a result, the VSS appears to devices in adjoining layers as a single switch with a single MAC address.

Step 6: Configure dual-active detection mechanism.

In the event that the VSL is severed (that is, all links are down), or for any reason communication is lost over the VSL (such as excessive high CPU utilization), both switches would assume the active control plane role, thus creating a dual-active condition, which can result in network instability. To prevent a dual-active scenario from causing an outage in the network, VSS supports multiple unique dual-active detection and recovery mechanisms.
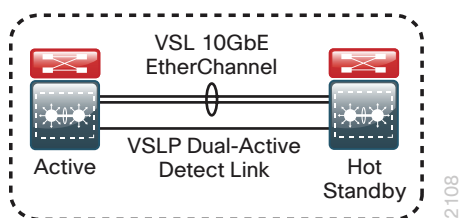
The dual-active detection mechanisms are used to trigger a VSS recovery mode. In the VSS recovery mode only one switch chassis is allowed to remain active, the other switch (the previous VSS active switch) enters recovery mode, and shuts down all of its interfaces except the VSL interfaces, thereby preventing instability in the network. Once the VSL is repaired, and communication over the VSL is reestablished, then the VSS reloads the switch in recovery mode and returns the VSS to a normal operating state.

You can use the following methods to detect this dual-active condition:

- Ethernet Fast-Hello (VSLP) link
- Enhanced Port Aggregation Protocol (PAgP) hellos with an adjacent switch

This design uses the Fast-Hello (VSLP) link for dual-active detection. To configure the link, use a Gigabit Ethernet interface on each VSS switch chassis and connect them together (similar to a VSL connection) in a back-to-back fashion. This link does not require high bandwidth because it is only a detection link with control plane hellos on it.

Figure 22 - VSLP


```
D6500-VSS(config)#interface range gigabit1/2/48, gigabit2/2/48
D6500-VSS(config-if-range)#dual-active fast-hello
D6500-VSS(config-if-range)#no shutdown
%VSDA-SW2_SPSTBY-5-LINK_UP: Interface Gi2/2/48 is now dual-active detection
capable
%VSDA-SW1_SP-5-LINK_UP: Interface Gi1/2/48 is now dual-active detection capable
```

Step 7: Configure the system virtual MAC address.

By default, the VSS system uses the default chassis-based MAC-address pool assigned to the switch that is resolved to be the active switch when the switches initialize. As a result of events such as stateful switchover, the MAC may change. Set a virtual MAC address for the VSS system so that either active supervisor will use the same MAC address pool, regardless of which supervisor is active, even across a system reload.

```
D6500-VSS(config)# switch virtual domain 100
D6500-VSS(config-vs-domain)# mac-address use-virtual
Configured Router mac address is different from operational value. Change will
take effect after the configuration is saved and the entire Virtual Switching
System (Active and Standby) is reloaded.
```

Step 8: Save and reload the switch.

Save the running configuration and then reload the entire system (both chassis).

```
copy running-config startup-config
reload
```

When the switches initialize after this final reload, the VSS configuration is complete.

Step 9: If you are configuring a VSS quad-supervisor VS4O system, allow the system to fully boot. Then insert the additional Supervisor Engine for each chassis, which has been previously prepared to boot the same version of code.

Because the Supervisor Engines have been prepared to boot the same version of code that is already active on the VSS, the Supervisor Engines boot and communicate with the In-Chassis Active (ICA) Supervisor Engine and receive the appropriate VSS variable configuration. A one-time reload automatically initiates on the inserted Supervisor Engines, and the modules boot as full participants in the VS4O platform.

After the Supervisor Engines have fully booted, you can observe the VS4O status.

```
C6500VSS#show switch virtual redundancy
```

The command output shows which switch and slot contains each Supervisor Engine in ACTIVE mode and STANDBY HOT mode, and the two Supervisors in the STANDBY HOT (CHASSIS) mode.

If you are using ports on the new Supervisor Engines for additional VSL connections as described in the tables before the first step, you now complete the connections, ensuring that you always have at least one VSL connection active. Once the physical connections are complete, add the connections into the existing VSL.

```
D6500-VSS(config)#interface range tengigabit 1/6/4-5
D6500-VSS(config-if)#channel-group 63 mode on
D6500-VSS(config-if)#no shutdown
D6500-VSS(config)#interface range tengigabit 2/6/4-5
D6500-VSS(config-if)#channel-group 64 mode on
D6500-VSS(config-if)#no shutdown
```

Verify that the status of all of the links in the VSL PortChannel are **P - bundled in port-channel** status.

```
D6500-VSS#show switch virtual link port-channel
```

Step 10: Configure QoS.

On the Cisco Catalyst 6500 Supervisor Engine 2T based switches and Cisco Catalyst 6880-X switches, QoS is enabled by default and policies for interface queuing are defined by attached service policies. The QoS policies are now defined using Cisco Common Classification Policy Language (C3PL), which is similar to Modular QoS CLI to reduce operational complexity.

All interface connections in the distribution and core are set to trust differentiated services code point (DSCP) markings. Even though this design is configured to trust DSCP markings, it is a best practice to ensure proper mapping of CoS to DSCP for VoIP. This mapping is accomplished by overriding the default mapping of CoS 5 "voice bearer traffic" to DSCP 40, with DSCP 46, which is the EF per-hop behavior for voice.

This egress QoS policy is configured to accommodate the 10-Gigabit and 40-Gigabit Ethernet ports on cards which use a 1P7Q4T queuing architecture.

```
! Enable port-based QoS
auto qos default
! Class maps for 1P7Q4T 10Gbps and 40Gbps ports service policy
class-map type lan-queuing match-any PRIORITY-QUEUE
  match dscp ef
  match dscp cs5
  match dscp cs4
  match cos 5
class-map type lan-queuing match-any CONTROL-MGMT-QUEUE
  match dscp cs7
  match dscp cs6
  match dscp cs3
  match dscp cs2
  match cos 3 6 7
class-map type lan-queuing match-any MULTIMEDIA-CONFERENCING-QUEUE
  match dscp af43
  match dscp af42
  match dscp af41
  match cos 4
class-map type lan-queuing match-any MULTIMEDIA-STREAMING-QUEUE
  match dscp af33
  match dscp af32
  match dscp af31
class-map type lan-queuing match-any TRANSACTIONAL-DATA-QUEUE
  match dscp af23
  match dscp af22
  match dscp af21
  match cos 2
class-map type lan-queuing match-any BULK-DATA-QUEUE
  match dscp af13
  match dscp af12
  match dscp af11
class-map type lan-queuing match-any SCAVENGER-QUEUE
  match dscp cs1
  match cos 1
!
policy-map type lan-queuing 1P7Q4T
 class PRIORITY-QUEUE
  priority
 class CONTROL-MGMT-QUEUE
```

```
    bandwidth remaining percent 14
    queue-buffers ratio 10
 class MULTIMEDIA-CONFERENCING-QUEUE
   bandwidth remaining percent 14
   queue-buffers ratio 10
   random-detect dscp-based
   random-detect dscp 38 percent 70 100
   random-detect dscp 36 percent 80 100
   random-detect dscp 34 percent 90 100
  class MULTIMEDIA-STREAMING-QUEUE
   bandwidth remaining percent 14
   queue-buffers ratio 10
   random-detect dscp-based
   random-detect dscp 30 percent 70 100
   random-detect dscp 28 percent 80 100
   random-detect dscp 26 percent 90 100
  class TRANSACTIONAL-DATA-QUEUE
   bandwidth remaining percent 14
   queue-buffers ratio 10
   random-detect dscp-based
   random-detect dscp 22 percent 70 100
   random-detect dscp 20 percent 80 100
   random-detect dscp 18 percent 90 100
  class BULK-DATA-QUEUE
   bandwidth remaining percent 6
   queue-buffers ratio 10
   random-detect dscp-based
   random-detect dscp 14 percent 70 100
   random-detect dscp 12 percent 80 100
   random-detect dscp 10 percent 90 100
  class SCAVENGER-QUEUE
   bandwidth remaining percent 2
   queue-buffers ratio 10
   random-detect dscp-based
   random-detect dscp 8 percent 80 100
 class class-default
   queue-buffers ratio 25
   random-detect dscp-based aggregate
   random-detect dscp values 0 1 2 3 4 5 6 7 percent 80 100
   random-detect dscp values 9 11 13 15 17 19 21 23 percent 80 100
   random-detect dscp values 25 27 29 31 33 35 37 39 percent 80 100
   random-detect dscp values 41 42 43 44 45 47 49 50 percent 80 100
   random-detect dscp values 51 52 53 54 55 57 58 59 percent 80 100
   random-detect dscp values 60 61 62 63 percent 80 100
 !
 table-map cos-discard-class-map
  map from 0 to 0
```

```
 map from 1 to 8
 map from 2 to 16
 map from 3 to 24
 map from 4 to 32
 map from 5 to 46
 map from 6 to 48
 map from 7 to 56
!
macro name EgressQoS
  service-policy type lan-queuing output 1P7Q4T
@
```

Step 11: If you are using Gigabit Ethernet cards supported in VSS mode on Cisco Catalyst 6500-XLand 6807-E Supervisor Engine 2T based switches, configure an additional QoS policy for the Gigabit Ethernet ports.

A separate egress QoS policy is configured to accommodate the Gigabit Ethernet cards, which use a 1P3Q8T queuing architecture supporting COS-based queuing. This policy does not apply to the Cisco Catalyst 6880-X platforms.

```
! Class maps for 1P3Q8T 1Gb ports service policy
class-map type lan-queuing match-any PRIORITY-QUEUE-GIG
 match cos 5 4
class-map type lan-queuing match-any CONTROL-AND-STREAM-MEDIA
 match cos 7 6 3 2
class-map type lan-queuing match-any BULK-DATA-SCAVENGER
 match cos 1
!
policy-map type lan-queuing 1P3Q8T
 class PRIORITY-QUEUE-GIG
  priority
  queue-buffers ratio 15
 class CONTROL-AND-STREAM-MEDIA
  bandwidth remaining percent 55
  queue-buffers ratio 40
class BULK-DATA-SCAVENGER
  bandwidth remaining percent 10
  queue-buffers ratio 20
  random-detect cos-based
  random-detect cos 1 percent 80 100
 class class-default
  queue-buffers ratio 25
  random-detect cos-based
  random-detect cos 0 percent 80 100
!
macro name EgressQoSOneGig
  service-policy type lan-queuing output 1P3Q8T
@
```

Option 2: Configure Cisco Catalyst 4500E VSS and 4500-X VSS platforms

The Cisco Catalyst 4500E and the Cisco Catalyst 4500-X Virtual Switching Systems merge two switches together as a single logical switch, using two Cisco Catalyst 4500E Series or two Cisco Catalyst 4500-X Series switches, respectively. The supervisor hardware on one of the switches acts as the active control plane for both switches by controlling protocols such as EIGRP, OSPF, Spanning Tree, CDP, and so forth, while the supervisor hardware in both switches actively switches packets.

The Cisco Catalyst 4500E VSS does not support Quad Supervisor VSS SSO in this release. However, for sparing purposes and for additional ports, a second Supervisor Engine can be installed in each VSS chassis. Though the Supervisor Engines cannot be active and participate in the control plane, the built-in Ethernet ports are active and available to be used for uplinks or other connectivity.

The following configuration example shows you how to convert two standalone Cisco Catalyst 4500E or Cisco Catalyst 4500-X switches to a Virtual Switching System (VSS). When you set up the Cisco Catalyst 4500E or Cisco Catalyst 4500-X Virtual Switching System, connect two 10-Gigabit Ethernet links between the chassis to provide the Virtual Switch Link (VSL). Always use at least two links.

You can use up to eight links as members of the VSL, and the links should be distributed across Supervisor Engines or line cards for resiliency. This design uses two 10-Gigabit Ethernet interfaces on each Supervisor Engine for the Cisco Catalyst 4500E, and two 10 Gigabit Ethernet ports on the Cisco Catalyst 4500-X. You connect the VSL interfaces together before you configure the VSS.

Step 1: Convert standalone Cisco Catalyst 4500E or 4500-X switches to VSS.

Configure a temporary hostname on each switch so you can keep track of your configuration steps. In a later step after the conversion is complete, you apply a replacement hostname to the merged VSS configuration.

On the Cisco Catalyst 4500E or 4500-X standalone switch #1:

```
Router#config t
Router#(config)#hostname VSS-Sw1
```

On the Cisco Catalyst 4500E or 4500-X standalone switch #2:

```
Router#config t
Router#(config)#hostname VSS-Sw2
```

To form a VSS pair, each switch in the pair must have a matching domain ID assigned. To support the interconnection of multiple VSS pairs, the domain ID selected for the pair should be unique. In this example, the domain number is 105. Each switch is also given a unique identifier within the domain, switch 1 or switch 2.

Figure 23 - VSS domain
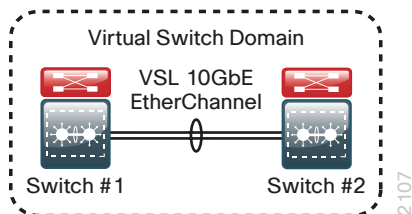
On the standalone switch #1:

```
VSS-Sw1(config)#switch virtual domain 105
VSS-Sw1(config-vs-domain)# switch 1
```

On the standalone switch #2:

```
VSS-Sw2(config)#switch virtual domain 105
VSS-Sw2(config-vs-domain)# switch 2
```

Step 2: Configure the Virtual Switch Link (VSL).

The VSL is a critical component of the Virtual Switching System. For each physical switch you must select a unique port-channel number identifying the same VSL. This allows the switch to maintain a separate identity for the interfaces used when making traffic forwarding decisions. This example uses port-channel number 63 on switch 1 and port-channel number 64 on switch 2. The PortChannel interface numbers are arbitrary and should be adapted to best suit your deployment. The chosen PortChannel numbers reflect the highest values supported across all VSS platforms validated in this release. You must configure **channel-group mode on** for the VSL port channel because it is an infrastructure link actively managed withing the VSS using Virtual Switch Link Protocol (VSLP).

On standalone switch #1:

```
VSS-Sw1(config)#interface port-channel 63
VSS-Sw1(config-if)#switchport
VSS-Sw1(config-if)#switch virtual link 1
VSS-Sw1(config-if)#no shutdown
VSS-Sw1(config)#interface range tengigabit 1/30-31
VSS-Sw1(config-if)#shutdown
VSS-Sw1(config-if)#channel-group 63 mode on
VSS-Sw1(config-if)#no shutdown
```

On standalone switch #2:

```
VSS-Sw2(config)#interface port-channel 64
VSS-Sw2(config-if)#switchport
VSS-Sw2(config-if)#switch virtual link 2
VSS-Sw2(config-if)#no shutdown
VSS-Sw2(config)#interface range tengigabit 1/30-31
VSS-Sw2(config-if)#shutdown
VSS-Sw2(config-if)#channel-group 64 mode on
VSS-Sw2(config-if)#no shutdown
```

The switches are not in VSS mode yet. Verify port-channel configuration on standalone switch #1.

```
VSS-Sw1# show etherchannel 63 sum
```

The command output includes output similar to the output below.

```
…
Group  Port-channel  Protocol    Ports
------+-------------+-----------+--------------------
63      Po63(SD)        —         Te1/30(w)   Te1/31(w)
```

Verify port-channel configuration on standalone switch #1.

```
VSS-Sw2# show etherchannel 64 sum
```

The command output includes output similar to the output below.

```
…
Group  Port-channel  Protocol    Ports
------+-------------+-----------+--------------------
64      Po64(SD)        —         Te1/30(w)   Te1/31(w)
```

The VSL port-channel ports will transition from waiting to be aggregated (w) mode to bundled in port-channel (P) mode only after the VSS conversion is complete in the next step.

Step 3: Enable virtual switch mode operation.

Now that a port-channel has been established between the switches, convert each switch to virtual switch mode operation. At the enable prompt (that is, not in configuration mode) on each switch, enter the following commands for each switch.

On standalone switch #1:

```
VSS-Sw1# switch convert mode virtual
```

On standalone switch #2:

```
VSS-Sw2# switch convert mode virtual
```

When asked if you want to proceed, answer yes.

```
This command will convert all interface names
to naming convention "interface-type switch-number/slot/port",
save the running config to startup-config and
reload the switch.

NOTE: Make sure to configure one or more dual-active detection methods
once the conversion is complete and the switches have come up in VSS mode.

Do you want to proceed? [yes/no]: yes
```

Each switch now renumbers its interfaces from interface y/z (where y is the slot number and z is the interface number) to interface x/y/z (where x is the switch number, y is the module number in that switch, and z is the interface on that module). This numbering scheme allows the two chassis to be addressed and configured as a single system from a single supervisor, which is the supervisor with the active control plane.

Once the configuration changes, it prompts you to save the configuration to bootflash. Press Return <CR> or Enter to accept the destination filename and location on each switch.

```
Converting interface names
Building configuration...
[OK]
Saving converted configuration to bootdisk: ...
Destination filename [startup-config.converted_vs-20140429-034524]?
9789 bytes copied in 0.680 secs (14396 bytes/sec)
*Apr 29 03:45:44.253: %SYS-5-RELOAD: Reload requested by Exec. Reload Reason:
Conversion.
Resetting .......
```

Both switches reload. The switch pair negotiates using VSLP over the VSL and becomes a VSS, with one of the switches resolved as the ACTIVE supervisor for the merged VSS switch. All configuration commands must now be entered on the single active switch console. The other physical chassis in the VSS pair contains the STANDBY HOT supervisor with a console port that displays the Standby prompt.

On the active switch console, use the following command to verify that both switches can see each other, that they are in SSO mode, and that the second supervisor is in STANDBY HOT status.

```
VSS-Sw1#show switch virtual redundancy
```

Confirm that the two Cisco Catalyst 4500 switches are now operating as a single VSS system by using configuration mode to rename the switch hostname.

```
VSS-Sw1(config)#hostname D4500-VSS
D4500-VSS(config)#
```

A critical aspect of the Cisco Catalyst VSS is the control plane and data plane operating models. From a control plane standpoint the VSS uses an active-standby operating model. This means that supervisor hardware on one chassis becomes the active control plane for the entire VSS while the other supervisor hardware on the paired chassis becomes the standby. The control plane handles protocol operations like IP routing, peering, route table updates, and spanning tree BPDUs. The dataplane handles the hardware forwarding of packets, and both switches are actively forwarding traffic in an active-active operating model.

The VSL allows the switches to communicate and stay in synchronization. The VSS uses the Stateful Switchover (SSO) redundancy facility to keep the control plane synchronized between the two switches. As a result, the VSS appears to devices in adjoining layers as a single switch with a single MAC address.

Step 4: Configure dual-active detection mechanism.

In the event that the VSL is severed (that is, all links are down), or for any reason communication is lost over the VSL (such as excessive high CPU utilization), both switches would assume the active control plane role, thus creating a dual-active condition, which can result in network instability. To prevent a dual-active scenario from causing an outage in the network, VSS supports multiple unique dual-active detection and recovery mechanisms.
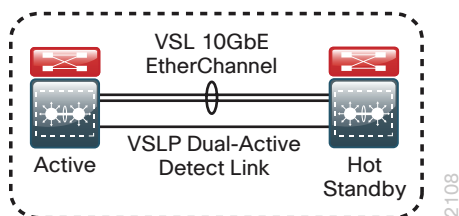
The dual-active detection mechanisms are used to trigger a VSS recovery mode. In the VSS recovery mode only one switch chassis is allowed to remain active, the other switch (the previous VSS active switch) enters recovery mode, and shuts down all of its interfaces except the VSL interfaces, thereby preventing instability in the network. Once the VSL is repaired, and communication over the VSL is reestablished, then the VSS reloads the switch in recovery mode and returns the VSS to a normal operating state.

You can use the following methods to detect this dual-active condition:

- Ethernet Fast-Hello (VSLP) link
- Enhanced Port Aggregation Protocol (PAgP) hellos with an adjacent switch

This design uses the Fast-Hello (VSLP) link for dual-active detection. To configure the link, use a Gigabit Ethernet interface on each VSS switch chassis and connect them together (similar to a VSL connection) in a back-to-back fashion. This link does not require high bandwidth because it is only a detection link with control plane hellos on it.

Figure 24 - VSLP

```
D4500-VSS(config)#interface range Ten1/1/32, Ten2/1/32
D4500-VSS(config-if-range)#dual-active fast-hello
D4500-VSS(config-if-range)#no shutdown
%VSDA-SW2_SPSTBY-5-LINK_UP: Interface Gi2/1/30 is now dual-active detection
capable
%VSDA-SW1_SP-5-LINK_UP: Interface Gi1/1/30 is now dual-active detection capable
```

The VSS configuration is complete.

---

i | **Tech Tip**

By default, at the time of virtual domain configuration, the Cisco Catalyst 4500 VSS
system uses a virtual MAC address for the VSS system so that either active supervisor
will use the same MAC address pool, regardless of which supervisor is active, even
across a system reload.

Step 5: If you are configuring a VSS system using Cisco Catalyst 4507R+E switches, each with two Cisco
Supervisor Engine 7-Es, the second Supervisor in each chassis can be installed as a spare Supervisor but the
Supervisor control plane cannot be active (that is, Catalyst 4500 VSS does not support Quad Supervisor VSS
SSO in this release). However, an installed but inactive second Supervisor has built-in Ethernet ports that are
active and available to be used for uplinks or other connectivity.

In order to use a system this way, repeat Step 1 through Step 4 in this procedure, with only the second
Supervisors installed. When the VSS configuration is complete, configure the following to ensure that the second
set of supervisors do not boot and become active.

```
D4500-VSS(config)#config-register 0x2100
D4500-VSS(config)#end
D4500-VSS#copy running-config startup-config
```

Press enter to accept the default destination filename.

```
D4500-VSS#redundancy reload peer
```

Press enter to confirm the peer switch reload.

```
D4500-VSS#reload
```

Press enter to confirm the switch reload.

Both Supervisors will reload and stop in ROMMON mode without booting IOS. The active Supervisors can now
be reinstalled, and configuration can continue once the VSS is active again.

Step 6: For each platform, define a macro that you will use in later procedures to apply the platform-specific
QoS configuration. This makes consistent deployment of QoS easier.

```
class-map match-any PRIORITY-QUEUE
  match  dscp ef
  match  dscp cs5
  match  dscp cs4
class-map match-any CONTROL-MGMT-QUEUE
  match  dscp cs7
  match  dscp cs6
  match  dscp cs3
  match  dscp cs2
```

```
class-map match-any MULTIMEDIA-CONFERENCING-QUEUE
  match  dscp af43
  match  dscp af42
  match  dscp af41
class-map match-any MULTIMEDIA-STREAMING-QUEUE
  match  dscp af33
  match  dscp af32
  match  dscp af31
class-map match-any TRANSACTIONAL-DATA-QUEUE
  match  dscp af23
  match  dscp af22
  match  dscp af21
class-map match-any BULK-DATA-QUEUE
  match  dscp af13
  match  dscp af12
  match  dscp af11
class-map match-any SCAVENGER-QUEUE
  match  dscp cs1
!
policy-map 1P7Q1T
 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
!
macro name EgressQoS
 service-policy output 1P7Q1T
@
```

Step 7: Save the running configuration.

```
copy running-config startup-config
```

Option 3: Configure the Cisco Catalyst 3850 platform

Step 1: To configure a Cisco Catalyst 3850 stack, use the CLI global exec mode (not configuration mode) to set the preferred active switch.

```
switch [switch number] priority 15
```

When there are multiple Cisco Catalyst 3850 Series Switches configured in a stack, one of the switches takes the ACTIVE switch role and another member takes the HOT-STANDBY role. Upon reload, the switch configured with the highest priority assumes the active role. If this is a new configuration, only the active switch console is active during the initial configuration. When two or more switches are configured as a stack, configure the active switch functionality on a switch of your preference.

Step 2: To make consistent deployment of QoS easier, each distribution platform defines a macro that will be used in later procedures to apply the platform specific QoS configuration.

> **i** **Tech Tip**
>
> The match statements in the class-maps can be combined into one line. However, listing them separately provides additional per-DSCP counters available from SNMP and from the following verification command.
>
> ```
> show policy-map interface
> ```

```
class-map match-any PRIORITY-QUEUE
  match  dscp ef
class-map match-any VIDEO-PRIORITY-QUEUE
  match  dscp cs5
  match  dscp cs4
class-map match-any CONTROL-MGMT-QUEUE
  match  dscp cs7
  match  dscp cs6
  match  dscp cs3
  match  dscp cs2
class-map match-any MULTIMEDIA-CONFERENCING-QUEUE
  match  dscp af41
  match  dscp af42
  match  dscp af43
class-map match-any MULTIMEDIA-STREAMING-QUEUE
  match  dscp af31
  match  dscp af32
  match  dscp af33
class-map match-any TRANSACTIONAL-DATA-QUEUE
  match  dscp af21
  match  dscp af22
  match  dscp af23
class-map match-any BULK-SCAVENGER-DATA-QUEUE
  match  dscp af11
  match  dscp af12
  match  dscp af13
```

```
      match  dscp cs1
!
policy-map 2P6Q3T
 class PRIORITY-QUEUE
    priority level 1
    police rate percent 10
 class VIDEO-PRIORITY-QUEUE
    priority level 2
    police rate percent 20
 class CONTROL-MGMT-QUEUE
    bandwidth remaining percent 10
    queue-buffers ratio 10
 class MULTIMEDIA-CONFERENCING-QUEUE
    bandwidth remaining percent 10
    queue-buffers ratio 10
    queue-limit dscp af43 percent 80
    queue-limit dscp af42 percent 90
    queue-limit dscp af41 percent 100
 class MULTIMEDIA-STREAMING-QUEUE
    bandwidth remaining percent 10
    queue-buffers ratio 10
    queue-limit dscp af33 percent 80
    queue-limit dscp af32 percent 90
    queue-limit dscp af31 percent 100
 class TRANSACTIONAL-DATA-QUEUE
    bandwidth remaining percent 10
    queue-buffers ratio 10
    queue-limit dscp af23 percent 80
    queue-limit dscp af22 percent 90
    queue-limit dscp af21 percent 100
 class BULK-SCAVENGER-DATA-QUEUE
    bandwidth remaining percent 5
    queue-buffers ratio 10
    queue-limit dscp values af13 cs1 percent 80
    queue-limit dscp values af12 percent 90
    queue-limit dscp values af11 percent 100
 class class-default
    bandwidth remaining percent 25
    queue-buffers ratio 25
!
macro name EgressQoS
 service-policy output 2P6Q3T
@
```

Option 4: Configure the Cisco Catalyst 3750-X platform

Step 1: When there are multiple switches configured in a stack, one of the switches controls the operation of the stack. This switch is called the stack master.

When three or more switches are configured as a stack, configure the stack master switch functionality on a switch that does not have uplinks configured.

```
switch [switch number] priority 15
```

If you configure stack master switch priority on Cisco 3750-X switch stack, a single reload is required to force the stack master to operate on the switch that you configured with the highest priority. Reload the switch stack after all of your configuration is complete for this entire "Configuring the Distribution Layer" process.

Step 2: By default, the newly active stack master switch assigns a new stack MAC address when the stack master switch fails. This new MAC address assignment can cause the network to reconverge because LACP and many other protocols rely on the stack MAC address and must restart. As such, you should use the **stack-mac persistent timer 0** command to ensure that the original master MAC address remains the stack MAC address after a failure.

```
stack-mac persistent timer 0
```

Step 3: To make consistent deployment of QoS easier, each distribution platform defines a macro that will be used in later procedures to apply the platform specific QoS configuration. Since AutoQoS might not be configured on this device, manually configure the global QoS settings by running the following commands.

```
mls qos map policed-dscp 0 10 18 to 8
mls qos map cos-dscp 0 8 16 24 32 46 48 56
mls qos srr-queue input bandwidth 70 30
mls qos srr-queue input threshold 1 80 90
mls qos srr-queue input priority-queue 2 bandwidth 30
mls qos srr-queue input cos-map queue 1 threshold 2 3
mls qos srr-queue input cos-map queue 1 threshold 3 6 7
mls qos srr-queue input cos-map queue 2 threshold 1 4
mls qos srr-queue input dscp-map queue 1 threshold 2 24
mls qos srr-queue input dscp-map queue 1 threshold 3 48 49 50 51 52 53 54 55
mls qos srr-queue input dscp-map queue 1 threshold 3 56 57 58 59 60 61 62 63
mls qos srr-queue input dscp-map queue 2 threshold 3 32 33 40 41 42 43 44 45
mls qos srr-queue input dscp-map queue 2 threshold 3 46 47
mls qos srr-queue output cos-map queue 1 threshold 3 4 5
mls qos srr-queue output cos-map queue 2 threshold 1 2
mls qos srr-queue output cos-map queue 2 threshold 2 3
mls qos srr-queue output cos-map queue 2 threshold 3 6 7
mls qos srr-queue output cos-map queue 3 threshold 3 0
mls qos srr-queue output cos-map queue 4 threshold 3 1
mls qos srr-queue output dscp-map queue 1 threshold 3 32 33 40 41 42 43 44 45
mls qos srr-queue output dscp-map queue 1 threshold 3 46 47
mls qos srr-queue output dscp-map queue 2 threshold 1 16 17 18 19 20 21 22 23
mls qos srr-queue output dscp-map queue 2 threshold 1 26 27 28 29 30 31 34 35
mls qos srr-queue output dscp-map queue 2 threshold 1 36 37 38 39
mls qos srr-queue output dscp-map queue 2 threshold 2 24
mls qos srr-queue output dscp-map queue 2 threshold 3 48 49 50 51 52 53 54 55
```

```
mls qos srr-queue output dscp-map queue 2 threshold 3 56 57 58 59 60 61 62 63
mls qos srr-queue output dscp-map queue 3 threshold 3 0 1 2 3 4 5 6 7
mls qos srr-queue output dscp-map queue 4 threshold 1 8 9 11 13 15
mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14
mls qos queue-set output 1 threshold 1 100 100 50 200
mls qos queue-set output 1 threshold 2 125 125 100 400
mls qos queue-set output 1 threshold 3 100 100 100 3200
mls qos queue-set output 1 threshold 4 60 150 50 200
mls qos queue-set output 1 buffers 15 25 40 20
mls qos
!
macro name EgressQoS
 mls qos trust dscp
 queue-set 1
 srr-queue bandwidth share 1 30 35 5
 priority-queue out
@
!
```

Procedure 2   Configure LAN switch universal settings

In this design, there are features and services that are common across all LAN switches, regardless of the type of platform or role in the network. These are system settings that simplify and secure the management of the solution.

This procedure provides examples for some of those settings. The actual settings and values will depend on your current network configuration.

Table 9 - Common network services used in the design examples

Setting	Value
Domain Name	cisco.local
Active Directory, DNS, DHCP Server	10.4.48.10
Authentication Control System	10.4.48.15
Network Time Protocol Server	10.4.48.17
EIGRP Named Mode Configuration Name	CAMPUS
EIGRP AS or OSPF AS	100
Multicast Range	239.1.0.0/16

Step 1: Configure the device hostname to make it easy to identify the device.

```
hostname [hostname]
```

Step 2: If the switch VTP mode has been changed from default, configure VTP transparent mode. This design uses VTP transparent mode because the benefits of dynamic propagation of VLAN information across the network are not worth the potential for unexpected behavior resulting from operational error.

VLAN Trunking Protocol (VTP) allows network managers to configure a VLAN in one location of the network and have that configuration dynamically propagate out to other network devices. However, in most cases, VLANs are defined once during switch setup with few, if any, additional modifications.

```
vtp mode transparent
```

Step 3: Enable Rapid Per-VLAN Spanning-Tree (PVST+). Rapid PVST+ provides an instance of RSTP (802.1w) per VLAN. Rapid PVST+ greatly improves the detection of indirect failures or linkup restoration events over classic spanning tree (802.1D).

Although this architecture is built without any Layer 2 loops, you should still enable spanning tree with the most up-to-date network safeguards. By enabling spanning tree, you ensure that if any physical or logical loops are accidentally configured, no actual layer 2 loops occur.

```
spanning-tree mode rapid-pvst
```

Step 4: Set the distribution layer switch to be the spanning-tree root for all VLANs on access layer switches or appliances that you are connecting to the distribution switch.

```
spanning-tree vlan 1-4094 root primary
```

Step 5: Enable Unidirectional Link Detection (UDLD) as the default for fiber ports.

UDLD is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. When UDLD detects a unidirectional link, it disables the affected interface and alerts you. Unidirectional links can cause a variety of problems, including spanning-tree loops, black holes, and non-deterministic forwarding. In addition, UDLD enables faster link failure detection and quick reconvergence of interface trunks, especially with fiber, which can be susceptible to unidirectional failures.

```
udld enable
```

Step 6: Enable the recovery mechanism to allow ports disabled as a result of errors to automatically clear the err-disable status and attempt a recovery to operational behavior and connected status. Enabling the recovery mechanism avoids having to intervene manually by using the CLI to shutdown and enable the port after the cause of the error is cleared. By default, the recovery mechanism waits five minutes to attempt clearing of the interface err-disable status.

```
errdisable recovery cause all
```

Step 7: Set EtherChannels to use the traffic source and destination IP address when calculating which link to send the traffic across. This normalizes the method in which traffic is load-shared across the member links of the EtherChannel. EtherChannels are used extensively in this design because of their resiliency capabilities.

```
port-channel load-balance src-dst-ip
```

Step 8: Configure DNS for host lookup.

At the command line of a Cisco IOS device, it is helpful to be able to type a domain name instead of the IP address for a destination.

```
ip name-server 10.4.48.10
```

Step 9: Configure device management protocols.

Secure HTTP (HTTPS) and Secure Shell (SSH) are more secure replacements for the HTTP and Telnet protocols. They use Secure Sockets Layer (SSL) and Transport Layer Security (TLS) to provide device authentication and data encryption.

The SSH and HTTPS protocols enable secure management of the LAN device. Both protocols are encrypted for privacy, and the unencrypted protocols, Telnet and HTTP, are turned off. Enabling HTTPS automatically generates a cryptographic key to use the service. When SSH is configured after HTTPS, you do not have to explicitly generate the cryptographic key that SSH requires, unless you wish to change the default key size.

Secure Copy (SCP) provides a secure and authenticated method for copying configuration and image files by making use of SSH as a secure transport. Enable SCP to allow secure file management with the device to avoid the use of less secure protocols such as TFTP and FTP.

Specify the transport preferred none on vty lines to prevent errant connection attempts from the CLI prompt. Without this command, if the ip name server is unreachable, long timeout delays may occur for mistyped commands.

```
no ip http server
ip http secure-server
ip domain-name cisco.local
ip ssh version 2
ip scp server enable
!
line vty 0 15
  transport input ssh
  transport preferred none
```

Step 10: Enable Simple Network Management Protocol (SNMP) in order to allow the network infrastructure devices to be managed by a Network Management System (NMS), and then configure SNMPv2c both for a read-only and a read-write community string.

```
snmp-server community [SNMP RO name] RO
snmp-server community [SNMP RW name] RW
```

Step 11: If your network operational support is centralized, you can increase network security by using an access list to limit the networks that can access your device. In this example, only devices on the 10.4.48.0/24 network will be able to access the device via SSH or SNMP.

```
access-list 55 permit 10.4.48.0 0.0.0.255
line vty 0 15
  access-class 55 in
!
snmp-server community [SNMP RO name] RO 55
snmp-server community [SNMP RW name] RW 55
```

The Cisco Catalyst 3850 Series Switch has an additional keyword to be added to the access-class, which allows console access from other switch members to not be affected.

```
line vty 0 15
  access-class 55 in vrf-also
```

> **! Caution**
>
> If you configure an access-list on the vty interface, you may lose the ability to use ssh to log in from one device to the next for hop-by-hop troubleshooting.

Step 12: Configure local login and password

The local login account and password provides basic device access authentication to view platform operation. The enable password secures access to the device configuration mode. By enabling password encryption, you prevent the use of plain text passwords when viewing configuration files. The **aaa new-model** command enables new access control commands and functions, and causes the local username and password on the router to be used in the absence of other AAA statements.

```
username admin secret [user password]
enable secret [enable password]
service password-encryption
aaa new-model
```

By default, https access to the switch will use the enable password for authentication.

Step 13: If you want to reduce operational tasks per device, configure centralized user authentication by using the TACACS+ protocol to authenticate management logins on the infrastructure devices to the AAA server.

As networks scale in the number of devices to maintain, there is an operational burden to maintain local user accounts on every device. A centralized AAA service reduces operational tasks per device and provides an audit log of user access for security compliance and root cause analysis. When AAA is enabled for access control, all management access to the network infrastructure devices (SSH and HTTPS) is controlled by AAA.

TACACS+ is the primary protocol used to authenticate management logins on the infrastructure devices to the AAA server. A local AAA user database is also defined on each network infrastructure device in order to provide a fallback authentication source in case the centralized TACACS+ server is unavailable.

```
tacacs server TACACS-SERVER-1
address ipv4 10.4.48.15
key [tacacs key]
!
aaa group server tacacs+ TACACS-SERVERS
  server name TACACS-SERVER-1
!
aaa authentication login default group TACACS-SERVERS local
aaa authorization exec default group TACACS-SERVERS local
aaa authorization console
ip http authentication aaa
```

Step 14: Configure a synchronized clock by programming network devices to synchronize to a local NTP server in the network. The local NTP server typically references a more accurate clock feed from an outside source. Configure console messages, logs, and debug output to provide time stamps on output, which allows cross-referencing of events in a network.

```
ntp server 10.4.48.17
ntp update-calendar
!
clock timezone PST -8
clock summer-time PDT recurring
!
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
```

The **ntp update-calendar** command configures the switch to update the hardware clock from the ntp time source periodically. Since not all switches have a hardware clock, this command is not supported by all devices.

Step 1: Configure BPDU Guard globally to protect PortFast-enabled interfaces.

In some scenarios, a service appliance that requires **spanning-tree portfast** may be connected to the distribution layer. When an interface is set for portfast, BPDU guard protects against an accidental connection of another switch into a PortFast-enabled interface, which could cause a catastrophic undetected spanning-tree loop.

If a PortFast-configured interface receives a BPDU, an invalid configuration exists, such as the connection of an unauthorized device. The BPDU guard feature prevents loops by moving a nontrunking interface into an errdisable state when a BPDU is received on an interface when PortFast is enabled.

Disable the interface if another switch is plugged into the PortFast-enabled interface.

```
spanning-tree portfast bpduguard default
```

On the Cisco Catalyst 6500 and Catalyst 6800 Series Switches, the global BPDU Guard command is slightly different.

```
spanning-tree portfast edge bpduguard default
```

Step 2: Configure an in-band management interface.

The loopback interface is a logical interface that is always reachable as long as the device is powered on and any IP interface is reachable to the network. Because of this capability, the loopback address is the best way to manage the switch in-band. Layer 3 process and features are also bound to the loopback interface to ensure process resiliency.

The loopback address is commonly a host address with a 32-bit address mask. Allocate the loopback address from the IP address block that the distribution switch summarizes to the rest of the network.

```
interface Loopback0
  ip address [ip address] 255.255.255.255
  ip pim sparse-mode
```

The need for the **ip pim sparse-mode** command will be explained further in Step 3 of Procedure 5, "Configure IP Multicast routing".

Step 3: Configure the system processes to use the loopback interface address for optimal resiliency:

```
snmp-server trap-source Loopback 0
ip ssh source-interface Loopback 0
ip pim register-source Loopback 0
ip tacacs source-interface Loopback 0
ntp source Loopback 0
```

The single logical distribution layer design, when configured with VSS, uses Stateful Switchover and Nonstop Forwarding to provide subsecond failover in the event of a supervisor data or control plane failure. This ability reduces packet loss in switchover to redundant logic and keeps packets flowing when the data plane is still intact to adjacent nodes. In the stack-based distribution layer approach, a single logical control point still exists and the master control plane in a stack can fail over to another member in the stack providing near-second or subsecond resiliency.

When the supervisor or master switch of a distribution platform switches over from the active to the hot-standby supervisor or switch, it will continue switching IP data traffic flows in hardware. However, the device in the active role requires time to reestablish control plane two-way peering with IP routing neighbors and avoid the peer router from tearing down adjacencies due to missed hellos that would cause a reroute and disruption of traffic. To allow this time for the device taking over the active role to recover, there is a Nonstop Forwarding (NSF) setting for the routing protocol to wait for the dual supervisor peer switch to recover. The neighboring router is said to be NSF-aware if it has a newer release of Cisco IOS Software that recognizes an NSF peer. All of the platforms used in this design are NSF-aware for the routing protocols in use.

The distribution layer switch is configured to enable NSF for the routing protocol in use so that it can signal a peer when it switches over from a previously active to a hot-standby device, to allow the peering neighbor time to reestablish the IP routing protocol relationship to that node. No tuning of the default NSF timers is needed in this network. Nothing has to be configured for an NSF-aware peer router.

Option 1: Configure EIGRP unicast routing

Enhanced Interior Gateway Routing Protocol (EIGRP) is the IP unicast routing protocol used in this design because it is easy to configure, does not require a large amount of planning, has flexible summarization and filtering, and can scale to large networks. If you use OSPF as an alternative to EIGRP, choose Option 2.

Step 1: Enable EIGRP named mode for the IP address space that the network will be using. If needed for your network, you can enter multiple network statements. Enable all routed links to be passive by default. The Loopback 0 IP address is used for the EIGRP router ID to ensure maximum resiliency. Because routing functionality is bounded at the distribution and not extended into the access layer, every distribution is configured as a stub network, optimizing performance. The summary keyword allows summary routes to be advertised and summarization is used whenever possible.

> **i**　**Tech Tip**
>
> There are situations where a layer-3 distribution switch cannot be an EIGRP stub network, such as in a WAN aggregation or for an Internet edge distribution. Those situations are addressed in the guides for those configurations.

Cisco Catalyst 6800 and 6500 Series Switches do not require the **ip routing** command because it is enabled by default on that platform.

```
ip routing
!
router eigrp LAN
 address-family ipv4 unicast autonomous-system 100
  af-interface default
   passive-interface
  exit-af-interface
  network 10.4.0.0 0.1.255.255
  eigrp router-id [ip address of loopback 0]
  eigrp stub summary
  nsf
 exit-address-family
```

Option 2: Configure OSPF unicast routing

Open Shortest Path First (OSPF) can be used instead of EIGRP for networks where OSPF is required for compatibility. If you configured EIGRP in the previous procedure, you can skip this option.

Step 1: Enable OSPF for the IP address space that the network will be using. If needed for your network, you can enter multiple network statements. Enable all routed links to be passive by default. The Loopback 0 IP address is used for the OSPF router ID to ensure maximum resiliency. Each distribution gets a unique non-zero area number, which is configured as a totally stubby area to optimize performance. An OSPF totally stubby area only has a single default route out to the rest of the network, which is the case for a distribution switch.

> **i** Tech Tip
>
> There are situations where a layer 3 distribution switch cannot be an OSPF totally stubby network, such as in a WAN aggregation. Those situations are addressed in the WAN guides for those configurations.

Cisco Catalyst 6800 and 6500 Series Switches do not require the **ip routing** command because it is enabled by default on those platforms.

```
ip routing
!
router ospf 100
 router-id [IP address of loopback 0]
 nsf
 area [unique area number] stub no-summary
 passive-interface default
 network 10.4.0.0 0.0.15.255 area [unique area number]
 network 10.4.40.0 0.0.0.255 area 0
```

IP Multicast allows a single IP data stream to be replicated by the infrastructure (that is, routers and switches) and sent from a single source to multiple receivers. Using IP Multicast is much more efficient than multiple individual unicast streams or a broadcast stream that would propagate everywhere. IP Telephony Music on Hold and IP Video Broadcast Streaming are two examples of IP Multicast applications.

To receive a particular IP Multicast data stream, end hosts must join a multicast group by sending an Internet Group Management Protocol (IGMP) message to their local multicast router. In a traditional IP Multicast design, the local router consults another router in the network that is acting as a Rendezvous Point (RP) to map the receivers to active sources so they can join their streams.

The RP is a control plane operation that should be placed in the core of the network or close to the IP Multicast sources on a pair of Layer 3 switches or routers. IP Multicast routing begins at the distribution layer if the access layer is Layer 2 and provides connectivity to the IP Multicast RP. In designs without a core layer, the distribution layer will perform the RP function.

Figure 25 - Rendezvous point placement in the network
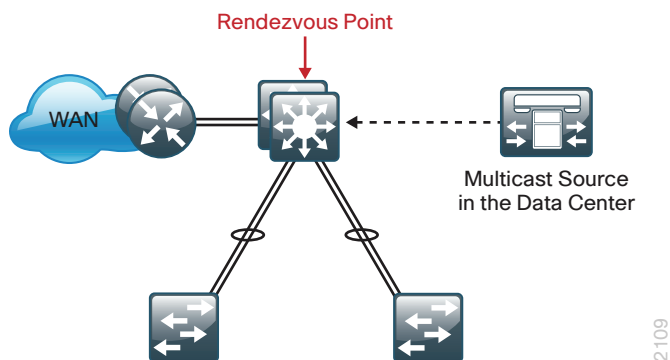
This design is based on sparse mode multicast operation.

Step 1: Configure IP Multicast routing on the platforms in the global configuration mode.

```
ip multicast-routing
```

Cisco Catalyst 3750 Series Switches instead require the **ip multicast-routing distributed** command.

Step 2: Configure the switch to discover the IP Multicast RP.

Every Layer 3 switch and router is configured to discover the IP Multicast RP with AutoRP in this design—other alternatives are not covered. Use the **ip pim autorp listener** command to allow for discovery across sparse mode links. This configuration provides for future scaling and control of the IP Multicast environment and can change based on network needs and design.

```
ip pim autorp listener
```

Step 3: Configure ip pim sparse-mode. All Layer 3 interfaces in the network should be enabled for sparse mode multicast operation.

```
ip pim sparse-mode
```

Example: Procedures 3-5 with EIGRP

```
spanning-tree portfast bpduguard default
!
interface Loopback 0
  ip address 10.4.15.254 255.255.255.255
  ip pim sparse-mode
!
snmp-server trap-source Loopback 0
ip ssh source-interface Loopback 0
ip pim register-source Loopback 0
ip tacacs source-interface Loopback 0
ntp source Loopback 0
!
ip routing
!
router eigrp LAN
 address-family ipv4 unicast autonomous-system 100
  af-interface default
   passive-interface
  exit-af-interface
  network 10.4.0.0 0.1.255.255
  eigrp router-id 10.4.15.254
  eigrp stub summary
  nsf
 exit-address-family
!
```

Example: Procedures 3-5 with OSPF

```
spanning-tree portfast bpduguard default
!
interface Loopback 0
  ip address 10.4.15.254 255.255.255.255
  ip pim sparse-mode
!
snmp-server trap-source Loopback 0
ip ssh source-interface Loopback 0
ip pim register-source Loopback 0
ip tacacs source-interface Loopback 0
ntp source Loopback 0
!
ip routing
!
router ospf 100
 router-id 10.4.15.254
 nsf
 area 0 authentication message-digest
 area 1 stub no-summary
```

```
 area 1 range 10.4.0.0 255.255.240.0
 passive-interface default
 no passive-interface Port-channel30
 network 10.4.0.0 0.0.15.255 area 1
 network 10.4.40.0 0.0.0.255 area 0
ip multicast-routing
ip pim autorp listener
!
```

<table>
<tr><td>**Procedure 6**</td><td>Configure IP Multicast RP</td></tr>
</table>

(Optional)

In networks without a core layer, the RP function can be placed on the distribution layer. If a core layer does exist, follow the IP Multicast Procedure 4 in the core layer section to configure the RP function.

Every Layer 3 switch and router must know the address of the IP Multicast RP, including the core switches that are serving as the RP. This design uses AutoRP to announce candidate RPs, which are the core switches, to the rest of the network.

Step 1: Configure loopback interface for RP.

Configure a second loopback interface to be used as the RP interface. The interface uses a host address mask (32 bits). All routers then point to this common IP address on **loopback 1** for the RP.

```
interface Loopback 1
 ip address 10.4.15.253 255.255.255.255
 ip pim sparse-mode
```

> **i** | **Tech Tip**
>
> Although you could use an existing loopback interface, adding a new interface increases the ability to rapidly adapt to future requirements which may drive a change in the location of the RP.
>
> For example, if your RP is currently configured on a distribution layer, you may want to move the RP when you add a core. Configuring the RP address on the loopback interface at the new location with the same IP address used on Loopback 1 in this procedure and establishing IP Multicast and MSDP peering enables the migration. All remote routers should still point to the same RP address, which simplifies the move and reduces disruption to the IP Multicast environment.

Step 2: Configure AutoRP candidate RP.

The **send-rp-announce** command in conjunction with the **group-list** option advertises the RP address, with the multicast range the device is willing to serve, as a candidate RP to the AutoRP mapping agents.

```
access-list 10 permit 239.1.0.0 0.0.255.255
ip pim send-rp-announce Loopback 1 scope 32 group-list 10
```

Step 3: Configure AutoRP mapping agent.

The AutoRP mapping agent listens for candidate RPs and then advertises to the rest of the network the list of available RPs. The **send-rp-discovery** command enables this switch to act as an AutoRP mapping agent.

```
ip pim send-rp-discovery Loopback0 scope 32
```

Procedure 7	Connect to access layer

The resilient, single, logical, distribution layer switch design is based on a hub-and-spoke or star design. The links to access layer switches and connected routers are Layer 2 EtherChannels. Links to other distribution layers, and the optional core are Layer 3 links or Layer 3 EtherChannels.

When using EtherChannel, the member interfaces should be on different switches in the stack or different modules in the modular switch for the highest resiliency.

The physical interfaces that are members of a Layer 2 EtherChannel are configured prior to configuring the logical port-channel interface. This allows for minimal configuration because most of the commands entered to a port-channel interface are copied to its members' interfaces and do not require manual replication.

Configure two or more physical interfaces to be members of the EtherChannel. It is recommended that they are added in multiples of two.

If this distribution layer will be used as a network-services aggregation block, you likely will not have an access layer to connect.

Step 1: Configure VLANs.

Configure all VLANs for the access layer switches that you are connecting to the distribution switch.

```
vlan [data vlan]
  name Data
exit
vlan [voice vlan]
  name Voice
exit
vlan [management vlan]
  name Management
exit
```

Step 2: If there is no external central site DHCP server in the network, you can provide DHCP service in IOS by configuring the IOS DHCP server. This function can also be useful at a remote-site where you want to provide local DHCP service and not depend on the WAN link to an external central site DHCP server.

```
ip dhcp excluded-address 10.4.100.1 10.4.100.10
ip dhcp pool access
  network 10.4.100.0 255.255.255.0
  default-router 10.4.100.1
  domain-name cisco.local
  dns-server 10.4.48.10
```

The example configuration provides IP addresses via the IOS based DHCP service for the subnet 10.4.100.0/24 and prevents the server from assigning reserved addresses .1-.10.

Step 3: Configure EtherChannel member interfaces.

This design uses Layer 2 EtherChannels to connect all access layer switches to the distribution layer and thereby create the hub-and-spoke resilient design that eliminates spanning-tree loops. Add links in multiples of two and distribute as much as possible across physical components of the platform. A configuration is shown using four member links for additional resiliency.

Connect the access layer EtherChannel uplinks to separate switches in the distribution layer Virtual Switching System or stack.

Also, apply the egress QoS macro that was defined in the platform configuration procedure to ensure traffic is prioritized appropriately.

Cisco Catalyst 4500 and 4500-X Series Switches do not use the **logging event bundle-status** command.

```
interface [interface type] [port 1]
  description Link to {your device here} Port 1
interface [interface type] [port 2]
  description Link to {your device here} Port 2
interface [interface type] [port 3]
  description Link to {your device here} Port 3
interface [interface type] [port 4]
  description Link to {your device here} Port 4
!
interface range [interface type] [port 1], [interface type] [port 2], [interface
type] [port 3], [interface type] [port 4]
  switchport
  channel-protocol lacp
  channel-group [number] mode active
  logging event link-status
  logging event bundle-status
  load-interval 30
  macro apply EgressQoS
  !6K platforms with 1-Gigabit Ethernet cards use alternative
  !macro apply EgressQoSOneGig
```

> **ℹ Tech Tip**
>
> The Cisco Catalyst 6500 and 6800 Series Switches have two egress QoS macros: EgressQoSOneGig, which is used for Gigabit Ethernet ports, and EgressQoS, which is used for 10-Gigabit or 40-Gigabit Ethernet ports. All other distribution layer platforms have a single egress QoS macro, which applies to all Ethernet ports, including Gigabit Ethernet and 10-Gigabit Ethernet.

Step 4: Configure the VLAN trunk interface to the access layer.

An 802.1Q trunk is used for the connection to the access layer, which allows the distribution switch to provide Layer 3 services to all the VLANs defined on the access layer switch. Prune the VLANs on the trunk to only the VLANs that are active on the access switch. When using EtherChannel the interface type will be port-channel and the number must match the channel group configured in Step 3 .

The Cisco Catalyst 3750 Series Switch requires the **switchport trunk encapsulation dot1q** command.

```
interface [port-channel] [number]
  description EtherChannel Link to {your device here}
  switchport trunk allowed vlan [data vlan],[voice vlan],[management vlan]
  switchport mode trunk
  logging event link-status
  logging event trunk-status
  logging event bundle-status
  load-interval 30
  no shutdown
  exit
```

If the interface type is not PortChannel, then the additional command macro apply EgressQoS must also be configured on the interface.

Next, mitigate VLAN hopping on the trunk for switch-to-switch connections.

There is a remote possibility that an attacker can create a double 802.1Q encapsulated packet. If the attacker has specific knowledge of the 802.1Q native VLAN, they could create a packet that when processed, removes the first or outermost tag when the packet is switched onto the untagged native VLAN. When the packet reaches the target switch, the inner or second tag is then processed and the potentially malicious packet is switched to the target VLAN.

At first glance, this appears to be a serious risk. However, the traffic in this attack scenario is in a single direction and no return traffic can be switched by this mechanism. Additionally, this attack cannot work unless the attacker knows the native VLAN ID.

Step 5: Configuring an unused VLAN on all switch-to-switch 802.1Q trunk links from access layer to distribution layer removes the remote risk of this type of attack. By choosing an arbitrary, non-default, unused VLAN assignment for the native VLAN, you reduce the possibility that a double 802.1Q-tagged packet can hop VLANs.

```
vlan 999
  name AntiVLANhopping
exit
!
interface [port-channel] [number]
  switchport trunk native vlan 999
```
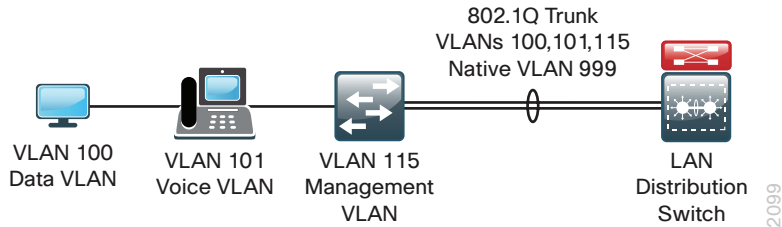
Step 6: Configure Layer 3.

Configure a VLAN interface (SVI) for every access layer VLAN so devices in the VLAN can communicate with the rest of the network.

Use the **ip helper-address** command to allow remote DHCP servers to provide IP addresses for this network. The address that the **helper** command points to is the central DHCP server. If you have more than one DHCP server, you can list multiple helper commands on an interface.

```
interface vlan [number]
  ip address [ip address] [mask]
  ip helper-address [dhcp server ip]
  ip pim sparse-mode
  no shutdown
```

If you configured the IOS DHCP server function on this distribution layer switch in Step 2 of this procedure, the **ip helper-address** is not needed on the VLAN interface.

Example: Access switch VLAN deployment

802.1Q Trunk
VLANs 100,101,115
Native VLAN 999

VLAN 100
Data VLAN

VLAN 101
Voice VLAN

VLAN 115
Management
VLAN

LAN
Distribution
Switch

```
vlan 100
  name Data
vlan 101
  name Voice
vlan 115
  name Management
vlan 999
  name AntiVLANhopping
spanning-tree vlan 1-4094 root primary
!
interface GigabitEthernet 1/1/1
  description Link to Access Switch Port 1
interface GigabitEthernet 2/1/1
  description Link to Access Switch Port 2
interface GigabitEthernet 1/1/2
  description Link to Access Switch Port 3
interface GigabitEthernet 2/1/2
  description Link to Access Switch Port 4
!
interface range GigabitEthernet 1/1/1, GigabitEthernet 2/1/1, GigabitEthernet
1/1/2, GigabitEthernet 2/1/2
  switchport
  channel-protocol lacp
  channel-group 10 mode active
  logging event link-status
  logging event trunk-status
  logging event bundle-status
  load-interval 30
  macro apply EgressQoS
  no shutdown
!
interface Port-channel 10
  description EtherChannel Link to Access Switch
  switchport trunk native vlan 999
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 100,101,115
```

```
    switchport mode trunk
    load-interval 30
    no shutdown
!
interface vlan 100
  ip address 10.4.0.1 255.255.255.0
  ip helper-address 10.4.48.10
  ip pim sparse-mode
!
interface vlan 101
  ip address 10.4.1.1 255.255.255.0
  ip helper-address 10.4.48.10
  ip pim sparse-mode
!
interface vlan 115
  ip address 10.4.15.1 255.255.255.128
  ip pim sparse-mode
```

Procedure 8   Connect to LAN core or WAN router

Any links to connected WAN routers or a LAN core layer should be Layer 3 links or Layer 3 EtherChannels. The LAN design does not extend Layer 2 VLANs beyond the distribution layer.

Option 1: Connect distribution layer switch to WAN router

When the LAN distribution layer connects to a WAN router this may present a number of scenarios:

- The distribution layer switch is a collapsed core HQ location connecting to one or more WAN headend routers.
- The distribution layer switch is collapsed core for a larger remote site with multiple WAN routers for survivability.
- The distribution layer switch is a WAN aggregation switch with a number of WAN headend routers connected to it for a modular block connecting to a LAN Core switch.

Because of the number of combinations, further investigation may be necessary to adjust for the LAN connectivity that matches your deployment scenario.

Option 2: Connect distribution layer switch to LAN core switch

Step 1: Configure the Layer 3 interface.

If you are using an EtherChannel to connect to the LAN core, the interface type is PortChannel and the number must match the channel-group number you will configure in Step 3. When configuring a Layer 3 EtherChannel, the logical port-channel interface is configured prior to configuring the physical interfaces associated with the EtherChannel.

```
interface [interface type] [number]
  description Link to {your device here}
  no switchport
  ip address [ip address] [mask]
  ip pim sparse-mode
  logging event link-status
  carrier-delay msec 0
  load-interval 30
  no shutdown
```

If the interface type is not a port-channel, then an additional command **macro apply EgressQoS** must also be configured on the interface.

Step 2: If the routing protocol you are using is OSPF, you add the router neighbor authentication configuration to the interface. The chosen password must match the neighbor peer, and you do additional OSPF authentication configuration in a later step.

```
interface [interface type] [number]
 ip ospf message-digest-key 1 md5 [neighbor key]
```

Step 3: If you want to run EtherChannel links to the core layer, configure the EtherChannel member interfaces.

Configure the physical interfaces to tie to the logical port-channel using the **channel-group** command. The number for the port-channel and channel-group must match.

Also, apply the egress QoS macro that was defined in the platform configuration procedure to ensure that traffic is prioritized appropriately.

Cisco Catalyst 4500 Series Switches do not use the **logging event bundle-status** command.

```
interface [interface type] [port 1]
  description Link to {your device here} Port 1
interface [interface type] [port 2]
  description Link to {your device here} Port 2
interface [interface type] [port 3]
  description Link to {your device here} Port 3
interface [interface type] [port 4]
  description Link to {your device here} Port 4
!
interface range [interface type] [port 1], [interface type] [port 2], [interface
type] [port 3], [interface type] [port 4]
  no switchport
  carrier-delay msec 0
  channel-protocol lacp
  channel-group [number] mode active
```

```
        logging event link-status
        logging event trunk-status
        logging event bundle-status
        load-interval 30
        macro apply EgressQoS
        no shutdown
```

> **i** **Tech Tip**
>
> The Cisco Catalyst 6500 Series Switches have two egress QoS macros:
> EgressQoSOneGig, which is used for Gigabit Ethernet ports, and EgressQoS, which is
> used for 10-Gigabit or 40-Gigabit Ethernet ports. All other distribution layer platforms
> have a single egress QoS macro, which applies to all Ethernet ports, including Gigabit
> Ethernet and 10-Gigabit Ethernet.

Step 4: Configure IP address summarization on the links to the core.

As networks grow, the number of IP subnets or routes in the routing tables grows as well. You configure IP
summarization on links where logical boundaries exist in order to reduce the amount of bandwidth, processor
speed, and memory necessary to carry large route tables and to reduce convergence time around a link failure.
If the connected device provides connectivity to another piece of the network (for example, the WAN, Internet, or
LAN core), configure summarization.

EIGRP Summarization
```
        router eigrp LAN
         address-family ipv4 unicast autonomous-system 100
          af-interface [interface type] [number]
           summary-address 10.4.0.0 255.255.240.0
          exit-af-interface
         exit-address-family
```

OSPF Summarization
```
        router ospf 100
         area [unique area number] range 10.4.0.0 255.255.240.0
```

> **i** **Tech Tip**
>
> If you have multiple address blocks that are not contiguous and can't be aggregated
> into a single summary, you can add additional EIGRP summary-address statements or
> OSPF area range statements, as appropriate. If this is a distribution switch at a WAN
> remote site, you need to complete additional steps on the remote-site router for any
> summary address, as described in the related WAN design guide.

Step 5: Configure router neighbor authentication, and override passive interface configuration for links to the
core.

After you have configured the Layer 3 interfaces and Layer 3 port-channels connecting to other Layer 3 devices,
allow the routing protocol to form neighbor relationships with MD5 authentication across these interfaces to
establish peering adjacencies and exchange route tables.

Unlike EIGRP named mode configuration, OSPF neighbor authentication also requires a configuration attached directly to the Layer-3 interfaces, thus additional router neighbor authentication configuration is completed as part of the prior Layer-3 interface configuration steps.

EIGRP Neighbor Authentication

```
key chain EIGRP-KEY
 key 1
  key-string [neighbor key]
!
router eigrp LAN
 address-family ipv4 unicast autonomous-system 100
  af-interface [interface type] [number]
   authentication mode md5
   authentication key-chain EIGRP-KEY
   no passive-interface
  exit-af-interface
 exit-address-family
```

OSPF Neighbor Authentication

```
router ospf 100
 area 0 authentication message-digest
 no passive-interface [interface type] [number]
```

Step 6: Save the running configuration that you have entered so it will be used as the startup configuration file when your switch is reloaded or power-cycled.

```
copy running-config startup-config
```

Example: Distribution to Core PortChannel configuration–EIGRP
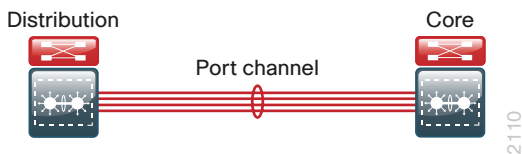


```
interface Port-channel 30
  description EtherChannel Link to Core Switch
  no switchport
  ip address 10.4.40.10 255.255.255.252
  ip pim sparse-mode
  load-interval 30
  no shutdown
!
interface range FortyGigabitEthernet 1/2/1, FortyGigabitEthernet 2/2/1,
FortyGigabitEthernet 1/3/1, FortyGigabitEthernet 2/3/1
  description EtherChannel Link to Core Switch
  no switchport
  carrier-delay msec 0
  channel-group 30 mode active
```

```
      logging event link-status
      logging event trunk-status
      logging event bundle-status
      load-interval 30
      macro apply EgressQoS
      no shutdown
 !
 key chain EIGRP-KEY
  key 1
   key-string [neighbor key]
 router eigrp LAN
  address-family ipv4 unicast autonomous-system 100
   af-interface default
    passive-interface
   exit-af-interface
   af-interface Port-channel30
    summary-address 10.4.0.0 255.255.240.0
    authentication mode md5
    authentication key-chain EIGRP-KEY
    no passive-interface
   exit-af-interface
   network 10.4.0.0 0.1.255.255
   eigrp router-id 10.4.40.254
   eigrp stub summary
   nsf
  exit-address-family
 !
```

Example: Distribution to Core PortChannel configuration—OSPF


```
 interface Port-channel 30
   description EtherChannel Link to Core Switch
   no switchport
   ip address 10.4.40.10 255.255.255.252
   ip pim sparse-mode
   ip ospf message-digest-key 1 md5 [neighbor key]
   no shutdown
 !
 interface range FortyGigabitEthernet 1/2/1, FortyGigabitEthernet 2/2/1,
 FortyGigabitEthernet 1/3/1, FortyGigabitEthernet 2/3/1
   description EtherChannel Link to Core Switch
   no switchport
   carrier-delay msec 0
```

```
        channel-group 30 mode active
        logging event link-status
        logging event trunk-status
        logging event bundle-status
        load-interval 30
        macro apply EgressQoS
        no shutdown
    !
    !
    router ospf 100
     router-id 10.4.15.254
     nsf
     area 0 authentication message-digest
     area 1 stub no-summary
     area 1 range 10.4.0.0 255.255.240.0
     passive-interface default
     no passive-interface Port-channel30
     network 10.4.0.0 0.0.15.255 area 1
     network 10.4.40.0 0.0.0.255 area 0
```
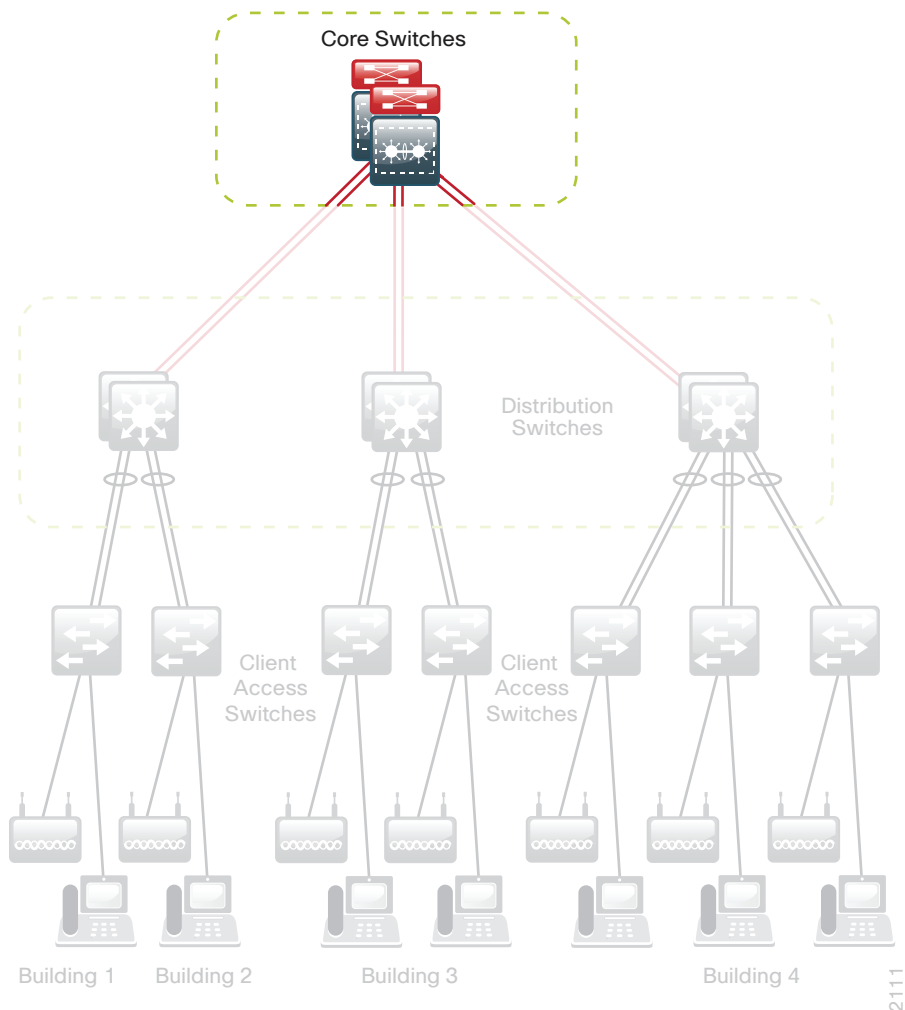
Core Layer

Design Overview

The core layer of the LAN is a critical part of the scalable network, yet by design, is one of the simplest. Like the distribution layer aggregates connectivity for multiple access layer switches, the core layer aggregates connectivity when there are multiple distribution blocks. As networks grow beyond three distribution blocks in a single location, a core layer should be used to optimize the design.

Beyond the simple aggregation of connectivity, the core layer serves to reduce the number of paths between distribution layers, which in turn lowers the time required to converge the network after a failure. By upgrading bandwidth between a distribution layer and the core, multiple distribution layer blocks can benefit from the increase versus the need to upgrade the bandwidth to every other device in a design without a core. The core layer is especially relevant to designs where the data center resources might be colocated with the LAN.

Figure 26 - Core layer overview

In large modular and scalable LAN designs, a core layer is used to aggregate multiple user connectivity distribution layer blocks and network-services distribution layer blocks. In designs with a colocated data center, the core provides high bandwidth fan-out connectivity to the rest of the network. The core layer also serves as the connection between the Wide Area Network (WAN) and Internet edge distribution layer blocks. Because of this central point of connectivity for all data flows, the core is part of the backbone IP routing address space and is designed to be highly resilient to protect from component-, power-, or operational-induced outages. The core layer should not contain highly complex or high touch services that require constant care and tuning, to avoid downtime required by complex configuration changes, increased software upgrades for new services, or links that toggle up/down as part of normal operations like user endpoint connectivity.

The core layer in this design is based on two physically separate switches which behave as a single logical device using Cisco Virtual Switching System (VSS). Although traditional core designs using two independent layer-3 platforms are valid design choices, using VSS in the core offers a number of optimizations for performance, confuration simplicity, and resiliency. For example, configuration of an RP in the core is inherently resilent, without the need to introduce additional protocols such as MSDP. Connectivity to and from the core should be Layer 3 only. No VLANs should span the core to drive increased resiliency and stability.

The core is built on dual switches to provide redundant logic, line cards, hardware, and power for the backbone operation. Each distribution layer block, router, or other appliance connecting to the core should be connected with an EtherChannel with at least one link to each core switch. This approach provides load sharing of IP traffic across links of the multichassis EtherChannel for traffic traversing the core, and fast failover based on EtherChannel without waiting for routing protocol topology changes to propagate the network.

The core is designed to be high bandwidth and provides for connectivity ranging from Gigabit Ethernet to 40-Gigabit Ethernet, and EtherChannel up to 40-Gigabit EtherChannel. The core can provide non-blocking bandwidth based on design and configuration. EtherChannel links homed to a switch should be spread across line cards when possible.

The supervisor modules for the core switches in a VSS pair operate in an active/standby mode for Stateful Switchover (SSO) operation to protect the core operation in the event that a control plane hardware or software failure occurs. The core switches are Nonstop Forwarding (NSF) aware to provide enhanced resilience for any dual supervisor connected devices and the VSS pair is NSF capable.

Core Layer Platforms

Cisco Catalyst 6807-XL VSS with Supervisor Engine 2T

- Cisco Catalyst 6807-XL VSS uses two physical chassis with Cisco Catalyst 6500 Supervisor Engine 2T, which offers a per slot switching capacity of 220 Gbps in the Cisco Catalyst 6807-XL chassis, and delivers hardware-enabled scalability and features. The performance is increased over the same Supervisor Engine in the 6500-E chassis, and enables the system to provide 40-Gigabit Ethernet connectivity, and growth capability for future 100-gigabit Ethernet connectivity.

- Adding an additional Cisco Catalyst 6500 Supervisor Engine 2T to each chassis in the VSS pair for a total of four supervisors creates a Quad-Supervisor SSO (VS4O) configuration, offering the ability to have an in-chassis standby supervisor capability. The in-chassis standby enables Enhanced Fast Software Upgrades (eFSU) for minimal downtime during software upgrades, along with the ability to recover from a degraded state of performance upon loss of a supervisor, without human intervention.

- Cisco 6500 Supervisor Engine 2T supports the line cards with Distributed Forwarding Card 4-E (DFC4-E), including the WS-X6816-10G, WS-X6908-10G, and WS-X6904-40G-2T, which provide enhanced hardware capabilities. The WS-X6908-10G provides eight 10-Gbps Ethernet ports with 1:1 oversubscription. The WS-X6904-40G-2T provides up to four 40-Gbps Ethernet ports or up to sixteen 10-Gbps Ethernet ports using modular adapters and can be programmed to run in 2:1 or 1:1 oversubscription mode.

- The Supervisor Engine 2T supports DFC4-A based line cards, including the WS-X6824 and WS-X6848, to provide gigabit Ethernet ports. The WS-X6724 and WS-X6748 gigabit Ethernet cards are also supported when installed with CFC or DFC4-A modules.

- The Cisco Supervisor Engine 2T-based switch enhances support for Cisco TrustSec (CTS) by providing MacSec encryption and role-based access control lists (RBACL), and delivers improved control plane policing to address denial-of-service attacks.

- VSS effectively allows the merging of two physical chassis into a logical entity that can be operated as a single device. This configuration provides redundant chassis, supervisors, line cards, and power supplies and can provide the highest density of the product options for Gigabit Ethernet, 10 Gigabit Ethernet, and 40-Gigabit EtherChannel uplinks using Cisco Multi-chassis EtherChannel (MEC).

- Provides Stateful Switch-Over (SSO) to synchronize infrastructure and forwarding state between chassis, along with Non-Stop Forwarding (NSF) for graceful-restart of L3 routing protocols, in the event of a chassis failure. Also allows Enhanced Fast Software Upgrades (EFSU) with In-Service Software Upgrades (ISSU) for minimizing downtime for system upgrades.

- The Cisco Catalyst 6807-XL chassis with the Supervisor Engine 2T is the premier core layer platform. It allows for high density aggregation of wiring closets connected with Gigabit Ethernet and 10-Gigabit Ethernet, while providing an advanced feature set and the highest resiliency available.

Cisco Catalyst 6500-E VSS with Supervisor Engine 2T

Cisco Catalyst 6500 VSS uses two physical chassis with Cisco Catalyst 6500 Supervisor Engine 2T, which offers a per slot switching capacity of 80 Gbps. The Cisco Catalyst 6500-E is an available alternative for the core layer VSS chassis with performance that can provide 40-Gigabit Ethernet connectivity.

Deployment Details

The core layer design uses a Cisco Catalyst 6500 VSS or Cisco Catalyst 6807-XL VSS with quad-supervisor SSO for resiliency.

PROCESS

Configuring the Core

1. Configure the platform
2. Configure LAN switch universal settings
3. Configure the core switch global settings and IP unicast routing
4. Configure IP Multicast routing
5. Connect to the distribution layer

Procedure 1   Configure the platform

Cisco Catalyst 6500-E and Cisco Catalyst 6807-XL Virtual Switching System merges two physical 6500-E or 6807-XL switches together as a single logical switch, using a single or optionally dual Cisco Supervisor Engine 2T modules in each physical switch One of the supervisors acts as the active control plane for both chassis by controlling protocols such as EIGRP and OSPF, Spanning Tree, CDP, and so forth, while both supervisors actively switch packets in each chassis.

The following configuration example shows you how to convert two standalone Cisco Catalyst 6500 or 6807-XL switches to a Virtual Switching System (VSS). If you are migrating your switches from an existing in-service dual chassis role to a VSS system, go to www.cisco.com and search on "Migrate Standalone Cisco Catalyst 6500 Switch to Cisco Catalyst 6500 Virtual Switching System" for information that describes how to do this migration. For an in-depth VSS configuration guide and configuration options, go to www.cisco.com/go/cvd/campus and, on the Cisco Validated Designs tab, look for the *Campus 3.0 Virtual Switching System Design Guide.*

When you set up the Virtual Switching System, connect 10-Gigabit or 40-Gigabit Ethernet links between the chassis to provide the Virtual Switch Link (VSL). Always use at least two links.

You can use up to eight links as members of the VSL, and the links should be distributed across Supervisor Engines or line cards for resiliency. This design uses the two 10-Gigabit Ethernet interfaces on each installed Supervisor Engine. You connect the VSL interfaces together before you configure the VSS.

To aid in understanding the connections between switches supporting the VSS configuration, the following tables can be used. The optional Fast-Hello connection is used in this design, but can be replaced or augmented by functionality available when implementing Ehanced PAgP, which is not covered in this release. The PortChannel interface numbers are arbitrary and should be adapted to best suit your deployment. The chosen PortChannel numbers reflect the highest values supported across all VSS platforms validated in this release.

Table 10 - Example VSS connections for the Cisco 6509-E chassis pair with two Supervisor Engine 2T

VSS connection	VSS Switch 1 Port (PortChannel)	VSS Switch 2 Port (PortChannel)
10-Gbps, VSL 1	Ten5/4 (Po63)	Ten5/4 (Po64)
10-Gbps, VSL 2	Ten5/5 (Po63)	Ten5/5 (Po64)
1-Gbps, Fast-Hello	Gig9/24	Gig9/24

For additional resiliency and to allow support for an Enhanced Fast Software Upgrade (eFSU), a VSS quad-supervisor Stateful Switchover (VS4O) setup can be created. The four Supervisors must have fully meshed connectivity to support minimal traffic disruption during software upgrades, as shown in the following example.

Table 11 - Example VSS connections for the Cisco 6807-XL chassis pair with four Supervisor Engine 2T, meshed VSL connections

VSS connection	VSS Switch 1 Port (PortChannel)	VSS Switch 2 Port (PortChannel)
10-Gbps, VSL 1	Ten3/4 (Po63)	Ten3/4 (Po64)
10-Gbps, VSL 2	Ten3/5 (Po63)	Ten4/4 (Po64)
10-Gbps, VSL 3	Ten4/4 (Po63)	Ten3/5 (Po64)
10-Gbps, VSL 4	Ten4/5 (Po63)	Ten4/5 (Po64)
1-Gbps, Fast-Hello	Gig7/48	Gig7/48

Step 1: If you are configuring a VSS quad-supervisor VS4O system, ensure that all four Supervisor Engines are available and the same version of code is installed and set to boot for each one of them. To make configuration easier by limiting the possible active console connections, leave only one Supervisor Engine fully inserted in each chassis.

⟨∞⟩ Reader Tip

The supported code used for this configuration and validation of all devices is listed in the appendix of this guide.

Step 2: If you are using a Cisco Catalyst 6800 or 6500 Series chassis with the Cisco Catalyst 6900 Series 40-Gigabit Ethernet Interface Module with FourX adapters to convert CFP ports into four 10-Gigabit Ethernet ports, configure the switch to enable the line card to use 10-Gigabit Ethernet functionality for the associated port-group. If you configure this after converting to VSS, the **switch** parameter is required. The port-group can be **1**, **2**, or **all**.

```
D6500-VSS(config)# hw-module switch [switch] slot [slot] operation-mode port-
group [port-group] TenGigabitEthernet
```

Example—command used before VSS conversion

```
D6500-VSS(config)#  hw-module slot 1 operation-mode port-group all
TenGigabitEthernet
```

Example—command used after VSS conversion

```
D6500-VSS(config)#  hw-module switch 1 slot 1 operation-mode port-group
TenGigabitEthernet
```

Confirm the change. The module resets and the port configuration is cleared.

Step 3: Convert standalone Cisco Catalyst 6807-XL or 6500-E Series Switches to VSS.

Configure a temporary hostname on each switch so you can keep track of your configuration steps. In a later step after the conversion is complete, you apply a replacement hostname to the merged VSS configuration.
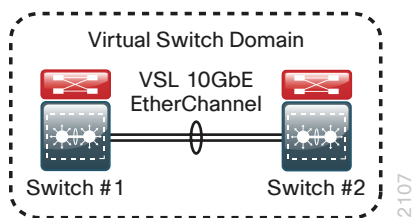
On the standalone switch #1:

```
Router#config t
Router#(config)#hostname VSS-Sw1
```

On the standalone switch #2:

```
Router#config t
Router#(config)#hostname VSS-Sw2
```

To form a VSS pair, each switch in the pair must have a matching domain ID assigned. To support the interconnection of multiple VSS pairs, the domain ID selected for the pair should be unique. In this example, the domain number is 1. Each switch is also given a unique identifier within the domain, switch 1 or switch 2.

Figure 27 - VSS domain

On the standalone switch #1:

```
VSS-Sw1(config)#switch virtual domain 101
VSS-Sw1(config-vs-domain)# switch 1
VSS-Sw1(config-vs-domain)# exit
VSS-Sw1(config)#
```

On the standalone switch #2:

```
VSS-Sw2(config)#switch virtual domain 101
VSS-Sw2(config-vs-domain)# switch 2
VSS-Sw2(config-vs-domain)# exit
VSS-Sw2(config)#
```

Step 4: Configure the Virtual Switch Link (VSL).

The VSL is a critical component of the Virtual Switching System. For each physical switch you must select a unique port-channel number identifying the same VSL. This allows the switch to maintain a separate identity for the interfaces used when making traffic forwarding decisions. This example uses port-channel number 63 on switch 1 and port-channel number 64 on switch 2. The PortChannel interface numbers are arbitrary and should be adapted to best suit your deployment. The chosen PortChannel numbers reflect the highest values supported across all VSS platforms validated in this release. You must configure **channel-group mode on** for the VSL port channel because it is an infrastructure link actively managed within the VSS using Virtual Switch Link Protocol (VSLP). This example uses the 10-Gigabit Ethernet interfaces on the supervisor of a Cisco Catalyst 6500-E Series switch for the EtherChannel member ports of the VSL.

On standalone switch #1:

```
VSS-Sw1(config)#interface port-channel 63
VSS-Sw1(config-if)#switch virtual link 1
VSS-Sw1(config-if)#no shutdown
VSS-Sw1(config)#interface range tengigabit 5/4-5
VSS-Sw1(config-if)#channel-group 63 mode on
VSS-Sw1(config-if)#no shutdown
```

On standalone switch #2:

```
VSS-Sw2(config)#interface port-channel 64
VSS-Sw2(config-if)#switch virtual link 2
VSS-Sw2(config-if)#no shutdown
VSS-Sw2(config)#interface range tengigabit 5/4-5
VSS-Sw2(config-if)#channel-group 64 mode on
VSS-Sw2(config-if)#no shutdown
```

At this point you should be able to see that port-channel 63 and 64 are up, and both links are active on standalone switch #1 and standalone switch #2, respectively. The switches are not in VSS mode yet.

```
VSS-Sw1# show etherchannel 63 port
VSS-Sw2# show etherchannel 64 port
```

The previous two commands show the same output below.

```
        Ports in the group:
        -------------------
Port: Te5/4
------------
Port state  = Up Mstr In-Bndl
…
Port: Te5/5
------------
Port state  = Up Mstr In-Bndl
…
```

Step 5: Enable virtual switch mode operation.

Now that a port-channel has been established between the switches, convert each switch to virtual mode operation. At the enable prompt (that is, not in configuration mode) on each switch, enter the following commands for each switch.

On standalone switch #1:

```
VSS-Sw1# switch convert mode virtual
```

On standalone switch #2:

```
VSS-Sw2# switch convert mode virtual
```

When asked if you want to proceed, answer yes.

```
This command will convert all interface names
to naming convention "interface-type switch-number/slot/port",
save the running config to startup-config and
reload the switch.

NOTE: Make sure to configure one or more dual-active detection methods
once the conversion is complete and the switches have come up in VSS mode.

Do you want to proceed? [yes/no]: yes
```

Each switch now renumbers its interfaces from interface y/z (where y is the slot number and z is the interface number) to interface x/y/z (where x is the switch number, y is the module number in that switch, and z is the interface on that module). This numbering scheme allows the two chassis to be addressed and configured as a single system from a single supervisor, which is the supervisor with the active control plane.

Once the configuration changes, it prompts you to save the configuration to bootflash. Accept the destination filename and location on each switch by pressing Return <CR> or Enter.

Both switches reload. The switch pair negotiates using VSLP over the VSL and becomes a VSS, with one of the switches resolved as the ACTIVE supervisor for the merged VSS switch. All configuration commands now must be entered on the single active switch console. The other physical chassis in the VSS pair contains the STANDBY HOT supervisor with a console port that displays the Standby prompt.

Verify that both switches can see each other, that they are in SSO mode, and that the second supervisor is in STANDBY HOT status.

```
VSS-Sw1#show switch virtual redundancy
```

Confirm that the two Cisco Catalyst 6500 Series Switches are now operating as a single VSS system by using configuration mode to rename the switch hostname.

```
VSS-Sw1(config)#hostname C6500-VSS
C6500VSS(config)#
```

A critical aspect of the Cisco Catalyst VSS is the control plane and data plane operating models. From a control plane standpoint the VSS uses an active-standby operating model. This means that supervisor hardware on one chassis becomes the active control plane for the entire VSS while the other supervisor hardware on the paired chassis becomes the standby. The control plane handles protocol operations like IP routing, peering, route table updates, and spanning tree BPDUs. The dataplane handles the hardware forwarding of packets, and both switches are actively forwarding traffic in an active-active operating model.

The VSL allows the switches to communicate and stay in synchronization. The VSS uses the Stateful Switchover (SSO) redundancy facility to keep the control plane synchronized between the two switches. As a result, the VSS appears to devices in adjoining layers as a single switch with a single MAC address.

Step 6: Configure dual-active detection mechanism.

In the event that the VSL is severed (that is, all links are down), or for any reason communication is lost over the VSL (such as excessive high CPU utilization), both supervisors would assume the active control plane role, thus creating a dual-active condition, that can result in network instability. To prevent a dual-active scenario from causing an outage in the network, VSS supports multiple unique dual-active detection and recovery mechanisms.

The dual-active detection mechanisms are used to trigger a VSS recovery mode. In the VSS recovery mode only one switch chassis is allowed to remain active, the other switch (the previous VSS active switch) enters recovery mode, and shuts down all of its interfaces except the VSL interfaces, thereby preventing instability in the network. Once the VSL is repaired, and communication over the VSL is reestablished, then the VSS reloads the switch in recovery mode and returns the VSS to a normal operating state.

You can use the following methods to detect this dual-active condition:

- Virtual Switch Link Protocol (VSLP) Ethernet Fast-Hello link
- Enhanced Port Aggregation Protocol (PAgP) hellos with an adjacent switch

This design uses the Fast-Hello (VSLP) link for dual-active detection. Configure the link by using a Gigabit Ethernet interface on each VSS switch chassis and cabling them together (similar to a VSL connection) in a back-to-back fashion. This link does not require high bandwidth because it is only a detection link with control plane hellos on it.

Figure 28 - VSLP
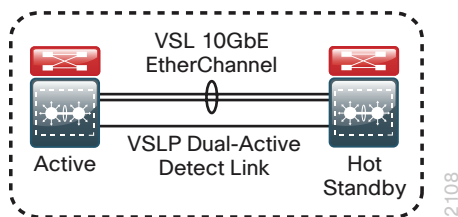
```
C6500-VSS(config)#interface range gigabit1/2/48, gigabit2/2/48
C6500-VSS(config-if-range)#dual-active fast-hello
C6500-VSS(config-if-range)#no shutdown
%VSDA-SW1-5-LINK_UP: Interface Gi1/1/24 is now dual-active detection capable
%VSDA-SW2_STBY-5-LINK_UP: Interface Gi2/1/24 is now dual-active detection capable
```

Step 7: Configure the system virtual MAC address.

By default, the VSS system uses the default chassis-based MAC-address pool assigned to the switch that is resolved to be the active switch when the switches initialize. As a result of events such as stateful switchover, the MAC may change. Set a virtual MAC address for the VSS system so that either active supervisor will use the same MAC address pool, regardless of which supervisor is active, even across a system reload.

```
C6500-VSS(config)# switch virtual domain 101
C6500-VSS(config-vs-domain)# mac-address use-virtual
Configured Router mac address is different from operational value. Change will
take effect after the configuration is saved and the entire Virtual Switching
System (Active and Standby) is reloaded.
```

Step 8: Save the running configuration, and then reload the entire system (both chassis).

```
copy running-config startup-config
reload
```

When the switches initialize after this final reload, the VSS configuration for dual Supervisor Engines is complete.

Step 9: If you are configuring a VSS quad-supervisor VS4O system, allow the system to fully boot. Then insert the additional Supervisor Engine for each chassis, which has been previously prepared to boot the same version of code.

Because the Supervisor Engines were prepared to boot the same version of code that is already active on the VSS, the Supervisor Engines boot and communicate with the In-Chassis Active (ICA) Supervisor Engine and receive the appropriate VSS variable configuration. A one-time reload automatically initiates on the inserted Supervisor Engines and the modules boot as full participants in the VS4O platform.

After the Supervisor Engines have fully booted, you can observe the VS4O status.

```
C6500VSS#show switch virtual redundancy
```

The command output shows which switch and slot contains each Supervisor Engine in ACTIVE mode and STANDBY HOT mode, and the two Supervisors in the STANDBY HOT (CHASSIS) mode.

If you are using ports on the new Supervisor Engines for additional VSL connections as described in the tables before the first step, you now complete the connections, ensuring that you always have at least one VSL connection active. Once the physical connections are complete, add the connections into the existing VSL.

```
D6500-VSS(config)#interface range tengigabit 1/6/4-5
D6500-VSS(config-if)#channel-group 63 mode on
D6500-VSS(config-if)#no shutdown
D6500-VSS(config)#interface range tengigabit 2/6/4-5
D6500-VSS(config-if)#channel-group 64 mode on
D6500-VSS(config-if)#no shutdown
```

Verify that the status of all of the links in the VSL PortChannel are **P - bundled in port-channel** status.

```
D6500-VSS#show switch virtual link port-channel
```

Step 10: Configure QoS.

On the Cisco Catalyst 6500 Supervisor Engine 2T based switches, QoS is enabled by default and policies for interface queuing are defined by attached service policies. The QoS policies are now defined using Cisco Common Classification Policy Language (C3PL), which is similar to Modular QoS CLI, in order to reduce operational complexity.

All interface connections in the distribution and core are set to trust differentiated services code point (DSCP) markings. Even though this design is configured to trust DSCP markings, it is a best practice to ensure proper mapping of CoS to DSCP for VoIP. This mapping is accomplished by overriding the default mapping of CoS 5 "voice bearer traffic" to DSCP 40, with DSCP 46, which is the EF per-hop behavior for voice.

This egress QoS policy is configured to accommodate the 10-Gigabit and 40-Gigabit Ethernet cards which use a 1P7Q4T queuing architecture.

```
! Enable port-based QoS
auto qos default
! Class maps for 1P7Q4T 10Gbps and 40Gbps ports service policy
class-map type lan-queuing match-any PRIORITY-QUEUE
  match dscp ef
  match dscp cs5
  match dscp cs4
  match cos 5
class-map type lan-queuing match-any CONTROL-MGMT-QUEUE
  match dscp cs7
  match dscp cs6
  match dscp cs3
  match dscp cs2
  match cos 3 6 7
```

```
class-map type lan-queuing match-any MULTIMEDIA-CONFERENCING-QUEUE
  match dscp af43
  match dscp af42
  match dscp af41
  match cos 4
class-map type lan-queuing match-any MULTIMEDIA-STREAMING-QUEUE
  match dscp af33
  match dscp af32
  match dscp af31
class-map type lan-queuing match-any TRANSACTIONAL-DATA-QUEUE
  match dscp af23
  match dscp af22
  match dscp af21
  match cos 2
class-map type lan-queuing match-any BULK-DATA-QUEUE
  match dscp af13
  match dscp af12
  match dscp af11
class-map type lan-queuing match-any SCAVENGER-QUEUE
  match dscp cs1
  match cos 1
!
policy-map type lan-queuing 1P7Q4T
 class PRIORITY-QUEUE
  priority
 class CONTROL-MGMT-QUEUE
  bandwidth remaining percent 14
  queue-buffers ratio 10
class MULTIMEDIA-CONFERENCING-QUEUE
  bandwidth remaining percent 14
  queue-buffers ratio 10
  random-detect dscp-based
  random-detect dscp 38 percent 70 100
  random-detect dscp-based
  random-detect dscp 36 percent 80 100
  random-detect dscp-based
  random-detect dscp 34 percent 90 100
 class MULTIMEDIA-STREAMING-QUEUE
  bandwidth remaining percent 14
  queue-buffers ratio 10
  random-detect dscp-based
  random-detect dscp 30 percent 70 100
  random-detect dscp-based
  random-detect dscp 28 percent 80 100
  random-detect dscp-based
  random-detect dscp 26 percent 90 100
 class TRANSACTIONAL-DATA-QUEUE
```

```
   bandwidth remaining percent 14
   queue-buffers ratio 10
   random-detect dscp-based
   random-detect dscp 22 percent 70 100
   random-detect dscp-based
   random-detect dscp 20 percent 80 100
   random-detect dscp-based
   random-detect dscp 18 percent 90 100
  class BULK-DATA-QUEUE
   bandwidth remaining percent 6
   queue-buffers ratio 10
   random-detect dscp-based
   random-detect dscp 14 percent 70 100
   random-detect dscp-based
   random-detect dscp 12 percent 80 100
   random-detect dscp-based
   random-detect dscp 10 percent 90 100
  class SCAVENGER-QUEUE
   bandwidth remaining percent 2
   queue-buffers ratio 10
   random-detect dscp-based
   random-detect dscp 8 percent 80 100
 class class-default
   queue-buffers ratio 25
   random-detect dscp-based aggregate
   random-detect dscp values 0 1 2 3 4 5 6 7 percent 80 100
   random-detect dscp values 9 11 13 15 17 19 21 23 percent 80 100
   random-detect dscp values 25 27 29 31 33 35 37 39 percent 80 100
   random-detect dscp values 41 42 43 44 45 47 49 50 percent 80 100
   random-detect dscp values 51 52 53 54 55 57 58 59 percent 80 100
   random-detect dscp values 60 61 62 63 percent 80 100
 !
 table-map cos-discard-class-map
  map from 0 to 0
  map from 1 to 8
  map from 2 to 16
  map from 3 to 24
  map from 4 to 32
  map from 5 to 46
  map from 6 to 48
  map from 7 to 56
 !
 macro name EgressQoS
   service-policy type lan-queuing output 1P7Q4T
 @
```

Step 11: If you are using Gigabit Ethernet cards supported in VSS mode on Cisco Catalyst 6500-E and 6807-XL Supervisor Engine 2T based switches, configure an additional QoS policy for the Gigabit Ethernet ports.

A separate egress QoS policy is configured to accommodate the Gigabit Ethernet cards which use a 1P3Q8T queuing architecture. This policy does not apply to the Cisco Catalyst 6880-X platforms.

```
! Class maps for 1P3Q8T 1Gb ports service policy
class-map type lan-queuing match-any PRIORITY-QUEUE-GIG
 match cos 5 4
class-map type lan-queuing match-any CONTROL-AND-STREAM-MEDIA
 match cos 7 6 3 2
class-map type lan-queuing match-any BULK-DATA-SCAVENGER
 match cos 1
!
policy-map type lan-queuing 1P3Q8T
 class PRIORITY-QUEUE-GIG
  priority
  queue-buffers ratio 15
 class CONTROL-AND-STREAM-MEDIA
  bandwidth remaining percent 55
  queue-buffers ratio 40
class BULK-DATA-SCAVENGER
  bandwidth remaining percent 10
  queue-buffers ratio 20
  random-detect cos-based
  random-detect cos 1 percent 80 100
 class class-default
  queue-buffers ratio 25
  random-detect cos-based
  random-detect cos 0 percent 80 100
!
macro name EgressQoSOneGig
  service-policy type lan-queuing output 1P3Q8T
@
```

In this design, there are features and services that are common across all LAN switches, regardless of the type of platform or role in the network. These are system settings that simplify and secure the management of the solution.

This procedure provides examples for some of these settings. The actual settings and values depend on your current network configuration.

Table 12 - Common network services used in the design examples

Setting	Value
Domain Name	cisco.local
Active Directory, DNS, DHCP Server	10.4.48.10
Authentication Control System	10.4.48.15
Network Time Protocol Server	10.4.48.17
EIGRP Named Mode Configuration Name	CAMPUS
EIGRP AS or OSPF AS	100
Multicast Range	239.1.0.0/16

Step 1: Configure the device hostname to make it easy to identify the device.

```
hostname [hostname]
```

Step 2: If the switch VTP mode has been changed from the default, configure VTP transparent mode. This design uses VTP transparent mode because the benefits of dynamic propagation of VLAN information across the network are not worth the potential for unexpected behavior resulting from operational error.

VLAN Trunking Protocol (VTP) allows network managers to configure a VLAN in one location of the network and have that configuration dynamically propagate out to other network devices. However, in most cases, VLANs are defined once during switch setup with few, if any, additional modifications.

```
vtp mode transparent
```

Step 3: Enable Rapid Per-VLAN Spanning-Tree (PVST+). Rapid PVST+ provides an instance of RSTP (802.1w) per VLAN. Rapid PVST+ greatly improves the detection of indirect failures or linkup restoration events over classic spanning tree (802.1D).

Although this architecture is built without any Layer 2 loops, you should still enable spanning tree with the most up-to-date network safeguards. By enabling spanning tree, you ensure that if any physical or logical loops are accidentally configured, no actual layer 2 loops occur.

```
spanning-tree mode rapid-pvst
```

Step 4: Enable Unidirectional Link Detection (UDLD) as the default for fiber ports.

UDLD is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. When UDLD detects a unidirectional link, it disables the affected interface and alerts you. Unidirectional links can cause a variety of problems, including spanning-tree loops, black holes, and non-deterministic forwarding. In addition, UDLD enables faster link failure detection and quick reconvergence of interface trunks, especially with fiber, which can be susceptible to unidirectional failures.

```
udld enable
```

Step 5: Enable the recovery mechanism to allow ports disabled as a result of errors to automatically clear the err-disable status and attempt a recovery to operational behavior and connected status. Enabling the recovery mechanism avoids having to intervene manually by using the CLI to shutdown and enable the port after the cause of the error is cleared. By default, the recovery mechanism waits five minutes to attempt clearing of the interface err-disable status.

```
errdisable recovery cause all
```

Step 6: Set EtherChannels to use the traffic source and destination IP address when calculating which link to send the traffic across. This normalizes the method in which traffic is load-shared across the member links of the EtherChannel. EtherChannels are used extensively in this design because of their resiliency capabilities.

```
port-channel load-balance src-dst-ip
```

Step 7: Configure DNS for host lookup.

At the command line of a Cisco IOS device, it is helpful to be able to type a domain name instead of the IP address for a destination.

```
ip name-server 10.4.48.10
```

Step 8: Configure device management protocols.

Secure HTTP (HTTPS) and Secure Shell (SSH) are more secure replacements for the HTTP and Telnet protocols. They use Secure Sockets Layer (SSL) and Transport Layer Security (TLS) to provide device authentication and data encryption.

The SSH and HTTPS protocols enable secure management of the LAN device. Both protocols are encrypted for privacy, and the unencrypted protocols, Telnet and HTTP, are turned off. Enabling HTTPS automatically generates a cryptographic key to use the service. When SSH is configured after HTTPS, you do not have to explicitly generate the cryptographic key that SSH requires, unless you wish to change the default key size.

Secure Copy (SCP) provides a secure and authenticated method for copying configuration and image files by making use of SSH as a secure transport. Enable SCP to allow secure file management with the device to avoid the use of less secure protocols such as TFTP and FTP.

Specify the transport preferred none on vty lines to prevent errant connection attempts from the CLI prompt. Without this command, if the ip name server is unreachable, long timeout delays may occur for mistyped commands.

```
no ip http server
ip http secure-server
ip domain-name cisco.local
ip ssh version 2
ip scp server enable
!
line vty 0 15
  transport input ssh
  transport preferred none
```

Step 9: Enable Simple Network Management Protocol (SNMP) in order to allow the network infrastructure devices to be managed by a Network Management System (NMS), and then configure SNMPv2c both for a read-only and a read-write community string.

```
snmp-server community [SNMP RO name] RO
snmp-server community [SNMP RW name] RW
```

Step 10: If your network operational support is centralized, you can increase network security by using an access list to limit the networks that can access your device. In this example, only devices on the 10.4.48.0/24 network will be able to access the device via SSH or SNMP.

```
access-list 55 permit 10.4.48.0 0.0.0.255
line vty 0 15
  access-class 55 in
!
snmp-server community [SNMP RO name] RO 55
snmp-server community [SNMP RW name] RW 55
```

> **! Caution**
>
> If you configure an access-list on the vty interface, you may lose the ability to use ssh to log in from one device to the next for hop-by-hop troubleshooting.

Step 11: Configure local login and password.

The local login account and password provides basic device access authentication to view platform operation. The enable password secures access to the device configuration mode. By enabling password encryption, you prevent the use of plain text passwords when viewing configuration files. The **aaa new-model** command enables new access control commands and functions, and causes the local username and password on the router to be used in the absence of other AAA statements.

```
username admin secret [user password]
enable secret [enable password]
service password-encryption
aaa new-model
```

By default, https access to the switch will use the enable password for authentication.

Step 12: If you want to reduce operational tasks per device, configure centralized user authentication by using the TACACS+ protocol to authenticate management logins on the infrastructure devices to the AAA server.

As networks scale in the number of devices to maintain, there is an operational burden to maintain local user accounts on every device. A centralized AAA service reduces operational tasks per device and provides an audit log of user access for security compliance and root cause analysis. When AAA is enabled for access control, all management access to the network infrastructure devices (SSH and HTTPS) is controlled by AAA.

TACACS+ is the primary protocol used to authenticate management logins on the infrastructure devices to the AAA server. A local AAA user database is also defined on each network infrastructure device to provide a fallback authentication source in case the centralized TACACS+ server is unavailable.

```
tacacs server TACACS-SERVER-1
 address ipv4 10.4.48.15
 key [tacacs key]
!
aaa group server tacacs+ TACACS-SERVERS
  server name TACACS-SERVER-1
!
aaa authentication login default group TACACS-SERVERS local
aaa authorization exec default group TACACS-SERVERS local
aaa authorization console
ip http authentication aaa
```

Step 13: Configure a synchronized clock by programming network devices to synchronize to a local NTP server in the network. The local NTP server typically references a more accurate clock feed from an outside source. Configure console messages, logs, and debug output to provide time stamps on output, which allows cross-referencing of events in a network.

```
ntp server 10.4.48.17
ntp update-calendar
!
clock timezone PST -8
clock summer-time PDT recurring
!
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
```

Procedure 3	Configure the core switch global settings and IP unicast routing

Step 1: Configure the in-band management interface.

The loopback interface for Cisco Layer 3 devices is a logical interface that is always reachable as long as the device is powered on and any IP interface is reachable to the network. Layer 3 process and features are also bound to the loopback interface to ensure resiliency of the processes. The loopback address is commonly a host address with a 32-bit address mask and has been allocated out of the core network address range. This example includes the **ip pim sparse-mode** command that will be explained further in Procedure 4.

```
interface loopback 0
 ip address [ip address] 255.255.255.255
 ip pim sparse-mode
```

Step 2: Configure the system processes to use the loopback interface address for optimal resiliency.

```
snmp-server trap-source Loopback 0
ip ssh source-interface Loopback 0
ip pim register-source Loopback0
ip tacacs source-interface Loopback0
ntp source Loopback0
```

Step 3: Configure IP unicast routing and neighbor authentication.

EIGRP Unicast Routing

Enable EIGRP for the IP address space that the network will be using. If needed for your network, you can enter multiple network statements. The Loopback 0 IP address is used for the EIGRP router ID to ensure maximum resiliency. You enable router authentication for all neighbors of the core. This allows all layer-3 devices attached to the core to form routing neighbor relationships.

```
key chain EIGRP-KEY
 key 1
  key-string [neighbor key]
router eigrp LAN
 address-family ipv4 unicast autonomous-system 100
  network 10.4.0.0 0.1.255.255
  eigrp router-id [ip address of loopback 0]
```

```
   nsf
   af-interface default
    authentication mode md5
    authentication key-chain EIGRP-KEY
   exit-af-interface
  exit-address-family
```

OSPF Unicast Routing

Enable OSPF for the IP address space that the network will be using. If needed for your network, you can enter multiple network statements. The Loopback 0 IP address is used for the OSPF router ID to ensure maximum resiliency. Unlike EIGRP named mode configuration, OSPF neighbor authentication also requires a configuration attached directly to the layer-3 interfaces, so additional router neighbor authentication configuration is done as part of a later step.

```
router ospf 100
 router-id [ip address of loopback 0]
 nsf
 area 0 authentication message-digest
 network 10.4.40.0 0.0.0.255 area 0
```

Procedure 4	Configure IP Multicast routing

IP Multicast allows a single IP data stream to be sent from a single source to multiple receivers and be replicated by the infrastructure (that is, routers and switches). Using IP Multicast is much more efficient than multiple unicast streams or a broadcast stream that would propagate everywhere. IP Telephony Music on Hold and IP Video Broadcast Streaming are two examples of IP Multicast applications.

To receive a particular IP Multicast data stream, end hosts must join a multicast group by sending an Internet Group Management Protocol (IGMP) message to their local multicast router. In a traditional IP Multicast design, the local router consults another router in the network that is acting as an RP to map receivers to active sources so they can join their streams.

The RP is a control-plane operation that should be placed in the core of the network or close to the IP Multicast sources on a pair of Layer 3 switches or routers. This design, which is based on the pim sparse mode multicast operation, uses the resiliency inherent in a VSS configuration for a simple yet scalable way to provide a highly resilient RP environment.

Step 1: Enable IP Multicast routing on the platform in the global configuration mode.

```
ip multicast-routing
```

Step 2: Configure a second loopback interface for RP functions on the core VSS switch. All routers point to this IP address on **loopback 1** for the RP. You configure the RP address from the core IP address space. Creating the RP on a second loopback interface allows for flexibility for potential RP migrations using Anycast RP operation. In the event you add a core layer to your existing network and the RP is currently configured on a distribution layer, you may want to move the RP to the core.

```
interface Loopback 1
 ip address 10.4.40.252 255.255.255.255
 ip pim sparse-mode
```

Step 3: Configure AutoRP candidate RPs.

The **send-rp-announce** command in conjunction with the **group-list** option advertises the RP address, with the multicast range the device is willing to serve, as a candidate RP to the AutoRP mapping agents.

```
access-list 10 permit 239.1.0.0 0.0.255.255
ip pim send-rp-announce Loopback1 scope 32 group-list 10
```

Step 4: Configure AutoRP mapping agent.

The AutoRP mapping agent listens for candidate RPs and then advertises to the rest of the network the list of available RPs. The **send-rp-discovery** command enables the core switches to act as AutoRP mapping agents.

```
ip pim send-rp-discovery Loopback1 scope 32
```

Step 5: Configure devices to listen to AutoRP announcements.

All Layer 3 switches and routers in the organization, including the RP switches, must be configured to listen to the AutoRP announcements from the mapping agents.

```
ip pim autorp listener
```

Devices other than the Cisco Catalyst 6800 and 6500 Series use the command **ip pim auto-rp listener**.

All Layer 3 interfaces in the network are enabled for sparse mode multicast operation.

```
C6500-VSS(config-if)#ip pim sparse-mode
```

Procedure 5   Connect to the distribution layer

In this design, links to the core VSS are configured as point-to-point Layer 3 routed EtherChannels. When using Cisco Catalyst 6800 or 6500 Series VSS system in the campus, all peer-connected links are EtherChannel links, with EtherChannel members distributed between the physical switches in the VSS. EtherChannel to the VSS provides for optimal forwarding because a packet that is received on the switch will be forwarded out a link on that same switch in normal operation instead of traversing the VSL connection.

Other benefits of EtherChannel to any single physical or logical device are that it makes it easier for bandwidth growth without changing the topology. A single link failure uses EtherChannel recovery versus using ECMP or a routing topology change to reroute the data flows for fastest recovery.

Since the core links are point-to-point routed links, use 30-bit IP address subnets and masks and do not use Switched Virtual Interfaces (SVI).

Step 1: Configure the Layer 3 interface.

If you use an EtherChannel to connect to a distribution layer platform, the interface type is PortChannel and the number must match the channel-group number you will configure in Step 2. When configuring a Layer 3 EtherChannel the logical port-channel interface is configured prior to configuring the physical interfaces associated with the EtherChannel.

```
interface [interface type] [number]
  description Link to {your device here}
  no switchport
  ip address [ip address] [mask]
  ip pim sparse-mode
  logging event link-status
  carrier-delay msec 0
```

```
load-interval 30
no shutdown
```

If the interface type is not a port-channel, then an additional command **macro apply EgressQoS** must also be configured on the interface.

Step 2: If the routing protocol you are using is OSPF, you add the router neighbor authentication to the interface. The chosen password must match the neighbor peer, and this OSPF interface authentication information is used in addition to the authentication information supplied in a previous step of the core layer configuration.

```
interface [interface type] [number]
 ip ospf message-digest-key 1 md5 [neighbor key]
```

Step 3: If you are connecting to the same distribution layer switch with multiple links, you can use a port channel for added bandwidth over a single logical link. Configure the physical interfaces to tie to the logical port channel by using the **channel-group** command. The number for the port channel and channel group are matched.

Also, apply the egress QoS macro that was defined in the platform configuration procedure to ensure traffic is prioritized appropriately.

> **i** **Tech Tip**
>
> The Cisco Catalyst 6500 and Catalyst 6800 Series Switches have two egress QoS macros: EgressQoSOneGig, which is used for Gigabit Ethernet ports, and EgressQoS, which is used for 10-Gigabit or 40-Gigabit Ethernet ports.

```
interface [interface type] [port 1]
  description Link to {your device here} Port 1
interface [interface type] [port 2]
  description Link to {your device here} Port 2
interface [interface type] [port 3]
  description Link to {your device here} Port 3
interface [interface type] [port 4]
  description Link to {your device here} Port 4
!
interface range [interface type] [port 1], [interface type] [port 2], [interface
type] [port 3], [interface type] [port 4]
  no switchport
  channel-protocol lacp
  channel-group [number] mode active
  logging event link-status
  logging event trunk-status
  logging event bundle-status
  load-interval 30
  macro apply EgressQoS
  no shutdown
```

Step 4: Save the running configuration that you have entered so it will be used as the startup configuration file when your switch is reloaded or power-cycled.

```
copy running-config startup-config
```

Example: Core to distribution port-channel configuration–EIGRP

Distribution — Port channel — Core

```
interface Port-channel 30
  description EtherChannel Link to Distribution Switch
  no switchport
  ip address 10.4.40.9 255.255.255.252
  ip pim sparse-mode
  load-interval 30
  no shutdown
!
interface range FortyGigabitEthernet 1/3/1, FortyGigabitEthernet 1/3/3,
FortyGigabitEthernet 2/3/1, FortyGigabitEthernet 2/3/3
  description EtherChannel Link to Distribution Switch
  no switchport
  carrier-delay msec 0
  channel-group 30 mode active
  logging event link-status
  logging event trunk-status
  logging event bundle-status
  load-interval 30
  macro apply EgressQoS
  no shutdown
!
key chain EIGRP-KEY
 key 1
  key-string [neighbor key]
router eigrp LAN
 !
 address-family ipv4 unicast autonomous-system 100
  !
  af-interface default
   authentication mode md5
   authentication key-chain EIGRP-KEY
  exit-af-interface
  !
  topology base
  exit-af-topology
  network 10.4.0.0 0.1.255.255
  eigrp router-id 10.4.40.254
  nsf
 exit-address-family
!
```

Example: Core to distribution port-channel configuration–OSPF
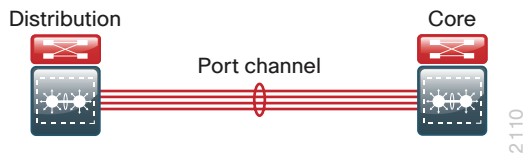

```
interface Port-channel 30
  description EtherChannel Link to Distribution Switch
  no switchport
  ip address 10.4.40.9 255.255.255.252
  ip pim sparse-mode
  ip ospf message-digest-key 1 md5 [neighbor key]
  load-interval 30
  no shutdown
!
interface range FortyGigabitEthernet 1/3/1, FortyGigabitEthernet 1/3/3,
FortyGigabitEthernet 2/3/1, FortyGigabitEthernet 2/3/3
  description EtherChannel Link to Distribution Switch
  no switchport
  carrier-delay msec 0
  channel-group 30 mode active
  logging event link-status
  logging event trunk-status
  logging event bundle-status
  no shutdown
  load-interval 30
  macro apply EgressQoS
!
router ospf 100
 router-id 10.4.40.254
 nsf
 area 0 authentication message-digest
 network 10.4.40.0 0.0.0.255 area 0
```

Appendix A: Product List

LAN Access Layer

Functional Area	Product Description	Part Numbers	Software
Modular Access Layer Switch	Cisco Catalyst 4500E Series 4507R+E 7-slot Chassis with 48Gbps per slot	WS-C4507R+E	3.3.1XO(15.1.1XO1) IP Base license
	Cisco Catalyst 4500E Supervisor Engine 8-E, Unified Access, 928Gbps	WS-X45-SUP8-E	
	Cisco Catalyst 4500E 12-port 10GbE SFP+ Fiber Module	WS-X4712-SFP+E	
	Cisco Catalyst 4500E 48-Port 802.3at PoE+ 10/100/1000 (RJ-45)	WS-X4748-RJ45V+E	
	Cisco Catalyst 4500E Series 4507R+E 7-slot Chassis with 48Gbps per slot	WS-C4507R+E	3.5.3E(15.2.1E3) IP Base license
	Cisco Catalyst 4500E Supervisor Engine 7L-E, 520Gbps	WS-X45-SUP7L-E	
	Cisco Catalyst 4500E 48 Ethernet 10/100/1000 (RJ45) PoE+,UPoE ports	WS-X4748-UPOE+E	
	Cisco Catalyst 4500E 48 Ethernet 10/100/1000 (RJ45) PoE+ ports	WS-X4648-RJ45V+E	
Stackable Access Layer Switch	Cisco Catalyst 3850 Series Stackable 48 Ethernet 10/100/1000 PoE+ ports	WS-C3850-48F	3.3.3SE(15.0.1EZ3) IP Base license
	Cisco Catalyst 3850 Series Stackable 24 Ethernet 10/100/1000 PoE+ Ports	WS-C3850-24P	
	Cisco Catalyst 3850 Series 2 x 10GE Network Module	C3850-NM-2-10G	
	Cisco Catalyst 3850 Series 4 x 1GE Network Module	C3850-NM-4-1G	
	Cisco Catalyst 3650 Series 24 Ethernet 10/100/1000 PoE+ and 2x10GE or 4x1GE Uplink	WS-C3650-24PD	3.3.3SE(15.0.1EZ3) IP Base license
	Cisco Catalyst 3650 Series 24 Ethernet 10/100/1000 PoE+ and 4x1GE Uplink	WS-C3650-24PS	
	Cisco Catalyst 3650 Series Stack Module	C3650-STACK	
	Cisco Catalyst 3750-X Series Stackable 48 Ethernet 10/100/1000 PoE+ ports	WS-C3750X-48PF-S	15.2(1)E3 IP Base license
	Cisco Catalyst 3750-X Series Stackable 24 Ethernet 10/100/1000 PoE+ ports	WS-C3750X-24P-S	
	Cisco Catalyst 3750-X Series Two 10GbE SFP+ and Two GbE SFP ports network module	C3KX-NM-10G	
	Cisco Catalyst 3750-X Series Four GbE SFP ports network module	C3KX-NM-1G	
	Cisco Catalyst 2960-X Series 24 10/100/1000 Ethernet and 2 SFP+ Uplink	WS-C2960X-24PD	15.0(2)EX5 LAN Base license
	Cisco Catalyst 2960-X FlexStack-Plus Hot-Swappable Stacking Module	C2960X-STACK	
Standalone Access Layer Switch	Cisco Catalyst 3650 Series 24 Ethernet 10/100/1000 PoE+ and 4x1GE Uplink	WS-C3650-24PS	3.3.3SE(15.01EZ3) IP Base license

LAN Distribution Layer

Functional Area	Product Description	Part Numbers	Software
Modular Distribution Layer Virtual Switch Pair	Cisco Catalyst 6800 Series 6807-XL 7-Slot Modular Chassis	C6807-XL	15.1(2)SY3 IP Services license
	Cisco Catalyst 6500 VSS Supervisor 2T with 2 ports 10GbE and PFC4	VS-S2T-10G	
	Cisco Catalyst 6500 4-port 40GbE/16-port 10GbE Fiber Module w/DFC4	WS-X6904-40G-2T	
	Cisco Catalyst 6500 4-port 10GbE SFP+ adapter for WX-X6904-40G module	CVR-CFP-4SFP10G	
	Cisco Catalyst 6500 CEF720 48 port 10/100/1000mb Ethernet	WS-X6748-GE-TX	
	Cisco Catalyst 6500 Distributed Forwarding Card 4	WS-F6K-DFC4-A	
	Cisco Catalyst 6500 Series 6506-E 6-Slot Chassis	WS-C6506-E	
	Cisco Catalyst 6500 VSS Supervisor 2T with 2 ports 10GbE and PFC4	VS-S2T-10G	
	Cisco Catalyst 6500 4-port 40GbE/16-port 10GbE Fiber Module w/DFC4	WS-X6904-40G-2T	
	Cisco Catalyst 6500 4-port 10GbE SFP+ adapter for WX-X6904-40G module	CVR-CFP-4SFP10G	
	Cisco Catalyst 6500 48-port GigE Mod (SFP)	WS-X6748-SFP	
	Cisco Catalyst 6500 Distributed Forwarding Card 4	WS-F6K-DFC4-A	
	Cisco Catalyst 6500 24-port GigE Mod (SFP)	WS-X6724-SFP	
	Cisco Catalyst 6500 Distributed Forwarding Card 4	WS-F6K-DFC4-A	
Extensible Fixed Distribution Layer Virtual Switch Pair	Cisco Catalyst 6800 Series 6880-X Extensible Fixed Aggregation Switch (Standard Tables)	C6880-X-LE	15.1(2)SY3 IP Services license
	Cisco Catalyst 6800 Series 6880-X Multi Rate Port Card (Standard Tables)	C6880-X-LE-16P10G	
Modular Distribution Layer Virtual Switch Pair	Cisco Catalyst 4500E Series 4507R+E 7-slot Chassis with 48Gbps per slot	WS-C4507R+E	3.5.3E(15.2.1E3) Enterprise Services license
	Cisco Catalyst 4500E Supervisor Engine 7-E, 848Gbps	WS-X45-SUP7-E	
	Cisco Catalyst 4500E 12-port 10GbE SFP+ Fiber Module	WS-X4712-SFP+E	
	Cisco Catalyst 4500E 48-Port 802.3at PoE+ 10/100/1000 (RJ-45)	WS-X4748-RJ45V+E	
Fixed Distribution Layer Virtual Switch Pair	Cisco Catalyst 4500-X Series 32 Port 10GbE IP Base Front-to-Back Cooling	WS-C4500X-32SFP+	3.5.3E(15.2.1E3) Enterprise Services license
Stackable Distribution Layer Switch	Cisco Catalyst 3850 Series Stackable Switch with 12 SFP Ethernet	WS-C3850-12S	3.3.3SE(15.0.1EZ3) IP Services license
	Cisco Catalyst 3850 Series 4 x 1GE Network Module	C3850-NM-4-1G	
	Cisco Catalyst 3850 Series 2 x 10GE Network Module	C3850-NM-2-10G	
	Cisco Catalyst 3750-X Series Stackable 12 GbE SFP ports	WS-C3750X-12S-E	15.2(1)E3 IP Services license
	Cisco Catalyst 3750-X Series Two 10GbE SFP+ and Two GbE SFP ports network module	C3KX-NM-10G	
	Cisco Catalyst 3750-X Series Four GbE SFP ports network module	C3KX-NM-1G	

LAN Core Layer

Functional Area	Product Description	Part Numbers	Software
Modular Core Layer Virtual Switch Pair	Cisco Catalyst 6800 Series 6807-XL 7-Slot Modular Chassis	C6807-XL	15.1(2)SY3 IP Services license
	Cisco Catalyst 6500 VSS Supervisor 2T with 2 ports 10GbE and PFC4	VS-S2T-10G	
	Cisco Catalyst 6500 4-port 40GbE/16-port 10GbE Fiber Module w/DFC4	WS-X6904-40G-2T	
	Cisco Catalyst 6500 16-port 10GbE Fiber Module w/DFC4	WS-X6816-10G-2T	
	Cisco Catalyst 6500 48-port GbE SFP Fiber Module w/DFC4	WS-X6848-SFP-2T	
	Cisco Catalyst 6500 Series 6506-E 6-Slot Chassis	WS-C6506-E	
	Cisco Catalyst 6500 VSS Supervisor 2T with 2 ports 10GbE and PFC4	VS-S2T-10G	
	Cisco Catalyst 6500 4-port 40GbE/16-port 10GbE Fiber Module w/DFC4	WS-X6904-40G-2T	
	Cisco Catalyst 6500 8-port 10GbE Fiber Module w/ DFC4	WS-X6908-10G-2T	
	Cisco Catalyst 6500 24-port GigE Mod (SFP)	WS-X6724-SFP	
	Cisco Catalyst 6500 Distributed Forwarding Card 4	WS-F6K-DFC4-A	

Appendix B:
Device Configuration Files

To view the configuration files from the CVD lab devices that we used to test this guide, please see the following:

http://cvddocs.com/fw/215-14b

Appendix C: Changes

This appendix summarizes the changes Cisco made to this guide since its last edition.

- We updated the document for readability, ease of configuration, and allow for better integration with other Enterprise CVDs.

- We updated all software versions as shown in Appendix A: Product List.

- We updated the hardware coverage.

   - In the access layer, we removed testing coverage for Cisco Catalyst 2960-S and 3560-X switches.

   - In the distribution layer, we added coverage for the Cisco Catalyst 3850 SFP-based switch, included the Cisco Catalyst 6807-XL and 6500-E VS4O (quad-supervisor SSO) configuration, and added the C6880-X-16P10G 16-port 10-Gigabit Ethernet card to the 6880-X configuration.

- We updated QoS polices.

   - Cisco Catalyst 3850 and 3650 Egress QoS policies were updated to take full advantage of the 2P6Q3T hardware capabilities and align to leading practices.

   - Cisco Catalyst 6807-XL and 6500-E QoS polices were updated to a simplified WRED aggregate class-default policy and WRED was removed from the control queues to align with leading practices.

   - QoS class-map match statements were separated to allow for better monitoring visibility.

- We added minor feature enhancements.

   - Secure Copy functionality is enabled as a default.

   - We use the more secure **username secret** password configuration.

   - We configure ports to automatically attempt to recover from errors with **errdisable recovery**.

   - We increase traffic visibility through a more granular **load-interval** configuration.

Feedback

Please use the feedback form to send comments and suggestions about this guide.

B-0000220-1 09/14