

Cisco Enterprise Campus Infrastructure

Best Practices Guide

December 2014

Contents

Executive Summary	3
Introduction	3
Enterprise Campus Network Design Alternatives	4
Campus Multitier Network Design Recommendations	4
Cisco Catalyst System-Level Design Best Practices	5
Access-Layer System Design Recommendation	5
Access-Layer System Redundancy Best Practices	6
Distribution-Layer System Design Recommendations	8
Distribution-Layer System Redundancy Best Practices	9
Distribution-Layer Network Design Recommendations	10
Distribution-Layer Network Design Alternative	10
Virtual Switching System Resiliency	11
Virtual Switching Domain and Best Practices.....	11
Virtual Switching Supervisor HA Best Practices	13
Virtual Switching Link Design and Best Practices	14
System and Network Connectivity Best Practices.....	17
Campus Network Oversubscription Best Practices	17
Access-Layer Network Connectivity Best Practices	18
Distribution-Layer Network Connectivity Best Practices	21
Cisco Multi-Chassis Layer 2 EtherChannel Best Practices	21
Multi-Chassis EtherChannel Best Practices	22
Campus Multilayer Network Design Best Practices	25
Multilayer VLAN Network Design Recommendations	25
Multilayer Network Protocols Best Practices.....	26
VLAN Trunking Protocol Recommendations	27
Dynamic Trunking Protocol (DTP) Recommendations	27
VLAN Trunk Design Recommendations	27
Spanning Tree Protocol Recommendations	29
Unidirectional Link Detection Recommendations	29
VSS MAC Address Table Synchronization Recommendations	30
Campus Core-Layer Network Design Best Practices	31
Core Uplink Design Recommendations	31
Cisco Multi-Chassis Layer 3 EtherChannel Best Practices	31
Enhanced Interior Gateway Routing Protocol Design Recommendations	32
Autonomous System and Network Best Practices	32
Secured Routing Best Practices.....	33
Network Route Summarization Best Practices	34
High-Availability Best Practices	34
Open Shortest Path First Routing Protocol Design Recommendations	35
Area and Network Design Best Practices.....	35
Secured Routing Best Practices.....	36
Network Route Summarization Best Practices	37
High-Availability Best Practices	37
Multicast Routing Protocol Recommendations.....	39
PIM Sparse Mode Best Practices	39
Secured Multicast Best Practices	40
High-Availability Best Practices	41
General Routing Recommendations	42
Equal Cost Multipath Routing Best Practices	42
Unicast IP Route Entry Purge Best Practices.....	43
IP Event Dampening	43
Summary	44
References	44

Executive Summary

Cisco® Unified Access establishes a framework that securely, reliably, and seamlessly connects anyone, anywhere, anytime, using any device to any resource. This framework empowers all employees with advanced services, taking advantage of an intelligent, enterprise-wide network to increase revenue, productivity, and customer satisfaction while reducing operational inefficiencies across the business. Cisco Unified Access includes services-rich network edge systems and combines a core network infrastructure embedded with integration of productivity-enhancing advanced technologies, including IP communications, mobility, security, video, and collaboration services.

Such mission-critical business application demands enterprises to implement a resilient and agile network to rapidly adapt to changing requirements and securely enable new and emerging services.

Introduction

This document consolidates the enterprise campus network design and deployment guidelines with various best practices from multiple deeply focused Cisco Validated Design Guides. The best practices conclusions are derived from thorough solution-level end-to-end characterization of various levels of system types, network design alternatives, and enterprise applications.

By following the best practices from this guide, the enterprise campus network can greatly simplify network operation, optimize application performance, and build resilience to operate networks in deterministic order during various types of planned and unplanned outages. This document limits the focus to construct a solid foundation and infrastructure between campus access, distribution, and core-layer systems. It covers the right set of recommendations to be applied on various types of platforms based on their roles in the network.

Figure 1. Large-Scale Enterprise Campus Distribution Network Design




Table 1 summarizes the hardware and software revisions that are addressed in this document.

Table 1. Cisco Catalyst Switches Hardware and Software Versions

Network Layer	Cisco Catalyst Switch	Software Version
Distribution	Cisco Catalyst 6800 Series Switches	15.1(2)SY2
Access	Cisco Catalyst 4500 Supervisor Engines 8-E, 7-E, and 7L-E	3.3.1.XO
	Cisco Catalyst 3850/3650 Series Switches	3.6.1.SE
	Cisco Catalyst 3750-X/3560-X Series Switches	3.6.1.SE
	Cisco Catalyst 2960 S/X/XR Series Switches	15.0.2-EX5

Enterprise Campus Network Design Alternatives

This section provides brief detailed network infrastructure guidance for each tier in the campus design model. Each design recommendation is optimized to keep the network simplified and cost-effective without compromising network scalability, security, and resiliency.

Campus Multitier Network Design Recommendations

The enterprise campus network deployment size and capacity vary broadly. Cisco offers a wide-ranging, rich Cisco Catalyst® switching portfolio that meets precise business and technical needs of individual customer requirements. With a variety of systems, offering variable port density, switching performance, scalability, and resiliency allows users to design and construct an end-to-end high-performance multitier network infrastructure.

Figure 2. Campus Multitier Network Deployment Models




Figure 2 illustrates multitier deployment models. Depending on the number of distribution-layer network blocks, scale, and performance requirements, the campus can be deployed either of these models.

As a best practice, Cisco recommends deploying a three-tier LAN design when numbers of distribution blocks are greater than two. Following are the primary benefits of deploying three-tier LAN networks:

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

Cisco Catalyst System-Level Design Best Practices

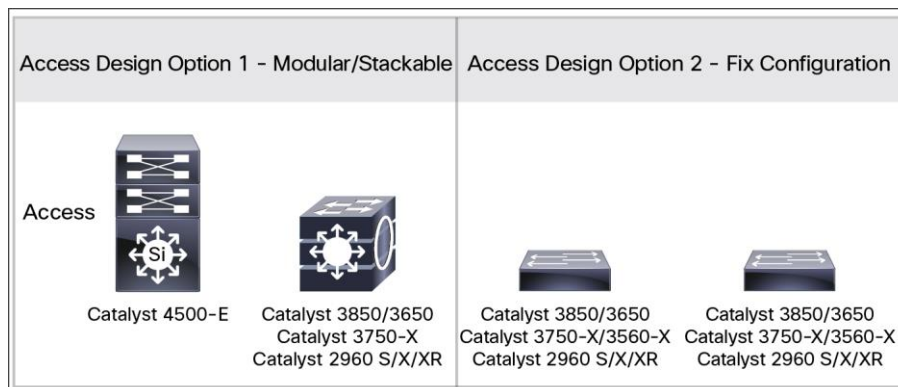
The enterprise campus network size broadly varies across different verticals and industries to enable communication infrastructure. The next-generation comprehensive Cisco Catalyst switching portfolio is designed to meet the scale of all deployment models. It is imperative to analyze business, technical, and application requirements to select the right products for unique and critical roles at different network tiers. This section provides product guidance and individual system-level best practices to construct end-to-end networks with more security, scalability, and resiliency.

Access-Layer System Design Recommendation

The access layer is the first tier or edge of the campus, where end devices such as PCs, printers, IP video surveillance cameras, Cisco TelePresence® devices, and so on attach to the wired portion of the campus network. It is also the place where IT managed devices are deployed that extend the network further out one more level, such as IP phones and wireless access points connecting wired or wireless end users. The wide varieties of possible types of devices that can connect and the various services and dynamic configuration mechanisms that are necessary make the access layer one of the most feature-rich parts of the enterprise campus network.

Based on the broad range of business communication devices and endpoints, network access demands, and capabilities, the following two types of access-layer design options can be deployed, as illustrated in Figure 3.

Figure 3. Access-Layer System Design Alternatives



Primary Benefits: Modular/Stackable System Design

- **Modular:** Provides investment protection by allowing seamless wiring-closet network expansion without infrastructure change.
- **Flexible:** Easy integration of new network module or stack-member switch without disrupting operations.
- **Resilient:**
 - Cisco Catalyst 4500E: Dual supervisor engine offers best-in-class system-level redundancy. Cisco recommended as best practices to deploy redundant supervisor to protect single-home endpoints using Stateful Switchover (SSO) technology. During planned outage such as In-Service Software Upgrade (ISSU) or abnormal supervisor failures, the network availability and capacity are fully protected for single-home devices.
 - Cisco Catalyst 3850/3650 StackWise[®]: The next-generation StackWise technology supports SSO to provide protocol redundancy and forwarding state machines with distributed forwarding architecture. However, the single-home endpoints will be affected during individual failure of stack-member switches.
 - Cisco Catalyst 3850 Cisco StackPower[®]: Network administrators must consider implementing power redundancy between the groups of Cisco Catalyst 3850s in a stack. The Cisco Catalyst 3850 provides nonstop forwarding even during catastrophic failures such as external power outage or the power supply unit failure.

Access-Layer System Redundancy Best Practices

The system-level redundancy support on modular versus fixed-configuration switches varies. When designed and deployed based on Cisco recommended best practices, it enables resilient infrastructure to maintain network communication for critical endpoint devices.

Supervisor and StackWise Best Practices

Table 2 provides best practices guidelines to deploy system-level redundancy with SSO technology on Cisco Catalyst 4500Es equipped with dual-supervisor engine modules as well as on next-generation Cisco Catalyst 3850/3650 Series fixed-configuration switches deployed in StackWise mode.

Table 2. Distribution-Layer System Resilient Best Practices

Best Practices	Cisco Catalyst 4500/3850/3650
Enable SSO on Cisco Catalyst 4500E system deployed with dual supervisor engine modules (default)	4500E(config)# redundancy 4500E(config-red)# main-cpu 4500E(config-r-mc)# mode sso
Enable SSO on 3850/3650 system deployed in StackWise (default)	3850-Stack(config)# redundancy 3850-Stack(config-red)# mode sso

Cisco Catalyst 3750-X StackWise-Plus and 2960 Series platforms in FlexStack and FlexStack Plus mode do not support SSO technology.

StackWise Software Autoupgrade Best Practices

The software resiliency in the Cisco Catalyst 3850/3650 is based on the Cisco IOS[®] Software high-availability framework when these switches are stacked together in StackWise mode. These next-generation fixed-configuration switches support 1+1 high-availability SSO function as modular-class platforms such as the 4500E. Thus it is imperative to have consistent Cisco IOS Software and license installed on the switches of each stack-member to provide 1+1 as well as N:1 ACTIVE stack system redundancy.

If new a 3850/3650 running an inconsistent software version joins the stack ring with the current running version, then such switch will force the stack ring down to an Route Processor Redundancy (RPR) state. In such a state the system remains completely down.

As a best practice, the newly joined switch can automatically receive consistent software versions from an ACTIVE switch and bring the system online without any user intervention. Table 3 illustrates simple command lines to automatically download consistent software versions to newly joined switches.

Table 3. Cisco Catalyst 3850/3650 Software Autoupgrade Best Practices

Best Practices	Cisco Catalyst 3850/3650: StackWise
Enable software autoupgrade on Cisco Catalyst 3850/3650 StackWise switch to automatically install consistent Cisco IOS Software on newly joined switch in stack ring	3850-Stack(config)#software auto-upgrade enable

Cisco StackPower Best Practices

The Cisco Catalyst 3850 and 3750-X Series platform supports innovative Cisco StackPower technology to provide power redundancy solutions for fixed-configuration switches. Cisco StackPower unifies the individual power supplies installed in the switches and creates a pool of power, directing that power where it is needed. Up to four switches can be configured in a Cisco StackPower stack with the special Cisco proprietary Cisco StackPower cable. The Cisco StackPower cable is different than the StackWise data cables and is available on all Cisco Catalyst 3850/3750X models. Cisco StackPower technology can be deployed in two modes:

- Sharing mode:** All input power is available to be used for power loads. The total aggregated available power in all switches in the power stack (up to four) is treated as a single large power supply. All switches in the stack can provide this shared power to all powered devices connected to Power over Ethernet (PoE) ports. In this mode, the total available power is used for power budgeting decisions without any power reserved to accommodate power supply failures. If a power supply fails, powered devices and switches could be shut down. This is the default mode of operation.
- Redundant mode:** The power from the largest power supply in the system is subtracted from the power budget and held in reserve. This reduces the total power available to PoE devices, but provides backup power in case of a power supply failure. Although there is less available power in the pool for switches and powered devices to draw upon, the possibility of having to shut down switches or powered devices in case of a power failure or extreme power load is reduced. Budgeting the required power and deploying each Cisco Catalyst 3850/3750-X switch in the stack with dual power supplies to meet demand are recommended. Enabling redundant mode offers power redundancy as a backup should one of the power supply units fail.

For better power redundancy across the stack ring, Cisco recommends deploying Cisco StackPower in redundant mode as the best practice.

Because Cisco StackWise-480 can group up to nine 3850 Series Switches in the stack ring, Cisco StackPower must be deployed with two power stack groups in order to accommodate up to four switches. The sample configuration in Table 4 demonstrates deploying Cisco StackPower in redundancy mode and grouping the stack members into power stack groups. To make the new power configuration effective, it is important that network administrator plan for network downtime because all the switches in the stack ring must be reloaded.

Table 4. Cisco Catalyst 3850/3750-X Cisco StackPower Best Practices

Best Practices	Cisco Catalyst 3850/3650: StackWise
Deploy Cisco StackPower technology to provide hitless power switchover on 3850 Series Switches. Redundant mode is recommended	3850-Stack(config)#stack-power stack PowerStack-1 3850-Stack(config-stackpower)#mode redundant 3850-Stack(config)#stack-power switch 1 3850-Stack(config-switch-stackpower)#stack PowerStack-1

Cisco StackWise and FlexStack Stack-MAC Best Practices

To provide a single unified logical network view in the network, the MAC addresses of Layer 3 interfaces on StackWise (physical, logical, Switch Virtual Interface (SVI), port channel) are derived from the Ethernet MAC address pool of the master switch in the stack. All Layer 3 communication from the StackWise switch to the endpoints (such as IP phones, PCs, servers, and core network system) is based on the MAC address pool of the master switch.

The stack-mac address on Cisco Catalyst 3850/3650 Series Switches deployed in StackWise mode maintains the stack-mac during ACTIVE stack switchover. By default, the stack-mac persistent timer is set to infinite, meaning do not modify the MAC address of Layer 3 interface. As best practices, retaining default settings and not modifying any stack-mac configuration are recommended.

Table 5. Cisco Catalyst 3850/3650 and 3750X Stack-MAC Best Practices

Best Practices	Cisco Catalyst 3850/3650: StackWise
Retain default stack-mac persistent setting on Cisco Catalyst 3850/3650 StackWise switches	3850-Stack(config)#default stack-mac persistent timer

By default the Cisco Catalyst 3750X StackWise-Plus and 2960 S/X/XR Series Switches do not protect the stack-mac address as do the Cisco Catalyst 3850/3650. Hence, as a best practice, setting the stack-mac persistent timer to zero (infinite) to prevent Address Resolution Protocol (ARP) and routing outages in the network is recommended.

Table 6. Cisco Catalyst 3750X and 2960-XR/S Stack-MAC Best Practices

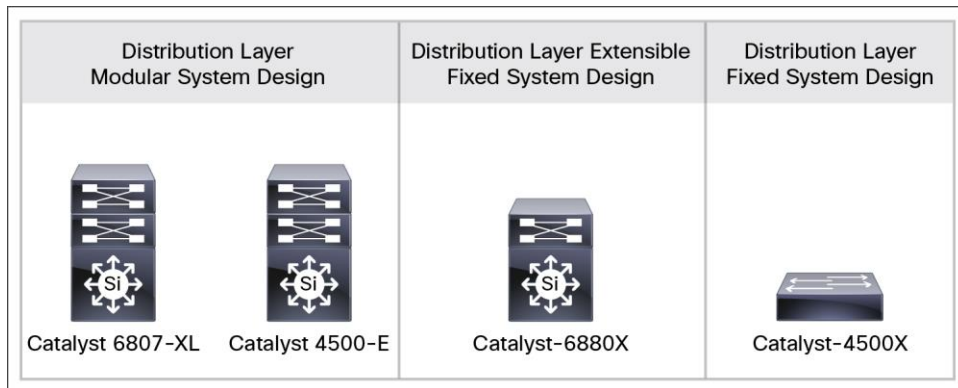
Best Practices	Cisco Catalyst 3750-X: StackWise
Modify default stack-mac persistent timer to infinite on Cisco Catalyst 3750X and 2960 S/X/XR Series Switches	3750-Stack(config)#stack-mac persistent timer delay 0

Distribution-Layer System Design Recommendations

The distribution or aggregation layer is the network demarcation boundary between the Layer 2 wiring closet network and the Layer 3 routed campus core network. The network operation, manageability, and application performance could become very complex with traditional Layer 2 technologies such as spanning-tree. The framework of the distribution-layer system must be designed with consideration of Cisco recommended best practices, which significantly reduce network complexities, increase reliability, and accelerate network performance. To build a strong campus network foundation with the three-tier model, the distribution layer has a vital role in consolidating networks and enforcing network edge policies.

The next-generation fixed and modular-class Cisco Catalyst switching portfolio enables a robust scale-up and scale-out networking architecture to build a high-performance infrastructure. It is imperative to analyze and deploy the right set of Cisco Catalyst switching products for building a mission-critical distribution-layer system. Figure 4 illustrates two system-level designs for distribution-layer deployments using enterprise-class networks.

Figure 4. Large Campus Distribution-Layer System Design Alternatives



Primary benefits: modular and extensible fixed system design:

- **Modular:** Provides investment protection with seamless network port and throughput expansion with multiterabit switching backplane without comprehensive infrastructure change.
- **Flexible:** Easy integration of new and mixed-medium network module in a switch without disrupting current network and business operations.
- **Resilient:**
 - **Supervisor module:** Enables nonstop network communication and uncompromised performance at distribution layer. Cisco SSO technology protects business continuity during planned outages such as software upgrade in real time or in unplanned outages such as catastrophic software failures. The Cisco Catalyst 6880X supports interchassis SSO with Cisco virtual switching system (VSS) technology.
 - **Network module:** Online insertion and removal procedure of network module is supported to complete new installation or migration without introducing system downtime.
 - **Power supplies:** Redundant power supplies provides environmental protection in system against power outages or supply unit failures. Modular power redundancy design provides flexibility to replace swap faulty unit without introducing network downtime.
 - **Fan trays:** Hot-swappable and redundant fan tray provides better cooling support and provides flexibility to replace faulty tray without introducing network downtime.

Alternatively, Cisco Catalyst 4500E/4500X or 3850 StackWise can be also deployed in distribution to meet the needs of small to medium-size network deployments. This best practices document primarily focuses on Cisco Catalyst 6800 Series systems for the distribution-layer role and providing best practices guidelines.

Distribution-Layer System Redundancy Best Practices

The Cisco Catalyst 6800 Series Switches are a leading enterprise-class high-performance and highly scalable system that is broadly deployed in campus distribution-layer and core-layer roles. The Cisco Catalyst 6880X switch is built upon the strong Cisco Catalyst 6500 Series Supervisor Engine 2T architecture but with a small form factor that expands consistent performance and scale for network deployments in a space constraint environment.

Table 7 provides best practices guidelines for Cisco Catalyst 6807-XL Series Switches deployed in recommended Cisco VSS mode or in traditional standalone mode and equipped with dual supervisor modules for additional system-level redundancy.

Table 7. Distribution-Layer System Resilient Best Practices

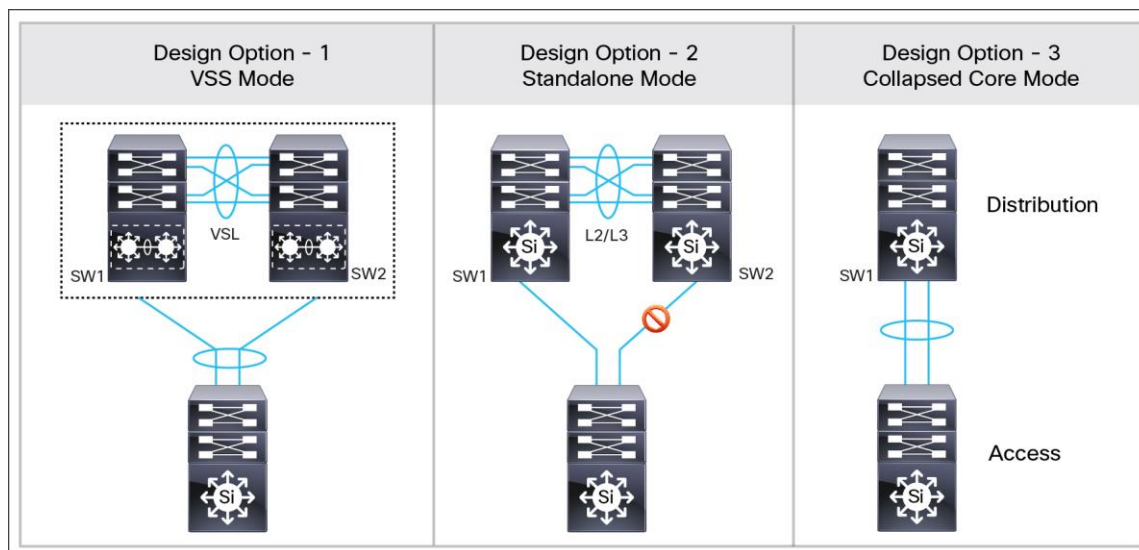
Best Practices	Cisco Catalyst 6800
Supervisor Module Enable SSO on system deployed with dual supervisor engine module (default).	Dist(config)# redundancy Dist(config-red)# main-cpu Dist(config-r-mc)# mode sso
Network Module Reduce higher data losses by powering down network module first prior to removing the module from service.	VSS Configuration Mode Dist(config)# no power enable switch [1 2] module [slot id] Standalone Configuration Mode Dist(config)# no power enable module [slot id]
Power Supply Redundancy Enable 1:1 power redundancy protection during catastrophic power outage or supply unit failure (default).	VSS Configuration Mode Dist(config)# power redundancy-mode switch [1 2] redundant Standalone Configuration Mode Dist(config)# power redundancy-mode redundant

Distribution-Layer Network Design Recommendations

Distribution-Layer Network Design Alternative

The enterprise campus distribution-layer network can be deployed in variety of designs to meet business, technical, and service needs. Most commonly it can be deployed in any of the design options as illustrated in Figure 5. Depending on network designs and primary technical requirements, network architects must make appropriate aggregation-layer design choices to enable end-to-end unified access network services.

Figure 5. Campus Distribution-Layer Network Design Alternatives



All of the preceding distribution design models offer consistent network foundation services, high availability, expansion flexibility, and network scalability at a system level. Each design option must be carefully evaluated to compare scale, performance, operation, and resiliency requirements:

- Distribution-layer design option 1: VSS mode
 This deployment model is intended for large to medium-size enterprise campus network deployments using Cisco VSS technology. Cisco Catalyst 6800 enables multiterabit high-performance aggregation blocks to build nonstop communication networks without compromising user-level complexity.

This is the primary recommended deployment mode for distribution-layer systems.

- Distribution-layer design option 2: standalone mode

The standalone mode is the default and has been a proven mode in enterprise networks for many years. The Cisco Catalyst 6800E system can be deployed in standalone mode with advanced fine tunings on distribution and access-layer switches to manually construct reliable networks based on Spanning Tree Protocol (STP). The network operation might become complex to deploy and troubleshoot when it expands significantly.

This is a secondary recommended deployment mode for distribution-layer systems. This document does not focus on this deployment mode.

- Distribution-layer design option 3: collapsed core/distribution mode

This deployment mode is sufficient for small campus or remote branch office networks. Any Cisco Catalyst Series Switches can be deployed in collapsed core and distribution role to provide demarcation between Layer 2 and Layer 3 network boundaries.

This document does not focus this deployment mode for distribution-layer systems.

Virtual Switching System Resiliency

Virtual Switching Domain and Best Practices

Domain ID Best Practices

Defining the VSS domain identifier (ID) is the initial premigration step toward implementing VSS with two standalone physical chassis. The VSS domain ID value ranges from 1 through 255. The virtual switch domain (VSD) includes two physical switches, and they must be configured with a common domain ID.

As a best practice, when implementing VSS in a multitier campus network design, the unique domain ID between different VSS pairs prevents network protocol conflicts and allows for simplified network operation, troubleshooting, and management.

Figure 6. VSS Domain ID Best Practices




Table 8. VSS Switch Domain and ID Best Practices

Cisco Catalyst 6800: SW1	Cisco Catalyst 6800: SW2
Distribution Layer: SW1 Dist-1(config)# switch virtual domain 10 Dist-1(config-vs-domain)# switch 1	Distribution Layer: SW2 Dist-2(config)# switch virtual domain 10 Dist-2(config-vs-domain)# switch 2
Core Layer: SW1 Dist-1(config)# switch virtual domain 20 Dist-1(config-vs-domain)# switch 1	Core Layer: SW2 Dist-2(config)# switch virtual domain 20 Dist-2(config-vs-domain)# switch 2

VSS Switch Priority Best Practices

When the two switches boot, switch priority is negotiated to determine control-plane ownership for the virtual switch. The virtual switch configured with the higher priority takes control-plane ownership, while the lower priority switch boots up in redundant mode. The default switch priority is 100; the lower switch ID is used as a tiebreaker when both virtual switch nodes are deployed with the default settings.

As a best practice, deploying both virtual switch nodes with identical hardware and software to take full advantage of the distributed forwarding architecture with a centralized control and management plane is recommended. Control-plane operation is identical on each of the virtual switch nodes.

Hence modifying the default switch priority is an optional setting and retains default values, as each of the virtual switch nodes can provide transparent operation to the network and the user.

Routed MAC Best Practices

The MAC address allocation for the interfaces does not change during a switchover event when the hot-standby switch takes over as the active switch. However, if both chassis are rebooted at the same time and the order of the active switch changes (the old hot-standby switch comes up first and becomes active), then the entire VSS domain uses that switch's MAC address pool. Any connected devices that do not support gratuitous ARP will be affected in network communication until the MAC address of the default gateway/interface is refreshed or timed out.

To avoid such a disruption, Cisco recommends using the configuration option provided with the VSS in which the MAC address for Layer 2 and Layer 3 interfaces is derived from the reserved pool. This takes advantage of the virtual switch domain identifier to form the MAC address. The MAC addresses of the VSS domain remain consistent with the usage of virtual MAC addresses, regardless of the boot order.

As a best practice, implementing routed MAC during the VSS migration process to avoid multiple system reboots is recommended. If VSS is already implemented without routed MAC address, then plan for downtime to get the new virtual MAC address effective on the next reboot cycle.

Table 9. VSS Routed MAC Best Practices

Cisco Catalyst 6800: SW1	Cisco Catalyst 6800: SW2
Distribution Layer: SW1 Dist-1(config)# switch virtual domain 10 Dist-1(config-vs-domain)# mac-address use-virtual	Distribution Layer: SW2 Dist-2(config)# switch virtual domain 10 Dist-2(config-vs-domain)# mac-address use-virtual

Standby Chassis Restoration Best Practices

The ACTIVE supervisor in a virtual switch will periodically update dynamic forwarding information to peer HOT_STANDBY supervisor module and local and remote distributed forwarding card (DFC) line cards to provide fully distributed forwarding capabilities.

This active synchronization provides rapid switchover to alternate connected paths, which helps protect user and application experience while the system undergoes the recovery process. The failed unit could be restoring for service, and during reinitialization state, the modules and ports become operational before actual forwarding information is resynchronized from the current ACTIVE switch.

Such a partial operational state of rejoined virtual switch will attract the data plane from peer devices while forwarding information is still pending. As a result, at the network level it will introduce higher data loss in the system recovery process than actual system loss.

As a best practice, Cisco recommends implementing a delay timer on standby chassis ports to become operational during bootup time. With this recommendation, the standby chassis will have sufficient time to first synchronize forwarding tables from the ACTIVE switch prior to becoming available for servicing user data from peer devices.


Table 10. VSS Routed MAC Best Practices

Best Practices	Cisco Catalyst 6800: VSS
Configure standby delay timer to 300 seconds	Dist(config)#switch virtual domain 10 Dist(config-vs-domain)#standby port delay 300 Dist(config-vs-domain)#standby port bringup 1 2

Virtual Switching Supervisor HA Best Practices

The Cisco VSS technology uses well-proven SSO technology to deliver supervisor module redundancy extending beyond a single system. The next-generation Supervisor Engine 2T in the Cisco Catalyst 6800 Series VSS system supports multidimension redundancy to protect against supervisor modules and chassis reset or failure. Figure 7 illustrates two supervisor redundancy alternatives to build simple and resilient distribution-layer networks.

Figure 7. Dual- and Quad-Sup Redundancy



- Dual-supervisor redundancy mode:** This deployment is equipped with a single Supervisor Engine 2T module on each Cisco Catalyst 6807-XL modular system. The system is limited to interchassis SSO redundancy where it provides nonstop communication to all dual-homed connected devices. Supervisor reset on any of the virtual switch systems will require complete chassis reset, including network modules causing performance to reduce by half and recovery time decrease depending on reset cause, failure type, and so on.

The fixed-configuration Cisco Catalyst 6880X in VSS mode can only support this redundancy mode because of its system architecture.

- Quad-supervisor redundancy mode:** This deployment is applicable for two 6807-XL systems deployed with two Supervisor Engine 2T modules per system in VSS mode. The system-level redundancy gets quadrupled with total four Supervisor Engine 2T modules to protect network availability and capacity during planned or unplanned network outages. This redundancy mode also offers redundancy to any single-homed network devices connected directly to Cisco Catalyst 6807-XL VSS systems.

As a best practice, quad-supervisor redundancy mode is recommended when the business and technical requirement demands a best-in-class resilient infrastructure to support network availability and capacity during planned or unplanned network outage.


Virtual Switching Link Design and Best Practices

To cluster two physical chassis into a single logical entity, Cisco VSS technology enables the extension of various types of single-chassis internal system components to the multichassis level. Each virtual switch must be deployed with direct physical links, which extend the backplane communication boundary (known as virtual switch links [VSLs]). VSLs can be considered Layer 1 physical links between two virtual switch nodes and are designed to operate without network control protocols. Therefore, VSLs cannot establish network protocol adjacencies and are excluded when building the network topology tables.

VSL Physical Link Connectivity Best Practices

The Cisco Catalyst 6807-XL or 6880X supports VSL capabilities on any supported 10G or 40G network modules. VSL physical connectivity between two systems must be deployed with special consideration to make sure it has the best level of redundancy to address various external or internal fault conditions. As a best practice, Cisco recommends equally diversifying VSL fiber connections between the Supervisor Engine 2T and network module. Figure 8 illustrates VSL connectivity best practices when a Cisco Catalyst 6800 Series system is deployed in VSS mode.

Figure 8. Cisco VSS VSL Best Practices



VSL Capacity Planning Best Practices

Cisco VSS forwards user data plane over VSL as a last resort when a virtual switch does not find a local forwarding path to switch the traffic. To minimize congestion on the VSL interface, as a best practice, the network administrator must also consider VSL capacity planning. This planning should make sure the VSL interface has sufficient aggregated bandwidth to reroute the user data plane in case of primary local path failure. These four major factors must be considered when calculating required VSL bandwidth capacity:

- The aggregated network uplink bandwidth capacity on per virtual switch node basis, for example, 2 x 10GE diversified to the same remote core system.

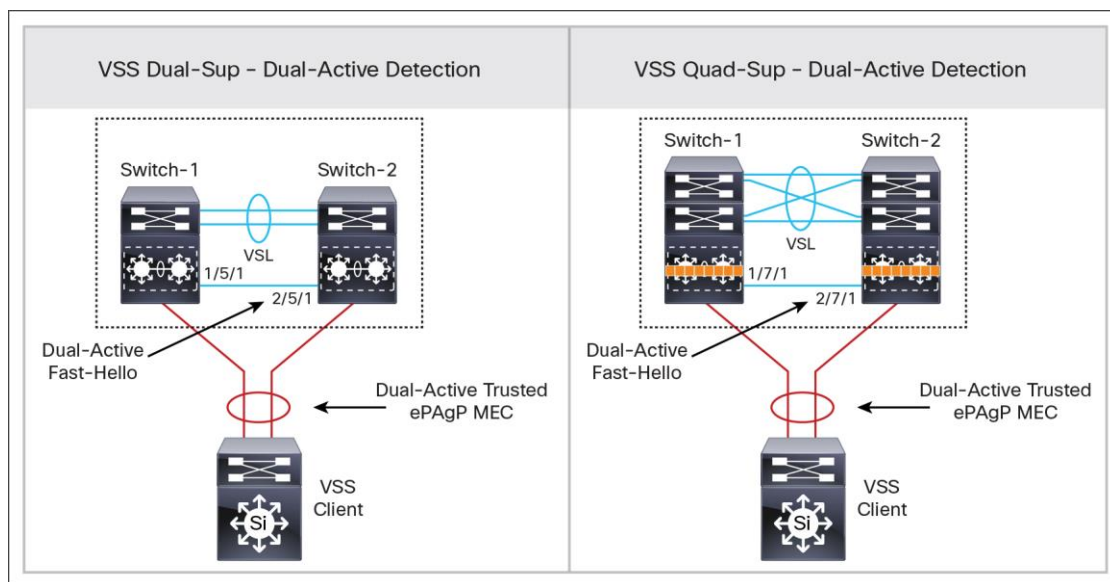
- Designing the network with single-homed device connectivity (no multichassis EtherChannel [MEC]). Single-homed network devices connectivity is highly discouraged.
- Remote Switched Port Analyzer (SPAN) from one switch member to another virtual switch system.
- If the 6807-XL system is deployed with integrated service module (that is, ASA-SM, WiSM2, and so on), then depending on the service module forwarding design and capacity, the user data may be carried over the VSL bundle.

Maximum VSL EtherChannel interface can carry up to 8 x 10G/40G ports. For optimal traffic load sharing between VSL member links, bundling VSL member links in powers of 2 (that is, 2, 4, and 8) is recommended.

VSL Redundancy Best Practices

The VSL EtherChannel functions as an extended backplane link that enables system virtualization by transporting interchassis control traffic, network control plane, and user data traffic. The state machine of the unified control-plane protocols and distributed forwarding entries gets dynamically synchronized between the two virtual switch nodes. Any fault triggered on a VSL component breaks the system virtualization and introduces a dual-active condition into the network. Figure 9 illustrates two techniques to detect such fault conditions and rapidly recover the system before the network and end-user applications are affected.

Figure 9. Dual Active Detection Best Practices



- Dual active fast hello best practices:
 - Up to four ports of any link type and speed can be used to pair up for dual active fast hello detection. Minimum one port is recommended.
 - Dual-active fast hello is supported on supervisor and network module ports. For dual-sup network design, onboard supervisor 1G port can be used for fast hello. For quad-sup network design one of the available ports from network module should be selected for better redundancy during failure of either supervisor module.
 - Dual-active fast hello can coexist with enhanced Port Aggregation Protocol (ePAGP). Having both mechanics implemented for better redundancy is recommended.

Table 11. Implementing Dual Active Fast Hello

Cisco Catalyst 6800: VSS
Dist(config)# interface range Gig 1/3/1, Gi2/3/1 Dist(config-if-range)# dual-active fast-hello

- Enhanced PAgP+ best practices:
 - Large number of Layer 2 or Layer 3 PAgP-enabled EtherChannels can be used for dual-active detection method. Minimum of two dual-active trusted MECs is recommended for better redundancy.
 - Trusted Layer 2 MEC could be an access-layer switch. Layer 3 MEC could be a core-layer switch. Both can coexist and are recommended.
 - Enabling or disabling dual-active trust setting on Layer 2 or Layer 3 requires interfaces to be in administrative shutdown state. Hence, as a best practice, planning a short downtime to implement the solution is recommended.

Table 12. VSS Dual-Active Trust on PAgP MEC Best Practice

Cisco Catalyst 6800: VSS	Dual-Active PAgP VSS Client
Dist(config)# interface range Port-Channel 101-102 Dist(config-if-range)# shutdown	No configuration required
Dist(config)# switch virtual domain 1 Dist(config-if-range)# dual-active detection pagp trust channel-group 101 Dist(config-if-range)# dual-active detection pagp trust channel-group 102	
Dist(config)# interface range Port-Channel 101-102 Dist(config-if-range)# no shutdown	

- Dual active exclude best practices:
 - During dual active state the recovery chassis automatically disables all ports to prevent network malfunction. As a best practice it recommended to exclude the Layer 3 out-of-band management port from each virtual switch chassis to remain operational for additional troubleshooting and debugging.
 - Any management port can be used from supervisor and network modules. For dual-sup network design onboard supervisor 1G port can be used. For quad-sup network design an available port from network module should be selected for better redundancy during failure of either supervisor module.
 - These excluded management ports must be reachable through separate Layer 2/Layer 3 network infrastructure instead from core routing space.
 - The static route from each Layer 3 excluded interface would be required to reach network beyond local subnet.

Table 13. VSS Dual-Active Trust Exclude Interface Best Practices

Cisco Catalyst 6800: VSS
Dist(config)# interface Gig1/3/2 Dist(config-if)# ip address <address> <mask> Dist(config-if)# no shutdown
Dist(config)# interface Gig2/3/2 Dist(config-if)# ip address <address> <mask> Dist(config-if)# no shutdown

Cisco Catalyst 6800: VSS

```
Dist(config)#switch virtual domain 1
Dist(config-if-range)#dual-active exclude interface Gi1/3/2
Dist(config-if-range)#dual-active exclude interface Gi2/3/2

Dist(config)#ip route <network> <mask> <gateway-1>
Dist(config)#ip route <network> <mask> <gateway-2>
```

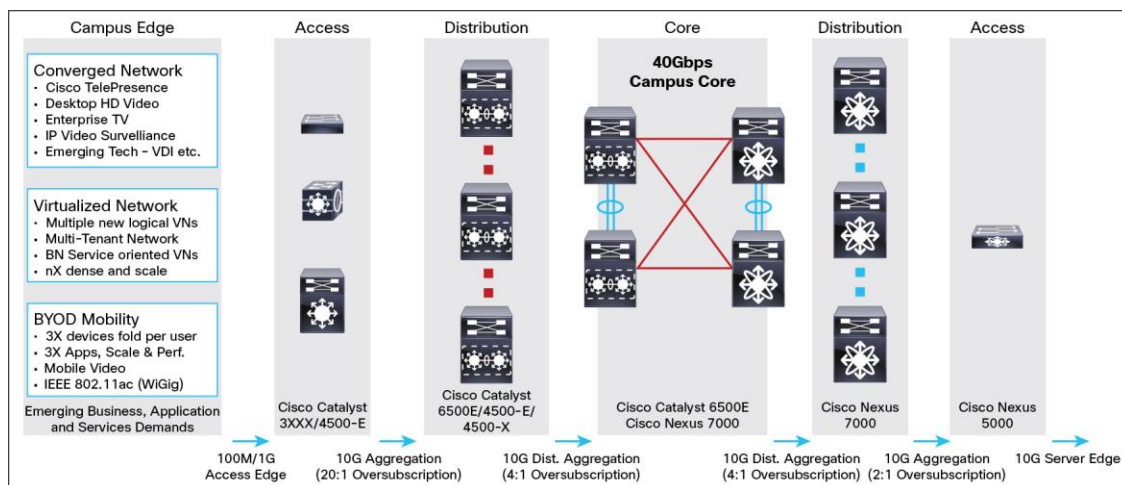
System and Network Connectivity Best Practices

Campus Network Oversubscription Best Practices

The switching performance demands can be combined with the scale factors at every network layer. Most large enterprise campus networks in multitier architectures are oversubscribed. The network expansions at the edge always demanded refreshing campus core switching capacity, but it is becoming increasingly challenging to maintain oversubscription ratios with prolific 10G use in next-generation wiring-closet networks. The high-speed switching capacity demands are fueled by several reasons: new devices such as 1G desktop, wireless 802.11AC, video, mobile device proliferations, and so on. To maintain consistent quality of experience (QoE) with increased scale and performance demands, IT needs to reevaluate two major bottleneck points in campus networks. These bottlenecks are at the campus distribution layer that aggregates 10G physical connection and core switching capacity to maintain 4:1 oversubscription ratios.

As a best practice, to mitigate campus and data center core scale and performance challenges and to protect 4:1 oversubscription ratios, the 40G Ethernet innovations in Cisco Catalyst 6800-E and Cisco Nexus® 7000 systems allow easy and seamless upgrade from existing 10G-based core infrastructure. The distribution layer should maintain a 20:1 oversubscription ratio in distribution-access networking blocks. If the access-layer switches in Intermediate Distribution Frame (IDF) are increasing capacity by tenfold, then network architects must make sure that aggregation switches, modules, and ports are able to meet the required performance. In modular distribution-layer design, the leading Cisco Catalyst 6800 Series system can be deployed with next-generation Supervisor Engine 2T with Cisco Catalyst 6904 or 6816 line-card modules for 10G port aggregations.

Figure 10. Network Oversubscription Best Practices



Access-Layer Network Connectivity Best Practices

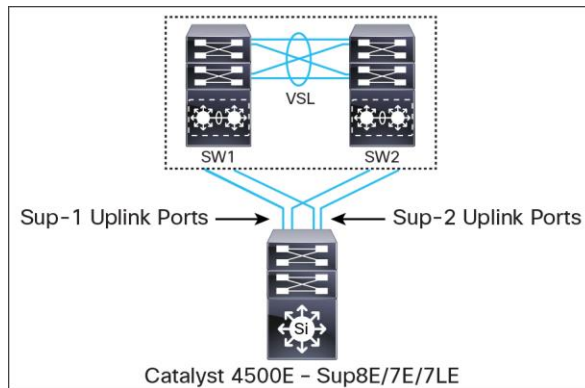
The uplink network connectivity design from a variety of access-layer model switches to distribution-layer systems is critically important for optimizing application performance and network resiliency. This best practice makes sure the user data plane is well load balanced to optimally utilize all system resources and that network protocols are building redundant forwarding paths to rapidly switch over during various types of planned and unplanned network outages.

Cisco Catalyst 4500-E Redundant Supervisor Uplink Recommendation

In redundant mode, the Cisco Catalyst 4507R+E or 4510R+E chassis is deployed with dual Supervisor Engine 8-E or 7-E modules in a redundant SSO configuration. The four 1G/10G uplink ports on each supervisor module are divided into two port groups: port group 1 (port 1 and 2) and 2 (port 3 and 4). Port group 2 becomes automatically inactive on both supervisor modules when the Cisco Catalyst 4500E system detects redundant modules installed in the chassis.

As illustrated in Figure 11, as a best practice, utilize all four uplink ports from both supervisors to a redundant Cisco Catalyst 6800 Series distribution-layer VSS system. Both supervisor modules can equally diversify port group 1 with the redundant upstream system for the same consistent bandwidth capacity, load balancing, and link redundancy as nonredundant mode.

Figure 11. 4500-E Redundant Supervisor Uplink Recommendations



Cisco StackWise and FlexStack Uplink Recommendation

Cisco Catalyst 3850/3650, 3750-X, and 2960 Series Switches support up to 4 physical uplink ports to connect distribution-layer switches. Typically up to 2 physical uplink interfaces are deployed from access-layer switches for optimal load balancing and redundancy in the wiring closet.

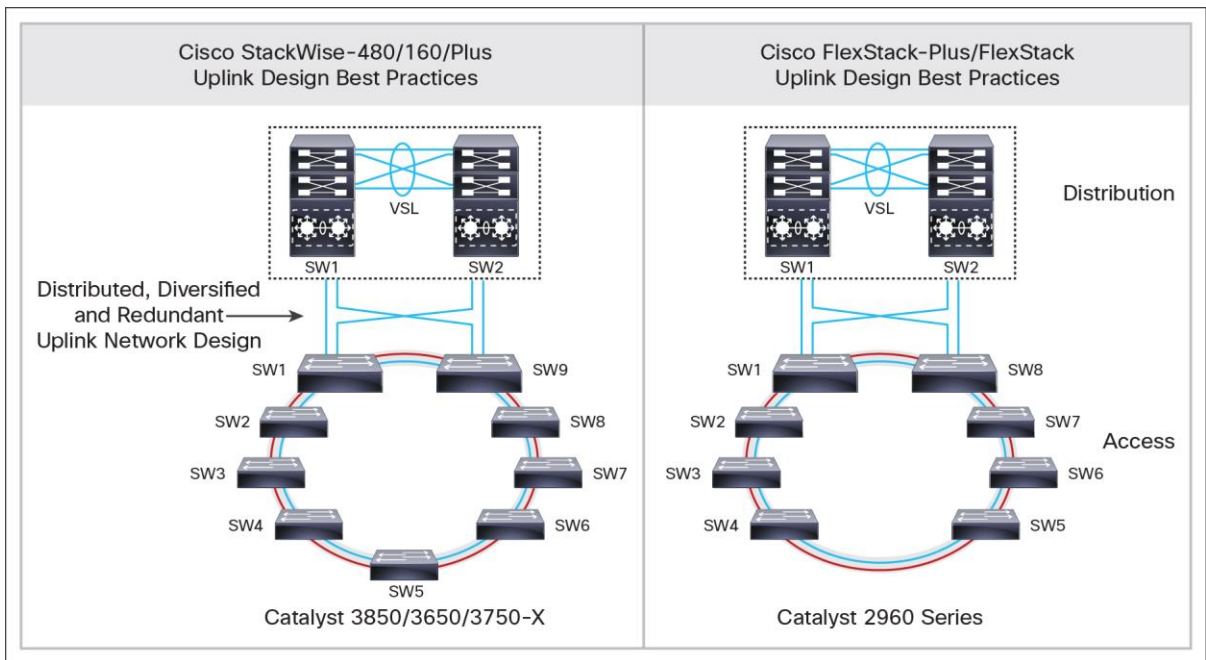
When these switches are deployed in stack configuration mode, maintaining the same uplink connection design principle as with a dual stack-member system is recommended. For example, nine Cisco Catalyst 3850 switches deployed in a stack ring would have 2 diversified uplink ports from switch 1 and 2 diversified uplink ports from switch 9. The remaining seven switches forward the data toward the core using a high-speed stack backplane.

This recommended uplink port design offers various benefits, from application performance to optimal user experience:

- Improved application performance by increasing aggregated stack switching capacity with multiple distributed high-speed 10Gbps uplinks between stack member Cisco Catalyst switches

- Enhanced bidirectional traffic engineering with intelligent network data load sharing within the stack ring and across all distributed uplink physical ports
- Improved system and application performance by utilizing the distributed forwarding architecture advantage of hardware resources: buffers, queues, ternary content-addressable memory (TCAM), and so on.
- Protection of stack and network-level redundancy and minimization of congestion between distributed aggregation systems caused during a major outage at the access or distribution layer

Figure 12. Cisco StackWise and FlexStack Uplink Recommendations



Cisco StackWise and FlexStack Switch Priority Recommendation

Cisco introduced various generations of StackWise and FlexStack technologies on different generations of switching portfolio. Every new generation of this technology offered advanced capabilities to optimize scale, performance, and resiliency while maintaining simplicity to the user and network.

The stack technology simplicity is attained by centralizing control and management planes from all stack-member switches to a dynamically elected single switch in a ring. The switching performance is optimized by building a fully distributed forwarding architecture to optimally utilize through stack backplane and all diversified system resources such as uplink fiber ports, TCAM, and so on.

As a best practice, decoupling the control and management-plane operation with uplink port operation in StackWise and FlexStack switches is imperative. This can be achieved by statically assigning switch IDs for each stack-member switch. The switch with higher priority is deterministically elected for ACTIVE/master role in a stack ring. For rapid reelection, a secondary switch can be preset with the next lower priority level. This best practice reduces network reconvergence time when the ACTIVE/master switch fails in the stack.

Figure 13. Cisco StackWise and FlexStack Switch Priority Recommendations

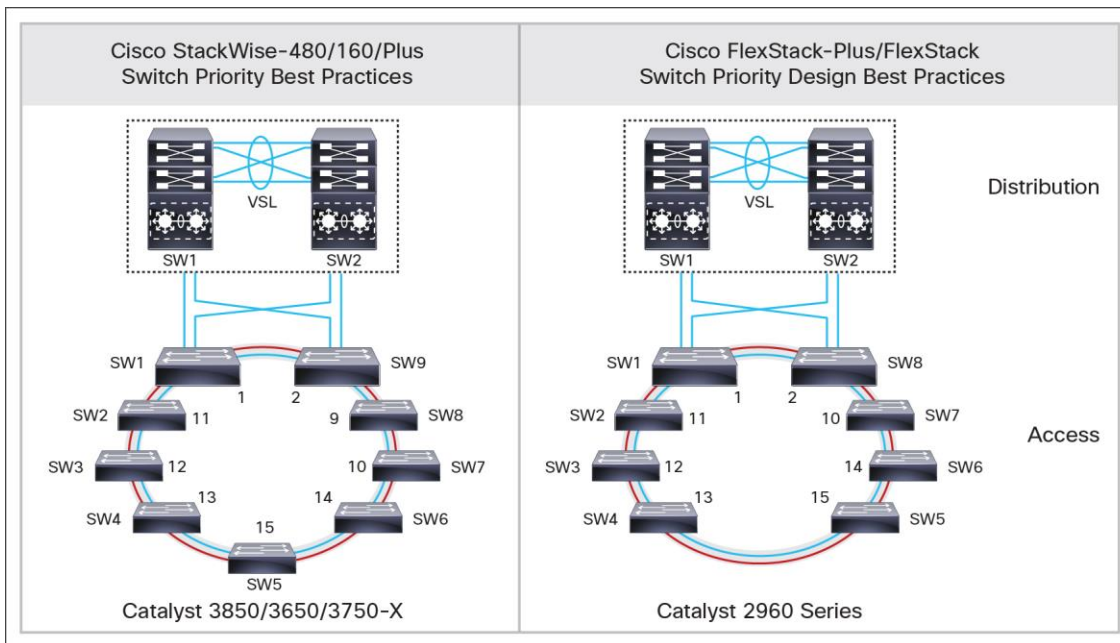


Table 14 provides best practices configuration to statically configure switch priority on Cisco Catalyst 3850/3650, 3750-X, and 2960 Series Switches.

Table 14. Cisco StackWise Switch Priority Best Practices

Cisco Catalyst StackWise and FlexStack Switch	Cisco StackWise and FlexStack Switch Priority
3850-Stack#switch 1 priority 1 ! Switch 1 with Uplink ports	3750X-Stack#config terminal
3850-Stack#switch 2 priority 11	3750X-Stack(config)#switch 1 priority 1
3850-Stack#switch 3 priority 12	! Switch 1 with Uplink ports
3850-Stack#switch 4 priority 13	3750X-Stack(config)#switch 2 priority 11
3850-Stack#switch 5 priority 15	3750X-Stack(config)#switch 3 priority 12
3850-Stack#switch 6 priority 14	3750X-Stack(config)#switch 4 priority 13
3850-Stack#switch 7 priority 10	3750X-Stack(config)#switch 5 priority 15
3850-Stack#switch 8 priority 9	3750X-Stack(config)#switch 6 priority 14
3850-Stack#switch 9 priority 2	3750X-Stack(config)#switch 7 priority 10
! Switch 9 with Uplink ports	3750X-Stack(config)#switch 8 priority 9
	3750X-Stack(config)#switch 9 priority 2
	! Switch 9 with Uplink ports

Table 15. Cisco FlexStack Switch Priority Best Practices


Cisco Catalyst StackWise and FlexStack Switch	Cisco StackWise and FlexStack Switch Priority
2960XR-Flex#config terminal	2960S-Flex#config terminal
2960XR-Flex(config)#switch 1 priority 1	2960S-Flex(config)#switch 1 priority 1
! Switch 1 with Uplink ports	! Switch 1 with Uplink ports
2960XR-Flex(config)#switch 2 priority 11	2960S-Flex(config)#switch 2 priority 14
2960XR-Flex(config)#switch 3 priority 12	2960S-Flex(config)#switch 3 priority 15
2960XR-Flex(config)#switch 4 priority 13	2960S-Flex(config)#switch 4 priority 2
2960XR-Flex(config)#switch 5 priority 15	! Switch 4 with Uplink ports
2960XR-Flex(config)#switch 6 priority 14	
2960XR-Flex(config)#switch 7 priority 10	
2960XR-Flex(config)#switch 8 priority 2	
! Switch 8 with Uplink ports	

Distribution-Layer Network Connectivity Best Practices

The campus distribution and core-layer systems typically have modular hardware with centralized processing on a supervisor module and distributed forwarding on high-speed network modules. The fundamentals of building a resilient campus network design with diversified, distributed, and redundant physical paths do not vary by role, system, or deployed configuration mode. The physical uplink connection from each distribution layer to the core can be with single or dual uplink Layer 3 interfaces. The single link builds a square physical topology that does not offer optimal load balancing and redundancy. It also affects application performance with slower network recovery during planned system or network-level failures.

Following common physical-layer network design principles as a best practice, Cisco recommends deploying full-mesh redundant physical connections at any level of network tiers as illustrated in Figure 14.

Figure 14. Full-Mesh Distribution Core-Layer Connectivity Best Practices



Cisco Multi-Chassis Layer 2 EtherChannel Best Practices

Traditionally campus networks were designed with standalone network systems and equal cost multipath (ECMP), which did not provide the flexibility to simplify network design with redundant devices or paths to logically act as a single entity. Campus network designs are evolving with Cisco's system virtualization innovation in Cisco Catalyst switching platforms, such as VSS, StackWise, and FlexStack, with redesign opportunities in all three tiers. While each of these virtualization techniques clusters multiple physical systems into a single large and unified logical system, the distributed multiple parallel physical paths can now be bonded into logical point-to-point EtherChannel interfaces between two systems. The principle of building a full-mesh physical campus network should not be changed when a campus device or link is implemented to operate in logical mode.

Designing a multilayer or campus backbone network with EtherChannel between two systems offers multiple benefits:

- **Simplify:** Bundling multiple ECMP paths into logical EtherChannel reduces redundant protocol adjacencies, routing databases, and forwarding paths.
- **Optimize:** Reduces the number of control-plane operations and optimizes system resources, such as CPU/memory utilization. Provides flexible Layer 2 to Layer 4 variables to intelligently load share and utilize all resources for network data traffic across each bundled interface.

- **Reduce complexity:** Simplifies network operational practice and reduces the amount of network configuration and troubleshooting required to analyze and debug problems.
- **Increase capacity:** Eliminates protocol-driven restrictions and doubles switching capacity in multilayer designs by utilizing all resources (bandwidth, queue, buffer, and so on) across all bundled Layer 2 uplinks.
- **Provide resilience:** Provides deterministic hardware-driven network recovery for seamless business operation. Minimizes routing database recomputation and topology changes during minor network faults, for example, link failure.


Multi-Chassis EtherChannel Best Practices

Cisco MEC technology is a breakthrough innovation that creates logical point-to-point EtherChannels distributed across multiple physical switches, which allows for a highly resilient virtual switch in the VSS domain. Deploying Layer 2 or Layer 3 MEC with VSS introduces the following benefits:

- In addition to all EtherChannel benefits, the distributed forwarding architecture in MEC helps increase network bandwidth capacity.
- Increases network reliability by eliminating the single-point-of-failure limitation compared to traditional EtherChannel technology.
- Simplifies the network control plane, topology, and system resources within a single logical bundled interface instead of multiple individual parallel physical paths.
- Independent of network scalability, MEC provides deterministic hardware-based subsecond network recovery.
- MEC technology on the Cisco Catalyst 6800 Series system in VSS mode remains transparent to remote peer devices.

Cisco recommends designing the campus network to bundle parallel paths into logical EtherChannel or MEC in all layers when possible. The campus network can be deployed in a hybrid EtherChannel and ECMP network design if any device cannot logically bind interfaces such as a standalone campus core-layer system. Figure 15 illustrates the recommended end-to-end EtherChannel/MEC network design that simplifies end-to-end network operation.


Figure 15. Multi-Chassis EtherChannel Recommendations



Cisco Enhanced PAgP and LACP Best Practices

The EtherChannel uses link-bundling protocols for various set of parameter checks on local physical ports and remote systems through per-port signaling mechanics. The member links of EtherChannel must join the port channel interface using Cisco enhanced PAgP or industry-standard IEEE Link Aggregation Control Protocol (LACP) port aggregation protocols. As illustrated in Figure 16, the EtherChannel implemented in distribution block or in core layer is recommended to implement with enhanced PAgP or LACP.

Figure 16. Cisco Enhanced PAgP and LACP Recommendation



Both protocols are designed to provide consistent link capabilities; however, the Cisco enhanced PAgP protocol provides an additional solution advantage, dual-active detection. Implementing these protocols provides the following additional benefits:

- Makes sure of link aggregation parameter consistency and compatibility between two the systems.
- Makes sure of compliance with aggregation requirements.
- Dynamically reacts to runtime changes and failures on local and remote EtherChannel systems.
- Detects and removes unidirectional links and multidrop Ethernet connections from the EtherChannel bundle.

All Cisco Catalyst switching platforms support Cisco enhanced PAgP and industry-standard LACP protocols. Table 16 illustrates configuration guidelines for Cisco enhanced PAgP and LACP protocols between two systems based on their software capabilities.

Table 16. Cisco Catalyst Enhanced PAgP and LACP Best Practices

Recommended Protocol	Distribution-Layer Switch	Access-Layer Switch
Enhanced PAgP	Cisco Catalyst 6800 VSS interface range <id> - <id> channel-protocol pagp channel-group <id> desirable !	Cisco Catalyst 4K/3K/2K interface range <id> - <id> channel-protocol pagp channel-group <id> desirable !
IEEE LACP	Cisco Catalyst 6800 VSS interface range <id> - <id> channel-protocol lacp channel-group <id> active !	Cisco Catalyst 3750-X Stack interface range <id> - <id> channel-protocol lacp channel-group <id> active !

Recommended Protocol	Distribution-Layer Switch	Access-Layer Switch
	Cisco Catalyst 6800 VSS interface range <id> - <id> channel-protocol lacp channel-group <id> active !	Cisco Catalyst 2960 Series FlexStack interface range <id> - <id> channel-protocol lacp channel-group <id> active !

Static EtherChannel Best Practices

Today most of the next-generation Cisco or other vendor networking devices support enhanced PAgP or IEEE LACP for dynamic EtherChannel link bundling. However, certain type of systems such as Cisco ISRs, Cisco 5508 WLC, and so on support EtherChannel capabilities in static or nonnegotiate configuration mode instead of dynamic. Static-mode EtherChannels cannot provide error detection on member-link configuration, parameter, and capability consistency, but they can provide load balancing and redundancy as dynamic EtherChannels.

As a best practice, Cisco recommends checking the update software version and latest EtherChannel capabilities of such devices to identify if PAgP or LACP is supported. If it is supported, then using dynamic mode instead of static mode is recommended. If it is not supported, then static mode EtherChannel configuration must be carefully evaluated on each end of interfaces and make sure the link type, configuration and their capabilities are consistent to provide the best level of load sharing and redundancy across the member links in EtherChannels.

Table 17. Static EtherChannel Best Practices

Recommended Protocol	Distribution-Layer Switch	Remote EtherChannel Device
Static	Cisco Catalyst 6800 VSS interface range <sw1-id> - <sw2-id> channel-group <id> on !	Remote Device interface range <id> - <id> channel-group <id> on ! Static EtherChannel configuration on remote device varies based on OS/CLI support

EtherChannel Load Balancing

The number of applications and their functions in a enterprise campus network design is highly variable, especially when the network is provided as a common platform for business operation, campus security, and open accessibility. It is important for the network to become more intelligence-aware, with data-link to transport-layer packet inspection and load sharing traffic to optimally utilize all available network resources.

Fine tuning to EtherChannel and MEC interface adds extra computing intelligence to the network to make protocol-aware egress forwarding decisions between multiple local member link paths. For each traffic flow, such tuning optimizes the egress path selection procedure with multiple levels of variable information that originated from the source host (that is, Layer 2 to Layer 4). EtherChannel load-balancing methods vary on Cisco Catalyst switching and routing platforms. Figure 17 illustrates the recommended end-to-end EtherChannel load-balancing method.

Figure 17. EtherChannel Load-Balancing Best Practices

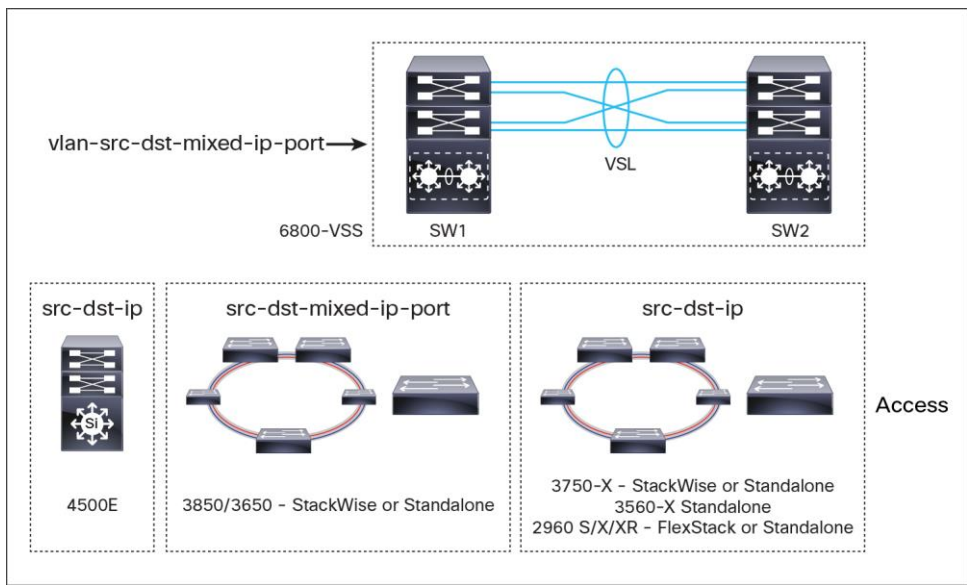


Table 18. EtherChannel Load-Balancing Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#port-channel load-balance vlan-src-dst-mixed-ip-port
Access	Cisco Catalyst 4500E 4500E(config)#port-channel load-balance src-dst-ip
	Cisco Catalyst 3850/3650 3850(config)#port-channel load-balance src-dst-mixed-ip-port
	Cisco Catalyst 3750-X/3560-X 3750(config)#port-channel load-balance src-dst-ip
	Cisco Catalyst 2960 S/X/XR 2960(config)#port-channel load-balance src-dst-ip

Campus Multilayer Network Design Best Practices

Multilayer VLAN Network Design Recommendations

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

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

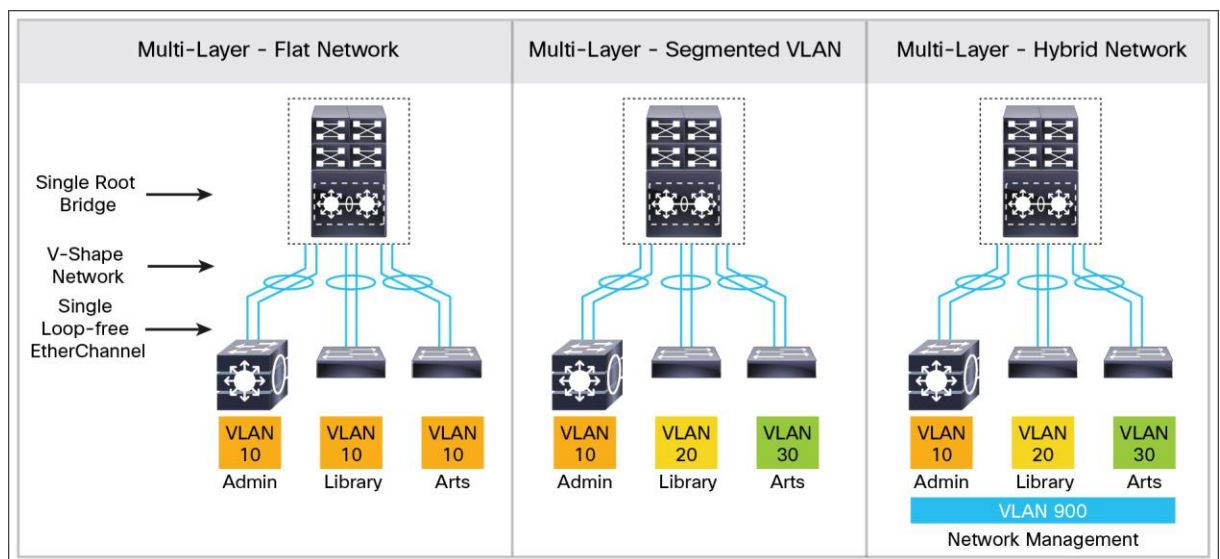
- **Flat:** Certain applications and types of user access require that the broadcast domain design span more than a single wiring closet switch. The multilayer network design provides the flexibility to build a single large broadcast domain with an extended star topology.

Such flexibility introduces scalability, performance, and security challenges and might require extra attention to protect the network against misconfiguration and miswiring, which can create spanning tree loops and destabilize the network.

- **Segmented:** Provides a unique VLAN for different organizational divisions and enterprise business functional segments to build a per-department logical network. All network communication between various enterprise and administrative groups passes through the routing and forwarding policies defined at the distribution layer.
- **Hybrid:** A hybrid logical network design segments VLAN workgroups that do not span different access-layer switches and allows certain VLANs (for example, that network management VLAN) to span across the access-distribution block. The hybrid network design enables flat Layer 2 communication without affecting the network and also helps reduce the number of subnets used.

Figure 18 illustrates the three unique VLAN design variations for the multilayer network.

Figure 18. Campus Multilayer VLAN Network Design Alternatives



As a best practice, it is recommended that the hybrid multilayer access-distribution block design provide a loop-free network topology with smaller broadcast and fault domains with limited VLANs that span across all access switches required for operational simplicity, such as the device management VLAN.

Multilayer Network Protocols Best Practices

The traditional multilayer network is composed with several protocols to build scalable and high-performance aggregation blocks in enterprise campus or branch office networks. Each network protocol is specifically designed to construct dynamical addressing, optimal forwarding tables, or high availability to provide resiliency for mission-critical networks.

This subsection provides best practices and deployment guidelines for deploying reliable Layer 2 network infrastructure in the distribution and access layer.

VLAN Trunking Protocol Recommendations

VLAN Trunking Protocol (VTP) is a Cisco proprietary Layer 2 messaging protocol that manages the addition, deletion, and renaming of VLANs on a networkwide basis. Cisco's VTP simplifies administration in a switched network. VTP can be configured in three modes: server, client, and transparent.

As a best practice, deploying VTP in transparent mode for better VLAN control, security, and manageability is recommended.

Table 19. Cisco VTP Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#vtp domain <name> 6800-VSS(config)#vtp mode transparent 6800-VSS(config)#vtp version 2 6800-VSS(config)#vtp password <password>
Access	Cisco Catalyst 4K/3K/2K 4500E(config)#vtp domain <name> 4500E(config)#vtp mode transparent 4500E(config)#vtp version 2 4500E(config)#vtp password <password>

Dynamic Trunking Protocol (DTP) Recommendations

Cisco Dynamic Trunk Protocol (DTP) is by default enabled on all Layer 2 Ethernet ports. Cisco DTP protocol make sure that the different parameters involved in sending IEEE 802.1Q frames, such as the configured encapsulation type, native VLAN, and hardware capability, are agreed upon by the switches at either end of a trunk. Cisco DTP also helps protect against nontrunk ports flooding tagged frames, a potentially serious security risk, by making sure that ports and their neighbors are in consistent states.

As a best practice, Cisco recommends retaining DTP settings at their default values on Layer 2 trunk ports. These recommendations should be maintained regardless of peering device type, that is, switch, router, wireless LAN controller, firewall, and so on.

Table 20. Cisco Dynamic Trunking Protocol Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#interface range <sw1-id> - <sw2-id> 6800-VSS(config-if)#switchport 6800-VSS(config-if)#switchport mode trunk 6800-VSS(config-if)#default switchport nonegotiate
Access	Cisco Catalyst 4K/3K/2K 4500E(config)#interface range <id> - <id> 4500E(config-if)#switchport 4500E(config-if)#switchport mode trunk 4500E(config-if)#default switchport nonegotiate

VLAN Trunk Design Recommendations

In a typical campus network design, a single access switch is deployed with more than a single VLAN, such as a data VLAN and a voice VLAN. The Layer 2 network connection between the distribution and access device is a trunk interface. A VLAN tag is added to maintain logical separation between VLANs across the trunk.

Implementing 802.1Q trunk encapsulation in static mode instead of negotiating mode to improve the rapid link bring-up performance is recommended.

Trunk VLAN Best Practices

Enabling the Layer 2 trunk on a port channel automatically enables communication for all of the active VLANs between access and distribution. This might adversely affect the large-scale network because the access-layer switch might receive traffic flood destined to another access switch. Hence it is important to limit traffic on Layer 2 trunk ports by statically allowing the active VLANs to make sure of efficient and secure network performance. Allowing only assigned VLANs on a trunk port automatically filters the rest.

Table 21. Layer 2 Trunk Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#interface range <sw1-id> - <sw2-id> 6800-VSS(config-if)#switchport trunk allowed vlan <range>
Access	Cisco Catalyst 4K/3K/2K 4500E(config)#interface range <id> - <id> 4500E(config-if)#switchport trunk allowed vlan <range>

Native VLAN Best Practices

By default on Cisco Catalyst switches, the native VLAN on each Layer 2 trunk port is VLAN 1 and cannot be disabled or removed from the VLAN database. The native VLAN remains active on all access switch Layer 2 ports. The default native VLAN must be properly configured to avoid several security risks: worms, viruses, or data theft. Any malicious traffic originated in VLAN 1 will span across the access-layer network. With a VLAN-hopping attack, it is possible to attack a system that does not reside in VLAN 1.

As a best practice, the best way to mitigate this security risk is to implement an unused and unique VLAN ID as a native VLAN on the Layer 2 trunk between the access and distribution switches. For example, configure VLAN 801 in access switch 1 and in the distribution switch. Then change the default native VLAN setting in both the switches. Thereafter, VLAN 801 must not be used anywhere for any purpose in the same access-distribution block.

Table 22. Native VLAN Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#interface range <sw1-id> - <sw2-id> 6800-VSS(config)#desc Connected to Access SW-1 6800-VSS(config-if)#switchport trunk native vlan <id-1> 6800-VSS(config)#interface range <sw1-id> - <sw2-id> 6800-VSS(config)#desc Connected to Access SW-2 6800-VSS(config-if)#switchport trunk native vlan <id-2>
Access (SW1)	Cisco Catalyst 4K/3K/2K and any other devices 4500E(config)#interface range <id> - <id> 4500E(config-if)#switchport trunk native vlan <id-1>

Network Layer	Cisco Catalyst Switch
Access (SW2)	Cisco Catalyst 4K/3K/2K and any other devices 3850(config)#interface range <id> - <id> 3850(config-if)#switchport trunk native vlan <id-2>

Spanning Tree Protocol Recommendations

Spanning Tree Protocol (STP) is a Layer 2 protocol that prevents logical loops in switched networks with redundant links. The borderless campus design uses an EtherChannel or MEC (point-to-point logical Layer 2 bundle) connection between access-layer and distribution switches, which inherently simplifies the STP topology and operation. In this point-to-point network design, the STP operation is done on a logical EtherChannel port instead of per-physical port; therefore, it will be assigned automatically in a forwarding state for all assigned VLANs. Over the years, STP protocols have evolved into the following versions:

- **Per-VLAN Spanning Tree Plus (PVST+):** Provides a separate 802.1D STP for each active VLAN in the network.
- **IEEE 802.1w-rapid PVST+:** Provides an instance of Rapid Spanning Tree Protocol (RSTP) (802.1w) per VLAN. It is easy to implement, proven in large-scale networks that support up to 3000 logical ports on Cisco Catalyst 6800 Series systems, and greatly improves network restoration time.
- **IEEE 802.1s Multiple Spanning-Tree (MST):** Provides up to 16 instances of RSTP (802.1w) and combines many VLANs with the same physical and logical topology into a common RSTP instance.

As a best practice, enabling the rapid PVST+ STP protocol in the multilayer network design on distribution and each access-layer systems is recommended. For large-scale distribution blocks, the network administrator can consider IEEE MST as an alternate solution to simplify spanning tree instances. The distribution-layer system, which provides an aggregation point of a physical and Layer 2 network of complete building blocks, should be statically designated as STP root switch for all extended VLAN ranges. The STP VLAN priority at the access layer should retain its default values. If default settings are modified, then it should be reverted to its default priorities, as illustrated in Table 23.

Table 23. Cisco Spanning Tree Protocol Mode and Root Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#spanning-tree mode rapid-pvst 6800-VSS(config)#spanning-tree vlan 1-4094 root primary
Access	Cisco Catalyst 4K/3K/2K 4500E(config)#spanning-tree mode rapid-pvst 4500E(config)#default spanning-tree vlan 1-4094 root

Unidirectional Link Detection Recommendations

Unidirectional Link Detection (UDLD) is a Layer 2 protocol that works with Layer 1 features to determine the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identity of neighbors and shutting down misconnected ports. When autonegotiation and UDLD are both enabled, the Layer 1 and Layer 2 detection methods work together to prevent physical and logical unidirectional connections and protocol malfunctions. The UDLD protocol functions transparently on Layer 2 or Layer 3 physical ports.

As a best practice at the protocol level, the unidirectional communication between two systems should be deployed based on these recommendations:

- **Layer 2 network:** In the multilayer standalone or EtherChannel-based network design, the UDLD protocol can be enabled on a per-trunk port level between the access and distribution switches.
- **Layer 3 ECMP:** In the Layer 3 ECMP-based campus core or in a routed access network design, the unidirectional communication between two systems can be detected by Layer 3 routing protocols instead UDLD as it operates on per-physical interface basis.
- **Layer 3 EtherChannel:** In a recommended EtherChannel-based network design, the UDLD should be implemented between two Layer 3 systems. Enabling UDLD on each Layer 3 member links can detect unidirectional forwarding paths and disables such ports from MEC.

UDLD operates in one of the following two modes globally at the system level:

- **Normal mode (recommended):** If bidirectional UDLD states that information times out, it is assumed there is no fault in the network, and no further actions are taken. The port state for UDLD is marked as undetermined, and the port behaves according to its STP state. Retaining UDLD message time at its default values to minimize false-positive situations during high CPU, interface congestion, or supervisor switchover conditions is recommended.
- **Aggressive mode:** If bidirectional UDLD states that information times out, UDLD attempts to reestablish the state of the port, provided it detects that the link on the port is operational. Failure to reestablish communication with the UDLD neighbor forces the port into the err-disable state, which must be manually recovered by either the user or the switch if it is be configured for autorecovery within a specified time interval.

Table 24. Cisco UDLD Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)# udld enable 6800-VSS(config)# default udld message
Access	Cisco Catalyst 4K/3K/2K 4500E(config)# udld enable 4500E(config)# default udld message

VSS MAC Address Table Synchronization Recommendations

A Cisco VSS ACTIVE system communicates with peering network devices and endpoints and continues to maintain a fully distributed forwarding architecture. Without any additional configuration and operator-level complexity, the MAC address information is globally synchronized across both chassis. This MAC address synchronization process is known as MAC out-of-band (OOB) synchronization. This forwarding information is dynamically programmed on DFCs of line cards and supervisor Policy Feature Card (PFC) for the local switching decision process. By default, the MAC address information is synchronized every 160 seconds to peer supervisor and DFC-enabled line-card modules.

As a best practice, retaining the default MAC OOB timer setting of 160 seconds with aging timer set to threefold at 480 seconds is recommended.

Table 25. Cisco VSS MAC OOB Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#mac address-table synchronize activity-time 160

Campus Core-Layer Network Design Best Practices

Core Uplink Design Recommendations

As described in the [Distribution-Layer Network Connectivity Best Practices](#) section, building a full-mesh physical network design between distribution and core-layer systems enables the best level of load balancing and resiliency in the campus backbone. It is imperative to design a highly reliable and stable core network infrastructure with the goal that it might have negligible application effects by completing nondisruptively during planned or unplanned network outages. The stringent high-availability requirement should also not compromise the operational complexity by overengineering the virtualized system and network, which deliver the same results with built-in various levels of intelligence in Cisco IOS Software.

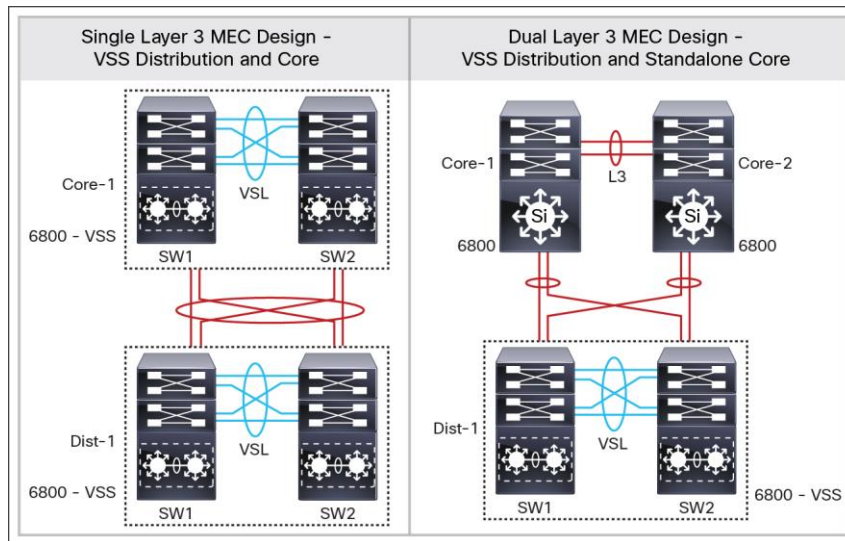
As a best practice, an additional equal set of full-mesh physical interfaces can be installed between core and distribution-layer systems for higher capacity and resiliency.

Cisco Multi-Chassis Layer 3 EtherChannel Best Practices

The Cisco VSS system is represented as a single large and logical system to the peering devices and networking protocols. For optimal load balancing and redundancy, the dynamic unicast and multicast routing protocols are required to operate individually on each parallel path installed between distribution and core-layer systems. The Cisco Express Forwarding load-balancing hash algorithm further computes a unique value combining Layer 3 and Layer 4 tuples of unicast IPv4 or IPv6 packets to determine the data forwarding results to each Layer 3 interface. This distribution and core network design can be greatly simplified and optimized with Layer 3 MEC as applied in a multilayer environment to simplify STP with Layer 2 MEC.

As a best practice, implementing Layer 3 MEC between distribution and core-layer systems is crucial to building a scalable, high-performance, resilient, and highly simple infrastructure. The Layer 3 MEC design depends on a core-layer system design whether it is in Cisco Catalyst 6800 Series VSS mode or in traditional standalone mode. Figure 19 illustrates Layer 3 MEC design recommendations for both system modes in core layer, while a Cisco Catalyst 6800 Series system is deployed in VSS mode in the campus distribution-layer network.

Figure 19. Multichassis Layer 3 EtherChannel Best Practices



Enhanced Interior Gateway Routing Protocol Design Recommendations

Enhanced Interior Gateway Routing Protocol (EIGRP) is a balanced hybrid routing protocol that builds neighbor adjacency and flat routing topology on a per-autonomous system basis. When designing and deploying in an enterprise campus network, it is imperative to build a common autonomous system and enable various best practices to optimize EIGRP performance and build secure routing adjacencies and resilient design for rapid fault recovery during failure conditions.

Autonomous System and Network Best Practices

As a best practice, Cisco recommends considering the following critical design tasks before implementing EIGRP in the campus distribution and core-layer network:

- **EIGRP autonomous system:** The enterprise campus infrastructure must be deployed in a single EIGRP autonomous system, which reduces operational tasks and prevents route redistribution, loops, and other problems that might occur because of misconfiguration. Multiple EIGRP autonomous systems and route redistribution must be avoided to build a reliable and scalable routing infrastructure.
- **EIGRP router ID:** Each EIGRP system in the network should be configured with a static and networkwide unique router ID. As a best practice, using one of the local loopback interfaces for better reliability and stability in an EIGRP-enabled network is recommended.
- **EIGRP autosummary:** All classless networks by default are advertised by an EIGRP system to neighbors without performing any automatic route summarization. As a best practice, preventing automatic route summarization that might create overlap summary entries in the network and effects to the users and applications is recommended. By default, autosummary is disabled and recommended to retain its default values.

Table 26. EIGRP Autonomous System Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#router eigrp <AS ID> 6800-VSS(config-router)#eigrp router-id <loopback_ip_address> 6800-VSS(config-router)#no auto-summary 6800-VSS(config-router)#eigrp log-neighbor-changes 6800-VSS(config-router)#network <address> <wildcard_mask>

Secured Routing Best Practices

As best practice, Cisco recommends securing EIGRP routing adjacency end to end to enable network-level protection. By securing the EIGRP domain, the routing adjacencies with trusted and managed peer devices are protected and operate in a controlled and deterministic order. The trusted peer's adjacency should establish the neighbor relationship with the message digest algorithm 5 (MD5) key to make sure that the communication between neighbors is encrypted and secured over the network. This increases network infrastructure efficiency and protection by securing the EIGRP adjacencies with internal systems.

- EIGRP neighbor control:** Block EIGRP neighbor processing with passive-mode configuration on physical or logical interfaces connected to non-EIGRP devices in the network, such as PCs, wireless LAN controllers, and so on. This best practice helps reduce CPU utilization and secures the network with unprotected EIGRP adjacencies with untrusted devices. Table 27 provides best practices guidelines to enable EIGRP protocol communication on trusted interfaces and blocks on all system interfaces. This recommended best practice must be enabled on all of the EIGRP Layer 3 systems in the network.

Table 27. EIGRP Neighbor Control Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#router eigrp <AS ID> 6800-VSS(config-router)#passive-interface default 6800-VSS(config-router)#no passive-interface <L3 Interface ID>

- Network security:** Each EIGRP neighbor in the LAN/WAN network must be trusted using MD5 authentication methods on each EIGRP-enabled system in the network. As a best practice, the recommended EIGRP MD5 adjacency authentication configuration on each non-passive EIGRP interface is to establish secure communication with remote neighbors. This recommended best practice must be enabled on all of the EIGRP systems in the network.

Table 28. EIGRP Neighbor Authentication Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#key chain <name> 6800-VSS(config-router)#key <ID> 6800-VSS(config-router)#key-string <password> 6800-VSS(config)#interface <ID> 6800-VSS(config-if)#ip authentication mode eigrp <AS> md5 6800-VSS(config-if)#ip authentication key-chain eigrp <AS> <key-chain-name>

Network Route Summarization Best Practices

The Cisco EIGRP routing protocol allows network administrators to summarize multiple individual and contiguous networks into a single summary network before advertising to the neighbor. Route summarization helps improve network performance, stability, and convergence by hiding the fault of an individual network that requires each router in the network to synchronize the routing topology.

As a best practice, Cisco recommends designing and deploying structured IP subnet plans for easy expansion, better operation, management, and troubleshooting. The campus distribution layer is the recommended tier in the campus network for aggregating multiple classless networks into a single and unique classful network to represent in the campus backbone network. Table 29 provides EIGRP route summarization deployment guidelines on each EIGRP running a Layer 3 interface of the distribution-layer system.

Table 29. EIGRP Route Summarization Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#interface <ID> 6800-VSS(config-if)#ip summary-address eigrp <AS> <network> <address>

High-Availability Best Practices

Traditionally, campus network resiliency highly depends on routing protocol resiliency, which can rapidly detect network faults, propagate topology changes, and reconverge the network. The default settings for any routing protocols are soft to protect against inducing false positives into the network during such failures; however, the side effect to default settings is that they lead to a slower recovery process. Network administrators require deep network design and protocol-level understanding to fine-tune advanced parameters to improve the recovery time.

As a best practice, when the campus network is deployed based on VSS, then Cisco recommends retaining the routing protocol settings at their default values. Cisco VSS includes built-in hardware-based intelligence to rapidly detect faults and self-initiate the recovery process without overengineering any protocol configuration.

EIGRP Protocol Timer

EIGRP supports hello and hold of timers to maintain routing adjacencies with each of its peers. In a recommended campus network design, the Layer 3 interfaces are directly attached between two EIGRP neighbors, and no intermediate devices are installed to deploy a disjointed backbone network. For such a topology, the fault detection and recovery on each system are hardware driven to detect local path failure and initiate the forwarding recovery process to the next preinstalled path in the routing table. Such a recovery process provides a deterministic subsecond network convergence time during failure.

As a best practice, Cisco recommends retaining and not modifying the EIGRP hello and hold timer to its default settings. Tweaking EIGRP timers to an aggressive rate might adversely affect the system recovery process during SSO on VSS chassis.

Table 30. EIGRP Hello and Hold Timer Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#interface <ID> 6800-VSS(config-if)# default ip hello-interface eigrp <AS> 6800-VSS(config-if)# default ip hold-time eigrp <AS>

EIGRP Graceful Restart for Nonstop Forwarding

When implementing nonstop forwarding (NSF) technology in systems using SSO redundancy mode, network disruptions are transparent to campus users and applications, and high availability is provided even during periods when the control-plane processing module (supervisor/route processor) is reset. During a failure, the underlying Layer 3 NSF-capable protocols perform graceful network topology resynchronization. The preset forwarding information on the redundant processor or distributed line-card hardware remains intact and continues to switch network packets. This service availability significantly lowers the mean time to repair (MTTR) and increases the mean time between failure (MTBF) to achieve the highest level of network availability.

As a best practice, for VSS-enabled systems, enabling the NSF function for unicast routing protocol is a foundation-level requirement.

Table 31. EIGRP Graceful Restart NSF Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#router eigrp <AS> 6800-VSS(config-router)#nsf

Open Shortest Path First Routing Protocol Design Recommendations

Open Shortest Path First (OSPF) is a widely deployed IETF standard link state and adaptive routing protocol for the heterogeneous vendor enterprise network environment. Unlike EIGRP, the OSPF network builds a structured network boundary into multiple areas, which helps in propagating summarized network topology information and rapidly performs OSPF database computations for intelligent forwarding decisions.

Area and Network Design Best Practices

OSPF divides the routing boundaries into nonbackbone areas that connect to a single core backbone area; such design helps simplify administration and optimizes network traffic and resource utilization. The OSPF protocol supports various types of areas; this best practices guide recommends the following two area types to implement in an enterprise campus network:

- **Backbone area:** The campus core layer is the heart of the network backbone, and as a best practice it must be configured with an OSPF backbone area. The campus aggregation-layer system must be implemented in the area border router (ABR) role, which interconnects to the core backbone area and to the access-layer nonbackbone area OSPF systems. Cisco recommends that OSPF routing and backbone area design be contiguous.
- **Stub/Totally Stub area:** The campus access-layer network requires concise network topology information and the default route from the distribution-layer systems to reach external networks. In a multilayer network, the default gateway to end station is serviced by a distribution-layer switch, whereas Layer 2 switches in between are simply transparent transit devices. Therefore, the non-OSPF backbone area between the distribution and access layer must be configured into stub area or totally stub area mode. Only the nonbackbone area can be deployed into stub or totally stub area mode.

OSPF supports several network types, each designed to operate optimally in various types of network connectivity and designs. The default network type for the OSPF protocol running over an Ethernet-based network is broadcast. Ethernet is a multiaccess network that provides the flexibility to interconnect several OSPF neighbors deployed in a single Layer 2 broadcast domain.

As a best practice campus network design, each OSPF system must statically assign router identification based on one of the local loopback interfaces to provide stability into the routing domain. In addition, the two Layer 3 OSPF systems interconnect directly to each other, thus forming direct and point-to-point communication. Cisco recommends modifying the default OSPF network type from broadcast to point-to-point on systems running Cisco IOS Software and any supporting peering devices. An OSPF point-to-point network optimizes adjacencies by eliminating DR/BDR processing and reducing routing complexities between all OSPF-enabled systems.

Table 32. OSPF Area and Network Design Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)# router ospf <ID> 6800-VSS(config-router)# router-id <Loopback_IP_Address> 6800-VSS(config-router)# network <core_network> <wildcard_mask> area 0 6800-VSS(config-router)# network <loopback_network> <wildcard_mask> area 0
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)# router ospf <ID> 6800-VSS(config-router)# network <access_network> <wildcard_mask> area <non-backbone-area-id> 6800-VSS(config-router)# area <non-backbone-area-id> stub no-summary
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)# interface <ID> 6800-VSS(config-if)# ip ospf network point-to-point

Secured Routing Best Practices

As described earlier in the [Secured Routing Best Practices](#) subsection in the [EIGRP Routing Design Recommendations](#) section, Cisco recommends securing OSPF routing adjacency end to end to enable network-level protection. When the OSPF domain is secured, the routing adjacencies with trusted and managed peer devices are protected and operate in a controlled and deterministic order. The trusted peer's adjacency should establish the neighbor relationship with the MD5 key to make sure the communication between neighbors is encrypted and secured over the network. This increases network infrastructure efficiency and protection by securing the OSPF adjacencies with internal systems.

- OSPF neighbor control:** Block OSPF neighbor processing with passive-mode configuration on physical or logical interfaces connected to non-EIGRP devices in the network, such as PCs, wireless LAN controllers, and so on. This best practice helps reduce CPU utilization and secures the network with unprotected OSPF adjacencies with untrusted devices. Table 32 provides best practices guidelines to enable OSPF protocol communication on trusted interfaces and blocks on all system interfaces. This recommended best practice must be enabled on all of the OSPF Layer 3 systems in the network.

Table 33. OSPF Neighbor Control Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)# router ospf <ID> 6800-VSS(config-router)# passive-interface default 6800-VSS(config-router)# no passive-interface <L3 Interface ID>

- **Network security:** Each OSPF neighbor in the LAN/WAN network must be trusted by implementing and validating the MD5 authentication methods on each OSPF-enabled system in the network. As a best practice, the recommended OSPF MD5 adjacency authentication configuration on all implemented OSPF areas is to establish secure communication with remote neighbors. This recommended best practice must be enabled on all of the OSPF Layer 3 systems in the network.

Table 34. OSPF Neighbor Authentication Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#router ospf <ID> 6800-VSS(config-router)#area 0 authentication message-digest 6800-VSS(config-router)#area <non-backbone-area-id> authentication message-digest 6800-VSS(config)#interface <ID> 6800-VSS(config-if)#ip ospf message-digest-key <key> <password>

Network Route Summarization Best Practices

As a best practice, the OSPF route summarization must be performed at the ABRs that connect the OSPF backbone and several aggregated nonbackbones; typically ABR routers are the campus distribution. The primary benefit of route summarization is that it summarizes multiple individual and contiguous networks into a single summary network before advertising into the OSPF backbone area. Route summarization helps improve network performance, stability, and convergence by hiding the fault of an individual network that requires each router in the network to synchronize the routing topology.

Table 35. OSPF Route Summarization Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#router ospf <ID> 6800-VSS(config-router)#area <non-backbone-area> range <network> <mask>

High-Availability Best Practices

OSPF Graceful Restart for NSF

The OSPF NSF capability and helper function in Cisco IOS Software-based systems make sure OSPF adjacencies and dynamically learned routes are protected when an NSF system undergoes the SSO switchover process. By default, every OSPF system running Cisco IOS Software is enabled with the NSF helper function. For the OSPF graceful restart recovery process, every VSS or dual-supervisor OSPF system must implement NSF capability under the OSPF routing process.

The NSF signaling can be exchanged in two modes: Cisco proprietary and IETF standard based. As a best practice, Cisco recommends enabling “nsf cisco” or “nsf” when a Cisco VSS system is pairing OSPF with a peer system running Cisco IOS Software. If the peer system does not support Cisco NSF but supports IETF NSF, then the network administrator must configure “nsf ietf” to successfully complete the graceful recovery process with a peering IETF-capable system.

Table 36. OSPF Graceful Restart for NSF Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS ! Configure "nsf cisco" if remote OSPF system is Cisco device 6800-VSS(config)#router ospf <ID> 6800-VSS(config-router)#nsf cisco
Distribution	Cisco Catalyst 6800 VSS ! Configure "nsf ietf" if remote OSPF system is Non-Cisco device and supports IETF NSF capabilities 6800-VSS(config)#router ospf <ID> 6800-VSS(config-router)#nsf ietf

OSPF Autocost and Static Interface Cost

The metric of an OSPF interface determines the best forwarding path based on lower metric or cost to the destination. By default, the metric or cost of an OSPF interface is automatically derived based on a fixed formula (108/bandwidth in bps) on Cisco Catalyst switches running Cisco IOS Software. For example, the OSPF cost for a 10Gbps link is computed as 1. In the Layer 3 EtherChannel/MEC-based network design, bundling multiple 10Gbps or 40Gbps links into a logical port channel interface dynamically increases the aggregated bandwidth; however, the OSPF cost remains 1 because of the fixed and static autocost reference bandwidth formula used globally across the entire system. Similarly, the default OSPF cost at interface level can be adjusted to generate a value different from the automatically computed value in the Cisco IOS Software.

As a best practice, retaining the autocost reference bandwidth at the global routing process level and on a per-port level at its default values is recommended. The maximum default 1G cost does not get adjusted or forces OSPF topology to recompute the database, which causes further propagation and in some cases develops a suboptimal forwarding path in the network.

Table 37. OSPF Autocost and Static Interface Cost Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#router ospf <ID> 6800-VSS(config-router)#default auto-cost
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#interface <id> 6800-VSS(config-if)#default ip ospf cost

OSPF Protocol Timer and Database Configuration

OSPF supports a variety of protocol and database timers for rapid fault detection, propagation, and recovery processes in the network. The default values are graceful to stabilize the OSPF routing domain and minimize false positive conditions during a system or interface-level stability issue in any part of the OSPF network.

By default, OSPF routers transmit hello packets every 10 seconds and terminate OSPF adjacency if the neighbor fails to receive the hello packet within four intervals or 40 seconds of dead time. In this optimized and best practice network design, Cisco recommends retaining default OSPF hello and hold timers on all OSPF-enabled platforms running Cisco VSS and peering devices. Implementing aggressive hello processing timers and dead times might adversely affect graceful recovery processes on any of the redundant campus-layer systems such as VSS or dual-supervisor standalone systems.

Table 38. OSPF Protocol Timer Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#interface <id> 6800-VSS(config-if)#default ip ospf hello-interval 6800-VSS(config-if)#default ip ospf dead-interval

The OSPF database settings should be kept at the default values because, as with Cisco VSS, the local failure detection and recovery process is hardware based instead of mechanically protocol-based. Configuring advance OSPF database features or parameters might become redundant because Cisco VSS provides deterministic subsecond convergence with default OSPF values.

Table 39. OSPF Database Timer Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#router ospf <ID> 6800-VSS(config-router)#default timers lsa arrival 6800-VSS(config-router)#default timers throttle lsa 6800-VSS(config-router)#default timers throttle spf

Multicast Routing Protocol Recommendations

Because unicast communication is based on the one-to-one forwarding model, it becomes easier in routing and switching decisions to perform destination address lookup, determine the egress path by scanning forwarding tables, and switch traffic. In the unicast routing and switching technologies discussed in the previous section, the network might need to be made more efficient by allowing certain applications in which the same content or application must be replicated to multiple users.

IP multicast delivers source traffic to multiple receivers using the least amount of network resources as possible without placing an additional burden on the source or the receivers. The multicast data replication in the network is done by the Protocol Independent Multicast (PIM)-enabled system that dynamically builds forwarding tables.

PIM Sparse Mode Best Practices

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

The enterprise campus design must be able to build packet distribution trees that specify a unique forwarding path between the subnet of the source and each subnet containing members of the multicast group. A primary goal in distribution tree construction is to make sure that no more than one copy of each packet is forwarded on each branch of the tree.

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

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

Selecting the PIM mode depends on the multicast applications that use various mechanisms to build multicast distribution trees. Based on the multicast scale factor and centralized source deployment design for one-to-many multicast communication in unified access campus network infrastructures, Cisco recommends deploying static PIM-SM because it is efficient and intelligent in building a multicast distribution tree.

Table 40. IP Multicast PIM Static Rendezvous Point Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	<p>Cisco Catalyst 6800 VSS</p> <pre>6800-VSS(config)#ip multicast-routing 6800-VSS(config)#ip pim rp-address <RP_IP_Address> 6800-VSS(config)#interface loopback <ID> 6800-VSS(config-if)#ip pim sparse-mode 6800-VSS(config)#interface <ID> 6800-VSS(config-if)#ip pim sparse-mode 6800-VSS(config)#interface VLAN <ID> 6800-VSS(config-if)#ip pim sparse-mode</pre>

Secured Multicast Best Practices

As a best practice, when designing and deploying IP multicast in the enterprise campus network, the network administrator must prevent the following two primary security concerns to build a secured multicast network infrastructure:

- **Rogue multicast source:** In a PIM-SM network, an unwanted traffic source can be controlled with the pim accept-register configuration. When the source traffic hits the first-hop router, the first-hop router (DR) creates the (S, G) state and sends a PIM source register message to the rendezvous point. If the source is not listed in the accept-register filter list (configured on the rendezvous point), the rendezvous point rejects the register and sends back an immediate register-stop message to the DR. The drawback with this method of source filtering is that with the pim accept-register command on the rendezvous point, the PIM-SM (S, G) state is still created on the first-hop router of the source. This can result in traffic reaching receivers local to the source and located between the source and the rendezvous point. Furthermore, because the pim accept-register command works on the control plane of the rendezvous point, this can be used to overload the rendezvous point with fake register messages and possibly cause a denial-of-service (DoS) condition. As a best practice, a simple access-control list (ACL) should be applied to the rendezvous point to filter on only the source address combining with the accept-register configuration. It is also possible to filter the source and the group using an extended ACL on the rendezvous point.

Table 41. IP PIM Rogue Multicast Source Security Best Practices

Network Layer	Cisco Catalyst Switch
PIM rendezvous point at core or any network layer	Cisco Catalyst 6800 VSS 6800-VSS(config)#ip access-list extended <ACL_NAME> 6800-VSS(config-ext-nacl)#permit ip <MCAST_SRC_IP_Subnet> <wildcard_mask> <MCAST_GRP_Address> <wildcard_mask> 6800-VSS(config-ext-nacl)#deny ip any any 6800-VSS(config)#ip pim accept-register list <ACL_NAME>

- **Rogue PIM rendezvous point:** Like the multicast source, any router can be misconfigured or can maliciously advertise itself as a multicast rendezvous point in the network with the valid multicast group address. With a static rendezvous point configuration, each PIM-enabled router in the network can be configured to use static rendezvous point for the multicast source and override any other automatic rendezvous point or Bootstrap Router (BSR) multicast router announcement from the network.

As a best practice, each PIM-enabled system in the campus network should accept PIM announcements only from the static rendezvous point and ignore dynamic multicast group announcements from any other unmanaged rendezvous point in the network.

Table 42. IP PIM Rogue Rendezvous Point Security Best Practices

Network Layer	Cisco Catalyst Switch
All Layer 3 systems	Cisco Catalyst 6800 VSS 6800-VSS(config)#ip access-list standard <ACL_NAME> 6800-VSS(config-std-nacl)#permit 224.0.1.39 6800-VSS(config-std-nacl)#permit 224.0.1.40 6800-VSS(config-std-nacl)#deny any 6800-VSS(config)#ip pim rp-address <RP_IP_Address> <ACL_NAME> override

High-Availability Best Practices

The Cisco Catalyst 6800 Series system deployed in VSS or dual-supervisor mode natively supports stateful IP multicast protocol redundancy with the same SSO technology as unicast routing protocols. The IP multicast protocol state machines and distributed forwarding entries information is hot synchronized with a secondary supervisor module to gracefully switch over the multicast flows during planned or unplanned network outage. For certain network designs and topologies, the PIM neighbor adjacency and routing timers can be adjusted for rapid fault detection and trigger a recovery process to rebuild multicast forwarding tables. However, in a VSS-enabled network infrastructure, the fault detection and recovery process works independently of any network protocols.

As a best practice, Cisco recommends IP multicast routing timers during SSO and neighbor query timers to be retained to their default values, allowing the VSS system to gracefully recover during supervisor switchover. While the new supervisor undergoes the recovery process, the multicast data continues to switch in the network with the last good known forwarding entries. Modifying default values to an aggressive level might induce false positives into the multicast network and might affect the application.

Table 43. IP PIM Timers Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)# default ip multicast redundancy routeflush maxtime
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)# interface <ID> 6800-VSS(config-if)# default ip pim query-interval

In a high-scale IP multicast environment, when a port carries a quantity of multicast group traffic to different VLAN interfaces, that creates a large multicast routing table because of its unique egress outgoing-interface list (OIL). When such physical ports join or are removed from a configured EtherChannel, that does not disrupt the IP PIM multicast forwarding topology; the data rerouting decision is hardware based, and internally the VSS system must recompute a new forwarding rule to take the next available alternate forwarding path into the port channel. This recomputation might delay the application convergence for few seconds, which might not be acceptable in a mission-critical IP multicast campus network environment.

As a best practice, Cisco recommends fine-tuning the multicast recovery process on port-channel interfaces to enable a scale-independent and deterministic multicast-enabled application network recovery process.

Table 44. IP Multicast Forwarding Fast-Redirect Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)# interface Port-Channel <ID> 6800-VSS(config-if)# platform multicast forwarding fast-redirect

General Routing Recommendations

Equal Cost Multipath Routing Best Practices

Enterprise campus networks are deployed in variety of IP routing designs. When Cisco VSS technology is deployed in the campus core layer, then the VSS in the distribution layer should be combined with a single Layer 3 MEC in the backbone because its architecture will inherently enable a significant number of benefits to the network administrator that simplify network operation and tremendously improve system, network, and application performance.

However in certain cases, when deploying Cisco VSS technology in the campus core is not an option, it will change the routing infrastructure design, because physically there are two standalone core systems into the network with traditional independent control and management planes. The network administrator can still deploy two separate Layer 3 MEC EtherChannels with diversified fiber connections from VSS to each upstream standalone core system. This network design creates an ECMP routing design, and the egress data forwarding decision from each Layer 3 MEC is in dual phase: Cisco Express Forwarding load balancing and MEC load balancing.

As a best practice for ECMP-based Layer 3 networks, Cisco recommends fine-tuning Cisco Express Forwarding load balancing to include Layer 3 and Layer 4 tuple inclusion to compute and derive the first phase of the optimal forwarding decision process between two upstream Layer 3 MEC interfaces. The best practices for the second phase Layer 3 MEC load balancing is described in the [Multi-Chassis EtherChannel Load Balancing](#) section.

Table 45. ECMP Cisco Express Forwarding Load-Balancing Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)#platform ip cef load-sharing full

Unicast IP Route Entry Purge Best Practices

In ECMP-based routing design, the route entry install and purge function from the routing table was traditionally handled by Cisco IOS Software operating systems. To address the requirement for faster convergence for certain network topologies, the default behavior was considered to be slow to install the next best path from the routing database into the routing information base (RIB) and subsequently in the forwarding information base (FIB). To optimize convergence in such network topologies, the default behavior has been modified in recent Cisco IOS Software releases.

When the Layer 3 physical or entire MEC interface goes down, the dynamic route entry removal process is now a routing protocol function. The network administrator would require fine-tuning the OSPF Link-State Advertisement (LSA)/Shortest Path First (SPF) timers, which the Link-State Database (LSDB) can compute and can propagate the topology change to the rest of the OSPF network rapidly. The overall unicast data-plane network recovery time would highly depend on advanced OSPF tunings applied on the routing process on every system into the network.

As recommended earlier, when the enterprise campus network design is physically full mesh and has logically virtualized the control and management plane with Cisco VSS and MEC technology, the forwarding and rerouting decision process becomes hardware driven. Hence, as a best practice, the campus network can be deployed with simple configuration and use built-in Cisco VSS resilient technology to provide deterministic subsecond convergence without any advanced fine-tuning into the network. The network administrator must apply the following best practices on all Layer 3 ECMP networks to revert traditional Cisco IOS Software-based route install and purge behavior in the campus network.

Table 46. Unicast IP Route Entry Purge Best Practices

Network Layer	Cisco Catalyst Switch
All ECMP routing systems	Cisco Catalyst 6800 VSS 6800-VSS(config)#no ip routing protocol purge-interface

IP Event Dampening

Unstable physical network connectivity with poor signaling or loose connections might cause continuous port flaps. When the enterprise campus network is not deployed using best practice guidelines to summarize the network boundaries at the aggregation layer, a single interface flap can severely affect the stability and availability of the entire campus network. Route summarization is one technique used to isolate the fault domain and contain local network faults within the domain.

To make sure of local network domain stability during port flaps, all Layer 3 interfaces can be implemented with IP event dampening, which uses the same fundamental principles as Border Gateway Protocol (BGP) dampening. Each time the Layer 3 interface flaps, IP dampening tracks and records the flap event. On multiple flaps, a logical penalty is assigned to the port, and it suppresses link status notifications to IP routing until the port becomes stable.

As a best practice, turning on IP event dampening on Layer 3 ports running routing protocols is recommended, as illustrated in Table 47.

Table 47. IP Event Dampening Best Practices

Network Layer	Cisco Catalyst Switch
Distribution	Cisco Catalyst 6800 VSS 6800-VSS(config)# interface Port-Channel <id> 6800-VSS(config-if)no switchport 6800-VSS(config-if) dampening
Core	Cisco Catalyst 6800 VSS 6800-VSS(config)# interface Port-Channel <id> 6800-VSS(config-if)no switchport 6800-VSS(config-if) dampening

Summary

Enterprise campus networks enabled by Cisco Unified Access are built-upon a Cisco next-generation architecture that delivers a new workspace experience securely, reliably, and smoothly, connecting anyone, anywhere, using any device, to any resource. This unified experience is only possible with a strong and resilient intelligent network that is designed to meet the needs of a global workspace. The network platform enabled by Cisco is the primary component of this architecture, providing borderless services such as mobility, security, media awareness, location, and Cisco EnergyWise[®], to deliver an optimal user experience. Building network designs with intelligence at the edge provides mobility and secures collaboration, as well as the overall infrastructure backbone to provide networkwide differentiated services for a consistent, highly available, and reliable user experience. Cisco Unified Access enables innovative business models, creating new user experiences that lead to increased customer satisfaction and loyalty.

References

Borderless Campus Design Guide

http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Campus/Borderless_Campus_Network_1-0/Borderless_Campus_1-0_Design_Guide.pdf

Cisco VSS Design Guide

http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Campus/VSS30dg/campusVSS_DG.pdf

Enterprise Campus Quality of Service

http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book.pdf

Security

<http://www.cisco.com/c/en/us/solutions/enterprise/design-zone-security/index.html>

Campus LAN Cisco Validated Design

<http://www.cisco.com/c/dam/en/us/td/docs/solutions/CVD/Aug2014/CVD-CampusWiredLANDesignGuide-AUG14.pdf>




Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)