



CCDP Learning

Authorized Self-Study Guide  
**Designing Cisco Network  
Service Architectures (ARCH)**

Second Edition

Foundation learning for ARCH exam 642-873

# Authorized Self-Study Guide Designing Cisco Network Service Architectures (ARCH), Second Edition

Keith Hutton  
Mark Schofield  
Diane Teare

Copyright © 2009 Cisco Systems, Inc.

Published by:  
Cisco Press  
800 East 96th Street  
Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America

First Printing December 2009

Library of Congress Cataloging-in-Publication Data:

Hutton, Keith.

Authorized self-study guide : designing Cisco network service architectures (ARCH) / Keith Hutton, Mark Schofield, Diane Teare. -- 2nd ed.

p. cm.

ISBN 978-1-58705-574-4 (hardcover)

1. Computer network architectures--Examinations--Study guides. 2. Computer networks--Design--Examinations--Study guides. 3. Internetworking (Telecommunication)--Examinations--Study guides. I. Schofield, Mark. II. Teare, Diane. III. Title. IV. Title: Designing Cisco network service architectures (ARCH).

TK5105.52.H98 2008

004.6'5--dc22

2008049128

ISBN-13: 978-1-58705-574-4

ISBN-10: 1-58705-574-0

## Warning and Disclaimer

This book is designed to provide information about designing Cisco network service architectures. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

The Cisco Press self-study book series is as described, intended for self-study. It has not been designed for use in a classroom environment. Only Cisco Learning Partners displaying the following logos are authorized providers of Cisco curriculum. If you are using this book within the classroom of a training company that does not carry one of these logos, then you are not preparing with a Cisco trained and authorized provider. For information on Cisco Learning Partners please visit: [www.cisco.com/go/authorizedtraining](http://www.cisco.com/go/authorizedtraining). To provide Cisco with any information about what you may believe is unauthorized use of Cisco trademarks or copyrighted training material, please visit: <http://www.cisco.com/logo/infringement.html>.



## Foreword

Cisco Certification Self-Study Guides are excellent self-study resources for networking professionals to maintain and increase internetworking skills and to prepare for Cisco Career Certification exams. Cisco Career Certifications are recognized worldwide and provide valuable, measurable rewards to networking professionals and their employers.

Cisco Press exam certification guides and preparation materials offer exceptional—and flexible—access to the knowledge and information required to stay current in one's field of expertise or to gain new skills. Whether used to increase internetworking skills or as a supplement to a formal certification preparation course, these materials offer networking professionals the information and knowledge required to perform on-the-job tasks proficiently.

Developed in conjunction with the Cisco certifications and training team, Cisco Press books are the only self-study books authorized by Cisco, and they offer students a series of exam practice tools and resource materials to help ensure that learners fully grasp the concepts and information presented.

Additional authorized Cisco instructor-led courses, e-learning, labs, and simulations are available exclusively from Cisco Learning Solutions Partners worldwide. To learn more, visit <http://www.cisco.com/go/training>.

I hope you will find this guide to be an essential part of your exam preparation and professional development, as well as a valuable addition to your personal library.

Drew Rosen  
Manager, Learning & Development  
Learning@Cisco  
September 2008

## Introduction

*Designing Cisco Network Service Architectures (ARCH)*, Second Edition, covers how to perform the conceptual, intermediate, and detailed design of a network infrastructure. This design supports network solutions over intelligent network services to achieve effective performance, scalability, and availability of the network. This book enables readers, applying solid Cisco network solution models and best design practices, to provide viable and stable enterprise internetworking solutions. In addition, the book has been written to help candidates prepare for the Designing Cisco Network Service Architectures Exam (642-873 ARCH). This exam is one of the requirements for the CCDP certification. This exam tests a candidate's knowledge of the latest development in network design and technologies, including network infrastructure, intelligent network services, and converged network solutions.

Since the first edition was published in 2004, the ARCH course has changed to reflect the new exam requirements. This led to the immediate need for an update to this examination preparation text. Readers of the previous edition of *Designing Cisco Network Architectures (ARCH)* can use this text to update their knowledge and skill sets.

## Goals of This Book

Upon completing this book, you will be able to meet these objectives:

- Introduce the Cisco Service-Oriented Network Architecture (SONA) framework, and explain how it addresses enterprise network needs for performance, scalability, and availability
- Describe how the Cisco Enterprise Architectures are used in the SONA framework for designing enterprise networks
- Create intermediate and detailed enterprise campus network, enterprise edge, and remote infrastructure designs that offer effective functionality, performance, scalability, and availability
- Create conceptual, intermediate, and detailed intelligent network service designs for network management, high availability, security, quality of service (QoS), and IP multicast
- Create conceptual, intermediate, and detailed virtual private network (VPN) designs
- Create conceptual, intermediate, and detailed voice over wireless network designs

## Prerequisite Knowledge

Although enthusiastic readers will tackle less-familiar topics with some energy, a sound grounding in networking is advised. To gain the most from this book, you should be familiar with internetworking technologies, Cisco products, and Cisco IOS Software features. You will find knowledge about the following topics helpful for your successful understanding of the material presented in this book:

- How to design the necessary services to extend IP addresses using variable-length subnet masking (VLSM), Network Address Translation (NAT), and route summarization
- How to implement appropriate networking routing protocols, such as Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), and Border Gateway Protocol (BGP) on an existing internetwork
- How to redistribute routes between different routing protocols
- The required Cisco products and services that enable connectivity and traffic transport for a multilayer campus network
- The necessary services at each layer of the network to enable all users to obtain membership in multicast groups in a working enterprise network
- How to control network traffic by implementing the necessary admission policy at each layer of the network topology
- How to identify the appropriate hardware and software solutions for a given set of WAN technology requirements, including access between a central campus, branch offices, and telecommuters
- The Cisco equipment to establish appropriate WAN connections
- How to use protocols and technologies that enable traffic flow between multiple sites while minimizing the amount of overhead traffic on each connection
- QoS capabilities to ensure that mission-critical applications receive the required bandwidth within a given WAN topology
- How to implement Cisco voice solutions
- How to implement Cisco wireless solutions
- How to implement basic security steps and mitigation techniques

## How This Book Is Organized

Of course, you can read the chapters in this book sequentially, but the organization also allows you to focus your reading on specific topics of interest. For example, if you want to focus on advanced routing design, you can skim Chapters 1 and 2 (which cover SONA and the elements of the enterprise campus network design), and then focus on the advanced IP addressing and routing topics in Chapter 3. Each chapter examines topics around a specific set of design issues. Specifically, the chapters in this book cover the following topics:

- Chapter 1, “Cisco SONA and the Cisco Enterprise Architecture,” introduces the hierarchical model. It reviews Cisco SONA framework. This chapter also introduces the Cisco Enterprise Campus Architecture and reviews the Cisco PPDI/O network lifecycle approach.
- Chapter 2, “Enterprise Campus Network Design,” reviews high-availability designs and how to implement optimal redundancy. An in-depth look at recommended practices for Layer 2 and Layer 3 design elements follows. A discussion of the Layer 2 to Layer 3 boundary designs and issues concludes with a number of considerations for supporting infrastructure services.
- Chapter 3, “Developing an Optimum Design for Layer 3,” begins by reviewing the importance of IP address planning, and then covers advanced routing elements. Discussions focus on scalable EIGRP, OSPF, and BGP designs.
- Chapter 4, “Advanced WAN Services Design Considerations,” covers advanced WAN service layers. This overview goes into more detail about the common WAN optical technologies of SONET, SDH, DWDM, and Resilient Packet Ring. A discussion about Metro Ethernet, VPLS, and MPLS VPN technologies follows (and includes an examination of a number of design considerations). The discussion then turns to implementing advanced WAN services.
- Chapter 5, “Enterprise Data Center Design,” focuses on the enterprise data center, and covers the data center architecture model and design consideration in the data center core, aggregation, and access layers. The discussion then turns to scaling, with a look at how to scale a three-layer data center architecture.
- Chapter 6, “SAN Design Considerations,” covers storage-area networks, from components and topologies to SAN technologies. SAN design factors center on port density and topology, with some discussion about extending the SAN with various protocols.
- Chapter 7, “E-Commerce Module Design,” begins with an e-commerce overview and a look at the components of high availability in this module. The chapter covers common e-commerce design components, designing an integrated e-commerce architecture, and how to fine-tune e-commerce designs.

- Chapter 8, “Security Services Design,” delves into designing firewall services in various scenarios. The chapter also covers network admission control services, with a review of Cisco NAC appliance fundamentals and NAS deployment options and designs. The discussion then turns to intrusion detection and prevention design.
- Chapter 9, “IPsec and SSL VPN Design,” examines remote-access VPN design. Site-to-site VPN designs are covered, too. This chapter also covers IPsec VPN technologies, including Cisco Easy VPN, GRE over IPsec, and DMVPN. Recommendations for managing VPNs and considerations for scaling VPNs conclude the chapter.
- Chapter 10, “IP Multicast Design,” covers IP multicast and multicast routing. Topics covered in this chapter include Protocol Independent Multicast (PIM), rendezvous points, and securing IP multicast.
- Chapter 11, “VoWLAN Design,” introduces the Cisco Unified Wireless Network and examines requirements for voice over WLAN in the enterprise network. This chapter also discusses VoWLAN coverage considerations and the site survey process.
- Chapter 12, “Network Management Capabilities with Cisco IOS Software,” examines Cisco network management capabilities embedded in Cisco IOS Software. This chapter also covers the syslog process, NetFlow, and NBAR, with a focus on the Cisco technologies themselves and how they enable other discovery tools, including Cisco AutoQoS. The chapter concludes with an overview of IP SLAs measurements.

This book also contains an appendix and an acronym list:

- Appendix A, “Answers to Review Questions,” provides the answers to all the chapter-ending review questions.
- “Acronyms and Abbreviations,” identifies abbreviations, acronyms, and initialisms used in this book.

**Note:** The website references in this book were accurate at the time of this writing. However, some might have changed since then. If a URL is unavailable, you can always search using the title as keywords in your favorite search engine.

# Enterprise Campus Network Design

---

The complexity inherent in today's campus networks necessitates a design process capable of separating solutions into basic elements. The Cisco hierarchical network model achieves this goal by dividing the network infrastructure into modular components. Each module is used to represent a functional service layer within the campus hierarchy.

---

## Designing High Availability in the Enterprise Campus

---

The Cisco hierarchical network model enables the design of high-availability modular topologies. Through the use of scalable building blocks, the network can support evolving business needs. The modular approach makes the network easier to scale, troubleshoot, and understand. It also promotes the deterministic traffic patterns.

This section reviews design models, recommended practices, and methodologies for high availability in the Cisco Enterprise Campus Architecture infrastructure.

## Enterprise Campus Infrastructure Review

The building blocks of the enterprise campus infrastructure are the access layer, the distribution layer, and the core layer. The principal features associated with each layer are hierarchical design and modularity. A hierarchical design avoids the need for a fully meshed network in which all nodes are interconnected. A modular design enables a component to be placed in service or taken out of service with little or no impact on the rest of the network. This methodology also facilitates troubleshooting, problem isolation, and network management.

### Access Layer


The access layer is the point of entry into the network for end devices, as illustrated in Figure 2-1.

The campus access layer aggregates end users and provides uplinks to the distribution layer. The access layer can support multiple features:

- **High availability:** At the access layer, high availability is supported through various hardware and software attributes. With hardware, system-level redundancy can be provided using redundant supervisor engines and redundant power supplies. It can also be provided by default gateway redundancy using dual connections from access switches to redundant distribution layer switches. With software, high availability is supported through the use of first-hop routing protocols (FHRP), such as the Hot



Standby Router Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP), and Gateway Load Balancing Protocol (GLBP).



**Figure 2-1** *Access Layer*

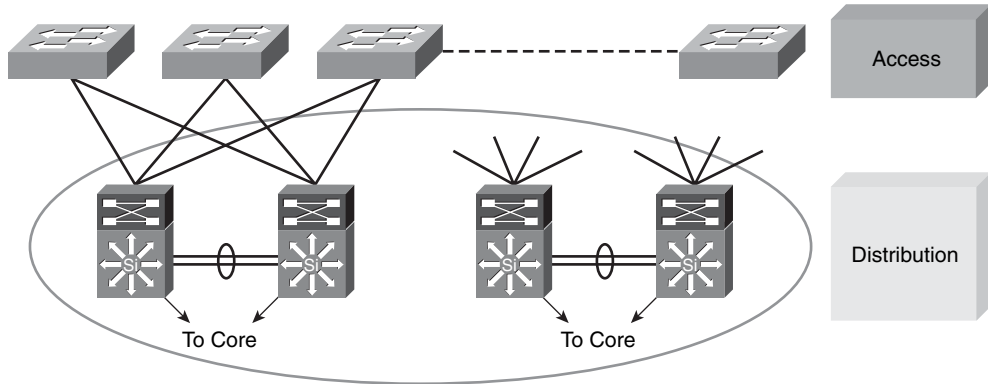
**Note:** Cisco offers a unique high-availability feature to its 3750 Workgroup Switch and Etherswitch Services Module called StackWise. StackWise technology enables switches to be interconnected to create a single logical unit through the use of special stack cables. The cables create a bidirectional path that behaves as a switch fabric for all the interconnected switches. The stack is managed as a single unit, eliminating the need for spanning tree and streamlining the interface to a single management session for all devices. For more information about StackWise, refer to Cisco.com.

**Note:** IOS Release 12.2(18) SXD extended high availability to the 6500/7600 series line of switches. It added services such as Control Plane Policing (CoPP), Nonstop Forwarding (NSF), Stateful Switchover (SSO), and Gateway Load Balancing Protocol (GLBP), which are discussed later in this chapter.

- **Convergence:** The access layer supports inline Power over Ethernet (PoE) for IP telephony and wireless access points, allowing customers to converge voice onto their data network and providing roaming wireless LAN (WLAN) access for users.
- **Security:** The access layer provides services for additional security against unauthorized access to the network through the use of tools such as IEEE 802.1x, port security, DHCP snooping, Dynamic ARP Inspection (DAI), and IP Source Guard.
- **Quality of service (QoS):** The access layer allows prioritization of mission-critical network traffic using traffic classification and queuing as close to the ingress of the network as possible. It supports the use of the QoS trust boundary.
- **IP multicast:** The access layer supports efficient network and bandwidth management using software features such as Internet Group Management Protocol (IGMP) snooping.

## Distribution Layer

The distribution layer aggregates traffic from all nodes and uplinks from the access layer and provides policy-based connectivity, as illustrated in Figure 2-2.



**Figure 2-2** *Distribution Layer*

Availability, load balancing, QoS, and provisioning are the important considerations at this layer. High availability is typically provided through dual paths from the distribution layer to the core and from the access layer to the distribution layer. Layer 3 equal-cost load sharing allows both uplinks from the distribution to the core layer to be used.


The distribution layer is the place where routing and packet manipulation are performed and can be a routing boundary between the access and core layers. The distribution layer represents a redistribution point between routing domains or the demarcation between static and dynamic routing protocols. The distribution layer performs tasks such as controlled routing and filtering to implement policy-based connectivity and QoS. To further improve routing protocol performance, the distribution layer summarizes routes from the access layer. For some networks, the distribution layer offers a default route to access layer routers and runs dynamic routing protocols when communicating with core routers.

The distribution layer uses a combination of Layer 2 and multilayer switching to segment workgroups and isolate network problems, preventing them from impacting the core layer. The distribution layer may be used to terminate VLANs from access layer switches. The distribution layer connects network services to the access layer and implements QoS, security, traffic loading, and routing policies. The distribution layer provides default gateway redundancy using an FHRP, such as HSRP, GLBP, or VRRP, to allow for the failure or removal of one of the distribution nodes without affecting endpoint connectivity to the default gateway.

**Note:** Cisco has introduced the Virtual Switching System (VSS), which can reduce or eliminate the need for FHRPs at the distribution layer. For more information about VSS, visit <http://www.cisco.com/go/vss>.

## Core Layer

The core layer provides scalability, high availability, and fast convergence to the network, as illustrated in Figure 2-3. The core layer is the backbone for campus connectivity, and is the aggregation point for the other layers and modules in the Cisco Enterprise Campus Architecture. The core provides a high level of redundancy and can adapt to changes quickly. Core devices are most reliable when they can accommodate failures by rerouting traffic and can respond quickly to changes in the network topology. The core devices implement scalable protocols and technologies, alternate paths, and load balancing. The core layer helps in scalability during future growth.

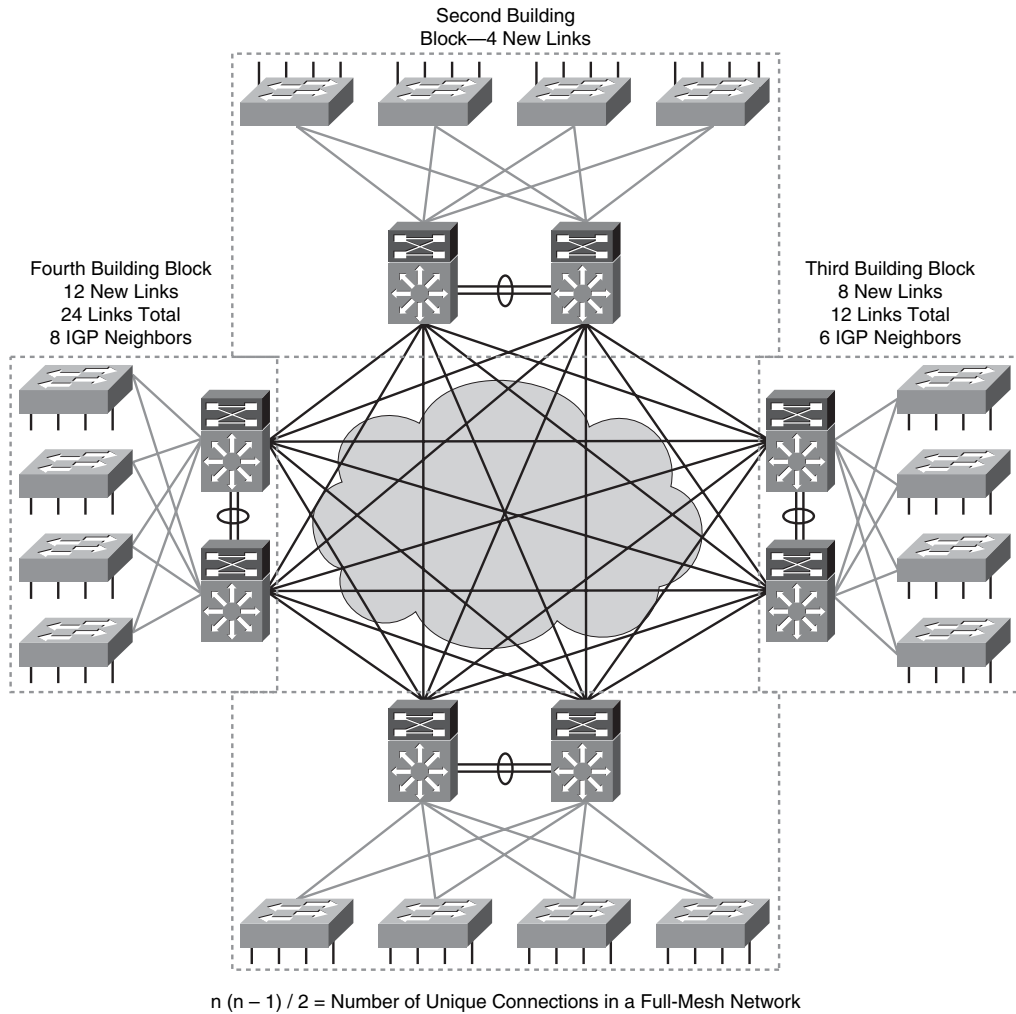


**Figure 2-3** Core Layer

The core is a high-speed, Layer 3 switching environment using hardware-accelerated services. For fast convergence around a link or node failure, the core uses redundant point-to-point Layer 3 interconnections because this design yields the fastest and most deterministic convergence results. The core layer is designed to avoid any packet manipulation, such as checking access lists and filtering, which would slow down the switching of packets.

Not all campus implementations require a campus core. The core and distribution layer functions can be combined at the distribution layer for a smaller campus.

Without a core layer, the distribution layer switches need to be fully meshed, as illustrated in Figure 2-4. This design can be difficult to scale, and increases the cabling requirements, because each new building distribution switch needs full-mesh connectivity to all the distribution switches. The routing complexity of a full-mesh design increases as new neighbors are added.



**Figure 2-4** *Is a Core Layer Needed?*

**Note:** Note that combining distribution and core layer functionality (collapsed core) requires a great deal of port density on the distribution layer switches. An alternative solution is a Layer 2 core with discrete VLANs on each core switch. This scenario requires only two ports per distribution layer switch—regardless of the number of buildings (switch blocks)—and so you can avoid the expense of multilayer core switches.

In Figure 2-4, a distribution module in the second building of two interconnected switches requires four additional links for full-mesh connectivity to the first module. A third distribution module to support the third building would require 8 additional links to support connections to all the distribution switches, or a total of 12 links. A fourth module supporting the fourth building would require 12 new links, for a total of 24 links

between the distribution switches. Four distribution modules impose eight Interior Gateway Protocol (IGP) neighbors on each distribution switch.

As a recommended practice, deploy a dedicated campus core layer to connect three or more buildings in the enterprise campus, or four or more pairs of building distribution switches in a very large campus. The campus core helps make scaling the network easier by addressing the requirements for the following:

- Gigabit density
- Data and voice integration
- LAN, WAN, and MAN convergence

## High-Availability Considerations

In the campus, high availability is concerned with minimizing link and node failures and optimizing recovery times to minimize convergence and downtime.


### Implement Optimal Redundancy

The recommended design is redundant distribution layer switches and redundant connections to the core with a Layer 3 link between the distribution switches. Access switches should have redundant connections to redundant distribution switches, as illustrated in Figure 2-5.

As a recommended practice, the core and distribution layers are built with redundant switches and fully meshed links to provide maximum redundancy and optimal convergence. Access switches should have redundant connections to redundant distribution switches. The network bandwidth and capacity is engineered to withstand a switch or link failure, supporting 120 to 200 ms to converge around most events. Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (EIGRP) timer manipulation attempt to quickly redirect the flow of traffic away from a router that has experienced a failure toward an alternate path.

In a fully redundant topology with tuned IGP timers, adding redundant supervisors with Cisco NSF and SSO may cause longer convergence times than single supervisors with tuned IGP timers. NSF attempts to maintain the flow of traffic through a router that has experienced a failure. NSF with SSO is designed to maintain a link-up Layer 3 up state during a routing convergence event. However, because an interaction occurs between the IGP timers and the NSF timers, the tuned IGP timers can cause NSF-aware neighbors to reset the neighbor relationships.

**Note:** Combining OSPF and EIGRP timer manipulation with Cisco NSF might not be the most common deployment environment. OSPF and EIGRP timer manipulation is designed to improve convergence time in a multiaccess network (where several IGP routing peers share a common broadcast media, such as Ethernet). The primary deployment scenario for Cisco NSF with SSO is in the enterprise network edge. Here, the data link layer generally consists of point-to-point links either to service providers or redundant Gigabit Ethernet point-to-point links to the campus infrastructure.

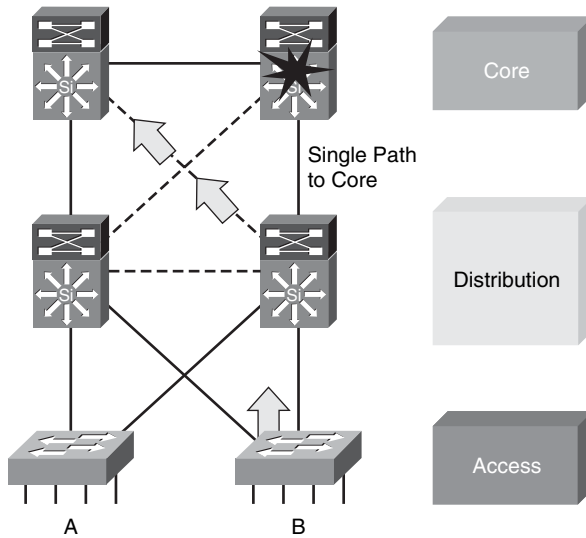


**Figure 2-5** *Optimal Redundancy*

In nonredundant topologies, using Cisco NSF with SSO and redundant supervisors can provide significant resiliency improvements.

### Provide Alternate Paths

The recommended distribution layer design is redundant distribution layer switches and redundant connections to the core with a Layer 3 link between the distribution switches, as illustrated in Figure 2-6.



**Figure 2-6** *Provide Alternate Paths*


Although dual distribution switches connected individually to separate core switches will reduce peer relationships and port counts in the core layer, this design does not provide sufficient redundancy. In the event of a link or core switch failure, traffic will be dropped.

An additional link providing an alternate path to a second core switch from each distribution switch offers redundancy to support a single link or node failure. A link between the two distribution switches is needed to support summarization of routing information from the distribution layer to the core.

### Avoid Single Points of Failure

Cisco NSF with SSO and redundant supervisors has the most impact in the campus in the access layer. An access switch failure is a single point of failure that causes outage for the end devices connected to it. You can reduce the outage to one to three seconds in this access layer, as shown in Figure 2-7, by using SSO in a Layer 2 environment or Cisco NSF with SSO in a Layer 3 environment.

**Note:** The SSO feature is available on the Catalyst 4500 and 6500/7600 switches.



**Figure 2-7** *Avoid Single Points of Failure*



## Cisco NSF with SSO

Cisco NSF with SSO is a supervisor redundancy mechanism in Cisco IOS Software that allows extremely fast supervisor switchover at Layers 2 to 4.

SSO allows the standby route processor (RP) to take control of the device after a hardware or software fault on the active RP. SSO synchronizes startup configuration, startup variables, and running configuration; and dynamic runtime data, including Layer 2 protocol states for trunks and ports, hardware Layer 2 and Layer 3 tables (MAC, Forwarding Information Base [FIB], and adjacency tables) and access control lists (ACL) and QoS tables.

Cisco NSF is a Layer 3 function that works with SSO to minimize the amount of time a network is unavailable to its users following a switchover. The main objective of Cisco NSF is to continue forwarding IP packets following an RP switchover. Cisco NSF is supported by the EIGRP, OSPF, Intermediate System-to-Intermediate System (IS-IS), and Border Gateway Protocol (BGP) for routing. A router running these protocols can detect an internal switchover and take the necessary actions to continue forwarding network traffic using Cisco Express Forwarding while recovering route information from the peer devices. With Cisco NSF, peer networking devices continue to forward packets while route convergence completes and do not experience routing flaps.

### Routing Protocol Requirements for Cisco NSF

Usually, when a router restarts, all its routing peers detect that routing adjacency went down and then came back up. This transition is called a *routing flap*, and the protocol state is not maintained. Routing flaps create routing instabilities, which are detrimental to overall network performance. Cisco NSF helps to suppress routing flaps.

Cisco NSF allows for the continued forwarding of data packets along known routes while the routing protocol information is being restored following a switchover. With Cisco NSF, peer Cisco NSF devices do not experience routing flaps because the interfaces remain up during a switchover and adjacencies are not reset. Data traffic is forwarded while the standby RP assumes control from the failed active RP during a switchover. User sessions established before the switchover are maintained.

The ability of the intelligent line cards to remain up through a switchover and to be kept current with the FIB on the active RP is crucial to Cisco NSF operation. While the control plane builds a new routing protocol database and restarts peering agreements, the data plane relies on pre-switchover forwarding-table synchronization to continue forwarding traffic. After the routing protocols have converged, Cisco Express Forwarding updates the FIB table and removes stale route entries, and then it updates the line cards with the refreshed FIB information.

**Note:** Transient routing loops or black holes may be introduced if the network topology changes before the FIB is updated.

The switchover must be completed before the Cisco NSF dead and hold timers expire; otherwise, the peers will reset the adjacency and reroute the traffic.

Cisco NSF protocol enhancements enable a Cisco NSF capable router to signal neighboring Cisco NSF-aware devices during switchover.

**Note:** A device is said to be Cisco NSF aware if it runs Cisco NSF-compatible software. A device is said to be Cisco NSF capable if it has been configured to support Cisco NSF. A Cisco NSF-capable device rebuilds routing information from Cisco NSF-aware or Cisco NSF-capable neighbors.

A Cisco NSF-aware neighbor is needed so that Cisco NSF-capable systems can rebuild their databases and maintain their neighbor adjacencies across a switchover.

Following a switchover, the Cisco NSF-capable device requests that the Cisco NSF-aware neighbor devices send state information to help rebuild the routing tables as a Cisco NSF reset.

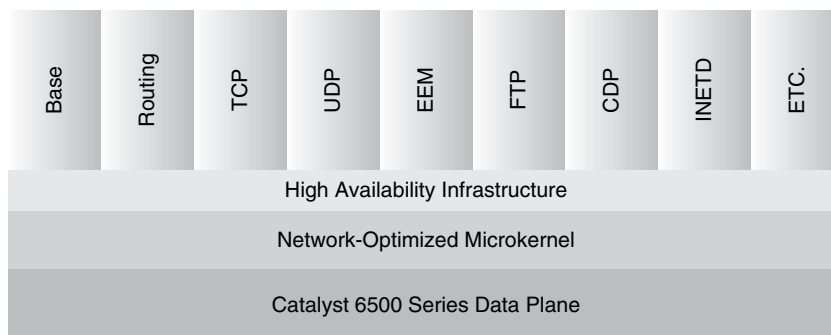
The Cisco NSF protocol enhancements allow a Cisco NSF-capable router to signal neighboring Cisco NSF-aware devices. The signal asks that the neighbor relationship not be reset. As the Cisco NSF-capable router receives and communicates with other routers on the network, it can begin to rebuild its neighbor list. After neighbor relationships are reestablished, the Cisco NSF-capable router begins to resynchronize its database with all of its Cisco NSF-aware neighbors.

Based on platform and Cisco IOS Software release, Cisco NSF with SSO support is available for many routing protocols:

- EIGRP
- OSPF
- BGP
- IS-IS

### Cisco IOS Software Modularity Architecture

The Cisco Catalyst 6500 series with Cisco IOS Software Modularity supports high availability in the enterprise. Figure 2-8 illustrates the key elements and components of the Cisco Software Modularity Architecture.



**Figure 2-8** Cisco IOS Software Modularity Architecture

When Cisco IOS Software patches are needed on systems without Cisco IOS Software Modularity, the new image must be loaded on the active and redundant supervisors, and the supervisor must be reloaded or the switchover to the standby completed to load the patch.

The control plane functions (that manage routing protocol updates and management traffic) on the Catalyst 6500 series run on dedicated CPUs on the multilayer switch forwarding card complex (MSFC). A completely separate data plane is responsible for traffic forwarding. When the hardware is programmed for nonstop operation, the data plane continues forwarding traffic even if there is a disruption in the control plane. The Catalyst 6500 series switches benefit from the more resilient control plane offered by Cisco IOS Software Modularity.

**Note:** Catalyst switch forwarding fabrics are broken down into three planes or functional areas, as follows:

- **Control plane:** The control plane is a logical interface that connects physical chassis components and software functions into a unified logical unit. The control plane connects the system controller functionality on the RP to the service processor (SP) module used to control each card and module in the chassis.
- **Data plane:** The data plane is where packet forwarding takes place. It is the path that packets take through the routing system from the physical layer interface module (PLIM) to the modular services card (MSC) to the switch fabric. On the 6500 series platforms, this would include the policy feature card (PFC) used for high-performance packet processing, and the distributed forwarding card (DFC), which provides local packet forwarding on select line cards.
- **Management plane:** The management plane is where control/configuration of the platform takes place.

The Cisco Catalyst 6500 series with Cisco IOS Software Modularity enables several Cisco IOS control plane subsystems to run in independent processes. Cisco IOS Software Modularity boosts operational efficiency and minimizes downtime:

- It minimizes unplanned downtime through fault containment and stateful process restarts, raising the availability of converged applications.
- It simplifies software changes through subsystem in-service software upgrades (ISSU), significantly reducing code certification and deployment times and decreasing business risks.
- It enables process-level, automated policy control by integrating Cisco IOS Embedded Event Manager (EEM), offloading time-consuming tasks to the network and accelerating the resolution of network issues. EEM is a combination of processes designed to monitor key system parameters such as CPU utilization, interface counters, Simple Network Management Protocol (SNMP), and syslog events. It acts on specific events or threshold counters that are exceeded.

**Note:** Embedded Event Manager is discussed in more detail in Chapter 12, “Network Management Capabilities with Cisco IOS Software.”

### Example: Software Modularity Benefits

Cisco IOS Software Modularity on the Cisco Catalyst 6500 series provides these benefits:

- **Operational consistency:** Cisco IOS Software Modularity does not change the operational point of view. Command-line interfaces (CLI) and management interfaces such as SNMP or syslog are the same as before. New commands to EXEC and configuration mode and new **show** commands have been added to support the new functionality.
- **Protected memory:** Cisco IOS Software Modularity enables a memory architecture where processes make use of a protected address space. Each process and its associated subsystems live in an individual memory space. Using this model, memory corruption across process boundaries becomes nearly impossible.
- **Fault containment:** The benefit of protected memory space is increased availability because problems occurring in one process cannot affect other parts of the system. For example, if a less-critical system process fails or is not operating as expected, critical functions required to maintain packet forwarding are not affected.
- **Process restartability:** Building on the protected memory space and fault containment, the modular processes are now individually restartable. For test purposes or nonresponding processes, the **process restart *process-name*** command is provided to manually restart processes. Restarting a process allows fast recovery from transient errors without the need to disrupt forwarding. Integrated high-availability infrastructure constantly checks the state of processes and keeps track of how many times a process restarted in a defined time interval. If a process restart does not restore the system, the high-availability infrastructure will take more drastic actions, such as initiating a supervisor engine switchover or a system restart.

**Note:** Although a process restart can be initiated by the user, it should be done with caution.

- **Modularized processes:** Several control plane functions have been modularized to cover the most commonly used features. Examples of modular processes include but are not limited to these:
  - Routing process
  - Internet daemon
  - Raw IP processing
  - TCP process
  - User Datagram Protocol (UDP) process
  - Cisco Discovery Protocol process
  - Syslog daemon
  - Any EEM components

- File systems
- Media drivers
- Install manager
- **Subsystem ISSU:** Cisco IOS Software Modularity allows selective system maintenance during runtime through individual patches. By providing versioning and patch-management capabilities, Cisco IOS Software Modularity allows patches to be downloaded, verified, installed, and activated without the need to restart the system. Because data plane packet forwarding is not affected during the patch process, the network operator now has the flexibility to introduce software changes at any time through ISSU. A patch affects only the software components associated with the update.

---

## Designing an Optimum Design for Layer 2

---

Layer 2 architectures rely on the following technologies to create a highly available, deterministic topology: Spanning Tree Protocol (STP), trunking (ISL/802.1q), Unidirectional Link Detection (UDLD), and EtherChannel.

The following section reviews design models and recommended practices for Layer 2 high availability and optimum convergence of the Cisco Enterprise Campus Infrastructure.

## Recommended Practices for Spanning-Tree Configuration

For the most deterministic and highly available network topology, the requirement to support STP convergence should be avoided by design. You may need to implement STP for several reasons:

- When a VLAN spans access layer switches to support business applications.
- To protect against user-side loops. Even if the recommended design does not depend on STP to resolve link or node failure events, STP is required to protect against user-side loops. There are many ways that a loop can be introduced on the user-facing access layer ports. Wiring mistakes, misconfigured end stations, or malicious users can create a loop. STP is required to ensure a loop-free topology and to protect the rest of the network from problems created in the access layer.
- To support data center applications on a server farm.

**Note:** Some security personnel have recommended disabling STP at the network edge. This practice is not recommended because the risk of lost connectivity without STP is far greater than any STP information that might be revealed.

If you need to implement STP, use Rapid Per-VLAN Spanning-Tree Plus (RPVST+). You should also take advantage of the Cisco enhancements to STP known as the Cisco STP toolkit.

## Cisco STP Toolkit

The Cisco enhancements to STP include the following. (Note that the enhancements marked with an \* are also supported with Rapid Per-VLAN Spanning-Tree Plus [RPVST+].)

- **PortFast\*:** Causes a Layer 2 LAN interface configured as an access port to enter the forwarding state immediately, bypassing the listening and learning states. Use PortFast only when connecting a single end station to a Layer 2 access port.
- **UplinkFast:** Provides from three to five seconds convergence after a direct link failure and achieves load balancing between redundant Layer 2 links using uplink groups.
- **BackboneFast:** Cuts convergence time by `max_age` for indirect failure. BackboneFast is initiated when a root port or blocked port on a network device receives inferior bridge protocol data units (BPDU) from its designated bridge.
- **Loop guard\*:** Prevents an alternate or root port from becoming designated in the absence of BPDUs. Loop guard helps prevent bridging loops that could occur because of a unidirectional link failure on a point-to-point link.
- **Root guard\*:** Secures the root on a specific switch by preventing external switches from becoming roots.
- **BPDU guard\*:** When configured on a PortFast-enabled port, BPDU guard shuts down the port that receives a BPDU.
- **Unidirectional Link Detection (UDLD):** UDLD monitors the physical configuration of fiber-optic and copper connections and detects when a one-way connection exists. When a unidirectional link is detected, the interface is shut down and the system alerted.

**Note:** The STP toolkit also supports the BPDU filter option, which prevents PortFast-enabled ports from sending or receiving BPDUs. This feature effectively disables STP at the edge and can lead to STP loops. It is not recommended.

## STP Standards and Features

STP enables the network to deterministically block interfaces and provide a loop-free topology in a network with redundant links. There are several varieties of STP:

- STP is the original IEEE 802.1D version (802.1D-1998) that provides a loop-free topology in a network with redundant links.
- Common Spanning Tree (CST) assumes one spanning-tree instance for the entire bridged network, regardless of the number of VLANs.
- Per VLAN Spanning Tree Plus (PVST+) is a Cisco enhancement of STP that provides a separate 802.1D spanning tree instance for each VLAN configured in the network.

The separate instance supports PortFast, UplinkFast, BackboneFast, BPDU guard, BPDU filter, root guard, and loop guard.

- The 802.1D-2004 version is an updated version of the STP standard.
- Multiple Spanning Tree (MST) is an IEEE standard inspired from the earlier Cisco proprietary Multi-Instance Spanning Tree Protocol (MISTP) implementation. MST maps multiple VLANs into the same spanning-tree instance. The Cisco implementation of MSTP is MST, which provides up to 16 instances of Rapid Spanning Tree Protocol (RSTP, 802.1w) and combines many VLANs with the same physical and logical topology into a common RSTP instance. Each instance supports PortFast, UplinkFast, BackboneFast, BPDU guard, BPDU filter, root guard, and loop guard.
- RSTP, or IEEE 802.1w, is an evolution of STP that provides faster convergence of STP.
- Rapid PVST+ (RPVST+) is a Cisco enhancement of RSTP that uses PVST+. It provides a separate instance of 802.1w per VLAN. The separate instance supports PortFast, UplinkFast, BackboneFast, BPDU guard, BPDU filter, root guard, and loop guard.

**Note:** When Cisco documentation and this course refer to implementing RSTP, they are referring to the Cisco RSTP implementation, or PVRST+.

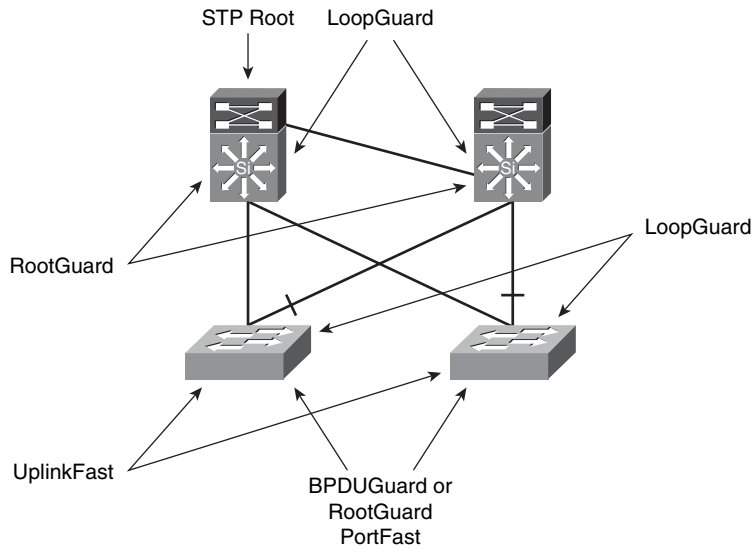
## STP Standards and Features

To configure a VLAN instance to become the root bridge, enter the **spanning-tree vlan *vlan\_id* root primary** command to modify the bridge priority from the default value (32768) to a significantly lower value. The bridge priority for the specified VLAN is set to 8192 if this value will cause the switch to become the root for the VLAN. If any bridge for the VLAN has a priority lower than 8192, the switch sets the priority to one less than the lowest bridge priority. Manually placing the primary and secondary bridges along with enabling STP toolkit options enables you to support a deterministic configuration where you know which ports should be forwarding and which ports should be blocking.

**Note:** Defining the root bridge under MST is done using the **spanning-tree mst *instance\_id* root primary**. When you use this command, the switch will review all bridge ID values it receives from other root bridges. If any root bridge has a bridge ID equal to or less than 24576, it will set its own bridge priority to 4096 less than the lowest bridge priority. To ensure that it will retain its position as the root bridge, you must also enable root guard.

Figure 2-9 illustrates recommended placements for STP toolkit features:

- Loop guard is implemented on the Layer 2 ports between distribution switches, and on the uplink ports from the access switches to the distribution switches.
- Root guard is configured on the distribution switch ports facing the access switches.
- UplinkFast is implemented on the uplink ports from the access switches to the distribution switches.



**Figure 2-9** *Layer 2 Hardening*

**Note:** When you are configuring MST, UplinkFast is not required as a feature on dual-homed switches. Rapid root port failover occurs as part of the default MST protocol implementation.

- BPDUGuard or root guard is configured on ports from the access switches to the end devices, as is PortFast.
- The UDLD protocol allows devices to monitor the physical configuration of the cables and detect when a unidirectional link exists. When a unidirectional link is detected, UDLD shuts down the affected LAN port. UDLD is often configured on ports linking switches.
- Depending on the security requirements of an organization, the port security feature can be used to restrict a port's ingress traffic by limiting the MAC addresses allowed to send traffic into the port.

## Recommended Practices for Trunk Configuration

A trunk is a point-to-point link between two networking devices that carry the traffic of multiple VLANs. Trunks are typically deployed on the interconnection between the access and distribution layers.

The current recommended practice is to use IEEE 802.1Q trunks. Cisco extensions to 802.1Q avoid security concerns related to the 802.1Q nontagged, native VLAN. The native VLAN is assigned to an unused ID, or the tagged, native VLAN option is used to avoid VLAN hopping.



**Note:** *VLAN hopping* is an attack using a double 802.1Q-encapsulated packet. If the attacker has specific knowledge of the 802.1Q native VLAN, a packet could be crafted that when processed, the first or outermost tag is removed when the packet is switched onto the untagged, native VLAN. When the packet reaches the target switch, the inner or second tag is then processed, and the potentially malicious packet is switched to the target VLAN. The traffic in this attack scenario is in a single direction, and no return traffic can be switched by this mechanism. In addition, this attack cannot work unless the attacker knows the native VLAN identity.

VLAN Trunking Protocol (VTP) is a protocol that enables network managers to centrally manage the VLAN database. VTP transparent mode is now a recommended practice because it decreases the potential for operational error.

By default, Cisco switches are configured as a VTP server with no VTP domain name specified.

Therefore, it is also recommended when configuring switches along with setting the mode to Transparent, to set the VTP domain name. This is important if you are connecting your switch to other domains, such as a service provider switch. Misconfiguration of the switch as a server or client with no VTP domain name will cause it to accept the domain name of an adjacent VTP server and overwrite the local VLAN database.

As a recommended practice, when configuring switch-to-switch interconnections to carry multiple VLANs, set Dynamic Trunking Protocol (DTP) to Desirable and Desirable with Encapsulation Negotiate to support DTP negotiation.

**Note:** Turning DTP to On and On with Nonnegotiate could save seconds of outage when restoring a failed link or node. However, with this configuration, DTP is not actively monitoring the state of the trunk, and a misconfigured trunk is not easily identified. One instance where you would use On with Nonnegotiate is if you are trunking between two different VTP domains. DTP includes the VTP domain name in its messages; and if the names do not match, the trunk will not come up if set to Desirable.

Another recommended practice is to manually prune unused VLANs from trunked interfaces to avoid broadcast propagation. You should avoid automatic VLAN pruning.

The final recommendation for trunk configuration is to disable trunks on host ports, because host devices will not need to negotiate trunk status. This practice speeds up Port-Fast and is also a VLAN-hopping security measure.

## VLAN Trunking Protocol

VTP version 3 supports centralized VLAN administration in a switched network. VTP runs only on trunks and provides the following four modes:

- **Server:** Updates clients and servers. The VTP server switch propagates the VTP database to VTP client switches.
- **Client:** Receives updates but cannot make changes.

- **Transparent:** Does not participate in the VTP domain. Lets updates pass through.
- **Off:** Ignores VTP updates.

With VTP, when you configure a new VLAN on a switch in VTP server mode, the VLAN is distributed through all switches in the VTP domain. This redistribution reduces the need to configure the same VLAN everywhere.

With hierarchical networks that do not span VLANs across the distribution layer, there is little need for a shared common VLAN database. In the recommended campus design, the same VLAN should not appear in two access layer switches. Adding and removing VLANs is generally not a frequent network management practice. In most cases, VLANs are defined once during switch setup with few, if any, additional modifications to the VLAN database in an access layer switch. The benefits of dynamic propagation of VLAN information across the network are not worth the potential for unexpected behavior due to operational error. For these reasons, VTP transparent mode is the recommended configuration option.

## Dynamic Trunking Protocol

DTP provides switch ports to negotiate the trunking method with another device and to automatically allow a link to become a trunk.

With Cisco devices, there are five Layer 2 port modes:

- **Trunk:** Puts the port into permanent trunking mode and negotiates to convert the link into a trunk link. The port becomes a trunk port even if the neighboring port does not agree to the change.
- **Desirable:** Actively attempts to form a trunk, subject to neighbor agreement. The port becomes a trunk port if the neighboring port is set to On, Desirable, or Auto mode.
- **Auto:** Makes the port willing to convert the link to a trunk link. The port becomes a trunk port if the neighboring port is set to on or desirable mode.
- **Access:** This is the access mode in Cisco IOS Software that specifies that the port never become a trunk, even if the neighbor tries. This mode puts the LAN port into permanent nontrunking mode and negotiates to convert the link into a nontrunking link.
- **Nonnegotiate:** Prevents the port from generating DTP frames. You must configure the neighboring port manually as a trunk port to establish a trunk link.

With Cisco devices, there are three Ethernet trunk encapsulation types:

- **ISL:** Uses Inter-Switch Link (ISL) encapsulation on the trunk link
- **Dot1q:** Uses 802.1Q encapsulation on the trunk link
- **Negotiate:** Specifies that the LAN port negotiate with the neighboring LAN port to become an ISL (preferred) or 802.1Q trunk, depending on the configuration and capabilities of the neighboring LAN port

The trunking mode, the trunk encapsulation type, and the hardware capabilities of the two connected LAN ports determine whether a link becomes an ISL or 802.1Q trunk.

A common practice is to configure both ends of the trunk to desirable. This has the operational benefit of providing a clear indication of a functional trunking connection with **show** commands, and is the general recommendation for DTP trunking.

An alternative practice is to set one side of the link (typically the access layer) to Auto and the other end (typically the distribution layer) to Desirable. This setting allows for automatic trunk formation, with DTP running on the interconnection to protect against some rare hardware failure scenarios and software misconfigurations.

For fastest convergence, a third configuration turns DTP to On and On with Nonnegotiate to save a few seconds of outage when restoring a failed link or node. However, DTP is not actively monitoring the state of the trunk with this configuration, and a misconfigured trunk is not easily identified. The Nonnegotiate setting can also cause loss of connectivity if the process is not performed in the correct order and there is no out-of-band connectivity to the farthest switch from where the in-band modifications are being made.

## Recommended Practices for UDLD Configuration

UDLD enables devices to monitor the physical configuration of the cables and detect when a unidirectional link exists where bidirectional communication has not been established.

UDLD is typically deployed on fiber topologies where physical misconnections can occur that enable a link to appear to be up/up when there is a mismatched set of transmit/receive pairs. UDLD supports both fiber-optic and copper Ethernet cables connected to LAN ports.

Each switch port configured for UDLD will send UDLD protocol hello packets at Layer 2 containing the device and port identifications of the port, and the device and port identifications of the neighbor as seen by UDLD on that port. Neighboring ports should see their own device and port identifications in the packets received from the other side. If the port does not see its own device and port identifications in the incoming UDLD packets for a specific duration of time, the link is considered unidirectional and is shut down. The default 15-second hello timers are the same for normal and aggressive UDLD. In normal mode, UDLD will error-disable only the end where the UDLD is detected; aggressive mode will error-disable both ends of a connection after aging on a previously bidirectional link in eight seconds.

A recommended practice is to enable UDLD aggressive mode in all environments where fiber-optic interconnections are used. UDLD is enabled globally on all fiber-optic LAN ports with the Cisco IOS software **udld {enable | aggressive}** command. UDLD is enabled on individual LAN ports with the **udld port [aggressive]** interface command.

**Note:** You should enable UDLD in global mode so that you do not have to enable it on every individual fiber-optic interface.

## Recommended Practices for EtherChannel

An EtherChannel bundles individual Ethernet links into a single logical link that provides the aggregate bandwidth of up to eight physical links.

EtherChannels are typically deployed between the distribution-to-core and core-to-core interconnections where increased availability and scaled bandwidth are required. EtherChannel link aggregation is used to provide link redundancy and prevent a single point of failure, and to reduce peering complexity because the single logical entity reduces the number of Layer 3 neighbor relationships as compared to multiple parallel links.

**Note:** EtherChannel also provides an optimization to STP by enabling all member ports to be placed in the forwarding mode. STP views the EtherChannel as a single logical link.

EtherChannels create channels containing up to eight parallel links between switches. If the channels are on interfaces that are on different physical line cards, there is increased availability because the failure of a single line card does not cause a complete loss of connectivity.

There are two variants for the control mechanism for EtherChannel: the prestandard Cisco implementation that uses Port Aggregation Protocol (PAgP), and the IEEE 802.3ad standards-based implementation that uses Link Aggregation Control Protocol (LACP). PAgP and LACP do not interoperate. You can manually configure a switch with PAgP on one side and LACP on the other side in the on/on mode. When this is done, ports configured in the on mode do not negotiate, and therefore there is no negotiation traffic between the ports. This configuration results in the EtherChannel being hard-coded.

When connecting a Cisco IOS Software device to a Catalyst operating system device, make sure that the PAgP settings used for establishing EtherChannels are coordinated. The defaults are different for a Cisco IOS Software device and a Catalyst operating system device. As a recommended practice, Catalyst operating system devices should have PAgP set to Off when connecting to a Cisco IOS Software device if EtherChannels are not configured. If EtherChannel PAgP is used, set both sides of the interconnection to Desirable.

Port aggregation should be disabled when not needed. Port aggregation can most effectively be controlled by disabling it on interfaces facing end users with the set port host macro on the Catalyst operating system or the switchport host macro on Cisco IOS Software. These macros disable both trunking and EtherChannel while enabling STP PortFast.

### Port Aggregation Protocol

PAgP is one of the control mechanisms for EtherChannel. PAgP has four modes related to the automatic formation of bundled, redundant switch-to-switch interconnections (see Figure 2-12):

- **On:** Mode that forces the LAN port to channel unconditionally. In the on mode, a usable EtherChannel exists only when a LAN port group in the on mode is connected to another LAN port group in the on mode. Because ports configured in the on mode

do not negotiate, no negotiation occurs traffic between the ports. You cannot configure the on mode with an EtherChannel protocol. If one end uses the on mode, the other end must also.

- **Desirable:** Places a port into an active negotiating state, in which the port starts negotiations with other ports by sending PAgP packets. This mode is not supported when the EtherChannel members are from different switches in the switch stack (cross-stack EtherChannel).
- **Auto:** Places a port into a passive negotiating state, in which the port responds to PAgP packets it receives but does not start PAgP packet negotiation. This setting minimizes the transmission of PAgP packets. This mode is not supported when the EtherChannel members are from different switches in the switch stack (cross-stack EtherChannel).
- **Off:** Do not become a member.

As with DTP, the long-standing practice for EtherChannel/PAgP has been to set one side of the interconnection (typically the access switch) to Auto and the other side (typically the distribution switch) to Desirable, or both sides to Desirable. In these configurations, an EtherChannel is established when configuration is complete, and connectivity to the remote switch is always available, even when the EtherChannel is not completely established.

**Note:** For Layer 2 EtherChannels, a desirable/desirable PAgP configuration is recommended so that PAgP is running across all members of the bundle, ensuring that an individual link failure will not result in a STP failure.

## Link Aggregation Control Protocol

LACP is another control mechanism for EtherChannel. LACP has four modes related to the automatic formation of bundled, redundant switch-to-switch interconnections:

- **On:** Mode that forces the LAN port to channel unconditionally. In the on mode, a usable EtherChannel exists only when a LAN port group in the on mode is connected to another LAN port group in the on mode. Because ports configured in the on mode do not negotiate, no negotiation traffic occurs between the ports. You cannot configure the on mode with an EtherChannel protocol. If one end uses the on mode, the other end must also.
- **Active:** LACP mode that places a port into an active negotiating state, in which the port initiates negotiations with other ports by sending LACP packets.
- **Passive:** LACP mode that places a port into a passive negotiating state, in which the port responds to LACP packets it receives but does not initiate LACP negotiation.
- **Off:** Do not become a member.

**Note:** The recommended practice for EtherChannel LACP is to set one side of the interconnection (typically the access switch) to Active and the other side (typically the distribution switch) to Passive, or both sides to Active. In these configurations, a channel is established when configuration is complete, and connectivity to the remote switch is always available, even when the channel is not completely established.

---

## Developing an Optimum Design for Layer 3

---

To achieve high availability and fast convergence in the Cisco enterprise campus network, the designer needs to manage multiple objectives, including the following:

- Managing oversubscription and bandwidth
- Supporting link load balancing
- Routing protocol design
- FHRPs

This section reviews design models and recommended practices for high availability and fast convergence in Layer 3 of the Cisco enterprise campus network.

### Managing Oversubscription and Bandwidth


Typical campus networks are designed with oversubscription, as illustrated in Figure 2-10. The rule-of-thumb recommendation for data oversubscription is 20:1 for access ports on the access-to-distribution uplink. The recommendation is 4:1 for the distribution-to-core links. When you use these oversubscription ratios, congestion may occur infrequently on the uplinks. QoS is needed for these occasions. If congestion is frequently occurring, the design does not have sufficient uplink bandwidth.

As access layer bandwidth capacity increases to 1 Gb/s, multiples of 1 Gb/s, and even 10 Gb/s, the bandwidth aggregation on the distribution-to-core uplinks might be supported on many Gigabit Ethernet EtherChannels, on 10 Gigabit Ethernet links, and on 10 Gigabit EtherChannels.

#### Bandwidth Management with EtherChannel

As bandwidth from the distribution layer to the core increases, oversubscription to the access layer must be managed, and some design decisions must be made.

Just adding more uplinks between the distribution and core layers leads to more peer relationships, with an increase in associated overhead.



**Figure 2-10** *Managing Oversubscription and Bandwidth*

EtherChannels can reduce the number of peers by creating single logical interface. However, you must consider some issues about how routing protocols will react to single link failure:

- OSPF running on a Cisco IOS Software-based switch will notice a failed link, and will increase the link cost. Traffic is rerouted, and this design leads to a convergence event.
- OSPF running on a Cisco Hybrid-based switch will not change link cost. Because it will continue to use the EtherChannel, this may lead to an overload in the remaining links in the bundle as OSPF continues to divide traffic equally across channels with different bandwidths.
- EIGRP might not change link cost, because the protocol looks at the end-to-end cost. This design might also overload remaining links.

The EtherChannel Min-Links feature is supported on LACP EtherChannels. This feature allows you to configure the minimum number of member ports that must be in the link-up state and bundled in the EtherChannel for the port channel interface to transition to the link-up state. You can use the EtherChannel Min-Links feature to prevent low-bandwidth LACP EtherChannels from becoming active.

## Bandwidth Management with 10 Gigabit Interfaces

Upgrading the uplinks between the distribution and core layers to 10 Gigabit Ethernet links is an alternative design for managing bandwidth. The 10 Gigabit Ethernet links can also support the increased bandwidth requirements.

This is a recommended design:

- Unlike the multiple link solution, 10 Gigabit Ethernet links do not increase routing complexity. The number of routing peers is not increased.
- Unlike the EtherChannel solution, the routing protocols will have the ability to deterministically select the best path between the distribution and core layer.

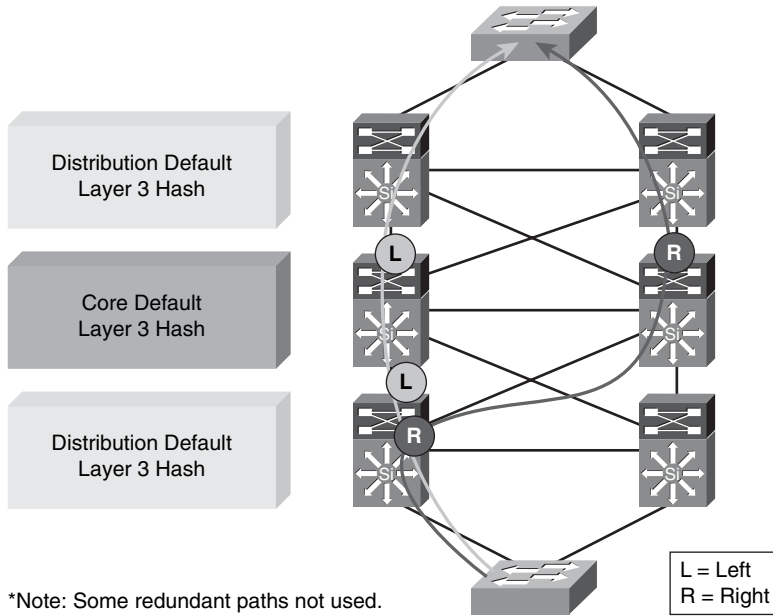
## Link Load Balancing

In Figure 2-11, many equal-cost, redundant paths are provided in the recommended network topology from one access switch to the other across the distribution and core switches. From the perspective of the access layer, there are at least three sets of equal-cost, redundant links to cross to reach another building block, such as the data center.

Cisco Express Forwarding (CEF) is a deterministic algorithm. As shown in the figure, when packets traverse the network that all use the same input value to the CEF hash, a “go to the right” or “go to the left” decision is made for each redundant path. When this results in some redundant links that are ignored or underutilized, the network is said to be experiencing CEF polarization.

To avoid CEF polarization, you can tune the input into the CEF algorithm across the layers in the network. The default input hash value is Layer 3 for source and destination. If you change this input value to Layer 3 plus Layer 4, the output hash value also changes.





**Figure 2-11** CEF Load Balancing (Default Behavior)

As a recommendation, use alternating hashes in the core and distribution layer switches:

- In the core layer, continue to use the default, which is based on only Layer 3 information.
- In the distribution layer, use the Layer 3 plus Layer 4 information as input into the CEF hashing algorithm with the command `Dist2-6500 (config)#mls ip cef load-sharing full`.

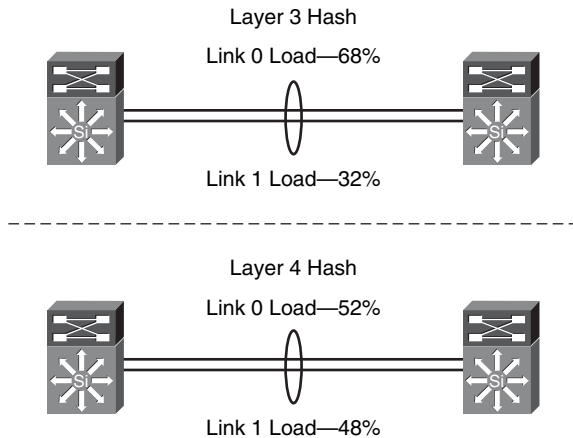
This alternating approach helps eliminate the always-right or always-left biased decisions and helps balance the traffic over equal-cost, redundant links in the network.

### Link Load Balancing

EtherChannel allows load sharing of traffic among the links in the channel and redundancy in the event that one or more links in the channel fail.

You can tune the hashing algorithm used to select the specific EtherChannel link on which a packet is transmitted. You can use the default Layer 3 source and destination information, or you can add a level of load balancing to the process by adding the Layer 4 TCP/IP port information as an input to the algorithm.

Figure 2-12 illustrates some results from experiments at Cisco in a test environment using a typical IP addressing scheme of one subnet per VLAN, two VLANs per access switch, and the RFC 1918 private address space. The default Layer 3 hash algorithm provided about one-third to two-thirds utilization. When the algorithm was changed to include Layer 4 information, nearly full utilization was achieved with the same topology and traffic pattern.



**Figure 2-12** *EtherChannel Load Balancing*

The recommended practice is to use Layer 3 plus Layer 4 load balancing to provide as much information as possible for input to the EtherChannel algorithm to achieve the best or most uniform utilization of EtherChannel members. The command **port-channel load-balance** is used to present the more unique values to the hashing algorithm. This can be achieved using the command **dist1-6500(config)#port-channel load-balance src-dst-port**.

To achieve the best load balancing, use two, four, or eight ports in the port channel.

## Routing Protocol Design

This section reviews design recommendations for routing protocols in the enterprise campus.

Routing protocols are typically deployed across the distribution-to-core and core-to-core interconnections.

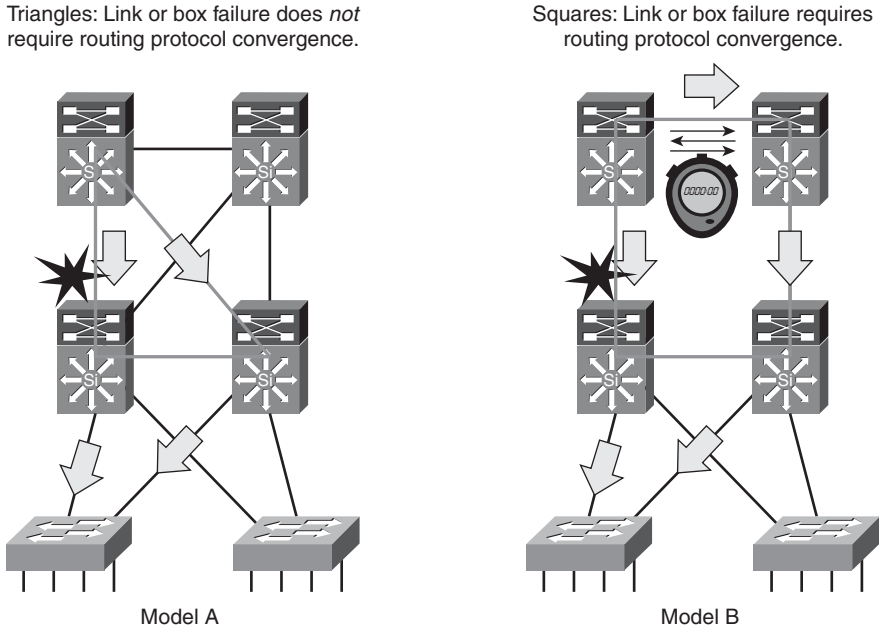
Layer 3 routing design can be used in the access layer, too, but this design is currently not as common.

Layer 3 routing protocols are used to quickly reroute around failed nodes or links while providing load balancing over redundant paths.

### Build Redundant Triangles

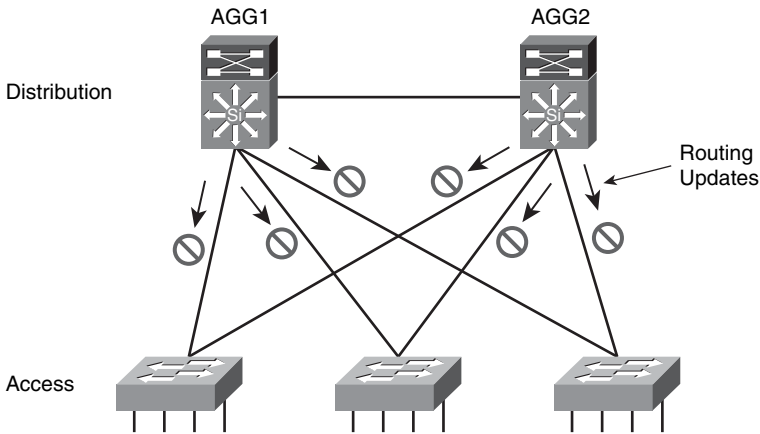
For optimum distribution-to-core layer convergence, build redundant triangles, not squares, to take advantage of equal-cost, redundant paths for the best deterministic convergence.

The topology connecting the distribution and core switches should be built using triangles, with equal-cost paths to all redundant nodes. The triangle design is shown in Figure 2-13 Model A, and uses dual equal-cost paths to avoid timer-based, nondeterministic convergence. Instead of indirect neighbor or route-loss detection using hellos and dead timers, the triangle design failover is hardware based and relies on physical link loss to mark a path as unusable and reroute all traffic to the alternate equal-cost path. There is no need for OSPF or EIGRP to recalculate a new path.



**Figure 2-13** *Build Redundant Triangles*

In contrast, the square topology shown in Figure 2-14 Model B requires routing protocol convergence to fail over to an alternate path in the event of a link or node failure. It is possible to build a topology that does not rely on equal-cost, redundant paths to compensate



**Figure 2-14** *Use Passive Interfaces at the Access Layer*

for limited physical fiber connectivity or to reduce cost. However, with this design, it is not possible to achieve the same deterministic convergence in the event of a link or node failure, and for this reason the design will not be optimized for high availability.

## Peer Only on Transit Links

Another recommended practice is to limit unnecessary peering across the access layer by peering only on transit links.

By default, the distribution layer switches send routing updates and attempt to peer across the uplinks from the access switches to the remote distribution switches on every VLAN. This is unnecessary and wastes CPU processing time.

Figure 2-14 shows an example network where with 4 VLANs per access switch and 3 access switches, 12 unnecessary adjacencies are formed. Only the adjacency on the link between the distribution switches is needed. This redundant Layer 3 peering has no benefit from a high-availability perspective, and only adds load in terms of memory, routing protocol update overhead, and complexity. In addition, in the event of a link failure, it is possible for traffic to transit through a neighboring access layer switch, which is not desirable.

As a recommended practice, limit unnecessary routing peer adjacencies by configuring the ports toward Layer 2 access switches as passive, which will suppress the advertising of routing updates. If a distribution switch does not receive routing updates from a potential peer on a specific interface, it will not need to process these updates, and it will not form a neighbor adjacency with the potential peer across that interface.

There are two approaches to configuring passive interfaces for the access switches:

- Use the **passive-interface default** command, and selectively use the **no passive-interface** command to enable a neighboring relationship where peering is desired.
- Use the **passive-interface** command to selectively make specific interfaces passive.

Passive interface configuration example for OSPF:

```
AGG1(config)#router ospf 1
AGG1(config-router)#passive-interface Vlan 99
! Or
AGG1(config)#router ospf 1
AGG1(config-router)#passive-interface default
AGG1(config-router)#no passive-interface Vlan 99
```

Passive interface configuration example for EIGRP:

```
AGG1(config)#router EIGRP 1
AGG1(config-router)#passive-interface Vlan 99
! Or
AGG1(config)#router EIGRP 1
AGG1(config-router)#passive-interface default
AGG1(config-router)#no passive-interface Vlan 99
```

You should use whichever technique requires the fewest lines of configuration or is the easiest for you to manage.

## Summarize at the Distribution Layer

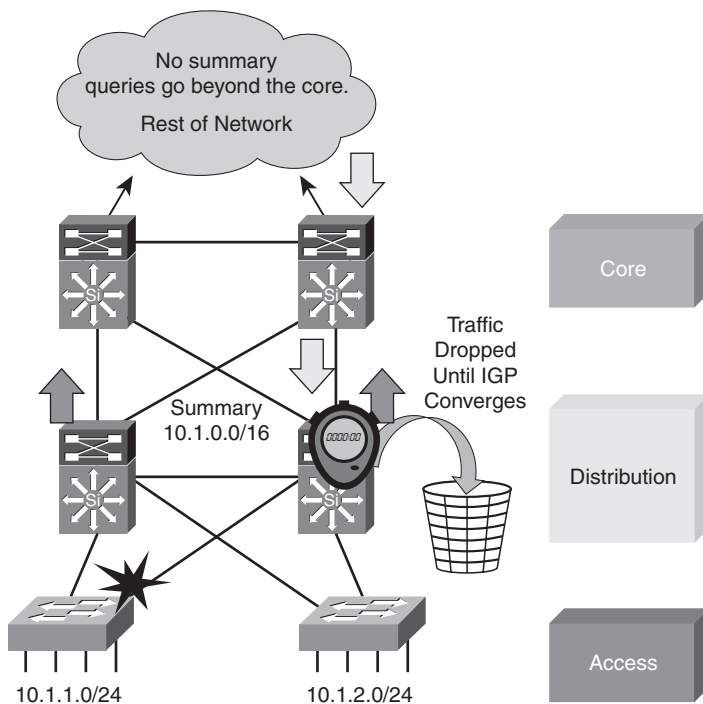
A hierarchical routing design reduces routing update traffic and avoids unnecessary routing computations. Such a hierarchy is achieved through allocating IP networks in contiguous blocks that can be easily summarized by a dynamic routing protocol.

It is a recommended practice to configure route summarization at the distribution layer to advertise a single summary route to represent multiple IP networks within the building (switch block). As a result, fewer routes will be advertised through the core layer and subsequently to the distribution layer switches in other buildings (switch blocks). If the routing information is not summarized toward the core, EIGRP and OSPF require interaction with a potentially large number of peers to converge around a failed node.

Summarization at the distribution layer optimizes the rerouting process. If a link to an access layer device goes down, return traffic at the distribution layer to that device is dropped until the IGP converges. When summarization is used and the distribution nodes send summarized information toward the core, an individual distribution node does not advertise loss of connectivity to a single VLAN or subnet. This means that the core does not know that it cannot send traffic to the distribution switch where the access link has failed. Summaries limit the number of peers that an EIGRP router must query or the number of link-state advertisements (LSA) that OSPF must process, and thereby speeds the rerouting process.

Summarization should be performed at the boundary where the distribution layer of each building connects to the core. The method for configuring route summarization varies, depending on the IGP being used. Route summarization is covered in detail in Chapter 3, “Developing an Optimum Design for Layer 3.” These designs require a Layer 3 link between the distribution switches, as shown in Figure 2-15, to allow the distribution node that loses connectivity to a given VLAN or subnet the ability to reroute traffic across the distribution-to-distribution link. To be effective, the address space selected for the distribution-to-distribution link must be within the address space being summarized.

Summarization relies on a solid network addressing design.

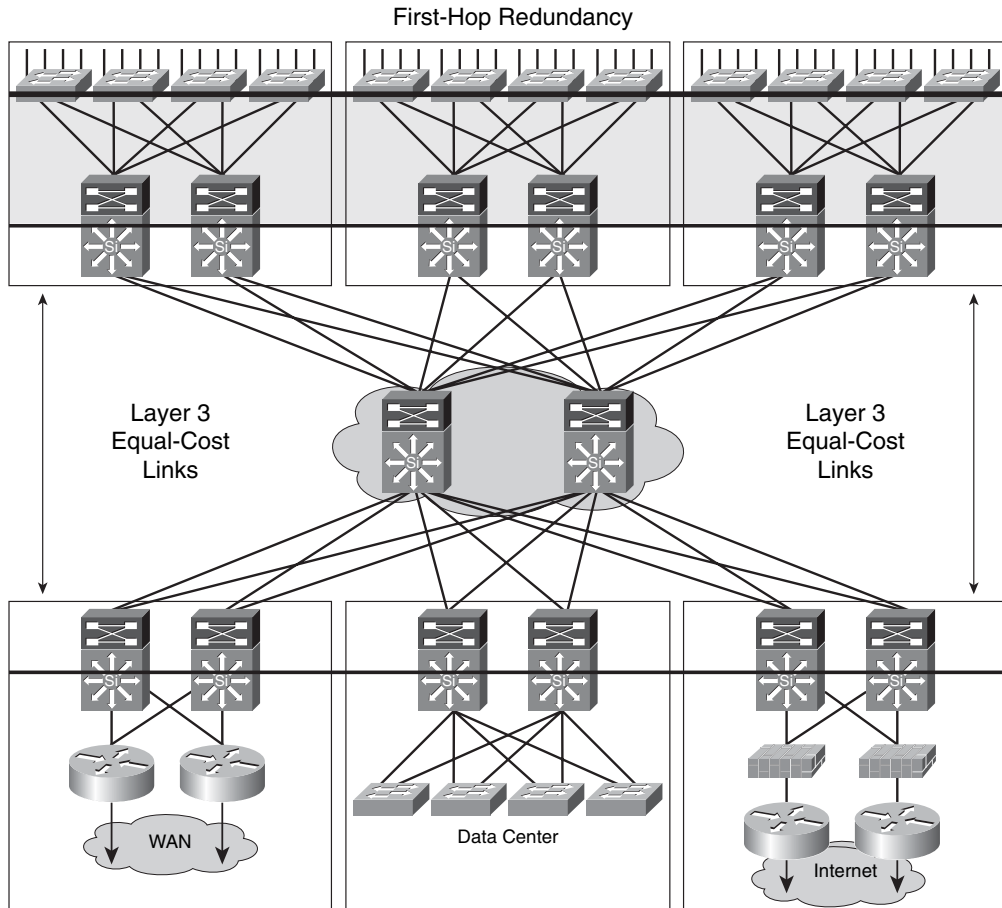


**Figure 2-15** Summarize at the Distribution Layer

## First-Hop Redundancy

First-hop redundancy or default-gateway redundancy is an important component in convergence in a highly available hierarchical network design.

First-hop redundancy allows a network to recover from the failure of the device acting as the default gateway for end nodes on a physical segment. When the access layer is Layer 2, the distribution layer switches act as the default gateway for the entire Layer 2 domain that they support, as illustrated in Figure 2-16.



**Figure 2-16** *First-Hop Redundancy*

A first-hop redundancy protocol is needed only if the design implements Layer 2 between the access switch and the distribution switch. If Layer 3 is supported to the access switch, the default gateway for end devices is at the access level.

In Cisco deployments, HSRP, developed by Cisco, is typically used as the FHRP. VRRP is an Internet Engineering Task Force (IETF) standards-based method of providing default-gateway redundancy. More deployments are starting to use GLBP, which can more

easily achieve load balancing on the uplinks from the access layer to the distribution layer, and first-hop redundancy and failure protection.

HSRP and VRRP with Cisco enhancements both provide a robust method of backing up the default gateway, and can provide subsecond failover to the redundant distribution switch when tuned properly. HSRP is the recommended protocol over VRRP because it is a Cisco-owned standard, which allows for the rapid development of new features and functionality before VRRP. VRRP is the logical choice over HSRP when interoperability with other vendor devices is required.

HSRP or GLBP timers can be reliably tuned to achieve 800-ms convergence for link or node failure in the Layer 2 and Layer 3 boundary in the building distribution layer. The following configuration snippet shows how HSRP can be tuned down from its default 3-second hello timer and 10-second hold timer in a campus environment to achieve subsecond convergence on aggregation switches:

```
interface Vlan5
ip address 10.1.5.3 255.255.255.0
ip helper-address 10.5.10.20
standby 1 ip 10.1.5.1
standby 1 timers msec 200 msec 750
standby 1 priority 150
standby 1 preempt delay minimum 180
```

### Preempt Delay Tuning

One important factor to take into account when tuning default gateway redundancy using HSRP or another protocol is its preemptive behavior.

Preemption causes the primary HSRP peer to re-assume the primary role when it comes back online after a failure or maintenance event. Preemption is the desired behavior because the RSTP root should be the same device as the HSRP primary for a given subnet or VLAN. However, if HSRP and RSTP are not synchronized after failure recovery, the interconnection between the distribution switches can become a transit link, and traffic takes a multihop Layer 2 path to its default gateway.

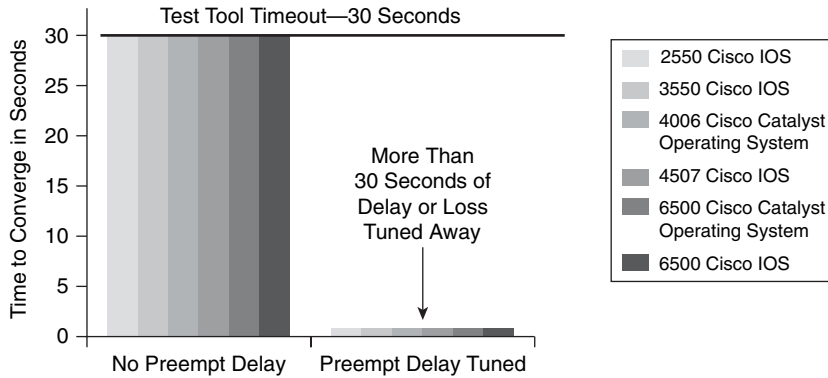
HSRP preemption needs to be aware of switch boot time and connectivity to the rest of the network. Preempt delay must be longer than the switch boot time:

- Layer 1 traffic forwarding on line cards
- Layer 2 STP convergence
- Layer 3 IGP convergence

It is possible for HSRP neighbor relationships to form and preemption to occur before the primary switch has Layer 3 connectivity to the core. If this happens, traffic from the access layer can be dropped until full connectivity is established to the core.

The recommended practice is to measure the system boot time, and set the HSRP preempt delay with the **standby preempt delay minimum** command to 50 percent greater than this value. This ensures that the HSRP primary distribution node has established full connectivity to all parts of the network before HSRP preemption is allowed to occur.

Figure 2-17 demonstrates the positive impact that proper HSRP tuning can have on network convergence.



**Figure 2-17** HSRP Preempt Delay Tuning

## Overview of Gateway Load Balancing Protocol

GLBP is a first-hop redundancy protocol designed by Cisco that allows packet load sharing between groups of redundant routers.

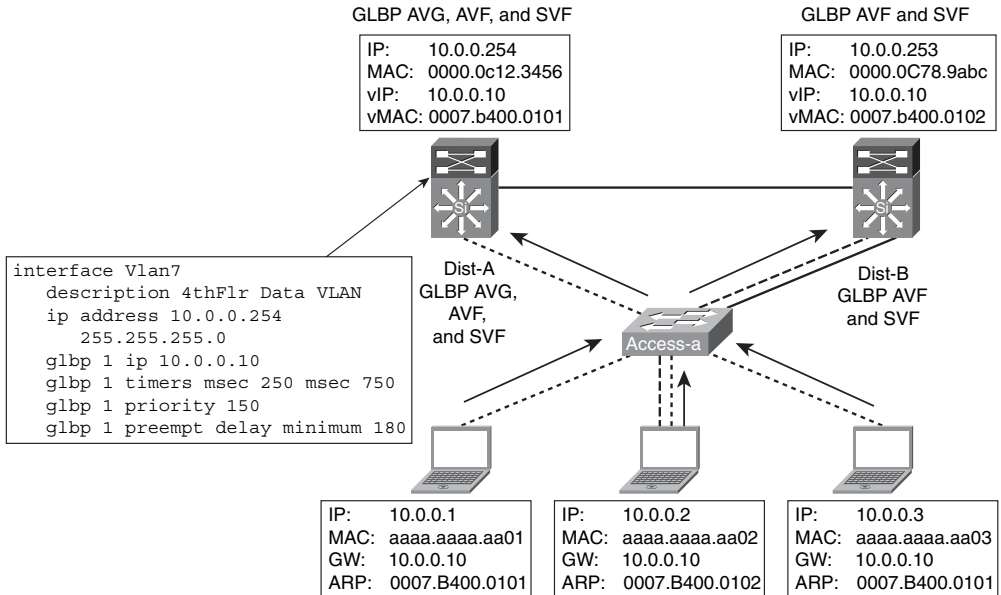
When HSRP or VRRP is used to provide default-gateway redundancy, the backup members of the peer relationship are idle, waiting for a failure event to occur before they take over and actively forward traffic. Methods to use uplinks with HSRP or VRRP are difficult to implement and manage. In one technique, the HSRP and STP or RSTP roots alternate between distribution node peers, with the even VLANs homed on one peer and the odd VLANs homed on the alternate. Another technique uses multiple HSRP groups on a single interface and uses DHCP to alternate between the multiple default gateways. These techniques work but are not optimal from a configuration, maintenance, or management perspective.

GLBP provides all the benefits of HSRP and includes load balancing, too. For HSRP, a single virtual MAC address is given to the endpoints when the endpoints use Address Resolution Protocol (ARP) to learn the physical MAC address of their default gateways. GLBP allows a group of routers to function as one virtual router by sharing one virtual IP address while using multiple virtual MAC addresses for traffic forwarding. Figure 2-18 shows a sample configuration supporting GLBP and its roles.

When an endpoint uses ARP for its default gateway, by default the virtual MACs are provided by the GLBP active virtual gateway (AVG) on a round-robin basis. These gateways that assume responsibility for forwarding packets sent to the virtual MAC address are known as active virtual forwarders (AVF) for their virtual MAC address. Because the traffic from a single common subnet goes through multiple redundant gateways, all the uplinks can be used.



**Note:** GLBP can also be configured to provide virtual MAC addresses based on a weighting value (if for example, two distribution layer devices have different capacities or utilization), or could be host based (where a particular end-station will use the same AVF MAC address every time), which would be required for tracking flows, such as Network Address Translation / Port Address Translation (NAT/PAT).



**Figure 2-18** Gateway Load Balancing Protocol

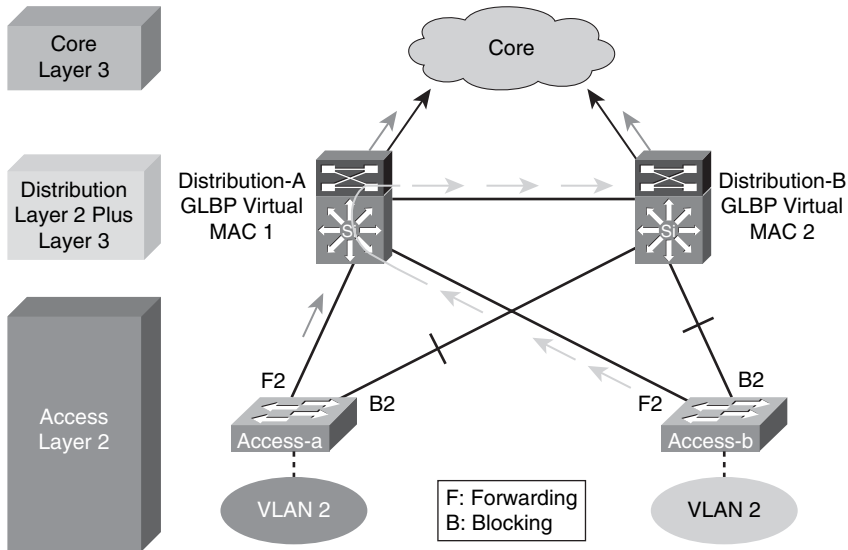
Failover and convergence in GLBP work in a similar fashion as HSRP. A secondary virtual forwarder (SVF) takes over for traffic destined to a virtual MAC impacted by the failure and begins forwarding traffic for its failed peer. The end result is that a more equal utilization of the uplinks is achieved with minimal configuration. As a side effect, a convergence event on the uplink or on the primary distribution node affects only half as many hosts with a pair of GLBP switches, giving a convergence event an average of 50 percent less impact.

Note that using GLBP in topologies where STP has blocked one of the access layer uplinks may result in a two-hop path at Layer 2 for upstream traffic, as illustrated in Figure 2-19.

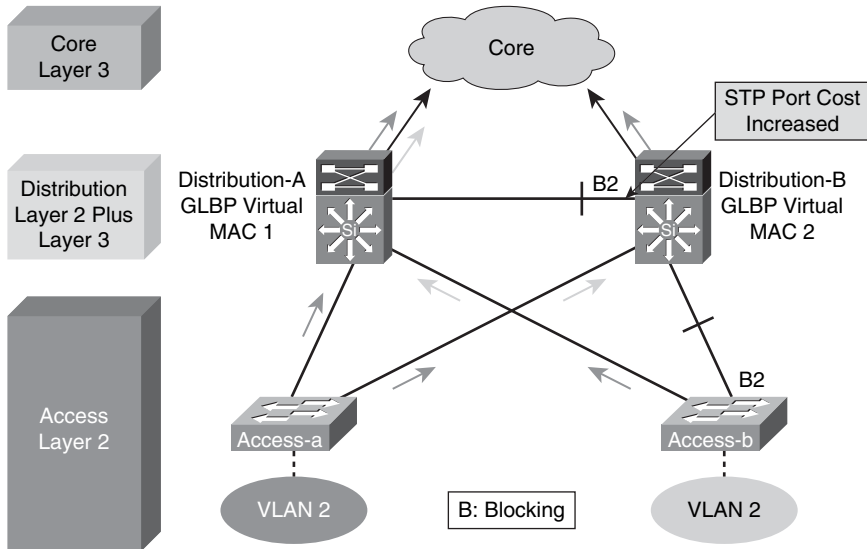
In environments where VLANs span across the distribution switches, HSRP is the preferred FHRP implementation.

In some cases, the STP environment can be tuned so that the Layer 2 link between the distribution switches is the blocking link while the uplinks from the access layer switches are in a forwarding state.

Figure 2-20 illustrates how you can tune STP by using the **spanning-tree cost interface** configuration command to change the port cost on the interface between the distribution layer switches on the STP secondary root switch. This option works if no VLANs span access switches.



**Figure 2-19** *GLBP VLAN Spanning*



**Figure 2-20** *GLBP and STP Tuning*

However, if the same VLAN is on multiple access switches, you will have a looped figure-eight topology where one access layer uplink is still blocking. The preferred design is to not span VLANs across access switches.

## Optimizing FHRP Convergence

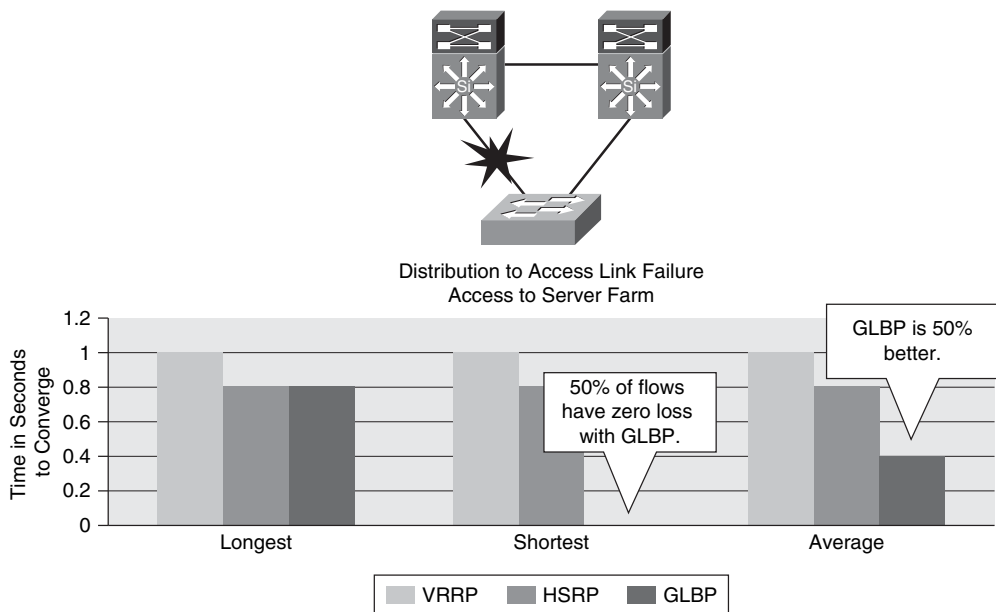
HSRP can be reliably tuned to achieve 800-ms convergence for link or node failure. With HSRP, all flows from one subnet go through the active HSRP router; so the longest, shortest, and average convergence times are the same and less than a second.

VRRP can be tuned with subsecond timers, although the results of this timer tuning is not known. With VRRP, all flows from one subnet go through the same VRRP master router, so the longest, shortest, and average convergence times are the same and about a second.

GLBP can also be reliably tuned to achieve 800-ms convergence for link or node failure. With GLBP, a convergence event on an uplink or on the primary distribution node affects only half as many hosts, so a convergence event has an average of 50 percent less impact than with HSRP or VRRP if the default round-robin load-balancing algorithm is used.

GLBP is currently supported on the Cisco Catalyst 6500 series switches and the Cisco Catalyst 4500 series switches.

Figure 2-21 illustrates the difference in convergence times between each of the respective FHRP when deployed on a distribution to access link in a server farm.



**Figure 2-21** *Optimizing FHRP Convergence*

## Supporting a Layer 2 to Layer 3 Boundary Design

This following section reviews design models and recommended practices for supporting the Layer 2 to Layer 3 boundary in highly available enterprise campus networks.

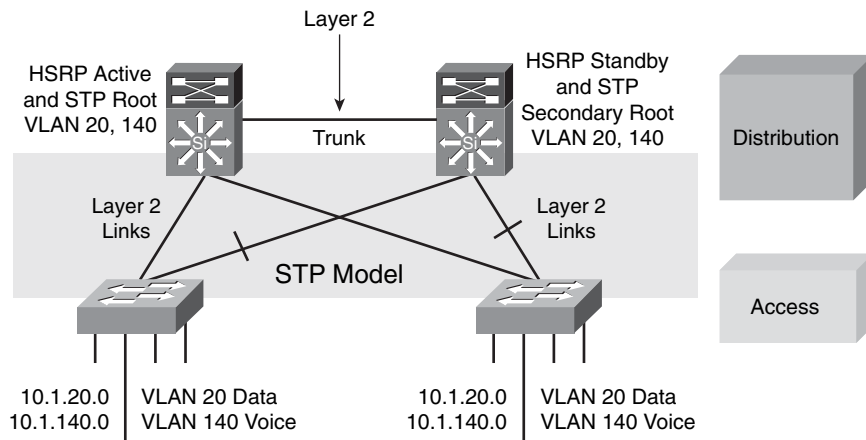
## Layer 2 to Layer 3 Boundary Design Models

There are several design models for placement of the Layer 2 to Layer 3 boundary in the enterprise campus.

### Layer 2 Distribution Switch Interconnection

If the enterprise campus requirements must support VLANs spanning multiple access layer switches, the design model uses a Layer 2 link for interconnecting the distribution switches.

The design, illustrated here in Figure 2-22, is more complex than the Layer 3 interconnection of the distribution switches. The STP convergence process will be initiated for uplink failures and recoveries.



**Figure 2-22** Layer 2 Distribution Switch Interconnection

You can improve this suboptimal design as follows:


- Use RSTP as the version of STP.

**Note:** RPVST+ is a Cisco enhancement of RSTP that uses PVST+. It provides a separate instance of 802.1w per VLAN. The separate instance supports PortFast, UplinkFast, BackboneFast, BPDU guard, BPDU filter, root guard, and loop guard. (RPVST+ is also known as PVRST+.)

- Provide a Layer 2 trunk between the two distribution switches to avoid unexpected traffic paths and multiple convergence events.
- If you choose to load balance VLANs across uplinks, be sure to place the HSRP primary and the STP primary on the same distribution layer switch. The HSRP and RSTP root should be collocated on the same distribution switches to avoid using the inter-distribution link for transit.

### Layer 3 Distribution Switch Interconnection (HSRP)

Figure 2-23 shows the model which supports a Layer 3 interconnection between distribution switches using HSRP as the FHRP.



**Figure 2-23** *Layer 3 Distribution Switch Interconnection*

In this time-proven topology, no VLANs span between access layer switches across the distribution switches. A subnet equals a VLAN, which equals an access switch. The root for each VLAN is aligned with the active HSRP instance. From a STP perspective, both access layer uplinks are forwarding, so the only convergence dependencies are the default gateway and return-path route selection across the distribution-to-distribution link.

This recommended design provides the highest availability.

With this design, a distribution-to-distribution link is required for route summarization. A recommended practice is to map the Layer 2 VLAN number to the Layer 3 subnet for ease of use and management.

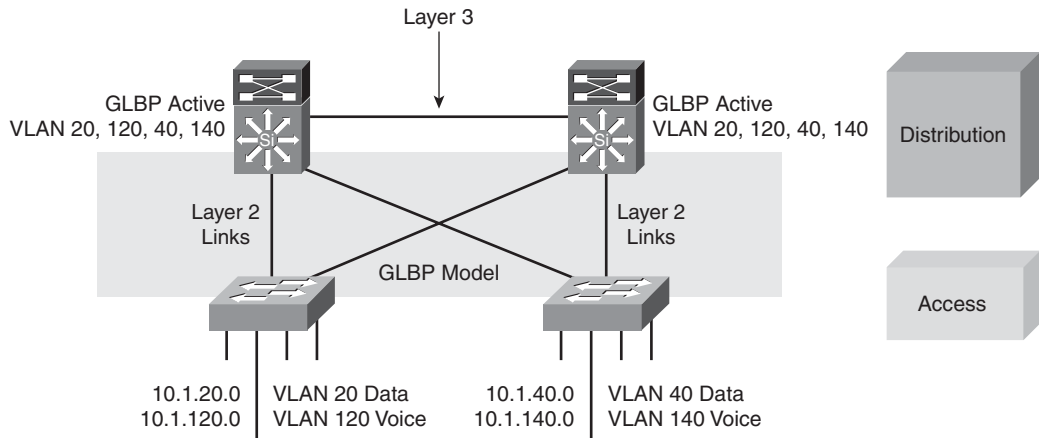
### Layer 3 Distribution Switch Interconnection (GLBP)

GLBP can also be used as the FHRP with the Layer 3 distribution layer interconnection model, as shown in Figure 2-24.

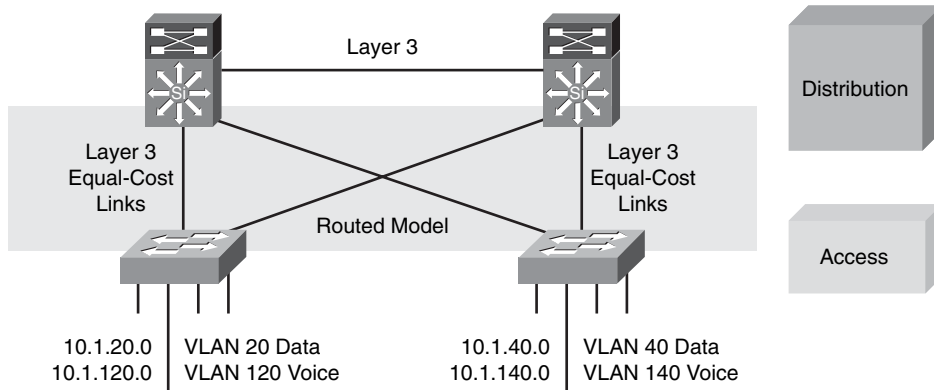
GLBP allows full utilization of the uplinks from the access layer. However, because the distribution of ARP responses is random, it is less deterministic than the design with HSRP. The distribution-to-distribution link is still required for route summarization. Because the VLANs do not span access switches, STP convergence is not required for uplink failure and recovery.

### Layer 3 Access to Distribution Interconnection

The design extending Layer 3 to the access layer, shown here in Figure 2-25, provides the fastest network convergence.



**Figure 2-24** Layer 3 Distribution Switch Interconnection with GLBP



**Figure 2-25** Layer 3 Access to Distribution Interconnection

A routing protocol such as EIGRP, when properly tuned, can achieve better convergence results than designs that rely on STP to resolve convergence events. A routing protocol can even achieve better convergence results than the time-tested design placing the Layer 2 to Layer 3 boundary at the distribution layer. The design is easier to implement than configuring Layer 2 in the distribution layer because you do not need to align STP with HSRP or GLBP.

This design supports equal-cost Layer 3 load balancing on all links between the network switches. No HSRP or GLBP configuration is needed because the access switch is the default gateway for the end users. VLANs cannot span access switches in this design.

The convergence time required to reroute around a failed access-to-distribution layer up-link is reliably under 200 ms as compared to 900 ms for the design placing the Layer 2 and Layer 3 boundary at the distribution layer. Return-path traffic is also in the sub-200 ms of convergence time for an EIGRP reroute, again compared to 900 ms for the traditional Layer 2 to Layer 3 distribution layer model.

Because both EIGRP and OSPF loads share over equal-cost paths, this design provides a convergence benefit similar to GLBP. Approximately 50 percent of the hosts are not affected by a convergence event because their traffic is not flowing over the link or through the failed node.

However, some additional complexity associated with uplink IP addressing and subnetting and the loss of flexibility is associated with this design alternative.

Routing in the access layer is not as widely deployed in the enterprise environment as the Layer 2 and Layer 3 distribution layer boundary model.

**Note:** Deploying a layer 3 access layer may be prohibited because of conformance with the existing architecture, price of multilayer switches, application, or service requirements.

### EIGRP Access Design Recommendations

When EIGRP is used as the routing protocol for a fully routed or routed access layer solution, with tuning it can achieve sub-200 ms convergence.

EIGRP to the distribution layer is similar to EIGRP in the branch, but it's optimized for fast convergence using these design rules:

- Limit scope of queries to a single neighbor:  
Summarize at the distribution layer to the core as is done in the traditional Layer 2 to Layer 3 border at the distribution layer. This confines impact of an individual access link failure to the distribution pair by stopping EIGRP queries from propagating beyond the core of the network. When the distribution layer summarizes toward the core, queries are limited to one hop from the distribution switches, which optimizes EIGRP convergence.  
  
Configure all access switches to use EIGRP stub nodes so that the access devices are not queried by the distribution switches for routes. EIGRP stub nodes cannot act as transit nodes and do not participate in EIGRP query processing. When the distribution node learns through the EIGRP hello packets that it is talking to a stub node, it does not flood queries to that node.
- Control route propagation to access switches using distribution lists. The access switches need only a default route to the distribution switches. An outbound distribution list applied to all interfaces facing the access layer from the distribution switch will conserve memory and optimize performance at the access layer.
- Set hello and dead timers to 1 and 3 as a secondary mechanism to speed up convergence. The link failure or node failure should trigger convergence events. Tune EIGRP hello and dead timers to 1 and 3, respectively, to protect against a soft failure in which the physical links remain active but hello and route processing has stopped.

EIGRP optimized configuration example:

```
interface GigabitEthernet1/1 ip hello-interval eigrp 100 2 ip hold-time
  eigrp 100 6
router eigrp 100 eigrp stub connected
```

**Note:** An EIGRP stub is included in the base image of all Cisco multilayer Catalyst switches.

## OSPF Access Design Recommendations

When OSPF is used as the routing protocol for a fully routed or routed access layer solution with tuning it can also achieve sub-200-ms convergence.

OSPF to the distribution layer is similar to OSPF in the branch, but it's optimized for fast convergence. With OSPF, summarization and limits to the diameter of OSPF LSA propagation is provided through implementation of Layer 2 to Layer 3 boundaries or Area Border Routers (ABR). It follows these design rules:

- Control the number of routes and routers in each area:
  - Configure each distribution block as a separate, totally stubby OSPF area. The distribution switches become ABRs with their core-facing interfaces in area 0, and the access layer interfaces in unique, totally stubby areas for each access layer switch. Do not extend area 0 to the access switch because the access layer is not used as a transit area in a campus environment. Each access layer switch is configured into its own unique, totally stubby area. In this configuration, LSAs are isolated to each access layer switch so that a link flap for one access layer switch is not communicated beyond the distribution pairs.
  - Tune OSPF millisecond hello, dead-interval, SPF, and LSA throttle timers as a secondary mechanism to improve convergence. Because CPU resources are not as scarce in a campus environment as they might be in a WAN environment, and the media types common in the access layer are not susceptible to the same half-up or rapid transitions as are those commonly found in the WAN, OSPF timers can safely be tuned, as shown in the configuration snippet here:

```
interface GigabitEthernet1/1 ip ospf dead-interval minimal
  hello-multiplier 4
router ospf 100 area 120 stub no-summary timers throttle spf 10 100 5000
timers throttle lsa all 10 100 5000 timers lsa arrival 80
```

**Note:** OSPF support is not included in the base image of all Cisco multilayer Catalyst switches, but it is available with the IP Services upgrade.

## Potential Design Issues

The following sections discuss potential design issues for placement of the Layer 2 to Layer 3 boundary in the enterprise campus.


### Daisy Chaining Access Layer Switches

If multiple fixed-configuration switches are daisy chained in the access layer of the network, there is a danger that black holes will occur in the event of a link or node failure.

In the topology in Figure 2-26, before failures no links are blocking from a STP or RSTP perspective, so both uplinks are available to actively forward and receive traffic. Both distribution nodes can forward return-path traffic from the rest of the network toward the access layer for devices attached to all members of the stack or chain.

Two scenarios can occur if a link or node in the middle of the chain or stack fails. In the first case, the standby HSRP peer can go active as it loses connectivity to its primary peer, forwarding traffic outbound for the devices that still have connectivity to it. The primary





**Figure 2-26** *Daisy Chaining Access Layer Switches*

HSRP peer remains active and also forwards outbound traffic for its half of the stack. Although this is not optimum, it is not detrimental from the perspective of outbound traffic.

The second scenario is the issue. Return-path traffic has a 50 percent chance of arriving on a distribution switch that does not have physical connectivity to the half of the stack where the traffic is destined. The traffic that arrives on the wrong distribution switch is dropped.

The solution to this issue with this design is to provide alternate connectivity across the stack in the form of a loop-back cable running from the top to the bottom of the stack. This link needs to be carefully deployed so that the appropriate STP behavior will occur in the access layer.


An alternate design uses a Layer 2 link between the distribution switches.

### Cisco StackWise Technology in the Access Layer

Cisco StackWise technology can eliminate the danger that black holes occur in the access layer in the event of a link or node failure. It can eliminate the need for loop-back cables in the access layer or Layer 2 links between distribution nodes.

StackWise technology, shown in the access layer in Figure 2-27, supports the recommended practice of using a Layer 3 connection between the distribution switches without having to use a loop-back cable or perform extra configuration.

The true stack creation provided by the Cisco Catalyst 3750 series switches makes using stacks in the access layer much less complex than chains or stacks of other models. A stack of 3750 switches appears as one node from the network topology perspective.



**Figure 2-27** *StackWise Technology*

If you use a modular chassis switch to support ports in the aggregation layer, such as the Cisco Catalyst 4500 or Catalyst 6500 family of switches, these design considerations are not required.

### Too Much Redundancy


Be aware that even if some redundancy is good, more redundancy is not necessarily better.

In Figure 2-28, a third switch is added to the distribution switches in the center. This extra switch adds unneeded complexity to the design and leads to these design questions:

- Where should the root switch be placed? With this design, it is not easy to determine where the root switch is located.
- What links should be in a blocking state? It is very hard to determine how many ports will be in a blocking state.
- What are the implications of STP and RSTP convergence? The network convergence is definitely not deterministic.

Too Much Redundancy Can Lead to Design Issues:

- Root Placement
- Number of Blocked Links
- Convergence Process
- Complex Fault Resolution



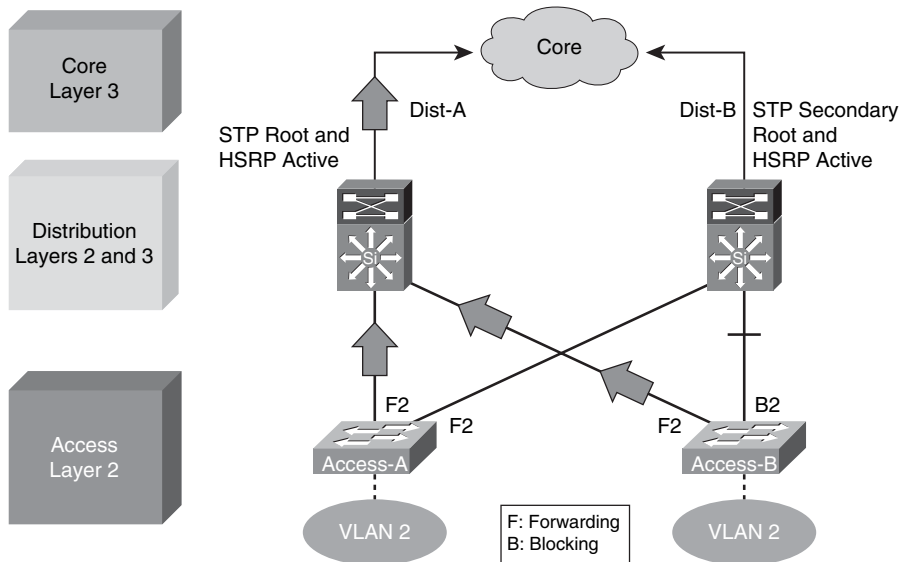
**Figure 2-28** *Too Much Redundancy*

- When something goes wrong, how do you find the source of the problem? The design is much harder to troubleshoot.

### Too Little Redundancy

For most designs, a link between the distribution layer switches is required for redundancy.

Figure 2-29 shows a less-than-optimal design where VLANs span multiple access layer switches. Without a Layer 2 link between the distribution switches, the design is a looped figure-eight topology. One access layer uplink will be blocking. HSRP hellos are exchanged by transiting the access switches.



**Figure 2-29** *Too Little Redundancy*

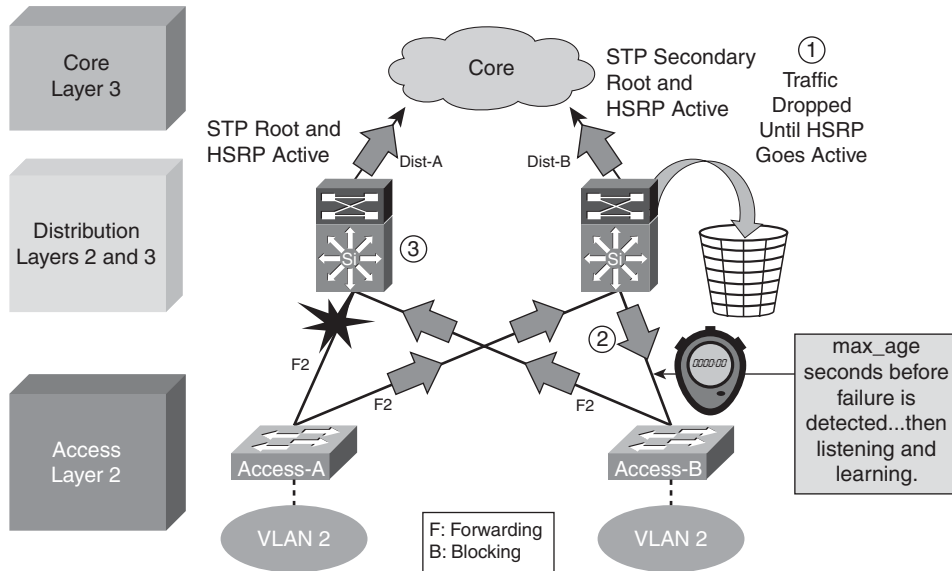
Initially, traffic is forwarded from both access switches to the Distribution A switch that supports the STP root and the primary or active HSRP peer for VLAN 2. However, this design will black-hole traffic and be affected by multiple convergence events with a single network failure.

#### Example: Impact of an Uplink Failure

This example looks at the impact of an uplink failure on the design when there is no link between the distribution layer switches.

In Figure 2-30, when the uplink from Access A to Distribution A fails, three convergence events occur:

1. Access A sends traffic across its active uplink to Distribution B to get to its default gateway. The traffic is black-holed at Distribution B because Distribution B does not initially have a path to the primary or active HSRP peer on Distribution A because of the STP blocking. The traffic is dropped until the standby HSRP peer takes over as the default gateway after not receiving HSRP hellos from Distribution A.



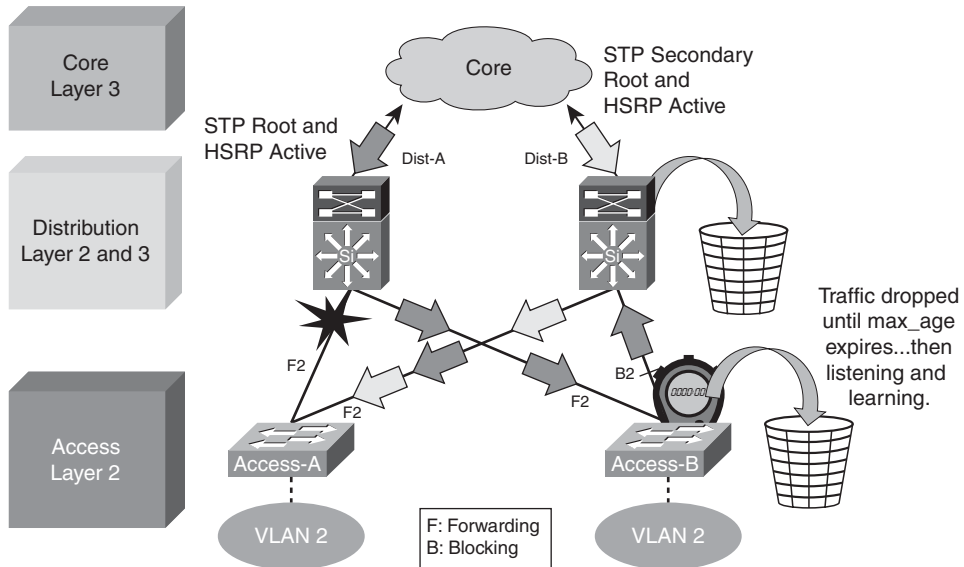
**Figure 2-30** *Impact of an Uplink Failure*

**Note:** With aggressive HSRP timers, you can minimize this period of traffic loss to approximately 900 ms.

2. The indirect link failure is eventually detected by Access B after the maximum-age (max\_age) timer expires, and Access B removes blocking on the uplink to Distribution B. With standard STP, transitioning to forwarding can take as long as 50 seconds. If BackboneFast is enabled with PVST+, this time can be reduced to 30 seconds, and RSTP can reduce this interval to as little as 1 second.
3. After STP and RSTP converge, the distribution nodes reestablish their HSRP relationships and Distribution A (the primary HSRP peer) preempts. This causes yet another convergence event when Access A endpoints start forwarding traffic to the primary HSRP peer. The unexpected side effect is that Access A traffic goes through Access B to reach its default gateway. The Access B uplink to Distribution B is now a transit link for Access A traffic, and the Access B uplink to Distribution A must now carry traffic for both the originally intended Access B and for Access A.

### Example: Impact on Return-Path Traffic

Because the distribution layer in Figure 2-31 is routing with equal-cost load balancing, up to 50 percent of the return-path traffic arrives at Distribution A and is forwarded to Access B. Access B drops this traffic until the uplink to Distribution B is forwarding. This indirect link-failure convergence can take as long as 50 seconds. PVST+ with UplinkFast reduces the time to three to five seconds, and RSTP further reduces the outage to one second. After the STP and RSTP convergence, the Access B uplink to Distribution B is used as a transit link for Access A return-path traffic.



**Figure 2-31** *Impact on Return Path Traffic*

These significant outages could affect the performance of mission-critical applications, such as voice or video. Traffic engineering or link-capacity planning for both outbound and return-path traffic is difficult and complex, and must support the traffic for at least one additional access layer switch.

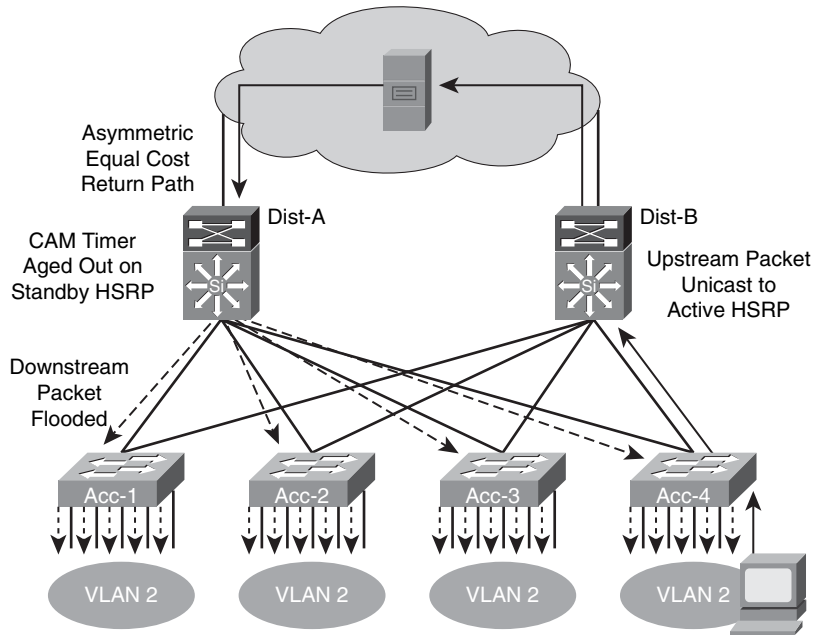
The conclusion is that if VLANs must span the access switches, a Layer 2 link is needed either between the distribution layer switches or the access switches.

### Asymmetric Routing (Unicast Flooding)

When VLANs span access switches, an asymmetric routing situation can result because of equal-cost load balancing between the distribution and core layers.

Up to 50 percent of the return-path traffic with equal-cost routing arrives at the standby HSRP, VRRP, or alternate, nonforwarding GLBP peer. If the content-addressable memory (CAM) table entry ages out before the ARP entry for the end node, the peer may need to flood the traffic to all access layer switches and endpoints in the VLAN.

In Figure 2-32, the CAM table entry ages out on the standby HSRP router because the default ARP timers are four hours and the CAM aging timers are five minutes. The CAM timer expires because no traffic is sent upstream by the endpoint toward the standby HSRP peer after the endpoint initially uses ARP to determine its default gateway. When the CAM entry has aged out and is removed from the CAM table, the standby HSRP peer must forward the return-path traffic to all ports in the common VLAN. The majority of the access layer switches do not have a CAM entry for the target MAC, and they broadcast the return traffic on all ports in the common VLAN. This unicast traffic flooding can have a significant performance impact on the connected end stations because they may receive a large amount of traffic that is not intended for them.



**Figure 2-32** *Asymmetric Routing*

### Unicast Flooding Prevention

The unicast flooding situation can be easily avoided by not spanning VLANs across access layer switches.

Unicast flooding is not an issue when VLANs are not present across multiple access layer switches because the flooding occurs only to switches supporting the VLAN where the traffic would have normally been switched. If the VLANs are local to individual access layer switches, asymmetric routing traffic is flooded on only the one VLAN interface on the distribution switch. Traffic is flooded out the same interface that would be used normally to forward to the appropriate access switch. In addition, the access layer switch receiving the flooded traffic has a CAM table entry for the host because the host is directly attached, so traffic is switched only to the intended host. As a result, no additional end stations are affected by the flooded traffic.

If you must implement a topology where VLANs span more than one access layer switch, the recommended workaround is to tune the ARP timer so that it is equal to or less than the CAM aging timer. A shorter ARP cache timer causes the standby HSRP peer to use ARP for the target IP address before the CAM entry timer expires and the MAC entry is removed. The subsequent ARP response repopulates the CAM table before the CAM entry is aged out and removed. This removes the possibility of flooding asymmetrically routed return-path traffic to all ports. You can also consider biasing the routing metrics to remove the equal cost routes.

## Supporting Infrastructure Services

This section reviews considerations for supporting infrastructure services in highly available enterprise campus networks. Considerations for building a converged network to support IP telephony are discussed. QoS attributes and aspects of the Cisco Catalyst Integrated Security features are also described.

### IP Telephony Considerations

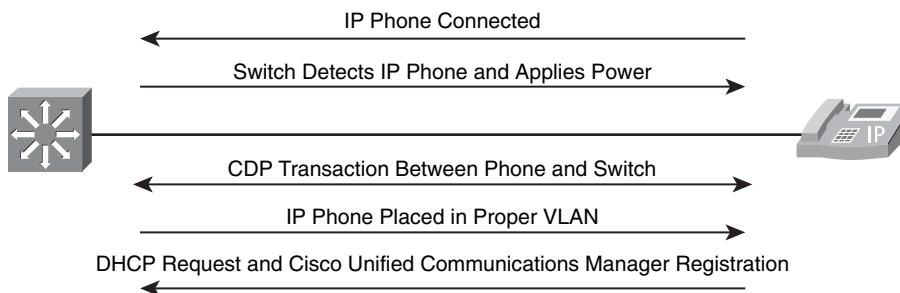
IP telephony services are supported at each layer of the campus network.

High availability, redundancy, and fast convergence needed by IP telephony services are supported throughout the enterprise campus network. QoS features are implemented throughout the network. The distribution layer typically supports policy enforcement.

However, because implementing IP telephony services extends the network edge, IP telephony has the most impact at the access layer of the network. The access layer supports device attachment and phone detection, inline power for devices, and QoS features including classification, scheduling, and the trust boundary.

#### IP Telephony Extends the Network Edge

Because the IP phone is a three-port switch, IP telephony services actually extend the network edge, as shown in Figure 2-33.



**Figure 2-33** *IP Telephony Extends the Network Edge*

The IP phone shown in Figure 2-33 contains a three-port switch that is configured in conjunction with the access switch and Cisco Unified Communications Manager:

- Power negotiation
- VLAN configuration
- 802.1x interoperation
- QoS configuration
- DHCP and Cisco Unified Communications Manager registration

When a Cisco IP phone is connected to the network, Cisco Catalyst multiservice switches detect and integrate the phone with the network. The switches can deliver Power over Ethernet (PoE) using existing copper cabling to power the IP phones. The switches place the IP phones and attached devices in the appropriate VLAN, often using 802.1x services. The switch supports the QoS configuration needed for the IP phones, and provides connection to DHCP servers and Cisco Unified Communications Manager systems for registration.

PoE is the ability for the LAN switching infrastructure to provide power over a copper Ethernet cable to an endpoint or powered device. This capability is also referred to as in-line power, and was originally developed by Cisco Systems in 2000 to support the emerging IP telephony deployments.

To support PoE delivery to power capable devices, a number of issues need to be resolved: phone detection, power delivery, power management, and cable and bandwidth management.

## PoE Requirements

There are two PoE implementations available, and two ways to provide power to the IP phones:

- Cisco line cards support prestandard PoE, IEEE 802.3af, and a mix of devices. IEEE 802.3af-only devices will not negotiate or receive power from an original Cisco PoE-only line card.
- Cisco devices use a bidirectional Cisco Discovery Protocol (CDP) exchange to negotiate the exact power requirements. Power negotiation optimizes power consumption by allowing the switch to reserve only the power needed for the device.

The earlier Cisco prestandard PoE devices initially receive 6.3W and then optionally negotiate their power requirements using CDP. Cisco prestandard devices use a relay in the powered device to reflect a special FastLink pulse for device detection.

The devices based on the IEEE 802.3af power standard initially receive 12.95 watts of power, unless a power-sourcing equipment (PSE) device can detect a specific powered device classification. An 802.3af PSE device applies a voltage in the range of  $-2.8$  to  $-10$  volts on the cable and then looks for a 25K ohm signature resistor in the powered device.

IEEE 802.3af power may be delivered using a PoE-capable Ethernet port, which is referred to as an endpoint PSE, or by a midspan PSE that can be used to deliver PoE in the event an existing non-PoE-capable Ethernet switch is used. An endpoint PSE, such as a PoE-capable Ethernet switch, can use either active data wires of an Ethernet port or spare wires to a powered device. Some midspan PSEs can only implement power over spare pairs of copper cabling and cannot be used to deliver PoE over 1000BASE-T connections.

The IEEE 802.3af power standard classes are shown here in Figure 2-34.

**Note:** A midspan PSE takes up rack space and adds a patch point to every PoE cable, increasing cabling costs and complexity.



IEEE 802.3af Power Classes

Class	Usage	Minimum Power Levels Output at the PSE	Maximum Power Levels at the Powered Device
0	Default	15.4W	0.44 to 12.95W
1	Optional	4.0W	0.44 to 3.84W
2	Optional	7.0W	3.84 to 6.49W
3	Optional	15.4W	6.49 to 12.95W
4	Reserved for Future Use	Treat as Class 0	Reserved for Future Use: A Class 4 Signature Cannot Be Provided by a Compliant Powered Device

**Figure 2-34** PoE Power Classes

## Power Budget and Management

Power budget planning is necessary to determine what devices can be supported today and in the future.

The switches manage power by what is allocated, not by what is currently used. However, the device power consumption is not constant:

- Cisco Unified IP Phone 7960G requires 7W when the phone is ringing at maximum volume.
- Cisco Unified IP Phone 7960G requires 5W when it is on or off hook.

Delivery of PoE using the IEEE 802.3af default classification may significantly increase the power requirements on both the PSE switch and the power infrastructure. To provide PoE in a cost-effective and efficient manner, Cisco Catalyst switches support Cisco Intelligent Power Management (Cisco IPM) in addition to IEEE 802.3af classification. This enables a powered device and PSE to negotiate their respective capabilities to explicitly manage how much power is required to power the device and also how the PSE-capable switch manages the allocation of power to individual powered devices. These Cisco IPM capabilities enable a network and facilities manager to effectively and economically manage the power resources within a wiring closet and help PSE-capable switches meet the objectives of the network.

Power management is complex. Power management can have significant ramifications with respect to the power supply required to drive all the powered devices and line cards, how power is delivered within the switch, how the switch manages power allocation, and finally, for the power delivery requirements of the wiring closet. You need to plan for maximum theoretical draw that so there will be sufficient power available to be allocated to end devices and the line cards in the switch. Even if the PSE and powered device support power classification, the classification ranges are fairly broad and can lead to wasted power budget allocation. When there is insufficient power in a chassis, the power management system will deactivate line cards.

Power requirements can be estimated using the Cisco Power Calculator found at <http://tools.cisco.com/cpc/launch.jsp>.



The Cisco Power Calculator enables you to estimate the power supply requirements for a specific PoE and line card configuration.

The Cisco Power Calculator requires a username and password. The tool allows a series of selections for the configurable products, and provides results showing the output current, output power, and system heat dissipation for a specific configuration.

The calculator is an educational resource and a starting point in planning power requirements; it does not provide a final power recommendation from Cisco.

The Cisco Power Calculator supports the following Cisco product series: Cisco Catalyst 6500, Catalyst 4500, Catalyst 3750, and Catalyst 3560 series switches, and the Cisco 7600 series router.

The Power Consumption Summary screen shown in Figure 2-35 displays the minimum power supply required for the selected configuration and percentage of power usage. The table displays output current (amperes), output power (watts), and heat dissipation (BTUs per hour).

Power Consumption/Heat Dissipation Summary			
Slot	Line Card	Optional DFC	Power over Ethernet Capabilities
1	WS-SUP32-GE-3B	---	---
2	WS-SUP32-GE-3B	---	---
3	-- EMPTY-SLOT --	---	---
4	-- EMPTY-SLOT --	---	---
Minimum Power Supply		Percentage of Power Used	
Single/Redundant PWR-2700-DC/4 with one input		28.05% 	
First Alternative Power Supply		Percentage of Power Used	
Combined PWR-2700-DC/4 with one input		16.80% 	
<b>Total Output Current</b>		<b>Total Output Power</b>	
<b>8.81 Amps</b>		<b>370.02 Watts</b>	
		<b>Total Heat Dissipation</b>	
		<b>1706.10 BTU/Hr</b>	

**Figure 2-35** *Power Consumption Summary*

The Cisco Power Calculator recommends the smallest power supply that meets the requirements of the configuration. The tool reports single and redundant power supply options, and also the combined power configuration mode as appropriate.

The power supply details area shown here in Figure 2-36 displays power utilization with various-sized power supplies.

The Configuration Details section of the Cisco Power Calculator output shown in Figure 2-37 displays the current, power, and heat dissipation for each component.

## Multi-VLAN Access Port

The concept of an access port has been extended to a multi-VLAN access port in the enterprise campus.

Power Supply Details				
Minimum Power Supply	Percentage of Power used	Total Output Current for this PSU (A)	Total Output Current Used (A)	Total Output Current Remaining (A)
<b>Combined WS-CAC-2500W</b>	<b>68.92%</b>	<b>45.77</b>	<b>31.55</b>	<b>14.22</b>
Other Power Supply Options	Percentage of Power used	Total Output Current for this PSU (A)	Total Output Current Used (A)	Total Output Current Remaining (A)
Combined WS-CAC-3000W	67.87%	46.48	31.55	14.93
Single/Redundant WS-CAC-6000W with Dual 110V inputs	49.58%	63.62	31.55	32.07
Combined WS-CAC-6000W with Dual 110V inputs	35.05%	90.00	31.55	58.45

[ Top ]

Configuration Details				
Slot	Line Card	Output Current (A)	Output Power (W)	Heat Dissipation (BTU/Hr)
FAN2	WS-C6K-9SLOT-FAN2	0.00	0.00	546.00

Figure 2-36 Power Supply Details

Configuration Details				
Slot	Line Card	Output Current (A)	Output Power (W)	Heat Dissipation (BTU/Hr)
FAN2	WS-C6K-9SLOT-FAN2	0.00	0.00	546.00
1	WS-X6548-GE-45AF	3.16	132.72	566.55
2	WS-X6548-GE-45AF	3.16	132.72	566.55
3	WS-X6548-GE-45AF	3.16	132.72	566.55
4	WS-X6148A-GE-AF	2.68	112.56	480.49
5	WS-SUP720-3BXL	7.82	328.44	1402.03
6	WS-SUP720-3BXL	7.82	328.44	1402.03
7	-- EMPTY-SLOT --	0.00	0.00	0.00
8	-- EMPTY-SLOT --	0.00	0.00	0.00
9	-- EMPTY-SLOT --	0.00	0.00	0.00
	Sub Total	27.80	1167.80	5530.19
PoE Device	Quantity	Output Current (A)	Output Power (W)	Heat Dissipation (BTU/Hr)
IEEE 802.3af Device - Class 2 (7W)	20	3.75	157.30	193.39
		Output Current (A)	Output Power (W)	Heat Dissipation (BTU/Hr)
	<b>Total</b>	<b>31.55</b>	<b>1324.90</b>	<b>5723.58</b>

Figure 2-37 Configuration Details

Multiservice switches support a new parameter for IP telephony support that makes the access port a multi-VLAN access port. The new parameter is called an auxiliary VLAN. Every Ethernet 10/100/1000 port in the switch is associated with two VLANs:

- A native VLAN for data service that is identified by the port VLAN ID (PVID)
- An auxiliary VLAN for voice service that is identified by the voice VLAN ID (VVID)
- During the initial CDP exchange with the access switch, the IP phone is configured with a VVID.
- The IP phone is also supplied with a QoS configuration using CDP. Voice traffic is separated from data and supports a different trust boundary.

Data packets between the multiservice access switch and the PC or workstation are on the native VLAN. All packets going out on the native VLAN of an IEEE 802.1Q port are sent untagged by the access switch. The PC or workstation connected to the IP phone usually sends untagged packets.

Voice packets are tagged by the IP phone based on the CDP information from the access switch.

The multi-VLAN access ports are not trunk ports, even though the hardware is set to the dot1q trunk. The hardware setting is used to carry more than one VLAN, but the port is still considered an access port that is able to carry one native VLAN and the auxiliary VLAN. The **switchport host** command can be applied to a multi-VLAN access port on the access switch.

**Note:** The switch downloads both the data (native) VLAN and the auxiliary (voice) VLAN to the phone. The IP phone marks any traffic on the voice VLAN by modifying the priority bits in the 802.1Q/p tag to CoS 5 (binary 111), which can later be easily mapped to a Layer 3 marking (for example, DSCP 46 or EF). The trust can also be extended to any CoS markings that may have been set by the attached PC (or can mark these values up or down as desired).

## QoS Considerations

Typical campus networks are built with oversubscription in mind. The network usually has multiple possible congestion points where important traffic may be dropped without QoS.

Most campus links are underutilized. Some studies have shown that 95 percent of campus access layer links are utilized at less than 5 percent of their capacity.

The rule-of-thumb recommendation for data oversubscription is 20:1 for access ports on the access-to-distribution uplink. The recommendation is 4:1 for the distribution-to-core links. When you use these oversubscription ratios, congestion may occur infrequently on the uplinks. QoS is needed for these occasions. If congestion is frequently occurring, the design does not have sufficient uplink bandwidth.

## Recommended Practices for QoS

QoS helps manage oversubscription and speed-transitions in the design. The following are recommended best practices for QoS:

- Deployed end-to-end to be effective
  - Ensures that mission-critical applications are not impacted by link or transmit queue congestion
  - Enforces QoS policies at aggregation and rate transition points
  - Uses multiple queues with configurable admission criteria, and scheduling effective
- QoS is deployed end-to-end with each layer supporting a role

Internet worms and denial-of-service (DoS) attacks can flood links even in a high-speed campus environment. QoS policies protect voice, video, and mission-critical data traffic while giving a lower class of service to suspect traffic.

Aggregation and rate transition points must enforce QoS policies to support preferred traffic and manage congestion. In campus networks, multiple queues with configurable admission criteria and scheduling are required on the LAN ports.

## Transmit Queue Congestion

The type of congestion that is most common in a campus network is called transmit-queue (Tx-queue) starvation.

Both LANs and WANs are subject to Tx-queue congestion:

- During a transition from LAN to WAN, a router has to make the rate transition from 10/100 Ethernet to WAN speeds. When this happens, the router must queue the packets and apply QoS to ensure that important traffic is transmitted first. Tx-queue starvation occurs when incoming packets are received faster than outgoing packets are transmitted. Packets are queued as they wait to serialize out onto the slower link.
- In the campus, as the LAN infrastructure transitions from 10 Gb/s or 1 Gb/s uplinks in the distribution layer to 10/100 Gb/s to the desktop, packets must be queued as they wait to serialize out the 10 or 100 Mb/s link.

The difference between a WAN router and a campus switch is the number of interfaces and the amount of memory associated with each. In the campus, the amount of Tx-queue space is much smaller than the amount of memory available in a WAN router. Because of the small amount of memory, the potential for dropped traffic because of Tx-queue starvation is relatively high.

## QoS Role in the Campus

QoS features are used to prioritize traffic according to its relative importance and provide preferential treatment using congestion management techniques.


Using QoS in the campus network design ensures that important traffic such as voice and video is placed in a queue that is configured so that it optimizes memory usage.

However, the network should provide an adequate level of service for all network traffic, including lower-priority, best-effort traffic under normal circumstances. For best-effort traffic, there is an implied good-faith commitment that there are at least some network resources available.

QoS is also needed to identify and potentially punish out-of-profile traffic such as potential worms, distributed denial of service (DDoS) attacks, and peer-to-peer media-sharing applications that may be placed in a scavenger class and marked with differentiated services code point (DSCP) class selector 1 (CS1). The scavenger class is intended to provide deferential services, or less-than best-effort services, to certain applications. During periods of congestion, scavenger-class traffic is the first to experience Tx-queue starvation and packet loss when the bandwidth is reserved for higher-priority traffic. As demand increases or capacity is reduced, best-effort traffic may also be affected. The minimum goal of high-availability network design is to ensure that high-priority voice, video, and mission-critical data applications are never affected by network congestion.

### Campus QoS Design Considerations

Campus QoS design is primarily concerned with classification, marking, and policing, as illustrated in Figure 2-38.



**Figure 2-38** *Campus QoS Design Considerations*

Queuing is enabled at any node that has the potential for congestion. The edge traffic classification scheme is mapped to the upstream queue configuration. The applications are classified and marked as close to their sources as technically and administratively feasible. Traffic flows are policed as close to their sources as possible.

Multiple queues are the only way to guarantee voice quality, protect mission-critical data, and throttle abnormal sources:

- Voice needs to be assigned to the hardware priority queue. VoIP deployments require provisioning-explicit priority servicing for VoIP traffic and a guaranteed bandwidth service for call-signaling traffic. Strict-priority queuing is limited to 33 percent of the capacity of the link.

- At least 25 percent of the bandwidth of the link is reserved for the default best-effort class, which is the default class for data traffic. Under normal circumstances, the network should provide an adequate level of service for best-effort traffic.
- Scavenger traffic needs to be assigned its own queue with a low threshold to trigger aggressive drops. Applications assigned to this class have little or no contribution to the organizational objectives of the enterprise. Assigning a minimal bandwidth queue to scavenger traffic forces it to be squelched to virtually nothing during periods of congestion, but allows it to be available if bandwidth is not being used for business purposes, which might occur during off-peak hours.

## Cisco Catalyst Integrated Security Features

The Cisco Catalyst Integrated Security capabilities provide campus security on the Cisco Catalyst switches through the use of integrated tools:

- Port security prevents MAC flooding attacks.
- DHCP snooping prevents client attacks on the DHCP server and switch.
- Dynamic ARP Inspection adds security to ARP using the DHCP snooping table to minimize the impact of ARP poisoning and spoofing attacks.
- IP Source Guard prevents IP spoofing addresses using the DHCP snooping table.

### Port Security Prevents MAC-Based Attacks

Port security can be used to prevent MAC-based attacks.

A MAC-based attack occurs when an attacker sends out floods of MAC addresses to a switch to overload the CAM table. When the CAM table limit is reached, the switch can no longer keep track of legitimate addresses and starts flooding all information to all ports.

Port security enables a network administrator to restrict the MAC addresses allowed or the maximum number of MAC addresses on a per-port basis. The allowed MAC addresses on a given port can be either statically configured by the administrator or dynamically learned by the switch. A security violation occurs when either the maximum number of MAC addresses on a given port is exceeded or a frame with a nonsecure source MAC address is seen on that port. The port is then shut down, or alternatively, a Simple Network Management Protocol (SNMP) trap is generated. Aging with either inactivity or a predefined time interval can be configured with port security for the dynamic or static secure MAC addresses.

**Note:** When a port security violation occurs, the port will take one of three actions, depending on how the port is configured: The port will be shut down (Shutdown), frames will simply be ignored (Protect), or the frames will be ignored and the violation counter incremented (Restrict).

## DHCP Snooping Protects Against Rogue and Malicious DHCP Servers

DHCP snooping can be used to protect against rogue and malicious DHCP servers.

In some cases, an intruder can attach a server to the network and have it assume the role of the DHCP server for that segment. This enables the intruder to give out false DHCP information for the default gateway and domain name servers, which points clients to the hacker's machine. This misdirection enables the hacker to become a "man in the middle" and to gain access to confidential information, such as username and password pairs, while the end user is unaware of the attack. DHCP snooping can prevent this. DHCP snooping is a per-port security mechanism used to differentiate an untrusted switch port connected to an end user from a trusted switch port connected to a DHCP server or another switch. It can be enabled on a per-VLAN basis. DHCP snooping allows only authorized DHCP servers to respond to DHCP requests and to distribute network information to clients. It also provides the ability to rate-limit DHCP request on client ports, thereby mitigating the effect of DHCP DoS attacks from an individual client or access port.

## Dynamic ARP Inspection Protects Against ARP Poisoning

Dynamic ARP Inspection can provide protection against ARP poisoning.

ARP does not have any authentication. It is quite simple for a malicious user to spoof addresses by using tools such as ettercap, dsniff, and arpspoof to poison the ARP tables of other hosts on the same VLAN. In a typical attack, a malicious user can send unsolicited ARP replies (gratuitous ARP packets) to other hosts on the subnet with the attacker's MAC address and the default gateway's IP address. Frames intended for default gateways sent from hosts with poisoned ARP tables are sent to the hacker's machine (allowing the packets to be sniffed) or an unreachable host as a DoS attack. ARP poisoning leads to various man-in-the-middle attacks, posing a security threat in the network.

Dynamic ARP Inspection helps prevent the man-in-the-middle attacks by not relaying invalid or gratuitous ARP replies out to other ports in the same VLAN. Dynamic ARP Inspection intercepts all ARP requests and all replies on the untrusted ports. Each intercepted packet is verified for valid IP-to-MAC bindings, which are gathered via DHCP snooping. Denied ARP packets are either dropped or logged by the switch for auditing, so ARP poisoning attacks are stopped. Incoming ARP packets on the trusted ports are not inspected. Dynamic ARP Inspection can also rate-limit ARP requests from client ports to minimize port scanning mechanisms.

## IP Source Guard Protects Against Spoofed IP Addresses

IP Source Guard is a unique Cisco IOS Software feature for Catalyst switches that helps mitigate IP spoofing.

IP Source Guard prevents a malicious host from attacking the network by hijacking its neighbor's IP address. IP Source Guard provides per-port IP traffic filtering of the assigned source IP addresses at wire speed. It dynamically maintains per-port VLAN ACLs based on IP-to-MAC-to-switch port bindings. The binding table is populated either by the DHCP snooping feature or through static configuration of entries. IP Source Guard is typically deployed for untrusted switch ports in the access layer.



### Example Catalyst Integrated Security Feature Configuration

This configuration snippet shows the commands to enable the Catalyst Integrated Security features.

```
ip dhcp snooping
ip dhcp snooping vlan 2-10
ip arp inspection vlan 2-10
!c
interface fastethernet3/1
switchport port-security
switchport port-security max 3
switchport port-security violation restrict
switchport port-security aging time 2
switchport port-security aging type inactivity
ip arp inspection limit rate 100
ip dhcp snooping limit rate 100
ip verify source port-security
!
interface gigabit1/1
ip dhcp snooping trust
ip arp inspection trust
```

---

## Summary

---

This chapter examined design models for high availability and fast convergence for the hierarchical layers of the Cisco Enterprise Campus Architecture. High availability in the campus minimizes convergence time after link and node failures with appropriate redundancy.

VLANs should not span access switches in the campus for predictable fast convergence. Layer 2 designs use RTSP when STP is required, define primary and secondary root switches, and use the Cisco STP toolkit to harden Layer 2. Trunks and channels are tuned for predictable fast convergence. Aggressive mode UDLD is configured on all fiber links.

Oversubscription and bandwidth are managed to minimize complexity and provide deterministic behavior. Layer 3 designs should load balance traffic over redundant equal-cost links built on triangles, not squares. Routing protocols should peer only on transit links, and summarize at the distribution layer. HSRP and GLBP support fast convergence for end devices.

The Layer 2 to Layer 3 boundary is typically at the distribution layer, but it can be placed at the access layer. Campus network designs should avoid daisy chaining access layer switches, provide appropriate redundancy, and avoid asymmetric flooding.

Infrastructure service considerations such as IP telephony and QoS impact the end-to-end network. The access layer supports device attachment, inline power for devices, and multi-VLAN access ports. End-to-end QoS helps manage oversubscriptions and network speed transitions. Tx-queue starvation is the most common campus congestion issue. Cisco Catalyst Integrated Security features provide security at the network edge.

## References

For additional information, refer to these resources:

- Cisco Systems, Inc. “Designing a Campus Network for High Availability,” at [http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/cdccont\\_0900aecd801a8a2d.pdf](http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/cdccont_0900aecd801a8a2d.pdf)
- Cisco Systems, Inc. “Hierarchical Campus Design at a Glance,” at [http://www.cisco.com/application/pdf/en/us/guest/netsol/ns24/c643/cdccont\\_0900aecd800d8129.pdf](http://www.cisco.com/application/pdf/en/us/guest/netsol/ns24/c643/cdccont_0900aecd800d8129.pdf)
- Cisco Systems, Inc. “High Availability Campus Network Design—Routed Access Layer using EIGRP or OSPF” at [http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/ccmigration\\_09186a00805fccbf.pdf](http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/ccmigration_09186a00805fccbf.pdf)
- Cisco Systems, Inc. “Enterprise QoS Solution Reference Network Design Guide,” at [http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/ccmigration\\_09186a008049b062.pdf](http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/ccmigration_09186a008049b062.pdf)
- Cisco Systems, Inc. “Cisco Nonstop Forwarding,” at <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122s/122snwft/release/122s20/fsnsf20s.pdf>
- Cisco Systems, Inc. “Cisco IOS Software Modularity on the Cisco Catalyst 6500 Series Switch,” at [http://www.cisco.com/application/pdf/en/us/guest/products/ps708/c1244/cdccont\\_0900aecd80313e09.pdf](http://www.cisco.com/application/pdf/en/us/guest/products/ps708/c1244/cdccont_0900aecd80313e09.pdf)

- Cisco Systems, Inc. “Cisco Catalyst Integrated Security-Enabling the Self-Defending Network,” at [http://www.cisco.com/application/pdf/en/us/guest/products/ps708/c1244/cdccont\\_0900aecd8015f0ae.pdf](http://www.cisco.com/application/pdf/en/us/guest/products/ps708/c1244/cdccont_0900aecd8015f0ae.pdf)
- Cisco Systems, Inc. “RST-2031: Multilayer Campus Architectures and Design Principles” Networkers 2006 presentation (accessible on a subscription basis), at <http://www.networkersonline.net>
- Cisco Systems, Inc. “RST-3363: Routed Fast Convergence and High Availability” Networkers 2006 presentation (accessible on a subscription basis), at <http://www.networkersonline.net>
- Cisco Systems, Inc. “RST-3466: Cisco IOS Software Modularity–Architecture and Deployment,” Networkers 2006 presentation (accessible on a subscription basis), at <http://www.networkersonline.net>

## Review Questions

Answer the following questions, and then refer to Appendix A, “Answers to Review Questions,” for the answers.

1. Which descriptions best define the core layer? (Choose two.)
  - a. It aggregates end users and supports a feature-rich environment.
  - b. It provides a high-speed, Layer 3 switching environment using hardware-accelerated services.
  - c. It performs high-speed routing and packet manipulations.
  - d. It provides scalability and fast policy-based connectivity.
  - e. It provides the backbone for campus connectivity.
2. What hardware supports Cisco IOS Software modularity? (Choose all that apply.)
  - a. Cisco Catalyst 3750 series
  - b. Cisco Catalyst 4500 series
  - c. Cisco Catalyst 6500 series
  - d. Cisco Catalyst XR series
  - e. All Cisco Catalyst series switches
3. Which statements are correct descriptions of NSF?
  - a. It allows the standby RP to take control of the device after a hardware or software fault on the active RP.
  - b. It is a Layer 3 function that works with SSO to minimize the amount of time a network is unavailable to its users following a switchover.
  - c. It is supported by the Cisco implementation of EIGRP, OSPF, RIP, and BGP protocols.

- d. It synchronizes startup configuration, startup variables, and running configuration.
  - e. The main objective of NSF is to continue forwarding IP packets following an RP switchover.
4. If you need to implement STP, which version is recommended for the enterprise campus?
- a. CST
  - b. HSRP
  - c. MST
  - d. PVST+
  - e. RSTP
5. What is the enterprise recommendation regarding UDLD?
- a. Adjust the default hello timers to three seconds for aggressive mode.
  - b. Enable it to create channels containing up to eight parallel links between switches.
  - c. Enable it in global mode and on every interface you need to support.
  - d. Enable it in global mode to support every individual fiber-optic and Ethernet interface.
  - e. Enable it in global mode to support every individual fiber-optic interface.
6. Which statements are correct descriptions of EtherChannels? (Choose two.)
- a. EtherChannels can reduce the number of peers by creating a single logical interface.
  - b. EtherChannels can increase the number of peers by creating multiple logical interfaces.
  - c. OSPF running on a Cisco IOS Software-based switch will not notice a failed link in a bundle.
  - d. EIGRP may not change the link cost if there is a failed link in a bundle.
  - e. EtherChannel Min-Links feature is supported on PAgP EtherChannels.
7. Which statements are correct descriptions of EtherChannel load balancing? (Choose three.)
- a. Load balancing using an alternate input hash can be tuned with the **cef port-channel load-balance** command.
  - b. Load balancing using an alternate input hash can be tuned with the **port-channel load-balance** command.
  - c. The default input hash value of Layer 3 for the source and destination does not load balance across the links.

- d.** The default input hash value of Layer 3 for source and destination and Layer 4 port does load balance across the links.
  - e.** To achieve the best load balancing, use alternating hashes in the core and distribution layer switches.
  - f.** To achieve the best load balancing, use two, four, or eight ports in the port channel.
- 8.** What are the reasons that passive interfaces should be implemented at distribution layer ports facing the access layer? (Choose two.)
  - a.** To limit unnecessary peering across the access layer switches when the Layer 2 to Layer 3 boundary is in the distribution layer
  - b.** To limit unnecessary peering across the access layer switches when the Layer 2 to Layer 3 boundary is in the access layer
  - c.** To provide high availability in the event of a link or node failure
  - d.** To support transit traffic through the access layer in the event of a link or node failure
  - e.** To avoid transit traffic through the access layer in the event of a link or node failure
- 9.** What are the advantages of GLBP in the distribution layer? (Choose two.)
  - a.** GLBP provides all the benefits of HSRP and includes load balancing when VLANs do not span the access switches.
  - b.** A convergence event on the uplink affects only half as many hosts as compared to HSRP when VLANs do not span the access switches.
  - c.** A convergence event on the uplink affects is processed in half the time as compared to HSRP when VLANs do not span the access switches.
  - d.** STP can block one of the access layer uplinks, and there is at most a two-hop Layer 2 path for upstream traffic when VLANs span access switches.
  - e.** STP can block one of the access layer uplinks, and there is at most a two-hop Layer 3 path for upstream traffic when VLANs span access switches.
- 10.** What is a potential issue when daisy chaining access layer switches?
  - a.** It is not easy to determine where the root switch is located.
  - b.** It is very hard to determine how many ports will be in a blocking state.
  - c.** The design will black-hole traffic and be affected by multiple convergence events with a single network failure.
  - d.** There is a danger that black holes will occur in the event of a link or node failure when the distribution interconnection is Layer 2.
  - e.** There is a danger that black holes will occur in the event of a link or node failure when the distribution interconnection is Layer 3.

- 11.** What is the best mechanism to prevent unicast flooding issues?
  - a.** Bias the routing metrics to remove equal-cost routes.
  - b.** Do not span VLANs across multiple access switches.
  - c.** Span VLANs across multiple access switches.
  - d.** Tune ARP timers so they exceed CAM timers.
  - e.** Tune CAM timers so they exceed ARP timers.
  
- 12.** What hardware is supported by the Cisco Power Calculator? (Choose all that apply.)
  - a.** Cisco Catalyst 3750 series
  - b.** Cisco Catalyst 4500 series
  - c.** Cisco Catalyst 6500 series
  - d.** Cisco Catalyst XR series
  - e.** All Cisco Catalyst series switches
  
- 13.** What features do Cisco Catalyst Integrated Security capabilities provide? (Choose three.)
  - a.** DHCP snooping prevents client attacks on the DHCP server and switch.
  - b.** Dynamic ARP Inspection adds security to ARP to minimize the impact of ARP poisoning and spoofing attacks.
  - c.** DHCP snooping prevents client attacks on the DHCP server and switch using the Dynamic ARP Inspection table.
  - d.** IP Source Guard prevents IP spoofing using the DHCP snooping table.
  - e.** IP Source Guard prevents IP spoofing using the Dynamic ARP Inspection table.

# Index

---

## Numerics

(\*, G), 445-447

(S, G), 444-447

### 1RU data center design

cabinet design, 218-219

server design, STP, 232

versus modular design, 217-218

802.3 MAC address format, 432

## A

ABRs (area border routers),  
connecting, 117-118

### access control

host receiver-side, 468

packet filter-based, 467

### access layer

Cisco Enterprise Campus  
Architecture, 23-24

data center model, 178, 193

*Layer 2 FlexLink design,*  
203-206

*Layer 2 loop-free design,*  
*topologies, 199-203*

*Layer 2 looped design,*  
*topologies, 195-197*

*Layer 3 designs, 208-210*

Hierarchical Network Model, 3

### access layer switches

daisy chaining, 63-64

passive interfaces, configuring, 51

ACL-friendly addressing,  
implementing, 89

active component alignment in data  
center model, 189

active/active firewall topology,  
338-339

asymmetric routing, 340-341

active/active service module design,  
data center model, 191

active/standby service module design,  
data center model, 190

adjacencies, effect on scalability,  
109-110

advanced WAN services, selecting,  
166

aggregation layer (data center model),  
178, 185

active component alignment, 189

active/active service module design,  
191

active/standby service module design,  
190

inbound path preference, establishing,  
192-193

integrated services modules, 188

scaling, 185-186

STP design recommendations, 186  
VRFs, 193

AH (Authentication Header), 400

alternate paths, providing in Cisco  
Enterprise Campus Architecture,  
30

alternative cell VoWLANs, AP placement, 516  
amplifiers, 145  
antennas (WLAN), 480  
  diversity, 483  
anycast RPs, 455  
application layer (Cisco SONA), 6, 10  
application optimization cycle, Cisco IOS Software tools, 528-529  
apply policy phase of Cisco IOS OER, 324  
APs (access points)  
  autonomous AP, 487, 520  
  BSS, 487  
  lightweight AP, 487-489, 520  
  nonoverlapping channels, 508-509, 511  
  placement in VoWLANs, 515-516  
  roaming, 488, 521  
  SSID, 488  
areas (OSPF)  
  designing, 110-111  
  filtering, 118-119  
  hub-and-spoke design requirements, 114  
ASDM (Cisco Adaptive Security Device Manager), 403, 411  
ASM (Any Source Multicast), 447  
assigning multicast addresses, 433

asymmetric routing, 68  
  with firewalls, 339-341  
  unicast flooding prevention, 69  
attacks in multicast environments, traffic forwarding, 465-466  
authentication for VoWLANs, 502-503  
Authorized Self-Study Guides  
  *Building Scalable Cisco Internetworks*, 297  
  *Designing for Cisco Internetwork Solutions (DESGN)*, Second Edition, 479  
Auto-RP, 456  
autonomous APs, 487, 520

## **B**

BackboneFast, 37  
backdoor routes (MPLS VPNs), 164  
bandwidth, managing with EtherChannel, 45-47  
base e-commerce module design, 299-301  
BB\_Credit flow control, 252  
best practices for scaling STP, 233  
best-effort delivery, UDP, 430  
BFD (Bidirectional Forwarding Detection), 123-124  
BGP (Border Gateway Protocol)  
  community attribute, 331  
  MED attribute, 331



- scaling, 124
  - full-mesh iBGP scalability*, 125-130, 133
- transit traffic, 97
- e-commerce designs, tuning, 315
- Bidir-PIM, 450-452**
  - BSR, 459-460
  - RPs, 454-456
- bit splitting, 90-91**
- black-holing packets, 317**
- blade servers, 210**
  - connectivity options, 212-213
  - failover features, 215
  - InfiniBand, 213-214
- border area filtering, 118**
- BPDU Guard, 37-39**
- BSCI (Building Scalable Cisco Internetworks) course, 94**
- BSR (Boot Strap Router), 459-460**
- BSS (Basic Service Set), 487**
- bump in the wire mode, 286**
- business drivers for SAN deployments, 246**
- business risk assessment for WAN services, 167**

## **C**

- cabinet design**
  - with 1RU switching, 218-219
  - with modular access switches, 221
- CAC, 501**
- cache management, NetFlow, 540**
- calculating**
  - oversubscription per uplink on access layer switches, 224
  - PPS rate, 413-414, 417
- CAM (content-addressable memory) table, 440**

- campus multicast protocols, 434**
  - CGMP, 439
  - ICMPv1, 436
  - IGMP, 436
  - IGMPv3, 437
- Catalyst switches, Cisco Catalyst Integrated Security, 78**
  - DAI, 79
  - DHCP snooping, 79
  - IP Source Guard, 79-80
  - Port Security, 78
- CCP (Cisco Configuration Professional), 411**
- CEF (Cisco Express Forwarding), 47**
- cell overlap guidelines, 509-510**
- cell phones versus VoWLAN solutions, 491**
- CGMP (Cisco Group Management Protocol), 439**
- Cisco ACE modules, 227, 290**
- Cisco ADSM (Adaptive Security Device Manager), 411**
- Cisco ASA 5500 series performance, 395-397**
- Cisco AutoQoS, 551, 554-556**
- Cisco Catalyst Integrated Security, 78**
  - DAI, 79
  - DHCP snooping, 79
  - IP Source Guard, 79-80
  - Port Security, 78
- Cisco Easy VPN, 384, 401**
- Cisco Enterprise WAN and MAN Architecture, 8**
- Cisco Enterprise Branch Architecture, 8**
- Cisco Enterprise Campus Architecture, 7**
  - access layer, 23-24
  - Cisco Catalyst Integrated Security, 78
    - DAI, 79
    - DHCP snooping, 79

- IP Source Guard*, 79-80
- Port Security*, 78
- core layer, 26-27
- distribution layer, 25
- high-availability considerations, 28
  - alternate paths, providing*, 30
  - Cisco IOS Software Modularity*, 33-36
  - optimal redundancy, implementing*, 28-30
  - single points of failure, avoiding*, 30
- IP telephony considerations, 70
  - multi-VLAN access port*, 73
  - PoE requirements*, 71
  - power budget planning*, 72-73
- QoS considerations, 75-77
  - transmit queue congestion*, 76
- Cisco Enterprise Data Center Architecture**, 8
- Cisco Enterprise Edge Architecture**, 7
- Cisco Enterprise Teleworker Architecture**, 8
- Cisco EPL (Ethernet Private-Line) service**, 151
- Cisco ERS (Ethernet Relay Service)**, 152-153
- Cisco EWS (Ethernet Wire Service)**, 153
- Cisco Fabric Manager**, 262
- Cisco IBNS (Identity-Based Networking Services Solution)**, 348
- Cisco IOS Flexible NetFlow**, 535
- Cisco IOS OER (Optimized Edge Routing)**, 320-321
  - apply policy phase, 324
  - learn phase, 322
  - measure phase, 322-323
  - optimize phase, 324
  - topologies, 324-325
  - verify phase, 324
- Cisco IOS Software**
  - Modularity Architecture, 33-36
  - network management
    - application optimization cycle*, 528-529
    - IP SLAs*, 556-567
    - NBAR*, 546-551, 554-556
    - NetFlow*, 534-537, 540-545
    - syslog*, 529-533
- Cisco ISRs (Integrated Services Routers)**, security performance, 395
- Cisco MDS 9000 SAN fabric family**, 253
  - SANTap, 256-258
- Cisco NAA (NAC Appliance Agent)**, 350, 366
- Cisco NAC Appliance**, 348
  - Cisco NAMs, 351
  - Cisco NAS
    - client access modes*, 354
    - deployment options*, 352-354
    - operating modes*, 354
    - physical deployment models*, 355
  - client security software applications, 366-368
  - components of, 349
  - Layer 2 in-band design, 355-357
  - Layer 2 out-of-band design, 358-359
  - Layer 3 in-band design, 359-360
  - Layer 3 out-of-band design, 360-362
  - policy updates, 350
  - posture validation, architectural components, 362
  - process flow, 351
- Cisco NAC Framework**, 348, 362-364
  - router platform support, 364-365
  - switch platform support, 366
- Cisco NAM (NAC Appliance Manager)**, 349-351

**Cisco NAS (NAC Appliance Server), 350**  
 client access modes, 354  
 deployment options, 352-354  
 operating modes, 354  
 physical deployment models, 355

**Cisco NSF (Nonstop Forwarding), 32-33**

**Cisco Optical Metro Ethernet Solution, 150**

**Cisco PIX Device Manager, 411**

**Cisco Power Calculator, 73**

**Cisco SDM (Security Device Manager), 410**

**Cisco Security Agent, 367**

**Cisco Security Manager, 411**

**Cisco SONA (Service-Oriented Network Architecture), 5**  
 application layer, 10  
 benefits of, 6  
 infrastructure services, 9-10

**Cisco StackWise technology, 64**

**Cisco STP Toolkit, 37**

**Cisco Trust Agent, 367**

**Cisco UWN (Unified Wireless Networks)**  
 architecture of, 488-489  
 mobility groups, 498-499  
 roaming, 494  
     *enhanced neighbor lists*, 499  
     *intercontroller roaming*, 495-497, 521, 524  
     *intracontroller roaming*, 495, 521

**CiscoView Device Manager, 411**

**CiscoWorks IPM (Internetwork Performance Monitor), 568**

**client access modes (Cisco NAS), 354**

**client NAT, 293**

**client security software applications (Cisco NAC Appliance), 366-368**

**clientless access for SSL VPN, 385-386**

**CNF (Cisco Nonstop Forwarding), 278**

**codecs, 512**

**collapsed-core design**  
 large-scale, dual-fabric core-edge design, 264-265  
 medium-scale, dual-fabric collapsed-core design, 263-264  
 single-switch collapsed-core design, 262  
 small-scale, dual-fabric collapsed-core design, 263

**commands**  
 debug ip nbar unclassified-port-stats, 555  
 default-information originate, 95  
 ip default-network, 96  
 ip igmp access-group, 468  
 ip multicast ttl-threshold, 467  
 ip pim accept-register, 470-471  
 ip pim autorp listener, 459  
 port-channel min-links, 225  
 process restart, 35  
 show auto discovery qos, 554  
 show interfaces fastethernet include rate, 417  
 show ip nbar protocol-discovery, 549  
 show spanning-tree summary total, 231  
 spanning-tree cost interface, 56

**community attribute (BGP), 331**

**community PVLAN ports, 344**

**community VLANs, 344-345**

**comparing**  
 Layer 2/Layer 3 access designs, 207-208, 216-217  
 modular versus 1RU data center designs, 217-218  
 multicast and unicast, 426-428  
 route reflectors and confederations, 133

**confederations**

- iBGP, scaling, 129-130, 133
- versus route reflectors, 133

**configuring**

## EtherChannel

- best practices*, 43
- LACP*, 44
- PAgP*, 43-44

## Layer 2

- best practices*, 36-37
- Cisco STP Toolkit*, 37
- standards and features*, 37

passive interfaces on access switches, 51

## trunks

- best practices*, 39
- DTP*, 41-42
- VTP*, 40

UDLD, *best practices*, 42

**connecting**

- ABRs, 117-118
- multiple ISPs, 295

connectivity options for blade servers,  
212-214

context load balancers, 288

contexts, 284

**control plane, 34**

modularized processes, 35-36

**convergence**

- EIGRP, 102-103
- OSPF, 121-122
  - BFD*, 123-124
  - iSPF*, 122-123

SAN design requirements, 260

**core layer**

Cisco Enterprise Campus Architecture,  
26-27

data center model, 177-181

*EIGRP design recommendations*,  
183

*OSPF design recommendations*,  
182

Hierarchical Network Model, 4

creating multicast distribution trees,  
442

CSM (Cisco Context Switching Module),  
289

CSS (Content Services Switch), 289
 

- one-armed SLB e-commerce module  
design with CSS, 313

CST (Common Spanning Tree), 37

CUWN (Cisco Unified Wireless  
Network), CAC, 501

CWDM (coarse wavelength division  
multiplexing), 143, 266

**D**

DAI (Dynamic ARP Inspection), 79

**daisy chaining**

- access layer switches, 63-64
- SCSI devices, 251

DAS (directly attached storage),  
245, 249

data center model. **See also** data centers  
access layer, 193

*Layer 2 FlexLink design*, 203-206

*Layer 2 loop-free design*, 199-203

*Layer 2 looped design*, 195-197

*Layer 3 designs*, 208-210

aggregation layer, 185

*active component alignment*, 189

*active/active service module  
design*, 191

*active/standby service module  
design*, 190

*inbound path preference,  
establishing*, 192-193

- integrated service modules, 188*
- scaling, 185-186*
- STP design recommendations, 186*
- VRFs, 193*
- benefits of, 179
- core layer, 179-181
  - EIGRP design recommendations, 183*
  - OSPF design recommendations, 182*
- data centers**
  - 1RU design
    - cabinet design, 218-219*
    - server design, STP, 232*
    - versus modular design, 217-218*
  - access layer switches, calculating oversubscription per uplink, 224
  - high availability, 233, 238
    - failover times, 236-237*
    - NIC teaming configurations, 234-235*
  - modular design
    - cabinet design, 221*
    - versus 1RU design, 217-218*
  - scalability, 228
    - server NIC density, 223*
  - service layer switches, 226
  - STPs, 228
    - logical interfaces, 230-231*
    - scaling, 229-232*
- data plane, 34**
- dB (decibels), 480**
- dBm (dB milliwatt), 480, 520**
- debug ip nbar unclassified-port-stats command, 555**
- default gateway redundancy, 54**
- default routes, 94-96**
- default-information originate command, 95**
- defensive filtering, 97**
- deploying**
  - IP SLA measurements, 564-566
  - NetFlow, 545
- deployment models**
  - IPs, 372-374
  - PIM
    - ASM, 447*
    - Bidir-PIM, 450-452*
    - PIM-SM, 448-450*
    - SSM, 452-453*
- design phase of PPDIIO network lifecycle approach, 11-13**
  - characterizing the existing network, 15
  - dividing network into areas, 16
  - identifying customer requirements, 14
  - top-down design practices, 16
- designing**
  - BGP, scalability, 124-130, 133
  - IP addressing
    - bit splitting, 90-91*
    - redistribution, 99-101*
    - route filtering, 96-97*
    - route summarization, 87-88*
    - for VPN clients, 91*
  - MPLS VPNs, customer considerations, 163-164
  - OSPF networks
    - ABRs, connecting, 117-118*
    - area filtering, 118-119*
    - areas, 110-111*
    - fast convergence, 121-124*
    - Flooding Reduction, 121*
    - full-mesh topology, 120-121*
    - hub-and-spoke design, 113-116*
    - mesh groups, 120-121*
    - required hierarchy, 112-113*
    - route summarization, 113*

- SANs
  - convergence, 260*
  - device oversubscription, 259-260*
  - fault isolation, 260*
  - stability, 260*
  - topology requirements, 258-259*
  - traffic management, 260*
- destination NAT, 92
- developments in SAN extension, 270
- device oversubscription requirements for SAN design, 259-260
- DHCP snooping, 79
- direct server traffic flows, one-armed SLB two-firewall layer design, 308
- disabling multicast groups for IPv6, 471
- distributed data centers, 298-299
- distribution layer
  - Cisco Enterprise Campus Architecture, 25
  - Hierarchical Network Model, 4
  - route summarization, 51-52
- distribution trees, 446
- DM fallback, 458-459
- DM flooding, 458-459
- DMVPN (Dynamic Multipoint VPN), 405-407
- DMZ (demilitarized zone), placing VPN device on, 398-399
- DNS-based site selection, 325-326
- documentation as tool in e-commerce design, 281
- DPT (Cisco Dynamic Packet Transport), 146
- DRs (designated routers), 436
- DTP (Dynamic Trunking Protocol), 401-402
- dual-homing, 295-296
- DVTIs (dynamic VTIs), 408
- DWDM (dense wavelength division multiplexing), 144, 266
- E**
- e-commerce designs
  - BGP tuning, 315
  - DNS-based site selection, 325-326
  - EOT (Enhanced Object Tracking), 317-320
  - firewalls, 281
    - e-commerce module topology, 282*
    - firewall contexts, 284*
    - firewall modes, 286-287*
    - server as application gateway, 282-284*
    - virtual firewall layers, 285*
  - high availability, 277
    - people component, 279-280*
    - processes, 280*
    - redundancy, 278*
    - technology, 278*
    - tools, 281*
  - integrated designs
    - base e-commerce designs, 299-301*
    - one-armed SLB two-firewall designs, 305, 308*
    - one-armed SLB with CSS, 313*
    - one-armed SLB with firewall contexts design, 308-310, 313*
    - with two firewall layers, 304-305*
  - OER, 320, 322
    - apply policy phase, 324*
    - learn phase, 322*
    - measure phase, 322-323*
    - optimize phase, 324*
    - topologies, 324-325*
    - verify phase, 324*

- SLBs, 288
  - inline bridge mode*, 291-292
  - one-armed mode*, 292-294
  - SLB router mode*, 290
- testing, 313-315
- topologies connecting multiple ISPs, 295
  - distributed data centers*, 298-299
  - one firewall per ISP*, 295
  - stateful failover with common external prefix*, 296
- e-commerce module topology, 282
- E-Ports, 253
- Easy VPN Server Wizard, 401-402
- ECMP (Equal Cost Multipath) routing, 342
- EDFA (erbium-doped fiber amplifier), 144
- EIGRP (Enhanced Interior Gateway Routing Protocol)
  - data center model, design recommendations, 183
  - default routes, 96
  - equal-cost routing, 102
  - fast convergence, 102-103
  - feasible successors, 102
  - Layer 3 access to distribution interconnection, 62
  - multi-AS systems, scaling, 104-108
  - scaling, 102
  - transit traffic, 97
- EISL (Enhanced Inter-Switch Link), 253
- EMS (Ethernet Multipoint Service), 150, 153
- encryption, VoWLANs, 502-503
- end-to-end QoS, 154, 501
- enhanced neighbor lists, 499
- Enterprise Campus Architecture, 23
- enterprise campus networks.
  - See also** enterprise networks
  - Cisco Catalyst Integrated Security, 78
    - DAI*, 79
    - DHCP snooping*, 79
    - IP Source Guard*, 79-80
    - Port Security*, 78
  - IP telephony
    - multi-VLAN access port*, 73
    - PoE requirements*, 71
    - power budget planning*, 72-73
  - QoS considerations, 75-77
- enterprise networks
  - RPR, 146
  - VLPS, 158
  - VoWLANs, 491
    - voice impact on WLANs*, 493
    - voice-ready architecture*, 492
- EOT (Enhanced Object Tracking), 317-320
- EPL (Ethernet Private-Line) service, 150-151
- equal-cost routing, 102
- ERMS (Ethernet Relay Multipoint Service), 150, 154
- error correction, 248
- error detection, 248
- ERS (Ethernet Relay Service), 150-153
- ESCON (Enterprise System Connection), 256
- ESM (Cisco IOS Embedded Syslog Manager), 531
- estimated PPS based on branch profile, 415
- EtherChannel
  - bandwidth, managing, 45-47
  - configuration best practices, 43
  - LACP, 44

- Min-Links, 225
- PAgP, 43-44
- EVCs (Ethernet Virtual Circuits), 150
- EWS (Ethernet Wire Service), 150, 153
- example of Hierarchical Network Model, 4
- excessive redundancy, dealing with, 65
- export versions, 534
  - NetFlow, 542

## F

- fabric, 245
- fabric routing, 254
- fabric virtualization, 253
- failover
  - blade servers features, 215
  - e-commerce designs, testing, 313
- fan-out ratio
  - calculating for large-scale, dual-fabric core-edge design, 265
  - SAN design requirements, 259
- fast convergence, OSPF, 121-124
- fault isolation requirements for SAN design, 260
- FCIP (Fibre Channel over IP), 254, 266-268
  - write acceleration, 270
- FCP (Fibre Channel Protocol), 252
- feasible successors, 102
- FHRPs (first-hop routing protocols), 23, 297. **See also** first-hop redundancy
  - convergence, optimizing, 58
- Fibre Channel, 245, 251
  - communications model, 252
  - FSPF, 254-255
- FICON (Fiber Connectivity), 256
- filtered redistribution, 99-100
- filtering

- EIGRP redistribution with route tags, 105
- EIGRP routing updates with route tags, 105-107
- OSPF areas, 118-119
- filters, 143**
- firewalls**
  - active/active topology, 338-339
  - asymmetric routing, 339-340
    - with active/active topology, 340-341*
  - contexts, 284
    - design considerations, 336-337*
  - e-commerce design, 281
    - e-commerce module topology, 282*
    - firewall contexts, 284*
    - firewall modes, 286-287*
    - server as application gateway, 282-284*
    - virtual firewall layers, 285*
  - integrating VPN with, 399
  - modes, 286-287, 333-334
  - multiple FWSMs, 341-344
    - community VLANs, 345*
    - isolated PVLAN ports, 344*
  - virtual firewalls, 335
    - firewall contexts, 336-337*
    - MSFC placement, 337*
  - zone-based policy firewalls, 346-347
- first-hop redundancy, 53**
  - FHRP convergence, optimizing, 58
  - GLBP, 55-56
  - HSRP preemption, 54
- Five Steps to Securing Your Wireless LAN and Preventing Wireless Threats, 504**
- Flexible NetFlow, 542-543
- FlexLink designs, 203-206
- Flooding Reduction (OSPF), 121



- flow monitors, 543
- flow records, NetFlow, 537, 540
- flows, NetFlow, 535
- FlowSets, 543
- FSPF (Fabric Shortest Path First), 254-255
- full-mesh iBGP scalability, 125
  - confederations, 129-130, 133
  - route reflectors, 126-128
- full-mesh OSPF topologies, 120-121
- FWSM (Firewall Services Module), 335, 388
  - multiple, load balancing, 341-343
    - community VLANs*, 345
    - isolated PVLAN ports*, 344

## G

- gain (RF), 480
- gateway modes for Cisco NAS, 353
- GDOI (Group Domain of Interpretation), 409
- GET (Cisco Group Encrypted Transport) VPNs, 409-410
- GLBP (Gateway Load Balancing Protocol), 24, 55-56
  - Layer 3 distribution switch interconnection, 60
- globally scoped IP addresses, 431
- GLOP addressing, 434
- graceful restart, 278
- GRE over IPsec, 403-405
- GRE tunnels, 391
- group-specific queries (IGMPv2), 437
- GSLB (Global Server Load Balancing), 325-326
- GSS (Cisco Global Site Selector),
  - DNS-based site selection, 325-326

## H

- H-VLPS, 159
- hard disks
  - interfaces, 247
  - SCSI, 250
- hardware-assisted data compression over FCIP, 270
- hardware-based zoning, 256
- HBA (host bus adapter), 246
- HCAAs (host channel adapters), 214
- headend VPN devices, 391
- hierarchical monitoring, IP SLA measurements, 567
- hierarchical network model, 3-4
- high availability
  - in Cisco Enterprise Campus Architecture, 28
    - alternate paths, providing*, 30
    - Cisco IOS Software Modularity*, 33-36
    - optimal redundancy, implementing*, 28-30
    - single points of failure, avoiding*, 30
  - in data centers, 233, 238
    - failover times*, 236-237
    - NIC teaming configurations*, 234-235
  - for e-commerce designs, 277
    - people component*, 279-280
    - processes component*, 280
    - redundancy*, 278
    - technology component*, 278
    - tools component*, 281
- high-availability SAN extension, 271
- HIPAA (Health Insurance Portability and Accountability Act), 392
- HIPSs (host-based intrusion prevention systems), 370

- host receiver-side access control, 468
- HSRP (Hot Standby Routing Protocol),
  - Layer 3 distribution switch interconnection, 60
- HSRP preemption, 54
- hub-and-spoke OSPF design, 113-116
- I**
- IANA (Internet Assigned Numbers Authority), 430
- iBGP (Internal BGP), scaling, 125
  - confederations, 129-130, 133
  - route reflectors, 126-128
- IBNS (Cisco Identity Based Networking Services), 348
- IDSs (intrusion detection systems), 368
  - deployment options, management interface deployment, 374
  - design considerations, 371
  - managing, 374-376
  - monitoring, 374-376
  - security management and monitoring infrastructure, 370
  - sensors, 369
- IEEE 802.11 operational standards, 481-483
- IEEE 802.11b, cell overlap guidelines, 509-510
- IEEE 802.11e, 500
- IEEE 802.11g, cell overlap guidelines, 509-510
- IGMP (Internet Group Management Protocol), 439
- IGMPv1, 436
- IMIX (Internet mix) traffic, 412
- implement phase of PPDIOO network lifecycle approach, 12
- implementing
  - ACL-friendly addressing, 89
  - role-based addressing, 90
- in-band traffic flow deployment model (Cisco NAS), 354
- inbound path preference, establishing in data center model, 192-193
- inbound route tags, filtering EIGRP routing updates, 105-107
- INCITS (InterNational Committee for Information Technology Standards), 245
- industrial, scientific, and medical bands, 480
- InfiniBand, 213-214
- infrastructure services, Cisco SONA, 9-10
- initiator/target communication, 251
  - zoning, 255
- inline bridge mode (SLBs), 291-292
- inline VLAN paring, 373
- insufficient redundancy, dealing with, 66-68
- integrated e-commerce designs
  - base e-commerce module design, 299-301
  - one-armed SLB two-firewall design, 305, 308
  - one-armed SLB with CSS design, 313
  - one-armed SLB with firewall contexts design, 308-310, 313
  - with two firewall layers, 304-305
- integrated services branch WANs, 414
- integrated services module, 188
- interactive services layer (Cisco SONA), 5
- interarea filtering, 119
- intercontroller roaming
  - Layer 2, 495-496, 521, 524
  - Layer 3, 497, 521
- interdomain multicast protocols, 434
  - PIM, 440
  - distribution trees*, 441

- shared distribution trees*, 445
  - source trees*, 443
- interfaces**
  - configuring passive interfaces on access switches, 51
  - on hard disks, 247
- intracontroller roaming, CUWN**, 495, 521
- IP addressing**
  - ACL-friendly, implementing, 89
  - NAT, 92-93
  - redistribution
    - filtered redistribution*, 99-100
    - migrating between routing protocols*, 101
  - role-based, 90
  - route filtering, 96-97
  - route summarization, 87, 93-94
    - bit splitting*, 90-91
    - default routes*, 95-96
    - originating default routes*, 94-95
    - summary address blocks*, 88
  - for site-to-site VPNs, 392-393
  - subnets, redesigning, 88
  - summary address blocks, 89
  - for VPN clients, designing, 91
- ip default-network command**, 96
- IP flows, NetFlow**, 536
- ip igmp access-group command**, 468
- IP multicast and VLPS**, 162
- ip multicast ttl-threshold command**, 467
- ip pim accept-register command**, 470-471
- ip pim autorp listener command**, 459
- IP SLAs (Service Level Agreements)**, 556
  - IPM support, 568
  - measurements, 557-558
    - deploying*, 564-566
    - hierarchical monitoring*, 567
  - network management applications, selecting, 568-569
  - operations, 560-561
  - SNMP features, 563-564
  - source/responder components, 560
  - timestamps, 562
- IP Source Guard**, 79-80
- IP telephony in enterprise campus networks**, 70
  - multi-VLAN access port, 73
  - PoE requirements, 71
  - power budget planning, 72-73
- IPM (CiscoWorks Internetwork Performance Monitor)**, 568
- IPsec VPNs**, 400
  - Cisco Easy VPN, 401
  - Cisco router performance with, 393-395
  - DMVPN, 405-407
  - GET VPNs, 409-410
  - GRE over IPsec, 403-405
  - routing protocol considerations, 417
  - VTI, 407-408
- IPSs (intrusion prevention systems)**, 368
  - deployment options, 372-373
    - management interface deployment*, 374
  - design considerations, 371
  - HIPSS, 370
  - managing, 370, 374-376
  - monitoring, 370, 374-376
  - sensors, 369
- IPv6 multicast groups, disabling**, 471
- iSCSI (Internet Small Computer Systems Interface)**, 267-269
  - scaling, 270

iSLB (iSCSI Server Load Balancing), 269  
 isolated PVLAN ports, 344  
 isolated VLANs, 344  
 iSPF (incremental SPF), 122-123  
 ISPs (Internet service providers), connecting multiple  
   distributed data centers, 298-299  
   one firewall per ISP, 295  
   stateful failover with common external prefix, 296  
 ISRs (integrated services routers), security performance, 395  
 IVR (Inter-VSAN Routing), 253-254

## J-K-L

JBOD (just a bunch of disks), 247  
 LACP (Link Aggregation Control Protocol), 43-44  
 large-scale, dual-fabric core-edge design, 264-265  
 large cell VoWLANs, AP placement, 515  
 Layer 2 access designs, comparing with Layer 3, 216-217  
 Layer 2 client access mode (Cisco NAS), 354  
 Layer 2 designs  
   EtherChannel  
     *configuration best practices*, 43  
     LACP, 44  
     PAGP, 43-44  
   Layer 2 FlexLink designs, 203-206  
   Layer 2 in-band design (Cisco NAC Appliance), 355-357  
   Layer 2 out-of-band design (Cisco NAC Appliance), 358-359  
 STP  
   *Cisco STP Toolkit*, 37  
   *configuration best practices*, 36-37  
   *standards and features*, 37

trunks  
   *configuring*, 39  
   DTP, 41-42  
   VTP, 40  
   UDLD, configuration best practices, 42  
 Layer 2 loop access layer model, 195-197  
 Layer 2 loop-free access layer model, 195, 199-200  
   loop-free U design, 201-203  
 Layer 2 switches  
   CGMP, 440  
   Cisco Catalyst switches, 440, 472  
   IGMP snooping, 439, 472  
   multicast switching, 438-440, 472  
 Layer 2 to Layer 3 boundary design, 59  
   access layer switches, daisy chaining, 63-64  
   asymmetric routing, 68-69  
   excessive redundancy, dealing with, 65  
   insufficient redundancy, dealing with, 66-68  
   Layer 2 distribution switch interconnection, 59  
   Layer 3 access to distribution interconnection, 60-63  
   Layer 3 distribution switch interconnection, 60  
 Layer 3 access designs, 208-210  
   comparing with Layer 2, 216-217  
   multicast source support, 209  
 Layer 3 access layer model, 195  
 Layer 3 client access mode (Cisco NAS), 354  
 Layer 3 designs, 45, 60  
   bandwidth management, 45-47  
   Cisco NAC Appliance  
     *in-band design*, 359-360  
     *out-of-band design*, 360-362

- first-hop redundancy, 53-54
    - FHRP convergence, optimizing*, 58
    - GLBP*, 55-56
    - HSRP preemption*, 54
  - link load balancing, 47-49
  - routing protocol design, 49-51
    - summarization*, 51-52
  - Layer 3 protocols, IGMP, 438
  - learn phase of Cisco IOS OER, 322
  - Leave Group message (IGMPv2), 437
  - lightweight APs, 487-489, 520
  - link load balancing, 47-49
  - load balancing multiple FWSMs, 341-343
    - community VLANs, 345
    - isolated PVLAN ports, 344
  - local network control block, 431
  - local scope IP multicast addresses, 431
  - logical interfaces (STP), 230-231
  - Loop Guard, 37-38
  - loop-free U designs, 200-203
  - loop-prevention mechanisms, 197
  - loss (RF), 480
- M**
- MAC addresses, 431
  - management interfaces, IDS/IPS
    - deployment options, 374
  - management plane, 34
  - managing
    - bandwidth with EtherChannel, 45-47
    - IPs/IDSs, 374-376
    - VPNs, 410-412
  - MANs (metropolitan area networks),
    - SAN extension, 266-268
  - mapping agents, 457
  - MARS (Cisco Security Monitoring, Analysis, and Response System), 411
  - measure phase of Cisco IOS OER, 322-323
  - measurements (IP SLAs)
    - deploying, 564-566
    - hierarchical monitoring, 567
  - MED attribute (BGP), 331
  - medium-scale, dual-fabric collapsed core design, 263-264
  - mesh groups, 120-121
  - messages, syslog, 531-533
  - metrics for EIGRP fast convergence, 103
  - Metro Ethernet, 147-148
    - LAN services, 150
      - EMS*, 153
      - end-to-end QoS*, 154
      - EPL*, 151
      - ERMS*, 154
      - ERS*, 152-153
      - EWS*, 153
      - selecting*, 156
    - service model, 147
  - Min-Links (EtherChannel), 225
  - mirroring, 248
  - mobility, CUWN, 494
    - intercontroller roaming, 495-497, 521, 524
    - intracontroller roaming, 495, 521
    - mobility groups, 498-499
  - modular data center design
    - cabinet design, 221
    - versus 1RU design, 217-218
  - modularized processes in control plane, 35-36
  - monitoring
    - IDSs/IPs, 374-376
    - SLAs, 171

**MPLS VPNs, 162**

- backdoor routes, 164
- customer considerations, 163-164
- Layer 2/3 characteristics, 163
- managed router service, 164-165

**MSFC (Multilayer Switch Feature Card), placement of, 337****MST (Multiple Spanning Tree), 38****MTBF (mean time between failure), 170****MTTR (mean time to repair), 170, 557****multi-AS EIGRP systems, scaling, 104-108****multicast, 429**

- address assignment, 433
- advantages of, 429
- attacks, traffic forwarding, 465-466
- campus multicast protocols, 434
- disadvantages of, 429-430
- distribution trees, 441
- IGMP, 436, 472
- IGMPv2, 436-437
- IGMPv3, group membership, 437
- interdomain multicast protocols, 434
- IP addresses, 431-433
- Layer 2 addresses, 431-433
- PIM, 440-442
  - Bidir-PIM, 450, 452*
  - BSR, 459-460*
  - PIM-SM, 448-450, 470-471*
  - RPs, 454-456*
  - shared distribution trees, 445*
  - source trees, 443*
  - SSM, 452-453*
- replication requirements, 464
- security
  - goals, 461*
  - host receiver-side access control, 468*

*packet filter-based access control, 467*

*scoped addresses, 466*

state requirements, 462

**multicast groups, 425**

disabling for IPv6, 471

**multiplexing, 145**

CWDM, 143

DWDM, 144

WDM, 142

**N****NAA (Cisco NAC Appliance Agent), 348****NAC (Network Admission Control), 89****NAC Framework, 348, 362-364**

router platform support, 364-365

switch platform support, 366

**NACs**

Cisco NAC Appliance, 348

*Cisco NAMs, 351*

*Cisco NAS, 352, 354-355*

*clients security software applications, 366-368*

*components of, 349*

*Layer 2 in-band design, 355-357*

*Layer 2 out-of-band design, 358-359*

*Layer 3 in-band design, 359-360*

*Layer 3 out-of-band design, 360-362*

*policy updates, 350*

*process flow, 351*

Cisco NAC Framework, 348, 362-364

router platform support, 364-365

switch platform support, 366

**NAT (network address translation), 93**

client NAT, 293

with external partners, 92

source NAT, 293

NAT gateway mode (Cisco NAS), 353

NBAR (Network-Based Application Recognition), 546

and AutoQoS, 551, 554-556

and NetFlow, 548

packet inspection, 546-547

Protocol Discovery, 548-550

NetFlow, 534-535

cache management, 540

collectors, 545

deploying, 545

export versions, 542

Flexible NetFlow, 542-543

flow records, 537, 540

flows, 535

IP flows, 536

and NBAR, 548

network infrastructure layer (Cisco SONA), 5

network lifecycle approach, PPDIOO design phase, 13-16

network management, application optimization cycle, 528-529. **See also**

network management tools

network management tools

IP SLAs, 556

*measurements, 557-558, 564-567*

*measurements, deploying*

*operations, 560-561*

*SNMP features, 563-564*

*source/responder components, 560*

*timestamps, 562*

NBAR, 546

*and Cisco AutoQoS, 551, 554-556*

*and NetFlow, 548*

*packet inspection, 546-547*

*Protocol Discovery, 548-550*

NetFlow, 534-535

*cache management, 540*

*collectors, 545*

*deploying, 545*

*export versions, 542*

*Flexible NetFlow, 542-543*

*flow records, 537, 540*

*flows, 535-536*

syslog, 529

*messages, 531-533*

NICs

requirements per data center server, 223

teaming configurations, 234

noise, 506

non-real-time applications, 428

NSF (Cisco Nonstop Forwarding), 237

NSSAs (not-so-stubby areas), 95

## O

OC (Optical Carrier) rates, 141

OER (Optimized Edge Routing), 320-322

apply policy phase, 324

learn phase, 322

measure phase, 322-323

optimize phase, 324

topologies, 324-325

verify phase, 324

one-to-many multicast applications, 427

one-armed mode (SLBs), traffic flows, 292

on misconfigured networks, 293

on properly configured networks, 294

one-armed SLB two-firewall e-commerce module design, 305, 308

one-armed SLB with CSS e-commerce design, 313

- one-armed SLB with firewall contexts
  - e-commerce design, 308, 310, 313
- operating modes (Cisco NAS), 354
- operation phase of PPDIIO network lifecycle approach, 12
- operations (IP SLAs), 560
- optimize phase of Cisco IOS OER, 324
- optimize phase of PPDIIO network lifecycle approach, 12
- originating default routes, 94-96
- OSPF (Open Shortest Path First)
  - ABRs, connecting, 117-118
  - areas
    - designing*, 110-111
    - filtering*, 118-119
  - data center model, design recommendations, 182
  - fast convergence, 121-122
    - BFD*, 123-124
    - iSPF*, 122-123
  - Flooding Reduction, 121
  - full-mesh topology, 120-121
  - hub-and-spoke design, 113-116
  - Layer 3 access to distribution interconnection, 63
  - mesh groups, 120-121
  - required hierarchy, 112-113
  - route summarization, 113
  - scalability, factors affecting, 108
    - adjacencies*, 109-110
    - routing information*, 110
  - stub areas, 95-96
  - transit traffic, 97
- out-of-band traffic flow deployment model (Cisco NAS), 354
- oversubscription
  - calculating per uplink on access layer switches, 224
  - managing, 45

## P

- packet filter-based access control, 467
- packet inspection, 546-547
- PAgP (Port Aggregation Protocol), 43-44
- Paquet, Catherine, 316
- passive interfaces, configuring on access switches, 51
- PBR (policy-based routing), 341
- PDLs (Packet Description Language Modules), 547
- PDM (Cisco PIX Device Manager), 411
- people component of high availability in e-commerce designs, 279-280
- performance
  - of Cisco ASA 5500 series, 395-397
  - of Cisco ISRs, 395
  - of site-to-site VPNs, 393-395
- physical deployment models (Cisco NAS), 355
- PIM (Protocol Independent Multicast).
  - See also** PIM-SM
  - deployment models
    - ASM*, 447
    - Bidir-PIM*, 450-452
    - PIM-SM*, 448-450
    - SSM*, 452-453
  - shared distribution trees, 445
  - source trees, 443
- PIM-SM, 448-452
  - BSR, 459-460
  - RPs, 454-456
  - source control, 470-471
- placement
  - of MSFC, 337
  - of VPN devices, 397
    - integrating with firewall*, 399
    - placing on firewall DMZ*, 398-399
    - placing parallel to firewall*, 398



plan phase of PPDIIO network lifecycle approach, 11

PoE (Power over Ethernet), requirements for IP telephony in enterprise campus networks, 71

point-of-sale WANs, 414

point-to-multipoint OSPF hub-and-network design, 116

point-to-point OSPF hub-and-spoke design, 116

policy updates (Cisco NAC Appliance), 350

Port Bandwidth Reservation, 265

port density requirements for SAN design, 258-259

port-channel min-links command, 225

PortFast, 37

power budget planning for IP telephony in enterprise campus networks, 72-73

PPDIIO, 10

benefits of, 12-13

design phase, 13

*characterizing the existing network, 15*

*dividing network into areas, 16*

*identifying customer requirements, 14*

*top-down design practices, 16*

implement phase, 12

operation phase, 12

optimize phase, 12

plan phase, 11

prepare phase, 11

PPS (packets per second)

calculating, 413-414, 417

estimating based on branch profile, 415

prepare phase of PPDIIO network lifecycle approach, 11

preventing unicast flooding, 69

process flow with Cisco NAC Appliance, 351

process restart command, 35

processes in e-commerce designs, 280

promiscuous PVLAN ports, 344

Protocol Discovery (NBAR), 548-550

“pull” model, 447

PVLANS (private VLANs)

isolated ports, 344

security, 342-344

PVST+ (Per VLAN Spanning Tree), 37

## Q

QinQ tunneling, 153

QoS (quality of service)

end-to-end, 501

in enterprise campus networks, 75-77

IEEE 802.11e, 500

query-interval response time (IGMPv2), 437

## R

RAID (redundant array of inexpensive disks), 248

Real IP gateway mode (Cisco NAS), 353

real-time applications, 428

redesigning IP addressing, 88

redistribution, 99

EIGRP, filtering with route tags, 105

filtered redistribution, 99-100

migrating between routing protocols, 101

redundancy

in e-commerce designs, 278

excessive, dealing with, 65

first-hop redundancy, 53-54

*FHRP convergence, optimizing, 58*

*GLBP, 55-56*

*HSRP preemption, 54*

- implementing in Cisco Enterprise Campus Architecture, 28-30
- insufficient, dealing with, 66-68
- regulatory encryption, deploying site-to-site VPNs, 392
- remote-access VPNs
  - Cisco Easy VPN, 384
  - designing, 387
    - access control*, 389
    - address assignment considerations*, 388
    - authentication*, 389
    - routing considerations*, 388
  - SSL VPN, 384-385
    - clientless access*, 385-386
    - thick clients*, 386-387
    - thin clients*, 386
- replication requirements for unicast/multicast, 464
- report suppression, 436
- responder component (IP SLAs), 560
- RF (radio frequencies), 480, 520
- RFC 1112, Host Extensions for IP Multicasting, 436
- RFC 4761, 157
- RFC 4762, 157
- RHI (route health injection), 299
- ROADM (reconfigurable OADM), 145
- roaming
  - enhanced neighbor lists, 499
  - intercontroller roaming, 495-497, 521, 524
  - intracontroller roaming, 495, 521
- role-based addressing, 90-91
- Root Guard, 37-38
- route filtering, 96-97
- route reflectors
  - iBGP, scaling, 126-128
  - versus confederations, 133

- route summarization, 87, 93-94
  - at distribution layer, 51-52
  - bit splitting, 90-91
  - default routes, 95-96
  - originating default routes, 94-95
  - OSPF, 113
  - stub areas, 95-96
  - summary address blocks, 88
- route tags, EIGRP
  - redistribution, filtering, 105
  - routing updates, filtering, 105-107
- routed mode (firewalls), 333
- router mode (SLBs), 290
- routing
  - in base e-commerce module design, 301
  - implications for VLPS, 161
- routing flaps, 32
- routing information, effect on OSPF scalability, 110
- routing protocols, 49-51
  - IPsec VPN considerations, 417
  - migrating between, 101
  - summarization, 51-52
- routing updates (EIGRP), filtering with route tags, 105, 107
- RPs (rendezvous points), 442, 454-459
- RPVST+ (Rapid PVST+), 38, 59
- RRI (Reverse Route Injection), 388
- RSSI (Received Signal Strength Indicator), 511
- RSTP (Rapid Spanning Tree Protocol), 38
- RTTMON (Cisco Round-Trip Time Monitor), 563

**S**

- SACK (selective acknowledgment), 252
- SAN consolidation, 261

**SAN extension, 266-268**

- developments in, 270
- high availability, 271

**SAN islands, 253****SANs (storage area networks), 245**

- business drivers, 246
- collapsed-core design
  - large-scale, dual-fabric core-edge design, 264-265*
  - medium-scale, dual-fabric collapsed core design, 263-264*
  - single-switch collapsed core design, 262*
  - small-scale, dual-fabric collapsed core design, 263*

- designing, 258-260

- fabric, 245

- Fibre Channel, 251-252

- hard disk interfaces, 247

- HBA, 246

- IVR, 254

- SCSI, 250-251

- security, 261

- simplified management, 262

- storage

- arrays, RAID, 248*

- devices, 247*

- topologies, 249-250*

- VSANs, 253

**SANTap, 256, 258****SAs (security associations), 400****scaling**

- BGP, 124

- full-mesh iBGP scalability, 125-130, 133*

- Cisco NAS, 351

- data centers, 228

- aggregation layer, 185-186*

- Cisco ACE modules, 227*

- server NIC density, 223*

- service layer switches, 226*

- EIGRP, 102-108

- IP SLA deployments, 566

- iSCSI, 270

- OSPF, 108-109

- adjacencies, 109-110*

- fast convergence, 110, 121-124*

- site-to-site VPNs, 393*

- STP

- best practices, 233*

- in data centers, 229-232*

- VLPS, 159-160

- VPNs, 412-414

- scoped addresses, 466

- SCSI (small computer systems interface), 250

- daisy-chaining, 251

- initiator/target communication, 251

- zoning, 255*

- SDH, 141

- SDM (Security Device Manager), Easy VPN Server Wizard, 401-402

- SDP (Session Description Protocol), 429

- security

- Cisco ASA 5500 series performance, 395, 397

- Cisco Catalyst Integrated Security

- DAI, 79*

- DHCP snooping, 79*

- IP Source Guard, 79-80*

- Port Security, 78*

- Cisco ISR performance, 395

- comprehensive SAN security solution, 261

- firewalls. *See also* firewalls

- active/active topology, 338-339*

- asymmetric routing*, 339-341
- modes*, 333-334
- multiple FWSMs*, 341-345
- virtual firewalls*, 335-337
- IDSs, 368
  - design considerations*, 371
  - security management and monitoring infrastructure*, 370
  - sensors*, 369
- IPs, 368
  - deployment options*, 372-374
  - design considerations*, 371
  - HIPs, 370
  - security management and monitoring infrastructure*, 370
  - sensors*, 369
- in multicast environments, 461
  - attack traffic forwarding*, 465-466
  - host receiver-side access control*, 468
  - packet filter-based access control*, 467
  - scoped addresses*, 466
- NACs, client security software applications, 366-368
- VoWLANs, 502-503
- selecting
  - advanced WAN services, 166
  - Metro Ethernet LAN service, 156
  - network management application for IP SLA deployment, 568-569
- semi-directional antennas (WLAN), 480
- servers as application gateway, 282-284
- service context, 189
- service layer switches, 226
- Session Directory, 429
- shared distribution trees, 445
- show auto discovery qos command, 554
- show interfaces fastethernet include rate command, 417
- show spanning-tree summary total command, 231
- simplified SAN management, 262
- simulcasts, 427
- single points of failure, avoiding in Cisco Enterprise Campus Architecture, 30
- single-switch collapsed core design, 262
- site surveys, conducting, 513-514, 517, 519
- site-to-site VPNs
  - designing, 391-3921
  - IP addressing considerations, 392-393
  - performance, 393-395
  - scaling, 393
  - topologies, 397
- six-phase network lifecycle, 3, 10-13
- SLAs (service level agreements), 170, 556
  - IPM support, 568
  - measurements, 557-558
    - deploying*, 564-566
    - hierarchical monitoring*, 567
  - monitoring, 171
  - operations, 560-561
  - SNMP features, 563-564
  - source/responder components, 560
  - timestamps, 562
- small cell VoWLANs, AP placement, 515
- small-scale, dual-fabric collapsed core design, 263
- SNMP features (IP SLAs), 563-564
- SNR (signal-to-noise ratio), 506-507
- software-based zoning, 256

- SONA (Cisco Service-Oriented Network Architecture), 5
  - application layer, 10
  - benefits of, 6
  - Cisco Enterprise Architecture, 7
  - Hierarchical Network Model, 3
    - access layer*, 3
    - core layer*, 4
    - distribution layer*, 4
    - example of*, 4
  - infrastructure services, 9-10
- SONET (synchronous optical networking), 141-142
- source NAT, 92, 293
- source trees, 443
- source/responder components (IP SLAs), 560
- SOX (Sarbanes-Oxley Act), 392
- spanning-tree cost interface command, 56
- spectrum analysis tool, 514
- SPF (Shortest Path First) algorithm, 122
- SRP (Cisco Spatial Reuse Protocol), 146
- SSC (Cisco Secure Services Client), 367
- SSIDs (Service Set Identifiers), 488
- SSL VPNs, 384-385
  - clientless access, 385-386
  - thick clients, 386-387
  - thin clients, 386
- SSM (Source Specific Multicast), 452-453
- SSO (Stateful Switchover), 278
- stability requirements for SAN design, 260
- StackWise technology, 24, 64
- state requirements for unicast/multicast, 462
- stateful failover with common external prefix, 296
- static RP addressing, 455
- statistical multiplexing, 150
- Steps to Success** partner program, 519
- storage arrays, 247-248
- storage devices, connecting to host CPUs, 247
- storage topologies
  - DAS, 249
  - NAS, 249-250
- STP (Spanning Tree Protocol)
  - Cisco STP Toolkit, 37
  - configuration best practices, 36-37
  - data center model design
    - recommendations, 186, 228
    - logical interfaces*, 230-231
    - scaling*, 229-230
  - scaling, 231-233
  - standards and features, 37
- striping, 248
- stub areas, 95-96
- subnets
  - ACL-friendly addressing, implementing, 89
  - IP addressing, redesigning, 88
- subnetting, bit splitting, 91
- summarization, 93
- summary address blocks, 88-89
- supported routers on NAC Framework, 364-365
- supported switches on NAC Framework, 366
- SVTIs (static VTIs), 408
- switches, daisy chaining, 63-64
- switches supporting CGMP and IGMP snooping, 439
- syslog, 529-533

**T**

tape acceleration, 270  
 targets, 251, 255  
 TCAs (target channel adapters), 214  
 Teare, Diane, 316  
 technology in e-commerce designs, 278  
 teleagent WANs, 414  
 teleworker WANs, 414  
 testing e-commerce designs, 313-315  
 thick clients, 386-387  
 thin clients, 386  
 timestamps (IP SLAs), 562  
 TLS (Verizon Transparent LAN Services), 154  
 TOE (TCP/IUP Offload Engine), 270  
 tool in e-commerce designs, 281  
 top-down design practices, 16  
 topologies
 

- firewall topologies, active/active, 338-341
- in Layer 2 loop-free access layer model, 200
  - loop-free U design*, 201-203
- in Layer 2 looped access layer design, 196-197
- SAN design requirements, 258-259
  - for site-to-site VPNs, 39

 traffic flows
 

- in base e-commerce module design, 301
- data center model, 182
- in one-armed SLB two-firewall layer design, 308
- in one-armed SLB with firewall contexts e-commerce design, 310-313
- in two-firewall layer e-commerce design, 305

 traffic management requirements for SAN design, 260

transit traffic, 96-97  
 transmit queue congestion, 76  
 transparent firewall mode, 287  
 transparent mode (firewalls), 333-334  
 transponders, 145  
 trunks, configuring
 

- best practices, 39
- DTP, 41-42
- VTP, 40

 tuning e-commerce designs, DNS-based site selection, 325-326  
 two-armed mode (SLBs), 293  
 two-firewall layer e-commerce design, 304-305

**U**

UDLD (Unidirectional Link Detection), 37-39, 42  
 UDP (User Datagram Protocol), 429
 

- jitter operation, 564

 unicast flooding, preventing, 69  
 unicast routing protocols, 440  
 UNII (Unlicensed National Information Infrastructure) band, 480  
 uplink failure, handling, 66  
 UplinkFast, 37-38  
 UWN (Unified Wireless Networks).  
**See** Cisco UWN

**V**

verify phase of Cisco IOS OER, 324  
 virtual firewalls, 335
 

- firewall contexts, 336-337
- layers, 285
- MSFC placement, 337

 virtual gateway mode (Cisco NAS), 353  
 VLAN hopping, 40  
 VLANs, community VLANs, 345

## VoWLANs (Voice over Wireless LANs)

- alternative cell deployment, AP placement, 516
  - authentication and encryption recommendations, 502-503
  - clients, Cisco Unified Wireless IP Phone 7921G, 504
  - coverage considerations
    - nonoverlapping channels*, 508-511
    - SNR*, 506-507
  - in enterprise networks, 491
    - voice impact on WLANs*, 493
    - voice-ready architecture*, 492
  - general recommendations for, 511-512
  - large cell deployment, AP placement, 515
  - security, 503
  - site surveys, 513
    - conducting*, 517-519
    - spectrum analysis tool*, 514
    - WCS planning tool*, 514
  - small cell deployment, AP placement, 515
  - Steps to Success* partner program, 519
- ## VPLS (Virtual Private LAN Service), 156
- availability, 162
  - in enterprise networks, 158
  - H-VPLS, 159
  - IP multicast, 162
  - QoS issues, 161
  - routing implications, 161
  - scaling, 159-160
- ## VPNs
- clients, designing IP address schemes, 91

- device placement, 397
    - integrating with firewall*, 399
    - placing parallel on firewall DMZ*, 398-399
    - placing parallel to firewall*, 398
  - IPsec VPNs, 400
    - Cisco Easy VPN*, 401
    - DMVPN*, 405-407
    - GET VPNs*, 409-410
    - GRE over IPsec*, 403-405
    - routing protocol considerations*, 417
    - VTI*, 407-408
  - managing, 410-412
  - remote-access
    - Cisco Easy VPN*, 384
    - designing*, 383-384, 387-389
    - SSL VPNs*, 384-387
  - scaling, 412-414
  - site-to-site, designing, 391-397
- ## VRF (Virtual Routing and Forwarding), data center model, 193
- ## VRRP (Virtual Router Redundancy Protocol), 269
- ## VSAN (virtual SANs), 253
- SAN consolidation, 261
- ## VSS (Virtual Switching System), 25
- ## VTIs (virtual tunnel interfaces), 407-408
- ## VTP (VLAN Trunking Protocol), 40

**W****WANs**

- advanced services, selecting, 166
- business risk assessment, 167
- enterprise optical connections, 141-142
- Metro Ethernet, 147-148
  - EMS, 153*
  - end-to-end QoS, 154*
  - EPL, 151*
  - ERMS, 154*
  - ERS, 152-153*
  - EWS, 153*
  - LAN services, 150-156*
  - service model, 147*
- requirements, 168
- RPR, 146
- SAN extension, 266-270
- VPLS, 156
  - and IP multicast, 162*
  - availability, 162*
  - H-VPLS, 159*
  - in enterprise networks, 158*
  - QoS issues, 161*
  - routing implications, 161*
  - scaling, 159-160*

**WCS (Wireless Control System)**

- planning tool, 514

**WDM (wavelength-division multiplexing), 142****wireless networks, 479**

- antennas, 480-483
- APs. *See* APs
- CUWN. *See* Cisco UWN
- IEEE 802.11 operational standards, 481-483
- RF, 480

**WLANs (Wireless Local Area Networks), 479****WLC (Wireless LAN Controllers), 487, 498-499****X-Y-Z****zone-based policy firewalls, 346-347****zoning, 255**