

**SWITCH**

---

# Implementing Cisco Switched Networks

---

## **Volume 2**

Version 1.0

## **Student Guide**

Text Part Number: 97-2835-01




**Americas Headquarters**  
Cisco Systems, Inc.  
San Jose, CA

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

**Europe Headquarters**  
Cisco Systems International BV  
Amsterdam, The Netherlands

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

 CCDE, CCENT, CCSI, Cisco Eee, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco.Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLNNX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0910R)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED "AS IS." CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

# Table of Contents

## Volume 2

<b><i>Implementing a Highly Available Network</i></b>	<b>5-1</b>
Overview	5-1
Module Objectives	5-1
<b>Understanding High Availability</b>	<b>5-3</b>
Overview	5-3
Objectives	5-3
Components of High Availability	5-4
Redundancy	5-6
Technology	5-7
People	5-8
Processes	5-10
Tools	5-12
Resiliency for High Availability	5-13
Network-Level Resiliency	5-14
High Availability and Failover Times	5-15
Optimal Redundancy	5-16
Provide Alternate Paths	5-17
Avoid Too Much Redundancy	5-18
Avoid Single Point of Failure	5-19
Cisco NSF with SSO	5-20
Routing Protocols and NSF	5-21
Summary	5-22
<b>Implementing High Availability</b>	<b>5-23</b>
Overview	5-23
Objectives	5-23
Distributed VLANs on Access Switches	5-24
Local VLANs on Access Switches	5-25
Layer 3 Access to Distribution Interconnection	5-26
Daisy-Chaining Access Layer Switches	5-27
Daisy-Chaining Access Layer Switch Issues	5-28
Cisco StackWise Access Switches	5-29
Too Little Redundancy	5-30
Impact of Uplink Failure	5-31
Summary	5-32
<b>Implementing Network Monitoring</b>	<b>5-33</b>
Overview	5-33
Objectives	5-33
Network Monitoring	5-34
Syslog	5-35
Syslog Features	5-36
Syslog Message Format	5-37
Syslog Message Log	5-38
Example: Syslog Messages	5-38
System Log Configuration	5-39
SNMP Overview	5-41
SNMPv1 vs. SNMPv2	5-42
About SNMPv3	5-44
SNMP Recommendations	5-45
SNMP Configuration	5-46
IP SLA Configuration	5-47
IP SLA Measurements	5-48
IP SLA Operations	5-49
IP SLA Source and Responder	5-50
IP SLA Operations with Responder	5-51
IP SLA Responder Time Stamps	5-53

IP SLA Configuration	5-54
IP SLA Verification	5-55
Summary	5-57
<b>Lab 5-1 Debrief</b>	<b>5-59</b>
Overview	5-59
Objectives	5-59
Review and Verification	5-60
Module Summary	5-63
Module Self-Check	5-65
Module Self-Check Answer Key	5-67
<b><i>Implementing Layer 3 High Availability</i></b>	<b>6-1</b>
Overview	6-1
Module Objectives	6-1
<b>Configuring Layer 3 Redundancy with HSRP</b>	<b>6-3</b>
Overview	6-3
Objectives	6-3
Describing Routing Issues	6-4
Using Proxy ARP	6-4
Using Default Gateways	6-5
Identifying the Router Redundancy Process	6-6
Router Redundancy Failover	6-7
HSRP Configuration	6-8
Virtual Router MAC Address	6-10
Forwarding Through Active Router	6-11
Virtual Router	6-11
Active Router	6-11
HSRP Active and Standby Router Interaction	6-12
HSRP State Table	6-13
HSRP State Transition	6-14
HSRP Priority and Pre-emption	6-15
HSRP Active Router and Spanning-Tree Topology	6-17
HSRP Authentication	6-18
HSRP and Timer Considerations	6-19
HSRP Timer Configuration	6-20
Subsecond Failover	6-21
Pre-empt Time Aligned with Router Boot Time	6-21
HSRP Versions	6-22
Displaying the Standby Status	6-23
HSRP Interface Tracking	6-24
HSRP and Tracking	6-26
HSRP and IP SLA Tracking	6-27
Multiple HSRP Groups	6-28
Multiple HSRP Group Configuration	6-29
HSRP Monitoring	6-30
Summary	6-32



<b>Configuring Layer 3 Redundancy with VRRP and GLBP</b>	<b>6-33</b>
Overview	6-33
Objectives	6-33
Virtual Router Redundancy Protocol	6-34
About the VRRP Operations Process	6-37
VRRP Transition Process	6-38
Configuring VRRP	6-39
Describing GLBP	6-41
GLBP Functions	6-42
GLBP Features	6-42
Identifying the GLBP Operations Process	6-44
GLBP Operation	6-45
GLBP Interface Tracking	6-47
GLBP Implementation	6-51
GLBP and VLAN Spanning	6-53
Summary	6-54
<b>Lab 6-1 Debrief</b>	<b>6-55</b>
Overview	6-55
Objectives	6-55
Review and Verification	6-56
<b>Lab 6-2 Debrief</b>	<b>6-59</b>
Overview	6-59
Objectives	6-59
Review and Verification	6-60
Module Summary	6-63
Module Self-Check	6-65
Module Self-Check Answer Key	6-67
<b><i>Minimizing Service Loss and Data Theft in a Campus Network</i></b>	<b>7-1</b>
Overview	7-1
Objectives	7-1
<b>Understanding Switch Security Issues</b>	<b>7-3</b>
Overview	7-3
Objectives	7-3
Overview of Switch Security Issues	7-4
Security Infrastructure Services	7-5
Reason for Internal Security	7-7
Unauthorized Access by Rogue Devices	7-8
Switch Attack Categories	7-9
MAC Flooding Attack	7-11
Suggested Mitigation for MAC Flooding Attacks	7-12
Port Security	7-13
Configure Port Security	7-15
Caveats to Port Security Configuration Steps	7-16
Verifying Port Security	7-17
Port Security with Sticky MAC Addresses	7-19
Authentication and Authorization Methods	7-20
Local User Authentication with AAA	7-21
802.1X Port-Based Authentication	7-23
Configuring 802.1X	7-25
Summary	7-26

<b>Protecting Against VLAN Attacks</b>	<b>7-27</b>
Overview	7-27
Objectives	7-27
Explaining VLAN Hopping	7-28
Switch Spoofing	7-28
VLAN Hopping With Double Tagging	7-30
Mitigating VLAN Hopping	7-31
VLAN Access Control Lists	7-32
Configure VACLs	7-33
Summary	7-35
<b>Protecting Against Spoofing Attacks</b>	<b>7-37</b>
Overview	7-37
Objectives	7-37
DHCP Spoofing Attacks	7-38
DHCP	7-40
DHCP Snooping	7-41
Configure DHCP Snooping	7-43
DHCP Snooping Verification	7-45
ARP Poisoning	7-46
Dynamic ARP Inspection	7-48
Dynamic ARP Inspection Configuration	7-50
IP Source Guard	7-51
IP Source Guard Configuration	7-54
Summary	7-56
<b>Securing Network Services</b>	<b>7-57</b>
Overview	7-57
Objectives	7-57
Vulnerabilities in Cisco Discovery Protocol	7-58
Neighbor Discovery Protocols	7-59
Cisco Discovery Protocol Configuration	7-60
LLDP Configuration	7-61
Cisco Discovery Protocol Vulnerabilities	7-62
Telnet Vulnerabilities	7-63
About SSH	7-64
Configuration of SSH	7-65
VTY ACLs	7-66
HTTP Secure Server Configuration	7-67
Switch Security Considerations	7-68
Organizational Security Policies	7-69
Secure Switch Devices	7-69
Switch Security Recommendations	7-71
Mitigating Compromises Launched Through a Switch	7-72
Summary	7-73
<b>Lab 7-1 Debrief</b>	<b>7-75</b>
Overview	7-75
Objectives	7-75
Review and Verification	7-76
Module Summary	7-79
Module Self-Check	7-81
Module Self-Check Answer Key	7-84

<b><i>Accommodating Voice and Video in Campus Networks</i></b>	<b>8-1</b>
Overview	8-1
Module Objectives	8-1
<b>Planning for Support of Voice in a Campus Network</b>	<b>8-3</b>
Overview	8-3
Objectives	8-3
Unified Communications	8-4
IP Telephony Components	8-5
Characteristics of Voice and Data	8-6
Video Applications	8-8
Voice and Video Traffic	8-9
Requirements for Voice, Data, and Video Traffic	8-10
Voice and Video in the Campus Network	8-11
Summary	8-12
<b>Integrating and Verifying VoIP in a Campus Infrastructure</b>	<b>8-13</b>
Overview	8-13
Objectives	8-13
Planning for VoIP Requirements	8-14
Integrating VoIP in the Campus	8-15
Voice VLANs	8-16
IP Telephony Extends the Network Edge	8-17
Multi-VLAN Access Port	8-18
Configuring and Verifying Voice VLANs	8-19
Power over Ethernet	8-20
Power over Ethernet 802.3af	8-22
New Power over Ethernet Developments	8-23
PoE Configuration	8-24
Switch Power Budget	8-25
PoE Verification	8-26
Adding Upper-Layer Services	8-27
Test Plan	8-28
Summary	8-29
<b>Working with Specialists to Accommodate Voice and Video on Campus Switches</b>	<b>8-31</b>
Overview	8-31
Objectives	8-31
High Availability for VoIP and Video	8-32
Example: Cisco Reliability and Availability	8-33
Building Voice/Video/Data Campus Networks	8-34
Quality of Service	8-35
Recommended Practices for QoS	8-37
QoS and Time-Sensitive Traffic	8-38
LAN-Based Classification and Marking	8-39
Layer 2 QoS Marking	8-40
Layer 3 QoS Marking	8-41
Describing QoS Trust Boundaries	8-42
Cisco Phone Connected to a Switch	8-43
Configuring QoS for Voice VLANs	8-44
Configuring Cisco AutoQoS	8-45
Cisco AutoQoS Configuration	8-46
Monitoring Cisco AutoQoS	8-47
Summary	8-48

<b>Lab 8-1 Debrief</b>	<b>8-49</b>
Overview	8-49
Objectives	8-49
Review and Verification	8-50
Module Summary	8-53
Module Self-Check	8-55
Module Self-Check Answer Key	8-58
<b><i>Integrating Wireless LAN into a Campus Network</i></b>	<b>9-1</b>
Overview	9-1
Module Objectives	9-1
<b>Comparing WLANs to Campus Networks</b>	<b>9-3</b>
Overview	9-3
Objectives	9-3
WLAN Overview	9-4
Wireless LAN	9-6
Wired and Wireless LAN	9-7
Similarities Between WLANs and LANs	9-8
Differences Between WLANs and LANs	9-9
Summary of WLAN and LAN Differences	9-10
WLAN Access Point Topology	9-11
About SSIDs	9-13
SSID and VLAN Support	9-14
Client Roaming	9-15
Layer 2 and Layer 3 Roaming	9-16
Security on WLAN and LAN	9-17
Summary	9-18
<b>Assessing the Impact of WLANs on Campus Networks</b>	<b>9-19</b>
Overview	9-19
Objectives	9-19
WLAN Implementations	9-20
Standalone WLAN Solution	9-21
Traffic Flow Between Clients—Standalone WLAN Solution	9-22
Controller-Based WLAN Solution	9-23
Controller-Based WLAN Solution	9-24
Traffic Flow Between Clients—Controller-Based WLAN Solution	9-26
Hybrid Remote Edge Access Points	9-27
Comparison of the WLAN Solutions	9-28
About WLCs	9-29
Connection of the Standalone Solution to the Network	9-30
SSIDs, VLANs, and Trunks in the Standalone Solution	9-31
Connection of the Controller-Based Solution to the Network	9-32
SSIDs, VLANs, and Trunks in the Controller-Based Solution	9-33
Controller-Based AP Protocol	9-35
Cisco WLC Ports and Protocols	9-36
Summary	9-37

<b>Preparing the Campus Infrastructure for WLANs</b>	<b>9-39</b>
Overview	9-39
Objectives	9-39
Access Point and Controller Placement	9-40
Distributed Controller Placement	9-41
Distributed WLC Design	9-41
Centralized Controller Placement	9-42
Centralized WLC Placement	9-42
WLAN Devices Connected to the LAN Switches	9-43
Configure Switches for WLAN Devices	9-44
Switch Configuration for Standalone APs and H-REAPs	9-45
Switch Configuration for Controller-Based APs	9-46
Switch Configuration for a WLC	9-47
Link Aggregation for 4400 Series Controllers	9-48
Link Aggregation	9-49
Switch Configuration for 4400 Series Controllers	9-50
Switch Configuration for Cisco WiSM Controllers	9-51
Switch Configuration for Cisco WiSM Controllers	9-52
Gathering Requirements	9-53
Planning the Integration	9-54
Creating a Test Plan	9-55
Summary	9-56
<b>Lab 9-1 Debrief</b>	<b>9-57</b>
Overview	9-57
Objectives	9-57
Review and Verification	9-58
Module Summary	9-61
Module Self-Check	9-63
Module Self-Check Answer Key	9-64



# Implementing a Highly Available Network

---

## Overview

A network with high availability provides an alternate means that allows constant access to all infrastructure paths and key servers. High availability is not only about adding redundant devices. It also implies planning to understand where the points of failure occur and to design the network so that an alternate solution exists to compensate for the loss of these points of failure.

## Module Objectives

Upon completing this module, you will be able to implement a highly available network. This ability includes being able to meet these objectives:

- Evaluate the uses, requirements, benefits, and performance expectations for high availability in a given enterprise network design
- Implement a high availability solution according to a given network design and requirements
- Construct implementation and verification plans to implement a highly available network solution, by monitoring the infrastructure resources that are affected, selecting the required tools and commands, and organizing the required tasks, for a provided Layer 3 network design and predefined infrastructure





## Lesson 1

---

# Understanding High Availability

---

## Overview

High availability is technology that enables networkwide resilience to increase IP network availability. Network applications must cross different network segments—from the enterprise backbone, enterprise edge, and service provider edge, through the service provider core. All segments must be resilient so that they recover quickly enough for faults to be transparent to users and to network applications. This lesson describes the high availability concept, and explains how resiliency is built and how the network is designed to always offer a path between any pair of endpoints.

## Objectives

Upon completing this lesson, you will be able to understand high availability. This ability includes being able to meet these objectives:

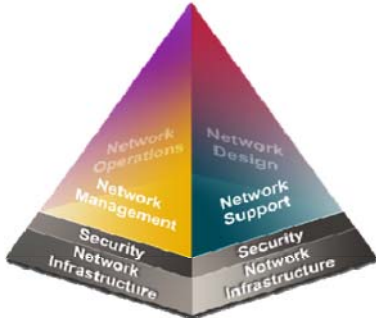
- Evaluate the uses, requirements, benefits, and performance expectations for high availability in a given enterprise network design
- Describe resiliency for high availability
- Design the network for optimal redundancy

# Components of High Availability

This topic describes high availability concept and components.

## Components of High Availability

- The objective of high availability is to prevent outages and minimize downtime.
- Achieving high availability integrates multiple components:
  - Redundancy
  - Technology
  - People
  - Processes
  - Tools
- The first two components are relatively easy to integrate.
- The last three components are usually where gaps lead to outages.



© 2009 Cisco Systems, Inc. All rights reserved. SWITCH-50

High availability is an organizational objective with the goal of preventing outages or at least minimizing downtime. Achieving high availability is hard work. It takes ongoing effort and iterated improvement. High availability is not something you have or do not have, but a skill that an organization achieves and perfects over time.

To start making progress on providing high availability requires integrating multiple components:

- Redundancy
- Technology (including hardware and software features)
- People
- Processes
- Tools

The network redundancy and technology components are relatively easy to accomplish because these elements can be purchased and deployed. A traditional network designer expects to be involved with these two aspects of high availability.

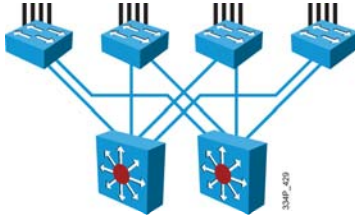
No matter how much and how well redundancy and technology are designed and deployed, high availability is not achieved unless the people component (sufficient labor pool with the right skills, training, and mindset), the process component (company expectations, change control process, and so on), and the tools component (network management, good documentation) are present. If any one of the last three high-availability components is missing, then incidents will happen and outages will occur. Initially, the network designer may not be able to fix the people, processes, and tools in an organization. Often it takes a consultant doing a post-outage design review to talk about these components and suggest changes.

# Redundancy

This subtopic describes redundancy as a component of high availability.

## Redundancy

- Redundancy is used to reduce or eliminate the effects of a failure.
- Design of redundancy attempts to eliminate single points of failure:
  - Avoid single causes of failure.
  - Use geographic diversity and path diversity.
  - Use dual devices and links.
  - Use dual WAN providers.
  - As appropriate, implement dual data centers.
  - As appropriate, use dual colocations, dual central office facilities, and dual power substations.
- Design of redundancy needs to trade off cost versus benefit:
  - Hours of downtime compared to the costs of redundancy, planning, etc.



© 2009 Cisco Systems, Inc. All rights reserved. RWT1111101-55

Redundancy designs attempt to eliminate single points of failure, where one failed device or design element brings down service.

A redundant design can use several mechanisms to prevent single points of failure:

- Geographic diversity and path diversity are often included.
- Dual devices and links are very common.
- Dual WAN providers are common.
- Dual data centers are sometimes used, especially for large companies and large e-commerce sites.
- Dual colocation facilities, dual phone central office facilities, and dual power substations can be implemented.

Redundant design must trade off cost versus benefit. It takes time to plan redundancy and to verify the geographic diversity of service providers. Additional links and equipment cost money to purchase and maintain. These options must be balanced against risks, costs of downtime, and so on. The time and money that are invested in redundancy designs must be spent where they will have the most impact. Consequently, redundancy is most frequently found in network, data center, or e-commerce module cores, and then in critical WAN links or Internet service provider (ISP) connections. Additional e-commerce module redundancy can double up elements in the path between users and applications, and the applications and back-end databases and mainframes.

# Technology

This subtopic describes technology as a component of high availability.

## Technology

- Cisco routing continuity options:
  - Cisco Nonstop Forwarding (NSF)
  - Stateful Switchover (SSO)
  - Catalyst 3750 Series Switches with Cisco StackWise technology
  - Catalyst 6500 VSS 1440
- Techniques for detecting failure and triggering failover:
  - Monitoring
  - IP SLAs and object tracking
- Other technologies:
  - Fast-routing convergence

© 2009 Cisco Systems, Inc. All rights reserved. SW111110-50

There are several Cisco routing continuity options, such as Cisco Nonstop Forwarding (NSF) with Stateful Switchover (SSO). Also, graceful restart capabilities improve availability. These technologies allow processor failover without a link flap, continued forwarding of packets, and maintenance of Border Gateway Protocol (BGP) adjacencies.

Techniques exist to detect failure and to trigger failover to a redundant device. These techniques include service monitoring for Cisco IOS IP service level agreements (SLAs) and object tracking.


Other technologies also contribute to high availability. For example, fast-routing convergence and server load balancers help maintain high availability. Firewall stateful failover can maintain user or application sessions across a firewall device failover.

# People

This subtopic describes people as components of high availability.

## People

- Staff work habits and skills can impact high availability.
  - Attention to detail.
  - Reliability and consistency.
- Good skills and ongoing technical training are needed:
  - Lab time working with technology, practical skills, troubleshooting challenging scenarios, etc.
  - Communication and documentation are important.
    - Define what other groups expect.
    - Define why the network is designed the way it is, how it is supposed to work.
- If people are not given the time to do the job right, they cut corners:
  - If the design target is just “adequate,” falling short leads to poor design.
- Staff team should align with services.
  - Owner and experts for each key service application and other components should be identified and included.



© 2009 Cisco Systems, Inc. All rights reserved.SWITCHING-55

Redundant equipment and links and advanced technology are just the beginning of high availability. In the Prepare, Plan, Design, Implement, Operate, and Optimize (PPDIOO) methodology, the people component is vitally important too. Staff work habits and skills can impact high availability. For example, attention to detail enhances high availability, while carelessness hurts availability. Reliable and consistent wiring and configurations are easier to manage and troubleshoot.

The level of staff skills and technical training are important elements of redundancy. Devices must be configured correctly. Lab testing is important in order to understand the circumstances that activate failover, and what failover will and will not accomplish. Thoroughness in lab testing often translates into less downtime in production. For example, nonstateful firewall failover may be adequate in terms of passing traffic. However, a practical understanding of the application can show that with nonstateful failover, application sessions will lock up for an extended period of time until an application timeout causes session reestablishment. Designs that include failover must be tested for the entire system, not just for individual components.

Good communication and documentation are also important. The network administrators need to communicate with other network, security, application, and server teams. The network documentation should cover why things are designed the way they are, and how the network is supposed to work. Failover behavior is complex enough that it is unwise to have to recapture failover logic and boundary conditions every time some part of the design changes.

Field experience leads to the observation that if people are not given time to do the job right, they will cut corners. Testing and documentation are often the first items to be eliminated. Lack of thorough testing and documentation will have long-term consequences on the ability to maintain, expand, and troubleshoot the network.

If the design target is just “adequate” coverage, falling short of that target can lead to a poor design. Designs should be better than adequate, to ensure that no part of the implementation or operation of the high-availability network is inadequate.


One other organizational recommendation is to align staff teams with services. If the corporate web page depends on staff who report to other managers, the manager of the e-commerce site may compete for staff time with the network engineering or operations manager. In most cases, the person who does the staff evaluation and provides the pay bonus generally gets most of the attention. This organizational structure can make it difficult to get routine testing or maintenance done for the e-commerce site if the staff does not report to the e-commerce manager. The owner or expert on key service applications and other components should be identified and included in design and redesign efforts.

# Processes

This subtopic describes processes as components of high availability.

## Processes

- Build repeatable processes.
  - Document change procedures, failover planning and lab testing, and implementation procedures.
- Use labs appropriately.
  - Lab equipment reflects the production network, failover mechanisms are tested and understood, and new code is validated before deployment.
- Use meaningful change controls.
  - Test all changes before deployment, use good planning with rollback plans, and conduct realistic and thorough risk analysis.
- Manage operation changes.
  - Perform regular capacity management audits, manage Cisco IOS versions, track design compliance as recommended practices change, and develop disaster recovery plans.



© 2009 Cisco Systems, Inc. All rights reserved. SWITCHING-55

Sound, repeatable processes can lead to high availability. Continual process improvement as part of the PPDIOO methodology plays a role in achieving high availability. Organizations need to build repeatable processes and gradually improve them. Tasks that are always implemented as a special one-time occurrence represent a lost opportunity to learn as an organization.

Organizations should build repeatable processes in the following ways:

- By documenting change procedures for repeated changes (for example, Cisco IOS Software upgrades)
- By documenting failover planning and lab testing procedures
- By documenting the network implementation procedure, so that the process can be revised and improved the next time that components are deployed

Organizations should use labs appropriately, as follows:

- Lab equipment should accurately reflect the production network.
- Failover mechanisms are tested and understood.
- New code is systematically validated before deployment.

Because staff members tend to ignore processes that consume a lot of time or that appear to be a waste of time, organizations also need meaningful change controls in the following ways:

- Test failover and all changes before deployment.
- Plan well, including planning rollbacks in detail.
- Conduct a realistic and thorough risk analysis.



The following management of operational changes is also important:


- Perform regular capacity management audits.
- Track and manage Cisco IOS versions.
- Track design compliance as recommended practices change.
- Develop plans for disaster recovery and continuity of operations.

# Tools

This subtopic describes tools that are used for high availability.

## Tools

- Monitor availability and key statistics for devices and links.
  - Use performance thresholds, Top N reporting, and trending to spot potential problems.
  - Monitor packet loss, latency, jitter, and drops.
- Good documentation is a powerful tool.
  - Maintain updated network diagrams.
  - Have network design write-ups.
  - Document key addresses, VLANs, and servers.
  - Tie services to applications, applications to virtual servers, and virtual servers to real server tables.



© 2009 Cisco Systems, Inc. All rights reserved. SWITCH-1701-50

Organizations are starting to monitor service and component availability. With proper failover, services should continue operating when single components fail. Without component monitoring, a failure to detect and replace a failed redundant component may lead to an outage when the second component subsequently fails.

Performance thresholds and reporting of the top N devices with specific characteristics (Top N reporting) are useful, both for noticing when capacity is running out, and also for correlating service slowness with stressed network or server resources. Monitoring packet loss, latency, jitter, and drops for WAN links or ISPs is also important. Those metrics can be the first indication of an outage or of a potential deterioration of a service level agreement (SLA) that could affect delivery of services.

Good documentation, such as the following, provides an extremely powerful set of tools:

- Network diagrams help in planning, and in fixing outages more quickly. Out-of-date documentation can lead to design errors, lack of redundancy, and other undesirable consequences.
- Documentation explaining how and why the network design evolved helps capture knowledge that can be critical when a different person needs to make design changes, reexamine how failover works, or make other changes.
- Key addresses, VLANs, and servers should be documented.
- Documentation tying services to applications and virtual and physical servers can be extremely useful when troubleshooting.

# Resiliency for High Availability

This topic describes resiliency as a component of high availability.

## Resiliency for High Availability

High availability is implemented with the following components:

### Network-level resiliency

- Redundant links
- Redundant devices

### System-level resiliency

- Integrated hardware resiliency
- Redundant power supply
- Stackable switches

### Management and monitoring

- Detection of failure

Supported features depend on switch family.

Network-level, and system-level resiliency and network monitoring are required components of high availability. High availability should be considered at every level of the network. In the context of this course, however, the focus is on network-level resiliency. You should still organize high availability at the system level. In a switched network, this means making sure that heavily solicited or key switches have redundant power supplies, or that duplicate devices are available to replace failed components of the network.

Another part of high availability is ensuring that you are informed when an element of your network fails. This is configured through monitoring and management features; the status or the failure of any element in your network should be immediately reported to a location, along with an immediate notification of the issue.

# Network-Level Resiliency

This subtopic describes high availability at the network level.

## Network-Level Resiliency

- Link redundancy
  - Redundant links
  - EtherChannel
- Fast convergence
  - Optimized link implementation
  - Tuning of Layer 2 and routing protocols
- Power redundancy
  - External redundant power supply
  - Uninterruptible power supply
- Monitoring
  - SNMP
  - Syslog
  - IP SLA
  - Time synchronization via NTP

© 2009 Cisco Systems, Inc. All rights reserved. SWITCHING-55

Network-level resiliency is built with device and link redundancy.

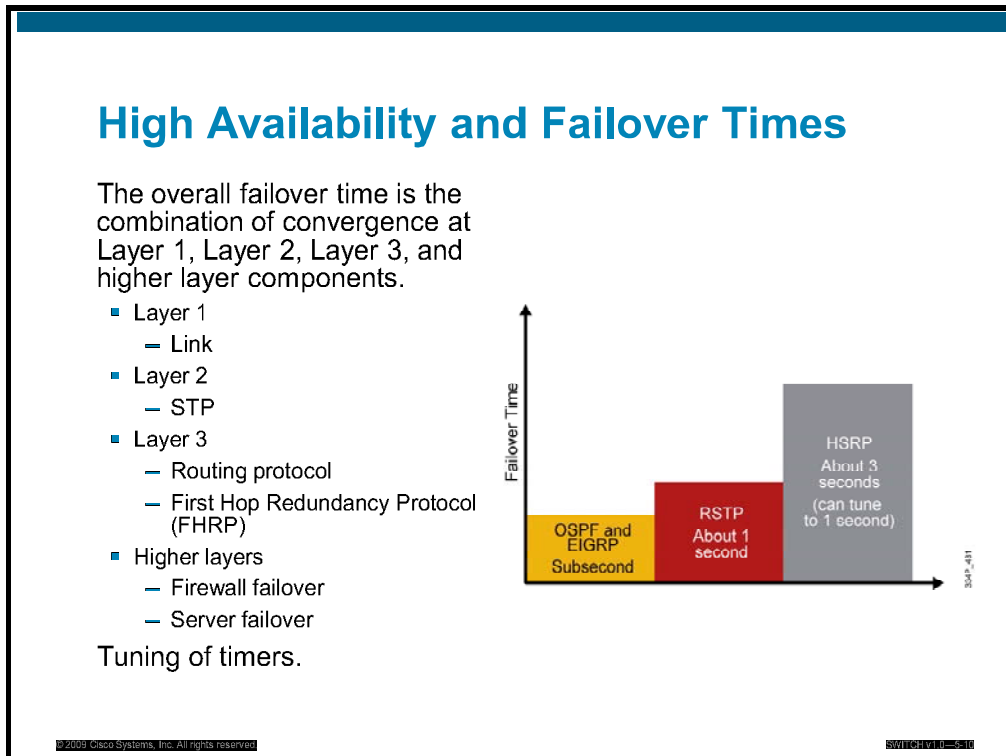
Link redundancy is an element that you have already configured throughout this course. When possible and necessary, duplicate links are installed between devices. If one physical link fails, the redundant one can hold the load while you replace the first one. These redundant links can be set in a standby mode, in which one link is active and the other one is blocked by Spanning Tree Protocol (STP), or in a load-balancing mode, with EtherChannel.

Another element of high availability at the network level is fast convergence. When a link fails, the redundant link or path should take precedence immediately, to avoid situations in which frames or packets are dropped due to slow convergence time. In this perspective, Rapid Spanning Tree Protocol (RSTP) is preferred over IEEE 802.1D STP. With the same logic in mind, fast convergence should apply to Layer 3 connections. Wherever possible, efficient routing protocols, such as the Enhanced Interior Gateway Routing Protocol (EIGRP), would be preferred to slower routing protocols like the Routing Information Protocol (RIP), to increase convergence speed.

Monitoring the various network elements involves several components. The first one is to synchronize time between interconnecting devices and the monitoring station. Knowing precisely when an event occurs is fundamental to managing failures and recoveries. The second element is to track events that are related to device status. You can do this using syslog and the Simple Network Management Protocol (SNMP). SNMP cannot monitor some elements. For example, your link to the next hop may be up, but a failure in the network renders your gateway unreachable. This event may be undetected by the local device-monitoring configuration. To circumvent this kind of issue, IP SLA is a protocol that is dedicated to testing connectivity between devices. It is an important addition for monitoring the network with increased accuracy.

# High Availability and Failover Times

This subtopic describes the impact of failover time on high availability.



The overall failover time in the data center is the combination of convergence at Layer 2, Layer 3, and Layer 4 components. The network components have different recovery times:

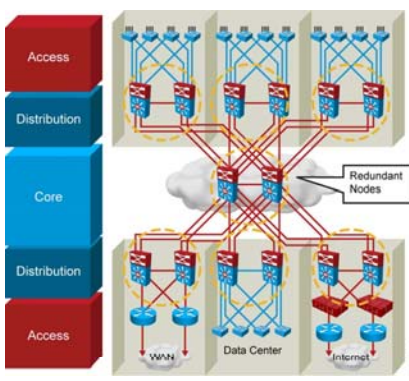
- Tuned routing protocols can fail over in less than 1 second. Open Shortest Path First (OSPF) and EIGRP can both achieve subsecond convergence time with recommended timer configurations.
- RSTP converges in about 1 second. RSTP permits subsecond convergence time for minor failures when logical ports are under watermarks, and can take 1 second to 2 seconds for major failure conditions.
- EtherChannel can fail over in about 1 second. When a link fails, Cisco EtherChannel technology redirects traffic from the failed link to the remaining links in less than 1 second.
- Default Hot Standby Router Protocol (HSRP) timers are 3 seconds for the hello time and 10 seconds for the hold time. A recommended practice is to configure the timers with a hello time of 1 second and a hold time of 3 seconds so that convergence occurs in less than 3 seconds. You can adjust the convergence time down to subsecond values, but you must consider the CPU load.
- Stateful service modules typically fail over within 3 to 5 seconds. The convergence time for Cisco Catalyst 6500 Series Firewall Services Module (FWSM) is about 5 seconds with recommended timers, and the Caching Services Module (CSM) is about 5 seconds with recommended timers. The Cisco Application Control Engine (ACE) can achieve failovers in about 1 second with its active/active configuration.
- The least tolerant TCP/IP stacks are the Windows Server and Windows XP client stacks, which have about a 9-second tolerance. Each of the TCP/IP stacks that are built into the various operating systems have a different level of tolerance for determining when a TCP session will drop. Other TCP/IP stacks, such as those found in Linux, Hewlett-Packard (HP), and IBM systems, are more tolerant and have a longer window before tearing down a TCP session.

# Optimal Redundancy

This topic describes how to implement optimal redundancy in a switched network.

## Optimal Redundancy

- Core and distribution have redundant switches and links.
- Access switches have redundant links.
- Network bandwidth and capacity can withstand single switch or link failure.
- Network bandwidth and capacity support 200–500 ms to converge around most events.



The diagram illustrates a multi-tier network architecture. On the left, a vertical stack of colored boxes represents the layers: Access (red), Distribution (blue), Core (blue), Distribution (blue), and Access (red). To the right, a detailed network topology shows multiple switches and routers interconnected. Red lines represent links between switches. A callout box labeled 'Redundant Nodes' points to a specific switch in the Core layer. The network is connected to external entities labeled 'WAN', 'Data Center', and 'Internet' at the bottom. The diagram demonstrates a highly redundant design with multiple paths between layers and within layers.

As a recommended practice, the core and distribution layers are built with redundant switches and fully meshed links that provide maximum redundancy and optimal convergence. Access switches should have redundant connections to redundant distribution switches. The network bandwidth and capacity are engineered to withstand a switch or link failure, usually recovering within 200 to 500 ms. OSPF and EIGRP timer manipulation quickly attempts to redirect the flow of traffic away from a router that has experienced a failure, and toward an alternate path.

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

---

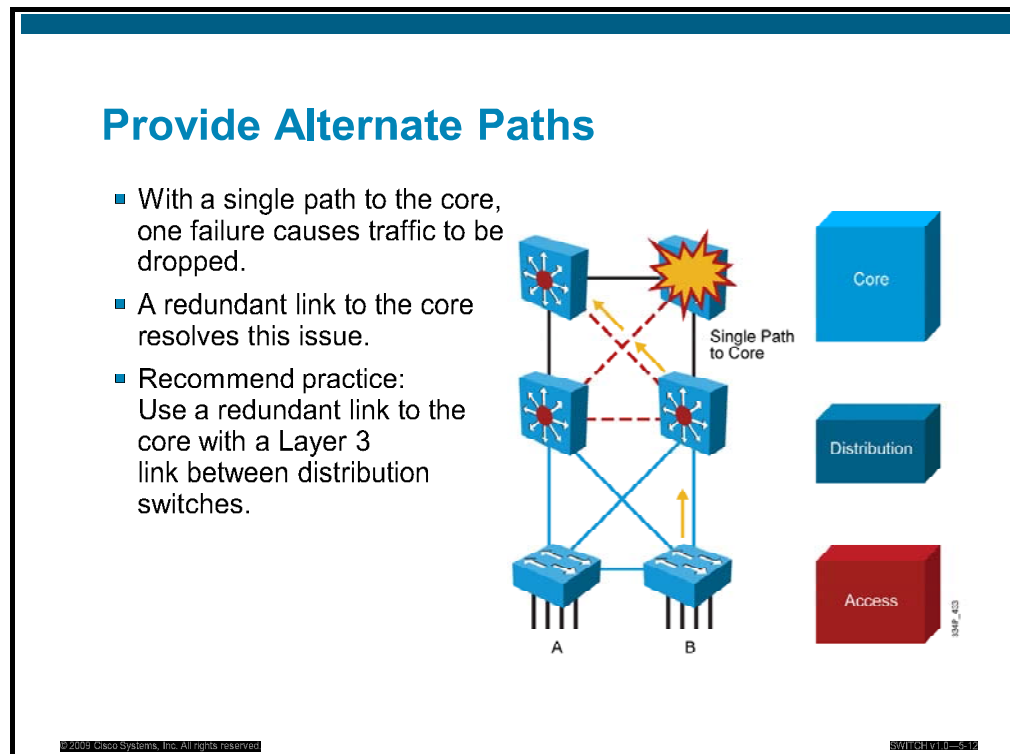
**Note** Combining OSPF and EIGRP timer manipulation with Cisco NSF might not be the most common deployment environment. OSPF and EIGRP timer manipulation is designed to improve convergence time in a multiaccess network (where several IGP routing peers share a common broadcast media, such as Ethernet). 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 to service providers or redundant Gigabit Ethernet point-to-point links to the campus infrastructure.

---

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

## Provide Alternate Paths

This subtopic describes how to provide alternate path in a highly available network.



Although dual distribution switches that are 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 is 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.

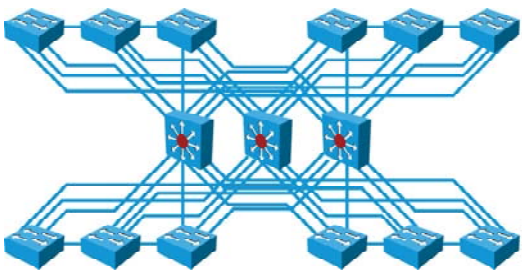
# Avoid Too Much Redundancy

This subtopic explains why too much redundancy does not achieve better availability.

## Avoid Too Much Redundancy

Too much redundancy can lead to design issues:

- Root placement
- Number of blocked links
- Convergence process
- Complex fault resolution
- Cost



© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH110-511

In the figure, a third switch is added to the distribution switches in the center. This extra switch adds unnecessary 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.
- Which links should be in a blocking state? It is very difficult 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.
- When something goes wrong, how do you find the source of the problem? The design is much more difficult to troubleshoot.




# Avoid Single Point of Failure

This subtopic describes how to avoid single point of failure in a switched network.

## Avoid Single Points of Failure

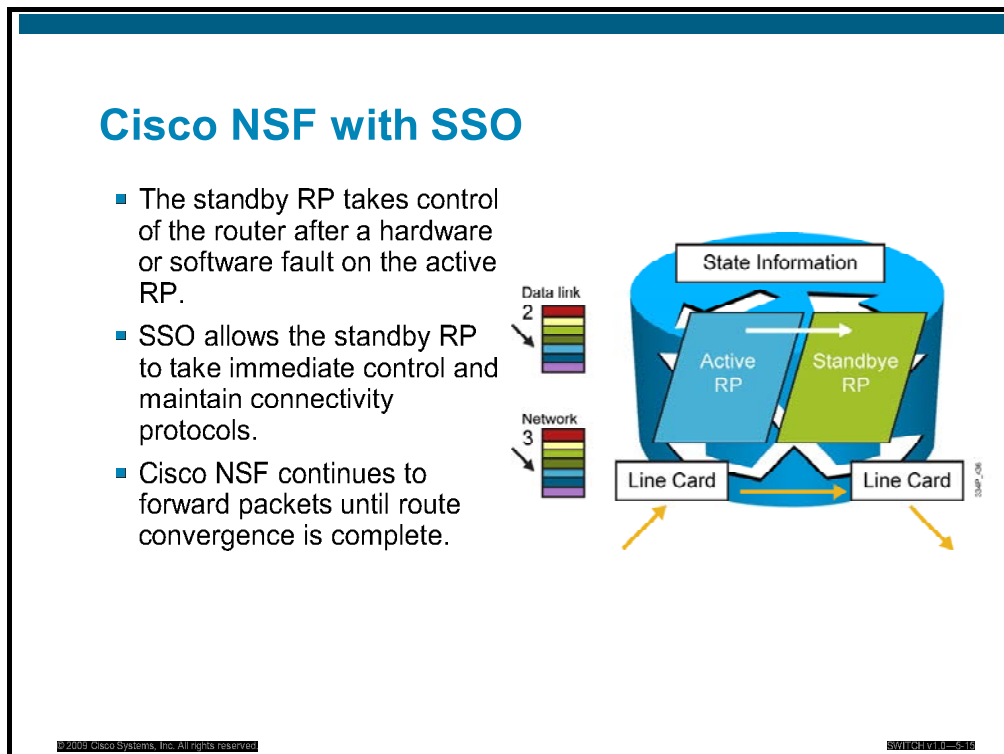
- The access layer is a candidate for supervisor redundancy.
- Layer 2 access layer SSO.
- Layer 3 access layer SSO and Cisco NSF.
- Reduces network outage to 1 to 3 seconds.
- Supported with Cisco Catalyst 4500 and 6500 Series Switches.



Avoiding single points of failure is another key element of high availability. If there is only one path to a device or a network, the loss of the path is unrecoverable. Redundancy is relatively easy to implement at the distribution or core layer, where duplicate links and duplicate devices can exist to interconnect the other elements of the network. Redundancy is more difficult to achieve at the access layer. An access switch failure is a single point of failure that causes an outage for the end devices that are connected to it. You can reduce the outage to 1 to 3 seconds in this access layer by utilizing SSO in a Layer 2 environment or Cisco NSF with SSO in a Layer 3 environment.

# Cisco NSF with SSO

This subtopic describes how Cisco NSF with SSO helps achieve high availability.



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 as well as 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) as well as access control list (ACL) and quality-of-service (QoS) tables.

Cisco NSF is a Layer 3 function that works with SSO to minimize the amount of time that 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 that is 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 is completed and do not experience routing flaps.

# Routing Protocols and NSF

This subtopic describes the routing protocol requirements for Cisco NSF.

## Routing Protocol Requirements for Cisco NSF

- Cisco NSF enhancements to routing protocols are designed to prevent routing flaps.
- Adjacencies must not be reset when switchover is complete; otherwise, protocol state is not maintained.
- FIB must remain unchanged during switchover.
  - Current routes are marked as stale during restart.
  - Routes are refreshed after Cisco NSF convergence is complete.
  - Transient routing loops or black holes may be introduced if the network topology changes before the FIB is updated.
- Switchover must be completed before dead or hold timer expires; otherwise, peers will reset the adjacency and reroute the traffic.
- Cisco NSF-capable routers are configured to support Cisco NSF.
- Routers that are aware of Cisco NSF know that Cisco NSF-capable router can still forward packets.
- Supported with EIGRP, OSPF, BGP, IS-IS.
- Supported with Cisco Catalyst 4550 and 6500 Series Switches.

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 do not reset. Data traffic is forwarded while the standby RP assumes control from the failed active RP during a switchover. User sessions that were established prior to the switchover are maintained.

It is crucial to Cisco NSF operation that the intelligent line cards can remain up through a switchover and can be kept current with the FIB on the active RP. While the control plane builds a new routing protocol database and restarts peering agreements, the data plane relies on preswitchover 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, or else the peers will reset the adjacency and reroute the traffic.

# Summary

This topic summarizes the key points that were discussed in this lesson.

## Summary

- High availability involves several elements: redundancy, technology, people, processes, and tools.
- At the network level, high availability involves making sure that there is always a possible path between two endpoints.
- High availability minimizes link and node failures to minimize downtime, by implementing link and node redundancy, providing alternate paths for traffic, and avoiding single points of failure.

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-5-1

## Lesson 2

---

# Implementing High Availability

---

## Overview

When designing a campus network, you need to plan the optimal use of the highly redundant devices. Carefully consider when and where to invest in redundancy to create a resilient and highly available network. This lesson describes these elements and explains how to implement them.

## Objectives

Upon completing this lesson, you will be able to implement high availability. This ability includes being able to meet these objectives:


- Implement high availability at the switch level
- Use Cisco StackWise technology on access switches
- Evaluate the impact of too little redundancy
- Assess the impact of uplink failure

# Distributed VLANs on Access Switches

This topic describes how to achieve redundancy at the switch level.

## Layer 2 Distributed VLANs on Access Switches

- Not a recommended design — slow convergence
- Use only if Layer 2 VLAN spanning flexibility is required
- Requires STP convergence for uplink failure and recovery
- More complex because the STP root and HSRP should match
- May be required for WLAN (standalone APs)



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. This design is more complex than the Layer 3 interconnection of the distribution switches. The Spanning Tree Protocol (STP) convergence process initiates for uplink failures and recoveries.

You should take the following steps to improve this suboptimal design:


- Use Rapid STP (RSTP) as the version of STP.
- Provide a Layer 2 trunk between the two distribution switches to avoid unexpected traffic paths and multiple convergence events.
- Place the Hot Standby Router Protocol (HSRP) primary router and the STP root bridge on the same distribution layer switch, if you choose to load-balance VLANs across uplinks. The HSRP and RSTP root should be colocated on the same distribution switches to avoid using the interdistribution link for transit.

## Local VLANs on Access Switches

This subtopic describes the impact of local VLANs on network redundancy and high availability.

### Layer 2 Local VLANs on Access Switches

- Recommended design, tried and true
- VLANs present on one access switch only
- Does not require STP convergence for uplink failure recovery
- Requires a distribution-to-distribution link for route summarization
- Can map Layer 2 VLAN numbers to Layer 3 subnets for ease of use and management



In this time-proven topology, no VLANs span between access layer switches across the distribution switches. A subnet equals a VLAN that, in turn, equals an access switch. The root for each VLAN is aligned with the active HSRP instance. From an 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.

---

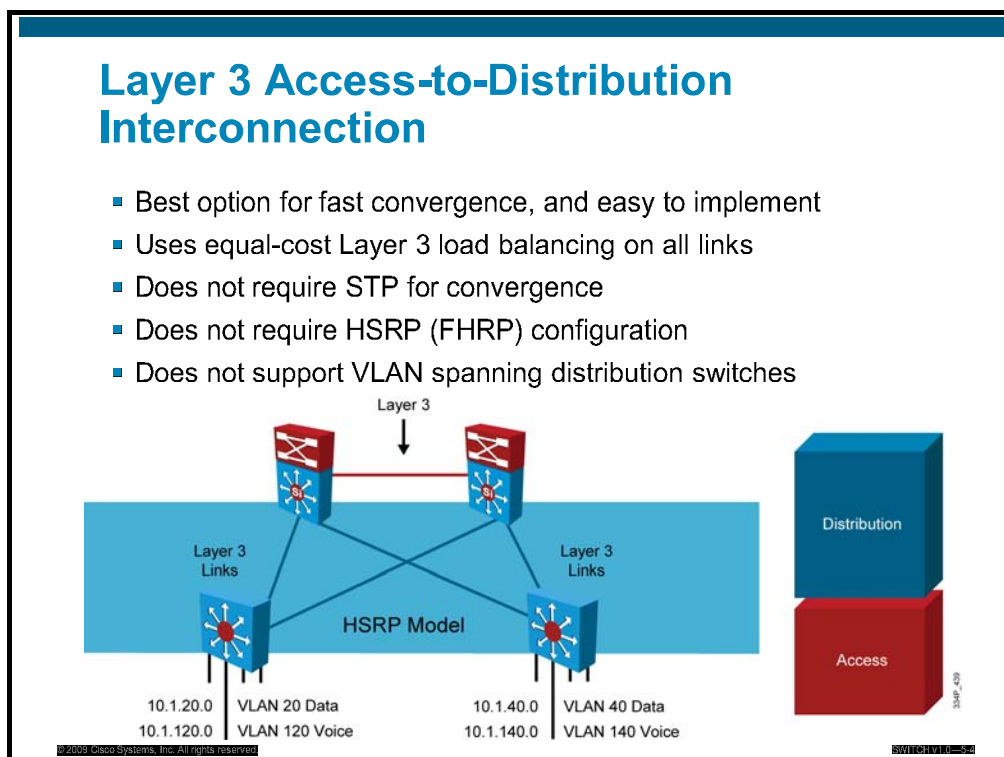
**Note** This recommended design provides the highest availability.

---

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 Access to Distribution Interconnection

This subtopic describes how high availability is achieved between access and distribution switches.



In this time-proven topology, no VLANs span between access layer switches across the distribution switches. A subnet equals a VLAN that, in turn, equals an access switch. The root for each VLAN is aligned with the active HSRP instance. From an 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.

---

**Note** This recommended design provides the highest availability.

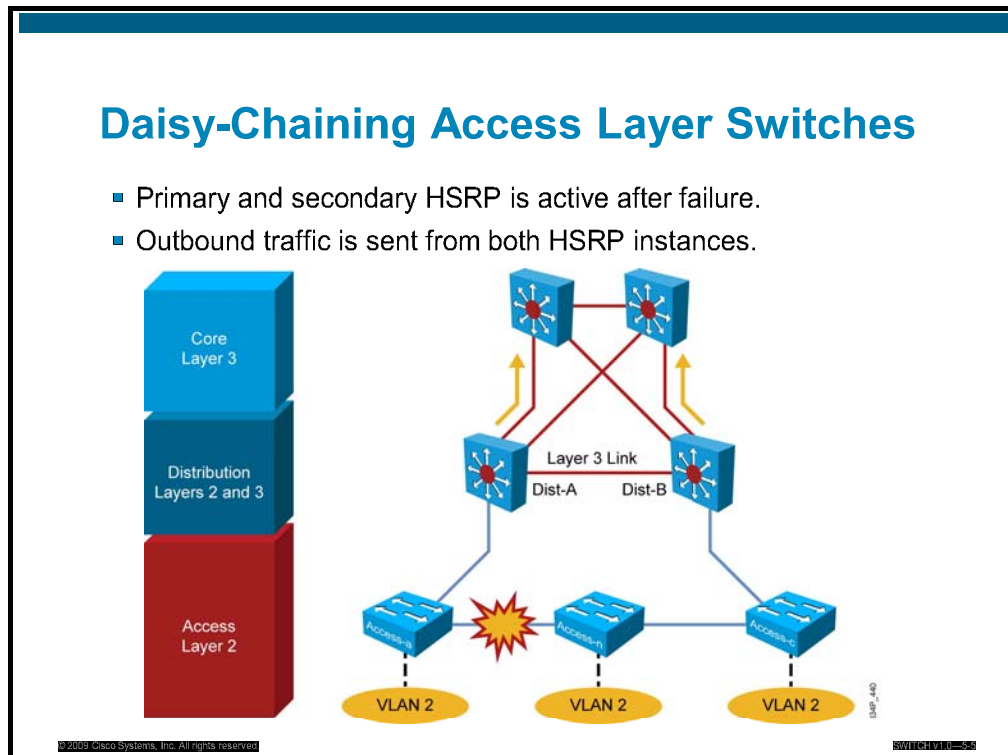
---

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.



# Daisy-Chaining Access Layer Switches

This subtopic describes how daisy-chaining access layer switches is a way to achieve redundancy.

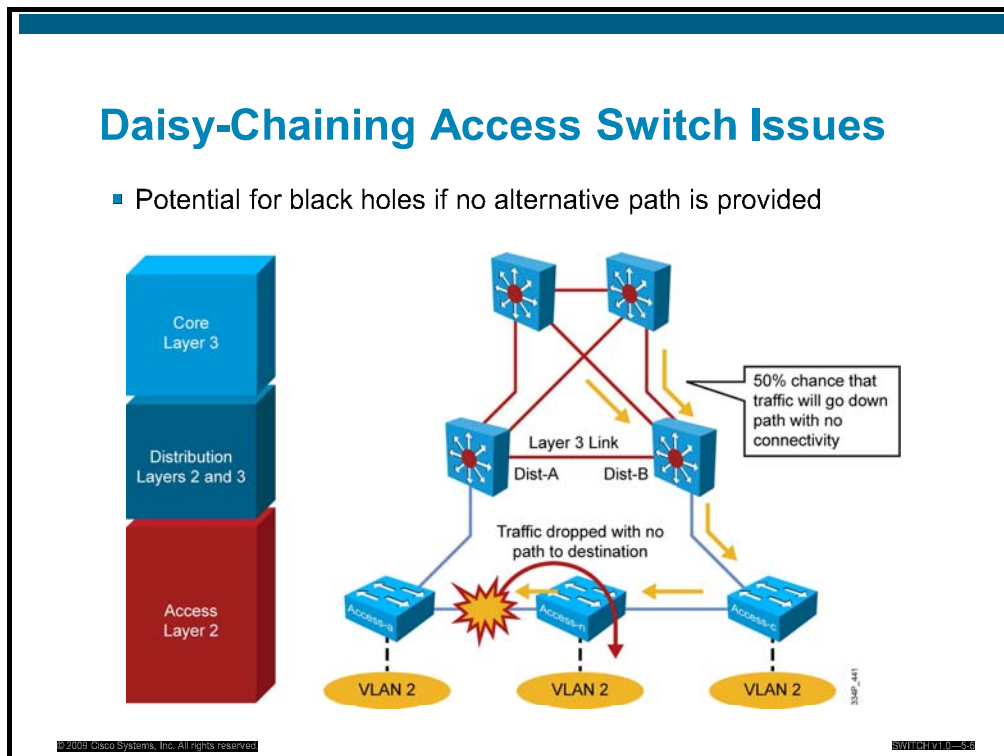


In the topology shown in the figure before failures, no links block from an STP or RSTP perspective. Both uplinks are available to actively forward and receive traffic. Both distribution nodes can forward return-path traffic from the rest of the network and toward the access layer for devices attached 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 scenario, 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 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.

# Daisy-Chaining Access Layer Switch Issues

This subtopic describes the limitations of access layer switches daisy chaining.



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 loopback cable that runs 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 Access Switches

This topic describes Cisco StackWise technology as a way to achieve redundancy.

## StackWise Technology Access Switches

- StackWise Technology eliminates the daisy-chain issue:
  - Loopback links are not required.
  - A Layer 2 link in the distribution is not required.
- StackWise switch provides redundancy.
- Uplinks can be on different switches within stack.
- Modular chassis-based switches can also eliminate the daisy-chain issue.



Cisco StackWise technology in the access layer supports the recommended practice of using a Layer 3 connection between the distribution switches without having to use a loopback cable or perform extra configuration.

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


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

# Too Little Redundancy

This topic explains the importance of avoiding too little redundancy.

## Avoiding Too Little Redundancy

- Looped figure-8 topology for VLANs spanning access switches
- Blocking on uplink from Access-b.
- Initially forwarding traffic from both access switches



The figure 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-8 topology. One access layer uplink will be blocking. HSRP hellos are exchanged by transiting the access switches.

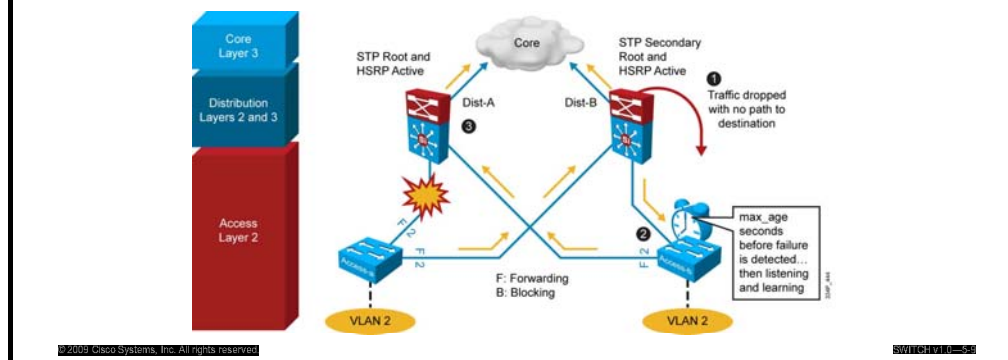
Initially, traffic is forwarded from both access switches to the Distribution A switch that supports the STP root and the primary or active HSRP peer for VLAN 2. However, this design will black-hole traffic and will be affected by multiple convergence events with a single network failure.

# Impact of Uplink Failure

This topic describes the impact of uplink failures.

## Impact of Uplink Failure

- 1 Traffic from Access A is dropped until HSRP goes active on Distribution B.
- 2 Blocking link on Access B takes 50 seconds to move to forwarding.
- 3 After STP converges, an HSRP preempt causes another transition. Access B is used as transit for Access A traffic.



In the figure, when the uplink from Access A to Distribution A fails, there are three convergence events:

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 due to 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.

---

**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 Per VLAN Spanning Tree Plus (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 Access A.

# Summary

This topic summarizes the key points that were discussed in this lesson.

## Summary

- Achieving redundancy can be achieved at Layer 2 and at Layer 3, by providing additional redundant paths between devices.
- StackWise Technology can be used to join several physical switches into one virtual switch.
- Redundancy is a balance between too much redundancy, which adds complexity to the network structure, and too little redundancy, which creates single points of failure.
- When uplinks fail, convergence path as well as convergence time must be taken into account to evaluate the impact of the failure on the network infrastructure.

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCHING-5-1

# Implementing Network Monitoring

---

## Overview

When designing a campus network, after redundancy is created, managing redundancy can be achieved by monitoring the network, through SNMP and syslog, and testing connectivity with IP service level agreements (SLAs). This lesson describes these various elements and explains how to implement them.

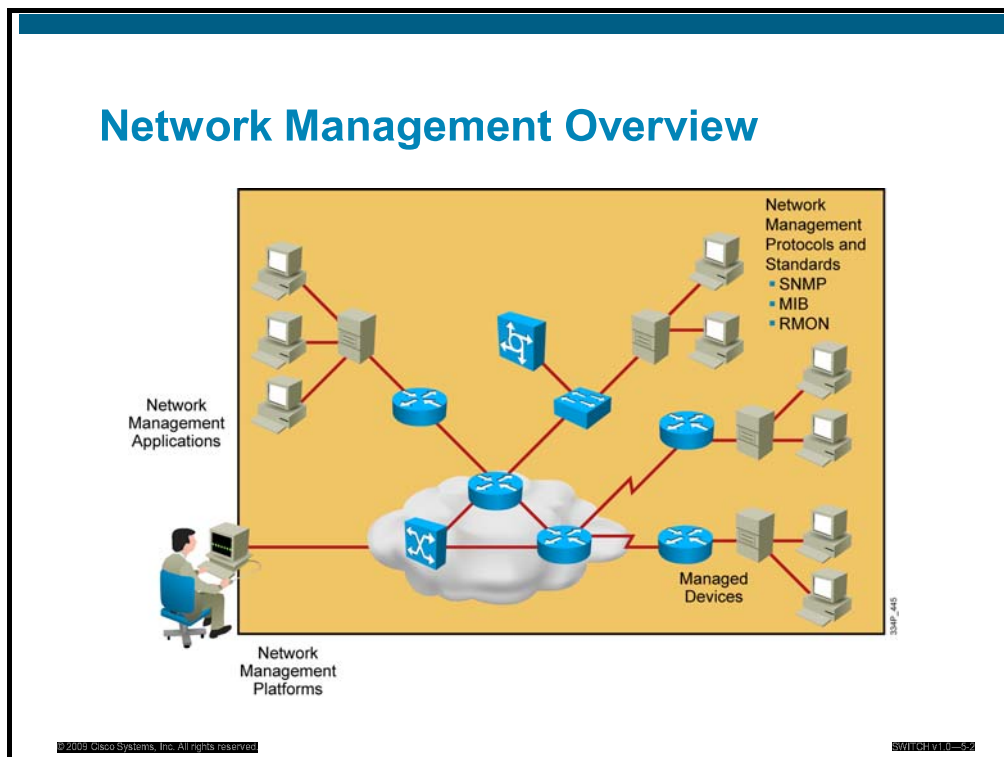
## Objectives

Upon completing this lesson, you will be able to implement network monitoring. This ability includes being able to meet these objectives:

- Implement network monitoringConfigure IP SLA technology

# Network Monitoring

This topic describes network monitoring and its configuration.



Network management is a set of tools and processes to help manage the network. Network administrators use network management so that they can be confident in the performance of the network:

- They use it to verify that the network is working well and behaving in the planned manner.
- They use it to characterize the performance of the network.
- They use it to understand how much traffic is flowing and where it is flowing in the network.
- They use it to provide tools and information to troubleshoot the network.




# Syslog

This subtopic describes syslog.

## Syslog Overview

- Allows software subsystems to report and save important error messages either locally or to a remote logging server.
- Can send messages on UDP port 514.
- Provides very comprehensive reporting mechanism in plain English text.
- Logging messages on console, terminal session, buffer, and syslog server.



The syslog process allows a device to report and save important error and notification messages, either locally or to a remote logging server. Syslog messages can be sent to local console connections, monitor connections (tty), the system buffer, or remote syslog servers. Syslog allows text messages to be sent to a syslog server using UDP port 514.

Syslog provides a very comprehensive reporting mechanism that logs system messages in plain English text. The syslog messages include both messages in a standardized format (called system logging messages, system error messages, or simply system messages) and output from debug commands. These messages are generated during network operation to assist with identifying the type and severity of a problem, or to aid users in monitoring router activity such as configuration changes.

# Syslog Features

This subtopic describes syslog features.

## Syslog Features

- Devices produce syslog messages.
- Syslog messages contain level and facility.
- Common syslog facilities:
  - IP
  - OSPF protocol
  - SYS operating system
  - IP Security (IPsec)
  - Route Switch Processor (RSP)
  - Interface
- Syslog levels:
  - Emergency (level 0, highest level)
  - Alert (level 1)
  - Critical (level 2)
  - Error (level 3)
  - Warning (level 4)
  - Notice (level 5)
  - Informational (level 6)
  - Debugging (level 7)

© 2009 Cisco Systems, Inc. All rights reserved. SWITCHING-50

The syslog is an essential component of any network operating system. It reports system state information to a network manager.

Cisco devices produce syslog messages as a result of network events. Every syslog message contains a severity level and a facility. Many networking devices support syslog, including routers, switches, application servers, firewalls, and other network appliances.

Syslog defines the levels that are listed in the figure. The smaller numerical levels are the more critical syslog alarms.

Syslog facilities are service identifiers that are used to identify and categorize system state data for error and event message reporting. Cisco IOS Software has more than 500 facilities. The most common syslog facilities are listed in the figure.

Other facilities include Cisco Discovery Protocol, Spanning Tree Protocol (STP), multicast, TCP, Border Gateway Protocol (BGP), RADIUS, Telnet, and those facilities that are related to quality-of-service (QoS) services.

More information about syslog is located at

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123sup/123sems/index.htm>

# Syslog Message Format

This subtopic describes the syslog message format.

## Cisco Syslog Message Standard

Documentation for each release explains the meaning of the messages.

`%FACILITY-SUBFACILITY-SEVERITY-MNEMONIC: Message-text`

`%SYS-5-CONFIG_I: Configured from console by cwr2000 on vty0 (192.168.64.25)`

The system messages begin with a percent sign (%) and are structured as shown in the figure and as described here:

- **Facility:** A code consisting of two or more uppercase letters that identify the hardware device, the protocol, or a module of the system software.
- **Severity:** A single-digit code from 0 to 7 that reflects the severity of the condition. The lower the number, the more serious the situation.
- **Mnemonic:** A code that uniquely identifies the error message.
- **Message-text:** A text string describing the condition. This portion of the message sometimes contains detailed information about the event, including terminal port numbers, network addresses, or addresses that correspond to locations in the system memory address space.

The figure shows a typical message that indicates that the operating system (facility = SYS) is providing a notification (SEVERITY = 5) that has been configured (MNEUMONIC = CONFIG). The message text indicates that a user on VTY0 from IP address 192.168.64.25 made this change.

---

**Note** The documentation for each Cisco IOS Software release, such as the information found at [http://www.cisco.com/en/US/products/hw/switches/ps4324/products\\_system\\_message\\_guides\\_list.html](http://www.cisco.com/en/US/products/hw/switches/ps4324/products_system_message_guides_list.html), explains the meaning of these messages.

---

# Syslog Message Log

This subtopic describes how syslog messages are stored and displayed.

## Example: Syslog Messages

```
08:01:13: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/5, changed state to up
08:01:23: %DUAL-5-NBRCHANGE: EIGRP-IPv4: (1) 1: Neighbor 10.1.1.1 (Vlan1) is up: new adjacency
08:02:31: %LINK-3-UPDOWN: Interface FastEthernet0/8, changed state to up
08:18:20: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/5, changed state to down
08:18:22: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/5, changed state to up
08:18:24: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/2, changed state to down
08:18:24: %ILPOWER-5-IEEE_DISCONNECT: Interface Fa0/2: PD removed
08:18:26: %LINK-3-UPDOWN: Interface FastEthernet0/2, changed state to down
08:19:49: %ILPOWER-7-DETECT: Interface Fa0/2: Power Device detected: Cisco PD
08:19:53: %LINK-3-UPDOWN: Interface FastEthernet0/2, changed state to up
08:19:53: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/2, changed state to up
08:27:42: %SYS-5-CONFIG_I: Configured from console by vty1 (10.1.1.24)
08:29:32: %ILPOWER-7-DETECT: Interface Fa0/3: Power Device detected: IEEE PD
08:29:36: %LINK-3-UPDOWN: Interface FastEthernet0/3, changed state to up
08:29:36: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/3, changed state to up
08:31:19: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/5, changed state to down
08:31:21: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/5, changed state to up
```

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH1-VLAN5-50

## Example: Syslog Messages

The figure here shows samples of syslog messages that Cisco IOS Software produces. The most common messages are link up and down messages and messages that a device produces when it exits from configuration mode. If access control list (ACL) logging is configured, the device generates syslog messages when packets match a parameter condition. ACL logging can be useful to detect packets that are denied access based on the security policy that is set by an ACL.

System log messages can contain up to 80 characters and a percent sign (%), which follows the optional sequence number or time-stamp information, if configured. Messages are displayed in this format:

- *seq no:timestamp: %facility-severity-MNEMONIC:description*

A sequence number appears on the syslog message if the **service sequence-numbers** global configuration command is configured.

The time stamp shows the date and time of the message or event if the **service timestamps log [datetime | log]** global configuration command is configured. The time stamp can have one of three formats:

- *mm/dd hh:mm:ss*
- *hh:mm:ss* (short uptime)
- *d h* (long uptime)

# System Log Configuration

This subtopic describes how system logs are configured on a switch.

## System Log Configuration

```
sw(config)# logging server ip address
sw(config)# logging trap level
```

- Configures a syslog server and the logging level

```
sw(config)# logging buffered [buffer_size] [alarm_level]
```

- Configures the system local log

```
sw# show logging ?
count      Show counts of each logging message
history    Show the contents of syslog history table
onboard    Onboard logging information
xml        Show the contents of XML logging buffer
|          Output modifiers
```

- Displays the local logs

To configure a syslog server, use the command **logging Ip address** *cf the syslog server*.

To configure from which severity level messages must be sent to the syslog server, use the command **logging trap level**:

```
sw(config)#logging trap ?
<0-7>      Logging severity level
alerts     Immediate action needed           (severity=1)
critical   Critical conditions               (severity=2)
debugging  Debugging messages               (severity=7)
emergencies System is unusable                (severity=0)
errors     Error conditions                  (severity=3)
informational Informational messages          (severity=6)
notifications Normal but significant conditions (severity=5)
warnings   Warning conditions                (severity=4)
```

Messages can be sent to a syslog server. They can also be kept on the local switch. To configure the local logs, use the command **logging buffered**. Valid parameters are the maximum local log size and the severity level that must be logged:

```
sw(config)#logging buffered ?
<0-7>      Logging severity level
<4096-2147483647> Logging buffer size
alerts     Immediate action needed           (severity=1)
critical   Critical conditions               (severity=2)
debugging  Debugging messages               (severity=7)
discriminator Establish MD-Buffer association
emergencies System is unusable                (severity=0)
errors     Error conditions                  (severity=3)
informational Informational messages          (severity=6)
notifications Normal but significant conditions (severity=5)
warnings   Warning conditions                (severity=4)
xml        Enable logging in XML to XML logging buffer
```

Use the **show logging** command to display the content of the local log files. When too many events are present in the log files, use the pipe argument (|) in combination with keywords such as *include* or *begin* to filter the output. For example, the following displays all events present in the local logs that involve an error report (severity level 3) about interface link status:

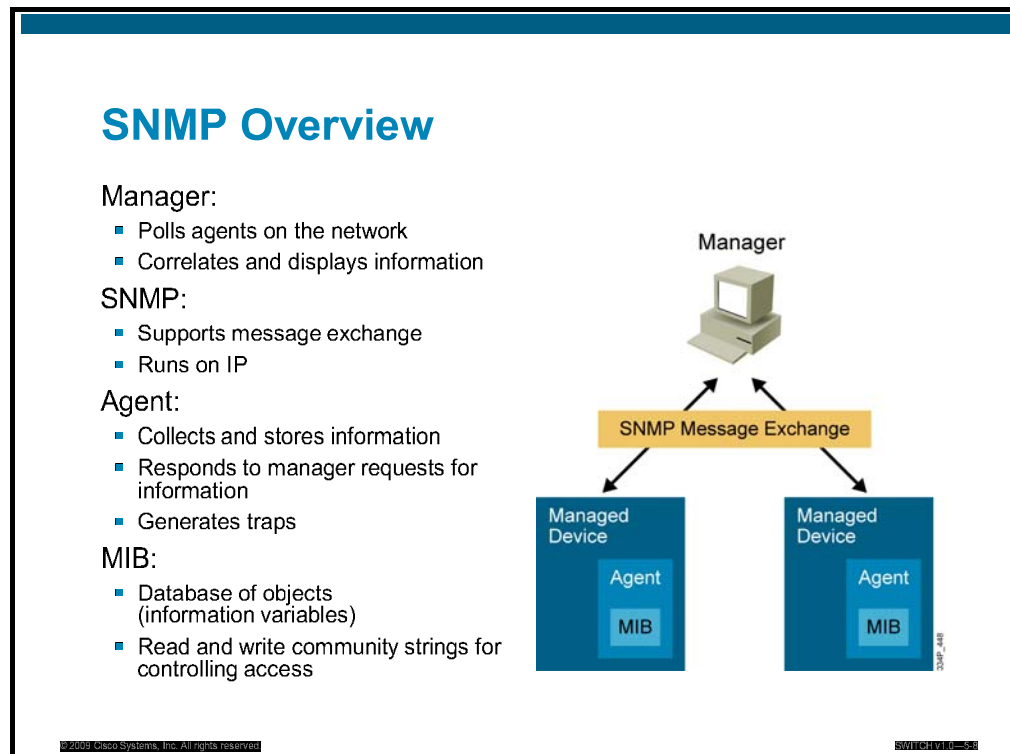
```
sw#show logging | inc LINK-3
2d20h: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
2d20h: %LINK-3-UPDOWN: Interface FastEthernet0/2, changed state to up
2d20h: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
```

Similarly, the following example displays all events that start with %DUAL (and are therefore reporting events that are related to the EIGRP DUAL algorithm):

```
sw#show logg | beg %DUAL
2d22h: %DUAL-5-NBRCHANGE: EIGRP-IPv4:(10) 10: Neighbor 10.1.253.13
(FastEthernet0/11) is down: interface down
2d22h: %LINK-3-UPDOWN: Interface FastEthernet0/11, changed state to down
2d22h: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/11,
changed state to down
```

# SNMP Overview

This subtopic describes the Simple Network Management Protocol (SNMP).



SNMP has become the standard for network management. SNMP is a simple solution that requires little code to implement and thus enables vendors to easily build SNMP agents for their products. Therefore, SNMP is often the foundation of a network management architecture.

SNMP defines how management information is exchanged between network management applications and management agents. A network management application periodically polls the SNMP agents that reside on managed devices by querying the device for data. The periodic SNMP polling has the disadvantage that there is a delay between the time that an event occurs and the time that it is noticed by the network management system (NMS). There is a trade-off between polling frequency and bandwidth usage.

A network management application can display the information in a GUI on the network manager.

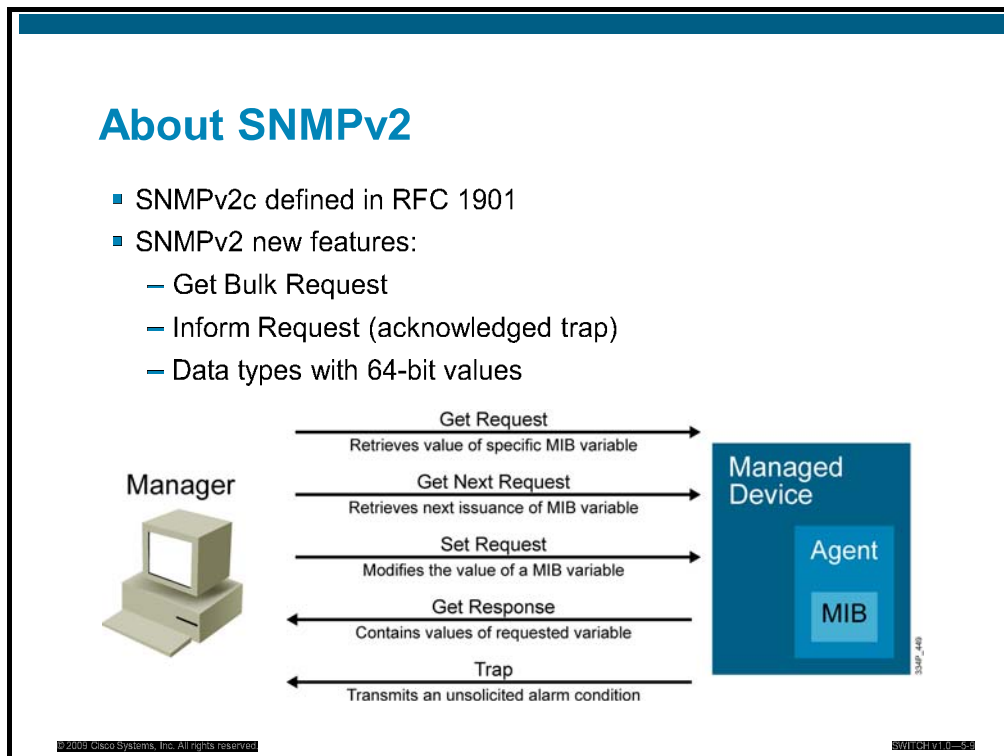
SNMP uses the User Datagram Protocol (UDP) transport mechanism of IP to retrieve and send management information, such as MIB variables.

SNMP management agents that reside on managed devices collect and store information about the device and its operation, respond to managerial requests, and generate traps to inform the manager of certain events. SNMP traps are sent by management agents to the network management system when certain events occur. Trap-directed notification can result in a substantial savings of network and agent resources by eliminating the need for some SNMP polling requests.

The management agent collects data and stores it locally in the MIB. Community strings control access to the MIB. To view or set MIB variables, the user must specify the appropriate community string for read or write access.

# SNMPv1 vs. SNMPv2

This subtopic compares SNMP v1 to SNMP v2.



The initial version of the SNMP standard (SNMP version 1, or SNMPv1) is defined in RFC 1157. There are five basic SNMP messages that the network manager uses to transfer data from agents that reside on managed devices:

- **Get Request:** Used to request the value of a specific MIB variable from the agent.
- **Get Next Request:** Used after the initial Get Request to retrieve the next object instance from a table or a list.
- **Set Request:** Used to set a MIB variable on an agent.
- **Get Response:** Used by an agent to respond to a Get Request or Get Next Request from a manager.
- **Trap:** Used by an agent to transmit an unsolicited alarm to the manager. An agent sends a Trap message when a certain condition occurs, such as a change in the state of a device, a device or component failure, or an agent initialization or restart.

SNMPv2 was introduced with RFC 1441, but members of the Internet Engineering Task Force (IETF) subcommittee could not agree on the security and administrative sections of the SNMPv2 specification. There were several attempts to achieve acceptance of SNMPv2 through the release of experimental modified versions.

Community-based SNMPv2 (SNMPv2c), defined in RFC 1901, is the most common implementation. SNMPv2C deploys the administrative framework defined in SNMPv1, which uses read/write community strings for administrative access.



SNMPv2 introduces two new message types:

- **Get Bulk Request:** This message type reduces repetitive requests and replies and improves performance when you are retrieving large amounts of data (for example, tables).
- **Inform Request:** Inform Request messages alert an SNMP manager of specific conditions. Unlike SNMP Trap messages, which are unconfirmed, the NMS acknowledges an Inform Request by sending an Inform Response message back to the requesting device.

SNMPv2 adds new data types with 64-bit counters, because 32-bit counters were quickly outmoded by fast network interfaces. On Cisco routers, SNMPv2 is implemented in Cisco IOS Software Release 11.3 and later.

---

<b>Note</b>	Neither SNMPv1 nor SNMPv2 offers security features. Specifically, SNMPv1 and SNMPv2 can neither authenticate the source of a management message nor provide encryption. Because of the lack of security features, many SNMPv1 and SNMPv2 implementations are limited to a read-only capability, reducing their utility to that of a network monitor.
-------------	--

---

## About SNMPv3

This subtopic describes SNMPv3.

### About SNMPv3

- RFCs 3410 through 3415
- Authentication and privacy
- Authorization and access control
- Usernames and key management
- Remotely configurable via SNMP operations

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-10-5-10

SNMPv3 is described in RFCs 3410 through 3415. It adds methods to ensure the secure transmission of critical data between managed devices.

SNMPv3 introduces three levels of security:

- **noAuthNoPriv:** No authentication is required, and no privacy (encryption) is provided.
- **authNoPriv:** Authentication is based on Hashed Message Authentication Code-Message Digest 5 (HMAC-MD5) or Hashed Message Authentication Code-Secure Hash Algorithm (HMAC-SHA). No encryption is provided.
- **authPriv:** In addition to authentication, Cipher Block Chaining-Data Encryption Standard (CBC-DES) encryption is used as the privacy protocol.

Security levels that are implemented for each security model determine which SNMP objects a user can access for reading, writing, or creating, and the list of notifications that its users can receive.

On Cisco routers, SNMPv3 is implemented in Cisco IOS Software Release 12.0 and later.

# SNMP Recommendations

This subtopic provides recommendations for SNMP implementation.

## SNMP Recommendations

Configure ACKs for SNMP community strings.

- Restricts SNMP traffic to addresses in ACL.

Use SNMPv3 if possible.

- Provides authentication and encryption.

SNMPv1 and SNMPv2 use community strings in clear text. As with all passwords, you should carefully choose these community strings to ensure that they are not trivial. Community strings should be changed at regular intervals and in accordance with network security policies. For example, the strings should be changed when a network administrator changes roles or leaves the company. If SNMP is used only to monitor devices, use read-only communities. Ensure that SNMP messages do not spread beyond the management consoles. You can use access lists to prevent SNMP messages from going beyond the required devices, and on the monitored devices to limit access for management systems only.

SNMPv3 is recommended because it provides authentication and encryption.

# SNMP Configuration

This subtopic describes how to configure SNMP basic features.

## SNMP Configuration

- Configure SNMP community strings.
- Configure SNMP access lists.
- Configure SNMP trap receiver.
- Configure SNMPv3 user.

```
sw(config)# access-list 100 permit ip 10.1.1.0 0.0.0.255 any
sw(config)# snmp-server community cisco RO 100
sw(config)# snmp-server community xyz123 RW 100
sw(config)# snmp-server trap 10.1.1.50
```

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH10-531

The first step necessary for SNMP configuration is to enable SNMP access. You do this by configuring community strings, which act somewhat like passwords. The difference is that there can be several community strings, and each one may grant different forms of access. In the preceding figure, **community cisco** grants read-only access to the local switch, while **xyz123** grants read and write access to the local switch.

The **100** at the end of the **snmp-server community** lines restricts access to sources that are permitted via standard access-list 100. In this case, all stations in subnet 10.1.1.0/24 can access the local switch with both **communities cisco** and **xyz123**.

The last line achieves two purposes: it configures the SNMP server and instructs the switch to send its traps to this server.

There are many other SNMP options. See this page for more information:

[http://www.cisco.com/en/US/docs/ios/netmgmt/configuration/guide/nm\\_cfg\\_snmp\\_sup\\_ps6350\\_TSD\\_Products\\_Configuration\\_Guide\\_Chapter.html](http://www.cisco.com/en/US/docs/ios/netmgmt/configuration/guide/nm_cfg_snmp_sup_ps6350_TSD_Products_Configuration_Guide_Chapter.html)

# IP SLA Configuration

This topic describes IP service level agreement (SLA) and its basic configuration.

## SLA Review

- Companies need predictability in IP services as networks becoming increasingly important.
- An SLA is a contract between the provider and its customers:
  - Provides a guarantee of service level.
  - Specifies connectivity and performance agreements for an end-user service.
  - Supports problem isolation and network planning.

The network has become increasingly critical for customers, and any downtime or degradation can adversely impact revenue. Companies need some form of predictability with IP services. An SLA is a contract between the network provider and its customers, or between a network department and internal corporate customers. It provides a form of guarantee to customers about the level of user experience.

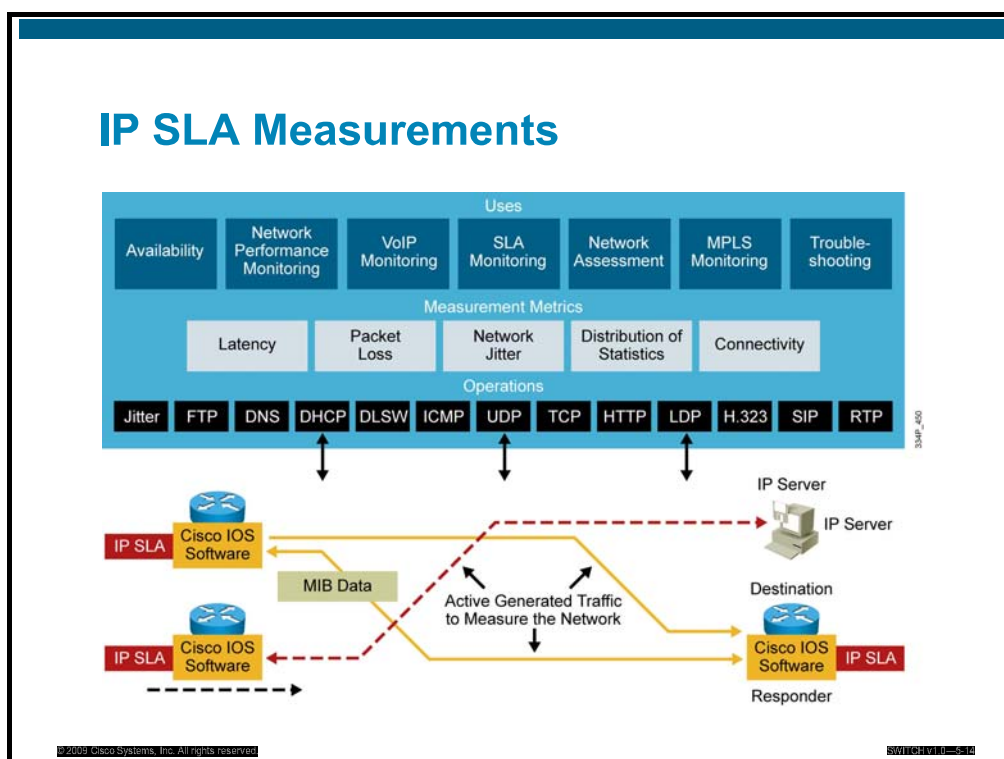
An SLA specifies connectivity and performance agreements for an end-user service from a service provider. The SLA will typically outline the minimum level of service and the expected level of service. The networking department can use the SLAs to verify that the service provider is meeting its own SLAs or to define service levels for critical business applications. An SLA can also be used as the basis for planning budgets and justifying network expenditures.

Administrators can ultimately reduce the mean time to repair (MTTR) by proactively isolating network issues. They can change the network configuration based on optimized performance metrics.

Typically, the technical components of an SLA contain a guaranteed level for network availability, network performance in terms of round-trip time (RTT), and network response in terms of latency, jitter, and packet loss. The specifics of an SLA vary depending on the applications that an organization is supporting in the network.

# IP SLA Measurements

This subtopic describes IP SLA measurements.



The IP SLA measurement functionality in Cisco IOS Software allows the configuration of a router to send synthetic traffic to a host computer or to a router that has been configured to respond. One-way travel times and packet loss data is gathered. Certain measurements also allow jitter data to be collected.

There are several common functions for IP SLA measurements:

- Edge-to-edge network availability monitoring
- Network performance monitoring and network performance visibility
- VoIP, video, and virtual private network (VPN) monitoring
- SLA monitoring
- IP service network health readiness or assessment
- Multiprotocol Label Switching (MPLS) network monitoring
- Troubleshooting of network operation

IP SLA measurement uses a variety of operations and actively generated traffic probes to gather many types of measurement statistics:

- Network latency and response time
- Packet loss statistics
- Network jitter and voice quality scoring
- Statistical end-to-end matrix of performance information
- End-to-end network connectivity

Multiple IP SLA operations (measurements) can be running in a network at one time. Reporting tools use SNMP to extract the data into a database and then report on it.

IP SLA measurements allow the network manager to verify service guarantees, which increases network reliability by validating network performance, proactively identifying network issues, and easing the deployment of new IP services.

## IP SLA Operations

This subtopic describes IP SLA operations.

### IP SLA Operations

- Operation is a measurement including protocol, frequency, traps, and thresholds.
- Network manager defines UDP or TCP port for each IP SLA measurement operation.
- IP SLAs can send traffic with different DSCP values.
- IP SLA control protocol is used between source and responder.
- MD5 authentication is supported between source and responder.
- Results are stored on IP SLA source in the IP SLA MIB.

The network manager configures a target device, protocol, and User Datagram Protocol (UDP) or TCP port number on the IP SLA source for each operation. The source uses the IP SLA control protocol to communicate with the responder before sending test packets. To increase security on IP SLA measurements control messages, the responder can utilize Message Digest 5 (MD5) authentication for securing the control protocol exchange. After the operation is finished and the response is received, the results are stored in the IP SLA MIB on the source, and are retrieved using SNMP.

IP SLA operations are defined to target devices. If the operation is something like Domain Name System (DNS) or HTTP, the target device might be any suitable computer. For operations such as testing the port that is used by a database, an organization might not want to risk unexpected effects and would use the IP SLA responder functionality to have a router respond in place of the actual database server. You can enable responder functionality in a router with only one command, and no complex or per-operation configuration is required.

## IP SLA Source and Responder

This subtopic describes the IP SLA source and responder roles.

### IP SLA Source and Responder

#### IP SLA source

- Cisco IOS Software device that sends data for operation.
  - Target device may or may not be a Cisco IOS Software device.
  - Some operations require an IP SLA responder.
- IP SLA source stores results in MIB.

#### IP SLA responder

- Greater measurement accuracy is available between an IP SLA source and responder.
- IP SLA responder is a Cisco IOS Software device that is configured to respond to IP SLA packets that are based on the **ip sla monitor responder** configuration command.

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-10-5-10

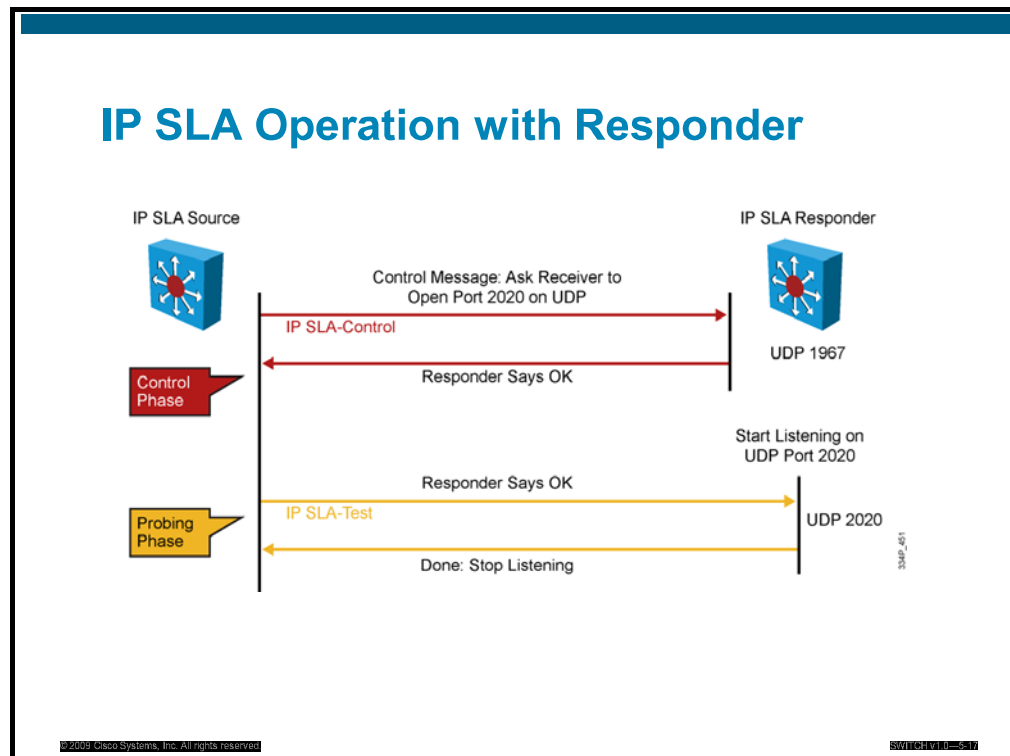
The IP SLA source is where all IP SLA measurement probe operations are configured either by the command-line interface (CLI) or through an SNMP tool that supports IP SLA operation. The source is also the Cisco IOS device that sends probe packets. The destination of the probe may be another Cisco router or another network target such as a web server or IP host.

Although the destination of the probe can be any IP device, the measurement accuracy is improved with an IP SLA responder. An IP SLA responder is a device that runs Cisco IOS Software and is configured as an IP SLA measurement responder with the **ip sla monitor responder** configuration command.



# IP SLA Operations with Responder

This subtopic describes IP SLA mechanism with source and responder.



The IP SLA source is where all IP SLA measurement probe operations are configured either by the command-line interface (CLI) or through an SNMP tool that supports IP SLA operation. The source is also the Cisco IOS Software device that sends probe packets. The destination of the probe may be another Cisco router or another network target such as a web server or IP host. Although the destination of the probe can be any IP device, the measurement accuracy is improved with an IP SLA responder. An IP SLA responder is a device that runs Cisco IOS Software and is configured as an IP SLA measurement responder with the **ip sla monitor responder** configuration command.

The network manager configures an IP SLA operation by defining a target device, protocol, and port number on the IP SLA source. The network manager can also configure reaction conditions. The operation is scheduled to be run for a period of time to gather statistics. The following sequence of events occurs for each IP SLA operation that requires a responder on the target:

1. At the start of the control phase, the IP SLA source sends a control message with the configured IP SLA operation information to IP SLA control port UDP 1967 on the target router. The control message carries information such as protocol, port number, and duration.
  - If MD5 authentication is enabled, MD5 checksum is sent with the control message.
  - If the authentication of the message is enabled, the responder verifies it; if the authentication fails, the responder returns an authentication failure message.
  - If the IP SLA measurement operation does not receive a response from a responder, it tries to retransmit the control message and eventually times out.

2. If the responder processes the control message, it sends an OK message to the source router and listens on the port that is specified in the control message for a specified duration. If the responder cannot process the control message, it returns an error. In the figure, UDP port 2020 is used for the IP SLA test packets.

---

<b>Note</b>	The responder is capable of responding to multiple IP SLA measurement operations that try to connect to the same port number.
-------------	---

---


3. If the return code of control message is OK, then the IP SLA operation moves to the probing phase, where it will send one or more test packets to the responder for response time computations. The return code is available in the **show ip sla statistics** command. In the figure, these test messages are sent on control port 2020.
4. The responder accepts the test packets and responds. Based on the type of operation, the responder may add an “in” time stamp and an “out” time stamp in the response packet payload to account for CPU time that is spent in measuring unidirectional packet loss, latency, and jitter to a Cisco device. These time stamps help the IP SLA source to make accurate assessments on one-way delay and the processing time in the target routers. The responder disables the user-specified port after it responds to the IP SLA measurements packet or when a specified time expires.

# IP SLA Responder Time Stamps

This subtopic describes IP SLA responder time stamps.

## IP SLA Responder Time Stamps

- IP SLA responder takes two time stamps (T2 and T3).
- IP SLA responder factors out destination processing time, making results highly accurate.
- IP SLA responder allows for one-way measurements for latency, jitter, and packet loss.



The figure illustrates the use of IP SLA responder time stamps in round-trip calculations. The IP SLA source will use four time stamps for the round-trip time (RTT) calculation. The IP SLA source sends a test packet at time T1.

The IP SLA responder includes both the receipt time (T2) and the transmitted time (T3). Because of other high-priority processes, routers can take tens of milliseconds to process incoming packets. The delay affects the response times because the reply to test packets might be sitting in a queue while waiting to be processed. This time stamping is made with a granularity of submilliseconds. At times of high network activity, an Internet Control Message Protocol (ICMP) ping test often shows a long and inaccurate response time, while an IP SLA-based responder shows an accurate response time. The IP SLA source subtracts T2 from T3 to produce the time spent processing the test packet in the IP SLA responder. This time is represented by a delta value.

The delta value is then subtracted from the overall RTT. The same principle is applied by the IP SLA source where the incoming T4 is also taken at the interrupt level to allow for greater accuracy, as compared with T5 when the packet is processed.

An additional benefit of two time stamps at the IP SLA responder is the ability to track one-way delay, jitter, and directional packet loss. These statistics are critical, because a great deal of network behavior is asynchronous. To capture one-way delay measurements, you must configure both the IP SLA source and the IP SLA responder with the Network Time Protocol (NTP).

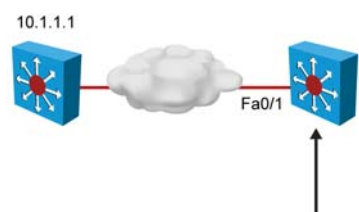
Both the source and the target must be synchronized to the same clock source. The IP SLA responder provides enhanced accuracy for measurements, without the need for dedicated third-party external probe devices. It also provides additional statistics, which are not otherwise available via standard ICMP-based measurements.

# IP SLA Configuration

This subtopic describes basic IP SLA configuration.

## IP SLA Configuration

- Configure IP SLA probe.
- Activate probe.
- Configure tracking object.
- Configure action on tracking object.
- Exact syntax depends on platform and on Cisco IOS version.



```
sw(config)# ip sla monitor 11
sw(config-sla)# type echo prot icmpEcho 10.1.1.1 source-int fa0/1
sw(config-sla)# frequency 10
sw(config)# ip sla monitor schedule 11 life forever start-time now
sw(config)# track 1 ip sla 11 reachability
```

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-10-5-10

There are many ways of implementing IP SLA. Different hardware platforms and different Cisco IOS versions may have a slightly different approach to IP SLA configuration. This is an example.

The first step is to use the command **ip sla monitor**, followed by a number to enter in IP SLA configuration mode.

In the preceding example, IP SLA test is done by sending an **icmpEcho** message to IP address 10.1.1.1 from the local interface f0/1. This message is sent every 10 seconds.

You can use the **ip sla responder** command to configure the 10.1.1.1 device to answer this message.

At this point, the type of message is configured, along with its frequency and target address. The next step is to decide when this test should start. You configure the start time by using the **ip sla monitor schedule** command. In this example, the test is to start immediately and is to last forever.

The IP SLA defines the test. Other commands are necessary to determine what action should be taken when the test result is received. In this example, the **track** command follows the IP SLA test result. Other commands can then use the track result to decrement the interfaces priority or to activate backup links.

# IP SLA Verification

This subtopic describes basic IP SLA verification.

## IP SLA Verification

Displays status of the IP SLA test, and its successes and failures.

**Show ip sla configuration** can also be used to display details of the IP SLA test that was conducted.

```
sw# show ip sla statistics
Round Trip Time (RTT) for Index 1
    Latest RTT: NoConnection/Busy/Timeout
Latest operation start time: 11:11:22.533 eastern Thu Jul 9 2010
Latest operation return code: Timeout
Over thresholds occurred: FALSE
Number of successes: 177
Number of failures: 6
Operation time to live: Forever
Operational state of entry: Active
Last time this entry was reset: Never
```

After IP SLA is configured, the test is conducted according to the schedule configuration. The test may succeed or fail. If you do not monitor the test results, it may fail silently. To display information about the test, use the **show ip sla statistics** command. It displays, among other parameters, the number of successes and the number of failures. It also shows whether the test is still being run. In the example shown in the figure, the test is still in the active state and has succeeded 177 times, but also has failed 6 times. Monitoring these statistics over time can tell you if there is a connection issue that has been discovered through the IP SLA test.

To get more information about a given IP SLA test configuration, use the **show ip sla configuration** command. This command displays how the test is run. For example:

```
sW#sh ip sla configuration
IP SLAs, Infrastructure Engine-II

Entry number: 1
Owner:
Tag:
Type of operation to perform: echo
Target address/Source address: 10.1.3.10/10.1.253.1
Type Of Service parameter: 0x0
Request size (ARR data portion): 28
Operation timeout (milliseconds): 5000
Verify data: No
Vrf Name:
Schedule:
  Operation frequency (seconds): 5
  Next Scheduled Start Time: Start Time already passed
  Group Scheduled : FALSE
  Randomly Scheduled : FALSE
  Life (seconds): Forever
  Entry Ageout (seconds): never
  Recurring (Starting Everyday): FALSE
  Status of entry (SNMP RowStatus): Active
Threshold (milliseconds): 5000
Distribution Statistics:
  Number of statistic hours kept: 2
  Number of statistic distribution buckets kept: 1
  Statistic distribution interval (milliseconds): 20
History Statistics:
  Number of history Lives kept: 0
  Number of history Buckets kept: 15
  History Filter Type: None
```

The result of this command gives additional information over the common **show running-config** display. With **show ip sla configuration**, you can also verify which IP address is used as a source, what is the size of each packet, and what are the default timeout and frequency for the test.

# Summary

This topic summarizes the key points that were discussed in this lesson.

## Summary

- Redundancy can be achieved at Layer 2 and at Layer 3 by providing additional redundant paths between devices.
- Syslog and SNMP are used to monitor device status.
- In IP SLA deployments, IP SLA measurements are performed between an IP SLA source and a destination (IP host or IP SLA responder).

© 2009 Cisco Systems, Inc. All rights reserved.

5-57





## Lesson 4

---

# Lab 5-1 Debrief

---

## Overview

In this lab, you have configured your pod switches and routers for high availability. You have configured and tested syslog, SNMP, and IP SLA. These three elements allow you to ensure that connectivity to an external device can be tested, and that chosen events that are relevant to your network will be sent to a syslog server or to an SNMP server.

During the lab debrief, the instructor will lead a group discussion in which you can present your solution. You will get an opportunity to verify your solution against a number of checkpoints provided by the instructor, and compare your solution to those of other students. The instructor will discuss alternative solutions and their benefits and drawbacks.

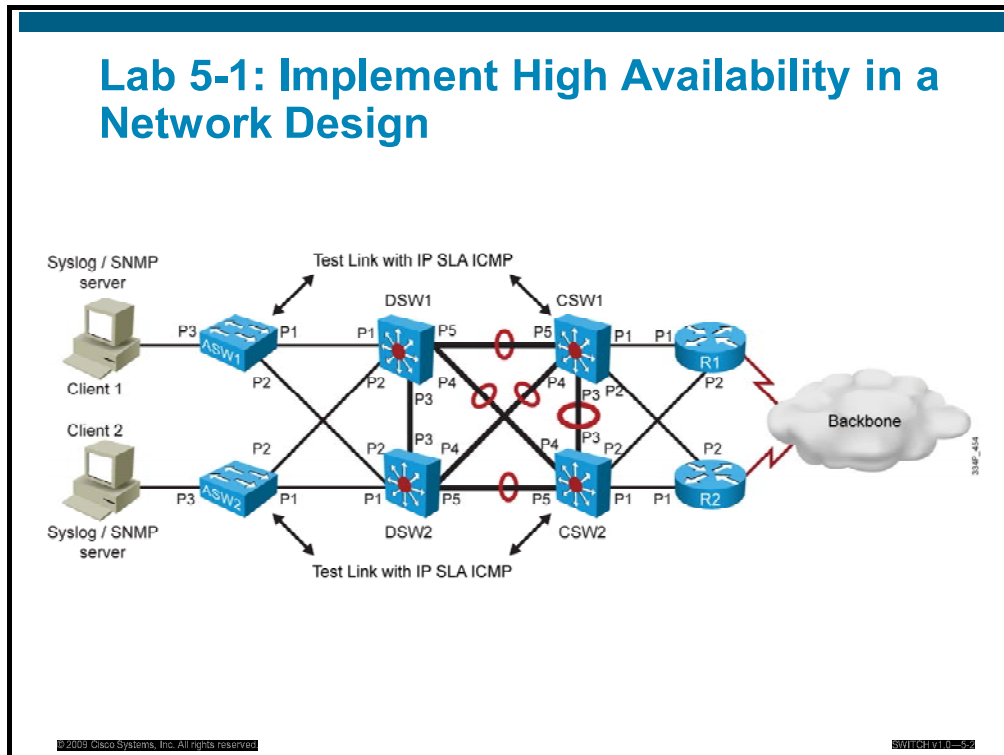
## Objectives

Upon completing this lesson, you will be able to configure logging to an external server, and configure SNMP parameters and IP SLA tests. This ability includes being able to meet these objectives:

- Review and verify your solution, as well as your findings and action log, against a set of checkpoints that are provided by the instructor
- Consolidate the lessons that you learned during the review discussions into a set of best-practice methods and commands to aid you in future deployment procedures

# Review and Verification

This topic describes the client requirements that were listed in Lab 5-1, asks how you can verify that you have identified the solution matching the client needs, and gives you an example of a possible solution.



Your lab consists of six switches and two routers. Your task is to configure these devices for global monitoring. In other words, in a large network, administrators do not have time to be connected to each device permanently to monitor all possible events. Global monitoring and management tools can be used, and useful information is extracted from networking devices and is sent to specific servers. In this lab, you configure syslog and SNMP parameters. You also test connectivity to a virtual point using IP SLA.

## Implement High Availability in a Network Design

Which items should be configured, and in which order?

- Where do you need syslog?
- How do you configure syslog?
- What needs to be configured for SNMP support?
- Which devices need SNMP configuration?
- Which devices are involved in IP SLA configuration?
- How did you configure IP SLA?

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCHING-55

These tasks are not related to each other. In the first task, you configure a syslog server and verify that messages are sent. In the second task, you configure an SNMP server and verify that traps are sent there. In the third task, you configure IP SLA. The IP SLA test could be used in combination with other configured items. For example, you could use the IP SLA test to verify that your switch can reach its gateway. If the connection is lost, the switch could revert to another gateway. This is the reason for high availability. When reverting to another gateway, the switch would also send a message to inform the administrator.



# Module Summary

This topic summarizes the key points that were discussed in this module.

## Module Summary

- High availability is a combination of several components, and involves network resiliency, device resiliency, power resiliency, and network monitoring.
- At the network level, you implement high availability by providing redundancy in connections paths between end devices.
- When high availability is in place, you can monitor device status states and connectivity by using tools such as SNMP, syslog, and IP SLA.



# Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which is a valid example of redundancy? (Source: Understanding High Availability)
- A) paths geographic diversity
  - B) server link isolation
  - C) duplicate IP addresses
  - D) several servers accessing the same router
- Q2) Which is an example of link redundancy? (Source: Understanding High Availability)
- A) redundant power supplies
  - B) redundant devices
  - C) redundant SNMP servers
  - D) redundant IP SLA tests
- Q3) Which is an example of system redundancy? (Source: Understanding High Availability)
- A) redundant power supplies
  - B) redundant devices
  - C) redundant network links
  - D) redundant IP SLA tests
- Q4) Which is an example of Layer 2 optimization for high availability? (Source: Understanding High Availability)
- A) replacing PVSTP with MSTP
  - B) disabling STP timers on nonroot bridges
  - C) replacing STP with RSTP
  - D) converting EtherChannel links to redundant links
- Q5) Which best describes Cisco NSF with SSO? (Source: Understanding High Availability)
- A) Cisco NSF with SSO is a routing protocol that allows extremely fast route update exchanges between Layer 3 switches and routers.
  - B) Cisco NSF with SSO is a CAM redundancy mechanism that allows extremely fast copy of CAM entries from one switch to the other.
  - C) Cisco NSF with SSO is a switch redundancy mechanism that allows extremely fast switchover between physical switches in standby mode.
  - D) Cisco NSF with SSO is a supervisor redundancy mechanism that allows extremely fast supervisor switchover at Layers 2 through 4.
- Q6) Why do Layer 2 distributed VLANs on access switches offer poor resiliency efficiency? (Source: Implementing High Availability)
- A) because first-hop redundancy cannot be implemented in this configuration
  - B) because STP convergence is slow
  - C) because STP cannot be configured in this configuration
  - D) because this configuration generates broadcast storms issues that cannot be solved with conventional STP mechanisms

- Q7) Which statement best describes daisy-chaining Layer 2 switches? (Source: Implementing High Availability)
- A) This configuration creates redundancy and additional security.
  - B) This configuration may create return traffic path issues.
  - C) This configuration allows isolating VLANs at the local switch level.
  - D) This configuration may create “train scenario” issues, in which traffic bounces across redundant links.
- Q8) Which command correctly defines the IP address of a syslog server listening at 10.1.1.1? (Source: Implementing Network Monitoring)
- A) **syslog 10.1.1.1 traps**
  - B) **monitor 10.1.1.1**
  - C) **logging 10.1.1.1**
  - D) **syslog-server 10.1.1.1 traps**
- Q9) Which command should you use to configure the read-only community xyz123? (Source: Implementing Network Monitoring)
- A) **snmp-community xyz123 RO**
  - B) **snmp-server community RO xyz123**
  - C) **snmp readonly-community xyz123**
  - D) **snmp-server community xyz123 RO**
- Q10) Which is the role of an IP SLA responder? (Source: Implementing Network Monitoring)
- A) answer to the IP SLA tests
  - B) store the IP SLA tests results
  - C) warn the administrator when the IP SLA test fails
  - D) send emergency and performance tests to other network devices



## Module Self-Check Answer Key

- Q1) A
- Q2) B
- Q3) A
- Q4) C
- Q5) D
- Q6) B
- Q7) B
- Q8) C
- Q9) D
- Q10) A



# Implementing Layer 3 High Availability

---

## Overview

A network with high availability provides an alternative means by which all infrastructure paths and key servers can be accessed at all times. The Hot Standby Router Protocol (HSRP) is one software feature that can be configured to provide Layer 3 redundancy to network hosts. HSRP optimization provides immediate or link-specific failover as well as a recovery mechanism. The Virtual Router Redundancy Protocol (VRRP) and Gateway Load Balancing Protocol (GLBP) are derivatives of HSRP, providing additional Layer 3 redundancy features such as load balancing.

## Module Objectives

Upon completing this module, you will be able to implement high-availability technologies and techniques using multilayer switches in a campus environment. This ability includes being able to meet these objectives:

- Configure HSRP to improve performance and resiliency in failover and recovery
- Configure VRRP and GLBP to improve performance and resiliency in failover and recovery



# Configuring Layer 3 Redundancy with HSRP

---

## Overview

Businesses and consumers that rely on intranet and Internet services for their mission-critical communications require and expect their networks and applications to be continuously available to them. Customers can satisfy their demands for near-100 percent network uptime if they use the Hot Standby Router Protocol (HSRP) in Cisco IOS Software. HSRP provides network redundancy for IP networks in a manner that ensures that user traffic immediately and transparently recovers from first-hop failures in network edge devices or access circuits. However, routing issues exist with various means of providing redundancy for the default gateway of each segment. Therefore, HSRP has very specific attributes that warrant further description, as does a delineation of HSRP operations on the network. HSRP interfaces transition through a series of states as they find their role in the capacity of active or standby HSRP router.

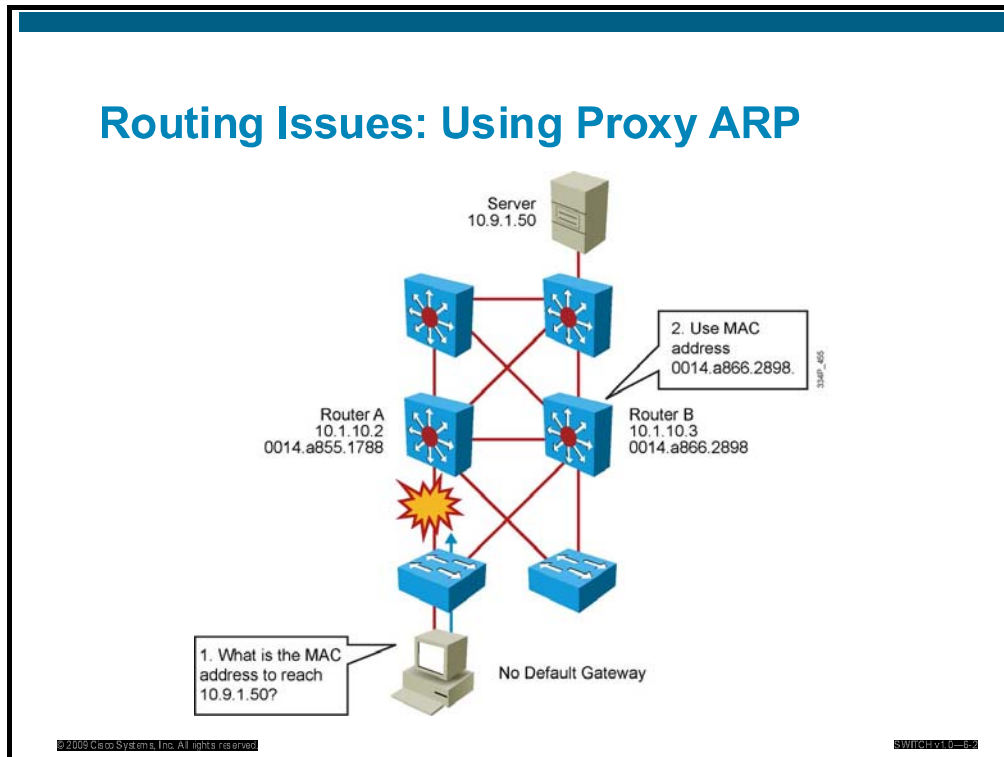
## Objectives

Upon completing this lesson, you will be able to configure Layer 3 redundancy with HSRP. This ability includes being able to meet these objectives:

- Describe routing issues
- Identify the router redundancy process
- Configure HSRP operations
- Describe and fine-tune HSRP
- Troubleshoot HSRP

# Describing Routing Issues

This topic describes routing issues that occur when you use default gateways and proxy Address Resolution Protocol (ARP).



## Using Proxy ARP

Before the default gateway was supported on most IP clients, networks were relying on the proxy ARP feature to reach IP devices outside the IP client subnet. Cisco IOS Software ran proxy ARP to enable hosts that had no knowledge of routing options to obtain the MAC address of a gateway that is able to forward packets off the local subnet.

For example, if the proxy ARP router receives an ARP request for an IP address that it knows is not on the same interface as the request sender, it will generate an ARP reply packet, giving its own local MAC address as the destination MAC address of the IP address that is being resolved. The host that sent the ARP request sends all packets that are destined for the resolved IP address to the MAC address of the router. The router then forwards the packets toward the intended host, perhaps repeating this process along the way. Proxy ARP is enabled by default.

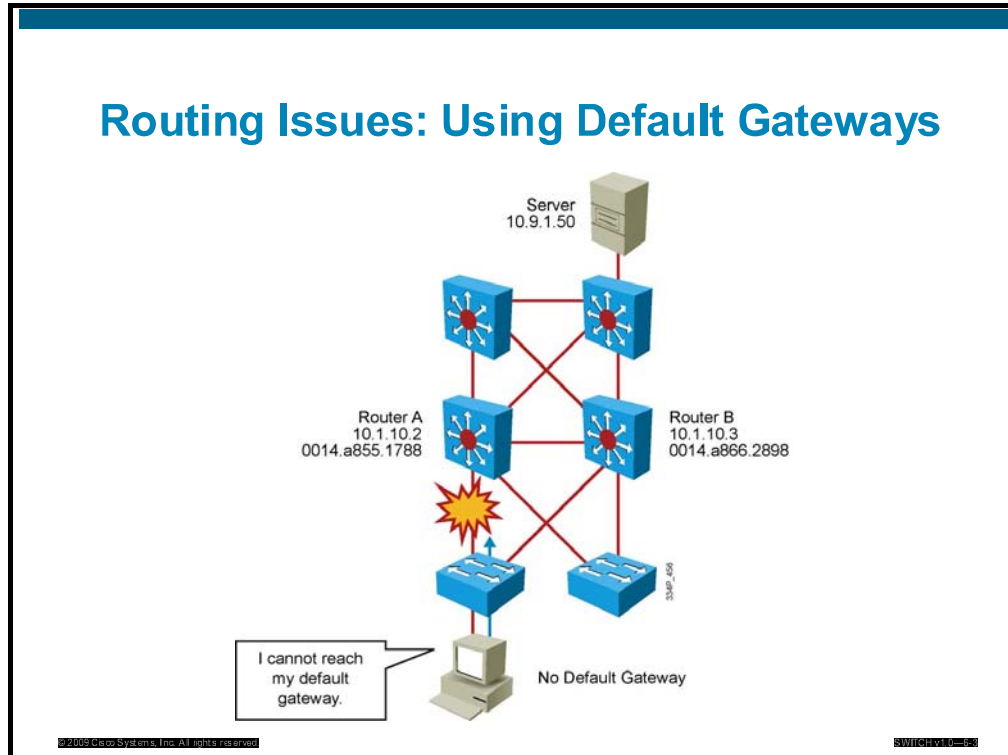
With proxy ARP, the end-user station behaves as if the destination device were connected to its own network segment. If the responsible router fails, the source end station continues to send packets for that IP destination to the MAC address of the failed router, and the packets are discarded.

Eventually, the proxy ARP MAC address will age out of the workstation's ARP cache. The workstation may eventually acquire the address of another proxy ARP failover router, but the workstation cannot send packets off the local segment during this failover time.

For further information on proxy ARP, refer to RFC 1027, *Using ARP to Implement Transparent Subnet Gateways*.

## Using Default Gateways

This subtopic describes routing issues that occur when you use default gateways.



Now that a default gateway is configured on most devices, the proxy ARP feature is not used anymore. Nevertheless, each client receives only one default gateway; there is no means by which to configure a secondary gateway, even if a second route exists to carry packets off the local segment.

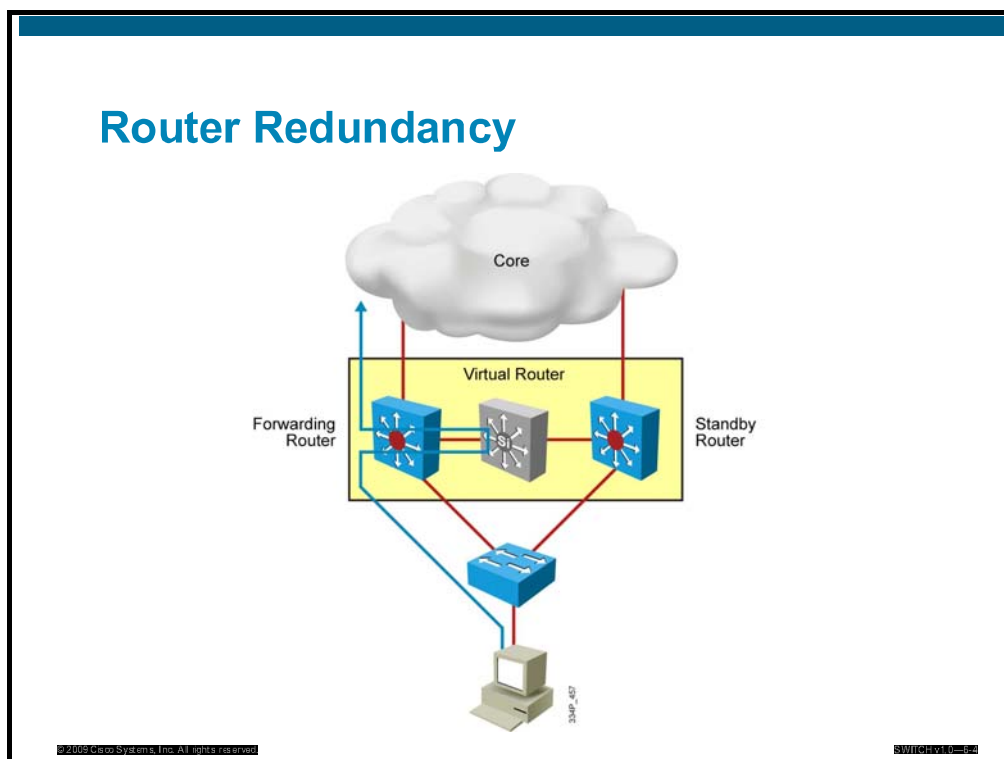
For example, primary and secondary paths between the building access submodule and the building distribution submodule provide continuous access in the event of a link failure at the building access layer. Primary and secondary paths between the building distribution layer and the building core layer provide continuous operations if a link fails at the building distribution layer.

In this example, Router A is responsible for routing packets for Subnet A, and Router B is responsible for handling packets for Subnet B. If Router A becomes unavailable, routing protocols can quickly and dynamically converge and determine that Router B will now transfer packets that would otherwise have gone through router A. Most workstations, servers, and printers, however, do not receive this dynamic routing information.

End devices are typically configured with a single default gateway IP address that does not change when network topology changes occur. If the router whose IP address is configured as the default gateway fails, the local device will be unable to send packets off the local network segment, effectively disconnecting it from the rest of the network. Even if a redundant router exists that could serve as a default gateway for that segment, there is no dynamic method by which these devices can determine the address of a new default gateway.

# Identifying the Router Redundancy Process

This topic describes how router device redundancy works.



With this type of router redundancy, a set of routers works together to present the illusion of a single virtual router to the hosts on the LAN. By sharing an IP address and a MAC (Layer 2) address, two or more routers can act as a single “virtual” router.

The IP address of the virtual router will be configured as the default gateway for the workstations on a specific IP segment. When frames are to be sent from the workstation to the default gateway, the workstation will use ARP to resolve the MAC address that is associated with the IP address of the default gateway. The ARP resolution will return the MAC address of the virtual router. Frames that are sent to the MAC address of the virtual router can then be physically processed by any active or standby router that is part of that virtual router group.

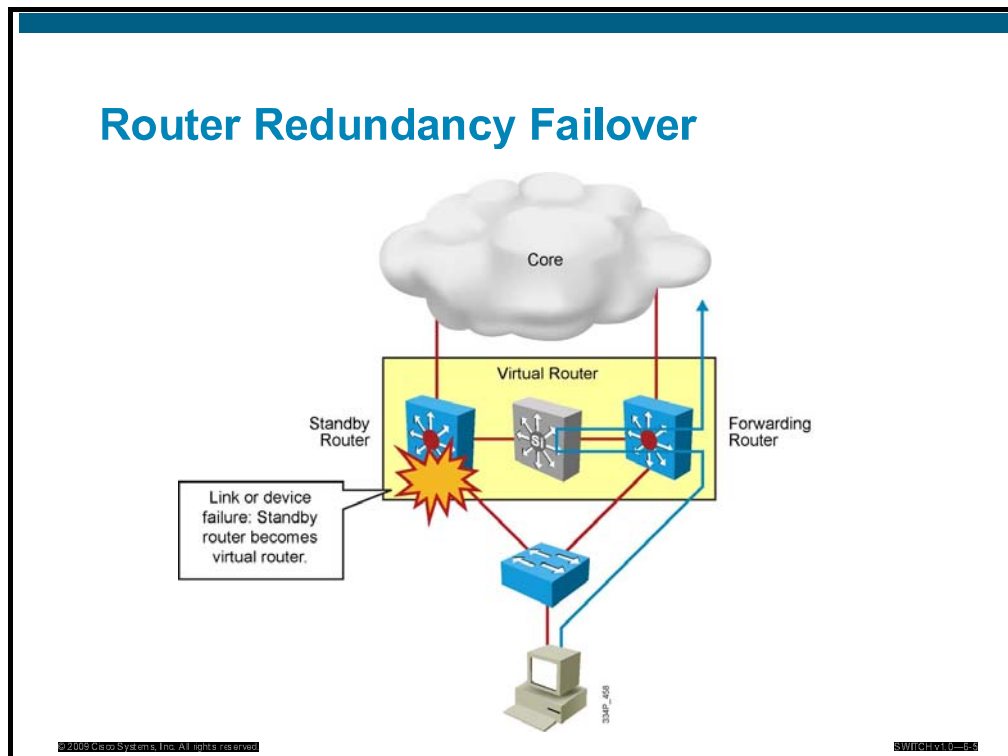
A protocol is used to identify two or more routers as the devices that are responsible for processing frames that are sent to the MAC or IP address of a single virtual router. Host devices send traffic to the address of the virtual router. The physical router that forwards this traffic is transparent to the end stations.

The redundancy protocol provides the mechanism for determining which router should take the active role in forwarding traffic and determining when that role must be taken over by a standby router. The transition from one forwarding router to another is transparent to the end devices.



# Router Redundancy Failover

This subtopic describes router device redundancy failover.



When the forwarding router or a link fails, this process occurs.

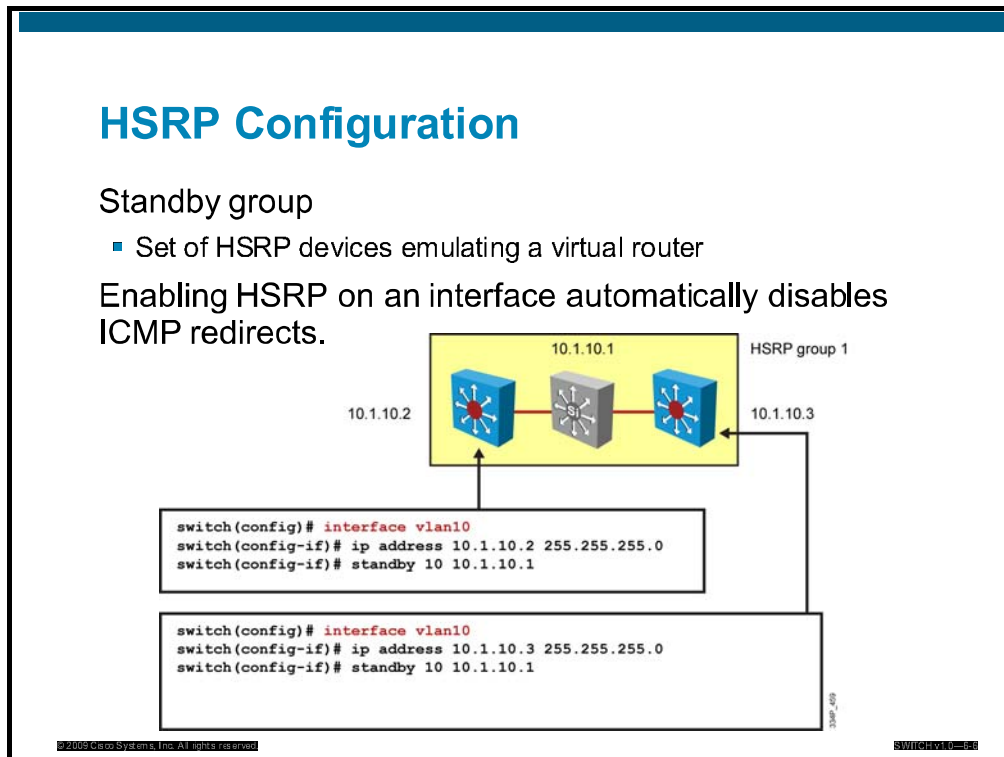
The table describes the steps that take place when a router fails.

## Router Redundancy Process

Step	Description
1.	The standby router stops seeing hello messages from the forwarding router.
2.	The standby router assumes the role of the forwarding router.
3.	Because the new forwarding router assumes both the IP and MAC addresses of the virtual router, the end stations see no disruption in service.

# HSRP Configuration

This topic describes HSRP configuration.



HSRP defines a standby group of routers, with one router as the active one. HSRP provides gateway redundancy by sharing IP and MAC addresses between redundant gateways. The protocol consists of virtual MAC and IP addresses that are shared between two routers that belong to the same HSRP group.

## HSRP Terminology

Term	Definition
Active router	The router that is currently forwarding packets for the virtual router
Standby router	The primary backup router
Standby group	The set of routers participating in HSRP that jointly emulate a virtual router

An HSRP group comprises these entities:

- One active router
- One standby router
- One virtual router
- Other routers

HSRP active and standby routers send hello messages to multicast address 224.0.0.2 User Datagram Protocol (UDP) port 1985.

## Commands Used to Configure and Verify HSRP

Command	Description
Switch(config-if)# <b>standby</b> <i>group-number</i> <b>ip</b> <i>ip-address</i>	Configures HSRP on this interface for this group number. IP address is that of the virtual gateway. Default group number is 0.
Switch(config-if)# <b>no standby</b> <i>group-number</i> <b>ip</b> <i>ip-address</i>	Disables HSRP on the interface.
Switch# <b>show running-config</b>	Displays HSRP parameters configured on each interface.
Switch# <b>show standby</b> [ <i>interface</i> ] [ <i>group</i> ] [ <i>brief</i> ]	The <b>show standby</b> command is all that is required. Use other commands to minimize output.

### Configuring HSRP Group on an Interface

This command enables HSRP on an interface:

```
Switch(config-if)#standby group-number ip ip-address
```

The following table describes the variables in the command that are used to configure an HSRP group on an interface.

#### HSRP Group Configuration Command

Variable	Definition
<i>group-number</i>	(Optional) Indicates the HSRP group to which this interface belongs. Specifying a unique group number in the <b>standby</b> commands enables the creation of multiple HSRP groups. The default group is 0.
<i>ip-address</i>	Indicates the IP address of the virtual HSRP router.

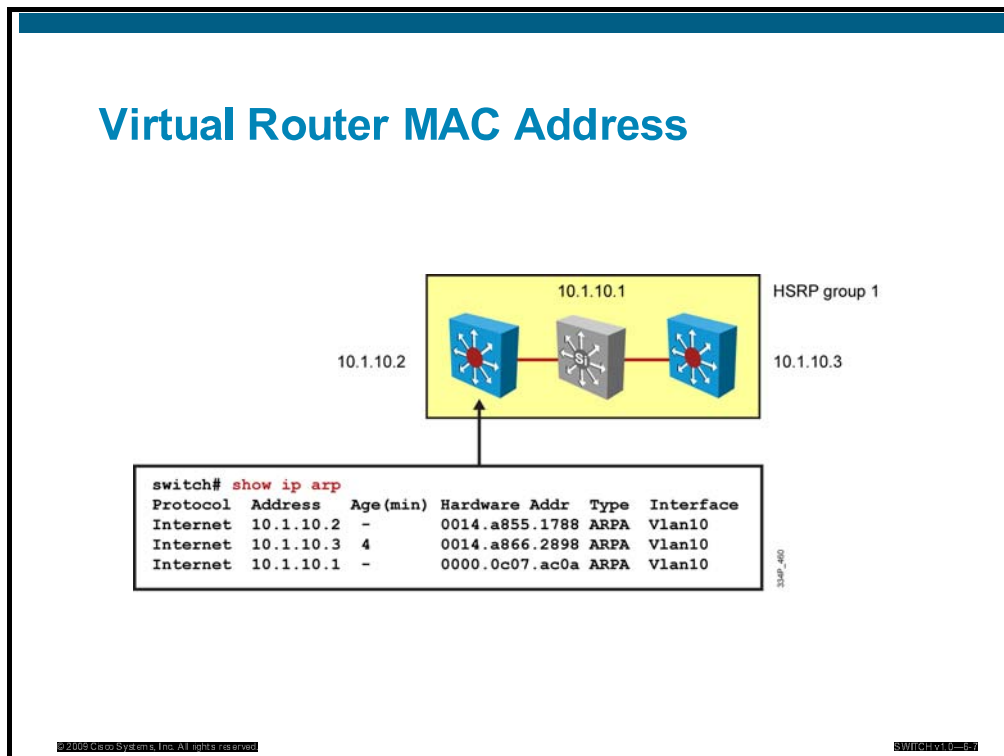
While running HSRP, the end-user stations must not discover the actual MAC addresses of the routers in the standby group. Any protocol that informs a host of a router actual address must be disabled. To ensure that the actual addresses of the participating HSRP routers are not discovered, you can enable HSRP on a Cisco router interface to automatically disable the Internet Control Message Protocol (ICMP) redirects on that interface.

After the **standby ip** command is issued, the interface changes to the appropriate state. When the router successfully executes the command, the router issues an HSRP message.

To remove an interface from an HSRP group, enter the **no standby group ip** command.

# Virtual Router MAC Address

This subtopic describes the virtual router MAC address.



The IP address and corresponding MAC address of the virtual router are maintained in the ARP table of each router in an HSRP group. As shown in the figure, the command **show ip arp** displays the ARP cache on a multilayer switch.

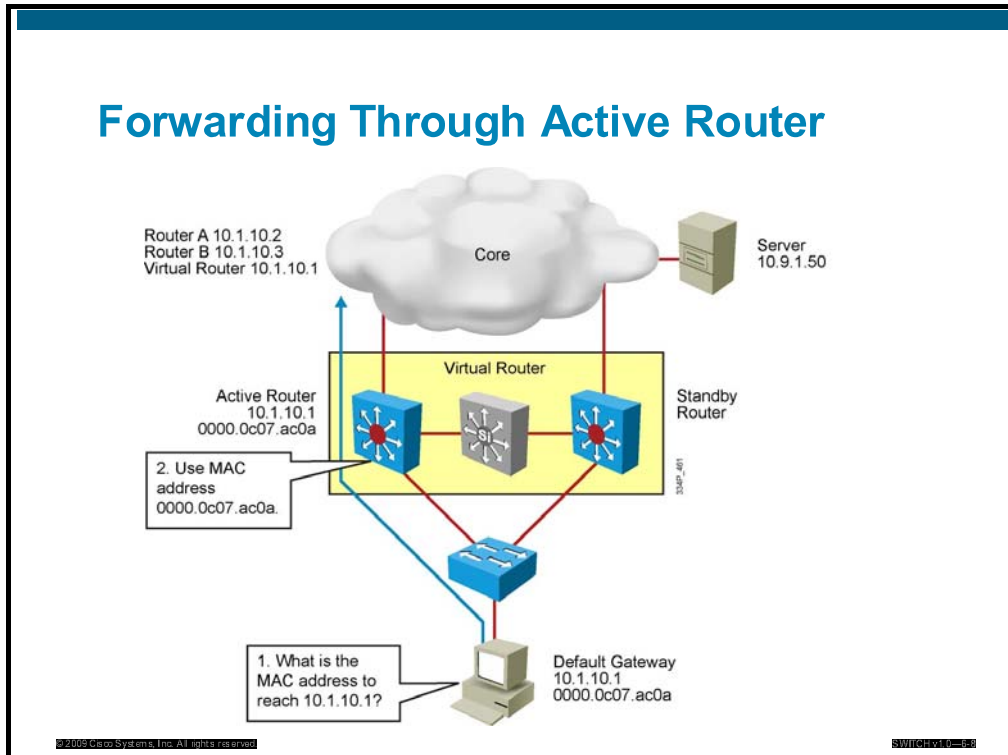
The table describes the command output for the **show ip arp** command.

## Interpretation of show ip arp Output

Field	Definition
Protocol	Protocol for network address in the Address field
Address	The network address that corresponds to the hardware address
Age (min)	Age, in minutes, of the cache entry
Hardware Addr	The MAC address that corresponds to the network address
Type	Type of encapsulation
Interface	Interface to which this address mapping has been assigned

# Forwarding Through Active Router

This subtopic describes how HSRP operates to provide a nonstop path redundancy for IP.



All the routers in an HSRP group have specific roles and interact in specific ways.

## Virtual Router

The virtual router is simply an IP and MAC address pair that end devices have configured as their default gateway. The active router will process all packets and frames sent to the virtual router address. The virtual router processes no physical frames.

## Active Router

Within an HSRP group, one router is elected to be the active router. The active router physically forwards packets that are sent to the MAC address of the virtual router.

The active router responds to traffic for the virtual router. If an end station sends a packet to the virtual router MAC address, the active router receives and processes that packet. If an end station sends an ARP request with the virtual router IP address, the active router replies with the virtual router MAC address.

In this example, router A assumes the active role and forwards all frames that are addressed to the well-known MAC address of 0000.0c07.acxx, where xx is the HSRP group identifier.

# HSRP Active and Standby Router Interaction

This subtopic describes the interaction between the active and standby routers.


## Active and Standby Routers

### Active router

- Responds to ARP requests of the default gateway with the MAC address of the virtual router
- Assumes the active forwarding of packets for the virtual router
- Sends hello messages
- Knows the virtual router IP address

### Standby router

- Listens for periodic hello messages
- Assumes the active forwarding of packets if no messages heard from active router



© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-001-001

When the active router fails, the other HSRP routers stop seeing hello messages from the active router. The standby router then assumes the role of the active router. If other routers are participating in the group, they then contend to be the new standby router.

In the event that both the active and standby routers fail, all routers in the group contend for the active and standby router roles.

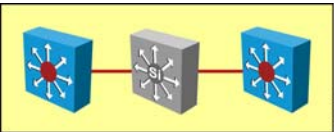
Because the new active router assumes both the IP and MAC addresses of the virtual router, the end stations see no disruption in service. The end-user stations continue to send packets to the virtual router MAC address, and the new active router delivers the packets to the destination.

# HSRP State Table

This subtopic describes the six HSRP states and their functions.

## HSRP States

An HSRP router can be in one of five states.



State	Definition
Initial	The state at the start. State after configuration change or when an interface first comes up.
Listen	The router knows the virtual IP address. It listens for hello messages from other routers.
Speak	The router sends periodic hello messages and actively participates in the election of the active or standby router.
Standby	The router is a candidate to become the next active router and sends periodic hello messages.
Active	The router currently forwards packets that are sent to the group virtual MAC address. The router sends periodic hello messages.

© 2009 Cisco Systems, Inc. All rights reserved.SWITCHING-10-151

A router in an HSRP group can be in one of five states: initial, listen, speak, standby, or active.

When a router exists in one of these states, it performs the actions required for that state. Not all HSRP routers in the group will transition through all states. For example, if there were three routers in the HSRP group, the router that is not the standby or active router will remain in the listen state.

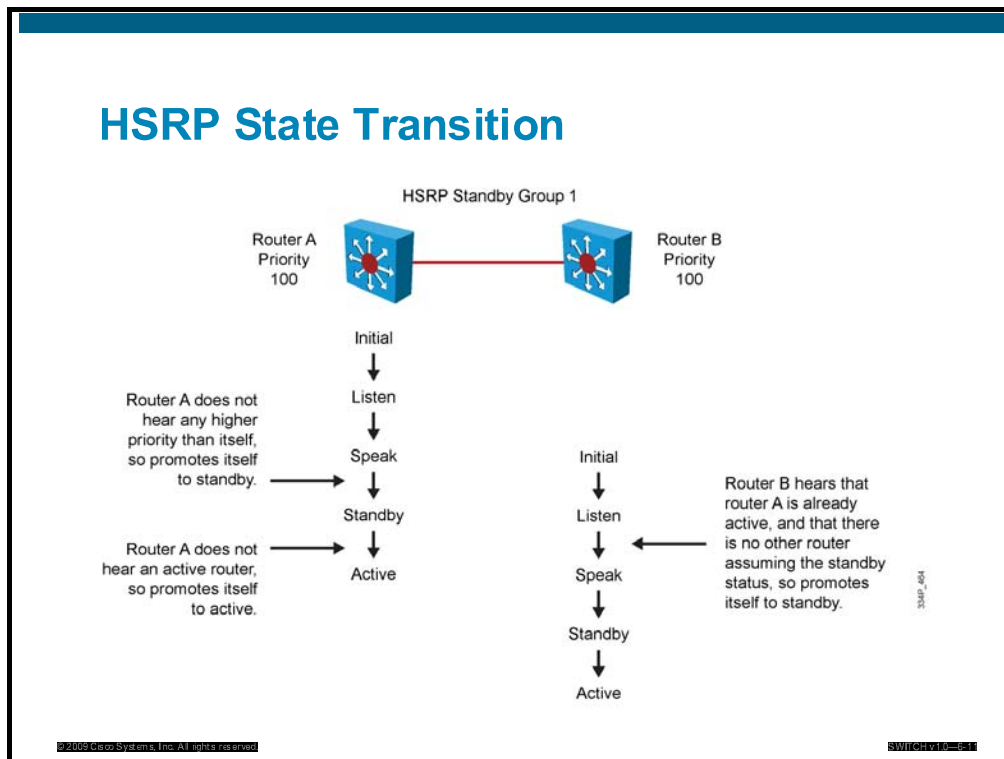
The table describes the different HSRP states.

## HSRP States

State	Definition
Initial	The beginning state. The initial state indicates that HSRP does not run. This state is entered via a configuration change or when an interface first comes up.
Listen	The router knows the virtual IP address, but the router is neither the active router nor the standby router. It listens for hello messages from those routers.
Speak	The router sends periodic hello messages and actively participates in the election of the active or standby router. A router cannot enter the speak state unless the router has the virtual IP address.
Standby	The router is a candidate to become the next active router and sends periodic hello messages. With the exclusion of transient conditions, there is, at most, one router in the group in the standby state.
Active	The router currently forwards packets that are sent to the group virtual MAC address. The router sends periodic hello messages. With the exclusion of transient conditions, there must be, at the most, one router in the active state in the group.

# HSRP State Transition

This subtopic describes the HSRP state transitions.



All routers begin in the initial state. This is the starting state and indicates that HSRP is not running. This state is entered via a configuration change, such as when HSRP is disabled on an interface, or when an HSRP-enabled interface is first brought up, such as when the **no shutdown** command is issued.

The purpose of the listen state is to determine if there are already active or standby routers for the group. In the speak state, the routers are actively participating in the election of the active router or standby router or both.

Each router uses three timers in HSRP. The timers time hello messages. When a timer expires, the router transitions to a new HSRP state.

In the example in the figure, router A starts. Because it is the first router for standby group 1 in the subnet, it transits through the listen and speak states, and then becomes the active router. Router B starts after A. While router B is in the listen state, router A is already assuming the standby role and then the active role. Because there is already an existing active router, router B assumes the standby role.

When two routers participate in an election process, a priority can be configured to determine which router should be active. Without specific priority configuration, each router has a default priority of 100, and the router with the highest IP address is elected as the active router.

Regardless of other router priorities or IP addresses, an active router will stay active by default. A new election will occur only if the active router is removed. When the standby router is removed, a new election is made to replace the standby. You can change this default behavior with the option **preempt**, which is examined later in this lesson.




# HSRP Priority and Pre-emption

This subtopic describes HSRP priority and pre-emption.

## HSRP Priority and Preemption

- The device with the highest priority in an HSRP group becomes the active router.
- The default priority is 100.
- In the case of a tie, the router with the highest configured IP address will become active.
- Preemption enables the higher priority device to become active.



```
switch(config)# interface vlan 10
switch(config-if)# ip address 10.1.1.2 255.255.255.0
switch(config-if)# standby 10 ip 10.1.1.1
switch(config-if)# standby 10 priority 110
switch(config-if)# standby 10 preempt
```

Each standby group has its own active and standby routers. The network administrator can assign a priority value to each router in a standby group, allowing the administrator to control the order in which active routers for that group are selected.

To set the priority value of a router, enter this command in interface configuration mode:

```
Switch(config-if)#standby group-number priority priority-value
```

The table describes the variables for the **standby** command.

### HSRP Standby Priority Configuration Commands

Variable	Definition
<i>group-number</i>	Indicates the HSRP group. This number can be in the range of 0 to 255.
<i>priority-value</i>	Indicates the number that prioritizes a potential hot standby router. The range is 0 to 255; the default is 100.

During the election process, the router with the highest priority in an HSRP group becomes the active router. In the case of a tie, the router with the highest configured IP address will become active.

To reinstate the default standby priority value, enter the **no standby priority** command.

If the routers do not have **preempt** configured, then a router that boots up significantly faster than the others in the standby group will become the active router, regardless of the configured priority. The former active router can be configured to resume the forwarding router role by pre-empting a router with a lower priority. To enable a router to resume the forwarding router role, enter this command in interface configuration mode:

```
Switch(config-if)#standby [group-number] preempt [{delay} [minimum  
delay] [sync delay]]
```

When you enter the **standby preempt** command, the interface changes to the appropriate state.

To remove the interface from pre-emptive status, enter the **no standby group preempt** command.


# HSRP Active Router and Spanning-Tree Topology

This subtopic describes the relationship between HSRP router election and spanning-tree root bridges.

## HSRP and STP

Configured active router should be the same as STP root bridge.

Blocked uplink caused traffic to take less-than-optimal path.



In a redundant spanning-tree topology, some links are blocked. The spanning-tree topology has no awareness of the standby configuration. There is no automatic relationship between the HSRP active router election process and the spanning-tree root bridge election process.

When configuring both the Spanning Tree Protocol (STP) and HSRP (or any other first-hop redundancy protocol), you should make sure that the active router is the same as the root bridge for the corresponding VLAN. When the root bridge is different from the HSRP active router, take some time to analyze the uplink path to the active router, to make sure that no suboptimal path is used.

# HSRP Authentication

This subtopic describes HSRP authentication.

## HSRP Authentication

- Configure authentication string (maximum of eight characters, default **cisco**) for standby group.

```
switch(config)# interface vlan 10
switch(config-if)# ip address 10.1.1.2 255.255.255.0
switch(config-if)# standby 10 ip 10.1.1.1
switch(config-if)# standby 10 priority 110
switch(config-if)# standby 10 preempt
switch(config-if)# standby 10 authentication xyz123
```

© 2009 Cisco Systems, Inc. All rights reserved.

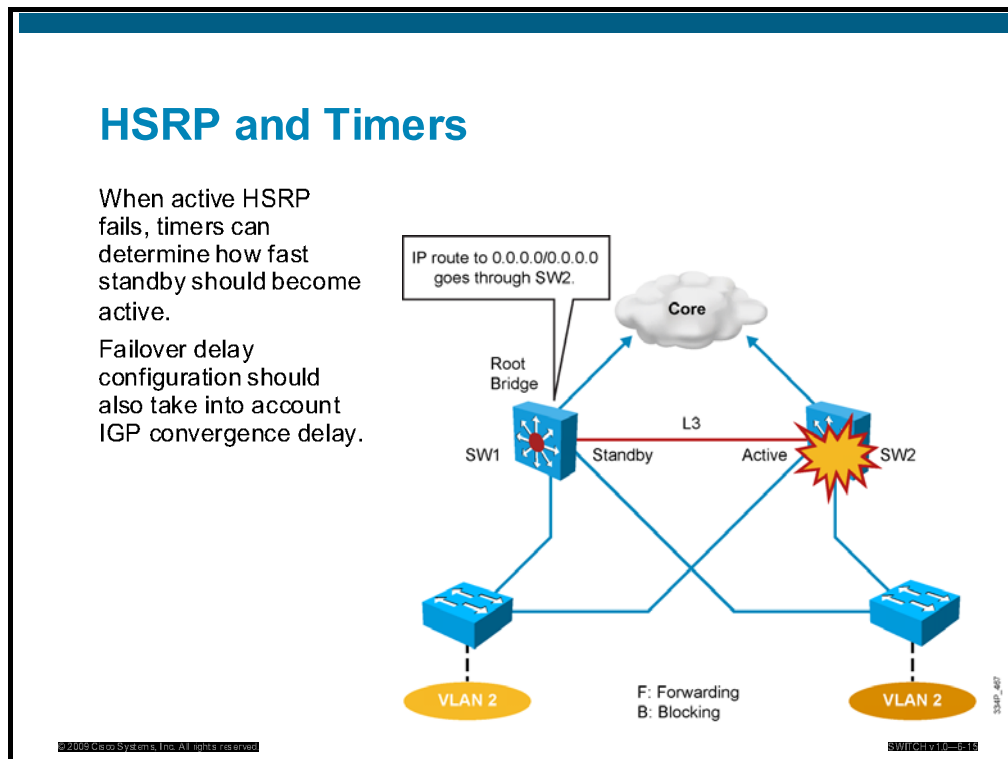
SWITCH-10-10

HSRP authentication prevents rogue routers on the network from joining the HSRP group.

You can enable HSRP authentication by configuring an authentication string on all member devices of the HSRP group. The string must be eight characters or fewer. If no string is configured, the default value of “cisco” is used.

# HSRP and Timer Considerations

This subtopic describes HSRP timers.



When an HSRP active router fails, the standby router detects the failure and assumes the active role. This mechanism relies on hello messages and holdtime intervals. The hello timer determines how often routers in the same standby group exchange messages.

The holdtime timer determines the time before the active or standby router is declared to be down.

Ideally, to achieve fast convergence, these timers should be configured to be as low as possible. Within milliseconds after the active router fails, the standby router can detect the failure, expire the holdtime interval, and assume the active role.

Nevertheless, timer configuration should also take into account other parameters that are relevant to the network convergence. For example, both HSRP routers may be running a dynamic routing protocol. The routing protocol probably has no awareness of the HSRP configuration, and sees both routers as individual hops toward other subnets. If HSRP failover occurs before the dynamic routing protocol converges, suboptimal routing information may still exist. In a worst-case scenario, the dynamic routing protocol continues seeing the failed router as the best next hop to other networks, and packets are lost. When configuring HSRP timers, make sure that they harmoniously match the other timers that can influence which path is chosen to carry packets in your network.

# HSRP Timer Configuration

This subtopic describes HSRP timer configuration.

## HSRP Timer Configuration

- Configure hello time and hold time to millisecond values.
- The hold time parameter value should be at least three times the value of the hello time parameter.
- Configure the preempt delay timer so that preemption occurs after the switch has fully rebooted and has established full connectivity to the network.

```
switch(config)# interface vlan 10
switch(config-if)# ip address 10.1.1.2 255.255.255.0
switch(config-if)# standby 10 ip 10.1.1.1
switch(config-if)# standby 10 priority 110
switch(config-if)# standby 10 preempt
switch(config-if)# standby 10 timers msec 200 msec 750
switch(config-if)# standby 10 preempt delay minimum 300
```

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-10-16-11

The hello message contains the priority of the router and hello time and holdtime parameter values. The hello time parameter value indicates the interval between the hello messages that the router sends. The holdtime parameter value indicates the amount of time that the current hello message is considered valid. The standby timer includes an **msec** parameter to allow for subsecond failovers. Lowering the hello timer results in increased traffic for hello messages and should be used cautiously.

If an active router sends a hello message, receiving routers consider that hello message to be valid for one hold time. The holdtime value should be at least three times the value of the hello time. The holdtime value must be greater than the value of the hello time.

You can adjust HSRP timers to tune the performance of HSRP on distribution devices, thereby increasing their resilience and reliability in routing packets off the local VLAN.

By default, HSRP hello time is 3 seconds and hold time is 10 seconds, which means that failover time could be as much as 10 seconds for clients to start communicating with the new default gateway. In some cases, this interval may be excessive for application support. The hello time and hold time parameters are configurable. To configure the time between hello messages and the time before other group routers declare the active or standby router to be nonfunctioning, enter this command in interface configuration mode:

```
Switch(config-if)# standby group-number timers [msec] hellotime
holdtime
```

---

**Note** Hello and dead timer intervals must be identical for all devices within the HSRP group.

---

The table describes the options for standby message timer configuration.

### Standby Message Timer Configuration Options

Variable	Description
<i>group-number</i>	(Optional) Group number on the interface to which the timers apply. The default is 0.
<i>hellotime</i>	Hello interval in seconds. This is an integer from 1 through 255. The default is 3 seconds.
<i>holdtime</i>	Time, in seconds, before the active or standby router is declared to be down. This is an integer from 1 through 255. The default is 10 seconds.

To reinstate the default standby-timer values, enter the **no standby group timers** command. The table above describes the HSRP timers.

## Subsecond Failover

The HSRP hello time and hold time can be set to millisecond values so that HSRP failover occurs in less than 1 second. Here is an example:

```
Switch(config-if)#standby 1 timers msec 200 msec 750
```

## Pre-empt Time Aligned with Router Boot Time

Pre-emption is an important feature of HSRP that allows the primary router to resume the active role when it comes back online after a failure or maintenance event. Pre-emption is a desired behavior because it forces a predictable routing path for the VLAN during normal operations and ensures that the Layer 3 forwarding path for a VLAN parallels the Layer 2 Spanning Tree Protocol (STP) forwarding path whenever possible.

When a pre-empting device is rebooted, HSRP pre-emption communication should not begin until the distribution switch has established full connectivity to the rest of the network. This situation allows the routing protocol convergence to occur more quickly, after the preferred router is in an active state.

To accomplish this, measure the system boot time and set the HSRP pre-emption delay to a value that is 50 percent greater than the boot time. This value ensures that the primary distribution switch establishes full connectivity to the network before HSRP communication occurs.

For example, if the boot time for the distribution device is 150 seconds, the **preempt** configuration would appear as follows:

```
standby 10 preempt
standby 10 preempt delay minimum 300
```

# HSRP Versions

This subtopic describes the HSRP versions.

## HSRP Versions

HSRPv1 (default).

- Group number from 0 to 255.
- Virtual MAC address 0000.0C07.ACXX (XX = HSRP group).
- Hello packets sent to multicast address 224.0.0.2.

HSRPv2 (Cisco IOS Release 12.2(46)SE and later).

- Group number from 0 to 4095.
- Virtual MAC address 0000.0C9F.FXXX (XXX = HSRP group).
- Hello packets sent to multicast address 224.0.0.102.
- HSRPv2 and HSRPv1 have different packet formats.

Configure same version on all devices of HSRP group.

```
switch(config-if)# standby 10 version 2
```

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-10-6-17

A new version of HSRP has been added to Cisco IOS Software.

The new version allows group numbers up to 4095. This allows you to use the VLAN number as the group number.

Therefore, the packet format, the MAC address of the virtual router, and the multicast address for the hello messages have been changed.

On all routers in an HSRP group the same version must be configured; otherwise, hello messages are not understood. Version 1 is the default.



# Displaying the Standby Status

This subtopic describes displaying the standby status.

## Displaying the Standby Status

```
switch# show standby brief
                P indicates configured to preempt.
                |
Interface   Grp  Pri P State   Active        Standby        Virtual IP
Vl10         10   110 P Active   local         10.1.10.3     10.1.10.1

switch# show standby
Vlan10 - Group 10
  State is Active
    1 state change, last state change 00:00:43
  Virtual IP address is 10.1.10.1
  Active virtual MAC address is 0000.0c07.ac0a
    Local virtual MAC address is 0000.0c07.ac0a (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.016 secs
  Preemption enabled
  Active router is local
  Standby router is 10.1.10.3, priority 90 (expires in 11.328 sec)
  Priority 110 (configured 110)
  Group name is "hsrp-Vl10-10" (default)
```

© 2009 Cisco Systems, Inc. All rights reserved.SWITCH10-151

The function of the HSRP standby router is to monitor the operational status of the HSRP group and to quickly assume packet-forwarding responsibility if the active router becomes inoperable. Both the active and standby routers transmit hello messages to inform all other routers in the group of their role and status. The routers use multicast address 224.0.0.2 UDP port 1985 for these messages.

An HSRP group may contain other routers that are group members but are not in an active or standby state. These routers monitor the hello messages that are sent by the active and standby routers to ensure that an active and standby router exists for the HSRP group of which they are members. These routers do forward packets addressed to their own specific IP addresses, but they do not forward packets that are addressed to the virtual router. These routers issue speak messages at every hello interval time.

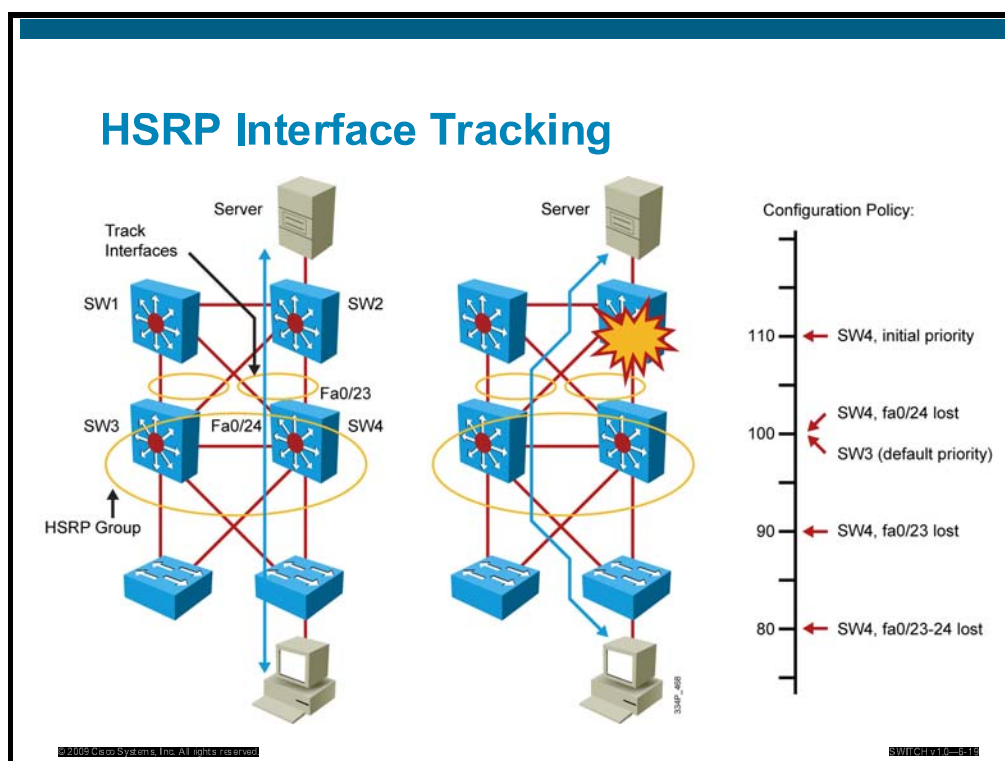
The table describes some of the terms that are used with HSRP.

## HSRP Terminology

Term	Definition
Hello interval time	The interval between successive HSRP hello messages from a given router. Default = 3 seconds.
Hold interval time	The interval between the receipt of a hello message and the presumption that the sending router has failed. Default = 10 seconds.

# HSRP Interface Tracking

This subtopic describes HSRP interface tracking.



In the figure, the distribution switches monitor the uplink to the core switches. The uplink between the active forwarding device for the standby group and the core experiences a failure. Without HSRP enabled, the active device would detect the failed link and send an Internet Control Message Protocol (ICMP) redirect to the other device. However, when HSRP is enabled, ICMP redirects are disabled. The left switch now has the better path to the server.

Interface tracking enables the priority of a standby group router to be automatically adjusted, based on the availability of the router interfaces. When a tracked interface becomes unavailable, the HSRP priority of the router is decreased. When properly configured, the HSRP tracking feature ensures that a router with an unavailable key interface will relinquish the active router role.

The HSRP group tracks the uplink interfaces. If the uplink to the core on the right switch fails, the router automatically decrements the priority on that interface and sends hello messages with the decremented priority. The switch on the left now has a higher priority and with pre-emption enabled becomes the active router.

A router can track several interfaces. In the figure, SW4 tracks both fa0/23 and fa0/24. The configuration policy states that SW4 initial priority should be 110. SW3 initial priority should be left to its default value, 100. If SW4 loses its link fa0/24 to SW1, SW4 priority should become the same as SW3 priority. If a new election has to occur, both multilayer switches have the same chances of becoming the active router. This decrement is made because fa0/24 is not the active link, but just a backup. If fa0/23 (the active uplink) is lost, then SW4 priority becomes lower than SW3 priority. If both fa0/23 and fa0/24 are lost, both decrements are applied and SW4 priority becomes 80.

## HSRP Interface Tracking

- Configure the standby group.
- Configure priority (default is 100).
- Configure preempt on all devices within the HSRP group.
- Configure the tracked interfaces and decrement (default decrement is 10).

```
switch(config)# interface vlan 10
switch(config-if)# ip address 10.1.1.2 255.255.255.0
switch(config-if)# standby 10 ip 10.1.1.1
switch(config-if)# standby 10 priority 110
switch(config-if)# standby 10 preempt
switch(config-if)# standby 10 track fastethernet0/23 20
switch(config-if)# standby 10 track fastethernet0/24
```

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-10-520

The table describes the variables in the HSRP configuration command.

### HSRP Tracking Configuration Arguments

Variable	Description
<i>group-number</i>	(Optional) Indicates the group number on the interface to which the tracking applies. The default number is 0.
<i>type</i>	Indicates the interface type (combined with the interface number) that will be tracked.
<i>number</i>	Indicates the interface number (combined with the interface type) that will be tracked.
<i>interface-priority</i>	(Optional) Indicates the amount by which the hot standby priority for the router is decremented when the interface becomes disabled. The priority of the router is incremented by this amount when the interface becomes available. The default value is 10.

To disable interface tracking, enter the **no standby group track** command.

The command to configure HSRP tracking on a multilayer switch is the same as on the external router, except that the interface type can be identified as a switch virtual interface (**vlan** followed by the vlan number that is assigned to that interface) or by a physical interface.

The internal routing device uses the same command as the external routing device to disable interface tracking.

You can apply multiple tracking statements to an interface. This may be useful if, for example, the currently active HSRP interface will relinquish its status only upon the failure of two (or more) tracked interfaces.

# HSRP and Tracking

This subtopic describes HSRP and tracking.

## Tracking Options

```
switch(config)# track 1 ?
  interface  Select an interface to track
  ip         IP protocol
  list       Group objects in a list
  rtr        Response Time Reporter (RTR) entry
```

**Standby** command can track interface or object.

Tracked objects are defined with **track object\_number**.

- **Track number interface** can verify line protocol or IP routing capability.
- **Track number ip route** can verify network reachability.
- **Track number list** can define complex conditions.
- **Track number rtr** can verify parameters such as IP SLA.

When configuring HSRP, you can define a track option to monitor the status of an interface. When the interface fails, the priority decreases. You can also track an object. When the conditions that are defined by this object are fulfilled, the router priority remains the same. As soon as the verification that is defined by the object fails, the router priority is decremented.

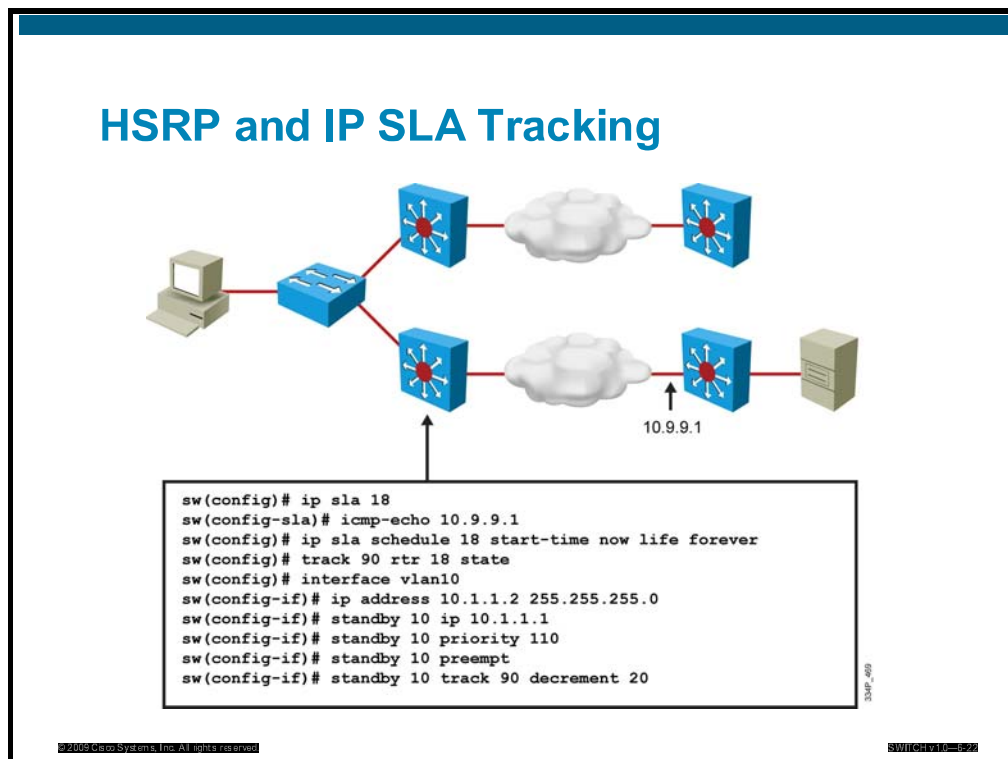
Tracked objects are defined in global configuration with the keyword **track**, followed by an object number. You can track up to 500 objects.

Tracked objects offer a vast group of possibilities. You can track the following:

- **An interface:** Just like the **standby track interface** command, a tracking object can verify the interface status (line protocol). You can also track IP routing on the interface. This option tracks whether IP routing is enabled, whether an IP address is configured on the interface, and whether the interface state is up before reporting to the tracking client that the interface is up.
- **IP route:** A tracked IP-route object is considered up and reachable when a routing-table entry exists for the route and the route is accessible. To provide a common interface to tracking clients, route metric values are normalized to the range of 0 to 255, where 0 is connected and 255 is inaccessible. You can track route reachability, or even metric values, to determine best-paths values to the target network. The tracking process uses a per-protocol configurable resolution value to convert the real metric to the scaled metric. The metric value that is communicated to clients is always such that a lower metric value is better than a higher metric value.
- **A list of objects:** You can track several objects and compare their results to determine if one or several of them should trigger the “success” or “fail” condition.
- **IP SLA:** This special case allows you to track advanced parameters such as IP reachability, delay, or jitter.

# HSRP and IP SLA Tracking

This subtopic describes HSRP and IP SLA tracking.



In the figure, a Cisco IOS IP Service Level Agreement (SLA) measurement is being run between two switches across a network cloud.

If the link fails, the priority of the active switch in the HSRP group is reduced and the other switch connection via the upper network becomes the active router to reach the server.

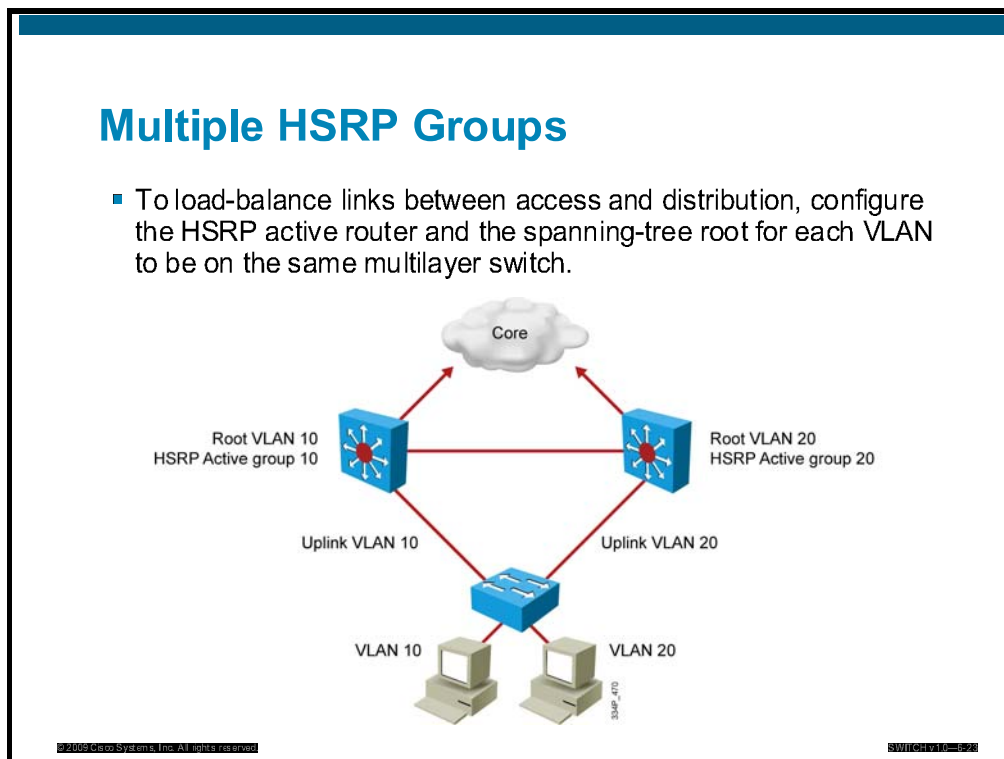
IP SLA tracking extends the HSRP interface tracking to allow you to track paths through the network.

Configuration steps are as follows:

- Step 1** Create an IP SLA process (18).
- Step 2** Schedule this IP SLA process.
- Step 3** Create an object (90) to track the state of this process.
- Step 4** Track the state of this object and decrement the HSRP device priority if the object fails.

# Multiple HSRP Groups

This subtopic describes multiple HSRP groups.



Routers can simultaneously provide redundant backup and perform load sharing across different IP subnets.

In the figure, two HSRP-enabled routers participate in two separate VLANs, using IEEE 802.1Q. Running HSRP over trunking allows users to configure redundancy among multiple routers that are configured as front ends for VLAN IP subnets.

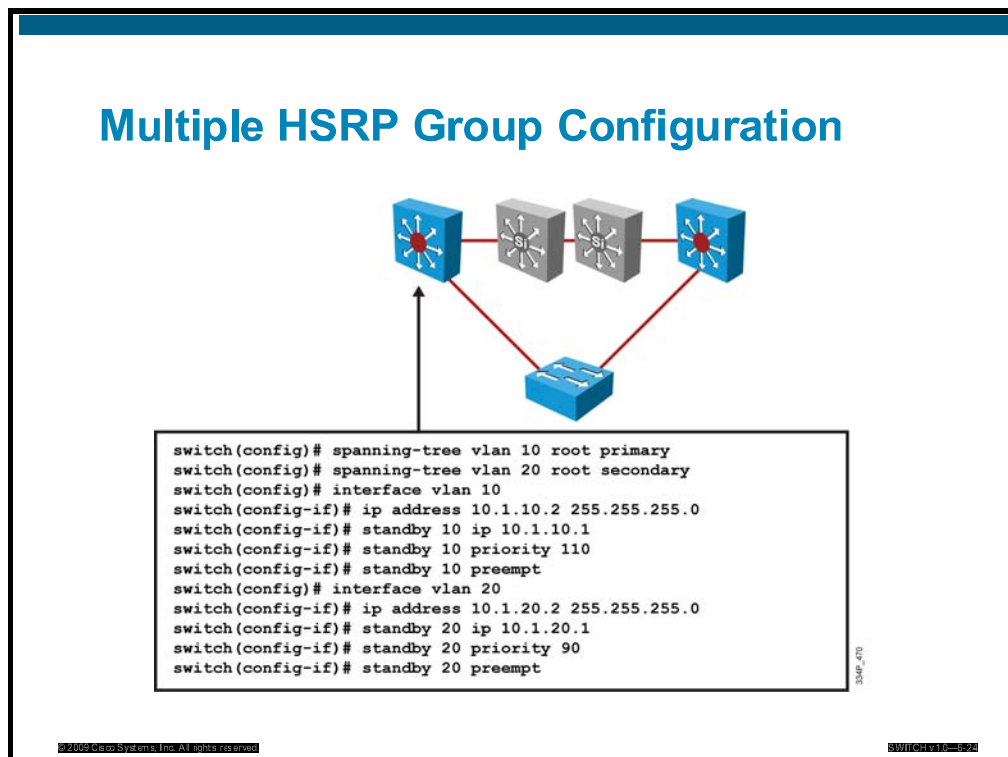
By configuring HSRP over trunks, you can eliminate situations in which a single point of failure causes traffic interruptions. This feature inherently provides some improvement in overall networking resilience by providing load-balancing and redundancy capabilities between subnets and VLANs.

For a VLAN, configure the same device to be both the spanning-tree root and the HSRP active router. This approach ensures that the Layer 2 forwarding path leads directly to the Layer 3 active router and so achieves maximum efficiency of load balancing on the routers and the trunks.

For each VLAN, a standby group, an IP address, and a single well-known MAC address with a unique group identifier is allocated to the group. Although up to 255 standby groups can be configured (4095 with version 2), it is advised that the actual number of group identifiers that are used be kept to a minimum. When you are configuring two distribution layer switches, typically you will require only two standby group identifiers, regardless of how many standby groups are created.

# Multiple HSRP Group Configuration

This subtopic describes multiple HSRP group configuration.



The figure shows the configuration for two HSRP groups for two VLANs and the corresponding STP root configuration.

The left switch is the root and is the active HSRP router for VLAN 10.

The corresponding configuration of the right switch has the switch as the root and as the active HSRP router for VLAN 20.


# HSRP Monitoring

This subtopic describes HSRP monitoring.

## Monitoring HSRP

```
SW1#show standby brief
          P indicates configured to preempt.
          |
Interface   Grp  Pri P State   Active    Standby    Virtual IP
Vl63        63   120 P Active local    10.1.63.2  10.1.63.254
Vl64        64   90  P Standby 10.1.64.1 local    10.1.64.254

SW1#show standby neighbor vlan64
HSRP neighbors on Vlan64
  10.1.64.1
    Active groups: 64
    No standby groups
```



Use the **show standby** family of commands to verify the HSRP state. Several arguments can be used. The command **show standby brief** simply displays a summary of the HSRP configurations. For each standby group, you can verify the local router neighbors.



The table that follows the output describes commands that are used to debug HSRP. When you simply type **show standby**, a complete display is provided. In this example, SW1 and SW2 are HSRP routers for VLANs 63 and 64. In each VLAN, SW1 has an IP address ending with .1, and SW2 has an IP address ending with .2. Both routers emulate the .254 gateway:

```
SW1#sh standby
Vlan63 - Group 63
  State is Active
  Virtual IP address is 10.1.63.254
  Active virtual MAC address is 0000.0c07.ac3f
    Local virtual MAC address is 0000.0c07.ac3f (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.248 secs
  Preemption enabled
  Active router is local
  Standby router is 10.1.63.2, priority 90 (expires in 10.096 sec)
  Priority 120 (configured 120)
    Track interface Port-channel31 state Up decrement 30
    Track interface Port-channel32 state Up decrement 30
  Group name is "hsrp-Vl63-63" (default)
Vlan64 - Group 64
  State is Standby
  Virtual IP address is 10.1.64.254
  Active virtual MAC address is 0000.0c07.ac40
    Local virtual MAC address is 0000.0c07.ac40 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.064 secs
  Preemption enabled
  Active router is 10.1.64.1, priority 120 (expires in 10.032 sec)
  Standby router is local
  Priority 90 (configured 90)
  Group name is "hsrp-Vl64-64" (default)
```

## HSRP Debug Commands

Command	Description
Switch# <b>debug standby [errors] [events] [packets]</b>	Displays all state changes to HSRP, including all hello packets. Arguments minimize output.
Switch# <b>debug standby terse</b>	Displays all HSRP errors, events, and packets, except hello and advertisement packets.

The output shows the sent (out) and received (in) hello messages for VLAN 10 and the active router for HSRP group 10.

# Summary

This topic summarizes the key points that are discussed in this lesson.

## Summary

- A single default gateway or proxy ARP does not provide the redundancy that is required in a campus network.
- HSRP provides router redundancy to end devices.
- HSRP is configured using the standby command for each interface.
- Preemption, timers, and interface tracking are options that can be configured to optimize HSRP and reduce failover time.
- Specific debug commands are used to view the HSRP state changes.

# Configuring Layer 3 Redundancy with VRRP and GLBP

---

## Overview

The Virtual Router Redundancy Protocol (VRRP) provides router interface failover in a manner similar to that of the Hot Standby Router Protocol (HSRP) but with added features and IEEE compatibility. The process by which VRRP operates is defined in this lesson. The Gateway Load Balancing Protocol (GLBP) and its operations will be defined and differentiated from both HSRP and VRRP. Specific commands are used to implement and verify VRRP and GLBP.

## Objectives

Upon completing this lesson, you will be able to configure Layer 3 redundancy with VRRP and GLBP. This ability includes being able to meet these objectives:

- Describe VRRP
- Identify the VRRP operations process
- Configure VRRP
- Describe GLBP
- Identify the GLBP operations process
- Configure GLBP

# Virtual Router Redundancy Protocol

This topic describes VRRP.

## HSRP vs. VRRP

HSRP	VRRP
Cisco proprietary, 1994.	IETF 1998–2005, RFC 3768.
16 groups max.	255 groups max.
1 active, 1 standby, several candidates.	1 active, several backups.
Virtual IP is different from active and standby real IP addresses.	Virtual IP address can be the same as the real IP address of one of the group members.
Uses 224.0.0.2.	Uses 224.0.0.18.
Can track interfaces or objects.	Can track only objects.
Default timers: hello, 3 sec; hold time, 10 sec.	Default timers: hello, 1 sec; hold time, 3 sec.
Authentication supported.	Authentication no longer supported.

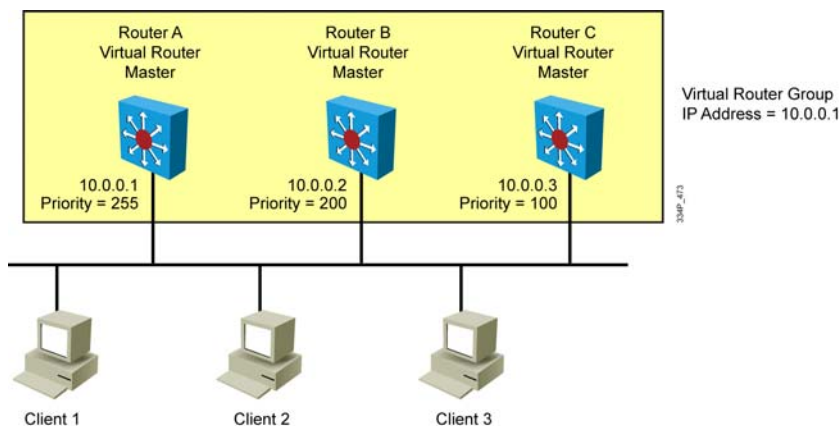
Like HSRP, VRRP allows a group of routers to form a single virtual router. In an HSRP or VRRP group, one router is elected to handle all requests sent to the virtual IP address. With HSRP, this is the active router. A HSRP group has one active router, one standby router, and perhaps many listening routers. A VRRP group has one master router and one or more backup routers.

The LAN workstations are then configured with the address of the virtual router as their default gateway. VRRP differs from HSRP in these ways:

- VRRP is an IEEE standard (RFC 2338 in 1998, and then RFC 3768 in 2005) for router redundancy; HSRP is a Cisco proprietary protocol, created in 1994 and formalized with the RFC 2281 in March 1998.
- In VRRP, the virtual router, representing a group of routers, is known as a VRRP group.
- In VRRP, the active router is referred to as the master virtual router.
- In VRRP, the master virtual router may have the same IP address as the virtual router group.
- In VRRP, multiple routers can function as backup routers.
- Intragroup communications use multicast IP address 224.0.0.2 for HSRP, and 224.0.0.18 for VRRP.
- Both HSRP and VRRP can track objects. HSRP can also directly track an interface status, whereas VRRP cannot directly track an interface status. Interfaces can be tracked with VRRP through a tracked object.
- The default timers are shorter in VRRP than HSRP. This fact often gave VRRP the reputation of being faster than HSRP. In reality, the convergence speed in the case of failover depends on the actual timer configuration.
- HSRP uses authentication within each group by default. When authentication is not configured, a default authentication, using “cisco” as the password, is actually used. VRRP used to support plaintext and Hashed Message Authentication Code-Message Digest 5 (HMAC-MD5) authentication methods (RFC 2338). The new VRRP RFC (RFC 3768) removes support for these methods. The consequence is that VRRP does not support authentication anymore. Nevertheless, current Cisco IOS Software still supports the RFC 2338 authentications mechanisms.

Beyond these differences, HSRP and VRRP are very similar in their features and behaviors. The main difference is that HSRP is a Cisco proprietary implementation whereas VRRP is an open standard. The consequence is that HSRP is usually found in Cisco networks. VRRP is used in multivendor implementations.

## About VRRP



In the example, routers A, B, and C are members of a VRRP group. The IP address of the virtual router is the same as that of the LAN interface of router A (10.0.0.1). Router A is responsible for forwarding packets that are sent to this IP address.

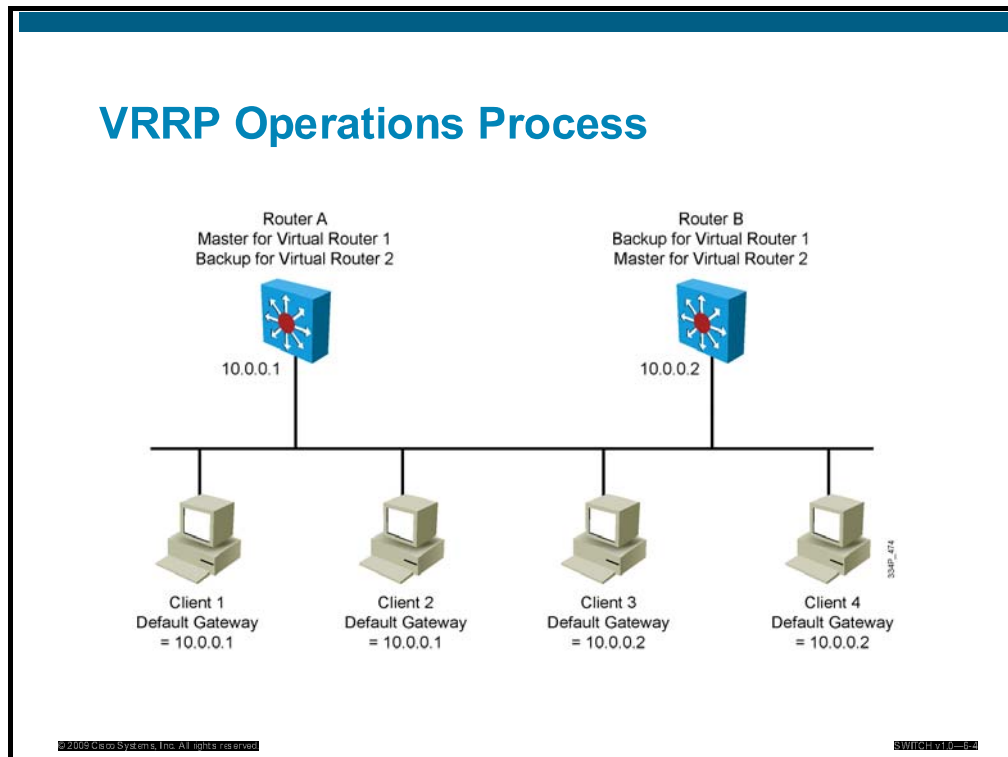
The clients have a gateway address of 10.0.0.1. Routers B and C are backup routers. If the master router fails, the backup router with the highest priority becomes the master router. When router A recovers, it resumes the role of master router.

VRRP offers these redundancy features:

- VRRP provides redundancy for the real IP address of a router or for a virtual IP address that is shared among the VRRP group members.
- If a real IP address is used, the router with that address becomes the master. If a virtual IP address is used, the master is the router with the highest priority.
- A VRRP group has one master router and one or more backup routers. The master router uses VRRP messages to inform group members that it is the master.

# About the VRRP Operations Process

This topic describes VRRP operations.



This figure shows a LAN topology in which VRRP is configured so that routers A and B share the load of being the default gateway for clients 1 through 4. Routers A and B act as backup virtual routers to one another should either one fail.

In this example, two virtual router groups are configured. For virtual router 1, router A is the owner of IP address 10.0.0.1 and is therefore the master virtual router for clients that are configured with that default gateway address. Router B is the backup virtual router to router A.

For virtual router 2, router B is the owner of IP address 10.0.0.2 and is the master virtual router for clients that are configured with the default gateway IP address 10.0.0.2. Router A is the backup virtual router to router B.

Given that the IP address of the VRRP group is that of a physical interface on one of the group members, the router owning that address will be the master in the group. Its priority is set to 255. Backup router priority values can range from 1 to 254; the default value is 100. The priority value 0 has special meaning, indicating that the current master has stopped participating in VRRP. This setting is used to trigger backup routers to quickly transition to the master without having to wait for the current master to time out.

With VRRP, only the master sends advertisements (the equivalent of HSRP hellos). The master sends the advertisement on multicast 224.0.0.18 protocol number 112 on a default interval of 1 second.

## VRRP Transition Process

The dynamic failover, when the active (master) becomes unavailable, uses three timers within VRRP: the advertisement interval, the master down interval, and the skew time:

- The advertisement interval is the time interval between advertisements (in seconds). The default interval is 1 second.
- The master down interval is the time interval for the backup to declare the master down (in seconds). The default is 3 x advertisement interval + skew time.
- The skew time,  $(256 - \text{priority} / 256)$  ms, ensures that the backup router with the highest priority becomes the new master.

The table lists the steps that are involved in the VRRP transition.

### VRRP Transition Process

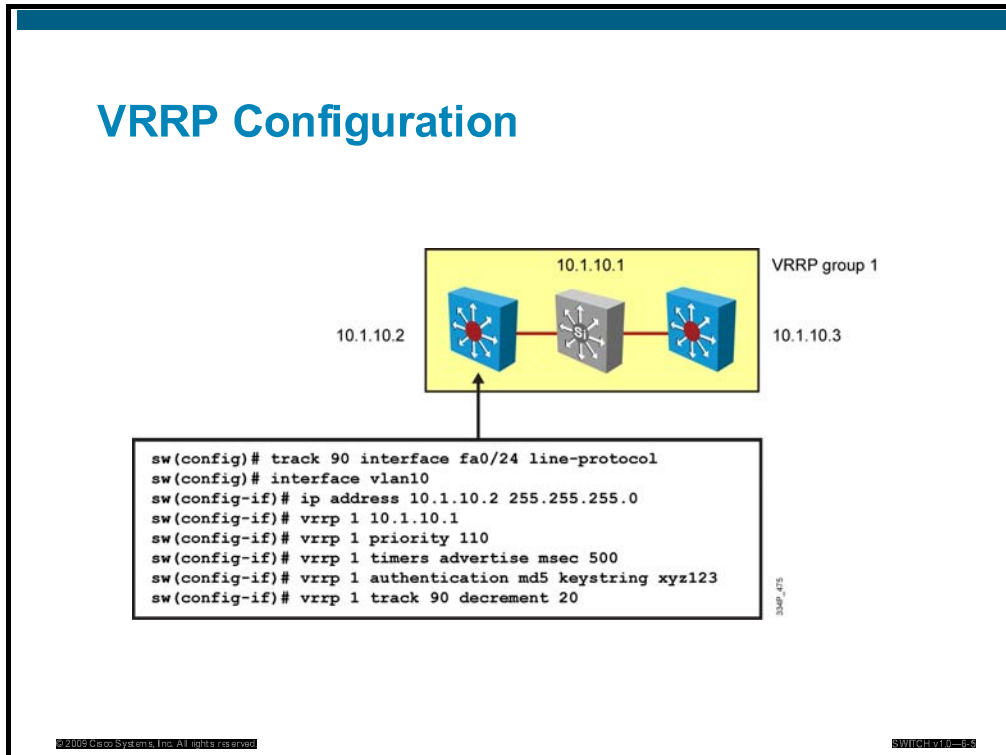
Step	Description	Notes
1.	Router A is currently the master, so it is sending advertisements by default every 1 second.	Router A is the only device sending advertisements.
2.	Router A fails.	Advertisements stop.
3.	Router B and router C stop receiving advertisements and wait for their respective master down interval to expire before transitioning to the master state.	By default, the master down interval is 3 seconds plus the skew time.
4.	Because the skew time is inversely proportional to priority, the master down interval of router B is less than that of router C.  Router B has a master down interval of approximately 3.2 seconds.  Router C has a master down interval of approximately 3.6 seconds.	The skew time for router B equals $(256 - 200) / 256$ , which is approximately equal to 0.2 seconds.  The skew time for router C equals $(256 - 100) / 256$ , which is approximately equal to 0.6 seconds.
5.	Router B transitions to the master state after 3.2 seconds and starts sending advertisements.	
6.	Router C receives the advertisement from the new master, so it resets its master down interval and remains in the backup state.	

<b>Note</b>	In the case of an orderly shutdown of the VRRP master, it sends an advertisement with a priority of 0. This priority setting then triggers the backup router to take over quicker by waiting only the skew time instead of the master down interval. Therefore, in the previous example, router B would have waited only 0.2 seconds to transition to the master state.
-------------	---



# Configuring VRRP

This topic describes the commands that are used to configure the VRRP and GLBP operations.



VRRP and GLBP are supported on select Cisco Catalyst platforms and, when supported, can be configured using these commands.

The following table describes VRRP command parameters.

## VRRP Commands

Command	Description
Switch(config-if)# <b>vrrp</b> group-number <b>ip</b> virtual-gateway-addr	Makes the interface a member of the virtual group identified with the IP virtual address.
Switch(config-if)# <b>vrrp</b> group-number <b>priority</b> priority_value	Sets the priority of this router. Highest value will win election as active router. Default is 100. If routers have the same VRRP priority, the gateway with the highest real IP address is elected to become the master virtual router.
Switch(config-if)# <b>vrrp</b> group-number <b>timers</b> advertise timer-value	Master router configures this parameter to advertise value to the other group members. Other group members configure timers that they have learned to accept.
Switch(config-if)# <b>vrrp</b> group-number <b>timers</b> learn	Configures nonmaster members to learn timer values from master.

The following table describes how to configure VRRP.

### VRRP Implementation

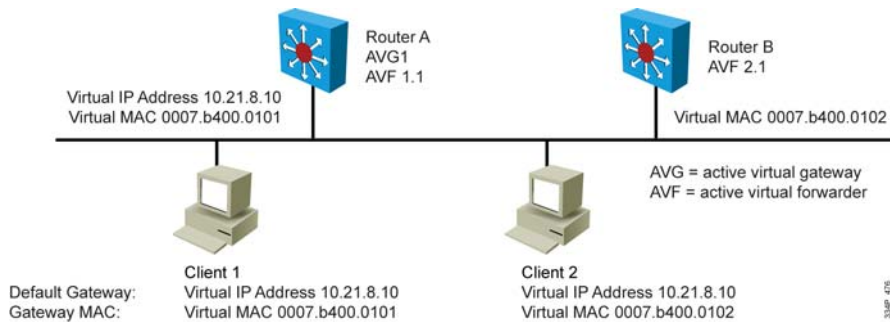
Step	Description
1.	To enable VRRP on an interface: <code>Switch(config-if)#vrrp group-number ip virtual-gateway-address</code>
2.	To set a VRRP priority for this router for this VRRP group: <code>Switch(config-if)#vrrp group-number priority priority-value</code>
3.	To change timer and indicate whether it should advertise (master) or learn (backup): <code>Switch(config-if)#vrrp group-number timers advertise timer-value</code> <code>Switch(config-if)#vrrp group-number timers learn</code>

# Describing GLBP

This topic describes GLBP.

## About GLBP

- Allows full use of resources on all devices without the administrative burden of creating multiple groups
- Provides a single virtual IP address and multiple virtual MAC addresses
- Routes traffic to single gateway distributed across routers
- Provides automatic rerouting in the event of any failure



Although HSRP and VRRP provide gateway resiliency, for the standby members of the redundancy group, the upstream bandwidth is not used while the device is in standby mode.

Only the active router for HSRP and VRRP groups forwards traffic for the virtual MAC. Resources that are associated with the standby router are not fully utilized. You can accomplish some load balancing with these protocols by creating multiple groups and assigning multiple default gateways, but this configuration creates an administrative burden.

GLBP is a Cisco proprietary solution that was created in 2005 to allow the automatic selection and simultaneous use of multiple available gateways in addition to automatic failover between those gateways. Multiple routers share the load of frames that, from a client perspective, are sent to a single default gateway address.

With GLBP, you can fully utilize resources without the administrative burden of configuring multiple groups and managing multiple default gateway configurations, as is required with HSRP and VRRP.

## GLBP vs. HSRP

HSRP	GLBP
Cisco proprietary, 1994.	Cisco proprietary, 2005.
16 groups max.	1024 groups max.
1 active, 1 standby, several candidates.	1 AVG, several AVFs; AVG load-balances traffic among AVFs and AVG.
Virtual IP is different from active and standby real IP addresses.	Virtual IP is different from AVG and AVF real IP addresses.
1 virtual MAC address for each group.	1 virtual MAC address per AVF or AVG in each group.
Uses 224.0.0.2.	Uses 224.0.0.102.
Can track interfaces or objects.	Can track only objects.
Default timers: hello, 3 sec; hold time, 10 sec.	Default timers: hello, 3 sec; hold time, 10 sec.
Authentication supported.	Authentication supported.

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-1000

## GLBP Functions

These are GLBP functions:

- **GLBP active virtual gateway (AVG):** Members of a GLBP group elect one gateway to be the AVG for that group. Other group members provide backup for the AVG in the event that the AVG becomes unavailable. The AVG assigns a virtual MAC address to each member of the GLBP group.
- **GLBP active virtual forwarder (AVF):** Each gateway assumes responsibility for forwarding packets that are sent to the virtual MAC address that is assigned to that gateway by the AVG. These gateways are known as AVFs for their virtual MAC address.
- **GLBP communication:** GLBP members communicate between each other through hello messages sent every 3 seconds to the multicast address 224.0.0.102, UDP port 3222.

## GLBP Features

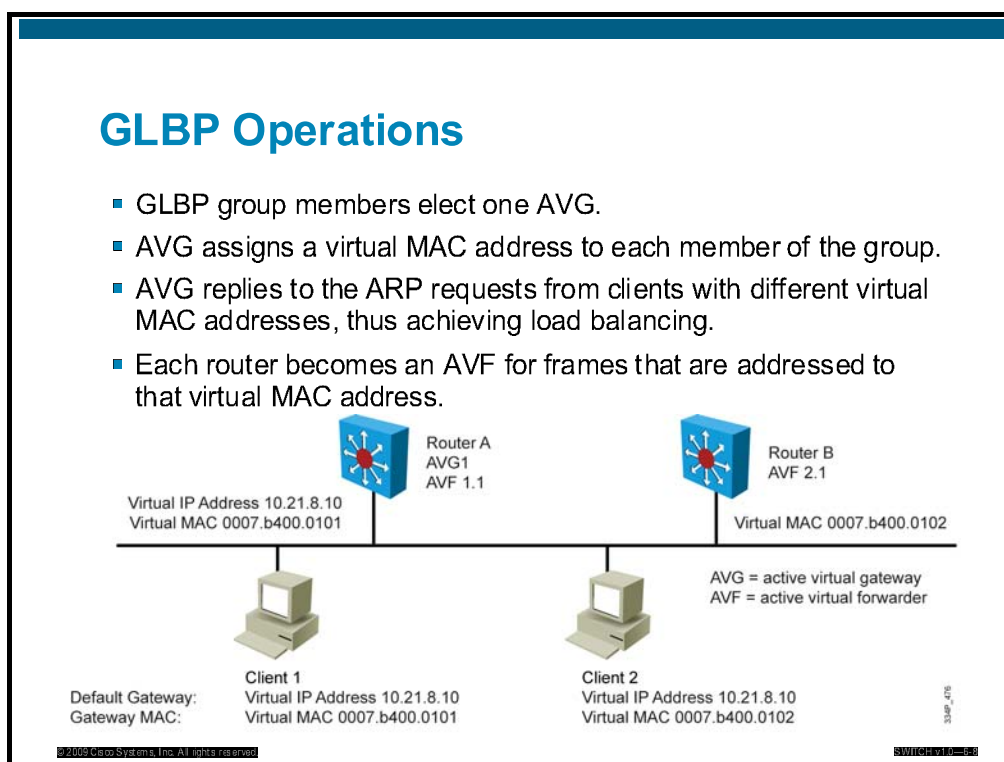
- GLBP features are as follows:
- **Load sharing:** You can configure GLBP in such a way that multiple routers can share traffic from LAN clients, thereby sharing the traffic load more equitably among available routers.
- **Multiple virtual routers:** GLBP supports up to 1024 virtual routers (GLBP groups) on each physical interface of a router and up to four virtual forwarders per group.
- **Pre-emption:** The redundancy scheme of GLBP enables you to pre-empt an AVG with a higher-priority backup virtual gateway that has become available. Forwarder pre-emption works in a similar way, except that forwarder pre-emption uses weighting instead of priority and is enabled by default.
- **Efficient resource utilization:** GLBP makes it possible for any router in a group to serve as a backup, which eliminates the need for a dedicated backup router, because all available routers can support network traffic.

GLBP provides upstream load sharing by utilizing the redundant uplinks simultaneously. It uses link capacity efficiently, thus providing peak-load traffic coverage. By making use of multiple available paths upstream from the routers or Layer 3 switches that are running GLBP, you can also reduce output queues.

Only a single path is used with HSRP or VRRP, while other paths are idle, unless multiple groups and gateways are configured. The single path may encounter higher output queue rates during peak times, which leads to lower performance from higher jitter rates. The impact of jitter is lessened and overperformance is increased because more upstream bandwidth is available and additional upstream paths are used.

# Identifying the GLBP Operations Process

This topic describes how GLBP provides balanced traffic on a per-host basis, using a round-robin scheme.



GLBP allows automatic selection and simultaneous use of all available gateways in the group. The members of a GLBP group elect one gateway to be the AVG for that group. Other members of the group provide backup for the AVG if it becomes unavailable. The AVG assigns a virtual MAC address to each member of the GLBP group. All routers become AVFs for frames addressed to that virtual MAC address. As clients send Address Resolution Protocol (ARP) requests for the address of the default gateway, the AVG sends these virtual MAC addresses in the ARP replies. A GLBP group can have up to four group members.

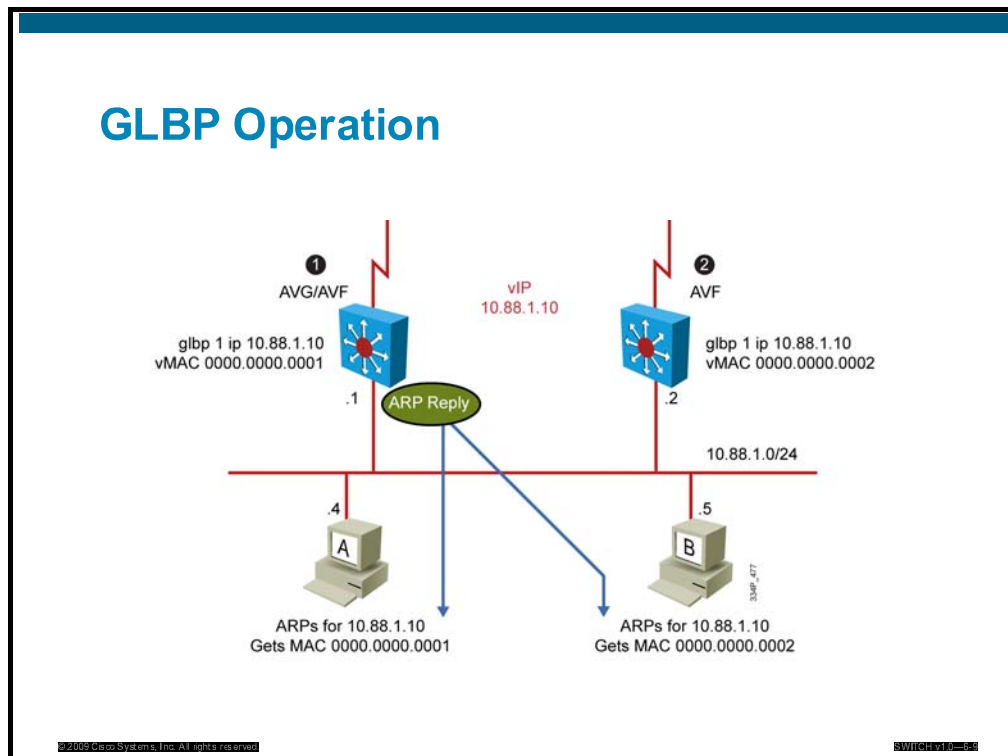
GLBP supports these operational modes for load balancing traffic across multiple default routers that are servicing the same default gateway IP address:

- **Weighted load-balancing algorithm:** The amount of load that is directed to a router is dependent upon the weighting value that is advertised by that router.
- **Host-dependent load-balancing algorithm:** A host is guaranteed the use of the same virtual MAC address as long as that virtual MAC address is participating in the GLBP group.
- **Round-robin load-balancing algorithm:** As clients send ARP requests to resolve the MAC address of the default gateway, the reply to each client contains the MAC address of the next possible router in round-robin fashion. The MAC addresses of all routers take turns being included in address resolution replies for the default gateway IP address.

GLBP automatically manages the virtual MAC address assignment, determines who handles the forwarding, and ensures that each station has a forwarding path in the event of failures to gateways or tracked interfaces. If failures occur, the load-balancing ratio is adjusted among the remaining AVFs so that resources are used in the most efficient way.

# GLBP Operation

This subtopic describes the GLBP operation.



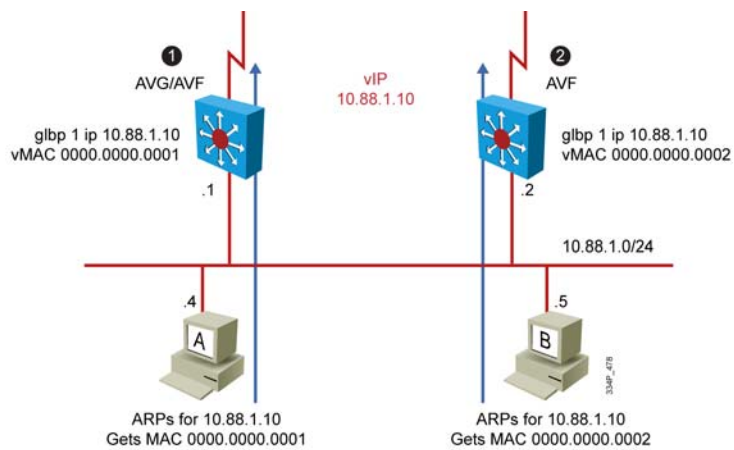
As shown in the figure, by default, GLBP will attempt to balance traffic on a per-host basis, using the round-robin algorithm.

The table describes how GLBP balances traffic using the round-robin algorithm.

## GLBP Per-Host Traffic Balancing

Step	Description
1.	When a client sends an ARP message for the gateway IP address, the AVG returns the virtual MAC address of one of the AVFs.
2.	When a second client sends an ARP message, the AVG returns the next virtual MAC address from the list.

## GLBP Operation (Cont.)



© 2009 Cisco Systems, Inc. All rights reserved.

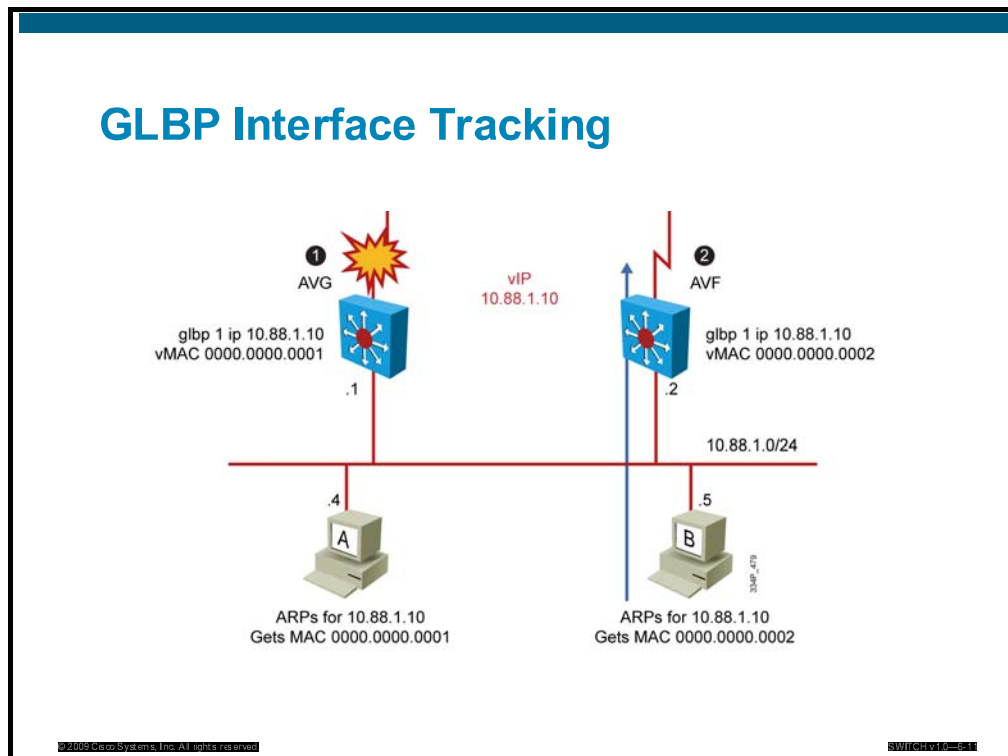
SWITCH v1.0

Having each resolved a different MAC address for the default gateway, clients A and B will send their routed traffic to separate routers, although they both have the same default gateway address configured. Each GLBP router is an AVF for the virtual MAC address to which it has been assigned.



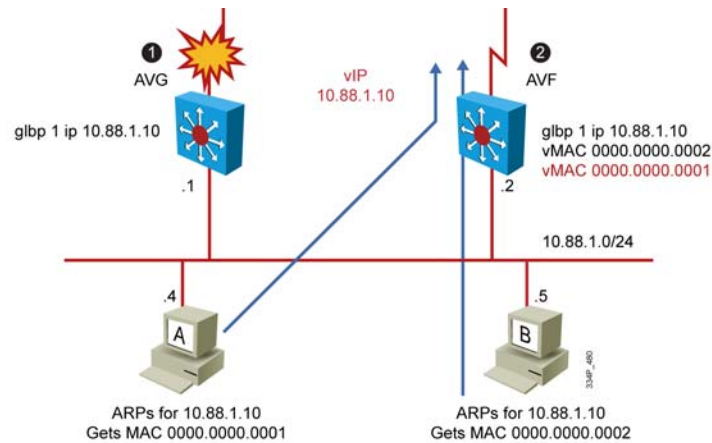
# GLBP Interface Tracking

This subtopic describes GLBP interface tracking.




Like HSRP, GLBP can be configured to track interfaces. In the figure, the WAN link from router R1 is lost. GLBP detects the failure. Just like HSRP, GLBP decrements the gateway priority when a tracked interface fails. The second gateway then becomes primary. This transition is transparent for the LAN client.

## GLBP Interface Tracking (Cont.)



Because interface tracking was configured on router R1, the job of forwarding packets for virtual MAC address 0000.0000.0001 will be taken over by the secondary virtual forwarder for the MAC, router R2. Therefore, the client sees no disruption of service nor does the client need to resolve a new MAC address for the default gateway.

## GLBP Weights and Decrements



GLBP weighting is used to determine whether a router can act as a virtual forwarder. You can set initial weighting values and specify optional thresholds. You can also track interface states and set a decrement value to reduce the weighting value if the interface goes down. When the GLBP router weighting drops below a specified value, the router will no longer be an active virtual forwarder. When the weighting rises above a specified value, the router can resume its role as an active virtual forwarder.

The weighting mechanism for GLBP is different from that of HSRP or VRRP. With HSRP and VRRP, one single threshold is defined. If the router priority (or weight) falls below the threshold, the router loses its active state. As soon as the router weight (or priority) exceeds the threshold, the router regains its active state. With GLBP, two thresholds are defined: one *lower threshold* that applies when the router loses weight, and one *upper threshold* that applies when the router regains weight. This double threshold mechanism allows for more flexibility than the single threshold system.

In the figure, SW4 is forwarding. Its initial weight (or priority) is 110. SW4 tracks both fa0/23 and fa0/24 interfaces. Fa0/23 is the active interface. Losing fa0/23 decrements SW4 by 20 points, thus bringing the SW4 weight down (from 110) to 90. Fa0/24 is a backup interface. Losing fa0/24 decrements SW4 by 10 points, thus bringing the SW4 weight down (from 110) to 100, which is the default weight of the other routers. Losing both fa0/23 and fa0/24 brings the SW4 weight down (from 110) to 80.

In that scenario, losing fa0/24 brings the SW4 weight to the same level as the weight of the other router. If an election occurs, SW4 may or may not be the AVG, but would still forward. Losing fa0/23 brings the SW4 weight below the weight of the other routers. If an election occurs, SW4 would not be elected as the AVG, but would still be an AVF. The lower threshold is set to 85, which can be reached only by losing either fa0/23 or fa0/24. In other words, losing fa0/23 or fa0/24 decreases the SW4 weight and may change its status from AVG to AVF, but will not prevent SW4 from being a forwarder. It is only when both interfaces are lost that SW4 stops forwarding.


Losing one interface is an issue but does not prevent SW4 from forwarding. Losing both interfaces is the sign of a major connectivity problem in the network. For that reason, the network administrator decided that, if SW4 lost both interfaces, it would not resume forwarding until both interfaces are back up. To implement this mechanism, the second threshold, the *upper threshold*, is set to 105. As long as the SW4 weight does not go below the lower threshold (85), the upper threshold is not called. As soon as SW4 goes below the lower threshold, SW4 stops forwarding and the upper threshold is called. It is then only when the SW4 weight will become higher than the upper threshold that SW4 will resume forwarding packets.

In the preceding example, when both interfaces are lost, SW4 goes below the lower threshold (85) to reach 80. Recovering fa0/23 or fa0/24 would add 20 or 10 points to the weight, but each interface weight is not enough to have SW4 exceed the upper threshold, 105. It is only when both interfaces get re-enabled that the weight exceeds the upper threshold and that SW4 resumes forwarding packets.

# GLBP Implementation

This topic describes the process that is used to implement GLBP.

## GLBP Configuration



```
SW4(config)# track 90 interface fa0/24 line-protocol
SW4(config)# track 91 interface fa0/23 line-protocol
SW4(config)# interface vlan10
SW4(config-if)# ip address 10.1.10.2 255.255.255.0
SW4(config-if)# glbp 1 10.1.10.1
SW4(config-if)# glbp 1 weighting 110 lower 85 upper 105
SW4(config-if)# glbp 1 timers msec 200 msec 700
SW4(config-if)# glbp 1 preempt delay minimum 300
SW4(config-if)# glbp 1 authentication md5 keystring xyz123
SW4(config-if)# glbp 1 weighting track 90 decrement 10
SW4(config-if)# glbp 1 weighting track 91 decrement 20
```

The table describes GLBP command parameters.

## GLBP Commands

Command	Description
Switch(config-if)# <b>glbp</b> group-number ip virtual-gateway-addr	Makes the interface a member of the virtual group that is identified with the IP virtual address.
Switch(config-if)# <b>glbp</b> group-number priority priority_value	Sets the priority of this router. Highest value will win election as active router. Default is 100. If routers have the same GLBP priority, the gateway with the highest real IP address will become the AVG.
Switch(config-if)# <b>glbp</b> group-number timers hello-value holdtime-value	Adjusts the hello timer and hold timer in seconds. Place the argument msec before the values to enter subsecond values.

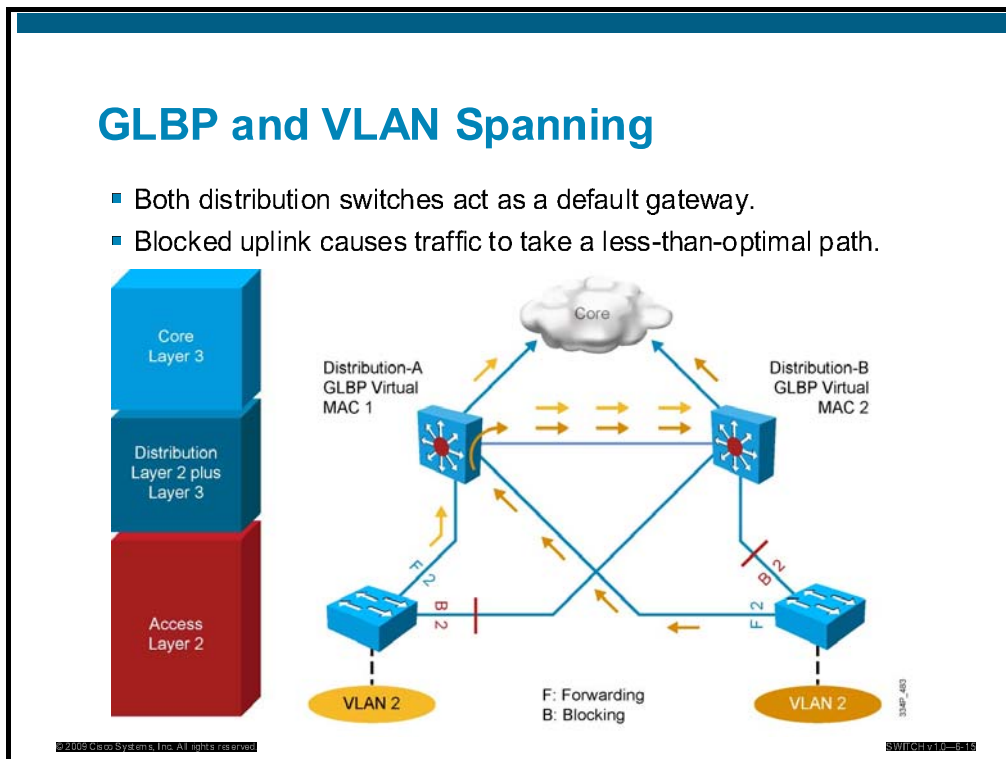
The following table describes the steps for configuring GLBP.

### GLBP Configuration

Step	Description
1.	Enable GLBP on an interface. <code>Switch(config-if)#<b>glbp</b> group-number <b>ip</b> virtual-gateway-address</code>
2.	Set a GLBP priority for this router for this GLBP group. <code>Switch(config-if)#<b>glbp</b> group-number <b>priority</b> priority-value</code>
3.	Change timer values for hello interval and hold time. <code>Switch(config-if)#<b>glbp</b> group-number <b>timers</b> hello holdtime</code>

# GLBP and VLAN Spanning

This subtopic describes GLBP with VLAN spanning access switches.



Topologies where STP has blocked one of the access uplinks may result in a two-hop path at Layer 2 for upstream traffic. In the figure, the Distribution A switch has a direct connection to the core and a redundant link to Distribution B. Because of spanning-tree operations, the interface that is directly linking to the core is in the blocking state. Although it is invisible and transparent to VLAN 2 clients, this state causes the frames that are coming from VLAN 2 to transit through Distribution A, and then actually go through Distribution B before being sent to the core.

In environments where VLANs span switches, HSRP is the recommended first-hop redundancy protocol implementation. In all cases, the active gateway should be configured to also be the root bridge for the VLAN in which first-hop redundancy is configured.

# Summary

This topic summarizes the key points that were discussed in this lesson.

## Summary

- VRRP provides router redundancy in a manner similar to that of HSRP.
- VRRP supports a master router and one or more backup routers.
- VRRP is configured per interface.
- GLBP provides router redundancy and load balancing.
- GLBP balances traffic by allocating a virtual MAC address to each AVF.
- The GLBP configuration steps are very similar to those of HSRP and VRRP.



## Lesson 3

---

# Lab 6-1 Debrief

---

## Overview

In this lab, you have configured your pod switches for HSRP. HSRP could be configured on any router or Layer 3 switch in your lab, if two HSRP-enabled Layer 3 devices face the same Layer 2 subnet. In this example, DSW1 and DSW2 were configured for HSRP for both VLAN 3 and VLAN 4.

During the lab debrief, the instructor will lead a group discussion in which you can present your solution. You will get an opportunity to verify your solution against a number of checkpoints that are provided by the instructor, and to compare your solution to those of other students. The instructor will discuss alternative solutions and their benefits and drawbacks.

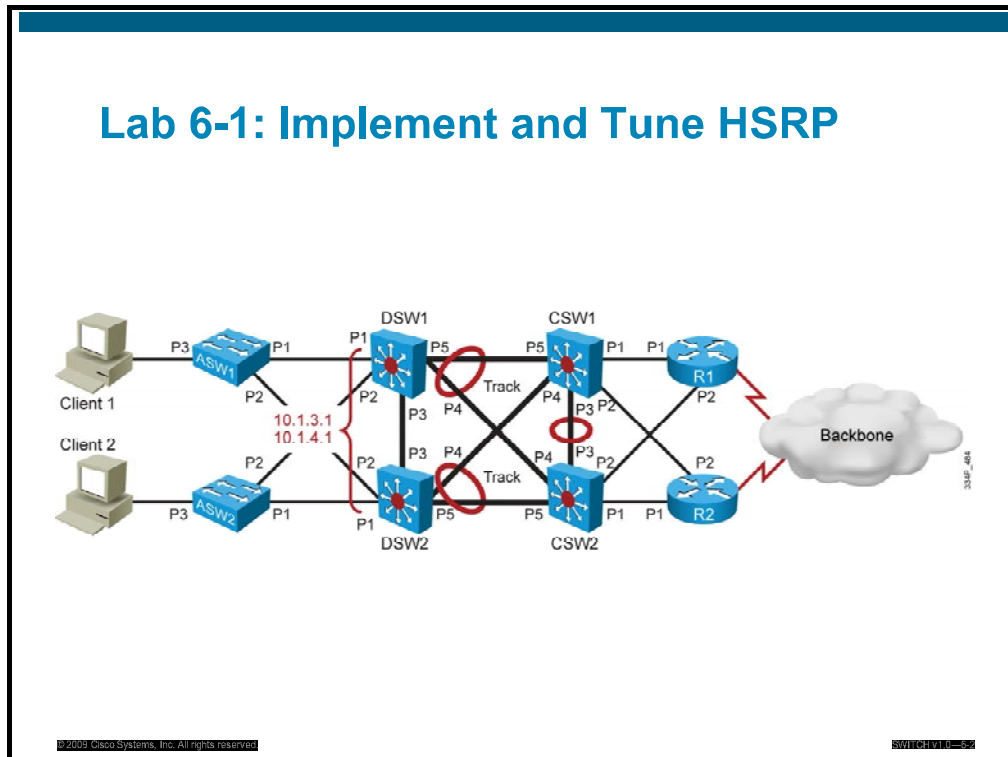
## Objectives

Upon completing this lesson, you will be able to configure, tune, and troubleshoot HSRP. This ability includes being able to meet these objectives:

- Review and verify your solution, as well as your findings and action log, against a set of checkpoints that are provided by the instructor
- Consolidate the lessons that you learned during the review discussions into a set of best-practice methods and commands to aid you in future deployment procedures

# Review and Verification

This topic describes the client requirements that were listed in Lab 6-1 and asks how you can verify that you have identified the solution matching the client needs.



Your lab consists of six switches and two routers. Your task is to configure two of these devices for HSRP. In this lab, you have one client in VLAN 3 (CLT1), and one client in VLAN 4 (CLT 2). You need to configure HSRP so that, for each client, two Layer 3 devices can be used as redundant gateways.

The ideal candidates are DSW1 and DSW2, which are the Layer 3 devices that are the closest to the clients. You could technically configure CSW1 and CSW2, or R1 and R2. You could even configure any pair of these Layer 3 devices, as long as the final client has Layer 2 reachability to both of these devices. For example, you could use DSW1 and R2, as long as CLT1 and CLT2 have Layer 2 reachability to both DSW1 and R2. It is still more logical to configure DSW1 and DSW2 from a design perspective, because they are the first Layer 3 possible hop.

## Design and Implementation Plan

Which items should be configured, and in which order?

- Which switches should be root, and for which VLANs?
- How did you configure HSRP?
- How did you implement priority?
- How did you implement tracking?
- Which tests did you conduct to verify HSRP operations?

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCHING-55

A successful implementation plan allows you to configure the devices without duplication and to reduce the risk of mistakes. Thus, you will need to make only a minimum set of configuration changes and will reduce the time for troubleshooting after the implementation. In other words, an implementation plan is efficient when you do not need to alter your previous configuration to implement new items. You should proceed in a logical order.



## Lesson 4

---

# Lab 6-2 Debrief

---

## Overview

In this lab, you have configured your pod routers for VRRP. HSRP could be configured on any router or Layer 3 switch in your lab, if two VRRP-enabled Layer 3 devices face the same Layer 2 subnet. Just as HSRP was configured on DSW1 and DSW2 in the previous lab, VRRP is configured here on R1 and R2 with CSW1 and CSW2 acting as clients.

During the lab debrief, the instructor will lead a group discussion in which you can present your solution. You will get an opportunity to verify your solution against a number of checkpoints that are provided by the instructor, and to compare your solution to those of other students. The instructor will discuss alternative solutions and their benefits and drawbacks.

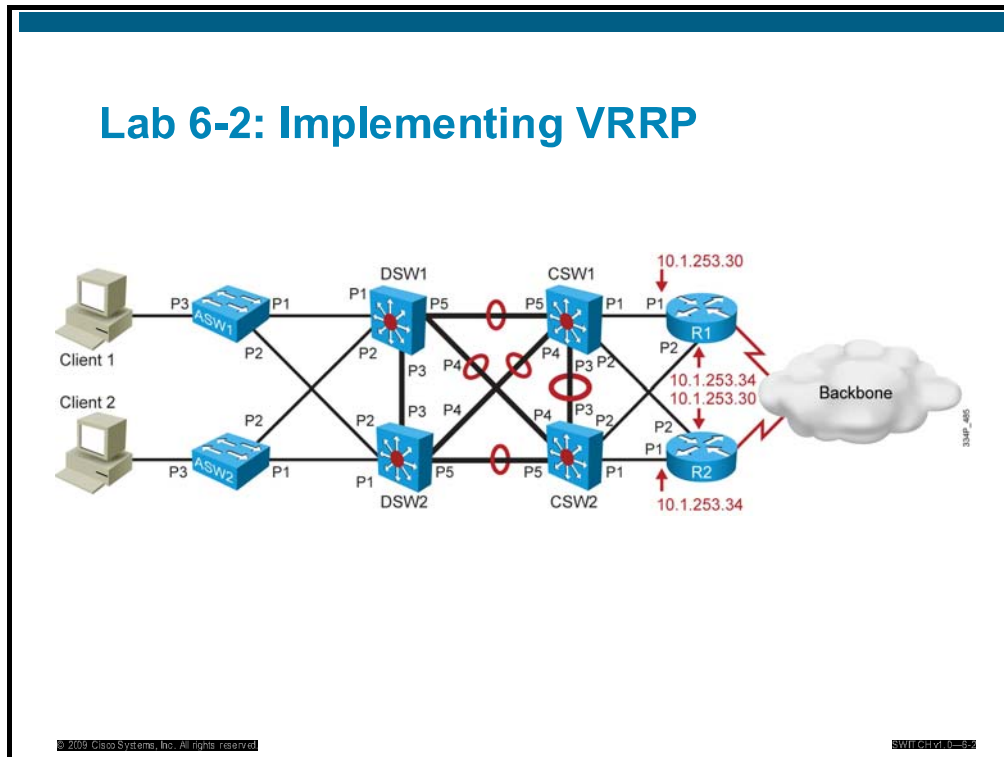
## Objectives

Upon completing this lesson, you will be able to configure, tune, and troubleshoot VRRP. This ability includes being able to meet these objectives:

- Review and verify your solution, as well as your findings and action log, against a set of checkpoints that are provided by the instructor.
- Consolidate the lessons that you learned during the review discussions into a set of best-practice methods and commands to aid you in future deployment procedures.

# Review and Verification

This topic describes the client requirements that were listed in Lab 6-2 and asks how you can verify that you have identified the solution matching the client needs.



Your lab consists of six switches and two routers. Your task is to configure two of these devices for VRRP. Because CLT1 and CLT2 are already using DSW1 and DSW2, respectively, as gateways, the lab instructs you to change the CSW1 and CSW2 configurations and create two SVIs on each switch. CSW1 then acts as a client for R1 and R2, which act as the first-hop redundant gateway in one VLAN, while CSW2 acts as another client, in another VLAN, with R1 and R2 still acting as a first-hop redundant gateway for this second VLAN.

## Design and Implementation Plan

Which items should be configured, and in which order?

- How should switches CSW1 and CSW2 be configured for this lab?
- How did you configure VRRP?
- How did you implement priority?
- Can you implement tracking?
- Which tests did you conduct to verify VRRP operations?

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCHING-55

A successful implementation plan allows you to configure the devices without duplication and to reduce the risk of mistakes. Thus, you will need to make only a minimum set of configuration changes and will reduce the time for troubleshooting after the implementation. In other words, an implementation plan is efficient when you do not need to alter your previous configuration to implement new items. You should proceed in a logical order.





# Module Summary

This topic summarizes the key points that were discussed in this module.

## Module Summary

- HSRP is enabled so that redundant routers can provide default gateway functionality.
- VRRP or GLBP can provide Layer 3 router failover in addition to load balancing at the distribution layer.



# Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) During which three HSRP states do routers send hello messages? (Choose three.)  
(Source: Configuring Layer 3 Redundancy with HSRP)
- A) initial
  - B) listen
  - C) speak
  - D) active
  - E) standby
- Q2) Which command enables HSRP? (Source: Configuring Layer 3 Redundancy with HSRP)
- A) **standby virtual ip 10.1.1.1**
  - B) **standby ip 10.1.1.1 group 1**
  - C) **hsrp 1 ip 10.1.1.1**
  - D) **standby 1 ip 10.1.1.1**
- Q3) During the election process, if all routers on VLAN10 have the same HSRP priority, which router will become the active router? (Source: Configuring Layer 3 Redundancy with HSRP)
- A) the router with the highest configured IP address on any interface
  - B) the router with the lowest configured IP address on interface VLAN10
  - C) the router with the highest configured IP address on interface VLAN10
  - D) the router with the lowest configured IP address on any interface
- Q4) Which command enables pre-emption with a delay of 2 minutes? (Source: Configuring Layer 3 Redundancy with HSRP)
- A) **standby 1 preempt delay minutes 2**
  - B) **hsrp 1 delay 120**
  - C) **standby 1 delay preempt 120**
  - D) **standby 1 preempt delay minimum 120**
- Q5) Which is the primary purpose of IP SLA when used in combination with HSRP?  
(Source: Configuring Layer 3 Redundancy with HSRP)
- A) to implement two active IP addresses per virtual interface
  - B) to change priority level based on reachability information
  - C) to disable the pre-empt function if a given IP address becomes unreachable
  - D) to report to a syslog server HSRP state changes
- Q6) With VRRP, which command should you use so that advertisements are generated every 3 seconds? (Source: Configuring Layer 3 Redundancy with VRRP and GLBP)
- A) **vrrp 1 timers 3 10**
  - B) **vrrp 1 timers advertise 3**
  - C) **vrrp 1 timers advertise msec 3**
  - D) **vrrp 1 timers 3**

- Q7) What is the difference between VRRP and HSRP? (Source: Configuring Layer 3 Redundancy with VRRP and GLBP)
- A) HSRP is a standard protocol, while VRRP is Cisco proprietary.
  - B) HSRP is configured at the global level, while VRRP is configured at the interface level.
  - C) VRRP offers Layer 2 first-hop redundancy, while HSRP offers Layer 3 first-hop redundancy.
  - D) VRRP is a standard protocol, while HSRP is Cisco proprietary.
- Q8) What is GLBP AVG? (Source: Configuring Layer 3 Redundancy with VRRP and GLBP)
- A) the gateway that is elected by the GLBP peers
  - B) the filtering feature that is embedded in the GLBP protocol
  - C) the name for the nonactive gateways
  - D) the name for the gateway protocol between GLBP and VRRP
- Q9) Which statement is a specificity of GLBP? (Source: Configuring Layer 3 Redundancy with VRRP and GLBP)
- A) GLBP offers one virtual IP address and one virtual MAC address.
  - B) GLBP works at Layer 2 only, while VRRP works at Layer 3 only.
  - C) GLBP cannot track interfaces, while VRRP can track interfaces.
  - D) GLBP allows for load sharing between peers.
- Q10) Which command enables GLBP on an interface? (Source: Configuring Layer 3 Redundancy with VRRP and GLBP)
- A) **glbp 1 10.1.1.1**
  - B) **glbp 1 ip 10.1.1.1 mac abcd.abcd.abcd**
  - C) **redundancy 1 10.1.1.1**
  - D) **standby 1 glbp 10.1.1.1**

## Module Self-Check Answer Key

- Q1) C, D, E
- Q2) D
- Q3) C
- Q4) D
- Q5) B
- Q6) B
- Q7) D
- Q8) A
- Q9) D
- Q10) A



# Minimizing Service Loss and Data Theft in a Campus Network

---

## Overview

This module defines the potential vulnerabilities relating to VLANs that can occur within a network. After the vulnerabilities are identified, solutions for each vulnerability are discussed, and configuration commands are defined.

The module also discusses port security for denial of MAC spoofing and MAC flooding, and the use of private VLANs (PVLANS) and VLAN access control lists (VACLs) to control VLAN traffic. VLAN hopping, DHCP spoofing, Address Resolution Protocol (ARP) spoofing, and Spanning Tree Protocol (STP) attacks are also explained. In addition, you will learn about potential problems and their solutions, and the method for securing the switch access with use of vty access control lists (ACLs), and implementing the Secure Shell (SSH) Protocol for secure Telnet access.

## Objectives

Upon completing this module, you will be able to implement security features in a switched network. This ability includes being able to meet these objectives:

- Explain the vulnerabilities of switches to network attacks
- Configure various features to prevent VLAN hopping and to address VLAN security issues
- Explain how to defend against spoof attacks with DAI, DHCP snooping, and IP Source Guard
- Secure network services by tuning Cisco Discovery Protocol, LLDP, Telnet, SSH, and web services





## Lesson 1

---

# Understanding Switch Security Issues

---

## Overview

You should take basic security measures to guard against a host of attacks that can be launched at a switch and its ports. You can take specific measures to guard against MAC flooding, which is a common Layer 2 malicious activity.

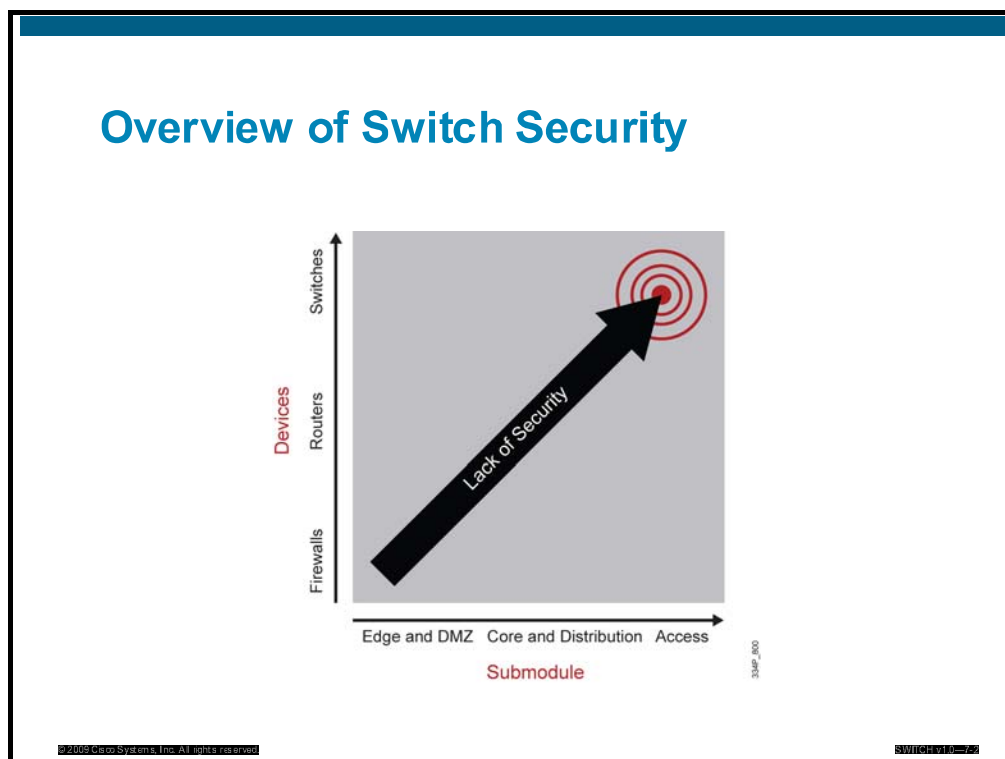
## Objectives

Upon completing this lesson, you will be able to describe and implement security features in a switched network. This ability includes being able to meet these objectives:

- Describe switch and Layer 2 security as a subset of an overall network security plan
- Describe how a rogue device gains unauthorized access to a network
- Categorize switch attack types and list mitigation options
- Describe how a MAC flooding attack works to overflow a CAM Campus Backbone Layer table
- Describe how port security is used to block input from devices based on Layer 2 restrictions
- Describe the procedure for configuring port security on a switch
- Describe the methods that can be used for authentication using AAA
- Describe port-based authentication using 802.1X

# Overview of Switch Security Issues

This topic describes switch and Layer 2 security as a subset of an overall network security plan.



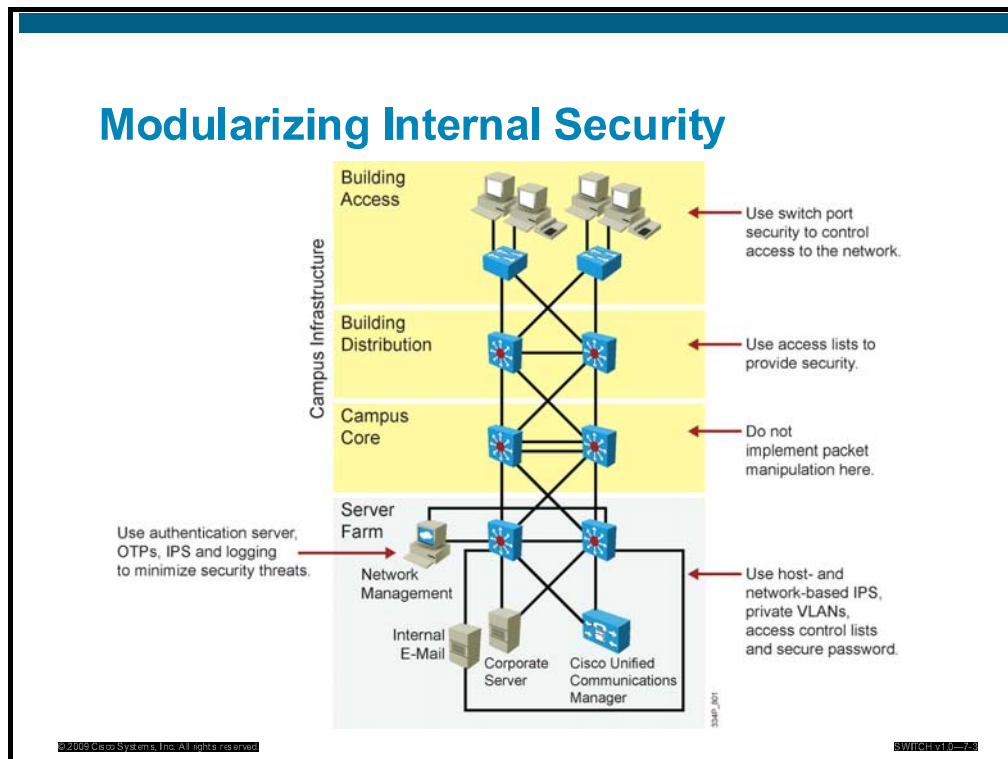
Much industry attention focuses on security attacks from outside the walls of an organization and at the upper Open Systems Interconnection (OSI) layers. Network security often focuses on edge routing devices and the filtering of packets that are based on Layer 3 and Layer 4 headers, ports, stateful packet inspection, and so forth. This includes all issues related to Layer 3 and above, as traffic makes its way into the campus network from the Internet. Campus access devices and Layer 2 communication are largely unconsidered in most security discussions.

The default state of networking equipment highlights this focus on external protection and internal open communication. Firewalls, placed at the organizational borders, arrive in a secure operational mode and allow no communication until they are configured to do so. Routers and switches that are internal to an organization and that are designed to accommodate communication, delivering needful campus traffic, have a default operational mode that forwards all traffic unless they are configured otherwise. Their function as devices that facilitate communication often results in minimal security configuration, and they become targets for malicious attacks. If an attack is launched at Layer 2 on an internal campus device, the rest of the network can be quickly compromised, often without detection.

Many security features are available for switches and routers, but they must be enabled to be effective. As with Layer 3, where security had to be tightened on devices within the campus as malicious activity that compromised this layer increased, now security measures must be taken to guard against malicious activity at Layer 2. A new security focus centers on attacks that are launched by maliciously using normal Layer 2 switch operations. Security features exist to protect switches and Layer 2 operations. However, as with access control lists (ACLs) for upper-layer security, a policy must be established and appropriate features configured to protect against potential malicious acts while maintaining daily network operations.

# Security Infrastructure Services

This topic describes the security design issues within an enterprise design network.



Security is an infrastructure service that increases the integrity of the network by protecting network resources and users from internal and external threats. Without a full understanding of the threats that are involved, network security deployments tend to be incorrectly configured, too focused on security devices, or lacking in the appropriate threat-response options.

You can evaluate and apply security on a module-by-module basis within the Cisco Enterprise Architecture. The following are some recommended-practice security considerations for each module:

- The campus core layer in the campus infrastructure module switches packets as quickly as possible. It should not perform any security functions, because these would slow down packet switching.
- The building distribution layer performs packet filtering to keep unnecessary traffic from the campus core layer. Packet filtering at the building distribution layer is a security function because it prevents some undesired access to other modules. Given that switches in this layer are usually Layer 3-aware multilayer switches, the building distribution layer is often the first location that can filter based on network layer information.
- At the building access layer, access can be controlled at the port level with respect to the data link layer information (for example, MAC addresses).
- The server farm module provides application services to end users and devices. Given the high degree of access that most employees have to these servers, they often become the primary target of internally originated attacks. Use host- and network-based intrusion prevention systems (IPSs), private VLANs, and access control to provide a much more comprehensive response to attacks. An onboard intrusion detection system (IDS) within multilayer switches can inspect traffic flows on the server farm module.

The server farm module typically includes a network management system that securely manages all devices and hosts within the enterprise architecture. Syslog provides important information regarding security violations and configuration changes by logging security-related events (authentication and so on). Other servers, including an authentication, authorization, and accounting (AAA) security server can work in combination with the one-time password (OTP) server to provide a very high level of security to all local and remote users. AAA and OTP authentication reduce the likelihood of a successful password attack.

# Reason for Internal Security

This subtopic describes reasons for internal security.

## Reasons for Internal Security

- The enterprise campus is protected by security functions in the enterprise edge:
  - If the enterprise edge security fails, the enterprise campus is vulnerable.
  - The potential attacker can gain physical access to the enterprise campus.
  - Some network solutions require indirect external access to the enterprise campus.
- All vital elements in the enterprise campus must be protected independently.

Several reasons exist for strong protection of the enterprise campus infrastructure, including security functions in each individual element of the enterprise campus:


- Relying on the security that has been established at the enterprise edge fails as soon as security there is compromised. Having several layers of security increases the protection of the enterprise campus, where usually the most strategic assets reside.
- If the enterprise allows visitors into its buildings, potentially an attacker can gain physical access to devices in the enterprise campus. Relying on physical security is not enough.
- Very often external access does not stop at the enterprise edge. Applications require at least an indirect access to the enterprise campus resources, which means that strong security is necessary.

# Unauthorized Access by Rogue Devices

This topic describes how a rogue device gains unauthorized access to a network.

## Rogue Devices

- Rogue network devices can be
  - Switches
  - Wireless access points
  - Hubs
- Connected to ports on access switches
- Connecting devices such as laptops or printers



Rogue access comes in several forms. For example, because unauthorized rogue access points are inexpensive and readily available, employees sometimes plug them into existing LANs and build ad hoc wireless networks without IT department knowledge or consent. These rogue access points can be a serious breach of network security because they can be plugged into a network port behind the corporate firewall. Because employees generally do not enable any security settings on the rogue access point, it is easy for unauthorized users to use the access point to intercept network traffic and hijack client sessions.

Malicious rogue access points, although much less common than employee-installed rogue access points, are also a security concern. These rogue access points create an unsecured wireless LAN connection that puts the entire wired network at risk. Malicious rogues present an even greater risk and challenge because they are intentionally hidden from physical and network view.

To mitigate Spanning Tree Protocol (STP) manipulation, use the **root guard** and the **BPDUGuard** enhancement commands to enforce the placement of the root bridge in the network and to enforce the STP domain borders. The RootGuard feature is designed to provide a way to enforce the root bridge placement in the network. The STP BPDUGuard is designed to allow network designers to keep the active network topology predictable. Although BPDUGuard may seem unnecessary, given that the administrator can set the bridge priority to zero, there is still no guarantee that it will be elected as the root bridge, because there might be a bridge with priority zero and a lower bridge ID. BPDUGuard is best deployed toward user-facing ports to prevent rogue switch-network extensions by an attacker.

# Switch Attack Categories

This topic categorizes switch attack types and lists mitigation options.

## Switch Attack Categories

- MAC address–based attacks
  - MAC address flooding
- VLAN attacks
  - VLAN hopping
- Spoofing attacks
  - Spoofing of DHCP, ARP, and MAC addressing
- Attacks on switch devices
  - Cisco Discovery Protocol
  - Management protocols

A device that is connected to the campus network typically launches Layer 2 malicious attacks. The attacks may originate from a physical rogue device that has been placed on the network for malicious purposes. The attack may also come from an external intrusion that takes control of, and launches attacks from, a trusted device. In either case, the network sees all traffic as originating from a legitimate connected device.

Attacks that are launched against switches and at Layer 2 can be grouped as follows:

- MAC layer attacks
- VLAN attacks
- Spoof attacks
- Attacks on switch devices

Significant attacks in these categories are discussed in more detail in subsequent sections of the course. Each attack method is accompanied by a standard measure for mitigating the security compromise.

The table describes attack methods and the steps to mitigation.

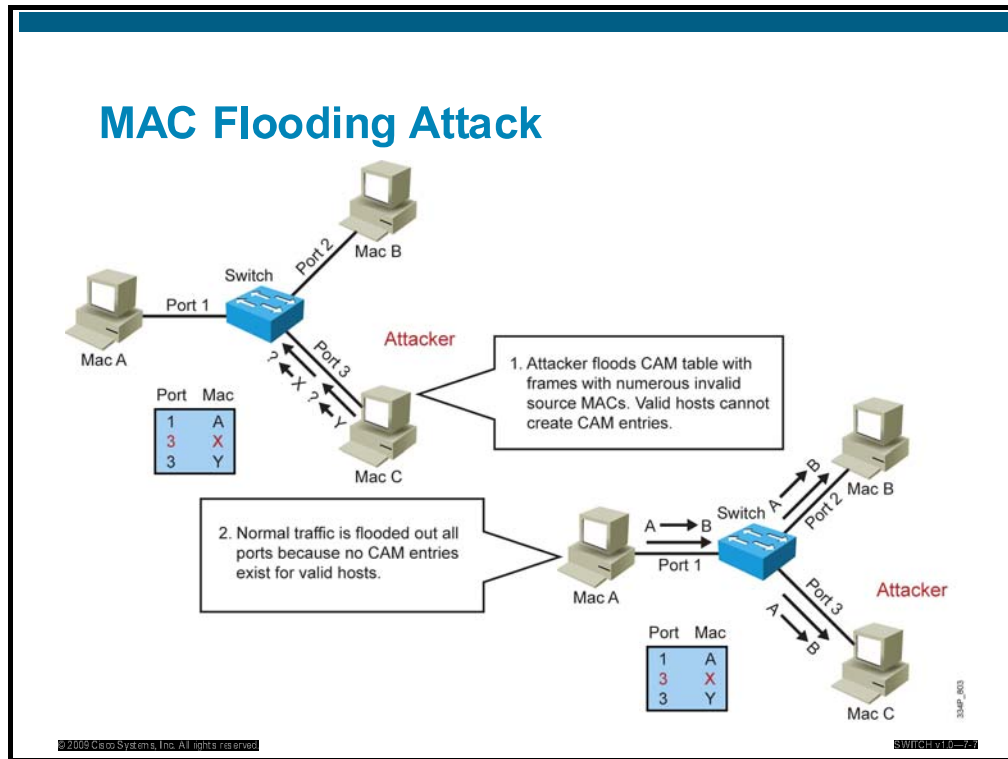
## Switch Security Concerns and Mitigation Steps

Attack Method	Description	Steps to Mitigation
<b>MAC Layer Attacks</b>		
MAC address flooding	Frames with unique, invalid source MAC addresses flood the switch, exhausting content-addressable memory (CAM) table space, disallowing new entries from valid hosts. Traffic to valid hosts is subsequently flooded out all ports.	Port security. MAC address VLAN access maps.
<b>VLAN Attacks</b>		
VLAN hopping	By altering the VLAN ID on packets that are encapsulated for trunking, an attacking device can send or receive packets on various VLANs, bypassing Layer 3 security measures.	Tighten up trunk configurations and the negotiation state of unused ports.  Place unused ports in a common VLAN.
Attacks between devices on a common VLAN	Devices may need protection from one another, even though they are on a common VLAN. This is especially true on service-provider segments that support devices from multiple customers.	Implement private VLANs (PVLANS).
<b>Spoofing Attacks</b>		
DHCP starvation and DHCP spoofing	An attacking device can exhaust the address space available to the DHCP servers for a period of time or establish itself as a DHCP server in man-in-the-middle attacks.	Use DHCP snooping.
Spanning-tree compromises	Attacking device spoofs the root bridge in the STP topology. If successful, the network attacker can see a variety of frames.	Proactively configure the primary and backup root devices.  Enable RootGuard.
MAC spoofing	Attacking device spoofs the MAC address of a valid host currently in the CAM table. Switch then forwards to the attacking device any frames that are destined for the valid host.	Use DHCP snooping, port security.
Address Resolution Protocol (ARP) spoofing	Attacking device crafts ARP replies intended for valid hosts. The MAC address of the attacking device then becomes the destination address found in the Layer 2 frames that were sent by the valid network device.	Use Dynamic ARP Inspection (DAI).  DHCP snooping, port security.
<b>Switch Device Attacks</b>		
Cisco Discovery Protocol manipulation	Information sent through Cisco Discovery Protocol is transmitted in clear text and unauthenticated, allowing it to be captured and to divulge network topology information.	Disable Cisco Discovery Protocol on all ports where it is not intentionally used.
Secure Shell (SSH) Protocol and Telnet attacks	Telnet packets can be read in clear text. SSH is an option but has security issues in version 1.	Use SSH version 2.  Use Telnet with vty ACLs.



# MAC Flooding Attack

This topic describes how port security is used to block input from devices based on Layer 2 restrictions.



A common Layer 2 or switch attack is MAC flooding, which results in an overflow of the CAM table of a switch. The overflow causes the flooding of regular data frames out all switch ports. This attack can be launched for the malicious purpose of collecting a broad sample of traffic or as a denial of service (DoS) attack.

The CAM tables of a switch are limited in size and therefore can contain only a limited number of entries at any one time. A network intruder can maliciously flood a switch with a large number of frames from a range of invalid source MAC addresses. If enough new entries are made before old ones expire, new valid entries will not be accepted. Then, when traffic arrives at the switch for a legitimate device that is located on one of the switch ports that was not able to create a CAM table entry, the switch must flood the frames to that address out all ports. This has two adverse effects:

- The switch traffic forwarding is inefficient and voluminous.
- An intruding device can be connected to any switch port and can capture traffic that is not normally detected on that port.

If the attack is launched before the beginning of the day, the CAM table would be full when the majority of devices are powered on. Then frames from those legitimate devices are unable to create CAM table entries as they power on. If this represents a large number of network devices, the number of MAC addresses that are flooded with traffic will be high, and any switch port will carry flooded frames from a large number of devices.

If the initial flood of invalid CAM table entries is a one-time event, the switch will eventually age out older, invalid CAM table entries, allowing new, legitimate devices to create entries.

Traffic flooding will cease and may never be detected, even though the intruder may have captured a significant amount of data from the network.

As the figure shows, MAC flooding occurs in several steps. The table describes the progression of a MAC flooding attack.

### MAC Flooding Attack Progression

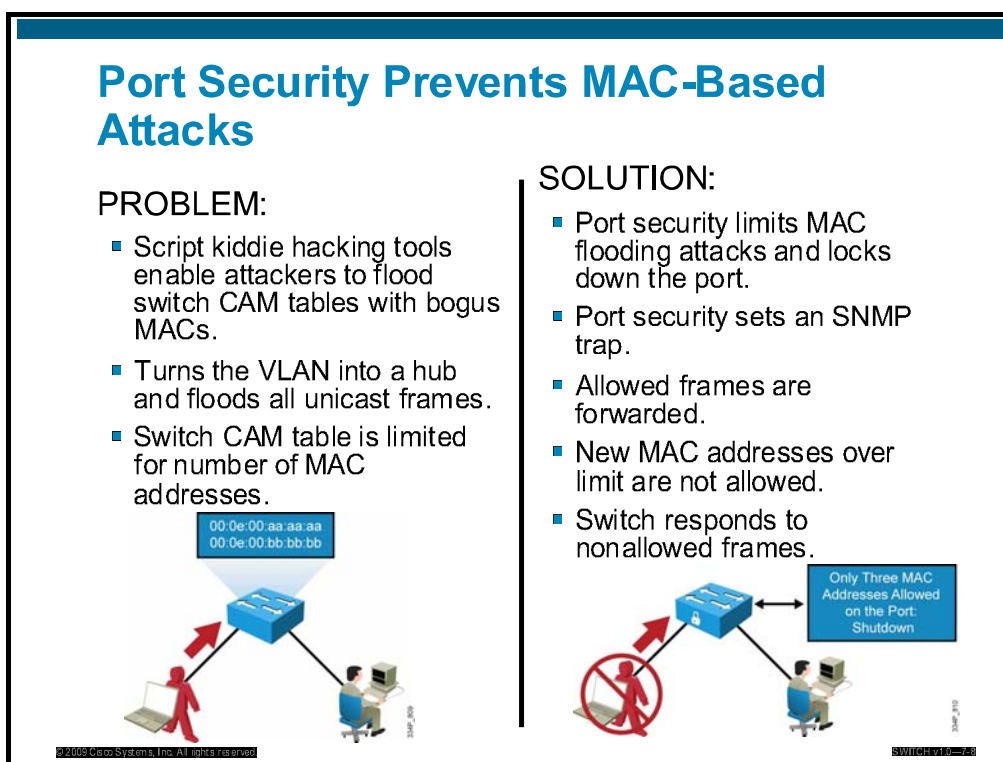
Step	Description
1.	Switch forwards traffic based on valid CAM table entries.
2.	Attacker (MAC address C) sends out multiple packets with various source MAC addresses.
3.	Over a short time period, the CAM table in the switch fills up until it cannot accept new entries. As long as the attack is running, the CAM table on the switch will remain full.
4.	Switch begins to flood all received packets out of every port so that frames sent from host A to host B are also flooded out of port 3 on the switch.

## Suggested Mitigation for MAC Flooding Attacks

Configure port security to define the number of MAC addresses that are allowed on a given port. Port security can also specify which MAC address is allowed on a given port.

# Port Security

This topic describes how port security is used to block input from devices based upon Layer 2 restrictions.



Port security, a feature that is supported on Cisco Catalyst switches, restricts a switch port to a specific set or number of MAC addresses. Those addresses can be learned dynamically or configured statically. The port will then provide access to frames from only those addresses. If, however, the number of addresses is limited to four but no specific MAC addresses are configured, the port will allow any four MAC addresses to be learned dynamically, and port access will be limited to those four dynamically learned addresses.

A port security feature called “sticky learning,” available on some switch platforms, combines the features of dynamically learned and statically configured addresses. When this feature is configured on an interface, the interface converts dynamically learned addresses to “sticky secure” addresses. This adds them to the running configuration as if they were configured with the **switchport port-security mac-address** command.

## Scenario

Imagine five individuals whose laptops are allowed to connect to a specific switch port when they visit an area of the building. You want to restrict switch port access to the MAC addresses of those five laptops and allow no addresses to be learned dynamically on that port.

## Process

The table describes the process that can achieve the desired results for this scenario.

### Implementing Port Security


Step	Action	Notes
1.	Configure port security.	Configure port security to allow only five connections on that port. Configure an entry for each of the five allowed MAC addresses. This configuration, in effect, populates the MAC address table with five entries for that port and allows no additional entries to be learned dynamically.
2.	Allowed frames are processed.	When frames arrive on the switch port, their source MAC address is checked against the MAC address table. If the frame source MAC address matches an entry in the table for that port, the frames are forwarded to the switch to be processed like any other frames on the switch.
3.	New addresses are not allowed to create new MAC address table entries.	When frames with a nonallowed MAC address arrive on the port, the switch determines that the address is not in the current MAC address table and does not create a dynamic entry for that new MAC address, because the number of allowed addresses has been limited.
4.	Switch takes action in response to nonallowed frames.	The switch will disallow access to the port and take one of these configuration-dependent actions: (a) the entire switch port can be shut down; (b) access can be denied for that MAC address only and a log error can be generated; (c) access can be denied for that MAC address but without generating a log message.
<b>Note</b> Port security cannot be applied to trunk ports where addresses might change frequently. Implementations of port security vary by Cisco Catalyst platform. Check documentation to determine whether particular hardware supports this feature and how the hardware supports the feature.		

# Configure Port Security

This topic explains the procedure for configuring port security on a switch.

## Configuring Port Security on a Switch

- Enable port security.
- Set MAC address limit.
- Specify allowable MAC addresses (optional).
- Define violation actions (shut down / protect / restrict).
- Configure address aging (optional).



```
switch(config)# interface fa0/1
switch(config-if)# description Access Port
switch(config-if)# switchport mode access
switch(config-if)# switchport access vlan 2
switch(config-if)# switchport port-security
switch(config-if)# switchport port-security maximum 2
switch(config-if)# switchport port-security mac-address 0000.1111.2222
switch(config-if)# switchport port-security mac-address 0000.1111.3333
switch(config-if)# switchport port-security violation restrict
switch(config-if)# switchport port-security aging time 60
switch(config-if)# switchport port-security aging type inactivity
```

To configure port security so that switch port access is limited to a finite number and a specific set of end-device MAC addresses, follow the steps in the table.

To configure port security, follow the steps listed in the table.

### Port Security Configuration Steps

Step	Description
1.	Enables port security. <code>Switch(config-if)#switchport port-security</code>
2.	Sets a maximum number of MAC addresses that will be allowed on this port. The default is one. <code>Switch(config-if)#switchport port-security maximum value</code>
3.	(Optional) Specifies which MAC addresses will be allowed on this port <code>Switch(config-if)#switchport port-security mac-address mac-address</code> <code>Switch(config-if)#switchport port-security mac-address mac-address</code>
4.	Defines what action an interface will take if a nonallowed MAC address attempts access. <code>Switch(config-if)#switchport port-security violation {shutdown   restrict   protect}</code>

## Caveats to Port Security Configuration Steps

There are some caveats to bear in mind:

**Step 1** Port security is enabled on a port-by-port basis.

**Step 2** By default, only one MAC address is allowed access through a given switch port when port security is enabled. This parameter increases that number. It implies no restriction on specific MAC addresses, but only on the total number of addresses that can be learned by the port. Learned addresses are not aged out by default but can be configured to do so after a specified time when you use the **switchport port-security aging** command. The *value* parameter can be any number from 1 to 1024, with some restrictions related to the number of ports on a given switch with port security enabled.

---

**Note** Be sure to set the *value* parameter to a value of **2** when you are configuring a port to support VoIP with a phone and computer that are accessible on the port. If the default value is used, a port-security violation will result.

---

**Step 3** Access to the switch port can be restricted to one or more specific MAC addresses. If the number of specific MAC addresses that are assigned when you use this command is lower than the *value* parameter that you set in Step 2, then the remaining allowed addresses can be learned dynamically. If you specify a set of MAC addresses that is equal to the maximum number allowed, access is limited to that set of MAC addresses.

**Step 4** By default, if the maximum number of connections is achieved and a new MAC address attempts to access the port, the switch must take one of these actions:

- **Protect:** Frames from the nonallowed address are dropped, but there is no log of the violation.

---

**Note** The *protect* argument is platform or version dependent.

---

- **Restrict:** Frames from the nonallowed address are dropped, a log message is created, and a Simple Network Management Protocol (SNMP) trap is sent.
- **Shut down:** If any frames are detected from a nonallowed address, the interface is errdisabled, a log entry is made, an SNMP trap is sent, and manual intervention or errdisable recovery must be used to make the interface usable.

# Verifying Port Security

This subtopic describes how to verify port security.

## Verifying Port Security

```
switch# show port-security [interface intf-id] [address]
```

```
switch# show port-security interface fastethernet0/1
Port Security           : Enabled
Port Status             : Secure-up
Violation Mode          : Restrict
Aging Time              : 60 mins
Aging Type              : Inactivity
SecureStatic Address Aging : Enabled
Maximum MAC Addresses   : 2
Total MAC Addresses     : 1
Configured MAC Addresses : 0
Sticky MAC Addresses    : 0
Last Source Address:Vlan : 001b.d513.2ad2:5
Security Violation Count : 0
```

You can use the **show port-security** command to verify the ports on which port security has been enabled. It also displays count information and security actions to be taken per interface.

The full command syntax is as follows:

Switch#**show port-security** [interface *intf\_id*] **address**

Arguments are provided to view the port security status by interface or to view the addresses that are associated with port security on all interfaces.

Use the *interface* argument to provide output for a specific interface.

## Verifying Port Security (Cont.)

```
switch# show port-security
Secure Port MaxSecureAddr CurrentAddr SecurityViolation Security Action
          (Count)          (Count)          (Count)
-----
Fa0/1      2                1                0                Restrict
-----
Total Addresses in System (excluding one mac per port) : 0
Max Addresses limit in System (excluding one mac per port) : 6144
```

```
switch# show port-security address
Secure Mac Address Table
-----
Vlan    Mac Address      Type                Ports    Remaining Age
      (mins)
-----
2       001b.d513.2ad2   SecureDynamic       Fa0/1    60 (I)
-----
Total Addresses in System (excluding one mac per port) : 0
Max Addresses limit in System (excluding one mac per port) : 6144
```

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-10-2-1

You can use the **show port-security** command to verify the ports on which port security has been enabled. It also displays count information and security actions to be taken per interface.

The full command syntax is as follows:

Switch#**show port-security** [interface *inif\_id*] **address**

Arguments are provided to view the port security status by interface or to view the addresses that are associated with port security on all interfaces.

Use the *address* argument to display MAC address table security information. The remaining age column is populated only when it is specifically configured for a given interface.

The example displays output from the **show port-security address** privileged EXEC command.



# Port Security with Sticky MAC Addresses

This subtopic describes the sticky MAC option with port security.

## Configuring Sticky MAC Addresses

```
switch(config)# interface fa0/1
switch(config-if)# switchport port-security mac-address sticky
```

```
switch# show port-security address
Secure Mac Address Table
-----
Vlan    Mac Address      Type                Ports    Remaining Age
(mins)
----    -
2       001b.d513.2ad2   SecureSticky        Fa0/1    -
```

```
switch# show running-config fastethernet 0/1
interface FastEthernet0/1
 switchport access vlan 2
 switchport mode access
 switchport port-security maximum 2
 switchport port-security
 switchport port-security violation restrict
 switchport port-security mac-address sticky
 switchport port-security mac-address sticky 001b.d513.2ad2
```

Port security can be used to mitigate spoof attacks by limiting access through each switch port to a single MAC address. This prevents intruders from using multiple MAC addresses over a short time period but does not limit port access to a specific MAC address. The most restrictive port security implementation would specify the exact MAC address of the single device that is to gain access through each port. Implementing this level of security, however, requires considerable administrative overhead.

Port security has a feature called sticky MAC addresses that can limit switch port access to a single, specific MAC address without the network administrator having to gather the MAC address of every legitimate device and manually associate it with a particular switch port.

When sticky MAC addresses are used, the switch port will convert dynamically learned MAC addresses to sticky MAC addresses and subsequently add them to the running configuration as if they were static entries for a single MAC address to be allowed by port security. Sticky secure MAC addresses will be added to the running configuration but will not become part of the startup configuration file unless the running configuration is copied to the startup configuration after addresses have been learned. If they are saved in the startup configuration, they will not have to be relearned upon switch reboot, and this provides a higher level of network security.

The command that follows will convert all dynamic port-security learned MAC addresses to sticky secure MAC addresses:

**switchport port-security mac-address sticky**

This command cannot be used on ports where voice VLANs are configured.

# Authentication and Authorization Methods

This topic describes security in a multilayer switched network.

## AAA Network Configuration

- Authentication
  - Verifies a user identify
- Authorization
  - Specifies the permitted tasks for the user
- Accounting
  - Provides billing, auditing, and monitoring



Authentication, authorization, and accounting (AAA) network security services provide the primary framework through which access control is set up on a switch. AAA is an architectural framework for configuring a set of three independent security functions in a consistent manner. AAA provides a modular way of performing these services. For purposes of this course, only authentication will be discussed.

Authentication is the way that a user is identified before being allowed access to the network and network services. You configure AAA authentication by defining a list of named authentication methods and then applying that list to various interfaces. The method list defines the types of authentication to be performed and the sequence in which they will be performed; it must be applied to a specific interface before any of the defined authentication methods will be performed.

The only exception is the default method list (which is named “default”). The default method list is automatically applied to all interfaces if no other method list is defined. A defined method list overrides the default method list.


In many circumstances, AAA uses protocols such as RADIUS, TACACS+, or IEEE 802.1X to administer its security functions. If the switch is acting as a network access server, AAA is the means through which the switch establishes communication between the network access server and the RADIUS, TACACS+, or 802.1X security server.

# Local User Authentication with AAA

This subtopic describes how to configure user authentication with AAA.

## Configuring User AAA Authentication

- Enable AAA.
- Configure RADIUS server.
- Configure authentication methods.
- Apply methods to interfaces.



```
sw(config)# username admin password cisco
sw(config)# aaa new-model
sw(config)# radius-server host 10.1.1.50 auth-port 1812 key xyz123
sw(config)# aaa authentication login default group radius local line
sw(config)# aaa authentication login NO_AUTH none
sw(config)# Line vty 0 15
sw(config-li)# login authentication default
sw(config-li)# password sanfran
sw(config-li)# line console 0
sw(config-li)# login authentication NO_AUTH
```

To configure user authentication, start by using the **aaa new-model** command. As soon as you enter this command, the default authentication methods do not apply any more, and you need to define how access is granted to the console or vty lines.

In this example, several methods are used.

The first method is the default method, specified by the line **aaa authentication login default**. This line defines that, by default, the RADIUS group should be used to authenticate users who are trying to log in to the local device. If the servers that are defined in the RADIUS group do not answer, local authentication is possible. If local authentication is not possible (because no local user was defined), line authentication is used. Several methods are allowed, but the main method here is RADIUS. Local authentication will be used only if the RADIUS servers are unreachable. Similarly, line authentication will be used only if local authentication is not possible. If RADIUS servers respond and authentication fails, the local or line method will not be used and the user will be denied access. If the RADIUS servers do not answer and a local user is defined, local authentication will be used (it may succeed or fail) and line authentication will not be used.

Because this default method defines “group RADIUS” as the first default method, one or several RADIUS servers must be defined. This example defines 10.1.1.50 as the RADIUS server, listening on UDP port 1812. The shared secret used to be allowed to query the RADIUS is xyz123.

Because the default method also allows local authentication, a local username and password pair are defined. You can create several users.

Because the default method also allows line authentication, a password is defined on the vty lines.

In a real implementation, you would probably use two of these methods but not all three. If local authentication is allowed, there is little purpose in defining a possible line authentication. All three methods are mentioned here to show that several backup methods can be used in combination with a primary authentication method. Each method requires specific configuration.

In the example, a second **aaa authentication login** mechanism is used. It is called **NO\_AUTH**, and is defined as needing no authentication when you use this method.

On vty lines, the default method is used, thus requiring RADIUS authentication, and then local or line authentication as secondary or tertiary methods.

On the console line, the **NO\_AUTH** method is called, which implies that no authentication is needed when you are connecting through the console.

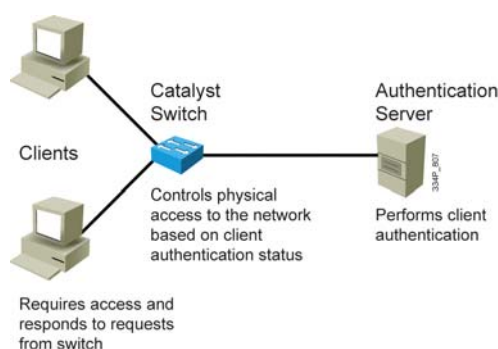
You can see in this example that “default” is the name of an authentication method. This name is special in the sense that all lines requiring authentication will use the default method if no other specific method is called on the corresponding line. This means, for example, that if you did not specify a method on the console line, it would have used the default method. You still have to specify login requirements. In the previous example, if you remove the line login authentication **NO\_AUTH** completely, no login requirement is specified for the console line, and no console connection is possible anymore. If you specify “login” without defining the method to use, the “default” method is called by default.

# 802.1X Port-Based Authentication

This topic describes IEEE 802.1X port-based authentication.

## 802.1X Port-Based Authentication

Network access through switch requires RADIUS authentication.



The 802.1X standard defines a port-based access control and authentication protocol that restricts unauthorized workstations from connecting to a LAN through publicly accessible switch ports. The authentication server authenticates each workstation that is connected to a switch port before making available any services that are offered by the switch or the LAN.

Until the workstation is authenticated, 802.1X access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the workstation is connected. After authentication succeeds, normal traffic can pass through the port.

With 802.1X port-based authentication, the devices in the network have specific roles, as follows:

- **Client:** The device (workstation) that requests access to LAN and switch services, and then responds to requests from the switch. The workstation must be running 802.1X-compliant client software, such as what is offered in the Microsoft Windows XP operating system. (The port that the client is attached to is the supplicant [client] in the 802.1X specification.)
- **Authentication server:** Performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. The RADIUS security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server.

- **Switch (also called the authenticator):** Controls physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client (supplicant) and the authentication server, requesting identifying information from the client, verifying that information with the authentication server, and relaying a response to the client. The switch uses a RADIUS software agent, which is responsible for encapsulating and decapsulating the EAP frames and interacting with the authentication server.

The switch port state determines whether the client is granted access to the network. The port starts in the unauthorized state. While in this state, the port disallows all ingress and egress traffic except for 802.1X protocol packets. When a client is successfully authenticated, the port transitions to the authorized state, allowing all traffic for the client to flow normally.

If the switch requests the client identity (authenticator initiation) and the client does not support 802.1X, the port remains in the unauthorized state, and the client is not granted access to the network.

In contrast, when an 802.1X-enabled client connects to a port and the client initiates the authentication process (supplicant initiation) by sending the EAPOL-start frame to a switch that is not running the 802.1X protocol, no response is received, and the client begins sending frames as if the port is in the authorized state.

You control the port authorization state by using the **dot1x port-control** interface configuration command and these keywords:

- **force-authorized:** Disables 802.1X port-based authentication and causes the port to transition to the authorized state without any authentication exchange required. The port transmits and receives normal traffic without 802.1X-based authentication of the client. This is the default setting.
- **force-unauthorized:** Causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.
- **auto:** Enables 802.1X port-based authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port. The authentication process begins when the link state of the port transitions from down to up (authenticator initiation) or when an EAPOL-start frame is received (supplicant initiation). The switch requests the identity of the client and begins relaying authentication messages between the client and the authentication server. The switch uniquely identifies each client that is attempting to access the network by using the client MAC address.

If the client is successfully authenticated (that is, if it receives an “accept” frame from the authentication server), the port state changes to authorized, and all frames from the authenticated client are allowed through the port.

If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can retransmit the request. If no response is received from the server after the specified number of attempts, authentication fails, and network access is not granted.


When a client logs out, it sends an EAPOL-logout message, causing the switch port to transition to the unauthorized state.

# Configuring 802.1X

This subtopic describes configuring 802.1X port-based authentication.

## Configuring 802.1X

- Enable AAA.
- Configure RADIUS server.
- Enable 802.1X globally.
- Configure interface for 802.1X.
- Define local user authentication.



```
sw(config)# aaa new-model
sw(config)# radius-server host 10.1.1.50 auth-port 1812 key xyz123
sw(config)# aaa authentication dot1x default group radius
sw(config)# dot1x system-auth-control
sw(config)# interface fa0/1
sw(config-if)# description Access Port
sw(config-if)# switchport mode access
sw(config-if)# dot1x port-control auto
```

The following procedure explains how to implement 802.1X port-based authentication.

### Implementing 802.1X Port-Based Authentication

Step	Description
1.	Enable AAA. <code>Switch(config)#aaa new-model</code>
2.	Create an 802.1X port-based authentication method list. <code>Switch(config)#aaa authentication dot1x {default} method1 [method2...]</code>
3.	Globally enable 802.1X port-based authentication. <code>Switch(config)#dot1x system-auth-control</code>
4.	Enter interface configuration mode and specify the interface to be enabled for 802.1X port-based authentication. <code>Switch(config)#interface type slot/port</code>
5.	Enable 802.1X port-based authentication on the interface. <code>Switch(config-if)#dot1x port-control auto</code>
6.	Return to privileged EXEC mode. <code>Switch(config)#end</code>

Be aware that entering the **aaa new-model** command disables the standard authentication process on the switch. You also need to redefine user login policies as specified in the earlier subtopic “Local User Authentication with AAA.”

# Summary

This topic summarizes the key points that were discussed in this lesson.

## Summary

- Layer 2 security measures must be taken as a subset of the overall network security plan.
- Rogue devices can allow access to the network and undermine the security.
- Switch attacks fall into four main categories.
- MAC flooding attacks are launched against Layer 2 access switches and can cause the CAM table to overflow.
- Port security can be configured at Layer 2 to block input from devices.
- Sticky MAC addresses allow port security to limit access to a specific, dynamically learned MAC address.
- AAA can be used for authentication on a multilayer switch.
- 802.1x port-based authentication can mitigate risk of rogue devices gaining unauthorized access.



## Lesson 2

---

# Protecting Against VLAN Attacks

---

## Overview

On networks using trunking protocols, there is a possibility of rogue traffic “hopping” from one VLAN to another, thereby creating security vulnerabilities. These VLAN hopping attacks are best mitigated by close control of trunk links.

You can configure close control of trunk links to mitigate VLAN hopping attacks and configure VLAN access control lists (VACLs) to filter traffic within a VLAN.

Private VLANs (PVLANS) can be configured to establish security regions within a single VLAN without subnetting, and VACLs can be used to filter traffic within a VLAN.

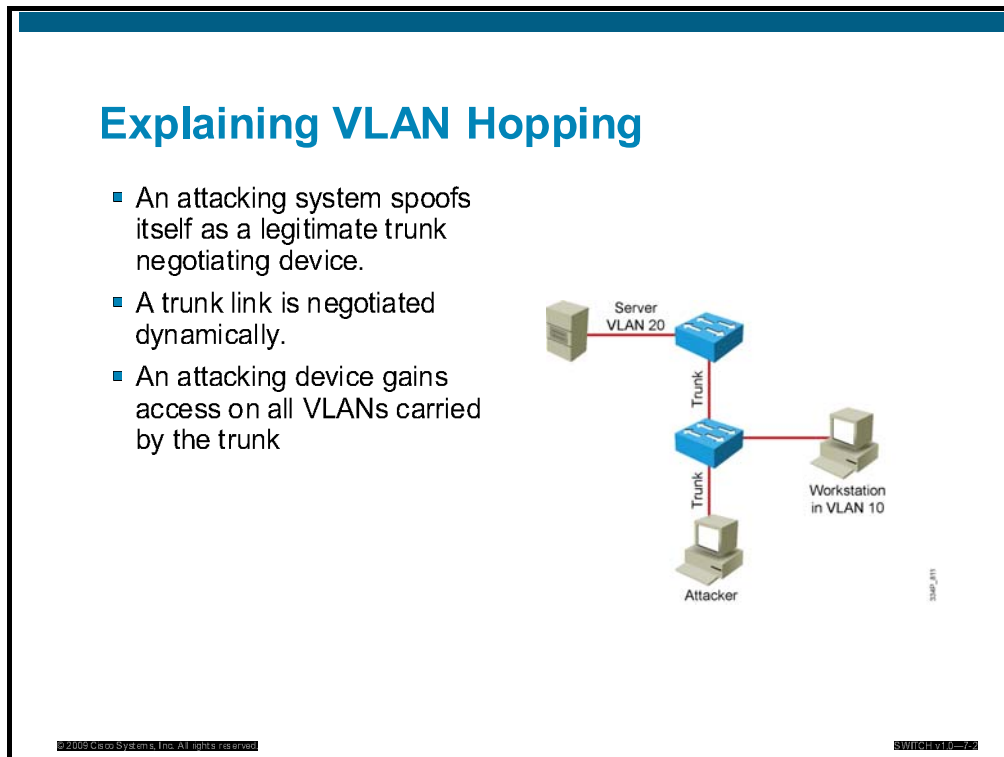
## Objectives

Upon completing this lesson, you will be able to configure various features to prevent VLAN hopping and to address VLAN security issues. This ability includes being able to meet these objectives:

- Describe how VLAN hopping occurs and why it is a security vulnerability
- Explain the procedure for configuring a switch to mitigate VLAN hopping attacks
- Describe VACLs and their purpose as part of VLAN security
- Explain the procedure for configuring VACLs

# Explaining VLAN Hopping

This topic describes how VLAN hopping occurs and why it is a security vulnerability.



VLAN hopping is a network attack whereby an end system sends packets to, or collects packets from, a VLAN that should not be accessible to that end system. This is accomplished by tagging the invasive traffic with a specific VLAN ID (VID) or by negotiating a trunk link to send or receive traffic on penetrated VLANs. VLAN hopping can be accomplished by switch spoofing or double tagging.

## Switch Spoofing

In a switch spoofing attack, the network attacker configures a system to spoof itself as a switch. The attack emulates Inter-Switch Link (ISL) or IEEE 802.1Q signaling along with Dynamic Trunking Protocol (DTP). This is signaling in an attempt to establish a trunk connection to the switch.

Any switch port configured as DTP auto, upon receipt of a DTP packet generated by the attacking device, may become a trunk port and thereby accept traffic that is destined for any VLAN supported on that trunk. The malicious device can then send packets to, or collect packets from, any VLAN that is carried on the negotiated trunk.

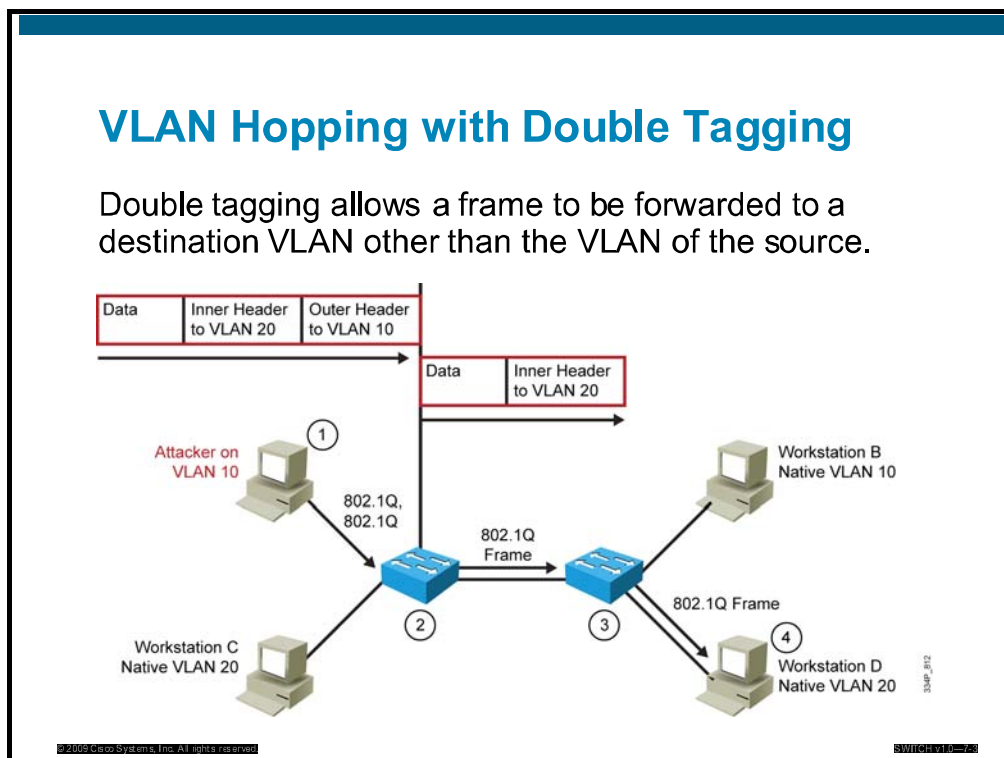
The table describes the switch spoofing sequence of events.

## Switch Spoofing Sequence of Events

Step	Description
1.	Attacker gains access to a switch port and sends DTP negotiation frames toward a switch with DTP running and autonegotiation turned on (often, the default settings).
2.	Attacker and switch negotiate trunking over the port.
3.	Switch allows all VLANs (default) to traverse the trunk link.
4.	Attacker sends data to, or collects data from, all VLANs carried on that trunk.

# VLAN Hopping With Double Tagging

This subtopic describes double tagging as a means of VLAN hopping.



Another method of VLAN hopping is for any workstation to generate frames with two 802.1Q headers to cause the switch to forward the frames onto a VLAN that would be inaccessible to the attacker through legitimate means.

The first switch to encounter the double-tagged frame strips the first tag off the frame, because the first tag (VLAN 10) matches the trunk port native VLAN, and then forwards the frame out.

The result is that the frame is forwarded, with the inner 802.1Q tag, out all the switch ports, including trunk ports that are configured with the native VLAN of the network attacker. The second switch then forwards the packet to the destination based on the VLAN ID in the second 802.1Q header. If the trunk does not match the native VLAN of the attacker, the frame is untagged and is flooded to only the original VLAN.

The table describes the double-tagging method of VLAN hopping.

## Double-Tagging Method of VLAN Hopping

Step	Description
1.	Workstation A (native VLAN 10) sends a frame with two 802.1Q headers to switch 1.
2.	Switch 1 strips the outer tag and forwards the frame to all ports within same native VLAN.
3.	Switch 2 interprets frame according to information in the inner tag marked with VLAN ID 20.
4.	Switch 2 forwards the frame out all ports associated with VLAN 20, including trunk ports.

# Mitigating VLAN Hopping

This topic describes how to mitigate VLAN hopping attacks.

## Mitigating VLAN Hopping

### Unused ports

- Shut down all unused ports.
- Configure all unused ports to access mode.
- Configure an access VLAN on all unused ports to an unused VLAN.
- Configure a native trunk VLAN on all unused ports to an unused VLAN.

### Trunk ports

- Configure a trunk port with trunk mode on, and disable trunk negotiation.
- Configure a native trunk VLAN on trunk ports to an unused VLAN.
- Configure the allowed VLANs on the trunk ports, and do not allow a native VLAN.

The measures for defending the network from VLAN hopping are a series of best practices for configuring all switch ports and parameters when establishing a trunk port.

- Configure all unused ports as access ports so that trunking cannot be negotiated across those links.
- Place all unused ports in the shutdown state and associate them with a VLAN that is designed for only unused ports, carrying no user data traffic.
- When establishing a trunk link, purposefully configure arguments to achieve the following results:
  - The native VLAN is different from any data VLANs.
  - Trunking is set up as “on,” rather than as “negotiated.”
  - The specific VLAN range is carried on the trunk.

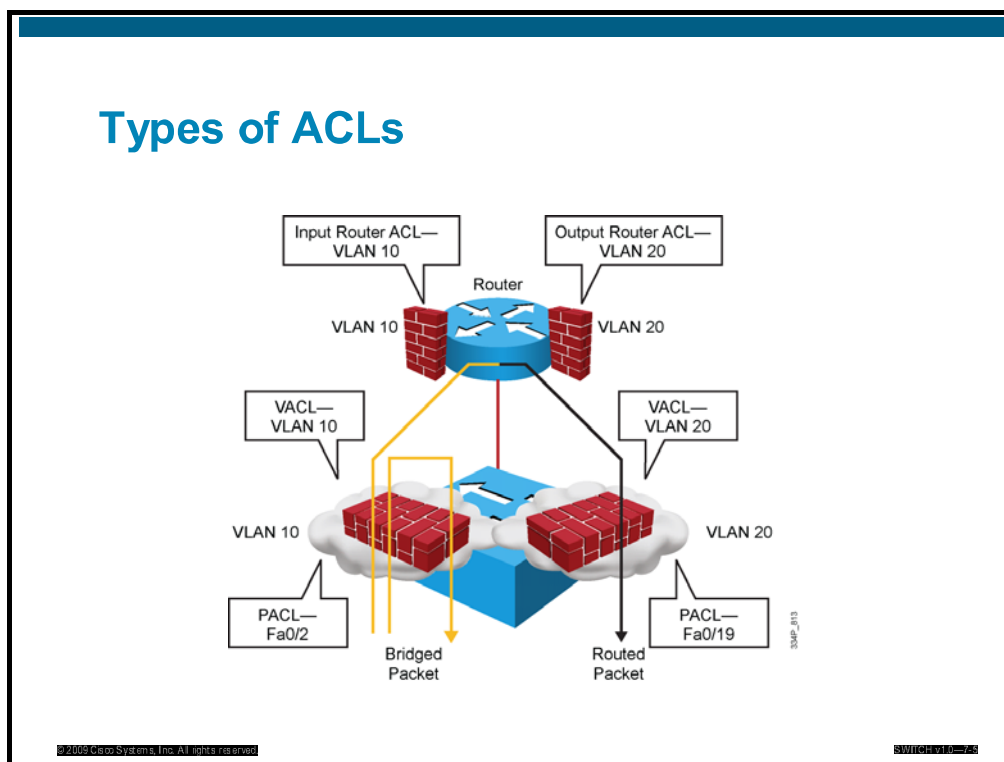
---

**Note** The configuration commands in the figure will not work on access ports that support VoIP because they will be configured as trunk ports. However, on all other access ports, it is best practice to apply these commands to mitigate VLAN hopping.

---

# VLAN Access Control Lists

Access control lists (ACLs) are useful for controlling access in a multilayer switched network. This topic describes VACLs and their purpose as part of VLAN security.



Cisco multilayer switches support three types of ACLs:

- **Router ACL:** Supported in the ternary content addressable memory (TCAM) hardware on Cisco multilayer switches. In Cisco Catalyst switches, a router ACL can be applied to any routed interface, such as a switch virtual interface (SVI) or a Layer 3 routed port.
- **Port access control list (PACL):** Filters traffic at the port level. PACLs can be applied on a Layer 2 switch port, trunk port, or EtherChannel port.
- **VACL:** Supported in software on Cisco multilayer switches.

Catalyst switches support four ACL lookups per packet: input and output security ACL, and input and output quality of service (QoS) ACL.

Catalyst switches use two methods of performing a merge: order independent and order dependent. With an order-independent merge, ACLs are transformed from a series of order-dependent actions to a set of order-independent masks and patterns. The resulting access control entry (ACE) can be very large. The merge is processor and memory intensive.

An order-dependent merge is a recent improvement on some Catalyst switches in which ACLs retain their order-dependent aspect. The computation is much faster and is less processor-intensive.

Router ACLs are supported in hardware through IP standard ACLs and IP extended ACLs, with permit and deny actions. ACL processing is an intrinsic part of the packet-forwarding process. ACL entries are programmed in hardware. Lookups occur in the pipeline, whether ACLs are configured or not. With router ACLs, access list statistics and logging are not supported.

# Configure VACLs

This topic describes how to configure VACLs.

## Configuring VACLs

- Create an access list.
- Configure an access map.
- Create a VLAN filter.
- Example: Drop all traffic from network 10.1.9.0/24 on VLAN 10 and 20, and drop all traffic to backup server 0000.1111.4444.

```
switch(config)# access-list 100 permit ip 10.1.9.0 0.0.0.255 any
Switch(config)# mac access-list extended BACKUP_SERVER
Switch(config-ext-mac)# permit any host 0000.1111.4444
switch(config)# vlan access-map 100 10
switch(config-map)# match ip address 100
switch(config-map)# action drop
switch(config-map)# vlan access-map 100 20
switch(config-map)# match mac address BACKUP_SERVER
Switch(config-map)# action drop
switch(config-map)# vlan access-map 100 30
switch(config-map)# action forward
switch(config)# vlan filter 100 vlan-list 10,20
```

VACLs (also called VLAN access maps in Cisco IOS Software) apply to all traffic on the VLAN. You can configure VACLs for IP and for MAC-layer traffic.

VACLs follow route-map conventions, in which map sequences are checked in order.

When a matching permit ACE is encountered, the switch takes the action. When a matching deny ACE is encountered, the switch checks the next ACL in the sequence or checks the next sequence.

Three VACL actions are permitted:

- **Permit** (with capture, Cisco Catalyst 6500 Series Switches only)
- **Redirect** (Catalyst 6500 Series Switches only)
- **Deny** (with logging, Catalyst 6500 Series Switches only)

The VACL capture option copies traffic to specified capture ports. VACL ACEs that are installed in hardware are merged with Router ACLs and other features.

Two features are supported on only the Catalyst 6500 Series Switches:

- **VACL capture:** Forwarded packets are captured on capture ports. The capture option is on only permit ACEs. The capture port can be an intrusion detection system (IDS) monitor port or any Ethernet port. The capture port must be in an output VLAN for Layer 3 switched traffic.
- **VACL redirect:** Matching packets are redirected to specified ports. You can configure up to five redirect ports. Redirect ports must be in a VLAN where a VACL is applied.

To configure VACLs, complete the following steps.

## Configuring VACLs

Step	Description
1.	Define a VLAN access map. <code>Switch(config)#<b>vlan access-map</b> map_name [seq#]</code>
2.	Configure a match clause. <code>Switch(config-access-map)#<b>action</b> {drop [log]}   {forward [capture]}   {redirect {{fastethernet   gigabitethernet   tengigabitethernet} slot/port}   {port-channel channel_id}}</code>
3.	Configure an action clause. <code>Switch(config-access-map)#<b>action</b> {drop [log]}   {forward [capture]}   {redirect {{fastethernet   gigabitethernet   tengigabitethernet} slot/port}   {port-channel channel_id}}</code>
4.	Apply a map to VLANs. <code>Switch(config)#<b>vlan filter</b> map_name vlan_list list</code>
5.	Verify the VACL configuration. <code>Switch#<b>show vlan access-map</b> map_name</code> <code>Switch#<b>show vlan filter</b> [ access-map map_name   vlan_id ]</code>



# Summary

This topic summarizes the key points discussed in this lesson.

## Summary

- VLAN hopping can allow Layer 2 unauthorized access to another VLAN.
- VLAN hopping can be mitigated by:
  - Properly configuring 802.1Q trunks
  - Turning off trunk negotiation
- Access lists can be applied to VLANs to limit Layer 2 access.
- VACLs can be configured on Cisco Catalyst switches.



# Protecting Against Spoofing Attacks

---

## Overview

Spoofing attacks can occur because several protocols allow a reply from a host even if a request was not received. By spoofing, or pretending to be another machine, the attacker can redirect part or all the traffic coming from, or going to, a predefined target. After the attack, all traffic from the device under attack flows through the computer of the attacker and then to the router, switch, or host.

A spoofing attack can affect hosts, switches, and routers that are connected to your Layer 2 network by sending false information to the devices that are connected to the subnet. Spoofing attacks can also intercept traffic that is intended for other hosts on the subnet. This lesson describes how to mitigate these attacks, and how to configure switches to guard against DHCP, MAC, and Address Resolution Protocol (ARP) threats.

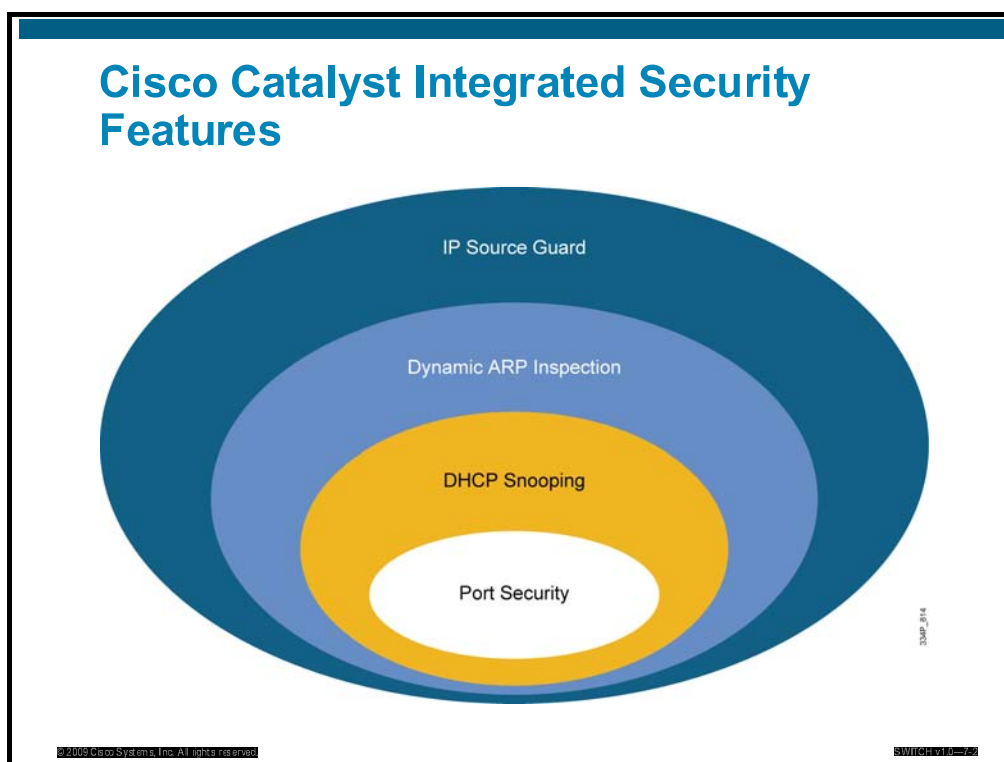
## Objectives

Upon completing this lesson, you will be able to protect your network against spoofing attacks. This ability includes being able to meet these objectives:

- Identify DHCP spoofing attacks
- Prevent attacks by using DHCP snooping
- Configure DHCP snooping
- Describe ARP poisoning
- Protect against ARP spoofing attacks with DAI

# DHCP Spoofing Attacks

This topic describes DHCP spoofing attacks. Protection against them is part of switch integrated security.




Cisco Catalyst integrated security capabilities provide campus security on the Cisco Catalyst switches using integrated tools:

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

## DHCP Spoofing Attacks

- An attacker activates a DHCP server on the VLAN.
- An attacker replies to a valid client DHCP request.
- An attacker assigns IP configuration information that establishes a rogue device as client default gateway.
- An attacker floods the DHCP server with requests.



One of the ways that an attacker can gain access to network traffic is to spoof responses that would be sent by a valid DHCP server. The DHCP spoofing device replies to client DHCP requests. The legitimate server may reply also, but if the spoofing device is on the same segment as the client, its reply to the client may arrive first.

The DHCP reply from the intruder offers an IP address and supporting information that designates the intruder as the default gateway or Domain Name System (DNS) server. In the case of a gateway, the clients will then forward packets to the attacking device, which will in turn send them to the desired destination. This is referred to as a man-in-the-middle attack, and it may go entirely undetected as the intruder intercepts the data flow through the network.

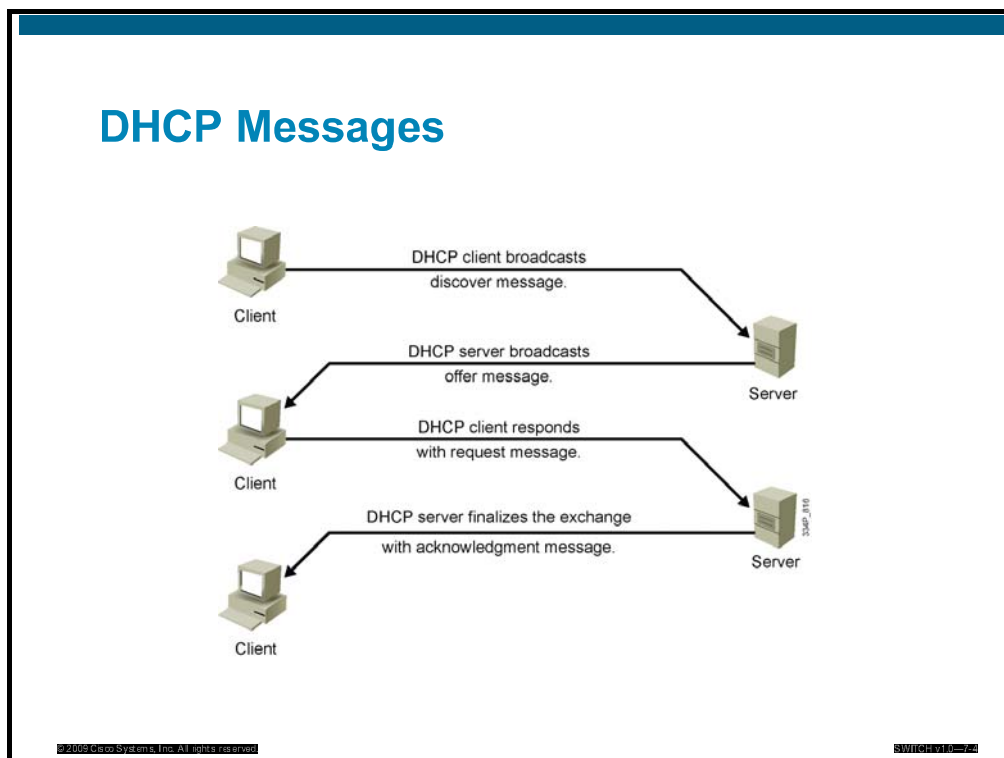
The table describes the DHCP spoofing-attack sequence, as shown in the figure.

### DHCP Spoofing Attack Sequence

Sequence of Events	Description
1.	Attacker hosts a rogue DHCP server off a switch port.
2.	Client broadcasts a request for DHCP configuration information.
3.	The rogue DHCP server responds before the legitimate DHCP server, assigning attacker-defined IP configuration information.
4.	Host packets are redirected to the attacker's address as it emulates a default gateway for the erroneous DHCP address that is provided to the client.

# DHCP

This subtopic describes DHCP.



DHCP uses four messages to provide an IP address to a client:


- DHCP discover broadcast from client
- DHCP offer broadcast to client
- DHCP unicast request from client
- DHCP unicast acknowledgment to client

# DHCP Snooping

This topic describes DHCP snooping.

## DHCP Snooping Protects Against Rogue and Malicious DHCP Servers


- DHCP requests (**discover**) and responses (**offer**) are tracked.
- Rate-limiting requests on untrusted interfaces limit DoS attacks on DHCP servers.
- Deny responses (**offers**) on untrusted interfaces to stop malicious or errant 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 allows the intruder to give out false DHCP information for the default gateway and domain name servers, which points clients to the hacker. This misdirection allows 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 situation. DHCP snooping is a per-port security mechanism that is used to differentiate an untrusted switch port that is connected to an end user from a trusted switch port that is connected to a DHCP server or to 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 requests on client ports, thereby mitigating the effect of DHCP denial-of-service (DoS) attacks from an individual client or access port.

## DHCP Snooping

- DHCP snooping allows the configuration of ports as trusted or untrusted.
- Untrusted ports cannot forward DHCP replies.
- Configure DHCP trust on the uplinks to a DHCP server.
- Do not configure DHCP trust on client ports.



© 2009 Cisco Systems, Inc. All rights reserved.

SWITCHING

DHCP snooping is a Cisco Catalyst feature that determines which switch ports can respond to DHCP requests. Ports are identified as trusted and untrusted. Trusted ports can source all DHCP messages, whereas untrusted ports can source requests only. Trusted ports host a DHCP server or can be an uplink toward the DHCP server. If a rogue device on an untrusted port attempts to send a DHCP response packet into the network, the port is shut down. This feature can be coupled with DHCP option 82, in which switch information, such as the port ID of the DHCP request, can be inserted into the DHCP request packet.

Untrusted ports are those that are not explicitly configured as trusted. A DHCP binding table is built for untrusted ports. Each entry contains the client MAC address, IP address, lease time, binding type, VLAN number, and port ID that are recorded as clients make DHCP requests. The table is then used to filter subsequent DHCP traffic. From a DHCP snooping perspective, untrusted access ports should not send any DHCP server responses, such as DHCP OFFER, DHCP ACK, or DHCP NAK.

Sequence of Configuration	Description
1.	Configure global DHCP snooping.
2.	Configure trusted ports.
3.	Configure Option 82 insertion off (default enabled by Step 2).
4.	Configure rate limiting on untrusted ports.
5.	Configure DHCP snooping for the selected VLANs.



# Configure DHCP Snooping

This topic describes DHCP snooping configuration.

## Configuring DHCP Snooping

- Enable DHCP snooping globally.
- Enable DHCP snooping on selected VLANs.
- Configure trusted interfaces (untrusted is default).
- Configure DHCP rate limit on untrusted interfaces.

```
switch(config)# ip dhcp snooping
switch(config)# ip dhcp snooping information option
switch(config)# ip dhcp snooping vlan 10,20
switch(config)# interface fastethernet 0/1
switch(config-if)# description Access Port
switch(config-if)# ip dhcp limit rate 50
switch(config)# interface fastethernet 0/24
switch(config-if)# description Uplink
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 10,20
switch(config-if)# ip dhcp snooping trust
```

To enable DHCP snooping, use the commands listed in the table.

## Steps for Enabling DHCP Snooping

Step	Comments
1. Enable DHCP snooping globally.  Switch(config)# <b>ip dhcp snooping</b>	By default, the feature is not enabled.
2. Enable DHCP option 82.  Switch(config)# <b>ip dhcp snooping information option</b>	This is optional for the forwarded DHCP request packet to contain information on the switch port where it originated.
3. Configure DHCP server interfaces or uplink ports as trusted.  Switch(config-if)# <b>ip dhcp snooping trust</b>	At least one trusted port must be configured. Use the <b>no</b> keyword to revert to untrusted.  By default, all ports are untrusted.
4. Configure the number of DHCP packets per second (p/s) that are acceptable on the port.  Switch(config-if)# <b>ip dhcp snooping limit rate</b> rate	Configure the number of DHCP p/s that an interface can receive. Normally, the rate limit applies to untrusted interfaces.  This step is used to prevent DHCP starvation attacks by limiting the rate of the DHCP requests on untrusted ports.
5. Enable DHCP snooping on specific VLAN(s).  Switch(config)# <b>ip dhcp snooping vlan</b> number [number]	This step is required for identifying VLANs that will be subject to DHCP snooping.
6. Verify the configuration.  Switch# <b>show ip dhcp snooping</b>	Verify the configuration.

# DHCP Snooping Verification

This subtopic describes DHCP snooping verification.

## Verifying DHCP Snooping

```
switch# show ip dhcp snooping
Switch DHCP snooping is enabled
DHCP snooping is configured on following VLANs:
10,20
DHCP snooping is operational on following VLANs:
10,20
DHCP snooping is configured on the following L3 Interfaces:

Insertion of option 82 is enabled
  circuit-id default format: vlan-mod-port
  remote-id: 001a.e372.ab00 (MAC)
Option 82 on untrusted port is not allowed
Verification of hwaddr field is enabled
Verification of giaddr field is enabled
DHCP snooping trust/rate is configured on the following Interfaces:
```

Interface	Trusted	Allow option	Rate limit (pps)
FastEthernet0/1	no	no	50
FastEthernet0/24	yes	yes	unlimited

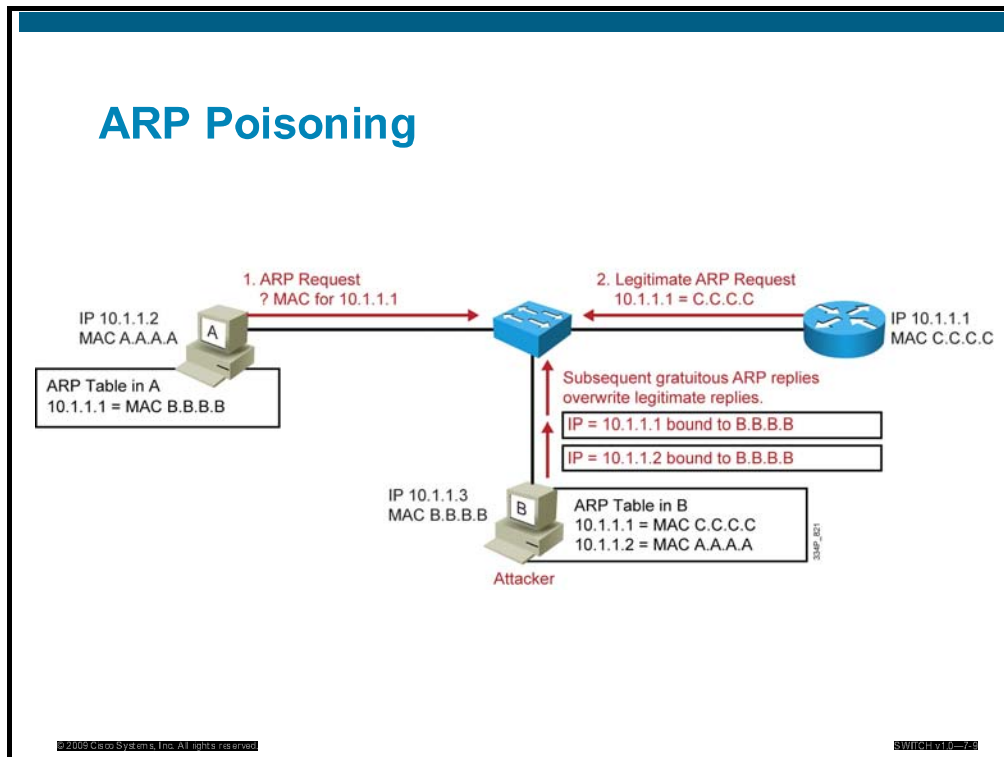
The **show ip dhcp snooping** family of commands is used to display information about the DHCP snooping configuration.

Only ports that are trusted or that have a rate limit applied will be shown in the output. All other ports are untrusted and are not displayed.

In the output, DHCP snooping is configured for VLANs 10 and 20, and is operational on both of them. Interface FastEthernet0/1 has its rate limited and is not trusted, while interface FastEthernet0/24 does not have any rate limitation and is trusted. All the other ports are untrusted and do not have rate limit. They are not displayed.

# ARP Poisoning

This topic describes ARP poisoning.



In normal ARP operation, a host sends a broadcast to determine the MAC address of a host with a particular IP address. The device at that IP address replies with its MAC address. The originating host caches the ARP response, using it to populate the destination Layer 2 header of packets that are sent to that IP address.

By spoofing an ARP reply from a legitimate device with a gratuitous ARP, an attacking device appears to be the destination host that is sought by the senders. The ARP reply from the attacker causes the sender to store the MAC address of the attacking system in its ARP cache. All packets that are destined for those IP addresses will be forwarded through the attacker system.

The figure illustrates the sequence of events in an ARP spoofing attack.

An ARP spoofing attack follows the sequence shown in the table.

### ARP Spoofing Attack

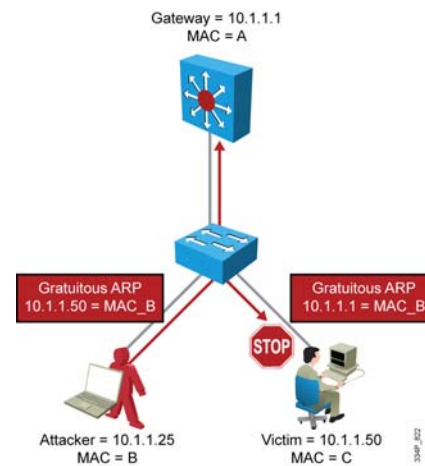
Step or Sequence Number	Description
1.	Host A sends an ARP request for MAC address of C.
2.	Router C replies with its MAC and IP addresses. C also updates its ARP cache.
3.	Host A binds MAC address of C to its IP address in its ARP cache.
4.	Host B (attacker) sends ARP binding MAC address of B to IP address of C.
5.	Host A updates ARP cache with MAC address of B bound to IP address of C.
6.	Host B sends ARP binding MAC address of B to IP address of A.
7.	Router C updates ARP cache with MAC address of B bound to IP address of A.
8.	Packets are diverted through attacker (B).

# Dynamic ARP Inspection

This topic describes Dynamic ARP Inspection (DAI).

## DAI Protection Against ARP Poisoning

- Protects against ARP poisoning (ettercap, dsniff, or arpspoof)
- Uses the DHCP snooping binding table
- Tracks IP-to-MAC bindings from DHCP transactions
- Drops gratuitous ARPs
- Stops ARP poisoning and man-in-the-middle attacks
- Rate-limits ARP requests from client ports; stops port scanning




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 MAC address of the attacker and the IP address of the default gateway. Frames that are intended for default gateways and are sent from hosts with poisoned ARP tables are sent to the hacker (allowing the packets to be sniffed) or to 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 that 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.

## About DAI

- DAI associates each interface with a trusted state or an untrusted state.
- Trusted interfaces bypass DAI.
- Untrusted interfaces undergo DAI validation.
- DHCP snooping is required to build a table with MAC-to-IP bindings for DAI validation.



To prevent ARP spoofing or poisoning, a switch must ensure that only valid ARP requests and responses are relayed. DAI prevents these attacks by intercepting and validating all ARP requests and responses. Each intercepted ARP reply is verified for valid MAC-address-to-IP-address bindings before it is forwarded to a PC to update the ARP cache. ARP replies coming from invalid devices are dropped.

DAI determines the validity of an ARP packet based on a valid MAC-address-to-IP-address bindings database that is built by DHCP snooping. In addition, to handle hosts that use statically configured IP addresses, DAI can validate ARP packets against user-configured ARP access control lists (ACLs).

To ensure that only valid ARP requests and responses are relayed, DAI performs these tasks:

- Forwards ARP packets that are received on a trusted interface without any checks
- Intercepts all ARP packets on untrusted ports
- Verifies that each intercepted packet has a valid IP-to-MAC address binding before forwarding packets that can update the local ARP cache
- Drops, logs, or drops and logs ARP packets with invalid IP-to-MAC address bindings

Configure all access switch ports as untrusted and all switch ports that are connected to other switches as trusted. In this case, all ARP packets that are entering the network would be from an upstream distribution or core switch, bypassing the security check and requiring no further validation.

You can also use DAI to rate-limit the ARP packets and then error-disable the interface if the rate is exceeded.

# Dynamic ARP Inspection Configuration

This subtopic describes DAI configuration.

## Configuring DAI

- Enable DHCP snooping globally.
- Enable DHCP snooping on selected VLANs.
- Enable ARP inspection on selected VLANs.
- Configure trusted interfaces (untrusted is default).

```
switch(config)# ip dhcp snooping
switch(config)# ip dhcp snooping vlan 10,20
switch(config)# ip arp inspection vlan 10,20
switch(config)# interface fastethernet 0/1
switch(config-if)# ip dhcp limit rate 50
switch(config)# interface fastethernet 0/24
switch(config-if)# description Uplink
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 10,20
switch(config-if)# ip dhcp snooping trust
switch(config-if)# ip arp inspection trust
```

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH10-7-10

The table describes the commands you use to configure DAI.

### DAI Commands

Command	Description
Switch(config)# <b>ip arp inspection vlan</b> <i>vlan_id</i> [, <i>vlan_id</i> ]	Enables DAI on a VLAN or range of VLANs
Switch(config-if)# <b>ip arp inspection trust</b>	Enables DAI on an interface and sets the interface as a trusted interface
Switch(config)# <b>ip arp inspection validate</b> { [ <i>src-mac</i> ] [ <i>dst-mac</i> ] [ <i>ip</i> ] }	Configures DAI to drop ARP packets when the IP addresses are invalid, or when the MAC addresses in the body of the ARP packets do not match the addresses that are specified in the Ethernet header

It is generally advisable to configure all access switch ports as untrusted and to configure all uplink ports that are connected to other switches as trusted.

This example of DAI implementation illustrates the configuration that is required on switch 2 with port FastEthernet 3/3 as the uplink port toward the DHCP server.

To mitigate the chances of ARP spoofing, these procedures are recommended:

**Step 1**     Implement protection against DHCP spoofing.

**Step 2**     Enable DAI.

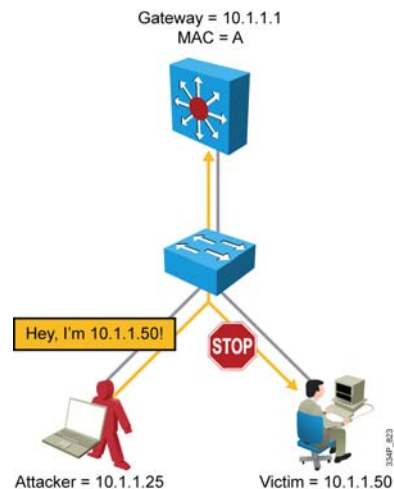


# IP Source Guard

This subtopic describes IP Source Guard.

## IP Source Guard Protection Against Spoofed IP Addresses

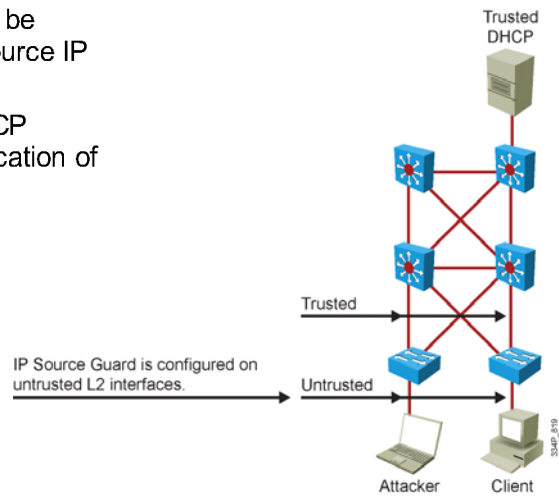
- Protects against spoofed IP addresses
- Uses the DHCP snooping binding table
- Tracks IP addresses to port associations
- Dynamically programs port ACLs to drop traffic not originating from an IP address assigned via DHCP



IP Source Guard prevents a malicious host from attacking the network by hijacking the IP address of its neighbor. 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.

## IP Source Guard

- DHCP snooping must be configured to verify source IP addresses.
- Port security with DHCP snooping allows verification of source IP and MAC addresses.



IP Source Guard is similar to DHCP snooping. You can enable IP Source Guard on a DHCP snooping untrusted Layer 2 port to prevent IP address spoofing. To start, all IP traffic on the port is blocked except for DHCP packets that are captured by the DHCP snooping process.

When a client receives a valid IP address from the DHCP server, or when a static IP source binding is configured by the user, a per-port and VLAN access control list (PVACL) is installed on the port.

This process restricts the client IP traffic to those source IP addresses that are configured in the binding; any IP traffic with a source IP address other than that in the IP source binding will be filtered out. This filtering limits the ability of a host to attack the network by claiming the IP address of a neighbor host.

---

**Note** If IP Source Guard is enabled on a trunk port with a large number of VLANs that have DHCP snooping enabled, you might run out of ACL hardware resources, and some packets might be switched in software.

---

IP Source Guard supports only the Layer 2 port, including both access and trunk. For each untrusted Layer 2 port, there are two levels of IP traffic security filtering, as follows:

- **Source IP address filter:** IP traffic is filtered based on its source IP address. Only IP traffic with a source IP address that matches the IP source binding entry is permitted.

An IP source address filter is changed when a new IP source entry binding is created or deleted on the port. The port PVACL will be recalculated and reapplied in the hardware to reflect the IP source binding change. By default, if the IP filter is enabled without any IP source binding on the port, a default PVACL that denies all IP traffic is installed on the port. Similarly, when the IP filter is disabled, any IP source filter PVACL will be removed from the interface.

- **Source IP and MAC address filter:** IP traffic is filtered based on its source IP address in addition to its MAC address; only IP traffic with source IP and MAC addresses that match the IP source binding entry are permitted.

# IP Source Guard Configuration

This subtopic describes IP Source Guard configuration.

## Catalyst Integrated Security Configuration

```
sw(config)# ip dhcp snooping
sw(config)# ip dhcp snooping vlan 10,20
sw(config)# ip arp inspection vlan 10,20
sw(config)# interface fastethernet 0/1
sw(config-if)# description Access Port
sw(config-if)# switchport mode access
sw(config-if)# switchport access vlan 10
sw(config-if)# switchport port-security maximum 2
sw(config-if)# switchport port-security violation restrict
sw(config-if)# switchport port-security
sw(config-if)# ip dhcp limit rate 50
sw(config-if)# ip verify source port-security
sw(config)# interface fastethernet 0/24
sw(config-if)# description Uplink
sw(config-if)# switchport mode trunk
sw(config-if)# switchport trunk allowed vlan 10,20
sw(config-if)# ip dhcp snooping trust
sw(config-if)# ip arp inspection trust
```

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH10-7-16

The table describes the procedure for enabling IP Source Guard.

DHCP snooping is required for learning valid IP address and MAC address pairs.

## IP Source Guard Configuration Commands

	Command	Purpose
<b>Step 1</b>	Switch(config)# <b>ip dhcp snooping</b>	Enables DHCP snooping, globally.  You can use the <b>no</b> keyword to disable DHCP snooping.
<b>Step 2</b>	Switch(config)# <b>ip dhcp snooping vlan</b> <i>number</i> [ <i>number</i> ]	Enables DHCP snooping on your VLANs.
<b>Step 3</b>	Switch(config)# <b>ip dhcp snooping vlan</b> <i>number</i> [ <i>number</i> ]	Configures the interface as trusted or untrusted.  You can use the <b>no</b> keyword to configure an interface to receive only messages from within the network.
<b>Step 4</b>	Switch(config-if)# <b>ip verify source vlan dhcp-snooping port-security</b>	Enables IP Source Guard, source IP, and source MAC address filtering on the port.
<b>Step 5</b>	Switch(config-if)# <b>switchport port-security limit rate invalid-source-mac</b> <i>N</i>	(Optional) Sets the rate limit for bad packets. This rate limit also applies to the port where DHCP snooping security mode is enabled as filtering the IP and MAC address.
<b>Step 6</b>	Switch(config)# <b>ip source binding ip-addr ip</b> <i>ip</i> <b>vlan</b> <i>number</i> <b>interface</b> <i>interface</i>	Configures a static IP binding on the port.
<b>Step 7</b>	Switch(config)# <b>end</b>	Exits configuration mode.
<b>Note</b>	The static IP source binding can be configured on a switch port only. If you issue the <b>IP source binding</b> VLAN interface command on a Layer 3 port, you will receive this error message: Static IP source binding can be configured on the switch port only.	

# Summary

This topic summarizes the key points that were discussed in this lesson.

## Summary

- DHCP spoofing attacks send unauthorized replies to DHCP queries.
- DHCP snooping is used to counter a DHCP spoofing attack.
- DHCP snooping is easily implemented on a Cisco Catalyst switch.
- ARP spoofing can be used to redirect traffic to an unauthorized device on the network.
- DAI in conjunction with DHCP snooping can be used to counter ARP spoofing attacks.

## Lesson 4

---

# Securing Network Services

---

## Overview

Each switch in the network provides a variety of services that span beyond frame switching. These services range from management web interface and Cisco Discovery Protocol neighbor discovery, to Telnet and preconfigured ports. Several of these services can be enabled by default to facilitate the switch initial integration into the network infrastructure. Some services may need to stay enabled, while other should be disabled when the integration is complete if they are not needed. This lesson describes how to secure Layer 2 devices by protecting physical and virtual ports, disabling unneeded services, forcing the encryption of sessions, and enabling logging at the device level.

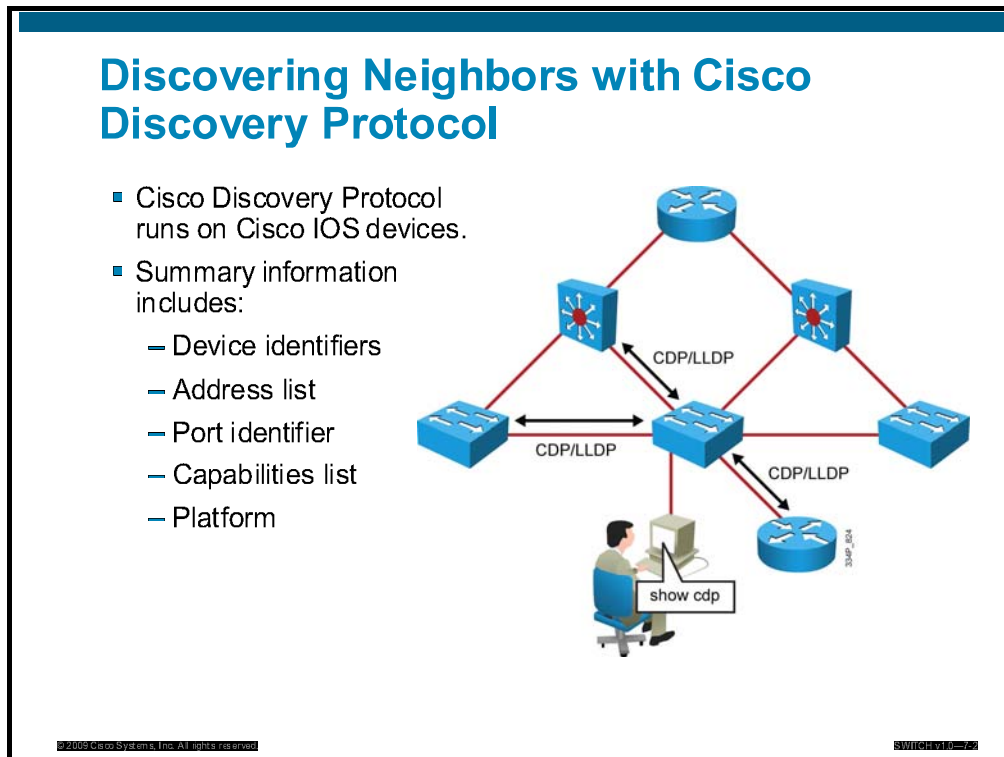
## Objectives

Upon completing this lesson, you will be able to secure network services. This ability includes being able to meet these objectives:

- Identify Cisco Discovery Protocol and LLDP vulnerabilities
- Identify Telnet protocol vulnerabilities
- Configure SSH
- Configure vty ACLs
- Configure Cisco IOS secure HTTP server
- Understand switch security considerations

# Vulnerabilities in Cisco Discovery Protocol

This topic describes the vulnerabilities in Cisco Discovery Protocol.



Cisco Discovery Protocol is a hello-based protocol, and all Cisco devices that run Cisco Discovery Protocol periodically advertise their attributes to their neighbors by using a multicast address. Cisco Discovery Protocol packets advertise a Time to Live (TTL) value in seconds, which indicates the length of time to retain the packet before discarding it. Cisco devices send Cisco Discovery Protocol packets with a TTL value that is not 0 after an interface is enabled. A TTL value of 0 is sent immediately before an interface is idled. Sending a Cisco Discovery Protocol packet with a TTL value of 0 allows a network device to quickly discover a lost neighbor.

By default, all Cisco devices receive Cisco Discovery Protocol packets and cache the information in the packet. The cached information is then available to a network management system (NMS) using the Simple Network Management Protocol (SNMP).

---

**Note** Cisco devices never forward a Cisco Discovery Protocol packet. Cisco Discovery Protocol can be disabled on an interface.

---

If any information changes from the last received packet, the device caches the new information and discards the previous information even if its TTL value has not yet expired.

For security reasons, you should block SNMP access to Cisco Discovery Protocol data (or to any other data) from outside your network and from subnets other than your management station subnet.

---

**Note** Do not run Cisco Discovery Protocol in these places:

1. On links that you do not want to be discovered, such as Internet connections
2. On links that do not go to Cisco devices

---



# Neighbor Discovery Protocols

This subtopic describes neighbor discovery protocols.

## Neighbor Discovery Protocols

- Cisco Discovery Protocol
  - Cisco Layer 2 protocol
  - Has additional capabilities (VLAN or PoE negotiation)
  - Enabled by default
- LLDP
  - Standard-based Layer 2 protocol
  - Disabled by default
- Provides a summary of directly connected switches, routers, and other Cisco devices
- Discovers neighbor devices regardless of which protocol suite they are running

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-010-02

Two protocols are available for neighbor discovery:

- Cisco Discovery Protocol
  - Cisco Layer 2 protocol
  - Enabled by default
- Link Layer Discovery Protocol (LLDP)
  - Standard-based Layer 2 protocol
  - Disabled by default

Neighbor discovery protocols provide a summary of directly connected switches, routers, and other Cisco devices. They also discover neighbor devices regardless of which protocol suite they are running.

# Cisco Discovery Protocol Configuration

This subtopic describes Cisco Discovery Protocol configuration.

## Cisco Discovery Protocol Configuration

```
switch(config)# [no] cdp run
switch(config-if)# [no] cdp enable
switch# show cdp neighbor [detail]
```

```
switch# show cdp neighbor
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,
                  D - Remote, C - CVTA, M - Two-port Mac Relay
Device ID        Local Intrfce    Holdtme    Capability  Platform  Port ID
c2960-8          Fas 0/8          168        S I        WS-C2960- Fas 0/8
```

© 2009 Cisco Systems, Inc. All rights reserved. SWITCH-100-20

Cisco Discovery Protocol is enabled by default:

- **no cdp run** disables Cisco Discovery Protocol globally
- **no cdp enable** disables Cisco Discovery Protocol on an interface

When Cisco Discovery Protocol is enabled, the command **show cdp neighbor** displays a summary of which devices are detected on which ports. You can get more detailed information about all neighbors with the command **show cdp neighbor detail**. If you need detailed information about only one neighbor on one link, you can also use (for example) **show cdp neighbor interface f0/8 detail**.

# LLDP Configuration

This subtopic describes LLDP configuration.

## LLDP Configuration

```
switch(config)# [no] lldp run
switch(config-if)# [no] lldp enable
switch# show lldp neighbor [detail]
```

```
switch# show lldp neighbor
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
  (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other
Device ID          Local Intf      Hold-time  Capability    Port
ID
c2960-8            Fa0/8          120        B             Fa0/8
Total entries displayed: 1
```

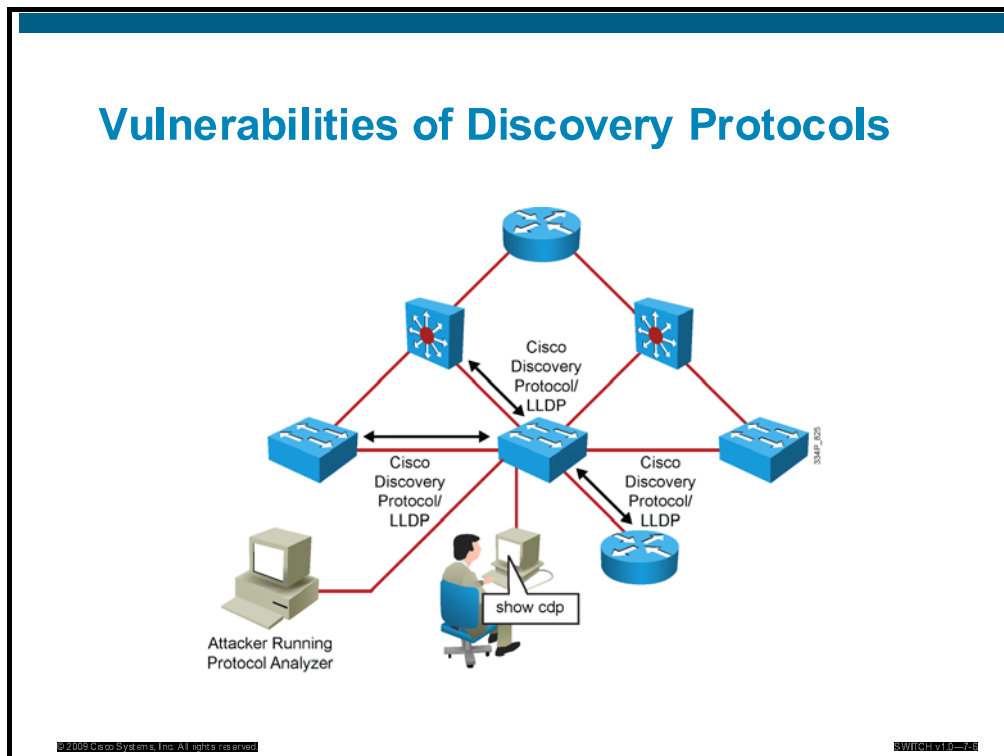
LLDP is disabled by default.

- **lldp run** enables LLDP globally.
- **lldp enable** enables LLDP on an interface.

Just like Cisco Discovery Protocol, when LLDP is enabled the command **show lldp neighbor** displays a summary of which devices are seen on which ports. You can get more detailed information about all neighbors with the command **show lldp neighbor detail**. If you need detailed information about only one neighbor on one link, you can also use (for example) **show lldp neighbor interface f0/8 detail**.

# Cisco Discovery Protocol Vulnerabilities

This subtopic describes Cisco Discovery Protocol vulnerabilities.



Attackers with knowledge of how Cisco Discovery Protocol works could find ways to take advantage of the cleartext Cisco Discovery Protocol packets to gain knowledge of the network. The protocol runs at Layer 2 and allows Cisco devices to identify themselves to other Cisco devices. However, the information sent through Cisco Discovery Protocol is transmitted in clear text and is unauthenticated. Utilizing a packet analyzer, attackers could glean information about the network device from Cisco Discovery Protocol advertisements.

Cisco Discovery Protocol is necessary for management applications and cannot be disabled without impairing some network-management applications. However, you can selectively disable it on interfaces where management is not being performed.

The table describes how Cisco Discovery Protocol can be used maliciously.

## Using Cisco Discovery Protocol Maliciously

Sequence of Events	Description
1.	System administrator uses Cisco Discovery Protocol to view neighbor information.
2.	Attacker uses a packet analyzer to intercept Cisco Discovery Protocol traffic.
3.	Attacker analyzes information in Cisco Discovery Protocol packets to gain knowledge of network address and device information.
4.	Attacker formulates attacks based on known vulnerabilities of network platforms.

# Telnet Vulnerabilities

This topic describes Telnet vulnerabilities.

## Vulnerabilities of the Telnet Protocol

The Telnet connection sends text unencrypted and potentially readable.



Known Telnet vulnerabilities are listed here:

- All usernames, passwords, and data that are sent over the public network in clear text are vulnerable.
- A user with an account on the system could gain elevated privileges.
- A remote attacker could crash the Telnet service, preventing legitimate use of that service.
- A remote attacker could find an enabled guest account that may be present anywhere within the trusted domains of the server.

# About SSH

This topic describes Secure Shell (SSH) Protocol.



SSH is a client and server protocol that is used to log in to another computer over a network, to execute commands in a remote machine, and to move files from one machine to another. It provides strong authentication and secure communications over insecure channels. It is a replacement for rlogin, remote shell (rsh), Remote Copy Protocol (RCP), and rdist, in addition to Telnet.

When you use the SSH login (instead of Telnet), the entire login session, including transmission of password, is encrypted; therefore, it is almost impossible for an outsider to collect passwords.

Although SSH is secured, the many implementations of SSH by various vendors contain vulnerabilities that could allow a remote attacker to execute arbitrary code with the privileges of the SSH process or to cause a denial of service. Most of the SSH vulnerabilities have been addressed in the latest Cisco IOS Software and in the SSH server and client software of other vendors.

---

**Caution** SSH version 1 implementations are vulnerable to various security compromises. Whenever possible, use SSH version 2 instead of SSH version 1.

---

To activate SSH on a vty interface, use the **transport input ssh** command.

# Configuration of SSH

This subtopic describes the configuration of SSH.

## Configuration of SSH

- Configure username and password.
- Configure domain name.
- Generate RSA keys.
  - SSH process is automatically started.
- Allow SSH protocol on vty lines.



```
switch(config)# username xyz password abc123
switch(config)# ip domain-name xyz.com
switch(config)# crypto key generate rsa
switch(config)# ip ssh version 2
switch(config)# line vty 0 15
switch(config-line)# login local
switch(config-line)# transport input ssh
```

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-01-00-7A

To configure SSH on a Cisco IOS switch, follow these steps:


- Step 1**     Configure a user with password.
- Step 2**     Configure hostname and domain name.
- Step 3**     Generate RSA keys.
- Step 4**     Allow SSH transport on the vty lines.

# VTY ACLs

This subtopic describes vty access lists.

## Configuration of vty ACLs

- Create standard or extended IP ACL.
- Configure **access-class** on line vty.



```
sw(config)# access-list 100 permit ip 10.1.1.0 0.0.0.255 any
sw(config)# line vty 0 15
sw(config-line)# access-class 100 in
```

Cisco provides access control lists (ACLs) to permit or deny Telnet access to the vty ports of a switch. Cisco devices vary in the number of vty ports that are available by default. When configuring vty ACLs, ensure that all default ports are removed or have a specific vty ACL applied.

Telnet filtering is normally considered an extended IP ACL function because it is filtering a higher-level protocol. However, because the **access-class** command is used to filter incoming Telnet sessions by source address and to apply filtering to vty lines, you can use standard IP ACL statements to control vty access. The **access-class** command also applies standard IP ACL filtering to vty lines for outgoing Telnet sessions that originate from the switch.

You can apply vty ACLs to any combination of vty lines. The same ACL can be applied globally to all vty lines, or separately to each vty line. The most common practice is to apply the same ACL to all vty lines.

To configure vty ACLs on a Cisco switch, create a standard IP ACL and apply it on the vty interfaces. Rather than apply the ACL to a data interface, apply it to a vty line or range of lines with the **access-class** command.



# HTTP Secure Server Configuration

This subtopic describes the switch Cisco IOS HTTP server configuration.

## Configuration of an HTTP Server

- Configure username and password.
- Configure domain name.
- Generate RSA keys.
- Enable HTTPS (SSL) server.
- Configure HTTP authentication.
- Configure an access list to limit access.



```
sw(config)# access-list 100 permit ip 10.1.9.0 0.0.0.255 any
sw(config)# username xyz password abc123
sw(config)# ip domain-name xyz.com
sw(config)# crypto key generate rsa
sw(config)# no ip http server
sw(config)# ip http secure-server
sw(config)# http access-class 100 in
sw(config)# http authentication local
```

A web interface is available to configure most switches. The main weakness of the web interface is that it is not encrypted, and part of it does not offer any filtering. By default, it can be viewed by any user entering the switch IP address in a web browser address bar.

To protect the web service, you can take several steps:

- Use HTTPS instead of the unprotected HTTP. You can do this by using the command **ip http secure server** instead of **ip http server**.
- Before you can use HTTPS, the web server residing on the switch must send a certificate. You can download this certificate to the switch if the network uses a PKI, or you can generate it on the switch locally with the command **crypto key generate rsa**. You must define a domain name to generate a certificate.
- Access to the web service should be filtered, so that only the administrator computer or subnet is allowed to gain access to the web interface. You do this by creating an access list. In this example, computers on subnet 10.1.9.0/24 have the right to access any IP address of the Layer 3 switch. You can then apply the access list filter to the web service with the command **http access-class**.
- Users accessing the web service may use it to configure the switch. These users should be authenticated. The **http authentication local** command determines that the switch will contain a local list of user credentials. The username xyz password abc123 creates a local user.

# Switch Security Considerations

This topic describes switch security considerations.

## Switch Security Recommendations

### Secure switch access

- Configure system passwords.
- Authenticate admin access via TACACS+ server.
- Configure encrypted or hashed passwords.
- Secure physical access to the console.
- Secure Telnet access with ACL.
- Use SSH when possible.
- Use HTTPS (SSL) when possible.
- Configure system-warning banners.
- Use syslog to log system messages.
- Disable unused services.

Network security vulnerabilities include loss of privacy, data theft, impersonation, and loss of integrity. It is important to take basic security measures on every network to mitigate the adverse effects of user negligence or acts of malicious intent.

Whenever new equipment is added to the network, follow these general security guidelines:

- Step 1** Consider or establish organizational security policies.
- Step 2** Secure the switch devices.
- Step 3** Secure the switch protocols.
- Step 4** Mitigate compromises launched through a switch.

# Organizational Security Policies

You should consider the policies of an organization when determining what level of security and what type of security should be implemented. You must balance the goal of reasonable network security against the administrative overhead that is clearly associated with extremely restrictive security measures.

A well-established security policy has these characteristics:

- Provides a process for auditing existing network security
- Provides a general security framework for implementing network security
- Defines disallowed behaviors toward electronic data
- Determines which tools and procedures are needed for the organization
- Communicates consensus among a group of key decision makers and defines responsibilities of users and administrators
- Defines a process for handling network security incidents
- Enables an enterprise-wide, all-site security implementation and enforcement plan

## Secure Switch Devices

Follow these best practices for secure switch access:

- **Set system passwords:** Use the **enable secret** command to set the password that grants enabled access to the Cisco IOS system. Because the **enable secret** command simply implements a Message Digest 5 (MD5) hash on the configured password, that password still remains vulnerable to dictionary attacks. Therefore, apply standard practices in selecting a feasible password.  
  
Try to pick passwords that contain both letters and numbers in addition to special characters, for example, “\$pecial\$” instead of “specials,” where the “s” has been replaced with “\$,” and the “l” has been replaced with “1” (one).
- **Secure access to the console:** Console access requires a minimum level of security both physically and logically. An individual who gains console access to a system will be able to recover or reset the system-enable password, thus allowing that person to bypass all other security implemented on that system. Consequently, it is imperative to secure access to the console.
- **Secure access to vty lines:** These are the minimum recommended steps for securing Telnet access.
  - Apply the basic ACL for in-band access to all vty lines.
  - Configure a line password for all configured vty lines.
  - If the installed Cisco IOS image permits, use SSH instead of Telnet to access the device remotely.
- **Use SSH:** The SSH protocol and application provide a secure remote connection to a router. Two versions of SSH are available: SSH version 1 and SSH version 2. SSH version 1 is implemented in Cisco IOS Software. It encrypts all traffic, including passwords, between a remote console and a network router across a Telnet session. Because SSH sends no traffic in clear text, network administrators can conduct remote access sessions that casual observers will not be able to view. The SSH server in Cisco IOS Software will work with publicly and commercially available SSH clients.

- **Configure system-warning banners:** For both legal and administrative purposes, configuring a system-warning banner to display before login is a convenient and effective way of reinforcing security and general usage policies. By clearly stating the ownership, usage, access, and protection policies before a login, you provide more solid backing for potential future prosecution.
- **Disable unneeded services:** By default, Cisco devices implement multiple TCP and User Datagram Protocol (UDP) servers to facilitate management and integration into existing environments. For most installations, these services are typically not required, and disabling them can greatly reduce overall security exposure. These commands will disable the services not typically used:
  - no service tcp-small-servers**
  - no service udp-small-servers**
  - no service finger**
  - no service config**
- **Disable the integrated HTTP daemon if not in use:** Although Cisco IOS Software provides an integrated HTTP server for management, it is highly recommended that it be disabled to minimize overall exposure. If HTTP access to the switch is required, use basic ACLs to permit access from only trusted subnets.
- **Configure basic logging:** To assist and simplify both problem troubleshooting and security investigations, monitor the switch subsystem information that is received from the logging facility. View the output in the on-system logging buffer memory. To render the on-system logging useful, increase the default buffer size.

# Switch Security Recommendations

This subtopic describes switch security recommendations.

## Switch Security Recommendations (Cont.)

### Secure switch protocols

- Trim Cisco Discovery Protocol and LLDP and use only as needed.
- Secure spanning tree.

### Mitigate compromises through a switch

- Take precautions for trunk links.
- Minimize physical port access.
- Establish standard access port configuration for both unused and used ports.
- Shut down unused ports.

Follow these best practices for switch security:

- **Cisco Discovery Protocol or LLDP:** Cisco Discovery Protocol or LLDP does not reveal security-specific information, but it is possible for an attacker to exploit this information in a reconnaissance attack, whereby an attacker learns device and IP address information for the purpose of launching other types of attacks. Here are two practical guidelines for Cisco Discovery Protocol and LLDP.
  - If Cisco Discovery Protocol or LLDP is not required, or if the device is located in an unsecure environment, disable Cisco Discovery Protocol or LLDP globally on the device.
  - If Cisco Discovery Protocol or LLDP is required, disable it on a per-interface basis on ports connected to untrusted networks. Because Cisco Discovery Protocol or LLDP is a link-level protocol, it is not transient across a network (unless a Layer 2 tunneling mechanism is in place). Limit it to run between trusted devices only, and disable it everywhere else. However, Cisco Discovery Protocol or LLDP is required on any access port when you are attaching a Cisco phone to establish a trust relationship.
- **Secure the spanning-tree topology:** It is important to protect the Spanning Tree Protocol (STP) process of the switches that compose the infrastructure. Inadvertent or malicious introduction of STP bridge protocol data units (BPDUs) could potentially overwhelm a device or pose a denial of service (DoS) attack. The first step in stabilizing a spanning-tree installation is to positively identify the intended root bridge in the design and to hard-set the STP bridge priority of that bridge to an acceptable root value. Do the same for the designated backup root bridge. These actions will protect against inadvertent shifts in STP caused by an uncontrolled introduction of a new switch.

On some platforms, the BPDUGuard feature may be available. If so, enable it on access ports in conjunction with the PortFast feature to protect the network from unwanted BPDU traffic injection. Upon receipt of a BPDU, the feature will automatically disable the port.

## Mitigating Compromises Launched Through a Switch

Follow these best practices to mitigate compromises through a switch:

- **Proactively configure unused router and switch ports.**
  - Use the **shut** command on all unused ports and interfaces.
  - Place all unused ports in a “parking-lot” VLAN that is used specifically to group unused ports until they are proactively placed into service.
  - Configure all unused ports as access ports, disallowing automatic trunk negotiation.
- **Considerations for trunk links:** By default, Cisco Catalyst switches running Cisco IOS Software are configured to automatically negotiate trunking capabilities. This situation poses a serious hazard to the infrastructure because an unsecured third-party device can be introduced to the network as a valid infrastructure component. Potential attacks include interception of traffic, redirection of traffic, DoS, and more. To avoid this risk, disable automatic negotiation of trunking and manually enable it on links that will require it. Ensure that trunks use a native VLAN that is dedicated exclusively to trunk links.
- **Physical device access:** Physical access to the switch should be closely monitored to avoid rogue device placement in wiring closets with direct access to switch ports.
- **Access port–based security:** Specific measures should be taken on every access port of any switch placed into service. Ensure that a policy is in place outlining the configuration of unused switch ports in addition to those that are in use.

For ports that are enabled for end-device access, there is a macro called **switchport host**, which, when executed on a specific switch port, does the following: sets the switch port mode to access, enables spanning-tree PortFast, and disables channel grouping.

---

<b>Note</b>	The <b>switchport host</b> macro disables EtherChannel, disables trunking, and enables STP PortFast.
-------------	--

---

The command is a macro that executes several configuration commands. There is no command such as **no switchport host** to revoke the effect of the **switchport host** command. To return an interface to its default configuration, use the **default interface** *interface-id* global configuration command. This command returns all interface configurations to the default.

# Summary

This topic summarizes the key points that were discussed in this lesson.

## Summary

- Cisco Discovery Protocol / LLDP packets can expose some network information.
- Authentication information and data carried in Telnet sessions is vulnerable.
- SSH provides a more secure option for Telnet.
- vty ACLs should be used to limit Telnet access to switch devices.
- Web service should be secured by using HTTPS and limiting who should access the web server and from where.
- Sound security measures and trimming of unused applications are recommended.





## Lesson 5

---

# Lab 7-1 Debrief

---

## Overview

In this lab, you have secured different aspects in your Layer 2 network. You have implemented solutions to control the Layer 2 access to switches, to manage the allowed traffic for the clients, thus protecting the network equipment from unauthorized access. You have secured the work of the Spanning Tree Protocol (STP) and prepared your network to defend against DHCP-based and Address Resolution Protocol (ARP) spoofing attacks. You have analyzed and designed the possible solutions to meet the business and technical requirements. You have created implementation plans for the different security solutions. When you were ready with the implementation steps, you connected to the remote lab and configured your switches to match the requirements. You then verified that your implementations respected the specific needs.

During the lab debrief, the instructor will lead a group discussion in which you can present your solution. You will get an opportunity to verify your solution against a number of checkpoints that are provided by the instructor, and to compare your solution to those of other students. The instructor will discuss the solutions and their benefits and drawbacks.

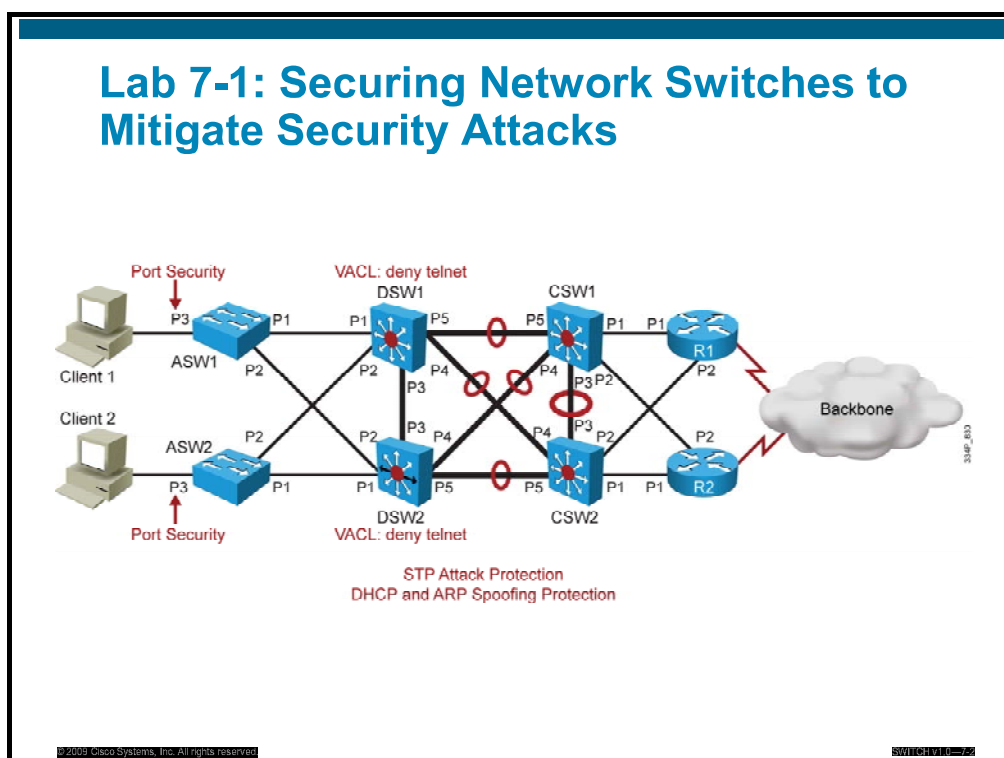
## Objectives

Upon completing this lesson, you will be able to produce a design plan and implementation plan for different security solutions in a Layer 2 network based on given business and technical requirements. You will then implement the plan and perform verifications while identifying checkpoints along the way. This ability includes being able to meet these objectives:

- Review and verify your solution, as well as your findings and action log, against a set of checkpoints provided by the instructor
- Consolidate the lessons that you learned during the review discussions into a set of best-practice methods and commands to aid you in future deployment procedures
- Perform a baseline assessment of network switch security settings
- Identify possible threats, points of attack, and vulnerability points in the network
- Write an implementation plan to implement security measures on network switches
- Write a plan to test and verify security threat mitigation measures for VLANs
- Configure port security, VACL, and other switch security features
- Verify correct implementation of security measures

# Review and Verification

This topic describes the requirements that were listed in Lab 7-1, asks how you can verify that you have identified the solution matching those needs, and gives you an example of a possible solution.



This lab consists of several tasks and the configuration changes that affect four of the switches—the two access switches and the two distribution switches. For the different solutions, you have to consider which devices will be affected and in what way. It is important when implementing security to have a detailed preliminary design and plan, because they may affect the functioning of core network services. Verification is therefore an extremely important step.

## Design and Implementation Plan

Which items should be configured, and in which order?

- Which switches will be configured?
- Which ports will be configured?
- Which port security options will be used?
- Which VLANs will be protected?
- Which ports will be configured as trusted?
- Which DHCP options will be used?
- Status verification after the implementation

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCHING-26

A successful implementation plan allows you to configure the devices without duplication, and it reduces the risk of mistakes. Thus, you will need to make only a minimum set of configuration changes, and you will reduce the time necessary for troubleshooting after the implementation. In other words, an implementation plan is efficient when you do not need to alter your previous configuration to implement new items. You should proceed in a logical order.



# Module Summary

This topic summarizes the key points that were discussed in this module.

## Module Summary

- Key switch security issues should be identified on a switched network and proper measures taken to mitigate known attacks.
- VLAN trunk links should be secured to defend against VLAN hopping attacks.
- Port security, DHCP snooping, and DAI are used to protect the network against spoofing attacks.
- When placed into service, switches should be configured according to recommendations to secure the switch device and its protocols from attacks that can be launched through a switch.



# Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which feature that is supported on Cisco Catalyst switches restricts a switch port to a specific set or number of MAC addresses? (Source: Understanding Switch Security Issues)
- A) port security
  - B) DHCP snooping
  - C) PVLAN
  - D) VACL
- Q2) At which layer should port security be implemented? (Source: Understanding Switch Security Issues)
- A) access
  - B) distribution
  - C) core
  - D) all three layers
- Q3) At which layer should packet manipulation usually be avoided? (Source: Understanding Switch Security Issues)
- A) access
  - B) distribution
  - C) core
  - D) none of the above
- Q4) What best describes a MAC flooding attack? (Source: Understanding Switch Security Issues)
- A) A device sends all its frames to the FF:FF:FF:FF:FF:FF address.
  - B) A switch CAM becomes overloaded with too many MAC addresses.
  - C) A device sends frames to too many destination MAC addresses.
  - D) A device sends frames at a rate faster than the switch link speed.
- Q5) What is the aim of the “sticky” option when used with port security? (Source: Understanding Switch Security Issues)
- A) A learned MAC address must stick to one single port.
  - B) A dynamically learned MAC address is considered like a statically learned MAC address.
  - C) For a given MAC address with the sticky option, the port security feature applies to whichever port the MAC address connects to.
  - D) A router on a stick can bypass the port security feature and use one MAC address per subinterface.
- Q6) Which command enables 802.1X globally on a switch? (Source: Understanding Switch Security Issues)
- A) Switch (config)# **dot1x enable**
  - B) Switch (config)# **switchport dot1x enable**
  - C) Switch (config)# **aaa dot1x enable**
  - D) Switch (config)# **dot1x system-auth-control**

- Q7) What is the role of the switch in a AAA architecture? (Source: Understanding Switch Security Issues)
- A) authentication server
  - B) supplicant
  - C) authenticator
  - D) RADIUS entry point
- Q8) What is one best practice for mitigating VLAN hopping? (Source: Protecting Against VLAN Attacks)
- A) Configure all unused ports as trunks.
  - B) Shut down all unused ports.
  - C) Set trunks to “negotiate” and not “on.”
  - D) Set the interface speed to 10 Mb/s.
- Q9) Which command should you use to begin VACL configuration? (Source: Protecting Against VLAN Attacks)
- A) **vlan access-list 100**
  - B) **vlan access filter**
  - C) **vlan map**
  - D) **vlan access-map**
- Q10) What is the purpose of the VLAN hopping with double tagging attack? (Source: Protecting Against VLAN Attacks)
- A) to attack a computer in a VLAN different from the attacker VLAN
  - B) to overload the switch with multiple VLAN tags to force it to flood the frame
  - C) to attack a computer on the native VLAN
  - D) to attack a switch or router on a trunk while keeping the frame tagged
- Q11) What would an attacker do to perform a switch spoofing attack? (Source: Protecting Against VLAN Attacks)
- A) Send random ARP requests to the connected stations.
  - B) Flush the switch CAM to force it to flood.
  - C) Source all its frames with the address FF:FF:FF:FF:FF:FF.
  - D) Mark its frames with ISL or 802.1Q.
- Q12) Which are three ways to protect against spoofing attacks? (Choose three.) (Source: Protecting Against Spoofing Attacks)
- A) DHCP snooping
  - B) port security
  - C) Dynamic ARP Inspection
  - D) IP HTTP secure server
- Q13) Which is the purpose of DHCP snooping? (Source: Protecting Against Spoofing Attacks)
- A) to protect against rogue DHCP clients
  - B) to protect against DHCP address reuse attacks
  - C) to protect against rogue DHCP servers
  - D) to protect against DHCP options malicious attacks



- Q14) With DHCP snooping, which port is “trusted”? (Source: Protecting Against Spoofing Attacks)
- A) The port to the known DHCP server is always trusted.
  - B) The port to the DHCP client is always trusted.
  - C) No ports are trusted when DHCP snooping is enabled.
  - D) Any port (to client and to server) can become trusted as soon as a DHCP transaction is secured.
- Q15) What is the purpose of the **ip arp inspection** command? (Source: Protecting Against Spoofing Attacks)
- A) to complement DHCP snooping by verifying the ARP table
  - B) to ensure that only one MAC address is associated with any given access port
  - C) to start the secure ARP proxy service on a multilayer switch
  - D) to statically define the IP address to MAC address pairs for key devices
- Q16) Which step is required before SSH can be enabled on a switch? (Source: Securing Network Services)
- A) Telnet must be disabled.
  - B) A domain name must be defined.
  - C) An upgraded Cisco IOS version must be loaded.
  - D) An SSH access list must be defined.

## Module Self-Check Answer Key

Q1)	A
Q2)	A
Q3)	C
Q4)	B
Q5)	B
Q6)	D
Q7)	C
Q8)	B
Q9)	D
Q10)	A
Q11)	D
Q12)	A, B, C
Q13)	C
Q14)	D
Q15)	A
Q16)	B

# Accommodating Voice and Video in Campus Networks

---

## Overview

When you are migrating to a VoIP network, all network requirements, including power and capacity planning, must be examined. In addition, you should implement congestion avoidance techniques. The same congestion considerations apply when deploying video applications, with an additional consideration related to bandwidth consumption. This module will highlight the basic issues and define the initial steps to take to ensure that a VoIP or video implementation works correctly.

## Module Objectives

Upon completing this module, you will be able to accommodate voice and video in campus networks. This ability includes being able to meet these objectives:

- Plan for support of voice in a campus network
- Implement and verify VoIP in a campus network
- Work with voice specialists to accommodate voice and video on campus switches



# Planning for Support of Voice in a Campus Network

---

## Overview

IP telephony services are often provided over the campus infrastructure. To have data and voice application traffic harmoniously coexist, mechanisms must be set in place to differentiate traffic and to offer priority processing to delay sensitive voice traffic. Quality of service (QoS) policies mark and qualify traffic as it traverses the campus switch blocks. Specific VLANs keep voice traffic separate from other data to ensure that it is carried through the network with special handling and with minimal delay. Specific design and implementation considerations should be made at all campus switches supporting VoIP.

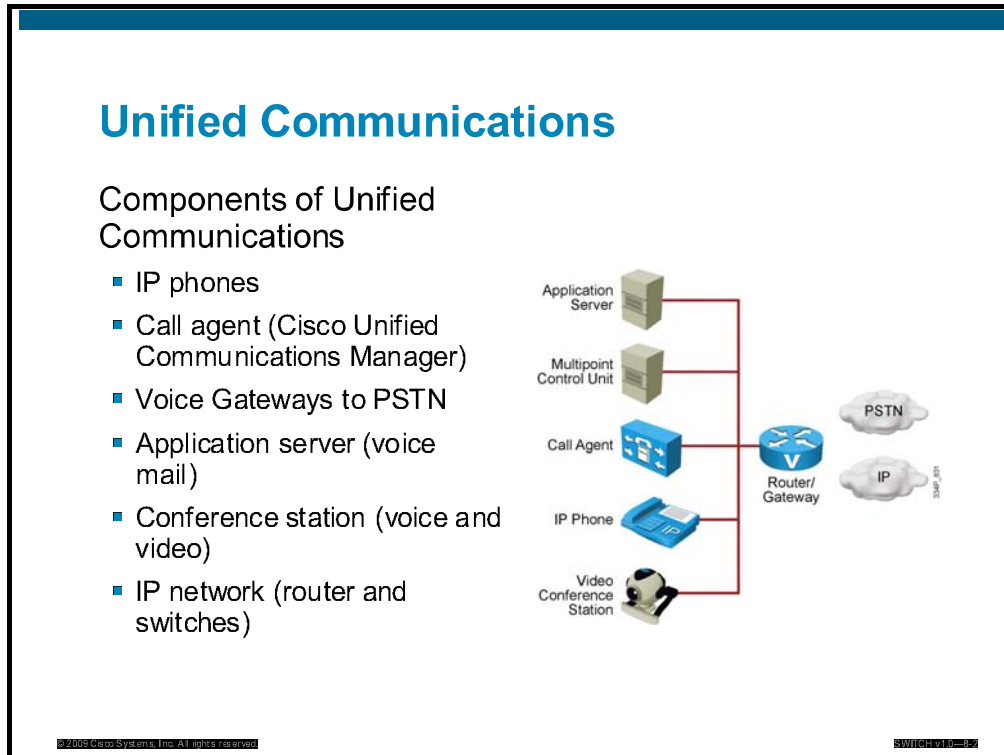
## Objectives

Upon completing this lesson, you will be able to describe the best practices for implementing voice in a campus network. This ability includes being able to meet these objectives:

- Describe the components of a VoIP network and the components of IP telephony
- Compare the uniform bandwidth consumption of voice traffic to the intermittent bandwidth consumption of data traffic
- Compare video bandwidth consumption to voice and data bandwidth consumption based on video application types
- Identify a solution for latency, jitter, bandwidth, packet loss, reliability, and security for voice and video traffic integration into a data network

# Unified Communications

This topic describes unified communications.



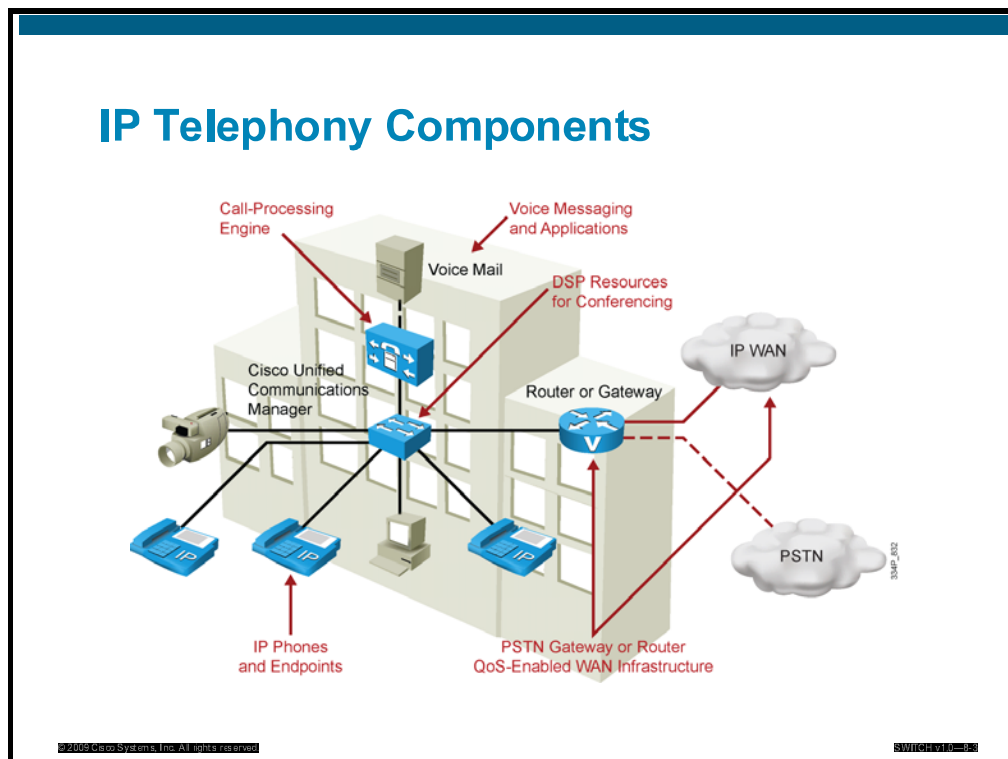
These are the basic components of a VoIP network:

- **IP phones:** Provide IP voice to the desktop.
- **Gatekeeper:** Provides connection admission control (CAC), bandwidth control and management, and address translation.
- **Gateway:** Provides translation between VoIP and non-VoIP networks, such as the public switched telephone network (PSTN). It also provides physical access for local analog and digital voice devices, such as telephones, fax machines, key sets, and PBXs.
- **Multipoint control unit:** Provides real-time connectivity for participants in multiple locations to attend the same videoconference or meeting.
- **Call agent:** Provides call control for IP phones, CAC, bandwidth control and management, and address translation.
- **Application servers:** Provide services such as voice mail, unified messaging, and Cisco Unified Communications Manager Attendant Console.
- **Videoconference station:** Provides access for end-user participation in videoconferencing. The videoconference station contains a video capture device for video input and a microphone for audio input. The user can view video streams and hear the audio that originates at a remote user station.

Other components, such as software voice applications, interactive voice response (IVR) systems, and softphones, provide additional services to meet the needs of enterprise sites.

# IP Telephony Components

This subtopic describes IP telephony components

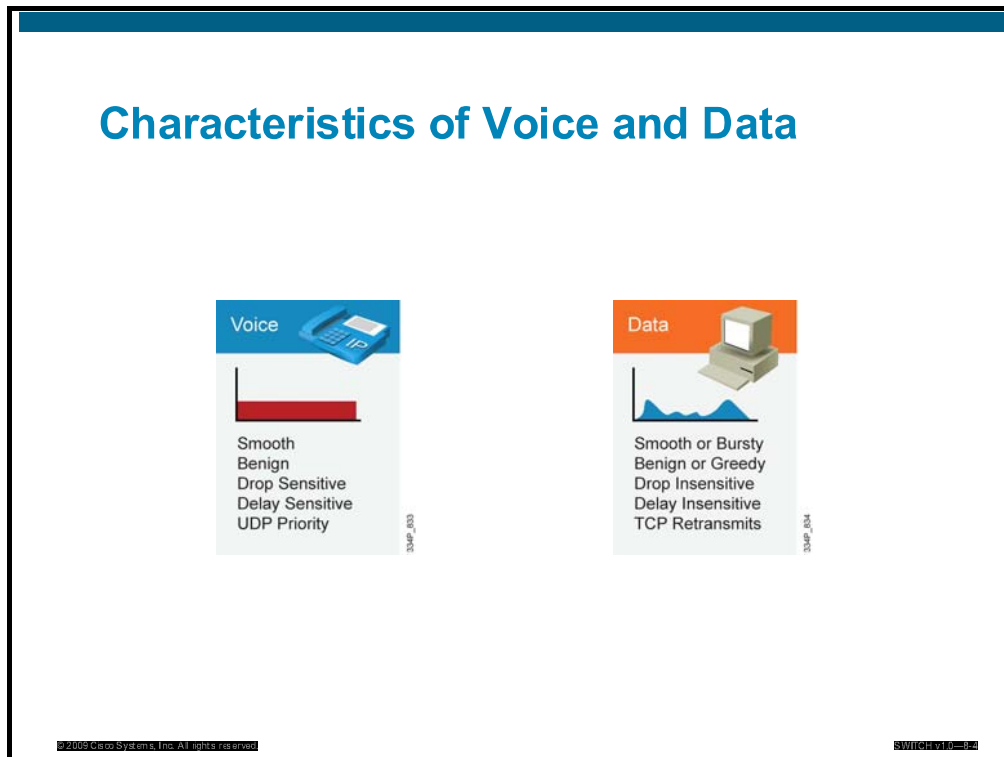


There are four main voice-specific components of the IP telephony network:

- **IP phones:** IP phones support calls in an IP telephony network. IP phones perform voice-to-IP (and vice versa) coding and compression using special hardware. IP phones offer services such as user directory lookups and Internet access for stock quotes. The telephones are active network devices and require power for their operation. Typically, a network connection or an external power supply provides the power.
- **Switches with inline power:** Switches with inline power enable a modular wiring-closet infrastructure to provide centralized power for Cisco IP telephony networks. These switches are similar to traditional switches, with an option to provide power to the LAN ports where IP phones are connected. In addition, these switches perform some basic QoS mechanisms, such as packet classification, which is a baseline for prioritizing voice through the network.
- **Call-processing manager:** Cisco Unified Communications Manager provides central call control and configuration management for IP phones. Cisco Unified Communications Manager provides the core functionality to initialize IP telephony devices and perform call setup and routing of calls throughout the network. Cisco Unified Communications Manager supports clustering, which provides a distributed, scalable, and highly available IP telephony model.
- **Voice gateway:** Voice gateways, also called voice-enabled routers or switches, provide voice services such as voice-to-IP coding and compression, PSTN access, IP packet routing, backup call processing, and voice services. Backup call processing allows voice gateways to take over call processing in case the primary call-processing manager fails. Typically, voice gateways support a subset of the call-processing functionality that is supported by Cisco Unified Communications Manager.

# Characteristics of Voice and Data

This topic describes voice and data traffic characteristics in the campus network.



Voice traffic has extremely stringent QoS requirements. Voice traffic generally generates a smooth demand on bandwidth and has minimal impact on other traffic, as long as voice traffic is managed.

Although voice packets are typically small (60 to 120 bytes), they cannot tolerate delay or drops. The result of delays and drops is poor, and often unacceptable, voice quality. Because drops cannot be tolerated, User Datagram Protocol (UDP) is used to package voice packets; TCP retransmit capabilities have no value.

For voice quality, the delay should be no more than 150 ms (one-way requirement) and less than 1 percent packet loss.

A typical voice call requires 17 kb/s to 106 kb/s of guaranteed priority bandwidth, plus an additional 150 b/s per call for voice-control traffic. Multiplying these bandwidth requirements by the maximum number of calls expected during the busiest time period indicates the overall bandwidth required for voice traffic.

The QoS requirements for data traffic vary greatly.

Different applications (for example, a human resources application versus an ATM application) may make greatly different demands on the network. Even different versions of the same application may have varying network traffic characteristics.

Data traffic can demonstrate either a smooth or bursty characteristic, depending upon the application, but it differs from voice and video in terms of delay and drop sensitivity. Almost all data applications can tolerate some delay and generally can tolerate high drop rates.



Because data traffic can tolerate drops, the retransmit capabilities of TCP become important and, as a result, many data applications use TCP.

In enterprise networks, important (business-critical) applications are usually easy to identify. Most applications can be identified based on TCP or UDP port numbers. Some applications use dynamic port numbers that, to some extent, make classifications more difficult. Cisco IOS software supports Network-Based Application Recognition (NBAR), which can be used to recognize dynamic port applications.

# Video Applications

This topic describes video applications.

## Video Applications

Video collaboration

Cisco TelePresence

IP surveillance

Digital media systems

Characteristics of video applications

- Most are interactive
- Video contains voice as well
- Requires a connection with little delay to prevent jitter (same as with voice)
- Requires a high-bandwidth connection depending on the video resolution
- Peer-to-peer traffic
- Video endpoints are connected to the access layer

For networking professionals, “Video” is often seen as just a Layer 7 application. Many different types of applications may be grouped under this generic term, and each type has its own set of usage and technical specifications. Video applications are usually seen as bandwidth intensive. Bandwidth-intensive applications are so named because they send a large number of frames onto the network. The content of the frame plays an important role in the overall bandwidth consumption. The content of the frame depends on the type of application for which video is used.


Some video applications are real time, such as video conferencing or Cisco TelePresence. The main concern for this type of video is the real-time requirement. Frames have to transit without delay, and QoS is a critical element. Voice may be contained in the global flow or sent as a distinct flow.

Some other video applications are not as sensitive to real-time issues. These video applications are usually one-way streams, where a client station is playing a video stream that is sent from a source. The client station may be able to buffer part of the video flow. The main concern is often quality, and quality depends on several factors, such as resolution, number of frames per second, type of codec, and so on.

Real-time video applications are often peer-to-peer. The consequence for the networking professional is that endpoints are connected at the access layer, in the same manner as a voice or data endpoint.

# Voice and Video Traffic

This subtopic describes voice and video traffic characteristics.



A video flow has different characteristics from a voice flow. Voice traffic is usually not very bandwidth intensive, but requires steady bandwidth. Very commonly, 50 voice packets need to be sent per second, each packet representing a few hundred bytes of information.

Because of the nature of the algorithm used to encode the flow, video traffic has a very different pattern. Traffic is often bursty, as each image or group of images needs several packets to be transmitted. Depending on the changes from one image to the next, there might be short time intervals without any network activity, or sudden bursts when the whole image needs to be changed.

Overall, video streams typically consume a lot of bandwidth and are often bursty, although they usually do not consume all the available bandwidth.

# Requirements for Voice, Data, and Video Traffic

This topic describes the requirements for voice, data, and video traffic.

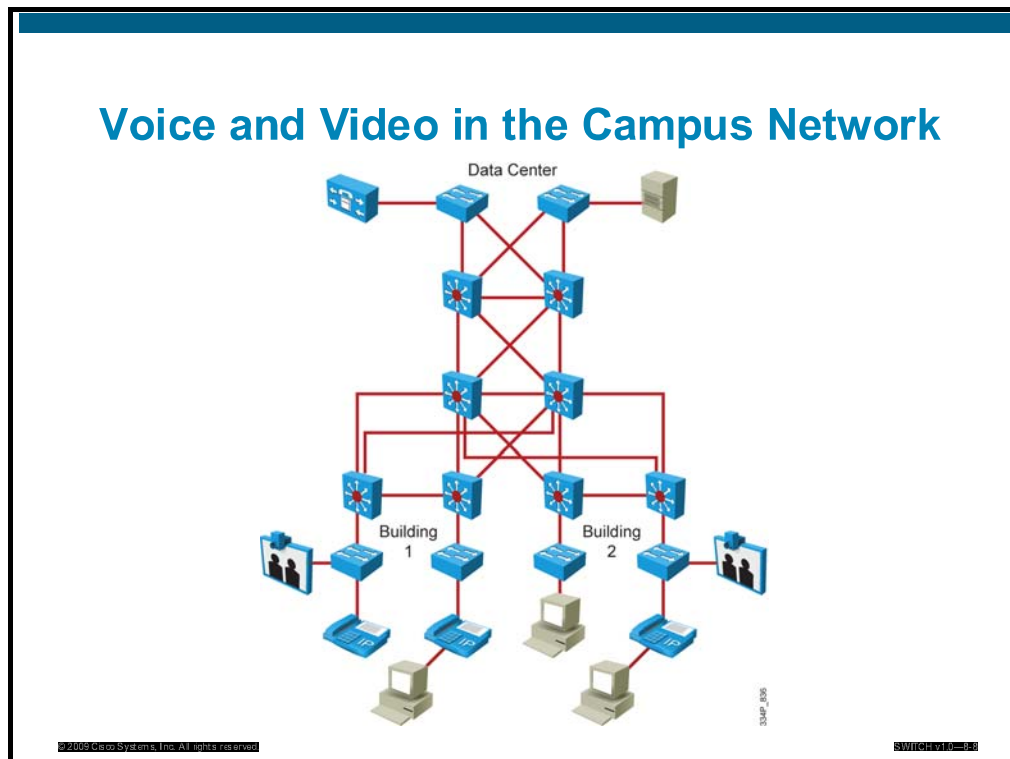
Requirements for Voice, Data, and Video Traffic			
Requirements	Voice traffic	Data traffic	Video traffic
Bandwidth	Low	High	High
Delay	Short	Mostly no problem	Constant or short
Jitter	Little	Mostly no problem	Little
Packet loss	Little	Mostly no problem	Little
Availability	High	High	Medium
Power	Yes	None	None
Security	Medium	High	Low or medium
Provisioning Management	High (DHCP)	Medium (DHCP)	Medium

The illustration shows the bandwidth, delay, jitter, packet loss, availability, power, security, and management requirements for voice, data and video traffic. The tolerance to delay and jitter of video traffic depends on the type of video flow: real-time video, such as video conferencing with embedded voice, has the same delay and jitter constraints as voice traffic. One-way video streams are less sensitive to delay and jitter issues. Most video-playing devices can buffer 5 seconds or more of video flow.

Video traffic also has little tolerance to packet loss. If too many frames are lost, the image does not refresh properly and the user experience degrades. Here again, the quality depends on the type of application used. Tolerance to packet loss is often higher for slow-rate and low-definition real-time interactive video flows than for one-way video streams. The issue is also related to the video codec. Some codecs can compensate for lost packets by interpolating the missing values, while others do not refresh the corresponding section of the image.

# Voice and Video in the Campus Network

This subtopic describes voice and video in the campus network.



Traffic flow for video applications, in the case of peer-to-peer interactive video, is very close to the voice and video flow. In the above illustration, two Cisco TelePresence stations are communicating. The flow goes from one station to the access switch, then the distribution switch in the building, then to the core switch before reaching the distribution Layer of the second building, then the access switch and the second station. This pattern is very close to the voice flow between two phones. Like a voice call, Cisco TelePresence stations may rely on a central server where information about the session may be obtained.

Data traffic does not often transit from one station to the other. Data clients usually communicate with data centers to upload or download data. Video streaming applications behave in the same manner as data applications: video streaming applications retrieve information from data centers and have little if any peer-to-peer interaction.

The result is that video traffic can contend for resources on uplinks between access, distribution, core, and server farm.

# Summary

This topic summarizes the key points that were discussed in this lesson.

## Summary

- IP phones, call agent, and application servers are used for unified communications.
- Voice sends a constant data stream while data traffic sends changing amounts of data (bursts).
- Video applications increase bandwidth requirements in the network.
- Video collaboration, digital media systems, and IP surveillance are emerging video applications.

## Lesson 2

---

# Integrating and Verifying VoIP in a Campus Infrastructure

---

## Overview

Implementing voice in a campus network requires close cooperation between the voice specialist and the network professional. Voice integration needs to be planned thoroughly to integrate seamlessly into the existing network. Once you reach the implementation phase, you must configure the access switches for VoIP support. This lesson describes the step involved in the configuration required to support VoIP on campus switches. It also describes a simple test plan that can be used to ensure that a phone properly communicates with and is integrated into the network infrastructure.

## Objectives

Upon completing this lesson, you will be able to integrate VoIP in a campus infrastructure and verify its integration. This ability includes being able to meet these objectives:

- Plan for VoIP requirements
- Describe voice VLANs
- Configure and verify voice VLANs
- Plan PoE requirements and configure PoE
- Provide additional services required by VoIP devices
- Create a test plan for VoIP integration

# Planning for VoIP Requirements

This topic describes voice requirements

Meeting the Requirements	
Requirements	Solutions
Provisioning Management Security	Voice VLAN DHCP
Bandwidth Delay / Jitter Packet loss	QoS
Power	Power over Ethernet UPS
Availability	High availability Redundancy

Voice traffic has specific requirements. The voice flow needs a steady pace and constant bandwidth to operate free of delay and jitter issues. Therefore, the first requirement for voice traffic integration is to implement quality of service (QoS). QoS addresses bandwidth, delay, jitter, and packet loss requirements. QoS does not completely solve bandwidth issues because QoS can only prioritize packets. When congestion is severe, QoS may not be sufficient anymore to guarantee voice traffic quality.

QoS does not solve single-point-of-failure issues. To address this concern, you must plan for redundancy.

Voice devices get their power from AC adapters, uninterruptible power supplies (UPSs), or from Power over Ethernet (PoE)-enabled switches. Voice VLAN and DHCP solve the provisioning, management, and security requirements.



# Integrating VoIP in the Campus

This subtopic describes the steps necessary to implement VoIP.

## Voice Implementation Steps

### Voice VLAN

- Configure voice VLANs on the access ports
- Add voice VLANs to the trunks
- Configure subnets for voice VLANs
- Configure DHCP for voice subnets (Cisco Unified Communications Manager IP address)
- Enable routing for voice VLANs

### PoE

- Provide PoE on access switches
- Implement UPS

### QoS

- Implement auto QoS for VoIP

### High Availability

- Optimize the topology, routing, and redundancy

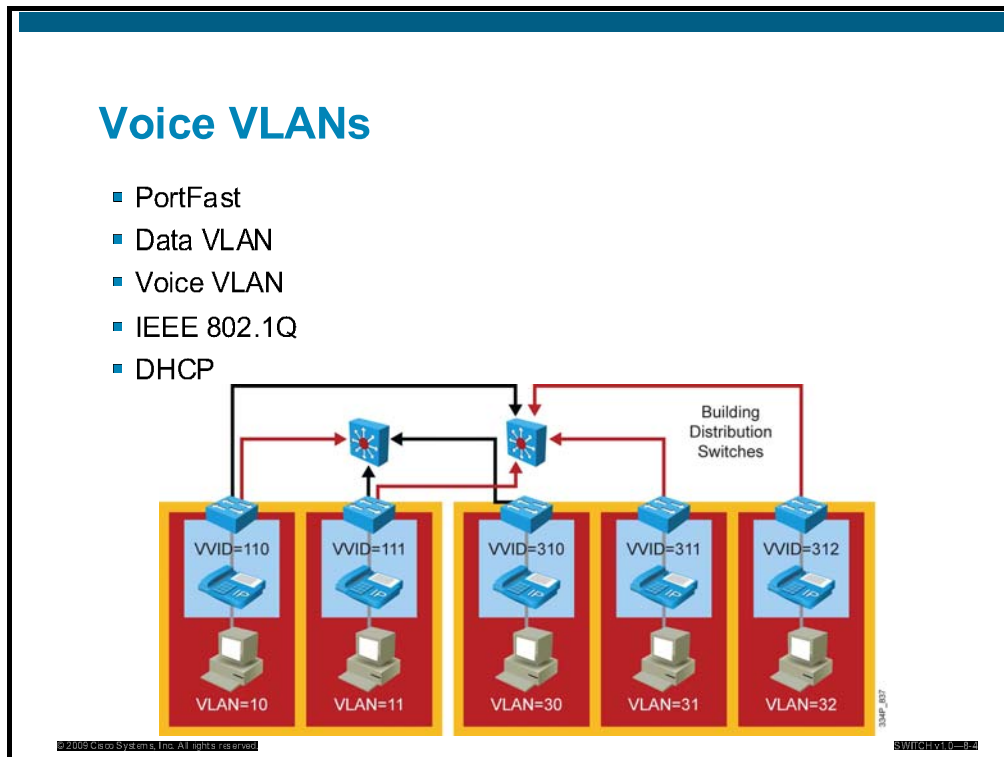
Because voice traffic is usually separated from data traffic and sent to a different VLAN, voice implementation is often very close to initial data VLAN deployment. Voice VLANs must be configured on the access switch ports. Voice VLANs must be added to the switch database and allowed to the trunks. At Layer 3, voice switch virtual interfaces (SVIs) might have to be configured, along with Layer 3 services such as DHCP and routing. A particularity of voice devices is that they need to register to the voice communication server, which in the Cisco solution is Cisco Unified Communications Manager or Cisco Unified Communications Manager Express. The address of this server is usually provided via an option (option 150) in the DHCP scope.

Once this framework is implemented, voice-specific configuration can be added. At the access layer, PoE is configured if needed. Then, from the access ports and throughout the network, QoS is configured to encompass the voice flow. Cisco AutoQoS can be used to configure QoS policies for VoIP on the access ports and the trunk links.

Because delay is a permanent concern, ensure that high availability is configured throughout the network, and that the failover timers are set to a small value to minimize the impact of a lost device on ongoing voice conversations.

# Voice VLANs

This topic describes voice VLANs.



Some Cisco Catalyst switches offer a unique feature called “auxiliary VLAN.” The auxiliary VLAN feature allows you to overlay a voice topology onto a data network. You can segment phones into separate logical networks, even though the data and voice infrastructure are physically the same.

The auxiliary VLAN feature places the phones into their own VLANs without any end-user intervention. These VLAN assignments can be seamlessly maintained, even if the phone is moved to a new location.

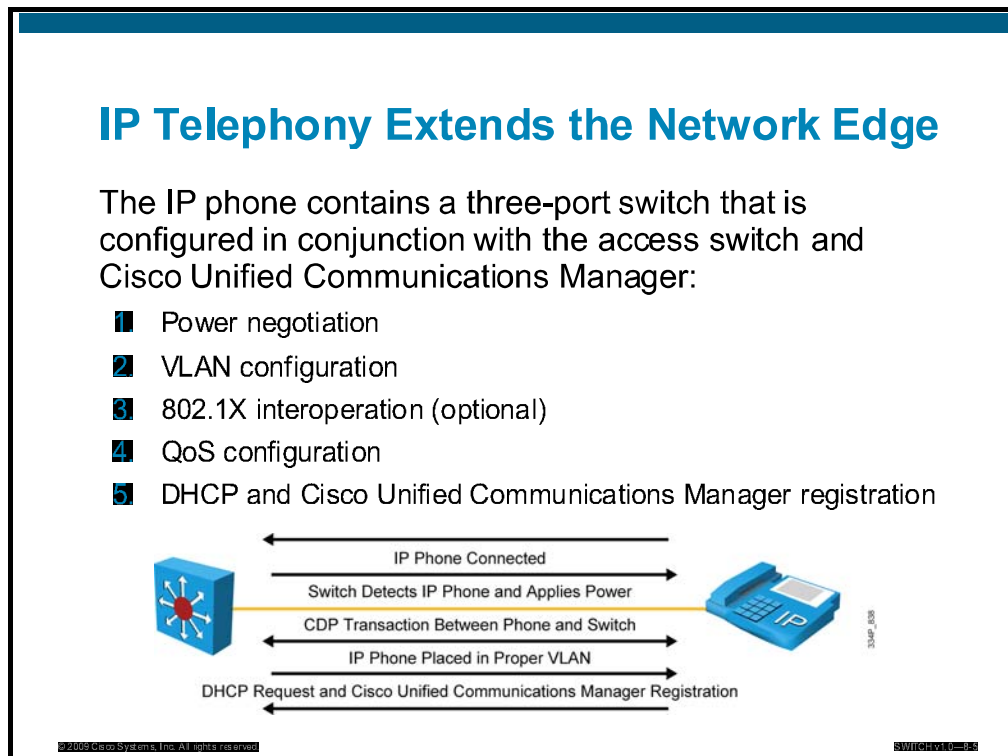
The user simply plugs the phone into the switch, and the switch will provide the phone with the necessary VLAN information. By placing phones into their own VLANs, network administrators gain the advantages of network segmentation and control. Furthermore, network administrators can preserve their existing IP topology for the data end stations. IP phones can be easily assigned to different IP subnets using standards-based DHCP operation.

With the phones in their own IP subnets and VLANs, network administrators can more easily identify and troubleshoot network problems. In addition, network administrators can create and enforce QoS or security policies.

With the auxiliary VLAN feature, Cisco enables network administrators to gain all the advantages of physical infrastructure convergence while maintaining separate logical topologies for voice and data terminals. This ability offers the most effective way to manage a multiservice network.

# IP Telephony Extends the Network Edge

This subtopic describes how IP telephony extends the network edge.



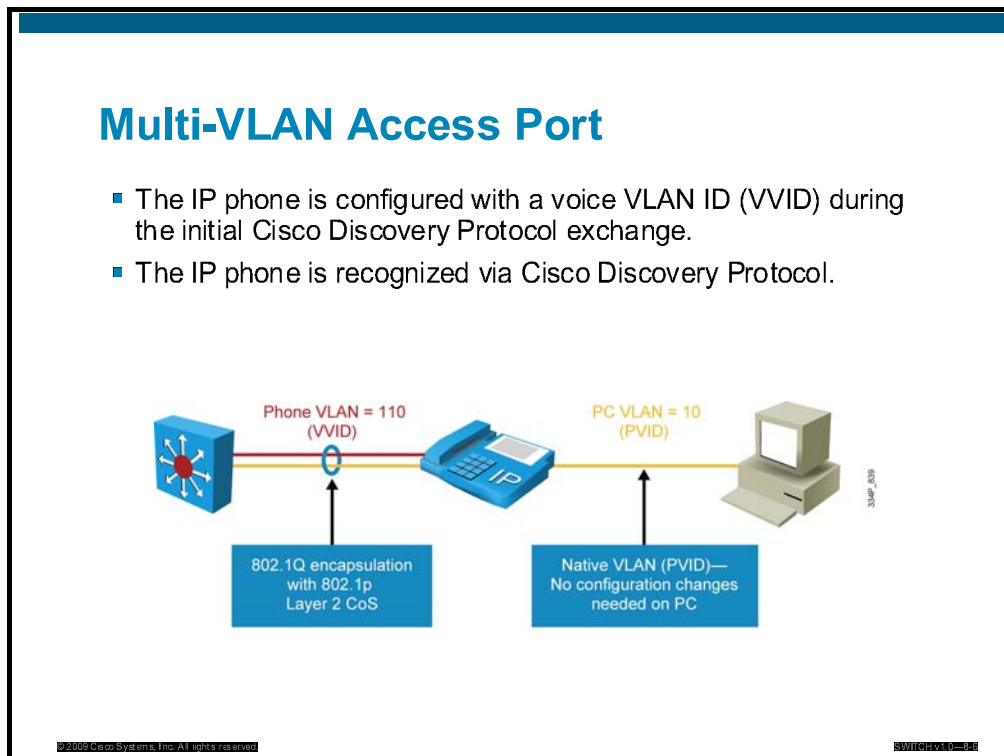
When a Cisco IP phone is connected to the network, Cisco Catalyst multiservice switches detect and integrate the phone with the network. The switches can deliver Power over Ethernet (PoE) using existing copper cabling to power the IP phones. The switches place the IP phones and attached devices in the appropriate VLAN, often using IEEE 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 “inline 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.

## Multi-VLAN Access Port

This subtopic describe the multi-VLAN access port for connecting an IP phone.



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 Cisco Discovery Protocol exchange with the access switch, the IP phone is configured with a VVID.
  - The IP phone is also supplied with a QoS configuration using Cisco Discovery Protocol. Voice traffic is separated from data traffic 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.

The IP phone tags voice packets based on the Cisco Discovery Protocol 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 two VLANs, 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.

# Configuring and Verifying Voice VLANs

This topic describes the configuration and verification of a voice VLAN.

## Voice VLAN Configuration

```
switch(config)# interface fastethernet 0/1
switch(config-if)# switchport mode access
switch(config-if)# switchport access vlan 10
switch(config-if)# switchport voice vlan 110
switch(config-if)# spanning-tree portfast
switch(config-if)# spanning-tree bpduguard enable
switch(config-if)# cdp enable
```

```
switch#show vlan
```

VLAN	Name	Status	Ports
1	default	active	Fa0/6, Fa0/7, Fa0/8, Fa0/9, Fa0/10
10	VLAN0010	active	Fa0/1
110	VLAN0110	active	Fa0/1

In the example above, interface Fa0/1 is configured to set data devices in data VLAN 10 and VoIP devices in voice VLAN 110. The switch uses Cisco Discovery Protocol and QoS marking to recognize the voice traffic. Because the port leads to an end device, PortFast is enabled. As a precaution, bridge protocol data unit (BPDU) guard is also enabled to avoid the Layer 2 loop issues created by switch rewiring.

When running the **show vlan** command, both the voice and the data VLAN are seen as applied to the interface Fa0/1.

# Power over Ethernet

This topic describes planning for PoE requirements and the configuration and verification of PoE.

## Power Sources for Access Points and IP Phones

<b>PoE switch</b> <ul style="list-style-type: none"><li>■ Benefits<ul style="list-style-type: none"><li>— Remote management</li><li>— Easy installation of UPS</li><li>— No additional power cabling</li></ul></li></ul>	<b>Power adapter</b> <b>Power injector</b> <ul style="list-style-type: none"><li>■ Disadvantages<ul style="list-style-type: none"><li>— Cannot be remotely managed</li><li>— May require additional configuration for access points</li><li>— Additional power cabling</li></ul></li></ul>
--	---

© 2009 Cisco Systems, Inc. All rights reserved.SWITCHING-99

All VoIP devices need a source of energy. When these devices are handheld mobile devices, the source of energy is usually a battery. When built for a static environment, on a desk and connected to an Ethernet cable, providing power through a battery is no longer the best solution for these devices. The VoIP phone can be plugged into an AC/DC outlet.

Although power is easily available in an office environment, using one AC/DC socket per VoIP may be considered to be a waste of physical resources. Because the phone has a cable connection to the Ethernet switch, a logical solution is to use this cable to provide power as well as connectivity.

This setting, called Power over Ethernet (PoE), implies that the power comes through the Ethernet cable. The source of this power may be the switch itself, if the switch is able to provide power to client devices. The switch is said to be “PoE able.”

If the switch itself cannot provide power, it is often possible to install an intermediate device between the switch and the VoIP phone. This device will receive power from a power outlet, and will connect to the switch with another cable. It connects to the port to which the phone will be connected. A third cable runs from this intermediate device to the phone, providing power along with data transmission to and from the switch. This device is called a “power injector.”

## Power over Ethernet (PoE)

- Sending operating power over Ethernet Category 5 cable
- Power-sourcing equipment (PSE)
  - Switches
  - Midspan power panels
- Powered devices
  - Access points
  - IP phones
  - IP surveillance cameras
- PoE supports distances up to 100 meters



To decrease the cost and complexity of the installation, powering the devices directly from the switch is often seen as the best solution.

A great advantage of PoE is that no electrician is required. Anyone qualified to run Category 5 cable can install the cabling that is required to power PoE-enabled devices. The standard Category 5 cable requirements still apply (maximum 328 feet or 100 meters).

Power-sourcing equipment (PSE) can be switches, routers with switch modules, and power injectors.

New PoE switches, such as the Cisco Catalyst 3560 24-port switch, can supply power of up to 15 W per port.

# Power over Ethernet 802.3af

This subtopic describes Power over Ethernet IEEE 802.3af.

## Power over Ethernet 802.3af

- Two incompatible PoE detection methods:
  - Cisco inline power (2000)
  - IEEE 802.3af standard (2003)
  - Supports up to 15.4 W per port
- New Cisco devices (PSEs and powered devices) support both PoE methods
  - Cisco Aironet 1131AG and 1242AG Access Points
  - Cisco Catalyst Switches: 3560, 3750, 4500, and 6500 line cards
  - Routers with PoE: Cisco 1812 Integrated Services router, switch modules
- Automatic detection is supported; no configuration is required.

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCHING-8-11

Two common Power over Ethernet methods exist: Cisco inline power and the IEEE 802.3af standard. Cisco inline power is prestandard because the IEEE 802.3 standard did not include specifications for Power over Ethernet. The 802.3af task force corrected this omission and a standard was realized in 2003. Cisco was actively involved in the 802.3af task force and support of the IEEE inline power standard. One of the main differences is that the 802.3af standard uses a power-detection mechanism that detects if the connected device needs power.

New Cisco devices (switches, access points) support both methods for backward compatibility. No specific configuration is required to choose the Cisco prestandard or the 802.3af standard.



# New Power over Ethernet Developments

This subtopic describes new Power over Ethernet developments.

## New PoE Developments

- New devices, such as access points and IP surveillance cameras, need more than 15 W of power.
- The new 802.3at standard is under development
  - Can supply up to 30 W of power
- Interim solution from Cisco
  - Enhanced PoE supplies up to 20 W of power
  - Is available on Cisco Catalyst 3560E and 3750E switches
  - Cisco Aironet 1250AG Access Point requires 18 W of power
  - Cisco Catalyst 3560 and 3750 switches with Cisco IOS Release 12.2(46)+ support the Aironet 1250AG Access Point

The IEEE 802.3af standard was revised and released as a new standard, IEEE 802.3at, which provides more power. At the time of this course development, this new standard was not implemented yet.

However, Cisco Systems provides an interim solution called “enhanced PoE,” which provides up to 20 W of power with the Cisco Catalyst E Series switches.

While most devices can use power supplied by a standard 802.3af source, some new devices require more power. For example, new access points such as the Cisco Aironet 1250 Access Point require up to 18 W. An 802.3af switch would not provide the required amount of power. With Cisco IOS Release 12.2 (46) and later, the Cisco Catalyst 3560-E and 3750-E switches can power the Aironet 1250AG Access Point, which requires 18 W for full operation. With this Cisco IOS access point and switch communicate power capabilities via Cisco Discovery Protocol, which allows the Aironet 1250AG Access Point to operate with reduced power of 15 W.

Power requirements for access points are specific. Cisco IP phones use less than 15 watts and can be powered from a standard 802.3af switch.

# PoE Configuration

This subtopic describes PoE configuration.

## PoE Switch

```
switch(config-if)# power inline {auto | never}
```

- PoE configuration

```
switch# show power inline [interface]
```

- Display PoE statistics

```
switch# show power inline
Available:124.0(w)  Used:91.2(w)  Remaining:32.8(w)
Interface Admin Oper      Power   Device                Class Max
              (Watts)
-----
Fa0/1      auto   on       15.0    AIR-LAP1242AG-E-K9    3      15.4
Fa0/2      auto   on       15.0    AIR-LAP1242AG-E-K9    3      15.4
Fa0/3      auto   on       15.0    AIR-AP1242AG-E-K9     3      15.4
Fa0/4      auto   on       15.4    AIR-LAP1142N-E-K9     3      15.4
Fa0/5      auto   on       15.4    AIR-AP1252AG-E-K9     3      15.4
Fa0/6      never  off      0.0     n/a                   n/a    15.4
Fa0/7      auto   off      0.0     n/a                   n/a    15.4
Fa0/8      auto   on       15.4    WS-C2960PD-8TT-L      3      15.4
```

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH10-8-19

Turning on PoE support is done at the port level. The **power inline auto** command is sufficient to enable PoE and autodetection of power requirements. A device not needing any PoE can still be connected to that port; power is supplied only if the device requires it. The amount of power supplied will be automatically detected. You still have to plan for the overall power consumed by all the devices connected to the PoE switch.

PoE is disabled with the **power inline never** command. Shutting the port down also stops the power supply.

The **show power inline** command displays the configuration and statistics about the power drawn by connected powered devices and the capacity of the power supply.

# Switch Power Budget

This subtopic describes the switch power budget.

## Switch Power Budget

```
switch# show power inline  
Available:124.0(w)   Used:91.2(w)   Remaining:32.8(w)
```

- Every switch has a dedicated amount of power available for PoE.
- Power used by the device is learned by the switch via Cisco Discovery Protocol.
- The power used by the device depends on the device 802.3af category or class.
- If the remaining power is less than 15 W, additional devices added will not get power.
- Not all switches have 15 W of power available for all ports.
- Plan for the power required for the number of connected powered devices.
- Example: Catalyst 3560 48-port switch with 370 W for a maximum of 24 PoE ports

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-10-0318

Every switch has a dedicated maximum amount of power available for PoE. The power used by each PoE port is deducted from the total available power. The power used by each device is dynamically learned via Cisco Discovery Protocol. This feature allows for optimization of the actual power consumption, as a PoE device does not always consume the maximum amount of power it may need. For example, an IP phone classified as 15 watts may use up to 15 watts when fully utilized, but might use only 6 or 7 watts while it is on hook.

Nevertheless, if a phone is said to require up to 15 watts, if the remaining power at the switch level is less than 15 W, the phone will not get power at all. Notice that not all switches have 15 watts of power available for all ports. Some switches have all ports supporting PoE with 15 watts each (which is the maximum possible power under the 802.3af protocol), while other switches have only a few ports supporting PoE, and not all of them offer 15 watts. For example, the Cisco Catalyst 3560 48-port switch (C3560-48PS) offers a maximum of 370 watts with 24 PoE ports.

# PoE Verification

This subtopic describes the verification of PoE switch port status.

## PoE Switch Port Status

Port	Description	Status	VLAN	Speed	Duplex	PoE
Fa0/1	LXP2	●	2	100	full	15.0 On
Fa0/2	LXP3	●	3	100	full	15.0 On
Fa0/3	LXP4	●	4	100	full	15.0 On
Fa0/4	LXP5	●	5	100	full	15.4 On
Fa0/5	LXP6	●	6	100	full	15.4 On
Fa0/6	Reserve	●	2			Off
Fa0/7	Reserve	●	2			0ff
Fa0/8	Router / Switch2	●	trunk	100	full	15.4 On
Gi0/1	Router	●	trunk	100	full	N/A

The Catalyst switch device manager displays the port status and the PoE statistics. The graphical interface is a very comfortable way of monitoring power consumption. Keep in mind that you can see the same information from the command-line interface (CLI) with the **show power** family of commands.


# Adding Upper-Layer Services

This topic describes Layer 3 to Layer 7 services required by VoIP devices.

## Additional VoIP Services

Once connected at Layer 2, VoIP devices need to:

- Obtain IP address
- Get firmware and VoIP configuration file from TFTP server
- Find and resolve Cisco Unified Communications Manager or Unified Communications Manager Express
- Register to Cisco Unified Communications Manager or Unified Communications Manager Express to obtain their extension



Once a VoIP device obtains power and accesses the network at Layer 2, it needs to obtain an IP address. Just like any other client device, it can obtain this IP address through a DHCP server.

Voice devices must download a specific configuration file that provides specific voice information, such as the codec and the location of Cisco Unified Communications Manager or Cisco Unified Communications Manager Express. This configuration file is downloaded using TFTP. The TFTP server IP address can be configured manually into each VoIP device, or provided as a DHCP option 150 (TFTP server address).

The VoIP device will then download its configuration file from the TFTP server, and also verify if a newer firmware is available. The VoIP device will then try to reach a Cisco Unified Communications Manager or Cisco Unified Communications Manager Express server. The IP address of this server can be provided within the configuration file, or can be provisioned through DNS, the resolution of the host CiscoCM1 that any Cisco VoIP devices attempts automatically. Many deployments use Cisco Unified Communications Manager or Cisco Unified Communications Manager Express as the TFTP server to simplify the overall procedure, thus removing the need for an external TFTP server and DNS resolution.

The phone contacts Cisco Unified Communications Manager or Cisco Unified Communications Manager Express, registers to it, obtains its line extension number, and is ready to place or receive calls.

# Test Plan

This topic describes a typical test plan for implementing VoIP.

## Test Plan

Does the IP phone receive an IP address?

Does the PC behind the IP phone receive an IP address?

Does the IP phone register with Cisco Unified Communications Manager or Cisco Unified Communications Manager Express?

Can you place a call to another IP phone?

To test voice implementation, it is recommended that you work through the logical boot process of the phone.

Like any other networking device, a phone needs to receive power upon boot up. It will then receive an IP address. If the IP address is received via DHCP, the DHCP option should also provide the Cisco Unified Communications Manager or Cisco Unified Communications Manager Express IP address. The phone will download its configuration and firmware from the voice server before registering, obtaining a phone extension, and be ready to place a call.

When testing a new VoIP implementation, a logical testing scheme is to check that the phone receives power, receives an IP address, actually tries to contact the Cisco Unified Communications Manager or Cisco Unified Communications Manager Express, and registers with the system. Once all these steps are complete, the phone should be able to place and receive calls.

# Summary

This topic summarizes the key points that were discussed in this lesson.

## Summary

- Voice integration implies separation of voice and data, power, QoS, and high availability considerations.
- Voice VLANs provide the ability to apply QoS to voice traffic without affecting the flow of data from the client PC.
- Configuration is necessary to separate VoIP traffic from data traffic coming to the same switch port.
- IP phones, access points, and IP surveillance cameras can be powered by Power over Ethernet.
- Once VoIP devices connect to the switch, they need to obtain specific VoIP-related services.
- A test can be conducted to verify whether the phone correctly connects and communicates with the infrastructure.





# Working with Specialists to Accommodate Voice and Video on Campus Switches

---

## Overview

Implementing voice and video in a campus network requires close cooperation between the voice specialist and the network professional. Voice and video are usually very sensitive to delay and jitter. Considerations such as requirements for quality of service (QoS) or high availability need to be discussed with the voice or video specialist to ensure an optimal flow throughout the network. This lesson describes the main points to consider and provides best practices for integrating voice and video in the corporate network.

## Objectives

Upon completing this lesson, you will be able to plan integration of VoIP and video traffic into a data network based on input from voice and video specialists. This ability includes being able to meet these objectives:

- Describe high availability applied to VoIP or video traffic
- Build an integrated voice/video/data campus network
- Explain the need for QoS for VoIP and video integration
- Configure basic QoS for voice and video VLANs

# High Availability for VoIP and Video

This subtopic describes the high availability requirements for voice and video traffic.

## High Availability for VoIP and Video

- Traditional telephony networks claim 99.999 percent uptime.
- Real-time video applications are very sensitive to packet loss and delay.
- Data networks must consider reliability and availability requirements when incorporating voice and video.
- Methods to improve reliability and availability include:
  - Redundant hardware
  - Redundant links
  - UPS
  - Proactive network management

To provide telephony users the same—or close to the same—level of service as they experience with traditional telephony, the reliability and availability of the data network takes on new importance.

Reliability is a measure of how resilient a network can be. Efforts to ensure reliability may include choosing hardware and software with a low mean time between failures, or installing redundant hardware and links. Availability is a measure of how accessible the network is to the users.

When a user wants to make a call, for example, the network should be accessible to that user at any time a call is required. Efforts to ensure availability may include installing proactive network management to predict failures before they happen and taking steps to correct problems in the design of the network as it grows.

When the data network goes down, it may not come back up for minutes or even hours. This delay is unacceptable for telephony users. Local users with network equipment, such as voice-enabled routers, gateways, or switches for IP phones, now find that their connectivity has been terminated. Administrators must provide an uninterruptible power supply (UPS) to these devices in addition to providing network availability.

Previously, depending on the type of connection the user had, users received their power directly from the telephone company central office or through a UPS that was connected to their keyswitch or PBX in the event of a power outage. Network devices must now have protected power to continue to function and provide power to the end devices.

Network reliability is based on incorporating redundancy into the network design. In traditional telephony, switches have multiple redundant connections to other switches. If either a link or a switch becomes unavailable, the telephone company can route the call in different ways. This ability is the reason that telephone companies can claim a high availability rate.

High availability encompasses many areas of the network. In a fully redundant network, the following components need to be duplicated:

- Servers and call managers
- Access layer devices, such as LAN switches
- Distribution layer devices, such as routers or multilayer switches
- Core layer devices, such as multilayer switches
- Interconnections, such as WAN links and public switched telephone network (PSTN) gateways, even through different providers
- Power supplies and UPSs

## **Example: Cisco Reliability and Availability**

In some data networks, a high level of availability and reliability is not critical enough to warrant financing the hardware and links required to provide complete redundancy. If voice is layered onto the network, these requirements need to be revisited.

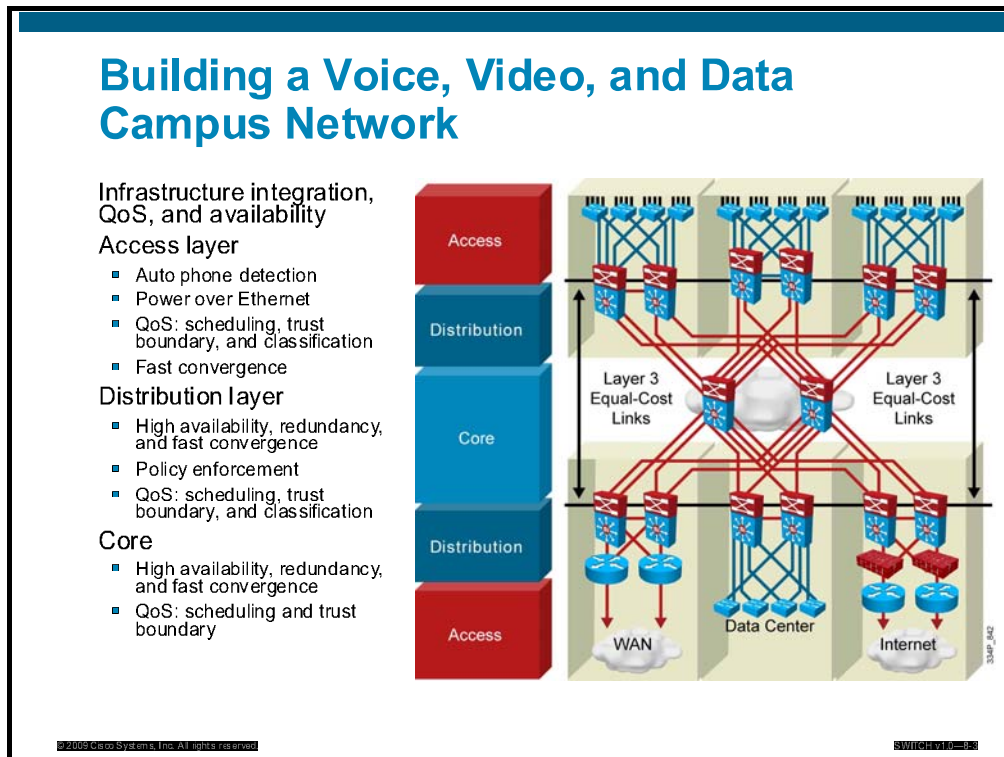
With Cisco Unified Communications system technology, the use of Cisco Unified Communications Manager clusters provides a way to design redundant hardware in the event of a Cisco Unified Communications Manager failure.

When using gatekeepers, you can configure backup devices as secondary gatekeepers in case the primary gatekeeper fails. You must also revisit the network infrastructure. Redundant devices and Cisco IOS services, such as Hot Standby Router Protocol (HSRP), can provide high availability.

For proactive network monitoring and trouble reporting, a network management platform such as CiscoWorks2000 provides a high degree of responsiveness to network issues.

# Building Voice/Video/Data Campus Networks

This topic describes building voice/video/data campus networks.



Each layer of a voice and video campus network has its own requirements, which is also true in data networks.

At the access layer, use Cisco Discovery Protocol to detect the IP phones. Power over Ethernet (PoE) can be enabled and set to auto, in order to provide power whenever needed. Video devices usually have their own power source and do not need PoE or Cisco Discovery Protocol discovery. The access layer is also where classification and marking is performed, to recognize and tag voice and video packets. These tags will be used throughout the network to prioritize the voice flow.

QoS policy is enforced at the distribution layer using sets of policy maps that will guarantee that voice or video traffic has the prioritization and bandwidth reservation it needs. At this layer, pay close attention to availability. Because it is common to oversubscribe the distribution layer links capacity, congestion can easily occur. This situation is true with voice traffic and even more delicate when bandwidth-intensive video traffic is involved. Verify that voice or video traffic can still get the bandwidth and delay it needs: QoS alone cannot solve issues involving lack of bandwidth.


The core layer is where packets are quickly switched from one network to the other. At this layer, no classification or marking occurs, and policy enforcement should not be needed any longer. Nevertheless, packets coming from foreign networks may enter close to the core layer. When unknown packets enter at the core layer, a trust boundary must be set there, and QoS scheduling may still have to take place. Oversubscription is still a common issue at the core layer, as it is at the distribution layer. Verify that voice and video have access to the bandwidth and delay needed.

# Quality of Service

This topic describes campus QoS for voice.

## Determining Equipment and Cabling Needs

Each link provides adequate bandwidth for traffic aggregation over that link.




Typical campus networks are designed with oversubscription. 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 usual 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. Use of these oversubscription ratios should avoid congestion on the uplinks, but QoS is needed for instances where congestion does occur. If congestion occurs frequently, the design does not have sufficient uplink bandwidth.

## Resource Contention

### Resource Contention for Voice Traffic

- Uplinks between the access and distribution layers
- Links to and from the core
- Links within the server farm



© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-10000

While these oversubscription ratios are generally valid when voice is added to a data network, you should carefully consider the traffic flows when video is added. Because video is bandwidth intensive, these ratios may starve video applications.

When voice calls are placed, control traffic flows to and from the Cisco Unified Communications Manager. As soon as phone-to-phone communication is established, the main stream flows directly between the phones. Only a small amount of data is exchanged between the phones and the Cisco Unified Communications Manager for all management and termination functions.

The same logic applies to peer-to-peer video conferencing. Traffic flows directly from one video endpoint to the other.

Streaming video usually displays a different pattern. The video source is stored in a data center, and traffic flows from the data center to the video client.


Monitoring the traffic flows and analyzing bandwidth consumption is the best way to determine the oversubscription ratio that best matches your network requirements.

# Recommended Practices for QoS

This subtopic describes recommended practices for QoS

## Recommended Practices: QoS

- Deployed in an end-to-end configuration
- 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



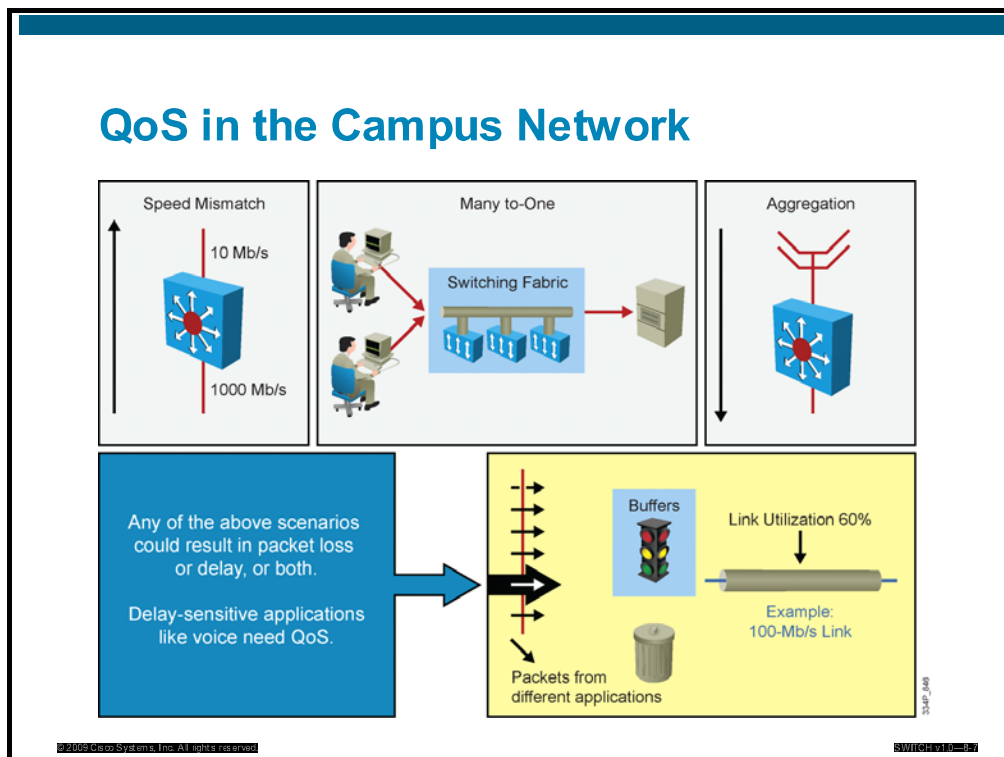
Effective QoS is deployed in an end-to-end configuration with each layer supporting a role.

Internet worms and denial of service (DoS) attacks have the ability to 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.

## QoS and Time-Sensitive Traffic

This subtopic describes how QoS is applied for time-sensitive voice and video traffic in the campus model.



Regardless of the speed of individual switches or links, speed mismatches, many-to-one switching fabrics, and aggregation may cause a device to experience congestion, which can result in latency. If congestion occurs and congestion management features are not in place, then some packets will be dropped, causing retransmissions that inevitably increase overall network load. QoS can mitigate latency caused by congestion on campus devices.

QoS is implemented by classifying and marking traffic at one device while allowing other devices to prioritize or to queue the traffic according to those marks applied to individual frames or packets. The table lists the campus devices involved in QoS marking or prioritizing.

### QoS Application in the Campus Network

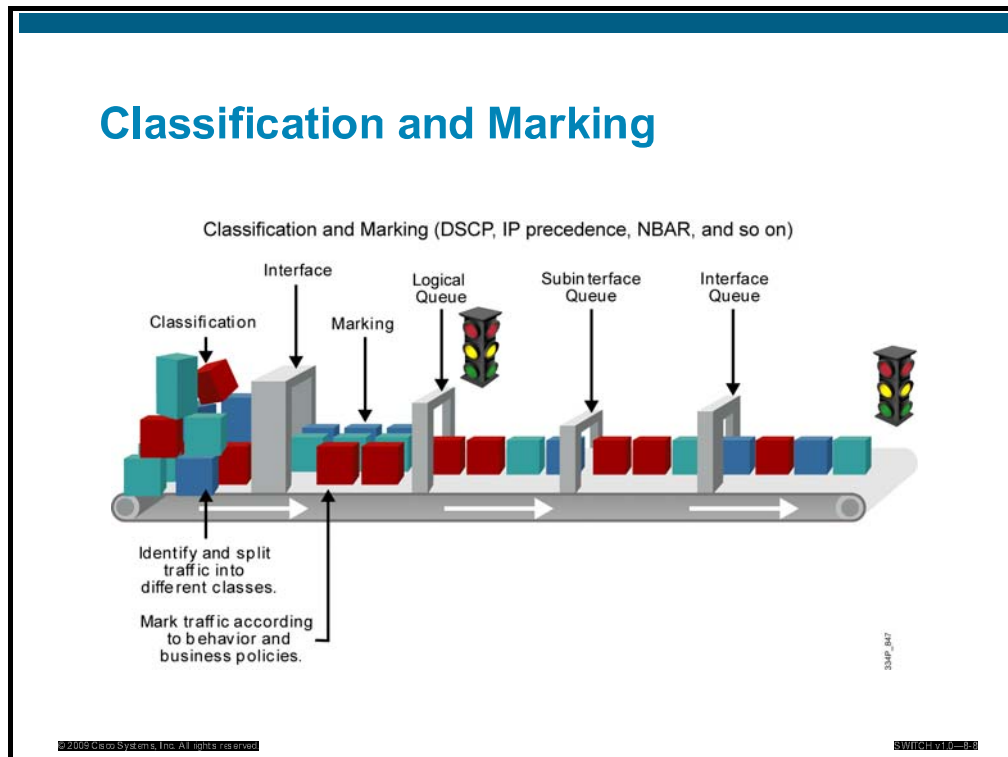
The table describes how QoS is applied in the campus network.

Campus Device	QoS Application
Access Layer	Initial point at which traffic enters the network. Traffic can be marked (or re-marked) at Layers 2 and 3 by the access switch as traffic enters the network or “trusted” that it is entering the network with an appropriate tag.
Distribution Layer	Marks of traffic inbound from the access layer can be trusted or reset, depending on the ability of the access layer switches. Priority access into the core is provided based on Layer 3 QoS tags.
Core	No traffic marking occurs at the core. Layer 2 or 3 QoS tags are trusted from distribution layer switches and used to prioritize and to queue the traffic as it traverses the core.



# LAN-Based Classification and Marking

This subtopic describes LAN-based classification and marking using a Layer 2 Cisco Catalyst workgroup switch.



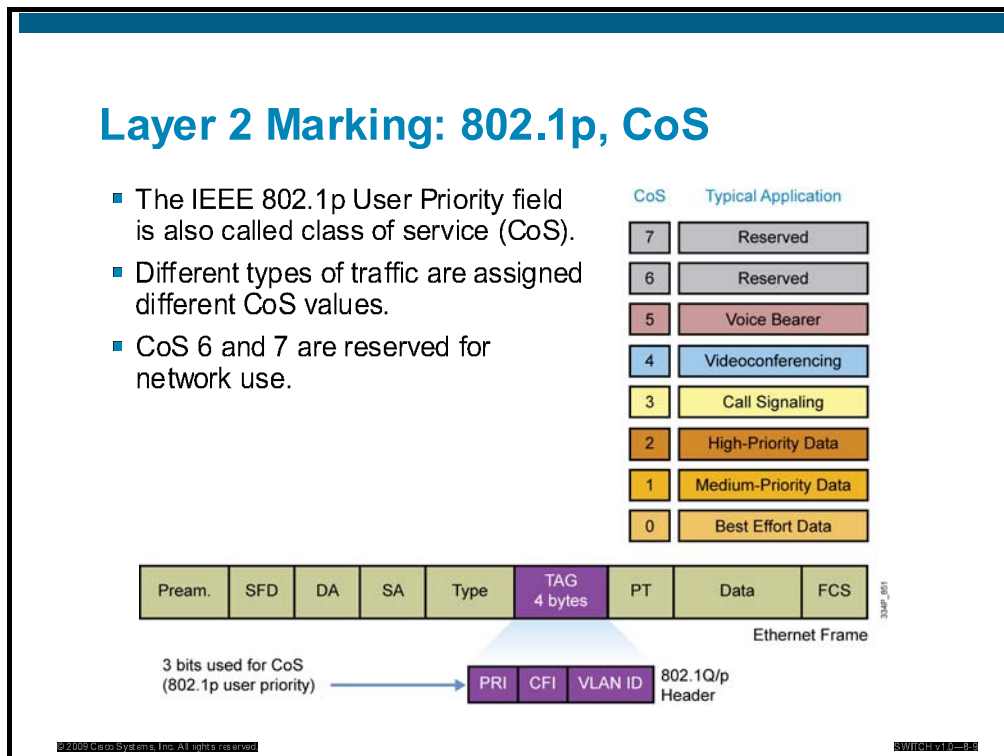
Classification and marking is the process of identifying traffic for proper prioritization as that traffic traverses the network. Traffic is classified by examining information at various layers of the Open Systems Interconnection (OSI) model. All traffic classified in a certain manner will receive an associated mark or QoS value. IP traffic can be classified according to any values configurable in an access control list (ACL) or any of these criteria:

- **Layer 2 parameters:** MAC address, Multiprotocol Label Switching (MPLS), ATM cell loss priority (CLP) bit, Frame Relay discard eligible (DE) bit, ingress interface
- **Layer 3 parameters:** IP precedence, DiffServe code point (DSCP), QoS group, IP address, ingress interface
- **Layer 4 parameters:** TCP or User Datagram Protocol (UDP) ports, ingress interface
- **Layer 7 parameters:** Application signatures, ingress interface

All traffic classified or grouped according to these criteria will be marked according to that classification. QoS marks or values establish priority levels or priority classes of service for network traffic as each switch processes the traffic. Once traffic is marked with a QoS value, then QoS policies on switches and interfaces will handle traffic according to the values contained in individual frames and packets. Because of classification and marking, traffic will be prioritized accordingly at each switch to ensure that delay-sensitive traffic receives priority processing as the switch manages congestion, delay, and bandwidth allocation.

## Layer 2 QoS Marking

This subtopic describes how QoS values are carried in Layer 2 headers.



QoS Layer 2 classification occurs by examining information in the Ethernet or 802.1Q header, such as the destination MAC address or VLAN ID. QoS Layer 2 marking occurs in the Priority field of the 802.1Q header. LAN Layer 2 headers have no means of carrying a QoS value, so 802.1Q encapsulation is required if Layer 2 QoS marking is to occur. The Priority field is 3 bits long and is also known as the 802.1p User Priority or class of service (CoS) value.

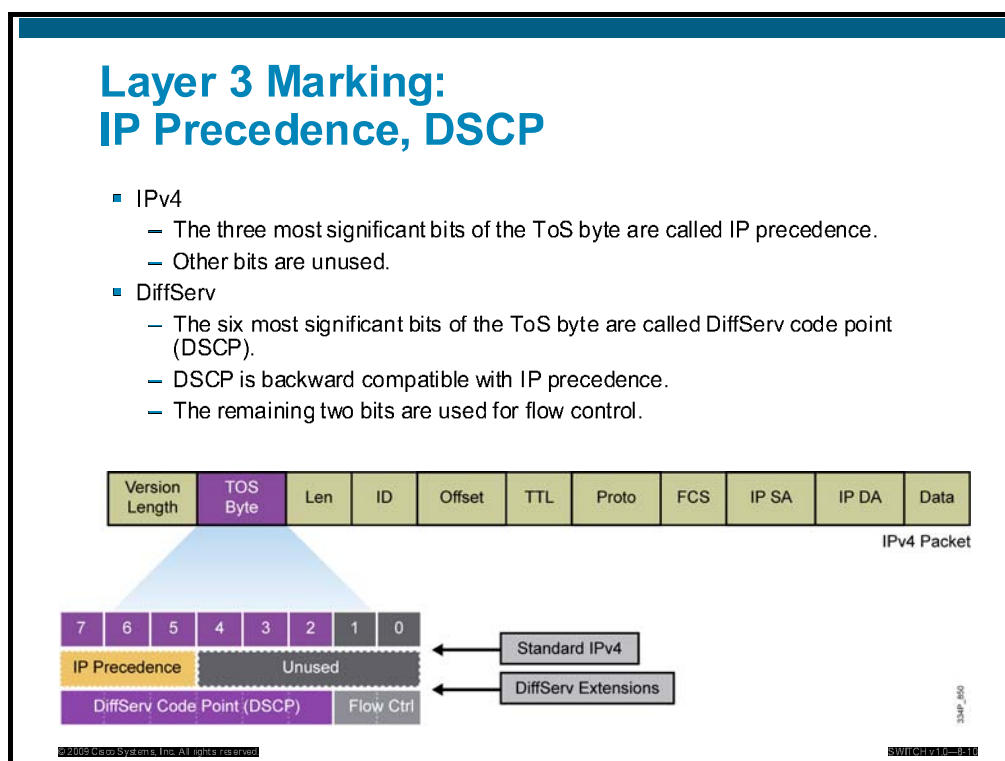
This 3-bit field supports CoS values ranging from 1 to 7; 1 is associated with delay-tolerant traffic such as TCP/IP. Voice traffic, which by nature is not delay tolerant, receives higher default CoS values, such as 3 for call signaling. A CoS value of 5 is given to voice-bearer traffic, which is the phone conversation itself, in which voice quality is impaired if packets are dropped or delayed.

As a result of Layer 2 classification and marking, these QoS operations can occur:

- **Input queue scheduling:** When a frame enters a port, it can be assigned to one of a number of port-based queues before being scheduled for switching to an egress port. Typically, multiple queues are used where traffic requires different service levels.
- **Policing:** Policing is the process of inspecting a frame to see if it has exceeded a predefined rate of traffic within a certain time frame, which is typically a fixed number internal to the switch. If a frame is determined to be in excess of the predefined rate limit, it can either be dropped, or the CoS value can be marked down.
- **Output queue scheduling:** The switch will place the frame into an appropriate outbound (egress) queue for switching. The switch will perform buffer management on this queue by ensuring that the buffer does not overflow.

# Layer 3 QoS Marking

This subtopic describes QoS information carried in Layer 3 headers.



QoS Layer 3 classification results from the examination of header values such as destination IP address or protocol. QoS Layer 3 marking occurs in the type of service (ToS) byte in the IP header. The first three bits of the ToS byte are occupied by IP precedence, which correlates to the three CoS bits carried in the Layer 2 header.

ToS Byte:	P2	P1	P0	T3	T2	T1	T0	Zero
DS Byte:	DS5	DS4	DS3	DS2	DS1	DS0	ECN1	ECN0
	(Class Selector)			(Drop Precedence)				

The ToS byte can also be used for DSCP marking. DSCP allows prioritization on a hop-by-hop basis as packets are processed on each switch and interface. The ToS bits are used by DSCP values as shown in the table. The first three DSCP bits, correlating to IP precedence and CoS, identify the DSCP CoS for the packet.

The next three DSCP bits establish a drop precedence value for the packet. Packets with a high DSCP drop precedence value will be dropped before those with a low value if a device or a queue becomes overloaded and must drop packets. Voice traffic will be marked with a low DSCP drop precedence value to minimize voice packet drop.


Each 6-bit DSCP value is given a DSCP name. DSCP classes 1 through 4 are Assured Forwarding (AF) classes.

# Describing QoS Trust Boundaries

This subtopic describes QoS trust boundaries.

## Classification Tools: Trust Boundaries

- A device is “trusted” if it correctly classifies packets.
- For scalability, classification should be performed as close to the edge as possible.
- The outermost trusted devices represent the “trust boundary.”
- ① and ② are optimal; ③ is acceptable (if the access switch cannot perform classification).



In a campus QoS implementation, boundaries are defined where the existing QoS values that are attached to frames and to packets are to be accepted or altered. These “trust boundaries” are established by configuring trust levels on the ports of key peripheral network devices where QoS policies will be enforced as traffic makes its way into the network. At these boundaries, traffic will be allowed to retain its original QoS marking or will be assigned new markings because of policies associated with its entry point into the network.

Trust boundaries establish a border for traffic entering the campus network. As traffic traverses the switches of the campus network, it is handled and prioritized according to the marks received or trusted when the traffic originally entered the network at the trust boundary.

At the trust boundary device, QoS values are trusted if they are considered to accurately represent the type of traffic and precedence processing that the traffic should receive as it enters the campus network.

If untrusted, the traffic is marked with a new QoS value that is appropriate for the policy that is in place at the point where the traffic enters the campus network. Ideally, the trust boundary exists at the first switch that receives traffic from a device or IP phone. It is also acceptable to establish the trust boundary at the point where traffic from an access switch enters a building distribution layer port.

---

**Note** Best practices suggest classifying and marking traffic as close to the traffic source as possible.

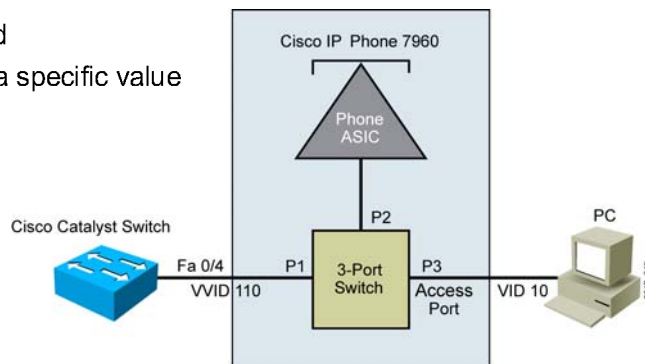
---

# Cisco Phone Connected to a Switch

This subtopic describes a Cisco IP phone connected to a switch.

## Cisco IP Phone Connected to a Switch

- Voice traffic is tagged for voice VLAN
- Data VLAN traffic from the PC can be one of the following:
  - Untrusted
  - Trusted
  - Set to a specific value



These commands are used to configure and to verify basic features used to manage voice traffic on Cisco Catalyst switch ports.

Step	Description
1.	Enable voice VLAN on a switch port and associate a VLAN ID. Switch(config-if)# <b>switchport voice vlan vlan-id</b>
2.	Trust the CoS value of frames as they arrive at the switch port. Switch(config-if)# <b>mls qos trust cos</b>
3.	Make a trust conditional on a Cisco IP phone being attached. Switch(config-if)# <b>mls qos trust device cisco-phone</b> Or Set the CoS value to frames coming from the PC attached to the IP phone. Switch(config-if)# <b>switchport priority extend cos cos_value</b>
3.	Display voice parameters configured on the interface. Switch# <b>show interfaces interface-id switchport</b>
4.	Display QoS parameters configured on the interface. Switch# <b>show mls qos interface interface-id</b>

# Configuring QoS for Voice VLANs

This topic describes the configuration of QoS for a voice VLAN.

## Voice VLAN Configuration

```
switch(config)# mls qos
switch(config)# interface fastethernet 0/1
switch(config-if)# switchport mode access
switch(config-if)# switchport access vlan 10
switch(config-if)# switchport voice vlan 110
switch(config-if)# mls qos trust cos
switch(config-if)# mls qos trust device cisco-phone
```

```
switch# show mls qos interface fastethernet 0/1
FastEthernet0/1
trust state: trust cos
trust mode: trust cos
COS override: dis
default COS: 0
pass-through: none
trust device: cisco-phone
```

To enable QoS on a switch, some hardware platforms require that you to enter a QoS-related command, while others need to have QoS support globally enabled through the **mls qos** command.

In the example in the figure above, QoS support is first enabled globally. On the first interface, a data VLAN 10 is configured. On the same link, voice traffic is sent to VLAN 100. The CoS values received from the phone are trusted.

The **show mls qos** family of commands can be used to display QoS configuration.

# Configuring Cisco AutoQoS

This subtopic describes Cisco AutoQoS commands.

## Cisco AutoQoS

- A single command at the interface level configures the interface and global QoS.
  - Support for Cisco IP Phone and Cisco IP Communicator.
    - Support for Cisco IP Communicator currently exists only on the Cisco Catalyst Series 6500 switch.
  - The trust boundary is disabled when a Cisco IP phone is moved.
  - Buffer allocation and egress queuing are dependent on the interface type (Gigabit Ethernet/Fast Ethernet).
- Supported on static, dynamic access, voice VLAN access, and trunk ports.
- Cisco Discovery Protocol must be enabled for Cisco AutoQoS to function properly.

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-10-5-12

To configure the QoS settings and the trusted boundary feature on the Cisco IP phone, Cisco Discovery Protocol version 2 or later must be enabled on the port. If the trusted boundary feature is enabled, a syslog warning message is displayed if Cisco Discovery Protocol is not enabled or if Cisco Discovery Protocol is running version 1.

Cisco Discovery Protocol needs to be enabled for only the **ciscoipphone** QoS configuration; Cisco Discovery Protocol does not affect other components of the automatic QoS features. When the **ciscoipphone** keyword is used with the port-specific automatic QoS feature, a warning is displayed if the port does not have Cisco Discovery Protocol enabled.

The trust device feature is enabled when the port-specific automatic QoS command is executed with the **ciscoipphone** keyword, but the trust option is not used. The trust device feature is dependent on Cisco Discovery Protocol. If Cisco Discovery Protocol is not enabled or not running version 2, a warning message is displayed, as follows:

```
Console> (enable) set port qos 4/1 autoqos voip ciscoipphone
Warning: CDP is disabled or CDP version 1 is in use. Ensure
that CDP version 2 is enabled globally, and also ensure that
CDP is enabled on the port(s) you wish to configure autoqos
on.
Port 4/1 ingress QoS configured for ciscoipphone.
It is recommended to execute the "set qos autoqos" global
command if not executed previously.
Console> (enable)
```

# Cisco AutoQoS Configuration

This subtopic describes the configuration of Cisco AutoQoS

## Cisco AutoQoS Configuration

```
switch(config-if)#  
auto qos voip trust
```

- The uplink interface is connected to a trusted switch or router, and the VoIP classification in the ingress packet is trusted.

```
switch(config-if)#  
auto qos voip cisco-phone
```

- Automatically enables the trusted boundary feature, which uses Cisco Discovery Protocol to detect the presence or absence of a Cisco IP phone.
- If the interface is connected to a Cisco IP phone, the QoS labels of the incoming packets are trusted only when the Cisco IP phone is detected.

© 2009 Cisco Systems, Inc. All rights reserved. SWITCH-10-8-10

When the Cisco AutoQoS feature is enabled on the first interface, QoS is globally enabled (**mls qos** global configuration command).

When the **auto qos voip trust** interface configuration command is entered, the ingress classification on the interface is set to trust the CoS QoS label received in the packet, and the egress queues on the interface are reconfigured. QoS labels in ingress packets are trusted.

When the **auto qos voip cisco-phone** interface configuration command is entered, the trusted boundary feature is enabled. The trusted boundary feature uses Cisco Discovery Protocol to detect the presence or absence of a Cisco IP phone.

When a Cisco IP phone is detected, the ingress classification on the interface is set to trust the QoS label received in the packet. When a Cisco IP phone is absent, the ingress classification is set to not trust the QoS label in the packet. The egress queues on the interface are also reconfigured. This command extends the trust boundary if an IP phone is detected.



# Monitoring Cisco AutoQoS

This subtopic describes monitoring of Cisco AutoQoS.

## Monitoring Cisco AutoQoS

Switch#

```
show auto qos [interface interface-id]
```

- Displays the Cisco AutoQoS configuration that was initially applied
- Does not display any user changes to the configuration that might be in effect

```
switch# show auto qos  
FastEthernet0/1  
auto qos voip cisco-phone
```

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-10-EST

To display the initial Cisco AutoQoS configuration, use the **show auto qos [interface *interface-id*]** privileged EXEC command. To display any user changes to that configuration, use the **show running-config** privileged EXEC command. The **show auto qos** and the **show running-config** command output can be compared to identify the user-defined QoS settings.

# Summary

This topic summarizes the key points discussed in this lesson.

## Summary

- High-availability networks must be created to avoid network congestion and overcome a lack of redundancy and poor engineering.
- Voice and video integration imply managing link bandwidth and packet prioritization.
- To ensure high quality VoIP, implementation of QoS is required.
- QoS can reduce latency in a campus network when VoIP is configured.

## Lesson 4

---

# Lab 8-1 Debrief

---

## Overview

In this lab, you have prepared your network for future VoIP solution implementation. You have gathered and analyzed the needed information. You have created an implementation and verification plan. You have connected to the remote lab and have made the needed reconfigurations.

During the lab debrief, the instructor will lead a group discussion where you can present your solutions. You will have an opportunity to verify your solution against a number of checkpoints provided by the instructor and compare your solution to the solutions of other learners. The instructor will discuss the solutions and their benefits and drawbacks.

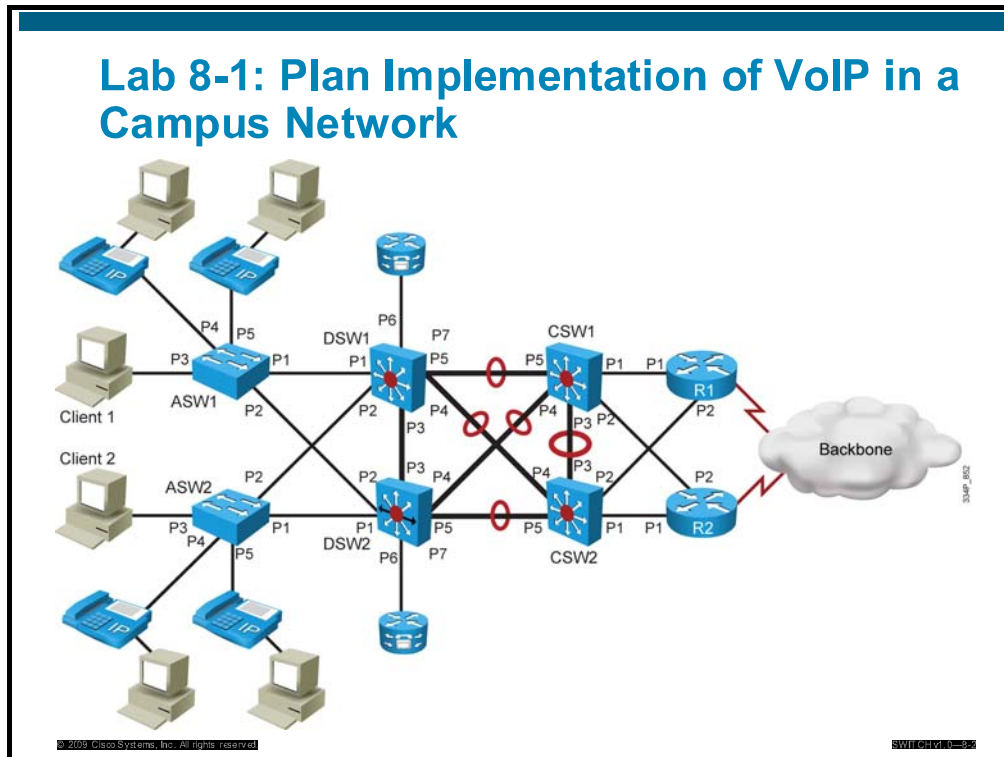
## Objectives

Upon completing this lesson, you will be able to produce a design plan and implementation plan for preparing the network for VoIP implementation based on given business and technical requirements. You will then implement the plan and perform verifications while identifying checkpoints along the way. This ability includes being able to meet these objectives:

- Review and verify your solution, as well as your findings and action log against a set of checkpoints provided by the instructor
- Consolidate the lessons learned during the review discussions into a set of best practice methods and commands to aid you in future deployment procedures

# Review and Verification

This topic describes the requirements that were listed in Lab 8-1, asks how you can verify that the solution you found matches the client needs, and gives you an example of a possible solution.



This lab is about the changes you must make in preparing a network to handle the impact from implementing a VoIP-based solution. You must gather and analyze information so that you can make decisions concerning the things that will need to be changed. After that, you must create an implementation plan and a verification plan for the changes that are needed.

## Design and Implementation Plan for VoIP in a Campus Network

Determine the items that should be configured, and the order in which the items should be configured.

- Which switches will be access switches for IP phones?
- Which switches will carry voice traffic ?
- How will QoS be configured in a campus network?
- Which router will be the DHCP server for IP phones?
- How will high availability be configured for VoIP?
- How will status be verified after the implementation?

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCHING-02-08

A successful implementation plan allows you to configure the devices without duplication and reduces the risk of mistakes. Thus, you will need to make only a minimum set of configuration changes, which will reduce the time for troubleshooting after the implementation. In other words, an implementation plan is efficient when you do not need to alter your previous configuration to implement new items. You should proceed in a logical order.



# Module Summary

This topic summarizes the key points discussed in this module.

## Module Summary

- Proper planning for support of voice in a campus network must take into account all aspects of network engineering when configuring a switch for VoIP.
- Voice VLAN and PoE configurations prepare switches for voice traffic.
- Using switch-based QoS policies and procedures in a VoIP network will ensure quality and reduce traffic.





# Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) When implementing VoIP, which design consideration is *not* an issue? (Source: Planning for Support of Voice in a Campus Network)
- A) Provision switches with inline power.
  - B) Ensure network bandwidth is adequate.
  - C) Determine if 800-number access is required.
  - D) Ensure that the physical plant can support VoIP.
- Q2) When you are installing a VoIP network, which method should you *not* use to improve reliability? (Source: Planning for Support of Voice in a Campus Network)
- A) redundant hardware
  - B) 24-hour staffing
  - C) redundant links
  - D) proactive network management
- Q3) Which response best describes the characteristics of VoIP traffic? (Source: Planning for Support of Voice in a Campus Network)
- A) smooth and delay insensitive
  - B) bursty and delay sensitive
  - C) smooth and sensible to jitter
  - D) bandwidth greedy and delay sensitive
- Q4) Which is the tolerable end-to-end delay for VoIP? (Source: Planning for Support of Voice in a Campus Network)
- A) 30 ms
  - B) 60 ms
  - C) 150 ms
  - D) less than 1 second
- Q5) Which is the function of a Call Processing Manager? (Source: Planning for Support of Voice in a Campus Network)
- A) provide call redirection when line is busy
  - B) perform call setup and routing
  - C) provide extension number to new registering phones
  - D) all of the above
- Q6) How does video traffic compare to VoIP traffic? (Source: Planning for Support of Voice in a Campus Network)
- A) Video is bandwidth intensive, just like voice.
  - B) Video is delay sensitive, just like voice.
  - C) Video needs steady packet transmission, just like voice.
  - D) All of the above are correct.

- Q7) Which transport layer protocol does VoIP use? (Source: Integrating and Verifying VoIP in a Campus Infrastructure)
- A) It uses TCP.
  - B) It uses ICMP.
  - C) It uses UDP.
  - D) It does not use a transport layer protocol; traffic goes directly from IP to the application.
- Q8) Which are two ways to accomplish QoS marking? (Choose two.) (Source: Working with Specialists to Accommodate Voice and Video on Campus Switches)
- A) using the Type field in the Ethernet header
  - B) using 802.1Q ToS bits
  - C) implementing DSCP at Layer 3
  - D) implementing DSCP at Layer 2
- Q9) In which location can trust boundaries *not* be created? (Source: Working with Specialists to Accommodate Voice and Video on Campus Switches)
- A) client IP phone
  - B) core switch
  - C) access switch
  - D) distribution switch
- Q10) Which protocol allows creating a CoS in the 802.1Q trunking protocol? (Source: Working with Specialists to Accommodate Voice and Video on Campus Switches)
- A) ISL
  - B) 802.1p
  - C) 802.1d
  - D) no protocol; CoS is part of 802.1Q
- Q11) Which is a common QoS configuration on access ports to IP phones? (Source: Working with Specialists to Accommodate Voice and Video on Campus Switches)
- A) **switchport trust cisco-phone**
  - B) **mls qos trust device cisco-phone**
  - C) **switchport qos trust cisco-phone**
  - D) **power inline trust autodetect**
- Q12) Which is the right configuration for IP phone access ports? (Source: Integrating and Verifying VoIP in a Campus Infrastructure)
- A) **switchport** in access mode, define data VLAN and define voice VLAN
  - B) **switchport** in voice access mode, define voice VLAN
  - C) **switchport** in trunk mode, allow voice VLAN and data VLAN
  - D) **switchport** in trunk mode, allow voice VLAN only
- Q13) Which are two ways in which an IP header can be configured for QoS? (Choose two.) (Source: Working with Specialists to Accommodate Voice and Video on Campus Switches)
- A) using IP precedence bits
  - B) using access lists
  - C) using resource reservation code points
  - D) using DSCPs

- Q14) How is PoE configured on a switch? (Source: Integrating and Verifying VoIP in a Campus Infrastructure)
- A) using the global command **poe enable**
  - B) using the global command **poe run** and the interface command **poe enable**
  - C) using the global command **poe autodetect**
  - D) using the interface command **power inline auto**

## Module Self-Check Answer Key

- Q1) C
- Q2) B
- Q3) C
- Q4) C
- Q5) D
- Q6) B
- Q7) C
- Q8) B, C
- Q9) B
- Q10) B
- Q11) B
- Q12) A
- Q13) A, D
- Q14) D

# Integrating Wireless LAN into a Campus Network

---

## Overview

This module introduces wireless LANs (WLANs). WLAN is an access technology that has an increasing significance for network access in offices, factories, hotels, and airports, and at home. This module explains the differences between wired and wireless LANs, describes basic WLAN topologies, and teaches you how to prepare the campus networks for the integration of WLANs.

## Module Objectives

Upon completing this module, you will be able to integrate WLANs into a campus network. This ability includes being able to meet these objectives:

- Compare the topologies and equipment of WLANs to wired campus networks
- Assess the impact of WLANs on campus infrastructure operations
- Create implementation and verification plans for preparing infrastructure devices to integrate WLANs



# Comparing WLANs to Campus Networks

---

## Overview

Wireless LANs (WLANs) are often compared to standard LANs and are often seen as “LANs without cables.” WLANs actually integrate into the LAN infrastructure to extend it. It does have similarities to wired LANs. Nevertheless, it also presents important differences that you need to understand to perform a successful integration. This lesson compares WLANs to wired LANs and describes how they differ.

## Objectives

Upon completing this lesson, you will be able to compare WLANs to wired LANs. This ability includes being able to meet these objectives:

- Describe WLANs
- Compare wired and wireless LANs
- Describe main wireless LAN topologies
- Describe the specific settings of WLANs, such as SSIDs, and WLAN-to-VLAN mapping

# WLAN Overview

This topic describes WLAN components.



The Cisco WLAN products that are shown in the figure support the five interconnecting elements of the Cisco Unified Wireless Network and business-class WLANs.

- **Client devices:** Cisco Compatible Extensions or Cisco Aironet client devices are strongly recommended for the Cisco Unified Wireless Network. With over 90 percent of shipping client devices certified as Cisco Compatible, almost any client device that you select should be Cisco Compatible certified.

Cisco Compatible client devices interoperate with and support innovative and unique Cisco Unified Wireless Network features, such as fast secure roaming, integrated intrusion prevention system (IPS), location services, and a variety of extensible authentication types.

- **Mobility platform:** Cisco offers access points and bridges for the carpeted enterprise, ruggedized environments, and challenging environments like the outdoors. Cisco Aironet lightweight access points are dynamically configured and managed through the Lightweight Access Point Protocol (LWAPP). Cisco Aironet Autonomous Access Points that have been converted to operate as lightweight access points running the LWAPP are supported.
- **Network unification:** The Cisco Unified Wireless Network uses the customer's existing wired network and investment in Cisco products. It supports a seamless network infrastructure across a range of platforms. Wired and wireless unification occurs with the Cisco Catalyst 6500 Series Wireless Services Module (WiSM), the 4400 Series and Cisco 2000 Series WLAN Controllers.



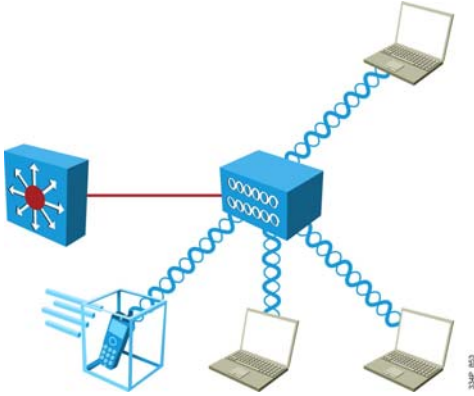
- **World-class network management:** Cisco delivers a world-class network management system (NMS) that visualizes and helps secure your air space. Cisco Wireless Control System (WCS) supports WLAN planning and design, RF management, location tracking, and IPS, in addition to WLAN systems configuration, monitoring, and management. This platform easily manages multiple controllers and their associated lightweight access points.
- **Unified advanced services:** Cisco provides unified support of leading-edge applications. Cisco offers advanced services that are industry leading, innovative, and comprehensive. Wireless lightweight access points, location appliances, and wireless IP phones deliver the Cisco Unified Wireless Network advanced services.

# Wireless LAN

This subtopic describes wireless LANs.

## Wireless LAN

- A WLAN is a shared network.
- An access point is a shared device that functions like a shared Ethernet hub.
- Data is transmitted over radio waves.
- Two-way radio communication (half-duplex) is used.
- The same radio frequency is used for transmitting and receiving (transceiving).



© 2009 Cisco Systems, Inc. All rights reserved. SWITCH-900

WLANs are similar to Ethernet networks in many ways. A WLAN is a shared network. An access point is a shared device that functions like a shared Ethernet hub. In the wireless cell, only one station can transmit at any time; all other stations listen. A station that wants to transmit must wait until the wireless media is not in use by another station.

This transmission setup is similar to that of a coaxial cable or half-duplex Ethernet and an Ethernet hub. Therefore, the performance of a wireless access point is similar to that of a hub. The average data rate per station is total bandwidth divided by the number of stations. The actual data throughput that is experienced by the wireless clients will be even less due to wireless-specific issues.


In WLANs, data is transmitted over radio waves. WLAN signals are similar to two-way radio communications. WLAN signals use the same frequency for transmitting and receiving (half-duplex). This setup means that a station that is transmitting cannot receive while it is transmitting. Therefore, only half-duplex transmission is possible. This transmission setup is similar to coaxial cable Ethernet.

# Wired and Wireless LAN

This topic describes WLANs and LANs.

## WLANs and LANs

A wireless LAN is an extension of the wired LAN.



Wired LANs require that users locate in one place and stay there. WLANs are an extension of the wired LAN network. They allow users to be mobile while using the mobile devices in different places without a wired network connection. A WLAN can be an overlay to, or substitute for, a traditional wired LAN network.

With Cisco Aironet WLANs, mobile users can do the following:

- Move freely around a facility
- Enjoy real-time access to the wired LAN at wired Ethernet speeds
- Access all the resources of wired LANs

## Similarities Between WLANs and LANs

This subtopic describes the similarities between WLANs and LANs.

### Similarities Between WLANs and LANs

- A WLAN is an 802.1X LAN.
  - Looks like a wired network to the user
  - Defines the physical and data-link layer
  - Uses a 48-bit MAC address (as with Ethernet)
- The same protocols and applications run over both WLANs and LANs.
  - IP (network layer)
  - IPsec VPNs (IP-based protocols)
  - Web, FTP, SMTP (IP-based applications)

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-98

WLANs are 802 LANs. The data in WLANs is sent over radio waves. In wired LANs, the data is sent over wires. However, the network interface of WLANs looks similar to that of wired LANs for the user.

Both WLANs and wired LANs define the physical and data-link layers and use MAC addresses. The same protocols and applications can be used over LANs and WLANs. Examples of such protocols are the IP and the IP Security (IPsec) protocol for virtual private networks (VPNs). Examples of applications are web, FTP, and Simple Network Management Protocol (SNMP) management.

# Differences Between WLANs and LANs

This subtopic describes the differences between WLANs and LANs.

## Differences Between WLANs and LANs

- WLANs use radio waves as the physical layer.
  - WLANs transmit data over the air instead of over wires.
  - WLANs use CSMA/CA instead of CSMA/CD to access the media.
- Radio waves have problems that are not encountered in wires.
  - Connectivity issues.
    - Coverage problems
    - Multipath issues
    - Interference, noise
  - Privacy issues.
- WLANs use mobile clients.
  - Battery-powered.
- WLANs must meet country-specific RF regulations.

Here is an explanation of how WLANs differ from LANs:

- In WLANs, radio frequencies are used as the physical layer of the network.
  - WLANs use Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) instead of Carrier Sense Multiple Access with Collision Detection (CSMA/CD), which is used by Ethernet LANs. Collision detection is not possible because a sending station cannot receive at the same time that it is transmitting and, therefore, cannot detect a collision. Instead, the Request To Send (RTS) and Clear to Send (CTS) protocols are used to avoid collisions.
  - WLANs use a different frame format from that of wired Ethernet LANs. Additional information for WLANs is required in the Layer 2 header of the frame.
- Radio waves have problems that are not found in wires.
  - Connectivity issues in WLANs can be caused by coverage problems, RF transmission, multipath distortion, and interference from other wireless services or other WLANs.
  - Privacy issues are possible because radio frequencies can reach outside the facility.
- In WLANs, mobile clients are used to connect to the network.
  - Mobile devices are often battery-powered.
- WLANs must meet country-specific RF regulations.
  - The aim of standardization is to make WLANs available worldwide. Because WLANs use radio frequencies, they must follow country-specific regulations for RF power and frequencies. This requirement does not apply to wired LANs.

## Summary of WLAN and LAN Differences

This subtopic summarizes the differences between WLANs and LANs.

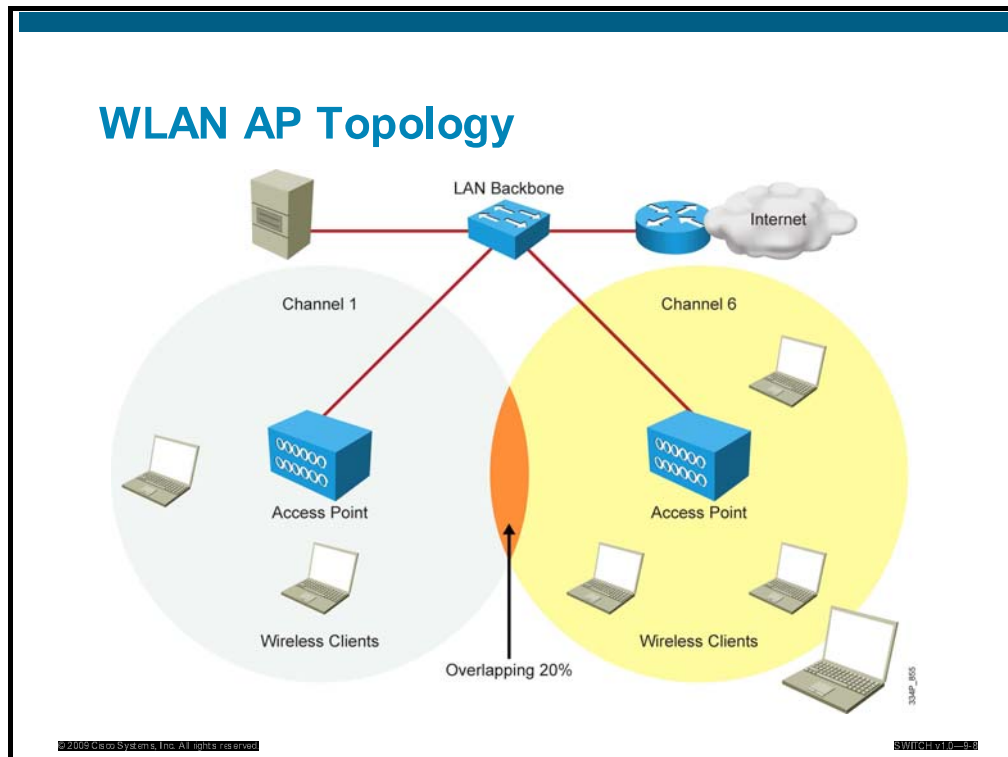
Summary of Differences Between WLANs and LANs		
	LAN	WLAN
Physical layer	Wire or fiber	RF
Data-link layer	802.3 Ethernet	802.11 WLAN
Media access	CSMA/CD	CSMA/CA
Transmission	Full-duplex	Half-duplex
Issues	Cable length	Distance, coverage, interference, power

© 2009 Cisco Systems, Inc. All rights reserved. SWITCH-99

This figure summarizes the differences between wired and wireless LANs. Because of the differences in the medium and the possible performances, the WLAN designs will be very different from traditional LANs designs.

# WLAN Access Point Topology

This topic describes the WLAN access point topology.



WLANs replace the Layer 1 transmission medium of a traditional wired network (usually Category 5 cable) with radio transmission over the air. Cisco Aironet WLAN products can be plugged into a wired network and can function as an overlay to traditional or wired LANs, or they can be deployed as standalone LANs where wired networking is not feasible. WLANs permit the use of desktop and portable computers or specialty devices in a system where connection to the network is essential.

A computer with a wireless network interface card (NIC) can connect to the wired LAN through the access point. Properly deployed WLANs can provide instant access to the network from anywhere in the facility. Users can roam without losing their network connection.

The basic service area is the area of RF coverage that is provided by an access point. To extend the basic service area, or to simply add wireless devices and extend the range of an existing wired system, you can add an access point. As the term “access point” indicates, this device is the point at which wireless clients can access the network.

The access point (AP) attaches to the Ethernet backbone and communicates with all the wireless devices in the cell area. The AP is the master for the cell and controls traffic flow to and from the network. The remote devices do not communicate directly with each other; they communicate with the AP.

If a single cell does not provide enough coverage, any number of cells can be added to extend the range. This range is known as an extended service area.

It is recommended that the extended service area cells have 20 percent overlap to allow remote users to roam without losing RF connections. Bordering cells should be set to different nonoverlapping channels for best performance.

More recently, wireless deployments have moved from microcell to picocell. Pico cells further reduce the AP coverage area by reducing power and increasing the total number of APs that are deployed.

It is important that not only the APs can reduce their transmit power settings, but also that the clients can reduce their transmit power. Both APs and clients should use a comparable transmit power so that the client associates with the nearest AP.

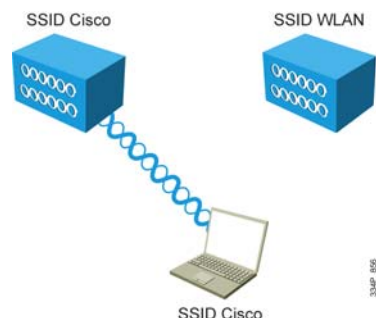


# About SSIDs

This subtopic describes the Service Set Identifier.

## About SSIDs

- An SSID (network name) is used to logically separate WLANs.
- APs are configured with SSIDs.
- An AP broadcasts the SSID.
- An SSID must match on the client and AP.
- A client can be configured without an SSID (hotspot mode).



© 2009 Cisco Systems, Inc. All rights reserved.

SWITCH-010-001

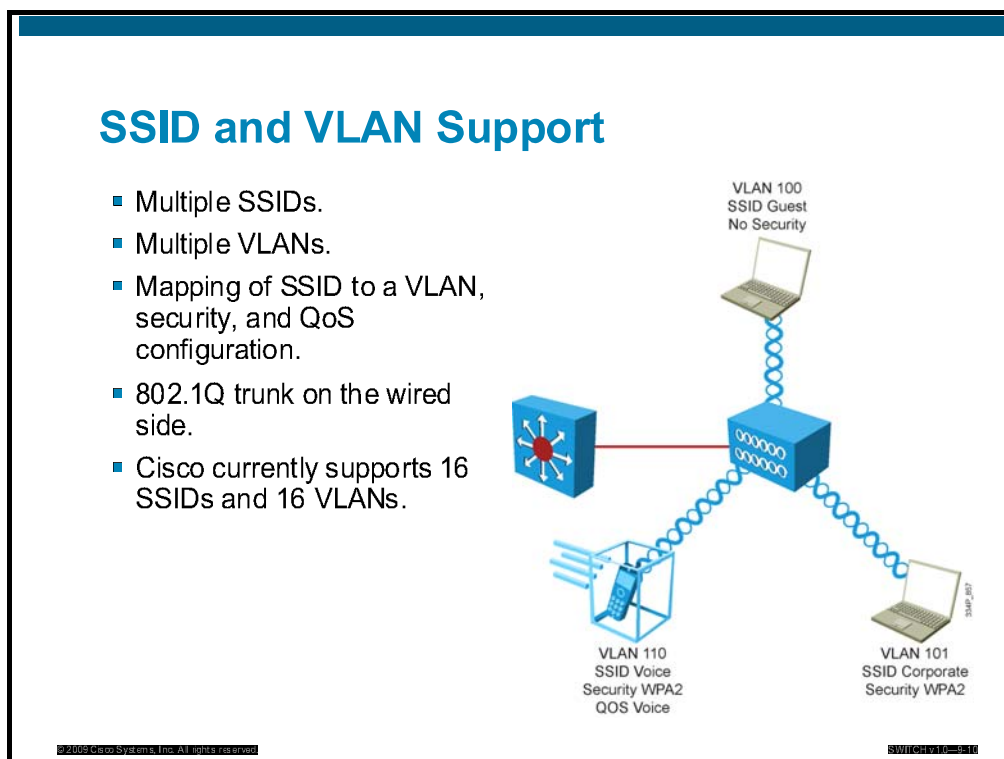
The Service Set Identifier (SSID) is the name of the wireless cell. It is used to logically separate WLANs. It must match exactly between the client and the AP.

The AP broadcasts the SSID in the beacons. Beacons are broadcasts that the APs send to announce the available services. Clients can be configured without an SSID (null SSID), detect all APs, and learn the SSID from the beacons of the AP.

SSID broadcasts can be disabled on the AP, but this approach does not work if the client needs to see the SSID in the beacon.

# SSID and VLAN Support

This subtopic describes SSID and VLAN support for WLANs.



LAN networks are increasingly being divided into workgroups that are connected via common backbones to form VLAN topologies. VLANs enable efficient traffic separation, provide better bandwidth utilization, and alleviate scaling issues by logically segmenting the physical LAN infrastructure into different subnets so that packets are switched only between ports within the same VLAN. When combined with central configuration management support, VLANs facilitate workgroups and client/server additions and changes.

Switches use VLANs to separate traffic. APs can extend VLANs to the wireless LAN by mapping VLANs to SSIDs. The wireless VLANs share the same wireless cell and channel. The result is a virtualization of the APs. The AP appears as multiple different APs.

- **VLAN 100:** Allows guests to connect directly to the Internet without having access to your enterprise servers. Without the VLAN function, two APs would be needed to provide isolated connectivity for the guest users and enterprise users.

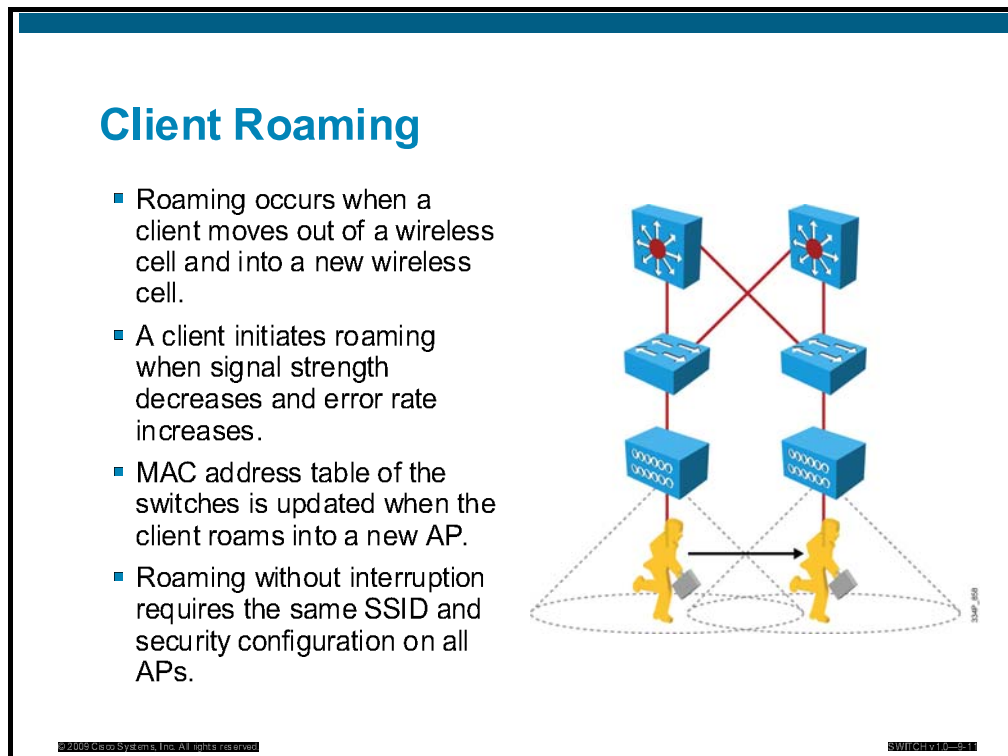
VLAN 100 would be configured with no security and would broadcast its SSID. An access control list (ACL) on the router could also be configured to ensure that traffic with VLAN 100 goes straight out the firewall.

- **VLAN 101:** Allows corporate users. This SSID is configured with Wi-Fi Protected Access 2 (WPA2) security.
- **VLAN 110:** Allows the enterprise voice applications. WPA2 security and quality of service (QoS) for voice is mapped to VLAN 110. This SSID is used by wireless IP phones.

The Cisco Aironet APs and WLAN controllers use the IEEE 802.1Q trunking protocol to connect to switches.

# Client Roaming

This subtopic describes client roaming.



Wireless clients associate with another AP if necessary. This process is called roaming between the wireless cells. The wireless client initiates the roaming if one of these conditions is detected:

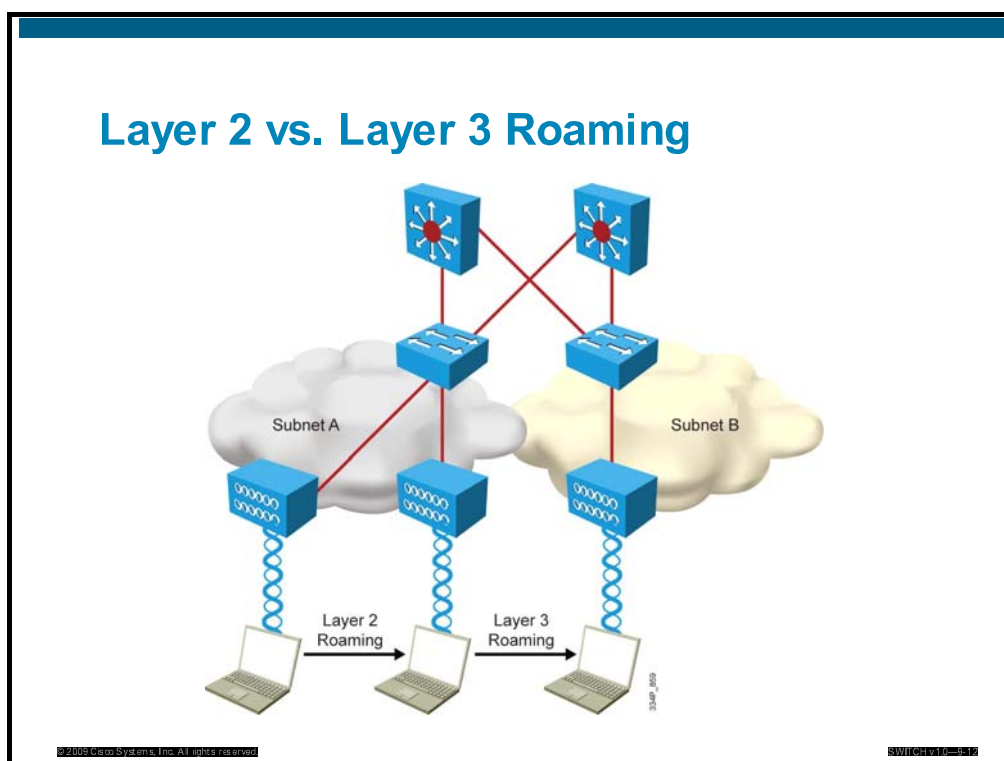
- The maximum data retry count is exceeded.
- The client has missed too many beacons from the AP.
- The client has reduced the data rate.
- The client intends to search for a new AP at periodic intervals.

Roaming without service interruption requires the identical configuration of SSID, VLANs, and IP subnets on all APs. Roaming is initiated by the client. The client searches for another AP with the same SSID and sends a reauthentication request to the new AP.

A short roaming time is important for delay-sensitive applications, such as voice and video.

## Layer 2 and Layer 3 Roaming

This subtopic describes Layer 2 and Layer 3 roaming.



Roaming maintains network connectivity while moving from one AP to another. Roaming between APs that reside on a single IP subnet (or VLAN) is considered Layer 2 (data-link layer) roaming. Roaming between APs that reside in different IP subnets is considered Layer 3 (network layer) roaming.

Roaming at Layer 2 is managed by the APs, using a combination of multicast packets that inform the switches in the network that the device has moved. The protocol between the APs is called the Inter-Access Point Protocol (IAPP).

Mobile IP is a technology that allows fixed IP addresses in an IP subnet of a network. It relies on devices such as routers, known as home agents and foreign agents, to tunnel traffic for a mobile device.

WLAN implementations allow Layer 3 roaming. Legacy Layer 3 roaming for WLANs was accomplished by Mobile IP, which has been replaced by the advanced feature set of lightweight APs in combination with WLAN controllers.

# Security on WLAN and LAN

This subtopic describes security on WLANs and LANs.

## Security on WLANs and LANs

- Data can be encrypted on the wireless link (WLAN).
- Data is unencrypted on wired link (LAN).
- Open wireless networks allow access and attacks to the wired network.
- Solutions:
  - Implement authentication to control access to the wireless network.
  - Encrypt data on the wireless link.
  - Implement firewall, IPS, and NAC to secure access to the network.
  - Use a VPN if encryption on the wired network is required.

Wireless LAN security is limited to the wireless space, which presents the following challenges:

- Data is encrypted on the wireless link (WLAN).
- Data is unencrypted on the wired link (LAN).
- Open wireless networks allow access and attacks to the wired network.

Solutions are as follows:

- Implement authentication to control access to the wireless network.
- Encrypt data on the wireless link.
- Implement firewall, intrusion prevention system (IPS), and Network Admission Control (NAC) to secure access to the network.
- Use VPN if encryption on the wired network is required.

# Summary

This topic summarizes the key points that were discussed in this lesson.

## Summary

- Wireless LANs are networks shared between mobile stations.
- WLANs and wired LANs have similarities and differences when compared with wired LANs.
- WLANs map SSIDs to VLANs and QoS, and support roaming of mobile stations between APs.
- WLAN traffic can be secured by WLAN standards on the wireless link but is unprotected on the wired network unless a VPN is used.

# Assessing the Impact of WLANs on Campus Networks

---

## Overview

Integrating wireless into the LAN network is a lot more than just merely connecting access points (APs) to access switches. In most implementations, APs will rely on wireless LAN controllers (WLCs), and the traffic flow from wireless clients will not stop at the AP level. This lesson describes the various possible WLAN implementations and their impact on the campus LAN infrastructure operations.

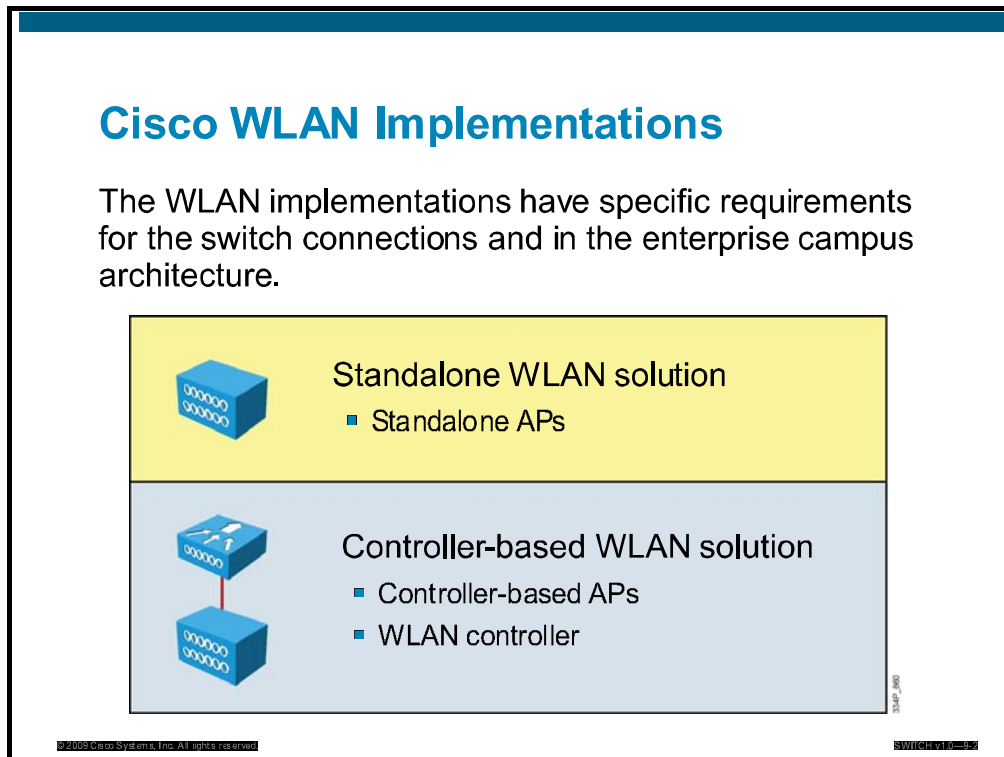
## Objectives

Upon completing this lesson, you will be able to assess the impact of WLANs on the campus infrastructure. This ability includes being able to meet these objectives:

- Describe WLAN implementations
- Compare WLAN solutions
- Assess traffic flow in an autonomous AP configuration and its impact on the campus LAN
- Assess traffic flow in a controller-based configuration and its impact on the campus LAN

# WLAN Implementations

This topic describes Cisco WLAN implementations



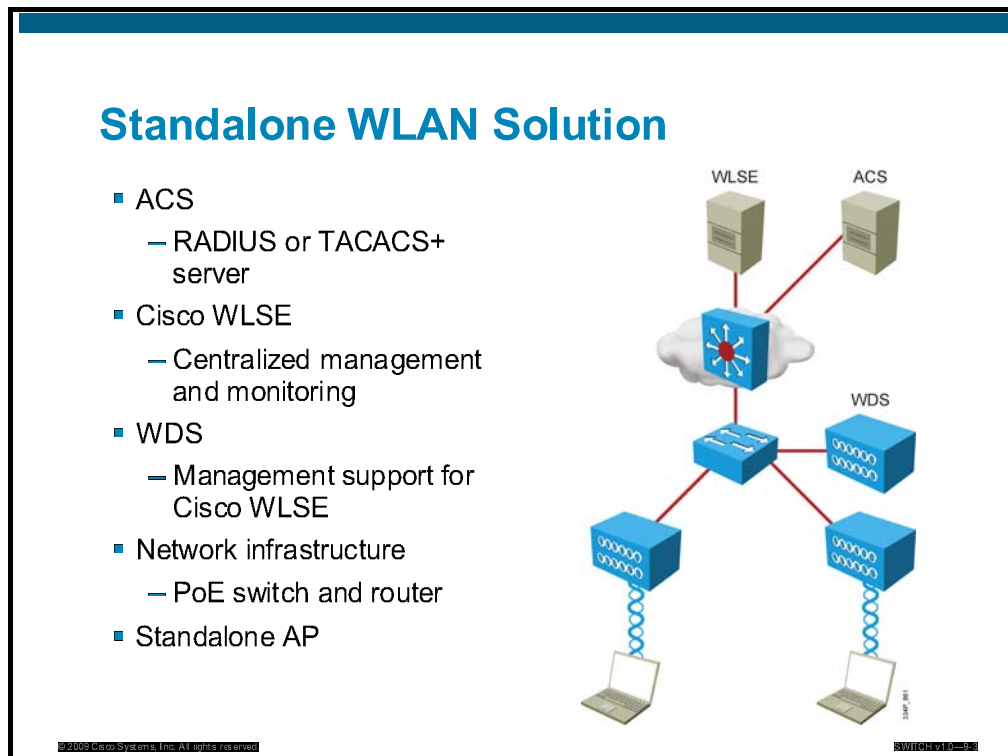
Cisco offers two WLAN implementations:

- The standalone WLAN solution is based on standalone APs.
- The controller-based WLAN solution is based on controller-based APs and WLCs.



# Standalone WLAN Solution

This subtopic describes the standalone WLAN solution.

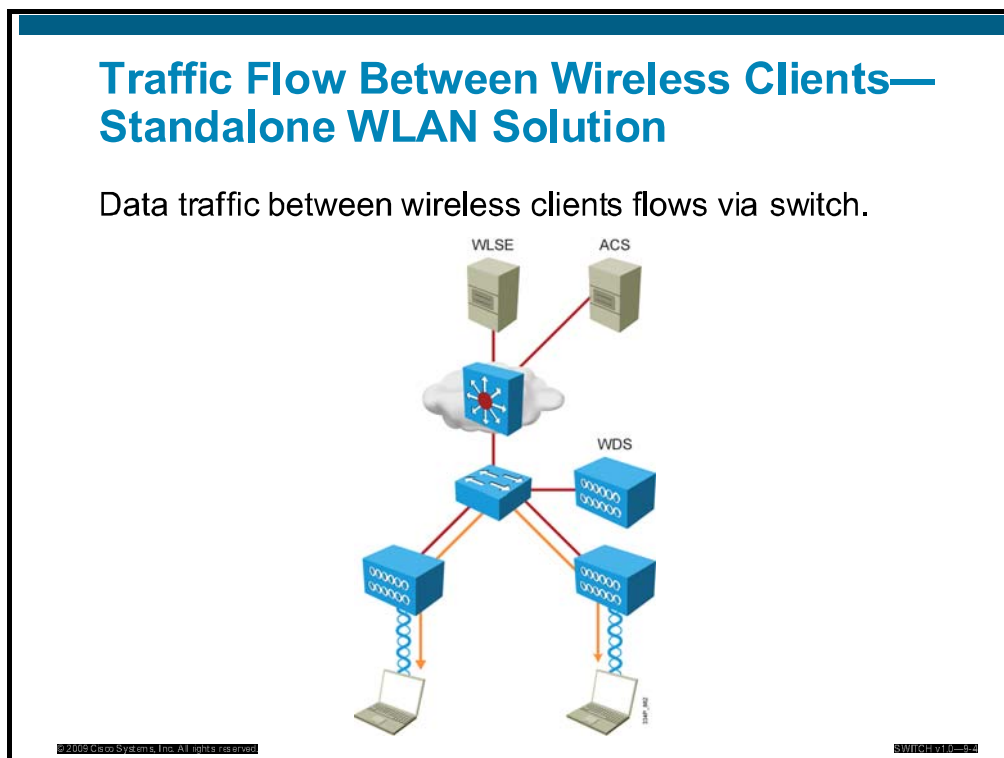


The figure shows the components of the distributed WLAN solution:

- Standalone APs using Cisco IOS Software.
- Network infrastructure with router and switches. Switches can be used to supply power to the APs (Power over Ethernet, or PoE).
- Wireless Domain Services (WDS) for RF management and fast, secure roaming.
- CiscoWorks Wireless LAN Solution Engine (CiscoWorks WLSE) for WLAN management
- Cisco Secure Access Control Server (Cisco Secure ACS) for wireless security and authentication using the RADIUS and TACACS+ protocols.

## Traffic Flow Between Clients—Standalone WLAN Solution

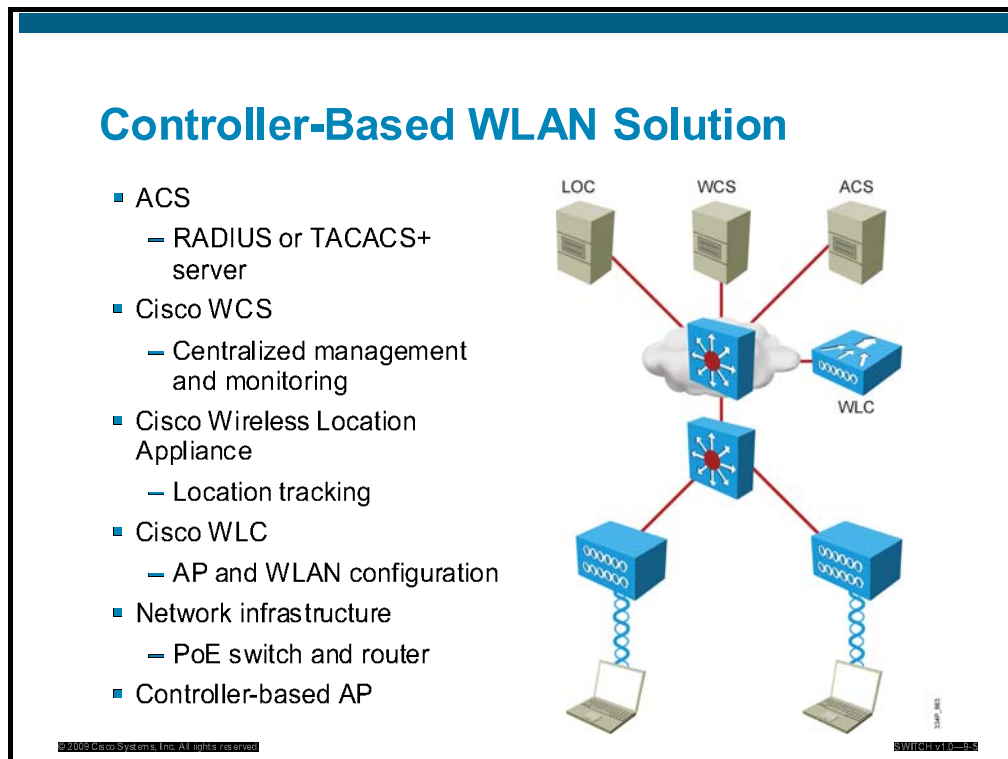
This subtopic describes the standalone WLAN solution to traffic flow between clients.



In the autonomous (or standalone) solution, each AP operates independently and acts as a transition point between the wireless media and the 802.3 media. The data traffic between two clients flows via the Layer 2 switch when on the same subnet from a different AP infrastructure. As the AP converts the IEEE 802.11 frame into an 802.3 frame, the wireless client MAC address is transferred to the 802.3 header and appears as the source for the switch. The destination, also a wireless client, appears as the destination MAC address. For the switch, the APs are relatively transparent.

# Controller-Based WLAN Solution

This subtopic describes the controller-based WLAN solution.




The figure shows the components of the controller-based WLAN solution:

- Controller-based access points.
- Network infrastructure with router and switches. Switches can be used to supply power to the APs (PoE).
- Cisco Wireless LAN Controller (Cisco WLC) for the control, configuration, and monitoring of the APs.
- Cisco Wireless Control System (Cisco WCS) for WLAN management (recommended).
- Cisco Wireless Location Appliance for location tracking (optional application server).
- Cisco Secure Access Control Server (Cisco Secure ACS) for wireless security and authentication using the RADIUS and TACACS+ protocols.

# Controller-Based WLAN Solution

This subtopic describes the wireless infrastructure based on controllers.



The controller-based architecture splits the processing of the 802.11 protocol between two devices: the AP and a centralized Cisco WLC. The processing of the 802.11 data and management protocols and the AP functionality is also divided between the two devices. This approach is called split MAC.

The AP handles the portions of the protocol that have real-time requirements:

- The frame exchange of a handshake between a client and AP when transferring a frame over the air
- The transmission of beacon frames
- The buffering and transmission of frames for clients in a power save operation
- The response to probe request frames from clients
- Forwarding notification of received probe requests to the controller
- Providing real-time signal quality information to the controller with every received frame
- Monitoring each radio channel for noise, interference, and the presence of other WLANs
- Monitoring for the presence of other APs

All remaining functionality is handled in the Cisco WLC, where time sensitivity is not a concern and controller-wide visibility is required.

- 802.11 authentication
- 802.11 association and reassociation (mobility)
- 802.11 frame translation and bridging

Data and control messages are encapsulated between the AP and the Cisco WLC using Control and Provisioning of Wireless Access Points (CAPWAP) or Lightweight Access Point Protocol (LWAPP).

Control traffic between the AP and the controller is encapsulated with the LWAPP or CAPWAP and is encrypted.

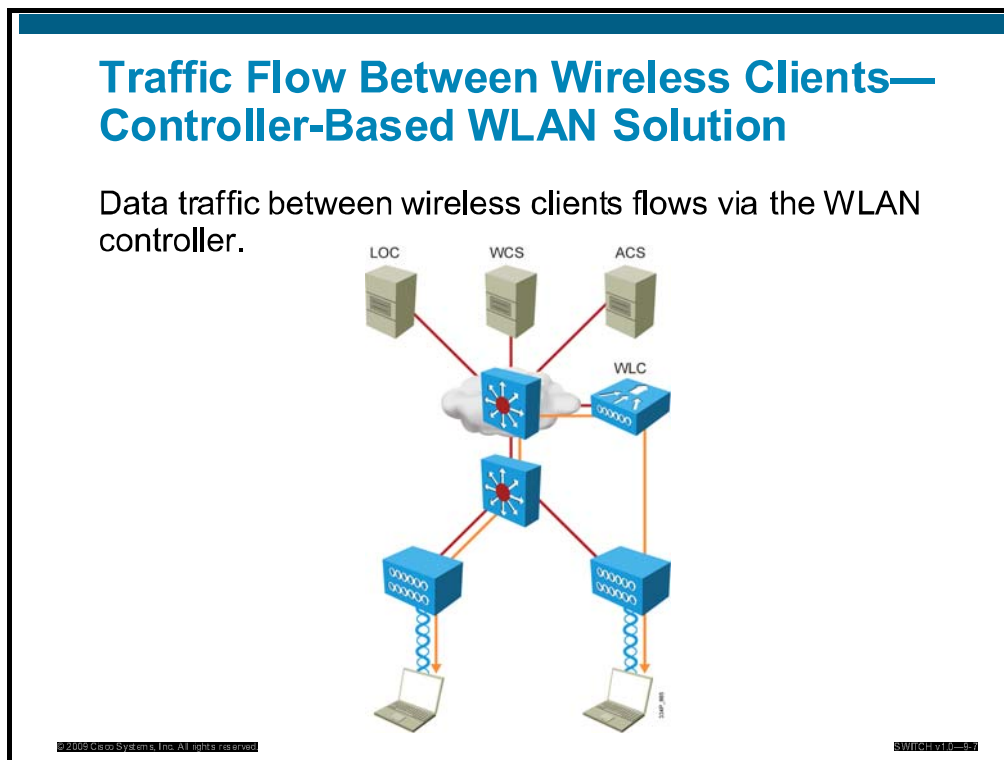
The data traffic between the AP and the controller is also encapsulated with LWAPP or CAPWAP. The data traffic is not encrypted. It is switched at the Cisco WLC, where VLAN tagging and quality of service (QoS) are also applied.

The AP accomplishes real-time frame exchange and certain real-time portions of MAC management. All client data traffic is sent via the Cisco WLC.

The WLC and the AP can be in the same or different broadcast domains and IP subnets. APs obtain an IP address via DHCP and then join a controller via the CAPWAP or LWAPP discovery mechanism.

## Traffic Flow Between Clients—Controller-Based WLAN Solution

This subtopic describes the controller-based WLAN solution to traffic flow between clients.

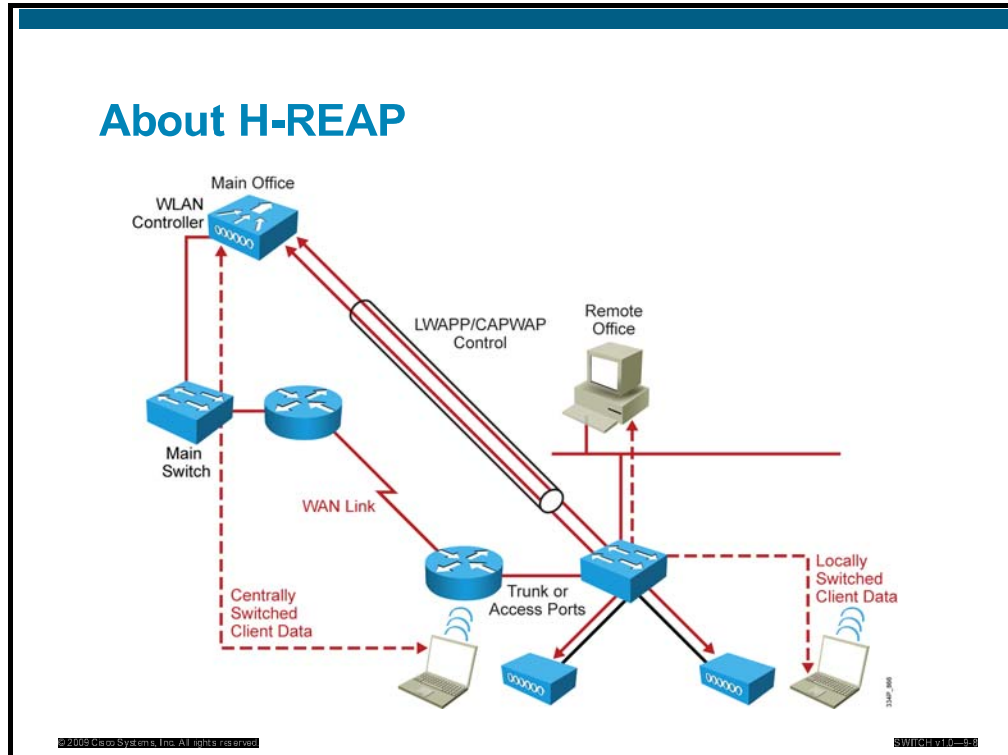


The traffic to and from the wireless clients always goes via the wireless LAN controller.

Traffic between the two wireless mobile stations is forwarded from the APs to the controller and then is sent to the mobile destination. For the infrastructure switch, there is traffic between two devices: the access point and the controller. The actual data information that is conveyed by this traffic is encapsulated and is therefore not directly visible in the outer header.

# Hybrid Remote Edge Access Points

This subtopic describes the specific case of the Hybrid Remote-Edge Access Point (H-REAP).



H-REAPs are controller-based APs that operate in a specific mode. The purpose of these APs is to continue to offer client connectivity when the connection to the controller is lost. A standard controller-based AP always needs connectivity to its controller. When this connectivity is lost, the AP stops offering wireless services and starts looking for a controller. It will restore wireless services only when it regains connectivity to its controller.

The H-REAP is a controller-based AP. It needs to reach a controller and register to it to offer wireless service. The main difference between a standard controller-based AP and the H-REAP is that the H-REAP can survive the loss of connectivity to its controller. The H-REAP then reverts to a mode that is close to autonomous APs and that still offers wireless access to its clients.

The H-REAP is adapted to remote offices where only one or two APs are to be deployed. The controller is not local but is accessed through the WAN. The H-REAP is also adapted to small offices with one controller and no backup controller.

Although the H-REAP is a controller-based solution, the switch port has to be configured the same way as for an autonomous AP.

# Comparison of the WLAN Solutions

This topic compares the WLAN solutions.

Comparison of WLAN Solutions		
	Standalone	Controller-Based
APs	Standalone Cisco IOS Software	Controller-based Cisco IOS Software
Configuration	AP	WLAN controller
Operation	Independent	Dependent on Cisco WLC
Management and monitoring	CiscoWorks WLSE	Cisco WCS
Redundancy	AP	AP WLAN controller

The two WLAN solutions have different characteristics and advantages. The APs use different Cisco IOS feature sets.

Standalone APs are configured per access point. Their Cisco IOS Software operates independently. You can perform centralized configuration, monitoring, and management via the CiscoWorks WLSE management system. You can install standalone APs with redundant APs.

Controller-based APs are configured via the WLC. They depend on the WLC for control and data transmission. The Cisco IOS Software on the controller-based AP is installed automatically by the WLC.

Monitoring and security is implemented by the WLC. You can perform centralized configuration, monitoring, and management via the Cisco WCS management system. You can install WLCs with redundant APs and controllers.



# About WLCs

This subtopic describes the types of WLAN controllers.


## Types of WLAN Controllers

### Appliance controllers

- 2100 Series
- 4400 Series
- 5500 Series

### Integrated controllers

- WLAN controller module for Cisco
- Cisco ISRs
- Catalyst 3750G Integrated Wireless LAN Controller
- Catalyst 6500 Series WiSM module for Catalyst 6500 Series Switches



© 2009 Cisco Systems, Inc. All rights reserved.SWITCH10-0511

Depending on the size of the campus and whether integration with Layer 3 infrastructure devices is desired, one of two categories of Cisco WLCs is typically deployed.

Cisco 5500 Series Wireless LAN Controllers support up to 250 access ports and 8-Gigabit Ethernet network ports. The Cisco 4400 Series Wireless LAN Controllers support between 12 and 100 APs.

The 4400 Series controllers are designed for medium to large facilities and can be used to support from 12 to 100 APs. These controllers can support from 40 to 2000 wireless mobile stations, depending on the mix of data and voice clients.

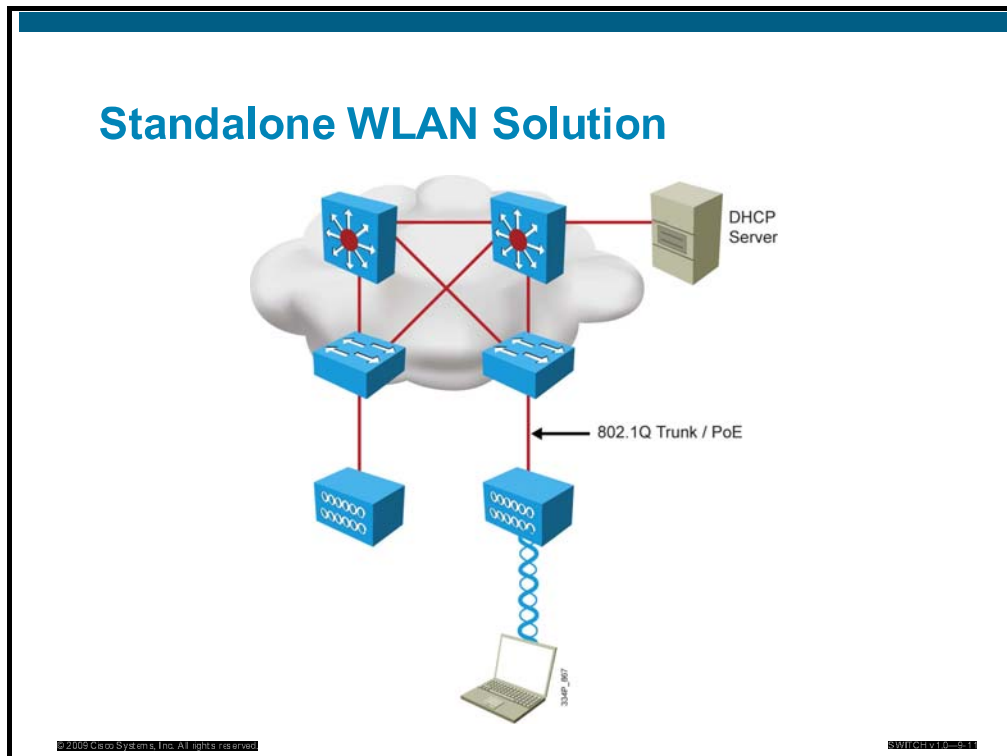
The Cisco 2100 Series Wireless LAN Controllers deliver WLAN services to small and medium-sized enterprise environments. It supports from 6 to 25 controller-based APs, making it a cost-effective solution for smaller buildings and branch offices within a distributed enterprise.

Controllers that are integrated into Layer 3 devices, such as the Cisco Catalyst 3750G Integrated Wireless LAN Controller or the Cisco Catalyst 6500 Series Wireless Services Module (Cisco Catalyst 6500 Series WiSM), support from 25 to 300 APs. The integrated controllers support Layer 2 connections internally and can use Layer 2 or Layer 3 connections to the wired enterprise network.

WLCs are also available for the Cisco Catalyst 6500 Series Switches and for Cisco Integrated Services Routers (Cisco ISRs).

# Connection of the Standalone Solution to the Network

This subtopic describes the connection of the standalone solution to the network.



The standalone APs are connected to trunk ports on switches with PoE. Management and data VLANs are connected via a trunk port to the standalone APs. The native VLAN is required to be used for management of the standalone AP.


The DHCP server provides IP addresses for the APs and the wireless clients.

# SSIDs, VLANs, and Trunks in the Standalone Solution

This subtopic describes Service Set Identifiers (SSIDs), VLANs, and trunks in the standalone solution.

## SSIDs, VLANs, and Trunks in the Standalone Solution

- Mapping of SSID, VLAN, and subnet at standalone AP.
- The client becomes a station within a VLAN connected to the AP.
- The client gets an IP address from a VLAN or subnet connected to the AP.
- The same VLANs or subnets on all APs.
- Layer 2 connection between APs.
- Layer 2 roaming only.

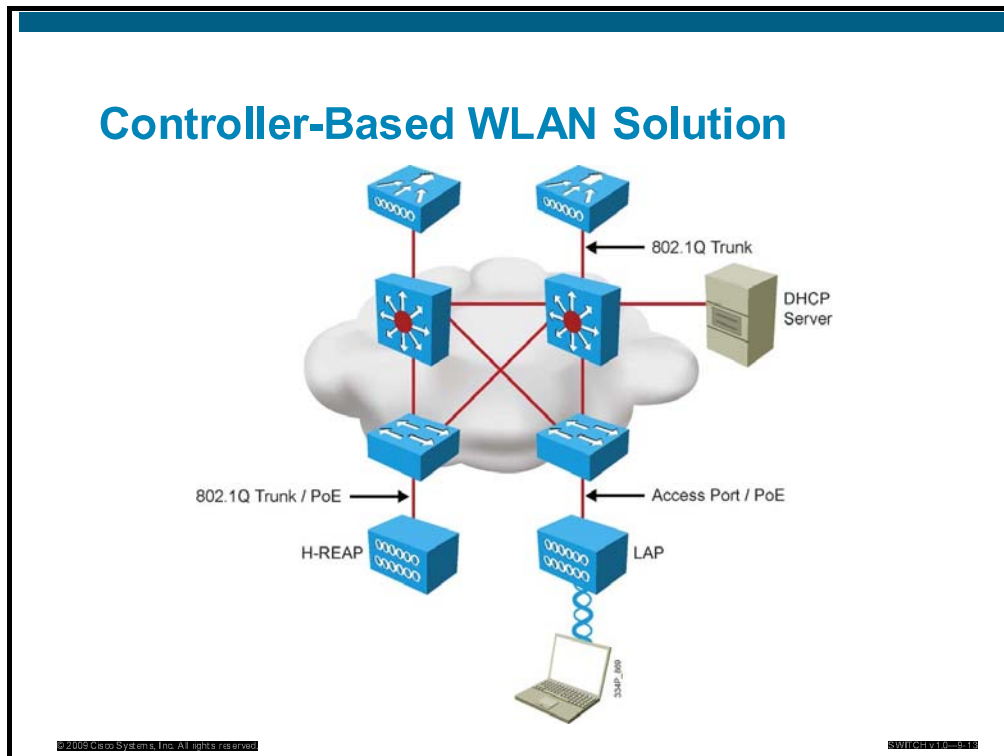


In the autonomous AP configuration, the AP is the transition point between the wireless side and the wired side. SSIDs are mapped to VLANs at the AP level. Upon associating with one of the available SSIDs, the wireless client becomes a station within a VLAN or subnet that is connected to the AP, and gets an IP address from a VLAN or subnet that is associated with the SSID within the AP configuration.

The AP connects to the switch through a trunk if different SSIDs and VLANs are configured. In this scenario, roaming between APs is possible as long as both APs offer the same SSID and are connected to the same Layer 2 network. Layer 3 roaming is not possible.

# Connection of the Controller-Based Solution to the Network

This subtopic describes the connection of the controller-based solution to the network.



The wireless LAN controller is connected to trunk ports on switches providing communication for management and data VLANs. The native VLAN is not required.

The controller-based APs are connected to access ports on switches with PoE. Only the AP VLAN is connected to the controller-based APs. The untagged access VLAN is used as the AP VLAN.

The AP VLAN is not required to connect to the controller, because the controller and the APs can be in different IP subnets. The AP can be in any IP subnet.

The H-REAP is connected to an 802.1Q trunk, just like an autonomous AP. The port native VLAN is used by the H-REAP to join its controller.


The DHCP server provides IP addresses for the APs and the wireless clients.

# SSIDs, VLANs, and Trunks in the Controller-Based Solution

This subtopic describes SSIDs, VLANs, and trunks in the controller-based solution.

## SSIDs, VLANs, and Trunks in the Controller-Based Solution

- Mapping of an SSID, VLAN, and subnet at the WLAN controller.
- The client becomes a station within a VLAN or subnet connected to the WLAN controller.
- The client gets an IP address from a VLAN or subnet connected to the WLAN controller.
- Any VLAN or subnet can be connected to the APs.
- APs and WLAN controller can be on same or different subnet.
- Layer 3 IP connection between APs and WLAN controller.
- Layer 2 and Layer 3 roaming are supported via WLAN controller.




In a controller-based configuration, the WLC maps the SSID, VLAN, QoS, and IP subnet. When a client associates with an SSID on an AP, the client becomes a station within a VLAN or subnet that is connected to the WLC, and gets an IP address from the VLAN or subnet that is connected to the WLC. The IP address is mapped to the SSID that is used by the client.

Any VLAN or subnet can be connected to the APs. All traffic that is arriving at the AP is encapsulated and is sent to the controller. Therefore, a clear differentiation is made between the client VLAN and the AP VLAN. APs and the WLC can be on the same or a different IP subnet, and there can be a Layer 3 IP connection between APs and the WLC.

In this configuration, Layer 2 and Layer 3 roaming are supported by the Cisco WLC, whether all the APs connect to the same controller or to different controllers that are part of the same group.

## SSIDs, VLANs, and Trunks with the H-REAP

- AP needs to connect to the WLC.
- Some WLANs are locally switched.
- Some WLANs are centrally switched.
- Trunk needs to allow locally switched VLANs.
- Native VLAN is the AP VLAN.



In the case of an H-REAP, some WLANs are centrally switched, which means that data for these WLANs is encapsulated into LWAPP or CAPWAP and are sent to the controller, just like a standard controller-based AP.

Some other WLANs are locally switched, which means that traffic is sent to the switch that is local to the H-REAP, and is not sent to the controller.

The consequence of this behavior is that the port to an H-REAP has to be an 802.1Q trunk. The native VLAN is the H-REAP VLAN, which is used to communicate with the controller.


On this 802.1Q trunk, the list of allowed VLANs matches the list of locally switched WLANs. Work with your wireless specialist to determine which VLANs should be locally switched.

# Controller-Based AP Protocol

This subtopic describes the controller-based AP protocol.

## Controller-Based AP Protocol

- Is the protocol between AP and WLC
- Has a split-MAC architecture
- AP
  - RF interface (radio)
  - MAC layer encryption
- WLC
  - Security policies
  - QoS policies
  - RF management
  - Mobility management



The diagram shows the traffic flow between a wireless client and a server on the network via an AP, a CAPWAP or LWAPP tunnel, and the WLC.

The CAPWAP or LWAPP tunnel connects the AP and the WLC.

The split-MAC architecture describes the task distribution between the AP and the controller.

The AP handles the radio interfaces and all receiving and transmission of wireless frames. Additionally, the AP does the RF monitoring and the MAC layer encryption and decryption of client data traffic.

The WLC is responsible for all security policies, QoS policies, RF management, client association, and mobility management.

# Cisco WLC Ports and Protocols

This subtopic describes ports and protocols that are used by the WLC.

## WLC Ports and Protocols

- The WLC uses these ports and protocols for communication with APs and management.
- These ports and protocols must be allowed in the ACLs and firewalls.
- Other ports and protocols may be used in the future.

<b>CAPWAP</b>	UDP 5246 UDP 5247
<b>LWAPP</b>	UDP 12222 UDP 12223
<b>HTTPS</b>	TCP 443
<b>SSH</b>	TCP 22
<b>RADIUS</b>	UDP 1812 UDP 1813
<b>SNMP</b>	UDP 161 UDP 162
<b>Mobility</b>	UDP 16666 UDP 16667 EoIP protocol

© 2009 Cisco Systems, Inc. All rights reserved.SWITCH v1.0-6-17

Access lists and firewalls on the network are required to allow the traffic between controllers, APs, and management stations for the successful operation of a wireless network.

CAPWAP and LWAPP require User Datagram Protocol (UDP) traffic between APs and the WLC.

The mobility traffic between WLCs requires UDP traffic and the Ethernet over IP (EoIP) protocol.

Secure Shell (SSH), Secure Sockets Layer (SSL) (HTTPS), and Simple Network Management Protocol (SNMP) are used to manage the WLC.



# Summary

This topic summarizes the key points that were discussed in this lesson.

## Summary

- Standalone and controller-based WLAN solutions are the Cisco implementations of WLAN.
- In the standalone configuration, the AP performs autonomously while in the controller-based solution; the AP receives its configuration from a central controller.
- In the standalone solution, the AP is a direct conversion point between the 802.11 environment and the 802.3 campus network.
- In the controller-based solution, the AP sends all traffic to the WLC, where the decision is made about how to forward the traffic to the LAN.



# Preparing the Campus Infrastructure for WLANs

---

## Overview

WLAN integration into the campus network is very different depending on whether an autonomous solution or a controller-based solution is expected. Device placements, port configurations, VLAN design, and configuration will depend heavily on the type of deployment. This lesson explains how to create implementation and verification plans for preparing infrastructure devices to integrate wireless LANs (WLANs), and how to configure the campus network accordingly.

## Objectives

Upon completing this lesson, you will be able to prepare the campus infrastructure for WLAN integration. This ability includes being able to meet these objectives:

- Decide on the best placement for APs and controllers
- Configure switches for WLAN devices
- Gather WLAN requirements
- Plan WLAN integrationCreate a test plan

# Access Point and Controller Placement

This topic describes access point (AP) and controller placement.

## AP and Controller Placement

- The APs are connected to access switches.
- The WLC is connected to the network at:
  - Distribution switches
  - Server farm or data center
- Centralized deployment is recommended.
- Minimize intercontroller roaming.
- Implement deterministic redundancy.
- Centralized deployment with the integrated platforms:
  - Catalyst 3750G Integrated Wireless LAN Controller for small-to-medium deployments
  - Catalyst 6500 Series WiSM for medium-to-large deployments
- Distributed deployment can be an alternative for existing networks.

As much as possible, controllers should be placed so that they minimize intercontroller roaming and latency of traffic flow over the wireless media. Lightweight Access Point Protocol (LWAPP) tunneling separates the physical controller placement from the subnets, so the wireless LAN controllers (WLCs) can be positioned where they are connected, secured, and powered and where traffic flows work well. When deploying controllers, you should use deterministic redundancy to avoid unnecessary intercontroller roaming that results from salt-and-pepper designs.

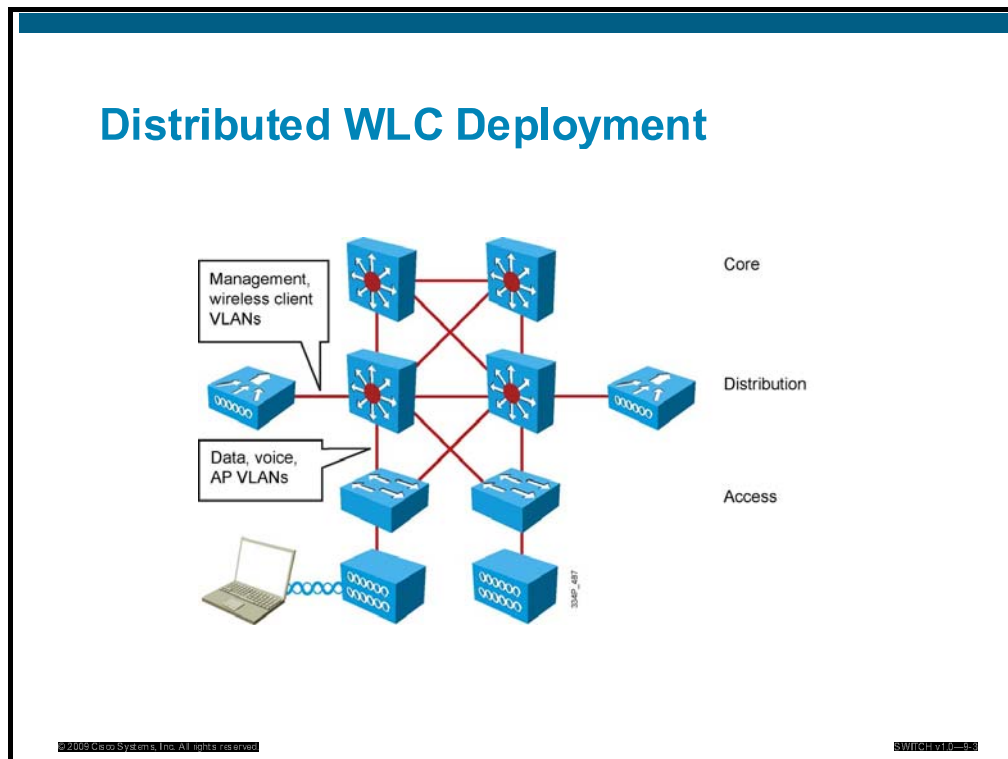
In the enterprise campus, centralized design supports the integrated controller platforms. Depending on the campus size and existing network devices, you can use the Cisco Catalyst 3750G Integrated Wireless LAN Controller for small-to-medium deployments or for an individual building, or the Cisco Catalyst 6500 Series Wireless Services Module (Catalyst 6500 Series WiSM) for medium-to-large deployments.

Distributed controller deployment may work well with existing networks or focused wireless coverage areas.

The general recommendation is to use a centralized design for controller placement to minimize operational complexity and support. However, this decision should be based on the ability of either design to support the current network and policies, as well as on plans for growth.

# Distributed Controller Placement

This subtopic describes distributed controller placement.



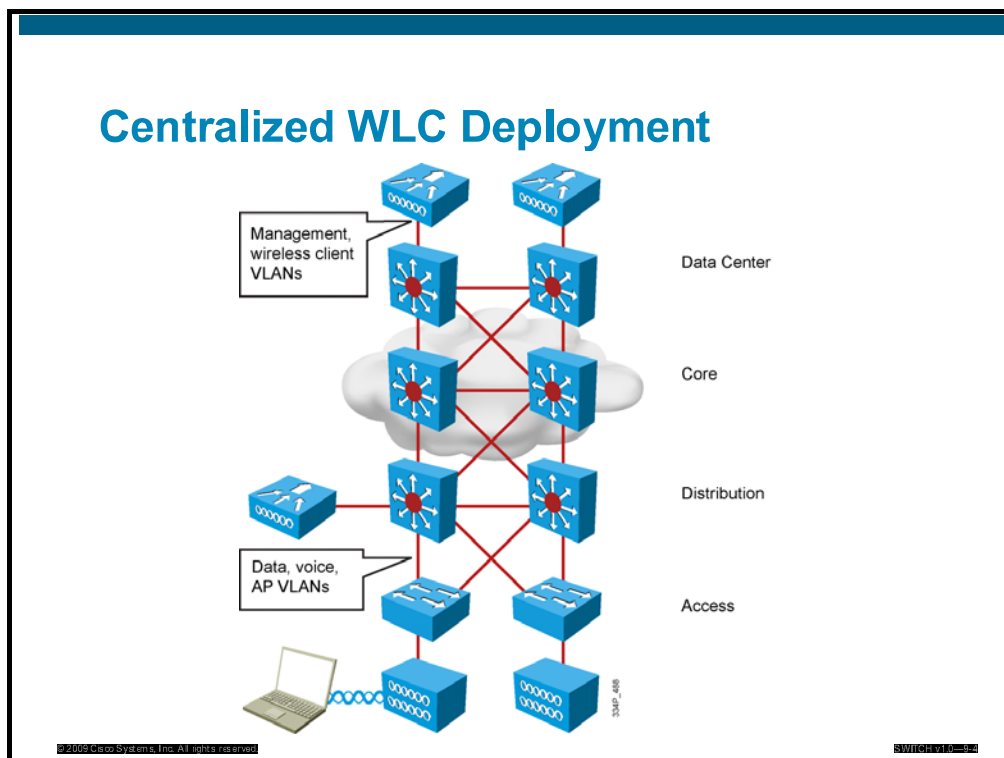
## Distributed WLC Design

The figure illustrates a distributed WLC design with the placement of APs in the access layer and WLCs in the distribution layer.

The distributed WLC design is well adapted for designs where coverage areas are isolated, and where mobility between covered areas is not implemented.

## Centralized Controller Placement

This subtopic describes centralized controller placement.



## Centralized WLC Placement

The figure illustrates a centralized WLC design with the placement of APs in the access layer and WLCs in a service block in the core layer.

The centralized WLC design supports simplified management with fewer endpoints and fewer locations to manage issues such as high availability, routing, and power needs. Centralized WLC design also supports the most efficient mobility.

# WLAN Devices Connected to the LAN Switches

This subtopic describes WLAN devices that are connected to the LAN switches.

## WLAN Devices Connected to Switches

### Standalone WLAN solution

- Standalone AP

### Controller-based solution

- Controller-based AP
- 2100 Series Wireless LAN Controllers
- 4400 Series Wireless LAN Controllers with Link Aggregation
- Catalyst 6500 Series WiSM module in Catalyst 6500 Series Switches

In a standalone solution, the only devices that connect to the access switches are APs. There might be management devices, such as the CiscoWorks Wireless LAN Solution Engine (CiscoWorks WLSE), but they usually connect like regular servers in a management VLAN. The main configuration task is to decide how the trunk to the APs should be configured.

In a controller-based solution, APs connect to the access switch. Depending on the role of the AP, it may connect to an access port or a trunk. Controllers also connect to the network. Depending on the model, they will need one or several ports, with or without link aggregation. While some controllers are independent appliances, some other controllers are modules that have been integrated into switches.

# Configure Switches for WLAN Devices

This topic describes configuration switches for WLAN devices.

## WLAN Device Connections

	Switch Port	QoS	Native VLAN	Management	Data
Standalone AP / Bridge	Trunk	Trust CoS	Management	Native VLAN	Local VLAN
Controller-Based AP	Access	Trust DSCP	AP IP Network	Via Controller	Via Controller
H-REAP	Trunk	Trust DSCP	AP IP Network	Via Controller	Local VLAN or via Controller
WLAN Controller	Trunk	Trust CoS	Not required	Mgmt VLAN	VLAN

The table shows an overview of the switch configuration for the AP and wireless LAN controllers. All ports to controller-based APs trust differentiated services code point (DSCP); all ports to controllers and autonomous APs trust class of service (CoS). The autonomous APs switch data traffic locally, while standard controller-based APs transmit all data traffic to the controller. H-REAP can switch locally or send traffic to the controller.



# Switch Configuration for Standalone APs and H-REAPs

This subtopic describes switch configuration for standalone APs and H-REAPs.

## Standalone AP and H-REAP

### 802.1Q trunk port

- Native management VLAN
- Data VLANs
- 802.1p QoS



```
switch(config)# interface fastethernet 0/1
switch(config-if)# switchport encapsulation dot1q
switch(config-if)# switchport trunk native vlan 10
switch(config-if)# switchport trunk allowed vlan 10,20
switch(config-if)# switchport mode trunk
switch(config-if)# spanning-tree portfast trunk
switch(config-if)# mls qos trust [cos | dscp]
```

The standalone AP is connected to a trunk port. If there is only one Service Set Identifier (SSID) and therefore no need for VLANs, the port may be converted to an access port.

When the AP connects to a trunk, a native (untagged) VLAN is required for management of the AP. When voice over wireless must be integrated, quality of service (QoS) must be configured for voice on the AP. CoS is trusted.

The H-REAP is connected to a trunk port. If there is only one SSID to switch locally and if this SSID-associated VLAN is the same as the H-REAP subnet, the port may be converted to an access port.

When the AP connects to a trunk, a native (untagged) VLAN is required for management of the AP. The H-REAP will use this subnet to communicate with its controller. When voice over wireless must be integrated, QoS must be configured for voice on the AP. DSCP is trusted both for locally switched traffic and for traffic that is sent to the controller.

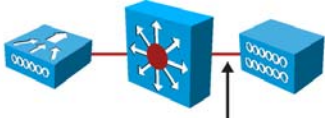
# Switch Configuration for Controller-Based APs

This subtopic describes switch configuration for controller-based APs.

## Controller-Based AP

Access port

- Native AP VLAN
- No data VLANs



```
switch(config)# interface fastethernet 0/2
switch(config-if)# switchport access vlan 10
switch(config-if)# switchport mode access
switch(config-if)# spanning-tree portfast
switch(config-if)# mls qos trust dscp
```

© 2009 Cisco Systems, Inc. All rights reserved. SWITCH-000000

The controller-based AP is usually connected to an access port. The native (untagged) VLAN is used for traffic to and from the controller. In a normal configuration, no traffic coming from or to a wireless client transits directly through the AP without going to the controller.

When voice over wireless is implemented, the AP will tag the frames that it sends to the controller to identify what type of traffic is being transmitted. In that case, QoS has to be configured to prioritize voice.

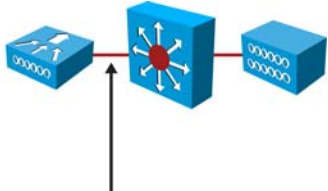
# Switch Configuration for a WLC

This subtopic describes the switch configuration for a WLC.

## WLAN Controller

802.1Q trunk port

- Management VLAN
- Data VLANs
- 802.1p QoS



```
switch(config)# interface fastethernet 0/3
switch(config-if)# switchport encapsulation dot1q
switch(config-if)# switchport trunk native vlan 99
switch(config-if)# switchport trunk allowed vlan 10,20
switch(config-if)# switchport mode trunk
switch(config-if)# spanning-tree portfast trunk
switch(config-if)# mls qos trust cos
```

© 2009 Cisco Systems, Inc. All rights reserved. SWITCH010-001

The WLC is connected to a trunk port. The allowed VLANs on the trunk are management and data VLANs. There may be a native VLAN, but not always. It is not a mandatory requirement.

Just like for APs, when voice over wireless is implemented, QoS has to be configured to prioritize voice traffic.

Notice that on the port to the AP, DSCP is trusted. Because the AP is in an access VLAN, its frame does not contain any 802.1D header, and therefore there is no CoS. Layer 3 QoS information must be used.


The controller connects to a trunk, and IEEE 802.1D is present in the header. Layer 2 QoS is to be trusted. This configuration requirement goes beyond the fact that the switch port is set as trunk and relies on the way in which QoS is tagged between the controller and the AP. Trusting DSCP on the controller port may have unpredictable consequences on the frame prioritization. CoS should be trusted.

# Link Aggregation for 4400 Series Controllers

This subtopic describes link aggregation for 4400 Series Controllers.

## 4400 Series Controller with Link Aggregation

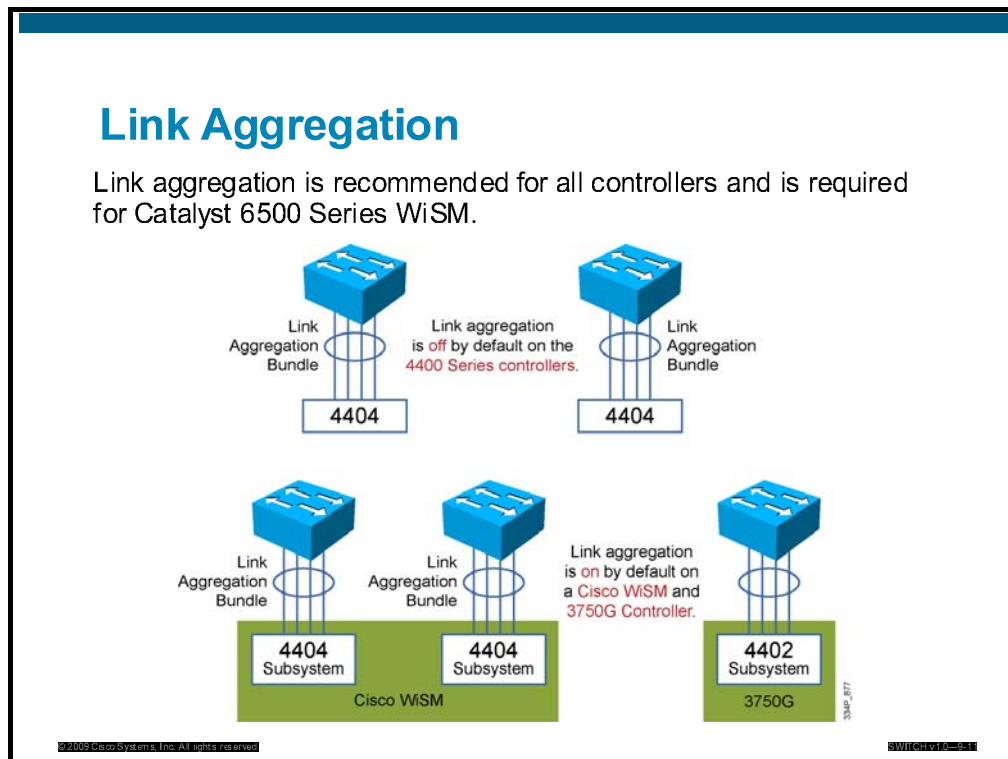
- One link aggregation group (LAG) per Cisco Wireless LAN Controller.
- Packets are forwarded from the controller on the same port on which they arrived.
- Load balancing is performed on the switch.
- A connection is made to a single switch or stack.
- EtherChannel configuration on switch is required.



Some controllers, such as the 4400 Series, have several links to the switch. These links can be bundled together. Link aggregation creates an EtherChannel between the controller and the switch. The EtherChannel provides additional bandwidth and link redundancy.

# Link Aggregation

This subtopic describes link aggregation on different WLCs.



Link aggregation is recommended on 4400 Series controllers. The Catalyst 6500 Series WiSM and the Catalyst 3750G Integrated Wireless LAN Controller require link aggregation be used. Link aggregation requires configuration of an EtherChannel on the switch. Notice that there is only one possible EtherChannel bundle for all links of the 4400 Series controllers. The Catalyst 6500 Series WiSM has two link bundles.

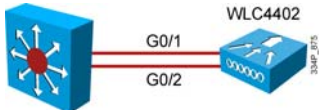
# Switch Configuration for 4400 Series Controllers

This subtopic describes the switch configuration for 4400 Series controllers.

## Switch Configuration for Link Aggregation

Gigabit Etherchannel 802.1Q trunk port

- Management VLAN
- AP VLAN
- Data VLANs
- 802.1p QoS



- Ports can be connected to different switches in a stacked switch or to different line cards in a modular switch.

```
switch(config)# interface gigabitethernet 0/1
switch(config-if)# channel-group 1 mode on
switch(config)# interface gigabitethernet 0/2
switch(config-if)# channel-group 1 mode on
switch(config)# interface port-channel 1
switch(config-if)# switchport encapsulation dot1q
switch(config-if)# switchport trunk native vlan 99
switch(config-if)# switchport trunk allowed vlan 10,20
switch(config-if)# switchport mode trunk
switch(config-if)# spanning-tree portfast trunk
switch(config-if)# mls qos trust cos
```

© 2009 Cisco Systems, Inc. All rights reserved. SWITCH-10-6-19

Link aggregation must be configured on the switch in a specific way. The WLC is connected to an EtherChannel trunk port. Use the **channel-group 1 mode on** command to configure a fixed, nonnegotiated EtherChannel. The controller is not able to negotiate the Bundle status, so neither Link Aggregation Control Protocol (LACP) nor Port Aggregation Protocol (PAgP) can be used.

# Switch Configuration for Cisco WiSM Controllers

This subtopic describes switch configuration for Cisco WiSM controllers.

## Cisco WiSM in Catalyst 6500 Series Switch



- Configure service VLAN and DHCP pool.
- Configure **wism service-vlan**.

```
c6500(config)# vlan 90
c6500(config)# interface vlan 90
c6500(config-if)# ip address 192.168.90.1 255.255.255.0
c6500(config)# ip dhcp pool WISM
c6500(config-dhcp)# network 192.168.90.0 255.255.255.0
c6500(config-dhcp)# default-router 192.168.90.1
c6500(config)# wism service-vlan 90
```

```
c6500# show wism status
Service Vlan: 90, Service IP Subnet: 192.168.90.1/255.255.255.0
WLAN
Slot Controller Service IP Management IP SW Version Status
-----+-----+-----+-----+-----+-----
3 1 192.168.90.2 169.254.1.1 4.2.176.0 Oper-Up
3 2 192.168.90.3 169.254.1.1 4.2.176.0 Oper-Up
```

The Cisco WiSM has two types of connection to the Catalyst 6500 Series platform. One type is the service port, used for local management of the Cisco WiSM, and the other type is the standard communication link to the network.

The first step is to configure the connection to the Cisco WiSM service port. You must create a specific service VLAN. The command **wism service-vlan** will affect that VLAN only to local communication with the Cisco WiSM. This VLAN should not be allowed on any link leaving the Catalyst 6500 Series switch.

In the service VLAN subnet, the Cisco WiSM will need to have an IP address, which can be statically defined or provided via DHCP.

# Switch Configuration for Cisco WiSM Controllers

This subtopic describes switch configuration for Cisco WiSM controllers.

## Cisco WiSM in Catalyst 6500 Series Switch (Cont.)

- Configure EtherChannel, VLANs, and QoS, using Cisco WiSM commands.
- Open a console session to configure the controllers.



```
c6500(config)# wism module slot# controller 1 native-vlan 99
c6500(config)# wism module slot# controller 2 native-vlan 99
c6500(config)# wism module slot# controller 1 allowed-vlan 10,20
c6500(config)# wism module slot# controller 2 allowed-vlan 10,20
c6500(config)# wism module slot# controller 1 qos trust cos
c6500(config)# wism module slot# controller 2 qos trust cos
```

```
c6500# session slot slot# processor controller#
```

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCHING-6500

The Cisco WiSM has eight ports connecting to the switch part of the Catalyst 6500 Series switch. These eight ports are organized into two EtherChannel bundles (the Cisco WiSM contains two distinct controllers). The **wism** family of commands is used to create the EtherChannel links between the Cisco WiSM and the switch.

If you need to connect to the Cisco WiSM command-line interface (CLI), use the **session** command, followed by the slot number in which the Cisco WiSM is set and the controller number to which you want to connect.



# Gathering Requirements

This topic describes gathering requirements.

## Gathering Requirements

- Controller-based or standalone solution?
- Number of APs?
- Where will the APs be installed?
- Switch ports for APs on access switches available?
- Is PoE on the access switches available?
- Are new access switches with PoE required?
- Is UPS required for APs?
- Where will the controllers be installed and connected?
- How do the APs and wireless clients get IP addresses (DHCP server)?
- VLANs or subnets for the APs, clients, and SSIDs?
- Is a RADIUS server required for security?
- Is a new access list required?
- Will a management system be installed?
- Is voice over WLAN planned (QoS)?
- What are the bandwidth requirements for wireless users?

The figure lists the standard questions that you need to answer before you can implement wireless integration into the campus network. The position and number of APs is, of course, a key question that you can resolve only by performing a professional wireless site survey. APs can use Power over Ethernet (PoE), or they may need AC adapters.

Depending on the number of APs, an autonomous or controller-based solution will be chosen. The traffic flow will depend on the chosen configuration. Another concern is security. Depending on the number and type of SSIDs, the traffic pattern, and therefore the VLAN requirements, will be very different.

# Planning the Integration

This subtopic describes planning the integration.

## Implementation Plan

- Collect the requirements.
- Check the existing network.
- Plan for additional equipment.
- Plan the implementation.
- Implement the new network components.
- Test the implemented network.

Follow these steps to create an implementation during design of the network:

- Collect requirements to determine how many ports of what type are needed and how they should be configured.
- Check the existing network to verify how the requirements can be integrated into the existing deployment. You will often find that, beyond the pure port count issue, the impact on bandwidth may imply additional connections.
- Plan additional equipment needed to fulfill the requirements.
- Plan the implementation.
- Implement new network components.

# Creating a Test Plan

This subtopic describes creating a test plan.

## Test Plan

- Can you reach the AP or WLC from management station?
- Can the AP reach the DHCP server?
- Does the AP get an IP address?
- Can the WLC reach RADIUS server?
- Do the clients get an IP address?
- Can the client access the network, server, Internet?

Implementation is not conducted without testing. When a wireless client starts, it first detects the wireless network and tries to connect at Layer 2 (using 802.11). This connection may imply 802.1X dialog between the AP (in autonomous mode) or the controller and a AAA server. When this step is complete, the wireless client tries to move to Layer 3 and get an IP address. The wireless client then has IP reachability to the network. Knowing these steps may help you troubleshoot wireless coactivity issues. For example, if the controller cannot communicate with the RADIUS server, the client may not be able to associate with the wireless network.

# Summary

This topic summarizes the key points that were discussed in this lesson.

## Summary

- WLAN devices are connected to switches.
- Access layer.
  - Standalone APs
  - Controller-based APs
- WLAN controllers are installed at the distribution layer or centralized in the data center and are connected to trunk ports or EtherChannel trunks.

## Lesson 4

---

# Lab 9-1 Debrief

---

## Overview

In this lab, you have configured your pod switches for theoretical planned wireless integration in your network. You have collected network design requirements from your clients, and then created an implementation plan with your team. After you determined which implementation plan was the most efficient, you connected to the remote lab and configured your switches to match your client requirements. You then verified that your implementation respected the client specific needs.

During the lab debrief, the instructor will lead a group discussion in which you can present your solution. You will get an opportunity to verify your solution against a number of checkpoints that are provided by the instructor, and to compare your solution to those of other students. The instructor will discuss alternative solutions and their benefits and drawbacks.

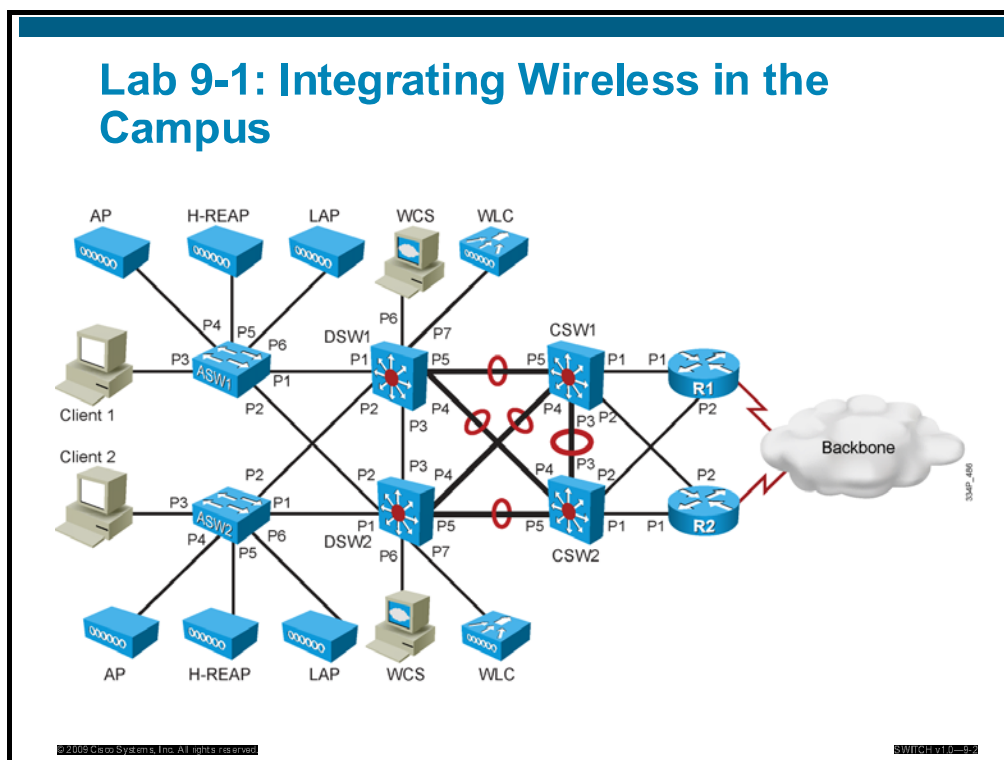
## Objectives

Upon completing this lesson, you will be able to produce a design plan and implementation plan for a Layer 2 network based on given business and technical requirements. You will then implement the plan and perform verifications while identifying checkpoints along the way. This ability includes being able to meet these objectives:

- Review and verify your solution, as well as your findings and action log, against a set of checkpoints that are provided by the instructor
- Consolidate the lessons that you learned during the review discussions into a set of best-practice methods and commands to aid you in future deployment procedures

# Review and Verification

This topic describes the client requirements that were listed in Lab 9-1, asks how you can verify that you have identified the solution matching the client needs, and gives you an example of a possible solution.



This lab consists of six switches that are to be configured in a coordinated manner to offer Layer 2 connectivity throughout your pod. Several devices are to be connected at a later time to your various switches, which means that your solution must be effective even though you will not be able to fully test all devices communications. Some devices are present in the topology to help you verify their connectivity. Your configuration tasks were focused on port configuration that would have to be set to specific VLANs or to a specific trunk. The difficulty of the exercise was to determine which ports must be trunks, which VLANs must be allowed, and which ports must be set as access ports.

## Implementation Plan

Which items should be configured, and in which order?

- VLAN assignment to ports?
- Trunk configuration?
- Layer 2 verification?
- VLAN pruning on trunks?
- VLAN creation?
- Layer 3 (ping) verification?

© 2009 Cisco Systems, Inc. All rights reserved.

SWITCHING-33

A successful implementation plan allows you to configure the devices with a minimum amount of duplication. In other words, an implementation plan is efficient when you do not need to alter your previous configuration to implement new items. You should proceed in a logical order. The list shown in the figure represents all the configuration tasks. List them in the most efficient order.





# Module Summary

This topic summarizes the key points that were discussed in this module.

## Module Summary

- WLANs are shared networks that provide access to networks for mobile users.
- Standalone and controller-based WLAN solutions are the Cisco WLAN implementations.
- Switches must be configured for the connection of APs and WLCs.



# Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which are two similarities between WLANs and wired LANs? (Choose two.) (Source: Comparing WLANs with Campus Networks)
- A) Both use MAC addresses.
  - B) Both use the same frame format.
  - C) Both can run the same applications.
  - D) Both use the same physical layer.
- Which are two differences between WLANs and wired LANs? (Choose two.) (Source: Comparing WLANs with Campus Networks)
- A) WLAN uses CSMA/CA, and wired LAN uses CSMA/CD.
  - B) WLANs have problems that are not found on wired LANs.
  - C) WLAN uses CSMA/CD, and wired LAN uses CSMA/CA.
  - D) WLANs and wired LANs run different applications.
- Q3) Which two statements are *not* true about SSIDs? (Choose two.) (Source: Comparing WLANs with Campus Networks)
- A) SSIDs on client and AP must match.
  - B) SSIDs are not case-sensitive.
  - C) A client can be configured without an SSID.
  - D) SSIDs on all APs must be identical.
- Q4) Which two wireless components are used for the standalone WLAN solution? (Choose two.) (Source: Assessing the Impact of WLANs on Campus Networks)
- A) WLC
  - B) ACS
  - C) CiscoWorks WLSE
  - D) WCS
- Q5) Which protocol supports split-MAC operation for the Cisco lightweight WLAN solution? (Source: Assessing the Impact of WLANs on Campus Networks)
- A) CCKM
  - B) LWAPP or CAPWAP
  - C) WLCCP
  - D) SNMP
- Q6) Which two devices can be used to configure a WLAN using controller-based APs? (Choose two.) (Source: Preparing the Campus Infrastructure for WLANs)
- A) Cisco WCS
  - B) Cisco WLSE
  - C) WLC
  - D) AP

## Module Self-Check Answer Key

- Q1) A, C
- Q2) A, B
- Q3) A, C
- Q4) B, C
- Q5) B
- Q6) A, C