



## CHAPTER 6

# Voice over WLAN Campus Test Architecture

---

This chapter describes the campus deployments of voice over wireless LAN (VoWLAN). The campus network used for the VoWLAN end-to-end testing was built based on the design recommendations documented in the existing campus CVD *Routed Access Design Guide*

Additional information on the existing campus CVD, including detailed test results, can be found under the campus heading at:

[http://www.cisco.com/en/US/netsol/ns742/networking\\_solutions\\_program\\_category\\_home.html](http://www.cisco.com/en/US/netsol/ns742/networking_solutions_program_category_home.html).

Elements from the following design guides were also used and documented in the network built for this chapter:

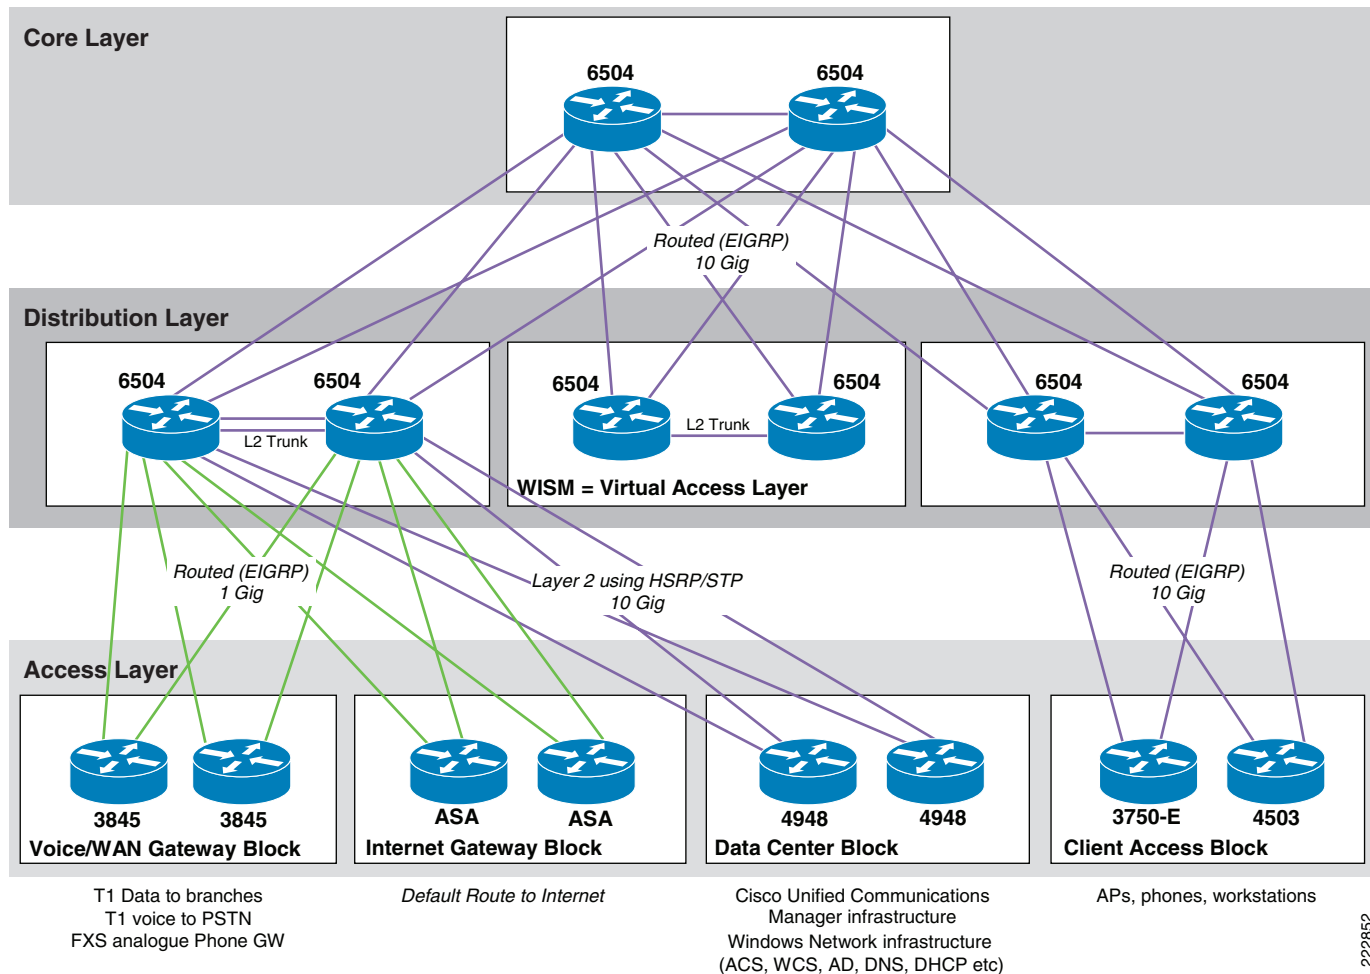
- *Enterprise QoS Solution Reference Network Design Guide Version 3.3*—  
[http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN\\_and\\_MAN/QoS\\_SRND/QoS-SRND-Book.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book.html)
- *Cisco AVVID Network Infrastructure IP Multicast Design (SRND)*—  
[http://www.cisco.com/application/pdf/en/us/guest/tech/tk363/c1501/ccmigration\\_09186a008015e7cc.pdf](http://www.cisco.com/application/pdf/en/us/guest/tech/tk363/c1501/ccmigration_09186a008015e7cc.pdf)
- *Cisco Unified Communications SRND Based on Cisco Unified Communications Manager 6.x*—  
[http://www.cisco.com/en/US/docs/voice\\_ip\\_comm/cucm/srnd/6x/uc6\\_0.html](http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/srnd/6x/uc6_0.html)
- *Enterprise Mobility 4.1 Design Guide*—  
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html>
- *Data Center Infrastructure Design Guide 2.5*—  
[http://www.cisco.com/en/US/docs/solutions/Enterprise/Data\\_Center/DC\\_Infra2\\_5/DCI\\_SRND\\_2\\_5\\_book.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_Infra2_5/DCI_SRND_2_5_book.html)

## Campus Design Overview

The campus design used for the VoWLAN testing is shown in [Figure 6-1](#). Most readers will be familiar with the hierarchical, access, distribution, core design used in this network. In the past, it was common to run Layer 2 links between the access layer and the distribution layer, and to use Spanning Tree Protocol (STP) to create a redundant, loop-free topology. Current design best practice replaces this Layer 2 STP configuration with a configuration that instead routes IP packets at the access layer. Routing in the access layer allows for faster re-convergence around failure, and simplifies the network by eliminating the need for STP and Hot standby Routing Protocol (HSRP). The campus CVD design documents referenced above provide additional information and detailed test results on the advantages of a routed access layer over a Layer 2 connection.

Data center devices have some unique network connectivity requirements that make it necessary to continue to connect to the access layer using Layer 2 and not routed IP. The unique requirements of the data center network connections are described in [Campus Layer 2 Spanning Tree Protocol Design for Data Center](#), page 6-9.

Figure 6-1 VoWLAN Campus Architecture



The architecture shown in [Figure 6-1](#) differs from the existing campus CVD *Routed Access Design Guide* in that multiple access layer blocks are connected to a single distribution block. In a production deployment, the campus CVD *Routed Access Design Guide* topology should be followed. This would result in each access layer block connecting to a dedicated distribution block, where policy appropriate to that type of access layer could be applied. Following the campus CVD *Routed Access Design Guide* would result in the voice/WAN gateway, the Internet Gateway, and the data center access blocks all connecting to dedicated distribution blocks rather than sharing a distribution block, as is shown in [Figure 6-1](#).

In accordance with the campus design recommendations, all router interconnections use fiber optic transport due to the superior link failure detection capabilities fiber links have when compared to wired Ethernet connections.

The [Configuration](#) section provides sample configurations that can be used as a basis for implementing the policies discussed in the following sections.

## Voice and Data VLAN Separation

Separate voice and data VLANs are used for the following reasons;

- Address space conservation and voice device protection from external networks—Private addressing of phones on the voice or auxiliary VLAN enables address conservation and ensures that phones are not accessible directly via public networks. PCs and servers are typically addressed with publicly routed subnet addresses; however, voice endpoints should be addressed using RFC 1918 private subnet addresses.
- QoS trust boundary extension to voice devices—QoS trust boundaries can be extended to voice devices without extending these trust boundaries and, in turn, QoS features to PCs and other data devices.
- Protection from malicious network attacks—Subnet access control, can provide protection for voice devices from malicious internal and external network attacks such as worms, denial-of- service (DoS) attacks, and attempts by data devices to gain access to priority queues.
- Ease of management and configuration—Separate VLANs for voice and data devices at the access layer provide ease of management and simplified QoS configuration.

## Campus IP Routing Design

Most of the information in this section was taken directly from the campus CVD *Routed Access Design Guide*. The campus CVD *Routed Access Design Guide* advocates the use of Layer 3 routing to the access layer in order to obtain the highest possible network availability. Test results in the *HA Campus Recovery Analysis* document (available at [http://www.cisco.com/en/US/netsol/ns815/networking\\_solutions\\_program\\_home.html](http://www.cisco.com/en/US/netsol/ns815/networking_solutions_program_home.html)) confirm the advantages of Layer 3 design over Layer 2 designs.

## EIGRP Routing Protocol

The campus documents referenced above describe the use of both OSPF and EIGRP as appropriate choices for the campus routing protocol. The test results show convergence for both of these routing protocols to be essentially equivalent. For this chapter, we chose to implement EIGRP due to its superior ease of deployment.

## Minimizing EIGRP Reconvergence Time

The length of time it takes for EIGRP or any routing protocol to restore traffic flows, after a network outage, within the campus is bounded by the following three main factors:

- The time required to detect the loss of a valid forwarding path
- The time required to determine a new best path
- The time required to update software and associated hardware forwarding tables

In the cases where the switch has redundant equal-cost paths, all three of these events are performed locally within the switch and controlled by the internal interaction of software and hardware. In the case where there is no second equal-cost path nor a feasible successor for EIGRP to use, the time required to determine the new best path is variable and primarily dependent on EIGRP query and reply propagation

across the network. To minimize the time required to restore traffic in the case where a full EIGRP routing convergence is required, it is necessary to provide strict bounds on the number and range of the queries generated.

Although EIGRP provides a number of ways to control query propagation, the two main methods are route summarization and the EIGRP stub feature. In the routed access hierarchical campus design, it is necessary to use both of these mechanisms.

## EIGRP Stub


Configuring EIGRP stub on the Layer 3 access switches prevents the distribution switch from generating downstream queries.

### Access Switch EIGRP Routing Process Stub Configuration

```
A4R-Top#sh run | beg router eigrp 100
router eigrp 100
  passive-interface default
  no passive-interface TenGigabitEthernet1/0/1
  no passive-interface TenGigabitEthernet2/0/1
  network 10.0.0.0
  no auto-summary
  eigrp router-id 10.33.9.19
  eigrp stub connected
```

By configuring the EIGRP process to run in “stub connected” state, the access switch advertises all connected subnets matching the network 10.0.0.0 0.255.255.255 range. It also advertises to its neighbor routers that it is a stub or non-transit router, and thus should never be sent queries to learn of a path to any subnet other than the advertised connected routes. With the design in Figure 6-2, the impact on the distribution switch is to limit the number of queries generated to 3 or less for any link failure.

**Figure 6-2** EIGRP Stub Limits the Number of Queries Generated to 3



To confirm that the distribution switch is not sending queries to the access switches, examine the EIGRP neighbor information for each access switch and look for the flag indicating queries being suppressed.

```
D3L#sh ip eigrp neigh det te 2/6
IP-EIGRP neighbors for process 100
H Address Interface Hold Uptime SRTT RTO Q Seq
(sec) (ms) Cnt Num
```

```

3 10.33.3.3 Te2/6 2 00:00:15 1004 5000 0 36
Version 12.2/1.2, Retrans: 0, Retries: 0
Stub Peer Advertising ( CONNECTED REDISTRIBUTED ) Routes
Suppressing queries

```

Configuring the access switch as a stub; router enforces hierarchical traffic patterns in the network. In the campus design, the access switch is intended to forward traffic only to and from the locally connected subnets. The size of the switch and the capacity of its uplinks are specified to meet the needs of the locally-connected devices. The access switch is never intended to be a transit or intermediary device for any data flows that are not to or from locally-connected devices. The hierarchical campus is designed to aggregate the lower speed access ports into higher speed distribution uplinks, and then to aggregate that traffic up into high-speed core links. The network is designed to support redundant capacity within each of these aggregation layers of the network, but not to support the re-route of traffic through an access layer. Configuring each of the access switches as EIGRP stub routers ensures that the large aggregated volumes of traffic within the core are never forwarded through the lower bandwidth links in the access layer, and also ensures that no traffic is ever mistakenly routed through the access layer, bypassing any distribution layer policy or security controls.

Each access switch in the routed access design should be configured with the EIGRP stub feature to aid in ensuring consistent convergence of the campus by limiting the number of EIGRP queries required in the event of a route recalculation, and to enforce engineered traffic flows to prevent the network from mistakenly forwarding transit traffic through the access layer.

For more information on the EIGRP stub feature, see the following URL:  
[http://www.cisco.com/en/US/docs/ios/12\\_0s/feature/guide/eigrpstb.html](http://www.cisco.com/en/US/docs/ios/12_0s/feature/guide/eigrpstb.html)

## Distribution Summarization

Configuring EIGRP stub on all of the access switches reduces the number of queries generated by a distribution switch in the event of a downlink failure, but it does not guarantee that the remaining queries are responded to quickly. In the event of a downlink failure, the distribution switch generates three queries; one sent to each of the core switches, and one sent to the peer distribution switch. The queries generated ask for information about the specific subnets lost when the access switch link failed. The peer distribution switch has a successor (valid route) to the subnets in question via its downlink to the access switch, and is able to return a response with the cost of reaching the destination via this path. The time to complete this event depends on the CPU load of the two distribution switches and the time required to transmit the query and the response over the connecting link. In the campus environment, the use of hardware-based CEF switching and GigE or greater links enables this query and response to be completed in less than a 100 msec.

This fast response from the peer distribution switch does not ensure a fast convergence time, however, EIGRP recovery is bounded by the longest query response time. The EIGRP process has to wait for replies from all queries to ensure that it calculates the optimal loop free path. Responses to the two queries sent towards the core need to be received before EIGRP can complete the route recalculation. To ensure that the core switches generate an immediate response to the query, it is necessary to summarize the block of distribution routes into a single summary route advertised towards the core.

```

D3L#sh run | begin TenGigabitEthernet2/1
interface TenGigabitEthernet2/1
description from D3L to CL
ip address 10.33.1.11 255.255.255.254
ip hello-interval eigrp 100 1
ip hold-time eigrp 100 3
ip authentication mode eigrp 100 md5
ip authentication key-chain eigrp 100 eigrp-chain
ip pim sparse-mode
ip summary-address eigrp 100 10.33.48.0 255.255.240.0 5
logging event link-status

```

```
load-interval 30
carrier-delay msec 0
mls qos trust dscp
```


The summary-address statement is configured on the uplinks from each distribution switch to both core nodes. In the presence of any more specific component of the 10.33.0.0/16 address space, it causes EIGRP to generate a summarized route for the 10.33.0.0/16 network, and to advertise only that route upstream to the core switches.

```
CL#sh ip route 10.33.50.0
Routing entry for 10.33.48.0/20
  Known via "eigrp 100", distance 90, metric 3328, type internal
  Redistributing via eigrp 100
  Last update from 10.33.1.11 on TenGigabitEthernet2/5, 4w3d ago
  Routing Descriptor Blocks:
  * 10.33.1.13, from 10.33.1.13, 4w3d ago, via TenGigabitEthernet2/6
    Route metric is 3328, traffic share count is 1
    Total delay is 30 microseconds, minimum bandwidth is 1000000 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 2
  10.33.1.11, from 10.33.1.11, 4w3d ago, via TenGigabitEthernet2/5
    Route metric is 3328, traffic share count is 1
    Total delay is 30 microseconds, minimum bandwidth is 1000000 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 2
```

With the upstream route summarization in place, whenever the distribution switch generates a query for a component subnet of the summarized route, the core switches reply that they do not have a valid path (cost = infinity) to the subnet query. The core switches are able to respond within less than 100 msec if they do not have to query other routers before replying back to the subnet in question.

Figure 6-3 shows an example of summarization toward the core.

Figure 6-3 Summarization toward the Core Bounds EIGRP Queries for Distribution Block Routes



Using a combination of stub routing and summarizing the distribution block routes upstream to the core both limits the number of queries generated and bounds those that are generated to a single hop in all directions. Keeping the query period bounded to less than 100 msec keeps the network convergence

similarly bounded under 200 msec for access uplink failures. Access downlink failures are the worst case scenario because there are equal-cost paths for other distribution or core failures that provide immediate convergence.

## Route Filters

The discussion on EIGRP stub above noted that in the structured campus model, the flow of traffic follows the hierarchical design. Traffic flows pass from access through the distribution to the core and should never pass through the access layer unless they are destined to a locally attached device. Configuring EIGRP stub on all the access switches aids in enforcing this desired traffic pattern by preventing the access switch from advertising transit routes. As a complement to the use of EIGRP stub, Cisco recommends applying a distribute-list to all the distribution downlinks to filter the routes received by the access switches. The combination of “stub routing” and route filtering ensures that the routing protocol behavior and routing table contents of the access switches are consistent with their role, which is to forward traffic to and from the locally connected subnets only.

Cisco recommends that a default route (0.0.0.0 mask 0.0.0.0) be the only route advertised to the access switches.

```
D3L#sh run | begin router eigrp 100
router eigrp 100
...
network 10.0.0.0
distribute-list only-default out TenGigabitEthernet2/4
distribute-list only-default out TenGigabitEthernet2/6
...
eigrp router-id 10.33.9.8
!
...
!
ip access-list standard only-default
permit 0.0.0.0
remark redistribute default route to access layer stub-routers
```

No mask is required in the configuration of this access list because the assumed mask, 0.0.0.0, permits only the default route in the routing updates.

In addition to enforcing consistency with the desire for hierarchical traffic flows, the use of route filters also provides for easier operational management. With the route filters in place, the routing table for the access switch contains only the essential forwarding information. Reviewing the status and/or troubleshooting the campus network is much simpler when the routing tables contain only essential information.

```
A4L#sh ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route
```

```
Gateway of last resort is 10.33.3.24 to network 0.0.0.0
```

```
2.0.0.0/32 is subnetted, 1 subnets
C    2.2.2.2 is directly connected, Loopback2
10.0.0.0/8 is variably subnetted, 5 subnets, 3 masks
C    10.33.49.0/24 is directly connected, Vlan49
C    10.33.48.0/24 is directly connected, Vlan48
C    10.33.3.4/31 is directly connected, TenGigabitEthernet1/1
```

```

C      10.33.3.24/31 is directly connected, TenGigabitEthernet1/2
C      10.33.9.20/32 is directly connected, Loopback1
D*EX 0.0.0.0/0 [170/3584] via 10.33.3.24, 3w4d, TenGigabitEthernet1/2
      [170/3584] via 10.33.3.4, 3w4d, TenGigabitEthernet1/1

```

If the network does not contain a default route, it may be acceptable to use an appropriate full network summary route in its place; that is, 10.0.0.0/8, or a small subset of summary routes that summarize all possible destination addresses within the network.

## CEF Switching

Per-destination load balancing is enabled by default when you enable CEF, and is the load balancing method of choice for most situations. Per-destination load balancing allows the router to use multiple paths to achieve load sharing. Packets for a given source-destination host pair are guaranteed to take the same path, even if multiple paths are available.

When you enable CEF or dCEF globally, all interfaces that support CEF are enabled by default.

## Adjusting EIGRP timers

The recommended best practice for campus design is to use point-to-point fiber connections for all links between switches. Link failure detection via 802.3z and 802.3ae remote fault detection mechanism provide for recovery from most campus switch component failures. Cisco recommends in the Layer 3 campus design that the EIGRP hello and dead timers be reduced to 1 and 3 seconds, respectively (see [Figure 6-4](#)). The loss of hellos and the expiration of the dead timer provides a backup to the L1/2 remote fault detection mechanisms. Reducing the EIGRP hello and hold timers from defaults of 5 and 15 seconds provides for a faster routing convergence in the rare event that L1/2 remote fault detection fails to operate, and hold timer expiration is required to trigger a network convergence because of a neighbor failure.

**Figure 6-4** Reducing EIGRP Hello and Dead Timers

```

interface TenGigabitEthernet4/3
description 10 GigE to Distribution 1
ip address 10.122.0.26 255.255.255.254
. . .
ip hello-interval eigrp 100 1
ip hold-time eigrp 100 3
. . .
interface TenGigabitEthernet2/1
description 10 GigE to Core 1
ip address 10.122.0.27 255.255.255.254
. . .
ip hello-interval eigrp 100 1
ip hold-time eigrp 100 3
. . .

```

Ensure Timers are consistent on both ends of the link

132710



## Campus Layer 2 Spanning Tree Protocol Design for Data Center

Much of the information in this section is summarized from the *Data Center Infrastructure 2.5 Design Guide*. For details, refer to this guide at:

[http://www.cisco.com/en/US/docs/solutions/Enterprise/Data\\_Center/DC\\_Infra2\\_5/DCI\\_SRND\\_2\\_5\\_book.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_Infra2_5/DCI_SRND_2_5_book.html).

### Data Center Layer Introduction

In the test network used for this document, the data center is used to connect network infrastructure devices such as Cisco Unified Communications Manager, Cisco Unified Unity, ACS, WCS, DNS, DHCP.

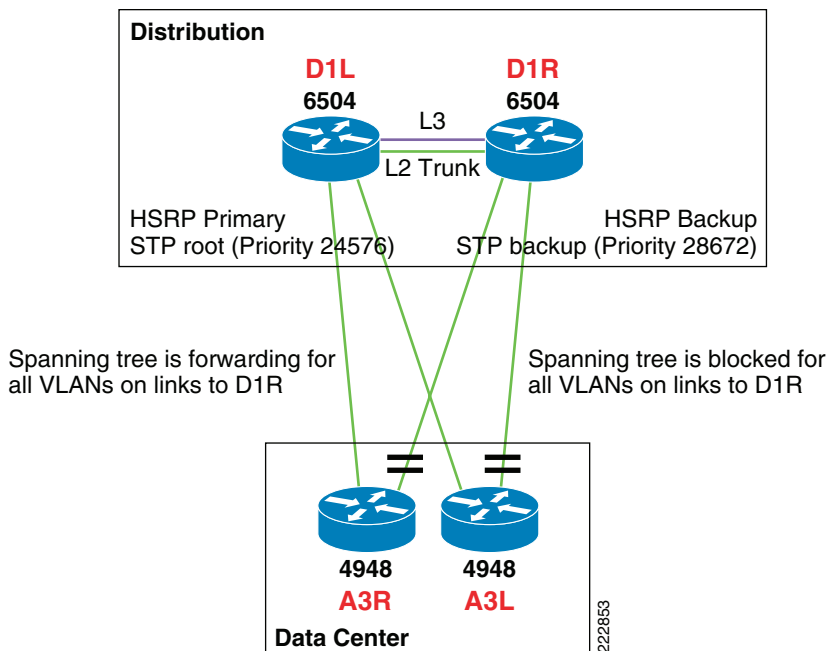
While the rest of the campus network uses Layer 3 routing (EIGRP) all the way to the access layer, the unique requirements of data center deployments require a Layer 2 connection to the data center. A Layer 2 access topology provides the following unique capabilities required in the data center:

- VLAN extension—The Layer 2 access topology provides the flexibility to extend VLANs between switches that are connected to a common aggregation module. This makes provisioning of servers to a particular subnet/VLAN simple, and without the worry of physical placement of the server in a particular rack or row.
- Layer 2 adjacency requirements—NIC teaming, high availability clusters, and database clusters are application examples that typically require NIC cards to be in the same broadcast domain (VLAN). The list of applications used in a clustered environment is growing, and Layer 2 adjacency is a common requirement.
- Custom applications—Many developers write custom applications without considering the Layer 3 network environment. This can create challenges in a Layer 3 IP access topology. These servers usually depend on Layer 2 adjacency with other servers and may require rewriting of code when changing IP addresses.
- Service modules—A Layer 2 access permits services provided by service modules or appliances to be shared across the entire access layer. Examples of this are when using the FWSM, CSM, and SSLSM. The active-standby modes of operation used by service modules require Layer 2 adjacency with the servers that use them.

### Spanning Tree Triangle Looped Topology

The triangle looped topology utilized in the network used for this document is currently the most widely implemented in the enterprise data center. This topology provides a deterministic design that makes it easy to troubleshoot while providing a high level of flexibility (see [Figure 6-5](#)).

Figure 6-5 Triangle Looped Access Topology



In a triangle looped access layer design, it is desirable to align the spanning tree root, HSRP default gateway, and active service modules on the same aggregation switch, as shown in Figure 6-5. Aligning the access layer switch uplink that is in the forwarding state directly to the same switch that is the primary default gateway and active service module/appliance optimizes the traffic flows. Otherwise, traffic flows can hop back and forth between aggregation switches, creating undesirable conditions and difficulty in troubleshooting.

### Lab Implementation Deviations from Standard Design Principals in the Data Center

In the network used for this chapter, multiple access layer blocks are connected to a single distribution block. This was done to reduce the number of distribution layer switches required. In a production deployment, the data center access-block would connect to a dedicated distribution block. In a production network, the distribution block dedicated to the data center access-block would run a variety of service modules supporting data center access. Example service modules include:

- Firewall Services Modules (FWSM)
- Secure Sockets Layer Services Modules (SSLSM)
- Content Switching Module (CSM)
- Intrusion Detection
- Network Analysis Module (NAM)
- Distributed denial-of-service attack protection (Cisco Guard)

While some combination of these would be expected in a production distribution layer block, our test network is not specifically focused on data center testing, and we have implement the network without any of these service modules.

# Campus Quality of Service (QoS)

Much of the information in this section is summarized from the *Enterprise QoS Solution Reference Network Design Guide Version 3.3*. For details refer to this guide at:

[http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN\\_and\\_MAN/QoS\\_SRND/QoS-SRND-Book.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book.html).

QoS refers to the set of tools and techniques used to manage network resources. QoS technologies allow different types of traffic to contend inequitably for network resources. Using QoS, some network traffic such as Voice, video, or critical data may be granted priority or preferential services from network devices; this prevents lower priority background traffic from degrading the quality of these strategic applications.

Until recently, QoS was not a great concern in the enterprise campus due to the large amount of available bandwidth as well as the asynchronous nature of data traffic and the ability of applications to tolerate the effects of buffer overflows and packet loss. However, with new applications such as voice and video, which are sensitive to packet loss and delay it is important to utilize quality of service features to protect high priority traffic.

## Campus QoS Services Required

Access switches require the following QoS policies:

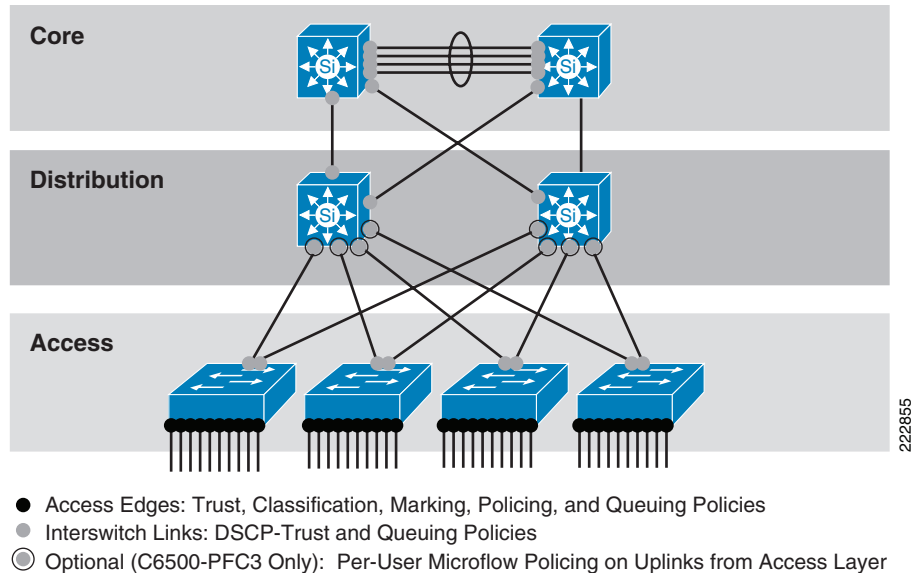
- Appropriate (endpoint-dependant) trust policies, and/or classification and marking policies
- Policing and markdown policies
- Queuing policies

Distribution and core switches require the following QoS policies:

- DSCP trust policies
- Queuing policies
- Optional per-user microflow policing policies (only on supported platforms)

These recommendations are summarized in [Figure 6-6](#).

Figure 6-6 Typical Campus Oversubscription Ratios




## Campus QoS and Interface Queuing

Many campus links are underutilized. Some studies have shown that 95 percent of campus access layer links are used at less than 5 percent of their capacity. This means that you can design campus networks to accommodate oversubscription between access, distribution and core layers. Oversubscription allows for uplinks to be used more efficiently and more importantly, reduces the overall cost of building the campus network.

Common campus oversubscription values are 20:1 for the access-to-distribution layers and 4:1 for the distribution-to-core layers, as shown in Figure 6-7.

Figure 6-7 QoS Requirement within the Campus



The potential for congestion exists in campus uplinks because of oversubscription ratios and speed mismatches in campus downlinks (for example, GigabitEthernet to FastEthernet links). The only way to provision service guarantees in these cases is to enable advanced interface queuing at these points.

For a given input or output interface, Cisco IOS can manage multiple queues. Traffic of a particular priority is mapped into the appropriate queue for that traffic class. Cisco IOS queue scheduling algorithms can then be configured to service those queues, giving more priority to servicing the queues containing higher priority traffic. In this way, higher priority traffic is protected from queue overruns caused by lower priority traffic.

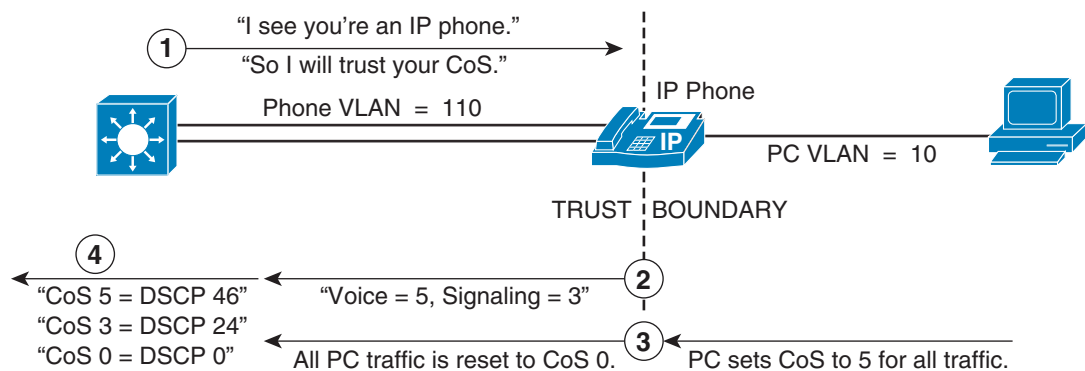
## QoS and Wired IP Phones

This section is specific to wired Ethernet-attached IP phone. Wireless LAN IP phones such as the Cisco 7921G use a different QoS procedure and are discussed in [Chapter 10, “Cisco Unified IP Phone 7921 Implementation for Voice over WLAN.”](#)

Cisco wired IP phones perform an intelligent exchange of information between the phone and the switchport it is plugged into using Cisco Discovery Protocol (CDP). When the switch discovers a Cisco IP phone, it can extend QoS trust to it dynamically.

[Figure 6-8](#) shows a conditional trust boundary extension granted to an IP phone that has passed a CDP exchange.

**Figure 6-8** Conditionally Trusted Endpoint Trust Boundary Extension and Operation



- ① Switch and phone exchange CDP; trust boundary is extended to IP phone.
- ② Phone sets CoS to 5 for VoIP and to 3 for Call-Signaling traffic.
- ③ Phone rewrites CoS from PC Port to 0.
- ④ Switch trusts CoS from phone and maps CoS → DSCP for output queuing.

222856

The sequence shown in [Figure 6-8](#) is the following:

- 
- Step 1** Switch and phone exchange CDP; trust boundary is extended to IP phone.
- Step 2** Phone sets CoS to 5 for VoIP and to 3 for call signaling traffic.
- Step 3** Phone rewrites CoS from PC to 0.
- Step 4** Switch trusts CoS from phone and maps CoS to DSCP for output queuing.
-

## AutoQoS and VoIP

When the main business objective of the QoS deployment is to enable QoS for IP Telephony only (i.e., without Scavenger-class QoS), then the network administrator may choose to take advantage of the Cisco AutoQoS VoIP feature.

AutoQoS VoIP is essentially an intelligent macro that enables an administrator to enter one or two simple AutoQoS commands to enable all the appropriate features for the recommended QoS settings for VoIP and IP telephony for a specific platform and/or a specific interface.

AutoQoS VoIP automatically configures the best-practice QoS configurations (based on previous Cisco Enterprise QoS SRNDs) for VoIP on Cisco Catalyst switches and IOS routers. By entering one global and/or one interface command (depending on the platform), the AutoQoS VoIP macro then would expand these commands into the recommended VoIP QoS configurations (complete with all the calculated parameters and settings) for the platform and interface on which the AutoQoS VoIP macro is applied.

For example, on Cisco Catalyst switches, AutoQoS performs the following automatically:

- Enforces a conditional-trust boundary with any attached Cisco IP phones
- Enforces a trust boundary on Catalyst switch access ports and uplinks/downlinks
- Modifies CoS-to-DSCP (and IP Precedence-to-DSCP) mappings, as required
- Enables Catalyst strict priority queuing for voice (CoS 5/DSCP EF) and preferential queuing for Call-Signaling traffic (CoS 3/DSCP CS3) \* Enables best-effort queuing for all other data (CoS 0/DSCP 0) traffic
- Modifies queue admission criteria (such as CoS-to-queue mapping)
- Modifies queue sizes and queue weights, where required

## QoS Policing

At the time of writing, classification using a port trust state (for example, `mls qos trust [cos | dscp | ip-precedence]`) and a policy map (for example, `service-policy input policy-map-name`) are mutually exclusive on the Catalyst 2970/3560/3750. The last one configured overwrites the previous configuration. This limitation is to be addressed; consult the latest Catalyst 2970/3560/3750 QoS documentation for updates on this limitation.

Because of the above limitation, policing will not be discussed in the configuration section of this document. For detailed guidance on policing configuration, refer to the *Enterprise QoS Solution Reference Network Design Guide Version 3.3* at the following URL:  
[http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN\\_and\\_MAN/QoS\\_SRND/QoS-SRND-Book.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book.html)

# Campus Multicast

For complete details on the recommended campus multicast design, see the Cisco AVVID Network Infrastructure IP Multicast Design SRND at the following URL:  
[http://www.cisco.com/application/pdf/en/us/guest/tech/tk363/c1501/ccmigration\\_09186a008015e7cc.pdf](http://www.cisco.com/application/pdf/en/us/guest/tech/tk363/c1501/ccmigration_09186a008015e7cc.pdf).

When designing multicast networking, a routed access design has advantages over a design that has Layer 2 to the distribution layer. In the Layer 2 design, there are two routers on the same subnet as the multicast hosts. This results in the following;

- One of the routers needs to be elected the PIM DR and IGMP querier
- One of the routers needs to be elected the HSRP active node
- One of the routers needs to be elected the Layer 2 root bridge

For each specific VLAN The root bridge, the HSRP active node, and the PIM DR should all be on the same distribution switch.

In the routed access design, this need for synchronized configuration is removed because there is only one router on the local segment, which by default results in synchronization of the unicast and multicast traffic flows. Additionally, with the migration of the multicast router from the distribution to the access, there is no longer a need to tune the PIM hello timers to ensure rapid convergence between the distribution nodes in the case of a failure. The same remote fault indicator mechanisms that trigger rapid unicast convergence drive the multicast software and hardware recovery processes, and there is no need for Layer 3 detection of path or neighbor failure across the Layer 2 access switch. The presence of a single router for each access VLAN also removes the need to consider non-reverse path forwarding (non-RPF) traffic received on the access side of the distribution switches. A multicast router drops any multicast traffic received on a non-RPF interface. If there are two routers for a subnet, the DR forwards the traffic to the subnet, and the non-DR receives that traffic on its own VLAN interface where it fails the RPF check and so must be dropped

## Multicast Traffic Flows and Router Functions

In the Layer 3 access design, there is a single router on the access subnet and no non-RPF traffic flows. Although the current generation of Cisco Catalyst switches can process and discard all non-RPF traffic in hardware with no performance impact or access list configuration required, the absence of non-RPF traffic simplifies operation and management.

The following summarizes the campus multicast configuration implementation:

- The access layer switches have IGMP snooping enabled.
- The RPs are located on the two core layer switches ( RPs should be close to the multicast sources. In our lab the core routers represented the best location.).
- PIM-SM is configured on all access layer, distribution layer, and core-layer switches.
- Anycast RP is configured for fast recovery of IP multicast traffic.
- PIM-SM and MSDP are enabled on all core layer switches. (auto RP is not used due to its reliance on PIM-sparse-dense)
- Each access layer switch points to the anycast RP address as its RP.
- MSDP is used to synchronize source active (SA) state between the core switches.

# Campus Security

The Campus Security section is focused on the Catalyst Integrated Security Features (CISF) features found in the access layer switches connecting end-user devices to the network.

Much of the information in this section is summarized from *Cisco Unified Communications SRND Based on Cisco Unified Communications Manager 6.x* at the following URL:

[http://www.cisco.com/en/US/docs/voice\\_ip\\_comm/cucm/srnd/6x/uc6\\_0.html](http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/srnd/6x/uc6_0.html)

Refer to the above document as well as the relevant access-layer switch documentation for more details.

## Catalyst Integrated Security Features (CISF)

Layer 2 switched environments can prove easy targets for security attacks. These attacks exploit normal protocol processing such as a switches' ability to learn MAC addresses, end-station Media Access Control (MAC) address resolution through Address Resolution Protocol (ARP) or Dynamic Host Configuration Protocol (DHCP) IP address assignments.

The rich set of integrated security features on Cisco Catalyst Switches (CISF) protect your critical network infrastructure with easy-to-use tools that effectively prevent the most common—and potentially damaging—Layer 2 security threats.

## Port Security

### Function

Shuts down MAC address-flooding attacks .


### How It Works

A classic attack on a switched network is a MAC content-addressable memory (CAM) flooding attack. This type of attack floods the switch with so many MAC addresses that the switch does not know which port an end station or device is attached to. When the switch does not know which port a device is attached to, it broadcasts the traffic destined for that device to the entire VLAN. In this way, the attacker is able to see all traffic that is coming to all the users in a VLAN.

Port security limits the number of MAC addresses allowed to access individual switch ports thus protecting against MAC flooding attacks from hacker tools such as macof (see [Figure 6-9](#)).



**Figure 6-9** Limited Number of MAC Addresses Prevents Rogue Network Extensions

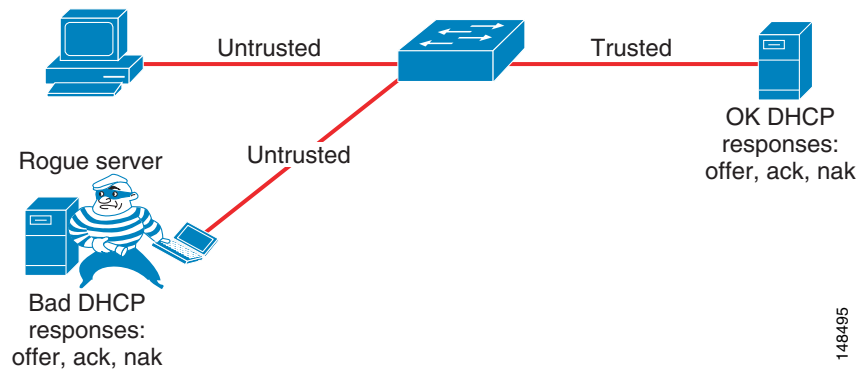


## DHCP Snooping

### Function

Prevent rogue DHCP server attacks, DHCP starvation attacks, create client binding information used by additional CISF tools

**Figure 6-10** Using DHCP Snooping



## DHCP Snooping: Prevent Rogue DHCP Server Attacks

### How It Works

Dynamic Host Configuration Protocol (DHCP) Snooping prevents a non-approved DHCP or rogue DHCP server from handing out IP addresses on a network by blocking all replies to a DHCP request unless that port is allowed to reply. When enabled, DHCP Snooping treats all ports in a VLAN as untrusted by default. An untrusted port is a user-facing port that should never make any reserved DHCP responses. If an untrusted DHCP-snooping port makes a DHCP server response, it will be blocked from responding. Therefore, rogue DHCP servers will be prevented from responding. However, legitimately attached DHCP servers or uplinks to legitimate servers must be trusted.

## DHCP Snooping: Prevent DHCP Starvation Attacks

### How It Works

DHCP address scope starvation attacks from tools such as Gobbler are used to create a DHCP denial-of-service (DoS) attack. Because the Gobbler tool makes DHCP requests from different random source MAC addresses, you can prevent it from starving a DHCP address space by using port security to limit the number of MAC addresses. However, a more sophisticated DHCP starvation tool can make the DHCP requests from a single source MAC address and vary the DHCP payload information. With DHCP Snooping enabled, untrusted ports will make a comparison of the source MAC address to the DHCP payload information and fail the request if they do not match.

## DHCP Snooping: Binding Information

### How It Works

Another function of DHCP Snooping is to record the DHCP binding information for untrusted ports that successfully get IP addresses from the DHCP servers. The binding information is recorded in a table on the Cisco Catalyst switch. The DHCP binding table contains the IP address, MAC address, lease length, port, and VLAN information for each binding entry. The binding information from DHCP Snooping remains in effect for the length of the DHCP binding period set by the DHCP server (that is, the DHCP lease time). The DHCP binding information is used to create dynamic entries for Dynamic ARP Inspection (DAI) to limit ARP responses for only those addresses that are DHCP-bound. The DHCP binding information is also used by the IP source guard to limit sourcing of IP packets to only those addresses that are DHCP-bound.

## Dynamic ARP Inspection

### Function

Adds security to ARP using DHCP snooping table.

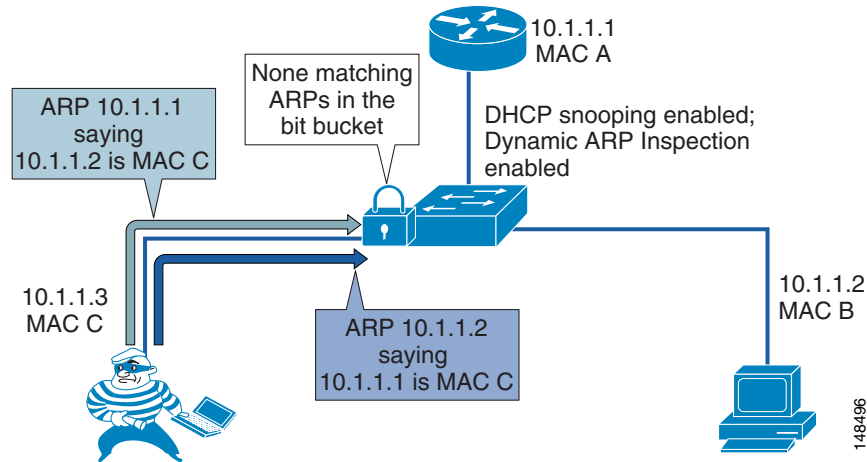
### How It Works

Dynamic Address Resolution Protocol (ARP) Inspection (DAI) is a feature used on the switch to prevent Gratuitous ARP attacks on the devices plugged into the switch and on the router.

Gratuitous ARP can be exploited by malicious programs that want to illegitimately take on the identity of another station. When a malicious station redirects traffic to itself from two other stations that were talking to each other, the hacker who sent the GARP messages becomes the man-in-the-middle. Hacker programs such as ettercap do this with precision by issuing "private" GARP messages to specific MAC addresses rather than broadcasting them. In this way, the victim of the attack does not see the GARP packet for its own address. Ettercap also keeps its ARP poisoning in effect by repeatedly sending the private GARP messages every 30 seconds.

Dynamic ARP Inspection (DAI) is used to inspect all ARP requests and replies (gratuitous or non-gratuitous) coming from untrusted (or user-facing) ports to ensure that they belong to the ARP owner. The ARP owner is the port that has a DHCP binding which matches the IP address contained in the ARP reply. ARP packets from a DAI trusted port are not inspected and are bridged to their respective VLANs.

**Figure 6-11** Using DHCP Snooping and DAI to Block ARP Attacks:



## IP Source Guard

### Function

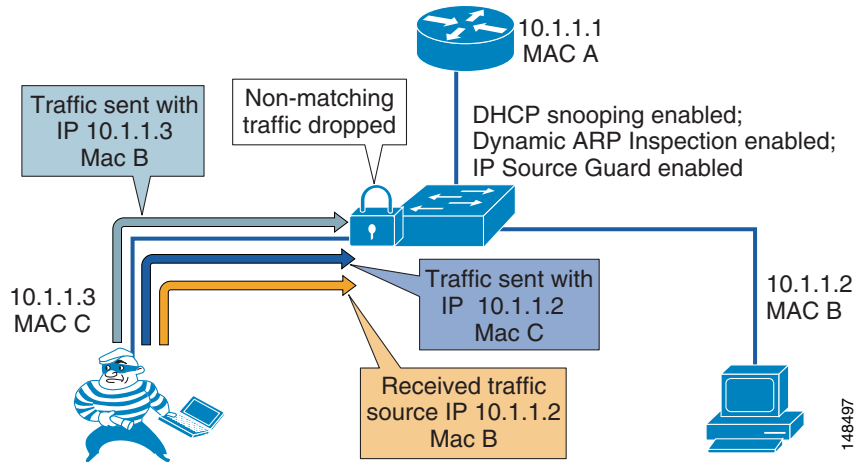
Prevents IP host spoofing.

### How It Works

IP address spoofing is commonly used to perform DoS attacks on a second party. A simple example of this occurs when an attacker pings a third-party system while sourcing the IP address of the second party that is being attacked (see [Figure 6-12](#)). The ping response will be directed to the second party from the third-party system. Aggressive SYN-flooding originating from spoofed IP addresses is another common type of attack used to overwhelm a server with TCP half-sessions.

The IP Source Guard (IPSG) feature, when invoked, dynamically creates an ACL based on the contents of the DHCP Snooping binding table. This ACL ensures that traffic is sourced from the IP address issued at DHCP binding time and prevents any traffic from being forwarded by other spoofed addresses. While DHCP Snooping is a prerequisite for IP Source Guard, DAI is not. However, Cisco recommend that you enable DAI in addition to IP Source Guard to prevent ARP-poisoning man-in-the-middle attacks in addition to IP address spoofing.

**Figure 6-12** Using IP Source Guard to Prevent Address Spoofing:



## Network Time Services

An essential element of network management, troubleshooting, and security operations is to have all network elements (including routers, switches, and servers) synchronized to a common clock source.

Network Time Protocol (NTP) is most commonly used to synchronize clocks in network equipment. NTP provides accuracies typically within a millisecond on LANs and up to a few tens of milliseconds on WANs.

By synchronizing the clocks across the network, it is possible to examine the exact sequence in which events occurred. This ability to analyze and correlate the sequence of events across multiple network elements makes it much easier to determine the root cause of network problems or security issues.

The configuration section of this document shows the commands used in our network to have NTP synchronize to an external time source. In most production environments, the external time source used would be redundant, dedicated hardware that synchronizes via a GPS receiver to a clock source that is itself directly synchronized to an atomic clock.

### NTP links

- <http://support.ntp.org/bin/view/Main/WebHome>
- [http://en.wikipedia.org/wiki/Network\\_Time\\_Protocol](http://en.wikipedia.org/wiki/Network_Time_Protocol)
- [http://www.cisco.com/en/US/tech/tk869/tk769/technologies\\_white\\_paper09186a0080117070.shtml](http://www.cisco.com/en/US/tech/tk869/tk769/technologies_white_paper09186a0080117070.shtml)
- [http://www.cisco.com/en/US/products/sw/voicewp/ps556/products\\_white\\_paper0900aecd8037fdb5.shtml](http://www.cisco.com/en/US/products/sw/voicewp/ps556/products_white_paper0900aecd8037fdb5.shtml)
- <http://www.symmetricom.com/>
- <http://www.meinberg.de/english/products/>

# UniDirectional Link Detection (UDLD)

Use UDLD to protect against one-way up/up connections. In fiber topologies where fiber optic interconnections are used, which is common in a campus environment, physical misconnections can occur that allow a link to appear to be up/up when there is a mismatched set of transmit/receive pairs. When such a physical misconfiguration occurs, protocols such as EIGRP or STP can cause network instability. UDLD detects these physical misconfigurations and disables the ports in question

## Configuration

### Software versions used in test network

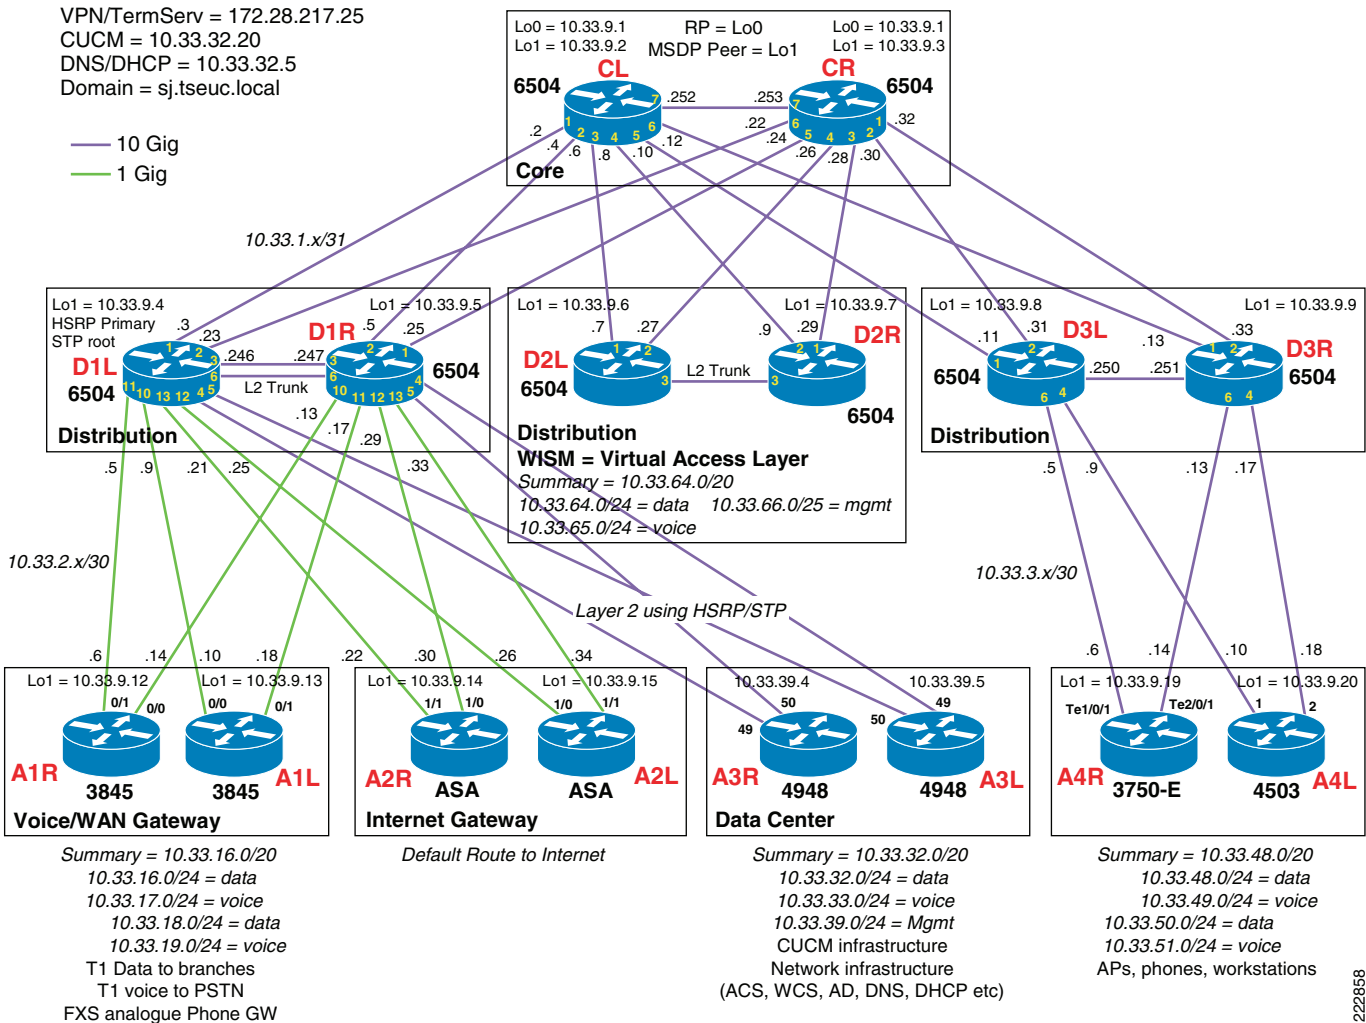
**Table 6-1**      *Software Releases Used in Testing*

Platform	Role	SW Version
6504	Campus Core	12.2(18)SXF9
6504	Campus Distribution	12.2(18)SXF9
6504	WLAN	12.2(18)SXF9
3845	Voice WAN gateway	12.4(15)T1
ASA	Internet gateway	7.2(2)
4948	Data center	12.2(25)EWA8
4503	Access	12.2(37)SG
3750-E	Access	12.2(37)SE1
2821	Branch Router	12.4(15)T1

### Lab Network used for Configuration Examples

Throughout the configuration section of this chapter, many configuration examples are shown from actual routers and switches used in the lab build-out used to test and write this document. [Figure 6-13](#) provides a detailed view of the lab network, that will assist in understanding the configuration examples below.

Figure 6-13 TSE VoWLAN campus Architecture Detailed



## Campus EIGRP Routing Configuration

### Campus EIGRP Routing Configuration Common to all Routers

Table 6-2 shows the configuration commands that are common to all routers/switches in the example network. The IP addressing is based on the network diagram shown in Figure 6-1.

Table 6-2 Configuration Common to all Routers

Configuration Command (Switch D3L used for example configuration)	Description of Configuration
ip cef distributed	Enable CEF
key chain eigrp-chain key 100 key-string cisco	Configure EIGRP authentication parameters (Ensure a more secure key is used in production deployments)

**Table 6-2 Configuration Common to all Routers (continued)**

interface loopback1 description unique /32 interface on every router ip address 10.33.9.8 255.255.255.255	Provides a unique IP address for every router. Can be used by services like NTP and multicast. Also a good telnet address.
spanning-tree mode rapid-pvst	Enable spanning tree as a fail-safe practice
router eigrp 100 passive-interface Loopback1 network 10.0.0.0 no auto-summary eigrp router-id 10.33.9.8	<ul style="list-style-type: none"> <li>* Enable EIGRP; this example uses AS 100, but any valid AS number can be substituted</li> <li>* Passive all interfaces not intended to form EIGRP neighbors. Add any other non-routing interfaces.</li> <li>* Define the networks to route using EIGRP</li> <li>* Ensure all subnets are routed unless explicitly summarized</li> <li>* Explicitly configure the EIGRP router id as a best practice. Use the Loopback1 address for this.</li> </ul>
<b>A similar configuration is required on all links between core and distribution or distribution and access layers</b> interface TenGigabitEthernet2/2 description from D3L to CR ip address . 10.33.1.31 255.255.255.254 ip hello-interval eigrp 100 1 ip hold-time eigrp 100 3 ip authentication mode eigrp 100 md5 ip authentication key-chain eigrp 100 eigrp-chain load-interval 30 carrier-delay msec 0 mls qos trust dscp logging event link-status	<ul style="list-style-type: none"> <li>* Use of /31 addressing on point to point links optimizes use of IP address space in the campus</li> <li>* Tune EIGRP timers on interfaces between the Core and Distribution layers for faster convergence. Reduce EIGRP hello and hold timers to 1 and 3 seconds. In a point-point L3 campus design the EIGRP timers are not the primary mechanism used for link and node failure detection (physical detection of fiber breaks are). They are intended to provide a fail-safe mechanism only.</li> <li>* enable eigrp authentication on this interface</li> <li>* Reduce load-interval to set the length of time (in seconds; default is 300) for which data is used to compute load statistics on this interface</li> <li>* Reduce carrier delay to 0. Tuning carrier delay no longer has an impact on GigE or 10GigE interfaces but is recommended to be configured as a best practice for network operational consistency</li> <li>* Configure trust DSCP to provide for maximum granularity of internal QoS queuing</li> </ul>

**Note**

The base EIGRP configuration shown in the section above is common to all routers.

## EIGRP Configuration Specific to Core Routers

There is no special routing policy applied to the core routers; their configuration is left as simple as possible so as to not interfere with their primary function of routing packets.

## EIGRP Configuration Specific to Distribution Routers

**Table 6-3 EIGRP Configuration Specific to Distribution Routers**

Configuration Command (Switch D3L used for example configuration)	Description of Configuration
ip access-list standard only-default permit 0.0.0.0	Access-list used by eigrp to only send default route to access layer

**Table 6-3 EIGRP Configuration Specific to Distribution Routers (continued)**

<pre>router eigrp 100 distribute-list only-default out TenGigabitEthernet2/4 distribute-list only-default out TenGigabitEthernet2/6</pre>	Apply a distribute list filtering all routes other than select default(s) to the access switches
<p><b>On the interfaces to the core routers only</b></p> <pre>Interface TenGigabitEthernet2/1 ip summary-address eigrp 100 10.33.48.0 255.255.240.0 5</pre>	Advertise a summary address for the entire distribution block upstream to the core The "5" parameter is the administrative distance which is used to advertise a summary without installing it in the routing table.

## EIGRP Configuration Specific to Access Routers

**Table 6-4 EIGRP Configuration Specific to Access Routers**

Configuration Command (Switch A4R used for example configuration)	Description of Configuration
<pre>router eigrp 100 passive-interface default no passive-interface TenGigabitEthernet2/0/1 no passive-interface TenGigabitEthernet2/0/1 eigrp stub connected summary</pre>	<p>* Configure EIGRP as an EIGRP stub router; advertising connected routes upstream to the distribution.</p> <p>* Make passive-interface the default and explicitly remove it from the links to the distribution</p>

## Campus Layer 2 Spanning Tree Protocol Design for Data Center Configuration

**Table 6-5 Layer 2 Configuration Specific to Distribution Switch/Routers**

Configuration Command (Switch D1L used for example configuration)	Description of Configuration
<pre>vtp domain datacenter vtp mode transparent</pre>	Define a VTP domain, and make all switches in it transparent mode
<pre>spanning-tree mode rapid-pvst no spanning-tree optimize bpdu transmission spanning-tree extend system-id spanning-tree pathcost method long spanning-tree vlan 1-4094 priority 24576</pre>	Configure Rapid Per-VLAN Spanning tree. Make one of the Distribution layer switches the STP root by configuring it with priority 24576. Make the other Distribution layer switch the backup STP root by configuring it with priority 28672
<pre>router eigrp 100 passive-interface</pre>	Ensure routing updates are not sent out the Layer 2 VLANs
<pre>vlan 32 name DataCenter-Data32</pre>	define each VLAN



**Table 6-5 Layer 2 Configuration Specific to Distribution Switch/Routers**

<pre>interface Vlan32 description DataCenter-Data32 ip address 10.33.32.2 255.255.255.0 no ip redirects standby 1 ip 10.33.32.1 standby 1 timers 1 3 standby 1 priority 120 standby 1 preempt delay minimum 60</pre>	Configure each of the VLANs servicing the datacenter access layer with IP address. Make the same router STP root and HSRP primary by configuring it with priority 120. Make the other Distribution layer switch the HSRP secondary by configuring it with priority 115
<pre>interface TenGigabitEthernet3/6 description Layer2_connection_TO_D1R switchport switchport trunk encapsulation dot1q switchport trunk native vlan 2 switchport trunk allowed vlan 2,32-35 switchport mode trunk no ip address logging event link-status</pre>	Configure the Layer 2 interface between the 2 Distribution layer switches and between the Distribution layer switches and the Data Centers access layer switches. Trunk the Data Center VLANs

**Table 6-6 Layer 2 Configuration Specific to Data Center L2 Access Switches**

Configuration Command (Switch A3R used for example configuration)	Description of Configuration
<pre>vtp domain datacenter vtp mode transparent</pre>	Define a VTP domain, and make all switches in it transparent mode
<pre>spanning-tree mode rapid-pvst spanning-tree loopguard default spanning-tree portfast bpduguard default no spanning-tree optimize bpdu transmission spanning-tree extend system-id spanning-tree pathcost method long</pre>	Configure spanning tree. The priority statements on the distribution layer switches will ensure the one of them will always be the root
<pre>vlan 32 name DataCenter-Data32</pre>	define each VLAN
<pre>interface GigabitEthernet1/1 description All Ports assigned to access VLAN 32 switchport access vlan 32 switchport mode access no cdp enable spanning-tree portfast</pre>	Assign each access port to a VLAN
<pre>interface TenGigabitEthernet1/49 description to_D1L switchport trunk encapsulation dot1q switchport trunk native vlan 2 switchport mode trunk logging event link-status</pre>	Configure each of the uplinks to the distribution layer
<pre>interface Vlan32 ip address 10.33.32.4 255.255.255.0 exit ip default-gateway 10.33.32.1 ip route 0.0.0.0 0.0.0.0 10.33.32.1</pre>	Give 1 of the VLANs an IP address so the switch can be managed via Telnet and SNMP, and so that the switch can connect to the NTP server Define a default gateway and default route so the switch can communicate with IP devices on different subnets

## Campus QoS Configuration

The **auto qos voip** commands shown below are macros. They generate many configuration statements that configure CoS-to-DSCP mappings, interface queue discard thresholds, the mapping of specific CoS and DSCP values to a particular queue and threshold, and queue buffer sizes as well as generating the specific interface QoS trust policies. The commands generated can be observed by entering the **debug auto qos** command before beginning the configuration.

**Table 6-7** Campus QoS Configuration Common to all Switch/Routers

Configuration Command	Description of Configuration
debug auto qos	Enable debugging for auto-QoS. When debugging is enabled, the switch displays the QoS configuration that is automatically generated when auto-QoS is enabled.
<b>For uplinks/downlinks (interfaces connected directly to other switch/routers)</b> interface <i>interface-id</i> auto qos trust dscp end	<ul style="list-style-type: none"> <li>* Specify the switch port identified as connected to a trusted switch or router, and enter interface configuration mode</li> <li>* Configure inter-switch links to trust the DSCP settings that have been marked by auto-qos at the network edge.</li> <li>* Return to privileged EXEC mode.</li> </ul>

**Table 6-8** Campus QoS Configuration for Access Layer Switch/Routers

Configuration Command	Description of Configuration
debug auto qos	Enable debugging for auto-QoS. When debugging is enabled, the switch displays the QoS configuration that is automatically generated when auto-QoS is enabled.
<b>For interfaces connected directly to Cisco IP Phones</b> cdp enable interface <i>interface-id</i> auto qos voip cisco-phone exit	<ul style="list-style-type: none"> <li>* Enable CDP globally. By default, it is enabled.</li> <li>* Specify the switch port connected to the Cisco IP Phone, and enter interface configuration mode.</li> <li>* Enable auto-QoS on the port, and specify that the port is connected to a Cisco IP Phone. The QoS labels of incoming packets are trusted only when the Cisco IP Phone is detected.</li> <li>* Return to global configuration mode.</li> <li>* Repeat for as many ports as are connected to Cisco IP Phones.</li> </ul>
<b>For interfaces connected to directly Cisco LWAPP APs</b> interface <i>interface-id</i> auto qos voip trust mls qos trust dscp exit	<ul style="list-style-type: none"> <li>* Specify the switch port connected to the Cisco LWAPP, and enter interface configuration mode.</li> <li>* Enable auto-QoS on the port, and specify that the port is to trust QoS and prioritize voice traffic.</li> <li>* With the auto qos voip trust macro, in addition to setting interface queuing parameters, the switch automatically sets the ingress classification to trust the CoS value received in the packet on a nonrouted port or to trust the DSCP value received in the packet on a routed port. We use this macro to have the interface queuing parameters automatically set, but an LWAPP AP sets its QoS values in DSCP bits not CoS bits, so we need to change the generated <i>mls qos trust cos</i> to <i>mls qos trust dscp</i></li> <li>* Return to global configuration mode.</li> <li>* Repeat for as many ports as are connected to Cisco LWAPP APs.</li> </ul>

## Campus Routed Multicast configuration

This section provides the configuration required for the campus network with the exception of the connection to the datacenter where slightly different configuration is required for the Layer 2 connections from the distribution to the access layers.

### Campus Multicast Configuration Common to all Multicast-enabled Routers

The voice/WAN Gateway routers and the Internet Gateway routers are not multicast enabled.

**Table 6-9** Campus Multicast Configuration Common to all Routers

Configuration Command (sample multicast groups used for illustration purposes)	Description of configuration
<code>ip multicast-routing</code>	Globally enable multicast
<code>interface Te1/0/1</code> <code>ip pim sparse-mode</code>	Enable PIM on all multicast-enabled interfaces (including VLAN interfaces)
<code>interface Loopback2</code> <code>description Garbage-CAN RP</code> <code>ip address 2.2.2.2 255.255.255.255</code>	Define a local loopback address to provide a sink hole route point for invalid multicast groups
<code>ip access-list standard GOOD-IPMC</code> <code>permit 239.1.2.3 0.0.0.0</code> <code>permit 230.230.0.0 0.0.255.255</code> <code>permit 239.192.240.0 0.0.3.255</code> <code>permit 239.192.248.0 0.0.3.255</code>	Explicitly define what multicast groups to be forwarded to the RP -- permit LWAPP multicast tunnel -- Permit Vocera traffic -- Permit music on hold traffic -- Permit IPTV medium-rate traffic
<code>ip pim rp-address 10.33.9.1 GOOD-IPMC</code> <code>override</code>	Send all traffic permitted by the GOOD-IPMC access-list to the anycast RP address
<code>ip pim rp-address 2.2.2.2</code>	Send all multicast traffic that does not match the multicast groups defined in GOOD-IPMC access-list to the locally-defined Garbage-Can RP.

### Campus Multicast Configuration for Core Routers


**Table 6-10** Campus Multicast Configuration for Core Routers

Configuration Command (Switch CL used for example configuration)	Description of Configuration
<code>interface Loopback0</code> <code>description shared MSDP local peer address</code> <code>ip address 10.33.9.1</code> <code>255.255.255.255</code>	Loopback0 is the shared IP address that is used by all routers providing RP functionality. All RPs can use this address concurrently. The MSDP protocol keeps the RPs in synch with SA information. The DRs use whichever RP its unicast routing protocol finds closest. The unicast routing protocol will automatically provide failover if a RP fails.
<code>interface Loopback1</code> <code>description unique router address (for MSDP/RP and other uses)</code> <code>ip address 10.33.9.2</code> <code>255.255.255.255</code> <code>ip pim sparse-mode</code>	Loopback1 is the unique IP address used to keep the redundant RPs in synch by exchanging Source Active (SA) information via the MSDP protocol.

**Table 6-10** Campus Multicast Configuration for Core Routers (continued)

ip msdp peer 10.33.9.3 connect-source Loopback1	Multicast Source Distribution Protocol (MSDP) is the protocol that ensures that all RPs which are sharing a single IP address receive updates from each other (as muticast sources join one or other of the RPs). The command defines the peer routers MSDP address (its loopback1) and uses this routers loopback1 as the source IP address
ip msdp cache-sa-state	Cache SA-pairs learnt from MSDP peer (even when there is no registered receiver for that multicast group); this sacrifices some router memory in order to reduce join latency.
ip msdp originator-id Loopback1	Tells an MSDP speaker that originates an SA message to use the IP address of the Loopback1 as the RP address in the SA message
ip access-list extended PERMIT-SOURCES permit ip 10.33.66.0 0.0.0.255 239.1.2.3 0.0.0.0 permit ip 10.33.65.0 0.0.0.255 230.230.0.0 0.0.255.255 permit ip 10.33.32.0 0.0.0.255 239.192.240.0 0.0.3.255 permit ip 10.33.32.0 0.0.0.255 239.192.248.0 0.0.3.255	Example list of multicast source/groups permitted to register with the RPs -- permit LWAPP multicast tunnel from controller management interfaces -- Permit Vocera traffic from WLAN Voice VLAN on controller -- Permit music on hold from the Data Center access VLAN -- Permit IPTV medium-rate traffic from the Data Center access VLAN
ip pim accept-register list PERMIT-SOURCES	Configure the RPs to only accept register messages from sources explicitly defined in the PERMIT-SOURCES access list

Figure 6-14 provides a logical view of the MSDP configuration of the core switches.

**Figure 6-14** Logical View of MSDP Configuration

## Campus Multicast Configuration for Distribution Routers

No multicast-specific configuration is required on distribution routers.

## Campus Multicast Configuration for Access Routers

**Table 6-11** Campus Multicast configuration for Access Routers and

Configuration Command (Switch A4R used for example configuration)	Description of Configuration
<code>ip pim spt-threshold infinity</code>	Reduces multicast state (S,G) from the leaf routers by keeping traffic on the shared tree. (Optional)
<code>interface VLAN 50</code> <code>ip pim query-interval 250 msec</code>	Increase the speed with which the router will detect other multicast routers on the same subnet; and then elect a designated router for IGMP queries and to send source registration messages to the rendezvous point (RP).
<code>IP igmp snooping VLAN 50</code> <code>immediate-leave</code>	Enables IGMP Fast-leave processing; upon receiving an "IGMP leave group" message, the switch immediately removes the interface from its Layer 2 forwarding table. Only enable this on vlans where only one host is connected to each Layer 2 LAN interface.

## Campus Multicast Configuration for Layer 2 Switches used for Data Center

This section provides the configuration required for the datacenter where slightly different configuration is required for the Layer 2 connections from the distribution to the access layers.

## Campus Multicast configuration for Layer 2 Distribution-Layer switches used for Datacenter

**Table 6-12** Campus Multicast configuration for Layer 2 Distribution-Layer Switches used for Data Center

configuration command (Switch D1L used for example config)	Description of configuration
<code>ip multicast-routing</code>	Globally enable multicast
<code>spanning-tree vlan 32 root primary</code> <code>interface VLAN 32</code> <code>standby 1 priority 120</code> <code>ip pim dr-priority 120</code>	The root STP bridge, the HSRP active node, and the PIM DR should all be on the same distribution switch for each specific VLAN.
<code>interface VLAN 32</code> <code>ip pim query-interval 250 msec</code>	Increase the speed with which the router will detect other multicast routers on the same subnet; and then elect a designated router for IGMP queries and to send source registration messages to the rendezvous point (RP).
<code>interface vlan32</code> <code>ip igmp snooping fast-leave</code>	Enables IGMP Fast-leave processing; upon receiving an "IGMP leave group" message, the switch immediately removes the interface from its Layer 2 forwarding table. Only enable this on vlans where only one host is connected to each Layer 2 LAN interface.
<code>interface VLAN 32</code> <code>ip pim sparse-mode</code>	Enable PIM on all multicast-enabled VLANs
<code>interface Loopback2</code> <code>description Garbage-CAN RP</code> <code>ip address 2.2.2.2 255.255.255.255</code>	Define a local loopback address to provide a sink hole route point for invalid multicast groups
<code>ip access-list standard GOOD-IPMC</code> <code>permit 239.1.2.3 0.0.0.0</code> <code>permit 230.230.0.0 0.0.255.255</code> <code>permit 239.192.240.0 0.0.3.255</code> <code>permit 239.192.248.0 0.0.3.255</code>	Explicitly define what multicast groups to be forwarded to the RP -- permit LWAPP multicast tunnel -- Permit Vocera traffic -- Permit music on hold traffic -- Permit IPTV medium-rate traffic

**Table 6-12** Campus Multicast configuration for Layer 2 Distribution-Layer Switches used for Data Center (continued)

ip pim rp-address 10.33.9.1 GOOD-IPMC override	Send all traffic permitted by the GOOD-IPMC access-list to the anycast RP address
ip pim rp-address 2.2.2.2	Send all multicast traffic that does not match the multicast groups defined in GOOD-IPMC access-list to the locally-defined Garbage-Can RP.

## Campus Multicast Configuration for Layer 2 AAccess-Layer Switches used for Data Center

No multicast-specific configuration is required on L2 access layer switches.

## Campus Multicast Configuration for Layer 2 Switches used for Data Center Verification

**Table 6-13** Verifying the L2 Multicast Configuration for the Data Center

configuration command	Description of configuration
show spanning-tree summary show standby brief show ip pim interface	These commands can be used to check that spanning-tree root, HSRP active, and PIM DR are all on the same router for a given VLAN

## Access Switch Catalyst Integrated Security Features (CISF) Configuration

The configuration below is applied to all access layer switches that will have end-users or APs connected to them. The port security configuration permits a maximum of 2 MAC addresses; this is enough for a Cisco IP phone with an attached PC to be connected. An LWAPP AP tunnels all client traffic to the LWAPP controller, so the switch it is connected to will only ever see the MAC and IP address of the AP on that port.

**Table 6-14** Switch Security Configuration

Configuration command	Description of Configuration
<b>DHCP Snooping</b> ip dhcp snooping vlan 50 51 no ip dhcp snooping information option ip dhcp snooping	Globally enable DHCP Snooping These commands will put every VLAN 50 and VLAN 51 port into "untrusted" state. If the path to the DHCP server is out one of these switch ports, that port should be manually set to "trusted"
<b>Dynamic Arp Inspection (DAI)</b> ip arp inspection vlan 50,51 ip dhcp snooping database tftp://10.33.32.10/tftpboot/cisco/a4r-dhcpd b	Globally enable Dynamic ARP Inspection (DAI), and make sure the DHCP snooping database is backed up. (DAI requires an accurate DHCP snooping database, and backing up to TFTP ensures its preservation - even if the switch reboots)

**Table 6-14** Switch Security Configuration (continued)

<b>Port Security</b> interface GigabitEthernet1/0/1 switchport access vlan 50 switchport mode access switchport voice vlan 51 switchport port-security switchport port-security maximum 2 switchport port-security violation restrict switchport port-security aging time 2 switchport port-security aging type inactivity	interface commands to enable Port Security
<b>IP Source Guard</b> interface GigabitEthernet1/0/1 ip verify source	Interface command to enable IP Source Guard
interface GigabitEthernet1/0/1 ip dhcp snooping limit rate 10	Interface command to limit the number of DHCP messages an interface can receive per second

## Router Time Services Configuration

**Table 6-15** Campus Router Time Services Configuration

Configuration Command	Description of Configuration
service timestamps debug datetime msec localtime service timestamps log datetime msec localtime	Add timestamps with millisecond resolution to debug and log messages; use the local timezone to represent the time
clock timezone PST -8 clock summer-time PDT recurring clock update-calendar	<ul style="list-style-type: none"> <li>* Internally, the router uses Universal Coordinated Time (AKA GMT); these commands configure the router to display time in a local time.</li> <li>* The <i>timezone command</i> uses a user-defined label and the local offset from UTC as parameters.</li> <li>* The <i>summer-time</i> command configures the router to automatically adjust the local-time displayed in accordance with daylight saving conventions; the default is for North America conventions, other regions daylight saving conventions can be explicitly configured if required</li> </ul>
ntp source Loopback1 ntp server 10.33.32.16	<ul style="list-style-type: none"> <li>* Source NTP Queries from Loopback1 interface instead of the interface that sends them. This ensures the NTP server always sees the same source IP and is useful if the NTP server limits the number of NTP Peers that can associate, or if it has Access Control Lists controlling what devices can sync with it.</li> <li>* The <i>NTP Server</i> command defines the IP address of the NTP server.</li> </ul>
show ntp status show clock details	Useful show commands for verifying correct NTP operation

## Uni-Directional Link Detection (UDLD) Configuration

**Table 6-16** *UniDirectional Link Detection (UDLD) Configuration*

configuration command	Description of configuration
uldp enable	Globally enables ULDP on all fiber-optic ports

UDLD is not supported on the ASAs or 3800s used in this document, and is enabled by default on the 4948s.