High Availability Campus Network Design - Routed Access Layer using EIGRP

Cisco Validated Design II

November 6, 2007

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Cisco Validated Design

The Cisco Validated Design Program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit www.cisco.com/go/validateddesigns.

High Availability Campus Network Design - Routed Access Layer using EIGRP

Preface

Document Purpose

This document presents recommendations and results for the CVDII validation of High Availability Campus Network Design - Routed Access Layer using EIGRP.

Definitions

This section defines words, acronyms, and actions that may not be readily understood.

Table 1 Acronyms and Definitions

CSSC	Cisco Secure Service Client
CTI	Common Test Interface
CUCM	Cisco Unified Communication Manager
CUWN	Cisco Unified Wireless Network
CVD	Cisco Validated Design
DR	Designated Router
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name Service
EIGRP	Enhanced Interior Gateway Routing Protocol
FTP	File Transfer Protocol
HA	High Availability
HTTP	Hyper Text Transfer Protocol
IAM	Information Access Manager
IGP	Interior Gateway Protocol
IGMP	Internet Group Management Protocol
LWAPP	Light Weight Access Point Protocol
MSDP	Multicast Source Discovery Protocol
NSITE	Network Systems Integration and Test Engineering

Table 1 Acronyms and Definitions

NTP	Network Time Protocol
PIM	Protocol Independent Multicast
PIM-Bidir	Protocol Independent Multicast - Bidirectional
PSQM	Perceptual Speech Quality Measurement
POP3	Post Office Protocol 3
QoS	Quality of Service
RP	Rendezvous Point
SCCP	Skinny Call Control Protocol
SPT	Shortest Path Tree
SIP	Session Initiation Protocol
TFTP	Trivial File Transfer Protocol
VLAN	Virtual Local Area Network
WLAN	Wireless Local Area Network
WLC	WLAN Controller
WiSM	Wireless Service Module for Catalyst 6500

CONTENTS

FIGURES

This page is intentionally left blank

TABLES

This page is intentionally left blank

CHAPTER 1

Cisco Validated Design Program

The Cisco® Validated Design Program (CVD) consists of systems and solutions that are designed, tested, and documented to facilitate faster, more reliable and more predictable customer deployments. These designs incorporate a wide range of technologies and products into a broad portfolio of solutions that meet the needs of our customers. There are two levels of designs in the program: Cisco Validated Design I and Cisco Validated Design II.

1.1 Cisco Validated Design I

Cisco Validated Design I are systems or solutions that have been validated through architectural review and proof-of concept testing in a Cisco lab. Cisco Validated Design I provide guidance for the deployment of new technology or in applying enhancements to existing infrastructure.

1.2 Cisco Validated Design II

The Cisco Validated Design II (CVD II) is a program that identifies systems that have undergone architectural and customer relevant testing. Designs at this level have met the requirements of a CVD I Validated design as well as being certified to a baseline level of quality that is maintained through ongoing testing and automated regression for a common design and configuration. Certified designs are architectural best practices that have been reviewed and updated with appropriate customer feedback and can be used in pre- and post-sales opportunities. Certified designs are supported with forward looking CVD roadmaps and system test programs that provide a mechanism to promote new technology and design adoption. CVD II Certified Designs advance Cisco System's competitive edge and maximize our customers' return on investment while ensuring operational impact is minimized.

A CVD II certified design is a highly validated and customized solution that meets the following criteria:

- Reviewed and updated for general deployment
- Achieves the highest levels of consistency and coverage within the Cisco Validated Design program
- Solution requirements successfully tested and documented with evidence to function as detailed within a specific design in a scaled, customer representative environment
- Zero observable operation impacting defects within the given test parameters , that is, no defects that have not been resolved either outright or through software change, redesign, or workaround (refer to test plan for specific details)
- A detailed record of the testing conducted is generally available to customers and field teams, which provides:

- Design baseline that provides a foundational list of test coverage to accelerate a customer deployment
- Software baseline recommendations that are supported by successful testing completion and product roadmap alignment
- Detailed record of the associated test activity that includes configurations, traffic profiles, memory and CPU profiling, and expected results as compared to actual testing results

For more information on Cisco CVD program, refer to:

http://www.cisco.com/go/cvd

Cisco's Network System Integration and Test Engineering NSITE team conducted CVD II testing for this program. NSITE's mission is to system test complex solutions spanning multiple technologies and products to accelerate successful customer deployments and new technology adoption.

CHAPTER 2

Executive Summary

This document validates the High Availability Campus Routed Access Design using EIGRP as IGP in the core, distribution and access layers and provides implementation guidance for EIGRP to achieve faster convergence.

Deterministic convergence times of less than 200 msec were measured for any redundant links or nodes failure in an equal-cost path in this design.

NSITE is currently validating OSPF as the IGP in routed access campus network and will publish details once validation is complete.

The aim of this solution testing is to accelerate customer deployments of this campus routed access design by validating in an environment where multiple integrated services like multicast, voice and wireless interoperate.

Extensive manual and automated testing was conducted in a large scale, comprehensive customer representative network. The design was validated with a wide range of system test types, including system integration, fault and error handling, redundancy, and reliability to ensure successful customer deployments. An important part of the testing was end-to-end verification of multiple integrated services like voice, and video using components of the Cisco Unified Communications solution. Critical service parameters such as packet loss, end-to-end delay and jitter for voice and video were verified under load conditions.

As an integral part of the CVDII program, an automated sustaining validation model was created for an on-going validation of this design for any upcoming IOS software releases on the targeted platforms. This model significantly extends the life of the design, increases customer confidence and reduces deployment time.

Table 2-1 CVDII Publication Status

Design Guide	Status
High Availability Campus Network Design - Routed Access Layer Using EIGRP	Passed

The following guide (CVD I) was the source for this validation effort:

High Availability Campus Network Design-Routed Access Layer using EIGRP or OSPF

This page is intentionally left blank

CHAPTER 3

High Availability Campus Routed Access with EIGRP

3.1 Test Coverage

3.1.1 Solution Overview

The hierarchical design segregates the functions of the network into separate building blocks to provide for availability, flexibility, scalability, and fault isolation. The distribution block provides for policy enforcement and access control, route aggregation, and the demarcation between the Layer 2 subnet (VLAN) and the rest of the Layer 3 routed network. The core layers of the network provide high capacity transport between the attached distribution building blocks.

Figure 3-1 High Availability Campus Routed Access Design - Layer 3 Access

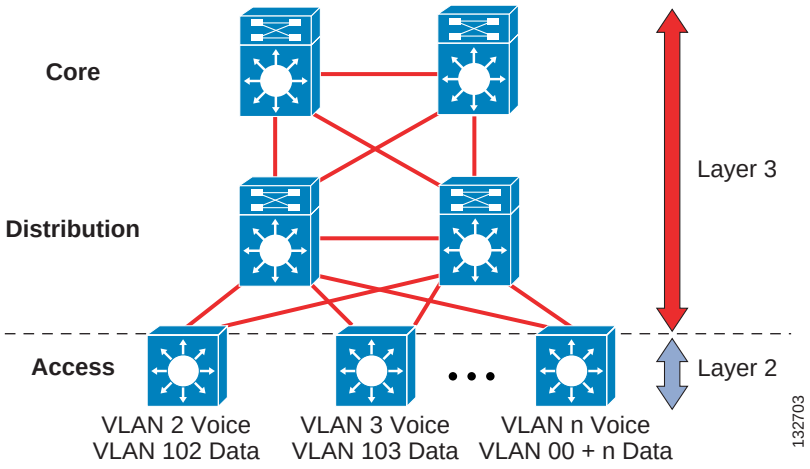

For campus designs requiring a simplified configuration, common end-to-end troubleshooting tools and fastest convergence, a distribution block design using Layer 3 switching in the access layer (routed access) in combination with Layer 3 switching at the distribution layer provides the fastest restoration of voice and data traffic flows.

Many of the potential advantages of using a Layer 3 access design include the following:

- Improved convergence

- Simplified multicast configuration
- Dynamic traffic load balancing
- Single control plane
- Single set of troubleshooting tools (eg. ping and traceroute)

Of these, perhaps the most significant is the improvement in network convergence times possible when using a routed access design configured with EIGRP or OSPF as the routing protocol. Comparing the convergence times for an optimal Layer 2 access design against that of the Layer 3 access design, four fold improvement in convergence times can be obtained, from 800-900msec for Layer 2 design to less than 200 msec for the Layer 3 access.

Figure 3-2 Comparison of Layer 2 and Layer 3 Convergence

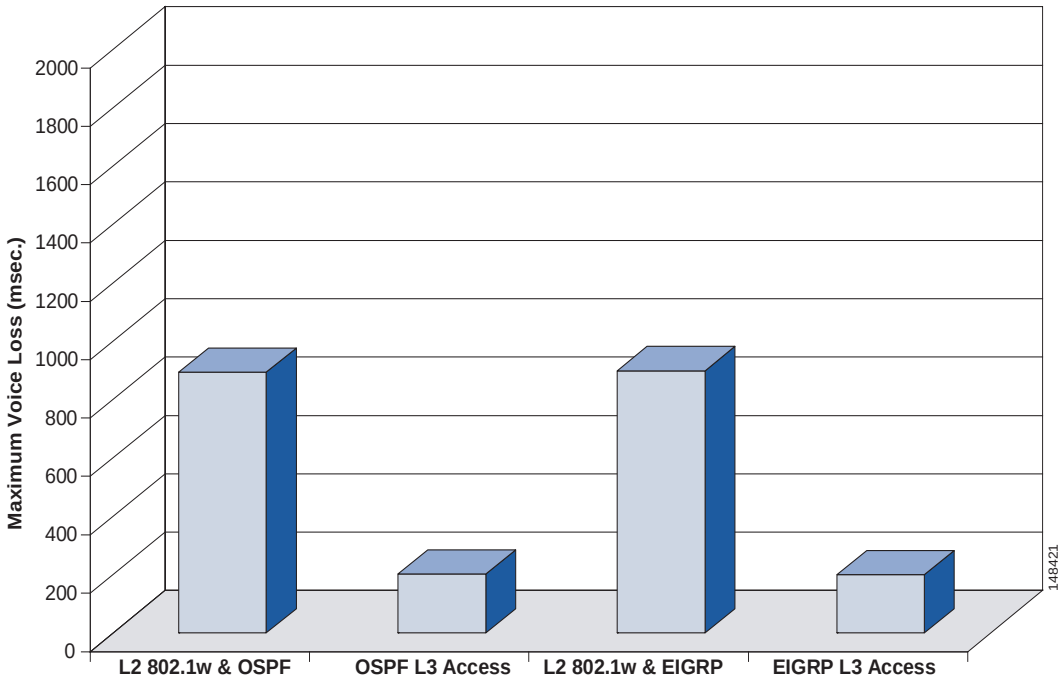

Note Convergence details in Figure 3-2 above are from the CVD-1 document. Hence, they include convergence times for EIGRP as well as OSPF.

In this phase, convergence time for EIGRP has been verified. NSITE is currently validating OSPF as the IGP in routed access campus network. Convergence time for OSPF will be confirmed once the validation is complete.

Although the sub-second recovery times for the Layer 2 access designs are well within the bounds of tolerance for most enterprise networks, the ability to reduce convergence times to a sub-200 msec range is a significant advantage of the Layer 3 routed access design.

For those networks using a routed access (Layer 3 access switching) within their distribution blocks, Cisco recommends that a full-featured routing protocol such as EIGRP or OSPF be implemented as the campus Interior Gateway Protocol (IGP). Using EIGRP or OSPF end-to-end within the campus provides

faster convergence, better fault tolerance, improved manageability, and better scalability than a design using static routing or RIP, or a design that leverages a combination of routing protocols (for example, RIP redistributed into OSPF).

3.1.2 Redundant Links

The most reliable and fastest converging campus design uses a tiered design of redundant switches with redundant equal-cost links. A hierarchical campus using redundant links and equal-cost path routing provides for restoration of all voice and data traffic flows in less than 200 msec in the event of either a link or node failure without having to wait for a routing protocol convergence to occur for all failure conditions except one (see section 3.1.3 Route Convergence, on page 11 for an explanation of this particular case).Figure 3-3 shows an example of equal-cost path traffic recovery.

Figure 3-3 Equal-cost Path Traffic Recovery

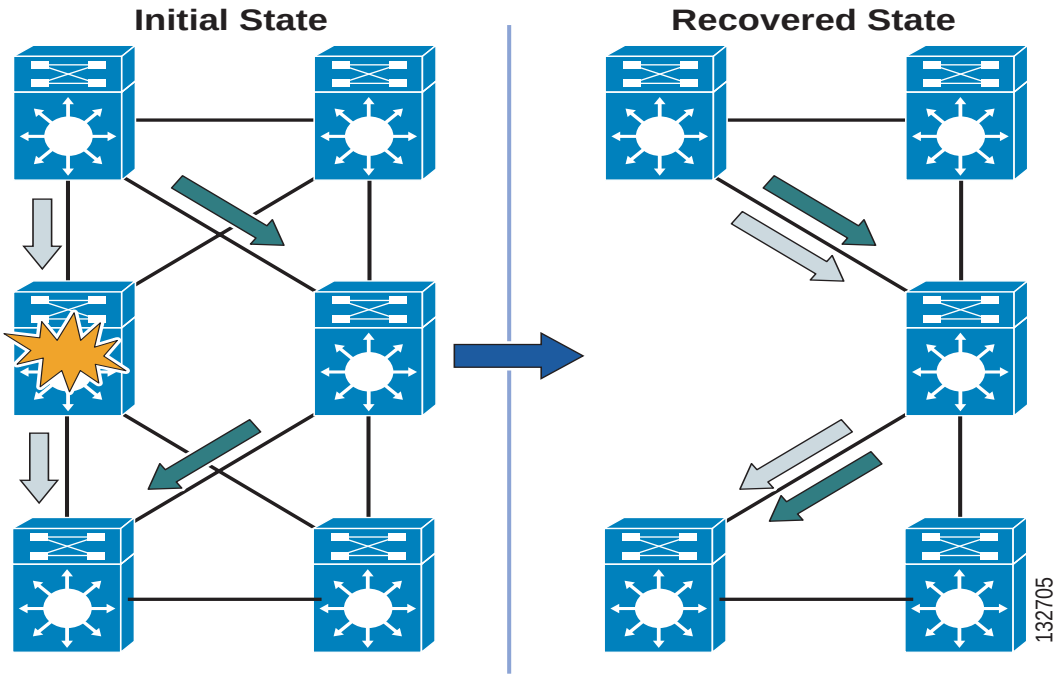

In the equal-cost path configuration, each switch has two routes and two associated hardware Cisco Express Forwarding (CEF) forwarding adjacency entries. Before a failure, traffic is being forwarded using both of these forwarding entries. On failure of an adjacent link or neighbor, the switch hardware and software immediately remove the forwarding entry associated with the lost neighbor. After the removal of the route and forwarding entries associated with the lost path, the switch still has a remaining valid route and associated CEF forwarding entry. Because the switch still has an active and valid route, it does not need to trigger or wait for a routing protocol convergence, and is immediately able to continue forwarding all traffic using the remaining CEF entry. The time taken to reroute all traffic flows in the network depends only on the time taken to detect the physical link failure and to then update the software and associated hardware forwarding entries.

Cisco recommends that Layer 3 routed campus designs use the equal-cost path design principle for the recovery of upstream traffic flows from the access layer. Each access switch needs to be configured with two equal-cost uplinks, as shown in Figure 4. This configuration both load shares all traffic between the two uplinks as well as provides for optimal convergence in the event of an uplink or distribution node failure.

In the following example, the Layer 3 access switch has two equal-cost paths to the default route 0.0.0.0

Figure 3-4 Equal-cost Uplinks from Layer3 Access to Distribution Switches

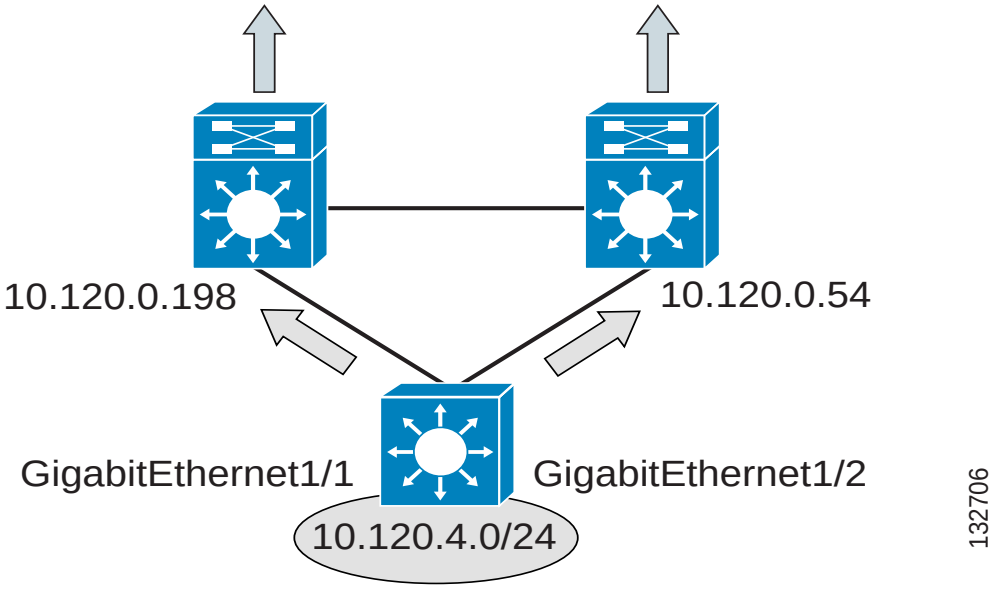

```
Layer3-Access#sh ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - OSPF, EX - OSPF external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       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 10.120.0.198 to network 0.0.0.0

     10.0.0.0/8 is variably subnetted, 5 subnets, 3 masks
C       10.120.104.0/24 is directly connected, Vlan104
C       10.120.0.52/30 is directly connected, GigabitEthernet1/2
C       10.120.4.0/24 is directly connected, Vlan4
C       10.120.0.196/30 is directly connected, GigabitEthernet1/1
D*EX 0.0.0.0/0 [170/5888] via 10.120.0.198, 00:46:00, GigabitEthernet1/1
               [170/5888] via 10.120.0.54, 00:46:00, GigabitEthernet1/2
```

3.1.3 Route Convergence

The use of equal-cost path links within the core of the network and from the access switch to the distribution switch allows the network to recover from any single component failure without a routing convergence, except one. As in the case with the Layer 2 design, every switch in the network has redundant paths upstream and downstream except each individual distribution switch, which has a single downstream link to the access switch. In the event of the loss of the fiber connection between a distribution switch and the access switch, the network must depend on the control plane protocol to restore traffic flows. In the case of the Layer 2 access, this is either a routing protocol convergence or a spanning tree convergence. In the case of the Layer 3 access design, this is a routing protocol convergence.

Figure 3-5 Traffic Convergence due to Distribution-to-Access Link Failure

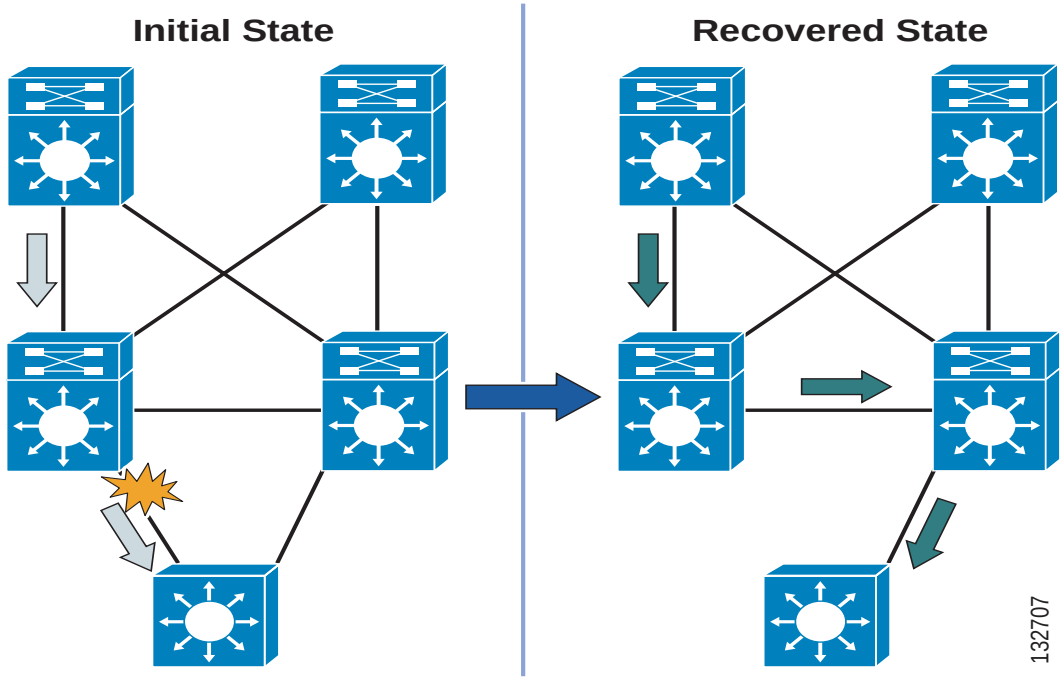

To ensure the optimal recovery time for voice and data traffic flows in the campus, it is necessary to optimize the routing design to ensure a minimal and deterministic convergence time for this failure case.

The length of time it takes for EIGRP, OSPF, or any routing protocol to restore traffic flows 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 (which is partially determined by the number of routers involved in determining the new path, or the number of routers that must be informed of the new path before the network can be considered converged).
- The time required to update software and associated CEF hardware forwarding tables with the new routing information.

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, EIGRP must determine a new route, and this process plays a large role in network convergence times.

In the case of EIGRP, the time is variable and primarily dependent on how many EIGRP queries the switch needs to generate and how long it takes for the response to each of those queries to return to calculate a feasible successor (path). The time required for each of these queries to be completed depends on how far they have to propagate in the network before a definite response can be returned. To minimize the time required to restore traffic flows, in the case where a full EIGRP routing convergence is required, it is necessary for the design to provide strict bounds on the number and range of the queries generated.

3.1.4 Link Failure Detection Tuning

The recommended best practice for campus design uses point-to-point fiber connections for all links between switches. In addition to providing better electromagnetic and error protection, fewer distance limitations and higher capacity fiber links between switches provide for improved fault detection. In a point-to-point fiber campus design using GigE and 10GigE fiber, remote node and link loss detection is normally accomplished using the remote fault detection mechanism implemented as a part of the 802.3z and 802.3ae link negotiation protocols. In the event of physical link failure, local or remote transceiver failure, or remote node failure, the remote fault detection mechanism triggers a link down condition that then triggers software and hardware routing and forwarding table recovery. The rapid convergence in the Layer 3 campus design is largely because of the efficiency and speed of this fault detection mechanism.

See IEEE standards 802.3ae and 802.3z for details on the remote fault operation for 10GigE and GigE respectively.

3.1.4.1 Carrier-delay Timer

Configure carrier-delay timer on the interface to a value of zero (0) to ensure no additional delay in the notification that a link is down. The default behavior for Catalyst switches is to use a default value of 0 msec on all Ethernet interfaces for the carrier-delay time to ensure fast link detection. It is still recommended as a best practice to hard code the carrier-delay value on critical interfaces with a value of 0 msec to ensure the desired behavior.

```
interface GigabitEthernet1/1
 description Uplink to Distribution 1
 ip address 10.120.0.205 255.255.255.252
 logging event link-status
 load-interval 30
 carrier-delay msec 0
```

Confirmation of the status of carrier-delay can be seen by looking at the status of the interface.

```
 GigabitEthernet1/1 is up, line protocol is up (connected)
 . . .
 Encapsulation ARPA, loopback not set
 Keepalive set (10 sec)
 Carrier delay is 0 msec
 Full-duplex, 1000Mb/s, media type is SX
 input flow-control is off, output flow-control is off
 . . .
```

Note On Catalyst 6500, "LINEPROTO-UPDOWN" message appears when the interface state changes before the expiration of the carrier-delay timer configured via the "carrier delay" command on the interface. This is an expected behavior on Catalyst 6500 and is documented in CSCsh94221. For details, refer to Appendix B.

3.1.4.2 Link Debounce Timer

It is important to review the status of the link debounce along with carrier delay configuration. By default, GigE and 10GigE interfaces operate with a 10 msec debounce timer that provides for optimal link failure detection. The default debounce timer for 10 / 100 fiber and all copper link media is longer than that for GigE fiber, and is one reason for the recommendation of a high-speed fiber deployment for

switch-to-switch links in a routed campus design. It is good practice to review the status of this configuration on all switch-to-switch links to ensure the desired operation via the command “show interfaces TenGigabitEthernet4/1 debounce.”

```
DistributionSwitch1#show interfaces tenGigabitEthernet 4/2 debounce

Port    Debounce time    Value(ms)
Te4/2    disable
```

The default and recommended configuration for debounce timer is "disabled", which results in the minimum time between link failure and notification of the upper layer protocols. Table 3.1 below lists the time delay that occurs before notification of a link change.

Table 3-1 Port Debounce Timer Delay Time

Port Type	Debounce Timer Disabled	Debounce Timer Enabled
Ports operation at 10 Mpbs or 100 Mpbs	300 milliseconds	3100 milliseconds
Ports operation at 1000 Mpbs or 10 Gbps over copper media	300 milliseconds	300 milliseconds
Ports operation at 1000 Mpbs or 10 Gbps over fiber media except WS-X6502-10GE	10 milliseconds	100 milliseconds
WS-X6502-10GE 10-Gigabit ports	1000 milliseconds	3100 milliseconds

For more information on the configuration and timer settings of the link debounce timer, see the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/122sx/swcg/intrface.htm#wp1044898

3.1.5 Features list

The validation coverage is outlined as follows:

- High Availability Campus Network design - Routed Access using EIGRP - EIGRP Stub, EIGRP timers tuning, EIGRP summarization, EIGRP route filters
- Multicast - PIM Sparse-mode, Static RP/Auto-RP, PIM bidir, MSDP, PIM Stub
- Wireless - Intra-controller and L2 Inter-controller Roaming, Voice over Wireless, dot1x authentication, WiSM
- Voice - SCCP, SIP, Delay/Jitter, PSQM
- Interoperability among multiple Cisco platforms, interfaces, and IOS releases
- Validation of successful deployment of actual applications (Cisco IP Telephony streams) in the network.
- End-to-end system validation of all the solutions together in a single integrated customer representative network

3.1.5.1 Implementing Routed Access using EIGRP

For those enterprise networks that are seeking to reduce dependence on spanning tree and a common control plane, are familiar with standard IP troubleshooting tools and techniques, and desire optimal convergence, a routed access design (Layer 3 switching in the access) using EIGRP as the campus routing protocol is a viable option. To achieve the optimal convergence for the routed access design, it is necessary to follow basic hierarchical design best practices and to use advanced EIGRP functionality, including stub routing, route summarization, and route filtering for EIGRP as defined in this document.

This section includes the following:

- EIGRP Stub
- Distribution Summarization
- Route Filters
- Hello and Hold Timer Tuning

3.1.5.1.1 EIGRP Stub

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 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 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.

```
router eigrp 100
    passive-interface default
    no passive-interface GigabitEthernet1/1
    no passive-interface GigabitEthernet1/2
    network 10.0.0.0
    no auto-summary
    eigrp router-id 10.120.4.1
    eigrp stub connected
```

The EIGRP stub feature when configured on all layer three access switches and routers prevents the distribution router from generating downstream queries.

By configuring the EIGRP process to run in the "stub connected" state, the access switch advertises all connected subnets matching the network 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 this design, the impact on the distribution switch is to limit the number of queries generated in case of a link failure.

3.1.5.1.2 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.

The summary-address statement is configured via command "**ip summary-address eigrp 100 10.120.0.0 255.255.0.0 5**" on the uplinks from each distribution switch to both core nodes. In the presence of any more specific route, say 10.120.1.0/24 address space, it causes EIGRP to generate a summarized route for the 10.120.0.0/16 network, and to advertise only that route upstream to the core switches.

```
interface TenGigabitEthernet4/1
 description Distribution 10 GigE uplink to Core 1
 ip address 10.122.0.26 255.255.255.254
 ip pim sparse-mode
 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
 ip summary-address eigrp 100 10.120.0.0 255.255.0.0 5
 mls qos trust dscp
```

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.

Summarization of directly connected routes is done on the distribution switches. Hence a layer3 link between the two distribution routers is required to exchange specific routes between them. This layer 3 link prevents the distribution switches from black holing traffic if either distribution switches lose the connection to the access switch.

Figure 3-6 Summarization towards the Core bounds EIGRP queries for Distribution block routes

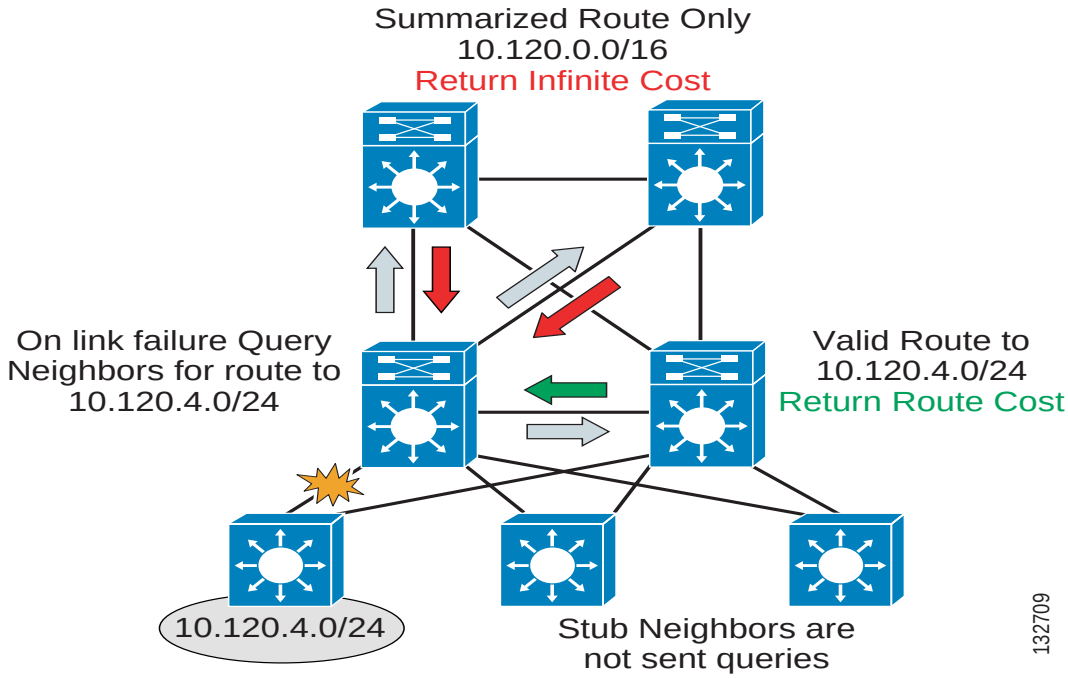

Using a combination of stub routing and summarizing the distribution block routes up-stream to the core both limits the number of queries generated and bounds those that are generated to a single hop in all directions. Keeping the query period bounded to less than 100 msec keeps the network convergence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

3.1.5.1.3 Route Filters

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 or "quad zero" route (0.0.0.0 mask 0.0.0.0) be the only route advertised to the access switches.

```
router eigrp 100
    network 10.120.0.0.0.255.255
    network 10.122.0.0.0.0.255
    ...
    distribute-list Default out GigabitEthernet3/3
    ...
    eigrp router-id 10.120.200.1

    !
    ip Access-list standard Default
        permit 0.0.0.0
```

3.1.5.1.4 Hello and Hold Timer Tuning

Cisco recommends in the Layer 3 campus design that the EIGRP hello and hold timers be reduced to one and three seconds, respectively. The loss of hellos and the expiration of the hold timer provide a backup to the L1/L2 remote fault detection mechanisms. Reducing the EIGRP hello and hold timers from defaults of five and fifteen seconds provides for a faster routing convergence in the rare event that L1/L2 remote fault detection fails to operate, and hold timer expiration is required to trigger a network convergence because of a neighbor failure.

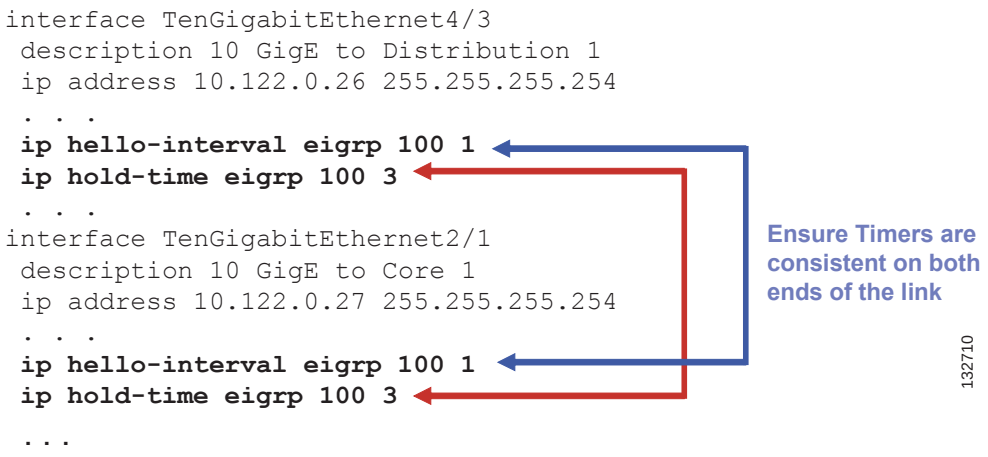

3.1.5.2 IP Multicast

IP multicast allows for a streamlined approach to data delivery whenever multiple hosts need to receive the same data at the same time. For example:

- When configured for IP multicast services, Music-on-Hold (MoH) can stream the same audio file to multiple IP phones without the overhead of duplicating that stream one time for each phone on hold.
- IP/TV allows for the streaming of audio, video, and slides to thousands of receivers simultaneously across the network. High-rate IP/TV streams that would normally congest a low-speed WAN link can be filtered to remain on the local campus network.

3.1.5.2.1 Multicast Forwarding

IP multicast delivers source traffic to multiple receivers using the least amount of network resources as possible without placing additional burden on the source or the receivers. Multicast packets are replicated in the network by Cisco routers and switches enabled with Protocol Independent Multicast (PIM) and other supporting multicast protocols.

Figure 3-7 Basic Multicast Service

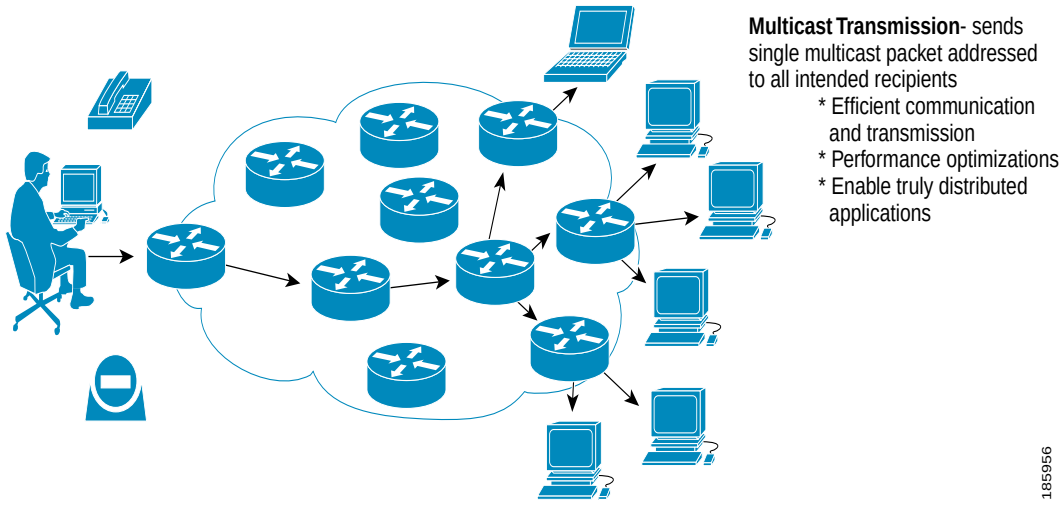

Multicast capable routers create "distribution trees" that control the path that IP Multicast traffic takes through the network in order to deliver traffic to all receivers. PIM uses any unicast routing protocol to build data distribution trees for multicast traffic.

The two basic types of multicast distribution trees are source trees and shared trees.

- Source trees-The simplest form of a multicast distribution tree is a source tree with its root at the source and branches forming a tree through the network to the receivers. Because this tree uses the shortest path through the network, it is also referred to as a shortest path tree (SPT).
- Shared trees-Unlike source trees that have their root at the source, shared trees use a single common root placed at some chosen point in the network. This shared root is called a Rendezvous Point (RP).

Figure 3-8 Shared Distribution Tree

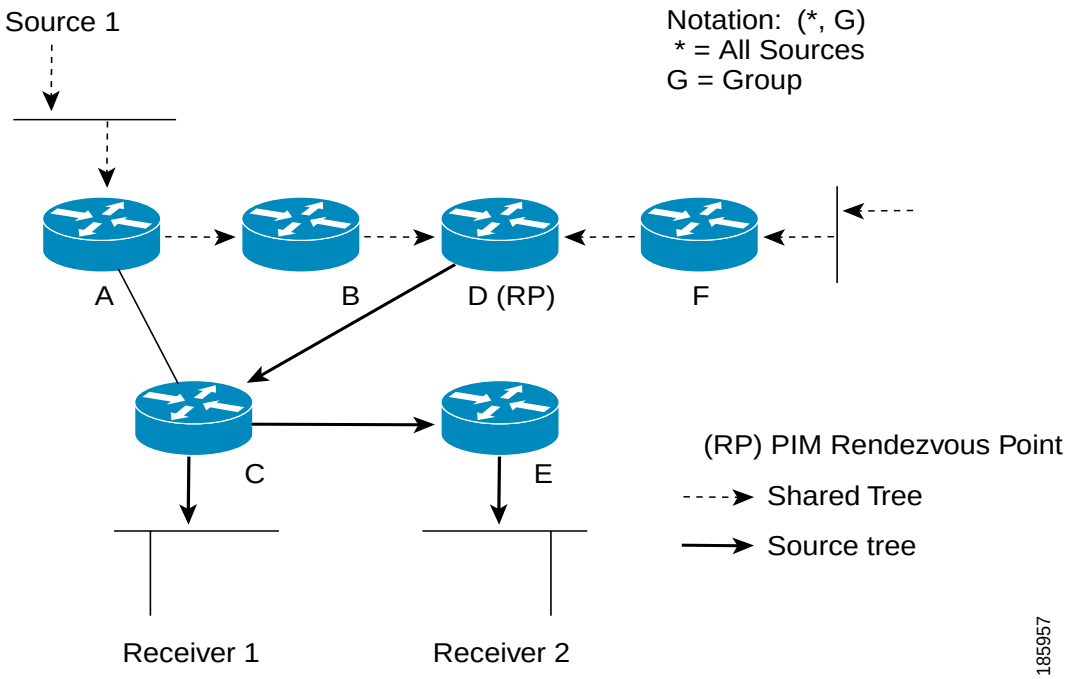

In the example above, the RP has been informed of Sources 1 and 2 being active and has subsequently joined the SPT to these sources.

PIM uses the concept of a designated router (DR). The DR is responsible for sending Internet Group Management Protocol (IGMP) Host-Query messages, PIM Register messages on behalf of sender hosts, and Join messages on behalf of member hosts.

3.1.5.2.2 Features of IP Multicast

The primary difference between multicast and unicast applications lies in the relationships between sender and receiver. There are three general categories of multicast applications:

- One to many, as when a single host sends to two or more receivers.
- Many-to-one refers to any number of receivers sending data back to a (source) sender via unicast or multicast. This implementation of multicast deals with response implosion typically involving two-way request/response applications where either end may generate the request.
- Many-to-many, also called N-way multicast, consists of any number of hosts sending to the same multicast group address, as well as receiving from it.

One-to-many are the most common multicast applications. The demand for many-to-many N-way is increasing with the introduction of useful collaboration and videoconferencing tools. Included in this category are audio-visual distribution, Webcasting, caching, employee and customer training, announcements, sales and marketing, information technology services and human resource information. Multicast makes possible efficient transfer of large data files, purchasing information, stock catalogs and financial management information. It also helps monitor real-time information retrieval as, for example, stock price fluctuations, sensor data, security systems and manufacturing.

3.1.5.2.3 PIM Sparse Mode

The PIM Sparse Mode is a widely deployed IP Multicast protocol and is highly scalable in Campus networks. *This mode is suitable for one-to-many (one source and many receivers) applications for Enterprise and Financial customers.*

PIM Sparse Mode can be used for any combination of sources and receivers, whether densely or sparsely populated, including topologies where senders and receivers are separated by WAN links, and/or when the stream of multicast traffic is intermittent.

- *Independent of unicast routing protocols* - PIM can be deployed in conjunction with any unicast routing protocol.
- *Explicit-join* - PIM-SM assumes that no hosts want the multicast traffic unless they specifically ask for it via IGMP. It creates a shared distribution tree centered on a defined "rendezvous point" (RP) from which source traffic is relayed to the receivers. Senders first send the data to the RP, and the receiver's last-hop router sends a join message toward the RP (explicit join).
- *Scalable* - PIM-SM scales well to a network of any size including those with WAN links. PIM-SM domains can be efficiently and easily connected together using MBGP and MSDP to provide native multicast service over the Internet.
- *Flexible* - A receiver's last-hop router can switch from a PIM-SM shared tree to a source-tree or shortest-path distribution tree whenever conditions warrant it, thus combining the best features of explicit-join, shared-tree and source-tree protocols.

In a PIM-SM environment, RPs (Rendezvous Point) act as matchmakers, matching sources to receivers. With PIM-SM, the tree is rooted at the RP not the source. When a match is established, the receiver joins the multicast distribution tree. Packets are replicated and sent down the multicast distribution tree toward the receivers.

Sparse mode's ability to replicate information at each branching transit path eliminates the need to flood router interfaces with unnecessary traffic or to clog the network with multiple copies of the same data. *As a result, PIM Sparse Mode is highly scalable across an enterprise network and is the multicast routing protocol of choice in the enterprise.*

For more details, refer to *Cisco AVVID Network Infrastructure IP Multicast Design*

http://www.cisco.com/application/pdf/en/us/guest/tech/tk363/c1501/ccmigration_09186a008015e7cc.pdf

3.1.5.2.4 PIM bidir

PIM bidir was simultaneously configured in addition to PIM-SM. Separate multicast streams for Bidir and PIM-SM were running at the same time and a few multicast receivers were configured to receive both, Bidir and PIM-SM streams.

In many-to-many deployments (many sources and many receivers) PIM bidir is recommended.

Bidir-PIM is a variant of the Protocol Independent Multicast (PIM) suite of routing protocols for IP multicast.

In bidirectional mode, traffic is routed only along a bidirectional shared tree that is rooted at the rendezvous point (RP) for the group. In bidir-PIM, the IP address of the RP acts as the key to having all routers establish a loop-free spanning tree topology rooted in that IP address. This IP address need not be a router, but can be any unassigned IP address on a network that is reachable throughout the PIM domain. Using this technique is the preferred configuration for establishing a redundant RP configuration for bidir-PIM.

Membership to a bidirectional group is signaled via explicit join messages. Traffic from sources is unconditionally sent up the shared tree toward the RP and passed down the tree toward the receivers on each branch of the tree.

Bidir-PIM is designed to be used for many-to-many applications within individual PIM domains. Multicast groups in bidirectional mode can scale to an arbitrary number of sources without incurring overhead due to the number of sources.

Bidir-PIM is derived from the mechanisms of PIM sparse mode (PIM-SM) and shares many shortest path tree (SPT) operations. Bidir-PIM also has unconditional forwarding of source traffic toward the RP upstream on the shared tree, but no registering process for sources as in PIM-SM. These modifications are necessary and sufficient to allow forwarding of traffic in all routers solely based on the (*, G) multicast routing entries. This feature eliminates any source-specific state and allows scaling capability to an arbitrary number of sources.

Figure 3-9 and Figure 3-10 show the difference in state created per router for a unidirectional shared tree and source tree versus a bidirectional shared tree.

Figure 3-9 Unidirectional Shared Tree and Source Tree

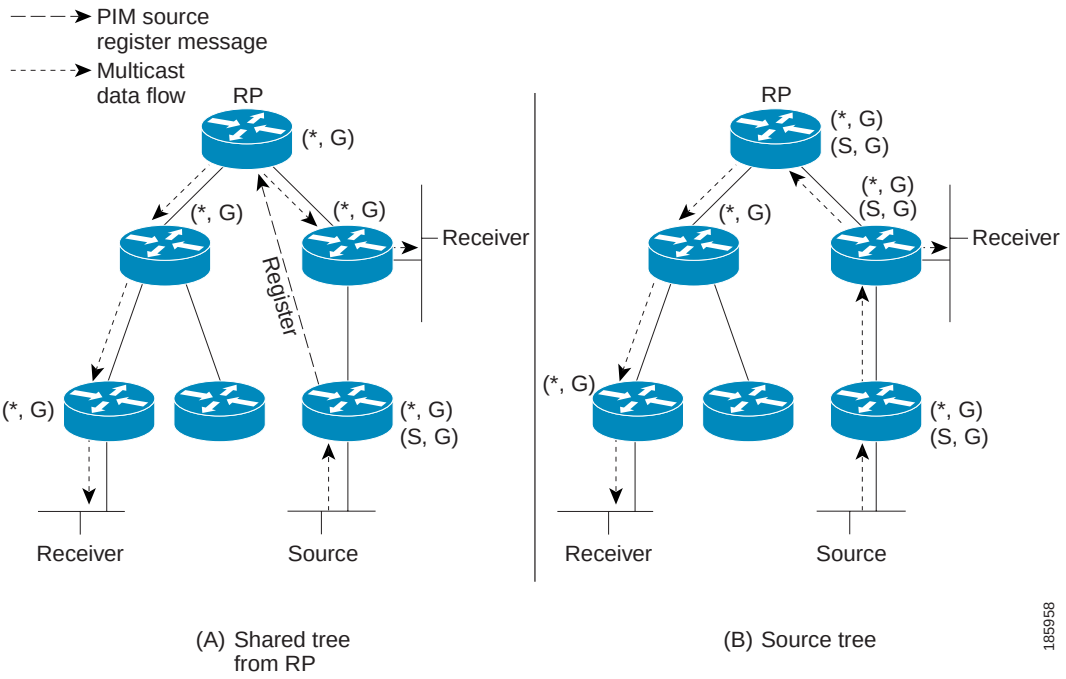

Figure 3-10 Bidirectional Shared Tree

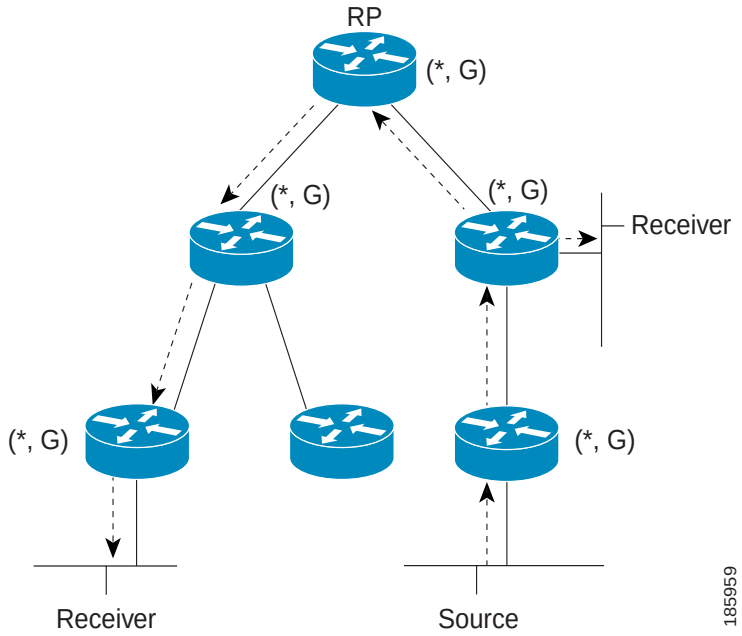

Main advantages with this mode are better support with intermittent sources and no need for an actual RP (works with Phantom RP). There is no need for MSDP for source information.

Note PIM bidir is currently not supported on Catalyst 4500 series.

For more details on PIM bidir refer to *Bidirectional PIM Deployment Guide.*

3.1.5.2.5 PIM Stub

Multicast control plane traffic is always seen by every router on a LAN environment. The Stub IP Multicast is used to reduce and minimize the unnecessary multicast traffic that is seen on LAN in the access layer and save the bandwidth on the media to forward multicast traffic to the upstream distribution/core layer.

In the Catalyst 3750 and 3560 Series Switches, the PIM Stub Multicast feature supports multicast routing between the distribution layer and access layer. This feature is currently available on Catalyst 3500/3700 platforms and restricts PIM control packets. This in turn helps reduce CPU utilization.

It supports two types of PIM interfaces: uplink PIM interfaces and PIM passive interfaces. In particular, a routed interface configured with the PIM Passive mode does not pass/forward PIM control plane traffic; it only passes/forwards IGMP traffic.

Complete these steps to configure PIM Stub Routing:

Step 1 Issue this command to enable multicast routing globally on thw switch or switch stack:

```
mix_stack(config)#ip multicast-routing distributed
```

Step 2 Issue this command to enable PIM SSM on the uplink:

```
mix_stack(config)#interface GigabitEthernet3/0/25
mix_stack(config-if)ip pim sparse-dense-mode
```

Step 3 Issue this command to enable PIM Stub Routing on the VLAN interface:

```
mix_stack(onfig)#inteface vlan100
mix_stack(config-if)ip pim passive
```

3.1.5.2.6 IGMP Snooping

IP multicast uses the host signaling protocol IGMP to indicate that there are multicast receivers interested in multicast group traffic.

Internet Group Management Protocol (IGMP) snooping is a multicast constraining mechanism that runs on a Layer 2 LAN switch. IGMP snooping requires the LAN switch to examine some Layer 3 information (IGMP join/leave messages) in the IGMP packets sent between the hosts and the router. When the switch hears the "IGMP host report" message from a host for a multicast group, it adds the port number of the host to the associated multicast table entry. When the switch hears the "IGMP leave group" message from a host, the switch removes the host entry from the table.

Note IGMP snooping is enabled by default and no explicit configuration is required.

IGMP v2 is widely deployed for PIM sparse as well as PIM bidir and therefore was implemented in our setup.

3.1.5.2.7 RP Deployment

Anycast RP is the preferred deployment model as opposed to a single static RP deployment. It provides for fast failover of IP multicast (within milliseconds or in some cases seconds of IP Unicast routing) and allows for load-balancing.

There are several methods for deploying RPs.

- RPs can be deployed using a single, static RP. This method does not provide redundancy or load-balancing and is not recommended.
- Auto-RP is used to distribute group-to-RP mapping information and can be used alone or with Anycast RP. Auto-RP alone provides failover, but does not provide the fastest failover nor does it provide load-balancing.
- Anycast RP is used to define redundant and load-balanced RPs and can be used with static RP definitions or with Auto-RP. *Anycast RP is the optimal choice as it provides the fast failover and load-balancing of the RPs.*

In the PIM-SM model, multicast sources must be registered with their local RP. The router closest to a source performs the actual registration. Anycast RP provides load sharing and redundancy across RPs in PIM-SM networks. It allows two or more RPs to share the load for source registration and to act as hot backup routers for each other (multicast only).

3.1.5.2.8 Anycast RP / MSDP

A very useful application of MSDP is Anycast RP. This is a technique for configuring a multicast Sparse Mode network to provide for fault tolerance and load sharing within a single multicast domain.

Two or more RPs are configured with the same IP address on loopback interfaces, say 10.0.0.1 for example:

Figure 3-11 Anycast RP

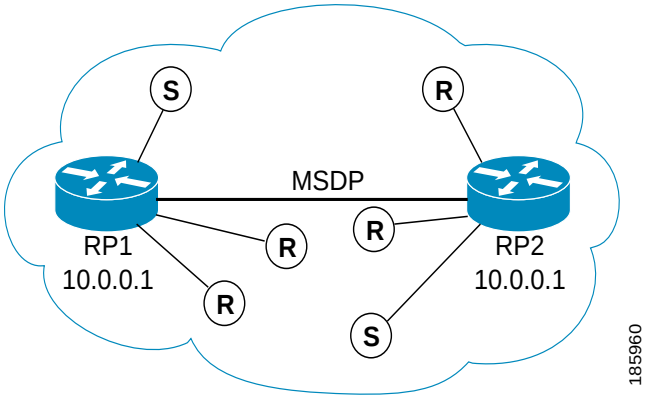

The loopback address should be configured as a 32 bit address. All the downstream routers are configured so that they know that their local RP's address is 10.0.0.1. IP routing automatically selects the topologically closest RP for each source and receiver. Since some sources might end up using one RP, and some receivers a different RP there needs to be some way for the RPs to exchange information about active sources. This is done with MSDP. All the RPs are configured to be MSDP peers of each other. Each RP will know about the active sources in the other RP's area. If any of the RPs was to fail, IP routing will converge and one of the RPs would become the active RP in both areas.

Note For Anycast RP configuration, create loopback1 interface for duplicate IP address on the RP routers and configure loopback0 interface with unique IP address used as router IDs, MSDP peer addresses etc.

MSDP

Multicast Source Discovery Protocol (MSDP) allows RPs to share information about active sources and is the key protocol that makes Anycast RP possible.

Sample configuration:

```
ip msdp peer 192.168.1.3 connect-source loopback 0
ip msdp cache-sa-state
ip msdp originator-id loopback0
```

3.1.5.2.9 Adjusting Timers for IP Multicast

Two timers can be adjusted to facilitate faster failover of multicast streams. The timers control the:

- PIM Query Interval
- Send-RP Announce Interval

PIM Query Interval

The **ip pim query-interval** command configures the frequency of PIM Router-Query messages. Router Query messages are used to elect a PIM DR. The default value is 30 seconds. For faster failover of multicast streams, Cisco recommends 1 second interval.

To verify the interval for each interface, issue the **show ip pim interface** command, as shown below.

```
svrL-dist#show ip pim interface
```

```
Address Interface Version/Mode Nbr   Query DR
Count Intvl
10.5.10.1 Vlan10 v2/Sparse0110.5.10.1
10.0.0.37 GigabitEthernet0/1 v2/Sparse1110.0.0.38
10.0.0.41 GigabitEthernet0/2 v2/Sparse1110.0.0.42
```

Send-RP Announce Interval

The **ip pim send-rp-announce** command has an interval option. Adjusting the interval allows for faster RP failover when using Auto-RP. The default interval is 60 seconds and the holdtime is 3 times the interval. So the default failover time is 3 minutes. The lower the interval, the faster the failover time.

Decreasing the interval will increase Auto-RP traffic but not enough to cause any kind of a performance impact. For faster failover time, Cisco recommends values of 3 to 5 seconds.

3.1.5.2.10 Multicast - Sources, Receivers, Streams

Multicast sources for PIM spare mode, PIM bidir and IPTV were connected to access switches in the Services Block.

Multicast receivers were connected to wiring closet switches.

Real IPTV streams (PIM-SM) and simulated multicast streams (PIM-SM and PIM-bidir) from traffic generators were part of the traffic profile during this validation.

3.1.5.3 Wireless

With WiSM and 4404 series as Cisco wireless controller, wireless deployment was verified in the HA Campus Routed Access environment with wireless AP's connected to Access switches. Clients authenticate using Dot1x with a Radius server as the authentication server. The Authenticator on the clients was Cisco Secure Service Client (CSSC).

The Cisco Unified Wireless Network (CUWN) architecture centralizes WLAN configuration and control into a device called a WLAN Controller (WLC). This allows the WLAN to operate as an intelligent information network and support advanced services, unlike the traditional 802.11 WLAN infrastructure that is built from autonomous, discrete entities. The CUWN simplifies operational management by collapsing large numbers of managed end-points-autonomous access points-into a single managed system comprised of the WLAN controller(s) and its corresponding, joined access points.

In the CUWN architecture, APs are "lightweight", meaning that they cannot act independently of a WLC. APs are "zero-touch" deployed and no individual configuration of APs is required. The APs learn the IP address of one or more WLC via a controller discovery algorithm and then establish a trust relationship with a controller via a "join" process. Once the trust relationship is established, the WLC will push firmware to the AP if necessary and a configuration. APs interact with the WLAN controller via the Lightweight Access Point Protocol (LWAPP).

3.1.5.3.1 Client Roaming

When a wireless client associates and authenticates to an AP, the AP's joined WLC places an entry for that client in its client database. This entry includes the client's MAC and IP addresses, security context and associations, and QoS context, WLAN and associated AP. The WLC uses this information to forward frames and manage traffic to and from the wireless client.

3.1.5.3.2 Intra-controller Roaming

The wireless client roams from one AP to another when both APs are associated with the same WLC. This is illustrated below.

Figure 3-12 Intra-controller roaming

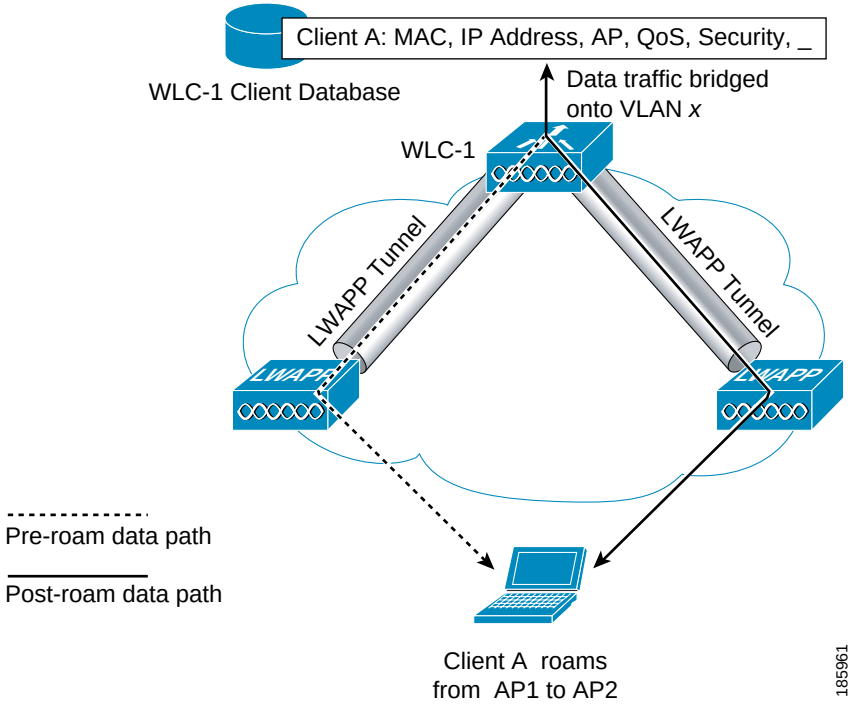

When the wireless client moves its association from one AP to another, the WLC simply updates the client database with the new associated AP.

3.1.5.3.3 Layer-2 Inter-Controller Roaming

The wireless client roams from an AP joined to one WLC and an AP joined to a different WLC.

Figure 3-13 L2 - Inter-controller roaming

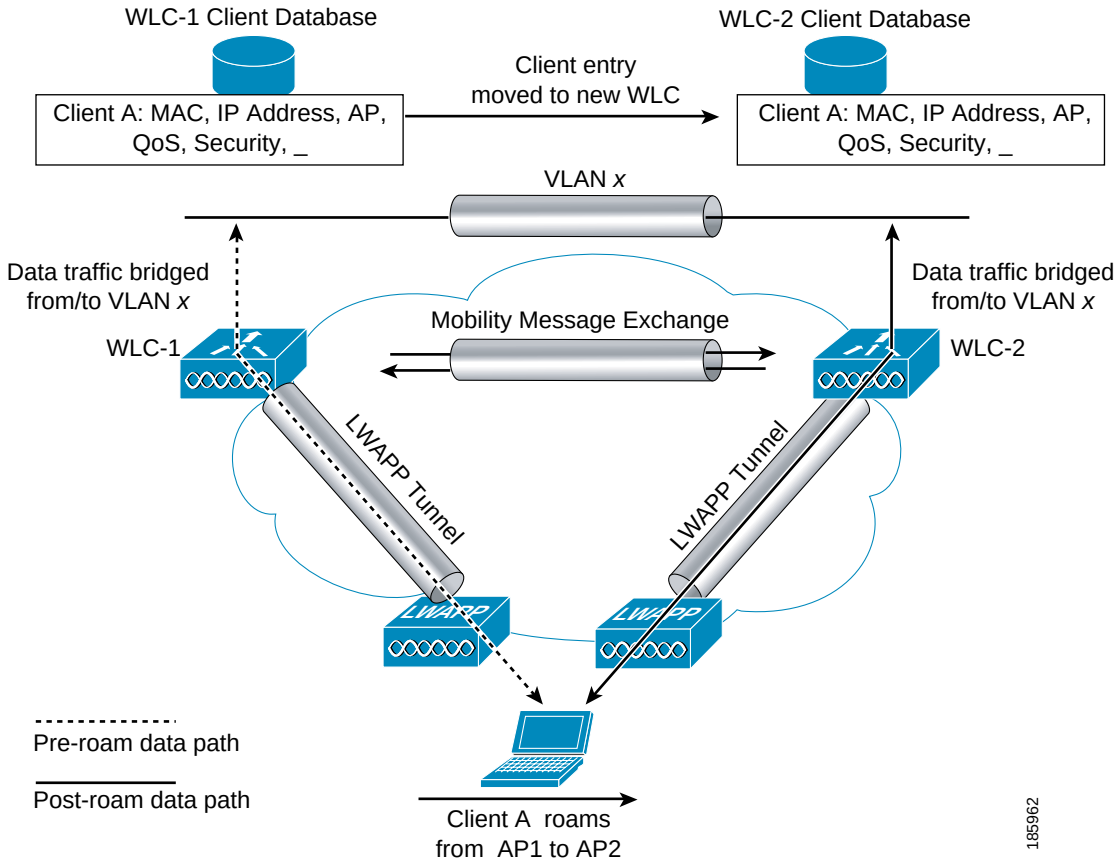

From the above, Layer 2 roam occurs when the controllers bridge the WLAN traffic on and off the same VLAN and the same IP subnet. When the client re-associates to an AP connected to a new WLC, the new WLC exchanges mobility messages with the original WLC and the client database entry is moved to the new WLC. New security context and associations are established if necessary and the client database entry is updated for the new AP. All of this is transparent to the end-user. Also, the client retains the IP address during this process.

3.1.5.3.4 WISM

The Cisco WiSM is a member of the Cisco wireless LAN controller family. It works in conjunction with Cisco Aironet lightweight access points, the Cisco WCS, and the Cisco wireless location appliance to deliver a secure and unified wireless solution that supports wireless data, voice, and video applications.

The Cisco WiSM consists of two Cisco 4404 controllers on a single module. The first controller is considered the WiSM-A card, while the second controller is considered WiSM-B card. Interfaces and IP addressing have to be considered on both cards independently. WiSM-A manages 150 access points, while WiSM-B manages a separate lot of 150 access points. These controllers can be grouped together in a mobility group, forming a cluster.

Wireless features were implemented in accordance with *Enterprise Mobility 3.0 Design Guide*.

Note Multicast over wireless was validated only using Cisco 4404 wireless controller. Due to problems described in CSCsj48453, Multicast over wireless using WiSM module for catalyst 6500 could not be verified. Due to this DDTS, Catalyst 6500 does not forward multicast traffic to WISM module when catalyst 6500 is configured in L3 mode. This is a severity 1 DDTS and is being worked on in the BU.

3.1.5.4 Voice over IP

The Cisco Unified Communications System delivers fully integrated communications by enabling data, voice, and video to be transmitted over a single network infrastructure using standards-based Internet Protocol (IP).

The foundation architecture for Cisco IP Telephony includes of the following major components:

3.1.5.4.1 Cisco IP Network Infrastructure

The Routed Access network design provides support for multiple clients such as hardware Cisco IP phones and video phones. Interface to PSTN network and traditional POTS phones are not in the scope of this design.

Campus LAN infrastructure design is extremely important for proper IP telephony operation on a converged network. Proper LAN infrastructure design requires following basic configuration and design best practices for deploying a highly available network.

Fast convergence of the network adds availability to the VoIP services.

- **Campus Access Layer**

  The access layer of the Campus LAN includes part of the network from the desktop port(s) to the wiring closet switch.

  Proper access layer design starts with assigning a single IP subnet per VLAN. Due to Routed Access network design, a VLAN cannot span multiple wiring closet switches; that is, a VLAN should have presence in one and only one access layer switch. More importantly, confining a VLAN to a single access layer switch also serves to limit the size of the broadcast domain. There is the potential for large numbers of devices within a single VLAN or broadcast domain to generate large amounts of broadcast traffic periodically, which can be problematic. A good rule of thumb is to limit the number of devices per VLAN to about 512, which is equivalent to two Class C subnets (that is, a 23-bit subnet masked Class C address). Typical access layer switches include the stackable Cisco Catalyst 3500 and 3700 series and the larger, higher-density Catalyst 4000 and 6000 switches.

- **Network Services**

  The deployment of an IP Communications system requires the coordinated design of a well structured, highly available, and resilient network infrastructure as well as an integrated set of network services including Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), Trivial File Transfer Protocol (TFTP), and Network Time Protocol (NTP).

3.1.5.4.2 Call Processing Agent

Cisco Unified Communications Manager (CUCM) is the core call processing software for Cisco IP Telephony solution. It builds call processing capabilities on top of the Cisco IP network infrastructure. CUCM software extends enterprise telephony features and capabilities to packet telephony network devices such as IP phones, media processing devices, voice gateways, and multimedia applications.

Typically, Cisco Unified Communications Manager cluster servers, including media resource servers, reside in a data center or services block or server farm environment.

- **Single-site Model**

In this testing, single-site call processing model was used.

The single-site model for IP telephony consists of a call processing agent located at a single site, or campus, with no telephony services provided over an IP WAN.

The single-site model has the following design characteristics:

- Single CUCM or CUCM cluster.
- Maximum of 30,000 Skinny Client Control Protocol (SCCP) or Session Initiation Protocol (SIP) IP phones or SCCP video endpoints per cluster.
- High-bandwidth audio (for example, G.711, G.722, or Cisco Wideband Audio) between devices within the site.
- High-bandwidth video (for example, 384 kbps or greater) between devices within the site. The Cisco Unified Video Advantage Wideband Codec, operating at 7 Mbps, is also supported.

3.1.5.4.3 Communication Endpoints

A communication endpoint is a user instrument such as a desk phone or even a software phone application that runs on a PC. In the IP environment, each phone has an Ethernet connection. IP phones have all the functions you expect from a telephone, as well as more advanced features such as the ability to access World Wide Web sites.

In this design, Cisco IP phones and Video telephones could be used for endpoints.

Under this infrastructure, SCCP phones were connected to access layer of the campus network, with IP addresses routable to CUCM. These phones get registered with CUCM which then devices a route plan to find and connect to the numbers dialed.

To implement voice that is representative of an enterprise customer network, following design guides were used together to reflect the complexity in the field. These guides were the design basis for the CVD II voice test suite. SCCP and SIP voice protocols were tested with Cisco Unified Call Manager.

- *Cisco Unified Communications SRND*

  http://www/en/US/products/sw/voicesw/ps556/products_implementation_design_guide_book09186a0080783d84.html

- *Guide to Cisco Systems' VoIP Infrastructure Solution for SIP*

  http://www/en/US/tech/tk652/tk701/technologies_configuration_guide_book09186a00800eaa0e.html

- *Cisco IOS SIP Configuration Guide*

  http://www/en/US/products/sw/iosswrel/ps5207/products_configuration_guide_book09186a00807517b8.html

3.2 HA Campus Routed Access Test Coverage Matrix - Features

Table 3-2 HA Campus Routed Access Test Coverage Matrix - Features

EIGRP as IGP	CVDI	CVDII
EIGRP Stub	✓	✓
Distribution Summarization	✓	✓
Route Filters	✓	✓
Hello and Hold timers tuning	✓	✓
EIGRP Authentication	✓	✓
MLS CEF loadsharing	✓	✓
Interface carrier-delay	✓	✓
PIM Sparse-mode	✓	✓
PIM bidir-mode		✓
IGMP Snooping		✓
Auto RP		✓
Accept-register filter		✓
Multicast limits		✓
MSDP / Anycast RP		✓
MSDP SA-filters		✓
Voice SCCP		✓
Voice SIP		✓
Wireless Dot1x		✓
Intra-controller Roaming		✓
L2 Inter-controller Roaming		✓
Voice over Wireless		✓

3.3 HA Campus Routed Access Test Coverage Matrix - Platforms

Table 3-3 HA Campus Routed Access Test Coverage Matrix - Platforms

Platform	Role	CVDI	CVDII
Cat6500	Core	✓	✓
Cat6500	Distribution	✓	✓
Cat4500	Distribution		✓
Cat6500	Access	✓	✓
Cat4500	Access		✓
Cat3750	Access	✓	✓
Cat3560	Access		✓
Cat3750E	Access		✓
Cat3560E	Access		✓

3.4 CVD II Test Strategy

Two sets of network device configurations were used to validate the HA Campus Routed Access design.

3.4.1 Baseline Configuration

The first set is the "Baseline configuration," as described in the *High Availability Campus Network Design-Routed Access Layer using EIGRP or OSPF* Design Guide recommendations.

Only convergence tests were executed with this set of baseline configuration.

3.4.2 Extended Baseline Configuration

The second set is "Extended Baseline configuration," which is the Baseline configuration with additional technologies such as Wireless and Multicast.

All tests, including convergence tests were executed using the Extended Baseline configuration.

3.4.3 Testbed Setup

The validation network consists of a three-layered Campus topology with access, distribution and core devices. The setup has three distribution blocks, user access and server access switches along with a set of core routers. The Cisco Unified Call Manager is connected to a server side distribution router for provisioning voice in the network. The wireless controller is also connected to the distribution router in services block. Wireless access points are connected to access switches.

Figure 3-14 High Availability Campus Routed Access design - Manual testbed

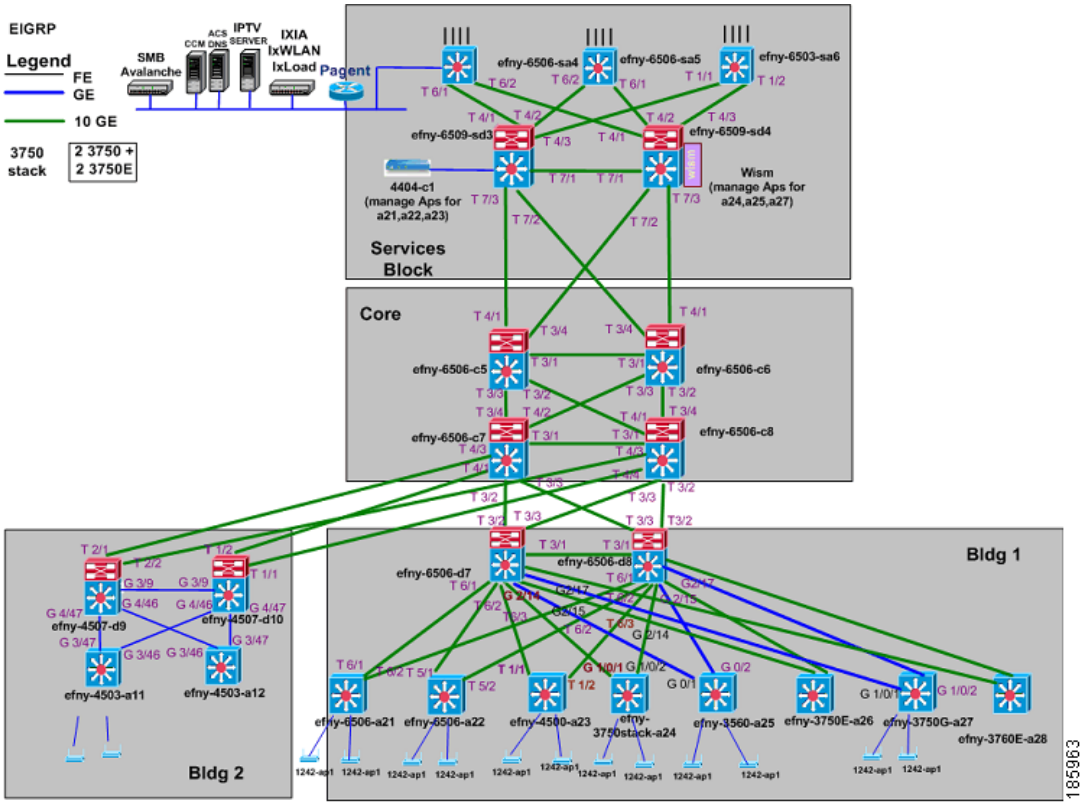

In addition, real world services such as the Cisco Call Manager, IP Phones, IPTV, Multicast, voice and Video were validated end-to-end.

3.4.3.1 Baseline Traffic

The following constitutes "Baseline traffic" that ran for every test that was executed:

- 3000 EIGRP routes (in Core)
- 5000 mroutes (in RP routers)
- 100 Stateful sessions per Access switch (TELNET + FTP + HTTP + DNS + POP3)
- 100 Mbps QoS traffic per Access switch (includes Voice, Multicast Video, Call Control, bulk data, critical data and best effort traffic) based on the *Enterprise QoS Solution Reference Network Design Guide*

3.4.4 Test Setup - Hardware and Software Device Information

Table 3-4 Hardware and Software Device Information

Hardware Platform	Role	DRAM	Software Version	Line Cards/Interfaces
Catalyst 6500 Sup720-3BXL	Core	RP - 1GB SP - 1GB	12.2(18)SXF8	WS-X6704-10GE - DFC-3BXL WS-X6748-GE-TX - DFC-3BXL
Catalyst 6500 Sup720-3BXL	Distribution	RP - 1GB SP - 1GB	12.2(18)SXF8	WS-X6704-10GE - DFC-3BXL WS-X6708-10GE - DFC-3BXL WS-X6724-SFP - DFC-3BXL WS-X6748-SFP - DFC-3BXL WS-SVC-WISM
Catalyst 6500 Sup720-3BXL	Access	RP - 1GB SP - 1GB	12.2(18)SXF8	WS-X6748-GE-TX - DFC-3BXL WS-X6704-10GE - DFC-3BXL
Catalyst 6500 Sup32-3B	Access	RP - 512MB SP - 512MB	12.2(18)SXF8	WS-X6548V-GE-TX
Catalyst 4500 Sup V-10GE	Access	512 MB	12.2(31)SGA	WS-X4548-GB-RJ45
Catalyst 3750 / 3750E and mixed stack	Access	256 MB	12.2(37)SE	WS-C3750-24P WS-C3750E-24P
Catalyst 3750G	Access	128 MB	12.2(37)SE	WS-C3750G-24PS
Catalyst 3560 / 3560E	Access	256 MB	12.2(37)SE	WS-C3560G-24PS WS-C3560E-24PD
AIR-(L)AP1242	Access-point		4.0.206.0	
AIR-WLC4404-K9	Wireless controller		4.0.206.0	
CCM	Cisco Call Manager		5.1	
ACS	Access control server		4.1	

3.4.5 Test Types

Validation tests are divided into the following categories:

- System Integration
- Negative/Redundancy
- Reliability

System integration and negative/redundancy tests were executed in manual as well as automated regression testing. Reliability tests were executed manually only.

3.4.5.1 System Integration Test

System Integration has two major components, feature combination and feature interaction.

Feature combination focuses on testing a feature when various combinations of other features are enabled.

Feature interaction tests were conducted to verify dependencies between features.

The following test suites were executed using extended baseline configuration:

- Routing - EIGRP
- Multicast
- Voice
- Wireless

The System Integration Tests combines all the features required for multiple features inter-operability. End-to-end service validation was performed for IP, Multicast, voice and video traffic. The services validated include Multicast using IPTV viewer, IP Telephony using Cisco IP Phones and data connectivity and Wireless using IXWLAN.

Health checks were performed before and after tests. These checks included memory and CPU utilization, tracebacks, memory alignment errors, deviations in number of routes and mroutes, interface errors, line card status and syslog messages.

All test cases under the System Integration Test were automated and test cases were executed in parallel.

3.4.5.2 Redundancy Test

Negative testing concerns error handling and robustness. Erroneous inputs were applied at the system level to verify behavior against error handling specifications. Unspecified inputs or conditions and faults were injected to evaluate system level robustness.

All the negative test cases were grouped together for better test management. During iterations of the negative tests, traffic was fully loaded and CPU and memory usage of the devices in the testbed were monitored.

The negative tests were categorized under the following failure scenarios:

- **Redundancy / High Availability:** Redundant router/link failover. (Reload the primary router and shut/no shut the links) for measuring convergence times with voice, Multicast and Unicast traffic streams.
- **Hardware:** Interface shut/no shut and monitor CPU spikes and memory utilization.
- **Control-plane:** Clear routing tables, flap routes, MSDP SA filters, Accept-Register filters and PIM Neighbor filters.

All test cases under the Redundancy Test were automated and test cases were executed serially.

3.4.5.3 Reliability Test

150-hour reliability test was executed for the entire testbed to ensure that the various solutions interoperate without memory or CPU issues or any operationally impacting defects. The total number of routes and mroutes are verified on UUT devices. Devices were monitored for tracebacks, alignment and interface errors, and syslogs for any error messages. End-to-end connectivity was maintained during this test.

During this phase of reliability testing, System Integration Test Suites were executed in parallel.

3.4.6 NSITE Sustaining Coverage

NSITE sustaining provides consistent, repeatable customer representative coverage of NSITE validated systems.

Except for reliability test cases all other test case were automated.

Sustaining test coverage included the following components:

- Automated test scripts for each automation test case
- Common library for managing the testbed, collecting and reporting test results
- Automated procedures to capture the manual execution results

All the real applications used in the manual validation phase, including IPTV server/client, Cisco Call Manager server and IP phones, were not automated. Instead, traffic tools were used to generate simulated voice and video traffic on the network.

3.5 CVD II - Feature Implementation Recommendations

These recommendations are based on:

- Cisco recommended Best Practices in various technologies.
- Years of field experience from Cisco engineers who work with complex networks and many of the largest customers.
- Issues encountered and successfully resolved during validation.

3.5.1 Routing

EIGRP Stub: Configure on all Access routers under **`router eigrp`**

Configuration:
```
router eigrp 100
 eigrp stub connected
```

Summarization: Route summarization is done at the distribution routers on interfaces (uplink to core router) connecting the distribution routers to core.

Configuration:
```
interface TenGigabitEthernet4/1
 ip summary-address eigrp 100 10.120.0.0 255.255.0.0 5
```

Note Summarization of directly connected routes is done on the distribution switches. Hence a layer3 link between the two distribution routers is required to exchange specific routes between them. This layer 3 link prevents the distribution switches from black holing traffic if either distribution switches lose the connection to the access switch.

Route filtering: Traffic flows pass from access through the distribution to the core and should never pass through the access layer unless they are destined to locally attached devices. Apply a distribute-list to all the distribution downlinks to filter the routes received by the access switches.

Configuration:
```
router eigrp 100
 distribute-list default out GigabitEthernet3/3
!
ip access-list standard default
 permit 0.0.0.0
```

Hello and Hold timer tuning: Tuning the EIGRP Hello and Hold timers provides for faster routing convergence in the rare event that L1 and L2 remote fault detection fails to operate. Hello and Hold timers are reduced to 1 and 3 seconds respectively.

Configuration:
```
interface TenGigabitEthernet4/3
 …
 ip hello-interval eigrp 100 1
 ip hold-time eigrp 100 3
```

Note Ensure that the timers are consistent on both ends of the link

CEF load balancing: To achieve best CEF load balancing, alternate L3 and L4 hashing on access, distribution and core routers.

Configuration: On access and core - `mls ip cef load-sharing simple`

On distribution - `mls ip cef load-sharing full`

3.5.2 Link Failure Detection

Carrier-delay timer: The default behavior for Catalyst switches is to use a default value of 0 msec on all Ethernet interfaces for the carrier-delay time to ensure fast link detection. It is still recommended as a best practice to hard code the carrier-delay value on critical interfaces with a value of 0 msec to ensure the desired behavior.

- Configuration: interface GigabitEthernet1/1

  carrier-delay msec 0

Note On Catalyst 6500, "LINEPROTO-UPDOWN" message appears when the interface state changes before the expiration of the carrier-delay timer configured via the "carrier delay" command on the interface. This is an expected behavior on Catalyst 6500 and is documented in CSCsh94221. For details, refer to Appendix B.

Link Debounce timer: By default, GigE and 10GigE interfaces operate with a 10 msec debounce timer that provides for optimal link failure detection. The default debounce timer for 10 / 100 fiber and all copper link media is longer than that for GigE fiber, and is one reason for the recommendation of a high-speed fiber deployment for switch-to-switch links in a routed campus design. It is good practice to review the status of this configuration on all switch-to-switch links to ensure the desired operation via the command "show interfaces TenGigabitEthernet4/1 debounce"

3.5.3 Multicast

PIM Spare mode: Configure PIM Sparse mode on all the interfaces.

Configuration: `ip pim sparse-mode`

RP: Configure routers in the core as Anycast RP.

PIM query interval: Configure PIM query-interval to 1 sec on interfaces to facilitate faster failover of Multicast streams.

Configuration: `ip pim query-interval 1`

Send-RP announce interval: For faster Anycast Auto-RP failover, configure send-rp announce interval to 5 sec.

Configuration: `ip pim send-rp-announce <interface> <RP announcement scope> interval 5`

Limit Multicast states: Configure mroute-limit and igmp-limit on all PIM routers.

Configuration: `ip multicast route-limit 11000`

`ip igmp limit 20`

Note Due to DDTS CSCsj48453, Catalyst 6500 does not forward multicast traffic to WISM module when catalyst 6500 is configured in L3 mode. This caveat is documented in section 3.1.5.3.4 WISM, page 3-22 as well as Appendix C of this document.

3.5.4 Wireless

Wireless controller: Connect wireless controller to distribution routers in Services Block.

Multicast over wireless: Do not use WiSM module for multicast over wireless. Instead, use Cisco 4404 wireless controller.

Note Due to DDTS CSCsj48453, Catalyst 6500 does not forward multicast traffic to WISM module when catalyst 6500 is configured in L3 mode. This caveat is documented in section 3.1.5.3.4 WISM, page 3-22 as well as Appendix C of this document.

3.5.5 Voice over IP

High Availability Campus Routed Access design provides a highly available, fault-tolerant infrastructure which is essential for easier migration to IP telephony, integration with applications such as video streaming and video conferencing in enterprise networks, and expansion of your IP telephony deployment across the WAN or to multiple CUCM clusters.

CHAPTER 4

Related Documents and Links

- *High Availability Campus Network Design-Routed Access Layer using EIGRP or OSPF*

  http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/ccmigration_09186a00805fccbf.pdf

- Cisco AVVID Network Infrastructure IP Multicast Design

  http://www.cisco.com/application/pdf/en/us/guest/tech/tk363/c1501/ccmigration_09186a008015e7cc.pdf

- *Bidirectional PIM Deployment Guide*

  http://www.cisco.com/warp/public/732/Tech/Multicast/docs/bidirdeployment.pdf

- Wireless features - *Enterprise Mobility 3.0 Design Guide*

  http://www.cisco.com/application/pdf/en/us/guest/netsol/ns279/c649/ccmigration_09186a00808118de.pdf

- *VOIP Infrastructure Solution Guide*

  http://www/en/US/tech/tk652/tk701/technologies_configuration_guide_book09186a00800eaa0e.html

- Cisco Unified Communications SRND

  http://www/en/US/products/sw/voicesw/ps556/products_implementation_design_guide_book09186a0080783d84.html

- *Cisco IOS SIP Configuration Guide*

  http://www/en/US/products/sw/iosswrel/ps5207/products_configuration_guide_book09186a00807517b8.html

This page is intentionally left blank

APPENDIX A

Test Cases Description and Test Results

A.1 Routing - IPv4

Table A-1 IPv4 Routing Test Cases

Test	Manual Test Case	Defects	Automation Test Case	Defects
System Integration Test Suite: This test suite is run in Campus Routed Access test network environment and all test cases within this test suite will run in parallel. Device configurations used for this test case will have feature combination and feature interaction with configurations from other test suites like Voice and Multicast test suites. This test suite will run with traffic streams flowing in the background that includes stateful traffic, stateless traffic, Multicast traffic, and voice traffic.				
EIGRP Neighbors Authentication This test case will setup and verify the EIGRP neighbor relationship is established between each router across the whole campus, using EIGRP MD5 authentication method. Specifically, in Routed Access Campus Design, configure the routers in access layer as layer 3 devices using EIGRP protocol. Verify that each access router has established the EIGRP neighbor relationship with routers in distribution layer.	Passed	—	Passed	—
EIGRP Stub This test case will set up and verify EIGRP stub routing on the access routers. In Routed Access Campus design, configuring the EIGRP stub feature on the layer 3 access routers prevents the distribution routers from sending downstream queries, which helps the convergence time.	Passed	—	Failed	CSCek78468
EIGRP Timers Tuning This test case will set up and verify the EIGRP hello, and dead timers are reduced to 1 and 3 seconds, in routed access design, reducing the EIGRP hello and holder timer from defaults of 5 and 15 second provides for a faster routing convergence in the rare event that L1/2 remote fault detection fails to operate.	Passed	—	Passed	—

EIGRP Summarization This test case will set up and verify the EIGRP route summarization on the interfaces of distribution routers to the core routers. Specifically in Routed Access Campus Design, configure the route summarization on distribution routers helps to reduce the convergence time by bounding those queries to a single hop in all direction.	Passed	—	Passed	—
EIGRP Route Filters This test case will set up and verify the EIGRP route filtering is applied on the distribution routers to ensure that the access router is only using the default routes (gateway) to the remote network as well as reduced routing table.	Passed		Passed	—

A.2 Convergence tests with Extended Baseline Configuration

Table A-2 Convergence Tests with Extended Baseline Configuration

Test	Manual Test Case	Defects	Automation Test Case	Defects
Negative Test Suite: Convergence Extended baseline Test Suite The Campus Routed Access Design provides the fastest convergence, the fastest restoration of voice and data traffic flows. With this design, the convergence time should be less than 200 msec. This test suite is run in Campus Routed Access test network environment setup with extended baseline configurations. All the test cases will be run in serial within this test suite. Device configurations used for this test suite will have feature combination and feature interaction with configurations from other test suites, Routing and Multicast test suites. This test suite will run with traffic streams flowing in the background that includes stateful traffic, stateless traffic, Multicast traffic and voice traffic.				
Convergence Link Failure for Voice Traffic The purpose of this test is to measure convergence times during link failures. Since campus design is totally redundant, it is expected to have convergence times that are in sub-second range. This test uses voice traffic only. Various instances are measured including link failures between core-distribution, core-core, and distribution-access devices. In this route convergence test, network convergence is measured from the data source to the receiver (End-to-end network convergence). SmartBits is used to send a data stream at a specified data rate. A corresponding convergence time is associated with each packet dropped during network convergence. For example, a data packet rate of 1000 pps corresponds to 1 millisecond (ms) convergence time for each packet dropped.	Passed	—	Passed	—

Convergence Redundant Router Failure for Voice traffic The purpose of this test is to measure convergence times during redundant router failures or reloads. Since campus design is totally redundant, it is expected to have convergence times in sub-seconds range. This test uses voice traffic only. Various instances are measured including core and distribution router reloads. In this route convergence test, network convergence is measured from the data source to the receiver (End-to-end network convergence). SmartBits is used to send a data stream at a specified data rate. A corresponding convergence time is associated with each packet dropped during network convergence. For example, a data packet rate of 1000 pps corresponds to 1 millisecond (ms) convergence time for each packet dropped.	Passed	—	Passed	—
Convergence Link Failure for Multicast Traffic The purpose of this test is to measure convergence times during link failures. Since campus design is totally redundant, it is expected to have convergence times that are in sub-second range. This test uses Multicast traffic only. Various instances are measured including link failures between core-distribution, core-core, and distribution-access devices. In this route convergence test, network convergence is measured from the data source to the receiver (End-to-end network convergence). SmartBits is used to send a data stream at a specified data rate. A corresponding convergence time is associated with each packet dropped during network convergence. For example, a data packet rate of 1000 pps corresponds to 1 millisecond (ms) convergence time for each packet dropped.	Passed	—	Passed	—

Convergence Redundant Router Failure for Multicast traffic The purpose of this test is to measure convergence times during redundant router failures or reloads. Since campus design is totally redundant, it is expected to have convergence times in sub-seconds range. This test uses Multicast traffic only. Various instances are measured including core and distribution router reloads. In this route convergence test, network convergence is measured from the data source to the receiver (End-to-end network convergence). SmartBits is used to send a data stream at a specified data rate. A corresponding convergence time is associated with each packet dropped during network convergence. For example, a data packet rate of 1000 pps corresponds to 1 millisecond (ms) convergence time for each packet dropped.	Passed	—	Passed	—
Convergence Link Failure for Unicast Traffic The purpose of this test is to measure convergence times during link failures. Since campus design is totally redundant, it is expected to have convergence times that are in sub-second range. This test uses Unicast traffic only. Various instances are measured including link failures between core-distribution, core-core, and distribution-access devices. In this route convergence test, network convergence is measured from the data source to the receiver (End-to-end network convergence). SmartBits is used to send a data stream at a specified data rate. A corresponding convergence time is associated with each packet dropped during network convergence. For example, a data packet rate of 1000 pps corresponds to 1 millisecond (ms) convergence time for each packet dropped.	Passed		Passed	—

Convergence Redundant Router Failure for Unicast traffic The purpose of this test is to measure convergence times during redundant router failures or reloads. Since campus design is totally redundant, it is expected to have convergence times in sub-seconds range. This test uses Unicast traffic only. Various instances are measured including core and distribution router reloads. In this route convergence test, network convergence is measured from the data source to the receiver (End-to-end network convergence). SmartBits is used to send a data stream at a specified data rate. A corresponding convergence time is associated with each packet dropped during network convergence. For example, a data packet rate of 1000 pps corresponds to 1 millisecond (ms) convergence time for each packet dropped.	Passed	—	Passed	—

A.3 Negative tests

Table A-3 Negative Tests

Test	Manual Test Case	Defects	Automation Test Case	Defects
Negative Test Suite: A negative test case will introduce certain conditions or failures to the network that will make other positive test cases fail in a system test environment; therefore, negative test cases are grouped together in a separate test suite for better test management.				
Interface Shut/Noshut The purpose of this test is to verify the effects of the interface of the core router/switch when it goes down and comes back up and what level of disruption of network connectivity does occur on the device and the network. The expected result is that after the interface is brought back up, the routing table will resume. In addition to that, the system experiencing an interface that goes down and up scenario should still have normal CPU utilization and there should not be any memory leak, traceback, or CPU hog etc. It is required that the core layer of the network provides the necessary scalability, load sharing, fast convergence, and high speed capacity.	Passed	—	Passed	—
Clear IP Routes The purpose of this test is to verify that router functionality is not interrupted and the ip routing table resumes after the routing table is reset by the clear ip route command.	Failed	CSCsk10711	Passed	—

EIGRP Flapping The purpose of this test is to verify that when there are small portions of EIGRP routes that are unstable and flapping, the core switch should be able to handle this behavior without negative impact on the network. The system should be up and continue to handle routing and the redirection of traffic properly. CPU utilization should remain in the normal range and there should be no memory leak, tracebacks, or CPU hogs.	Passed	—	Passed	—
PIM Neighbor Filters The purpose of this test is to verify PIM Neighbor Filters operation. The "pim neighbor-filter" command is to limit and prevent the other neighbors to form PIM neighbors. This can help to achieve stub Multicast routing.	Passed	—	Passed	—
MSDP SA filters The purpose of this test is to verify MSDP SA filters.With a default configuration, MSDP exchanges SA messages without filtering them for specific source or group addresses. To improve the scalability of MSDP in the native IP Multicast Internet, and to avoid global visibility of domain local (S,G) information, MSDP SA filters are used to reduce unnecessary creation, forwarding, and caching of some of these well-known domain local sources.	Passed		Passed	—
Accept Register filters Use this command to prevent unauthorized sources from registering with the RP. If an unauthorized source sends a register message to the RP, the RP will immediately send back a register-stop message. Accept register filter feature can be used on the RP to limit the sources that register for a particular Multicast group. This also only gives limited security since the source traffic may flow down the *,G or S,G tree of an active flow without registering at the RP.	Passed	—	Passed	—

A.4 Multicast tests

Table A-4 Multicast Test Cases

Test	Manual Test Case	Defects	Automation Test Case	Defects
System Integration Test Suite: Multicast Tests The purpose of this test is to provision and test Multicast in Routed Access setup. This test suite was run in Campus Routed Access test network environment setup and all testcases will be run in parallel with other test cases within this test suite. Device configurations used for this test case will have feature combination and feature interaction with configurations from other test suites, Routing, Wireless test suites. Test cases will run with traffic streams flowing in the background that includes stateful traffic, stateless traffic, Multicast traffic and voice traffic.				
PIM Sparse-Mode PIM Sparse-mode is able to support full spectrum of Multicast applications either one-to-many or many-to-many. For running Multicast traffic with PIM Sparse mode, static RPs are defined on core routers. Separate RPs are defined for the publish groups (Source in Services Block) and the subscribe groups (Feedback traffic sourced in user distribution). The Shortest Path Tree (SPT) threshold value was set to infinity to ensure that Multicast traffic used only the shared tree.	Passed	—	Passed	—

PIM bidir-mode The purpose of this testcase is to provision PIM bidir in Routed Access test environment and verify that it inter-operation with PIM Sparse-mode. PIM bidir helps deploy emerging communication and financial applications that rely on many-to-many model. This test will verify basic functionality of bidirectional PIM groups, mode flags, and the designated forwarder (DF) mode. SupervisorÃ' EngineÃ' 720 (Sup720) supports hardware forwarding of IPv4 bidirectional PIM groups. To support IPv4 bidirectional PIM groups, the Sup720 implements a new mode called designated forwarder (DF) mode. The DF is the router elected to forward packets to and from a segment for an IPv4 bidirectional PIM group. In DF mode, the supervisor engine accepts packets from the RPF and from the DF interfaces. When the supervisor engine is forwarding IPv4 bidirectional PIM groups, the RPF interface is always included in the outgoing interface list of (*, G) entry, and the DF interfaces are included depending on IGMP/PIM joins. If the route to the RP becomes unavailable, the group is changed to dense mode. Should the RPF link to the RP become unavailable, the IPv4 bidirectional PIM flow is removed from the hardware forwarding information base (FIB).	Passed	—	Passed	—
IGMP Snooping The purpose of this test is to verify the functionality of the IGMP snooping feature. This test configures a switch to use IGMP snooping in subnets that receive IGMP queries from either IGMP or IGMP snooping. IGMP snooping constrains IPv4 Multicast traffic at Layer 2 by configuring Layer 2 LAN ports dynamically to forward IPv4 Multicast traffic only to those ports that want to receive it.	Passed	—	Passed	—
Multicast Limits The purpose of this test is to verify the Multicast route limit, the PIM register rate limit, the MSDP SA limit, and the IGMP limit in test environment deployed with Multicast traffic and states. These will help safe guard network against anamolies in Multicast states. There are many commands that can limit the amount of state that can be created by Multicast traffic. This test uses the Multicast route-limit, pim register-rate-limit, msdp sa-limit, and igmp limits commands.	Passed	—	Passed	—

Auto RP The purpose of this test is to verify the functionality of the auto-RP and auto-RP listener features when static RP's also defined. Auto RP is defined for different set of Multicast groups than static RP groups.	Passed		Passed	—
MSDP/ Anycast RP The purpose of this test is to verify the basic functionality of MSDP and MSDP/Anycast. For running Multicast traffic, static rendezvous points (RPs) are defined. Redundant RPs are configured with MSDP to facilitate Anycast-RP. Separate RPs are defined for the publish groups and the feedback groups building floors.	Passed	—	Passed	—

PIM Stub The PIM stub feature supports Multicast routing between distribution layer and access layer. The PIM stub router contains two types of PIM interfaces: Uplink PIM interfaces and PIM passive interfaces. The uplink PIM interfaces have full PIM functionality and are used to connect to distribution routers. The PIM passive interfaces are connected to layer 2 access domains (for example, VLANs). The PIM stub feature provides the following restricted Multicast routing support: (1) The PIM stub router does not route the transit traffic between distribution routers. This behavior is enforced by Unicast (EIGRP) stub routing. The proper Unicast stub routing configuration is required to assist this PIM stub router behavior. The PIM stub feature does not prevent router administrator configuring RIP, static routes or PIM RP to bypass this restriction. (2) Only direct-connected Multicast (IGMP) receivers and sources are allowed in the layer 2 access domains. The PIM protocol is not supported in access domains. The PIM passive interface do not send or process any received Multicast control packets include PIM, DVMRP messages. Those Multicast control packets come in from PIM passive interfaces are ignored and dropped. The non-RPF traffic from PIM passive interface is dropped. (3) The redundant PIM stub router topology is not supported. The redundant topology here means that more than one PIM router forward Multicast traffic to a single access domain. Because of blocking PIM messages, the PIM assert and DR election mechanisms are not supported on PIM passive interface. Only the non-redundant access router topology is supported by PIM stub feature. The PIM passive interface assumes that it is the only interface and DR (Designated Router) on that access domain.	Passed	—	Passed	—

A.5 VoIP Tests

Table A-5 VoIP Test Cases

Test	Manual Test Case	Defects	Automation Test Case	Defects
System Integration Test Suite: VoIP This testcase is run in Campus Routed Access test network environment setup and will run in serial with "SCCP to SCCP", "SCCP to SIP", "Quality of Voice", "SRTP", and "Video Telephony" test cases within this test suite. Device configurations used for this testcase will have feature combination and feature interaction with configurations from other test suites, routing, and Multicast test suites. This test case will run with traffic streams flowing in the background that includes stateful traffic, stateless traffic, and Multicast traffic.				
SIP to SIP The purpose of this test is to provision SIP-to-SIP calls and verify calls be will be successful. In this test, ABACUS VoIP traffic/quality generation/verification test tool will be used to emulate 1,000 SIP client end-points (IP Phones) and to generate 500 VoIP call signaling and RTP (Real-Time Transport Protocol) traffic streams over the campus for 10 times. Total CSR (Call Success Rate) should be greater than 99.9%. In order to achieve the goal, all technology aspects of layer2, IP routing, QoS, Security in campus should be working and functioning properly. The Session Initiation Protocol (SIP) is an application-layer control (signaling) protocol for creating, modifying, and terminating sessions with one or more participants. These sessions include Internet telephone calls, multimedia distribution, and multimedia conferences. SIP has the following features: - Lightweight, in that SIP has only six methods, reducing complexity. - Transport-independent, (SIP can be used with UDP, TCP, ATM, etc.). - Text-based, allowing for humans to read SIP messages. SIP clients use TCP or UDP typically using port 5060 to connect to SIP servers and other SIP endpoints. SIP is primarily used in setting up and tearing down voice or video calls. However, it can be used in any application where session initiation is a requirement. These include Event Subscription and Notification, Terminal mobility and so on.	Passed	—	Passed	—

SCCP to SCCP The purpose of this test is to provision SCCP-toSCCP calls and verify calls be will be successful. In this test, ABACUS VoIP traffic/quality generation/verification test tool will be used to emulate 1,000 SCCP client end-points (IP Phones) and to generate 500 VoIP call signaling and RTP (Real-Time Transport Protocol) traffic streams over the campus for 10 times. Total CSR (Call Success Rate) should be greater than 99.9%. In order to achieve the goal, all technology aspects of layer2, IP routing, QoS, Security in campus should be working and functioning properly. SCCP (Skinny Client Control Protocol) is a proprietary terminal control protocol owned and defined by Cisco Systems, Inc. as a messaging set between a skinny client and the Cisco CallManager. Examples of skinny clients include the Cisco 7900 series of IP phone and the 802.11b wireless Cisco 7920. Skinny is a lightweight protocol that allows for efficient communication with Cisco Call Manager. Call Manager acts as a signalling proxy for call events initiated over other common protocols such as H.323, SIP, ISDN and/or MGCP. A skinny client uses TCP/IP to and from one or more Call Managers in a cluster. RTP/UDP/IP is used to and from a similar skinny client or H.323 terminal for the bearer traffic (real-time audio stream). SCCP is a stimulus-based protocol and is designed as a communications protocol for hardware endpoints and other embedded systems, with significant CPU and memory constraints.	Passed	—	Passed	—

SIP to SCCP and SCCP to SIP The purpose of this test is to verify SCCP-to-SIP and SIP-to-SCCP calls. In this test, ABACUS VoIP traffic/quality generation/verification test tool will be used to emulate 200 SCCP and 200 SIP client end-points and to generate 100 SIP-to-SCCP and 100 SCCP-to-SIP VoIP call signaling and RTP (Real-Time Transport Protocol) traffic streams over the campus for 10 times. Total CSR (Call Success Rate) should be greater than 99.9%. In order to achieve the goal, all technology aspects of layer2, IP routing, QoS, Security in campus should be working and functioning properly. The Session Initiation Protocol (SIP) is an application-layer control (signaling) protocol for creating, modifying, and terminating sessions with one or more participants. These sessions include Internet telephone calls, multimedia distribution, and multimedia conferences. SIP has the following features: - Lightweight, in that SIP has only six methods, reducing complexity. - Transport-independent, (SIP can be used with UDP, TCP, ATM, etc.). - Text-based, allowing for humans to read SIP messages. SIP clients use TCP or UDP typically using port 5060 to connect to SIP servers and other SIP endpoints. SIP is primarily used in setting up and tearing down voice or video calls. However, it can be used in any application where session initiation is a requirement. These include Event Subscription and Notification, Terminal mobility, and so on. SCCP (Skinny Client Control Protocol) is a proprietary terminal control protocol owned and defined by Cisco Systems, Inc. as a messaging set between a skinny client and the Cisco CallManager. Examples of skinny clients include the Cisco 7900 series of IP phone and the 802.11b wireless Cisco 7920. Skinny is a lightweight protocol that allows for efficient communication with Cisco Call Manager. Call Manager acts as a signalling proxy for call events initiated over other common protocols such as H.323, SIP, ISDN and/or MGCP. A skinny client uses TCP/IP to and from one or more Call Managers in a cluster. RTP/UDP/IP is used to and from a similar skinny client or H.323 terminal for the bearer traffic (real-time audio stream). SCCP is a stimulus-based protocol and is designed as a communications protocol for hardware endpoints and other embedded systems, with significant CPU and memory constraints.	Passed	—	Passed	—

Media Transport Delay The purpose of this test is to measure delay. In this test, ABACUS VoIP traffic/quality generation/verification test tool will be used to emulate 200 SIP client end-points and to generate 100 VoIP call signaling and RTP (Real-Time Transport Protocol) traffic streams over the campus for 10 times. Average one-way delay should be less than 100 ms. In order to achieve the goal, all technology aspects of layer2, IP routing, QoS, Security in campus should be working and functioning properly. This testcase will run with traffic streams flowing in the background that includes stateful traffic, stateless traffic, and Multicast traffic. **One-Way Delay (OWD)** One-way delay measures time interval between the time a voice pattern leaves the transmitting device and the time reaches the receiving device. The accuracy of this measurement is ±2 ms, and the resolution is 1. In Simplex mode, the number of measurements for one-way delay on a channel should equal the number of PSQM values.	Passed	—	Passed	—
Jitter The purpose of this test is to measure RTP jitter which is measured during the interval between the sending of two RTCP packets. In this test, ABACUS VoIP traffic/quality generation/verification test tool will be used to emulate 200 SIP client end-points and to generate 100 VoIP call signaling and RTP (Real-Time Transport Protocol) traffic streams over the campus for 10 times. Average RTP jitter should be less than 50 ms. In order to achieve the goal, all technology aspects of layer2, IP routing, QoS, Security in campus should be working and functioning properly.	Passed		Passed	—

Quality of Voice The purpose of this test is to measure Quality of Voice using PSQM model. PSQM uses a psychoacoustic model that aims to mimic the perception of sound in real life. Although it was originally developed to test Codecs, it was widely used for testing VoIP systems. The algorithm tested Codec by comparing the signal after it has been through the coder and decoder process with the original signal. A network can be similarly tested by replacing the coder-decoder elements with an SUT. PSQM provides an output in the range 0 to 6.5, where 0 indicates a good channel, and 6.5 indicates a very poor channel. PSQM Value: In this test, ABACUS VoIP traffic/quality generation/verification test tool will be used to emulate 200 SIP client end-points and to generate 100 VoIP call signaling and RTP (Real-Time Transport Protocol) traffic streams over the campus for 10 times.	Passed	—	Passed	—
Video Telephony The purpose of this test is to provision H.264 and verify call success. H.264/AVC/MPEG-4 Part 10 contains a number of new features that allow it to compress video much more effectively than older standards and to provide more flexibility for application to a wide variety of network environments. H.264 is supported with Cisco 7900 Video IP Phones. In this test, ABACUS VoIP traffic/quality generation/verification test tool will be used to emulate 200 SIP client end-points and to generate 100 VoIP call signaling and H.264 RTP (Real-Time Transport Protocol) traffic streams over the campus 10 times. Total CSR (Call Success Rate) should be greater than 99.9%. In order to achieve the goal, all technology aspects of layer2, IP routing, QoS, Security in campus should be working and functioning properly.	Passed	—	Passed	—

A.6 Wireless Tests

Table A-6 Wireless Test Cases

Test	Manual Test Case	Defects	Automation Test Case	Defects
System Integration Test Suite: Wireless This test suite is targeted to deploy and verify a campus wireless solution with a Routed Access Campus Topology using the WiSM, 4404 Wireless controller and Light Weight Access Points. This Test Suite is to verify that all traffic (data, voice, video, Multicast) can successfully coexist between the wireless clients and a wireless campus network.				
Wireless Controller System This test case will configure the WiSM controller and verify that it functions as expected. The Cisco WiSM is a Cisco Wireless LAN 4404 Controller Module in a Catalyst 6k. It works in conjunction with the Cisco Lightweight Access Points (LWAPP) protocol to support wireless data, voice, and video applications. This test case will verify that the WiSM can manage the LWAPP Access Points. This test will also run on the Cisco 4404 wireless controllers in the manual testbed only.	Passed	—	Passed	—
Dot1x Authentication This test case will verify that a wireless PC Client employee account can associate successfully with a WLAN using Dot1x Authentication and can transmit/receive wireless data traffic across campus network. The Cisco Secure Services Client will be used as the dot1x supplicant.	Passed	—	Passed	—
Voice over Wireless This test case will verify VoIP on wireless network. The Cisco Unified Wireless IP Phone 7920 is an easy to use IEEE 802.11b wireless IP phone that provides comprehensive voice communications in conjunction with Cisco Unified CallManager and the Cisco wireless infrastructure. This test case will verify that a Cisco Unified Wireless IP Phone 7920 transparently delivers voice traffic across campus network over the 4404/WISM Controllers.	Passed	—	Passed	—
Intra-controller Roaming Intra-controller roaming enables a client to change its connection between access points in the same subnet (Intra-controller roaming) to support time-sensitive applications such as VoIP, video streaming, and client/server-based applications. This test case will verify the ability of a wireless client Intra-controller roaming between the APs.	Passed	—	Passed	—

L2 Inter-controller Roaming L2 inter-controller roaming enables a client to change its connection between access points between subnets to support time-sensitive applications such as voice/video streaming and client/server-based applications. L2 inter-controller roaming includes two features -access-point-assisted channel scanning and fast IEEE 802.1X rekeying. This test case will verify the ability of a wireless client L2 inter-controller roaming between the APs.	Passed		Passed	—
Wireless control system verification Cisco Wireless Control System (WCS) is the industry-leading platform for wireless LAN planning, configuration, and management. It works in conjunction with Cisco Aironet Lightweight Access Points, Cisco wireless LAN controllers and the Cisco Wireless Location Appliance. With Cisco WCS, network administrators have a single solution for RF prediction, policy provisioning, network optimization, troubleshooting, user tracking, security monitoring, and wireless LAN systems management. Robust graphical interfaces make wireless LAN deployment and operations simple and cost-effective. Cisco WCS includes tools for wireless LAN planning and design, RF management, location tracking, Intrusion Prevention System (IPS), and wireless LAN systems configuration, monitoring, and management. This test case will verify that WCS can manage the wireless networks across the campus network	Passed	—	Passed	—
Multicast over wireless This test case will verify that IP/TV Multicast traffic can be successfully delivered to Wireless PC clients.	Passed	—	NA	—

This page is intentionally left blank

APPENDIX B

Defects

Six software defects were identified during test execution.

B.1 CSCek78468

When **`eigrp stub connected`** is configured under router EIGRP process, this configuration changes to **`eigrp stub connected summary`** after reload.

Severity:Moderate

Workaround: Manually configure “eigrp stub connected” after reload.

Status: Resolved. Fix not yet been integrated in any release.

B.2 CSCek75460

Loopback interface appears under show IP protocol even after it is deleted.

Severity: Minor

Workaround: Removing and configuring EIGRP protocol clears this problem.

Status: Resolved

Fix has been integrated in the following releases:

12.0(32.03)S06 12.2(32.08.11)XIB62.16 12.4(16.09)T

B.3 CSCsk10711

Installation of static routes into the routing table takes about 5 to 6 seconds after executing **`clear ip route *`**. This delay was noticed during negative testing.

Severity: Moderate

Workaround: None.

Status: Closed. This issue will be resolved via routing infrastructure changes in future IOS releases.

B.4 CSCsh94221

The “LINEPROTO-UPDOWN” message appears when the interface state changes before the expiration of the carrier-delay timer configured via the `carrier delay` command on the interface On Catalyst 6500.

Severity: Moderate

Workaround: None.

Status**:** Closed since this is an expected behavior on Catalyst 6500.

B.5 CSCsk01448

Tracebacks in the syslogs due to CPU hog by PIM process during router reload.

Severity: Moderate

Workaround: None.

Status: New

B.6 CSCsj48453

Catalyst 6500 does not forward multicast traffic to WISM module when catalyst 6500 is configured in L3 mode.

Severity: Catastrophic

Workaround: None.

Status: Assigned

APPENDIX C

Technical Notes

C.1 Technical Note 1:

Wireless controller software should be upgraded in sequence.

It is a good practice to back up the configuration file before upgrading or downgrading the software to avoid losing all or part of the configuration stored in NVRAM. While upgrading the software on wireless controller, follow the upgrade path outlined below.

Table C-1 Wireless Controller Upgrade Path

Desired Upgrade	Upgrade Path
3.2 or earlier to 4.1	Use the following steps to upgrade to 4.1: **Step 1Upgrade to software release 4.0.** **Step 2Upgrade to release 4.1.**
4.0 to 4.1	No restriction. Upgrade directly to 4.1

Note In some cases, when upgrading to software release 4.1 directly from release 3.2 or earlier, upgrade may not be successful or some features may not function as expected.

C.1 Technical Note 1: