StadiumVision

For Authorized Cisco StadiumVision Partners Only

Cisco Connected Stadium Design Guide

Version 2.9

June 2016

Cisco Systems, Inc.
www.cisco.com

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Contents

About this Guide

This document provides a detailed description of the Cisco Connected Stadium Solution. This solution provides the wired infrastructure is specifically design to support the various applications used in Sports and Entertainment venues. As such, it describes the design decisions including relevant samples of configuration and accompanying descriptions of the features within the network elements.

Document Audience

This document is intended for Cisco engineers and product managers and Cisco partners. Additionally, technical sales and marketing people can use this document as a master reference guide when helping customers understand what components they need for implementing the Cisco Connected Stadium Solution.

Related Documentation

The Cisco StadiumVision design is outside the scope of this document, however, you can find detailed information about the Cisco StadiumVision Solution at the following link: Cisco StadiumVision.

Document History

Table 1. Revision History

Date	Rev	Authors	Comments
6/20/16	2.9	Steve Schubert	Updated the Wireless Service Block section to align with the Connected Stadium Wi-Fi design guide. Added multicast information for in-suite video. Updated Option 60 syntax in StadiumVision section of the Cisco Solution Integration chapter. Updated the IOS DHCP Server recommendations. Added Appendix B – Cisco Connected Stadium Recommended Equipment Added Appendix C – Cisco Connected Stadium Product Options
9/15/15	2.8	Steve Schubert	Updated Figures 4, 8, 11, 12, 13, and 14. Added details concerning the Wireless Service Block. Updated Tables 7, 8, and 12. In the Cisco Solution Integration chapter, added Figure 39 and further details concerning StadiumVision multicast, added SV-4K PTP Figure 43 and corrected the DMP SV-4K DHCP options configuration steps.
3/18/15	2.7	Steve Schubert	Corrected MSDP originator-id configuration. Added the SV-4K and DHCP Option 43 Appendix.
12/17/14	2.6	Steve Schubert	Added new PoE section in Access Layer section of Chapter 1. Added new DMP LLDP Requirements in Chapter 6.
6/3/13	2.3	Steve Schubert	Added StadiumVision Multi-Venue Support in the Integrating StadiumVision on the Connected Stadium Network chapter.
4/1/13	2.2	Steve Schubert	Updated Wireless Service Block design.
3/25/13	2.4	Steve Schubert	Added new reference documents. Added new Core switch to Table 2.

Date	Rev	Authors	Comments
			Added new Access Layer switches to Table 3. Added new Data Center switches to Table 4. Added new Video Distribution switches to Table 5. Added new Internet Edge switches to Table 6. Added new WSB switches in Table 7. Removed Table 8 and added reference to Connected Stadium Wi-Fi design guide for WLC recommendations. Add clarification on MC Address Table 12 on page 44. Added note in Security Overview Chapter 4 page 77. Added footnotes for Prime Infrastructure version recommendations. Added clarifying language to the StadiumVision Server Network Configuration section and updated Figure 41. StadiumVision Server Configuration Overview.
2/1/13	2.1	Steve Schubert	Updated Service Block hardware lists, Wireless Services Block design, NCS & LMS to Prime Infrastructure.
8/6/12	2.0	Steve Schubert	Complete revamp of the content with updated with new information and consolidated existing content.
9/14/11	1.3	Steve Schubert	Incorporated SV Director VLAN and spanning-tree access layer recommendations.
11/17/10	1.2	Debbie Morrison	Incorporated IPMc comments from Vijay.
3/26/10	1.1	Gary M Davis	Changed references to Cisco Connected Stadium.

Date	Rev	Authors	Comments
10/22/09	1.0	Debbie Morrison, Trish McBride	Reformat and Edit.
10/08/09	0.1	Kevin Turek, Gary M Davis, Steven Fly, Matt Swartz, Chris Kodadek, Jay Standley, Arturo Molina, Ric Daza, Matt Soderlund, Kyle Prevey, Cary Stotland, Anisha Lalani	First draft.

Related Documents

- *Borderless Campus Network Design Guide*
- *Medianet Campus QoS Design 4.0*
- *Cisco SAFE Reference Guide*
- *SAFE - Network Foundation Protection*
- *Cisco Design Zone* (Top-Level Design Home Page)
- *Cisco Catalyst 3750-E and 3560-E Switch Software Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide*
- *Cisco Catalyst 6500 Switch Software Configuration Guide*
- "Cisco StadiumVision Director Server Redundancy" in the *Cisco StadiumVision Director Server Architecture* module of the *Cisco StadiumVision Director Server Administration Guide*.
- *Cisco Prime Infrastructure User Guide, Release 2.0*
- *Cisco Connected Stadium Wi-Fi Design Guide* (available to qualified Cisco StadiumVision partners)

StadiumVision

Network Architecture

Architecture Overview

The Cisco Connected Stadium network is built around a scalable tiered hierarchical and modular design including collapsed core/distribution and access layers, as shown in the figure below. The design follows Cisco's Borderless Network architecture with some customization to fit the Sports and Entertainment vertical market.

The architecture uses a modular architecture where functions are separated into functional blocks and dual-homed into a redundant, collapsed core/distribution pair of switches. Applying this modular approach to the WLAN places the WLAN Controllers and associated components like access switches, firewalls, and associated servers into a wireless service block.

Figure 1. Cisco Connected Stadium Network Topology

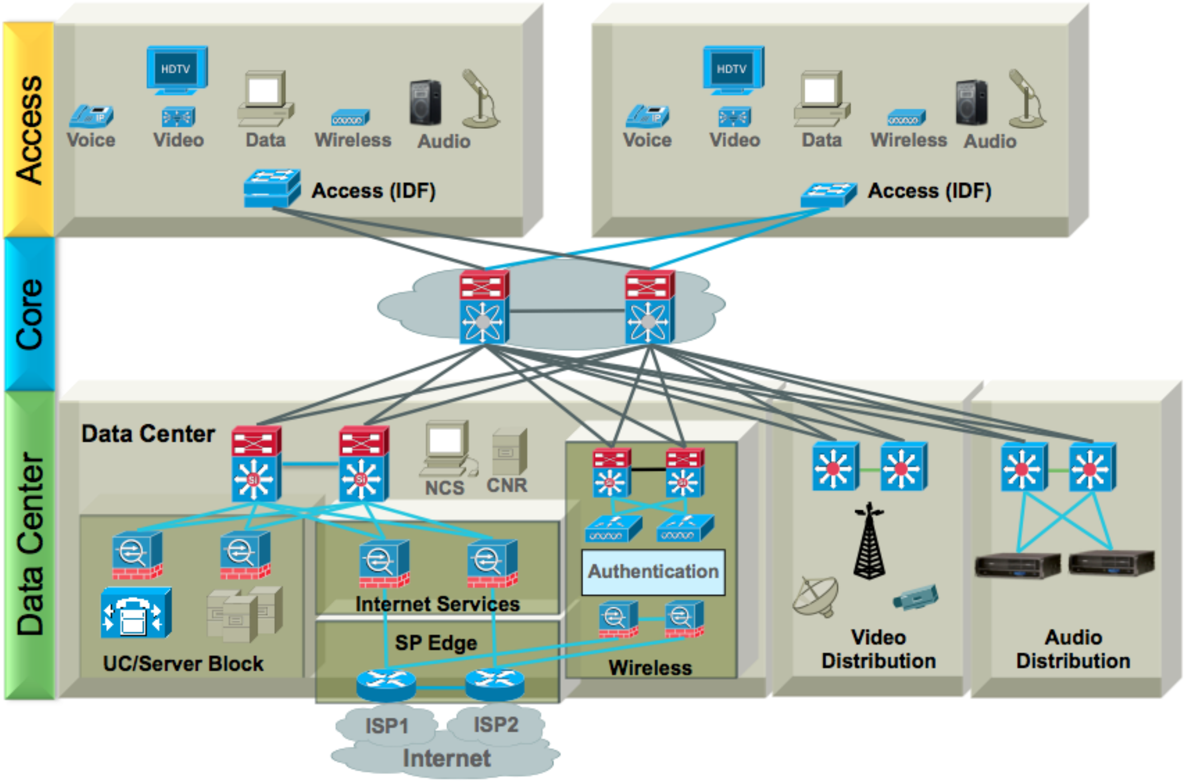

Attributes of the Connected Stadium network are as follows:

- Two-tiered hierarchy—Collapsed core/distribution layer and access layer.
- Modular topology—Building (aggregation) blocks provide flexibility during changes and upgrades. For example, the wireless services block is connected into the network via dual 1-Gbps or 10-Gbps uplinks.
- Use of VLANs and Layer 3 to the Access Layer to limit fault domains and provide clear demarcations and service isolation.
- Core designed with redundant hardware, software and links to provide maximum redundancy and optimal convergence.
- Load-balancing redundant uplinks from the core to the access layer.
- Uses EIGRP (recommended) or OSPF for unicast routing.

- Two-tiered hierarchy—Collapsed core/distribution layer and access layer.
- Modular topology—Building (aggregation) blocks provide flexibility during changes and upgrades. For example, the wireless services block is connected into the network via dual 1-Gbps or 10-Gbps uplinks.
- Use of VLANs and Layer 3 to the Access Layer to limit fault domains and provide clear demarcations and service isolation.
- Core designed with redundant hardware, software and links to provide maximum redundancy and optimal convergence.
- Load-balancing redundant uplinks from the core to the access layer.
- Uses EIGRP (recommended) or OSPF for unicast routing.

NOTE:OSPF is hierarchical and is typically used for three-tiered network designs. EIGRP is better suited for smaller two-tier networks like those used in Sports and Entertainment venues.

- Uses PIM Sparse Mode for multicast routing with redundant Rendezvous Points.

Core/Distribution Layer

The Core layer of the network provides the high-speed switching and aggregation of Access Layer switches including those used in the Service blocks. The Connected Stadium Core design consists of three fully-redundant options, Layer 3 to the Access (dual-chassis), L2 to the Access (dual-chassis) and a single chassis version of each.

Figure 2. Core Design Options

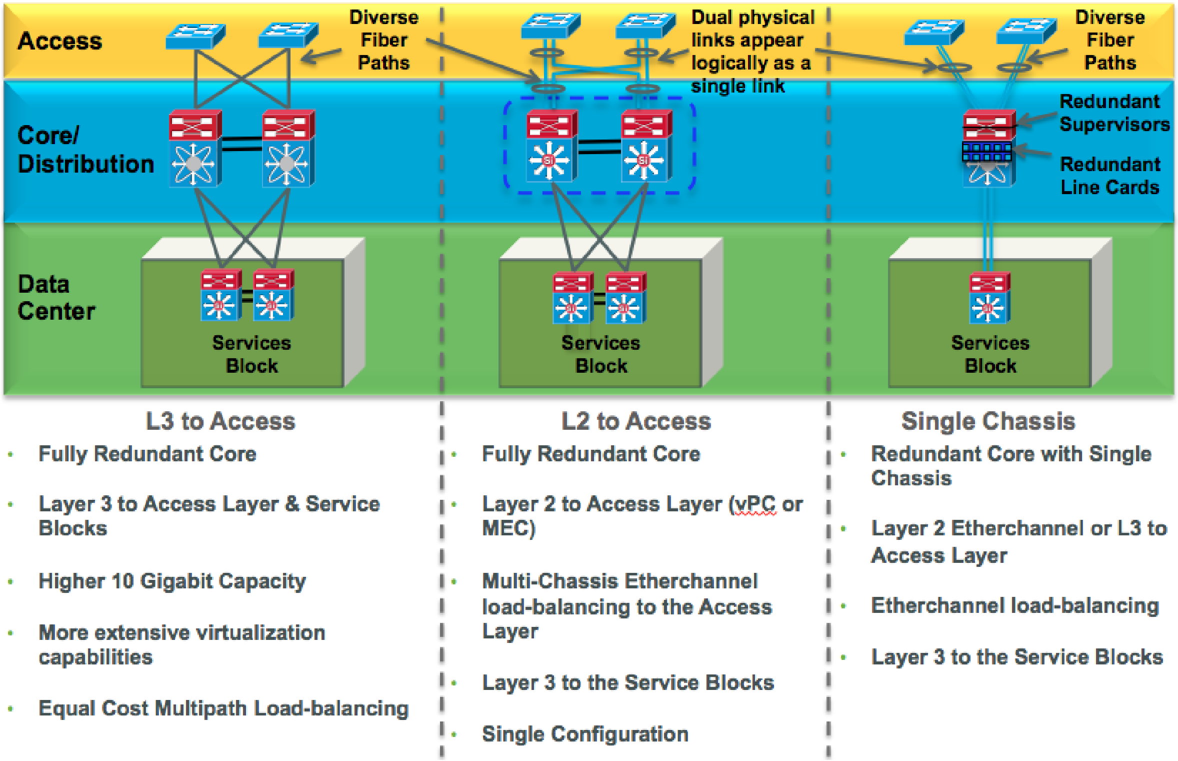

L3 to Access

- Fully Redundant Core
- Layer 3 to Access Layer & Service Blocks
- Higher 10 Gigabit Capacity
- More extensive virtualization capabilities
- Equal Cost Multipath Load-balancing

L2 to Access

- Fully Redundant Core
- Layer 2 to Access Layer (vPC or MEC)
- Multi-Chassis Etherchannel load-balancing to the Access Layer
- Layer 3 to the Service Blocks
- Single Configuration

Single Chassis

- Redundant Core with Single Chassis
- Layer 2 Etherchannel or L3 to Access Layer
- Etherchannel load-balancing
- Layer 3 to the Service Blocks

The Catalyst 6500 VSS is the most versatile and least complex dual-chassis core switching option. The attributes that set it apart are it's single configuration file and scalable non-spanning tree Layer 2 to the Access layer capability.

The Nexus 7000 dual-chassis core option provides higher 10 Gigabit density at a more economical price point than the Catalyst 6500 VSS option. The Nexus 7000 is typically used in larger venues. As compared to the Catalyst 6500 VSS, the Nexus 7000 dual-chassis core option provides similar functionality with its support of virtual Port Channel (vPC) technology for providing a non-spanning tree L2 to the Access layer topology. However, additional complexity is incurred due to the individual core switch configurations and the requirement for Hot Standby Routing Protocol (HSRP) for gateway redundancy for those devices located in the access layer.

NOTE: When using Nexus F2 linecards, the recommendation is to use Layer 3 to the Access Layer switches. The F2 linecard ARP table capacity limits the number of directly connected Layer 3 hosts. See for recommended limits.

If rack space is limited, a single chassis option using the Catalyst 6500, Nexus 7000 or Catalyst 4500 or dual-chassis 4500X VSS (typically used for small venues) is a good choice. To achieve the highest possible level of availability in the single chassis option, redundant components should be used. That is, the chassis should be populated with dual Supervisors, redundant power supplies, redundant fabric modules, and redundant line cards.

Access Layer

The Access Layer switches that are primarily used in sports and entertainment venues are low-profile (1 Rack Unit (RU)) stackable switches. A number of switches are interconnected to form a ‘stack’ – a single logical switch with the resilience and redundancy of a multi-chassis solution. The interconnection of the switches is done using a dedicated ring technology that is self-healing in the event of switch or cable failure allowing the stack to continue to provide network access. As described below, the stack typically consists of two to four interconnected switches but larger stacks may be required within specific IDFs.

Switch numbering and priority within the stack is recommended to follow the physical layout of the stack to ensure smooth troubleshooting and replacement.

Figure 3. Access Layer Switch Design

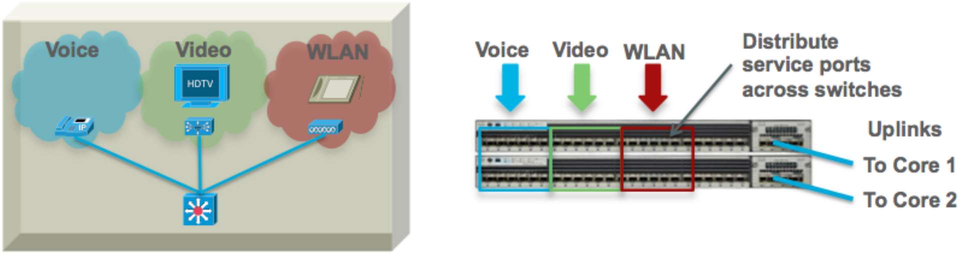

Stack Scaling

The stadium may require an adaptable yet manageable model for handling the different port density demands of each area of the venue. Initially a single four-switch stack model was developed to handle the predominant demands of the access IDFs but a lower density, and larger multi-stack option may be required. The size and number of uplinks will vary depending on the projected bandwidth requirements.

NOTE: The Access layer switches should be upgradeable to 10 Gigabit uplinks if 1 Gigabit is deemed sufficient for the intial deployment.

Below is a description of how to determine uplink bandwidth requirements.

Figure 4. Sizing the Uplinks to the Core

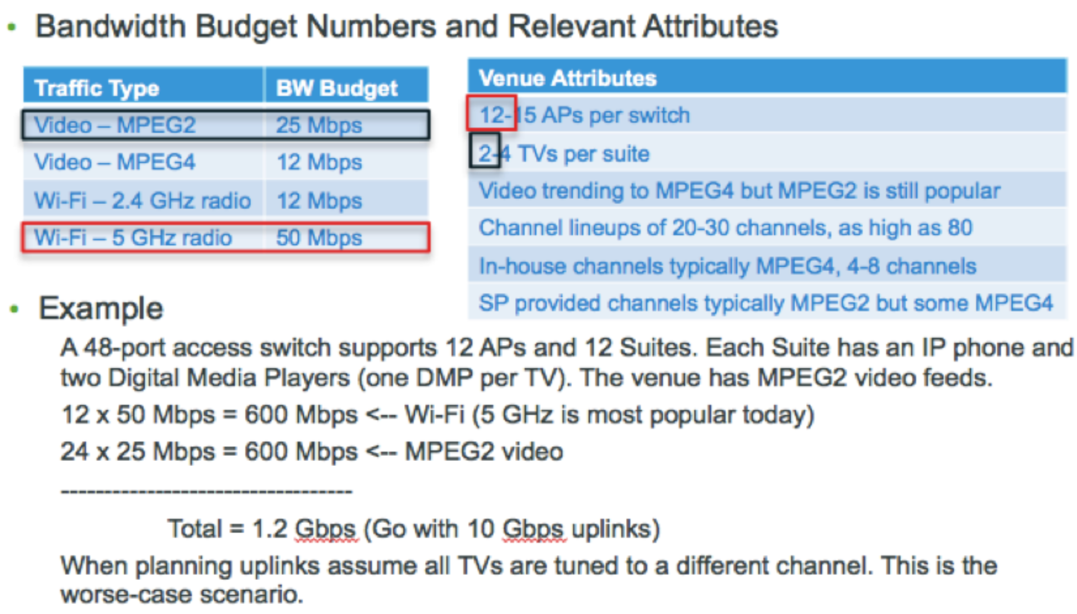

- Bandwidth Budget Numbers and Relevant Attributes

Traffic Type	BW Budget
Video – MPEG2	25 Mbps
Video – MPEG4	12 Mbps
Wi-Fi – 2.4 GHz radio	12 Mbps
Wi-Fi – 5 GHz radio	50 Mbps

Venue Attributes
12-15 APs per switch
2-4 TVs per suite
Video trending to MPEG4 but MPEG2 is still popular
Channel lineups of 20-30 channels, as high as 80
In-house channels typically MPEG4, 4-8 channels
SP provided channels typically MPEG2 but some MPEG4

- Example

A 48-port access switch supports 12 APs and 12 Suites. Each Suite has an IP phone and two Digital Media Players (one DMP per TV). The venue has MPEG2 video feeds.

12 x 50 Mbps = 600 Mbps <-- Wi-Fi (5 GHz is most popular today)

24 x 25 Mbps = 600 Mbps <-- MPEG2 video

Total = 1.2 Gbps (Go with 10 Gbps uplinks)

When planning uplinks assume all TVs are tuned to a different channel. This is the worse-case scenario.

Below are some suggested stack configurations.

Low Density Model (2 x 10 GE or 1 GE uplinks, 2-3 switches)

The Low Density model uses 2 switches to form a stack, using a single uplink from each switch for sufficient bandwidth and redundancy via diverse paths. An additional switch may be added to the stack for additional port density (this additional center switch does not have its own uplinks but relies on the stack ring to maintain connectivity to the stack and to the core network).

High Density Model (4 x 10 GE or 1 GE uplinks, 4-6 switches)

The High Density model offers a good balance of available bandwidth and redundancy, port density and cost. The stack may have additional switches added to the stack to increase port density without requiring additional uplinks but careful traffic/bandwidth monitoring should be done to decide whether a move to a multi stack solution is more appropriate. The stack may be expanded physically to nine switches if required.

Multi Stack Model (4 x 10GE or 1GE, 6+ switches)

For high density demands in the access layer, inserting a new stack of switches provides dedicated bandwidth and resilience in a simple way to understand and troubleshoot. An example is the smaller multi-stack configuration used within the HD Video Room to provide sufficient capacity and resilience for the High Definition video distribution throughout the stadium. (Although this smaller multi-stack configuration actually relies on 2 x 2 switch stacks, the model still applies). Each stack may be hybridly expanded using additional switches as described in the High Density model. In the case where more than six switches are deployed in an IDF, two separate stacks are deployed.

Port Assignment

To physically spread the load of traffic, sequential port allocation or ‘striping’ of access ports across the switches within the stack avoids bottlenecks within a single overloaded switch and is recommended for this stack environment. See Figure 3 for an example of how ports of a VLAN are distributed (or striped) across a two-switch stack.

Fiber Uplinks

The Access layer stack of switches is connected via either 2 x 10 GE or 1 GE or 4 x 10 GE or 1 GE fiber cables to the core switches. The fibers should be routed via two diverse paths to avoid catastrophic fiber failure in any one fiber run. The fibers are connected to the Access stack in alternate switches to provide redundancy in case of switch failure (into two separate switches for a Low Density IDF or into four separate switches for High Density or each multi-stack IDF). Small /30 subnets are used for these uplinks to provide routed EIGRP dual paths and manageability for each individual fiber link or in the case of a VSS core, Mutli-chassis Etherchannel is used to bundle the fiber uplinks into a single logical uplink. In either uplink configuration, traffic is load-balanced across all links. It’s important in bandwidth planning that traffic can be handled by remaining active links when there is a link failure.

Uni-Directional Link Detection

Uni-Directional Link Detection (UDLD) is used to detect and avoid RX/TX single fiber failures affecting the stability of the routing and switching environment. UDLD is configured on the 10 GE fiber uplinks to avoid such problems.

VLANs

The following services are examples of VLANs that are defined on access switches.

- IP Phone VLAN.
- HD Video VLAN for Digital Media Players (DMPs).
- Security Camera VLAN.
- Point of Sale (PoS) VLAN for connecting PoS terminals.
- Internal VLAN for attaching internal venue PCs.
- BMS VLAN for connecting Building Management System devices.
- Wireless VLAN for attaching wireless APs.

Spanning Tree and Protection

Spanning Tree is enabled in Rapid PVST mode on the Access stack to ensure loops from external devices are not introduced to the Layer 3 access network.

Portfast

The Spanning Tree feature portfast is configured on all access ports on the Access switch stacks to allow host ports to move quickly from Blocking to Forwarding.

BPDUGuard

BPDUGuard is also enabled on all access ports to ensure ports are automatically disabled if they receive BPDUs from miscabled connection to external switches which could cause Spanning Tree disruption within the access layer and potential switching loops. When BPDUs are seen on such access ports the port is err-disabled to avoid disruption and messages appear on the NMS systems to alert Operations staff to investigate the issue.

Storm-control

The storm-control feature is configured on all access uplinks to limit broadcast traffic to less than 2 Kbps. With this feature enabled, CPU utilization will remain around 10% if a broadcast storm is introduced.

Quality of Service

End-to-end quality of service is required to support the converged applications within the stadium. As a result, the Access stacks distributed throughout the stadium perform classification, marking and queuing functions, in order to enforce the required behavior as packets traverse the network infrastructure. Refer to the chapter on *Quality of Service (QoS)* for the recommended QoS settings and policies.

DHCP Helper

Some VLANs require DHCP to provide IP addresses for hosts and devices. IP helper services enable on the SVI for a particular VLAN forwards DHCP requests to a central DHCP server within the network.

NMS

The *Network Management* chapter lists and describes the commands recommended for deployment within the Cisco Connected Stadium network. For additional information concerning network management security please refer to the *Security* chapter.

Wireless Access

802.11 wireless Access Points may be used for extending the Access layer. Wireless APs are attached to the Access stack via the Wireless VLAN. Wireless clients and devices attach via designated SSIDs and the traffic is delivered via CAPWAP tunnel from the AP to the WLAN Controllers located in the Wireless Services block. From there, client traffic exits the WLAN controller on a dedicated VLAN and is routed to their appropriate destination. See the Wireless Services section for more details how that part of the network operates and is designed.

Power over Ethernet

Access Layer switches should support the higher power IEEE 802.1at Power over Ethernet, also known as PoE+, which supports up to 30W per port. The newer 3700 Series Access Points and Digital Media Players will require higher power to take advantage of new capabilities. The Access Layer switches should also always be equipped with the highest wattage power supplies and careful consideration must be made when choosing a switch model to ensure the switch can support the required number of PoE+ ports.

Service Blocks

The Connected Stadium network design uses a modular architecture where servers and/or specialized hardware devices that provide a service to the network are connected via a dedicated redundant pair of switches. By using this modular design approach, specific services like the video headend may be upgraded or troubleshot independent of the rest of the network and therefore, minimizing any negative impact these actions may have on other services or locations in the Connected Stadium network.

Data Center

The Data Center Service Block uses a dedicated pair of switches to provide fully redundant connections to server farms and the network.

Service blocks for the Internet Edge, Wireless Services, Server Farms, Unified Communications may be consolidated and aggregated up to the Data Center service switches or possibly broke out into their own Service Blocks with dedicated network hardware. The following sections describe this dedicated Service block design approach.

Figure 5. Data Center Service Block

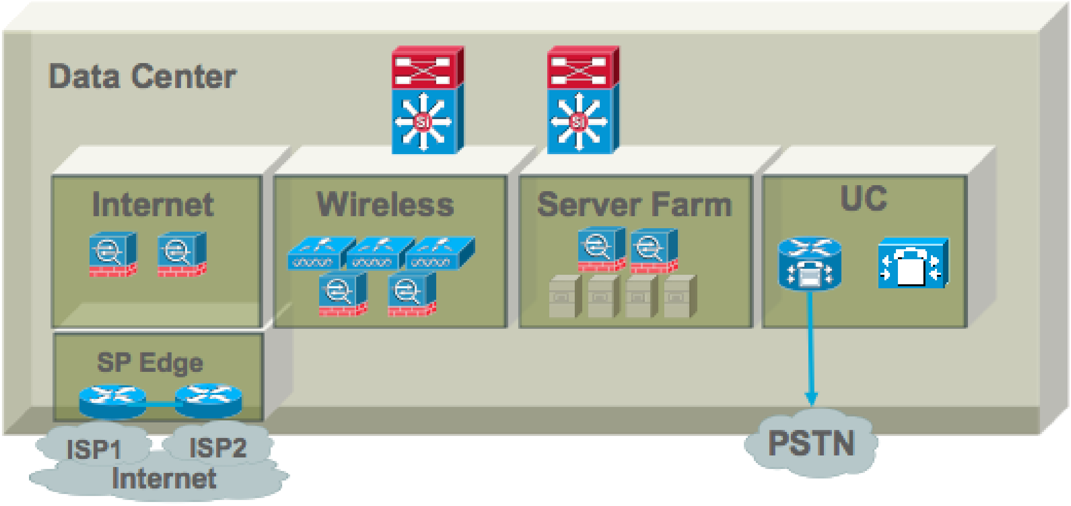

Video Distribution

The Video Distribution block (also commonly referred to as the Video Headend) provides the aggregation and distribution point for the IP video sources for the venue. It is usually located in or near the video broadcast room where the in-house feeds (for example, field camera feeds), Off-air (for example, local TV channels), Cable or Satellite TV provider feeds are groomed and distributed. The figure below shows a fully-redundant Video Distribution block with two Video Distribution Service (VDS) switches. Each switch has a set of identical video feeds. Although, each set of feeds are transmitting within the Video Distribution block, only a single switch is forwarding multicast video into the network at any one time. Reasons why and details on how this achieved will be provided in the Multicast Routing section.

Figure 6. Video Distribution Service Block

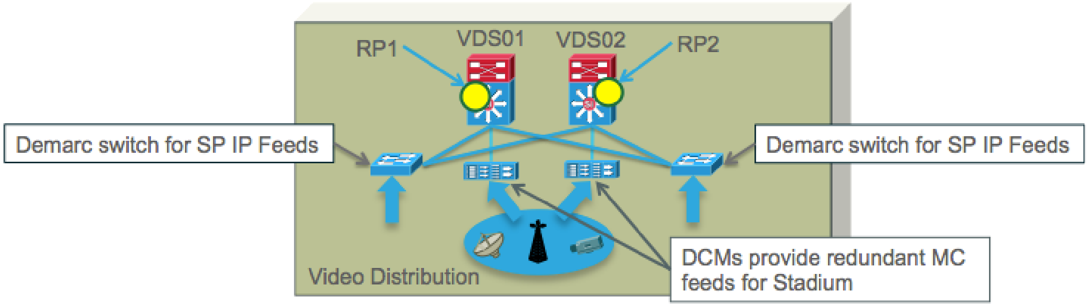
RP1
VDS01
VDS02
RP2
Demarc switch for SP IP Feeds
Demarc switch for SP IP Feeds
Video Distribution
DCMs provide redundant MC feeds for Stadium

Attributes of the Video Distribution Service block are as follows:

- Serves as the multicast point of origin for video traffic within the venue.
- Uses PriorityCast for redundancy with Rendezvous Points (RPs) on Video Distribution switches.
- Uses a separate switch as demarcation point for incoming Service Provider IP Video feeds. Demarc switches may use IP Base and configured as EIGRP stub.
- Acts as the trust boundary for QoS and Network Security policies to be applied.
- Video Distribution Switches are full IGP routing peers to the Core (not stub) and therefore, use IP Services or higher featured IOS images.

Internet Edge

The Internet Edge provides the secure access to the Internet and also, if required, Virtual Private Network (VPN) remote access for internal stadium personnel.

Figure 7. Internet Edge

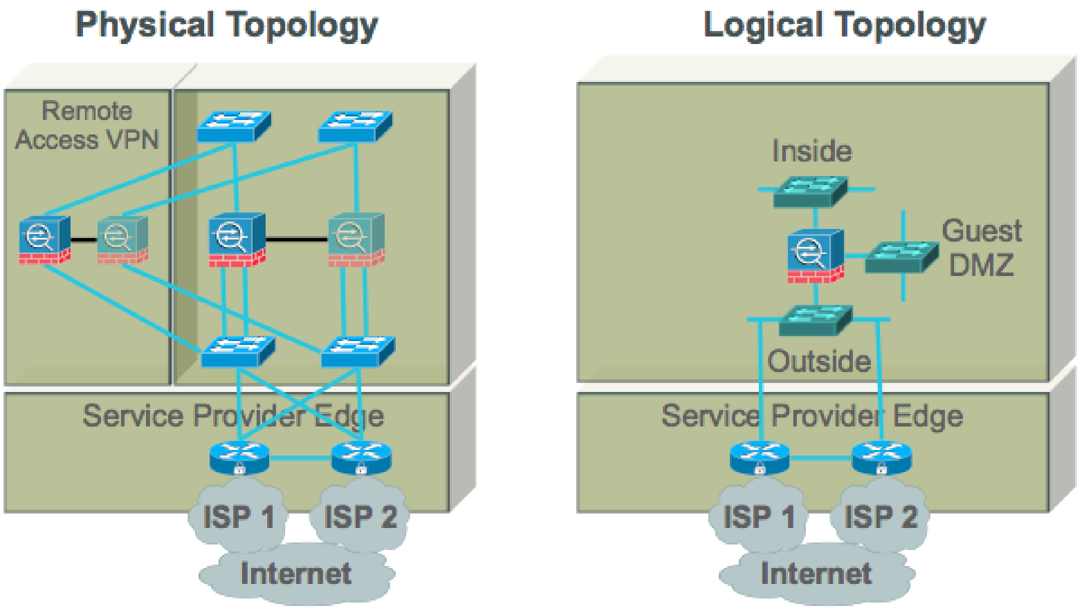

Attributes of the Internet Edge are as follows:

- Internet Edge Switches are full IGP routing peers to the Core (not stub)
- VLANs or separate switches may be configured as Inside (towards Core), Guest DMZ for wired guests, and Outside (towards Internet) Interfaces
- Internet Edge connects to the Service Provider Edge routers

Wireless Services Block

Similar to other service blocks in the Connected Stadium network, the Wireless Services Block (see Figure 8) uses a pair of switches that act as the aggregation point for the venue's wireless equipment (i.e., Wireless LAN controllers, Firewalls and Authentication equipment). This modular design approach makes it possible to perform software and hardware upgrades and troubleshoot with little disruption to the rest of the network.

NOTE: WLAN controllers may connect directly to the Core/Distribution switches. However, the ARP and MAC capacity requirements and design should be followed.

Wireless Service Block Design and Switch Recommendations

Figure 8. Wireless Service Block: Physical Design

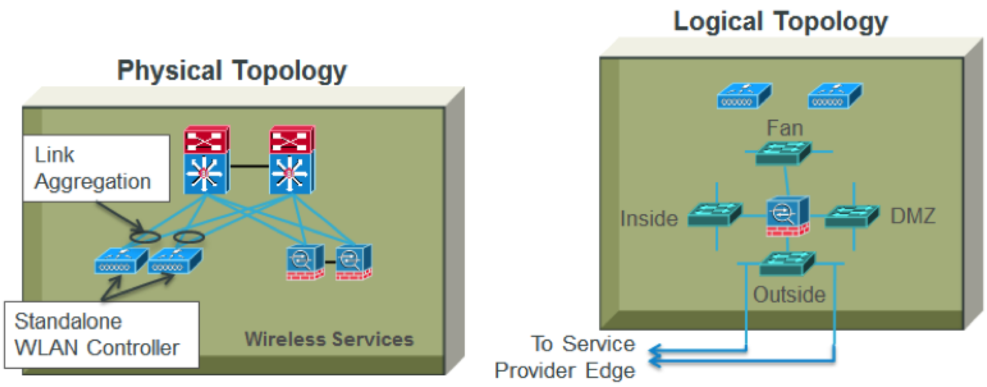

Attributes of the Wireless Services block are as follows:

- Wireless Service switches act as the aggregation point of Wireless LAN controllers, Firewalls and Authentication equipment (if authentication is required).
- Standalone Wireless LAN controllers use Link Aggregation across switches using virtual Port Channel (Nexus) or Multi-chassis Etherchannel (6500 VSS[1]) for network connection redundancy. Integrated WLAN controllers can be used within 6500 switches.
- Exiting the WLAN controller, VLANs are used to segment the different WLANs (i.e., SSIDs) from each other. A Virtual Routing Forwarding (VRF) instance is used to isolate the unsecured Guest traffic from the venue's internal global routing table. A Switched Virtual Interface (SVI) is configured within the VRF to act as the default gateway for the Wi-Fi clients.
- Since the Wireless Service Block (WSB) switches act as the default gateway for the connected Wi-Fi clients, the WSB switches must have the MAC address and ARP table capacity to support very high numbers of directly-connected Layer 3 devices. See for recommended WSB switches and their associated client capacity.

[1] For more information about the Cisco Virtual Switching System, visit: http://www.cisco.com/en/US/products/ps9336/index.html

- A dedicated pair of active/standby firewalls is recommended to provide the guest network protection from the Internet and to perform the Network Address Translation (NAT) function for outgoing client traffic. The firewall consist of Inside (towards Core), DMZ for Network servers (e.g., StadiumVision Mobile Reporter), Guest for Wi-Fi guests, and Outside (towards Internet) interfaces.
- Wireless Service switches are full IGP routing peers to the Core (not stub).
- Most stadiums will use 802.1X authentication for internal Wireless LAN users with an AAA server like Cisco's Identity Services Engine (ISE).
- For on-boarding the fan Wi-Fi clients, CMX Connect or Enterprise Mobility Services Platform (EMSP) are recommended. If using on-site CMX Connect servers, they would be placed on a DMZ interface behind a firewall for protection.

Wireless Service Block Design Details

The Wireless Service Block aggregates the wireless equipment into their own set of redundant switches providing a self-contained module that can be more easily upgraded and serviced without interrupting the whole network. These switches provide the default gateway for the WLAN device and routes their traffic according to their particular user profile (e.g., stadium admin are routed to internal servers, PoS and Ticketing scanners to their servers, fans to the Internet). Because Sports and Entertainment venues host thousands or tens of thousands of Fan Wi-Fi devices, these switches must have the capacity to handle the load. Therefore, MAC and ARP table capacity must be capable of handling the number Wi-Fi devices the network is designed for.

NOTE: Typical take-rates are calculated as 30% of the venue's seating capacity.

For non-Fan WLANs, traffic exits the controllers via a dedicated VLAN. A Switch Virtual Interface (SVI) is configured in the WSB switch and is routed into the network via the global routing table.

By far, access for Fan devices presents the vast majority of the devices the WLAN must support. The Fan WLAN consists of a very large /16 subnet providing approximately 64000 IP addresses. Note that to support this large of a subnet; the WLAN must have broadcast forwarding disabled and peer-to-peer blocking set to drop. Since the Fan

WLAN is an unsecured network, the SVI that service these Fan clients are placed into a Virtual Route Forward (VRF). This removes access to the global routing table of the stadium network and the associated mission-critical resources used in venue operations. A VRF default route steers Fan traffic to the ASA Firewall and to the Internet and the ASA firewall has a route pointing back to the Fan Wi-Fi VLAN within the VRF as shown in Figure 9.

Figure 9. Wireless Service Block Design Details

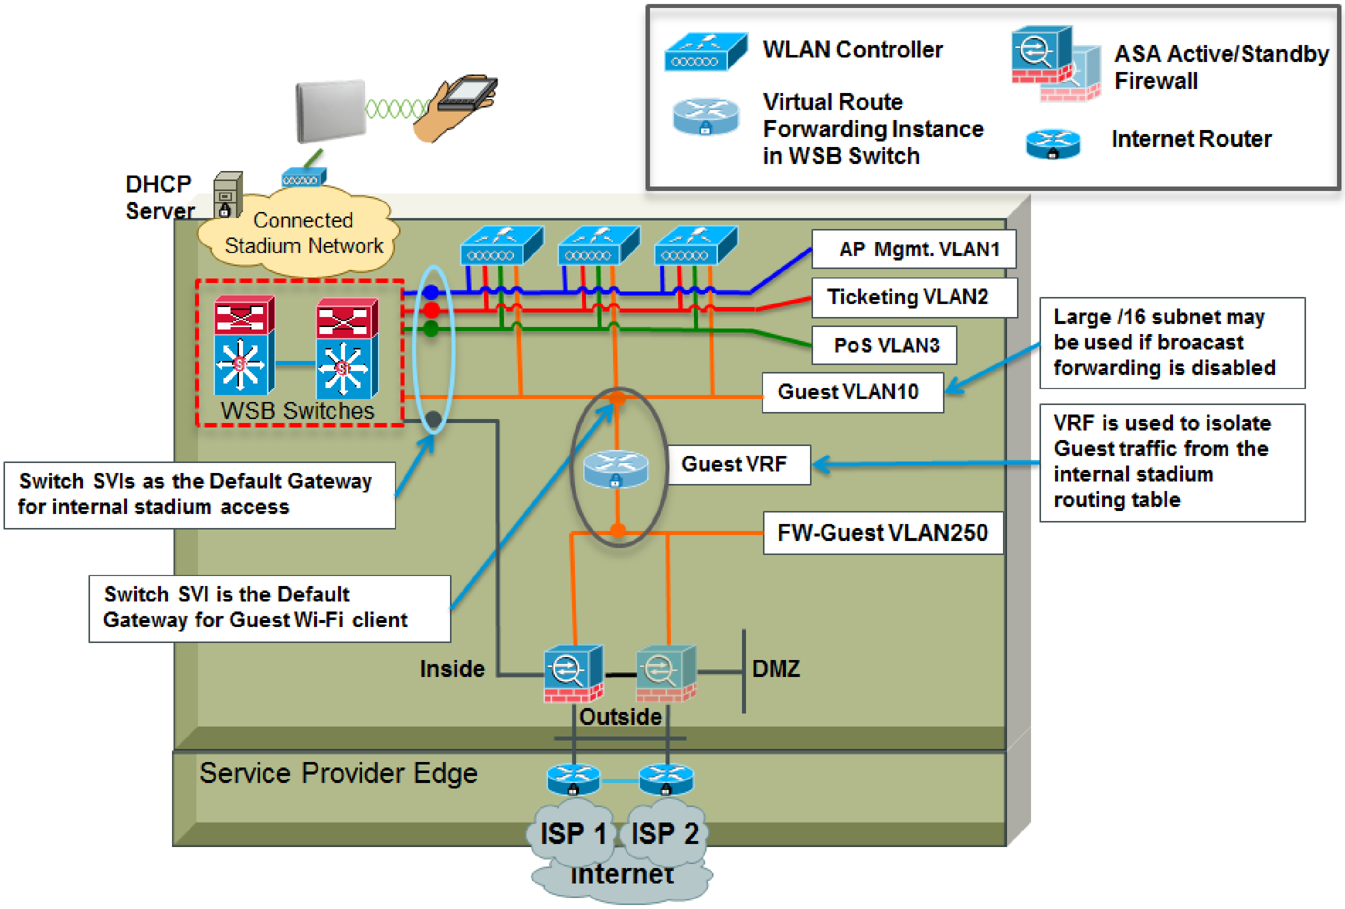
WLAN Controller
Virtual Route Forwarding Instance in WSB Switch
ASA Active/Standby Firewall
Internet Router
DHCP Server
Connected Stadium Network
WSB Switches
AP Mgmt. VLAN1
Ticketing VLAN2
PoS VLAN3
Guest VLAN10
Guest VRF
FW-Guest VLAN250
Large /16 subnet may be used if broacast forwarding is disabled
VRF is used to isolate Guest traffic from the internal stadium routing table
Switch SVIs as the Default Gateway for internal stadium access
Switch SVI is the Default Gateway for Guest Wi-Fi client
Inside
DMZ
Outside
Service Provider Edge
ISP 1
ISP 2
Internet

Fan traffic flows within the Wireless Services block is shown in Figure 10. Notice that WLAN clients require obtaining an IP address via DHCP. Access to a DHCP server is provided by an IP helper address in the VRF and then through the firewall.

NOTE: This DHCP service for WLAN clients is often overlooked until late in the design process and results in a last-minute configuration scramble and/or a less-than-optimal design. Understanding these linkages between what is traditionally provided on the wired network side and what is required on the wireless network side is all part of prudent network design.

NOTE: Although the VRF isolates the unsecured traffic from the internal networks, security best practices should also be followed to ensure the infrastructure is protected. That is, the appropriate Access Lists are placed on the VRF interface to only allow authorized traffic through and deny direct access to the interface. Below is a simplified example of an ACL used on the Fan VLAN interface.

```
interface Vlan10
 description VRF_FAN_WLAN1
 ip vrf forwarding Fan
 ip address 10.172.0.1 255.255.0.0
 ip access-group protect-infrastructure in
!
ip access-list protect-infrastructure
  10 deny ip any 10.172.0.0/16
  20 deny ip any 192.168.32.0/24
  30 permit ip any any
```

For more information on Infrastructure Security see the *iACL Sample Configurations in the SAFE - Network Foundation Protection Design Guide* on Cisco.com available at: http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Security/Baseline_Security/securebasebook.html

Figure 10. Fan Traffic Flows through the Wireless Service Block

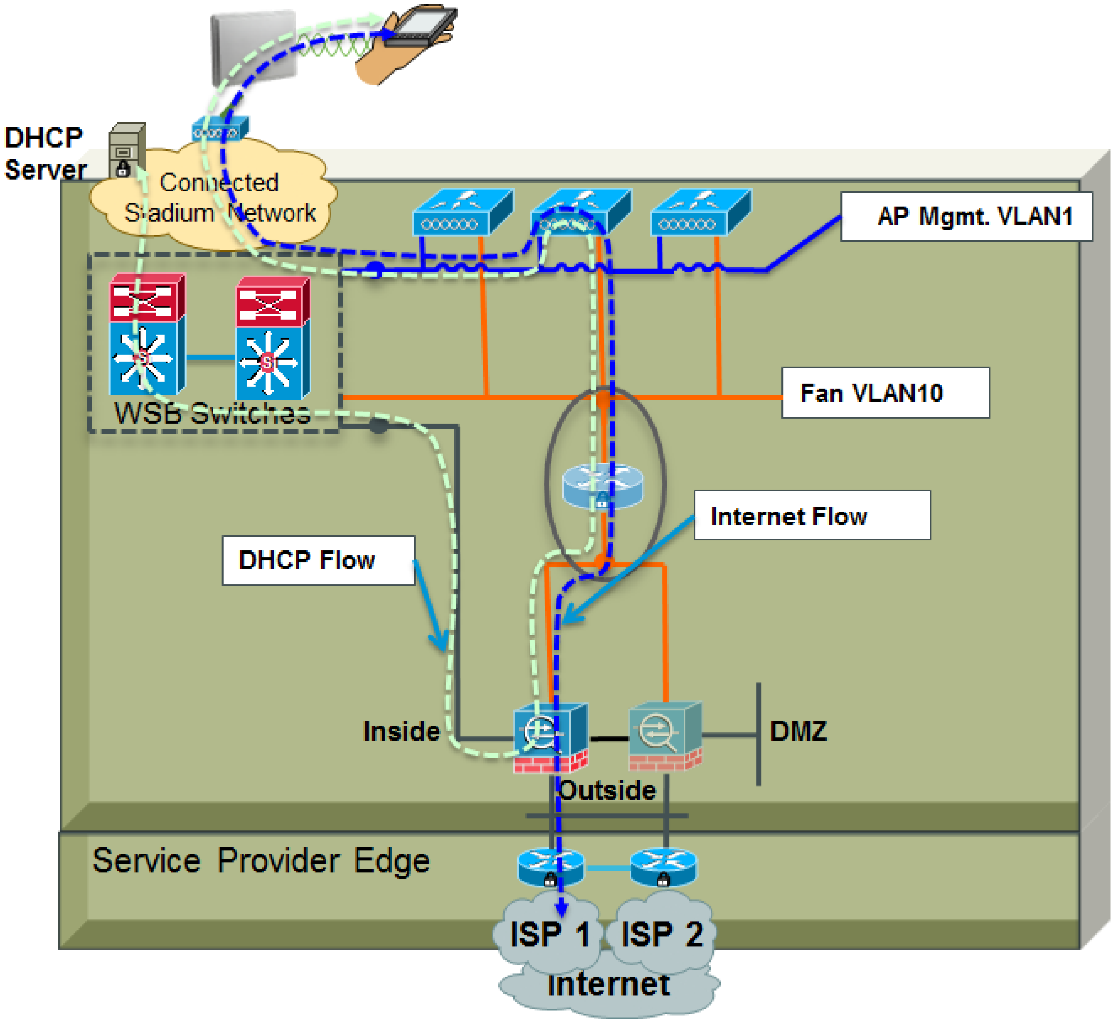

Choosing a Wireless Service Block Switch

When choosing a WSB switch, use the Wi-Fi client take rate as your guide and leave room for growth.

NOTE: Do not choose a WSB switch where the expected number of connected Wi-Fi clients are at the switch's capacity limit.

Figure 11. Choosing the Wireless Service Block Switch

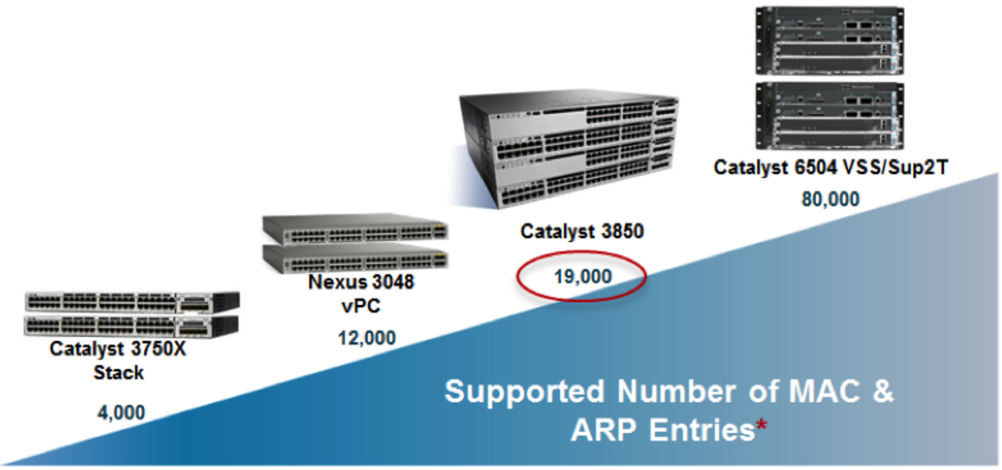

Formula	Example
Seating Capacity x expected Wi-Fi take rate = Number of Simultaneous Wi-Fi Devices	60,000 x .30 = 18,000

** 80% of lowest supported number of MAC or ARP table entries*

Wireless LAN Components

WLAN Controllers

Choosing the right type and number of WLAN controllers and Firewalls to be used in the Wireless Service block is described below.

The recommended WLAN controllers are the 5500 and 8500 series standalone WLAN controllers and the WISM2 integrated WLAN controller for the 6500. Key attributes to consider when choosing the WLAN controller is how many APs and concurrent client sessions the controller can support.

In most cases, the number of concurrent client sessions is the driving attribute. Below are the WLAN controller types and key design attributes.

Figure 12. WLAN Controller Redundancy Options

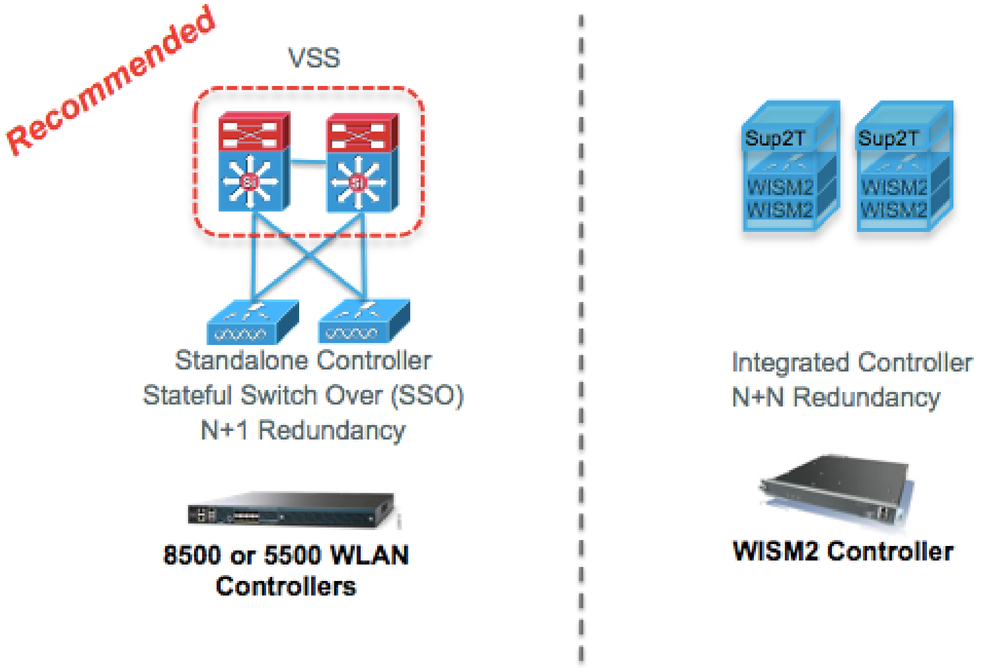

NOTE: The High Availability Stateful Switchover (SSO) model provides Box-to-Box redundancy with one controller in active state and a second controller in hot standby state. To simplify management, choose a controller that can support the total number of AP's at the venue. This reduces the configuration to only that of the active WLAN controller.

Determining the Number of WLAN Controllers

The previous section discussed the two methods of connecting WLAN controllers and this section will describe how to determine the number of WLAN controllers required for a specific sized venue.

The example below assumes a venue that seats sixty thousand visitors and 30% of those visitors will be Wi-Fi users and will be connected simultaneously. The Wireless LAN team has determined that this venue requires 500 APs for full stadium coverage.

TIP: 20-30% is a good range to use as a take rate.

Figure 13. WLAN Controller Design

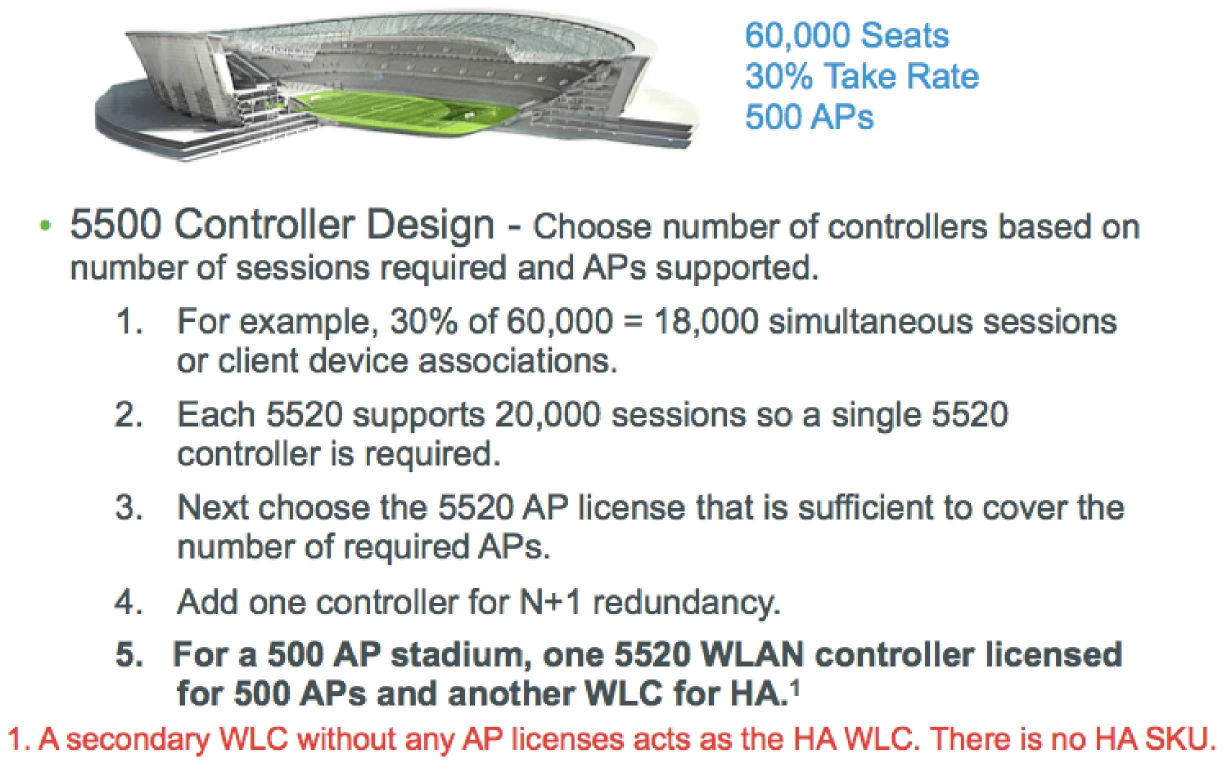

60,000 Seats
30% Take Rate
500 APs

- 5500 Controller Design - Choose number of controllers based on number of sessions required and APs supported.
  1. For example, 30% of 60,000 = 18,000 simultaneous sessions or client device associations.
  2. Each 5520 supports 20,000 sessions so a single 5520 controller is required.
  3. Next choose the 5520 AP license that is sufficient to cover the number of required APs.
  4. Add one controller for N+1 redundancy.
  5. **For a 500 AP stadium, one 5520 WLAN controller licensed for 500 APs and another WLC for HA.[1]**

[1] A secondary WLC without any AP licenses acts as the HA WLC. There is no HA SKU.

Sizing the Firewall for Guest Access

The section above described how to calculate the number of Wireless LAN controllers and now, we need to determine the size of the ASA 5500 firewalls required to support this venue. First, let’s take a look at the traffic flow through the Wireless Service Block. There are a couple of things that need to be highlighted here.

First, is that wireless clients will require an IP address via DHCP. You will need to be clear on how that is accomplished. The Wireless LAN controller can act as a DHCP server or as a DHCP proxy. Cisco recommends that the controller not do either, but rather allow the DHCP request to pass through to the network to be handled via the Wireless Service Block switch via an IP helper configured on an SVI for the Guest VLAN.

Because this DHCP traffic is coming from an unsecured network, using a firewall as a DHCP proxy with tight access control lists is a secure way to pass that DHCP traffic to the DHCP server. Once the client has an IP address, it is able to access the Internet.

Another thing to note is that the firewall will be performing thousands of Network Address Translations, which must be considered when determining the number of public IP addresses required in the NAT pool and the required capacity of the firewall. So let’s take a look at determining which ASA firewall should be used.

A set of redundant ASA 5500 series firewalls act as the secure gateway to the Internet for guests. 0 shows how traffic flows through the Wireless Service Block. As with determining the number of WLAN controllers required to provide WLAN access, a similar calculation is required for choosing the right Firewall.

Using the same take rate calculation of 18,000, this is the number of simultaneous Wireless LAN devices associated to and using the network. We can estimate the number of sessions or connections through the firewall. It’s important to realize that applications use several connections when accessing web pages or other data sources on the Internet. And it’s likely that devices like smartphones will have more than one application running at one time. So we need to compensate for that fact and use a multiplier for each device. We recommend using a minimum of 20 connections per device as the multiplier. As you can see in Figure 14, this gives us 360,000 simultaneous connections through the firewall.

To choose a firewall with sufficient capacity, we’ll want to find one that this number of connections does not exceed 80% of the firewall’s connection capacity. So doing the math, we need a firewall that can handle at least 450,000 connections.

For this venue, a good choice is the ASA 5525-x that can support up to 500,000 connections.

NOTE: Once the estimate of the number of connections is made, that number can be used to determine the number of Internet-routable addresses that are required in the NAT pool. Since there’s approximately 65,000 udp and tcp ports available and when overloading the number of internal addresses to external addresses is assumed, the number of connections divided by 65,000 will be the number of external IP addresses required in the NAT pool. In the case of 450,000 connections, the NAT pool will require at least, 7 external Internet-routable IP addresses.

Figure 14. Sizing the Firewall for the Wireless Services Block

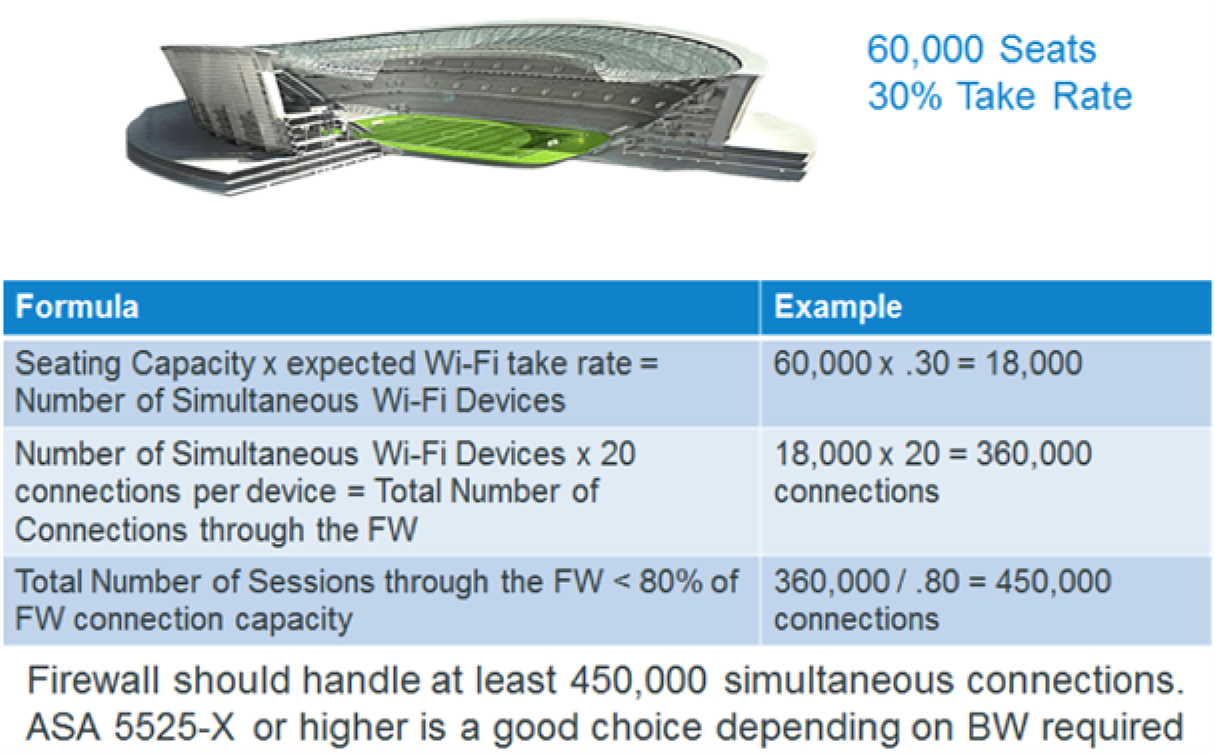

60,000 Seats
30% Take Rate

Formula	Example
Seating Capacity x expected Wi-Fi take rate = Number of Simultaneous Wi-Fi Devices	60,000 x .30 = 18,000
Number of Simultaneous Wi-Fi Devices x 20 connections per device = Total Number of Connections through the FW	18,000 x 20 = 360,000 connections
Total Number of Sessions through the FW < 80% of FW connection capacity	360,000 / .80 = 450,000 connections

Firewall should handle at least 450,000 simultaneous connections. ASA 5525-X or higher is a good choice depending on BW required

When choosing a firewall for a venue, the attributes you want to look for are the number of connections, whether it supports high availability, and throughput. As Wi-Fi gets better with technology, the amount of bandwidth consumption grows. Be sure to take that into account by accommodating the bandwidth of the current Internet circuits and allowing additional capacity for future growth.

Table 2. Firewalls for the Wireless Services Block

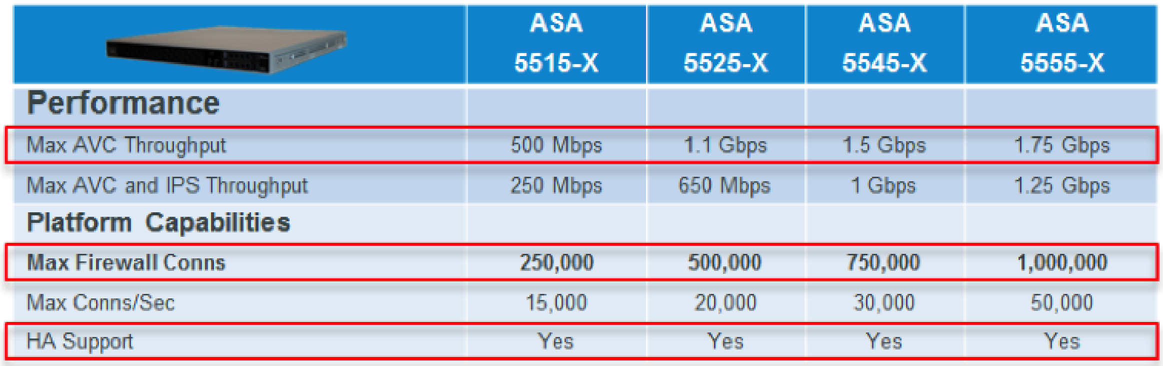

	ASA 5515-X	ASA 5525-X	ASA 5545-X	ASA 5555-X
Performance				
Max AVC Throughput	500 Mbps	1.1 Gbps	1.5 Gbps	1.75 Gbps
Max AVC and IPS Throughput	250 Mbps	650 Mbps	1 Gbps	1.25 Gbps
Platform Capabilities				
Max Firewall Conns	**250,000**	**500,000**	**750,000**	**1,000,000**
Max Conns/Sec	15,000	20,000	30,000	50,000
HA Support	Yes	Yes	Yes	Yes

NOTE: Means reader take note. Notes contain helpful suggestions or references to material not covered in the document.

Figure 15. Wireless Service Switch to 5508 Standalone WLAN Controller Configurations

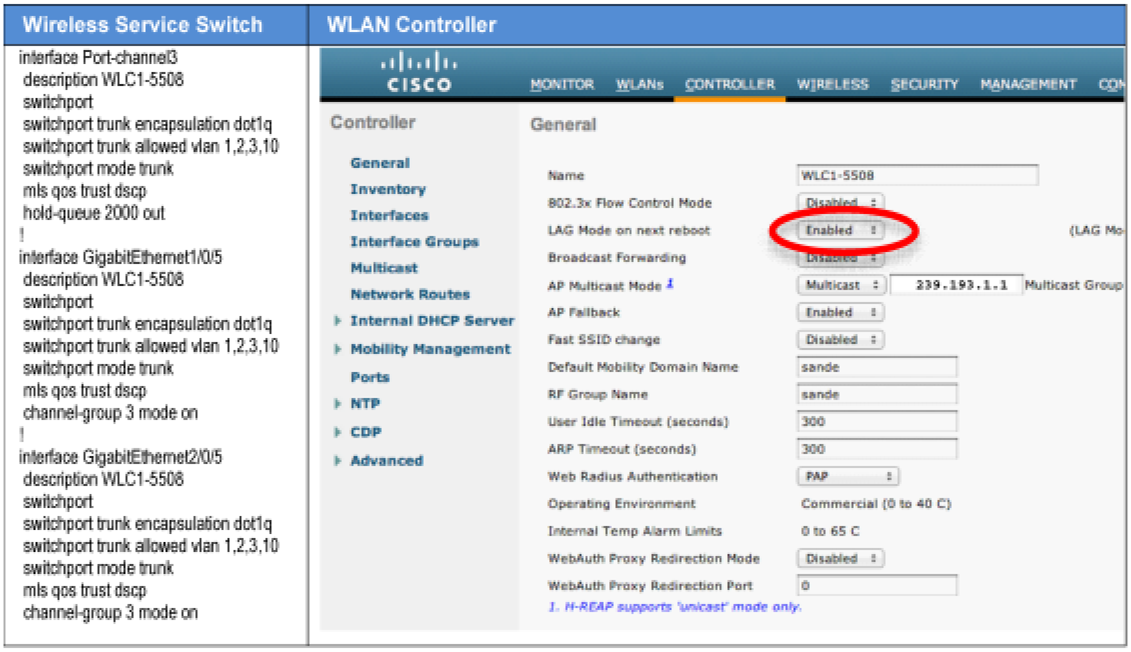

```
interface Port-channel3
 description WLC1-5508
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,2,3,10
 switchport mode trunk
 mls qos trust dscp
 hold-queue 2000 out
!
interface GigabitEthernet1/0/5
 description WLC1-5508
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,2,3,10
 switchport mode trunk
 mls qos trust dscp
 channel-group 3 mode on
!
interface GigabitEthernet2/0/5
 description WLC1-5508
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,2,3,10
 switchport mode trunk
 mls qos trust dscp
 channel-group 3 mode on
```

General

Name: WLC1-5508
802.3x Flow Control Mode: Disabled
LAG Mode on next reboot: Enabled (LAG Mo
Broadcast Forwarding: Disabled
AP Multicast Mode [1]: Multicast 239.193.1.1 Multicast Group
AP Fallback: Enabled
Fast SSID change: Disabled
Default Mobility Domain Name: sande
RF Group Name: sande
User Idle Timeout (seconds): 300
ARP Timeout (seconds): 300
Web Radius Authentication: PAP
Operating Environment: Commercial (0 to 40 C)
Internal Temp Alarm Limits: 0 to 65 C
WebAuth Proxy Redirection Mode: Disabled
WebAuth Proxy Redirection Port: 0

1. H-REAP supports 'unicast' mode only.

NOTE: Repeat the configuration shown above for each WLAN controller.

Switch Port Configuration

Below is an example of a switch configuration required for connecting WLC.

```
interface Port-channel<GRP1>
description <WLC name> Port Channel
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan <vlan-list>
switchport mode trunk
switchport nonegotiate
mls qos trust cos
spanning-tree portfast trunk
interface GigabitEthernet<port>
description <WLC name> Port x
switchport
```

```
channel-group <GRP1> mode on
```

To correctly distribute the AP to WLC CAPWAP tunnels plus the client traffic, the switch Etherchannel algorithm must be configured to perform load balancing based on source and destination IP addressing using the following global configuration command:

```
port-channel load-balance src-dst-ip
```

Other Design Considerations

Below are other design considerations for designing and deploying a Connected Stadium network.

Device Naming Convention

The following naming convention represents one option, if the floor name is included in the room numbering scheme:

<ORGANIZATION>_<DEVICE TYPE>-[ZZ]_<ROOM DESIGNATION>_YY

ORGANIZATION	A unique name used by the venue owner to identify venue
DEVICE TYPE	The model or type of device
ROOM DESIGNATION	The room ID where the equipment is deployed
YY	The IDF closet number
ZZ	An optional value in the event an IDF has multiple switch stacks or devices deployed

The following hostname demonstrates how the naming convention would be applied to one of multiple Catalyst 3750 stacks deployed on the Event level deployed in a room with the room ID EVT100 in IDF 05.

STAD_Catalyst 3750-01_EVT100_05

IP Addressing

In dividing the IP address space within the stadium apply the following general principles:

- Only use the Official league-recognized (e.g., 10.x.x.x) address space for the venue's systems and hosts that might require league access.
- Note and integrate existing IP address space within the schema.

- Use RFC 1918 private address spaces for all areas and NAT/PAT publicly through the SP Internet gateways.
- Use the RFC 2365, Administratively Scoped IP Multicast, address space. This address space assignment has become the defacto “private” space for use by enterprise environments.
- Use up to 127 separate IDFs within the stadium required subnetting using 7 bits of address space within the 172.16-31 and the league address space, while accommodating high concentrations of devices such as IP Phones, DMPs, wireless APs and PoS terminals where /24 subnets are needed.
- Address space was also carved out for loopbacks, point-to-point Layer 3 links, DC server farm subnets (including multiple subnets for VMware frontend/backend/management and heartbeat VLANs), wireless access subnets and guest networks across the Layer 2 environment.

Virtual LAN (VLAN)

VLANs are used throughout the network design. Follow the below general principles when assigning VLANs within the Connected Stadium network.

- Assign a VLAN per service type. For example, the IP telephones should be assigned a Voice VLAN, APs an AP VLAN, Digital Media Players (DMPs) a Video VLAN, etc.
- VLANs should be not be propagated across IDF closets or across the network if possible. There will be exceptions for applications that have endpoints located throughout the venue that require Layer 2 connectivity (e.g., Cobranet). These applications can be best accommodated in the VSS architecture where spanning-tree is minimized while providing uplink load-balancing and essentially active/active uplink redundancy.

Routing

Unicast

The recommended routing protocol for the Connected Stadium network is Enhanced Interior Gateway Routing Protocol (EIGRP).

EIGRP is an interior gateway protocol suited for many different topologies and media. It is well-understood, simple to configure, incurs very low-overhead on the network, and offers fast converge in case of network failure.

Below are the attributes of the Connected Stadium EIGRP design.

- Access Layer Switches are configured as EIGRP Stub. The access edge interfaces on the switch stack are configured as stub interfaces within EIGRP. This means that only the core switches are EIGRP peers for exchanging routing information and no routes will be exchanged with devices connecting to the access edge ports.
- Service Block Switches are full EIGRP routing peers
- Default EIGRP Timers are used
- EIGRP routing neighbor authentication is used
- Passive-Interface default is configured and only the interfaces required for exchanging routing information (e.g., uplinks to the Core) are enabled
- Specifically configure router ID to more easily identify the router
- Log-neighbor-changes is configured to detect routing instabilities
- EIGRP is configured not to auto-summarize routes

EIGRP Router ID

EIGRP identifies each router with a unique router ID. By default, Cisco IOS assigns the highest IP address of the loopback interface(s) if available or the highest IP address of an interface as the Router ID. It is recommended to configure the router ID deterministically using the command “router-id <ip_address>” under EIGRP process.

NOTE: EIGRP neighbors reset automatically when this command is configured.

The following EIGRP router configurations are recommended:

1. Configure an EIGRP router-id on each EIGRP device:

```
Catalyst 3750# conf t
```

2. Enter configuration commands, one per line. End with CNTL/Z.

```
Catalyst 3750(config)# router eigrp <AS Number>
Catalyst 3750(config-router)# router-id <Loopback IP Address>
```

Stub Routing

Core to access routing is configured using stubs to allow summarization (using suitable IP address blocks) and reduce convergence impact throughout the network in the event of failure. Stub routing is commonly used in a hub and spoke topology. Stub routing is employed in the Cisco Connected Stadium network to simplify routing decisions and convergence within the network and is available in the lower-cost IP Base IOS version of software used in the Access layer switches.

The access layer switch stacks are dual-homed to two core switches. This configuration adds redundancy and introduces unique issues, and the stub feature helps to address some of these issues. The stub feature ensures routes learned by the access switch from the core are not advertised to another switch and avoid becoming a transit for traffic destined for the network core.

Example 1. Multi-Homed Access Switch and Stub - Example

```
Access Switch: Connecting to Core Device
!
router eigrp <AS Number>
```

```
 no auto-summary
 eigrp router-id <Loopback IP Address>
 eigrp stub connected
 network <subnet> <mask>
!
```

Passive Interface

The “passive-interface” command disables the transmission of EIGRP hello packets on an interface. Unlike IGRP or RIP, EIGRP sends hello packets in order to form and sustain neighbor adjacencies. Without a neighbor adjacency, EIGRP can't exchange routes with a neighbor. If an unexpected neighbor appears on a non-passive interface this may confuse the routing decisions created by the design.

The more deterministic way to define the routing behavior is to use the passive-interface default command. The following commands demonstrate the use of the “passive-interface” command:

```
!
router eigrp <AS Number>
 passive-interface default
 no passive-interface TenGigabitEthernet1/0/1
 no passive-interface TenGigabitEthernet2/0/1
 no passive-interface TenGigabitEthernet3/0/1
 no passive-interface TenGigabitEthernet4/0/1
 no auto-summary
!
```

“Passive-interface default” disables all interfaces except those explicitly configured to work. This ensures that the routing protocol works only on the interfaces desired.

EIGRP Log Neighbor Changes

The following EIGRP router configurations are recommended:

- “eigrp log-neighbor-changes” enables the logging of neighbor adjacency changes to monitor the stability of the routing system and to help detect problems

NOTE: Although Catalyst and Nexus CLI are very similar, there are differences on how configuration is done. See the following document for those differences: http://docwiki.cisco.com/wiki/Cisco_NX-OS/IOS_EIGRP_Comparison.

Multicast

Overview

The Connected Stadium network uses IP Multicast for delivering a number of services. To satisfy the unique requirements of these services the Connected Stadium network employs Protocol Independent Multicast (PIM) Sparse mode routing with two redundancy strategies.

Below are the attributes of the Connected Stadium Multicast network design.

- Uses PIM Sparse Mode multicast routing protocol
- Uses a set of Rendezvous Points (RP) on the Core switches for general stadium MC support
- Uses a set of Rendezvous Points (RP) on the Video Distribution switches for stadium video MC support
- For general stadium MC, uses Anycast RP & Multicast Source Discovery Protocol (MSDP) for RP redundancy if using Nexus Core switches
- Anycast RP provides an active/active redundancy strategy
- For stadium video MC, uses Prioritycast RP for source and RP redundancy
- Prioritycast RP provides an active/standby redundancy strategy
- Use ACLs and MC Boundaries to limit MC to their designated areas

Figure 16. Connected Stadium Multicast Architecture Overview

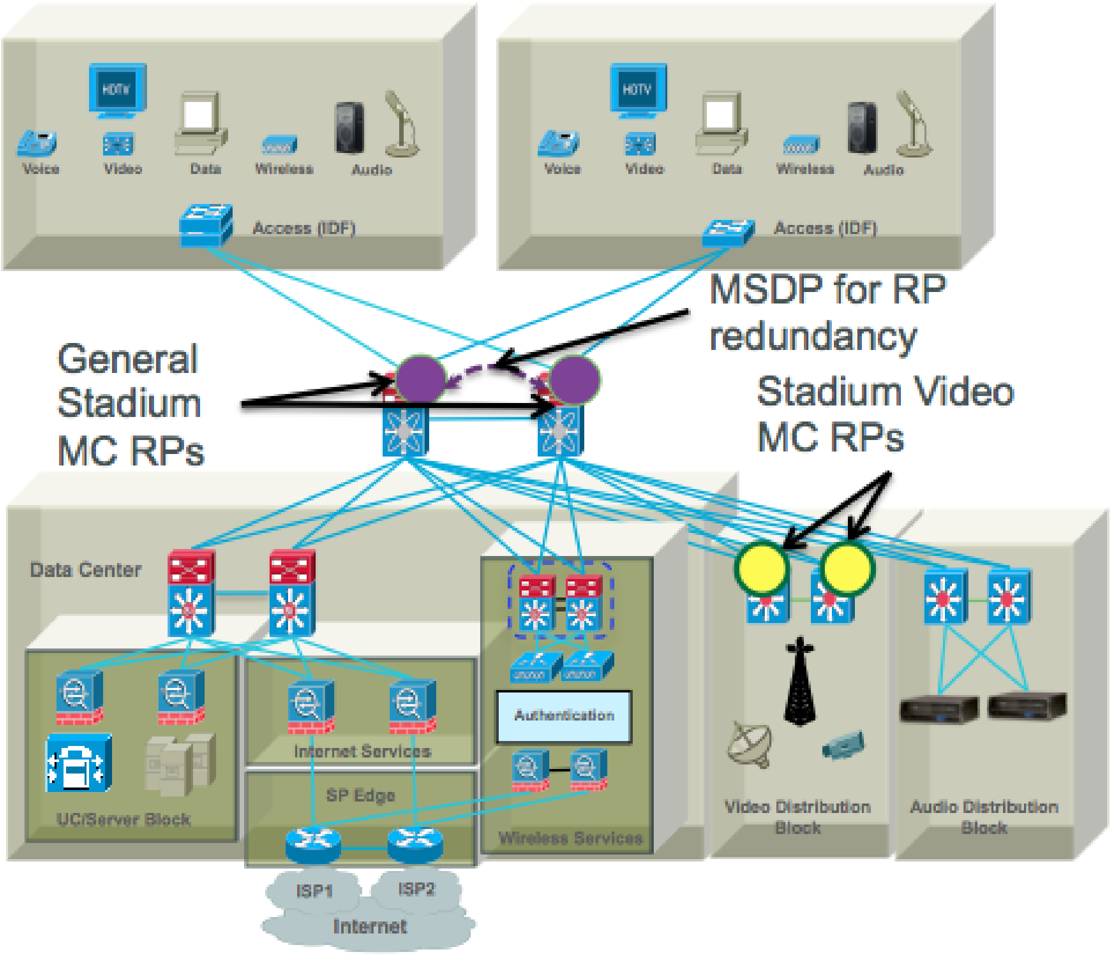

Multicast Redundancy Strategies

Why two redundancy strategies?

Anycast and Prioritycast RP redundancy strategies have some telling attributes about how multicast traffic is handled and the impact it has on the network and endpoints. Below are the attributes of each strategy.

Anycast

- Provides Source and RP Redundancy in an Active/Active redundancy model
- Sub-second failover
- Simple to implement. Sources can be anywhere in the network with no special IP addressing

- An RP is configured with the same address and netmask on each of the Core switches
- Multicast Source Discovery Protocol (MSDP) is configured between the Nexus Core switches and their respective RPs to share source information

Prioritycast

- Provides Source and RP Redundancy in an Active/Standby redundancy model
- Sub-second failover
- More Complex to implement. Redundant sources must use duplicate addressing with different masks
- Provides a single source stream on the network at a time. This is important to reduce the amount of traffic on the network. Especially heavy traffic like video
- Because the network controls what source traffic is allowed on the network, no vendor proprietary source sync protocol is required between sources to trigger the backup source to start streaming
- Because only a single stream is on the network at any one time, the video endpoints do not have to arbitrate between two duplicate video streams. This means lower endpoint complexity and processing power are required.

Prioritycast

Because PIM Sparse mode operates in an on-demand fashion, receivers must request a video stream using an IGMP join request. This request is received by the receiver’s local switch and is directed to a pre-configured Rendezvous Point (RP). The RP is where sources and receivers register and is how they find each other in the network. Once registered, a tree is built to connect sources and receivers that the MC stream will traverse. Reverse Path Forwarding (RPF) using the network’s unicast routing table is used to derive the shortest paths (or branches of the tree) between sources and receivers. Prioritycast uses unicast routing mechanisms to have the network act as the arbiter of what source streams traverse the network and when. This is how the active/standby redundancy strategy is implemented. Below is a description of how this is accomplished.

1. Prioritycast uses duplicate multicast video sources, each source connected to a separate VDS switch

2. Each primary multicast source & Rendezvous Point (RP) and it's backup use identical IP addresses with differing network masks
3. The primary MC source & RP uses the longest network mask and therefore, is the active source & RP on the network
4. If the primary VDS switch or uplinks or primary MC video source Ethernet link fail, the network will converge and place the backup MC video source onto the network. The transition is transparent to the video receivers due to identical source IP addresses.

Figure 17. Prioritycast in Action

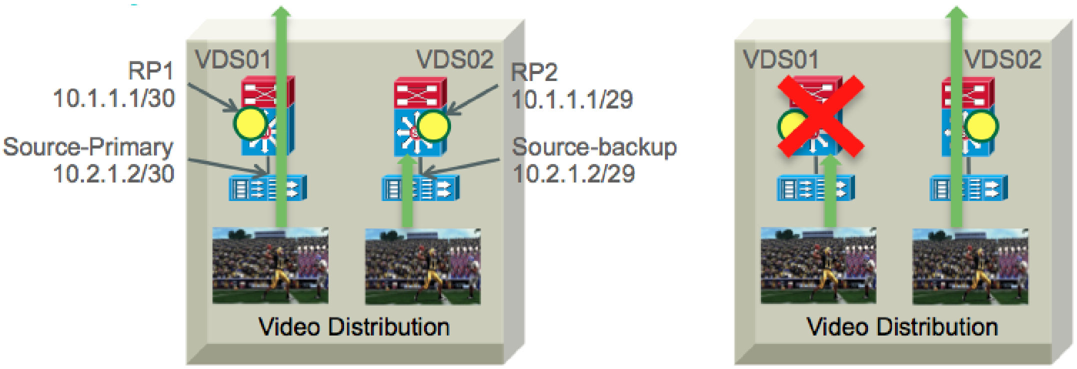

Configuration Examples for the Core, Access, and Video Distribution Switches

Basic IP Multicast on the Nexus 7000

```
feature pim
interface Ethernetx/y
ip pim sparse-mode
ip pim hello-interval 1000
!
ip pim rp-address 10.1.1.1 group-list 239.192.0.0/24
```

Basic IP Multicast on the Access Layer switch

```
ip multicast-routing distributed
!
interface TenGigabitEthernetx/0/z
description ** Up Link **
ip pim sparse-mode
!
```

```
interface GigabitEthernetx/0/z
switchport access vlan x
switchport voice vlan y
spanning-tree portfast
!
interface Vlanx
ip pim passive
!
ip pim rp-address 10.1.1.1 prioritycast-grp-acl override
!
ip access-list standard prioritycast-grp-acl
permit 239.192.0.0 0.0.0.255
```

VDS01 – Primary (Active)

```
interface Loopback1
description prioritycast-RP Primary address for Global SV HD IP
Multicast
ip address 10.1.1.1 255.255.255.252
ip pim sparse-mode
!
ip pim rp-address 10.1.1.1 prioritycast-grp-acl override
!
ip access-list standard prioritycast-grp-acl
permit 239.192.0.0 0.0.0.255
!
interface GigabitEthernetx/y
description Channel MMM, Primary Source 10.2.1.2
ip address 10.2.1.1 255.255.255.252
ip pim sparse-mode
!
interface GigabitEthernetx/y
description Primary DCM Mgmt Port
ip address 192.168.30.1 255.255.255.248
!
interface TenGigabitEthernet0/x
```

```
description ** Up Link to n7k
ip pim query-interval 1
ip pim sparse-mode
```

VDS02 – Secondary (Standby)

```
interface Loopback1
description prioritycast-RP Secondary address for Global SV HD IP Multicast
ip address 10.1.1.1 255.255.255.248
ip pim sparse-mode
!
ip pim rp-address 10.1.1.1 prioritycast-grp-acl override
!
ip access-list standard prioritycast-grp-acl
permit 239.192.0.0 0.0.0.255
!
interface GigabitEthernetx/y
description Primary DCM Mgmt Port
 ip address 192.168.30.2 255.255.255.248
!
interface GigabitEthernetx/y
description Channel MMM, Secondary Source 10.2.1.2
ip address 10.2.1.1 255.255.255.248
ip pim sparse-mode
!
interface TenGigabitEthernet0/x
description ** Up Link to n7k
ip pim query-interval 1
ip pim sparse-mode
```

Anycast

General IP Multicast traffic is handled by PIM Sparse-mode and the Anycast redundancy strategy. The use of PIM Sparse-mode is consistent with the Prioritycast stategy. However, Multicast Source Discovery Protocol (MSDP) is used to support RP redundancy. In Anycast RP, all the RPs are configured to be MSDP peers of each other. When a source registers with one RP, an Source Advertisement (SA) message will be sent to the other RPs informing them that there is an active source for a particular multicast group. The result is that each RP will know about the active sources in the area of the other RPs. If any of the RPs were to fail, IP routing would converge and one of the RPs would become the active RP in more than one area. New sources would register with the backup RP. Receivers would join toward the new RP and connectivity would be maintained.

NOTE: The RP is normally needed only to start new sessions with sources and receivers. The RP facilitates the shared tree so that sources and receivers can directly establish a multicast data flow. If a multicast data flow is already directly established between a source and the receiver, then an RP failure will not affect that session. Anycast RP ensures that new sessions with sources and receivers can begin at any time.

Figure 18. Anycast in Action

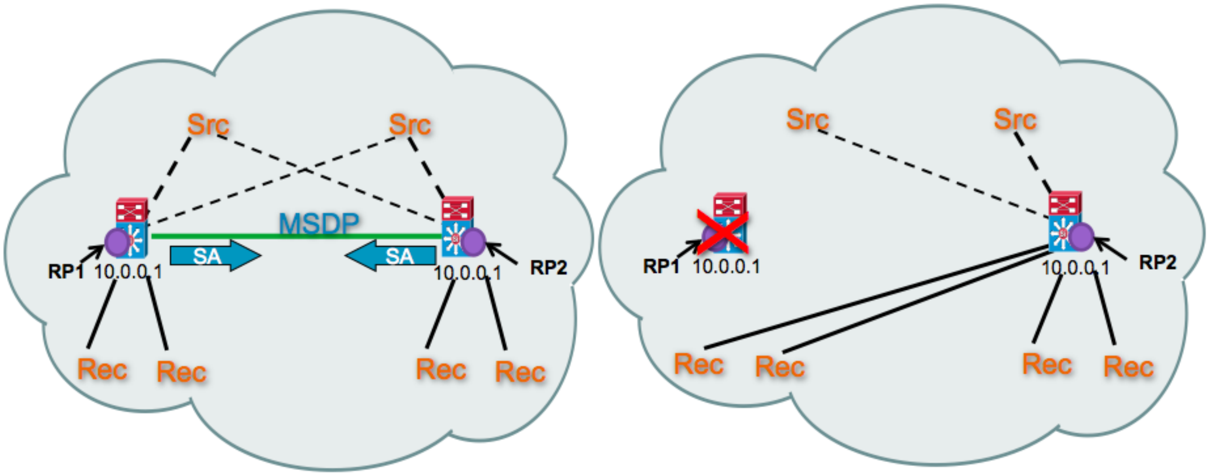

Configuration Examples for Core and Access switches

Nexus 7000 to support General IP Multicast

```
interface loopback0    ← Must be the same IP Address on both n7k routers
ip address 10.0.0.1/32
ip router eigrp 100
!
interface loopback1
ip address 10.2.2.1/32
ip router eigrp 100
!
feature msdp <<< Must be the same on both n7k routers
ip msdp originator-id 10.2.2.1
ip msdp peer 10.2.2.2
connect-source 10.2.2.1
ip msdp reconnect-interval 1
ip msdp group-limit 800 source 0.0.0.0/0
ip msdp sa-limit 10.2.2.2 2000
!
ip pim rp-address 10.0.0.1 group-list 239.193.0.0/20
```

Catalyst 3750 to Support General IP Multicast

```
ip pim rp-address <anycast-rp-addr> anycast-grp-acl override
ip access-list standard anycast-grp-acl
permit 239.193.0.0 0.0.15.255
```

Considerations when DMPs are Attached the VDS

Although it is not an intended design, the VDS switch sometimes ends up hosting some number of DMPs, depending on the situation.

Given that VDS1 is the primary (/30) and VDS2 (/29) is secondary, if a DMP is hosted on VDS2, it will not receive the multicast streams. However, a DMP on VDS1 will. This is because all traffic sourced on VDS2 will be RPF dropping against the /30 route. Also, a DMP hosted on VDS1 will not receive any streams if VDS1 loses the link to the DCM.

As an option, you can add a link between the VDS1 and VDS2. Further configuration will be required on the devices.

An separate management connection is required to manage the video sources on VDS2. This is for the same reason mentioned above about traffic sourced on VDS2 will be RPF dropped and therefore, inaccessible for managing the device.

For more information about Anycast RP see the following link.
http://www.cisco.com/en/US/docs/ios/solutions_docs/ip_multicast/White_papers/anycast.html

Multicast Stadium Applications

There are five applications that use IP multicast and leverage these two different redundancy strategies:

Cisco StadiumVision

- Up to 80 video channels streamed to Digital Media Players (DMPs) with attached HD screens distributed throughout the stadium. Uses Prioritycast IP Multicast Topology.
- Multicast control of DMP states (i.e., what is displayed on the TVs). Uses Anycast IP Multicast Topology.

Point of Sale Radiant (POS) System

- Some PoS systems use an IP Multicast based solution for server/terminal discovery communications.
- Uses Anycast IP Multicast Topology.

Security Video

- Integrated video surveillance network solution using IP Multicast to view video sourced from the video servers.
- Uses Anycast IP Multicast Topology.

VOIP Music on Hold (MoH)

- IP Multicast for delivery of Music on Hold.
- Uses Anycast IP Multicast Topology.

StadiumVision In-Suite Video

- IP Multicast video streamed from an SV-4K Digital Media Player.

NOTE: To control the distribution of video, use a TTL=1 to limit the video to the local VLAN or use ACL's or multicast boundaries to limit what VLANs can request the video.

Multicast Address Assignment

Class D address assignments (IP Multicast) utilize the Administrative Scoped IP Multicast (RFC-2365) address space. The table below shows multicast group addresses assigned to stadium applications.

Table 3. Multicast Address Assignments by Application Type

Type	Required Addresses	Range	RP Redundancy Strategy
StadiumVision Video	50-80	239.192.0.0/24	Prioritycast
StadiumVision Control	1	239.193.0.0/24	General
AP-to-WLC CAPWAP Tunnel	1 per WLC	239.193.1.0/24	General
Radiant POS	250	239.193.2-3.0/24	General
IP Telephony Music on Hold	10	239.193.5-6.0/24	General
Security	160-400	239.193.8-10.0/24	General
StadiumVision In-suite Video Source (video sourced from access switches)	1 unique address per DMP-encoded channel	239.193.20.0/24	General

Quality of Service (QoS)

Overview

The Connected Stadium network provides end-to-end Quality of Service support. Traffic is classified, marked and policed as it enters the network. This traffic is then queued according to administratively configured priority as it is carried through the network. This purposeful handling of the traffic guarantees that the network meets the requirements of the applications the stadium uses to conduct business and offers to the fans.

Reference Materials

Enterprise Quality of Service Design 4.0
http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSIntro_40.html

Campus QoS Design 4.0 http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSCampus_40.html

WLAN QoS Design (BYOD CVD)
http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Borderless_Networks/Unified_Access/BYOD_Design_Guide/BYOD_AVC.html

QoS Design Strategy

The methodology used to determine when and how to use QoS is as follows:

1. Determine what applications are relevant to the business.
2. Classify each application’s traffic relative to the classes described in RFC-4594 for optimal servicing and required bandwidth allotments.
3. Determine queuing strategy (e.g., 1P7Q3T).

4. Configure QoS at each Policy Enforcement Point (PEP) and possible congestion point along the traffic path.
   a. Use Cisco's Campus and Data Center QoS policies described in the Cisco Validated Design (CVD) guides as a foundation.
   b. Customize policies as needed.

Figure 19. Connected Stadium QoS Implementation

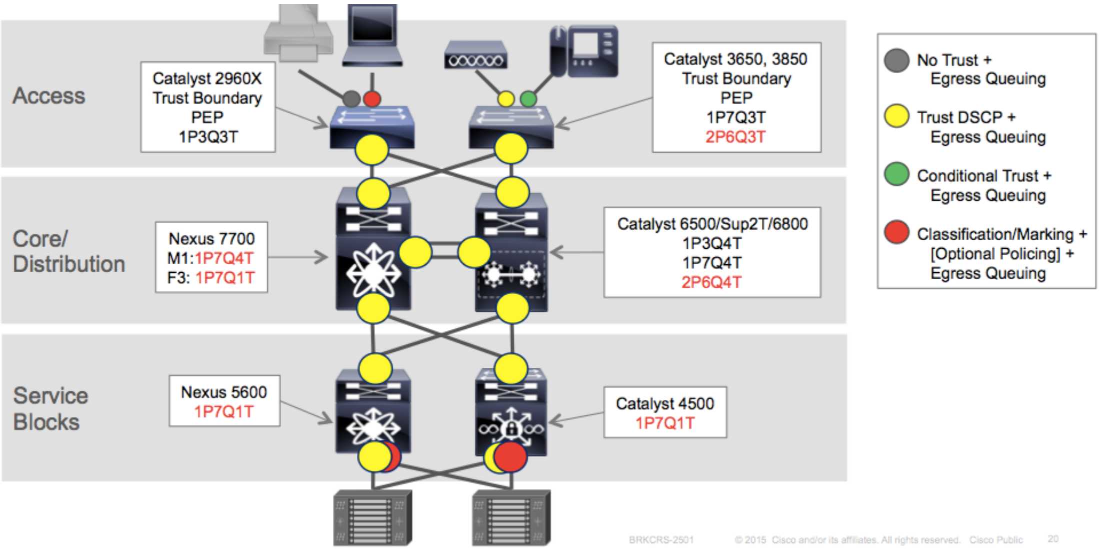

NOTE: The diagram above shows the variety of switching platforms to depict within a single diagram the different queuing strategies each switch type supports. Real world designs would show common platforms used at the different layers of the network.

Business Relevant Applications

Below are a set of applications relevant to a Sports and Entertainment customer and are used in example configurations.

- IP Telephony
- StadiumVision (Video & IEEE-1588 for time synchronization)
- Video Surveillance
- Network Audio (audio streams & IEEE-1588 for time synchronization)

- Point of Sale
- Ticketing
- Building Management Systems (HVAC, ACU)

General Switch Platform Characteristics

Cisco's switching platforms have evolved over time to use a more consistent set of configuration using the Modular QoS Command Line Interface (MQC). However, there are still differences to be aware of as mentioned below.

- Catalyst 2960-X / 3560 / 3750 are the last platforms to use Multilayer Switch QoS (MLS QoS).
    - QoS is disabled by default and must be globally enabled with mls qos command.
    - Once enabled, all ports are set to an untrusted port-state.
- Catalyst 3650/3850 and 4500 use IOS Modular QoS Command Line Interface (MQC).
    - QoS is enabled by default.
    - All ports are trusted at layer 2 and layer 3 by default.
- Catalyst 6500/6800 use Cisco Common Classification Policy Language (C3PL) QoS.
    - QoS is enabled by default (Sup2T) – Disabled by default (Sup720).
    - All ports are trusted at layer 2 and layer 3 by default.
- C3PL presents queuing policies similar to MQC, but as a defined "type" of policy.
    - Nexus 7000/7700 use NX-OS QoS.
    - QoS is enabled by default.
    - All ports are trusted at layer 2 and layer 3 by default.
- NX-OS presents queuing policies similar to MQC, but as a defined "type" and with default class-map names.

Ingress QoS Models

The ingress QoS model applies either a port trust state or an explicit classification and marking policy to the switch ports (or VLANs, in the case of VLAN-based QoS), as well as optional ingress policers and ingress queuing (as required and supported).

Trust Boundary

The Network Edge is where the Trust Boundary is important to establish and determines how endpoint traffic is handled as it enters the network.

Trust has the following modes:

- **No Trust**—COS and DSCP are marked to 0 for Best Effort traffic handling.

  Port goes to a no trust state after a Conditional Trust fails to see the Cisco device via CDP.

- **Conditional Trust**—After the switch identifies the Cisco device via Cisco Discovery Protocol (CDP), endpoints are trusted and consequently the device traffic QoS markings are honored.

```
interface range GigabitEthernet1/0/x - y
 description Network Edge Port
switchport access vlan 102
 switchport mode access
 switchport voice vlan 103
 trust device cisco-phone
```

NOTE: Do not use the ‘trust device media-player’ for the Cisco SV-4K or SV-2K Digital Media Players. These new devices are not recognized by switches and therefore will put the port into a no trust state and not invoke attached Service-policies.

- **Trust**—Endpoint COS and/or DSCP markings are honored.

  Ports are trusted by default. No need for a command.

Once ingress traffic has been trusted, classified, and (optionally) policed at the network edge, then the ingress QoS model for all inter-switch ports can be set to trust the DSCP markings of all incoming packets.

NOTE: Most of the modern Cisco switch platforms have QoS enabled by default with basic queuing and all ports are trusted.

Figure 20. Trust Boundaries – Access Layer

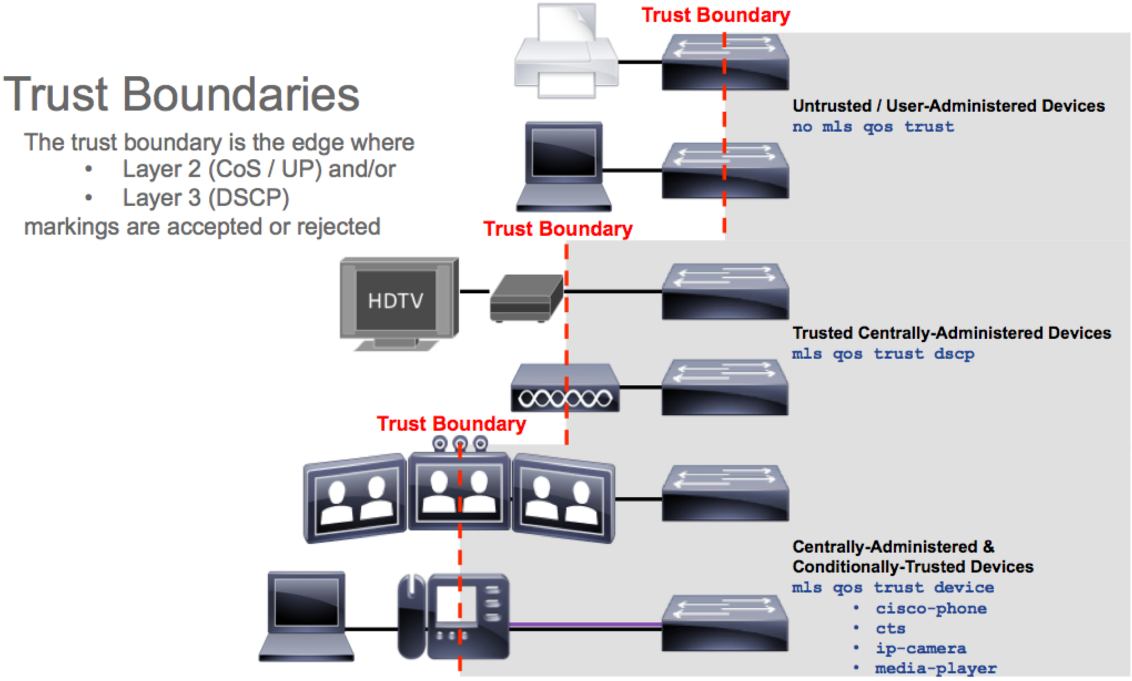
Trust Boundaries
The trust boundary is the edge where
• Layer 2 (CoS / UP) and/or
• Layer 3 (DSCP)
markings are accepted or rejected
Trust Boundary
Untrusted / User-Administered Devices
no mls qos trust
Trust Boundary
HDTV
Trusted Centrally-Administered Devices
mls qos trust dscp
Trust Boundary
Centrally-Administered &
Conditionally-Trusted Devices
mls qos trust device
• cisco-phone
• cts
• ip-camera
• media-player

Application Classifying, Marking and Policing

Business applications are mapped to specific application classes based on industry standards and best practices and the traffic's specific bandwidth, packet loss, latency and jitter requirements. Where applicable, policing is used to protect priority queues from devices over-utilizing their assigned bandwidth allocation and as a rudimentary call admission control mechanism.

Classify and Mark

The Cisco Connected Stadium network will use Cisco's Campus QoS framework that follows the RFC-4594 QoS model and application classes as shown in Table 4. Table

5 shows the representative set of relevant applications used in stadiums and their associated QoS markings and where they're mapped into the framework.

Application traffic can be classified a few ways. It can be classified by their CoS or DSCP values marked by the endpoints or with extended access lists (ACLs) which can match packets on Source/Destination IP address, protocol type, and UDP/TCP port numbers. Once classified, traffic can be, if required, policed and queued according to its traffic profile requirements.

Classification & Marking Recommendations

- Always enable QoS policies in hardware—rather than software—whenever a choice exists.
- Classify and mark applications as close to their sources as technically and administratively feasible.
- Use DSCP marking whenever possible.
- Follow standards-based DSCP PHB markings to ensure interoperability and future expansion.

Table 4. RFC 4594-based 12-Class QoS Model and Application Classes

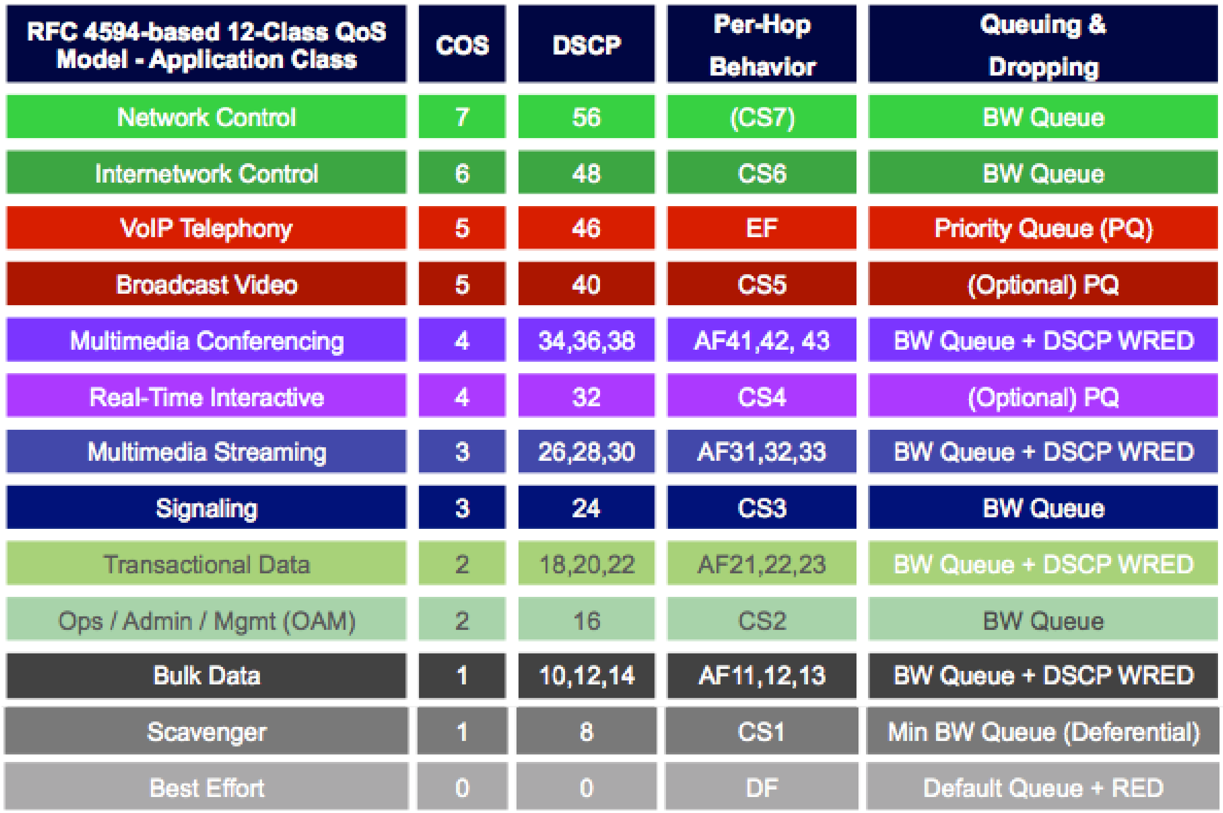

RFC 4594-based 12-Class QoS Model - Application Class	COS	DSCP	Per-Hop Behavior	Queuing & Dropping
Network Control	7	56	(CS7)	BW Queue
Internetwork Control	6	48	CS6	BW Queue
VoIP Telephony	5	46	EF	Priority Queue (PQ)
Broadcast Video	5	40	CS5	(Optional) PQ
Multimedia Conferencing	4	34,36,38	AF41,42, 43	BW Queue + DSCP WRED
Real-Time Interactive	4	32	CS4	(Optional) PQ
Multimedia Streaming	3	26,28,30	AF31,32,33	BW Queue + DSCP WRED
Signaling	3	24	CS3	BW Queue
Transactional Data	2	18,20,22	AF21,22,23	BW Queue + DSCP WRED
Ops / Admin / Mgmt (OAM)	2	16	CS2	BW Queue
Bulk Data	1	10,12,14	AF11,12,13	BW Queue + DSCP WRED
Scavenger	1	8	CS1	Min BW Queue (Deferential)
Best Effort	0	0	DF	Default Queue + RED

Table 5. Connected Stadium Applications and Their QoS Mapping

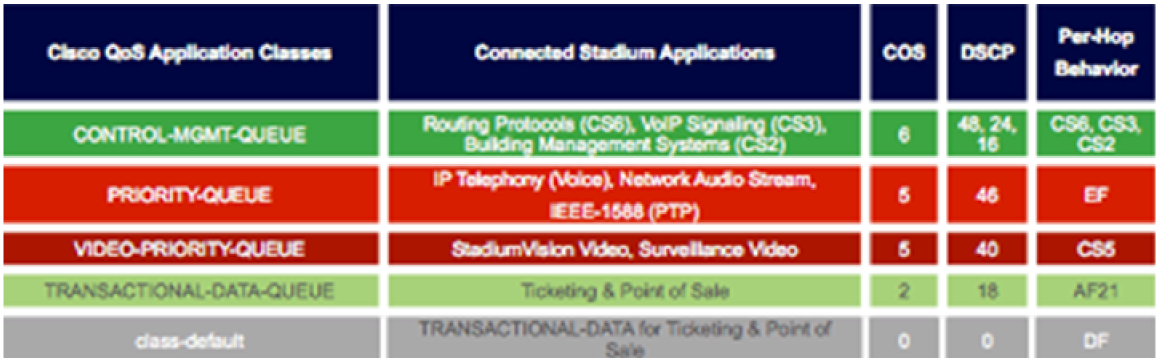

Cisco QoS Application Classes	Connected Stadium Applications	COS	DSCP	Per-Hop Behavior
CONTROL-MGMT-QUEUE	Routing Protocols (CS6), VoIP Signaling (CS3), Building Management Systems (CS2)	6	48, 24, 16	CS6, CS3, CS2
PRIORITY-QUEUE	IP Telephony (Voice), Network Audio Stream, IEEE-1588 (PTP)	5	46	EF
VIDEO-PRIORITY-QUEUE	StadiumVision Video, Surveillance Video	5	40	CS5
TRANSACTIONAL-DATA-QUEUE	Ticketing & Point of Sale	2	18	AF21
class-default	TRANSACTIONAL-DATA for Ticketing & Point of Sale	0	0	DF

NOTE: Network audio vendors have different DSCP value recommendations for audio streams classification than shown above but all recommend that traffic to be placed in the priority queue. Cisco recommends using DSCP EF for classifying audio.

Classify and Mark – Examples

Catalyst Switches

Class-maps are used to classify traffic.

Classify by incoming COS value

```
class-map match-any CISCO-PHONE-VOICE
match cos 5
class-map match-any CISCO-PHONE-SIGNALING
match cos 3
```

Classify with an ACL – Catalyst Switch

```
ip access-list extended IEEE-1588
remark PTP for DMPs and Audio Systems
permit udp any host 224.0.1.129
permit udp any host 224.0.1.130
permit udp any host 224.0.1.131
permit udp any host 224.0.1.132
!
class-map match-any IEEE-1588
match access-group name IEEE-1588
```

Classify with an ACL – Nexus Switch

```
ip access-list CRM-SERVER
permit ip 50.1.1.0/24 any
!
class-map type qos CRM-CLASS
```

NOTE: NX-OS uses the “type qos” form of the class-map, policy-map and service-policy commands.

Police, Remark and Drop

Policing is used to limit the bandwidth allocation to those applications using precious priority queue resources.

For example, due to the strict packet loss, latency and jitter requirements, IP Telephony uses the priority queue and therefore is policed on ingress to the network.

The bandwidth required by voice bearer and call signaling traffic originating from a Cisco IP Telephone is the following.

- 128 kbps for Voice traffic, marked CoS 5/DSCP EF
- 32 kbps for call signaling traffic marked CoS 3/DSCP CS3

To ensure the voice bearer and call signaling bandwidth usage are within profile, policing is configured to remark or drop any traffic exceeding the above allocations.

Policing and Remarking Recommendations

- Police traffic flows as close to their source as possible
- Whenever possible, markdown according to standards-based rules. For Example: Assured Forwarding Traffic (AF21 example)
    - Conforming AF21 traffic is marked/remarked AF21
    - Exceeding AF21 traffic is remarked AF22
    - Violating AF21 traffic is remarked AF23

Police and Remark/Drop – Examples

Policy-maps are used to apply policies on classified traffic. In the case of policing, dropping and remarking. The policies will be applied on traffic entering the interface.

Police and Remark

```
! Use a table-map to specify how to remark traffic
!
table-map policed-dscp
map from 0 to 8
map from 10 to 8
map from 18 to 8
map from 24 to 8
map from 46 to 8
!
```

Within the policy-map, set the DSCP value, police voice traffic and remark it down before transmitting it if it exceeds the specified bandwidth

```
!
policy-map CISCO-PHONE
class CISCO-PHONE-VOICE
set dscp ef
  police cir 128000 bc 8000
  conform-action transmit
  exceed-action set-dscp-transmit dscp table policed-dscp
class CISCO-PHONE-SIGNALING
set dscp cs3
  police cir 32000 bc 8000
  conform-action transmit
  exceed-action set-dscp-transmit dscp table policed-dscp
class class-default
set dscp default
!
```

Police and Drop

```
!
policy-map IEEE-1588
class IEEE-1588
  set dscp cs5
   police cir 1000000 bc 8000 conform-action transmit exceed-action
drop
class class-default
  set dscp default
```

Port-Based Versus VLAN-Based Policies

Policies deployed in the Cisco Connected Stadium network may be configured in one of three ways.

Port-based QoS—When a QoS policy is applied on a per-port basis, it is attached to a specific physical switch port and is active on all traffic received on that specific port (only). QoS policies are applied on a per-port basis, by default.

VLAN-based QoS—When a QoS policy is applied on a per-VLAN basis, it is attached to a logical VLAN interface and is active on all traffic received on all ports that are currently assigned to the VLAN.

Per-port/per-VLAN-based QoS—When a QoS policy is applied on a per-port/per-VLAN basis, it is attached to specific VLAN on a trunked port and is active on all traffic received from that specific VLAN from that specific trunked port (only). Per-port/per-VLAN QoS is not supported on all platforms and the configuration commands are platform-specific, so review the platform documentation before applying these policies.

Refer to Figure 21.

Figure 21. Port-Based Versus VLAN-Based Policies

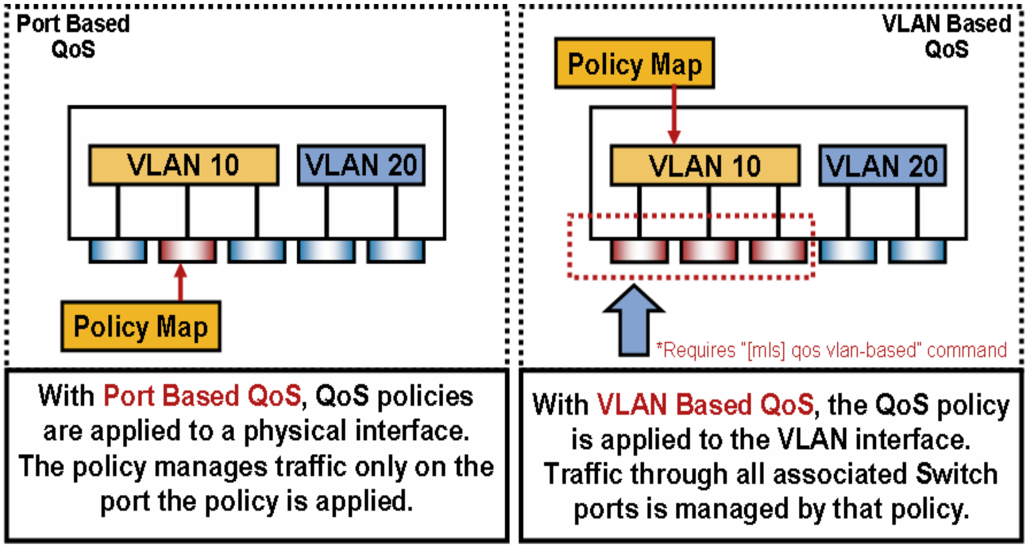
Port Based QoS
VLAN 10
VLAN 20
Policy Map
VLAN Based QoS
Policy Map
VLAN 10
VLAN 20
*Requires "[mls] qos vlan-based" command
With Port Based QoS, QoS policies are applied to a physical interface. The policy manages traffic only on the port the policy is applied.
With VLAN Based QoS, the QoS policy is applied to the VLAN interface. Traffic through all associated Switch ports is managed by that policy.

NOTE: Connected Stadium designs using modern Catalyst platforms would use per port configuration in most cases.

Port-Based QoS Policies

Service-policy commands are used to apply policy-maps to ports in both an input and output basis.

```
interface GigabitEthernet1/0/2
description Customer Port
service-policy input CISCO-PHONE
service-policy output 2P6Q3T-OUT
```

VLAN-Based QoS Policies

Service-policy commands are used to apply policy-maps to VLANs via its associated Switch Virtual Interface (SVI) in both an input and output basis.

```
interface Vlan 10
description Customer Port
service-policy input CISCO-PHONE
service-policy output 2P6Q3T-OUT
```

Ingress Queuing

Once traffic is classified, marked and policed on ingress to the network, next queuing is used to maintain the traffic’s loss, latency and jitter requirements as it egresses each interface along its path. Bandwidth is allotted to each queue and traffic is steered and handled according to its priority relative to other traffic sharing the links as it makes its way to its destination.

NOTE: Catalyst switches primarily use Egress queuing due to a deep set of egress buffers and Nexus switches use both Ingress and Egress queuing in combination. This is commonly referred to as Virtual Output Queuing (VOQ). Central arbitration is used to handle traffic on ingress, as well as, on egress.

Queuing Recommendations

- Enable queuing policies at every node that has the potential for congestion
- Whenever possible, assign each application class to its own dedicated queue

NOTE: Although the different classes of traffic should be handled the same throughout the network, the number of queues at each hop don't necessarily have to be the same. That is, Core switches may only have 4 queues while Access switches use 8. Traffic with similar transport requirements may be consolidated into a lower number of queues. For example, EF traffic in the priority queue, AF, DF, and Scavenger each in their own queue.

- EF Queue Recommendations
  - Limit the amount of strict priority queuing to 33% of link bandwidth capacity
  - Govern strict-priority traffic with an admission control mechanism
  - Do not enable WRED on the priority queue (EF Traffic)
- AF Queue Recommendations
  - Provision guaranteed bandwidth allocations according to application requirements
  - Enable DSCP-based WRED on this queue(s)
- DF Queue Recommendations
  - Provision at least 25 percent of link bandwidth for the default Best Effort class
  - Enable WRED (effectively RED) on the default class
- Scavenger Queue Recommendations
  - Assign minimum bandwidth to the Scavenger-class queue
  - WRED is not required on the Scavenger-class queue
- Use only platforms and/or service providers that offer a minimum of four standards-based queuing behaviors:
  - An RFC 3246 Expedited Forwarding Per-Hop Behavior
  - An RFC 2597 Assured Forwarding Per-Hop Behavior
  - An RFC 2474 Default Forwarding Per-Hop Behavior
  - An RFC 3662 Lower Effort Per-Domain Behavior

- Enable DSCP-based WRED on AF queues
- Enable WRED on the DF queue
- Do not enable WRED on control traffic application class queues
- WRED is not required on the Scavenger queue
- Optional: Tune WRED thresholds consistently—for example:

  Set the minimum WRED thresholds for AFx3 to 60% of the queue depth

  Set the minimum WRED thresholds for AFx2 to 70% of the queue depth

  Set the minimum WRED thresholds for AFx1 to 80% of the queue depth

  Set all maximum WRED thresholds to 100%

Figure 1 shows the traffic's assigned queues and their bandwidth allocations.

NOTE: Queue bandwidth allocations should be configured to accommodate the application needs of the venue.

Figure 22. Overlaying Connected Stadium Applications on Cisco's QoS Template

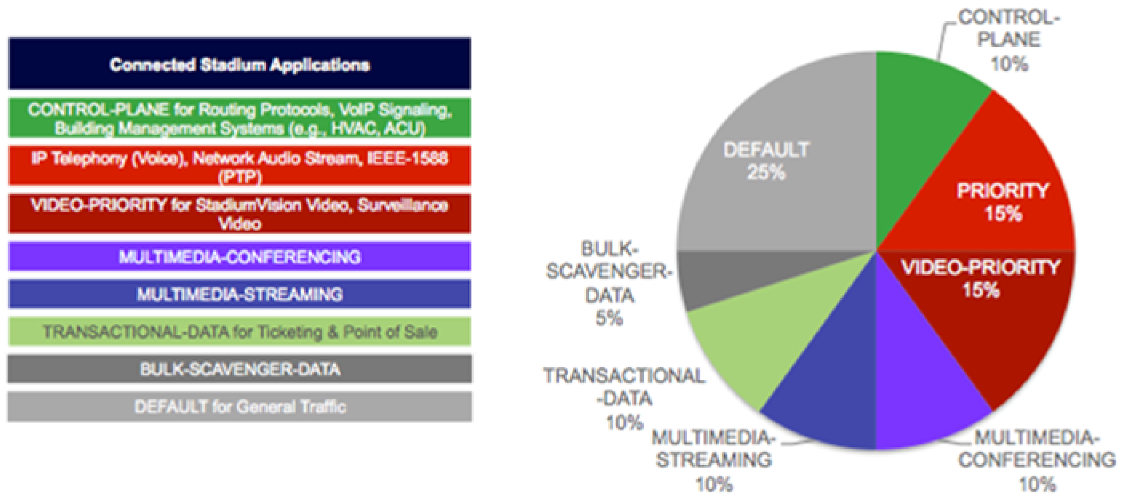

Egress QoS Models

Egress QoS models primarily deal with queuing and dropping policies (although additional egress QoS features—such as egress policing—are supported on some platforms).

NOTE: Catalyst switches primarily use egress queuing due to a deep set of egress buffers and Nexus switches use both Ingress and Egress queuing in combination. This is commonly referred to as Virtual Output Queuing (VOQ). Central arbitration is used to handle traffic on ingress, as well as, on egress.

Egress Queuing—Examples

```
! Class-maps to classify traffic to be queued before it exits an interface
!
class-map match-any PRIORITY
match dscp ef
class-map match-any VIDEO-PRIORITY
match dscp cs5
match dscp cs4
class-map match-any CONTROL-PLANE
match dscp cs7
match dscp cs6
match dscp cs3
match dscp cs2
class-map match-any MULTIMEDIA-CONFERENCING
match dscp af41
match dscp af42
match dscp af43
class-map match-any MULTIMEDIA-STREAMING
match dscp af31
match dscp af32
match dscp af33
class-map match-any TRANSACTIONAL-DATA
match dscp af21
match dscp af22
match dscp af23
class-map match-any BULK-SCAVENGER-DATA
match dscp af11
match dscp af12
match dscp af13
match dscp cs1
!
! Policy-map to assign bandwidth allocation and drop thresholds for
traffic exiting an interface. A Service-policy output command will
apply the policy-map to the interface.
!
policy-map 2P6Q3T-OUT
class PRIORITY
priority level 1
police rate percent 15
class VIDEO-PRIORITY
priority level 2
police rate percent 15
class CONTROL-PLANE
bandwidth remaining percent 10
queue-buffers ratio 10
class MULTIMEDIA-CONFERENCING
bandwidth remaining percent 10
queue-buffers ratio 10
queue-limit dscp af43 percent 80
queue-limit dscp af42 percent 90
queue-limit dscp af41 percent 100
class MULTIMEDIA-STREAMING
bandwidth remaining percent 10
queue-buffers ratio 10
queue-limit dscp af33 percent 80
queue-limit dscp af32 percent 90
queue-limit dscp af31 percent 100
class TRANSACTIONAL-DATA
bandwidth remaining percent 10
queue-buffers ratio 10
queue-limit dscp af23 percent 80
queue-limit dscp af22 percent 90
queue-limit dscp af21 percent 100
```

```
class BULK-SCAVENGER-DATA
bandwidth remaining percent 5
  queue-buffers ratio 10
  queue-limit dscp values af13 cs1 percent 80
queue-limit dscp values af12 percent 90
queue-limit dscp values af11 percent 100
class class-default
bandwidth remaining percent 25
queue-buffers ratio 25
```

Applying QoS Policies

The service-policy command is used to apply policies (i.e., policy-maps) to interfaces both in the input and output direction of traffic flow. The service-policy command can be configured on a port, VLAN or port-channel in the case of an etherchannel.

NOTE: The words input and output are used interchangeably with ingress and egress.

Connection Types

Access to the Cisco Connected Stadium network can be broken down into two basic types, Network Edge and Interswitch.

Network Edge

The Network Edge port class refers to interfaces that connect to endpoint devices and are where trust boundaries are established.

Inter-switch

The Inter-switch class refers to interfaces between switches along the path between endpoint devices. Trust DSCP is typical on inter-switch links.

Input Policies

Input policies are used on the Network Edge to classify, mark, and police traffic as it enters the network and sometimes on the inter-switch ports to queue traffic.

Figure 23. Platform Trust Assignments

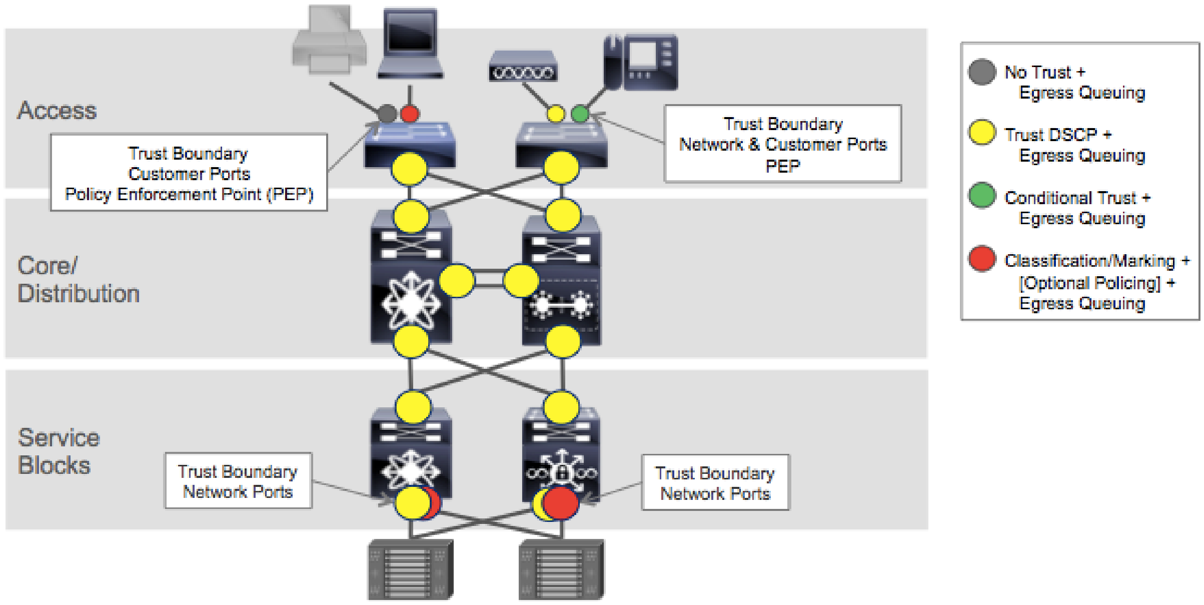

Output Policies

Output policies are used on both the network edge and inter-switch ports to insure traffic is delivered to the next hop and end device in the order of priority required by the specific application.

Figure 24. Platform Egress Queuing Assignments

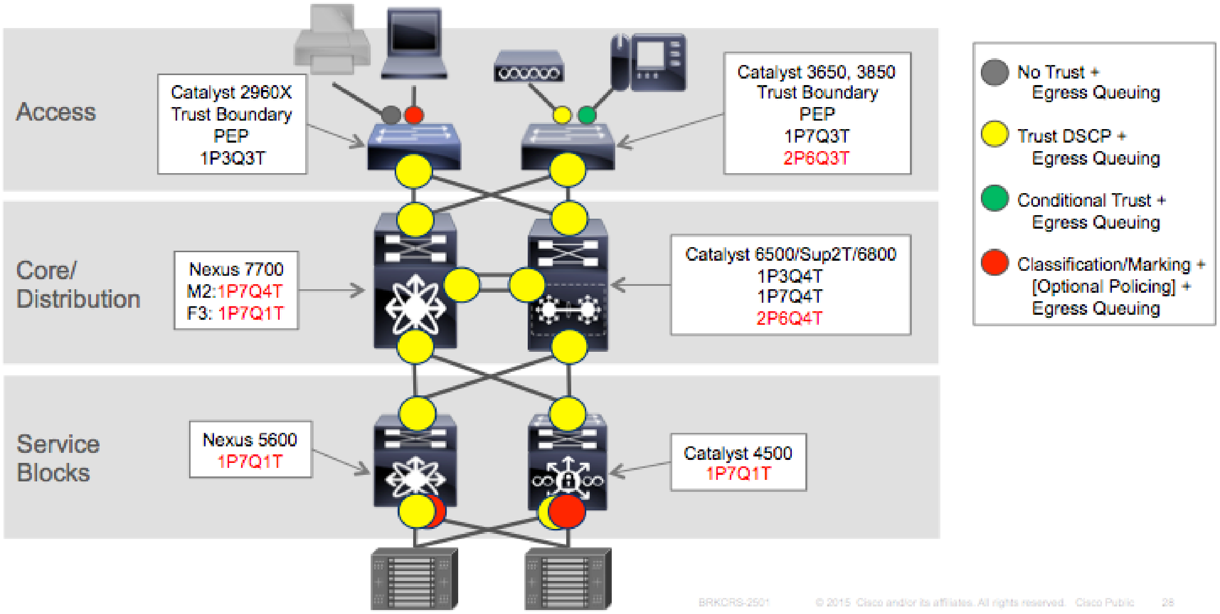

NOTE: 1P7Q4T means:

- 1 Priority Queue +
- 7 Queues with 4 Thresholds

Etherchannel-Based QoS Policies

Service-policy commands are used to apply policy-maps to virtual port-channel interfaces on the input and on the physcal port-member interfaces on the output.

```
port-channel load-balance src-dst-ip
!
interface port-channel1
description Uplink Port-Channel
service-policy input INGRESS-POLICY
!
interface GigabitEthernet range 1/0/1 - 2
description Uplink Port-Channel1 port-member
switchport mode trunk
channel-group 1
service-policy output 2P6Q3T-OUT
```

NOTE: The Cisco Catalyst 2960-X series is the exception to the rule where both the ingress and egress service-policies are on the physical port-member interfaces.

Configuration Example: Network Edge

Below are the port configuration snippets used for the different application endpoints.

NOTE: Only MQC examples are shown for the 3850. See Cisco Validated Design Guide Campus Wired LAN - Configuration Files Guide for other switch configuration examples.

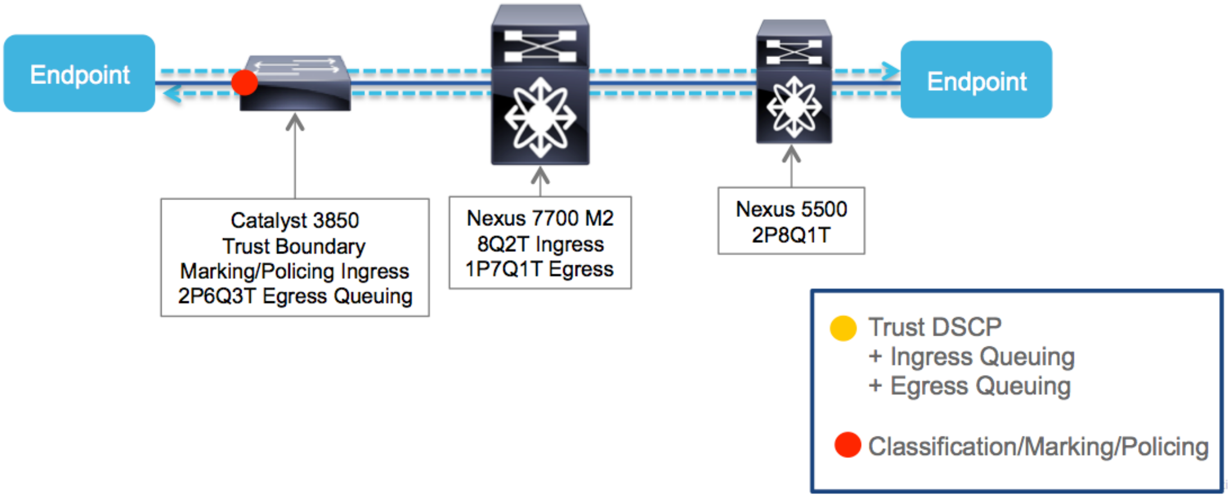

Cisco IP Phone

```
class-map match-any CISCO-PHONE-VOICE
match cos 5
class-map match-any CISCO-PHONE-SIGNALING
match cos 3
!
table-map policed-dscp
map from 0 to 8
map from 10 to 8
map from 18 to 8
map from 24 to 8
map from 46 to 8
!
policy-map CISCO-PHONE
class CISCO-PHONE-VOICE
set dscp ef
  police cir 128000 bc 8000
  conform-action transmit
  exceed-action set-dscp-transmit dscp table policed-dscp
class CISCO-PHONE-SIGNALING
set dscp cs3
  police cir 32000 bc 8000
  conform-action transmit
  exceed-action set-dscp-transmit dscp table policed-dscp
class class-default
  set dscp default
!
interface range GigabitEthernet1/0/x - y
description Access Port
switchport access vlan 102
switchport mode access
switchport voice vlan 103
trust device cisco-phone
service-policy input CISCO-PHONE
```

Cisco SV-4K Digital Media Player (DMP)

```
ip access-list extended IEEE-1588
remark PTP for DMPs and Audio Systems
permit udp any host 224.0.1.129
permit udp any host 224.0.1.130
permit udp any host 224.0.1.131
permit udp any host 224.0.1.132
!
ip access-list extended DMP-AS-A-VIDEO-SOURCE
remark ACL used to mark a SV-4K sourced video stream
```

```
permit udp <DMP Subnet/xx> host <239.193.20.x>
!
table-map policed-dscp
map from 0 to 8
map from 10 to 8
map from 18 to 8
map from 24 to 8
map from 46 to 8
!
class-map match-any IEEE-1588
match access-group name IEEE-1588
class-map match-any DMP-AS-A-VIDEO-SOURCE
match access-group name DMP-AS-A-VIDEO-SOURCE
!
policy-map CISCO-SV-DMP
class IEEE-1588
set dscp ef
  police cir 1000000 bc 300000
  conform-action transmit
  exceed-action set-dscp-transmit dscp table policed-dscp
class DMP-AS-A-VIDEO-SOURCE
  set dscp cs5
  police cir 25000000 bc 5000000
  conform-action transmit
  exceed-action set-dscp-transmit dscp table policed-dscp
class class-default
  set dscp default
!
interface range GigabitEthernet1/0/x - y
description SV-DMP Port
switchport mode access
switchport access vlan 110
service-policy input CISCO-SV-DMP
```

NOTE: The SV-4K and DMP-2Ks use IEEE-1588 Precision Timing Protocol (PTP) for synchronizing content display among multiple screens. The SV-4K can also be a multicast video source.

Cisco Digital Content Manager (DCM)

```
ip access-list extended DCM-DIRECTTV
permit ip any 239.192.0.0 3.255.255.255
```

or

```
ip access-list extended DCM-DIRECTTV
   permit ip host <source DCM> any
!
class-map match-any DCM-DIRECTTV
match access-group name DCM-DIRECTTV
!
policy-map POLICE-DCM
class DCM-DIRECTTV
  set dscp cs5
police 4000000000 1000000 exceed-action drop
class class-default
set dscp default
!
interface GigabitEthernetx/y
description DCM output, Primary Source 10.2.1.2
ip address 10.2.1.1 255.255.255.252
service-policy input POLICE-DCM
```

Point of Sale

```
ip access-list extended POS
permit udp <SOURCE POS SUBNET> <DESTINATION POS SUBNET> [range of Port#s]
permit tcp <SOURCE POS SUBNET> <DESTINATION POS SERVERS> [range of Port#s]
permit ip <SOURCE POS SUBNET> <DESTINATION POS SERVERS>
!
class-map match-any POS
  match access-group name POS
!
```

```
policy-map POS
class POS
  set dscp af21
class class-default
  set dscp default
!
interface range GigabitEthernet1/0/x - y
description POS Port
switchport mode access
switchport access vlan 120
service-policy input POS
```

Ticketing

```
ip access-list extended TICKET-MASTER
permit tcp <SOURCE TICKETING SUBNET> <DESTINATION TICKETING SUBNET>
[eq | range] <Port#s>
permit udp <SOURCE TICKETING SUBNET> <DESTINATION TICKETING SUBNET>
[eq | range] <Port#s>
permit ip <SOURCE TICKETING SUBNET> <DESTINATION TICKETING SUBNET>
!
class-map match-any TICKET-MASTER
match access-group name TICKET-MASTER
!
policy-map TICKET-MASTER
class TICKET-MASTER
  set dscp af21
class class-default
set dscp default
!
interface range GigabitEthernet1/0/x - y
description TICKET-MASTER Port
switchport mode access
switchport access vlan 120
service-policy input TICKET-MASTER
```

Network Audio

```
ip access-list extended IEEE-1588
remark PTP for the Audio System time synchronization
permit udp any host 224.0.1.129
permit udp any host 224.0.1.130
permit udp any host 224.0.1.131
permit udp any host 224.0.1.132
!
```

NOTE: Network Audio Streams require Priority Queue treatment for to maintain high-quality.

```
!
ip access-list extended NETWORK-AUDIO-STREAM
remark ACL used to mark NETWORK AUDIO streams
permit udp <SOURCE AUDIO SUBNET> <DESTINATION DESTINATION SUBNET>
[eq | range] <Port#s>
!
table-map policed-dscp
map from 0 to 8
map from 10 to 8
map from 18 to 8
map from 24 to 8
map from 46 to 8
!
class-map match-any IEEE-1588
match access-group name IEEE-1588
class-map match-any NETWORK-AUDIO-STREAM
match access-group name NETWORK-AUDIO-STREAM
!
policy-map NETWORK-AUDIO
class IEEE-1588
set dscp ef
  police cir 1000000 bc 300000
  conform-action transmit
  exceed-action set-dscp-transmit dscp table policed-dscp
class NETWORK-AUDIO-STREAM
  set dscp EF
class class-default
   set dscp default
!
interface range GigabitEthernet1/0/x - y
description NETWORK-AUDIO Port
switchport mode access
switchport access vlan 110
service-policy input NETWORK-AUDIO
```

Cisco Video Surveillance Camera

Below is an example of how the QoS can be configured. If using Cisco Video Surveillance follow the CVD at http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Video/IPVS/IPVS_DG/IPVS-DesignGuide.pdf

```
ip access-list extended IPVS-CAMERA
permit ip <IPVS CAMERA SUBNET> <MASK> <IPVS SERVER SUBNET> <MASK>
!
class-map match-any IPVS-CAMERA
match access-group name IPVS-CAMERA
!
policy-map MARK-IPVS
  class IPVS-CAMERA
    set dscp cs5
    police 4000000 125000 exceed-action drop
  class class-default
   set dscp default
!
interface GigabitEthernetx/0/y
description IPVS-CAMERA
switchport
```

```
switchport access vlan <VLAN #>
service-policy input MARK-IPVS
```

Video Surveillance Server

```
ip access-list extended IPVS-SERVER
 permit ip <IPVS SERVER SUBNET> <MASK> <IPVS MULTICAST GROUP> <MASK>
or
permit ip <IPVS SERVER SUBNET> <MASK> <Monitoring Dest IP/Subnet>
<MASK>
!
class-map match-any IPVS-SERVER
  match access-group name IPVS-SERVER
!
policy-map IPVS-SERVER
  class IPVS-SERVER
set dscp cs5
    police 4000000 125000 exceed-action drop
  class class-default
   set dscp cs3
!
interface GigabitEthernetx/y
description IPVS-SERVER
switchport
switchport access vlan <VLAN #>
service-policy input IPVS-SERVER
```

Building Management Systems Servers (ACUs, HVAC units)

```
ip access-list extended BMS
permit udp <ACU SERVER SUBNET> any [eq | range Port#s]
permit tcp <ACU SERVER SUBNET> any [eq | range Port#s]
permit ip <ACU SERVER SUBNET> [any | <ACU SUBNET>
permit udp <HVAC SUBNET> any [eq | range Port#s]
permit tcp <HVAC SUBNET> any [eq | range Port#s]
permit ip <HVAC SUBNET> [any | <HVAC SUBNET>
```

```
class-map match-any BMS-SERVERS
match access-group name BMS

policy-map MARK-BMS
class BMS-SERVERS
  set dscp cs2
class class-default
  set dscp default
!
!
interface GigabitEthernetx/y
description IPVS-SERVER
switchport
switchport access vlan <VLAN #>
service-policy input [MARK-BMS]
```

Configuration Example: Inter-Switch Links

3850 Egress to N7K M2 Ingress Queuing

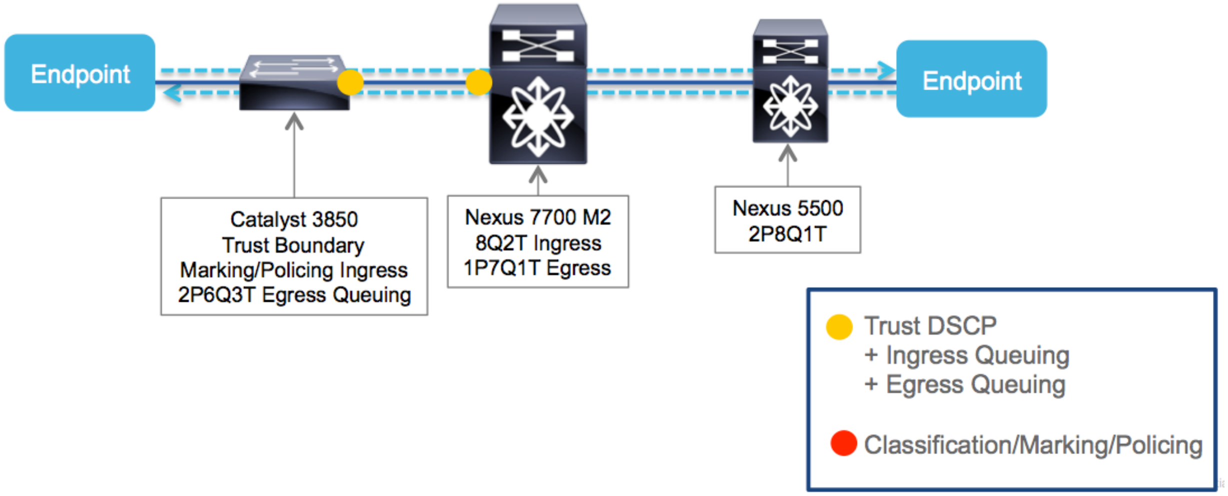

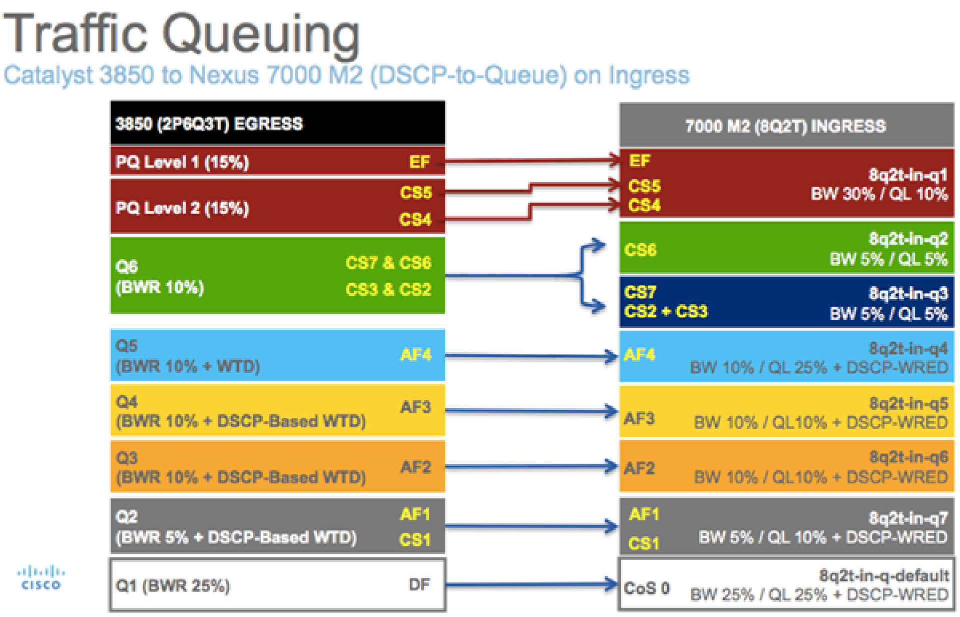

Traffic Queuing

Catalyst 3850 to Nexus 7000 M2 (DSCP-to-Queue) on Ingress

3850 (2P6Q3T) EGRESS		7000 M2 (8Q2T) INGRESS	
PQ Level 1 (15%)	EF	EF CS5 CS4	8q2t-in-q1 BW 30% / QL 10%
PQ Level 2 (15%)	CS5 CS4		
Q6 (BWR 10%)	CS7 & CS6 CS3 & CS2	CS6	8q2t-in-q2 BW 5% / QL 5%
		CS7 CS2 + CS3	8q2t-in-q3 BW 5% / QL 5%
Q5 (BWR 10% + WTD)	AF4	AF4	8q2t-in-q4 BW 10% / QL 25% + DSCP-WRED
Q4 (BWR 10% + DSCP-Based WTD)	AF3	AF3	8q2t-in-q5 BW 10% / QL10% + DSCP-WRED
Q3 (BWR 10% + DSCP-Based WTD)	AF2	AF2	8q2t-in-q6 BW 10% / QL10% + DSCP-WRED
Q2 (BWR 5% + DSCP-Based WTD)	AF1 CS1	AF1 CS1	8q2t-in-q7 BW 5% / QL 10% + DSCP-WRED
Q1 (BWR 25%)	DF	CoS 0	8q2t-in-q-default BW 25% / QL 25% + DSCP-WRED

Class Maps — 8 Queues	
Catalyst 3850 (2P6Q3T) EGRESS	Nexus 7000 M2 (8Q2T) INGRESS
class-map match-any VOICE-PQ1 match dscp ef class-map match-any VIDEO-PQ2 match dscp cs4 match dscp cs5 class-map match-any CONTROL-PLANE match dscp cs2 match dscp cs3 match dscp cs6 match dscp cs7 class-map match-any MULTIMEDIA-CONFERENCING match dscp af41 match dscp af42 match dscp af43 class-map match-any MULTIMEDIA-STREAMING match dscp af31 match dscp af32 match dscp af33 class-map match-any TRANSACTIONAL-DATA match dscp af21 match dscp af22 match dscp af23 class-map match-any BULK-SCAVENGER match dscp af11 match dscp af12 match dscp af13 match dscp cs1	! DSCP classification available from NX-OS 6.2.2 on hardware qos dscp-to ingress module-type all ! PRIORITY class-map type queuing match-any 8q2t-in-q1 no match cos 0-7 no match dscp 0-63 match dscp 32, 40, 46 match cos 5 ! CONTROL-PLANE class-map type queuing match-any 8q2t-in-q2 no match cos 0-7 no match dscp 0-63 match dscp 48 match cos 6 ! CONTROL-PLANE class-map type queuing match-any 8q2t-in-q3 no match cos 0-7 no match dscp 0-63 match dscp 16, 24, 56 match cos 7 ! MULTIMEDIA-CONFERENCING class-map type queuing match-any 8q2t-in-q4 no match cos 0-7 no match dscp 0-63 match dscp 34, 36, 38 match cos 4 ! MULTIMEDIA-STREAMING class-map type queuing match-any 8q2t-in-q5

Class Maps — 8 Queues	
Catalyst 3850 (2P6Q3T) **EGRESS**	**Nexus 7000 M2 (8Q2T) INGRESS**
! CLASS-DEFAULT automatically added	no match cos 0-7 no match dscp 0-63 match dscp 26, 28, 30 match cos 3 ! TRANSACTIONAL-DATA class-map type queuing match-any 8q2t-in-q6 no match cos 0-7 no match dscp 0-63 match dscp 18, 20, 22 match cos 2 ! BULK-SCAVENGER-DATA no match cos 0-7 no match dscp 0-63 match dscp 8, 10, 12, 14 match cos 1 ! CLASS-DEFAULT automatically added

Policy Maps — 8 Queues	
Catalyst 3850 (2P6Q3T) EGRESS	**Nexus 7000 M2 (8Q2T) INGRESS**
policy-map 2P6Q3T-OUT class PRIORITY priority level 1 police rate percent 15 queue-buffers ratio 5 class VIDEO-PRIORITY priority level 2 police rate percent 15 queue-buffers ratio 5 class CONTROL-PLANE bandwidth remaining percent 10 queue-buffers ratio 5 class MULTIMEDIA-CONFERENCING bandwidth remaining percent 10 queue-buffers ratio 10 queue-limit dscp af43 percent 80 queue-limit dscp af42 percent 90 queue-limit dscp af41 percent 100 class MULTIMEDIA-STREAMING bandwidth remaining percent 10 queue-buffers ratio 10 queue-limit dscp af33 percent 80 queue-limit dscp af32 percent 90 queue-limit dscp af31 percent 100 class TRANSACTIONAL-DATA bandwidth remaining percent 10 queue-buffers ratio 10 queue-limit dscp af23 percent 80	policy-map type queuing 8Q2T-IN ! PRIORITY + VIDEO PRIORITY (COS 5) class type queuing 8q2t-in-q1 bandwidth percent 30 queue-limit percent 10 ! CONTRL-PLANE (COS 7) class type queuing 8q2t-in-q2 bandwidth percent 5 queue-limit percent 5 ! CONTROL-PLANE (COS 6) class type queuing 8q2t-in-q3 bandwidth percent 5 queue-limit percent 5 ! MULTIMEDIA-CONFERENCING (COS 4) class type queuing 8q2t-in-q4 bandwidth percent 10 queue-limit percent 25 random-detect dscp-based random-detect dscp 34 minimum-threshold percent 80 maximum-threshold percent 100 random-detect dscp 36 minimum-threshold percent 80 maximum-threshold percent 100 random-detect dscp 38 minimum-threshold percent 80 maximum-threshold percent 100 ! MULTIMEDIA-STREAMING (COS 3) class type queuing 8q2t-in-q5 bandwidth percent 10 queue-limit percent 10 random-detect dscp-based random-detect dscp 26 minimum-threshold

Policy Maps — 8 Queues	
Catalyst 3850 (2P6Q3T) EGRESS	**Nexus 7000 M2 (8Q2T) INGRESS**
queue-limit dscp af22 percent 90 queue-limit dscp af21 percent 100 class BULK-SCAVENGER-DATA bandwidth remaining percent 5 queue-buffers ratio 20 queue-limit dscp values af13 cs1 percent 80 queue-limit dscp values af12 percent 90 queue-limit dscp values af11 percent 100 class class-default bandwidth remaining percent 25 queue-buffers ratio 25 ! interface TenGigabitEthernet3/1/1 description Link to Core Nexus 7k_M1 service-policy output 2P6Q3T-OUT	percent 80 maximum-threshold percent 100 random-detect dscp 28 minimum-threshold percent 80 maximum-threshold percent 100 random-detect dscp 30 minimum-threshold percent 80 maximum-threshold percent 100 ! TRANSACTIONAL-DATA (COS 2) class type queuing 8q2t-in-q6 bandwidth percent 10 queue-limit percent 10 random-detect dscp-based random-detect dscp 18 minimum-threshold percent 80 maximum-threshold percent 100 random-detect dscp 20 minimum-threshold percent 80 maximum-threshold percent 100 random-detect dscp 22 minimum-threshold percent 80 maximum-threshold percent 100 ! BULK-SCAVENGER-DATA (COS 1) class type queuing 8q2t-in-q7 bandwidth percent 5 queue-limit percent 10 random-detect dscp-based random-detect dscp 8 minimum-threshold percent 80 maximum-threshold percent 100 random-detect dscp 10 minimum-threshold percent 80 maximum-threshold percent 100 random-detect dscp 12 minimum-threshold percent 80 maximum-threshold percent 100 random-detect dscp 14 minimum-threshold

Policy Maps — 8 Queues	
Catalyst 3850 (2P6Q3T) EGRESS	**Nexus 7000 M2 (8Q2T) INGRESS**
	percent 80 maximum-threshold percent 100 ! CLASS-DEFAULT (COS 0) class type queuing 8q2t-in-q-default bandwidth percent 25 queue-limit percent 25 random-detect dscp-based random-detect dscp 0 minimum-threshold percent 80 maximum-threshold percent 100 ! interface Ethernet 4/1-24 description Link to 3850 Access service-policy type queuing input 8Q2T-IN

N7K M2 Ingress to Egress Queuing

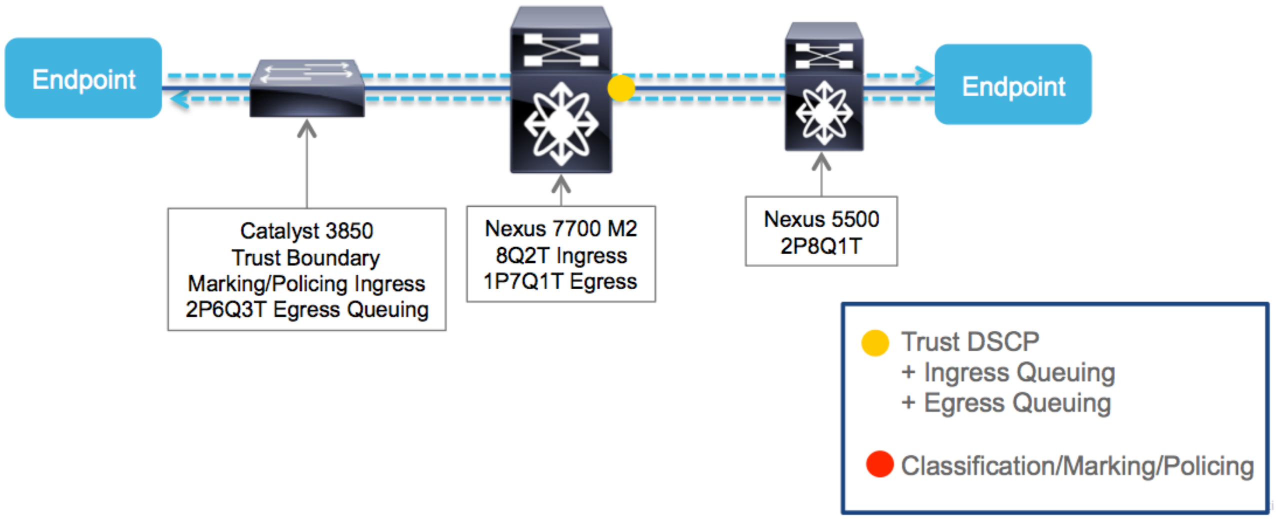

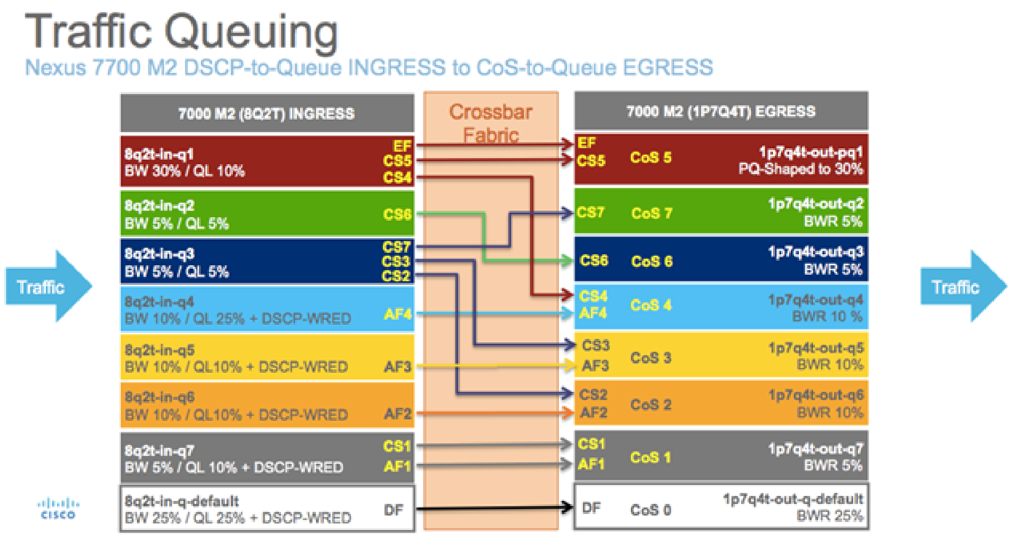

Traffic Queuing

Nexus 7700 M2 DSCP-to-Queue INGRESS to CoS-to-Queue EGRESS

NOTE: When using CoS-to-Queue, the three Most Significant Bits of the DSCP value determines the CoS queue used on egress.

Class Maps — 8 Queues
Nexus 7000 M2 (1P7Q4T) EGRESS (CoS-to-Queue)
class-map type queuing match-any 1p7q4t-out-pq1
match cos 5
class-map type queuing match-any 1p7q4t-out-q2
match cos 7
class-map type queuing match-any 1p7q4t-out-q3
match cos 6
class-map type queuing match-any 1p7q4t-out-q4
match cos 4
class-map type queuing match-any 1p7q4t-out-q5
match cos 3
class-map type queuing match-any 1p7q4t-out-q6
match cos 2
class-map type queuing match-any 1p7q4t-out-q7
match cos 1

Policy Maps — 8 Queues
Nexus 7000 M2 (1P7Q4T) EGRESS (CoS-to-Queue)

```
 policy-map type queuing 1P7Q4T-OUT
  class type queuing 1p7q4t-out-pq1
   priority
   shape average percent 30
   queue-limit percent 10
  class type queuing 1p7q4t-out-q2
   bandwidth remaining percent 5
   queue-limit percent 5
  class type queuing 1p7q4t-out-q3
   bandwidth remaining percent 5
   queue-limit percent 5
  class type queuing 1p7q4t-out-q4
 bandwidth remaining percent 10
   queue-limit percent 25
   random-detect cos-based
   random-detect cos 4 minimum-threshold percent 80 maximum-threshold percent 100
  class type queuing 1p7q4t-out-q5
   bandwidth remaining percent 10
   queue-limit percent 10
   random-detect cos-based
   random-detect cos 3 minimum-threshold percent 80 maximum-threshold percent 100
  class type queuing 1p7q4t-out-q6
   bandwidth remaining percent 10
   queue-limit percent 10
   random-detect cos-based
   random-detect cos 2 minimum-threshold percent 80 maximum-threshold percent
100
  class type queuing 1p7q4t-out-q7
   bandwidth remaining percent 5
```

Policy Maps — 8 Queues
Nexus 7000 M2 (1P7Q4T) EGRESS (CoS-to-Queue)

```
    queue-limit percent 10
    random-detect cos-based
    random-detect cos 1 minimum-threshold percent 80 maximum-threshold percent 100
  class type queuing 1p7q4t-out-q-default
   bandwidth remaining percent 25
    queue-limit percent 25
    random-detect cos-based
    random-detect cos 0 minimum-threshold percent 80 maximum-threshold percent 100
!
interface Ethernet 4/1-24
service-policy type queuing output 1P7Q4T-OUT
```

N7K M2 Egress to N5500 Ingress Queuing

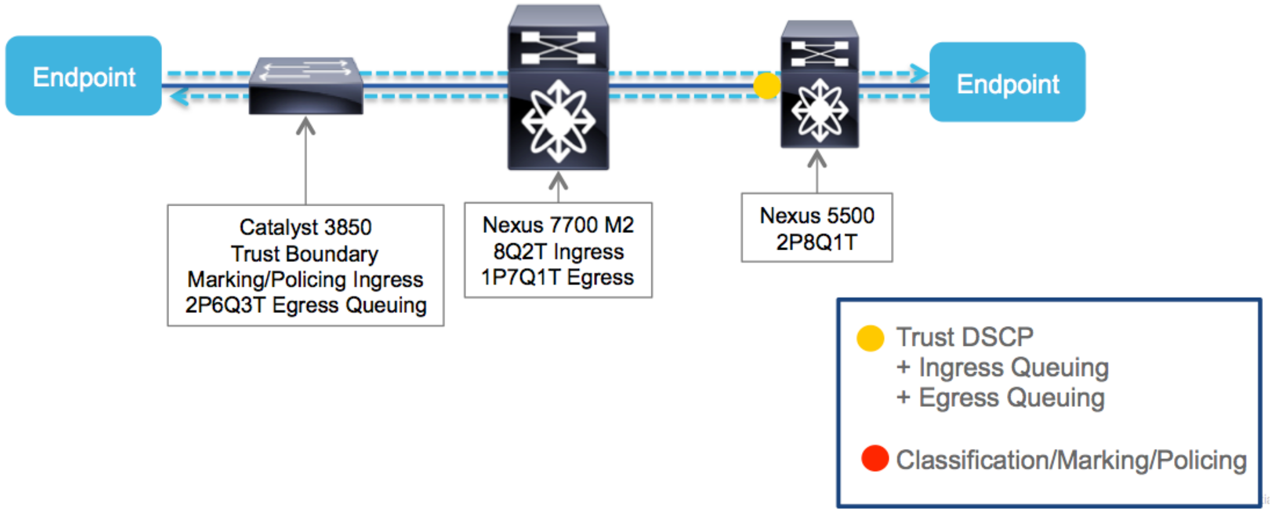

Traffic Queuing

Nexus 7000 M2 CoS-to-Queue EGRESS to Nexus 5500 INGRESS

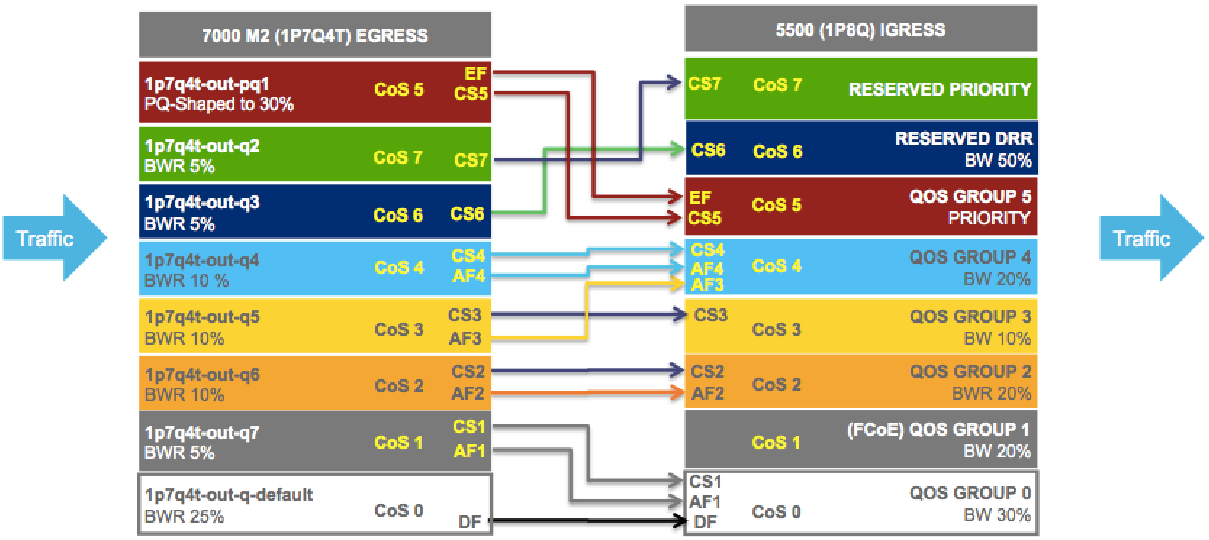

7000 M2 (1P7Q4T) EGRESS		
1p7q4t-out-pq1 PQ-Shaped to 30%	CoS 5	EF CS5
1p7q4t-out-q2 BWR 5%	CoS 7	CS7
1p7q4t-out-q3 BWR 5%	CoS 6	CS6
1p7q4t-out-q4 BWR 10 %	CoS 4	CS4 AF4
1p7q4t-out-q5 BWR 10%	CoS 3	CS3 AF3
1p7q4t-out-q6 BWR 10%	CoS 2	CS2 AF2
1p7q4t-out-q7 BWR 5%	CoS 1	CS1 AF1
1p7q4t-out-q-default BWR 25%	CoS 0	DF

5500 (1P8Q) IGRESS		
CS7	CoS 7	RESERVED PRIORITY
CS6	CoS 6	RESERVED DRR BW 50%
EF CS5	CoS 5	QOS GROUP 5 PRIORITY
CS4 AF4 AF3	CoS 4	QOS GROUP 4 BW 20%
CS3	CoS 3	QOS GROUP 3 BW 10%
CS2 AF2	CoS 2	QOS GROUP 2 BWR 20%
	CoS 1	(FCoE) QOS GROUP 1 BW 20%
CS1 AF1 DF	CoS 0	QOS GROUP 0 BW 30%

NOTE: The Nexus 5500 only offers up to six customizable queues (including the FCoE and default classes). The configuration below adapts the 8-class QoS model to the available queues in the Nexus 5500 switch.

In NX-OS, CoS or DSCP values are mapped to QoS groups, not directly to queues.

In most IOS platforms, the policy map is attached to an interface. However, in NX-OS, it is also possible to attach policy maps to the system itself. In this way, system qos is a type of MQC target. From an architectural perspective, the system qos target is the crossbar fabric itself. You use a service-policy statement to associate a policy map with the system qos target. A system qos policy applies to all interfaces on the switch unless a specific interface has an overriding service policy configuration. The system qos policies are used to define system classes, the classes of traffic across the entire switch, and their attributes.

To ensure QoS consistency (and for ease of configuration), the switch distributes the system qos parameter values to all its attached converged network adapters (CNAs) using the Data Center Bridging Exchange (DCBX) protocol.

If service policies are configured at the interface level, the interface-level policy always takes precedence over system class configuration or defaults.

On the Cisco Nexus 5500 series switch, a system class is uniquely identified by a qos-group value. A total of six system classes are supported. Class 0 is reserved as the default class, but five additional system classes (Classes 1 to 5) can be created by the administrator and assigned to various traffic types. If Fibre Channel over Ethernet (FCoE) is in use, this is assigned as Class 1, leaving four configurable classes.

The configurations below for the Nexus 5500 support this system QoS design model.

The QoS configuration can be accomplished by following these four steps:

1. Configure the qos type class and policy maps.
2. Configure the network-qos type class and policy maps.
3. Configure the queuing type class and policy maps.
4. Apply all the service policies to the system QoS class.

QoS Class, Policy Maps, and Service Policies

Nexus 5500 (4Q1T) INGRESS

```
! Nexus 5500 has 4 user-configurable queues. This also assumes Fiber-Channel is configured.


! Configure the qos type class maps
class-map type qos match-any PRIORITY
  match dscp cs5,ef
  match cos 5
class-map type queuing match-any MULTIMEDIA-CONFERENCING
  match dscp cs4,af41,af42,af43
  match cos 4
class-map type qos match-any MULTIMEDIA-STREAMING
  match dscp af31,af32,af33
  match cos 3
class-map type qos match-any SIGNALING
match dscp cs3
class-map type qos match-any TRANSACTIONAL-DATA
match dscp cs2,af21,af22,af23
match cos 2
! CLASS-DEFAULT automatically added
!
! Configure the qos type policy maps
policy-map type qos 8-CLASS-QOS-POLICY
class PRIORITY
  set qos-group 5
class MULTIMEDIA-CONFERENCING
  set qos-group 4
class MULTIMEDIA-STREAMING
  set qos-group 4
class SIGNALING
```

QoS Class, Policy Maps, and Service Policies
Nexus 5500 (4Q1T) INGRESS

```
  set qos-group 3
class TRANSACTIONAL-DATA
  set qos-group 2
class class-fcoe
  set qos-group 1
!
! Configure the network-qos type class maps
class-map type network-qos PRIORITY
match qos-group 5
class-map type network-qos MULTIMEDIA
match qos-group 4
class-map type network-qos SIGNALING
match qos-group 3
class-map type network-qos TRANSACTIONAL-DATA
match qos-group 2
!
! Configure the network-qos type policy maps
policy-map type network-qos 8-CLASS-NQ-POLICY
class type network-qos PRIORITY
class type network-qos MULTIMEDIA
class type network-qos SIGNALING
class type network-qos TRANSACTIONAL-DATA
class type network-qos class-fcoe
!
! Configure the queuing type class maps
class-map type queuing PRIORITY
match qos-group 5
class-map type queuing MULTIMEDIA
```

QoS Class, Policy Maps, and Service Policies

Nexus 5500 (4Q1T) INGRESS

```
match qos-group 4
class-map type queuing SIGNALING
match qos-group 3
class-map type queuing TRANSACTIONAL-DATA
match qos-group 2
!
! Configure the queuing type policy maps
policy-map type queuing 8-CLASS-GLOBAL-QUEUING-POLICY
class type queuing PRIORITY
 priority
class type queuing MULTIMEDIA
 bandwidth percent 20
class type queuing SIGNALING
 bandwidth percent 10
class type queuing TRANSACTIONAL-DATA
 bandwidth percent 20
class type queuing class-fcoe
 bandwidth percent 20
class type queuing class-default
 bandwidth percent 30
!
! Apply all the service policies to the system QoS class
system qos
service-policy type qos input 8-CLASS-QOS-POLICY
service-policy type queuing output 8-CLASS-GLOBAL-QUEUING-POLICY
service-policy type network-qos 8-CLASS-NQ-POLICY
```

Nexus Classification, Marking and Policing

Figure 25. N5500 Network Edge

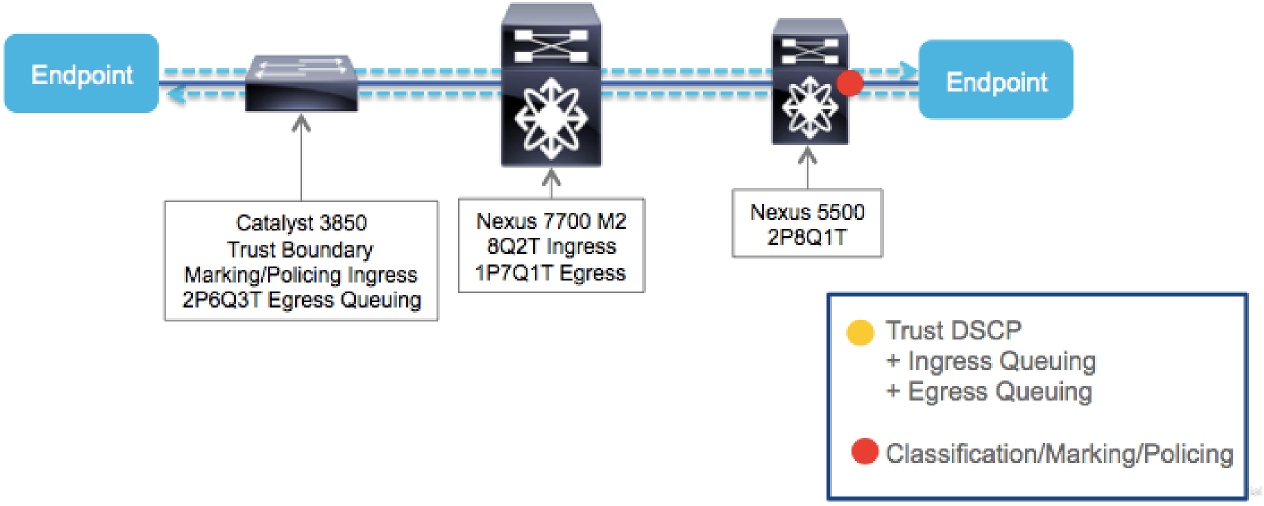

By default, all Ethernet interfaces on the Nexus platforms are trusted. In other words, the DSCP and CoS values are preserved unless marking policies are configured to specifically overwrite the QoS values. Therefore, if the servers are trusted, no explicit configuration is required.

When the server is untrusted, the objective is to reset the QoS markings to zero and to ensure that all traffic from such a server is sent to the default class. In NX-OS lingo, on the Nexus 5500, this means that all traffic from an untrusted server is mapped to QoS group 0 and the DSCP is set to 0 regardless of the original packet QoS markings.

```
N5K(config)# policy-map type qos NO-TRUST
N5K(config-pmap-qos)# class type qos class-default
! For matched traffic (all traffic), the DSCP is to 0
N5K(config-pmap-c-qos)# set dscp 0
!
! This command enables system qos (global) configuration mode
N5K(config)# system qos
! Applies the NO-TRUST policy globally
N5K(config-sys-qos)# service-policy type qos input NO-TRUST
! This section applies the NO-TRUST policy to an interface
```

```
N5K(config)# interface ethernet 1/1-15
N5K(config-if-range)# service-policy type qos input NO-TRUST
ERROR: policy already attached at system level
```

Network Audio (Core Server)

```
ip access-list IEEE-1588
remark PTP for the Audio System time synchronization
permit udp any host 224.0.1.129
permit udp any host 224.0.1.130
permit udp any host 224.0.1.131
permit udp any host 224.0.1.132
!
```

NOTE: Network Audio Streams require Priority Queue treatment for to maintain high-quality.

```
!
ip access-list extended NETWORK-AUDIO-STREAM
remark ACL used to mark NETWORK AUDIO streams
permit udp <SOURCE AUDIO SUBNET> <DESTINATION DESTINATION SUBNET>
[eq | range] <Port#s>
!
class-map type qos IEEE-1588
match access-group name IEEE-1588
class-map type qos NETWORK-AUDIO-STREAM
match access-group name NETWORK-AUDIO-STREAM
!
policy-map type qos NETWORK-AUDIO
class type qos IEEE-1588
  set qos-group 5
set dscp ef
class type qos NETWORK-AUDIO-STREAM
  set qos-group 5
  set dscp ef
!
interface range GigabitEthernet1/0/x - y
```

```
description NETWORK-AUDIO Port
switchport mode access
switchport access vlan 110
service-policy type qos input NETWORK-AUDIO
```

Now that the classification and marking portions of the configuration are finished, the next steps involve enabling the QoS group to which this traffic is added. If this step is missed, the queue for this class is not enabled and QoS does not work. It is necessary to first activate and configure the new QoS group before it can be used.

NOTE: The QoS group may have been previously configured as shown above in Section titled QoS Class and Policy Maps + Service Policies in this Nexus 5500 section.

```
N5K(config)# class-map type network-qos PRIORITY
N5K(config-cmap-nq)# match qos-group 5
! Match on all traffic mapped to QoS group 5
! Step 2 – configure the network-qos policy map
N5K(config)# policy-map type network-qos 8-CLASS-NQ-POLICY
N5K(config-pmap-nq)# class type network-qos PRIORITY
! Enable the class – essentially turning on the queue
N5K(config-pmap-nq)# set cos 5
! Manually set the cos value for this class
! Step 3 – attach the policy map to the system class
N5K(config-pmap-nq-c)# system qos
N5K(config-sys-qos)# service-policy type network-qos 8-CLASS-NQ-
POLICY
```

Security

Overview

Sports and Entertainment venues are high profile installations with extensive media coverage surrounding both the events and the stadium itself. Because of this publicity these venues make for attractive targets for malicious network breaches.

The Security model of the Connected Stadium is more similar to a Service Provider model than that of an Enterprise where the network is generally viewed as untrusted and perimeter security is implemented at the network and Data Center edges.

NOTE: This design guide shows multiple layers of Firewalls used throughout the design. These layers maybe consolidated and handled by a single pair (e.g., deployed in active/standby redundancy) of Firewalls. Careful consideration must be taken when sizing the Firewall and understanding the different traffic flows to be protected using this consolidated approach.

Figure 26. Connected Stadium Security Model

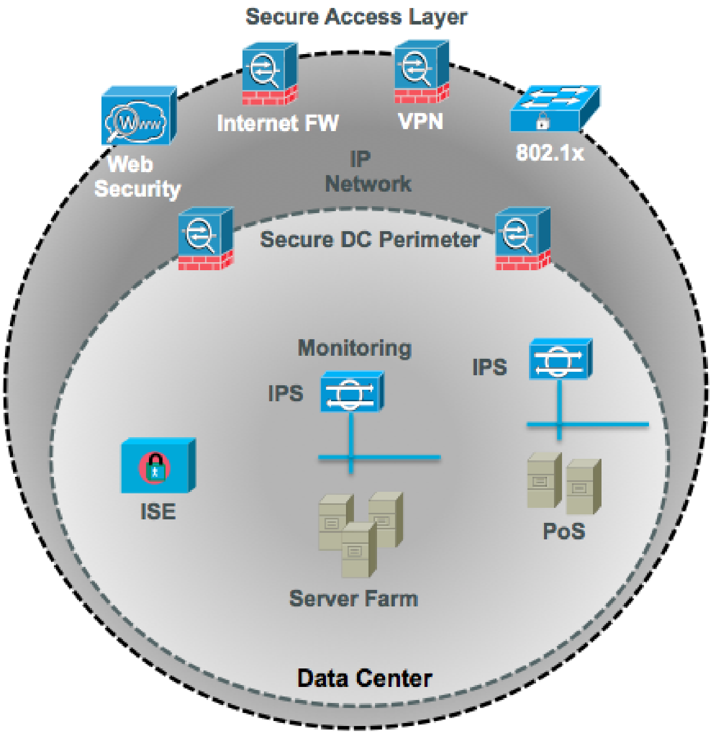

Internet Edge Design

This section describes the requirements, design, and configuration of the Internet firewall services in the Cisco Connected Stadium.

Figure 27. Internet Edge Design

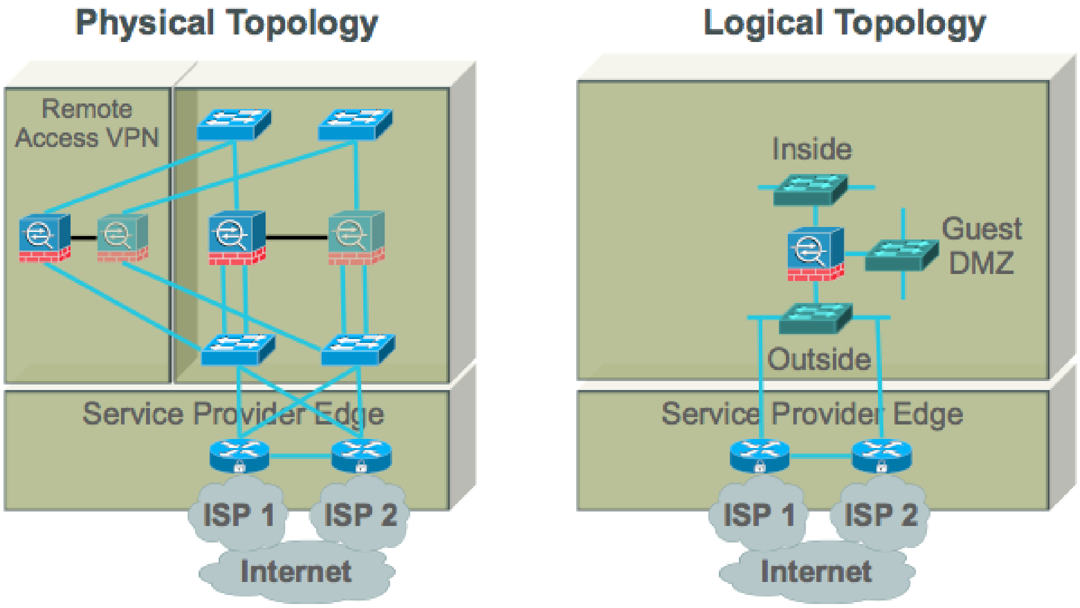

Firewall Design Attributes

The following internet firewall design attributes apply to the Cisco Connected Stadium:

- Internet access is through a connection provided to the stadium via a redundant pair of links provided by Internet Service Providers (ISPs). These links terminate in a pair of redundant routers in the Service Provider Edge of the network.
- The Internet ASAs are configured in a redundant active/standby, stateful failover design.
- The Internet ASAs provide Internet access to guests on the DMZ. The Internet ASAs deny these hosts access to the internal network.
- The Internet ASAs are managed via SSH/HTTPS - telnet is disabled.
- The Internet ASAs are configured to authenticate, authorize, and account all device access to the ISE server via RADIUS.
- The Internet ASAs are configured to use NTP.
- The Internet ASAs are configured to log events to a logging server.
- The Internet ASA outside interfaces are connected to external Catalyst 3750 switches.

- The Internet ASA inside interfaces are connected to the data center switches or has a dedicated pair of switches.
- The Internet ASA firewalls attach their DMZ interfaces to the VLAN that provides guest wired access via the data center switches or Core/Distribution switches.
- The Firewall divides the network into security zones that are logically separated from one another.

From a security perspective, the various functional areas on the network are separated into security zones. Communication between zones are limited to the specific traffic required for an application to function.

Ports and Security Level Assignment

Table 6 lists the applicable Security zones that are used for the interfaces of the ASAs (shown in Figure 28).

Table 6. Internet ASA Port/Security Level Assignment

Interface Name	Physical Interface	Security Level	Description
outside	GigabitEthernet0/0	0	Connection to the Internet routers
inside	GigabitEthernet0/1	100	Connection to the internal stadium network
GuestDMZ	GigabitEthernet0/2.X	50	Connection to the DMZ for Guest Access
failover	GigabitEthernet0/3	n/a	Connection between ASA pair for failover communications

Figure 28. Internet ASA Interfaces

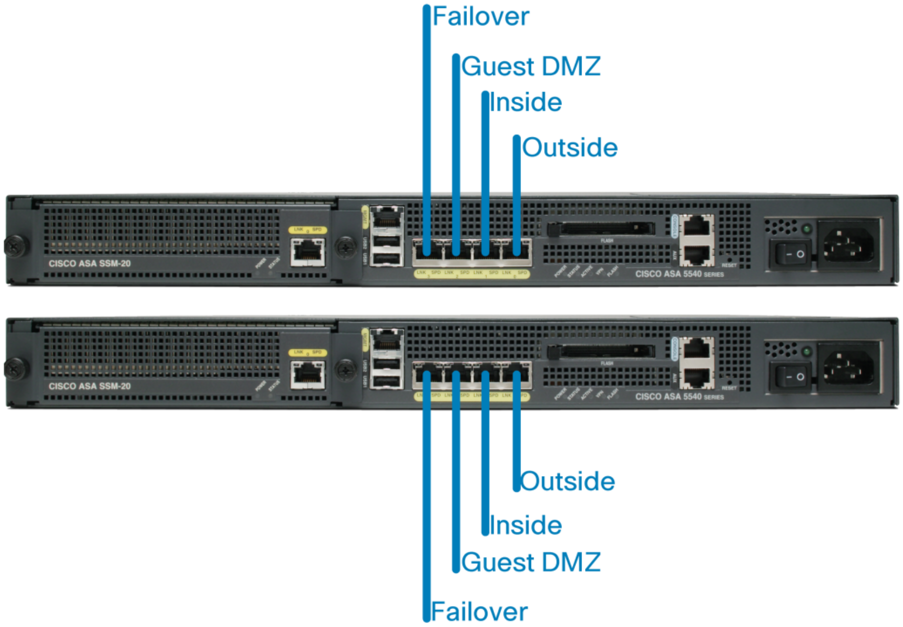

Example 2. Internet ASA Configuration Example

```
hostname SV-4K
domain-name
stadium.com
enable password
<removed>
passwd <removed>
names
!
interface
GigabitEthernet0/0
  nameif outside
  security-level 0
  ip address A.B.C.D
<mask> standby
A.B.C.E
!
interface
GigabitEthernet0/
  nameif inside
  security-level 100
  ip address A.B.C.D
<mask> standby
A.B.C.E
!
interface
GigabitEthernet0/2
  no nameif
  no security-level

  no ip address
!
interface gig-
abitethernet0/2.X
  description <DMZ
interface>
  vlan X
  nameif SV-4K
  security-level 50
  ip address A.B.C.D
<mask> standby
A.B.C.E
!
interface Gig-
abitEthernet0/3
 description
LAN/STATE Failover
Interface
!
interface Man-
```

```
!
ftp mode passive
dns server-group DefaultDNS
 domain-name stadium.com
<insert ACLs here>
pager lines 24
logging enable
logging asdm informational
logging trap informational
logging host <Logging server IP> inside
no logging message 305009
no logging message 305010
no logging message 305011
no logging message 305012
no logging message 302014
no logging message 302016
mtu outside 1500
mtu inside 1500
mtu <dmz> 1500
failover
failover lan unit primary
failover lan interface failover Gig-
abitEthernet0/3
failover polltime unit 2 holdtime 6
failover polltime interface 4 holdtime
20
failover link failover Gig-
abitEthernet0/3
failover interface ip failover
192.168.255.1 255.255.255.252 standby
192.168.255.2
icmp unreachable rate-limit 1 burst-
size 1
icmp permit any inside
asdm image disk0:/asdm-613.bin
no asdm history enable
arp timeout 14400
<insert global config here>
<insert nat config here>
<insert static config here>
<insert access-group config here>
route outside 0.0.0.0 0.0.0.0 <DS3
router HSRP> 1
<insert other static routes here>
timeout xlate 3:00:00
timeout conn 1:00:00 half-closed
0:10:00 udp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00
h225 1:00:00 mgcp 0:05:00 mgcp-pat
```

```
aaa local authen-
tication
attempts max-fail 3
<insert AAA config
here>
http server enable
<insert CSM/ASDM
access here>
<insert SNMP server
configuration here>
snmp-server enable
traps snmp
authentication linkup
linkdown coldstart
<insert SSH allowed
subnets here>
ssh timeout 5
console timeout 0
!
threat-detection
basic-threat
threat-detection stat-
istics access-list
no threat-detection
statistics tcp-inter-
cept
<insert local
usernames here>
!
class-map inspection_
default
 match default-inspec-
tion-traffic
!
policy-map type
inspect dns
preset_dns_map
 parameters
  message-length max-
imum 1024
policy-map global_
policy
 class inspection_
default
  inspect dns preset_
dns_map
  inspect ftp
  inspect h323 h225
```

agement0/0
shutdown
  no nameif
  no security-level
  no ip address
	0:05:00
	timeout sip 0:30:00 sip_media 0:02:00
	sip-invite 0:03:00 sip-disconnect
	0:02:00
	timeout sip-provisional-media
	0:02:00
	uauth 0:05:00 absolute
		inspect h323 ras
		  inspect rsh
		  inspect rtsp
		  inspect esmtp
		  inspect sqlnet
		  inspect skinny
		  inspect sunrpc
		  inspect xdmcp
		  inspect sip
		  inspect netbios
		  inspect tftp
		!
		service-policy
		global_policy global
		prompt hostname con-
		text

Data Center Edge

This section describes the requirements, design, and configuration of the Data Center firewall services in the Cisco Connected Stadium.

Figure 29. Cisco Connected Stadium Data Center Design

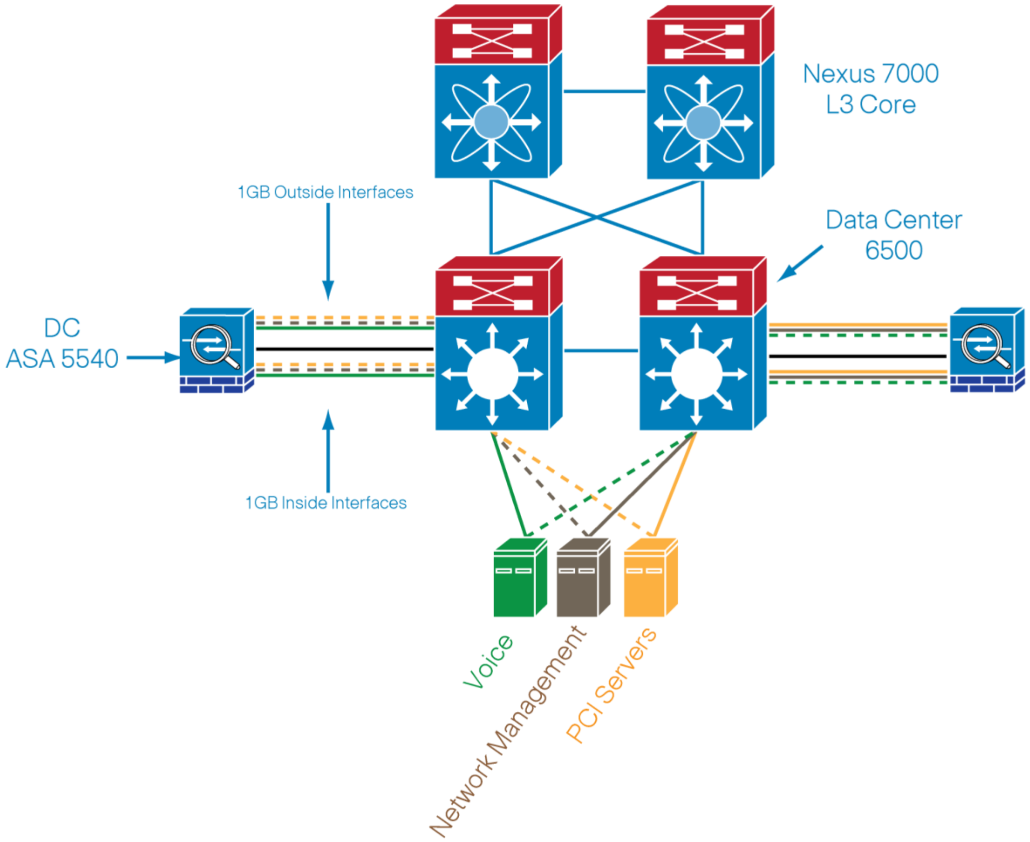

It is recommended to expand the protection the Data Center firewalls provide to cover all critical business application servers.

Firewall Design Considerations

The following data center firewall design considerations apply to the Cisco Connected Stadium:

- The Data Center firewalls are configured for multiple contexts. The contexts are:
  - uc—Firewalling for Unified Communications servers
  - pci—Firewalling for servers that fall under the PCI data security standard
  - nms-sec—Firewalling for NMS and Security management servers

NOTE: Multiple context mode does not support the following features: Dynamic routing protocols (security contexts support only static routes), VPN, and Multicast.

- The Data Center firewalls are configured in transparent mode, allowing support for multicast traffic. Multicontext mode supports bridging multicast.
- The Data Center firewalls are managed via SSH/HTTPS, telnet is disabled.

NOTE: In multiple context mode, the configuration of SSH access should be done in context level. Once you successfully connect and authenticate to the admin context or the system execution space, you can configure and monitor any of the contexts as necessary without opening sessions to the individual virtual firewalls

- The Data Center firewalls are configured to authenticate, authorize, and account all device access to the ISE server via RADIUS.
- The Data Center firewalls are configured to use NTP.
- The Data Center firewalls are configured to log events to a logging server.
- The Data Center firewalls are configured in an Active/Active design to maximize traffic flow.
- The Data Center firewalls are configured with individual management IP addresses per context.

Firewall Architecture

The data center firewall architecture used in this example (as shown in Figure 30) consists of:

- The Data Center ASA 5540 redundant pair, with an additional 4GE-SSM module.
- The Data Center Catalyst 6500 switches where the Data Center ASA firewalls attach their inside and outside interfaces for each context. The failover interface is also connected to these switches.

Figure 30. Data Center Firewall Architecture

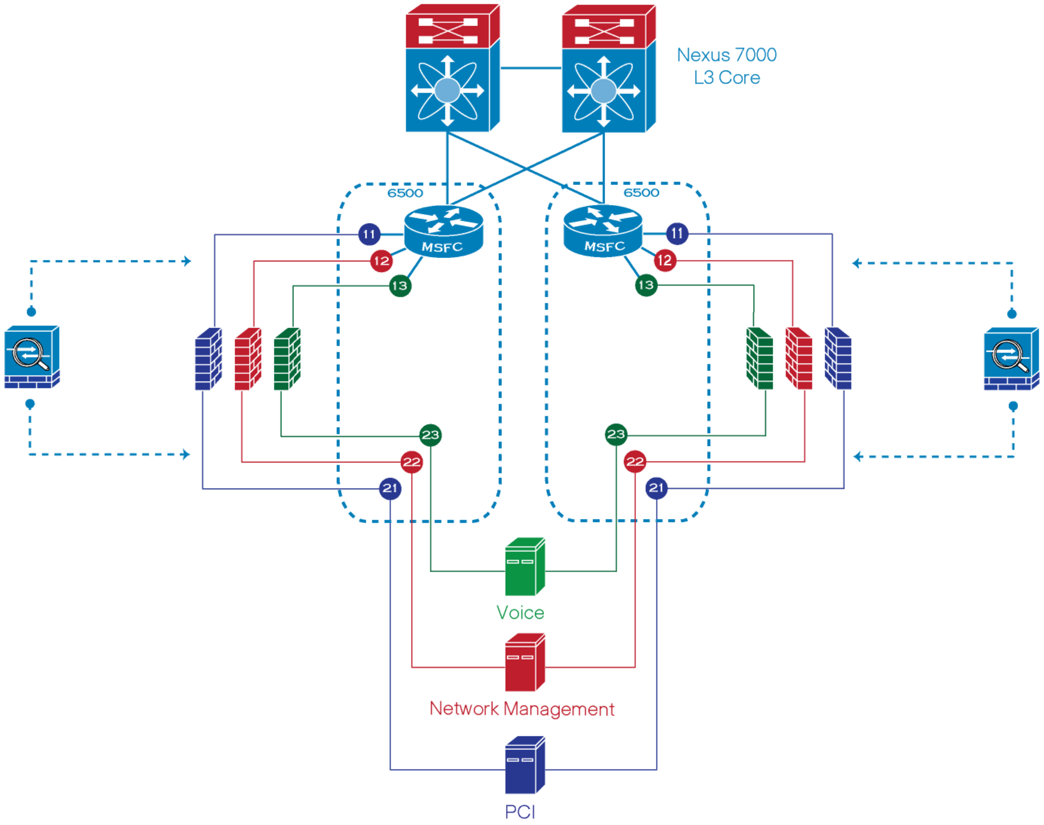

Each security context is configured to have an ‘inside’ and an ‘outside’ interface on the ASA. These interfaces are physical GigabitEthernet links. A given security context bridges the protected servers on the VLAN connected to the inside interface to the gateway for that subnet, on the VLAN connected to the outside interface. All traffic through a security context is subject to stateful inspection and TCP normalization, along with ACLs applied to the interfaces.

Initially all traffic is allowed between the inside and outside interfaces of a given context. IT staff can refine their ACLs and add additional policies as needed after the network has been verified and handed over.

Ports and Security Level Assignment

Figure 31. Data Center ASA Interfaces

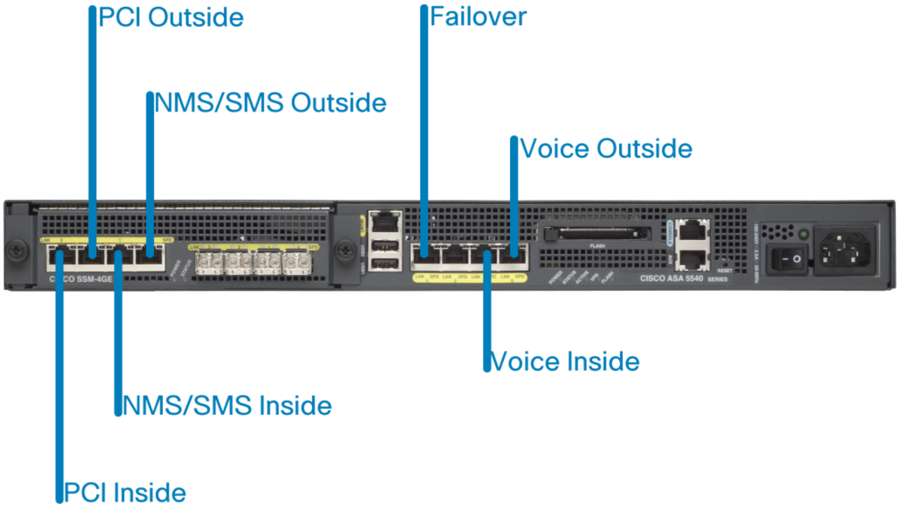

Table 7 lists the ASA interface configuration values.

Table 7. ASA Interface Configuration

Context	Interface	VLAN ID	Port/IF Connection	Security Level	Description
sec-nms	outside	405	GigabitEthernet0/0	0	Outside interface for sec-nms context.
sec-nms	inside	404	GigabitEthernet0/1	100	Inside interface for sec-nms context.
			GigabitEthernet0/2		
	failover		GigabitEthernet0/3	n/a	Connection between ASA pair for failover communications. Configured in the system context.
uc	outside	403	GigabitEthernet1/0	0	Outside interface for uc context.
uc	inside	402	GigabitEthernet1/1	100	Inside interface for uc context.
pci	outside		GigabitEthernet1/2	0	Outside interface for pci context.
pci	inside		GigabitEthernet1/3	100	Inside interface for pci context.

Example 3. Data Center ASA Configuration Example

Primary DC System Context:

```
firewall transparent
hostname DC-ASA
enable password <removed>
password <removed>
asdm image <image>
boot system <image>
mac-address auto
interface gigabitethernet0/0
   no shutdown
interface gigabitethernet0/1
   no shutdown
interface gigabitethernet0/2
   shutdown
interface gigabitethernet0/3
   no shutdown
interface gigabitethernet1/0
   no shutdown
interface gigabitethernet1/1
   no shutdown
interface gigabitethernet1/2
   no shutdown
interface gigabitethernet1/3
   no shutdown
failover
failover lan unit primary
failover lan interface
failover gigabitethernet0/3
failover polltime unit 2
holdtime 6
failover polltime interface 4
holdtime 20
failover link failover
gigabitethernet0/3
failover interface ip failover
192.168.255.5 255.255.255.252
standby 192.168.255.6
failover group 1
   primary
   preempt
failover group 2
   secondary
   preempt
admin-context sec-nms
context sec-nms
   description Network
Management and Security
   allocate-interface
```

```
context uc
   description Voice Servers
   allocate-interface
gigabitethernet1/0
   allocate-interface
gigabitethernet1/1
   config-url flash:/uc.cfg
   join-failover-group 2
context pci
   description PCI servers
   allocate-interface
gigabitethernet1/2
   allocate-interface
gigabitethernet1/3
   config-url flash:/pci.cfg
   join-failover-group 2
<insert local usernames here>
```

Secondary DC System Context:

```
firewall transparent
failover
failover lan unit secondary
failover lan interface
failover gigabitethernet0/3
failover link failover
gigabitethernet0/3
failover interface ip failover
192.168.255.5 255.255.255.252
standby 192.168.255.6
```

Primary DC Sec NMS Context:

```
enable password
<password>
password
<password>
hostname sec-nms
interface
gigabitethernet0/0
   nameif outside
   security-level
0
interface
gigabitethernet0/1
   nameif inside
   security-level
100
access-list
inside_acl
extended permit ip
any any
access-list
outside_acl
extended permit ip
any any
access-list
etherlist
ethertype permit
any
access-group
etherlist in
interface inside
access-group
etherlist in
interface outside
access-group
inside_acl in
interface inside
access-group
outside_acl in
interface outside
logging enable
logging asdm
informational
logging trap
informational
logging host
<LOGGING SERVER
IP> inside
```

gigabitethernet0/0 allocate-interface gigabitethernet0/1 config-url flash:/sec-nms.cfg join-failover-group 1		no logging message 305009 no logging message 305010 no logging message 305011 no logging message 305012 no logging message 302014 no logging message 302016 ip address <IP> <subnet> standby <standby IP> monitor-interface outside monitor-interface inside route <Interface> 0.0.0.0 0.0.0.0 <gateway> 1 <insert AAA config here> http server enable <insert CSM/ASDM access here> <insert SNMP server configuration here> <insert SSH allowed subnets here> <insert local usernames here> <insert qos class/policies here>

Primary DC US Context:	Primary DC PCI Context:	Secondary DC PCI Context:
enable password <password> password <password> hostname voice interface gigabitethernet1/0 nameif outside security-level 0 interface gigabitethernet1/1 nameif inside security-level 100 access-list inside_acl extended permit ip any any access-list outside_acl extended permit ip any any access-list etherlist ethertype permit any access-group etherlist in interface inside access-group etherlist in interface outside access-group inside_acl in interface inside access-group outside_acl in interface outside logging enable logging asdm informational logging trap informational logging host <LOGGING SERVER IP> inside no logging message 305009 no logging message 305010 no logging message 305011 no logging message 305012 no logging message 302014 no logging message 302016 ip address <IP> <subnet> standby <standby IP> monitor-interface outside monitor-interface inside route <Interface> 0.0.0.0 0.0.0.0 <gateway> 1 <insert AAA config here> http server enable <insert CSM/ASDM access here> <insert SNMP server configuration here <insert SSH allowed subnets	enable password <password> password <password> hostname pci interface gigabitethernet1/2 nameif outside security-level 0 interface gigabitethernet1/3 nameif inside security-level 100 access-list inside_acl extended permit ip any any access-list outside_acl extended permit ip any any access-list etherlist ethertype permit any access-group etherlist in interface inside access-group etherlist in interface outside access-group inside_acl in interface inside access-group outside_acl in interface outside logging enable logging asdm informational logging trap informational logging host <LOGGING SERVER IP> inside no logging message 305009 no logging message 305010 no logging message 305011 no logging message 305012 no logging message 302014 no logging message 302016 ip address <IP> <subnet> standby <standby IP> monitor-interface outside monitor-interface insid route <Interface> 0.0.0.0 0.0.0.0 <gateway> 1 <insert AAA config here> http server enable	enable password <password> password <password hostname pci interface gigabitethernet1/2 nameif outside security-level 0 interface gigabitethernet1/3 nameif inside security-level 100 access-list inside_acl extended permit ip any any access-list outside_acl extended permit ip any any access-list etherlist ethertype permit any access-group etherlist in interface inside access-group etherlist in interface outside access-group inside_acl in interface inside access-group outside_acl in interface outside logging enable logging asdm informational logging trap informational logging host <LOGGING SERVER IP> inside no logging message 305009 no logging message 305010 no logging message 30501 no logging message 30501 no logging message 30201 no logging message 30201 ip address <IP> <subnet> standby <standby IP monitor-interface outsid monitor-interface inside route <Interface> 0.0.0.0 0.0.0.0

here> <insert local usernames here> <insert qos class/policies here>	<insert CSM/ASDM access here> <insert SNMP server configuration here> <insert SSH allowed subnets here> <insert local usernames here> <insert qos class/policies here>	<gateway> 1 <insert AAA config here> http server enable <insert CSM/ASDM access here> <insert SNMP server configuration here> <insert SSH allowed subnets here> <insert local usernames here> <insert qos class/policies here>

Virtual Private Network (VPN)

This section describes the requirements, design, and configuration of the Virtual Private Network (VPN) services in the Cisco Connected Stadium.

VPN IPSec / SSL

From a VPN perspective, there are different methods that can be used to establish a VPN connection using the ASA. IPSec and SSL VPN are the two methods described in this document. This document does not cover clientless and thin-client SSL modes.

IPSec

The ASA provides the option of using IPsec for client-to-LAN VPN connections. In IPsec terminology, a peer is a remote-access client or another secure gateway. For both connection types, the security appliance supports only Cisco peers. Because Cisco adheres to VPN industry standards, ASAs may work with other vendors' peers; however, Cisco does not support them.

During tunnel establishment, the two peers negotiate security associations that govern authentication, encryption, encapsulation, and key management. These negotiations involve two phases: first, to establish the tunnel (the IKE SA); and second, to govern traffic within the tunnel (the IPsec SA).

SSL

The Cisco AnyConnect SSL VPN Client provides secure SSL connections to the security appliance for remote users. Without a previously-installed client, remote users enter in their browser the IP address of an interface configured to accept SSL VPN connections. Users must enter the URL in the form https://<address>. After connecting to the URL, a client downloads that matches the operating system of the remote computer. After downloading, the client installs and configures itself, establishes a secure SSL connection and either remains or uninstalls itself.

VPN Design Considerations

The following VPN design considerations apply to the Cisco Connected Stadium:

- ASA VPN provides secure access protected by inspection rules.
- ASA VPN is managed via SSH/HTTPS, telnet is disabled.
- ASA VPN is configured to authenticate, authorize, and account all device access to the ISE server via RADIUS.
- ASA VPN is configured to use NTP.
- ASA VPN is configured to log events to a logging server.

VPN Architecture

The VPN architecture (shown in Connected Stadium VPN Architecture) consists of the following components:

- Pair of ASA 5540s.
- The external Catalyst 3750 switches where the data center VPN ASAs attach their outside interfaces.
- The Catalyst 6500 data center switches where the data center VPN ASAs attach to their inside.

Figure 32. Connected Stadium VPN Architecture

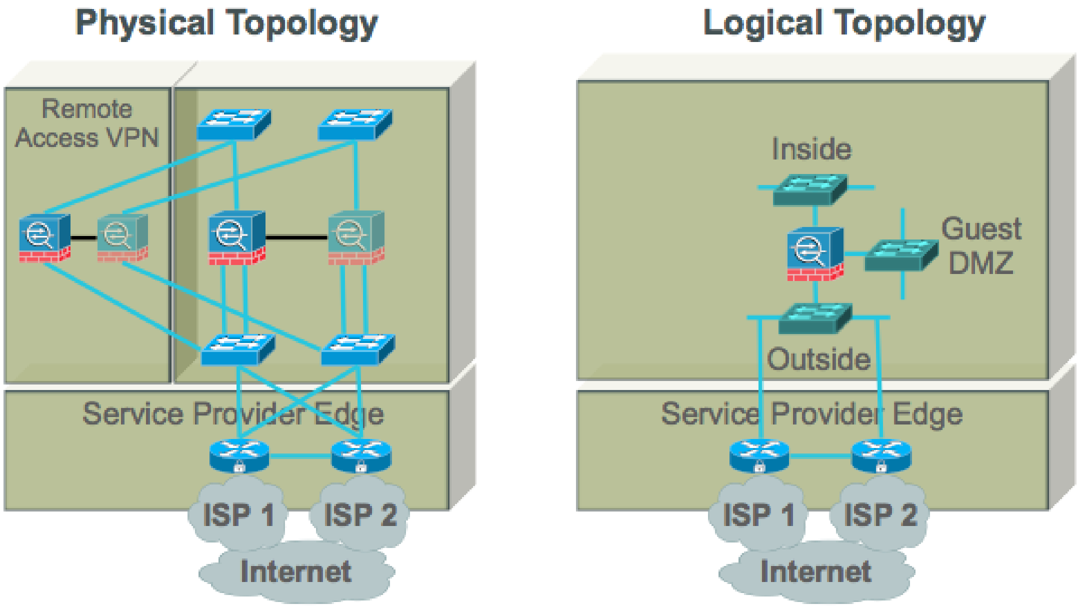

Ports and Security Level Assignment

Table 8 lists the applicable VLANs that are used for the interfaces of the ASAs (shown in Figure 33).

Table 8. VPN ASA VLAN Assignment

Interface Name	Physical Interface	Security Level	Description
outside	GigabitEthernet0/0	0	Connection to the Internet routers
inside	GigabitEthernet0/1	100	Connection to the internal stadium network
failover	GigabitEthernet 0/3	n/a	Failover Interface

Figure 33. VPN ASA Interfaces

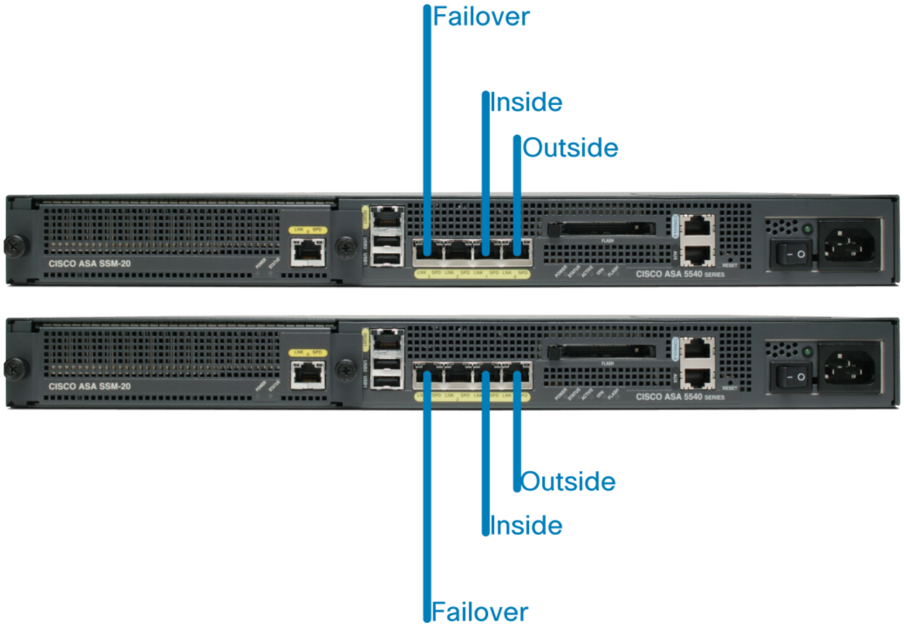

Load Balancing

Load balancing is the ability to have Cisco VPN Clients shared across multiple ASA units without user intervention. Load-balancing ensures that the public IP address is highly available to users. For example, if the Cisco ASA that services the public IP address fails, another ASA in the cluster assumes the public IP address.

VPN ASA Load Balancing Configuration

Both VPN appliances need the same configuration and, when complete, will participate in what is called a cluster. The cluster load balances clients amongst themselves with the intent to not overload one of the appliances.

```
VPN-ASA(config)#vpn load-balancing
VPN-ASA(config-load-balancing)#priority 10
VPN-ASA(config-load-balancing)#cluster key cisco123
VPN-ASA(config-load-balancing)#cluster ip address 172.16.172.54
VPN-ASA(config-load-balancing)#cluster encryption
VPN-ASA(config-load-balancing)#participate
```

AnyConnect

The Cisco AnyConnect 2.0 client can be used on a variety of operating systems, such as Windows 2000, XP, Vista, Linux (Multiple Distros) and MAC OS X. The AnyConnect client can be installed manually on the remote PC by the system administrator. It can also be loaded onto the security appliance and made ready for download to remote users. After the application is downloaded, it can automatically uninstall itself after the connection terminates, or it can remain on the remote PC for future SSL VPN connections. The following example makes the AnyConnect client ready to download upon successful browser-based SSL authentication.

1. Enable the AnyConnect client.

```
VPN_ASA(config)#webvpn
VPN_ASA(config-webvpn)#enable outside
VPN_ASA(config-webvpn)#svc enable
VPN_ASA(config)#sysopt connection permit-vpn
```

2. Generate a crypto key complete certificate enrollment.

```
VPN_ASA(config)#crypto key generate rsa label sslvpnkeypair
VPN_ASA(config)#crypto ca trustpoint localtrust
VPN_ASA(config-ca-trustpoint)#enrollment self
VPN_ASA(config-ca-trustpoint)#fqdn sslvpn.cisco.com
VPN_ASA(config-ca-trustpoint)#subject-name CN=sslvpn.stadium.com
VPN_ASA(config-ca-trustpoint)#keypair sslvpnkeypair
VPN_ASA(config-ca-trustpoint)#crypto ca enroll localtrust
noconfirm
VPN_ASA(config)# ssl trust-point localtrust outside
```

3. Create tunnel group policies.

```
VPN_ASA(config)#tunnel-group SSLClientProfile type remote-access
VPN_ASA(config)#tunnel-group SSLClientProfile general-attributes
VPN_ASA(config-tunnel-general)#default-group-policy
SSLCLientPolicy
VPN_ASA(config-tunnel-general)#tunnel-group SSLClientProfile
webvpn-attributes
VPN_ASA(config-tunnel-webvpn)#group-alias SSLVPNClient enable
VPN_ASA(config-tunnel-webvpn)#webvpn
VPN_ASA(config-webvpn)#tunnel-group-list enable
```

VPN ASA Configuration

To begin configuring the VPN ASA:

1. Connect to Primary ASA Unit via console port.
2. At the prompt enter **enable**.

```
ciscoasa>enable
ciscoasa#
```

3. At the enable mode prompt enter **conf term**.

```
ciscoasa#conf term
ciscoasa(config)#
```

Hostname Configuration

To configure the system hostname using the hostname command.

```
ciscoasa(config)#hostname VPN_ASA

VPN_ASA(config)#
```

Notice that the prompt changes to reflect just configured name.

Interface Configuration

Table 9 lists the ASA interface configuration values.

Table 9. ASA Interface Configuration

Interface	VLAN ID	Port/IF Connection	IP Address & Prefix / Standby IP	Sec Level
outside		GigabitEthernet 0/0		0
inside		GigabitEthernet 0/1		100
failover		GigabitEthernet 0/3		n/a

The following basic interface configuration should be done in single context mode.

1. Configure the physical interfaces. The physical interfaces should be configured according to following template.

```
VPN_ASA(config)# interface gigabitethernet x/x
VPN_ASA(config-if)#description interface_description
VPN_ASA(config-if)# nameif interface_name
VPN_ASA(config-if)# security-level level
VPN_ASA(config-if)# ip address ip_address subnet_mask standby
secondary_ip_address
```

```
VPN_ASA(config-if)# no shutdown
```

2. Configure the subinterfaces. An interface with one or more VLAN subinterfaces is automatically configured as an 802.1Q trunk. The physical interface must be enabled before any traffic can pass through an enabled subinterface.

```
VPN_ASA(config)# interface gigabitethernet x/y.z
VPN_ASA(config-subif)# vlan A
VPN_ASA(config-subif)# no shutdown
```

   When using multiple context mode, configuring subinterfaces should be done in system configuration mode.

3. Configure the management interface.

```
VPN_ASA(config)# interface management y/y
VPN_ASA(config-if)# no shutdown
VPN_ASA(config-if)# ip address ip_address subnet_mask standby
secondary_ip_address
VPN_ASA(config-if)# management-only
VPN_ASA(config-if)# exit
```

Saving Configurations

To save the configuration, enter the following command.

```
VPN_ASA# write memory
```

ASA Management

To manage the ASA, you must enable the password and the desired type of remote access. To set the enable password for privileged EXEC mode, use the enable password command in global configuration mode.

```
VPN_ASA(config)# enable password password [level level] [encrypted]
```

SSH Access

To allow SSH access to the ASA:

1. Specify the domain name to which the ASA belongs.

```
VPN_ASA(config)# domain-name domain_name
```

2. To generate an RSA key pair, which is required for SSH, enter the following command.

```
VPN_ASA(config)# crypto key generate rsa modulus modulus_size
```

3. To save the RSA keys to persistent Flash memory, enter the following command.

```
VPN_ASA(config)# write mem
```

4. To identify the IP addresses from which the security appliance accepts connections, enter the following command for each address or subnet.

```
VPN_ASA(config)# ssh source_IP_address mask source_interface
```

5. To set the duration for how long an SSH session can be idle before the security appliance disconnects the session, enter the following command.

```
VPN_ASA(config)# ssh timeout minutes
```

HTTPS Access for CSM and ASDM

To allow HTTPS access to the ASA:

1. To identify the IP addresses from which the security appliance accepts HTTPS connections, enter the following command for each address or subnet.

```
VPN_ASA(config)# http source_IP_address mask source_interface
```

2. To enable the HTTPS server, enter the following command.

```
VPN_ASA(config)# http server enable
```

3. To specify the location of the ASDM image, enter the following command.

```
VPN_ASA(config)# asdm image disk0:/asdmfile
```

AAA for System Administrators

To authenticate users who access the CLI, enter the following commands.

```
VPN_ASA(config)# aaa new-model
VPN_ASA(config)# aaa authentication {telnet | ssh | http | serial}
console {LOCAL | server_group [LOCAL]}
```

NOTE: Use a server group followed by LOCAL to use the local database as a backup.

To authenticate users who enter the enable command, enter the following commands.

```
VPN_ASA(config)# aaa new-model
VPN_ASA(config)# aaa authentication enable console {LOCAL | server_
group [LOCAL]}
```

Login Banner

You can configure a message to display when a user connects to the security appliance, before a user logs in, or before a user enters privileged EXEC mode.

```
VPN_ASA(config)# banner {exec | login | motd} text
```

Adds a banner to display at one of three times: when a user first connects (message-of-the-day (motd)), when a user logs in (login), and when a user accesses privileged EXEC mode (exec). When a user connects to the security appliance, the message-of-the-day banner appears first, followed by the login banner and prompts. After the user successfully logs in to the security appliance, the exec banner id displayed.

Network Time Protocol Settings

To use the clock from another device in the network, use following commands.

```
VPN_ASA(config)# ntp server ip-address [version number] [key keyid]
[source interface] [prefer]
```

By default, this synchronization is unauthenticated. To add authentication to the clock sources, use the following commands.

1. Enable the NTP authentication feature.

   ```
   VPN_ASA(config)# ntp authenticate
   ```

2. Define the authentication keys. Each key has a key number, type and value.

   ```
   VPN_ASA(config)# ntp authentication-key number md5 value
   ```

3. Define which authentication keys are trusted. If a key is trusted, the system synchronizes to the clock of the source that originates this key in its NTP packets.

   ```
   VPN_ASA(config)# ntp trusted-key key-number
   ```

SNMP

The ASA provides support for network monitoring using SNMP v1 and v2c. The ASA supports traps and SNMP read access, but does not support SNMP write access.

You can configure the ASA to send traps (event notifications) to an NMS, or you can use the NMS to browse the MIBs on the ASA.

To enable the SNMP agent and identify an NMS that can connect to the ASA:

1. Ensure that the SNMP server on the ASA is enabled.

   ```
   VPN_ASA(config)# snmp-server enable
   ```

2. Identify the IP address of the NMS that can connect to the ASA.

   ```
   VPN_ASA(config)# snmp-server host interface_name ip_address [trap
   | poll] [community text] [version 1 | 2c] [udp-port port]
   ```

   Specify trap or poll if you want to limit the NMS to receiving traps only or browsing (polling) only. By default, the NMS can use both functions. SNMP traps are sent on UDP port 162 by default. You can change the port number using the udp-port keyword.

3. To specify the community string, enter the following command.

   ```
   VPN_ASA(config)# snmp-server community key
   ```

4. Enable the ASA to send traps to the NMS.

   ```
   VPN_ASA(config)# snmp-server enable traps [all | syslog | snmp
   [trap] [...] | entity [trap] [...] | ipsec [trap] [...] | remote-
   access [trap]
   ```

   Enter this command for each feature type to enable individual traps or sets of traps, or enter the all keyword to enable all traps.

Logging

When using the ASA in single mode, logging is enabled using the following command.

```
VPN_ASA(config)# logging enable
```

Output to Syslog Server

To configure the ASA to send logs to a Syslog server:

1. Define syslog server to receive logs.

   ```
   VPN_ASA(config)# logging host interface_name ip_address [tcp
   [/port] | udp[/port]] [format emblem]
   ```

If you want to designate more than one syslog server as an output destination, enter a new command for each syslog server.

2. Specify which system log messages should be sent to the syslog server.

```
VPN_ASA(config)# logging trap {severity_level | message_list}
```

3. (optional) Continue TCP logging when the syslog server is down.

```
VPN_ASA(config)# logging host interface_name server_ip [tcp/port]
[permit-hostdown]
```

4. (optional) Set the logging facility to a value other than its default of 20.

```
VPN_ASA(config)# logging facility number
```

Output to Console

To configure the ASA to send logs to the console port, use the following command.

```
VPN_ASA(config)# logging console {severity_level | message_list}
```

It is recommended to disable logging to console.

Output to ASDM

To configure the ASA to send logs to an ASDM:

1. Specify the system log messages that should go to ASDM.

```
VPN_ASA(config)# logging asdm {severity_level | message_list}
```

2. Define the number of system log messages retained in the ASDM log buffer.

```
VPN_ASA(config)# logging asdm-buffer-size num_of_msgs
```

3. (optional) To erase the current contents of the ASDM log buffer, enter the following command.

```
VPN_ASA(config)# clear logging asdm
```

QoS

The ASA achieves QoS by allowing two types of traffic queues for each interface: a low-latency queue (LLQ) and a default queue. Only the default traffic is subject to rate limiting.

To configuring a traditional QoS policy for the security appliance:

1. Defining traffic classes (**class-map** command).

```
VPN_ASA(config)# class-map class-map-name-1
VPN_ASA(config-cmap)# match match-criteria-1
VPN_ASA(config-cmap)# exit
VPN_ASA(config)# class-map class-map-name-2
VPN_ASA(config-cmap)# match match-criteria-2
VPN_ASA(config-cmap)# exit
```

2. Associating policies and actions with each class of traffic (**policy-map** command).

```
VPN_ASA(config)# policy-map policy-map-name
VPN_ASA(config-pmap)# class class-map-name-1
VPN_ASA(config-pmap-c)# policy-1
VPN_ASA(config-pmap-c)# policy-n
VPN_ASA(config-pmap-c)# exit
VPN_ASA(config-pmap)# class class-map-name-2
VPN_ASA(config-pmap-c)# policy-1
VPN_ASA(config-pmap-c)# policy-n
```

3. Attaching policies to logical or physical interfaces (**service-policy** command).

```
VPN_ASA(config)# service-policy policy-map-name interface
[outside|inside]
```

In addition, if you are differentiating between priority traffic and best-effort traffic, you must define a low-latency queue (**priority-queue** command) on each named, physical interface transmitting prioritized traffic.

The following example enables a default priority-queue with the default queue-limit and tx-ring-limit.

```
VPN_ASA(config)# priority-queue name_interface
```

The size that you specify for the priority queue affects both the low latency queue and the best-effort queue. The **queue-limit** command specifies a maximum number of packets that can be queued to a priority queue before it drops data. This limit must be in the range of 0 through 2048 packets.

```
VPN_ASA(config)# priority-queue name_interface
VPN_ASA(config-priority-queue)# queue-limit number
```

The **tx-ring-limit** command lets you configure the maximum number of packets (that is, the depth) allowed to be queued in the Ethernet transmit driver ring at any given time.

```
VPN_ASA(config)# priority-queue name_interface
VPN_ASA(config-priority-queue)# tx-ring-limit number
```

Access Lists

Object groupings simplify the implementation of an access list. By grouping similar objects together, object groups can be used in an ACE (Access Control Entry) instead of having to enter an ACE for each object separately. Following types of object groups can be created:

- Protocol
- Network
- Service
- ICMP type

To add a protocol group, enter the following commands.

```
VPN_ASA(config)# object-group protocol grp_id
VPN_ASA(config-protocol)# description text
VPN_ASA(config-protocol)# protocol-object protocol
```

To add a network group, enter the following commands.

```
VPN_ASA(config)# object-group network grp_id
VPN_ASA(config-network)# description text
VPN_ASA(config-network)# network-object {host ip_address | ip_address mask}
```

To add a service group, enter the following commands.

```
VPN_ASA(config)# object-group service grp_id {tcp | udp | tcp-udp}
VPN_ASA(config-service)# description text
VPN_ASA(config-service)# port-object {eq port | range begin_port end_port}
```

To add an ICMP type group, enter the following commands.

```
VPN_ASA(config)# object-group icmp-type grp_id
VPN_ASA(config-icmp-type)# description text
VPN_ASA(config-icmp-type)# icmp-object icmp_type
```

Once the object groups have been defined, the following command is used to configure the access list.

```
VPN_ASA(config)# access-list id [line line-number] [extended] {deny
| permit} {protocol | object-group protocol_obj_grp_id} {src_ip mask
| interface ifc_name | object-group network_obj_grp_id} [operator
port | object-group service_obj_grp_id] {dest_ip mask | interface
ifc_name | object-group network_obj_grp_id} [operator port | object-
group service_obj_grp_id | object-group icmp_type_obj_grp_id][log
[[level] [interval secs] | disable | default]][inactive |time-range
time_range_name]
```

Then to apply access lists on ingress on the firewall interfaces, use the following command.

```
VPN_ASA(config)#access-group access-list in interface interface_name
```

Inspection

When a user establishes a connection, the security appliance checks the packet against access lists, creates an address translation, and creates an entry for the session in the fast path, so that further packets can bypass time-consuming checks. However, the fast path relies on predictable port numbers and does not perform address translations inside a packet.

Many protocols open secondary TCP or UDP ports. The initial session on a well-known port is used to negotiate dynamically assigned port numbers. Other applications embed an IP address in the packet that needs to match the source address that is normally translated when it goes through the security appliance.

If you use applications like these, then you need to enable application inspection.

Default Inspection Policy

By default, the configuration includes a policy that matches all default application inspection traffic and applies inspection to the traffic on all interfaces (a global policy).

The default policy configuration follows.

```
class-map inspection_default
 match default-inspection-traffic
policy-map type inspect dns preset_dns_map
 parameters
  message-length maximum 512
```

```
policy-map global_policy
  class inspection_default
   inspect dns preset_dns_map
   inspect ftp
   inspect h323 h225
   inspect h323 ras
   inspect rsh
   inspect rtsp
   inspect esmtp
   inspect sqlnet
   inspect skinny
   inspect sunrpc
   inspect xdmcp
   inspect sip
   inspect netbios
   inspect tftp
service-policy global_policy global
```

Application Inspection Configuration

If you want to match non-standard ports, then you need to create a new class map. You can combine multiple class maps in the same policy if desired, so you can create one class map to match certain traffic, and another to match different traffic. However, if traffic matches a class map that contains an inspection command, and then matches another class map that also has an inspection command, only the first matching class is used.

1. Identify traffic which you want to inspect.

   ```
   VPN_ASA(config)# class-map class_map_name
   VPN_ASA(config-cmap)# match access-list acl_inspect
   ```

2. Add or edit a Layer 3/4 policy map that sets the actions to take with the class map traffic.

   ```
   VPN_ASA(config)# policy-map name
   VPN_ASA(config-pmap)#
   ```

   The default policy map is called "global_policy." This policy map includes the default inspections list. If you want to modify the default policy (for example, to add or delete an inspection, or to identify an additional class map for your actions), then enter global_policy as the name.

3. Identify the class map from step 1 to which you want to assign an action.

```
VPN_ASA(config-pmap)# class class_map_name
VPN_ASA(config-pmap-c)#
```

   If you are editing the default policy map, it includes the inspection_default class map. You can edit the actions for this class by entering inspection_default as the name. To add an additional class map to this policy map, identify a different name. You can combine multiple class maps in the same policy if desired, so you can create one class map to match certain traffic, and another to match different traffic.

4. Enable application inspection.

```
VPN_ASA(config-pmap-c)# inspect protocol
```

5. Activate the policy map on one or more interfaces.

```
VPN_ASA(config)# service-policy policymap_name {global |
interface interface_name}
```

Where global applies the policy map to all interfaces, and interface applies the policy to one interface. By default, the default policy map, "global_policy," is applied globally. Only one global policy is allowed. You can override the global policy on an interface by applying a service policy to that interface. You can only apply one policy map to each interface.

TCP Normalization

The TCP normalization feature uses Modular Policy Framework, so that implementing TCP normalization consists of identifying traffic, specifying the TCP normalization criteria, and activating TCP normalization on an interface.

1. To specify the TCP normalization criteria that you want to look for, create a TCP map by entering the following command.

```
VPN_ASA(config)# tcp-map tcp-map-name
```

2. Configure the TCP map criteria by entering commands for one or more of the following options:

   ○ Prevent inconsistent TCP retransmissions.

```
VPN_ASA(config-tcp-map)# check-retransmission
```

- Verify the checksum.

```
VPN_ASA(config-tcp-map)# checksum-verification
```

- Allow packets whose data length exceeds the TCP maximum segment size.

```
VPN_ASA(config-tcp-map)# exceed-mss {allow | drop}
```

- Set the maximum number of out-of-order packets that can be queued for a TCP connection.

```
VPN_ASA(config-tcp-map)# queue-limit pkt_num
```

- Clear reserved bits in the TCP header, or drop packets with reserved bits set.

```
VPN_ASA(config-tcp-map)# reserved-bits {allow | clear | drop}
```

- Drop SYN packets with data.

```
VPN_ASA(config-tcp-map)# syn-data {allow | drop}
```

- Clear the selective-ack, timestamps, or window-scale TCP options.

```
VPN_ASA(config-tcp-map)# tcp-options {selective-ack | timestamp |
window-scale} {allow | clear}
```

- Disable the TTL evasion protection.

```
VPN_ASA(config-tcp-map)# ttl-evasion-protection
```

- Allow the URG pointer.

```
VPN_ASA(config-tcp-map)# urgent-flag {allow | clear}
```

- Drop a connection that has changed its window size unexpectedly.

```
VPN_ASA(config-tcp-map)# window-variation {allow | drop}
```

3. Identify the traffic to which you want to apply TCP normalization, add a class map using the class-map command.

```
VPN_ASA(config)# class-map class-map-name
VPN_ASA(config-cmap)# match match-criteria
VPN_ASA(config-cmap)# exit
```

4. Add or edit a policy map that sets the actions to take with the class map traffic and apply the TCP map.

```
VPN_ASA(config)# policy-map name
VPN_ASA(config-pmap)# class class_map_name
VPN_ASA(config-pmap-c)# set connection advanced-options tcp-map-
name
```

Other options to apply to the connections are:

- Set maximum connection limits or whether TCP sequence randomization is enabled.

```
VPN_ASA(config-pmap-c)# set connection {[conn-max n] [embryonic-
conn-max n] [per-client-embryonic-max n] [per-client-max n]
[random-sequence-number {enable | disable}]}
```

- Set connection timeouts.

```
VPN_ASA(config-pmap-c)# set connection timeout {[embryonic
hh:mm:ss] {tcp hh:mm:ss [reset]] [half-closed hh:mm:ss] [dcd
hh:mm:ss [max_retries]]}
```

NOTE: Embryonic connections are half-open connections.

5. Activate the policy map on one or more interfaces

-
  ```
  VPN_ASA(config)# service-policy policymap_name {global |
  interface interface_name}
  ```

Network Address Translation

Network Address Translation (NAT), as well as Port Address Translation (PAT), can be dynamic or static.

Dynamic NAT and PAT

For dynamic NAT and PAT, first configure a **nat** command identifying the real addresses on a given interface that you want to translate. Then configure a separate global command to specify the mapped addresses when exiting another interface (in the case of PAT, this is one address). Each **nat** command matches a global command by comparing the NAT ID, a number that you assign to each command.

Basic NAT

To identify addresses on one interface that are translated to mapped addresses on another interface, use the nat command in global configuration mode. This command configures dynamic NAT, where an address is translated to one of a pool of mapped addresses.

```
VPN_ASA(config)# nat (real_interface) nat_id access-list acl_name
[dns] [outside] [norandomseq] [[tcp] tcp_max_conns [emb_limit]] [udp
udp_max_conns]
```

To create a pool of mapped addresses for NAT, use the global command in global configuration mode.

```
VPN_ASA(config)# global (mapped_if) nat_id {mapped_ip[-mapped_ip]
[netmask mask] | interface}
```

For example, the following commands translate the 10.1.1.0/24 network on the inside interface.

```
VPN_ASA(config)# nat (inside) 1 10.1.1.0 255.255.255.0
VPN_ASA(config)# global (outside) 1 209.165.201.1-209.165.201.30
```

Basic PAT

PAT translates multiple real addresses to a single mapped IP address. Specifically, the security appliance translates the real address and source port (real socket) to the mapped address and a unique port above 1024 (mapped socket). Each connection requires a separate translation, because the source port differs for each connection.

To identify addresses on one interface that are translated to a mapped address and port on another interface, use the nat command in global configuration mode. This command configures dynamic NAT or PAT, where an address is translated to the mapped address.

For example, the following commands identify a pool of addresses for dynamic NAT as well as a PAT address for when the NAT pool is exhausted.

```
VPN_ASA(config)# nat (inside) 1 10.1.1.0 255.255.255.0
VPN_ASA(config)# global (outside) 1 209.165.201.5
VPN_ASA(config)# global (outside) 1 209.165.201.10-209.165.201.20
```

Translation from Lower Security Level to Higher

To translate the lower security DMZ network addresses so they appear to be on the same network as the inside network, which can simplify routing, enter the following command.

```
VPN_ASA(config)# nat (real_interface) nat_id access-list acl_name
[dns] [outside] [norandomseq] [[tcp] tcp_max_conns [emb_limit]] [udp
udp_max_conns]
```

NOTE: The outside argument has been removed starting from 8.0 software release.

For example:

```
VPN_ASA(config)# nat (dmz) 1 10.1.2.0 255.255.255.0 outside
VPN_ASA(config)# global (inside) 1 10.1.1.45
```

Static NAT

Static NAT creates a fixed translation of real address(es) to mapped address(es).With dynamic NAT and PAT, each host uses a different address or port for each subsequent translation. Because the mapped address is the same for each consecutive connection with static NAT, and a persistent translation rule exists, static NAT allows hosts on the destination network to initiate traffic to a translated host (if there is an access list that allows it).

For policy static NAT, enter the following command.

```
VPN_ASA(config)# static (real_interface, mapped_interface) {mapped_
ip | interface} access-list acl_name [dns] [norandomseq] [[tcp] tcp_
max_conns [emb_limit]] [udp udp_max_conns]
```

For example, the following command maps an inside IP address (10.1.1.3) to an outside IP address (209.165.201.12).

```
VPN_ASA(config)# static (inside,outside) 209.165.201.12 10.1.1.3
netmask 255.255.255.255
```

The following command maps the outside address (209.165.201.15) to an inside address (10.1.1.6).

```
VPN_ASA(config)# static (outside,inside) 10.1.1.6 209.165.201.15
netmask 255.255.255.255
```

The following command statically maps an entire subnet.

```
VPN_ASA(config)# static (inside,dmz) 10.1.1.0 10.1.2.0 netmask
255.255.255.0
```

Static PAT

Static PAT lets you translate the real IP address to a mapped IP address, as well as the real port to a mapped port. You can choose to translate the real port to the same port, which lets you translate only specific types of traffic, or you can take it further by translating to a different port.

For policy static PAT, enter the following command.

```
VPN_ASA(config)# static(real_interface,mapped_interface) {tcp | udp}
{mapped_ip | interface} mapped_port access-list acl_name [dns]
[norandomseq] [[tcp] tcp_max_conns [emb_limit]] [udp udp_max_conns]
```

For example, the following commands redirect HTTP traffic initiated from hosts on the 10.1.3.0 network from the security appliance outside interface (10.1.2.14) to the inside host at 10.1.1.15.

```
VPN_ASA(config)# access-list HTTP permit tcp host 10.1.1.15 eq http
10.1.3.0 255.255.255.0 eq http
VPN_ASA(config)# static (inside,outside) tcp 10.1.2.14 http access-
list HTTP
```

Policy NAT

Policy NAT lets you identify real addresses for address translation by specifying the source and destination addresses in an extended access list. You can also optionally specify the source and destination ports.

```
ASA_name(config)# nat (real_ifc) nat_id access-list access_list_name
[dns] [outside] [tcp] tcp_max_conns [emb_limit]] [udp udp_max_conns]
[norandomseq]
```

To configure policy static identity NAT, enter the following command.

```
VPN_ASA(config)# static (real_interface,mapped_interface) real_ip
access-list acl_id [dns] [norandomseq] [[tcp] tcp_max_conns [emb_
limit]] [udp udp_max_conns]
```

Bypassing NAT When NAT Control is Enabled

The following are ways to bypass NAT when NAT control has been enabled.

Identity NAT/ NAT 0

Identity NAT translates the real IP address to the same IP address. Only "translated" hosts can create NAT translations, and responding traffic is allowed back. Use the following command to configure identity NAT.

```
VPN_ASA(config)# nat (real_interface) 0 real_ip [mask] [dns]
[outside] [norandomseq] [[tcp] tcp_max_conns [emb_limit]] [udp udp_
max_conns]
```

Static Identity NAT

Static identity NAT translates the real IP address to the same IP address. The translation is always active, and both "translated" and remote hosts can originate connections. Use the following command to configure regular static identity NAT.

```
VPN_ASA(config)# static (real_interface,mapped_interface) real_ip
real_ip [netmask mask] [dns] [norandomseq] [[tcp] tcp_max_conns
[emb_limit]] [udp udp_max_conns]
```

NAT Exceptions

```
VPN_ASA(config)#access-list no_nat extended permit
 ip host 192.168.50.5 192.168.25.0 255.255.255.0
VPN_ASA(config)#nat (inside) 0 access-list no_nat
```

Network Management

Overview

This section describes the configuration commands required in order to manage and monitor each device in the network infrastructure for the Cisco Connected Stadium network. It provides a single source of information for analyzing and troubleshooting device connectivity to the Network Management Systems. Also, it enables a better understanding of the changes required and the impact in the NMS systems of any device operating system or software change.

The NMS functional areas involved are:

- Fault Management
- Configuration/Compliance Management
- Performance Management

The NMS applications included in the solution are:

- Cisco Prime LAN Management Solution
- Cisco Prime Network Control Solution

The network devices to be managed and monitored directly are:

- Cisco Catalyst and Nexus series switches

Logging

If a logging server (e.g. Cisco Monitoring, Analysis and Response System) is available, below are the recommendations for configuring the Cisco switches for logging.

Access (3750 12.2(44)SE4)

This section describes the commands required to configure system message logging on the both Catalyst and Nexus switches.

Enabling Sequence Numbers in Log Messages

Because there is a chance that more than one log message can have the same time stamp, you can display messages with sequence numbers so that you can unambiguously see a single message.

```
service sequence-numbers
```

Enabling Time Stamps on Log Messages

Enabling time stamps on log messages is important for troubleshooting. The time stamp includes the date, time in milliseconds relative to the local time-zone, and the time zone name.

```
service timestamps log datetime msec localtime show-timezone
```

Disabling Console Logging

Logging messages to the console are disabled. This can be enabled later for troubleshooting.

```
no logging console
```

Logging Messages to an Internal Buffer

This command enables logging messages to an internal buffer on the on the stack master switch. It also defines the size of the buffer and the level of messages to be stored.

```
logging buffered 32768 debugging
```

Storing Log Messages in a File

If the stack master switch fails, the log file is lost unless you previously saved it to flash memory. This command stores log messages in a file in flash memory on the stack master.

```
logging file flash:logging 32768 debugging
```

Defining the Message Severity Level

Log messages sent to the NMS server are limited by specifying the severity level to informational (level 6). This is also the default configuration.

```
logging trap informational
```

Limiting Syslog Messages Sent to the History Table and to SNMP

The level of messages sent and stored in the switch history table is set to informational (level 6) and the amount of events stored is set to 500.

```
logging history informational
logging history size 500
```

Enabling the Configuration-Change Logger

The following commands enable tracking of configuration changes made with the command-line interface (CLI). Also, the amount of configuration changes entries retained is set to 500. This configuration is done under the configuration change logger configuration mode found under the archive configuration mode.

```
(config)# archive
(config-archive)# log config
(config-archive-log-cfg)# logging enable
(config-archive-log-cfg)# logging size 500
```

Specifying the Source of Messages

All syslog messages are sourced by the Loopback 0 interface in all devices.

```
logging source-interface Loopback 0
```

Setting the Message Display Destination Device

```
logging <IP address of logging system>
```

Core (Nexus 7000 NxOS v4.1)

This section describes the commands required to configure system message logging on the “Core MDF Switching” Nexus 7000.

Enabling Time Stamps on Log Messages

Enabling time stamps on log messages is important for troubleshooting. The time stamp includes the date, time in milliseconds relative to the local time-zone, and the time zone name.

```
logging timestamps milliseconds
```

Disabling Console Logging

Logging messages to the console are disabled. This can be enabled later for troubleshooting.

```
no logging console
```

Logging Messages to a Device File, Set the Severity

This command enables logging messages to an internal file and defines the logging level.

```
logging logfile logging 7 size 32768
```

Setting the Message Display Destination Device

```
logging <IP address of logging system>
```

Core (6500 12.2(33)SHX4S)

This section describes the commands required to configure system message logging on the Cross Connect and Server Rack Switches.

Enabling Sequence Numbers in Log Messages

Because there is a chance that more than one log message can have the same time stamp, you can display messages with sequence numbers so that you can unambiguously see a single message.

```
service sequence-numbers
```

Enabling Time Stamps on Log Messages

Enabling time stamps on log messages is important for troubleshooting. The time stamp includes the date, time in milliseconds relative to the local time-zone, and the time zone name.

```
service timestamps log datetime msec localtime show-timezone
```

Disabling Console Logging

Logging messages to the console are disabled. This can be enabled later for troubleshooting.

```
no logging console
```

Logging Messages to an Internal Buffer

This command enables logging messages to an internal buffer on the on the stack master switch. It also defines the size of the buffer and the level of messages to be stored.

```
logging buffered 32768 debugging
```

Specifying the Source of Messages

All syslog messages are sourced by the Loopback 0 interface in all devices.

```
logging source-interface Loopback 0
```

Defining the Message Severity Level

Log messages sent to the NMS server are limited by specifying the severity level to informational (level 6). This is also the default configuration.

```
logging trap informational
```

Limiting Syslog Messages Sent to the History Table and to SNMP

The level of messages sent and stored in the switch history table is set to informational (level 6) and the amount of events stored is set to 500.

```
logging history informational
logging history size 500
```

Enabling the Configuration-Change Logger

The following commands enable tracking of configuration changes made with the CLI. Also, the amount of configuration changes entries retained is set to 500. This configuration is done under the configuration change logger configuration mode found under the archive configuration mode.

```
(config)# archive
(config-archive)# log config
(config-archive-log-cfg)# logging enable
(config-archive-log-cfg)# logging size 500
```

Setting the Message Display Destination Device

```
logging <IP address of logging system>
```

SNMP Configuration

Access (3750 12.2(44)SE4)

This section describes the commands required to configure SNMP on Catalyst and Nexus switches.

Specifying the SNMP Trap Source Interface

The following command specifies the source interface, which provides the IP address for the trap message.

```
snmp-server trap-source Loopback0
```

Setting the Agent Contact and Location Information

The following commands set the SNMP agent contact and location information, making it available for NMS applications polling information from devices.

```
snmp-server location DEVICE-LOCATION
snmp-server contact CONTACT-INFORMATION
```

Interface Identification Persistence for SNMP

The following command enables SNMP interface index values that remain constant across reboots only on a specific interface.

```
snmp ifmib ifindex persist
```

CPU Thresholding Notification

The following command sets CPU thresholding notification types and values; so when CPU resources rise to 90% of utilization or falls below 80% during a 60 second interval, a notification gets triggered.

```
process cpu threshold type process rising 90 interval 60 falling 80
interval 60
```

Define the Community Strings

```
snmp-server community public ro 10
snmp-server community private rw 10
```

Specifying SNMP Trap Recipient

The following commands specify the recipients of the SNMP traps, using version 2c and a public community string defined by the stadium. It is highly recommended not to use ‘public’ as the public community string. (Cisco LMS Server: 10.x.246.21)

```
snmp-server host 10.x.246.21 version 2c public
```

Specifying SNMP Agent Traps Notifications

The following commands globally enable the mechanism for the specified trap notification.

```
snmp-server enable traps snmp authentication linkdown linkup
coldstart warmstart
snmp-server enable traps transceiver all
snmp-server enable traps tty
snmp-server enable traps eigrp
snmp-server enable traps license
snmp-server enable traps cluster
snmp-server enable traps fru-ctrl
```

```
snmp-server enable traps entity
snmp-server enable traps cpu threshold
snmp-server enable traps power-ethernet group 1-9
snmp-server enable traps power-ethernet police
snmp-server enable traps vtp
snmp-server enable traps vlancreate
snmp-server enable traps vlandelete
snmp-server enable traps flash insertion removal
snmp-server enable traps port-security
snmp-server enable traps envmon fan shutdown supply temperature
status
snmp-server enable traps stackwise
snmp-server enable traps cef resource-failure peer-state-change
peer-fib-state-change inconsistency
snmp-server enable traps config-copy
snmp-server enable traps config
snmp-server enable traps config-ctid
snmp-server enable traps bridge newroot topologychange
snmp-server enable traps stpx inconsistency root-inconsistency loop-
inconsistency
snmp-server enable traps syslog
snmp-server enable traps mac-notification change move threshold
snmp-server enable traps vlan-membership
snmp-server enable traps errdisable
```

Enabling MAC Address Table Notifications

The following commands enable notifications traps via SNMP or Syslog regarding MAC address table changes, moves and threshold, as follows:

Changes

- When enabling MAC address table change notification, traps gets generated for dynamic changes to the MAC address table.
- This feature is not enable by default, the commands included here allows notifications to be sent if this feature is enabled.

Moves

- MAC-move notification generates a syslog message whenever a MAC address or host moves between different switch ports.
- This feature is not enabled by default, the commands included here allows notifications to be sent if this feature is enabled.

Threshold

- When enabling content-addressable memory CAM table usage monitoring, the number of valid entries in the CAM table are counted and if the percentage of the CAM utilization is higher or equal to the specified threshold, a notification is sent. The percentage limit for CAM utilization is 50 percent by default and notifications are sent every 120 seconds.

```
mac address-table notification change interval 60
mac address-table notification change history-size 100
mac address-table notification change
mac address-table notification mac-move
mac address-table notification threshold
```

Core (Nexus 7000 NxOS v4.1)

This section describes the commands required to configure SNMP on the “Core MDF Switching” Nexus 7000.

Setting the Agent Contact and Location Information

The following commands set the SNMP agent contact and location information, making it available for NMS applications polling information from devices.

```
Snmp-server location DEVICE-LOCATION
Snmp-server contact CONTACT-INFORMATION
```

Enable the SNMP Agent

```
snmp-server protocol enable
```

Define the Community Strings

```
snmp-server community public ro
snmp-server community private rw
```

In order to restrict read-only (RO) SNMP access, additional entries must be configured in the nms-access-acl. Applying access-lists directly to the snmp-server community command is not supported.

Specifying SNMP Trap Recipient

The following commands specify the recipients of the SNMP traps, using version 2c and a public community string defined by the stadium. It is highly recommended not to use 'public' as the public community string. (Cisco LMS Server: 10.x.246.21)

```
snmp-server host 10.x.246.21 version 2c public
```

Define the traps to be sent

```
snmp-server enable traps license
snmp-server enable traps entity fru
snmp-server enable traps link
snmp-server enable traps aaa server-state-change
snmp-server enable traps snmp authentication
```

Core (6500 12.2(33)SHX4S)

This section describes the commands required to configure SNMP on the Cross Connect and Server Rack Switches.

Specifying the SNMP Trap Source Interface

The following command specifies the source interface, which provides the IP address for the trap message.

```
snmp-server trap-source Loopback0
```

Setting the Agent Contact and Location Information

The following commands set the SNMP agent contact and location information, making it available for NMS applications polling information from devices.

```
snmp-server location DEVICE-LOCATION
snmp-server contact CONTACT-INFORMATION
```

Interface Identification Persistence for SNMP

The following command enables SNMP interface index values that remain constant across reboots only on a specific interface.

```
snmp ifmib ifindex persist
```

CPU Thresholding Notification

The following command sets CPU thresholding notification types and values; so when CPU resources rise to 90% of utilization or falls below 80% during a 60 second interval, a notification gets triggered.

```
process cpu threshold type process rising 90 interval 60 falling 80
interval 60
```

Define the Community Strings

```
snmp-server community public ro 10
snmp-server community private rw 10
```

Specifying SNMP Trap Recipient

The following commands specify the recipients of the SNMP traps, using version 2c and a public community string defined by the stadium. It is highly recommended not to use ‘public’ as the public community string. (Cisco LMS Server: 10.x.246.21)

```
snmp-server host 10.x.246.21 version 2c public
```

Define the Traps to be Sent

```
snmp-server enable traps snmp authentication linkdown linkup
coldstart warmstart
snmp-server enable traps chassis
snmp-server enable traps module
snmp-server enable traps transceiver all
snmp-server enable traps casa
snmp-server enable traps tty
snmp-server enable traps ipmulticast
snmp-server enable traps MAC-Notification change move threshold
snmp-server enable traps pim neighbor-change rp-mapping-change
invalid-pim-message
```

```
snmp-server enable traps bridge newroot topologychange
snmp-server enable traps stpx inconsistency root-inconsistency loop-
inconsistency
snmp-server enable traps syslog
snmp-server enable traps memory bufferpeak
snmp-server enable traps fru-ctrl ps-output-change status
snmp-server enable traps entity
snmp-server enable traps module-auto-shutdown sys-action
snmp-server enable traps power-ethernet group 1-48
snmp-server enable traps cpu threshold
snmp-server enable traps vlancreate
snmp-server enable traps vlandelete
snmp-server enable traps flash insertion removal
snmp-server enable traps c6kxbar intbus-crcexcd intbus-crcrcvrd
swbus
snmp-server enable traps dot1x no-guest-vlan no-auth-fail-vlan
snmp-server enable traps envmon fan shutdown supply temperature
status
snmp-server enable traps port-security
snmp-server enable traps alarms
snmp-server enable traps vlan-mac-limit
```

NOTE: There is no “enable traps” command for EIGRP protocol.

Enabling MAC Address Table Notifications

The following commands enable notifications traps via SNMP or Syslog regarding MAC address table changes and moves, as follows:

Changes

- When enabling MAC address table change notification, traps gets generated for dynamic changes to the MAC address table. The change history is set to 100 entries and notifications are sent every 60 seconds.
- This feature is not enable by default, the commands included here allows notifications to be sent if this feature is enabled.

Moves

- MAC-move notification generates a syslog message whenever a MAC address or host moves between different switch ports.
- This feature is not enabled by default, the commands included here allows notifications to be sent if this feature is enabled.

```
mac address-table notification change interval 60
mac address-table notification change history-size 100
mac address-table notification change
mac address-table notification mac-move
```

NMS Access Control Lists

The NMS tools require various access methods to the devices. NMS access is controlled via the following Access Control Lists:

- SNMP NMS Servers ACL (standard ACL 10): This ACL is applied to the **snmp-server community** command and restricts the host that is able to poll information from the devices via snmp.
- (Optional) Additional Access NMS Servers ACL (nms-access-acl): This access list is applied to the device interfaces in order to allow and limit the NMS servers access via other protocols, such as:
  - SSH
  - Telnet
  - FTP
  - TFTP
  - ICMP – Traceroute & Ping
- The NMS servers IP addresses and applications are:
  - LMS Server: 10.x.246.21
  - NCS Server: 10.x.246.22

The following commands create the required access control lists:

```
access-list 10 permit host 10.x.246.21
access-list 10 permit host 10.x.246.22
ip access-list extended nms-access-acl
permit tcp host 10.x.246.21 any eq ftp-
data
 permit tcp host 10.x.246.21 any eq ftp
 permit tcp host 10.x.246.21 eq ftp any
 permit tcp host 10.x.246.21 eq ftp-data
 any
 permit tcp host 10.x.246.21 any eq
 telnet
 permit tcp host 10.x.246.21 any eq 107
 permit tcp host 10.x.246.21 eq 107 any
 permit tcp host 10.x.246.21 eq telnet
 any
 permit udp host 10.x.246.21 any eq tftp
 permit udp host 10.x.246.21 eq tftp any
 permit udp host 10.x.246.21 eq 1758 any
 permit udp host 10.x.246.21 any eq 1758
 permit tcp host 10.x.246.21 any eq 22
 permit tcp host 10.x.246.21 eq 22 any
 permit icmp host 10.x.246.21 any ttl-
 exceeded
 permit icmp host 10.x.246.21 any port-
 unreachable
 permit icmp host 10.x.246.21 any echo
 permit icmp host 10.x.246.21 any echo-
 reply
```

```
permit tcp host 10.x.246.22 any eq ftp-
data
 permit tcp host 10.x.246.22 any eq ftp
 permit tcp host 10.x.246.22 eq ftp any
 permit tcp host 10.x.246.22 eq ftp-data
 any
 permit tcp host 10.x.246.22 any eq
 telnet
 permit tcp host 10.x.246.22 any eq 107
 permit tcp host 10.x.246.22 eq 107 any
 permit tcp host 10.x.246.22 eq telnet
 any
 permit udp host 10.x.246.22 any eq tftp
 permit udp host 10.x.246.22 eq tftp any
 permit udp host 10.x.246.22 eq 1758 any
 permit udp host 10.x.246.22 any eq 1758
 permit tcp host 10.x.246.22 any eq 22
 permit tcp host 10.x.246.22 eq 22 any
 permit icmp host 10.x.246.22 any ttl-
 exceeded
 permit icmp host 10.x.246.22 any port-
 unreachable
 permit icmp host 10.x.246.22 any echo
 permit icmp host 10.x.246.22 any echo-
 reply
```

Prime Infrastructure Management

The Cisco Prime Infrastructure provides management of the WLC and their access points, as well as spectrum analysis to allow further tuning of the wireless infrastructure after initial deployment.

Prime Infrastructure is a network management tool to manage multiple controllers as a network of controllers. Prime Infrastructure includes the same configuration, performance monitoring, security, fault management, and accounting options used at the controller level and adds a graphical view of multiple controllers and managed access points. This tool provides the following capabilities.

- Devices management – Device management is the core task of managing Cisco controllers, APs and Wireless devices. Prime Infrastructure provides the ability to discover and manage devices, inventory devices, and organize devices into easily manageable groups.

- Fault monitoring – Prime Infrastructure interrogates managed devices for specific data and alerts the administrator via an SNMP trap and/or a syslog message when irregularities or faults are discovered. The data interrogated by the fault monitoring feature includes RF, rogue APs, security configuration parameters (security policy monitoring) and SNMP parameters used for performance related faults.
- Simplified controller and AP configuration via configuration templates – Prime Infrastructure enables the creation of configuration templates that can be applied to a group of controllers and APs that have unique set of configuration parameters.
- Manage controller and AP images – Prime Infrastructure can keep a library of controller and AP images. Prime Infrastructure can push out controller and AP upgrades and report status of upgrades.
- Wireless LAN Planning and Design –Prime Infrastructure provides integrated RF prediction tools that can be used to create a detailed wireless LAN design, including AP placement, configuration, and performance and coverage estimates. Floor plans can be imported into Prime Infrastructure and assign RF characteristics to building components to increase design accuracy. Graphical heat maps help visualize anticipated wireless LAN behavior for easier planning and faster rollout.
- Manage Device Reports – Prime Infrastructure maintains historical data and performance statistics of Cisco Wireless devices.

Due to the scale and complexity of deploying Wi-Fi in sports and entertainment properties, Prime Infrastructure is required to effectively manage the Cisco Connected Stadium Wi-Fi Solution.

Figure 34. Cisco Prime Infrastructure

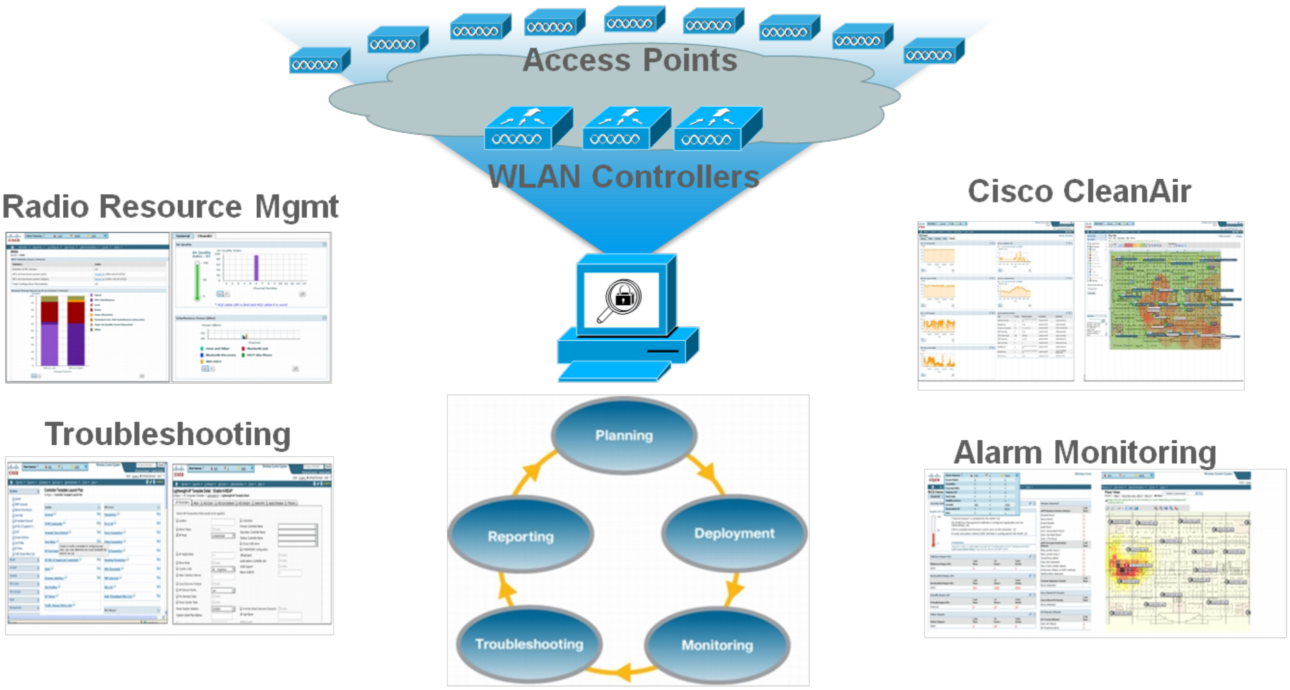

Table 10. Recommended Prime Infrastructure Version[1]

Description	Version
Cisco Prime Infrastructure	1.2 or later[2]

To learn more about the Cisco Prime Infrastructure, visit: http://www.cisco.com/en/US/products/ps6305/index.html

NOTE: For location support, the Cisco Mobility Service Engine must be used. For more information, visit: http://www.cisco.com/en/US/products/ps9742/index.html

[1] Use the Prime Infrastructure version required to support the specific equipment used. If unsure, contact your Cisco account team.

[2] If deploying a Connected Stadium Wi-Fi, see the Connected Stadium Wi-Fi design guide for the most up-to-date Cisco Prime Infrastructure version recommendations.

Cisco Solution Integration

Cisco StadiumVision

This section of the design document provides a brief overview of the elements of the Cisco StadiumVision solution.

Overview

Figure 35. Cisco StadiumVision Overview

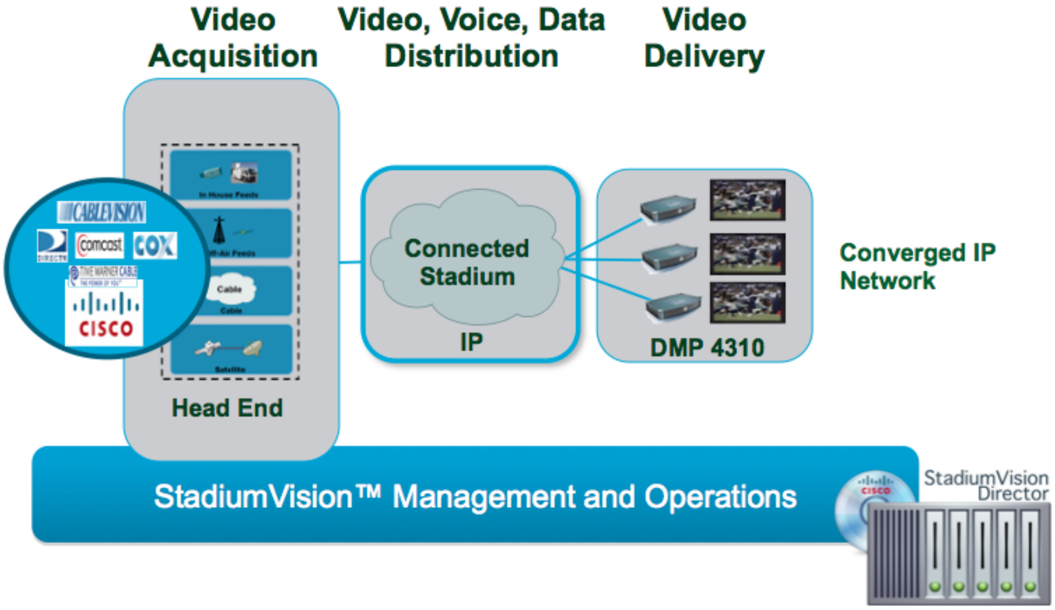

The Cisco StadiumVision solution includes a combination of video and digital signage content that is displayed on TVs throughout the stadium.

The digital signage content is pre-positioned on the Digital Media Players (DMPs) prior to an event. The only impact to the network is when the content is being pre-loaded on the DMPs. This is a non-real time event in which the task is completed well in advance of the content being displayed.

Full motion video, on the other hand, has more of an impact on the Cisco Connected Stadium network as this content is live streaming video that is distributed to TVs in real time. The streaming of the video occurs during an event and the network needs to accommodate for this accordingly.

Figure 36 depicts the architecture of the Cisco StadiumVision solution with the video head end equipment, including video distribution switches and Digital Content Managers (DCMs), located in the video room and the video endpoints, Digital Media Players (DMPs) and TVs, connected to the access IDF switches throughout the stadium.

Figure 36. Cisco StadiumVision Architecture

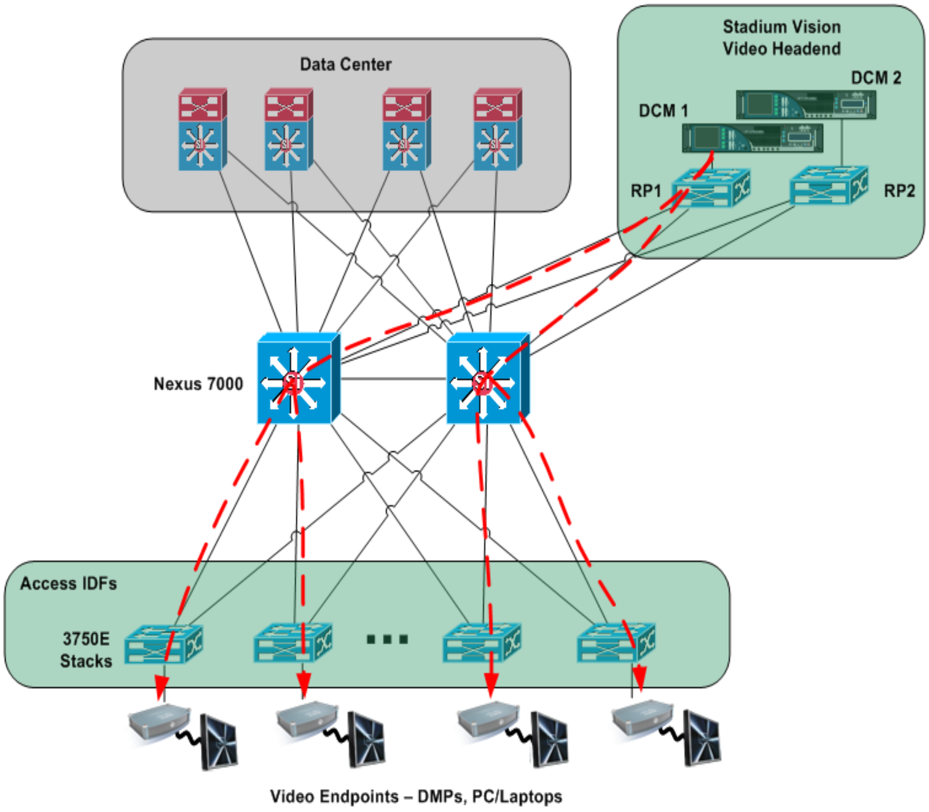

Video Redundancy

As illustrated in Figure 36, video redundancy is accomplished by having redundant DCMs streaming identical video channels. Prioritycast is used for controlling what video stream is propagated on the network and controlling the fail-over to the video streams provided by the standby DCM.

Video Head End

The primary source for distributing Cisco StadiumVision High Definition (HD) streams to the TVs located throughout the stadium is a pair of Digital Content Managers (DCMs). The DCMs are configured to distribute the TV channels made available for viewing within the stadium. They convert incoming video streams from various sources (DirecTV channels, Off-air channels, In-house cameras, etc) into IP multicast streams and distribute the streams to the TVs throughout the stadium. The Cisco StadiumVision video traffic traverses the Connected Stadium network to the DMPs, which are connected to the access IDF switches.

Each TV channel that is distributed across the network is a unique multicast stream and requires up to 20 Mbps of bandwidth depending on the compression standard used (i.e., MPEG2 or MPEG4). The Cisco StadiumVision solution is designed to support up to 250 IPTV channels. The majority of the Channels will be HD.

Video Delivery

Cisco StadiumVision supports up to 5000 TVs deployed within the stadium. Each of these TVs has an associated DMP that is used to control the content displayed on them.

As previously stated each IPTV channel is sent across the stadium network via IP multicast. The DMPs use IGMP to join the appropriate IP multicast group, as defined at the head end, to receive the video stream for the desired IPTV channel. The DMPs receive the multicast streams, decode them, and display them on the TV to which they are connected.

The DMPs are managed by a centralized application called Cisco StadiumVision Director. Through event scripts, Cisco StadiumVision Director can control which IP multicast channel each DMP receives. Cisco StadiumVision Director also allows for a Cisco IP phone, IR remote, or third-party touch panel to control which IPTV channel is displayed on a TV.

DMP Control

Cisco StadiumVision Director uses both unicast and multicast for controlling DMPs and distributing content. The figure below describes how StadiumVision Director uses the Anycast RPs to enable the DMPs to join the StadiumVision Directors multicast DMP control channel.

Figure 37. Cisco StadiumVision Director Controlling DMPs via Multicast

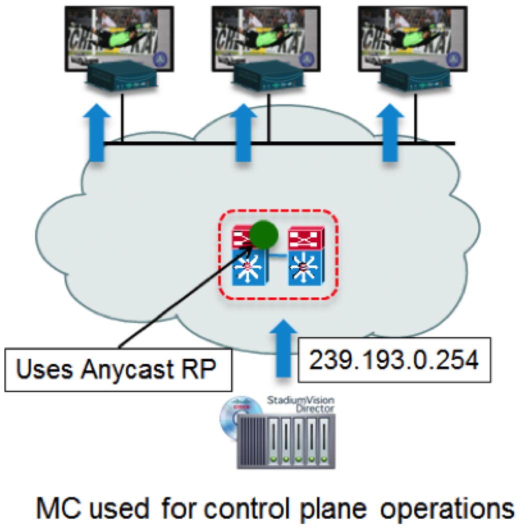

Integrating Cisco StadiumVision on the Connected Stadium Network

The Role CDP and LLDP Play in Cisco StadiumVision

Cisco 4310 Digital Media Player uses CDP

The DMP, starting with firmware version SE 2.2.1 build 1932, supports Medianet services, which includes Cisco Discovery Protocol (CDP). This capability allows the switch and DMP to learn about each other by exchanging CDP messages. At the switch CLI, the show cdp neighbor command displays the DMPs that are connected to the switch. This information can be very useful when troubleshooting DMP issues.

```
Training_SW#show cdp neighbor

Capability Codes: R - Router, T - Trans Bridge, B - Source Route
Bridge, S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone, D -
Remote, C - CVTA, M - Two-port Mac Relay
```

Device ID	Local Intrfce	Holdtme	Capability	Platform	Port ID
SEP081735D5E79A	Gig 0/1	173	H P M	IP Phone	Port 1
SEP081735D5E72F	Gig 0/9	174	H P M	IP Phone	Port 1
00:0f:44:01:b5:50	Gig 0/7	171	H	DMP 4310G	eth0
SEP081735D5E749	Gig 0/6	174	H P M	IP Phone	Port 1
sande-cucm	Gig 0/24	137	H	VMware Vi	eth0
SEP081735D5E6DF	Gig 0/8	175	H P M	IP Phone	Port 1
00:0f:44:01:a8:d7	Gig 0/10	175	H	DMP 4310G	eth0
00:0f:44:01:b6:c3	Gig 0/5	175	H	DMP 4310G	eth0
00:0f:44:01:bf:60	Gig 0/11	175	H	DMP 4310G	eth0

CDP information pertaining to the connected switch port can be viewed in the Cisco StadiumVision Director Management Dashboard.

Figure 38. CDP Information Display in Cisco StadiumVision Director

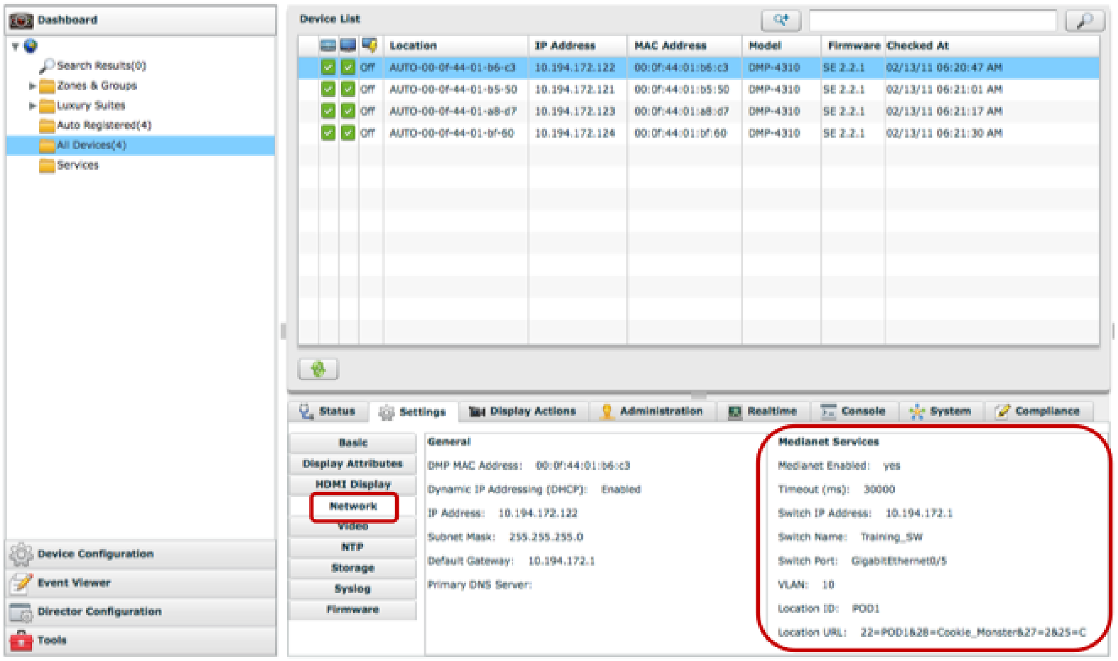

Device List

			Location	IP Address	MAC Address	Model	Firmware	Checked At
✓	✓	Off	AUTO-00-0f-44-01-b6-c3	10.194.172.122	00:0f:44:01:b6:c3	DMP-4310	SE 2.2.1	02/13/11 06:20:47 AM
✓	✓	Off	AUTO-00-0f-44-01-b5-50	10.194.172.121	00:0f:44:01:b5:50	DMP-4310	SE 2.2.1	02/13/11 06:21:01 AM
✓	✓	Off	AUTO-00-0f-44-01-a8-d7	10.194.172.123	00:0f:44:01:a8:d7	DMP-4310	SE 2.2.1	02/13/11 06:21:17 AM
✓	✓	Off	AUTO-00-0f-44-01-bf-60	10.194.172.124	00:0f:44:01:bf:60	DMP-4310	SE 2.2.1	02/13/11 06:21:30 AM

General

DMP MAC Address: 00:0f:44:01:b6:c3

Dynamic IP Addressing (DHCP): Enabled

IP Address: 10.194.172.122

Subnet Mask: 255.255.255.0

Default Gateway: 10.194.172.1

Primary DNS Server:

Medianet Services

Medianet Enabled: yes

Timeout (ms): 30000

Switch IP Address: 10.194.172.1

Switch Name: Training_SW

Switch Port: GigabitEthernet0/5

VLAN: 10

Location ID: POD1

Location URL: 22=POD1&28=Cookie_Monster&27=2&25=C

SV-4K and DMP-2K DMPs use LLDP

The SV-4K and DMP-2K DMP supports standard Link Layer Discovery Protocol (LLDP). This capability allows the switch and DMP to learn about each other by exchanging LLDP messages. At the switch CLI and similar to the show cdp neighbor command, the show lldp neighbor command displays the DMPs that are connected to the switch. This information can be very useful when troubleshooting DMP issues.

To configure LLDP on a Catalyst switch perform the following commands.

SUMMARY STEPS

1. configure terminal
2. lldp run
3. interface *interface-id*
4. lldp transmit
5. lldp receive
6. end
7. show lldp

Configuring Civic Location in the Switch

IOS civic location is a collection of labels that can be configured on each switch port, and then communicated to the DMP via CDP. One use case for civic location is jack ID, hence allowing the DMP to learn what Ethernet jack it is connected to. The DMP reports any civic location information it learns back to Cisco StadiumVision Director, where it can be retrieved and displayed by Cisco StadiumVision Director as shown above. It can also be viewed directly in the DMP’s web interface.

The following example shows the civic location being configured and applied to interface gigabit Ethernet 0/5:

```
Training_SW#configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.

Training_SW(config)#location civic-location identifier 1

Training_SW(config-civic)#? Type ? to see all the options of
configuring the location.

Training_SW(config-civic)#additional-location-information POD1
Training_SW(config-civic)#building C

Training_SW(config-civic)#floor 2

Training_SW(config-civic)#room Cookie_Monster

Training_SW(config-civic)#end


Training_SW#conf t

Enter configuration commands, one per line.  End with CNTL/Z.

Training_SW(config)#int g0/5

Training_SW(config-if)#location civic-location-id 1

Training_SW(config-if)#end
```

Figure 39. Civic Location Information Displayed in the DMP

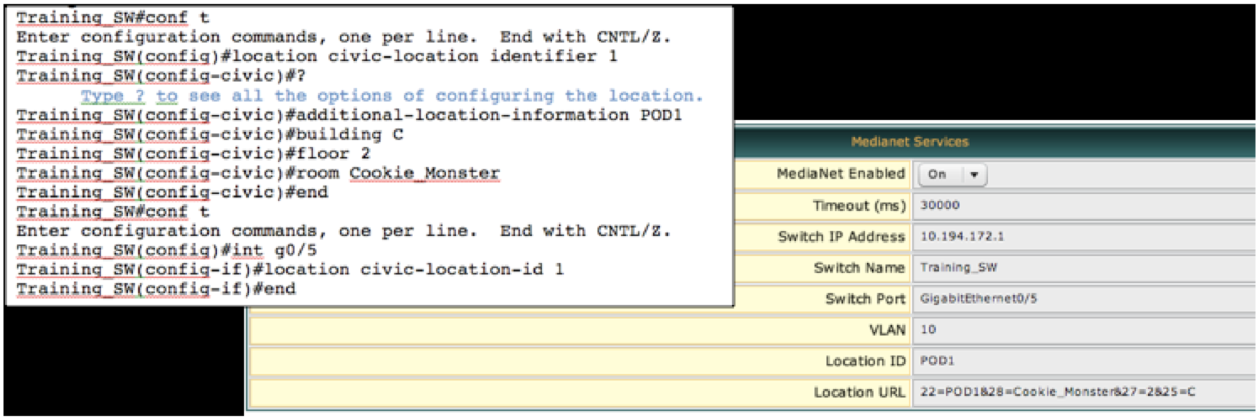

```
Training_SW#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Training_SW(config)#location civic-location identifier 1
Training_SW(config-civic)#?
      Type ? to see all the options of configuring the location.
Training_SW(config-civic)#additional-location-information POD1
Training_SW(config-civic)#building C
Training_SW(config-civic)#floor 2
Training_SW(config-civic)#room Cookie_Monster
Training_SW(config-civic)#end
Training_SW#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Training_SW(config)#int g0/5
Training_SW(config-if)#location civic-location-id 1
Training_SW(config-if)#end
```

Medianet Services

MediaNet Enabled	On ▾
Timeout (ms)	30000
Switch IP Address	10.194.172.1
Switch Name	Training_SW
Switch Port	GigabitEthernet0/5
VLAN	10
Location ID	POD1
Location URL	22=POD1&28=Cookie_Monster&27=2&25=C

NOTE: The civic location feature is only supported in 2xxx and 3xxx series switches, and only when they are running IOS 12.2.55SE1 or later. It is currently not supported in the 65xx switches.

IOS civic location information can be viewed in the Management Dashboard, but Cisco StadiumVision Director does not use the civic location information to automatically link the DMP to its target SV Location.

Configuring DHCP Options in the Switch and DHCP Server to support Cisco StadiumVision Director Auto-Registration

The Cisco Digital Media Player (DMP) receives firmware and configuration from StadiumVision Director. The DMP finds the StadiumVision Director server through the use of DHCP options (specifically options 60 and 43). Option 60 is used to signal the

SV-4K and DMP-2K media player that the option 43 content is meant for it to be received.

NOTE: The SV-4K DMP is introduced in Cisco StadiumVision Director Release 4.0. The DMP-2K is supported beginning in release 4.1.

- If you are supporting a deployment with mixed models, such as the Cisco DMP 4310G and SV-4K devices, you will need to configure an Option 43 string for each model.
- If the DHCP server is limited to a single Option 43 string per DHCP pool, then be sure to configure a separate VLAN & DHCP scope for the SV-4K media players.

As mentioned above, the DMP uses DHCP Options 60 and 43 to auto-register with Cisco StadiumVision Director.

To configure DHCP option 60, complete the following steps:

1. Configure the DHCP Option 60, Vendor Class Identifier string:
   - DMP-2K string for new, factory-shipped devices: "Cisco DMP-2K"
   - SV-4K string for North America: "Cisco SV-4K-NA"
   - SV-4K string for all other regions: "Cisco SV-4K-ROW"
2. Configure the converted DHCP Option 43, Vendor Specific Option URL:

IMPORTANT: The option 43 string must be converted to TLV format for compatibility with the SV-4K and DMP-2K.

The URL is:

http://*x.x.x.x*:8080/StadiumVision/dmp_v4/scripts/boot.brs

where “*x.x.x.x*” is the IP address of the Cisco StadiumVision Director server.

The encoded syntax in Hex:

55	Length (number of characters in the URL)	URL as shown above

For example,

- Option 43 (URL in ASCII):

  http://10.194.175.122:8080/StadiumVision/dmp_v4/scripts/boot.brs
- Option 43 (Hex encoded)

  To determine the length, count the number of ASCII characters in the URL and add 1 for the null terminator (00) then convert that decimal number to hex:

  5540.6874.7470.3a2f.2f31.302e.3139.342e.3137.352e.3132.323a.3830.3830.2f5
  3.7461.6469.756d.
  5669.7369.6f6e.2f64.6d70.5f76.342f.7363.7269.7074.732f.626f.6f74.2e62.7273

NOTE: For additional configuration information and guidelines, refer to the "Appendix A: Configuring an IOS DHCP Server to Support the SV-4K and DMP-2K" on page 159.

Configuring Options in Cisco Network Registrar

1. Log in to CNR, select **Advanced**, then **DHCP** as shown in Figure 40.

Figure 40. Cisco Network Registrar Interface

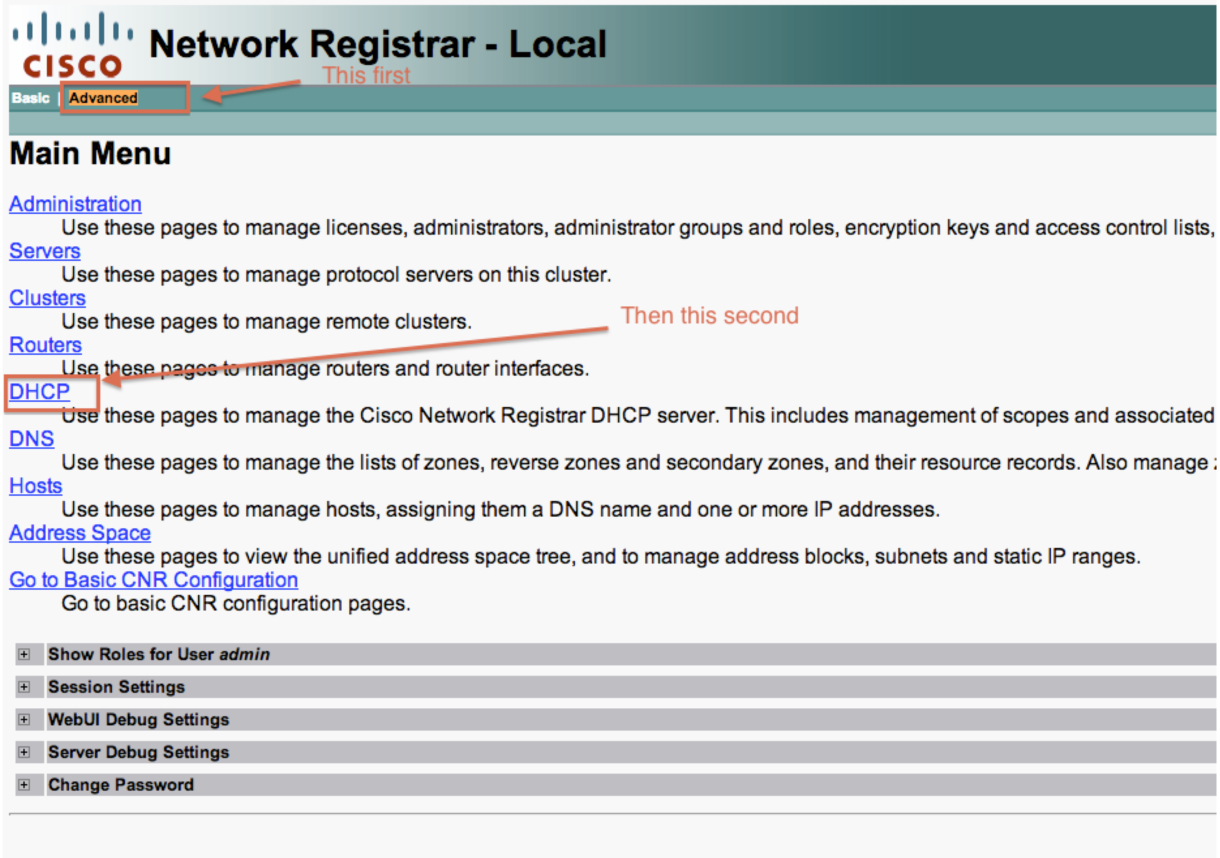

This first

Main Menu

Administration
Use these pages to manage licenses, administrators, administrator groups and roles, encryption keys and access control lists,

Servers
Use these pages to manage protocol servers on this cluster.

Clusters
Use these pages to manage remote clusters.

Then this second

Routers
Use these pages to manage routers and router interfaces.

DHCP
Use these pages to manage the Cisco Network Registrar DHCP server. This includes management of scopes and associated

DNS
Use these pages to manage the lists of zones, reverse zones and secondary zones, and their resource records. Also manage

Hosts
Use these pages to manage hosts, assigning them a DNS name and one or more IP addresses.

Address Space
Use these pages to view the unified address space tree, and to manage address blocks, subnets and static IP ranges.

Go to Basic CNR Configuration
Go to basic CNR configuration pages.

2. Select the **Options** settings, then **Add a New Option**.

Attribute	Value	Comment
Name	SV-4K Options	Can be any string.
DHCP Type	V4	
Description	options for BrightSign 4k1440	Can be any string.
Vendor Option String	Cisco SV-4K-NA	If DMP model is SV-4K in North America

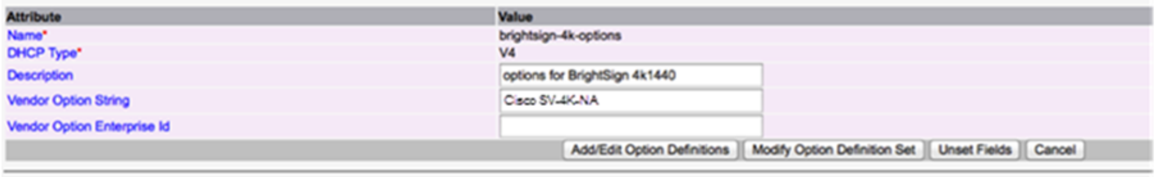

Attribute	Value
Name*	brightsign-4k-options
DHCP Type*	V4
Description	options for BrightSign 4k1440
Vendor Option String	Cisco SV-4K-NA
Vendor Option Enterprise Id	

NOTE: You won't be able to create multiple options for same Vendor Option String.

3. Click **Add Option Definition Set**, then specify the option 43 value type you need.

   c. Click the name you previously created, then **Add/Edit Option Definitions**.

   d. On next screen, click **Add Option Definitions**.

Number	43	
Name	brightsign-4k-option-43	could be anything
Description	option 43 items for brightsign 4k boxes	could be anything
type	binary	
repeat	[0]	

   e. Click **Add Option Definition**.

   f. Click **Modify Option Definition Set**.

NOTE: If you don't click Modify Option Definition Set here, it won't be saved.

4. Assign the new option to the right policies.

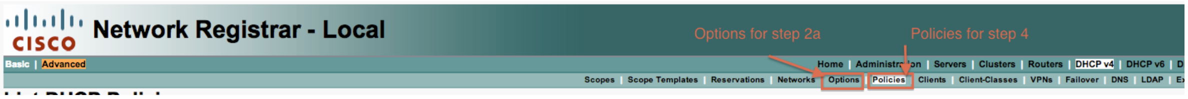
CISCO
Network Registrar - Local
Basic | Advanced
Home | Administration | Servers | Clusters | Routers | DHCP v4 | DHCP v6 | D
Scopes | Scope Templates | Reservations | Networks | Options | Policies | Clients | Client-Classes | VPNs | Failover | DNS | LDAP | Ex

Options for step 2a

Policies for step 4

   a. Click **Policies** in the upper menu (2nd level) and select a policy from the list.

   b. Under **DHCPv4 Vendor Options**, select the appropriate name you created in the steps above, then click **Select**.

   c. In the second dropdown, select the name you created.

   d. Paste the option 43 binary field computed above into the "Value" field.

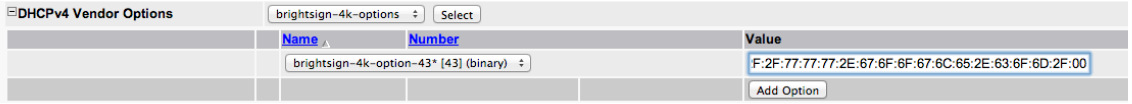

DHCPv4 Vendor Options: brightsign-4k-options

Name	Number	Value
brightsign-4k-option-43* [43] (binary)		F:2F:77:77:77:2E:67:6F:6F:67:6C:65:2E:63:6F:6D:2F:00

   e. Click **Add Option**.

   f. Click **Modify Policy** at the bottom of the screen.

NOTE: Do not forget to click **Modify Policy**. If you do not click it, your change will not be saved.

Precision Time Protocol (PTP) for SV-4K and DMP-2K DMP Synchronization

The IEEE 1588 Precision Time Protocol (PTP) is a configurable synchronization option in the SV-4K DMP to synchronize clocks among DMPS driving the display of time-critical content like for video walls.

Figure 41. SV-4K PTP Synchronization

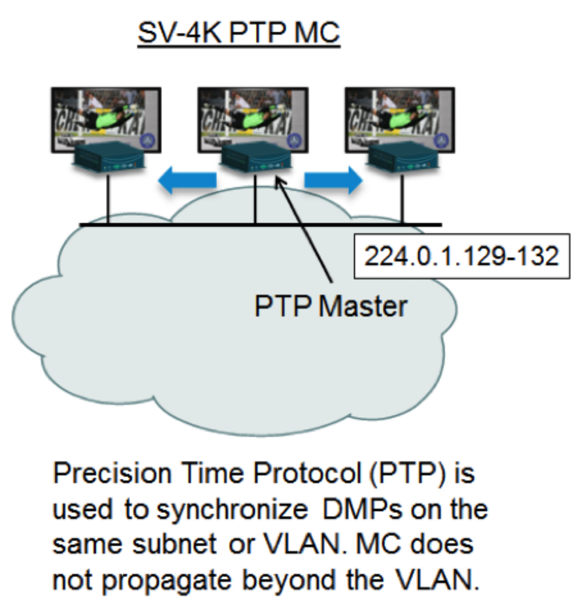

Precision Time Protocol (PTP) is used to synchronize DMPs on the same subnet or VLAN. MC does not propagate beyond the VLAN.

When using PTP, one DMP is designated as the domain master clock. It will synchronize with a Network Time Protocol (NTP) reference clock and then act as the reference point for a set (also known as a domain) of slave DMP clocks. The protocol provides the means for slave DMPs to determine the path delay incurred from the master to themselves. This time delay is then incorporated in the slave’s time to allow for highly precise time synchronization to the master DMP clock.

IEEE-1588 PTP uses multicast messages for communication with the following addresses:

224.0.1.107

224.0.1.129 – Default Domain 0

224.0.1.130 – Alternate Domain 1

224.0.1.131 – Alternate Domain 2

224.0.1.132 – Alternate Domain 3

NOTE: The SV-4K and DMP-2K use a TTL of 1 by default, meaning PTP multicast is confined to the local subnet or VLAN. Cisco StadiumVision Director Release 4.1 allows the TTL to be configured to a value greater than 1. Careful consideration should be used when configuring TTL > 1 to traverse multiple hops due to the increased latency incurred. This may negatively effect synchronization, and is not recommended for video walls. In addition, use of TTL >1 requires anycast rp multicast configuration rather than prioritycast.

Cisco StadiumVision Server Network Configuration

The Cisco StadiumVision Director servers control the content disposition and displays throughout the venue. The Cisco StadiumVision Director servers are configured in a active/standby redundancy model. Both servers must reside in the same VLAN and optimally connecting to their own switch as shown in the diagram below. HSRP should be configured to provide default gateway redundancy. Cisco StadiumVision Director servers would typically be installed in the Data Center.

NOTE: Although connecting servers to the Core switches is not typically recommended. There are instances where this maybe done when Data Center switches are not used in the network. The main requirement is having the Layer 2 connection between the two switches where SV Director's are connected.

In addition, SV Directors may be connected to a VDS switch but this allows the VDS switch as a single point of failure and is not a recommended connection scenario.

Configuration information can be found at the following URL:

http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750x_3560x/software/release/12-2_55_se/configuration/guide/3750xscg/swhsrp.html

Figure 42. Cisco StadiumVision Server Configuration Overview

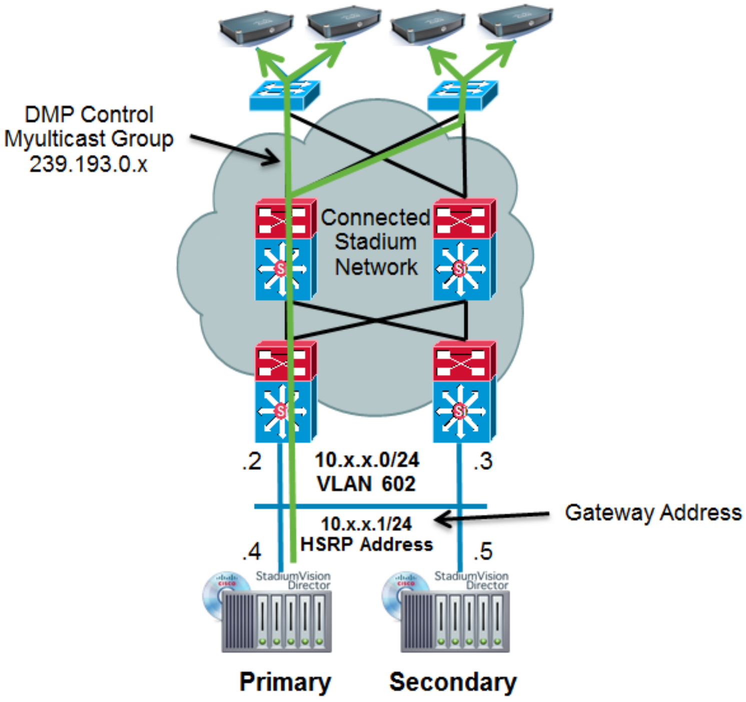

Cisco StadiumVision Video Injection Protection

Cisco StadiumVision utilizes multicast distributed video streams between the headend's DCMs and the DMPs. To prevent rogue devices attached at the access layer or elsewhere within the network from injecting their own video streams, suitable access lists are created on the Catalyst 3750 access switches.

Switch Access Control Lists to Prevent Video Injection

In order to prevent video injection from all access ports, the following ACL should be applied on all interfaces

```
ip access-list extended dmp-ipmc-block
 permit ip 10.x.0.0 0.0.255.255 224.0.0.0 0.0.0.255 -> Access for
the DMP VLANs
 permit igmp 10.x.0.0 0.0.255.255 any -> Needed this line as without
it the IGMP stream would timeout within 2 mins
 deny   ip any 224.0.0.0 15.255.255.255
 permit ip any any
```

And then the ACL can be applied to all switch ports on access-layer Catalyst 3750s:

```
interface GigabitEthernet1/0/1
 switchport access vlan 102
 switchport mode access
 ip access-group dmp-ipmc-block in
```

Additionally, the following ACL should be applied to all Catalyst 3750 uplinks:

```
ip access-list extended uplink-ipmc
 permit ip any 224.0.0.0 0.0.0.255 -> Allow link layer multicast
 permit ip any 239.193.0.0 0.0.255.255 -> Allow the Radiant,
StadiumVision Control, and WLC-to-AP MC traffic
 permit udp 172.x.x.0 0.0.0.255 239.192.0.0 0.0.255.255 -> Allow the
DCM subnets
 deny   ip any 224.0.0.0 15.255.255.255 -> Deny all other multicast
traffic
 permit ip any any
```

The implementation of these ACLs will prohibit anything plugged in at the access layer from injecting a video stream as the only authorized locations to multicast video are from secured locations (i.e. the datacenter).

Cisco StadiumVision Multiple Venue Support

Beginning in Cisco StadiumVision Director Release 3.1, a centralized server site can be deployed with multiple remote sites in a multi-venue architecture. This section describes how to configure the Connected Stadium WAN to accommodate this multi-venue architecture.

See "Configuring Cisco StadiumVision Director for Multiple Venue Support" in the *Cisco StadiumVision Director Server Administration Guide* for more information.

Figure 43 shows the multicast communication between the Cisco StadiumVision Director (SVD) and SVD Remote. Multicast is used to more efficiently send common messaging between the central SVD and the remote SVD servers. Unicast is used for remote SVD-specific messaging. Notice that all remote servers use the same multicast group at their respective venues (239.193.1.1:7778 by default). Therefore, this multicast group must be site-localized and not propagated across the WAN.

> **NOTE:** If your network configuration requires you to change the default multicast addressing used for Cisco StadiumVision Director Remote server deployment, see "Configuring Cisco StadiumVision Director for Multiple Venue Support" in the *Cisco StadiumVision Director Server Administration Guide* for configuration instructions.

NOTE: For the in-venue multicast LAN configuration, follow the multicast design recommendations described in the previous chapters of this document.

Figure 43. SVD to SVD Remote Multicast Communication

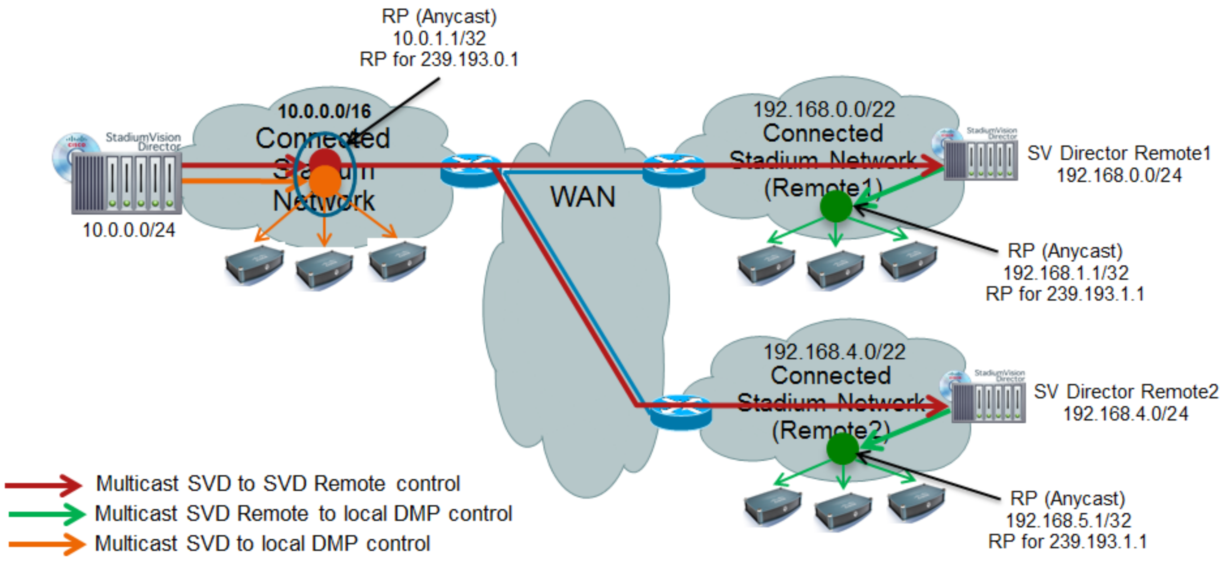

SVD Remote local Multicast (shown in green) should not be propagated outside the site.

Figure 44 shows the commands required to enable multicast over a point-to-point WAN link and the additional commands required in the switches at the remote venue.

Figure 44. SVD Multicast Routing over Point-to-Point WAN link

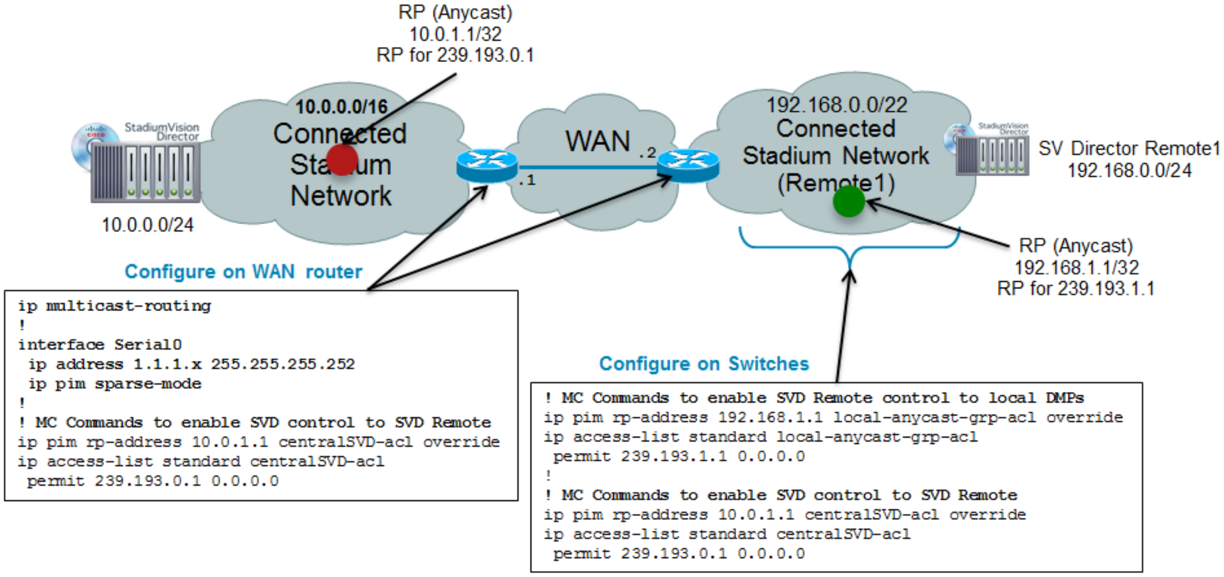

```
ip multicast-routing
!
interface Serial0
 ip address 1.1.1.x 255.255.255.252
 ip pim sparse-mode
!
! MC Commands to enable SVD control to SVD Remote
ip pim rp-address 10.0.1.1 centralSVD-acl override
ip access-list standard centralSVD-acl
 permit 239.193.0.1 0.0.0.0
```

```
! MC Commands to enable SVD Remote control to local DMPs
ip pim rp-address 192.168.1.1 local-anycast-grp-acl override
ip access-list standard local-anycast-grp-acl
 permit 239.193.1.1 0.0.0.0
!
! MC Commands to enable SVD control to SVD Remote
ip pim rp-address 10.0.1.1 centralSVD-acl override
ip access-list standard centralSVD-acl
 permit 239.193.0.1 0.0.0.0
```

Unicast and multicast routing must be configured within each location (i.e., Central and Remote sites)

Figure 45 shows the specific commands to enable unicast routing between the central and remote SVD servers and between the remote venue and the central PIM Rendevous Point (RP) that enables multicast routing.

Figure 45. SVD Unicast Routing over Point-to-Point WAN Link

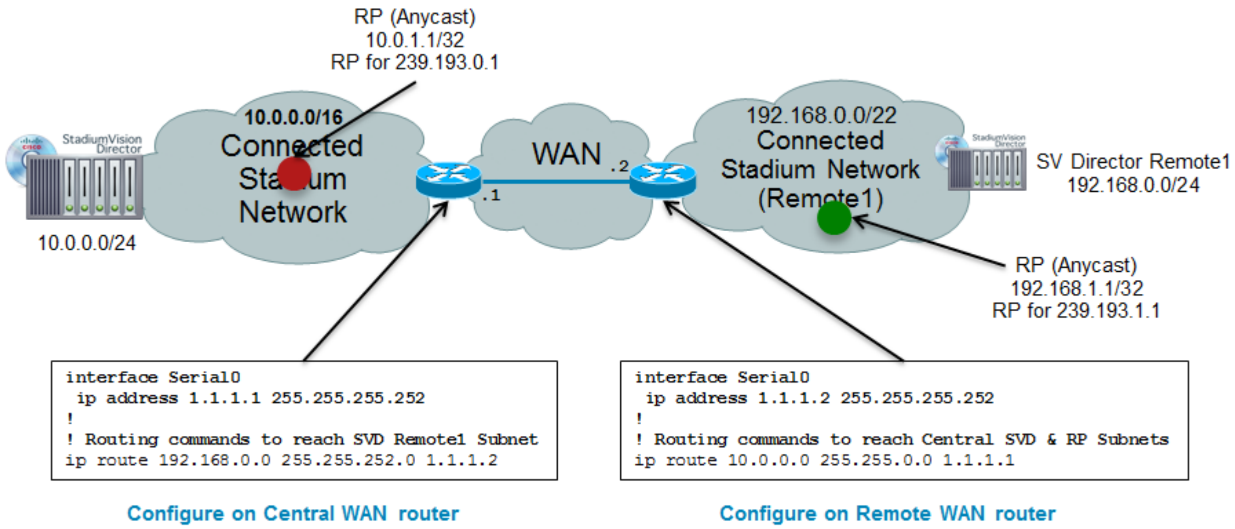

```
interface Serial0
 ip address 1.1.1.1 255.255.255.252
!
! Routing commands to reach SVD Remote1 Subnet
ip route 192.168.0.0 255.255.252.0 1.1.1.2
```

Configure on Central WAN router

```
interface Serial0
 ip address 1.1.1.2 255.255.255.252
!
! Routing commands to reach Central SVD & RP Subnets
ip route 10.0.0.0 255.255.0.0 1.1.1.1
```

Configure on Remote WAN router

Note: A dynamic routing protocol may be used in place of the static routes shown.

Figure 46 shows the commands used to enable both unicast and multicast routing over a routed WAN (e.g., private MPLS network). A Multipoint Generic Route Encapsulation (mGRE) tunnel is used to create a VPN across the WAN to enable multicast between the central and remote venues.

NOTE: A mGRE tunnel between the central and remote venue is not required if multicast is supported on the routed WAN.If the WAN is a public network (e.g., Internet), IPSec may be used for data privacy.

Figure 46. SVD Unicast/Multicast Routing over Routed WAN

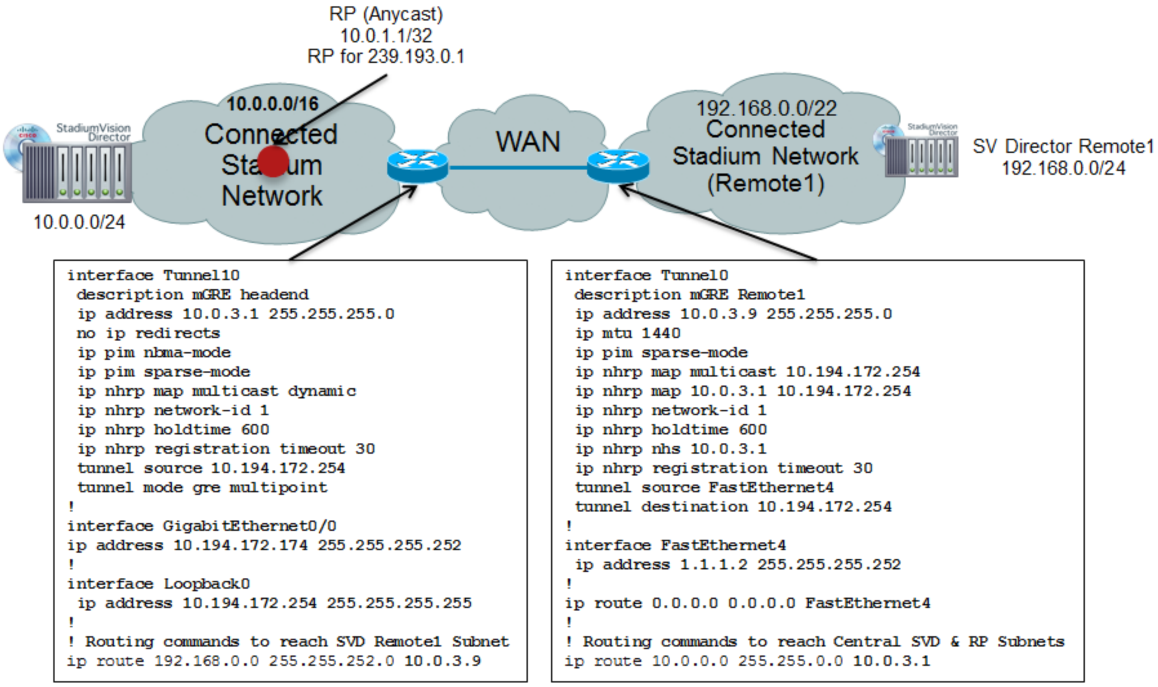

```
interface Tunnel10
 description mGRE headend
 ip address 10.0.3.1 255.255.255.0
 no ip redirects
 ip pim nbma-mode
 ip pim sparse-mode
 ip nhrp map multicast dynamic
 ip nhrp network-id 1
 ip nhrp holdtime 600
 ip nhrp registration timeout 30
 tunnel source 10.194.172.254
 tunnel mode gre multipoint
!
interface GigabitEthernet0/0
ip address 10.194.172.174 255.255.255.252
!
interface Loopback0
 ip address 10.194.172.254 255.255.255.255
!
! Routing commands to reach SVD Remote1 Subnet
ip route 192.168.0.0 255.255.252.0 10.0.3.9
```

```
interface Tunnel0
 description mGRE Remote1
 ip address 10.0.3.9 255.255.255.0
 ip mtu 1440
 ip pim sparse-mode
 ip nhrp map multicast 10.194.172.254
 ip nhrp map 10.0.3.1 10.194.172.254
 ip nhrp network-id 1
 ip nhrp holdtime 600
 ip nhrp nhs 10.0.3.1
 ip nhrp registration timeout 30
 tunnel source FastEthernet4
 tunnel destination 10.194.172.254
!
interface FastEthernet4
 ip address 1.1.1.2 255.255.255.252
!
ip route 0.0.0.0 0.0.0.0 FastEthernet4
!
! Routing commands to reach Central SVD & RP Subnets
ip route 10.0.0.0 255.255.0.0 10.0.3.1
```

Appendix A: Configuring an IOS DHCP Server to Support the SV-4K and DMP-2K

NOTE: An IOS Switch is not recommended as a production DHCP server for a venue. This example is provided to allow for simple DHCP support for DMPs primarily used for testing and in a lab.

IMPORTANT: If a Cisco IOS DHCP Server is used for IP address allocation for DMP's in a production network, we strongly recommend using database agents to maintain the DMP address bindings in case the switch is rebooted. See IP Addressing: DHCP Configuration Guide, Cisco IOS XE Release 3SE (Catalyst 3850 Switches) for further IOS DHCP Server configuration.

In this example, a Cisco IOS DHCP server is being used. Only a single pairing of options 60 and 43 can be used within a single scope (requiring two different IP scopes if there is a need to support both the Cisco DMP-4310G and the SV-4K media players, for example). This module only provides the procedure for the SV-4K media player.

Example 4. SV-4K DHCP Scope

```
! DHCP Database Agent CLI example

ip dhcp database ftp://user:password@172.16.4.253/router-dhcp write-
delay 120

!

ip dhcp pool SV-4K

network 10.10.1.0 255.255.255.0

default-router 10.10.1.1

option 60 ascii Cisco SV-4K-NA

option 43 hex
5540.6874.7470.3A2F.2F31.302E.3139.342E.3137.352E.3132.323A.3830.383
0.2F53.7461.6469.756D.5669.7369.6F6E.2F64.6D70.5F76.342F.7363.7269.7
074.732F.626F.6F74.2E62.7273
```

NOTE: The periods in the option 43 string are automatically created by the IOS and *are not typed* when creating the ASCII string. An example of the raw ASCII string value as typed can be found below.

Example 5. Creating Option 60

Type the line as shown above. Possible variations are as follows:

Some IOS versions may require the text enclosed in quotes “Cisco SV-4K-ROW”. If the IOS doesn't accept a leading quote in the string, then quotations are not needed.

SV-4K string for North America DMPs: "Cisco SV-4K-NA" SV-4K string for all other region DMPs: "Cisco SV-4K-ROW"

Example 6. Creating Option 43

Option 43 is a vendor-specific option. For the SV-4K and DMP-2K media player, it needs to point to a specific URL to execute auto-registration, auto-provisioning, and to receive initial firmware and subsequent upgrades, etc. That URL is:

http://*x.x.x.x***:8080/StadiumVision/dmp_v4/scripts/boot.brs**

where *x.x.x.x* represents the IPv4 address of the Cisco StadiumVision Director server.

Different devices require option 43 information be presented in different formats. The SV-4K and DMP-2K media player requires type-length-value (TLV) format for the data string. Specifically, the TLV format is constructed in the following manner:

- The string is built using hex values.
- The string begins with a two-character hex representation of the option 43 Type (an option 43 sub-option).
- The second two-character hex representation is the length of the information string, expressed in the number of ASCII characters of the string.
- Following the length value, the ascii string is typed out by using the two-character hex representation of each character in the string.
- For the SV-4K media player, the type designation is type 85 (decimal), expressed as type 55 (hex).

Example 7. Option 43

For the following URL string:

```
http://10.194.175.122:8080/StadiumVision/dmp_v4/scripts/boot.brs
```

The option 43 string would be as follows (except all on one line, without carriage returns) in the following format:

```
option 43 hex value
```

As shown below, the resulting URL above would be represented by the string below:

option 43 hex (there is a space here between ‘hex’ and the remainder of the string)

```
a.3830.3830.2f53.7461.6469.756d.5669.7369.6f6e.2f64.6d70.5
f76.342f.7363.7269.7074.732f.626f.6f74.2e62.7273
```

Using the description above and the ASCII lookup table below, the first characters of the string would be explained in the following way:

```
5540687474703A2F2F
```

55 – Type code in hex

40 – Length of the characters to the right of the length value, expressed by counting the number of two-character ascii characters represented (form the chart below). The numbers if characters in the example above is 64.

Decimal 64 = 40 hex.

And the first few characters, one at a time, would be as follows:

```
68 74 74 70 3A 2F 2F
```

H t t p : / /

The IP address is then represented the same way:

```
31 30 2E 31 39 34 2E 31 37 35 2E 31 32 32
```

10.19 4.17 5 .122

The remainder of the string is the text of the remainder of the URL.

> **NOTE:** As a Linux OS device, the Cisco StadiumVision Director server is case sensitive, and the correct letter case must be represented in the ASCII text strings.

TIPs:

- To make creating the string easier, start the string with 55 and copy the remainder of the ASCII string past the IP address, as these values never change.
- The example string had a length value of 3E (62 represented as a hex value).
- In this value, *do not* represent 62 as the hex character values of 6 and 2: instead, take the decimal value of 62 and convert to its equivalent hex (base 16) value (62 was converted to 3E in the example).
- Decimal-to-hex converters are easily found on the Internet, and most computer calculators have a decimal-to-hex function available, possibly under an advanced feature.

The example IP address was 12 characters long which leaves the remainder of the hex string to be 50 characters long. Add the number of ASCII characters in your IP address (including the dots!) to the known 50 characters and create your string length value to replace the 3E.

This technical information can be used to program other 3rd-party DHCP servers.

Figure 47. ASCII Character lookup table

DEC	HEX	Symbol	Description
32	20		Space
33	21	!	Exclamation mark
34	22	"	Double quotes (or speech marks)
35	23	#	Number
36	24	$	Dollar
37	25	%	Procenttecken
38	26	&	Ampersand
39	27	'	Single quote
40	28	(	Open parenthesis (or open bracket)
41	29	)	Close parenthesis (or close bracket)
42	2A	*	Asterisk
43	2B	+	Plus
44	2C	,	Comma
45	2D	-	Hyphen
46	2E	.	Period, dot or full stop

DEC	HEX	Symbol	Description
47	2F	/	Slash or divide
48	30	0	Zero
49	31	1	One
50	32	2	Two
51	33	3	Three
52	34	4	Four
53	35	5	Five
54	36	6	Six
55	37	7	Seven
56	38	8	Eight
57	39	9	Nine
58	3A	:	Colon
59	3B	;	Semicolon
60	3C	<	Less than (or open angled bracket)
61	3D	=	Equals
62	3E	>	Greater than (or close angled bracket)
63	3F	?	Question mark
64	40	@	At symbol
65	41	A	Uppercase A
66	42	B	Uppercase B
67	43	C	Uppercase C
68	44	D	Uppercase D
69	45	E	Uppercase E
70	46	F	Uppercase F
71	47	G	Uppercase G
72	48	H	Uppercase H
73	49	I	Uppercase I
74	4A	J	Uppercase J
75	4B	K	Uppercase K
76	4C	L	Uppercase L
77	4D	M	Uppercase M
78	4E	N	Uppercase N
79	4F	O	Uppercase O
80	50	P	Uppercase P
81	51	Q	Uppercase Q
82	52	R	Uppercase R
83	53	S	Uppercase S
84	54	T	Uppercase T
85	55	U	Uppercase U

DEC	HEX	Symbol	Description
86	56	V	Uppercase V
87	57	W	Uppercase W
88	58	X	Uppercase X
89	59	Y	Uppercase Y
90	5A	Z	Uppercase Z
91	5B	[	Opening bracket
92	5C	\	Backslash
93	5D	]	Closing bracket
94	5E	^	Caret - circumflex
95	5F	_	Underscore
96	60	`	Grave accent
97	61	a	Lowercase a
98	62	b	Lowercase b
99	63	c	Lowercase c
100	64	d	Lowercase d
101	65	e	Lowercase e
102	66	f	Lowercase f
103	67	g	Lowercase g
104	68	h	Lowercase h
105	69	i	Lowercase i
106	6A	j	Lowercase j
107	6B	k	Lowercase k
108	6C	l	Lowercase l
109	6D	m	Lowercase m
110	6E	n	Lowercase n
111	6F	o	Lowercase o
112	70	p	Lowercase p
113	71	q	Lowercase q
114	72	r	Lowercase r
115	73	s	Lowercase s
116	74	t	Lowercase t
117	75	u	Lowercase u
118	76	v	Lowercase v
119	77	w	Lowercase w
120	78	x	Lowercase x
121	79	y	Lowercase y
122	7A	z	Lowercase z
123	7B	{	Opening brace
124	7C	|	Vertical bar

DEC	HEX	Symbol	Description
125	7D	}	Closing brace
126	7E	~	Equivalency sign - tilde
127	7F		Delete

Appendix B: Cisco Connected Stadium Recommended Equipment

Core/Distribution

Table 11. Recommended Equipment for Core/Distribution

Hardware	Feature Set
Catalyst 6500 Series, Sup720, Sup2T Catalyst 6800 Series[1]	IP Services or higher IP Services or higher
Nexus 7000 Series – M & F2 Series Modules[2]	Enterprise Services Package
Nexus 7700 Series – F3 Series Modules[3]	Enterprise Services Package
Catalyst 4500E Series, Sup7E[4]	Enterprise Services
Catalyst 4500X, VSS[5]	Enterprise Services

[1] Caveat: Instant Access architecture has not been tested.

[2] If M Series and F2 Series modules are both present in the Nexus chassis. VDCs must be used for each module type. Also F2 series has a limited ARP table capacity which must be considered when using Layer 2 to the Access Layer switches.

[3] Caveat: The 7700 Series switches have not been lab tested but are approved for use due to their common hardware architecture and software as that of the 7000 Series switches.

[4] Requires 3.4.0SG for VSS support

[5] Requires 3.4.0SG for VSS support

Access

Table 12. Recommended Equipment for Access Layer

Hardware	Feature Set
Catalyst 2960X[1]/XR (PoE)	LAN Base/IP Lite[2]
Catalyst 3560X (PoE)[3]	IP Base
Catalyst 3650 (PoE)[4]	IP Base
Catalyst 3750X (PoE)	IP Base
Catalyst 3850 (PoE)[5]	IP Base
Catalyst 4500E Series (PoE)	IP Base

Data Center

Table 13. Recommended Switches for Data Center

Hardware	Software Feature Set
Catalyst 3750-X	IP Services or higher
Catalyst 4500X	Enterprise Services
Catalyst 4500, Sup7E	Enterprise Services
Catalyst 6509E, Supervisor 720	IP Services or higher
Catalyst 6509E, Supervisor 2T	IP Services or higher
Catalyst 6800 Series	IP Services or higher
Nexus Data Center Series[6]	Enterprise LAN Services

[1] The 2960-X does not have hot-swappable, dual power supply and fan modules

[2] IP Lite can be used for L3 to the Access Layer architectures

[3] The 3560-X is not stackable.

[4] Integrated WLAN functionality is not used.

[5] Integrated WLAN functionality is not used.

[6] If Service Blocks (e.g., Video Distribution, Wireless Service) are integrated into the Data Center Block, careful consideration must be made to accommodate the specific design requirements for the other Service Blocks. For example, MAC/ARP table capacity must support the number of Wi-Fi clients if integrating the Wireless Service Block.

Table 14. Recommended Servers for Data Center

Server Hardware
Cisco UCS B and C Series

Video Distribution

Table 15. Recommended Equipment for Video Distribution

Hardware	Software Feature Set
Catalyst 3560X (non-PoE)	IP Services or higher
Catalyst 3650 (non-PoE)	IP Services or higher
Catalyst 3750X/E (non-PoE)	IP Services or higher
Catalyst 3850 (non-PoE)	IP Services or higher

Internet Services

Table 16. Recommended Equipment for Internet Services

Hardware	Software Feature Set
Catalyst 3560X (non-PoE)	IP Services or higher
Catalyst 3650 (non-PoE)	IP Services or higher
Catalyst 3750X/E (non-PoE)	IP Services or higher
Catalyst 3850 (non-PoE)	IP Services or higher
ASA 5500 & 5500-X Series Firewalls	8.0 or higher
ASR/ISR Series (Service Provider Edge)	Advanced IP Services

Wireless Services

Table 17. Wireless Service Block Switches and their Connected Host Support

Hardware	Software Feature Set	Connected Host Support[1]
Catalyst 3850	IP Services or higher	19,000
Catalyst 6500, Supervisor 720	IP Services or higher	24,000
Catalyst 4500, Sup7E[2]	Enterprise Services	38,000
Catalyst 4500X[3]	Enterprise Services	38,000
Catalyst 6500, Supervisor 2T	IP Services or higher	80,000
Catalyst 6807	IP Services or higher	80,000
Catalyst 6816, 32, 24, 40, 80	IP Services or higher	100,000
Nexus 7000 Series, F2 Line Card[4]	Enterprise LAN Services	12,000
Nexus 7000 Series, M1 or M2 Line Card[5]	Enterprise LAN Services	80,000
Nexus 7700 Series, F3 Line Card	Enterprise LAN Services	50,000

[1] Wi-Fi Client Support is based on the switch ARP table limits. The listed numbers are approximately 20% lower than the absolute table maximum to allow for non-Wi-Fi client connections (e.g, devices connected to other ports). Note that MAC table limits may be much higher than ARP capacity.

[2] Requires 3.4.0SG for VSS support.

[3] Requires 3.4.0SG for VSS support.

[4] The Nexus with the F2 line card should not be used for connecting WLAN controllers. The MAC and ARP tables are very limited for supporting high numbers of Wi-Fi devices.

[5] Note Nexus M1 and M2 Line cards provide much higher capacity than F2 Line cards.

Table 18. Recommended Wireless LAN Controllers

Cisco WLAN Models	Description
Cisco 8500 Series	Support from 100 to 6000 access points and 64,000 clients with 8540
Cisco 5500 Series	Support from 1 to 1500 access points and 20,000 clients with 5520
Wireless Services Module Version 2 (WISM2)	Support from 300 to 1000 access points and 15,000 clients with WISM2

Table 19. Recommended Software Release Versions for Wireless LAN Controllers

Controller	Software Version	Comments
8500 Series	8.1.131.0	Always ask your Cisco technical representative to verify the latest recommended version.
5500 Series	8.1.131.0	
WISM2	8.1.131.0	

Appendix C: Cisco Connected Stadium Product Options

This appendix shows Cisco equipment that has been recently released or has not gone through functionality testing in the SESG lab, has very limited field exposure, or is not typically used in stadiums but are approved products for use.

This appendix provides a future look at equipment anticipated to be added to the design guide over time. This will enable the sale of new equipment with the assurance that it is intended to be supported and added to this design guide as it's solution tested and/or deployed in similar network enviornments demonstrating compatibility and stabliity.

Product	Place in the Network	Description
Cisco Catalyst 6800IA	Access Layer	As part of the Instant Access Solution, The Cisco Catalyst 6800ia is an extension of the Cisco Catalyst 6500/6800 parent switch
Cisco HyperFlex HX-Series Nodes	Data Center	Cisco HyperFlex HX-Series combine compute, storage, and networking into an easy-to-use system.
Nexus 2000, 3000, 6000, 9000	Core, Service Blocks, Data Center	Datacenter switches
Web Security Appliance (WSA/vWSA)	Data Center, Internet Edge	An appliance that provides advanced threat defense, advanced malware protection, application visibility and control, insightful reporting, and secure mobility
Cisco Content Security Management	Data Center, Internet Edge	The Cisco Content Security Management Appliance (SMA) centralizes management and reporting functions across multiple Cisco email and web security appliances.

Product	Place in the Network	Description
Appliance (SMA/SMAv)		
Cisco FirePOWER 7000/8000 Series Appliances	Data Center, Internet Edge	Next-generation intrusion prevention system (NGIPS) solution