Cisco Vision Network, Server, and Video Headend Requirements Guide

Dynamic Signage Director

First Published: 2020-05-18

Cisco Systems, Inc. www.cisco.com

Contents

About This Guide

This document provides a detailed description of the Cisco Vision Dynamic Signage Solution. This solution provides the wired infrastructure specifically designed 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.

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 Vision Dynamic Signage Solution.

This section includes the following topics:

Document Revision History

Table 1 lists the technical changes made to this document since it was first published.

Table 1 Document Revision History

Date	Change Summary
2020-05-18	Initial publication of Release 6.2.

Document Organization

This guide includes the following modules:

Chapter	Description
Solution Component Overview	Defines a brief overview of the components and operation of the Cisco Vision Dynamic Signage Solution.
Solution Operations and Deployment Requirements	Describes the network architecture requirements and the component design and deployment requirements.
Deployment and Requirements	Describes in deeper detail multicast applications and strategies, video wall synchronizing, wireless access and the Digital Media Players.
Headend Section	Describes various components of the headend, the encoders, encrypted video streams, and video monitoring.
Appendix A: Standards	Describes the Serial digital interface and Society of Motion Picture and Television Engineers standards.
Appendix B: Bill of Materials	Describes the Bill of Material for standard and small server deployments.

Related Documentation and Resources

Release-Specific Documents

- Release Notes for Cisco Vision Dynamic Signage Director Release 6.2
- For the listing page of all Cisco Vision documentation, go to:

  http://www.cisco.com/c/en/us/support/video/stadiumvision/tsd-products-support-series-home.html

Cisco Vision Dynamic Signage Director Documentation Go URL

For more information about Cisco Vision Dynamic Signage Director hardware and software installation, configuration, and operation, see the documentation available on Cisco.com at:

www.cisco.com/go/ciscovisiondocs

Cisco Vision Documentation Notifications

You can receive periodic emails that summarize new and changed information in Cisco Vision documentation by subscribing to the sv-doc-notify@external.cisco.com email alias.

Contact us at ciscovisiondocs@external.cisco.com to request this notification service.

Cisco Vision Documentation Team Email Contact Information

Submit questions, suggestions, or other feedback to us at ciscovisiondocs@external.cisco.com.

Cisco Vision Dynamic Signage Solution Component Overview

This section provides a brief overview of the components and operation of the Cisco Vision Dynamic Signage Solution.

The Cisco Vision Dynamic Signage Solution enables the integration and automated delivery of customized and dynamic content from multiple sources to different areas of the venue in standard definition (SD), high definition (HD), to ultra high definition (UHD). The solution is designed to enhance the visitor experience and provide the venue with additional revenue streams through targeted advertising via engaging, moving, dynamically updating content.

Four major components that constitute a Cisco Vision Dynamic Signage Solution (Figure 1 on page 9):

- Cisco Vision Dynamic Signage Director for centralized content management and operations
- Digital Media Player for content playback
- Cisco Digital Network, the IP infrastructure foundation for content transport
- Video Headend for video aggregation and distribution

Figure 1 Cisco Vision Dynamic Signage Solution Components

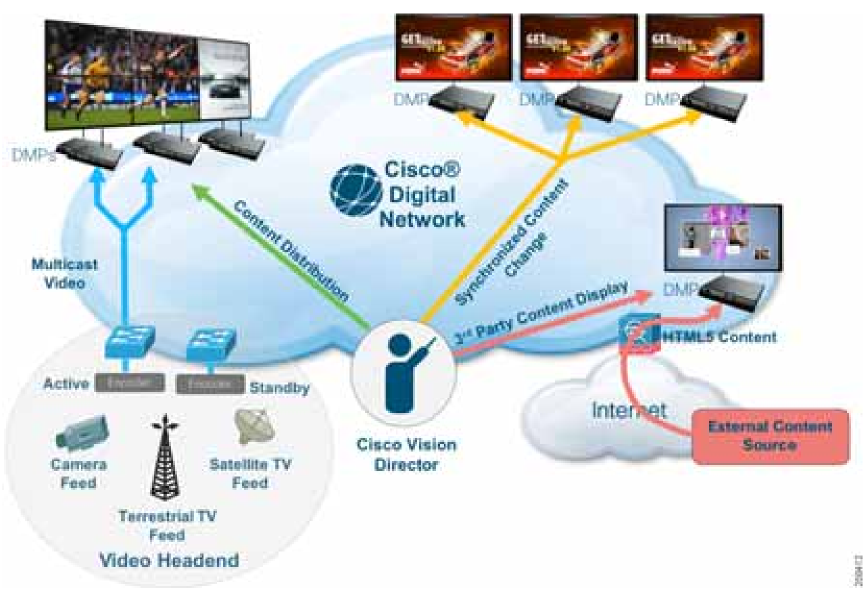

Cisco Vision Dynamic Signage Director

The Cisco Vision Dynamic Signage Director provides centralized management and operations for the Cisco Vision Dynamic Signage Solution. It acts as a single point of control for managing all Digital Media Player (DMP) endpoints, for placing and delivering content (video, graphics, and external content), for defining unique display areas (zones and groups), as well as for the creation of entitlement areas (bars, restaurants, clubs, and suites). It also provides the interface to third-party applications and devices, scoreboards and statistics systems, external contact closure and IP triggering systems, and third-party touch panels (for local, display control).

The capability requirements of the virtual server handling the Cisco Vision Dynamic Signage Director are based on the size and complexity of the deployment. There are two design size classifications defined in this document, primarily categorized by number of DMPs in the deployment: Standard and Small. However, other operational factors should also be taken into account.

Table 1 Director Server Classification

Server Classification	Number of DMPs	Operational Limits
Standard	5000	See product documentation[1]
Small	1250	See product documentation[1]

1. Refer to Release 6.2: Cisco Vision Dynamic Signage Director Operations Guide.

For detailed server specifications refer to Cisco Vision Dynamic Signage Director Solution Requirements, page 22.

Sample Bill of Materials (BOM) for servers fitting various deployment scales of Cisco Vision Dynamic Signage Director are located in Appendix B: Bill of Material, page 51 of this document.

Digital Media Player (DMP)

The DMP renders and displays static and dynamic content on each of the venue's connected displays. In addition to the support of UHD video resolution, the DMP can be powered by 802.3at Power over Ethernet (PoE) and supports dual video regions, video wall and virtual ribbon-board synchronization, and the rendering of HTML5 content. The DMP also supports Live TV playback via the HDMI 2.0a input to play content from any broadcast channel - even protected HDCP content.

For solution DMP deployment components, see Cisco Vision Product Deployment Requirements.

Cisco Digital Network

Cisco Digital Network is the foundational IP infrastructure that not only connects the video headend with the DMPs but typically interconnects all building IP endpoints to each other and to the outside world. The Cisco Vision Dynamic Signage Solution requires a converged, highly scalable, secure digital network designed specifically for low latency and redundancy to bring together all forms of access, communications, entertainment, and operations. This infrastructure is designed to enable the delivery of high-quality video, using advanced features of IP multicast and quality of service (QOS). This network also acts as the foundation to enable other services within the venue such as wireless communications, physical security, IP telephony, and network audio. The Cisco Digital Network is depicted in Figure 2 on page 11.

Non-Cisco network deployment is possible if all documented requirements are satisfied. See Cisco Vision Dynamic Signage Solution Operation and Network Requirements guide.

Figure 2 Cisco Digital Network

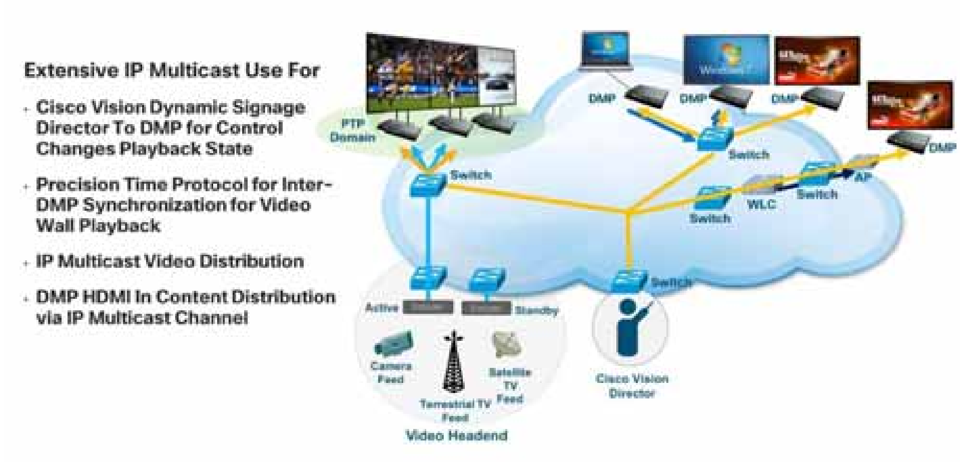

Video Headend

The headend is where video is received from various sources such as in-house feeds (through the venue video control room), over-the-air channels (typically from local over-the-air broadcast networks), and broadcast channels from cable or satellite providers. It is responsible for placing the video feeds onto the IP network with minimal latency. Video feeds may be provided in Ultra HD or HD resolution and are in encrypted or unencrypted formats.

The headend of the Cisco Vision Dynamic Signage Solution is designed to accommodate all of these feeds and perform the necessary encoding, transcoding, and extracting to create H.264 (MPEG-4, Part 10), H.265 (HEVC), or legacy H.262 (MPEG-2) encoded streams. The headend then takes the processed streams, assigns a unique IP multicast address to each, and places it on the IP network to be joined by the DMP endpoints as a channel.

Figure 3 Video Headend Overview

Deployment Models

On-Premise

The on-premise deployment model resembles an Enterprise client-server model with a server and its associated endpoints connected to an Enterprise campus network. The Cisco Vision Dynamic Signage Director server typically resides in the Data Center or in a Video Distribution service block (i.e., Video Headend) usually located near the broadcast room where all the various video input source feeds enter the venue.

Multiple Venue

The multiple venue deployment model is one where the Dynamic Signage Director is located at a central location and DMPs are distributed locally and across a WAN to remote locations. The WANs supported here include Cisco ISR (Integrated Services Router) which supports multicast, or Meraki-based WAN (unicast support only).

Figure 4 Deployment Model Overview

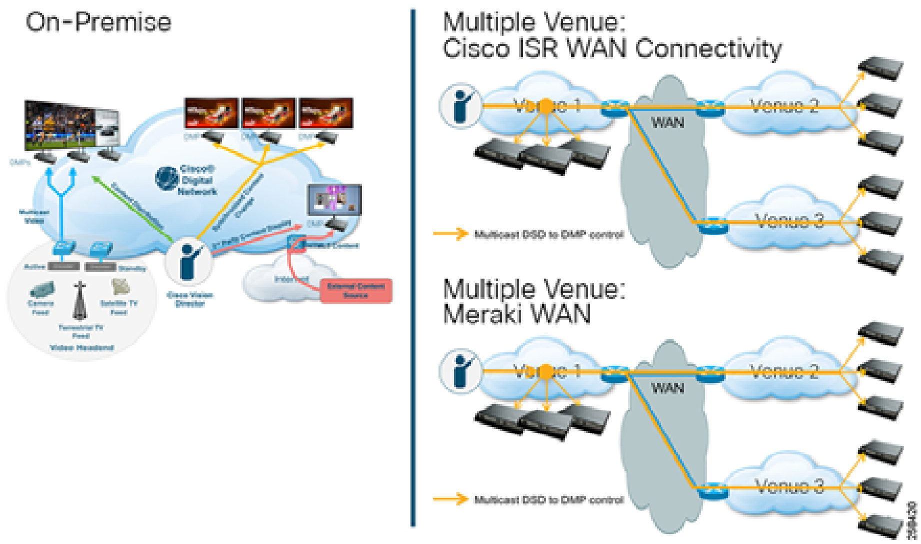

Solution Operations and Deployment Requirements

Network Architecture Role in Cisco Vision Deployments

Cisco Vision Dynamic Signage Solution is a proven solution for video and content delivery in large venue deployments where high performance video display formats, low latency, resiliency, along with a comprehensive feature set for configuring a variety of TV/monitor presentation modes is needed. However, Cisco Vision Dynamic Signage Solution is also used in deployments where a smaller feature set and with less critical timing constraints suffice for the customer engagement and interactivity needs.

The underlying network is a key determinant in the type of performance and feature set that can be delivered and supported by Cisco Vision Dynamic Signage solution. This section will highlight the network salient characteristics, prevalent best practices that have proven to work in large scale and demanding deployments, and alternative network features that can support the requirements of smaller deployments that have different performance delivery profiles.

Even though Cisco Vision Dynamic Signage Solution is a network-platform agnostic solution, the cited network architecture elements called out in this chapter will reference proven Cisco and Meraki network architectures and features. References to product family names, like Catalyst or Nexus, are meant to be illustrative and not prescriptive. It is more important to follow the relevant requirements that correspond to the scale/needs of the deployment. For convenience, these deployments will be referenced as performance-oriented and management-oriented. This categorization is not indicative of size of deployment or range of features, since scalability from single, large venue to large distributed deployments can be accommodated from the same Cisco Vision product.

In the sections that follow, requirements will be categorized as Mandatory or Recommended. Mandatory requirements are those which must be followed in the design. The Recommended requirements align the deployment with proven best practices and delivery of highest performance of the solution's available feature. Deviation from these guidelines may impact solution performance or its supported design attributes.

Network Architecture Requirements

Cisco's Borderless Network architecture, with some customization, provides the best practice blueprint for building a scalable, tiered, hierarchical, and modular design including collapsed core/distribution and access layers that is suited to large venue deployments of Cisco Vision Director.

The architecture modularizes functions into separate blocks which are dual-homed into a redundant, collapsed core/distribution pair of switches (Figure 1 on page 14).

Figure 1 Cisco Vision Network Topology

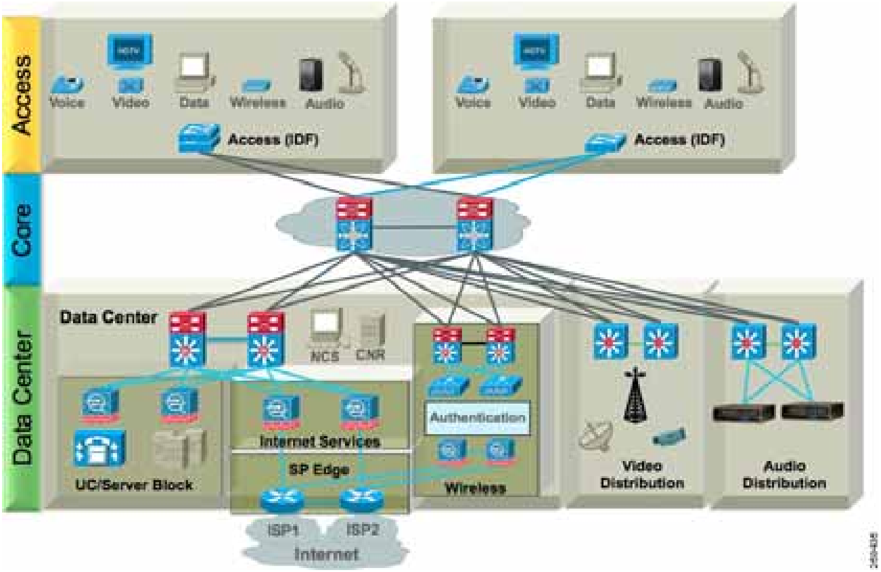

The Core layer of the network provides the high-speed, and redundant switching and aggregation of Access Layer switches including those used in the Service blocks. Building (aggregation) blocks provide flexibility during changes and upgrades.

Figure 2 Core Design Options

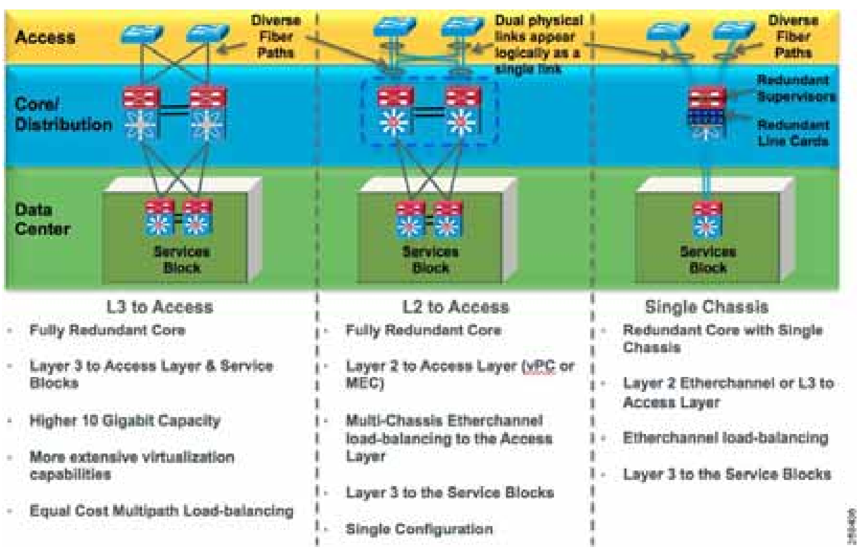

The attributes of the core designs are highlighted in the graphics above. This design can be typically implemented with Cisco Nexus or Catalyst switches or with Cisco Meraki switches.

Currrently, the collapsed core architecture is far more prevalent. The full Cisco Vision Director deployment in a Connected Stadium scenario is represented below.

Figure 3 Architecture for Connected Stadium

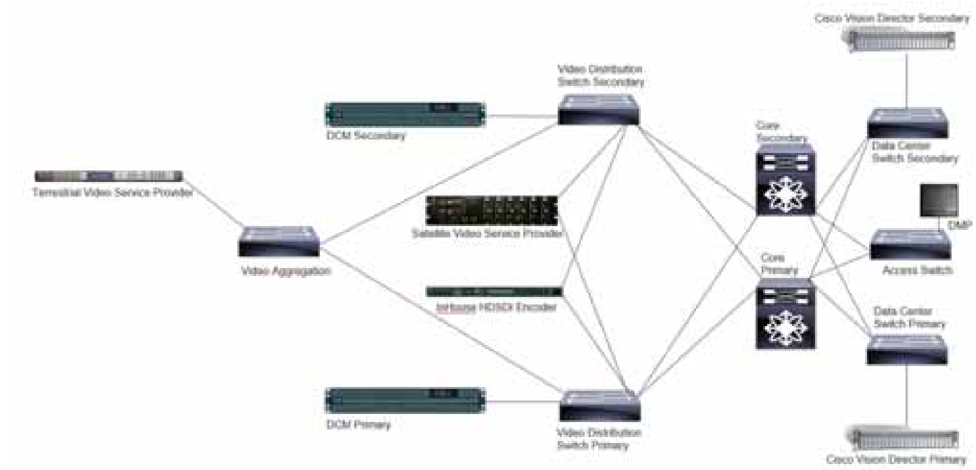

Cisco Software-Defined Access (SDA) as part of the Cisco Digital Network Architecture is supported for Cisco Dynamic Signage Director deployments.

Figure 4 SDA Architecture with Cisco Vision Dynamic Signage Director

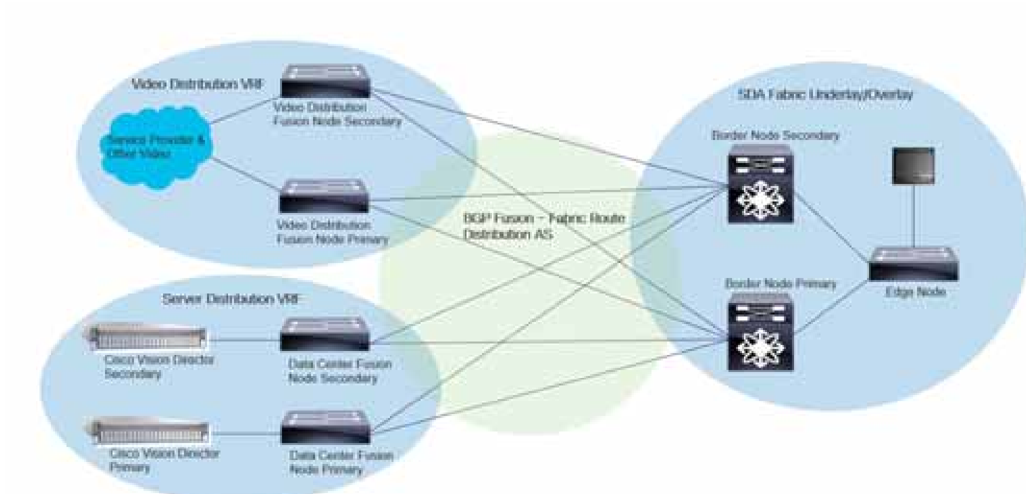

See Software Defined Access (SDA) Specific Guidelines, page 16 for other relevant requirements for this architecture.

Note: Mandatory requirements cited in this document must be followed in the design. The **Recommended** requirements align the deployment with proven best practices. Deviation from these guidelines may impact solution performance or its supported design attributes.

Multicast Role in Cisco Vision Director Solution Deployments

When Cisco Vision Dynamic Signage Solution is deployed in environments that require the highest scale, fastest performance, and full range of features, it requires IP multicast for the following functions:

- DMP control and zone-based content synchronization.
- Precision Time Protocol (PTP) for DMP-to-DMP synchronization.

- Encode and transmit IP multicast - The DMPs may take an HDMI input, or their own screen rendering, and provide an IP multicast stream source on the network.
- Joining multicast video channels coming out from the video headend.

Figure 5 IP Multicast Overview

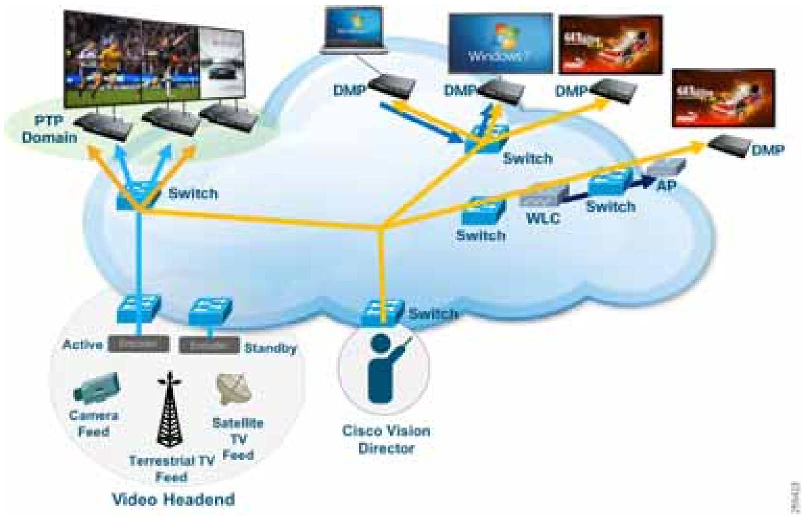

In third-party networks that do not support multicast (e.g., Cisco Meraki WAN), unicast configurations can be used from the Cisco Vision Dynamic Signage Director to direct DMPs to video sources.

Since multicast implementations, based on type of network equipment in use, can have several variations based on the level of features that are supported, the following section will highlight the key attributes. Find a more detailed review of the multicast design considerations in different operation scenarios in the Deployment and Requirements - Expanded Topics, page 25.

Software Defined Access (SDA) Specific Guidelines

These guidelines are specific to SDA, in addition to other Cisco Vision Solution network guidelines.

1. PTP has been tested in two methods with SDA/DNAC and the Cisco Vision Solution.
   a. PTP TTL (time to live) set to 1 (system default).
      With a PTP TTL set to 1, DMP synchronization will not function outside a single switch stack due to TTL being decremented across the VXLAN tunnel. All DMPs that require content synchronization, such as video walls, must be on the same switch stack.
   b. PTP TTL set to 4.
      With a PTP TTL set to 4, DMP synchronization will function across the entire VLAN.
      PTP TTL set to 4 is the preferred configuration. This may not be possible due to SDA limitations.
2. VLAN - Best practices for the Cisco Vision Director Solution is to have no more than 500 DMPs per VLAN. This is due to the tremendous amount of PTP traffic generated from DMPs on the VLAN.
   a. If PTP TTL is set to 1, then one large VLAN per site may be used because the PTP traffic is isolated per switch stack. Deviation from the known best practice of no more than 500 DMPs per VLAN has not been tested.
   b. If PTP TTL is set to 4, with synchronization on a per VLAN basis, the 500 DMP limit per VLAN should be adhered to. However, even though setting PTP to 4 is preferred for scalability, there may be practical hardware constraints related to SDA's architecture ability to accommodate a large number of potential multicast routes in the SDA fabric that should be considered.

3. QOS Video Requirements - QOS must be implemented throughout the fabric. Video traffic and DSD control traffic must be in a priority queue. All other traffic other than Voice must not take precedence in the fabric. QOS may be otherwise implemented according to SDA/DNAC (digital network architecture center) best practices.

4. Video Distribution Switch (VDS) Requirements

   a. Border Gateway Protocol (BGP) - VDS switches are considered “Fusion Nodes” in the SDA/DNAC architecture. Route protocol must be External Border Gateway Protocol (EBGP). Virtual Routing and Forwarding (VRF) for the overlay must be utilized and match that of the overlay fabric of SDA.
   The significant delta between standard design and SDA is the addition of BGP and tuned timers to facilitate quicker failover.
   Each VDS switch must peer with each SDA Border Node. Underlay global VRF peering may or may not be used.
   A network statement or redistribution of connection may be used for the Cisco Vision VRF.

   b. Source-specific multicast (SSM) must be configured on each VDS Switch.

   c. External Fabric Rendezvous Point (RP) any-source multicast (ASM) - VDS Fusion Nodes must be configured with an external PriorityCast RP address and associated access control list (ACL).

5. Server Switch / Datacenter Gateway Requirements

   a. BGP & VRF Route Leaking - The Server Switch is considered a “Fusion Node” in the SDA/DNAC architecture. Route protocol must be EBGP. VRF for the overlay must be utilized and matching that of the overlay fabric of SDA. VRF for the underlay must also be used for the purpose of route distribution between underlay where network services such as Cisco Vision DSD, NTP, and DHCP may be redistributed from the underlay or other network into the Cisco Vision VRF.

   The significant delta between standard design and SDA is the addition of BGP and tuned timers to facilitate quicker failover, as well as route leaking between VRFs.
   The Server switch must peer with each SDA Border Node.
   A network statement or redistribution of connection may be used for the Cisco Vision VRF or underlay.
   No ASM is supported on the Server switches.

   b. SSM must be configured on the Server switch.

6. SDA Edge Node

   a. External Fabric RP (ASM) - SDA Edge Node must be configured with an external PriorityCast RP address and associated ACL.

   b. LLDP must be configured on each edge node to facilitate power negotiation.

   c. SSM must be configured on the edge nodes.

7. SDA Border Node

   a. External Fabric RP (ASM) - SDA Edge Node must be configured with an external PriorityCast RP address and associated ACL.

   b. BGP - VDS switches are considered “Fusion Nodes” in the SDA/DNAC architecture. Route protocol must be EBGP. VRF for the overlay must be utilized and matching that of the overlay fabric of SDA.

   The significant delta between standard design and SDA is the addition of BGP and tuned timers to facilitate quicker failover.
   Each VDS switch must peer with each SDA Border Node. Underlay global VRF peering may or may not be used.
   A network statement or redistribution of connected may be used for the Cisco Vision VRF.

   c. SSM must be configured on the border nodes.

Multicast – Mandatory Requirements (Performance-Oriented Deployments)

The latest releases of Cisco Vision Dynamic Signage Solution support SSM and Internet Group Management Protocol (IGMP)v3 architectures. SSM allows for efficient data delivery in one-to-many communications. This means that Anycast and Prioritycast protocols referenced in the Mandatory listing are no longer required, even though the general design remains. If the network has implemented IGMPv2 [Protocol Independent Multicast (PIM) sparse and dense mode] then RP will be required for the Anycast and Prioritycast protocols to allow discovery and joining of multicast groups to their destinations.

Here are the MANDATORY requirements:

1. Implement PIM Sparse mode per VLAN Switch Virtual Interface (SVI) or routed interface must be used for the distribution or intermediate relay of Cisco Vision Dynamic Signage Director or IPTV Video Headend equipment traffic.
2. IGMPv2 on every DMP VLAN SVI should be implemented. IGMPV3 is optional but is mandatory for SSM deployments.
3. Enable IP Multicast in the Global Routing Table, VRF and/or virtual device context (VDC) where Cisco DMP, Cisco Vision Dynamic Signage Director and IP Multicast Video traffic will traverse.
4. Implement SSM for DMP control information. If not available, Anycast RP should be implemented in the Core/Spine switches for the Cisco Vision Dynamic Signage Director IP Multicast control information and the network Core/Spine switches must be the root of the Anycast RP tree. In addition, Multicast Source Discovery Protocol (MSDP) must be used to exchange source or group information between the Anycast RPs of any non VSS core/spine pair.
5. Implement SSM for video sourcing connection (Prioritycast). If not available, Prioritycast RP must be implemented in the IPTV Video Headend Video Distribution Switches, with PIM RP-Mapping configuration on all L3 networks and devices between the Cisco DMP VLANS, originating multicast video, Cisco Vision Dynamic Signage Director or other critical traffic.
6. Network support for IEEE 1588 Precision Time Protocol on every DMP VLAN Switch Virtual Interface.
7. The network Multicast Routing Information Base and Multicast Forwarding Information Base must support the total sum or greater of Multicast streams in the Intranetwork including Data, Video and PTP multicast.

Unicast – Considerations (Management-Oriented Deployments)

In networks that do not have the required multicast features, or where some DMPs are deployed at locations that cannot be reached via multicast, unicast settings in Cisco Vision Director allow remote control and operation of DMPs and the following restrictions should be noted:

1. State synchronization across multiple DMPs – unicast state changes cannot provide the same synchronization that multicast messaging offers.
2. Zone-based video wall synchronization is not supported.
3. Certain data integration with external sources that are normally done via multicast have to be implemented in a work-around fashion.

Routing/Switching - Mandatory Network Design Requirements (All Deployments)

1. Hierarchical model consisting of Core, Distribution and Access Layer, or a Collapsed Core model consisting of Spine and Leaf topology.
2. Power over Ethernet compliant to IEEE 802.1at (POE+ 30 watts per DMP interface for indicated models) provided by the network Access layer/Leaf switches.

The following table shows the power requirements of the Series 2-4 DMPs:

Table 1 Power Requirements of the Series 2 - 4 DMPs

Power Requirement	Series 2		Series 3		Series 4	
	DMP - 2K	SV - 4K	CV - HD	CV - UHD	CV - HD2	CV - UHD2
PoE	15W	Note[1]	15W	Note[1]	15W	Note[1]
PoE+		30W		30W		30W

Note [1]: When only 15W is available, it may appear that the DMP is partially working, but some features including display of video and HTML5 issues can occur Powering the DMP in this mode is not supported and it should not be deployed in this fashion.

3. Use redundant power supplies at the network Access layer/Leaf switches capable of supporting all connected DMPs.
4. Implement IEEE 802.1ab (LLDP) on all Access layer/Leaf switches physically connected to each DMP that requires 30 watts (e.g., UHD and UHD2 models).
5. Use unique IP DHCP scopes, providing infinite leases using option 43, for each Cisco DMP VLAN. (DHCP option 60 is optional).
6. Use a preserved state database with synchronization between redundant DHCP servers for each Cisco DMP VLAN.
7. Use Spanning Tree Portfast or equivalent on all host access interfaces.
8. BPDU Guard or equivalent on all host access interfaces.
9. Employ VLAN pruning on all trunk interfaces.
10. Use 802.1w Rapid Spanning Tree on all Intra-network trunk interfaces.
11. Do not exceed 500 DMPs per VLAN.

Use no security measures or mechanisms which may prevent a DMP from obtaining and maintaining a lease in the current provisioned VLAN. These may reallocate provisioned VLANS, intercept or respond to the DHCP ARP verification process or other security technologies, which may prevent access to the network in a timely manner or require authentication.

Use no traffic filtering device, or security or provisioning device, in any part of the end-to-end L2/L3 network path between Cisco DMP VLANS, IPTV Video Headend equipment, Cisco Vision Dynamic Signage Director and/or other critical traffic that would disrupt the normal flow of data traffic between the mentioned devices.

In SDA Architectures, video walls and synchronization are only applicable within a specific node infrastructure. This is because TTL is decremented between nodes and increasing PTP TTL is not recommended as a method to compensate for it.

Routing/Switching - Recommended Network Design Requirements (All Deployments)

1. At least 4 x 10-Gigabit uplinks between each Core or Spine switch.
2. Cisco Catalyst switches for the network Access layer/Leaf switches.
3. Cisco Catalyst or Nexus switches (redundant configuration) for the network Core/Spine switches.
4. Use virtual port channel (VPC), VSS or similar technologies in the network Core/Spine switches.

5. Use rapid per VLAN spanning tree (RPVST) on Intra-network trunk interfaces.
6. Two 10 Gigabit Ethernet links per network Access layer/Leaf and Distribution layer switches, with each link connected to an adjacent hierarchical switch.
7. Use fiber-based uplinks between the network Access layer/Leaf and Distribution layer to both Core/Spine switches.
8. Do not cascade network Access layer/Leaf switches.
9. Enhanced Interior Gateway Routing Protocol (EIGRP) is recommended for unicast routing in two-tier networks. Note the following exceptions:
   a. Network may use open shortest path first (OSPF) protocol. OSPF is hierarchical and typically used for three-tiered network designs.
   b. Network may use border gateway protocol (BGP) for connectivity to fusion nodes in SDA architectures.
10. Implement default gateway redundancy. Use Gateway Load Balancing Protocol (GLBP) which provides for router load balancing as well as redundancy among master and standby routers, or, for active/standby router redundancy only, implement either Virtual Router Redundancy Protocol (VRRP) or Hot Standby Router Protocol (HSRP).
11. Use dual redundant external Internet access for external content integration to sources beyond the Intra-network.
12. Use a separate Cisco DMP VLAN per intermediate distribution frame (IDF) switch stack.

Quality of Service (QOS) - Mandatory Requirements (All Deployments)

In Cisco Vision, end-to-end QOS support is needed for delivery of video traffic. Video 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 performance requirements of video delivered for display on DMP connected monitors, whether it is sourced from the video headend, or from a DMP that is used as an encoded source on the network. Here are the requirements:

1. Do not use Auto QOS.
2. Use Differentiated Services Code Points (DSCP) and Per-Hop Behaviors (PHB).
3. Implement traffic classification with appropriate DSCP marking applied at ingress to the network for all IP Multicast video, Cisco Vision Dynamic Signage Director data traffic, and other critical traffic.
4. Use DSCP classification of CS5 for all Cisco DMP IP Multicast, IP Multicast Video, and Cisco Vision Dynamic Signage Director IP Multicast traffic.
5. Implement an egress queue policy-map which includes a priority queue that reserves 15% of total 10-Gigabit uplink interface bandwidth for IP Multicast Video traffic on all interfaces between the Cisco DMP VLANS, IPTV Video Headend equipment, Cisco Vision Dynamic Signage Director, and/or other critical traffic.
6. Use QOS trust boundaries on all Intra-network interfaces between the DMP VLANS and the originating multicast video, Cisco Vision Dynamic Signage Director, or other critical traffic.

For additional design information, refer to the Enterprise QoS Solution Reference Network Design Guide.

Video Headend and Video Delivery Requirements

Follow these requirements for the transport and delivery of IP video architecture and design. Areas covered within the section include:

- Video Sources (VS)
- Video Encoding (VE)

- Video Transport (VT) and Delivery

Video Sources (VS) – Mandatory Requirements

1. For any In-House generated video, Terrestrial/Satellite/Over The Air video service providers and all other third-party video service providers, ensure the network supports all customer-selected UHD, HD and/or SD programs as provisioned in total aggregate throughput by the end customer and/or their authorized agents for input, and end-to-end in the Cisco Vision Dynamic Signage Director solution.
2. Ensure the IPTV Video Headend encapsulates all encoded MPEG-2, MPEG-4, and MPEG-H IPTV traffic in ISO 13818 MPEG-TS (transport streams) for transport throughout the network.
3. IPTV Video Headend must have redundant network connections between the video distributions switches and the Core/Spine switches.
4. Network infrastructure incorporates supported multicast routing for distribution of all IPTV channels.
5. For IGMPV2 implementations - IPTV Video Headend Video Distributions Switches must be the root of the Prioritycast RP Tree.
6. IPTV Video Headend Video Distribution Switches must not be used as access switches for DMPs, except for the purpose of monitoring video traffic.
7. IP Multicast Video must not be encrypted with any encryption algorithm except those explicitly outlined as supported in the associated Cisco Vision Dynamic Signage Director documentation.
8. IP Multicast Video Transport Streams must be received by DMPs as Single Program Transport Streams.
9. Locally played content must be encoded using a constant bit rate for use in video wall applications.
10. Ensure all video to be used in a video wall application complies with a 40 ms PCR interval and +/- 500 ms jitter/accuracy.

Video Sources (VS) – Recommended Requirements

1. For IGMPV2 implementations – if the IPTV Video Headend is implemented with a primary and redundant Video Multiplexer then Prioritycast RP should be used for Redundancy.
2. IP Multicast Video Transport Streams bit rate should be 10-25 mbps per HD stream and 25-35 mbps per UHD stream.
3. IPTV Video Headend should include a primary and redundant Video Distribution Switch (VDS).

Video Encoding (VE) – Mandatory Requirements

1. IPTV Video Headend must support MPEG-2, MPEG-4 or high efficiency video coding (HEVC) internet protocol television (IPTV) traffic in ISO 13818 MPEG-TS encapsulated in IP Multicast frames/packets.
2. IPTV Video Headend must support the encoding of uncompressed HD-SDI In-House video sources as MPEG-2, MPEG-4 or HEVC IPTV transport streams encapsulated in IP Multicast frames/packets.

Video Encoding (VE) – Recommended Requirements

1. Employ Cisco-qualified encoders in the IPTV Video Headend to inter-operate with the Cisco Vision Dynamic Signage Director solution.

2. Use standards-based MPEG multiplexer where needed to groom the stream in some supported method or to separate Multiple Program Transport Stream (MPTS) to Single Program Transport Stream (SPTS).

Cisco Vision Dynamic Signage Director Solution Requirements

The requirements captured within this section are generally derived from business requirements around the operation and maintenance of the Cisco Vision Dynamic Signage Director solution, including:

- Content Transformation (CT)
- Cisco Vision Dynamic Signage Director (DSD)

Cisco Vision Dynamic Signage Director (DSD) Server Requirements

Cisco Vision Director is designed to run on a virtual machine (VM) provisioned on an ESX server. For vSphere version compatibility, consult the Cisco Vision Software Installation and Upgrade Guide: Release 6.2.

Cisco Vision Director is available as an ISO image, where Release 6.2 ships with Red Hat Enterprise Linux (RHEL)7, while 6.1 (and below) ships with RHEL5.

Table 2 on page 22 lists the VM hardware and OS guest specifications for Cisco Vision Director Release 6.2.

Table 2 Virtual Machine Hardware and OS Specifications

System Component	Specification
VM Version	Release 6+
Guest Operating System	Red Hat Enterprise Linux 7 (64-bit)
Network Adapter	VMXNET3[1]
SCSI Controller	LSI Logic Parallel or LSI Logic SAS
Disk Provisioning	Thick

1. E1000 is used for RHEL 5 and VMNET adapter for RHEL 7. E1000 may be used for small configuration where the overall size of content to be distributed is small. VMXNET3 may not work for Cisco Vision Director versions prior to 6.2.

Two memory profiles are supported:

- Standard
- Small

Note: Support for two configurations started in Release 6.1.

Since Release 6.1, support for two configurations was implemented. A full installation of Cisco Vision Director will automatically choose the configuration based on the amount of RAM allocated to the VM. *Small* configuration is chosen by the installer when the detected RAM allotted to the VM is less than the minimum required for a *Standard* configuration. Choose the configuration based on the size and scope of the overall signage solution.

Note: For vSphere version compatibility, consult the Cisco Vision Software Installation and Upgrade Guide: Release 6.2.

Table 3 Minimum Virtual Machine System Requirements for Standard Configuration

System Component	Minimum Requirement
Processor	>2.5 GHz, 6 Core, 19.25 MB Cache (Intel 6128)[2]
Forward write operations per second	12 Gbit/s SATA SSD, Raid 10[3]
Virtual CPUs[1]	24
Virtual Disk Space	900 GB
Virtual RAM (VRAM)	32 GB

1. Hyperthreading can be used. Be sure that the BIOS is properly configured to enable it.
2. When selecting among CPUs with similar cost, choose higher CPU clock rate instead of additional cores.
3. For storage area network (SAN) implementations, a performance of 10K input/output operations/second (IOPS) is required.

Table 4 Minimum Virtual Machine System Requirements for Small Configuration

System Component	Minimum Requirement
Processor	2.5 GHz, 6 Core, 19.25 MB Cache (Intel 6128)[2]
Forward write operations per second	6 Gbit/s SATA SSD, Raid 10[3]
Virtual CPUs[1]	4 – 6, depending on capacity
Virtual Disk Space	225 GB
Virtual RAM (VRAM)	8 GB

1. Hyperthreading can be used. Be sure that the BIOS is properly configured to enable it.
2. When selecting among CPUs with similar cost, choose higher CPU clock rate instead of additional cores.
3. For storage area network (SAN) implementations, a performance of 10K input/output operations/second (IOPS) is required.

Cisco Vision Dynamic Signage Director (DSD) – Mandatory Requirements

1. Customer must adhere to the best practices outlined in the latest realease of the Cisco Vision Dynamic Signage Director Administration Guide, Cisco Vision Dynamic Signage Director Operations Guide, Configuring Cisco Vision Cisco Vision Dynamic Signage Director for External Triggers, Cisco Vision Director Data Integration Guide, and Cisco Vision Dynamic Signage Director Release Notes. Here is a link to the Product Page: https://www.cisco.com/c/en/us/support/video/stadiumvision/tsd-products-support-series-home.html

2. The server for Cisco Vision Director must meet the minimum system virtual machine requirements for CPU, Memory, Drive space, and read/write IOPS listed in Table 3 on page 23 and Table 4 on page 23.

3. The deployment must include a primary and secondary Cisco Vision Director server.

4. The primary and secondary Cisco Vision Director VMs must be in the same VLAN.

5. Do not stream video over 802.11a, b, or n WIFI to the DMP.

6. Use an internal NTP source accurately synced to a minimum of Stratum 3 to an NTP master clock. The NTP source must not be virtualized.

7. Cisco Vision Director VM must not be used by DMPs for NTP synchronization.

8. Customer must use the required firmware for each DMP model as specified in the Cisco Vision Dynamic Signage Director Release Notes.

9. Cisco Vision Director must have at least 1 Gigabit ethernet connection to the LAN. A 2x1Gigabit port channel or 2x10 Gigabit port channel is recommended for redundancy.

Director Server Ports

Note: For a complete port reference for Cisco Vision Dynamic Signage Director servers, see the “Port Reference” module of the Cisco Vision Software Installation and Upgrade Guide: Dynamic Signage Director for your release.

DMP Ports

While the DMP is a separate product, you can refer to the input/output ports on the DMP to ensure the proper communication with Cisco Vision Director and other external systems..

DMP – Bandwidth Considerations for WAN Deployments

For deployments where the DMPs are distributed, as in the case of Multiple Venue Deployments, bandwidth capacity should be properly considered even if live content is not streamed across the WAN. For example, DMP firmware upgrades will be needed from time to time and these files are generally in the 130-150 MB files size range. The transfer of firmware file may fail if due to bandwidth constraints it takes longer than 30 minutes (current time-out period. Refer to the latest Release Notes for updates) to complete.

To estimate the minimum length of time required for a file transfer:

1. Divide usable WAN bandwidth (in Mbps) by (Firmware upgrade file size in MB X 8 bits/Byte).

   The result will be in minutes and must not exceed 30.

Note: Actual transfer time will normally be higher due to packet headers, retransmissions due to high latency, etc. The DMPs are upgraded individually and consume the link until the upgrades are complete.

Deployment and Requirements - Expanded Topics

The topics covered in this chapter expand upon the checklist of requirements provided earlier and provides details on more technical and operational aspects of the networking requirements.

Multicast

The digital signage network uses IP Multicast for delivering a number of services. The unique requirements of these services are met using Protocol Independent Multicast (PIM) Sparse mode routing with two redundancy strategies. Earlier in this document differences in implementations between source-specific multicast (SSM) and rendezvous point (RP) were pointed out, and in this section, we'll provide a more detailed explanation into how they work in normal and failover scenarios. If implementing a new install, it is preferable to use the SSM guidelines instead of configuring RPs which is a supported legacy design architecture. Either method is very similar on how addressing is done on the video ingress, but they are different with regards to PIM communications and how the traffic gets forwarded from the source to the receiver.

Multicast Applications

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

Cisco Vision Director

- Video channels streamed to Digital Media Players (DMPs) attached. Uses Prioritycast RP Multicast Topology.
- Multicast control of DMP states (i.e., what is displayed on the TVs). Uses Anycast RP Multicast Topology.

Cisco Vision Director In-Suite Video

- IP Multicast video streamed from a Digital Media Player.

Note: To control the distribution of video, use a TTL=1 to limit the video to the local VLAN or use access control lists (ACLs) or multicast boundaries to limit what VLANs can request the video.

Protocol Independent Multicast

The IP Multicast design employs Protocol Independent Multicast (PIM) Sparse mode routing with Rendezvous Point (RP) redundancy in the following manner:.

- Because PIM Sparse mode operates in an on-demand fashion, receivers must request a video stream using an internet group management protocol (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 multicast 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.

Anycast

The Anycast strategy uses two or more RPs with the same IP address and mask and Multicast Source Discovery Protocol (MSDP) to distribute multicast source registration information among the RPs. MSDP is not used in virtual switching system (VSS) core designs. It is only used in the collapsed core and access designs (e.g., Cisco Nexus) or with other manufacturers that have a single control plane. In SSM implementations, MSDP is retained but RPs are not needed.

This methodology ensures that each RP knows about all sources and can facilitate building an RP tree between the source and receiver.

Figure 1 Anycast Overview

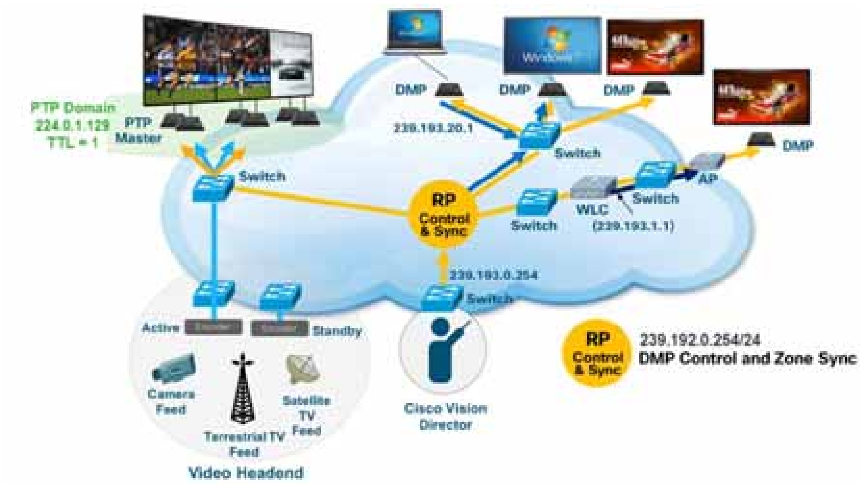

Operation and Configuration

1. In Anycast RP, all the RPs are configured to be MSDP peers of each other.
2. When a multicast source (e.g., Cisco Vision Dynamic Signage Director) registers with one RP, a Source Advertisement (SA) message will be sent to the other RPs informing them that there is an active source for a particular multicast group.
3. The result is that each RP will know about the active sources in the area of the other RPs.
4. If any of the RPs were to fail, IP routing would converge.
5. New sources would register with the next closest RP.
6. 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 2 Anycast in Action

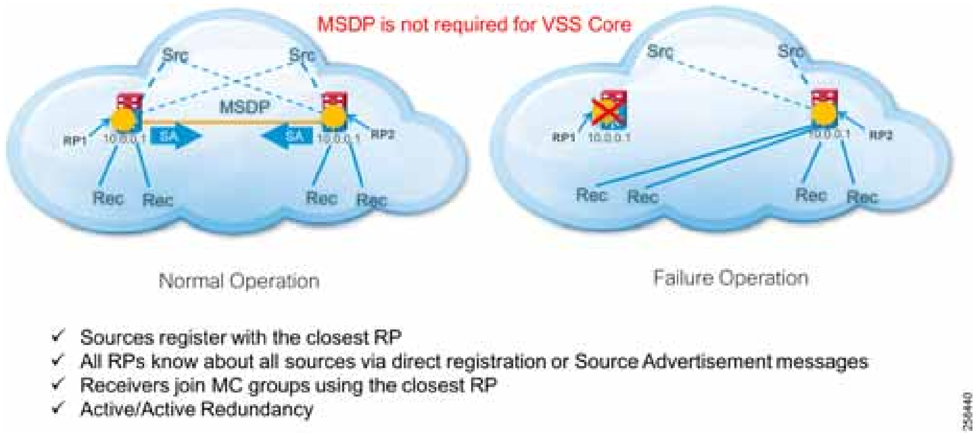

General IP Multicast traffic is handled by PIM Sparse-mode and the Anycast redundancy strategy. The use of PIM Sparse-mode is consistent with the Prioritycast stategy. However, MSDP is used to support RP redundancy.

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 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, RP, and it's backup use identical IP addresses with differing network masks.
3. The primary MC source and RP uses the longest network mask and is the active source and 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 3 Prioritycast in Action

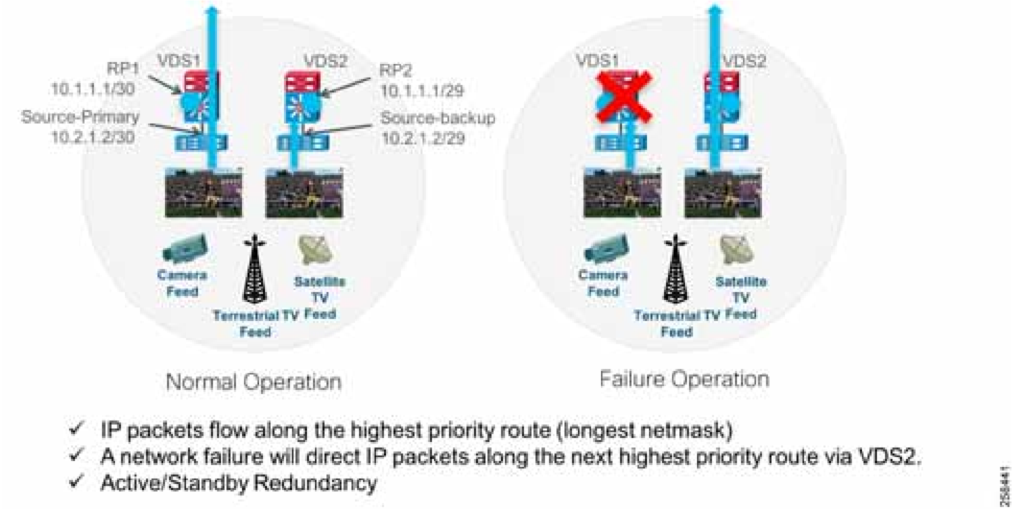

Multicast design uses these attributes:

- Uses PIM Sparse Mode multicast routing protocol.
- Uses a set RP on the Core switches for general multicast support.
- Uses a set of RP on the VDS for video multicast support.
- For general multicast DMP control, Anycast RP and MSDP for RP redundancy is used (Nexus core switches). For SSM, MSDP is still used for Anycast but without RPs.
- Anycast RP provides an active/active redundancy strategy.
- Uses Prioritycast RP for source and RP redundancy for video multicast.
- Prioritycast RP provides an active/standby redundancy strategy.
- Uses ACLs and MC Boundaries to limit MC to their designated areas.

Figure 4 Connected Stadium Multicast Architecture Overview

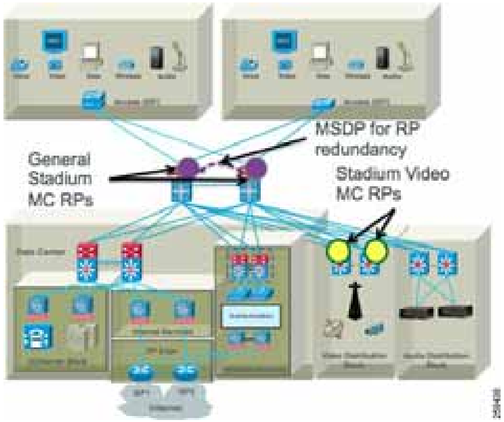

Multicast Redundancy Strategies

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 on each of the Core switches.
- MSDP netmask 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.

Analysis of Multicast and Failover Scenarios

This topic will illustrate in detail how multicast behaves normally and during failover of source or RP with Prioritycast RP and multicast sources.

Background

The RP is a loop back address defined on each VDS but with different subnet mask, /30 (primary VDS), /29 (secondary VDS). The rp-address is defined on each VDS, and the core VSS switch. The rp-address statement has an access list associated with it that defines what multicast addresses are associated with that RP address. Each VDS has a directly connected source via routed interface or VLAN with a L3 netmask of /30 (primary VDS), and /29 (secondary VDS).

Each loopback on the VDS switches has the same IP address as the other, and only the netmask differs. In most circumstances, this should also be true for the multicast source unicast IP address. They should be the same on each VDS switch. The SVI/routed interface should also have the same IP address but with different netmasks.

Normal Operation

The VDS prioritycast loopback RP addresses are up on both VDS switches. There are multicast packets coming into each VDS switch via the routed/VLAN with netmask of /30, /29, respectively. The source IP address of the multicast is the same for each unique multicast addresses on each VDS switch (from the same source on both VDS1 and VDS2).

1. The Route table on each VDS switch will show a directly connected route for loopback and each routed/SVI source network. Each switch will also learn of the other networks via routing protocol, but the connected routes will be preferred because their administrative distance (AD) is 0 in the routing table.
2. Route table on the VSS core will only know about the routing protocol advertisements and will learn of the RP and routed/SVI source networks. The VSS core will insert the network with the longest mask into the routing table for both RP and routed/SVI networks.
3. VSS core will map via the rp-address statement multicast source address as defined by the associated access list.
4. Then an IGMP Join comes in from a receiver into the VSS switch.
    a. The Join contains a multicast group to be joined.
    b. VSS core will see if the joined group matches the access list associated with the rp-address statement.
    c. VSS core finds a match defined to an rp-address statement and the associated multicast range defined by the access list. The address defined by the rp-address statement is the prioritycast loopback address of each VDS switch.
    d. VSS core looks up the address in its routing table and finds two matches for the loopback address. It will see both /29 and /30 routing table entries.
    e. Switch sends the pim join upstream to the RP and the unicast route with the longest match, out the interface to VDS1 is selected because it has the /30 mask.
    f. VDS1 already has all multicast routes built in its routing table for each incoming group coming from the source attached to the /29,/30 routed or SVI interface.
    g. VDS1 has a (*,G) and source specific (S,G) built for each source. The (*,G) is the anysource definition and the (S,G) is the source specific definition for each multicast address.
    h. VDS1 after receiving the join from the VSS core also looks up the rp-address statement as well as the associated access list. The VDS1 finds that the RP is a directly connected loopback address and has a match for the requested group in the access list associated with the rp-address statement. VDS1 looks up the requested group in its multicast routing table and finds a match.
    i. VDS1 adds the outgoing interface to the *,G for the interface that the join came in on towards the VSS core and forwards one packet to the VDS core.
    j. VDS1 will also build the same for the (*,G) entry. With the one packet it received, it will inspect the source IP of the multicast packet then send a shortest path tree (SPT) join towards the network of the source IP (VDS1).

k. VSS core adds the outgoing interface to the *,G for the interface that the join came in on towards the receiver VLAN. VSS core will also build the same for the S,G entry.

l. VSS core then floods the requested group to the receiver VLAN.

One important thing to note here is the *,G versus the S,G because it will play an important part of the decision making process in any failover configuration. There are two important concepts at play here. One is SPT and the other is rendezvous point tree (RPT). SPT is the source specific tree that gets built, this is the S,G entry (this only happens after RPT in non SSM, (ASM configurations). The S,G is the entry associated with the best unicast reverse path to the multicast source IP itself and is source specific multicast (SSM). The RPT is the shared anysource tree or RP tree, anysource multicast (ASM). The *,G RPT entry is associated with the best unicast reverse path to the RP address as defined by the rp-address statement and associated ACL.

In an ASM (anysource *,G) environment, it's important to understand that the RP is the single point of registration for both multicast sources and receivers. This is because joins do not contain any source information, the RP is tasked with keeping track of all incoming multicast sources as well as all incoming requests from downstream PIM neighbors (receivers - DMPs). The RPT mechanism for joins is always the first to occur in RP designs because receivers and downstream PIM neighbors do not know the unicast source IP of the multicast address. After the join reaches the RP via the RPT from a downstream PIM neighbor, a switchover occurs. The RPT will forward an ASM multicast packet to the downstream PIM neighbor tree that made the request. Once each downstream L3 neighbor receives this packet, it then builds its own SPT/SSM entry in its multicast routing table and the downstream neighbor no longer needs the RPT or RP to receive the multicast because it knows the unicast path towards the source and multicast path in which to send PIM joins towards. This is called the RPT to SPT cutover. The cutover in modern IOS happens with a 0 bit threshold, which means that the neighbor (VSS core) will switch over to the SPT (SSM) path upon receiving the first packet from the RP, or a PIM message from the RP with the source IP of the source multicast group, and it does this based on the unicast route table for the unicast source IP of the group. After this point, RP is not used and why you always see both *,G and S,G entry in all the downstream multicast routing tables.

Failure Case: Source Network Fails; RP Does Not Fail

In this scenario, the same overall steps from above still occur until we get to step E. Let's take into account three sub scenarios:

1. one with an existing *,G/S,G already built before the failure

2. a new one to be built after the failure of the source VLAN SVI/routed interface but before the multicast routes on VDS1 timeout

3. one join requested after the multicast routes on VDS1 timeout

Any multicast routes already built will continue to work. Because the route is already installed on the VSS core to the source multicast network, when the unicast source information disappears for the original path to VDS1, the secondary path to VDS2 will replace it in the switches unicast routing table. The VSS core will send PIM joins to the secondary VDS switch. The secondary VDS switch will receive the source specific join and will start to forward traffic. One thing to note is that the *,G entry will remain on VDS1, because this is where the RP is located. The S,G entry will be on VDS2. No *,G entry built with outgoing interface list (OIL) will be built on VDS2 because the group is strictly source specific at this point. The streams will stay built towards VDS2.

After the failure, there is a period of time in which the S,G multicast routes on VDS still exist. However, no multicast source is coming into VDS1, but the entries will remain until the expire timer removes them. During this time, VDS1 still knows the source IP for the requested group. The mroutes (multicast routes) on VDS1 will show the incoming interface change from the directly connected interface, to the interface that faces VDS2 (VSS core's interface), it learns this path from it's own unicast routing table. If a join comes in during this time from the VSS core, the VDS1 switch will reply to the join request with a PIM message that contains the source IP of the group that it still has cached in the multicast routing table VSS core will then build the stream, but forward PIM joins towards the known source from it's unicast routing table, which would be on VDS2. The stream will be built and stay built towards VDS2.

After the failure, VDS1's multicast routing entries for S,G will expire. A join from a receiver off of the VSS core will still be forwarded towards the RP on VDS1. VDS1 will build a *,G anysource entry in it's routing table with OIL of the VSS Core, but no traffic will flow from VDS1 to the VSS core because VDS1 at this point has no S,G. In the PIM join reply to the VSS core from VDS1, there will be no source information included, so this time the VSS core has no knowledge of the source IP of the multicast group and will not forward any PIM join requests to VDS2. The only entries that exist for this join will be on VSS core and VDS1 and it will be an ASM join (*,G) only. The stream will not be built as S,G and no data will flow.

Failure Case: RP Fails, Source Network Does Not Fail

This is the least likely scenario to happen. The RP loopback address should never go down. If it does go down from the perspective of the VSS core, it's likely due to a configuration issue, or the links between VDS1 and the core fail. In the instance of the links failing, the source network would also likely fail as well.

After the failure of the RP on VDS1, any existing multicast will continue to play as the VSS core has already sent source specific join to VDS1. As long as the source network doesn't go down, their streams will continue to be forwarded from VDS1 even though VDS2 is the RP. Any new streams will have the join sent to VDS2. VDS2 will send 1 packet down to the VSS core and the VSS core will be sent a source specific join towards VDS1 for the source network.

Complete Failure of RP and Source Network on VDS1

This should behave as you might expect. Both RPT and SPT joins will go to VDS2.

Summary

To summarize, the behavior of RP up, source network up, will flood from VDS1 only. If Source network goes down, but RP stays up on VDS1, then existing multicast on VDS1 will switch over to VDS2 and new multicast routes will continue to build until the mroute tables entries expire, after which point no new multicast routes will be built.

There are three ways to mitigate this. One is to just get rid of RP and use SSM. With the current releases of Cisco Vision Director this would be a good idea. You could also use MSDP between VDS1 and VDS2 to share SA's between the two. The third method is to use an embedded event manager (EEM) script to bring down the RP if the source network fails. Both MSDP and EEM would violate our design best practices but could be used as a one off-if the customer accepts the risks.

Network and Precision Time Protocols (NTP & PTP)

A common time source for synchronization is important in any network, but especially for delivery of video content that needs to seamlessly appear across multiple DMPs and monitors. Use the check list below to understand the requirements and caveats for provisioning Cisco Vision Dynamic Signage Director servers and DMPs to use NTP and PTP for synchronization.

- Network Time Protocol (NTP) service is required in Cisco Vision Solution on the following devices:
  - Cisco Vision Dynamic Signage Director servers
  - Series 2-4 DMPs that are designated as the Precision Time Protocol (PTP) master device
- By default, both NTP and PTP services are automatically enabled for the DMPs.
- An NTP source also must be used to provide initial clocking to the devices that are elected PTP masters in the network.
- Only the DMP PTP masters derive a clock using NTP.
- Do not use Cisco Vision Dynamic Signage Director as an NTP source for other devices in your network.

- If deploying Cisco Vision Dynamic Signage Director as a virtual machine, configure Cisco Vision Director to use a reliable NTP server running on a bare metal server rather than a source from the local VM environment.
- Verify Cisco Vision Dynamic Signage Director and DMPs can reach the NTP source.

The DMPs must not reference an NTP server pool. If the Cisco Vision Dynamic Signage Director server references an NTP server pool (the default), then select a specific server from that same pool as the NTP server for the DMPs.

- Only IPv4 is supported for the NTP server address on the DMPs.
- The NTP server for the DMPs must not be a load-balanced server.
- The Cisco Vision Dynamic Signage Director network must be configured to allow bidirectional transmission of UDP messages on port 123 for NTP messages between the NTP source and DMPs.

PTP for Video Wall Synchronization

NTP provides reliable clocking for your Cisco Vision network and helps ensure synchronicity between redundant servers and between the Cisco Vision Dynamic Signage Director and the DMPs. The IEEE 1588 Precision Time Protocol (PTP) is used between DMPs.

Figure 5 The Role of NTP and PTP for Synchronization

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

Figure 6 DMP PTP Synchronization

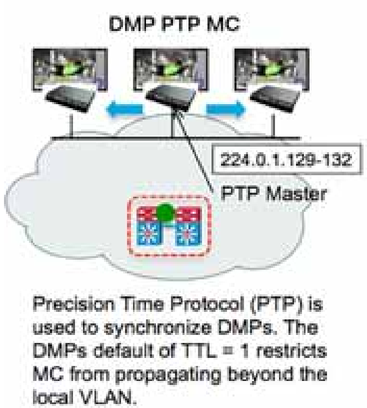

When using PTP, one DMP is designated as the domain master clock. It will synchronize with a NTP reference clock and then act as the reference point for a set (subdomain only) of slave DMP clocks. The protocol provides the means for slave DMPs to determine the path delay incurred from the master to themselves. 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:

224.0.1.129 – Default subdomain 0 (only subdomain supported)

Note: The DMPs use a TTL of 1 default, meaning PTP multicast is confined to the local subnet or VLAN. The TTL may be changed to greater than 1 to traverse a number of layer 3 hops. Careful consideration should be used when configuring TTL > 1 to traverse multiple hops due to the increased latency incurred, and the potential to exhaust the FIB database. This may negatively affect synchronization. Also, the multicast routing in the network must be configured for the PTP group addresses mentioned above.

In SDA architectures, TTL is decremented between nodes even in the same VLAN, so PTP will only function within a node or switches cascaded from a node. It is not recommended to compensate by increasing TTL since it might exhaust the multicast route table on the underlay. Given this, video walls and synchronization will only be applicable within a specific node infrastructure.

Fiber Uplinks

Note: The more common design implemented today is collapsed core, using trunk connections from the access layer back to the core, and with all layer 3 back in the core.

Access layer stack of switches are connected via 2 x 10 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. 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, Multi-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 the 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.

Spanning Tree and Protection

Spanning Tree Rapid PVST is enabled by default on many Cisco switches and provides per-VLAN spanning tree protection. It is preferred over MST since that only provides one spanning tree domain. Use RPVST or 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

Bridge Control Data Unit Guard (BPDUGuard) enabled on all access ports ensures that 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.

Wireless Access

Certain models of DMPs supported Wi-Fi connectivity. For example, Series 2 DMPs and some specific hardware models in the Series 3 DMPs supported Wi-Fi. For those use cases, if content synchronization is needed, then the Wi-Fi network must support multicast to the edge. Refer to the relevant Wi-Fi design guides for further information.

Unicast control is currently supported in Cisco Vision Director releases. However, that feature does not support content synchronization.

Figure 7 Multicast over Wi-Fi Overview

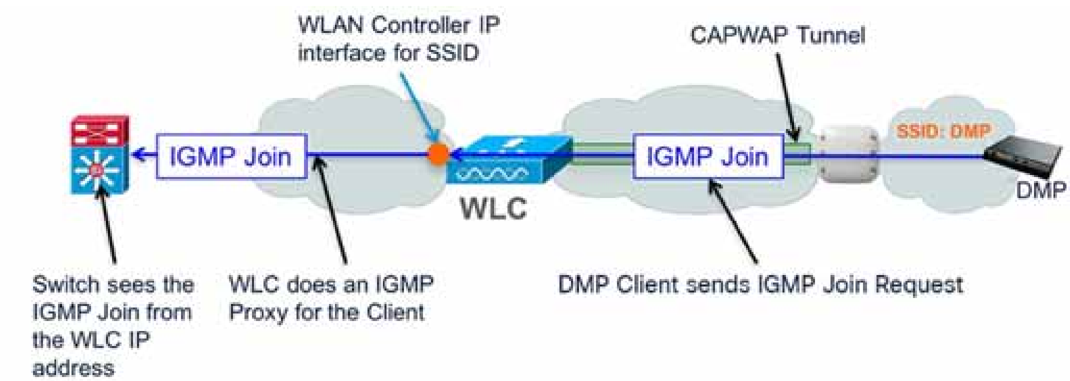

Digital Media Player Topics

The Role of LLDP in DMP Connectivity

The Cisco Vision DMPs support standard Link Layer Discovery Protocol (LLDP). This capability allows the switch and DMP to learn about each other by exchanging LLDP messages and to negotiate 802.3at power over the Ethernet connection. Cisco Vision Dynamic Signage Director uses LLDP information for populating the switch information in the Management Dashboard. This aids in troubleshooting.

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 new DMPs require higher power to take advantage of new capabilities. The following table shows the power requirements of the DMPs.

Table 1 Power Requirements of the DMPs

Power Requirement	Series 2		Series 3		Series 4	
	DMP - 2K	SV - 4K	CV - HD	CV - UHD	CV - HD2	CV - UHD2
PoE	15W	Note[1]	15W	Note[1]	15W	Note[1]
PoE+		30W		30W		30W

Note[1]: When only 15W is available, it may appear that the DMP is partially working, but some features including display of video and HTML5 issues can occur, so powering this DMP in this mode is not supported and it should not be deployed in this fashion.

The Access Layer switches should also always be equipped with redundant as well as 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/PoE+ ports.

Figure 8 The Role of 802.3at Power over Ethernet and LLDP

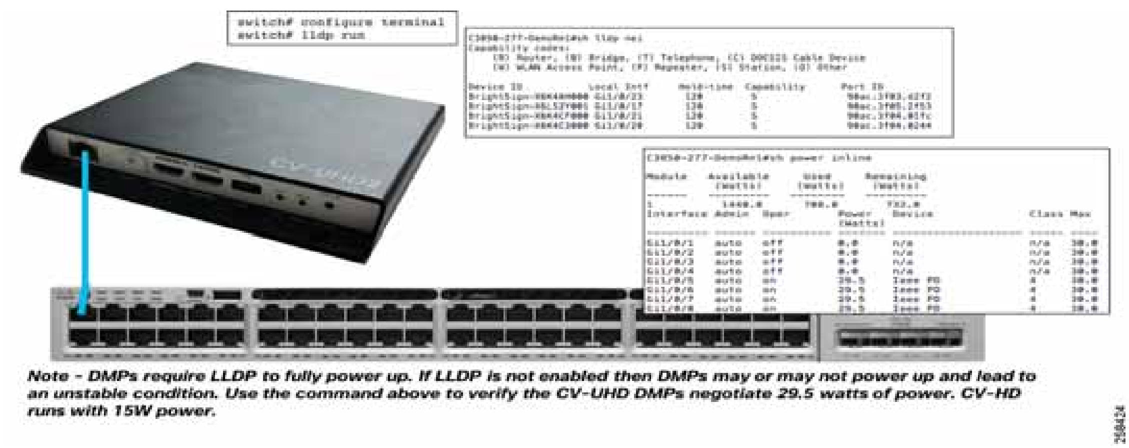

The Role of Switch Port Civic Location

IOS civic location is a collection of labels that can be configured on each switch port, and then communicated to the DMP via LLDP. 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 Vision Dynamic Signage Director.

The Role of DHCP Options for DMP Auto-Registration

The DMP receives firmware and configuration from Cisco Vision Dynamic Signage Director. The DMP finds the Cisco Vision Dynamic Signage Director server using DHCP options (option 43). Historically, option 60 was used with early DMPs to signal the DMP for the correct option 43 content meant for it. Option 60 is no longer used or required for current DMP models.

Figure 9 DHCP and DMP Auto Registration

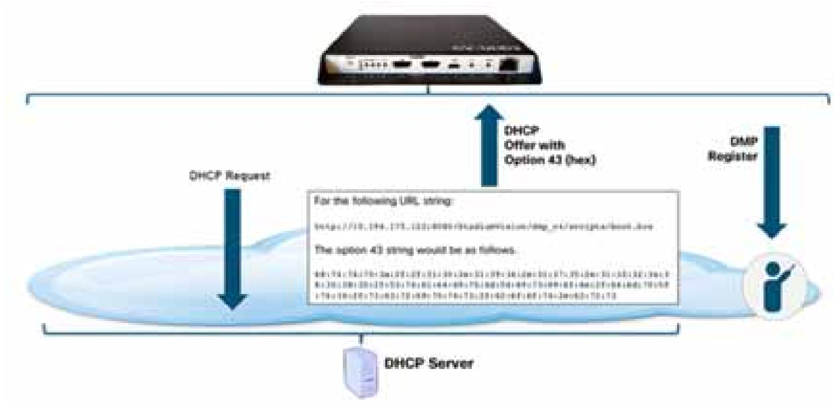

Configuration of DHCP Option 43 TLV String

This section highlights some details on option 43 use and format only. For full configuration details, refer to the appropriate Cisco Vision Deployment Guide for the specific Series and/or models of Digital Media Players. The guides are listed here:

https://www.cisco.com/c/en/us/support/video/stadiumvision/products-maintenance-guides-list.html

DHCP Option 43 provides a means of delivering a vendor specific URL to the DMP. Depending on the switch/IOS version, the DMP may require type-length-value (TLV) format for the option 43 data string, otherwise the “ascii” option can be used in some IOS versions. Specifically, the TLV format is constructed in the following manner:

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

For the following URL string, where 10.194.175.122 is the IP address of the Director server:

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

The option 43 string would be as follows.

Hint: Use an ascii-to-hex conversion tool to simplify creating the hex string.

```
68:74:74:70:3a:2f:2f:31:30:2e:31:39:34:2e:31:37:35:2e:31:32:32:3a:38:30:38:30:2f:4
3:69:73:63:6f:56:69:73:69:6f:6e:2f:64:6d:70:5f:76:34:2f:73:63:72:69:70:74:73:2f:62
:6f:6f:74:2e:62:72:73
```

Next, you place in front of this string the hex representation for <decimal type code>:<decimal number of characters in the string>

Note: In Microsoft Word, you can carefully highlight the string and then click Tools > Word Count to get the number of characters in the string.

The type code is 55 in hex and in the above URL example, there are 62 characters in the string. Decimal 62 is equal to 3E in hex.

```
55:3E:68:74:74:70:3a:2f:2f:31:30:2e:31:39:34:2e:31:37:35:2e:31:32:32:3a:38:30:38:3
0:2f:43:69:73:63:6f:56:69:73:69:6f:6e:2f:64:6d:70:5f:76:34:2f:73:63:72:69:70:74:73
:2f:62:6f:6f:74:2e:62:72:73
```

Connecting the DMP to the Wi-Fi Network

This section does not apply to the current DMP models and is left for reference for specific older DMP models that support Wi-Fi connectivity and are deployed where there is no existing Ethernet cabling, or simply as an alternative to Ethernet network connectivity.

Refer to the appropriate Cisco Vision Deployment Guide for the specific Series and/or models of Media Players. The guides are listed here:

https://www.cisco.com/c/en/us/support/video/stadiumvision/products-maintenance-guides-list.html

Using a DMP as an IP Multicast Source

Two modes of streaming from the DMP are supported:

1. HDMI-In streaming: In Release 4.1 and later releases, Cisco Vision Dynamic Signage Solution supports streaming audio/video from a laptop or other supported device connected to the HDMI-In port on the SV-4K, CV-UHD, or CV-UHD2 media players to be played as a multicast-based channel over the wired Ethernet port.
2. Display streaming: The entire video composition (without audio) on the DMP and presented on a TV connected to its HDMI-Out port will be encoded and streamed out from the DMP as a multicast-based channel over the wired ethernet port. This is useful to create a stream source that is a composition of multiple elements on the screen, for example a HTML5 page, and some other video source (or even HDMI-In video).

Note: Proper QOS DSCP classification must be set for DMP IP Multicast traffic on the network.

Figure 10 Using the DMP as an IP Multicast Source

Note: If you want to maintain privacy of channels, create a DMP-encoded channel per suite with a unique multicast address (from 239.192.20.0/24 range), and create a separate channel guide per suite. For example, if you have 10 suites–create 10 separate DMP-encoded channels with unique multicast addresses, create 10 different channel guides for each DMP-encoded channel, and assign each suite to a different channel guide.

For more information about configuring this feature, see Release 6.2: Cisco Vision Dynamic Signage Director Operations Guide.

DMP Control and Content Synchronization

Cisco Vision Dynamic Signage Director uses IP multicast to send messages to control DMPs and synchronized content. Be sure to use the value that is configured in your network for transport of Cisco Vision Dynamic Signage Director control messages. Typically, the Multicast RP used for DMP control and synchronization is on the network's core switches. If SSM is used instead of RP, then for address usage adhere to RFC 4607.

Figure 11 Using IP Multicast for DMP Control and Content Synchronization

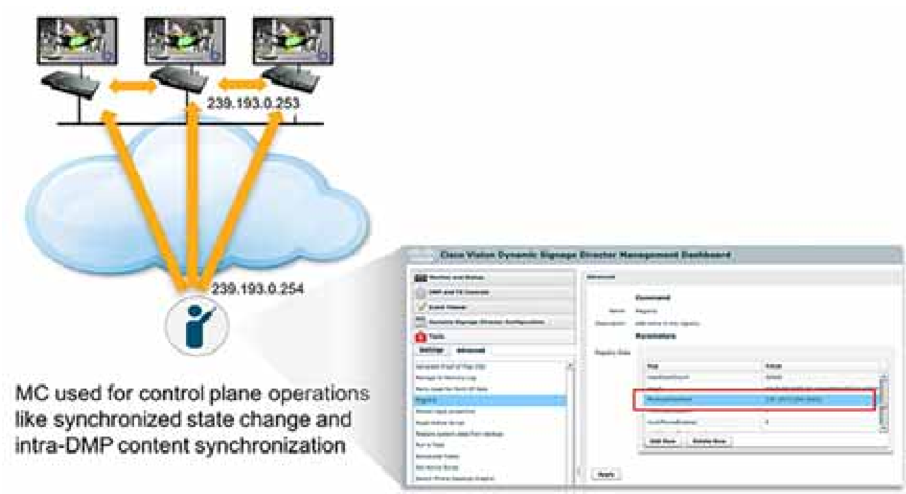

Table 2 Multicast Addresses used by the Cisco Vision Solution

Table 3

Multicast Address	Default Value	Description
239.192.0.0/24	239.192.0.254 DMP Control 239.193.0.253 DMP synchronization	For example, 239.192.0.254 - DMP Control from Cisco Vision Director 239.193.0.253 Zone-based Synchronization (TTL=1) between DMPs
239.192.0.0/24	Configured in the Video Headend Per RFC 5771 Administratively Scoped Blocks	Video MC Channels
239.192.20.0/24	Needs to be configured in Cisco Vision Director Per RFC 5771 Administratively Scoped Blocks	DMP as MC Source
224.0.1.129	224.0.1.129 Default, Zone 0*	PTP for Synchronization

Note: As of Release 6.2, only PTP domain 0 is supported. See latest Release Notes for updates.

Cisco Vision Dynamic Signage Director on the Network - Failover

For redundancy, Cisco Vision Dynamic Signage Director is installed on two virtual servers, where one of the servers operates as the primary active server and the other server operates as a secondary backup server. If a failure occurs, you can configure the backup server to become the active server, but the failover process is not automatic.

Both servers must reside in the same VLAN, have the same hardware configuration, and optimally connecting to their own switch as shown in the diagram below. HSRP should be configured to provide default gateway redundancy. Cisco Vision Dynamic Signage Director servers would typically be installed in the Data Center.

Figure 12 Cisco Vision Server Configuration Overview

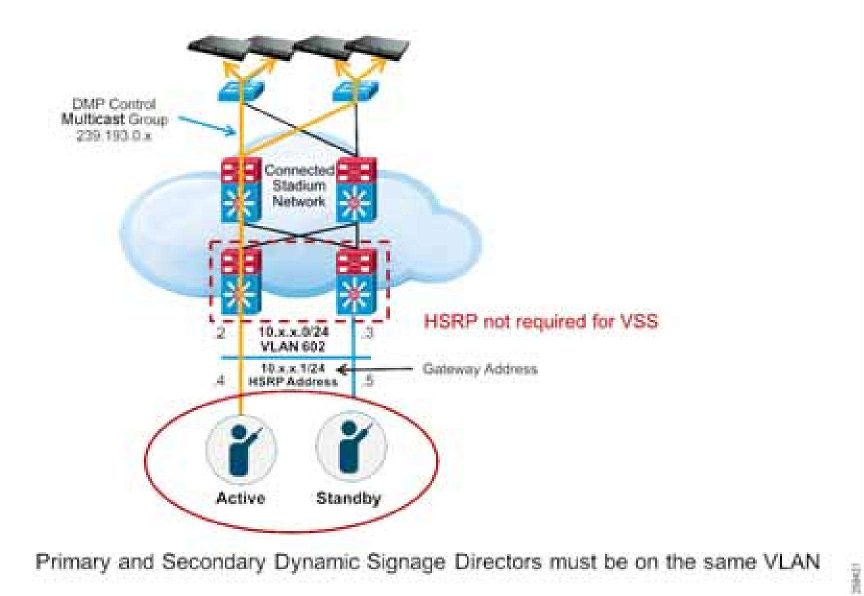

The primary and secondary servers are addressed as independent hosts with two different IP addresses on the same subnet.

The secondary server is only connected to the network to be made available as a backup to the primary should a failure occur. In addition, the secondary server can (and should) be configured to be backed up with data from the primary server on a scheduled basis so that it can be ready as a warm standby.

When the primary server fails, a manual process is used to restore the secondary server from a backup, shut down the primary server, change the secondary server's IP address to that of the primary and then to bring the secondary server into service.

Note: Although connecting servers to the Core switches is not typically recommended. There are instances where this may be 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 the Cisco Vision Dynamic Signage Directors are connected.

For HSRP configuration, find the appropriate IOS software release guide for your switch.

Headend Section

Cisco Vision Director is a proven, end-to-end, high-definition IPTV solution that provides advanced video content management and delivery. It is a centrally managed, video processing and distribution solution that enables the integration and automated delivery of customized and dynamic content from multiple sources to different monitor displays in High Definition (HD), or Ultra High Definition (UHD) local video.

Cisco Vision Director is purpose-built for large venues, retailers, hospitality providers, and transportation hubs which have extensive video systems deployed throughout and is designed to enhance the viewing of live events, multimedia information, and dynamic content. In addition, it leverages video systems in restaurants, clubs, and luxury suites to allow customers to view both in-house programming as well as external network channels.

Cisco Vision Director comprises four major components, as shown in Figure 1 on page 41.

- Video acquisition (or video headend)
- Converged voice, video, and data high-speed IP network
- Video delivery (and digital signage playback)
- Centralized management and operations

Figure 1 Cisco Vision Major Components

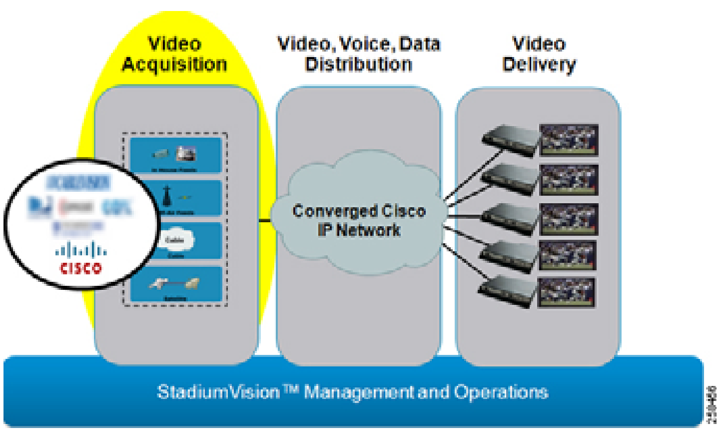

Functional Overview

The headend is designed to acquire, process, and encode the video content used in the Cisco Vision Director solution. Figure 2 on page 42 provides a simplified view of the video headend design, incorporating multiple types of video sources.

Figure 2 Headend Functional Elements

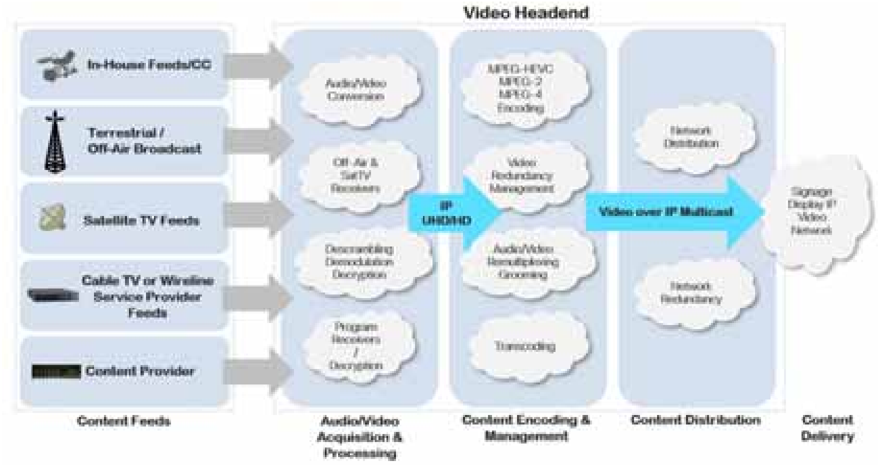

In the headend, the video feed is:

- Provided by multiple sources.
- Acquired and processed through the appropriate receivers and decrypters.
- Encoded using an HD (H.264/MPEG-4) or UHD (H.265) encoder.
- Groomed and aggregated using a standards-based MPEG multiplexer
- Sent using multicast to the IP network.

Once the multicast stream is on the network, the DMP can join it via IGMP and display it on the corresponding TV.

The Cisco Vision headend provides support for:

- High-definition (HD) and Ultra High Definition (UHD) channel lineup with delivery of H.262 (MPEG-2, Part 2) and H.264 (MPEG-4, Part 10) and H.265 (HEVC) into the video distribution plant.
- In-house video feeds.
- Terrestrial TV (also called Off-Air or Free-to-Air) video feeds.
- Encrypted video feeds from a cable or satellite provider.
- Direct IP feeds from cable provider or from DirecTV receivers (as in North America).
- Encoded external video sources. The DMP as an encoder can provide this function
- Support for standards-based MPEG multiplexer and the native features it offers for changing video and audio feeds in the headend core.
- A fixed channel lineup where each video channel is set to an IP Multicast address.

Headend Reference Architectures

Cisco Vision Director includes two options for headend architecture:

- Standard (recommended) architecture, which provides redundancy with failover to ensure high availability.
- Baseline architecture, which provides a lower-cost entry point that can be modified at a later date to incorporate the redundancy and failover of the standard architecture.

Note: Generally speaking, this document assumes a standard (redundant) architecture.

Baseline Architecture

The baseline architecture (Figure 3 on page 43) provides an entry-level solution for cost- sensitive venues for North America and internationally. In this architecture:

- There is no redundancy. However, the design allows this to be added at a later date.
- The recommended encoding for in-house feeds (HD/UHD) is H.264/H.265.
- No provisions are made for to accommodate legacy analog TV or RF plants.
- The video distribution switch (VDS) should have dual 10 Gigabit connections to the core switch.

Standard Architecture

- The standard architecture (Figure 3 on page 43 and Figure 4 on page 43), is the recommended architecture for Cisco Vision Director deployments. Beyond the baseline architecture, it incorporates redundancy for high availability, as well as support for additional features.

Figure 3 Standard Reference Architecture for North America

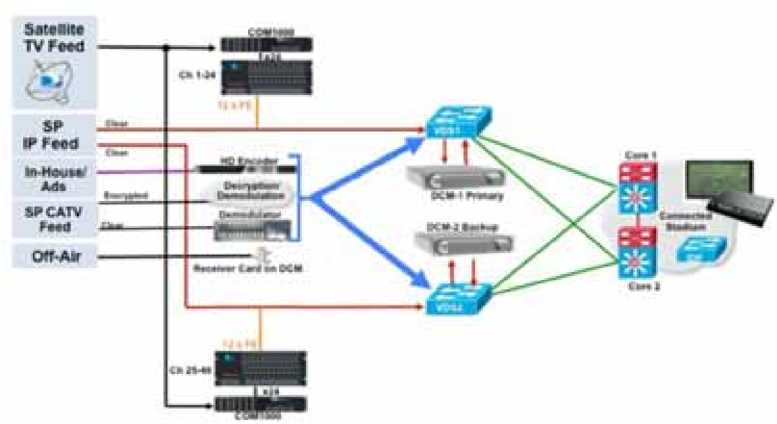

Figure 4 Standard Reference Architecture for International Video Headend Implementations

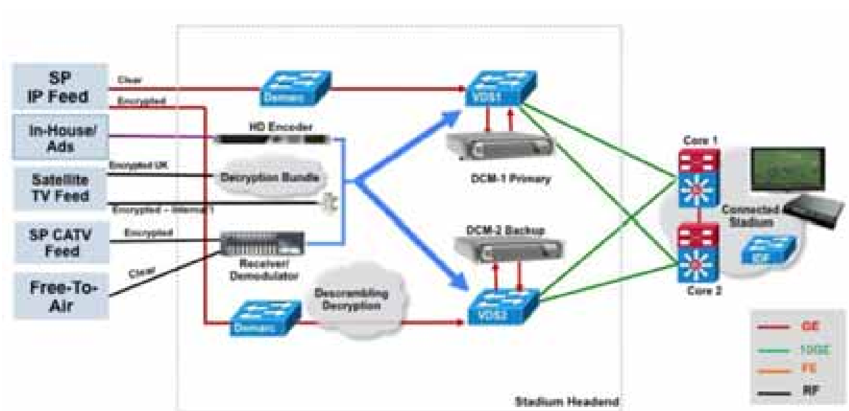

In the standard architecture:

- Feeds, acquisition devices, processing devices, and video distribution switches are redundant.
- The recommended encoding for in-house feeds (HD/UHD) is H.264/H.265.
- The VDS should be deployed in a redundant fashion with dual links to each of the core switches

Standards-Based MPEG Multiplexer and Headend Network Infrastructure

Standards-based MPEG multiplexer and the video distribution switches (VDSs) are at the heart of the headend in the Cisco Vision Director solution. While other components of the headend may vary depending on the type of feed, these components remain relatively constant.

- In the baseline head-end design, the core MPEG multiplexer and video distribution switches can be deployed in a non-redundant configuration.
- For improved availability, in the standard headend design the MPEG Multiplexer and the video distribution switches are deployed in a redundant configuration, as previously shown in Figure 3 on page 43.

Standards-based MPEG Multiplexer

The standards-based MPEG multiplexer, like Synamedia's DCM D9902, takes video feeds from the various sources and sends the feeds to the video distribution switch with fixed IP Multicast addresses. The core MPEG Multiplexer can receive either:

- H.262, H.264, or H.265 feeds over Asynchronous Serial Interface (ASI) connections from encoders or demodulators.
- H.262, H.264, or H.265 via direct IP feeds (unicast/multicast) from a video provider or from a local source. In this case, the MPEG Multiplexer serves as a demarcation point between the carrier and encoded feeds and the IP video distribution network.
- Video sources encoded by Harmonic CP9000 encoder should not be passed through the MPEG Multiplexer since it adds 100ms of latency to the path of multicast video.

An example of MPEG Multiplexer is Synamedia DCM D9902, but any standards-based MPEG Multiplexer is acceptable.

Video Distribution Switch

The video distribution switch provides the connection between the headend and the converged Cisco IP network. These switches support the advanced features (quality of service, IP Multicast, strict priority queuing) and the performance required for distribution of video streams over an IP network.

The choice of video distribution switch should be based on the required port density, as well as throughput requirements. While PoE is not required for standard operations, it may be needed for video troubleshooting and test points for DMPs located at the headend.

For general switch requirements see Solution Operations and Deployment Requirements, page 13.

Redundancy for MPEG Multiplexers in Headend Streams

For redundancy, each HD video channel is sourced from two MPEG Multiplexers, which are configured in an Active-Active setup. The two MPEG Multiplexers act as redundant multicast sources for each channel and connect to separate access routers on Layer 3 interfaces. Although both MPEG Multiplexers are actively sending video streams of the same content, the video streams of the secondary MPEG Multiplexers will only be forwarded on the network in the event of a failure of the primary video stream.

The network is setup for automatic failover of the primary video feed to redundant paths. As shown in Figure 5, if the video feed port on the MPEG Multiplexers Primary fails, the network will converge to the alternate active video feed emitted from the MPEG Multiplexers Secondary (backup). The link between video distribution switches is set with a relatively high EIGRP Delay (in the event of a core-video distribution switch link failure, to force an RPF interface change to the core router).

In the Cisco Vision design, an IP Multicast technique called Prioritycast is used for video multicast. This enables both MPEG Multiplexers to send out exact replicas of the channel lineup using the same IP multicast addresses. Prioritycast is an implementation strategy that provides load sharing and redundancy in Protocol Independent Multicast sparse mode (PIM-SM) networks.

With IGMP-V3 enabled, Source-Specific Multicast is the preferred method which does not require Rendezvous Points (RP) to be configured in the network. But where that option is not available, Prioritycast allows two or more rendezvous points (RPs) to share the load for source registration and the ability to act as hot backup routers for each other. Unlike Anycast, in which clients connect to the closest instance of redundant IP address, with Prioritycast, clients connect to the highest-priority instance of the redundant IP address. This allows for greater control of designating the preferred source.

In the Cisco Vision Director IP multicast implementation:

- The source is the MPEG Multiplexer and the router is the video distribution switch.
- Each multicast channel is assigned the source IP address of the MPEG Multiplexer.
- Each channel is assigned a source prefix of /29 in the IP Multicast address scheme, further divided into a /30. The source (channel) is assigned the lowest of the two host addresses and the router interface is assigned the highest of the two host addresses.
- Each MPEG Multiplexer uses the same source IP address but with different prefixes (i.e., /29 and /30).
- The MPEG Multiplexer with the most specific (highest priority) prefix (i.e., /30) is the primary MPEG Multiplexer serving the network.
- The primary servers and RPs are on the same video distribution switch using the /30 network mask.

In this design, the Reverse Path Forwarding (RPF) interfaces are determined using the installed unicast route of the longest advertised RP and sources prefixes. Thus, the rendezvous point tree (RPT) and shortest path tree (SPT)-switchover trees are on the same path.

Additional advantages of this design include:

- One hop convergence for source redundancy due to network failure that results in the loss of the RP/Source prefixes.
- No secondary source traffic exists in the network while providing optimum network failover to the secondary source.
- Shortest path trees are built to the installed longest available prefix for the source IP address

Note: If there is a switch failure or MPEG Multiplexer failure, the secondary MPEG Multiplexer will take over in under a few seconds. However, there will be a momentary freeze-frame and glitch seen on the video screen at the DMP that will recover if the IP packet loss is not unusually high. The secondary stream source is also assumed to be healthy for this failover to be successful.

Figure 5 Example Core MPEG Multiplexers (shown as DCM) and VDS Redundant Connections with the In-House Video Encoder

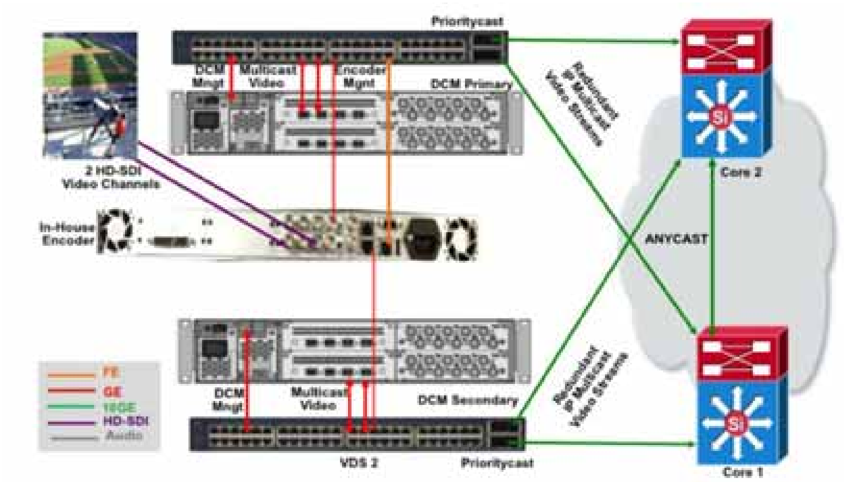

In this example (for RP implementations):

- An MPEG-4 encoder for an in-house feed is shown.
- The outputs from the encoder are connected to both MPEG Multiplexers.
- Each MPEG Multiplexer has GbE connections to the corresponding video distribution switch.
- Each video distribution switch has a 10 GbE connection to both core switches.
- A 1:1 active/active Prioritycast model is used between the video distribution switches and the core MPEG Multiplexers.
- All (*,G) pairs are sourced from the Primary source on MPEG Multiplexer1.
- The VDS2 ports connected to the MPEG Multiplexer2 will be activated only if there is a failure with the primary MPEG Multiplexer1 or VDS1.
- Upon the failure of MPEG Multiplexer1 or VDS1, the subnet will no longer be available, and the network will converge on MPEG Multiplexer2.
- In the core switches, unicast route is used for route convergence towards the Prioritycast RP.

Note: The above is an example and in some encoding cases, like when using the low-latency Harmonic CP9000 encoder, the encoder will connect directly to the VDS switches.

Encrypted Video Streams

Video streams can be encrypted in AES-128 format. Settings in the Cisco Vision Dynamic Signage Director allows to configure that as a global setting for the DMPs.

When a video feed is not encrypted, the decryption key passed to the DMPs will simply ignore it.

DMPs support AES-128 encryption.

Video Encoding

Video encoders are commonly used to take uncompressed source in-house video and convert it to the proper IPTV format and encapsulate in multicast frames.

Harmonic CP9000 Video Encoder

The Harmonic CP9000 is a HD ultra-low latency streaming encoder from Harmonic. After configuration per Harmonic's configuration guide, the output can be readily placed into the VDS as source multicast streams. Here are some of the highlights:

- HD and UHD formats
- AVC and HEVC (H.264/H.265) codecs
- Low latency depending on network performance from 266 ms – 566 ms end to end.

Note: Output from this encoder is IPTV and should not be passed through a MPEG Multiplexer since this will unnecessarily add latency to the stream. This encoder features dual GbE interfaces which should be connected to the VDS switches directly.

Service Provider Demarcation Switch Examples

The demarcation (demarc) switch is used with clear and encrypted IP feeds from a service provider. A few examples of the typical setup, including one used with DirecTV bundle will be shown here. Other configurations are possible.

Demarc Switch: Clear IP Feeds Connections

Figure 6 Demarc Connections–Clear IP Feeds

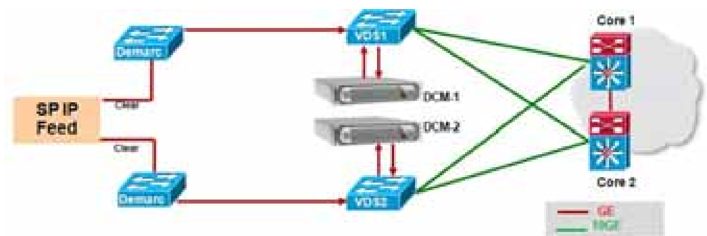

On ingress, each demarcation switch is connected to the source via GbE going into a VLAN.

The service provider typically floods multicast to the demarc switches. If the SP does not provide the demarc, then the feeds will terminate at the VDS1/2 directly.

On egress, each demarcation switch sends an identical copy of the feed to each of the video distribution switches via GbE.

Typically, there is L2 on the demarc, it's all VLAN/Trunk with IGMP snooping disabled on demarc for content VLANs.

If MPEG Multiplexers (shown as DCMs in Figure 6 on page 47) are not used, then the provider's L3 addressing must be known by the global/VRF routing table.

DirecTV / Technicolor Connections (North America)

The Technicolor COM3000 can accommodate up to six COM51 blades in the COM400 chassis, with each blade able to tune up to 23 high-definition channels as well as tuning DirecTV UHD content.

Figure 7 Connections

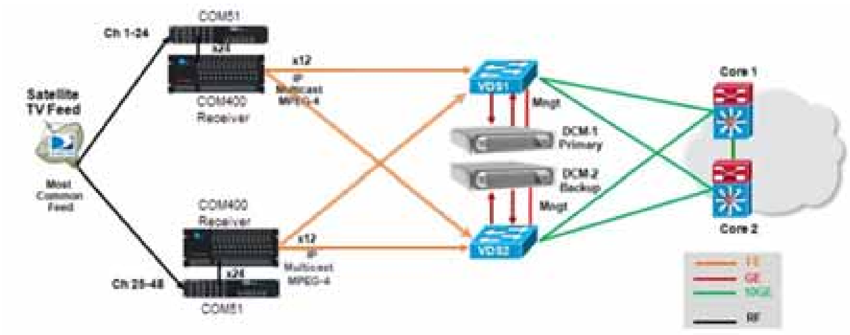

Typically no demarc switch is used in this configuration unless provided by the service provider, so the COM400 will terminate directly to the VDS1/2. Here's a link to the datasheet.

If MPEG Multiplexer (Figure 7 on page 48 1 as DCM) is used, then this and any SP feeds should be terminated on the multiplexer via L2 VLAN without IGMP snooping.

Video Monitoring

Our recommendation is to use a MPEG Analyzer for comprehensive monitoring.

For basic troubleshooting, sometimes a DMP off the demarcation switch or VDS switch for troubleshooting clear IP video sources at the service provider's drop-off point can be used. Ensure that POE and IP Services license are supported on primary demarcation or VDS switch to use DMP for troubleshooting. For baseline design or cost-sensitive venues with redundancy, plan connections and configurations to go on VDS.

In addition, a HD-SDI TV that supports input, audio and closed caption can be used to monitor video feed.

Figure 8 Monitoring incoming HD-SDI feeds via HD-SDI TV

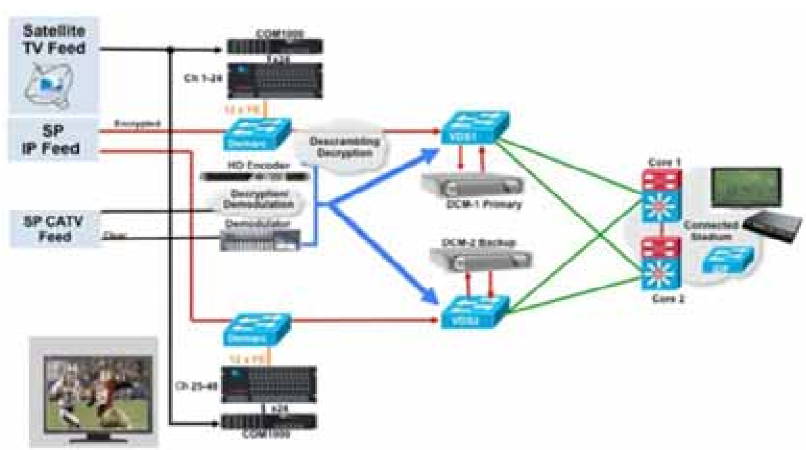

Appendix A: Standards

SDI

Serial digital interface (SDI) refers to a family of video interfaces standardized by SMPTE. For example, ITU-R BT.656 and SMPTE 259M define digital video interfaces used for broadcast-grade video. A related standard, known as high-definition serial digital interface (HD-SDI), is standardized in SMPTE 292M. SDI uses Non-Return-to- Zero Inverted (NRZI) encoding.

Interface cables use BNC connectors and RG-59 cables.

SMPTE

The Society of Motion Picture and Television Engineers (SMPTE), is an international professional association of engineers working in the motion imaging industries. An internationally recognized standards developing organization, SMPTE has over 400 standards that define the Recommended Practices and Engineering Guidelines for television, motion pictures, digital cinema, audio and medical imaging.

Table 1 Select SMPTE Standards

Standard	Name	Bit Rates	Example Video Formats
SMPTE 259M	SD-SDI	270 Mbps, 360 Mbps, 143 Mbps, and 177 Mbps	480i, 576i
SMPTE 344M		540 Mbps	480p, 576p
SMPTE 292M	HD-SDI	1.485 Gbps, and 1.485/1.001 Gbps	720p, 1080i
SMPTE 372M	Dual Link HD-SDI	2.970 Gbps, and 2.970/1.001 Gbps	1080p
SMPTE 424M	3G-SDI	2.970 Gbps, and 2.970/1.001 Gbps	1080p
SMPTE 2082	12G-SDI	12 Gbps	2160P60

SMPTE

Appendix B: Bill of Material

Standard Server Bill of Material

Note: The Bill of Material (BOM) only shows the server configuration and excludes the Cisco Vision Dynamic Signage Director, DMPs, and other system licenses.

In this BOM, RAID 5 setting instead of RAID 10 is shown, but this is due to limited selection options in the ordering tool. During deployment, use RAID 10, which can be enabled via the MegaRAID BIOS directly (not through CIMC).

Table 1 Bill of Material for Standard Server

Part Number	Description	Qty
UCSC-C220-M5SX	UCS C220 M5 SFF 10 HD w/o CPU, mem, HD, PCIe, PSU	1
CON-SNT-C220M5SX	SNTC 8X5XNBD UCS C220 M5 SFF 10 HD w/o CPU, mem, HD, PCIe, P	1
UCS-MR-X16G1RS-H	16GB DDR4-2666-MHz RDIMM/PC4-21300/single rank/x4/1.2v	4
UCS-SD400G123X-EP	400GB 2.5in Enterprise Performance 12G SAS SSD(3X endurance)	6
UCS-SD-32G-S	32GB SD Card for UCS servers	2
CIMC-LATEST	IMC SW (Recommended) latest release for C-Series Servers	1
UCSC-PSU1-1050W	Cisco UCS 1050W AC Power Supply for Rack Server	2
CAB-9K12A-NA	Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	2
UCSC-RAILB-M4	Ball Bearing Rail Kit for C220 & C240 M4 & M5 rack servers	1
UCS-MSTOR-SD	Mini Storage Carrier for SD (holds up to 2)	1
UCSC-BBLKD-S2	UCS C-Series M5 SFF drive blanking panel	4
UCSC-HS-C220M5	Heat sink for UCS C220 M5 rack servers 150W CPUs & below	1
CBL-SC-MR12GM52	Super Cap cable for UCSC-RAID-M5 on C240 M5 Servers	2
UCSC-SCAP-M5	Super Cap for UCSC-RAID-M5, UCSC-MRAID1GB-KIT	1
UCS-CPU-6128	3.4 GHz 6128/115W 6C/19.25MB Cache/DDR4 2666MHz	2
UCSC-RAID-M5	Cisco 12G Modular RAID controller with 2GB cache	1
UCS-SID-INFR-UNK	Unknown	1
UCS-SID-WKL-UNK	Unknown	1
R2XX-RAID5	Enable RAID 5 Setting	1
UCSC-SW-C220M5-P01	Performance Optimized setting for C220 M5 servers	1
VMW-VCS-STD-1A=	VMware vCenter 6 Server Standard, 1 yr support required	2
CON-ISV1-VCXSTD1A	VCenter Server STD for vSphere 1-Inst; ANNUAL List 1-YR Reqd	2
UCS-VMW-TERMS	Acceptance of Terms, Standalone VMW License for UCS Servers	2

Small Server Bill of Material

Note: The BOM only shows the server configuration and excludes the Cisco Vision Dynamic Signage Director, DMPs, and other system licenses.

Table 2 Bill of Material for Small Server

Part Number	Description	Qty
UCSC-C220-M5SX	UCS C220 M5 SFF 10 HD w/o CPU, mem, HD, PCIe, PSU	1
CON-SNT-C220M5SX	SNTC 8X5XNBD UCS C220 M5 SFF 10 HD w/o CPU, mem, HD, PCIe, P	1
UCS-MR-X16G1RS-H	16GB DDR4-2666-MHz RDIMM/PC4-21300/single rank/x4/1.2v	4
UCS-SD240GM1X-EV	240GB 2.5 inch Enterprise Value 6G SATA SSD	6
UCS-SD-32G-S	32GB SD Card for UCS servers	2
CIMC-LATEST	IMC SW (Recommended) latest release for C-Series Servers	1
UCSC-PSU1-1050W	Cisco UCS 1050W AC Power Supply for Rack Server	2
CAB-9K12A-NA	Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	2
UCSC-RAILB-M4	Ball Bearing Rail Kit for C220 & C240 M4 & M5 rack servers	1
UCSC-SW-C220M5-P01	Performance Optimized setting for C220 M5 servers	1
UCS-SID-INFR-UNK	Unknown	4
UCS-SID-WKL-UNK	Unknown	1
UCS-MSTOR-SD	Mini Storage Carrier for SD (holds up to 2)	2
UCSC-BBLKD-S2	UCS C-Series M5 SFF drive blanking panel	1
UCSC-HS-C220M5	Heat sink for UCS C220 M5 rack servers 150W CPUs & below	2
CBL-SC-MR12GM52	Super Cap cable for UCSC-RAID-M5 on C240 M5 Servers	1
UCSC-SCAP-M5	Super Cap for UCSC-RAID-M5, UCSC-MRAID1GB-KIT	1
UCS-CPU-6128	3.4 GHz 6128/115W 6C/19.25MB Cache/DDR4 2666MHz	1
UCSC-RAID-M5	Cisco 12G Modular RAID controller with 2GB cache	1
R2XX-RAID10	Enable RAID 10 Setting	1
VMW-VCS-STD-1A=	VMware vCenter 6 Server Standard, 1 yr support required	2
CON-ISV1-VCXSTD1A	VCenter Server STD for vSphere 1-Inst; ANNUAL List 1-YR Reqd	2
UCS-VMW-TERMS	Acceptance of Terms, Standalone VMW License for UCS Servers	2