CHAPTER 2

Network Design

This chapter describes the different network topologies for the Cisco TV CDS, the different network connections of the CDS servers, the CDS workflow, and network configuration considerations. The topics covered in this chapter include:

Overview

The TV CDS enables cable operators and multiple service operators (MSOs) to offer VOD and MediaX services to consumer customers over their existing hybrid fiber coaxial (HFC) network, with existing next-generation digital STBs. The TV CDS solution uses a Gigabit Ethernet (GE) transport network from the headend to the distribution hub, where the HFC network terminates.

TV CDS grows seamlessly from a single server implementation to multiple servers. As growth continues, TV CDS allows operators to install distributed servers to address concentrations of subscribers while leaving content ingest and management centralized.

Stream Groups can be distributed close to the subscriber and linked back to the central Vault locations by way of the Cisco Cache Control Protocol (CCP). Cisco CCP automatically ensures that any new content that is required by a customer edge device is transferred within a maximum of a 250-millisecond delay to the appropriate edge location; as a result, all content appears local to each edge site, even though most content is stored at the central Vault location.

The TV CDS offers different configurations with regards to network topology, business management systems (BMSs), and streaming modes.

CDS with Vaults and Streamers

In a TV CDS with Vaults and Streamers, MPEG-2 transport stream (TS) video is stored on the Vaults with the associated trick-mode files. Content is transported from the Vaults to the Streamers as needed, by using CCP over Gigabit Ethernet networks. Content is sent unicast from the Streamers and delivered

to the quadrature amplitude modulation (QAM) devices over Gigabit Ethernet or asynchronous serial interface (ASI), and then modulated onto the HFC plant to the subscriber's set-top box (STB) for viewing.

CDS with ISVs

For the smallest networks, Cisco packages the CDS in a single server, the Integrated Streamer-Vault (ISV), offering a solution for VOD services with large content libraries but small stream counts.

In a TV CDS with ISVs, MPEG-2 TS video is stored on the ISV servers with the associated trick-mode files. Content is sent unicast from the ISVs and delivered to the QAM devices over a Gigabit Ethernet network, and then is modulated onto the HFC plant to the subscriber's STB for viewing.

CDS with Caching Nodes

For larger networks, Cisco offers the CDS with Caching Nodes in the Virtual Video Infrastructure (VVI). In a VVI, Caching Nodes are the intermediary fill source for Streamers, which removes a large portion of the distribution traffic from the Vaults.

In a TV VVI, MPEG-2 TS video is stored on the Vaults with the associated trick-mode files. Content is transported from the Vaults to the Caching Nodes as needed, by using CCP over Gigabit Ethernet networks. Content is distributed from the Caching Nodes to the Streamers as needed, by using CCP over Gigabit Ethernet networks, or by using HTTP over Gigabit Ethernet networks. Content is sent unicast from the Streamers and delivered to the QAM devices over a Gigabit Ethernet network, and then is modulated onto the HFC plant to the subscriber's STB for viewing.

TV CDS Topologies and VVI Topologies

The TV CDS (using Vaults and Streamers, or ISVs) and the TV VVI (using Vaults, Caching Nodes, and Streamers), supports centralized, decentralized, and hybrid Gigabit Ethernet network designs. Because the use of Vaults and Streamers separates storage from streaming, streaming requirements can be satisfied on an "as needed" basis and the streaming can be centralized or distributed among multiple locations. Caching Nodes separate the ingest and storage of content from the distribution of content, offering greater flexibility and network efficiency.

The TV CDS topology and TV VVI topology can change with the evolving needs of the system operator. If the need to decentralize becomes evident, you can move the Streamers or Vaults to remote hubs without disrupting service. The VVI offers additional flexibility in designing your network. Vaults can be centrally located at a national network, and content may be classified by market (city, state, or a broader region) depending on the AMS or BMS used. Caching Nodes can be located centrally, or distributed closer to the regional networks where the Streamers are located. Using Caching Nodes in the network design takes the distribution traffic off the network backbone.

Caution All Cisco servers are connected through a switch. Because all Vaults, CCP Streamers, and Caching Nodes in the same array exchange heartbeat messages through the cache interfaces, it is important to ensure there is enough bandwidth among switches involved in delivering cache traffic, as well as to support the same aggregated amount of traffic on all cache interfaces.

Note When using ISVs, with the Vault and Streamer functions contained in one server, the only topology possible is centralized.

Centralized Topology

In a centralized topology, both Vault and Streamer servers are located in either a single video headend or a remote hub. This is the right solution for certain situations, for instance very small starting systems or where a large amount of bandwidth is available. A centralized topology has advantages in reducing operational cost by placing equipment in one physical location. Figure 2-1 illustrates the centralized topology for Vaults and Streamers.

Figure 2-1 Centralized Topology with Vaults and Streamers

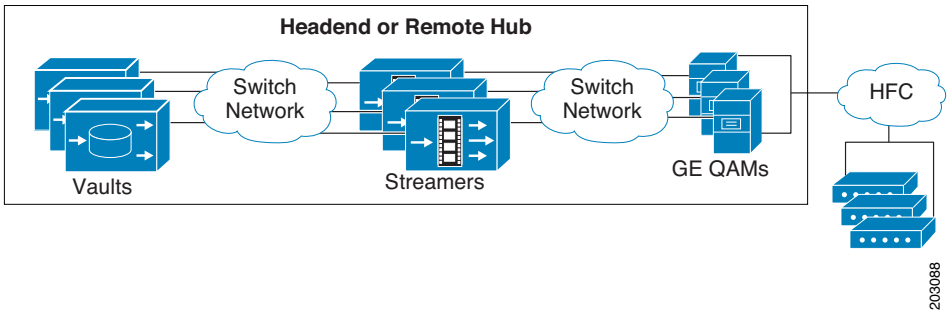

Figure 2-2 illustrates the centralized topology for ISVs.

Figure 2-2 Centralized Topology with ISVs

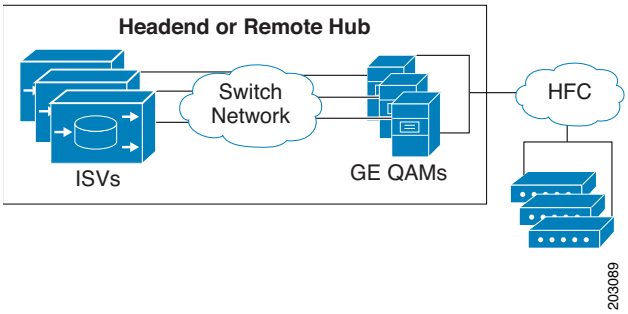

Figure 2-3 illustrates the centralized topology for a VVI.

Figure 2-3 Centralized Topology with Caching Nodes

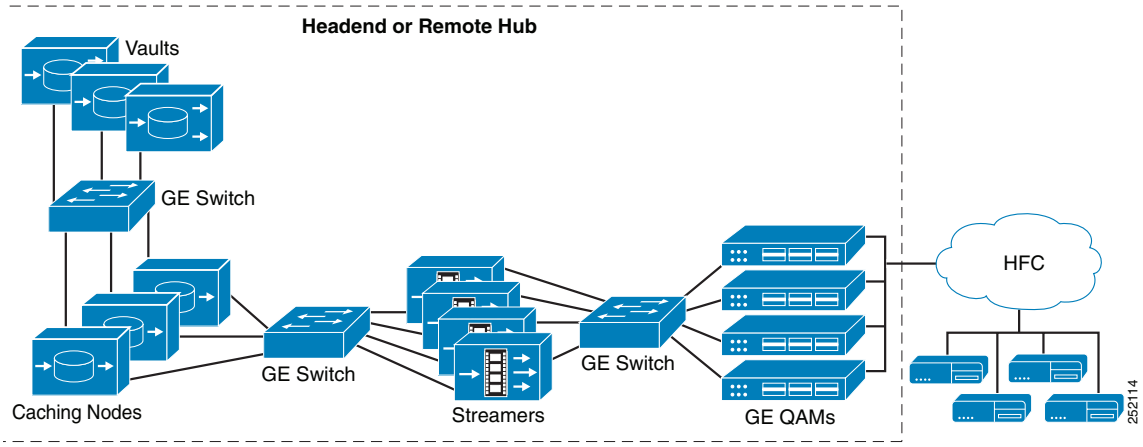

Decentralized Topology

The decentralized topology is a hub-and-spoke topology between the headend site and multiple hub sites, where the Vault servers are located at the headend and the Streamer servers are in the hub sites. For a VVI, a decentralized topology provides a three-tiered approach by having the Vaults located in the headend, the Caching Nodes in intermediary sites, and the Streamers in the hub sites. The decentralized topology works well for distributing Streamer Groups close to subscribers. A decentralized topology has advantages in reducing the amount of long-haul fiber transport bandwidth needed—typically by a factor of ten or better. Figure 2-4 illustrates the decentralized topology.

Figure 2-4 Decentralized Topology

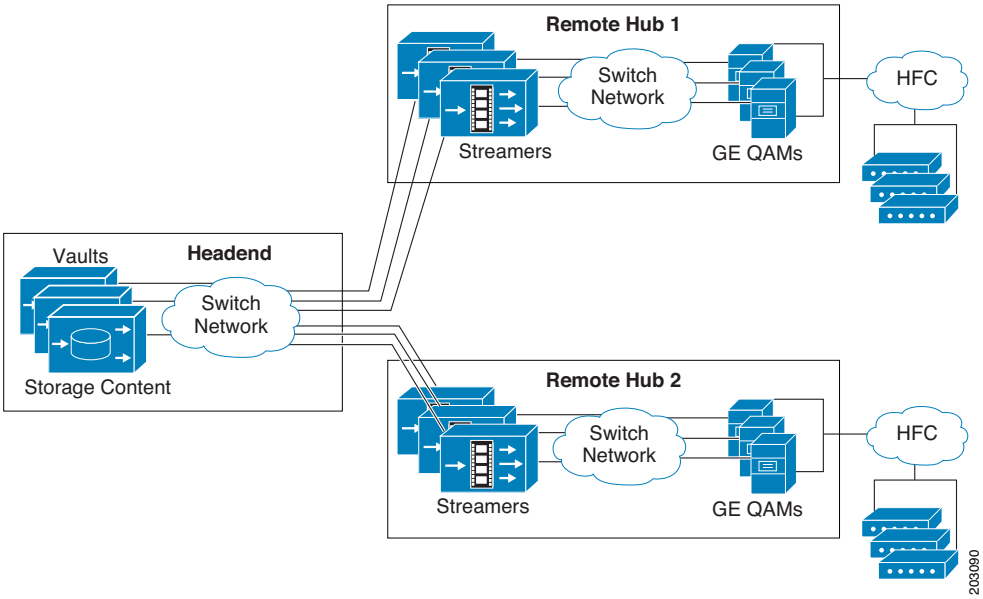

Figure 2-5 illustrates the decentralized topology with Caching Nodes.

Figure 2-5 Decentralized Topology with Caching Nodes

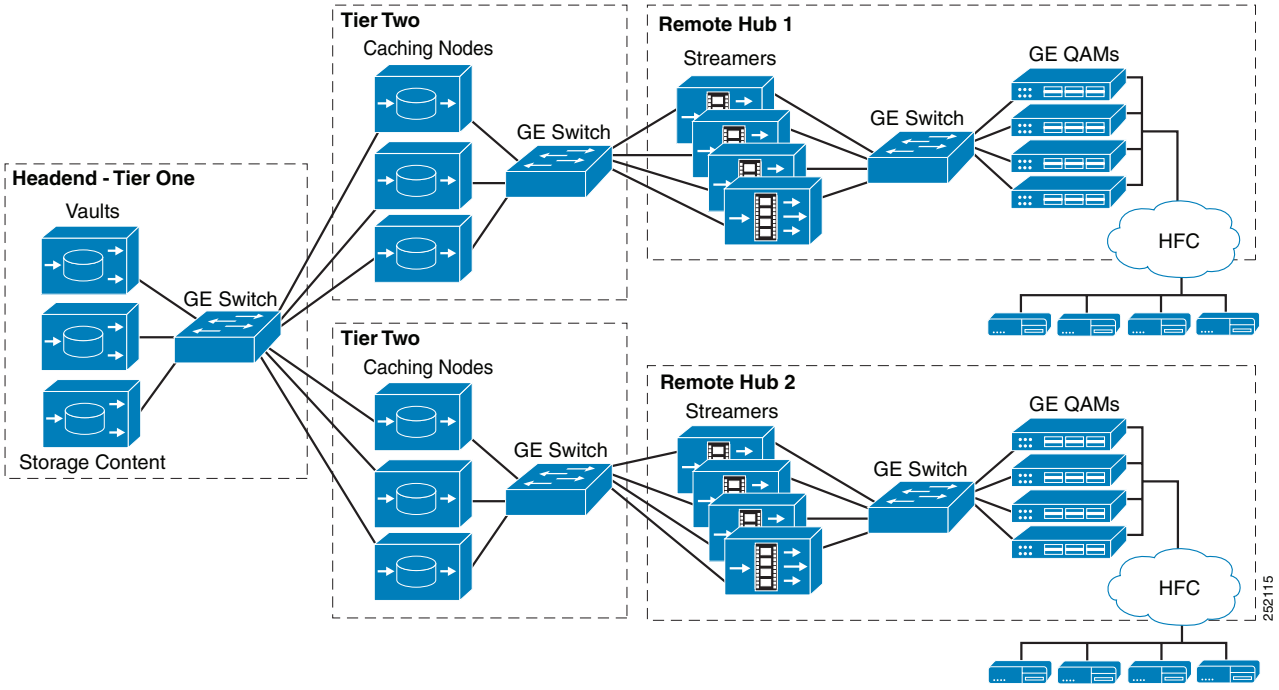

Hybrid Topology

In a hybrid topology, the Vault servers and backup Streamer servers are located at the headend, with the active Streamers at a remote hub site. If the remote hub site goes down, the Streamers at the headend take over. A hybrid topology blends the advantages of centralized and decentralized topologies that is based on needs of the system implemented. Figure 2-6 illustrates the hybrid topology.

Figure 2-6 Hybrid Topology

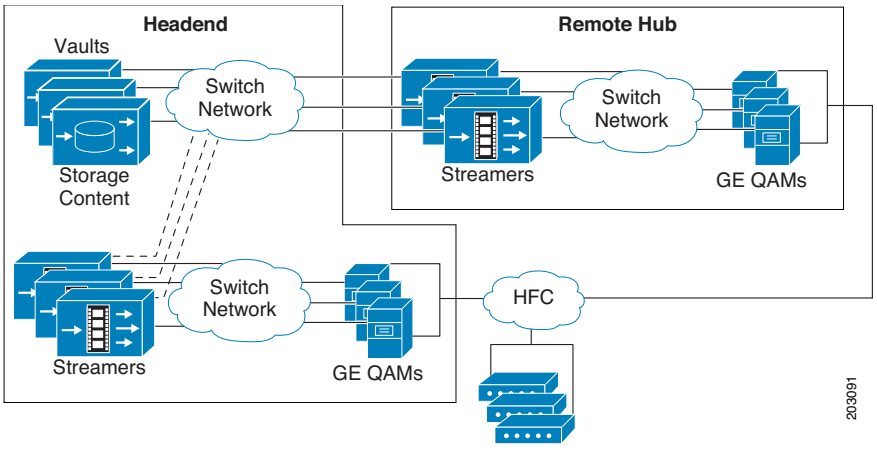

Figure 2-7 illustrates the hybrid topology with Caching Nodes.

Figure 2-7 Hybrid Topology with Caching Nodes

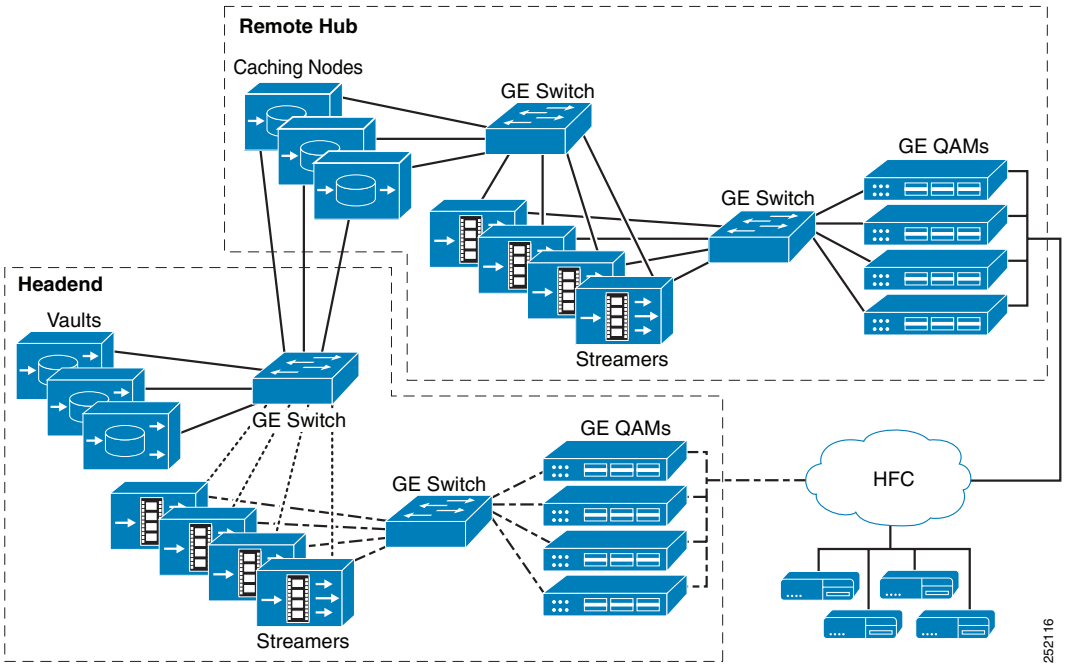

TV VVI Management

The TV VVI offers two types of management, centralized and split-domain.

In a CDS, Streamers cannot communicate with Streamers in other groups. In a VVI, Streamers in other groups can communicate with each other on an as-needed basis.

All Vaults, Streamers, and Caching Nodes are identified by an array ID, a group ID, and a server ID. In the CDSM GUI, the array ID identifies servers that are part of the same system. The group ID identifies servers that are part of the same group (Vault Group, Cache Group, and Stream Group), and the server ID is a unique number that identifies the server. Table 2-1 lists the CDSM GUI ID names and maps them to the CServer names in the setupfile and .arroyorc files.

Table 2-1 ID Names in the CDSM GUI and CServer Files

CDSM GUI ID Name	CServer Files ID Name
Array ID on the Array Name page	groupid
Group ID on the Server-Level pages	groupid
Stream Group ID on the Server Setup page	arrayid
Cache Group ID on the Server Setup page	arrayid
Vault Group ID on the Server Setup page	arrayid
Stream Group ID on the Configuration Generator page	arrayid

Note VVI is available in an RTSP environment, and in an ISA environment that uses the Shared Content Store feature and CCP Streamers.

Centralized Management

Centralized management uses one Virtual Video Infrastructure Manager (VVIM) to manage the Vaults, Caching Nodes, and Streamers in a VVI.

Note Centralized management is only in an RTSP environment.

Split-Domain Management

Split-domain management uses one VVIM to manage the domain of Vaults and Caching Nodes, and separate managers, the Stream Managers, to manage each domain of Streamers. The Stream Managers communicate with the VVIM over port 80. If port 80 is not open for communication, the managers cannot communicate with each other and configuration settings need to be uploaded to the Stream Managers from information downloaded from the VVIM.

In a split-domain VVI that uses HTTP for communication between the Caching Nodes and Streamers, and in a split-domain VVI that uses CCP in an RTSP environment, the databases for each domain are separate. The information stored in each database is not shared with the servers in the other domains.

In an ISA environment with a split-domain VVI that uses CCP for communication between the Caching Nodes and Streamers, the database is replicated among all servers in the Vault/Cache domain and the Stream domains. Because the VVI allows intercommunication among different Cache Groups and Stream Groups when CCP Streamers are used, the server ID and group ID must be unique across the system.

Note Split-domain management is supported in an RTSP environment and an ISA environment with the Shared Content Store feature and CCP Streamers.

CDS Workflow

Content is ingested and stored in the Vault array. The Vault array consists of two or more Vault Groups, which in turn consists of two or more Vaults that are either colocated or distributed to multiple locations across an Ethernet network. Content ingest is initiated by the backoffice based on a subscriber request, and based on schedule or barker channel content. Manual ingest, which is operator initiated, is also offered as an optional feature.

As the content is ingested into the Vault, any necessary trick-mode files are created. The content and trick-mode files are then mirrored within the same Vault or across the Vault array. The replication of content allows for data recovery should a Vault undergo a failure.

Content is delivered from the Vault array to the Streamer array in response to cache-fill calls from the Streamers in order to fulfill subscriber requests for VOD content. Content is also distributed across the network in response to scheduled or barker stream content fulfillment.

If a VVI is deployed, content is delivered from the Vault Group to the Cache Group in response to cache-fill calls from the Streamers. The Caching Nodes are explained in more detail in the "Caching Node Workflow" section on page 2-9.

Within the Streamer array are one or more Stream Groups. The following section describes how the Stream Groups deliver streams to the subscriber STBs.

Note All servers can be on different subnetworks. However, given current backoffice restrictions, the externalized IP address is constrained to migrate among servers on the same subnetwork. This means the content store server in an Interactive Services Architecture (ISA) environment can migrate only among Vaults that are on the same subnet, and the Setup and Control servers can migrate only among Streamers on the same subnet.

Streamer Workflow

A Stream Group is a configurable group of Streamers that are designated to serve specified QAM devices, and subsequently, specific service groups. From a session setup and control perspective, there are three logical types of servers in a Stream Group:

- Setup server
- Control server
- Play server

The Setup and Control servers have both a primary and a backup server. The primary server services all messages, while the backup server simply maintains states. If a primary server is unreachable, the backup server takes over control and creates another backup server. Thus, there is always a primary and backup pair of servers for setup and control. The Play server does not have a backup server. However, the Control server selects a new Play server in the event of a failure of the existing Play server.

Note The ability to have both a primary and backup server depends on the number of Streamers in the Stream Group.

The Setup and Control server IP addresses are configurable. For an ISA environment, the Setup IP address is the same as the Stream Master IP address. For RTSP, the Setup server and Control server must be the same server. For both ISA and RTSP environments, the Stream Service selects a Streamer in the Stream Group to be the Setup server, and another Streamer (sometimes the same Streamer) to be the Control server.

Setup Server

A Streamer designated as the Setup server interfaces with the backoffice and forwards the setup messages to the appropriate Stream Group that is assigned to the destination service group. One Streamer in the Stream Group that is collocated with the backoffice server is assigned as the primary Setup server. The Setup server receives the setup request from the backoffice and maps the service group.

The Setup server returns the IP address of the Control server, and the STB issues subsequent control messages to this IP address

Control Server

The Control server assigns requests to specific Streamers and dynamically migrates streams between Streamers based upon changes in stream states (for example, content splice boundaries, maintenance trickle down, or server failures). One server in the Stream Group is assigned as the primary Control server. The Control server runs the Lightweight Stream Control Protocol (LSCP) proxy in an ISA environment and the Real-Time Streaming Protocol (RTSP) proxy in an RTSP environment.

For each and every setup message received from the backoffice, a CCP message is generated and sent to the Control server. In the initial setup request, the Control server receives the setup parameters but does not choose a Play server. Once a control message is received from the STB, the Control server gets performance information (for example, server load) from the potential Play servers within the Stream Group and sends a CCP message to the best candidate. Subsequent control messages, whether from the STB or from the Setup server, are forwarded to the chosen Play server.

Play Server

The Play server is the Streamer that is assigned to play the stream. This Streamer acquires the content, whether in RAM, a local disk, or a Vault, and ensures guaranteed service delivery of the stream. Every Streamer in a Stream Group is a possible candidate to be the Play server.

Caching Node Workflow

A Cache Group is a configurable group of Caching Nodes that serve content to specified Stream Groups. When a content request is received by a Streamer, the Streamer first checks to see if the content is stored locally, which includes DRAM, disk cache, and Streamers in the same Stream Group. Content on the Streamers is always the most popular content, so user requests are generally served from local storage.

Streamers send cache-fill calls to remote servers for content that is not found locally. The remote servers can be Streamers in other Stream Groups and Caching Nodes in Cache Groups. The cache-fill source selected, whether another Streamer or a Caching Node, is based on the network capacity and fill-source capacity (disk and memory), as well as the preference configured for that group of servers. Caching Nodes could respond to the request with a message stating the content is not currently cached, but there are other fill sources the Caching Nodes can contact (Caching Nodes in other Cache groups, and Vaults).

The Caching Nodes use CCP to communicate with the Vaults, and use either CCP or HTTP to communicate with Streamers.

Note HTTP Streamers are supported in an ISA environment and an RTSP environment. However, HTTP Streamers are not supported in an ISA environment using the Shared Content Store feature.

HTTP Streamers

HTTP can be used for communication between the Caching Nodes and the Streamers. The HTTP Streamer communicates with a proxy for locating a fill source and pulling content.

A locate service serves as a proxy for a group of Caching Nodes and Vaults. The service is accessed by way of a highly available secondary IP address hosted by the Caching Node. The secondary IP address is bound to a fill port (Locate Port).

HTTP Streamers request content by means of HTTP GET requests to the proxy service (the server with the locate service). The proxy server checks its own storage and peer fill sources (servers in the same group) for the content using extended-CCP. If the content is found, the best source is chosen based on capacity and a redirect response is sent to the chosen server. If the content is not found, a cache-fill request is sent to the remote servers.

Once the best server is chosen to send the content to the HTTP Streamer, a single cache-fill port on that server is chosen for the HTTP transfer of the content. This is different from CCP transfers, which could potentially use all cache-fill ports to deliver the content.

HTTP Locate Port

With respect to resiliency, the Locate Port service is similar to the Setup and Control servers,. The primary server of the Locate Port service has the locate port IP address bound to an interface. The backup server becomes the primary if the primary fails.

Peer Caching Nodes advertise among themselves about the ability to host the HTTP Locate Port Service, this includes primary, backup, available, and not usable states. Available means the Caching Node can be either a primary or backup if needed. Not usable means that the server cannot host the service; for the HTTP Locate Port this typically means that there are no usable network ports for the service.

A dedicated network port on the Caching Node is used solely for the HTTP Locate Port service. The primary server determines service availability based on the link status of the dedicated network port. Failover of the service occurs if the network port loses link status. A reestablished link results in the server becoming available.

CCP Streamers

The CCP Streamers use CCP to communicate with the Caching Nodes. They do not use the proxy address that was assigned to the Locate Port for HTTP Streamers. CCP Streamers load-balance locate requests across fill sources.

The Streamer or Caching Node sends a locate-and-request message from the proxy server. The Proxy server sends a message to the best source to fill the request.

Streamers or Caching Nodes needing content first query peer sources (servers within the same group). Streamers also query local Streamers, if the content is not found, then a request to the remote sources is sent. Remote sources are queried based on a preference list. Sources are grouped and preferences are assigned for each group.

Vault Workflow

The Vaults ingest content using three different methods:

- FTP pull
- FTP push
- Live capture of MPEG-2 transport streams over UDP

With FTP pull, the original content is kept on an FTP server (catcher), for a period of time and mechanisms are in place to restart ingests until they have successfully completed.

With FTP push, only a window of data is buffered by a device that grooms the live (broadcast) feed and pushes the data to the Vault.

With live capture over UDP, the Vault captures the live multicast feed directly, and a failed ingest would result in the recording not being recoverable. It is therefore very important for the Vault not to drop any packets because of its own resource constraints. For recommended settings, see the "Configuring the Vault Ingest Interfaces for Live Capture over UDP" section on page 4-60.

BMS Considerations

The TV CDS integrates with Interactive Services Architecture (ISA) used in business management systems (BMSs) such as the Tandberg OpenStream and the RTSP used in BMSs such as ARRIS nABLE, as well as in environments that are a combination of both ISA and RTSP. The BMS determines the roles and responsibilities of the TV CDS.

OpenStream ISA Integration

The OpenStream BMS is built on Common Object Request Broker Architecture (CORBA) and provides naming and notification services. The Naming Service allows the TV CDS to locate objects in the system such as content, equipment, assets, and so on. The Notification Service allows the TV CDS to listen for important events in the system as well as to send events to the OpenStream BMS and other components in the system.

Figure 2-8 illustrates how the TV CDS integrates with the OpenStream BMS.

Figure 2-8 TV CDS Integration into the OpenStream BMS

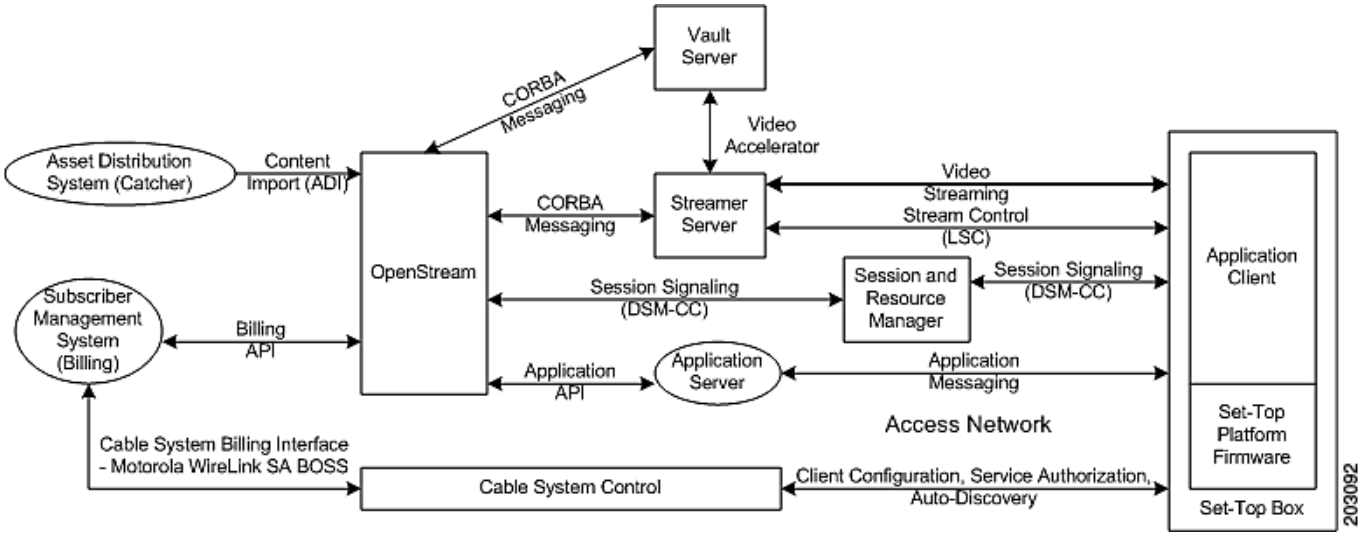

Streaming Mode

OpenStream uses a session-based approach to handle resource requirements and allocation. In the course of setting up a session, a QAM device is specified that has available capacity and connectivity to the Cisco Streamer and the STB requesting the service. Typically, the Session and Resource Manager (SRM) is responsible for the allocation of network resources. OpenStream uses the Digital Storage Media-Command and Control (DSM-CC) session management protocol to request resources from the SRM.

When using Gigabit Ethernet for streaming, OpenStream communicates with the SRM to negotiate network resources and allocation for sessions.

When using Asynchronous Serial Interface (ASI) for streaming, the Cisco Streamer performs the role of the SRM by managing and allocating the access network resources and providing this information to the OpenStream BMS.

Shared Content Store

The Shared Content Store (SCS) feature works with a single, centralized AMS and catcher, through which all initiation for content ingest and content deletion is sent. The SCS handles ingest and deletion requests from multiple backoffices by way of the central AMS. The scenario of backoffices independently ingesting and deleting content through their local AMS is not supported.

Note The Shared Content Store feature requires the Virtual Video Infrastructure feature with Caching Nodes.

Figure 2-9 shows a high-level view of the SCS and a single, centralized AMS for multiple video hub offices (VHOs). A VHO is a local deployment that includes the video backoffice, Streamers, application servers, QAM devices, and other headend equipment.

Figure 2-9 Shared Content Store For Multiple Video Headends

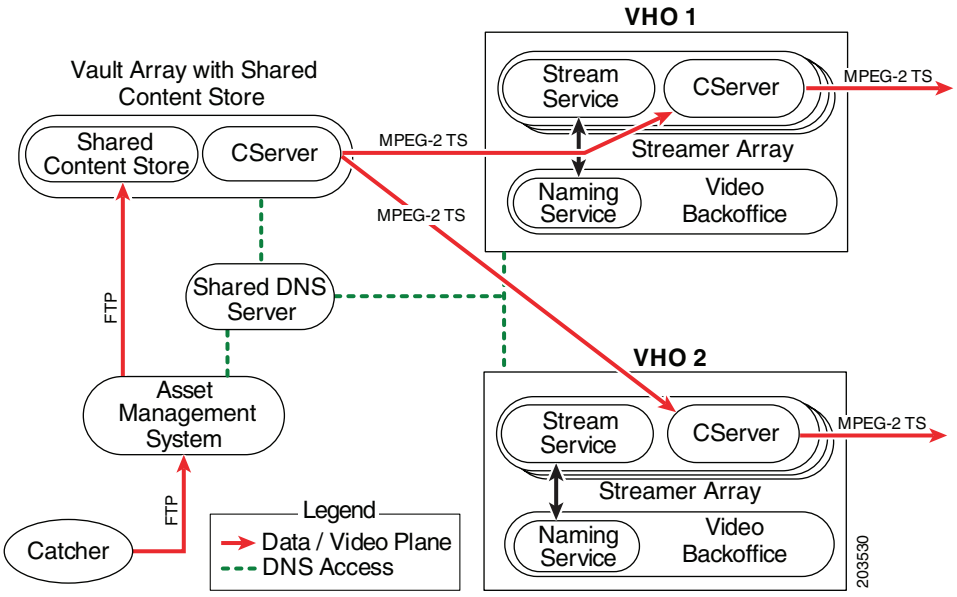

The DNS server typically runs on the BMS server. The Naming Service is part of the video backoffice (VBO). All CORBA components, including the AMS, Stream Service, and Content Store, need to register with the Naming Service. The catcher receives or "catches" new content assets from an external communication device such as a satellite transmission or FTP server. Once the package is received completely by the catcher, it sends the package by way of FTP to the AMS. The package consists of video and image content assets, as well as ADI metadata.

Following are the requirements for the SCS feature:

- A single, shared DNS server with all devices registering their hostnames to it. A central, shared DNS is required to resolve multiple Naming Services belonging to the different VHOs.
- Hostnames must be unique for all devices. This is required for the Naming Service discovery.
- Each VHO has its own Naming Service to which the ISA components of the VHO register.

- AMS controls the ingest and deletion of content.
- The Vault array has one SCS.
- SCS registers with each Naming Service.

A VVI with SCS must be initialized in the following order:

1. The shared DNS server needs to be up and running before starting up the shared AMS, SCS, and VHO devices.
2. SCS successfully registers with the Naming Service for each VBO.
3. Each VHO Stream Service registers with its respective Naming Service.

Ingesting Content with the Shared Content Store

Upon receiving the content package, the AMS schedules it for ingest by informing the Package Factory in each participating VBO of the content package, and passing the pertinent information (the ADI metadata, the URL where the content package can be accessed in the AMS, and the verb *ingest*).

The SCS creates one interoperable object reference (IOR) for each content package. The IOR is returned to all video backoffice (VBO) Package Factories that request it, including any that requested it at the time the IOR was being created.

Deleting Content with the Shared Content Store

To delete content that was ingested for more than one VBO, the AMS is used to send the *export package delete* request to each VBO. The content is deleted from the Vault array only when all VBOs have requested the deletion. If one or more VBOs have not requested that the content be deleted, the content remains in the Vault array.

nABLE Integration

The nABLE BMS uses a combination of eXtensible Markup Language (XML) over Hypertext Transfer Protocol (HTTP) and RTSP for communication between nABLE Headquarters (HQ) and Real-time (RT) components and the CDS. The HQ communicates file-related requests by using XML/HTTP to the Vault server, as well as server status information requests to both the Streamer and Vault servers. The RT communicates with the Streamer server by way of RTSP to establish session setups for multiple, interchangeable VOD flows (RTSP or DSM-CC).

Note Currently, configuring the CDS for integration with the nABLE BMS is performed by Cisco field engineers. For more information on integration of the CDS with the nABLE BMS, contact the Cisco technical support department.

Figure 2-10 illustrates how the CDS integrates with the nABLE BMS.

Figure 2-10 TV CDS Integration into the nABLE BMS

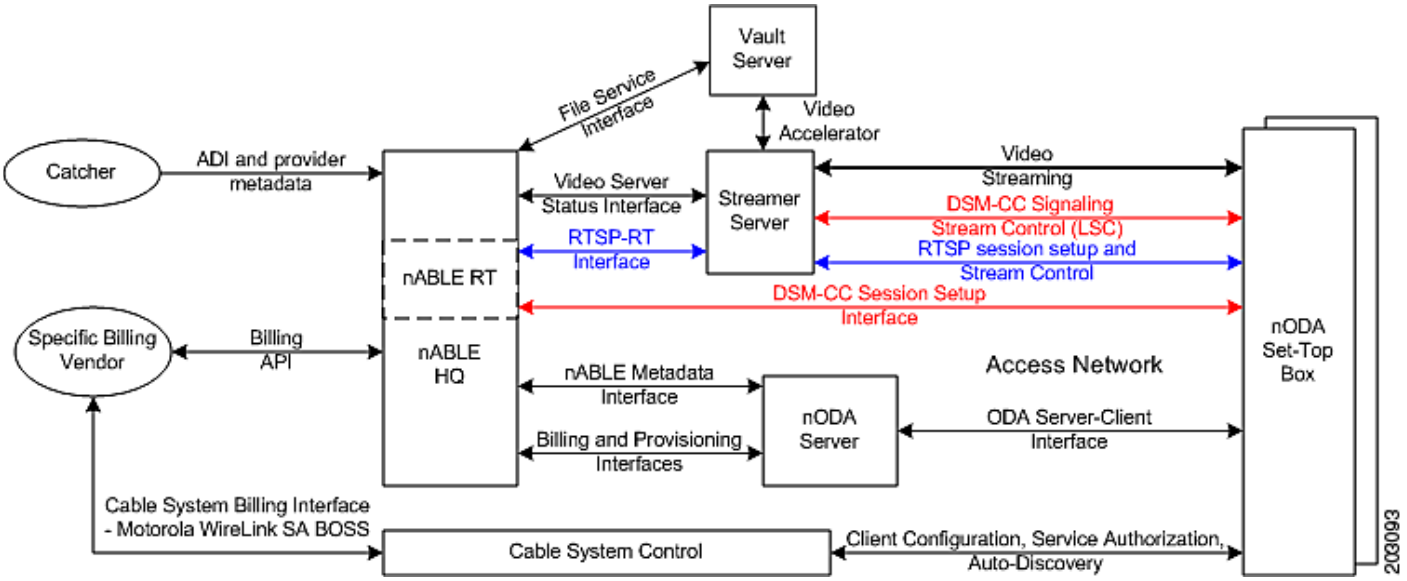

Network Connections

The network connections for a TV CDS with Vaults and Streamers, a TV CDS with ISVs, and a TV VVI with Caching Nodes all have different network connections. Table 2-2 lists the different required interfaces for each CDS server. The interfaces are described in the following sections. Figure 2-11 illustrates a TV CDS with Vaults and Streamers. Figure 2-12 illustrates a TV CDS with ISVs. Figure 2-13 illustrates a TV VVI with Caching Nodes.

Table 2-2 CDS Interfaces

Interface	Vault	Streamer	ISV	Caching Node
Management	1	1	1	1
Ingest	1	—	1	—
Cache	1 to 8	1 to 13	1 to 4[1]	1 to 12
Stream	—	1 to 13	1 to 4	—

1. The cache interfaces on an ISV are used for content mirroring among ISVs.

Note Table 2-2 lists the mandatory interfaces for each CDS server. If HTTP Streamers are used in a VVI, each Caching Node must have one interface designated as the Locate interface. Stream Control is an optional interface function. For more information, see the "Configuring the Interfaces" section on page 4-51.

Figure 2-11 shows the different logical networks of a CDS consisting of Vaults and Streamers. The ingest network receives content from the content source by way of an FTP staging server or FTP catcher and the content is ingested by the Vaults. The management network consists of communication between the CDSM and the BMS, as well as communication to the Vaults, Streamers QAM devices, and STBs. The cache network consists of Vaults and Streamers.

Figure 2-11 Vault and Streamer Network Connections

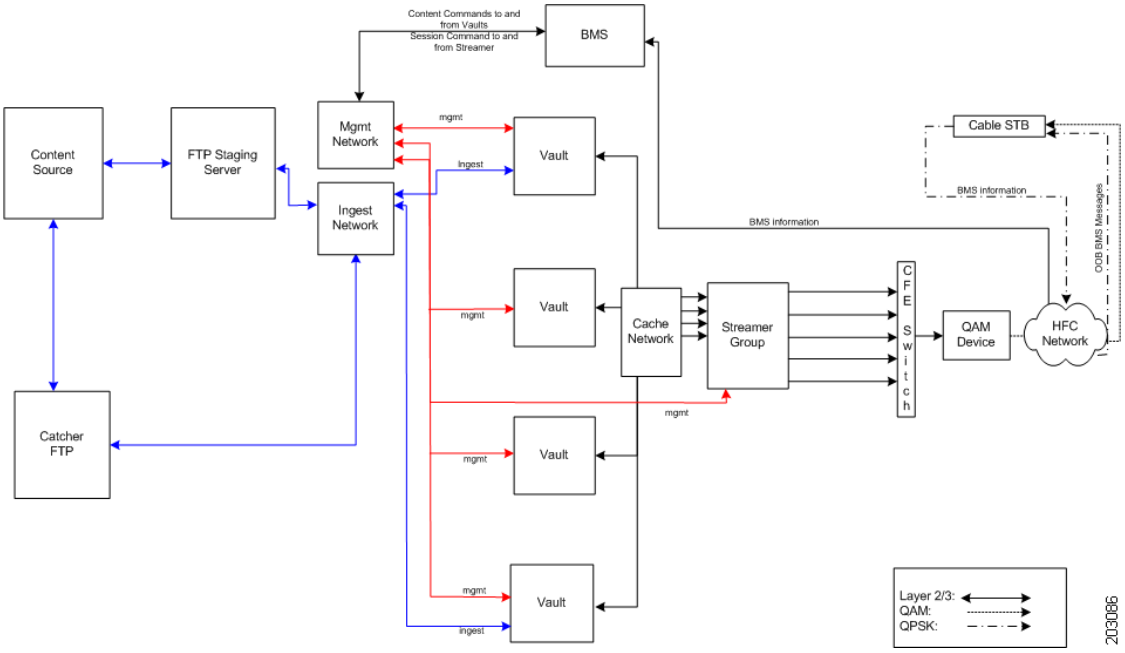

Figure 2-12 shows the different logical networks of a CDS consisting of ISVs. The ingest network receives content from the content source by way of an FTP staging server or FTP catcher and the content is ingested by the ISVs. The management network consists of communication between the CDSM and BMS, as well as communication to the ISVs, QAM devices, and STBs.

Figure 2-12 ISV Network Connections

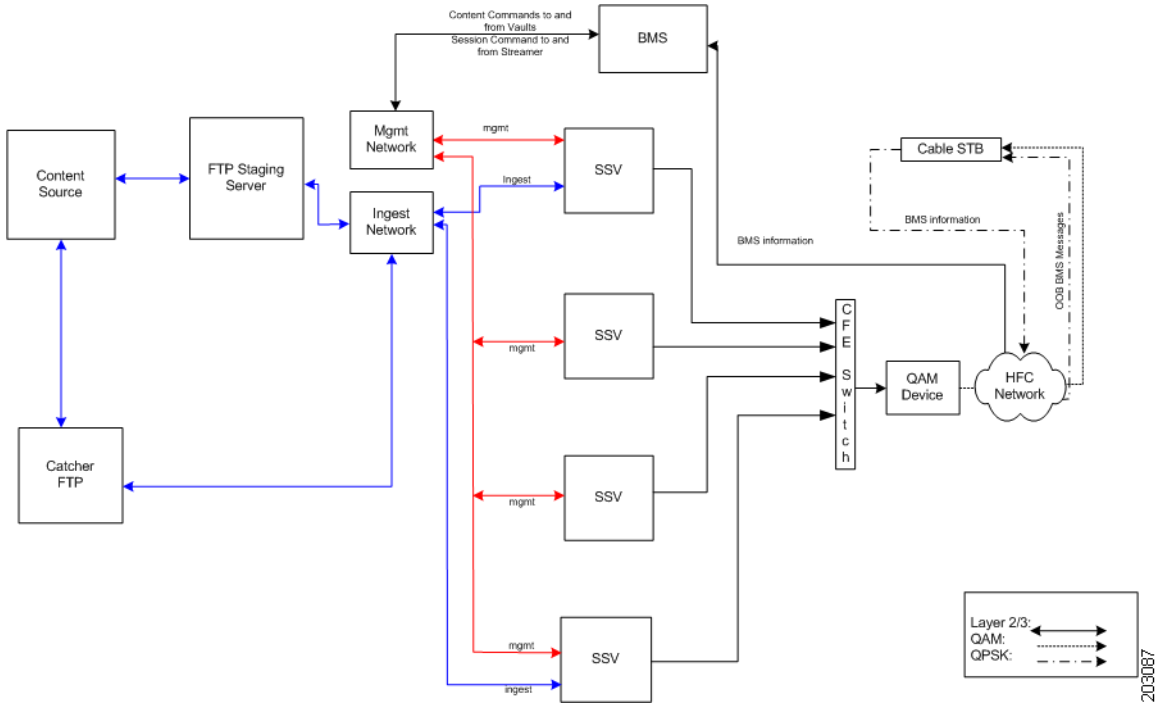

Figure 2-13 shows the different logical networks of a VVI. The ingest network receives content from the content source by way of an FTP staging server or FTP catcher where it is ingested by the Vaults. The management network consists of communication between the CDSM and BMS, as well as communication to the Vaults, Streamers, Caching Nodes, QAM devices, and STBs.

Figure 2-13 VVI Network Connections

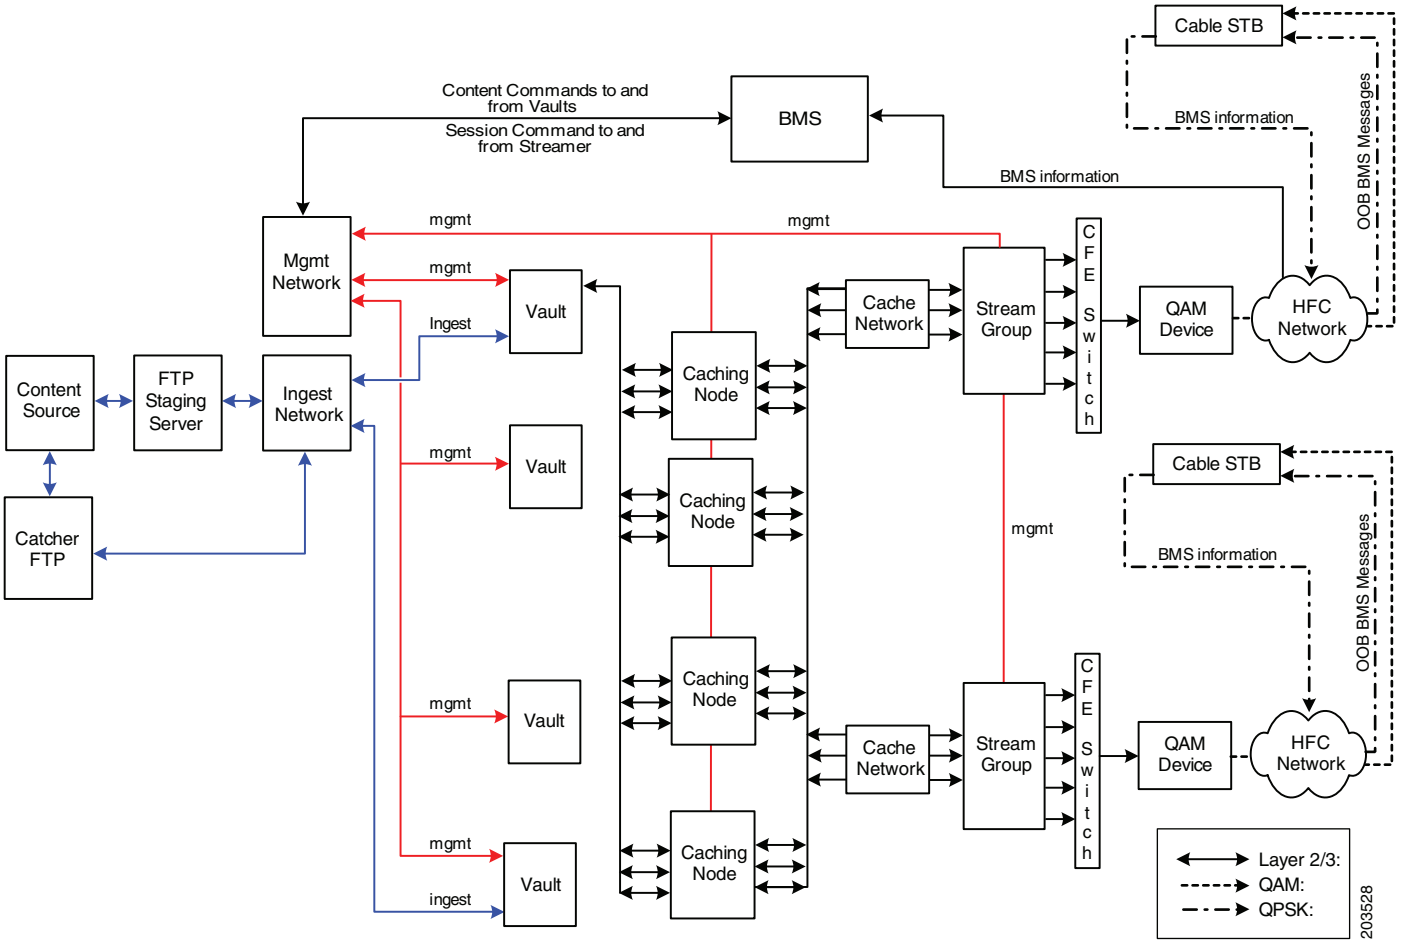

Ingest Interface

The ingest interface takes in FTP traffic from the content provider at a maximum rate of one gigabit per second. After the Vault server receives URL information about the content from the BMS by using the management interface, the ingest interface either (1) receives FTP traffic by acting as an FTP client, or (2) receives live data upon receiving a request to act as the FTP server.

When using Layer 2 packet forwarding, in order to segregate all ingest traffic through the switching fabric, we recommend the use of a port-based VLAN.

Management Interface

The management interface communicates with the network management system (NMS) by way of SNMP, the BMS by way of ISA commands and also RTSP, and with all Vaults, Caching Nodes, and Streamers in the same array. Information shared among servers in the same array includes the following:

- Host service information
- Domain Name System (DNS) service information
- QAM gateway information
- All ISA information

Management traffic is low volume; however, when using Layer 2 packet forwarding, we recommend using a port-based VLAN to ensure delivery of critical management communications.

Cache Interfaces

The CCP uses the cache interfaces on the Vaults, Caching Nodes, and Streamers to transmit the following data among servers in the same array:

- Content sent to the Streamers
- Content mirrored among the Vaults
- Messages containing information used for performance optimization exchanged among all the CDS servers

Note All Cisco CDS servers are connected through a switch fabric. Because all Vaults, Caching Nodes, and Streamers in the same array exchange heartbeat messages through the cache interfaces, it is important to ensure there is enough bandwidth among switches involved in delivering cache traffic and to support the same aggregated amount of traffic on all cache interfaces.

When using Layer 2 packet forwarding for cache traffic, we recommend the use of a port-based VLAN.

Cache/Stream Interfaces

The cache/stream interfaces on the Streamer server can be used for both cache and streaming traffic. The number of interfaces designated for each traffic type is configurable. If an interface is configured for both cache and streaming traffic, priority is given to the higher-bandwidth stream traffic provided cache traffic is able to transmit on other interfaces.

When using Layer 2 packet forwarding for cache and stream traffic, we recommend the use of a port-based VLAN.

Streaming Interface

The streaming interface delivers streaming traffic consisting of MPEG-2 transport streams to STBs by way of QAM devices.

If an interface is configured for both stream and cache traffic, and the jumbo frames feature is not enabled for stream traffic while jumbo frames is enabled for cache traffic, stream traffic uses 1500-byte packets while cache traffic uses jumbo frames.