



CHAPTER 1

Product Overview

This chapter provides a brief introduction to the Cisco TV Content Delivery System for an Interactive Services Architecture (ISA) environment. This chapter covers the following major topics:

- [Overview, page 1-1](#)
- [Content Delivery System Architecture, page 1-10](#)

Overview

The Cisco TV Content Delivery System (CDS) is a distributed network of Content Delivery Engines (CDEs) running Content Delivery Applications (CDAs) that collaborate with each other to deliver personalized entertainment and interactive media to subscribers.

The Cisco TV CDS has a variety of mechanisms to accelerate the distribution and delivery of content. The CDS interoperates with electronic program guides (EPGs), set-top boxes (STBs), and backoffice applications, offering an end-to-end solution for video delivery systems.

The Cisco TV CDS functionality can be separated into five areas:

- Ingest
- Storage
- Caching
- Streaming
- Management

Each CDE in the CDS contributes to one or more of these functions as determined by the CDAs running on it. [Table 1-1](#) describes the relationship between the CDA names and the names the TV Content Delivery System Manager (CDSM) uses.

Table 1-1 CDA Mapping to Functionality and CDSM

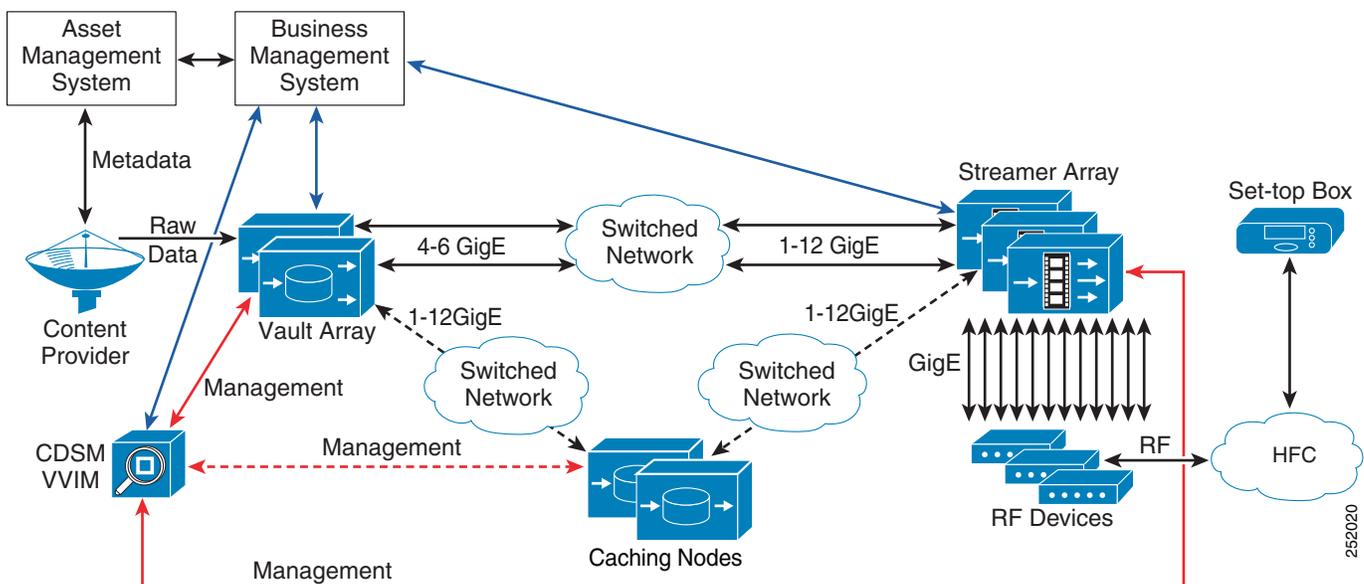
CDA Name	Functionalities	CDSM Device Name
Vault	Ingest and storage	Vault
Content Cache	Content distribution between Vaults and Streamers	Caching Node
TV Streamer	Content caching, personalization, and streaming to STBs	Streamer

Table 1-1 CDA Mapping to Functionality and CDSM (continued)

CDA Name	Functionalities	CDSM Device Name
TV MediaX Suite	Aids content ingest workflow and scheduling tasks for both asset-based and real-time content	CDSM
TV Content Delivery System Manager	Management	CDSM

Figure 1-1 illustrates how a TV CDS network can be deployed. A business management system (BMS), commonly called a backoffice, enables service providers to deploy on-demand services using video on demand (VOD) servers, networks, billing systems and other system components. The asset management system (AMS) manages the content on headend and node servers, while the BMS handles functions related to pitching and catching. Sometimes there is some overlap of functionality between the BMS and the AMS.

There are two types of systems available with the TV CDS: a CDS with an array of Vaults and Streamers, and a Virtual Video Infrastructure (VVI) with an array of Vaults, Caching Nodes, and Streamers. The CDSM manages the Vaults and Streamers in a CDS. The VVIM manages the Vaults, Caching Nodes, and Streamers in a VVI with centralized management. For more information about network design and VVI management, see the “TV CDS and VVI Topologies” section on page 2-2. Figure 1-1 shows a high-level view of both a CDS and a VVI.

Figure 1-1 High-Level System View of Content Delivery System and Virtual Video Infrastructure

The Cisco TV CDS solution has three major elements:

- One or more Vault Groups consisting of one or more Vaults. The Vaults are responsible for ingest and reliable storage of VOD content. The number of Vaults in the Vault Group, and the number of Vault Groups is driven by the amount of content that the system offers and the degree of redundancy.

- One or more Stream Groups each consisting of one or more Streamers. The Stream Group is responsible for the personalization and streaming of content in response to user requests. The number of Streamers and Stream Groups is determined by the number of streams deployed and by the topology that best suits your individual network and redundancy requirements.
- The CDSM is used to manage the Vaults, Streamers, and Caching Nodes in the same array, collects event logs, and provides reporting tools.

**Note**

In smaller systems, the Integrated Streamer-Vault (ISV) server can be used, where the Vault and Streamer functionalities exist in one ISV server.

The Cisco TV VVI solution has four major elements:

- One or more Vault Groups consisting of one or more Vaults. The Vaults are responsible for ingest and reliable storage of video on demand (VOD) content. The number of Vaults in the Vault Group, and the number of Vault Groups is driven by the amount of content that the system offers and the degree of redundancy.
- One or more Cache Groups, consisting of one or more Caching Nodes. The Caching Nodes provide more flexibility in designing a multi-tiered Virtual Video Infrastructure (VVI) by acting as a tier between the Vaults and the Streamers. The Caching Nodes facilitate content distribution and remove distribution traffic from the network backbone.
- One or more Stream Groups each consisting of one or more Streamers. The Stream Group is responsible for the personalization and streaming of content in response to user requests. The number of Streamers and Stream Groups is determined by the number of streams deployed and by the topology that best suits your individual network and redundancy requirements.
- The CDSM is used to manage the Vaults, Streamers, and Caching Nodes in the same array, collect event logs, and provide reporting tools. In a split-domain management system configuration, there is a Stream Manager that manages all the Streamers, and a Virtual Video Infrastructure Manager (VVIM) that manages all the Vaults and Caching Nodes.

TV CDS Software

The Cisco TV CDS kernel software, known as the CServer, creates a logical network that pools, load balances, and coordinates the physical resources of the CDEs, so that the whole network operates and is managed as if it is a single resource.

The CServer facilitates the rapid movement of content between Vaults and Streamers while keeping required bandwidth to a minimum. To accomplish this, the Cisco TV CDS software uses a proprietary protocol, the Cache Control Protocol (CCP), across the gigabit Ethernet networks. All content is held reliably on the Vault servers and a large amount, but not all, of the content is also contained on the Streamer servers. Cisco CCP, a multilayered caching architecture, along with associated software algorithms ensures that content segments are delivered only to the Streamers where there is demand for that content. The TV CDS software monitors the frequency of subscriber demand and places content appropriately in either the dynamic random access memory (DRAM) or disk cache of the serving Streamer.

Content is delivered across the network in response to cache-fill calls from the Streamers in an opportunistic manner, depending on the availability of bandwidth; delivery can be faster than real-time delivery where bandwidth allows. The TV CDS software ensures content on the Streamer servers is always the most popular content; that is, the content requested by the largest number of subscribers. User requests are generally served from the cache on the Streamer. Requests for content that are not already

in the local cache of the Streamer are pulled from the Vault, cached on the Streamer, and streamed to the subscriber. Wherever the content is stored relative to the point of playout, all content appears as if it is local to the Streamer and the streaming of any content is nearly instantaneous.

Caching Nodes

A Caching Node is an intermediary fill source for the Streamers. Caching Nodes are deployed in Virtual Video Infrastructures (VVI). The VVI is a deployment type of the TV CDS. In a CDS, servers cannot communicate with servers in other groups. In a VVI, servers in other groups can communicate with each other on an as needed basis. Streamers and Caching Nodes dynamically discover fill sources within other groups. Streamers send cache-fill calls to remote servers (Streamers in other Stream Groups and Caching Nodes) for content that is not found locally (DRAM, disk cache, or peer Streamers). In a VVI, the Caching Nodes can communicate with the Streamers by using CCP or HTTP. For more information on how a Caching Node interfaces with a CCP Streamer and an HTTP Streamer, see the [“Caching Node Workflow” section on page 2-10](#).

Streamer Load Balancing

To ensure that new streams are distributed to the best Streamer in the group, each Stream Group runs a load distribution protocol among its members. The best Streamer is the Streamer that has the requested content in the highest-performing cache resource (DRAM or disk) or that has the most unused capacity. In this way, new Streamers are brought into operation hitlessly—because after a new server is in service, fresh streams are automatically allocated to it. Furthermore, the cache capacity of the group is the sum of the caches of all Streamers in the group, which provides the most optimal system operation and the highest cache-hit rate.

CServer Functionality

The CServer is responsible for the following:

- Storing content
- Streaming content
- Managing bandwidth usage for ingests
- Managing bandwidth usage for streaming
- Mirroring content among Vault servers
- Making decisions on content retention on Streamer servers

Streamer Content Delivery Applications

On top of the CServer, and taking advantage of the services it offers, a variety of applications deliver individual personalized entertainment services. Cisco currently offers the following applications:

- TV Streamer delivering VOD and network personal video recorder (nPVR) services
- TV MediaX Suite for simplifying ingest and workflow scheduling tasks for asset-based and real-time content

In a full TV CDS network, the Vault, TV Streamer, and CDSM are required. The TV MediaX Suite is an optional CDA. In a smaller TV CDS network, the ISV can be used in place of the Vault and TV Streamer.

TV Streamer CDA

The TV Streamer CDA is used for VOD delivery systems. TV Streamers are responsible for personalizing content and playing that content out under subscriber control.

TV MediaX Suite CDA

The TV MediaX Suite CDA offers a set of tools that simplify content ingest workflow and scheduling tasks for both asset-based and real-time content. The TV MediaX Suite CDA consists of the following features:

- Publisher—Coordinates the ingest of pre-encrypted content.
- Scheduler—Schedules real-time content or imports the schedule from an electronic program guide (EPG).

For information on configuring TV MediaX, see the [“TV MediaX Configuration Workflow” section on page 3-11](#).

TV Playout CDA

The TV Playout feature is for ISA environments and includes Public, Education, and Government (PEG) channels and Barker Streams. PEG channels differ from traditional broadcast channels in that the service provider itself must ingest and stream the content rather than receiving and forwarding a satellite feed.

For information on configuring TV Playout, see the [“TV Playout Configuration Workflow” section on page 3-12](#).

Content Delivery

The CDS delivers real-time, time-shifted, and on-demand video content to set-top boxes, personal computers, or any other device accessible through a Service Provider network.

The Cisco VVI allows service providers to support a broad range of services. For example, with the ability to distribute content from anywhere to anywhere, operators can provide user-generated and online video just as easily as any other on-demand title. The ability to deliver content with sub-second latency also lets service providers dramatically expand the video library that can be made immediately accessible to customers, allowing them to access content that resides in a different state or country virtually instantly.

Operators can also support popular real-time and time-shifted services, such as letting viewers tuning into a program in progress and restart it from the beginning, or providing network-based personal video recorder (nPVR) functions such as the ability to pause, fast forward, and rewind live TV. The Cisco VVI's centralized storage and localized streaming architecture also distributes screen-formatting processes to the network edge.

The key content delivery capabilities include the following:

- Supports multiple content formats (high-definition and standard-definition content, multiple video codec formats, multiple media file types, and so on)
- Supports ingest and streaming of real-time video services, VOD services, and Internet video
- Supports advertising content distribution and streaming
- Supports nPVR capabilities to provide a digital video recorder (DVR)-like experience with the network

- Provides a single content delivery network for serving set-top boxes (STBs), PCs, and mobile devices
- Supports content security and encryption
- Supports narrowcast service such as VOD, time-shifted TV, and switched digital video (SDV) sharing the same infrastructure
- Supports both traditional and next-generation STBs and headends

Real-Time Splicing of MPEG-2 Transport Streams

The ISA Stream Extensions feature allows real-time splicing of MPEG-2 transport streams as identified by the Society of Telecommunications Engineers (SCTE) 35 standard. The embedded SCTE 35 cue messages contain information for digital program insertion (including advertisement insertion) in live content as well as content recorded for the purpose of enabling time-shifted on-demand services.

Pre-roll, post-roll, and mid-roll placements of digital program insertion, that is based on a playlist structure, is supported on a CDS in an ISA environment. The Vault detects the SCTE-35 cues and processes them at the time of ingest. The StreamExtChannel event channel on the CORBA NotificationService is used to send ContentSignalingEvents that contain the SCTE-35 cue information to the backoffice.



Note

The SCTE-35 cue message cannot be greater than 400 bytes.

CDSM Configuration

To configure this feature set the **TME/SCE** field to Enable for **MystroMDN**.

The **TME/SCE** field is located on different CDSM pages, depending on the type of system configured (VVI or CDS).

- In a VVI with split-domain management, on the VVIM GUI, choose **Configure > System Level > Distributed ISA Setup** or **Configure > System Level > Shared ISA Setup**
- In a VVI with split-domain management, on the Stream Manager, choose **Configure > Array Level > VHO ISA Setup**
- In a VVI with central management or a legacy CDS, choose the **Configure > Array Level > Streamer BMS**



Note

The configuration change for the ISA Stream Extensions feature requires that the ISA service is restarted on both the master Vault and the Master Streamer. To identify the master Streamer and master Vault, use the CDSM Monitor Services page to find the Streamer running the master stream service and the Vault running the Content Store master. See the [“Services Monitor” section on page 5-40](#) for more information. To restart the ISA service, choose **Maintain > Services**, select the check box for ISA, and click **Submit**.

Dynamic Modification of Playlists

The following dynamic playlists modifications can be performed on playlists that have been defined and created:

- Delete_Segment—Remove a segment from the playlist
- Replace_Segment—Replace a segment in the playlist with one or more segments

- `Splice_Segment`—Insert one or more segments at a specified NPT start value or NPT end value within an existing playlist segment
- `Add_Segment`—Add one or more segments after a segment in the playlist

**Note**

Each playlist can have up to 64 content elements.

The CDSM provides augmentations to the Stream Play History report. The Stream Play History report first displays the Session ID Summary. When a session ID is clicked, if a playlist was streamed for the session, the Session Playlist History report is displayed.

**Note**

The Trick Mode Capture feature must be enabled to access the Stream Play History reports.

Content Chunking

For DVD on Demand solutions and long recordings, Release 2.5.2 supports ingest and streaming of assets up to 120 GB in size and recordings that last longer than 12 hours. This is accomplished by dividing the asset into multiple chunks of approximately 16 GB each.

**Note**

The Content Chunking feature is disabled by default. All the CDS servers in a deployment must be upgraded before enabling this feature. To enable, the following line must be added to the setupfile of each CDS server and the server must be rebooted: **content id type 2**

Trick-Mode Restriction

Restriction of trick-mode controls (pause, rewind, fast-forward) per playlist segment is supported.

If a client issues a trick-mode command for a locked-out playlist segment or attempts to bypass a trick-mode restricted segment by jumping to the next segment, an `LSC_NOT_PERMITTED` response is sent to the set-top box. If a client has sent a fast-forward trick-mode command and a restricted segment is reached, the stream continues at normal play speed and an `LSC_DONE` response is sent to the set-top box with the NPT of the beginning locked out segment. An `LSC_NOT_PERMITTED` response is also sent to indicate that the `LSC_DONE` is due to a locked out trick-mode segment.

The CDSM GUI provides the ability to configure these settings on the MPEG Tuning page (**Configure > System Level > MPEG Tuning**).

HTTP Live Streaming

HTTP Live Streaming is fully supported; similar to live streaming over Cache Control Protocol (CCP). The enhancements to HTTP Live Streaming consist of the following:

- [Catch-Up to Live](#)
- [Play While Ingesting the Same Content](#)

Catch-Up to Live

A video player can play live content close to the live point, within 2.5 seconds of the live point, without macroblocking or leaving artifacts on the screen of the player.

If play starts at 0 or some point before the live point, then the Catch-up to Live feature allows the end-user to fast-forward to the live point and resume normal play at the live point. The play point will be within 2.5 seconds of the live point.

Play While Ingesting the Same Content

While ingesting the content, a STB can request the content play start at 0, at “play now,” or at any specific normal play time (NPT) value between 0 and the live point; and the content will begin playing at the requested point of play.

When a set-top box (STB) sends a “play now” request, meaning the STB is requesting that the play begin at the live point, the “play now” point is within 2.5 seconds of the live point.

VOD Error Repair

The VOD Error Repair feature retransmits lost packets to improve the quality of the end-user video experience. The VOD Error Repair feature uses negative acknowledgement (NACK) retransmission methods to implement retransmission-based error repair.



Note

VOD Error Repair is supported on ISA environments that use the Cisco (RTSP) setting as the LSCP Client Protocol, and RTSP environments that use the Cisco RTSP deployment type.

In addition to UDP streaming, unicast Realtime Transport Protocol (RTP) with Realtime Transport Control Protocol (RTCP) streaming, as well as Error Repair (ER) are supported.

The client dictates which streaming protocol is used by way of the RTSP SETUP message. The following streaming protocols are supported in the same system with simultaneous streams of each type:

- UDP
- RTP
- UDP with NAT traversal (Interactive Connectivity Establishment [ICE])
- RTP with NAT traversal (ICE)
- RTP with retransmission-based error repair
- RTP with NAT traversal (ICE) and retransmission-based error repair

For sessions that use UDP, aside from RTSP messages, only the media server sends packets.

For sessions that use RTP, RTCP packets may be sent from the server to the client or from the client to the server. The client must be aware of the server’s IP address and ports for receiving these packets.

For sessions that use NAT, the server sends its own IP address and ports as ICE candidates.

For sessions that do not use NAT, the transport header must include a “server ports” parameter.

For sessions that use RTP retransmission-based error repair, a client sends a second SETUP request to the CDS Control server, which requires a total of four open ports. The first SETUP message has two ports (one for RTP and one for RTCP). and the second SETUP message has two ports that carry two ICE candidates. The URLs used for the retransmission stream are appended with the “/rtx” ending.

Following is an example of the first SETUP message:

```
SETUP rtsp://192.0.2.100/movie.mpg RTSP/1.0<CRLF>
CSeq: 2 <CRLF>
Transport: RTP/AVPF/UDP; unicast; destination=54.0.1.1; client_port=8998-7123,
          MP2T/DVBC/UDP; unicast; destination=54.0.1.1; client_port=8998<CRLF>
```

```
RTSP/1.0 200 OK<CRLF>
CSeq: 2<CRLF>
Session: 12345678<CRLF>
Transport: RTP/AVPF/UDP; unicast; destination=54.0.1.1; client_port=8998-7123;
source=101.1.2.3; server_port=50236-50237<CRLF>
```

Following is an example of the second SETUP message:

```
SETUP rtsp://192.0.2.100/movie.mpg/rtx RTSP/1.0<CRLF>
Session: 12345678 <CRLF>
CSeq: 2 <CRLF>
Transport: RTP/AVPF/UDP; unicast; destination=54.0.1.1; client_port=8999-7124<CRLF>
<CRLF>
```

```
RTSP/1.0 200 OK<CRLF>
CSeq: 2<CRLF>
Session: 12345678<CRLF>
Transport: RTP/AVPF/UDP; unicast; destination=54.0.1.1; client_port=8999-7124;
source=101.1.2.3; server_port=50238-50239<CRLF>
<CRLF>
```


Note

Retransmission-based Error Repair is only available with RTP streaming.

Background

RTP packets include sequence numbers that are used to detect missing packets and reorder out-of-order packets. RTCP is the control protocol for RTP and is used to send receiver reports from the client to the server that include monitoring information, to send sender reports from the server to the client, and to request retransmission, which is the RTCP NACK packet that includes the RTP sequence number.

The Streamer receives the retransmission RTCP NACK request. Each NACK request identifies one or more missing RTP packets. The Streamer keeps a small buffer of recently transmitted packets and the missing packets are retransmitted based on how many packets the buffer maintains.

Error Repair Client on STB

VOD Error Repair feature requires that the STB have the Cisco Visual Quality Experience Client (VQE-C) software running on it. The VQE-C is the error-repair client software, which has the following capabilities:

- Receives RTP video packets
- Detects missing packets
- Requests retransmission of missing packets
- Merges retransmitted packets with original stream
- Collects statistics and counters for monitoring

The VQE-C is a software development kit (SDK) that is available for download through the open-source program. Additionally, the STB must comply with the Cisco RTSP syntax for VOD Error Repair.

Monitoring

The play management application (PMA) log file, vqe.log, is located in the /arroyo/log directory. To check for PMA errors, enable the PMA debug flag for the vqe_cp facility on the Logging page in the CDSM.

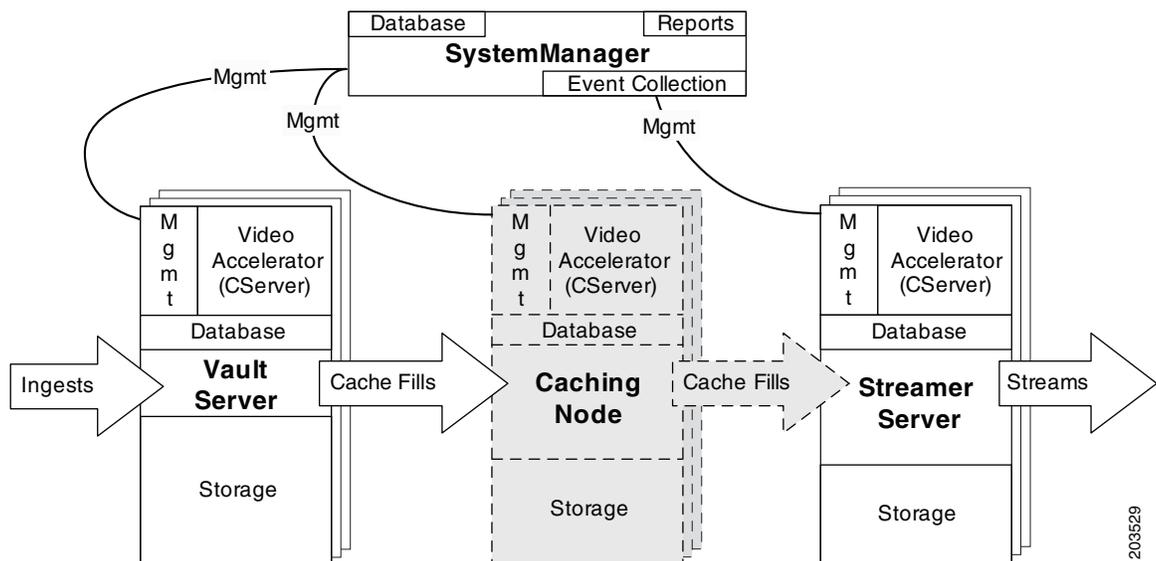
AMT

Application Monitoring Tool (AMT) runs a web application on each Streamer and provides several troubleshooting tools. For more information, see [Appendix E, “Using the TV CDS Streamer Application Monitoring Tool.”](#)

Content Delivery System Architecture

Vaults and Streamers have different but important functions that are required for the TV CDS software to run efficiently. The Integrated Streamer-Vault (ISV) server combines the functionality of both the Vault and Streamer for smaller networks. The Content Delivery System Manager provides a browser-based user interface for configuration, monitoring, maintenance, and reports of the TV Content Delivery System solution. In a VVI, the Caching Nodes provide a pure caching layer for a multi-tiered VVI. [Figure 1-2](#) shows the different elements of the TV Content Delivery System and the TV Virtual Video Infrastructure with the addition of the Caching Nodes.

Figure 1-2 High-Level View of the Content Delivery System and Virtual Video Infrastructure



[Table 1-2](#) describes the system elements shown in [Figure 1-2](#).

Table 1-2 High-Level Description of the TV CDS and TV VVI

Content Delivery System Element	Description
CServer	The CServer is the kernel software that handles bandwidth management, storage decisions, Real Time Streaming Protocol (RTSP) and Lightweight Stream Control Protocol (LSCP) and stream processing on the TV Content Delivery System.
Database	The database stores information about the system, including current states of all ingests and streams, configuration settings, and system statistics. Some database elements are global among all servers and some are local. For example, statistics are stored on the local server and the Content Delivery System Manager only. States about stream objects are replicated on all Streamer servers. The Content Delivery System Manager stores a superset of all database elements.
Management	There are two types of management: <ul style="list-style-type: none"> • Content Delivery System Manager—Browser-based user interface • SNMP agent—Network Management System (NMS) interface
Storage	There are four levels of storage (or cache): <ul style="list-style-type: none"> • All content is stored on the Vault server, as well as mirrored to other Vaults. • Requested content is stored on the Caching Nodes. • Recently requested content, or popular content is stored on the hard drive of the Streamer. • Currently requested content, or popular content, is stored in the random access memory (RAM) of the Streamer.
Event Collection	The Content Delivery System Manager collects logged events for reporting purposes as well as for third-party applications
Reports	The Content Delivery System Manager provides a reporting tool to aid performance trending and analysis of streams, popular content, bandwidth usage, and more.

Vault

The Vault ingests content delivered over a standard interface (for example, using FTP to receive content from a catcher), performs whatever processing is required (for example, generating trick-play files), and stores the processed content reliably on disk. A Vault Group consists of a scalable number of Vaults that divide the responsibility for ingest and storage among the members of the group. Vault servers can be colocated or distributed to multiple locations across an IP or Ethernet network. Each Vault can simultaneously ingest up to 160 channels of MPEG-2 transport stream (TS) content and store up to 6000 hours of MPEG-2 TS standard definition content with two mirrored copies of the content and one or two trick files.

Streamer

A Streamer server receives content from the Vault and delivers that content to subscribers. Streamers can be of different capacity, depending on the needs of the network, and have different applications, depending on the type of content being delivered. Currently, the highest-capacity Streamer can simultaneously stream approximately 2500 streams of MPEG-2 TS standard definition VOD. Streamers can be colocated with Vaults or distributed to remote locations. The Stream Group is responsible for the personalization and streaming of content in response to user requests.

Caching Node

The Caching Node provides a 10-Gbps throughput to facilitate the distribution of content from the Vaults to the Streamers. The Caching Nodes allow for the ability to create a tier-based hierarchy in the CDS. Caching Nodes are deployed in VVIs. Vaults can be strategically located for storing content on a national network, while the Streamers are located in a regional network. The Caching Node can be colocated with the Vaults or distributed closer to regional locations across an IP or Ethernet network. A Cache Group consists of several Caching Nodes that divide the responsibility for distribution among the members of the group.

The Caching Nodes use CCP to communicate with the Vaults and Streamers. Alternatively, the Caching Nodes can use HTTP instead of CCP to communicate with Streamers.

Integrated Streamer-Vault

The Integrated Streamer-Vault (ISV) server offers the functionality of both a Vault and Streamer in one server.

The ISV server ingests content delivered over a standard interface, performs whatever processing is required, and stores the processed content reliably on disk. An ISV array consists of a scalable number of ISV servers that divide the responsibility for ingest, storage, and streaming among the members of the array.

Content Delivery System Manager and Virtual Video Infrastructure Manager

The Content Delivery System Manager (CDSM) and the Virtual Video Infrastructure Manager (VVIM) are each a browser-based user interface accessible by a web browser program and designed to manage a TV CDS or a TV VVI network.

The CDSM provides centralized management functions for the TV CDS, including configuration, monitoring, troubleshooting, reporting, and maintenance.

The VVIM provides centralized management function for the TV VVI, including configuration, monitoring, troubleshooting, reporting, and maintenance. The VVIM in a centralized domain management configuration manages the Vaults, Caching Nodes, and Streamers, The VVIM in a split-domain management configuration manages the Vaults and Caching Nodes, while the Streamers are managed by the Stream Manager. For more information about split-domain management, see the [“TV VVI Management” section on page 2-6](#).

In both the CDS and VVI, all Vaults and Streamers are identified by an array ID, a group ID, and a server ID. 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 or Stream Group), and the server ID is a unique number that identifies the server. [Table 1-3](#) lists the CDSM GUI ID names and maps them to the CServer names in the setupfile and .arroyorc files.

Table 1-3 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

In a VVI with CCP Streamers, similar to a CDS, all Vaults, Streamers, and Caching Nodes are identified by an array ID, a group ID, and a server ID. The group ID and server ID in a VVI with CCP Streamers must be unique among other groups and servers in the same system.

In a VVI with HTTP Streamers, the Vaults, Streamers, and Caching Nodes still use an array ID, a group ID, and a server ID for identification, but there is additional functionality that allows the Vaults and Caching Nodes to communicate using CCP, while the Caching Nodes communicate with the Streamers using HTTP. It is not required that the group ID and server ID be unique, but it is recommended.

The CDSM and VVIM (as well as the Stream Manager) have three configuration and monitoring levels: system, array, and server. System-wide configuration affects all servers managed by that manager. The array-level configuration affects all the servers of the specified array or group, and the server-level configuration applies changes to a specific server.

The CDSM and VVIM offer a drill-down approach to show the status of any stream or ingest point, or the physical status of any piece of hardware.

The CDSM reporting helps operators manage all aspects of the TV CDS. Information on stream traffic, content statistics, and server data are gathered from all servers in the network and correlated automatically, showing at a glance the status of the network and reporting on statistics such as content popularity, stream usage, and bandwidth usage for each service group.

The VVIM monitoring and reporting helps operators manage all aspects of the TV VVI in either a centralized management capacity or a split-domain management capacity. In a split-domain capacity, the VVIM monitors the ingests and the Stream Manager monitors the streams of the Streamers in its domain.

[Figure 1-3](#) shows the system monitoring page of the CDSM.

Figure 1-3 Content Delivery System Manager User Interface

The screenshot shows the Cisco Content Delivery System Manager interface. The top navigation bar includes 'Home', 'Help', and 'Logout'. Below this is a menu with 'Configure', 'Monitor', 'Report', and 'Maintain'. Under 'Monitor', there are sub-menus for 'System Level', 'Array Level', and 'Server Level'. The main content area is titled 'System Health MONITOR' and contains a table of server health data. The table has columns for IP address, overall health, Network, disk, Services, and health. All servers shown have green status indicators.

IP Address	overall	Network	disk	Services	health
10.22.216.119	Green	Green	Green	Green	Green
10.22.216.123	Green	Green	Green	Green	Green
10.22.216.148	Green	Green	Green	Green	Green

CDSM Time 14:19:24

Resiliency and Redundancy

The TV Content Delivery System is designed to have no single point of failure. The TV Content Delivery System incorporates redundancy at several levels within the architecture. These levels of redundancy eliminate any customer impact from potential failures of Vault disks, Vault servers, Streamer disks, Streamer servers, ISV servers, Ethernet connections, processors, and power supplies.

Each server constantly monitors the state of its peers. The TV CDS unique resource pooling and auto-failover techniques allow all servers in the network to actively contribute to satisfying storage and streaming demand at all times. If a server fails, the load is instantaneously redistributed among the surviving servers, ensuring continuity of service.

Vault Disk Redundancy

The Vault server protects content through full 1:N redundancy. If a disk fails, the data is available from a redundant server, spreading the load and optimizing the bandwidth. Additionally, the regeneration of the redundant content utilizes the bandwidth of the whole Vault array rather than just the disk bandwidth available inside a particular server, significantly reducing the rebuild time. The need to replace the failed drive is not time critical in the least, making quarterly replacement of any failed Vault drives feasible.

Mirroring

The primary method to protect the content against loss because of hardware failure is mirroring. Content is stored on a Vault and, based on the policy, it is mirrored to other locations in the Vault array. The number of mirrored copies is configurable. There are three types of mirroring:

- Local mirroring
- Mirroring within an array
- Array mirroring (from Vault Group to Vault Group)

Local Mirroring

Local mirroring defines the number of copies of each content object to maintain on the unique drives of a single Vault. Local mirroring allows resiliency for a small installation (for example one Vault). Local mirroring guards against a single drive failure, but does not protect against service interruption or potential data loss in the event of a complete server failure.

Local mirroring is not configured by default, and is generally only used when there is a single Vault in a system. Local mirroring is configured in the **Configure > Server Level > Server Setup** page with the **Vault Local Copies** field, which corresponds to the tunable “vault local copy count” in CServer. Up to four local copies are supported.

Mirroring within an Array

Mirroring within an array defines the number of copies of each content object in an array to maintain across the Vaults within that array or site. Mirroring within an array guards against a single drive failure or the failure of an entire server. The number of copies to maintain within that array is configurable in the **Configure > Server Level > Server Setup** page with the **Vault Mirror Copies** field, which corresponds to the tunable “vault mirror copies” in CServer. Up to 10 copies within an array are supported.

Array Mirroring

Array Mirroring (from Vault Group to Vault Group) specifies that each content object on all of the Vaults in one group has at least one copy on a Vault in the mirrored Vault Group. Array Mirroring is only responsible for ensuring that a single copy of each content exists in the mirrored Vault Group. If more than one copy of each content object is required within an array, Mirroring within an Array (not Array Mirroring) is responsible for this task. Array Mirroring is configured in the **Configure > Array Level > Vault Redundancy Map** page, which corresponds to the tunables “allow vault array mirroring” and “vault array mirror” in CServer. Each Vault Group can have up to 3 mirrored Vault Groups configured.



Note

Array Mirroring is part of the Vault Groups feature and is only available if Vault Groups is enabled on the CDSM Setup page. For more information, see the [“Vault Groups” section on page F-5](#).

Vault Server Resiliency

The Cisco TV CDS can handle the loss of an entire Vault server without impacting the subscriber. The communication with the backoffice suite is performed by a Vault server that is designated as the Vault master. If the Vault master fails, one of the remaining slave Vault servers in the Vault array transparently takes over as the master. The remaining Vaults detect the loss of a Vault server, run a check of all stored content, and regenerate redundant content that was affected by the lost Vault server. This regeneration runs in the background, utilizing spare system bandwidth that is not consumed by subscriber load, resulting in the shortest possible regeneration window possible without compromising performance to the subscriber.

Vault Master

The Vault master, designated by a virtual IP address on its management interface, is used as the representative of the Vault array to the backoffice and handles the ingest of new content.

Vault Group Redundancy

In addition to the Vault server redundancy, the Cisco TV CDS offers redundancy for Vault Groups. When the CDS is configured with Vault Group redundancy and at least two Vault Groups are configured, the system handles the loss of an entire Vault Group without impacting the subscriber experience. Content is mirrored among as many as four Vault Groups (one Vault Group ingests the content and up to three Vault Groups mirror the content), which may be in different geographic regions. If the primary Vault Group becomes unavailable, because of network, power, or other catastrophic problems, any Streamer or Caching Node that was requesting content from that Vault Group would fail over to the other Vault Group until the primary Vault Group came back online and could again respond to cache-fill requests for content.

With Vault redundancy, at least one copy of each content within a group is mirrored to a configured peer group. Vault Group mirroring runs as a low-priority process, so as not to impact the performance of the guaranteed streaming delivery.

**Note**

The maximum number of Vault Groups is 20.

Streamer Disk Redundancy

The disks in the Streamer are not used for full content storage as in most VOD implementations. Rather, the Streamer disks are part of the TV CDS multilevel caching architecture. If a disk is lost on a Streamer, the only impact is a marginal loss of caching capability for the system. Any content that was cached on that Streamer disk is retrieved again from the Vault. The RAM on the Streamer has enough content cached for streaming to the subscriber, so that this refetch of content from the Vault occurs without impacting the subscribers. For example, for a Streamer array of five Streamers with sixteen hard drives each, a lost drive only reduces the total caching capability by less than 1.25 percent. The need to replace the failed drive is not time critical in the least, making quarterly replacement of any failed Streamer drives feasible.

Streamer Server Resiliency

The Cisco TV CDS architecture allows for failed Streamer servers as well. If any Streamer server fails, the communication to the backoffice is transparently handed off to another Streamer. With the TV CDS software, if a Streamer server fails the other Streamers recognize that failure and continue streaming to that subscriber.

Caching Node Disk Redundancy

The disks in the Caching Node are not used for full content storage like most VOD implementations. Rather, the Caching Node disks are part of the TV CDS multilevel caching architecture. If a disk is lost on a Caching Node, the only impact is a marginal loss of caching capability for the system. Any content that was cached on that Caching Node disk is retrieved again from the Vault.

Caching Node Resiliency

The Cisco TV CDS architecture allows for failed Caching Nodes as well. If a Caching Node fails, any cache-fill transmissions that were in process at the time of the failure are re-requested by the Streamer, and any new requests are responded to by the remaining Cache Nodes in the Cache Group.

CDSM Redundancy

The Cisco TV CDS offers 1+1 redundancy for CDSMs. The primary CDSM, designated by a virtual IP address on the management interface, is used as the representative of the CDSMs to the web browser and northbound integrations, such as HTML API calls and SNMP calls.

All CDS servers keep track of a controller IP address in the `.arroyorc` file. With CDSM redundancy, both management IP addresses are specified in the `.arroyorc` file on each CDS server, except the CDSM, which only has the other CDSM IP address.

The `statsd` process is configured with a virtual IP address that can move from one CDSM to the other. If the primary CDSM becomes unavailable, because of network, power, or other catastrophic problems, the secondary CDSM takes over the virtual IP address and the administrator can connect to the secondary CDSM within 15 seconds.

Login information is not shared between CDSMs. If the administrator is logged in and a failover occurs, the administrator has to log in again to the other CDSM.

The CDS servers (Vault, Caching Node, Streamer, and ISV) participate in replication with both the primary and secondary CDSM in the same manner as occurred without redundancy, including synchronization of tables. However, the CDS servers can only retain up to one hour of reporting data, so if a CDSM is down for over an hour, when the CDSM recovers, it only is able to get the last hour of reporting data from each CDS server, which means the reporting data is not synchronized between the primary and secondary CDSMs. Reporting data is archived in comma-separated value (CSV) files every 24 hours and these CSV files are deleted when they are older than 30 days.

Ethernet Link Resiliency

All Ethernet links used within the Cisco TV CDS architecture incorporate link failure detection with automatic failover. This includes the interconnections between the Vault array and the Streamer array for cache-fill, and the Ethernet links that carry the subscriber streams to the transport networks.

Scalability

The Cisco TV CDS has separated streaming and storage, which enables a cable operator to add storage without affecting streaming counts, to add streaming without affecting storage, and in VVIs, to add distribution nodes without directly affecting storage or streaming. This flexibility allows cable operators to grow according to the needs of customers and to scale the system on an as-needed basis. For example, if more storage is required, the cable operator adds a Vault server without taking the system offline, and in Layer 2 networks the new device is automatically discovered within the architecture and the new resources are automatically utilized by the system. If additional streaming is required, the content provider either purchases more streaming licenses within the current servers, or a Streamer server is added to the system without the need to take the system offline.

