



CHAPTER 1

Product Overview

This chapter provides a brief introduction to the Cisco TV Content Delivery System for a Real Time Streaming Protocol (RTSP) environment. This chapter covers the following major topics:

- [Overview, page 1-1](#)
- [Content Delivery System Architecture, page 1-5](#)

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

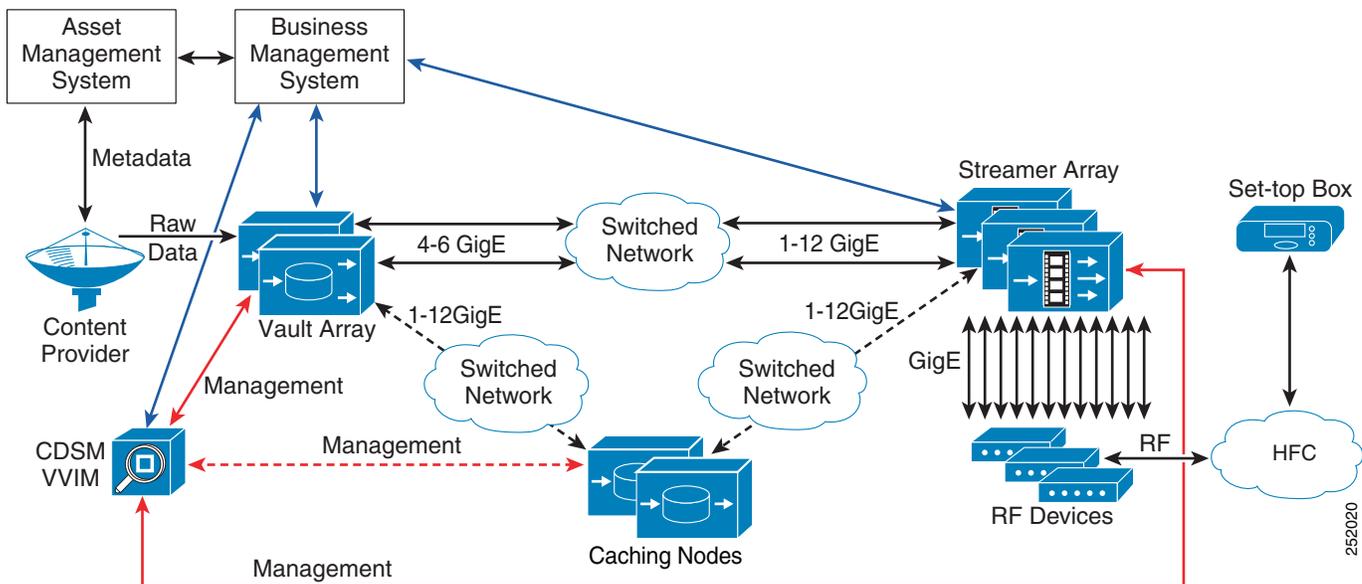
CDA Name	Functionalities	CDSM Device Name
Vault	Ingest and storage	Vault
Content Cache	Content distribution between Vaults and Streamers	Caching Node
TV Streamer	Content caching, personalization, and streaming to STBs	Streamer

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 Topologies 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:

- A Vault array consisting of one or more Vault servers. The Vault array is responsible for ingest and reliable storage of video on demand (VOD) content. The number of Vault servers in the Vault array is driven by the amount of content that the system offers and the degree of redundancy.

- One or more Streamer arrays each consisting of one or more Streamer servers. The Streamer array is responsible for the personalization and streaming of content in response to user requests. The number of Streamer servers and Streamer arrays is determined by the number of streams deployed and by the topology that best suits your individual network and redundancy requirements.
- A CDSM server. The Content Delivery System Manager is used to manage the Vault and Streamer servers, collect event logs, and provide 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 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 on 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 that 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 on 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 2-9](#).

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

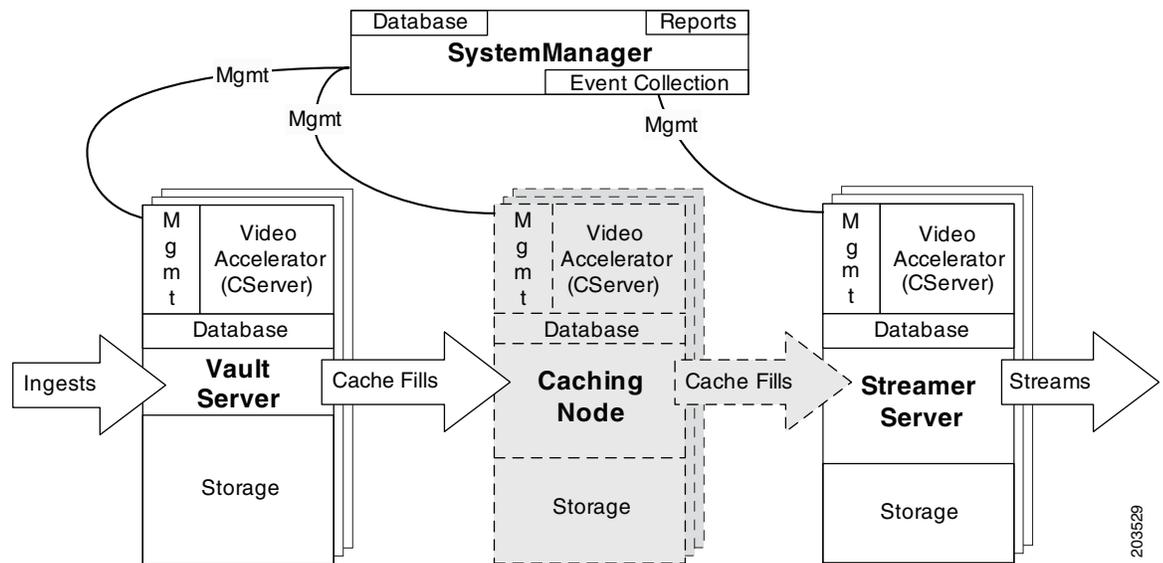
The TV MediaX Suite CDA offers a set of tools that simplify content ingest workflow and scheduling tasks for both asset-based and real-time content. The TV MediaX Suite CDA consists of the following features:

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

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 on the Streamer. • Currently requested content, or popular content, is stored in the random access memory (RAM) on 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 collocated 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 collocated 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 collocated 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. 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 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 displays the Cisco Content Delivery System Manager (CDSM) user interface. The top navigation bar includes 'Home', 'Help', and 'Logout'. Below this is a menu with 'Configure', 'Monitor', 'Report', and 'Maintain'. The 'Monitor' menu is expanded to show 'System Level', 'Array Level', and 'Server Level'. The main content area is titled 'System Health MONITOR' and contains a table of server health data. A sidebar on the left lists various system and content objects.

System Health MONITOR						
Below is a breakdown of all servers that are part of the Cisco CDSM. If one of the monitored areas is yellow or red you may click that item to view the specific reason for the warning.						
	overall	Network	disk	Services	health	
10.22.216.119	■					
10.22.216.123	■					
10.22.216.148	■					

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 two types of mirroring:

- Remote mirroring
- Local mirroring

When remote mirroring is used, copies of the content are mirrored to drives on other Vaults, based on the number of Vault mirror copies configured.

When local mirroring is used, copies of the content are mirrored across all the available drives on the same Vault, so that the content can be recovered from another drive if one of the drives fails.

Local mirroring is not turned on by default, and is generally only used when there is a single Vault in a system.

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 1:1 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. Content is mirrored between two Vault Groups, 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 the 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 for this release is two.

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 is only 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.