



CHAPTER 1

Introduction to Cisco CDA Visual Quality Experience Application

This chapter provides information on Cisco CDA Visual Quality Experience (VQE) Application and contains the following major topics:

- [VQE Overview, page 1-1](#)
- [VQE Benefits, page 1-8](#)
- [VQE-S, page 1-9](#)
- [VQE-C, page 1-11](#)
- [VQE Deployment Options and Requirements, page 1-14](#)
- [VCPT and Channel Information, page 1-15](#)
- [VCDS, page 1-17](#)
- [VQE-S AMT, page 1-21](#)
- [VQE-S RTCP Exporter for Video-Quality Monitoring, page 1-23](#)
- [VCDS AMT, page 1-24](#)
- [Content Delivery Engine 110, page 1-26](#)

VQE Overview

This VQE overview has the following sections:

- [Introduction to VQE, page 1-2](#)
- [VQE Major Software Components, page 1-4](#)
- [VQE Hybrid Error Repair, page 1-4](#)
- [VQE RCC, page 1-5](#)
- [RTP and RTCP, page 1-7](#)
- [VQE Web Browser-based Tools, page 1-7](#)
- [Lookaside Mode and the Cisco CDE110, page 1-8](#)

Introduction to VQE

Cisco CDA Visual Quality Experience (VQE) Application offers service providers a set of technologies and products associated with the delivery of Internet Protocol television (IPTV) video services. VQE is designed to improve the quality of IPTV services and viewing experiences of the subscriber. VQE is part of a Cisco end-to-end solution that builds video awareness into the network infrastructure. For Cisco VQE Release 3.5, VQE technology is intended for wireline operators who offer managed broadcast (multicast) IPTV services using xDSL.

IPTV subscribers expect high video quality. Because many subscribers are migrating from existing analog or digital cable services, their quality expectation has already been set. To attract subscribers, IPTV providers must meet or exceed the video experience of existing services. VQE technology and products provide that capability.

Video is less tolerant of network factors such as jitter, delay, and especially packet loss because a single IP packet carries up to seven MPEG transport frames. Therefore, IP networks require additional functionality to deliver the video quality expected by subscribers. The accepted industry benchmark for quality is to deliver a maximum of one video *artifact* or perceived distortion during the viewing of a one-hour movie. This level of quality translates into a network-layer requirement of a loss of fewer than $7.8E-7$ video packets. Most xDSL networks are not optimized to deliver such low levels of packet loss.

The second issue that affects IPTV video quality is channel change time (CCT). Consumers have become used to the current CCT—which is less than 2 seconds—offered by digital cable and digital satellite services. Several factors can contribute to longer IPTV channel change times, including Internet Group Management Protocol (IGMP) delays, MPEG decoding delays, and I-frame acquisition time. Non-optimized IPTV channel change times can take several seconds. Channel change times as long as 5 seconds have been observed.

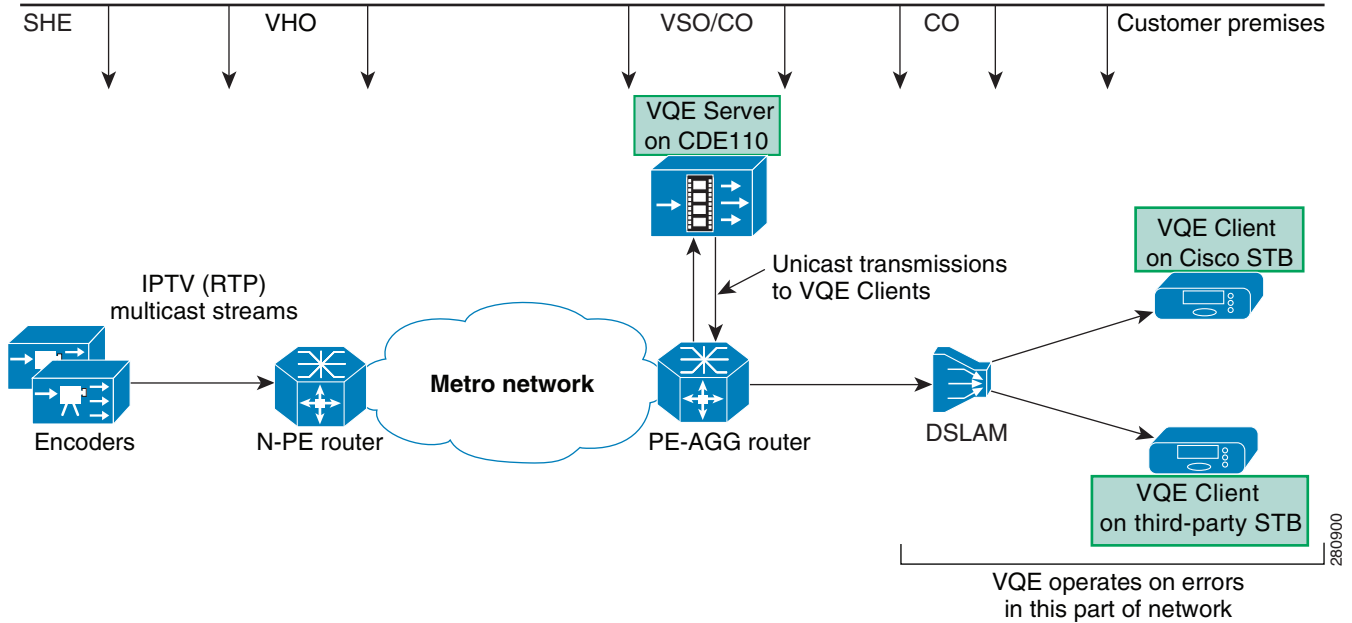
VQE addresses the issue of video quality from both a network infrastructure and a video technology perspective. VQE provides the linkage to optimize video delivery over next-generation carrier networks. Based on industry standards, including Real-Time Transport Protocol (RTP) and RTP Control Protocol (RTCP), Cisco VQE provides these mechanisms to help in delivering entertainment-grade services to subscribers:

- Unicast Retransmission—Optimized, selective retransmission of dropped IPTV packets caused by noisy DSL lines or errors in the home network caused by poor quality wiring. The set-top box (STB) receiver, the VQE Client (VQE-C), sends non-acknowledgement (NACK) packets to the VQE-S to request retransmission of the lost packets.
- Forward Error Correction (FEC)—Extra information is sent along with the video data at the application layer. The additional information is used by the VQE-C on the STB to detect and correct lost packets.
- Rapid Channel Change (RCC)—When the subscriber requests a channel change, the VQE-C on the STB sends the VQE-S a request for the IPTV packets of the new channel. The VQE-S sends the VQE-C an optimized unicast burst of IPTV packets and other channel information for the new channel from the cached video data of the VQE-S. This greatly reduces the time needed to display the new channel.
- IPTV Packet Loss Monitoring—Facilities such as VQE-S RTCP Exporter help operators measure, baseline, and pinpoint problem areas of the video infrastructure, including transmission lines and home networks.

The two error-repair options, Unicast Retransmission and Forward Error Correction (FEC), can be used separately or together for Hybrid Error Repair (see [“VQE Hybrid Error Repair” section on page 4](#)). Each repair mechanism has its own advantages and limitations. Whether Unicast Retransmission or FEC or both are used, the subscriber does not detect the error repair and no video artifact results.

Figure 1-1 shows the location of the major network and VQE components in the service-provider network and in the customer premises equipment (CPE) of the subscriber.


Figure 1-1 VQE Major Components



If FEC is used for error repair, an SMPTE 2022-compatible FEC stream generator, such as the Cisco Digital Content Manager (DCM), is required. The FEC stream generator, shown in Figure 1-2, is often located in the headend and sends application-layer FEC packets. RTP encapsulation is a prerequisite for FEC.

The FEC stream generator, which is connected to the real-time encoders, subscribes to video channels from the headend and encodes FEC for each channel. It is responsible for originating FEC packets associated with individual channels and with primary multicast streams. For each FEC-enabled channel, a primary media stream and two encoded FEC streams are sent over multicast addresses (see Figure 1-2).

Figure 1-2 FEC Stream Generator



VQE Major Software Components

The two major VQE software components that implement Unicast Retransmission, FEC, RCC, and IPTV Packet Loss Monitoring are:

- VQE-C—Software embedded in the CPE of the subscriber—typically a STB. The VQE-C provides the CPE interface to the VQE Server (VQE-S) to support Unicast Retransmission, RCC, and IPTV Packet Loss Monitoring statistics. The VQE-C receives the primary media data packets and, if FEC is enabled, one or two streams with FEC packets. When the VQE-C software detects packet loss on a channel that is configured for FEC and Unicast Retransmission, the following occurs:
 - If there are packet losses in the primary media stream, the VQE-C first tries to repair the lost packets using the FEC streams.
 - If some packet losses cannot be corrected by FEC, the VQE-C requests a Unicast Retransmission of the missing packets from the VQE-S.
- VQE-S—Software that runs on a Linux-based Cisco Content Delivery Engine 110 (CDE110) appliance located in the intelligent edge of the service provider’s network. For Unicast Retransmission and RCC, the VQE-S caches primary video packets from an encoder or other headend device.
 - For Unicast Retransmission, working with the VQE-C, the VQE-S monitors the subscriber’s reception of video packets and uses its cached video data to service Unicast Retransmission requests from the VQE-C on the STB.
 - For RCC, when the subscriber requests a channel change, the VQE-C on the STB sends a request for a new channel to the VQE-S. To service the RCC request, the VQE-S sends the VQE-C a unicast burst of video packets from its cached video data for the channel and also sends some MPEG priming information to facilitate immediate decoding.


With both Unicast Retransmission and FEC, the missing packets are resequenced by the STB without interruption.

The VQE-C is available with certain Cisco STBs running Scientific-Atlanta IPTV Layer (SAIL) 1.x and 2.x. Please contact your Cisco sales representative for further information. The VQE-C can also be integrated into a STB from third-party vendors. The VQE-C code and Software Development Kit (SDK) is available to third-party vendors through an open-source program.

VQE Hybrid Error Repair

VQE Hybrid Error Repair occurs when Unicast Retransmission and FEC are used together. The devices used for Hybrid Error Repair are shown in [Figure 1-3](#), which for simplicity omits some of the network elements. Hybrid Error Repair allows the service provider to customize VQE error repair to match the error characteristics of a given access network.

Figure 1-3 Hybrid Error Repair



In an IPTV system, video data is very sensitive to packet losses because of the interdependence of the encoded data. One packet loss could cause quality degradation of several successive frames in a group of pictures (GOP). Selective retransmission and FEC are two methods to protect channels from packet losses.

With selective retransmission (Unicast Retransmission), RTCP NACK compound packets are sent to the VQE-S to request retransmission of the lost packets whenever the receiver (VQE-C) detects dropped packets. Selective retransmission is very efficient in terms of bandwidth utilization in a unicast scenario, but may flood the network with control packets when errors are highly correlated in nature. Starting with Cisco VQE Release 3.5.5, it is possible to configure the VQE-C on a per-channel basis to send RTCP NACK compound packets without Receive Reports (RRs) to reduce the bandwidth consumed by RTCP NACK compound packets on the access link.

The application-layer FEC method of error repair is usually used for controlling errors in a one-way communication system. The FEC stream generator (for example, the Cisco DCM) sends FEC information along with the primary media stream. The VQE-C uses the FEC data to detect and correct the lost packets. No feedback to the VQE-S is needed. FEC is optimized for non-bursty, correlated errors. Unicast Retransmission is better for bursty, uncorrelated errors.

In addition to providing a flexible solution for error repair, VQE Hybrid Error Repair is also able to correct more lost-packet errors than when FEC is used alone. If the VQE-C tries but is unable to retrieve lost packets using FEC, it can request that the VQE-S selectively retransmit the dropped packets.

The VQE Hybrid Error Repair solution provides the flexibility to customize an error-repair scheme that best suits the network error characteristics and available access-link bandwidth.

VQE RCC

Channel change time is defined as the time from when a subscriber initiates a channel change to the time the video for the new channel is displayed. The RCC functionality built into the VQE-S and the VQE-C reduces channel change time from several seconds to approximately one second.

There are several factors that contribute to channel change time:

- Multicast latency for *leaving* the old channel (IGMP leave latency).
- Time it takes for the STB to receive the information it needs to begin demultiplexing, decoding, decrypting, and displaying the video stream. This time includes how long it takes to receive the following:

- Program Association Tables (PAT), Program Map Tables (PMT), and Entitlement Control Messages (ECM) if decryption is required.
- Program Clock Reference (PCR) and sequence header information (for example, frame rate).
- Random access point (such as I-frame) acquisition delay.
- Network buffer delays, including delays caused by error-mitigation techniques (Unicast Retransmission or FEC).
- MPEG decoder buffer delay.
- Multicast latency for *joining* the new channel (IGMP join latency).


Not all of the above factors need to be addressed by RCC functionality. For example, many Digital Subscriber Line Access Multiplexers (DSLAMs) can provide fast leave for IGMP, which addresses the multicast leave latency factor.

The goal of RCC is to reduce or eliminate the main sources of channel change delay. With RCC, the resulting channel change delay should be similar to or better than the delay observed in a typical digital broadcast.

For RCC, the VQE-S caches multicast IPTV packets corresponding to each channel (see [Figure 1-4](#)). Each cache holds a few GOPs of video data as well as PAT, PMT, ECM, PCR, and sequence information.

When the subscriber requests a channel change, the VQE-C on the STB requests the IPTV packets for the new channel from its target the VQE-S, using a specific RTCP message. After it has sent the RTCP message, the VQE Client issues an IGMP join for the new channel at an optimum point in time.

Figure 1-4 VQE Server Caches Multicast IPTV Packets for RCC



Upon receipt of the request, the VQE-S locates the appropriate channel cache, identifies the location of the IPTV packet carrying the beginning of a recent I-frame, and sends a short unicast burst of packets starting with the I-frame to the requesting VQE Client.

Before sending the unicast burst, the VQE-S originates an RTCP message to the requesting VQE Client *priming* it with video parameters such as PAT, PMT, and ECM.

Because the incoming burst from the VQE-S contains an I-frame, the STB decoder can immediately start processing the MPEG information. This greatly reduces the time a subscriber waits before the image is rendered on the TV screen. Other VQE optimizations for RCC, such as Fast Decoder Buffer Fill, shorten channel change time by reducing MPEG decoder buffer delay.

After a short period of time, multicast packets for the new channel start to arrive at the STB. The VQE-C monitors RTP sequence numbers from both unicast and multicast streams. It is likely that the VQE-C sees a few packets with duplicate RTP sequence numbers before the unicast stream ends. During this period, the VQE-C only forwards one copy of the RTP packet to the MPEG demultiplexing stage.

The VQE-C is responsible for managing the seamless transition between unicast and multicast IPTV packets. The RCC unicast burst continues for up to the full duration of the IGMP join or until it is explicitly stopped by a message sent from the VQE-C to the VQE-S.

RTP and RTCP

VQE relies on RTP and RTCP. RTP is used to carry video packets over multicast streams from the video headend to the VQE-Cs on the STBs. It is also used to transport specific video packets between the VQE-S and a VQE-C. RTCP is a signaling protocol used between VQE devices. RTP is the transport baseline for application-layer FEC, Unicast Retransmission, and RCC.

RTP encapsulation is typically the responsibility of real-time encoders and specialized video products, such as a Cisco DCM. These devices often reside in the video headend office (VHO) or super headend (SHE). RTP sequence numbers are assigned to IPTV packets and are unique within a given multicast group or channel.

A growing number of real-time encoders support native RTP and SMPTE 2022 FEC. For those that support UDP only or RTP encapsulation, the Cisco DCM product is recommended. For FEC support, the DCM requires the GbE I/O Board with a FEC daughter card, and DCM software version 5 or higher.

VQE Web Browser-based Tools

VQE also provides three web browser-based tools: VQE Channel Provisioning Tool (VCPT), VQE-S Application Monitoring Tool (AMT), and VQE Client Configuration Delivery Server (VCDS) AMT.

- VCPT is an optional channel-provisioning utility to aid with the channel lineup configuration required by both the VQE-S and the VQE-C. The channel information is in Session Description Protocol (SDP) format. The VCPT sends the channel information to the VQE-Ss and VCDSs. The VCDS servers provide the channel lineup or network configuration to the VQE-Cs on the STBs.
- VQE-S AMT is a browser-based GUI that displays configuration, status, and statistics on the VQE-S processes, the channel lineup, Unicast Retransmission, RCC, Ethernet interfaces, and VQE-S RTCP Exporter. The VQE-S AMT also allows you to configure debugging and logging facilities.
- VCDS AMT is a browser-based GUI that displays configuration, status, and statistics on the VCDS and VQE Tools server. The VCDS AMT also allows you to configure debugging and logging facilities.

The VQE-S and the VCPT are bundled software and hardware solutions. The VQE-S and the VCPT run on separate Cisco Content Delivery Engine 110 (CDE110) appliances. A typical network might consist of multiple CDE110s hosting the VQE-S and one or two CDE110s hosting the VCPT, the VCDS AMT, and the VCDS.

Lookaside Mode and the Cisco CDE110

The VQE-S software functions in *lookaside mode* where the VQE-S is not directly in the video data path. Lookaside mode has these advantages:

- If the Cisco CDE110 appliance or the VQE-S software should fail, there is no loss of IPTV service to the customer.
- Because no multicast forwarding is provided by the VQE-S, more IPTV customers can be serviced by each VQE-S instance.
- Number of Cisco CDE110 appliances hosting the VQE-S can be easily scaled up as subscriber count and the demand for VQE services increase.

The Cisco Content Delivery Engine 110 is a standalone, carrier-hardened appliance running the Linux operating system. The Cisco CDE110 is NEBS-compliant and suitable for central office lights-out locations. The Cisco CDE110 comes with the required software preinstalled: VQE-S or VCPT, Linux, Apache web server, and other software.

VQE Benefits

Cisco VQE Release 3.5 provides the following benefits to the service provider:

- Supports Hybrid Error Repair allowing use of Unicast Retransmission or FEC or both
- Addresses noise issues associated with lossy DSL lines and home network wiring
- Reduces or eliminates the need for outside plant optimization, such as pair swapping, joint renewals, and drop-cable reruns
- Increases addressable market consumer base because Hybrid Error Repair technology enables video service over noisier transmission lines, thus extending available footprint
- Reduces or eliminates the need to fragment video service into consumers that can and cannot receive service based on line quality attributes
- Supports RCC to effectively reduce or eliminate the main sources of channel change delay
- Employs much of the same infrastructure and the same technology (RTCP) for Unicast Retransmission and RCC—thereby reducing service-provider training time
- Reduces or eliminates quality-related service-center calls
- Establishes a video quality baseline for all consumers with granularity per STB
- Provides an end-to-end view of network characteristics from an IPTV delivery perspective
- Uses open, standards-based protocols

Cisco VQE also provides the subscriber with an enhanced video experience with higher and consistent visual and audio quality. Subscribers with noisier transmission lines or longer loop lengths can take advantage of the video offerings, bundles, and unique content of the service provider.

VQE-S

The VQE-S software is hosted on a Cisco CDE110 appliance running a standard Linux operating system. The Cisco CDE110 comes with the required software preinstalled: VQE-S, VQE-S AMT, Linux, Apache web server, and other software.

The VQE-S is responsible for the following functions:

- Creating a channel configuration database using the channel configuration information sent by the VCPT
- Maintaining per-channel and per-component state information
- Handling Unicast Retransmission by caching RTP data streams for channels and sending repair packets to the requesting VQE-Cs on the STBs
- Handling RCC by caching RTP data streams as well as PAT, PMT, ECM, PCR, and sequence information and sending RCC unicast bursts to requesting VQE-Cs when a channel change occurs
- Load balancing VQE-S services across the Cisco CDE110 Ethernet or bond interfaces
- Providing detailed statistics on IPTV delivery down to the STB VQE-C level
- Monitoring the health of VQE-S application processes

Like a regular IP host, the VQE-S joins multicast groups using Internet Group Management Protocol (IGMP). The VQE-S maintains a dedicated buffer for each channel. The VQE-S receives the multicast stream for each channel from upstream, caching a few seconds of the most recently received program content from each. The VQE-S can use the same cache of video to service both Unicast Retransmission requests and RCC requests.

For Unicast Retransmission, when a VQE-C requests retransmission of missing packets, the VQE-S locates them in its cache and, if found, delivers them to the STB through an associated RTP retransmission stream. See [Table H-3 on page H-2 in Appendix H, “VQE Server Performance and Scaling Limits”](#) for information on the VQE-S inbound repair request limits.

Several factors affect how many errors a single VQE-S can repair for a single VQE-C. These factors include the distribution of errors, the bandwidth of the channel, and the size of the jitter buffer in the STB. The VQE-S software includes global and per-client policers to provide sensible rate limits for the traffic associated with Unicast Retransmission.

For RCC, the VQE-S uses much the same infrastructure as Unicast Retransmission with some major additions. The VQE-S can use the same cache of video data for both Unicast Retransmission and RCC. Instead of requesting repair of a specific packet, the VQE-C sends the VQE-S an RTCP message requesting a unicast burst of IPTV packets for a new channel.

In addition, the VQE-S uses its MPEG Parser component to choose the start point of the RCC burst for the new channel. To properly form the RCC burst, the MPEG Parser examines the video data for several pieces of crucial information, caches the information, and provides it to other VQE-S components upon receipt of an RCC request. Before play-out can begin, the VQE-C must receive this information (for example, PAT, PMT, and PCR) to properly display the video. This additional data is sent out of band of the RCC burst and is used to prime the MPEG decoder on the client side.

The VQE-S separates the bandwidth resources that are dedicated to Unicast Retransmission and RCC. A VQE-S configuration parameter (`vqe.vqes.reserved_er_bw`) controls the bandwidth resources that are dedicated to Unicast Retransmission. The parameter allows the amount of resources dedicated to Unicast Retransmission to be reduced so that the resources are available for RCC instead. The VQE-S provides capacity statistics to indicate when Unicast Retransmission and RCC capacity limits have been exceeded. For more information on the VQE-S capacity statistics, see the [“VQE-S Capacity Statistics” section on page H-3 of Appendix H, “VQE Server Performance and Scaling Limits”](#).

The VQE-S supports setting Differentiated Services Code Point (DSCP) values on IPTV-related packets so that specific levels of service can be assigned to RTCP and RTP packets. For RTP traffic, the service provider can configure separate DSCP values for Unicast Retransmission and RCC. DSCP values are set using the VQE Configuration Tool.

For VQE-S performance information, see [Appendix H, “VQE Server Performance and Scaling Limits”](#).

VQE-S High Availability

The VQE-S provides several high-availability mechanisms for resiliency and redundancy:

- VQE-S processes (Control Plane, Data Plane, Multicast Load Balancer, and STUN Server) are constantly monitored by the VQE monitoring process—Process Monitor. If a VQE-S process fails, the Process Monitor automatically attempts to restart it.
- If the Process Monitor itself fails, the Linux initialization process detects this failure and restarts the VQE-S service (Process Monitor).
- If a Cisco CDE110 running a VQE-S fails and there is a redundant, backup CDE110 running a VQE-S that is configured to receive the same multicast streams, the backup VQE-S takes over servicing the streams.
- Use of anycast IP addresses and equal-cost multi-path routing allows multiple VQE-Ss in a single facility to be load balanced between servers and to provide failover protection in case of a server failure. For more information, see the [“Load Balancing and Redundancy with Multiple VQE-S” section on page 2-21](#).

For Multicast Load Balancing, when a multicast stream used for caching on the VQE-S host starts or stops, the VQE-S determines the best Cisco CDE110 Ethernet interface on which to join or leave the multicast group. The VQE-S software distributes the joins across available interfaces to avoid oversubscription. The VQE-S also monitors the status of these interfaces, moving the streams to other interfaces in case of interface failure.

The VQE-S uses Linux system and network services for system initialization, network access, interface status monitoring, and multicast stream reception.

Support for All NAT Mapping Types

VQE supports all Network Address Translation (NAT) mapping types, including address and port-dependent mapping (symmetric NAT). Symmetric NAT is the most restrictive form of NAT behavior.

Deployments where a CPE device is behind a NAT device require a NAT transversal mechanism, such as a Simple Traversal of UDP (User Data Protocol) through NATs (STUN) Server. The STUN Server is included with the VQE-S software, and a STUN Client is included with the VQE-C. When the VQE-C tunes to a new channel, it sends STUN binding requests to the STUN Server on the VQE-S host. The VQE-C sends the requests to the feedback target IP address of the channel and RTP and RTCP retransmission ports as configured on the VQE-S. The VQE-C uses the STUN Server responses to determine whether it is behind a NAT device and what type of NAT device it is.

The VQE-C is optimized to handle a variety of NAT configurations. For example, if the VQE-C determines from the initial STUN responses that it is not behind a NAT device, it turns off NAT mode so that the VQE-C does not send further STUN messages.

The STUN Server that is included with the VQE-S can be turned on or off using a configurable option in the VQE-S configuration file. Unless it is certain that no STBs being serviced by VQE-S are behind a NAT device, we recommend that you enable the STUN Server. For information on enabling the STUN Server, see the [“VQE STUN Server Is Enabled By Default” section on page D-20](#).

VQE-C

The VQE-C software runs on customer premises equipment (CPE), such as a STB. The VQE-C supports Unicast Retransmission, FEC, RCC, and video-quality statistics by providing the following:

- CPE interface to the VQE-S for Unicast Retransmission and RCC
- FEC functionality to receive and decode FEC packets for error repair
- RTP packet reordering
- RTP data plane jitter buffer and de-jittering
- IPTV Packet Loss Monitoring

For Unicast Retransmission, if an error in video transmission occurs, the VQE-C software detects the packet loss and requests a retransmission by sending a RTCP NACK compound packet to the VQE-S while holding the video sequence in queue. The VQE-S automatically repairs the error by transmitting the missing packet, which is resequenced by the STB without interruption. The entire error-repair cycle is imperceptible to the viewer. The RTCP NACK compound packet consists of an RR (Receiver Report) message, an SDES (source description) message, a PUBSPORTS (publish ports) message, if applicable, and a generic RTP feedback (RTPFB) NACK message. Starting with Cisco VQE Release 3.5.5, the RR packets may be omitted to reduce the bandwidth utilized by the RTCP NACK compound packets on the access link.

For RCC, when a subscriber selects a new channel, the VQE-C sends to the VQE-S a special RTCP packet requesting a unicast burst of the IPTV packets for the new channel. As soon as the unicast IPTV packets arrive, the VQE-C software is responsible for sending the packets to the decoder. When the multicast IPTV packets for the new channel begin to arrive, the VQE-C manages the seamless transition between unicast packets from the VQE-S and multicast packets from the headend encoder.

The VQE-Cs can get channel configuration information and per-client network configuration from a VCDS or from a centralized network management/configuration server that supports the DESCRIBE request of the RTSP protocol. For information on the interactions between a VQE-C and a VCDS, see the [“VCDS” section on page 17](#) section.

Each VQE-C can communicate client-specific information to a VQE-S so that some VQE-S services can be optimized for the particular client. Two examples of this capability are as follows:

- **Maximum Receive Bandwidth**—Each VQE-C can communicate a maximum receive bandwidth on the STB to the VQE-S so that VQE-S can precisely set the excess bandwidth available for Unicast Retransmission and RCC. This information is specified with VQE-C system parameters and is used to determine the rate at which VQE-S sends packets to this VQE-C.
- **Decoder Buffer Burst Shaping**—Each VQE-C can communicate the size of the STB decoder hardware buffer using the VQE-C system parameter `max_fastfill`. This size is used for RCC and allows a STB with limited memory to perform RCC. Using the buffer size provided by the VQE-C, the VQE-S modifies the shape of the RCC burst to take advantage of fast decoder fill capabilities.

The VQE-C supports setting Differentiated Services Code Point (DSCP) values on IPTV-related packets it sends so that specific levels of service can be assigned to packets. DSCP values can be configured with the `rtcp_dscp_value` parameter in the VQE-C system configuration file. The setting of `rtcp_dscp_value` applies to both the RTCP messages and the STUN messages. For information on VQE-C system parameters, see *Cisco CDA Visual Quality Experience Client System Configuration Guide*.

Starting with Cisco VQE Release 3.5.5, the VQE-C provides a data path for packets from UDP MPEG2-TS streams and RTPv2 streams that are not configured in the VQE-C channel line up. No VQE services (that is, Error Repair and RCC) are provided for these streams. For UDP streams, the VQE-C acts as a pass-through device. For RTP streams not configured in the channel line up, the VQE-C provides reordering of packets.

The VQE-C supports the monitoring and reporting of a subset of TR-135 statistics. The technical report *TR-135—Data Model for a TR-069 Enabled STB* defines a hierarchical model of statistics that may be collected by remote management stations for the purposes of monitoring and troubleshooting STB services. The VQE-C is in an ideal position to support the collection of the subset of these statistics that correspond to the IPTV and VoD streams that the VQE-C buffers. For more information on configuring the VQE-C to collect TR-135 statistics, see *VQE-C System Integration Guide*. For information on TR-135 data model statistics, the TR-135 data model can be found at this url: <http://www.broadband-forum.org/technical/download/TR-135.pdf>

The VQE-C supports statistics counters that can be utilized to implement SNMP MIB views. The VQE-C provides APIs to export tuner and channel cumulative MIB-compatible counters. For more information on creation of SNMP MIB Views of VQE-C counters, see *VQE-C System Integration Guide*.

The VQE-C is available in two deployment models:

- VQE-C is integrated into selected Cisco STBs.
- VQE-C code is offered through an open-source program for integration with STBs from third-party vendors.

For information on VQE-C deployment models, see “[VQE Deployment Options and Requirements](#)” section on page 14.

VQE-C Forward Error Correction

The VQE-C supports both one-dimension FEC (one FEC stream) and two-dimension FEC streams (two FEC streams). The Cisco DCM or other SMPTE 2022 compliant FEC stream generator sends one primary media stream and one or two FEC streams over different UDP ports to VQE-C. The VQE-C receives and processes these FEC streams to provide packet-level error repair.

Before receiving the FEC packets, the VQE-C must learn some basic information about the FEC session, such as the IP address and port numbers of FEC streams. This information is obtained through the channel configuration file that is sent to the VQE-C from the channel provisioning server, such as VCDS. The FEC streams are configured on a per-channel basis.

When a VQE-C detects packet loss on a channel that is configured for FEC and Unicast Retransmission, the following occurs:

- VQE-C first tries to repair the lost packets using the FEC streams.
- If some packet losses cannot be corrected by FEC, the VQE-C requests a Unicast Retransmission of the missing packets from the VQE-S.

The use of one-dimension (1-D) or two-dimension (2-D) FEC is configured when the channel is defined. 1-D FEC uses one FEC stream, and 2-D FEC uses two FEC streams. While 2-D FEC can correct more packet losses than 1-D FEC, 2-D FEC sometimes requires more intensive processing of both FEC

streams to maximize the number of packets recovered. FEC bandwidth overheads may be too high in some deployments for full FEC-based error repair deployment. FEC can be turned on or off on a per-channel basis through the VQE-C channel configuration file.

The VQE-C supports several extensions to the SMPTE 2022 standard. The VQE-C autodetects L (column) and D (row) values, and allows any combination of L and D sizes where $L * D \leq 256$. This is an extension of the SMPTE 2022 limit of $L * D \leq 100$. The VQE-C also allows any payload value to identify FEC packets where the standard says the value must be equal to 96.

The VQE-C supports SMPTE 2022 Annex A and Annex B stream orderings. For detailed information on 1-D and 2-D FEC, see SMPTE standard *Forward Error Correction for Real-Time Video/Audio Transport Over IP Networks* (SMPTE 2022-1-2007). The standard is available for purchase at this URL:

<http://www.smpte.org/standards>

VQE-C IPTV Packet Loss Monitoring

When used with VQE-S RTCP Exporter, the VQE-C software also provides the instrumentation for IPTV Packet Loss Monitoring and valuable *last hop* analysis. The VQE-C generates RTP packet-level statistics for packet loss, jitter, delay, and other quality measurements. The VQE-C provides statistics on both Unicast Retransmission and FEC. For the RTCP reports, the VQE-C sends RTCP compound packets to their target VQE-S. Each compound packet contains an RTCP receiver report as well as other information.

For more information on IPTV Packet Loss Monitoring, see the “[VQE-S RTCP Exporter for Video-Quality Monitoring](#)” section on page 23.

VQE-C Software Development Kit and Documentation

The VQE-C consists of a software development kit (SDK), which can be used for VQE-C integration into STBs from third-party vendors. The VQE-C code and SDK is available to third-party vendors through an open-source program. The VQE-C code resembles a standard Linux software component. The VQE-C source code is currently supported for the Linux operating system. For information on support for other operating systems, contact your Cisco account representative.

The VQE-C library provides a set of high-level APIs designed to support easy integration into an existing STB software base. The programmatic interface provides a *socket replacement* interface, which is used to get packets from a repaired VQE-S server-enhanced video stream. The VQE-C also provides APIs for updating its channel configuration data and acquiring statistics on error repairs.

The integrator configures the VQE-C through system configuration file parameters. The parameters allow customizing of many elements of the VQE-C system (for example, number of concurrent streams and client policing). The configuration of a VQE-C must be coordinated with the configuration of the VQE-S. Certain features are operational only when they are enabled on both the VQE-C and the VQE-S.

The VQE-C command-line interface (CLI), based on the open source library `libcli` (<http://sourceforge.net/projects/libcli>), is designed primarily for testing and debugging the VQE-C software on the STB. The scope of the CLI is limited to the VQE-C software only. The CLI is accessible by Telnet.

The VQE-C SDK and documentation can be downloaded from Cisco.com. See the “[Related Documentation](#)” section on page -xi for a list of the VQE-C documentation that is available.

VQE Deployment Options and Requirements

The two basic deployment options for VQE are as follows:

- VQE reference architecture model—For existing IPTV deployments or new IPTV deployments that do not use Cisco STBs
- Cisco end-to-end IPTV solution model—For new IPTV opportunities

With both deployment options, the VQE-S is deployed on a Cisco CDE110 appliance running Linux.

VQE Reference Architecture Model

The VQE reference architecture model is designed for existing IPTV deployments or new IPTV deployments that do not use Cisco STBs. Cisco offers the VQE-C as open-source software. The VQE-C is implemented so that service providers and CPE device vendors can integrate the VQE-C software with third-party STBs. Appropriate development-level documentation is available along with the VQE-C code.

With this model, the service provider uses the VCPT to define channels and servers and to create channel lineups for different subscriber regions. The VCPT sends the channel information to the VQE-Ss, and to the VCDSs or a remote server from which each VQE-C gets its channel information. The channel information is in the Session Description Protocol (SDP) format required by the VQE-S and the VCDS.

Cisco End-to-End IPTV Solution Model

The Cisco end-to-end IPTV solution model is designed for new or *greenfield* IPTV opportunities. VQE technology is included as an integral part of the Cisco end-to-end video solution.

In this model, the VQE-C is integrated with selected Cisco STBs. The main difference between the Cisco end-to-end IPTV solution and the VQE reference architecture models lies with the integration responsibility of the VQE-C.

- For the Cisco end-to-end IPTV solution model, Cisco is responsible for the integration and testing of the VQE-C.
- For the VQE reference architecture model, the third-party vendor is responsible for the integration and testing of the VQE-C.

For the Cisco end-to-end IPTV solution, contact your Cisco representative for details of the STB models and software versions supported.

VQE Deployment Requirements

To deploy VQE, the following prerequisites must be met:

- RTP support—Video streams from the headend must be encapsulated in RTP. Service providers can deploy products such as the Cisco Digital Content Manager (DCM) to provide RTP encapsulation capabilities at the video headend.
- VQE-C and CPE integration— VQE-C must be integrated with the software of the CPE device (STB).

- TV channel information—VQE-C and the VQE-S require details of network-level TV channel lineup information. This includes per-channel IP multicast addresses, port numbers, and some other parameters. The information must be presented to the VQE-C and the VQE-S components in Session Description Protocol (SDP) format. The VCPT is designed specifically for this purpose.
- If forward error correction is used, an SMPTE 2022 compatible device, such as an encoder or Cisco DCM, is required to send FEC streams to the VQE-C.
- VQE-S network connectivity—VQE-S requires a connection to the edge router for the purposes of joining and receiving Internet Group Management Protocol (IGMP) multicast groups (channels). A direct, Layer 3 connection between the VQE-S and the edge router is preferred.
- VQE-C network connectivity—VQE-C requires an IP unicast path to and from its designated VQE-S. The path is used for RTP Control Protocol (RTCP) signaling between the VQE-C and the VQE-S, and for sending RTP data packets from the VQE-S to the VQE-C.
- All versions of IGMP are supported by the VQE-S and the VQE-C.

VCPT and Channel Information

The VCPT is responsible for the creation, maintenance, and distribution of the channel information containing channel-lineup data. The VCPT includes a browser-based GUI that allows the service provider to provision the following:

- Channel definitions—Information on the channels that is serviced by VQE
- Server definitions—Information on each VQE-S, VCDS, and Remote Server that receives the channel information
- Channel lineups—Associations between channels, the VQE-Ss, and the VCDSs

Figure 1-5 shows the details that the service provider defines for each channel using VCPT. The channel details include information that is used for VQE error repair—both Unicast Retransmission and FEC—and for RCC. The option to enable reduced-size RTCP messages is available in Cisco VQE Release 3.5.5 and later releases.

Figure 1-5 VCPT Channel Definition

The screenshot displays the 'Details' window for a VCPT channel definition. The window is divided into two main sections: 'Channel Details' and 'Channel Features'.

Channel Details:

- Channel Name:** Channel 14
- Feedback Target IP:** 10.1.7.1
- Original Stream:**
 - Multicast IP: 224.1.1.7
 - Source IP: 10.10.10.5
 - RTP Port: 1025
 - RTCP Port: 1026
 - Bit Rate: 4000 (kbps)
- Unicast Retransmission Stream:**
 - Source IP: 10.1.7.1
 - RTP Port: 50012
 - RTCP Port: 50013
- FEC 1 Stream:**
 - Source IP: 10.10.10.5
 - Multicast IP: 224.1.8.1
 - RTP Port: 1027
- FEC 2 Stream:**
 - Source IP: 10.10.10.5
 - Multicast IP: 224.1.8.1
 - RTP Port: 1028

Channel Features:

- Enable RTCP
 - Enable Error Repair
 - Enable Rapid Channel Change
 - Enable Extended RTCP Reports
 - Enable Reduced Size RTCP
- Enable FEC 1 Stream
- Enable FEC 2 Stream

At the bottom right of the window, there are 'Update' and 'Cancel' buttons. A small vertical number '254375' is visible on the right edge of the window frame.

The VCPT GUI has a clone capability to simplify and expedite channel information. When the service provider uses the VCPT to define the set of the VQE-Ss that receive the channel information, the VQE-Ss can be grouped based on channel lineups. Using separate VCPT configuration files makes it possible to manage multiple deployments. For example, one VCPT configuration file might be for the channel lineup in one metro region, and another VCPT configuration file might be for the channel lineup in another metro region.

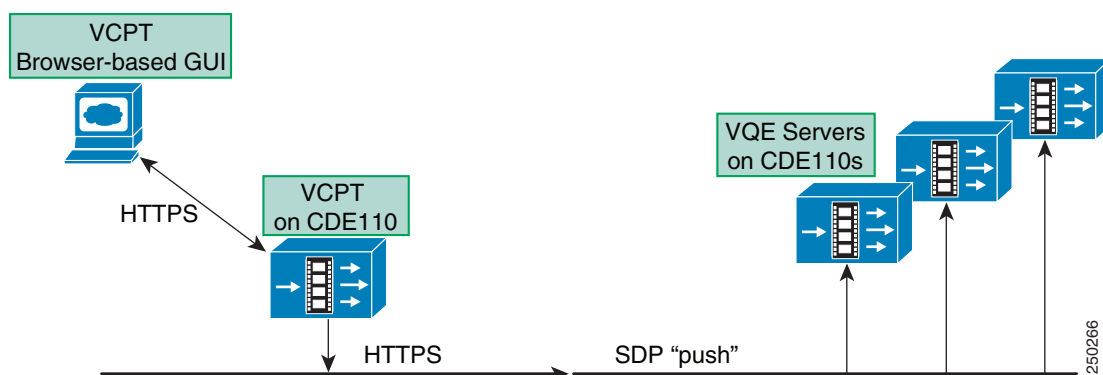
The VCPT GUI has the capability to import a configuration from a file. The file must be in either XML or CSV format, and must comply with the data rules and definitions specified in [Appendix I, “VCPT Configuration Files”](#). The VCPT GUI also has the capability to export a configuration to an XML or CSV file.

The VCPT channel-provisioning process creates a persistent local database, which is stored on the Cisco CDE110 appliance. When the Cisco CDE110 or VCPT is restarted, channel data and server grouping information is read from the local database.

When the user completes channel, server, and channel-lineup configuration and starts the VCPT send operation, the VCPT sends the channel information in Session Description Protocol (SDP) format to the set of VQE-Ss, to the VCDS, or to a Remote Server.

The VCPT sends or *pushes* the channel information to all VQE-Ss that are defined in the current VCPT configuration file. The channel information is sent to the VQE-Ss over secure HTTPS. The VCPT contains a secure HTTPS client, and each VQE-S has an embedded web server running. Each VQE-S stores its own local copy of the channel information. [Figure 1-6](#) shows the interactions between the VCPT and the VQE-Ss. For information on the VQE-Ss, see the “VQE-S” section on page 9.

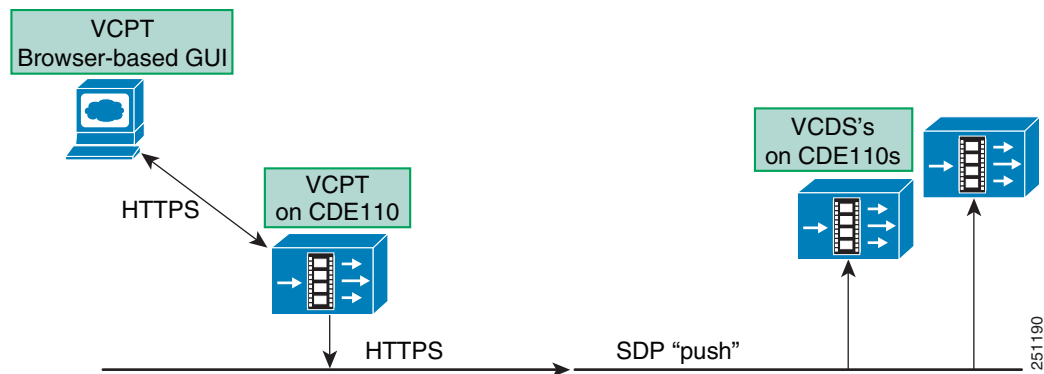
Figure 1-6 VCPT: Sending Channel Information to VQE Servers



The VCPT is also responsible for sending channel information to each VCDS. The VCDS is a software component installed on each VQE Tools server, the Cisco CDE110 that also hosts the VCPT. When the service-provider operator initiates the VCPT send operation, the VCPT *pushes* the channel information to one or more VCDSs. The VCPT sends the channel information in SDP format through HTTPS similar to the way it is sent to the VQE-Ss.

[Figure 1-7](#) shows the interactions between the VCPT and the VCDSs. Each VQE Tools server includes both the VCPT and the VCDS. The VCPT can also send channel information to the VCDS on the same CDE110 on which the VCPT resides. This interaction within a single VQE Tools server (CDE110) is not shown in [Figure 1-7](#).

Figure 1-7 VCPT: Sending Channel Information to VCDSs



An alternative VQE-C Configuration Delivery Server to the VCDS can be used to deploy network and channel information to the access network. The service-provider may have a multicast-based STB file delivery solution already in place. The VCPT can send channel information to a remote server, once the service-provider defines the remote server in the VCPT by providing the following information:

- Management IP address of the remote server.
- Transfer port number to use.
- Full path and filename of the configuration file on the remote server.
- Name of the person with the authority to access this path.

The remote server is treated as another server role, with channels being associated with the remote server. The VCPT *pushes* a channel configuration to the remote server using the password-less secure copy protocol (SCP). From the remote server, you can distribute network and channel configuration data to STBs by integrating your deployment with customized middleware from a third-party vendor.

VCDS

The VCDS is a software component on the VQE Tools server that can be used to deliver the channel configuration file to the VQE-Cs on the STBs, and per-client network configuration files to the VQE-Cs on the STBs. This VCDS function is explained in the following section:

- [VCDS Delivery of the Channel Configuration File, page 1-17](#)
- [VCDS Delivery of the Per-Client Network Configuration File, page 1-18](#)


VCDS Delivery of the Channel Configuration File

When the VCDS receives information on the channel lineup from the VCPT, the VCDS creates a channel configuration file and stores it on the local disk. Each VQE-C running in a CPE device, such as a STB, uses a Real Time Streaming Protocol (RTSP) *pull* operation to receive the channel configuration file from the VCDS. The VCDS is a simplified RTSP server that supports the VQE-C RTSP *pull* operation. The VQE-C learns the name of the VCDS through a Domain Name System (DNS) server using a SRV lookup.

After the VQE-C learns the VCDS name, it sends out an RTSP DESCRIBE request asking for the change status of the channel configuration. If there is new information on channels, the VQE-C asks the VCDS to send the channel configuration file. In response, the VCDS sends the channel configuration file for the entire channel set in SDP format. Figure 1-8 shows the interaction between a VCDS and the VQE-Cs on the subscriber STBs.

For information on the interactions between the VCDS and the other components in the channel delivery infrastructure, see Appendix E, “Configuring DHCP and DNS Servers for VCDS.”

Figure 1-8 VQE-C on STB: Receiving Channel Information File from VCDS



The VQE-C requests the channel configuration file from the VCDS when the STB is started and the VQE-C is initialized. After the initialization, the VQE-C can receive an updated channel configuration file by the following mechanisms:

- **Periodic polls**—VQE-C periodically polls the VCDS to determine whether a new channel configuration file is available. If a new version of the channel configuration file is detected during a poll, the VQE-C requests the updated channel configuration file.
- **Triggered polls**—When the STB attempts to tune to a channel that is not present in the VQE-C channel lineup, the VQE-C attempts to retrieve a new channel configuration file from the VCDS. As it does in a periodic-poll update, the VQE-C first determines whether an updated version of the file is available. If so, the VQE-C requests the updated channel configuration file.

With both periodic polls and triggered polls, the VQE-C uses the new channel configuration file to update the channel lineup that is in use and writes the channel configuration to the file specified in the VQE-C channel-lineup system configuration.

VCDS Delivery of the Per-Client Network Configuration File


The VCDS can provide customized per-client network configuration to the VQE-C on the STB after the VQE-C has been initialized. The VQE-C network configuration might, for example, specify whether RCC or Forward Error Correction is enabled on the VQE-C, or specify some configuration parameters for the VQE-C.

Updates to the VQE-C network configuration occur because of periodic polls. The VQE-C polls the VCDS to check whether the VQE-C network configuration file has been updated. If a new network configuration file is detected during a poll, the VQE-C requests the file from VCDS and overwrites the network configuration file it has stored in flash memory. The new network configuration does not apply until the next STB reboot.

The three major components involved in the delivery of per-client network configuration are the VQE-C system configuration provisioning server, the VCDS, and the VQE-C on the STB. [Figure 1-9](#) shows the interactions between these components. A simplified description of these interactions is as follows:

- VQE-C system configuration provisioning server sends or pushes the client database file and group attribute file to the VCDSs.
- In response to a request from the VQE-C on the STB, the VCDS sends the updated per-client network configuration file to the VQE-C.
- VQE-C receives the updated network configuration file, merges the attributes into its running configuration, and initializes the system using the merged running configuration on the next reboot.

Figure 1-9 Major Components for Delivering the Per-Client Network Configuration File



The following sections provide information on the role of the major components used to deliver the per-client network configuration file for VQE-C system configuration:

- [VQE-C System Configuration Provisioning Server Role, page 1-20](#)
- [VCDS Role, page 1-20](#)
- [VQE-C Role, page 1-20](#)

For information on the interactions between the VCDS and the other components in the VQE configuration delivery infrastructure, see [Appendix E, “Configuring DHCP and DNS Servers for VCDS.”](#)

An alternative to VCDS and CDI may be required for deploying configuration to the access network if the service provider already has a solution in place for multicasting network and channel configuration files to the access network or if the service provider wishes to unicast per-subscriber configuration changes to STBs.

The delivery of the network and channel configuration files may be handled externally to the VQE-C on the STB. The VQE-C provides a set of APIs to receive per-client network configuration data, channel configuration data, and per-subscriber configuration data (that is, an override configuration) from an external client. Deploying network, channel and override configurations to the VQE-C without using VCDS and CDI requires integrating the Cisco’s VQE solution with customized middleware from a third-party vendor.

VQE-C System Configuration Provisioning Server Role

The VQE-C system configuration provisioning server sends or pushes the client database file and group attribute file to one or more VCDSs using HTTPS. The service provider can provide a VQE-C system configuration provisioning server that meets the specific requirements of the deployment.

As an alternative to a VQE-C system configuration provisioning server, the VQE Tools software includes the `vcds_send_file` command that can be used to send the client database file and group attribute file to a VCDS. For information on this command, see [Appendix F, “Using the `vcds_send_file` Command.”](#)

The VQE-C system configuration provisioning server has a different role than the channel-provisioning server (for example VCPT or ISDS). The VQE-C system configuration provisioning server delivers the following files to the VCDSs:

- *Client database file*—In this file, each STB VQE-C has a unique identity (for example, MAC address) that is associated with a group identifier (group ID) for the specific network configuration that should be applied to the VQE-C system. Each VQE-C is associated with one group ID. The client database file can be a complete set or superset of all STBs or a subset which adds to or modifies a previous set that exists on the VCDS.
- *Group attribute file*—In this file, different sets of per-client network configurations are defined. Each set of attributes is identified by a group ID.

The client database file and the group attribute file are XML-based and follow a schema specified by a specification defined by Cisco. On a VQE Tools server, the XML schema and examples of the two files are included with the software. See [Appendix F, “Using the `vcds_send_file` Command”](#) for the locations of the files on a VQE Tools server.

For information on getting specifications for the client database file, group attribute file, and the VQE-C system configuration provisioning server, contact your Cisco account representative.

VCDS Role

When the VQE-C system configuration provisioning server sends or pushes the client database file and the group attribute file to one or more VCDSs, each VCDS does the following:

- Accepts a push of the client database file and the group attribute file from the VQE-C system configuration provisioning server
- Parses the files for validity
- Updates its local copies of the client database file and group attribute file
- Services VQE-C requests for an updated per-client network configuration file

VQE-C Role

To obtain and process the per-client network configuration file, the VQE-C on the STB performs the following tasks:

- By a periodic poll, determines whether an updated network configuration file is available
- Retrieves the updated network configuration file from the VCDS
- Writes the updated network configuration file to local flash memory
- Derives a *running configuration* in memory by merging its existing network configuration file using the attributes in the updated network configuration file
- On the next STB reboot, initializes the VQE-C network configuration using the merged running configuration

For more information on VQE-C's use of a per-client network configuration file, see *Cisco CDA Visual Quality Experience Client System Configuration Guide*.

VQE-S AMT


The VQE-S AMT is a browser-based GUI that allows the service-provider operator to do the following:

- Monitor the health of the VQE-S processes
- View channel configuration details, status, and statistics
- Monitor statistics for Unicast Retransmission and RCC
- Monitor statistics for STUN Server usage
- View configuration details, status, and statistics for:
 - Multicast Load Balancer
 - VQE-S RTCP Exporter
- Change VQE-S logging levels and debugging options

The next paragraphs provide a few examples of VQE-S AMT functionality.

When you log in to VQE-S AMT, the initial window (see [Figure 1-10](#)) shows the health of VQE-S processes and other status information.

Figure 1-10 Monitoring VQE-S Processes



253690

With VQE-S AMT, you can view the channel lineup that was sent from the VCPT. Figure 1-11 shows a partial example of the channel lineup with usage statistics that VQE-S AMT displays.

Figure 1-11 Viewing Channel Lineups with Usage Statistics

Channel Lineups <small>Last update: 2008-07-08T14:44:14, Total Channels: 250, Active Channels: 240 Aggregated Bandwidth: 1215000(kbits/sec), Total Receivers: 1040</small>							
Name	Status	Original Multicast	Feedback Target	Bit Rate (Kbps)	Repair Enabled	RCC Enabled	Member Receiver Population
Channel 229.1.1.1	✓	229.1.1.1:53170	11.112.1.1:53171	4000	true	true	400
Channel 229.1.1.10	✓	229.1.1.10:53206	11.112.1.10:53207	13000	true	true	0
Channel 229.1.1.11	✓	229.1.1.11:53210	11.112.1.11:53211	4000	true	true	0
Channel 229.1.1.12	✓	229.1.1.12:53214	11.112.1.12:53215	4000	true	true	0
Channel 229.1.1.13	✓	229.1.1.13:53218	11.112.1.13:53219	4000	true	true	0
Channel 229.1.1.14	✓	229.1.1.14:53222	11.112.1.14:53223	4000	true	true	0
Channel 229.1.1.15	✓	229.1.1.15:53226	11.112.1.15:53227	2000	true	true	0

In Figure 1-11, the channel-lineup summary data indicates when the lineup was last updated (for example, with the VCPT) and provides totals for all channels and active channels as well as aggregated bandwidth and total receivers:

Last update: 2008-07-08T14:44:14, Total Channels: 250, Active Channels: 240
Aggregated Bandwidth: 1215000 (kbits/sec), Total Receivers: 1040

In the channel-lineup summary data, the rightmost column, Member Receiver Population, is the number of VQE-Cs that are currently receiving this multicast stream.

The VQE-S AMT uses the Unicast Retransmission and RCC counters kept by the VQE-S to display a variety of data. For Unicast Retransmissions, the counters include non-acknowledgement (NACK) messages received from VQE-Cs, RTP packets requested and sent, error repair rates, and requests exceeding Unicast Retransmission capacity limits. For RCC, the counters include the number of RCCs requested, accepted, and refused, as well as the number of requests exceeding RCC capacity limits.

The Cisco VQE-S AMT provides limited configuration capabilities. The items that can be configured with the VQE-S AMT include parameters for the following:

- Logging priority level for the VQE-S processes
- Debugging options for the VQE-S-related functions, including RTP/RTCP packets, events, and errors

These configuration changes are temporary in nature, and only remain in effect until the VQE-S restarts.

In addition, the VQE-S AMT allows you to reset a subset of VQE-S counters to zero and subsequently, to restore these counters to their true values.

The VQE-S channel lineup is stored locally on the Cisco CDE110 appliance. If the VQE-S is restarted, the channel lineup is read from the local repository. The VQE-S counters for statistics that the VQE-S AMT displays are reset to zero when the VQE-S is restarted.

The VQE-S AMT is a web application that uses the application server and web server that are preinstalled on the Cisco CDE110 where the VQE-S runs. The VQE-S AMT has an XML-RPC client that communicates with multiple internal applications, such as VQE-S processes, to send and receive application management data.

VQE-S RTCP Exporter for Video-Quality Monitoring

The VQE-S provides a variety of data for monitoring IPTV packet delivery and for isolating faults. The VQE-S receives RTCP reports from the VQE-Cs on the CPE devices and from reports generated by the VQE-S itself. In those reports, the VQE-Cs provide statistics for RTP packet loss, jitter, delay, UDP packet loss, and other quality measurements. The VQE-Cs also provide statistics on Unicast Retransmission.

For the RTCP reports, each VQE-C periodically sends RTCP compound packets to its target VQE-S. Each compound packet contains an RR, an SDES, PUBPORTs, if present, and a generic RTP feedback NACK, if present. The VQE-C sends additional RTCP reports every time a Unicast Retransmission request is made. Starting with Cisco VQE Release 3.5.5, the service provider may omit RRs from the RTCP NACK compound packet for each channel using the VCPT. In addition to the RTCP compound packet, each VQE-C may transmit pre- and post-repair RTCP XR reports to its target VQE-S if the option to enable XR reports is enabled for each channel using the VCPT.

The service provider can use VQE-S RTCP Exporter to export the RTCP compound packets and the XR reports to a video-quality monitoring (VQM) application. Starting with Cisco VQE Release 3.5.5, the service provider can exclude RTCP NACK compound packets from being exported to the VQM application by setting the VCDB parameter `vqe.vqes.exporter_filter_nack` to TRUE. The VQM application collects the exported data in a database for use in video-quality analysis. The video-quality monitoring application is outside the scope of the VQE solution. The VQE documentation set includes detailed information on the data collected and the formats used in the RTCP reports.

The VQE-S RTCP Exporter is responsible for sending the RTCP compound packets to an external device, which typically hosts the video-quality monitoring application. The compound packets are sent over a TCP socket to a configurable location. The monitoring application is identified by IP address or Internet domain name and a TCP port number.

The data in the RTCP compound packets are very useful for determining the quality of video service and for isolating faults. The data help the service provider to measure, baseline, and pinpoint problem areas of the video infrastructure, including transmission lines and home networks. The granularity of the data is per STB. The data could be stored in a database and searched for answers to questions of interest, such as whether packet loss and jitter events have occurred in the network and, if so, where and when the events have occurred.

On the VQE-S, RTCP Extended Reports and the Extended Report (XR) packet type are supported. The following XR report block types are supported:

- Loss RLE (run-length encoded)
- Statistics Summary
- Post-Repair Loss RLE
- Multicast Acquisition
- Diagnostic Counters (available in Cisco VQE Release 3.5.5 and later releases)

Extended Reports provide information that supplements the statistics contained in the report blocks used by the RTCP sender and receiver reports. For example, the Loss RLE report block type provides much more detailed reporting on individual packets and loss events than is provided in standard RTCP reports. The VCPT allows the service provider to specify whether or not RTCP Extended Reports are used for each channel.

The following documents, which are available from <http://www.faqs.org/faqs/>, provide more information on RTCP reports:

- RTCP reports are described in RFC 3550.
- RTCP Extended Reports are described in RFC 3611.

The Post-Repair Loss RLE report block type contains information on individual packet receipt and loss events, after packet recovery techniques (Unicast Retransmission or FEC or both) have been applied. For more information, see *Post-Repair Loss RLE Report Block Type for RTCP XR* at:

<http://tools.ietf.org/html/draft-begen-avt-post-repair-rtcp-xr-01>

The Multicast Acquisition report block type is also capable of supporting information on the total channel change time, the expected RCC presentation timestamp and the actual RCC presentation timestamp. The availability of this timing information is dependent on the specific VQE-C platform integration. For more information, see *Multicast Acquisition Report Block Type for RTCP XR* at:

<http://tools.ietf.org/id/draft-begen-avt-rapid-sync-rtcp-xr-01.txt>

Starting with Cisco VQE Release 3.5.5, the Diagnostic Counters XR report block type contains information on cumulative overruns and underruns, post-repair loss events, late packet drops, and output queue drops per channel.

VCDS AMT

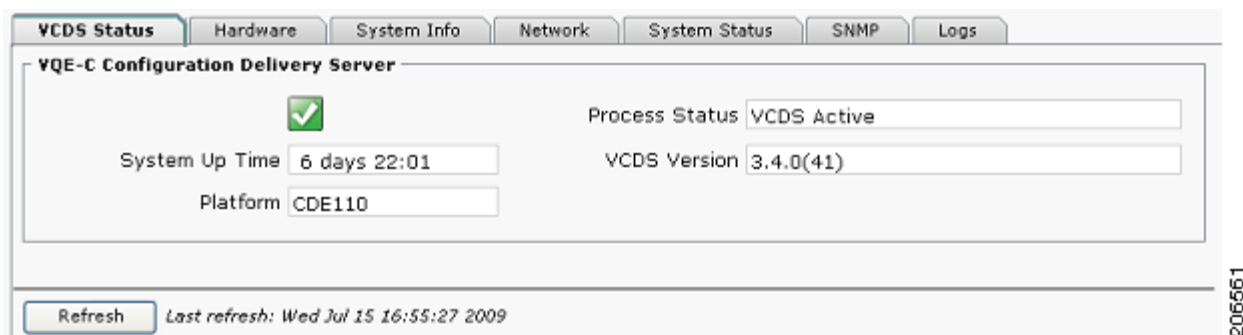
The VCDS AMT is a browser-based GUI that allows the service-provider operator to do the following:

- Monitor the health of the VCDS process and the VQE Tools Server
- View configuration details and statistics for the VCDS
- Configure VCDS logging levels and debugging options

The next paragraphs provide a few examples of the VCDS AMT functionality.

When you log into VCDS AMT, the initial window (see [Figure 1-12](#)) shows the health of the VCDS process, the VQE Tools server, and other status information.

Figure 1-12 VCDS Status Information



The VCDS AMT displays status and configuration on the VCDS configuration, and on the VQE-C channel configuration, the client database, and the group attribute files (see [Figure 1-13](#)).

Figure 1-13 VCDS Configuration Window

Section	Parameter	Value
VCDS Configuration	TCP Port	8554
	Max Connections Allowed	2000
	Inactive Connection Timeout (sec)	10
VQE-C Channel Lineup	Validation Check	Pass
	No. of Channels in Lineup	1
	Lineup Filename	/etc/opt/vqes/vqec_channels.cfg
	Last Modified	Thu, 9 Jul 2009 15:27:15 EDT
VQE-C Group Map	Validation Check	Pass
	No. of CNames in Map	2
	Map Filename	/etc/opt/vqes/vqec_group_map
	Last Modified	Mon, 13 Jul 2009 16:26:27 EDT
VQE-C Network Configuration	Validation Check	Pass
	No. of Groups in Network Configuration File	2
	Network Configuration Filename	/etc/opt/vqes/vqec_attributes.cfg
	Last Modified	Mon, 13 Jul 2009 17:20:30 EDT

The VCDS AMT displays configuration and status information for the channel configuration, the client database, and the group attribute files. For the channel configuration file, the VCDS AMT indicates whether the file contents are valid. It displays the full pathname of the file, the number of channels contained in the channel lineup and the timestamp of the last modification to the file. For the client database file, the VCDS AMT indicates whether the file contents are valid. It displays the full pathname of the file, the number of CNames (VQE-C unique identifier) in the file and the timestamp of the last modification to the file. For the group attribute file, the VCDS AMT indicates whether the file contents are valid. It displays the full pathname of the file, the number of attribute groups in the file and the timestamp of the last modification to the file.

The VCDS AMT provides limited configuration capabilities. The items that can be configured with the VCDS AMT include parameters for the following:


- Logging priority level for VCDS processes
- Debugging options for VCDS-related functions

The VCDS AMT is a web application that uses the application server and web server that are preinstalled on the Cisco CDE110 where the VCDS runs. The VCDS AMT has an XML-RPC client that communicates with multiple internal applications, such as VCDS processes, to send and receive application management data.

Content Delivery Engine 110

Each VQE-S runs on one Cisco Content Delivery Engine 110 (CDE110). If the VCPT and the VCDS are used, another Cisco CDE110 hosts these two facilities. The Cisco CDE110, shown in [Figure 1-14](#), comes with the Red Hat Enterprise Linux Release 5.1 operating system and either the VQE-S or the VCPT and the VCDS software preinstalled.

Figure 1-14 Content Delivery Engine 110



Note

The descriptions of hardware given in the following paragraphs are for the latest CDE110 models (CDE111-2-146TXA-K9 and CDE111-2-146TXD-K9). For complete information on all CDE110 models, see the *Cisco Content Delivery Engine 110 Hardware Installation Guide*.

The Cisco CDE110 appliance is a NEBS-3 and ETSI-compliant carrier-grade rack server. It is powered by two 64-bit Quad-Core Intel Xeon L5410 processors with 12 MB of shared Layer 2 cache. For maximized bandwidth, it contains 8 GB of dual-channel Fully Buffered DIMM (FB-DIMM) memory at 667 MHz. For storage, the Cisco CDE110 has one 36-GB simple-swap, serial attached SCSI (SAS) hard disk drive. The optical drive is a CD/DVD RW combination drive.

The Cisco CDE110 has six 10/100/1000 Mb Ethernet ports and a serial port for the system console. Earlier CDE110 models have four 10/100/1000 Mb Ethernet ports. The Ethernet ports can be load-balanced for incoming multicast IPTV streams and for outgoing Unicast Retransmission and RCC streams to the VQE-Cs.



Note

The VQE-S supports Gigabit Ethernet (that is, 1000 Mb Ethernet) interfaces only.

The Cisco CDE110 has a 1-RU form factor and is available with redundant AC or DC hot-swappable power supplies. The Cisco CDE110 Telco Alarm Management features provide visual, audible (optional), and SNMP event indications of faults, consistent with the rigid requirements of the telecom central office environment.

