

Delivering Video Quality in Your IPTV Deployment

Defining Video Quality

As service providers begin to deploy IPTV and interactive video applications, they must not only meet but surpass the existing video quality expectations of their subscribers to acquire and retain market share. Although this benchmark may vary as a function of geography, the accepted industry benchmark is to deliver a maximum of 1 video “artifact” or perceived distortion during the viewing of a 2-hour movie. This very stringent “application” layer requirement can be mapped to a network layer requirement of less than 1.0E-6 video packet loss. It is critical that all components in the overall video delivery system provide comprehensive end-to-end quality of experience (QoE) for all subscribers.

Key Factors That Affect Video Quality of Experience

One of the factors that affect video quality is the reliability and robustness of the video delivery network. Measurable packet loss can occur in numerous video network-level failure scenarios. A critical goal for delivery of high-quality video is consistent network recovery in all failure scenarios that aim to minimize the associated effects of packet loss.

Beyond controlling the effects of network-level failures, Cisco® has developed a series of innovations instilling greater service intelligence capabilities in the network to further improve the video quality experience. The Cisco Visual Quality of Experience (VQE) technology offers service providers additional capabilities to deliver entertainment-grade services to their subscribers. This advanced technology helps mitigate the following challenges faced by IPTV service providers:

- Network congestion
- Bit errors on access lines (for example, DSL) and optical trunks
- Rapid channel change time (CCT)
- Video monitoring and reporting

The following sections explore each of the primary factors that affect the overall video QoE; recommendations to enable service providers to optimize their deployments are included.

Network-Level Outages

The service provider’s goal is to deliver competitive service offerings while meeting their customers’ service-level agreements (SLAs). Although service provider network deployments can be designed with highly reliable components and many degrees of redundancy, it is impossible to completely eliminate all potential network-level failures. As such, recovery from network-level outages must not only minimize effects on video applications but also be consistent and deterministic in nature.

Many network-level outages can affect the delivery path of broadcast (multicast) and video-on-demand (VoD) traffic. A comprehensive solution must be able to manage all these potential failures, in parallel, to retain the integrity of the visual experience. These potential network outages include:

- Scheduled upgrades

- Trunk or ring link failure
- Video Headend office (VHO) node failure
- Video Switching office (VSO) node failure
- Super Headend (SHE) outages

Designers can take a variety of approaches when developing a video delivery and distribution network. A popular approach is to deliver video with Layer 3 forwarding intelligence across the entire aggregation network. Employing Layer 3 in the video distribution network relies on Protocol Independent Multicast (PIM) with Source Specific Multicast (SSM) and Interior Gateway Protocol (IGP) with enhancements that enable fast network-level convergence. These Layer 3 fast-convergence techniques provide consistent subsecond network and application recovery for all failure modes.

An alternative to employing Layer 3 for video services in the aggregation network advocates the use of Layer 2 Hierarchical Virtual Private LAN Service (H-VPLS). In this case, Multiprotocol Label Switching (MPLS) Fast Reroute (FRR) is employed to recover from trunk or link failures. However, trunk or ring link failures account for only a very small percentage of the overall network outages. In ring-based failure scenarios, the H-VPLS nodes loop all traffic back around the ring, using twice the bandwidth and potentially causing additional network congestion.

Although MPLS FRR can support recovery from single link failures, it does not provide the network recovery for failures either in the VSO or, even more significantly, in the regional VHO. For these sorts of unscheduled system failures, a network based on H-VPLS must rely on interactions between the Layer 2 H-VPLS nodes and the Layer 3 PIM nodes. In effect, a video network based on H-VPLS relies on the Internet Group Management Protocol (IGMP) designated router election process for most failure modes, resulting in a 2–3 second video application outage that is highly visible to many thousands of subscribers.

The Cisco solution for video delivery is based on the Cisco IP NGN Carrier Ethernet Design, which promotes the delivery of video with Layer 3 forwarding intelligence all the way to the aggregation nodes connecting the access devices (for example, DSL access multiplexer [DSLAM]). Network outages of small durations typically cause small amounts of video artifacts (pixelization, macro-blocking, and predictor errors). However, longer network outages tend to cause more severe video artifacts (for example, image freeze) that are more noticeable to many subscribers. When looking at the combination of all potential network failures, Cisco IP NGN Carrier Ethernet Design provides sub-second network recovery that delivers significantly more robust and reliable video delivery to all subscribers. In addition to the consistent network recovery capabilities, the Cisco IP NGN Carrier Ethernet design provides more efficient video delivery by optimizing the use of bandwidth even during link, node, and video super head-end failures. Network congestion due to network failures may also cause packet loss, thereby affecting video network quality. Predicting the amount of VoD traffic required by the video distribution network may not meet consumer peak usage rates and will potentially create network congestion. Employing Layer 3 intelligence in the aggregation network provides further safeguards that use the innovations of Cisco VQE technology.

Cisco VQE Technology

Cisco Integrated Video Admission Control

Service providers that deploy video application services must optimize their network deployments to serve subscribers with a quality experience while managing their operational and capital cost structure. As a result, service providers often oversubscribe their networks; the network is

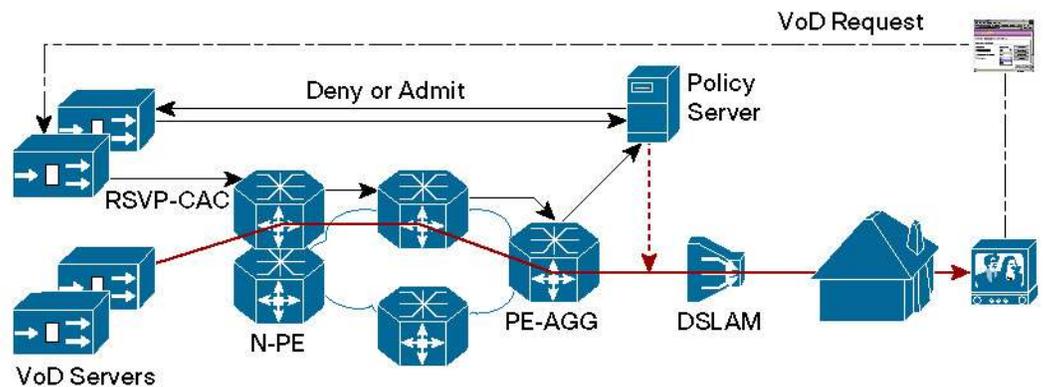
engineered for a peak bandwidth less than the sum of bandwidth available to individual subscribers – with the assumption that not all subscribers will need their bandwidth, even during peak periods. Although this design is cost-effective, it introduces the possibility of congestion. Service providers need to apply various traffic-management techniques to assure that service quality is maintained for all subscribers. Cisco VQE offers advances that providers can take advantage of to avoid network congestion by using a Cisco Integrated Video Admission Control solution.

Service providers typically engineer their networks with a set amount of bandwidth for broadcast TV on each segment, including both standard-definition (SD) and high-definition (HD) television channels. VoD traffic is estimated at a peak concurrency rate but is somewhat more difficult to predict. Given the success of VoD services, the growth of high-definition VoD services, and collaborative over-the-top (OTT) video traffic, the bandwidth required to transport the video services will continue to increase and can result in unexpected congestion in the network.

With transport-based services (where the user SLA is defined by transport parameters such as 5-Mbps high-speed Internet [HSI] service with a potentially different network QoS SLA from user to user), per-subscriber shaping can be used to control network congestion. However, in reality, video, which is a managed application service (where the SLA is defined by the application experience such as 200 channels of digital broadcast video), cannot be shaped because shaping may result in packet loss. Congestion due to network oversubscription is not acceptable for entertainment-grade video. For example, on a Friday night, the typical VoD peak, admitting just one incremental stream to a network has the potential to cause congestion and could degrade the VoD and broadcast quality for *all* users who are in the middle of using a service. Again, unlike High Sped Internet (HSI), video traffic cannot be shaped when congestion occurs and need to be prevented to benefit the largest number of subscribers.

The optimal way to solve this type of network congestion is to use application intelligence to deny a new subscriber's VoD request when the network is operating at full capacity – to prevent congestion in the first place. This solution affects only this one subscriber by delaying the request until sufficient bandwidth becomes available and simultaneously retains the viewing experience for all those that were subscribed to the service. Because video cannot be policed or shaped, per-subscriber QoS adds unnecessary complexity and the solution proposed treats all VoD traffic with one level of differentiated services code point (DSCP)-based QoS. The critical innovation is enabling the video distribution network elements to communicate directly with the VoD server and perform video Call Admission Control (CAC) to prevent network congestion associated with admitting too many video streams. This integrated admission control solution designed primarily for VoD services must interoperate with complex network topologies that have redundant and load-sharing paths in the transport layer of the network. The on-path video CAC mechanisms must also work with access links and business policies that may be enforcing other types of constraints on a subscriber's service. The Cisco Integrated Video Admission Control solution for VoD works in coordination with network routers, policy, and VoD servers to collectively improve the visual QoE.

Two important simultaneous processes combine to enable Cisco Integrated Video Admission Control. The first process is an on-path signaling procedure that performs admission control for the core and distribution layers, as illustrated in Figure 1.

Figure 1. Cisco Integrated Video Admission Control (On-Path)

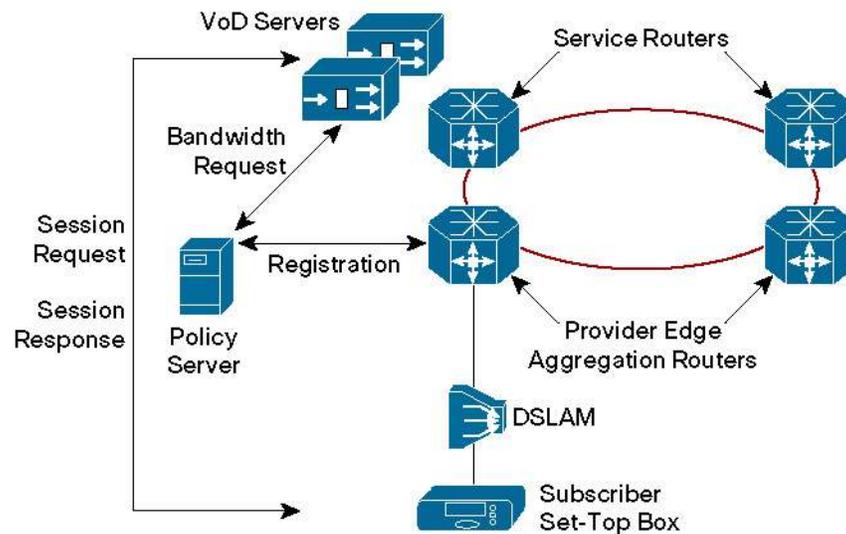
The on-path procedure uses Resource Reservation Protocol (RSVP) for signaling, sent by the VoD server (or a component on its behalf) before the beginning of the video session.¹

Subscribers press the remote control for their set-top box (STB) to select an on-demand movie. The request for a movie goes through the provider edge aggregation router and on to the video server. The RSVP message from the VoD server follows the video entitlement management message (EMM) from the middleware application server to the subscriber and traverses the same path that the VoD session will use, while tracking any real-time changes in the complex network topologies in both the core and distribution layers. Along the path, Cisco routers perform a network bandwidth accounting function. The routers allow the video session for the on-demand movie if enough bandwidth is available and deny the request if insufficient bandwidth is available to support the video stream. When a stream is denied, an RSVP-CAC message is sent back to the VoD server, which in turns sends the subscriber a message or a busy signal, in concert with the video middleware application client-server software.

It is critical that native Layer 3 forwarding intelligence be present on every network element, from the video distribution router connecting the VoD server(s) to the aggregation router in the central offices, to enable this on-path admission control solution. All provider edge aggregation routers (PE-AGG) illustrated in Figure 1 support this level of Layer 3 forwarding and application intelligence. Network delivery architectures using Layer 2 mechanisms (for example, H-VPLS and IEEE 802.1ad) in aggregation networks cannot support this type of dynamic intelligent on-path video admission control.

A complementary process of integrated connection admission control is designed to prevent a VoD video stream from being transported if the access link to the subscriber lacks the capacity to carry the stream. To provide this additional subscriber safeguard, the VoD server (in concert with appropriate middleware application components) sends a request to the Cisco Business Policy Manager (BPM) (or equivalent policy server), which is keeping track of each access link to the subscriber. Subscriber access links are usually configured in a simple and semistatic manner, highly appropriate for an off-path admission control approach illustrated in Figure 2.

¹ RSVP in the context of a metropolitan network with tens of thousands of streams does not have the same scaling problems that might have occurred when RSVP was in the core for millions of telephony streams.

Figure 2. Access Link Admission Control (Off-Path)

The policy server can check to see if the access link has enough unused bandwidth and can also check to see if business policies are in place (that is, the subscriber is entitled or is paying for this service) to allow the stream to be supported. The requested session is then either allowed or denied on the access link. This off-path portion of the solution is not recommended to perform admission control for the distribution and aggregation layers, because the policy server-based approach to tracking any real-time changes in these complex topologies would not meet the response requirements of the consumers. However, combination of an integrated on-path admission control function in the aggregation network and an off-path policy server query at the network edge has been shown to be the most reliable and efficient mechanism for video admission control.

Access Network – Video Error Repair

One of the most significant concerns for IPTV providers delivering video to subscribers over copper lines is the amount of overall line errors that can occur on these subscriber access lines (for example, DSL). In fact, the effect of DSL access-line errors on video network quality far exceeds the combined effect of all network outages described earlier. Most providers have built out their networks to deliver 10-4 packet loss across their broadband subscriber access lines. These access networks were initially deployed to support HSI services that are much more tolerant of packet loss and can take advantage of TCP retransmission characteristics.

As mentioned earlier, quality video is intolerant of packet loss and requires a packet loss of 10⁻⁶ delivered to the video decoder contained in the STB or other video decoding device. Bit errors on access lines actually cause significant video packet loss, and with a bit error rate of 10E-6, one may see an average of 1 artifact per minute. In fact, DSL line errors of 10e-6 are a couple of orders of magnitude below the desired threshold for high quality video delivery. As such, these line errors limit the deployment of high-quality video to a percentage of available access loops in order to meet the packet-loss requirements. As a result, wireline service providers are faced with a smaller addressable market offering fewer subscriber acquisitions – a challenge when marketing the service to all subscribers. Even if market geographies are certified for IPTV service delivery, the operator still runs the risk of DSL line errors causing video quality degradation on many of their loops, which can be the result of temperature, humidity, or intermittent noise that can affect the

overall video experience and further affect subscriber retention or cause negative word-of-mouth discussions that result in lost opportunities.

The Cisco VQE technology innovations offer providers an error-repair capability to address these line faults. Cisco has taken advantage of the end-to-end video aggregation solution, which links the aggregation network to the STB, to mitigate access-line errors. The ingredients of this solution are the STB, its buffer, its monitoring capability, and an intelligent video delivery network to deliver 10-6 packet loss over a larger portion of the subscriber access lines. The solution allows the STB to monitor video packet loss using the standardized Real-Time Transport Control Protocol (RTCP). When packet loss is detected, the STB makes a packet retransmission request for the lost packet. The video delivery network intercepts this request and delivers a replacement packet to the STB before the location of the packet (that is, its place in the packet order) in the STB buffer has been reached by the video decoder. In effect, this solution combines innovative packet monitoring and intelligent packet retransmission mechanisms to improve consumer visual QoE.

The Cisco Visual Quality of Experience (VQE) technology significantly improves the effective video network quality by reducing effective recurring video outages caused by access-line errors. These capabilities translate into numerous benefits for the video service provider:

- Increased data rates over a larger percentage of the subscriber access lines, resulting in a larger addressable market, more subscribers, increased revenue, and simplified marketing planning
- Dramatically improved video quality over the entire video aggregation and access network, resulting in a marketing advantage over the competition and enhanced subscriber retention

In effect, the Cisco VQE technology for error recovery improves the effective network video quality but also ensures that all subscribers are provided with a consistent entertainment experience.

Standards-Based Rapid Channel Change

Service providers are continually providing consumers with more video content by increasing the available video channels for broadcast and on-demand viewing. As a result, the consumers are continuously searching for new content to view next, increasing the amount of channel change requests received from the remote control of each STB. In addition, consumers have become used to the current CCT – which is virtually instant – offered by analog, digital cable, and digital satellite services. In general, consumers are impatient and expect immediate response to their viewing requests.

CCT is defined as the amount of time it takes from the initiation of a channel change request by the subscriber to the time the video of the new channel is displayed. In an IPTV video network, each TV channel is carried in a separate multicast stream. The IP-aware STBs use the IGMP protocol to switch between the channels. When a subscriber requests a channel change, the STB starts by sending an IGMP “leave” message upstream toward the multicast source. The STB then sends an IGMP “join” message upstream to request transmission of the “new” channel. After some time the IP Multicast stream corresponding to the new channel starts to flow toward the STB. When the STB receives all the necessary information required to decode the underlying MPEG stream, it begins to display the video for the new channel on the TV display.

Many factors contribute to the CCT behavior:

- IGMP signaling delay – The IGMP leave and join messages sent by the STB either flow directly to the aggregation router or are processed by an intervening IGMP snooping device

such as a DSLAM and should support “fast” leave to minimize IGMP signaling delay. Irregardless, IGMP signaling is a very minor component of channel change delay.

- MPEG decoding delay – After the STB starts to receive the new multicast stream containing MPEG Transport Stream (TS) packets, it needs to acquire program-specific information (PSI) frames in order to determine the desired TV channel. The time it takes for the decoder to receive the program-allocation-table (PAT) frame determines the time it takes for it to display programs. MPEG specifications indicate a decoder must receive a PAT every 0.5 seconds.
- I-frame acquisition delay – To reduce the amount of bandwidth required for digital video transmission, compression methods such as MPEG formulate the video frames of a digital video stream into various types of standardized MPEG frames. These frame types are known as I-frames, P-frames, and B-frames, named for the type of MPEG frame prediction used in performing their encoding. The I-frames do not rely on previous or future frames for encoding and thus uses intra-frame coding. A group of pictures (GOP) is a collection of these MPEG frames (typically 12 to 15 frames). The I-frame delay is proportional to the size of the GOP. In typical digital broadcast, this delay is approximately 0.5 seconds. IPTV systems tend to have larger I-frame delay due to the higher compression to reduce bit rates, which increases the size of the GOP. I-frame delay is the most prominent factor in channel change delay.
- Conditional-access-system (CAS) key acquisition delay – In a typical CAS, the encryption of digital services can be achieved by using entitlement control messages (ECMs) and EMMs. A STB must receive and decrypt the correct ECMs and EMMs in order to generate the final keys needed to decrypt a particular video stream. ECMs carry control words containing keys required to decrypt an MPEG stream. The STB needs to acquire the relevant ECMs before it can decrypt the MPEG stream and display the video. This CAS key acquisition delay can also add to the channel change delay.

In order to expedite the CCT, the Cisco VQE offers *rapid channel change* (RCC), which effectively reduces or eliminates the main sources of channel change delay:

- Multicast leave latency – leaving the current TV channel
- The time it takes for the STB to receive information it needs to begin demultiplexing, decoding, decrypting, and displaying the video stream, including:
 - PAT, Program Map Table (PMT), and ECM (if decryption is required)
 - Program Clock Reference (PCR) and sequence header information (for example, frame rate, etc.)
 - First I-frame, filling the jitter buffer
- Multicast join latency – to select the new TV channel

The RCC solution employs standard Real-Time Transport Protocol (RTP) and RTCP protocols to perform the appropriate signaling between the STB and the video aggregation network that will optimize the effective CCT. In addition, the RCC capabilities employ the same standards-based technology (RTCP) as the error recovery mechanisms described previously. As a result, service providers benefit from innovative technology for both aspects of video quality experience, reducing their training requirements. This solution is highly scalable in terms of the number of concurrent subscribers supported, enabling service providers to control their capital and operational expenditures associated with their video service delivery.

Advanced Video Quality Monitoring and Reporting

While service providers continue to strive for increased service revenue associated with video services, they must also constrain their ongoing operational expenditures associated with the acquisition of new subscribers. In addition to reducing the amount of time spent on getting new subscribers, service providers must ensure that the customers are retained. One of the aspects of customer retention is to continually deliver high-quality video services. However, service providers must also have advanced video monitoring and reporting tools to proactively determine the QoE being viewed by their subscribers.

The Cisco VQE technology suite offers monitoring capabilities that employ standards-based RTCP capabilities in the STB. These capabilities are linked with the video delivery network application intelligence to provide per-subscriber QoE information. These capabilities enable the service provider to proactively monitor the viewing experience of each subscriber and take ameliorative action without requiring onsite technicians to diagnose the problems. As a result, service providers can reduce their overall operations costs by reducing call center calls and eliminating complete equipment upgrades while increasing customer satisfaction.

Conclusion

Service providers must address numerous large system-level concerns to ensure a high Quality of Experience (QoE) for subscribers receiving IPTV over service provider network infrastructure.

These concerns include:

- Minimizing packet loss across all network failure scenarios to reduce the cumulative visual effect on all subscribers – The preferred solution is to use Layer 3 video distribution with enhanced PIM SSM and IGP for consistent sub-second Layer 3 convergence in all network outage and failure scenarios. Reduction of the visual effects associated with the video service delivery is critical to customer retention.
- Avoiding network congestion associated with concurrency rate of VoD through a robust Cisco Integrated Video Admission Control solution – The optimal solution is a system that delivers both off- and on-path admission control. On-path video admission control is essential to remain aware of topology changes in dynamic networks and meet the viewing expectations of the consumers without degrading service to all subscribers.
- Providing a solution for packet loss associated with bit errors on the DSL access line to increase addressable market and deliver quality video – An emerging visual VQE technology promises to take advantage of a linkage between the STB and an intelligent network to provide standards-based video packet retransmission and forward error concealment.
- Improving the expected CCT associated with employing IPTV services by eliminating the main sources of channel change delay with a highly scalable, standards-based RTCP signaling solution
- Providing advanced video QoE monitoring capabilities on a per-subscriber basis to reduce equipment upgrades associated with video access network quality, network congestion, and outages

Cisco has developed numerous infrastructure advances and has provided service intelligence innovations with the Cisco VQE technology – and these advances and innovations help to overcome many of the challenges of IPTV service delivery. Cisco VQE is integrated with the video delivery network, taking advantage of standard protocols (RTP), balanced with an understanding of

QoS, IP Multicast, and the high-availability requirements necessary to deliver the entertainment experiences demanded by consumers. The Cisco IP NGN Carrier Ethernet Design provides the optimal delivery of video with the intelligence and agility to support these new VQE innovations. This commitment to the delivery of entertainment-grade video is essential to meet the demands of the empowered consumer and enable the success of providers delivering interactive multimedia applications over their converged Carrier Ethernet infrastructure.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IQ Expertise, the IQ logo, IQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0807R)