

Preserving Video Quality in IPTV Networks

Javed Asghar, Ian Hood and Francois Le Faucheur

Abstract – Growing numbers of service providers are implementing video services over IP networks, and discovering the unique challenges of providing a high quality of experience (QoE) to subscribers. Delivering consistent QoE in packet-switched networks can be a complex proposition due to the high sensitivity of video traffic to packet loss, as well as to delay and jitter. Preserving video quality in IP Television (IPTV) networks that mostly rely on copper access lines poses an even greater challenge. Service providers need intelligent mechanisms in core and distribution networks to prevent congestion that deteriorates video quality, as well as intelligence between the aggregation networks and IP set-top box (STB) to repair packet losses, speed up channel change times and monitor the quality of the subscriber experience. This paper discusses a series of core and aggregation-layer approaches, collectively referred to as Visual Quality of Experience (VQE) to address these issues. The VQE approach encompasses on-path video connection admission control (CAC) employing Resource ReSerVation Protocol (RSVP) in the core and aggregation layers, and a real-time signaling mechanism operating between the provider aggregation edge and the IP STBs to address packet losses, long channel change times and quality monitoring.

Index Terms – IPTV, RSVP, call admission control, DSL line errors, channel change time, IPTV quality monitoring.

Introduction

Consumer consumption of network bandwidth continues to grow exponentially. The Internet Innovation Alliance in Washington, D.C., suggests that, by 2010, the average U.S. household will use 1.1 terabytes of bandwidth each month – meaning that 20 homes will generate more traffic than the entire Internet did in 1995 [1]. The most significant factor driving this bandwidth growth is the increased transport of video over IP networks.

Growing numbers of service providers are turning to video services as a means to expand their revenues and market share. As more service providers roll out video offerings, the ability to meet and exceed customer expectations for video quality will become a critical service differentiator. Delivering high-quality video, however, is a complex proposition; particularly for service providers relying on packet-switched networks. This is due to the uniquely resource-intensive and packet loss-sensitive nature of the video traffic.

Broadcast IPTV video as well as video-on-demand (VoD) services demand high-quality real-time connections that can carry considerably high bandwidth. An IPTV stream requires approximately 2-4 Mbps for standard-definition (SD) video using the MPEG4 compression standard, and 6-12 Mbps for MPEG4 high-definition (HD) video. To deliver a high-quality user experience, the European Telecommunications Standards Institute's (ETSI) Digital Video Broadcast (DVB) standard [2] recommends a maximum of one video artifact per two-hour movie, which translates to a packet loss rate in the order of 10^{-7} . IP video distribution, however, is volatile, error-prone, and subject to a wide range of distortions, artifacts, and degradations during acquisition, compression, processing, reproduction and transmission.

IP video streams are also highly sensitive to packet loss. Even a very minimal packet loss in an IPTV stream can result in a significant degradation in video quality. A single dropped I-frame in an IP video stream causes the user's screen to literally freeze for a few seconds, until the next I-frame arrives and the screen can be refreshed. Many IPTV networks rely on digital subscriber line (DSL) access links, which employ twisted copper in the last mile to the customer home. Copper access networks are highly susceptible to interference from many sources, and thus, are subject to significant packet loss.

While a variety of issues can affect video quality in IPTV networks, we identify four issues as the key causes of unsatisfactory subscriber QoE: network congestion, bit errors on access lines, slow channel change times and limited ability to monitor per-subscriber video quality. We discuss a set of standards-based core and aggregation-layer techniques to address each of these issues, collectively referred to as Visual Quality of Experience (VQE). The VQE approach encompasses two key components to address the quality issues associated with both VoD and multicast video services:

On-path VoD connection admission control (CAC) in the network core and aggregation layers to prevent network congestion

A standards-based signaling mechanism operating between the provider edge (PE) aggregation router and the IP set-top box (STB) to perform loss repair, accelerate channel changes and provide quality monitoring and QoE reporting

Together, these VQE techniques can provide an end-to-end quality control and monitoring capability for IPTV service providers that extends from the core network to the aggregation edge, and all the way to the IP STBs. Ultimately, these capabilities allow service providers to:

- Improve video quality, increase customer satisfaction and strengthen the service provider's brand
- Reduce network congestion and guarantee a consistent QoE to all subscribers
- Make efficient use of all available bandwidth in the network
- Reduce access line signal-to-noise ratio and bit-error rate requirements for IPTV services expanding the addressable market for their services
- More accurately and proactively isolate quality issues, reducing costly help-desk calls and truck rolls

Video Connection Admission Control

Background

One of the most significant challenges in delivering high-quality IP video is ensuring adequate bandwidth in the aggregation network to support the traffic load without congestion. Reserving bandwidth for broadcast (multicast) IPTV traffic is a relatively simple task, as bandwidth requirements are determined solely by the number of channels offered, independent of the number of subscribers viewing those channels. Supporting VoD sessions, however, in which each unicast stream requires dedicated bandwidth, is a more complex problem.

Unlike conventional Internet/web data traffic, unicast VoD streams are inelastic and cannot readily adapt to network congestion, whether due to oversubscription of a link or to an outage. As discussed, VoD streams are also large (requiring as much as 12 Mbps for HD video) and intolerant to packet loss and jitter. Since all subscribers on an aggregation link typically share a class-based

Quality-of-Service (QoS) queue, the moment the aggregate VoD load exceeds capacity – even by a single session – video quality degrades for all subscribers on that link. This problem is compounded in resilient aggregation networks in that, in the event of a failure (whether due to a trunk or ring failure, a node failure, or even a scheduled upgrade), all sessions converge on and oversubscribe the remaining reduced-capacity path.

It is simply not cost-effective for a service provider to engineer a congestion-free residential metro or aggregation network that can support concurrent VoD connectivity to, for example, 25,000 customers connected to an aggregation router. At the same time, given the delay-sensitive nature of video traffic, service providers cannot oversubscribe the aggregation network as would be possible when supporting other types of traffic, such as high-speed Internet services.

Service providers must therefore engineer the network based on expected peak VoD load. The highly dynamic nature of subscriber VoD usage, however, makes such predictions difficult. Even using sound estimates based on past usage, situations are likely to arise (whether due to unexpected usage spikes or to a network failure that reduces capacity) in which more subscribers are requesting VoD sessions than current network capacity can support. As network capacity approaches oversubscription, some mechanism is needed to prevent new users from pulling more VoD streams until the capacity is available to support them, and to provide a graceful denial to subscribers, comparable to a busy signal on a telephone network.

The most effective solution is to employ some signaling intelligence in the network to perform per-flow admission control and deny the initiation of new VoD sessions until the necessary network bandwidth can be re-captured (i.e., until a current VoD user goes offline), and then allocate those resources dynamically to new users. Such a mechanism also should ensure that in the event of a network failure that causes re-routed sessions to oversubscribe the remaining path, some subset of those established sessions are torn down to preserve the quality of the remaining sessions.

Ideally, this admission control mechanism should meet six key requirements:

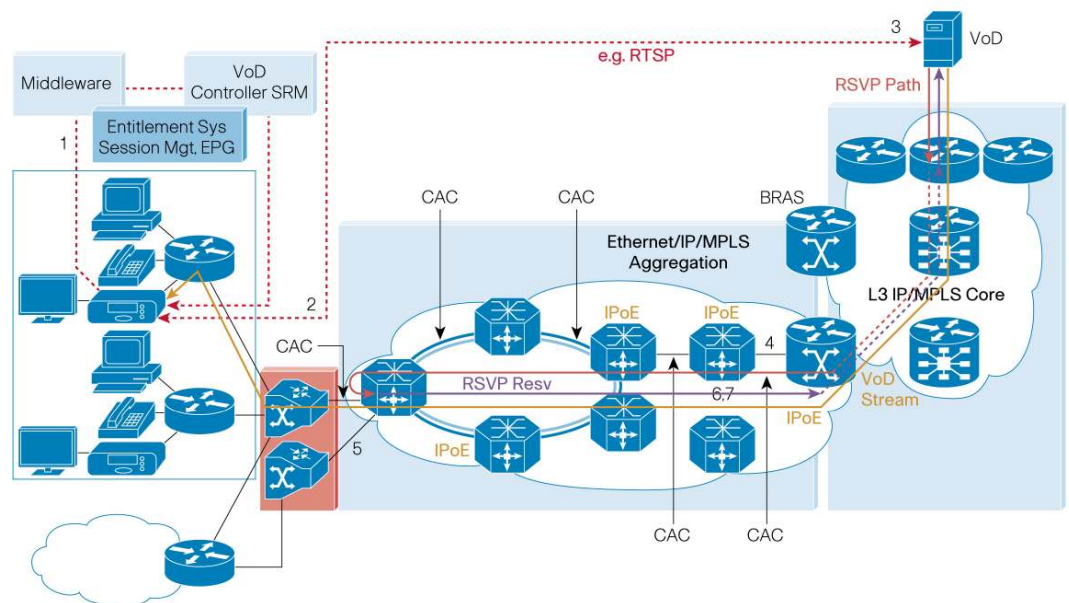
1. It should be accurate, allowing distribution network elements to communicate directly with the VoD server and make CAC decisions based on real-time bandwidth availability on the actual path that will support the VoD session.
2. It should be efficient, allowing service providers to allocate all available bandwidth to VoD sessions without requiring unused capacity to be held in reserve.
3. It should employ intelligent resiliency mechanisms that adapt to failures gracefully, re-establishing admission-controlled VoD sessions on new paths without tearing down large numbers of sessions unnecessarily.
4. It should be topology-independent, allowing it to interoperate with complex network topologies that have redundant and load-sharing paths in the transport layer of the network.
5. It should be network-based and STB-independent, interoperating with deployed STBs without requiring them to run any CAC-specific protocol.
6. It should be standards-based, relying on standardized Layer-3 forwarding mechanisms that do not require extensive integration with system middleware and can interoperate with any standards-based equipment deployed in the access layer of the network.

We propose an on-path VoD CAC approach based on the Internet Engineering Task Force (IETF) Resource ReSerVation Protocol (RSVP) [3] to perform essential video admission control functions and meet these requirements.

On-Path RSVP CAC

The proposed CAC solution employs standardized RSVP for signaling, sent by the VoD server (or a component on its behalf) before the beginning of the video session. At a high level, this approach operates as follows: The subscriber issues a command to the STB to select an on-demand video stream. This is translated into a request towards the VoD middleware that authorizes it and then returns to the STB information about the VoD server to contact for streaming. The STB then issues a streaming request to the VoD server. In turn, the VoD server issues an RSVP message (Path message), which follows the exact path that the VoD session will use, reflecting any real-time changes in the aggregation network. Along the path, each video distribution element performs a network bandwidth accounting function. The routers allow the VoD session if sufficient bandwidth is available to support the video stream, and deny the request if it is not. When a stream is accepted, the router sends an RSVP message (Resv message) to the VoD Server, which in turn starts video streaming. When a stream is denied, the router sends an RSVP error message back to the VoD server, which in turn sends the subscriber a graceful denial message (e.g., a busy signal) in concert with the video middleware application client-server software. Fig. 1 illustrates the on-path approach in a typical service provider network.

Figure 1. On-Path RSVP CAC Signaling for VoD Session



This model employs an on-path CAC approach. “On-path” refers to the fact that the admission control request and confirmation are explicitly signaled in the network (via RSVP) along the actual path that the new VoD session would follow. RSVP-based on-path CAC can be seen as a fully distributed admission control method, in that admission control decisions are made by the network elements themselves, without any external mechanisms required. The on-path RSVP CAC approach contrasts with “off-path” admission control methods that concentrate CAC intelligence in a centralized server.

On-path RSVP CAC offers the following advantages:

- *Accuracy* – The bandwidth reservation for each session occurs along the exact Interior Gateway Protocol (IGP) path the VoD session will use
- *100 percent bandwidth efficiency* – If a link has the capacity for 1000 VoD sessions, for example, RSVP will always accept the 1000th session and reject the 1001st.
- *Resiliency* – In the event of a topology change (due to network failures, router upgrades, addition of links, etc.) reservations automatically re-establish along the newly re-routed IGP path of the VoD session, without having to re-signal CAC or rebuild the entire connection. RSVP on-path VoD CAC co-exists with Multiprotocol Label Switching (MPLS) Fast Re-Route (FRR) link and node protection to restore connectivity (including CAC reservations) within 50 milliseconds (ms) of link or node failure. In networks that do not run FRR, RSVP CAC provides native Fast Local Repair features to dynamically re-establish reservations after a failure.
- *Flexibility* – RSVP CAC allows service providers to extend CAC end-to-end from the video head-end (VHE) to the last-mile router, within any network topology, regardless of where video content is sourced.
- *STB interoperability* – RSVP CAC functions entirely within the core and aggregation layers. It requires no integration with STBs, and can be implemented even in environments with older, previously deployed STBs.
- *Minimal processing requirements* – RSVP CAC has minimal impact on router central processing units (CPUs), making it suitable for very large deployments. For example, test on a commercially available router reveal processing time of an RSVP message of a few milliseconds and over 100,000 CACed sessions consuming less than 10% CPU in steady state.
- *Ease of implementation* – RSVP CAC functions are localized to the network and require only that network elements support a subset of the fully standardized RSVP protocol.

RSVP CAC Operations

The following section details the steps involved in a typical RSVP CAC operation for both a successful and unsuccessful VoD session request:

1. The STB relays the customer request for a new VoD session via an application-layer Real Time Streaming Protocol (RTSP) [4] setup command to the VoD server. (Depending on the network, other steps may occur before this, such as the STB communicating with the VoD middleware or session manager to obtain the address and other parameters of the VoD server.)
2. Before accepting the VoD session request at the application layer, the VoD server requests the corresponding bandwidth reservation in the network by generating RSVP Signaling (e.g., an RSVP Path message) which travels downstream toward the STB. The Path message traffic specification (TSPEC) contains the bandwidth required for the VoD session (e.g., 4 Mbps).
3. At each hop, IGP routes the Path message downstream toward the requesting STB. RSVP-enabled routers at each hop process the path message and install a corresponding Path soft state.
4. When the Path message reaches the aggregation PE router, the router activates the RSVP Receiver Proxy function on its DSL access multiplexer (DSLAM)-facing interface to terminate the Path message and generate a corresponding RSVP reservation request (Resv) message.

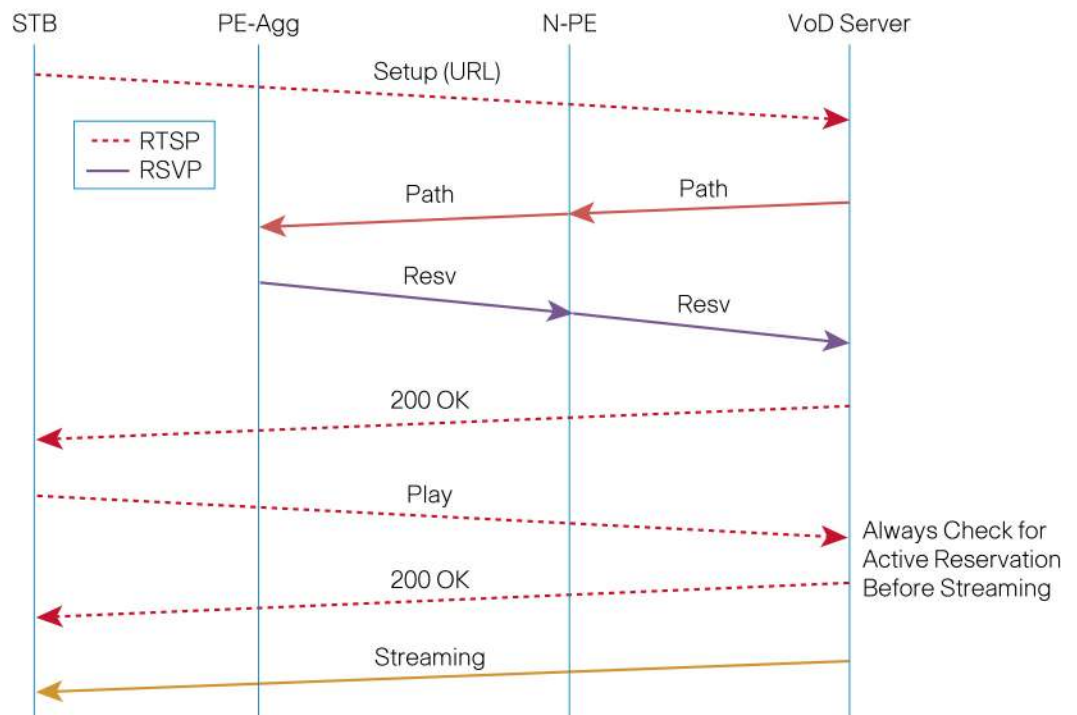
(Note: at this point, the RSVP path message is not forwarded to the DSLAM or requesting STB. It is terminated on the egress interface of the aggregation PE router.) The Resv message contains the VoD session bandwidth encoded in the Resv flow specification (FLOWSPEC), which is the same value as in the Path TSPEC (e.g., 4 Mbps).

5. The RSVP Receiver Proxy function hands the Resv message over to the regular RSVP module on the PE aggregation router. The RSVP module then performs processing of the Resv message exactly as if it had been received from the downstream STB, and performs admission control on the egress, or DSLAM-facing interface. If this admission control decision is successful, RSVP creates a corresponding Resv soft state and forwards the Resv message upstream to the previous RSVP hop.
6. As per normal RSVP operations, the Resv message gets routed back upstream toward the VoD server along the exact same path as was followed by the Path message. (To ensure this, each RSVP router explicitly addresses the Resv message to the RSVP Previous hop, based on the previously established Path state.)
7. At each RSVP hop, the router performs admission control on the downstream link from that hop and makes an admission control decision based on the available bandwidth on that egress interface. The router allocates the requested bandwidth (e.g., 4 Mbps) from the global RSVP pool of the interface and installs the Resv soft state.

If the admission control decision is successful at every hop, the following occurs:

1. The VoD server receives the Resv message confirming the successful reservation of the bandwidth (e.g., 4 Mbps) for the VoD session across the aggregation network.
2. The VoD server then confirms to the requesting STB, at the application layer, that the new VoD session is accepted (sending a RTSP status code 200, i.e., "OK," to confirm RTSP setup).
3. The STB responds with a RTSP "Play" command to the VoD server.
4. The VoD server recognizes that the corresponding RSVP reservation is already in place and responds again with a RTSP status code 200 indicating the "Play" command was executed successfully.
5. The VoD server begins unicasting the corresponding media stream to the subscriber.

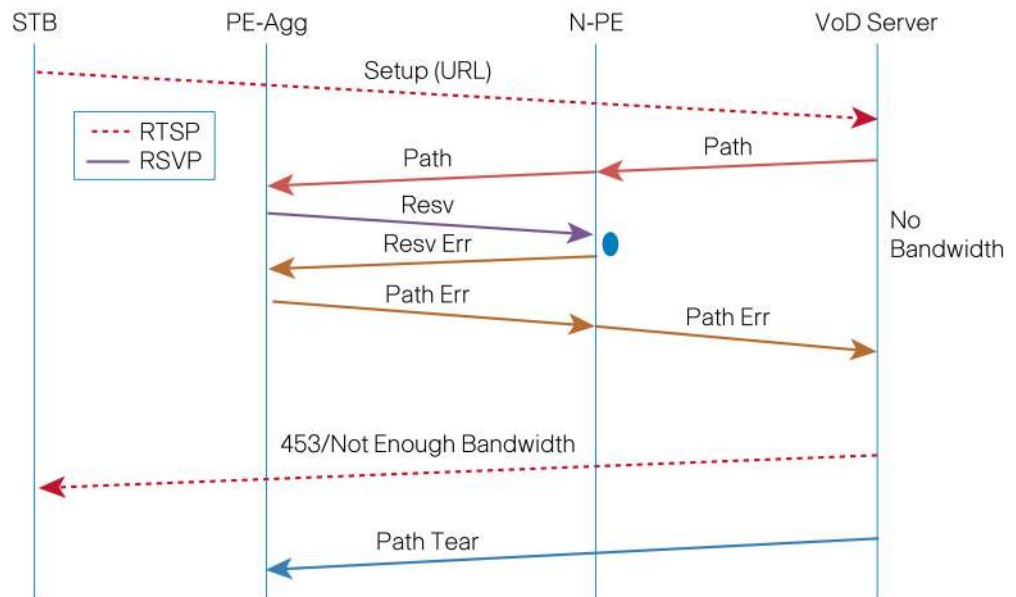
Fig. 2 illustrates a successful admission control decision.

Figure 2. Successful On-path RSVP CAC Signaling

Alternatively, if a router at any hop upstream from the PE aggregation router performs admission control on its downstream link and finds insufficient remaining bandwidth for the new session, the admission control for the new VoD session fails, and the router does not install this reservation. In this scenario, the following occurs:

1. The router sends a Resv Error message back downstream toward the PE aggregation router indicating that admission control for the reservation has failed. (All other existing reservations are not affected by this failure and remain in place.)
2. The RSVP Receiver Proxy function in the PE aggregation router processes the Resv Error message, and sends a Path Error message containing the error code "Admission Control Failed" back upstream to the VoD server to provide explicit notification of the CAC failure. The regular RSVP operations are entirely receiver-driven, so the initial RSVP standards only allowed the use of this error code in a downstream message (Resv) and not in an upstream message (Path Error). The Transport Area (TSV) Working Group of the IETF has standardized an extension to RSVP to support this functionality in the presence of an RSVP receiver Proxy [5].
3. On receipt of the Path Error message, the VoD server recognizes that admission control failed, and sends a RTSP status code 453 (i.e., "Not Enough Bandwidth") to the STB.
4. The STB recognizes that the VoD request was denied because the network is busy, and the VoD middleware displays a meaningful, user-friendly message to the subscriber, such as "Sorry, the network is currently busy. Please try your request again later."
5. The VoD server injects a Path Tear message (addressed to the STB, just as the original Path message was) to immediately clean up all the RSVP soft states in the network for that failed reservation and free up any bandwidth which may have been reserved by the RSVP Receiver Proxy function on the downstream interface of the PE aggregation router.

Fig. 3 illustrates an unsuccessful admission control decision.

Figure 3. Unsuccessful On-path RSVP CAC Signaling

Through this process, RSVP VoD CAC provides a highly efficient mechanism for extending admission control intelligence from the VHE to the last-mile router, and preventing network congestion to preserve subscriber video quality.

Responding to Network Failures

In the event of a network failure that causes re-routed VoD sessions to oversubscribe the remaining path, RSVP CAC provides features to dynamically tear down a subset of established sessions to preserve the quality of the remaining sessions. In addition to meeting the sub-50-ms convergence times that service providers require, RSVP provides the intelligence to tear down only those reservations that cannot fit on the new path and preserve those that can. RSVP dynamically re-establishes reservations around the failed link without requiring all VoD sessions to be torn down and rebuilt end-to-end.

RSVP CAC co-exists with MPLS Traffic Engineering (TE) FRR, allowing new paths to be established dynamically after a failure within 50 ms, with no impact to CAC. For networks that do not run TE FRR (or as a backup resiliency feature in networks that do), RSVP also provides a Fast Local Repair feature, which is enabled by default once RSVP is enabled.

To accomplish this intelligent resiliency, RSVP re-subjects every reservation affected by a route change during reconvergence to admission control. RSVP maintains sessions that can fit on the new path and tears down sessions that cannot, and in both cases, provides proper notification to the VoD server. This process occurs within a few seconds of reconvergence – meaning that users are subjected to a maximum of a few seconds of QoE degradation as a result of oversubscription after reconvergence.

RSVP Pre-Emption Features

A basic RSVP CAC implementation will effectively control VoD sessions to eliminate network congestion. Some service providers, however, may wish to exercise additional control over the priority in which requested VoD sessions are accepted or denied, or established VoD sessions are torn down in the event of a reconvergence. RSVP VoD CAC supports this pre-emption priority feature [6].

RSVP CAC pre-emption features allow each reservation to be associated with a preemption priority. In the event that not all reservations can be established through a given link, the RSVP admission control algorithm will admit reservations in accordance with their preemption priority. For example, if all the RSVP bandwidth on a link is already reserved by established reservations with a low preemption priority, and an establishment request arrives for a new reservation with a high preemption priority, then RSVP will bump one low priority reservation (or as many as needed) to make room for the new reservation. Similarly, if the RSVP bandwidth on an interface is reduced for some reason and all currently established reservations no longer fit, RSVP will bump as many reservations as needed, starting with those with the lowest pre-emption priority.

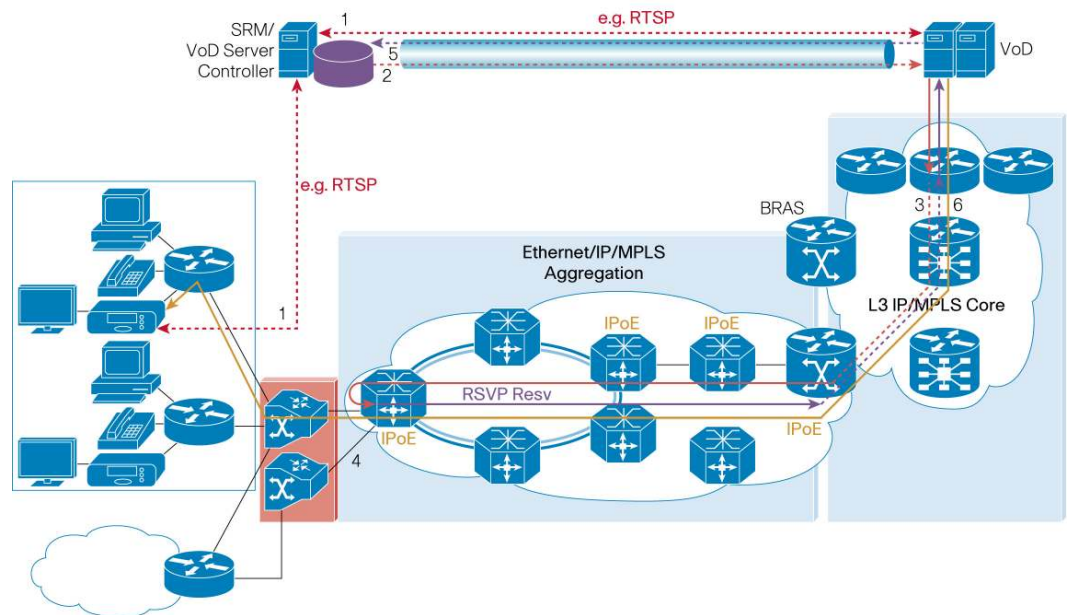
Service providers employing RSVP VoD CAC can use these preemption features to provide granular prioritization across different types of content and build a system that intelligently responds to resource shortages according to a variety of operational priorities. For example, service providers can associate a lower pre-emption priority with free VoD content and a higher priority with paid content. In normal circumstances, both free and paid content would be supported. In special circumstances, however, when bandwidth becomes scarce, free VoD sessions would always make room for revenue-generating paid sessions.

To use RSVP preemption, the VoD gear responsible for initiating RSVP signaling (either the VoD server or the session manager) need only include a POLICY_DATA object containing a Preemption Priority Policy Element in the Path messages. This POLICY_DATA object will be echoed back inside the Resv by the RSVP Receiver Proxy and then honored by every RSVP router processing the Resv message.

Implementing RSVP on VoD Equipment

Service providers have two options for implementing the RSVP protocol stack on VoD equipment to perform RSVP admission control. The first option is to deploy RSVP directly on the VoD server. In this scenario, RSVP messages are synchronized with VoD streaming. (The example scenario described above is based on this implementation.)

If deploying RSVP directly on the VoD server is not feasible (for example, if the VoD server cannot easily support RSVP), service providers can deploy RSVP on a centralized session resource manager (SRM). In this scenario, the SRM is remotely connected via a Generic Routing Encapsulation (GRE) tunnel to the first-hop router (FHR) that is directly connected to the VoD server. Fig. 4 illustrates the RSVP CAC process in a network with RSVP deployed on a centralized SRM.

Figure 4. RSVP on SRM

Since Path messages generated by the remote session manager must follow the same path through the network as their corresponding VoD sessions, the session manager injects Path messages into the network at the FHR by tunneling them into the GRE tunnel to the FHR. To the FHR then, it appears that the RSVP Path message has been generated by the VoD server, and RSVP VoD CAC proceeds as described in the example above.

Deploying On-Path RSVP VoD CAC

To support on-path RSVP CAC, IPTV networks must meet two key requirements. First, every network element, from the video distribution router connecting the VoD servers to the aggregation routers in central offices, must support native Layer-3 forwarding intelligence. Second, because RSVP is an IP-based protocol and follows the exact path of VoD streams, the network must transport VoD sessions natively over IP or over MPLS Label Distribution Protocol (LDP) in the Global Table.

IPTV services typically are delivered over “triple play” (voice, broadband Internet, and video) networks. As a result, service providers must account for the handling of both VoD traffic, which demands more rigorous treatment, and other types of Internet and business traffic which do not. Following are three RSVP CAC deployment models describing a native IP aggregation network, a hybrid network that transports video as native IP and MPLS/LDP for non-video traffic, and a network that transports all services (including video) over MPLS/LDP:

- *Pure native-IP aggregation network* – In this scenario, the network carries all traffic as native IP and applies CAC to VoD sessions.
- *VoD over IP + Internet/business traffic over MPLS* – In this scenario, the network also carries VoD traffic as native IP and applies CAC. Non-VoD traffic, however, is carried over MPLS/LDP either in Virtual Routing and Forwarding (VRF) tunnels or in the Global Table. MPLS/LDP selective advertisement ensures that prefixes used for VoD traffic are not label-switched.

- *VoD over MPLS* – In this scenario, the network uses MPLS/LDP for all traffic, including business/Internet traffic and VoD sessions. The network carries VoD traffic over MPLS/LDP in the Global Table and applies RSVP CAC. However, the network transmits the IP version 4 (IPv4) Path messages for the corresponding VoD flows as native IP packets without MPLS encapsulation. This allows for RSVP processing at every hop, even if the VoD media packets themselves are encapsulated in MPLS.

RSVP CAC Scalability

On-path RSVP-based CAC is highly scalable. While RSVP has had a reputation for scalability issues in the past, this perception is due to the previous generation of RSVP's model for QoS enforcement (RSVP IntServ). RSVP IntServ required maintenance of a separate scheduling state (e.g., a separate queue) in the data path. The current RSVP CAC approach is based on the "RSVP over Diffserv" model. In this model, the router maintains a soft state per reservation in the control plane, but relies purely on normal Diffserv mechanisms in the data path, allowing it to scale to hundreds of thousands of reservations per device [7].

With RSVP over Diffserv, RSVP does not perform any policing nor separate queuing (neither per flow nor aggregate) for admitted flows. It operates purely in the control plane and is solely responsible for performing admission control of flows over a configured bandwidth pool.

Extending RSVP CAC to Core Networks

Large service providers tend to distribute VoD content to multiple locations and often overprovision their IP/MPLS core networks, so RSVP VoD CAC is not typically required in the core network. RSVP CAC can, however, extend across the entire end-to-end network and deliver benefits even in core networks. For example, while oversubscription may not be a concern in the immediate future, service providers typically have little visibility into video traffic in core networks, and no efficient means of accurately monitoring VoD bandwidth utilization. Extending RSVP CAC to the core network can provide an intelligent mechanism for accomplishing this, while supporting VoD bandwidth control capabilities that may be employed in the future.

The chief requirement for providing RSVP CAC in core networks is core platform support for IPv4 RSVP CAC. If the core network meets this requirement, RSVP supports aggregation of RSVP reservations over MPLS TE tunnels, and extend CAC end-to-end across the provider network [8]. RSVP can also be supported when transported over an MPLS VPN network [10].

Packet Loss Repair, Rapid Channel Change and Quality Monitoring

Loss Repair

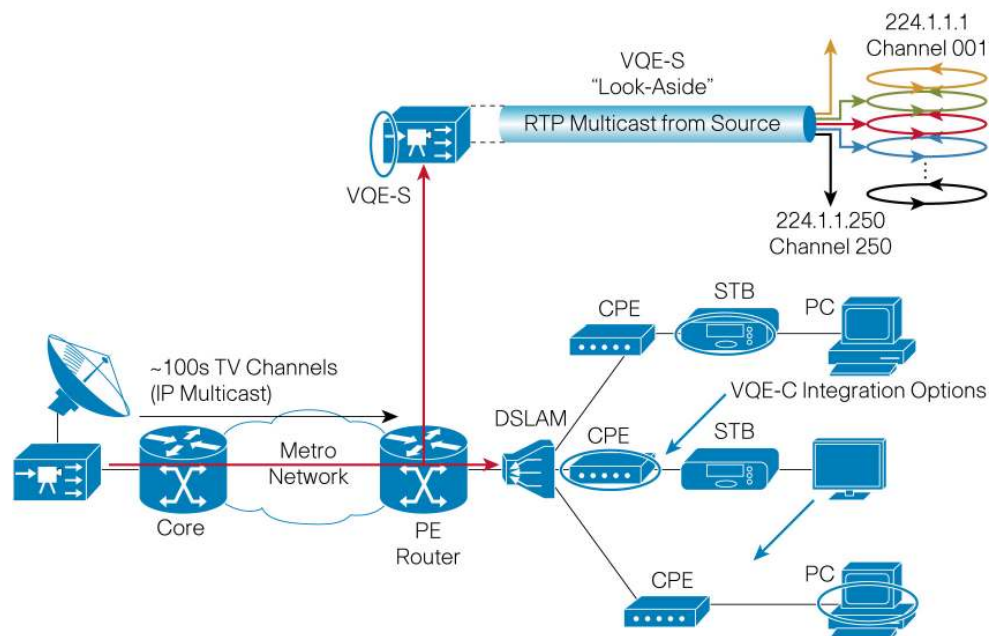
Service providers face challenges in delivering a high QoE for multicast TV, especially over the copper wiring that is common in the last mile of IPTV networks. The most significant concern for such service providers is the amount of errors that can occur on subscriber access lines. Typical DSL bit-error rates of 10^{-7} translate to packet loss rates in the order of 10^{-3} , which approximately produce an artifact in every few minutes. Audio/video packets need to be protected against such errors and any packet loss needs to be repaired before it is received by the audio/video decoders.

Incoming IPTV packets will be dropped if they fail cyclic redundancy checks (CRC) on the receiving DSL modem, routing gateway or STB. Bit errors can be caused by electrical impulse noise as packets traverse DSL transmission lines, or may result from poor building wiring, especially common in aging multi-tenant dwellings. Although DSL offers physical-layer error-

mitigation features such as Forward Error Correction (FEC) and interleaving, there may still remain unrepaired bit errors and they will cause packet drops at higher layers [11].

In light of these issues, service providers need an intelligent signaling mechanism operating between the STB and the PE aggregation network that can monitor for packet losses and repair them. As part of the VQE approach, we discuss a client-server approach that employs Real-time Transport Protocol (RTP) video encapsulation, application-layer FEC and selective retransmission to provide these services [9]. Drawing on intelligence localized in both the subscriber STB and the service provider's intelligent network edge, such a mechanism provides an ideal solution for dynamically recognizing and addressing access line errors. The server portion of the VQE solution is deployed at the edge of the network, i.e., the Video Switching Office (VSO) or Central Office (CO). The client software resides in the subscriber STBs. Fig. 5 illustrates a typical VQE deployment.

Figure 5. VQE Client-Server Model



A fundamental requirement in the VQE approach is that the media streams are encapsulated in RTP. This allows each IPTV packet within a given multicast group or channel to be assigned a unique RTP sequence number. This sequence number gives VQE the ability to identify the missing packets and place the recovered packets in the correct order upon their arrival.

To support VQE loss-repair functions, the VQE server is configured to receive all the broadcast channels that are currently being watched by the downstream STBs. The server joins the respective multicast groups using Internet Group Management Protocol (IGMP) "join" requests, like any other IP host. The VQE server maintains a dedicated circular buffer for each channel/multicast group and caches a few seconds of program content of each channel. This cached data is used to service loss-repair requests from the downstream VQE clients.

The loss-repair operation proceeds as follows:

1. The video source transmits multicast packets encapsulated in RTP to the subscriber STB.
2. The VQE client software in the subscriber STB tracks the sequence numbers of incoming RTP video packets. When it detects a missing sequence number (for example, if a corrupted packet has failed the CRC check and been discarded), the client first tries to repair the missing packet by using the application-layer FEC data (if available) [11]. If FEC fails, the client requests retransmission of that packet from its associated VQE server by sending a Real-Time Transport Control Protocol (RTCP) message [9]. The client may request a single packet or multiple packets in a single RTCP message.
3. Upon receipt of the request, the VQE server pulls up the requested packet(s) from its cache and transmits them to the requesting client.
4. The receiving client splices the retransmitted packet(s) into its de-jittering buffer according to the RTP sequence numbers.

Fig. 6 illustrates the VQE client detecting a missing packet while Fig. 7 represents this packet retransmission and splicing into the de-jitter buffer.

Figure 6. Packet Loss Repair – Loss Detection

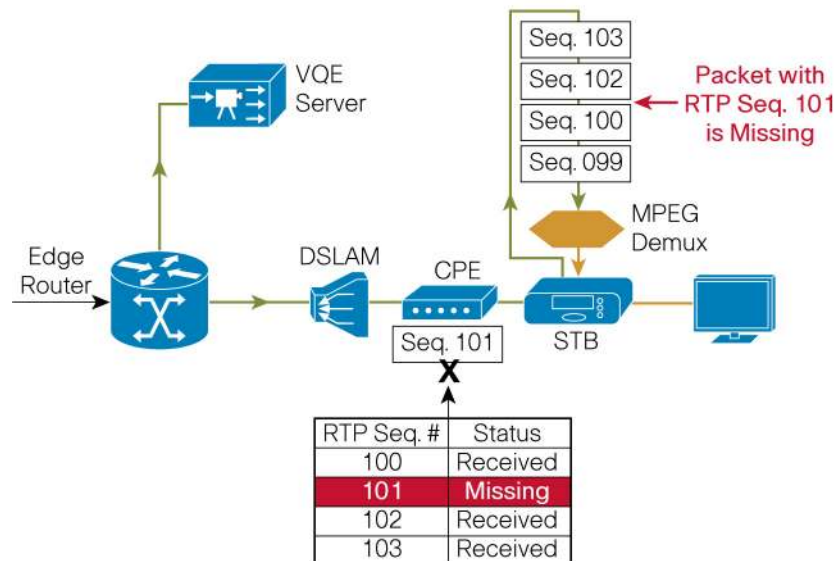
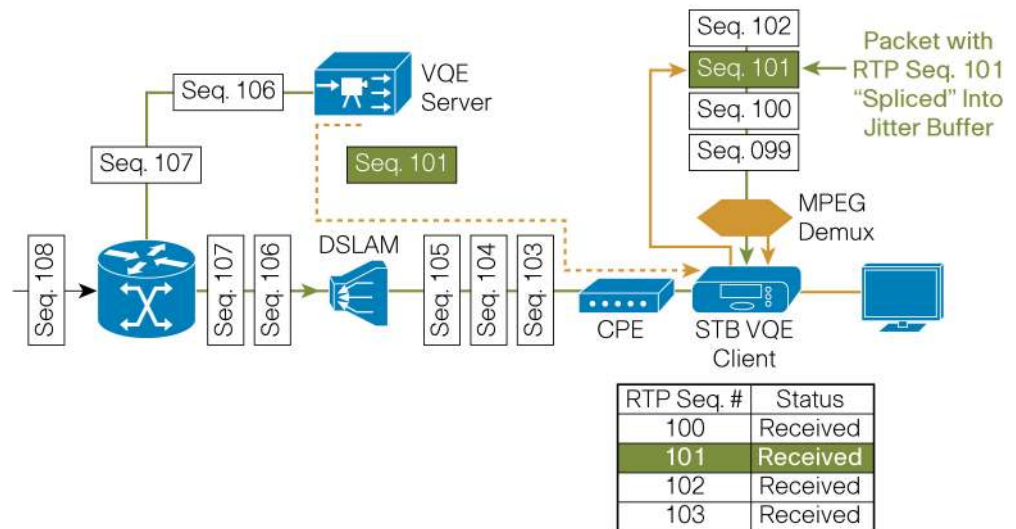


Figure 7. Packet Loss Repair – Retransmission and Splicing

The entire loss-repair process should complete in approximately 100 ms, although the recovery time depends on the round-trip delay between the VQE client and server. Since the operation happens prior to the MPEG TS de-multiplexing/decoding stage and since typical IP STBs employ 250-ms de-jittering buffers, the transaction should be transparent to the subscriber from an audio/video quality perspective.

Ultimately, this VQE function results in a highly reliable loss-repair capability, allowing service providers to address one of the chief causes of poor IPTV video quality and deliver a better QoE to more subscribers.

Rapid Channel Change

As service providers continually add channels and VoD content to compete for subscribers, users switch between channels more frequently than in the past. Subscribers expect to receive the same, virtually instantaneous channel change time (CCT) to which they have grown accustomed with analog TV broadcast, and they expect an immediate response to their viewing requests. CCT delays in IPTV systems, however, can be both longer and more variable than conventional systems – both of which issues can detract from the overall subscriber QoE.

These increased CCT delays can result from a number of factors, including [12]:

- IGMP signaling delays, i.e., multicast leave and join delays
- Transport stream random access point (TSRAP) acquisition delay, i.e., program specific information (PSI) including program association table (PAT) and program map table (PMT) acquisition delays, I-frame acquisition delay and conditional access system (CAS) key acquisition delay
- Loss-repair and de-jittering buffer delays in the STB
- MPEG buffering delay

Unlike the traditional broadcast model, in which the STB receives all offered channels concurrently, IP STBs typically must communicate with the upstream multicast routers for each channel change. The IGMP “leave” and “join” messages sent by the STB either flow directly to the PE aggregation router or are processed by an intervening IGMP snooping device such as a DSLAM. IPTV networks should support “fast leave” to minimize IGMP signaling delay. A common

misconception is that IP multicast signaling delays are the main contributors to increased or variable CCT delay. In fact, multicast “leave” and “join” signaling typically represent a relatively modest share of the total CCT. On the other hand, TSRAP acquisition and MPEG buffering delays are usually the largest contributors to the CCT [12].

The VQE client-server intelligence operating between the STB and the provider edge used to repair packet losses can be extended to circumvent I-frame acquisition delays and reduce CCT, and ultimately, to improve overall subscriber QoE.

As with loss repair, the VQE rapid channel change transactions employ RTP and RTCP to perform signaling between the STB and the video aggregation network, and optimize CCT. The key to accelerating CCT in this model is that the VQE server begins unicasting the media packets to the client STB at the same time as the network is processing the IGMP “leave” and “join” requests to begin normal multicast streaming of the new channel. The STB can begin processing the unicast packets immediately, and then synchronize the display with the multicast stream once it becomes available. This capability greatly reduces the time a subscriber waits before the image is rendered on the TV screen. Early results produced from a testing platform show that channel changes of less than a second are achievable with the VQE channel change approach [12].

The rapid channel change operates as follows:

1. When the subscriber changes channels, the STB-based VQE client first issues and IGMP leave from old channel and requests IPTV packets for the new channel from its target VQE server, again using RTCP messages.
2. Upon receipt of the rapid CCT request, the VQE server locates the appropriate channel cache, identifies the location of IPTV packets carrying a recent I-frame for that channel, and transmits a short unicast burst of packets, starting with the I-frame, to the requesting client.
3. Because the incoming burst from the VQE server contains an I-frame, the STB decoder can immediately begin processing the MPEG information upon receipt of the unicast burst. In parallel, the STB VQE client will use an IGMP “join” for the new multicast stream
4. After a short time, multicast packets begin arriving at the STB. The STB VQE client monitors the RTP sequence numbers from both the unicast and multicast streams. After a short overlap, the VQE Server will stop the unicast stream that is no longer useful.

To minimize the risks of induced congestion in the access and relevant parts of the aggregation network, the VQE server performing the CCT transaction shapes the unicast stream such that the combination of the unicast and corresponding multicast stream do not exceed a pre-defined rate. Also, the CCT unicast streams need to be taken into account by the RSVP VoD admission control solution discussed in section II. This may be achieved by factoring the CCT bursting behavior when allocating video bandwidth to IPTV and thereby determining the remaining bandwidth that can be used by RSVP CAC for VoD.

QoE Monitoring and Reporting

It is not enough to deploy mechanisms for delivering higher video quality if a service provider has no means of measuring subscriber QoE. This concern is more than academic; accurate video quality information is critical to effective traffic modeling and traffic engineering. Implementing mechanisms to monitor video quality, however, has traditionally been a complicated endeavor. Moving forward, to satisfy demanding customers and differentiate their offerings based on video quality, service providers will require straightforward and easy-to-obtain information about per-

subscriber video flows. As in the loss repair and rapid CCT scenarios, service providers can benefit from RTCP to address this requirement.

Since each STB-based VQE client supports RTP, each client supports RTP's rich packet-level statistics-gathering capabilities. Statistics include cumulative information on loss, jitter and delay of the RTP streams. The VQE mechanism draws on those statistics for standards-based monitoring and reporting. In operation, the VQE client transmits receiver and extended reports [9] to a target VQE server (not to the RTP source). These reports are sent periodically. The VQE server in turn sends the summarized receiver reports via a TCP interface to a network analysis tool where they can be processed in detail.

With these continuous RTP signaling capabilities, service providers can proactively monitor per-subscriber video quality from the video headends to the STB, and generate accurate reports of per-subscriber QoE information. This allows service providers to isolate problems, such as narrowing down the source of a quality issue to the DSLAM, the external wiring plant or the in-house wiring. Ultimately, these capabilities allow service providers to respond to quality issues proactively, and take ameliorative action for many issues without requiring an onsite diagnosis.

Conclusion

As more service providers incorporate video into their service portfolios and IP communication networks, they will face a continually growing challenge to preserve video quality and deliver excellent subscriber QoE. The VQE approach described in this paper provides a standards-based, highly efficient set of techniques to inject granular QoE control mechanisms into IPTV networks. With these techniques, service providers can extend intelligent quality controls from the core network through the aggregation layer, and from the provider edge to the subscriber home to deliver consistently superior QoE.

References

- [1] Internet Innovation Alliance, "Bringing Up Broadband: Higher Traffic; Higher Costs," *USA Today*, April 27, 2008.
- [2] ETSI EN 300 429 V1.2.1 (1998-04), "Digital Video Broadcasting (DVB); Framing structure, channel coding and modulation for cable systems."
- [3] IETF RFC 2205, "Resource ReSerVation Protocol (RSVP) Version 1 Functional Specification". [Online]. Available: <http://www.ietf.org/rfc/rfc2205.txt>
- [4] IETF RFC 2326, "Real Time Streaming Protocol" [Online]. Available: <http://www.ietf.org/rfc/rfc2326.txt>
- [5] IETF draft-ietf-tsvwg-rsvp-proxy-approaches, "RSVP Proxy Approaches" draft-ietf-tsvwg-rsvp-proxy-protocol Extensions for RSVP Receiver Proxy" [Online]. Available: <https://datatracker.ietf.org/drafts/draft-ietf-tsvwg-rsvp-proxy-approaches/> and <https://datatracker.ietf.org/drafts/draft-ietf-tsvwg-rsvp-proxy-protocol/>
- [6] IETF RFC 3181, "Signaled Preemption Priority Policy Element." [Online]. Available: <http://www.ietf.org/rfc/rfc3181.txt>
- [7] IETF RFC2998, "Integrated Services Operation Over Diffserv Networks." [Online]. Available: <http://www.ietf.org/rfc/rfc2998.txt>
- [8] RFC4804, "Aggregation of RSVP Reservations Over MPLS TE/DS-TE Tunnels." [Online]. Available: <http://www.ietf.org/rfc/rfc4804.txt>
- [9] IETF RFC 3550, "RTP: A Transport Protocol for Real-Time Applications." [Online]. Available: <http://www.ietf.org/rfc/rfc3550.txt>
- [10] IETF draft-ietf-tsvwg-rsvp-l3vpn "Support of RSVP in Layer 3 VPNs" [Online]. Available: <https://datatracker.ietf.org/drafts/draft-ietf-tsvwg-rsvp-l3vpn/>
- [11] A. C. Begen, "Error control for IPTV over xDSL networks," in IEEE Consumer Communications and Networking Conf. (CCNC), 2008
- [12] A. C. Begen, N. Glazebrook, and W. V. Steeg, "A unified approach for repairing packet loss and accelerating channel changes in multicast IPTV," in IEEE Consumer Communications and Networking Conf. (CCNC), 2009

Ian Hood, Professional Engineer and Cisco Senior Product Marketing Manager, is responsible for driving the global market expansion of Cisco Carrier Ethernet and IPTV/Quad Play solutions based on the Cisco IP Next-Generation Network (NGN) architecture. He is an expert on a broad portfolio of intelligent Carrier Ethernet service delivery platforms and solutions designed to meet service provider needs. With more than 20 years of experience in networking, Ian has held prominent roles leading the innovative and successful IPTV and Carrier Ethernet business operations for Motorola and Nortel Networks. Ian hails from Canada as a licensed Professional Engineer in Ontario, and holds a B.A.Sc. in Electrical Engineering from the University of Waterloo. He is a lifetime honored member of the International Executive Guild.

Javed Asghar, Cisco Technical Marketing Architect, is responsible for driving Next-Generation IPTV/Quad-Play solutions and Carrier Ethernet architecture in the Edge Routing Business Unit (ERBU). He is a technical expert in advanced Video/IPTV technologies and Carrier Ethernet service delivery platforms and solutions designed to meet service and cable provider needs. He has actively participated in design and deployment of many IPTV and L2VPN networks in production today. Javed has also held prominent roles leading hardware architecture, software development, and test projects in the development of IPTV and Carrier Ethernet technologies within Cisco for six years before his current role. He is a frequent speaker on video and L2VPN topics at various industry conferences and a lifetime member of IEEE.

Francois Le Faucheur is a Cisco Distinguished Engineer. He started his career working for service providers in France (France Telecom) and Australia (Telecom Australia aka Telstra). He joined Cisco Systems in 1995, where he first provided technical expertise for Service Providers on ATM, IP and MPLS backbone evolution. In 1998, Francois moved to Cisco's engineering team as a technical leader in the area of QoS, MPLS as well as Voice and Video Admission Control in Next-Generation Networks. Francois is very active in the IETF in the areas of MPLS, Quality of Service, IPv6 over MPLS and RSVP. He has many patents in those areas and co-authored a book on MPLS Designs.



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, Cisco WebEx, 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, CCIP, 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, 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. (0809R)