CISCO SYSTEMS

White Paper

Integrated Video Admission Control for the Delivery of a Quality Video Experience

WHY ADMISSION CONTROL IS ESSENTIAL FOR DELIVERING IPTV

Delivering entertainment-grade video over IP, also often referred to as IPTV, poses significant challenges as service providers scale solutions to manage millions of subscribers, withstand periods of peak demand, and deliver a superior quality of experience while balancing network capacity and efficient capital investment. In the Cisco® IP Next-Generation Network (IP NGN) architecture, Cisco Systems® offers Integrated Video Admission Control to deliver a high-quality user experience even when demand oversubscribes a network.

This paper details Cisco Integrated Video Admission Control solutions for video on-demand and IPTV in carrier-class, IP Next-generation networks offering triple-play services.

SUMMARY

Cisco works closely with service providers on Cisco IP NGN technologies to ensure that video reaches users with high quality, based on policies administered in the network's service layer. Given the bandwidth demands and quality requirements of video services, Cisco developed an admission-control solution that consists of two highly reliable components: video-on-demand (VoD) admission control and broadcast video admission control.

Performing admission control is necessary to preserve a high-quality video experience. The Cisco admission control solution takes into account complex network topologies that have redundant and load-sharing paths in the transport network as well as access link utilization policies and business policies that may be enforcing other constraints on subscriber services. Cisco admission control solution can differentiate between paid VoD, free VoD, specific high-demand broadcast channels, and many other variables, giving service providers enhanced choice and control over the IPTV experience. Cisco routers, in coordination with the Cisco Broadband Policy Manager (BPM) and on-demand servers and managers, perform the Integrated Video Admission Control function collectively.

CHALLENGE

Service providers offering high-speed Internet service typically define the service by the speed of upstream and downstream throughput, and thus the subscriber's service-level agreement (SLA) is defined directly by transport parameters. By contrast, with IPTV the service provider defines the service by the application experience based -- the number of channels provided, the quality of the video, the size of the video library, the user interface, the video recording capabilities, and the interactive features provided in the video broadcast service. The SLA for IPTV is not defined by network transport parameters but by the quality of experience (QoE). The network, in turn, must provide the appropriate quality of service (QoS) to support this experience-oriented SLA. Because of these distinct differences, Cisco has architected its IP NGN to enable optimum transport for each distinct service type. Each distinct service can be protected by independent mechanisms, and SLAs can be enforced differently.

High-speed Internet is treated as a *transport service*, and Internet traffic passes through a broadband remote access server (BRAS), which provides per-subscriber QoS and authentication enforcement. Per-subscriber QoS is critical so that the service can be accurately delivered and metered for each subscriber – for instance, a subscriber might select the turbo button to increase bandwidth for a period of time, and the QoS must be changed only for this subscriber. By contrast, video traffic is treated as an *application service*. Network QoS for application services is the same for all users, and QoS can therefore be delivered on a per-service basis (aggregate QoS). This simplifies

operations as the user is provisioned at the video application layer by the middleware, and the video service does not have to be provisioned again at the transport layer. However, IPTV services pose many other challenges to service providers.

IPTV is very intolerant to packet loss. Because IPTV is highly compressed, losing a video packet may result in the loss of valuable encoded information and a visible degradation of video quality, including macroblocking, pixelation, and even loss of a picture frame. Industry norms have defined the acceptable video quality of experience to be no more than one visible degradation per 2-hour movie. The corresponding network QoS for this results in an allowed packet loss rate of approximately 1 in 1 million (10^{-6}). For service providers delivering IPTV, this 10^{-6} maximum loss is considered a baseline requirement in the market.

Putting pressure on this required loss threshold is the continued growth of on demand content which delivers a dedicated stream to each viewer, consuming large amounts of bandwidth. On the business side, service providers looking for ways to differentiate their video offering can offer more broadcast content, more VoD content, and network-based personal video recorder features. However, service providers are challenged by the vast amount of content available from these on-demand and broadcast sources. As users get used to having the content they desire at their fingertips whenever they want it, they are not likely to accept only basic broadcast viewing. A recent report by Horowitz Associates found that the great majority of users with access to VoD used the service at least weekly, and approximately 20 percent used it daily (Figure 1).

Figure 1.   Video On-Demand Usage Among Cable Subscribers[1]
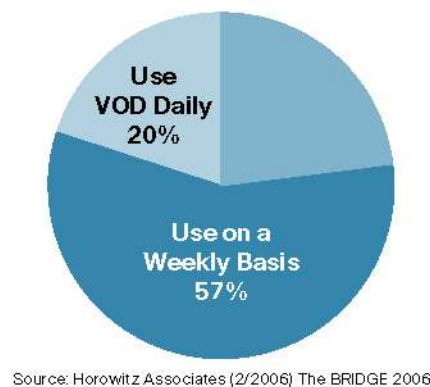
Source: Horowitz Associates (2/2006) The BRIDGE 2006

Additionally, according to Bernstein Research, VoD usage among Comcast subscribers grew by 400 percent between the summer of 2004 and 2005. In mature Comcast markets like Philadelphia, and among Time Warner Cable systems in New York, usage is approaching one selection per day, on average, per subscriber.

Providers typically set up their networks with a set amount of bandwidth for broadcast TV on each segment, including both standard and high-definition TV (HDTV). Broadcast TV is considered the foundation of any IP broadcast service and is protected at all costs. For example, a loss of video service during the Super Bowl in the United States would be a severe customer service issue. However, users are typically slightly more tolerant of VoD blocking, and therefore the network is designed to use the remaining bandwidth with VoD services after bandwidth is allocated for broadcast.

Given the success of VoD services and the likelihood of increasing high-definition VoD services, the bandwidth required to transport the video services will continue to increase and the possibility of video degradation will become more challenging. Because video cannot tolerate packet loss, congestion due to oversubscription is not acceptable – admitting just one more stream to a network near peak capacity could degrade the VOD and broadcast quality for *all* users. Exacerbating this risk, for the bandwidth-consuming on-demand services, the one metric that is the toughest to predict is the peak concurrency rate: the number of subscriber video-consuming devices (such as set-top

---

[1] Michael Hopkins, The Bridge, 'From Channel Surfing to Content Searching,' February 17, 2006, No. 6.1, Vol. 5

boxes) receiving an on-demand stream at the same time. With the addition of premium channel content available on demand and bundled inside a subscription service, peak concurrency rates experienced on Friday and Saturday nights have already climbed to as much as 20 percent of the capacity of all set-top boxes in many U.S. markets.

Using the example in Figure 2, the distribution network that is designed to serve on-demand content to 4000 subscribers out of one central office (CO) with an assumed 20 percent peak concurrency requires about 4 Gbps of capacity for the on-demand standard definition (SD) service alone.

Figure 2.   One Telco's Video Deployment Plans and Assumptions

Bandwidth from Video Hub Office to Central Office
VoD
• 4 million homes served by 400 COs = 10,000 average homes per CO
• IPTV service take rate = 40% = 4000 video subscribers per CO
• 2 TVs per subscriber x 20% VoD peak concurrency = 1600 streams per CO
• 94% SD VoD at 2 Mbps per stream + 6% HD VoD at 8 Mbps per stream
• **Approximately 4 Gbps per CO**
Broadcast
• 300 channels with 20% HDTV at 8 Mbps for each and 80% SD at 2 Mbps for each
• **Approximately 1 Gbps per CO**

As video, and especially VoD, can quickly consume much more network bandwidth than high-speed Internet access and voice-over-IP (VoIP) services, it is easy to see why network design must be optimized to support it with the utmost efficiency.

In deploying video services that are efficient, dependable, and scalable, service providers can either provision their networks with the maximum bandwidth required to handle worst-case *estimated* peak concurrency, including in network failure scenarios, or deploy an admission control solution to manage peaks that occasionally exceed the bandwidth available. The first option would be prohibitively expensive and wasteful. The second option, which balances capital investment with delivering a high-quality user experience, has become a top priority among service providers as video joins data and voice as a network service in high demand.

In the earlier example, on a busy Saturday night peaks as high as 20 percent of all set-top boxes will be viewing one VoD stream. However, once in a great while certain events, such as the release of a new movie, will result in even more subscribers in front of their TVs. Also, the failure of a network link may sometimes drive peak demand higher than the network capacity. As end-to-end IPTV streams grow as expected, admission control is vital to avoid oversubscription with the ensuing congestion, packet loss, and possible degradation of the user experience.

Due to its high use of bandwidth, VoD is the most important service to control, but broadcast TV is also important, especially as more and more high-definition channels are added. To date, most service providers have not imposed limits on the number of channels or bandwidth used by broadcast services. In today's IPTV networks, there are limited admission control functions in place for VoD to protect against oversubscription. VoD streams that do get set up can experience unacceptably poor quality if some packets are dropped during congestion periods caused by oversubscription. As demand for more VoD and broadcast high-definition channels increases, the prospect of degraded video transmissions is troubling and the need for admission control for the video service is more pressing.

SOLUTION

The Service Exchange Framework (SEF) in the Cisco IP NGN architecture provides robust linkages between applications and the network layer. In general, this is where diverse services are controlled in a converged, integrated platform where multidimensional identity is

maintained, policies and sessions are managed, and monitoring and billing functions are administered. It is in the SEF layer where the Cisco Integrated Video Admission Control solution for VoD and broadcast video services is based to ensure the highest possible QoE.

Figure 3.   Cisco IP Next-Generation Network Layers
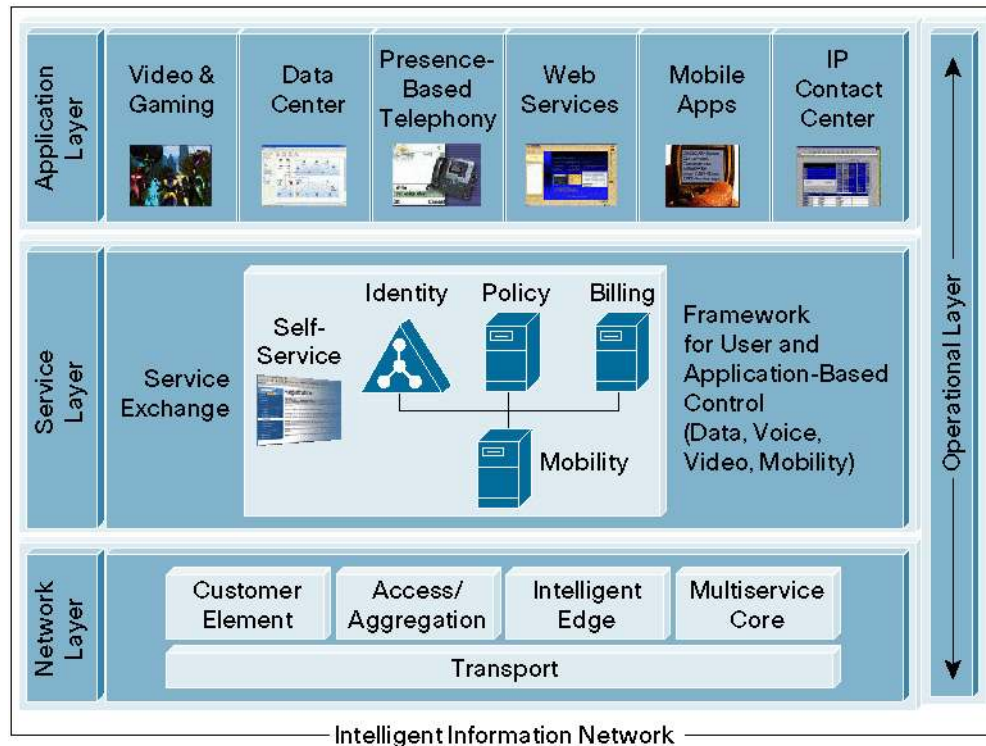
CISCO INTEGRATED VIDEO ADMISSION CONTROL SOLUTION FOR VoD

An admission control solution for VoD must interoperate with complex network topologies that have redundant and load-sharing paths in the transport layer of the network. It must also work with access links and business policies that may be enforcing other types of constraints on a subscriber's service. The Cisco admission control solution for VoD works in coordination with network routers, the Cisco BPM, and on-demand servers to collectively perform intelligent admission control for video.

The Cisco solution includes two important simultaneous processes that together are called Integrated Video Admission Control. The first is an on-path signaling procedure that performs admission control for the core and distribution layers (Figure 4).

Figure 4. Cisco VoD Integrated Video Admission Control with Both Off-Path and On-Path Functions
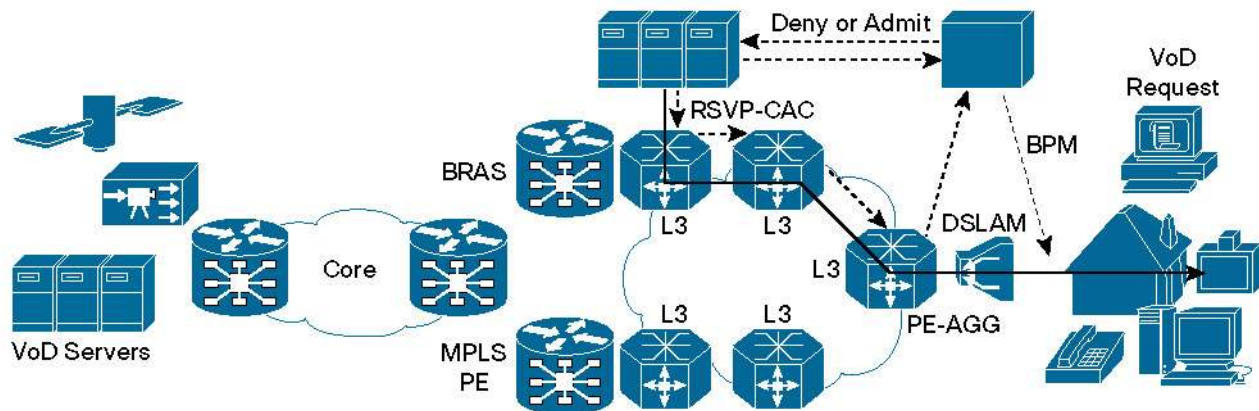
The on-path procedure uses Resource Reservation Protocol (RSVP) for signaling, sent by the VoD server (or a component on its behalf) before the beginning of the VoD session. In Figure 4, a set-top box selects a VoD movie. The request for a movie goes through the provider edge aggregation router (a.k.a "PE-AGG") and on to the video server. The RSVP message from the VoD server traverses the same path the VoD session will use, tracking in real time any changes in the complex network topologies in the core and distribution layers. Along the path, Cisco routers perform a bandwidth accounting function. They allow the session if bandwidth is available for that VoD stream and deny it if it is not. When a stream is denied, a RSVP-CAC message is sent back to the VoD server, which in turns sends the subscriber a message or a busy signal, generally through the video middleware.

It is critical that native Layer 3, for example, PIM SSM for multicast and IGP for unicast, be present on every network element, from the distribution router connecting to the VoD server complex down to the aggregation router in the central offices, to make this on-path admission control solution possible. The provider edge aggregation routers in Figure 4 are Layer 3 capable. Architectures using Layer 2 aggregation networks cannot support on-path admission control.

The second, complementary process of admission control is designed to prevent a video on-demand stream from being sent if the access link to the subscriber lacks the capacity to carry the stream. To provide a safeguard, the VoD server (or equivalent network component) sends a request to the Cisco BPM (policy server), which is keeping track of each access link to the subscriber. Access links are usually simple and semistatic, highly appropriate for an off-path approach (Figure 5).

Figure 5. Cisco Off-Path Signaling for VoD Admission Control
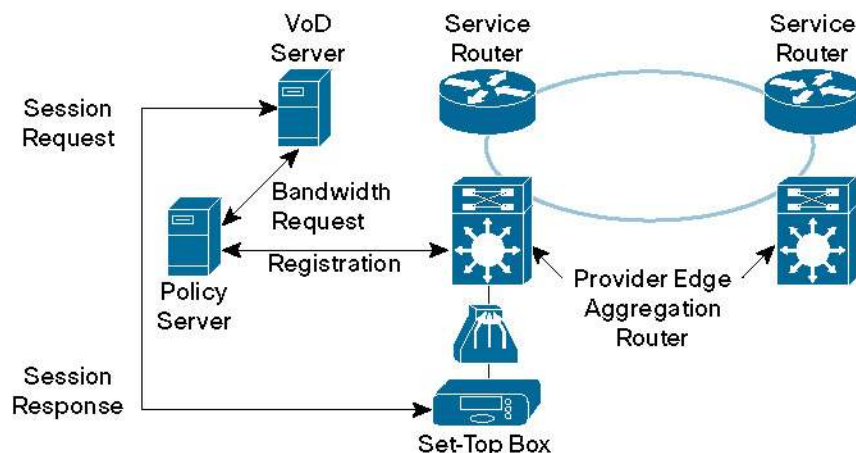
The policy server can check to see if the access link has enough unused bandwidth and can also check to see if business policies are in place (such as paid VoD versus free VoD) to allow the stream to be supported. The session is then either allowed or denied. This off-path portion of the solution is not recommended solely to perform admission control for the distribution and aggregation layers, because an off-path, policy server only-based approach to tracking any real-time changes in these complex topologies would be suboptimal at best. But the combination of an on-path admission control function and an off-path policy server query at the network edge has been shown to be the most reliable and efficient mechanism for VoD admission control.

The Cisco Integrated Video Admission Control solution is added to an architecture where aggregate QoS is used to ensure that video meets the required 10^{-6} maximum packet drop rates required for a high-quality video experience. Alternative approaches advocate per-subscriber shaping of video, which has two key problems:

- Video cannot be shaped because this may make some packets eligible for packet drop, which would make it difficult to meet the 10^{-6} target. Thus, the added complexity of hierarchical per-subscriber QoS provides no benefit in the context of video.
- For a video service, the subscriber is provisioned at the application layer over middleware. Having to provision the subscriber's bandwidth allocation for video again at the network layer is a second unnecessary operational step.

A robust Integrated Video Admission Control solution is required because either the network can support 10^{-6} packet loss or better and the stream should be admitted, or it cannot due to the potential congestion that would lead to packet loss, and the stream should be denied. Once video admission control is implemented, a simplified aggregate QoS queuing strategy that allows all video traffic to share the same queue avoids complex and costly hierarchical queuing strategies.

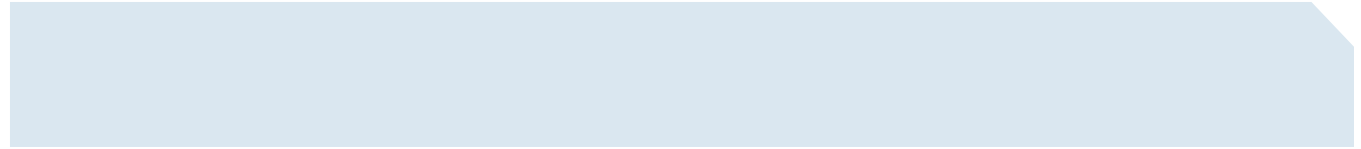### Cisco Integrated Video Admission Control for Broadcast Video

Cisco Integrated Video Admission Control for broadcast video benefits from the rich set of multicast forwarding and routing technologies developed by Cisco. Multicast replication and forwarding is performed in hardware at data rates exceeding tens to hundreds of millions of packets per second in Cisco switches and high-end routers.

Cisco Integrated Video Admission Control for broadcast video uses Internet Group Management Protocol (IGMP) limits to trigger a bandwidth availability check (Figure 6). Routers in the path are configured to maintain certain limits on broadcast bandwidth and can limit the number of channels simultaneously sent to the DSL access multiplexer (DSLAM).

Figure 6. Cisco Broadcast Admission Control
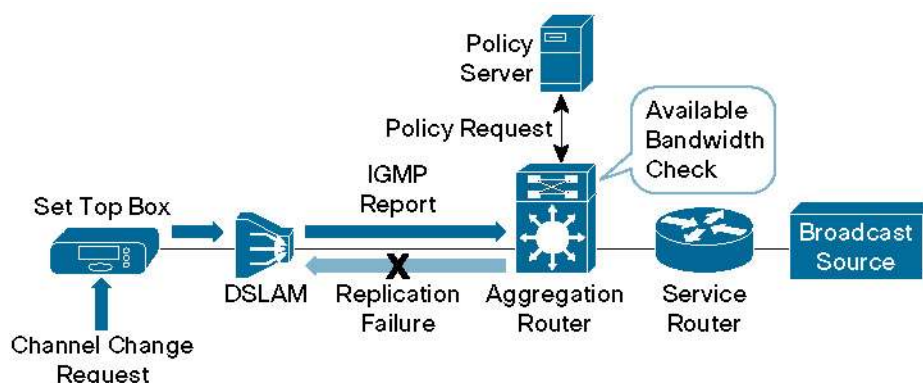
The multicast destination maps to the broadcast channel's bandwidth, and the bandwidth is subtracted from the link the request was received on. In the rare case that there is not enough bandwidth to support the 10^{-6} packet loss or better for the next broadcast channel, the router does not replicate a channel's multicast group out of a particular downstream interface.

Among aggregation network alternative approaches that operate at Layer 2, IGMP snooping is a feature used to determine which multicast packets get replicated to various access devices (such as DSLAMs). With this approach all multicast traffic is replicated throughout the aggregation network, and the granular control offered by native Layer 3 multicast is not available. In turn, this prevents the implementation of multicast admission control in a Layer 2 aggregation network.

CONCLUSION

For emerging IPTV providers, delivering a high-quality experience superior to the competition is paramount to achieving success and maintaining their customer base. Video is intolerant of packet loss, and industry norms dictate a 10^{-6} packet loss or better. Delivering that high level of quality while maintaining a sensible level of investment in network infrastructure requires a robust video admission control solution. Cisco Integrated Video Admission Control, combining an on-path admission control function based on RSVP with an off-path admission control function based on the Cisco BPM, promises to be the most reliable and efficient method yet. Only with native Layer 3, for example, PIM SSM for multicast and IGP for unicast, in the aggregation network can providers deploy this sort of robust Integrated Video Admission Control solution. Furthermore, Cisco Integrated Video Admission Control is combined with aggregate DiffServ QoS to deliver the required quality while minimizing operational complexity, so that all video traffic shares the same queue and avoids hierarchical queuing for video services.

Cisco Integrated Video Admission Control is a key function in the Service Exchange Framework of the Cisco IP NGN, realizing the high-quality experience needed for the video application in terms of the network QoS requirements. Cisco is the first to provide this linkage between applications and networks to help providers optimize their network investments while ensuring a high-quality user experience.

More than 10 million subscribers access IPTV services over Cisco IP NGNs. These networks are extending intelligence such as Layer 3 functionality into the aggregation network to deliver video with efficiency and premium quality. The transformation of networks from transport-oriented environments to service- and application-oriented environments will enhance the competitiveness of early adopters. With Cisco Integrated Video Admission Control, service providers around the world are deploying and benchmarking admission control solutions for evolving VoD, broadcast, and other applications to provide high-quality, dependable services through highly intelligent, adaptable infrastructures that can differentiate between and cater to individual subscribers.

FOR MORE INFORMATION

For more information about the Cisco Integrated Video Admission Control solution, visit:
http://www.cisco.com/en/US/netsol/ns610/networking_solutions_solution_category.html

CISCO SYSTEMS

llllnnllllnn.®

Printed in USA   C11-350531-01 06/06