IP/MPLS Networks: Optimize Video Transport for Service Providers

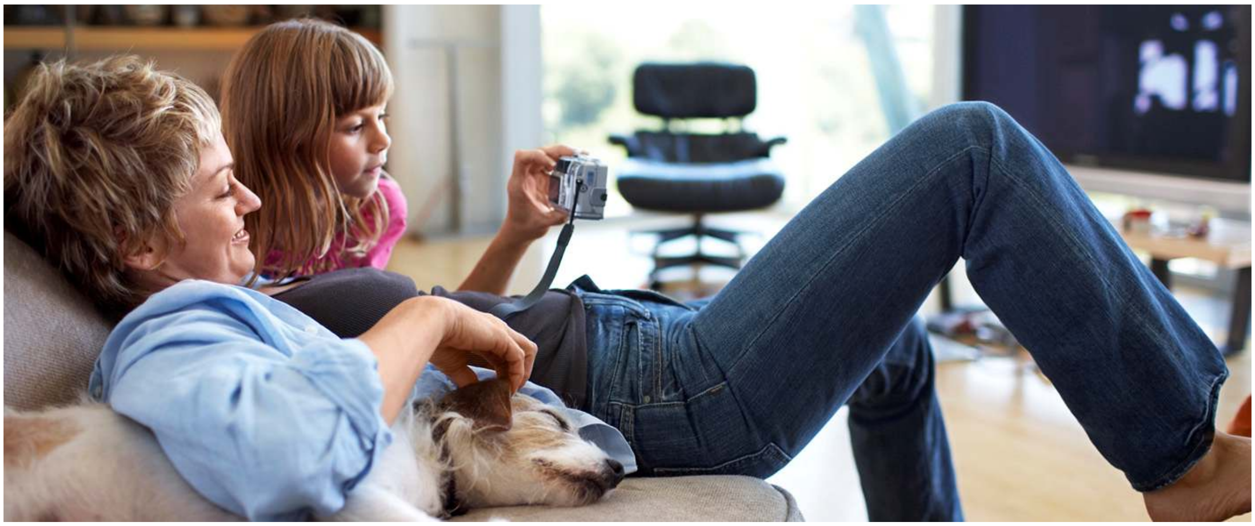

What You Will Learn

- Evolutionary trends of video transport over IP/MPLS networks
- Characteristics of packet video traffic in IP/MPLS networks and effects on user Quality of Experience
- Service Provider specific IP/MPLS video network design options
- Optimization techniques for transport of packet video in IP/MPLS networks

Introduction

IP is becoming the prevalent network technology for video transport, and video is becoming an increasingly significant component of IP network traffic. Network infrastructures that were designed for best-effort delivery, with relatively low bandwidths and high latencies, will not be able to meet the exacting requirements of next-generation video services. The requirements for supporting multiple services, including voice and video, on converged IP and Multiprotocol Label Switching (IP/MPLS) networks have promoted developments in quality of service (QoS), resiliency, availability, and scalability.

As a result, IP/MPLS networks can now deliver the service quality demanded by the highest-quality video services. IP is becoming the preferred converged network technology capable of meeting the requirements of premium video services, both in primary video distribution and video contribution environments, and in secondary video distribution networks such as Internet Protocol Television (IPTV) over residential broadband (for example, DSL and cable). IP/MPLS transport networks are now the primary means for the delivery of next-generation video services.

Cisco combines industry-leading experience in deploying carrier-class IP/MPLS networks with a deep understanding of video services gained from the acquisition of Scientific Atlanta. As a result, Cisco is exceptionally capable of meeting the requirements of service providers for transporting premium video services. The acceleration in video deployments increases the need to address the primary challenges facing broadcasters and service providers looking for IP-based video transport solutions, including how to deliver video service-level requirements and how to effectively manage and monitor the video transport service. Cisco has used its combined network and video application expertise to meet this challenge by developing network architectures and technologies for Video

Optimized Network transport, incorporated into the industry's leading portfolio of video-enabled products and solutions.

This portfolio includes the Cisco® CRS-1 and CRS-3 Carrier Routing System products, Cisco ASR 9000 Series Aggregation Services Routers, Cisco 7600 Series Routers, and Cisco Digital Content Manager (D9900) intelligent video multiplexers. These products provide the following features to cost-effectively support the most stringent service-level agreements (SLAs) and enhanced video monitoring capabilities demanded for the highest-quality video transport services:

- Fabric-based multicast replication, rather than line-card-based replication, giving outstanding multicast scale and performance for point-to-multipoint services
- Multipriority differentiated services schedulers, delivering the lowest possible delay and jitter in multiservice scenarios
- Point-to-multipoint MPLS traffic engineering for admission control and bandwidth reservation with fast reroute, providing 50-ms recovery times in both link and line-card failures
- Multicast only fast reroute (MoFRR), providing rapid recovery in case of node failure
- IP over dense wavelength-division multiplexing (IPoDWDM), allowing reductions in capital expenditures (CapEx) and operating expenses (OpEx), and the fastest possible detection of network failures, for even faster failure recovery
- Embedded video flow monitoring, providing pervasive router support for ubiquitous and cost-optimized detection and identification of network problems that reduce video quality

In combination with the router-embedded video flow monitoring capabilities, the Cisco Video Assurance Management Solution (VAMS) addresses video service management and monitoring, delivering an end-to-end service-aware view of the condition of the video transport network.

Also, the Cisco video adapter and encoder products, developed by the Cisco Service Provider Video Technology Group (SPVTG, previously Scientific Atlanta), allows the Cisco combined network and video solution to deliver lossless video transport for premium contribution and distribution services.

This suite of capabilities is the result of network-level and system-level development focused explicitly on the requirements of video transport and broadcasters, led by the Cisco Service Provider System Unit (SPSU), which has initiated a customer-inclusive systems-based development approach.

Developments in Video Transport

Video transport has become the largest source of network growth worldwide, and IP/MPLS is the prevalent network technology for video transport, so video is becoming an increasingly significant component of IP-based network traffic. The Cisco Visual Networking Index, which tracks IP traffic growth and the effect of visual network applications, predicts that by 2012 the Internet will be 75 times larger than it was in 2002, and that a video will account for almost 90 percent of all consumer traffic on the Internet.

Figure 1. Cisco Forecasts 44 Exabytes per Month of IP Traffic in 2012

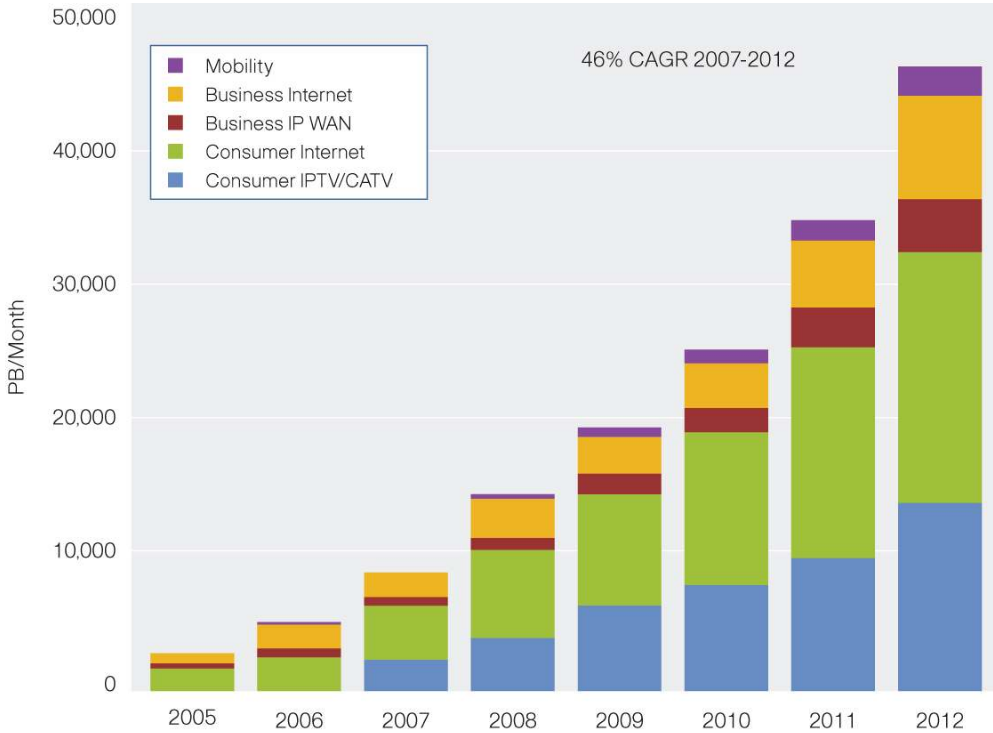

Major findings of the Cisco Visual Networking Index include the following.

- Global IP traffic will quadruple from 2009 to 2014. Overall, IP traffic will grow at a compound annual growth rate (CAGR) of 34 percent. Annual global IP traffic will exceed 3/4 of a zettabyte (767 exabytes) in four years. Global IP traffic grew 45 percent during 2009 to reach an annual run rate of 176 exabytes per year or 15 exabytes per month. In 2014, global IP traffic will reach 767 exabytes per year or 64 exabytes per month. The average monthly traffic in 2014 will be equivalent to 32 million people continuously streaming the movie Avatar in 3D for the entire month.
- Global Internet video traffic will surpass global peer-to-peer (P2P) traffic by the end of 2010. For the first time since 2000, P2P traffic will not be the largest Internet traffic type. The global online video community will surpass 1 billion users by the end of 2010. This number of people is exceeded only slightly by the populations of China (1.3 billion) and India (1.1 billion), making this user group equivalent to the third largest country in the world.
- P2P traffic is growing in volume, but declining as a percentage of overall IP traffic. P2P file-sharing networks now carry 3.5 exabytes per month and will continue to grow at a moderate pace with a CAGR of 16 percent from 2009 to 2014. Other means of file sharing, such as one-click file hosting, will grow rapidly at a CAGR of 47 percent and will reach 4 exabytes per month in 2014. Despite this growth, P2P as a percentage of consumer Internet traffic will drop to 17 percent by 2014, down from 39 percent at the end of 2009.
- Internet video is now over one-third of all consumer Internet traffic, and will approach 40 percent of consumer Internet traffic by the end of 2010, not including the video exchanged through P2P file sharing.
- The sum of all forms of video (TV, Video-on-Demand [VoD], Internet, and P2P) will continue to exceed 91 percent of global consumer traffic by 2014. Internet video alone will account for 57 percent of all consumer Internet traffic in 2014.
- Advanced Internet video (3D and high-definition video) will increase 23-fold between 2009 and 2014. By 2014, 3D and HD Internet video will comprise 46 percent of consumer Internet video traffic.

- The growth of video communications traffic is accelerating. Though still a small fraction of overall Internet traffic, video instant messaging and video calling are experiencing high growth. Video communications traffic will increase sevenfold from 2009 to 2014.
- Real-time video is growing in importance. By 2014, Internet TV will be over 8 percent of consumer Internet traffic, and ambient video will be an additional 5 percent of consumer Internet traffic. Live TV has gained substantial ground in the past few years. Globally, P2P TV traffic is now over 280 petabytes per month.
- VoD traffic will double every 2.5 years through 2014. Consumer IPTV and cable television (CATV) traffic will grow 33 percent between 2009 and 2014.

With the increased deployment of technologies such as IPTV, services that were previously delivered through traditional formats are now being delivered by IP/MPLS networks, which are also used for the delivery of conventional Internet services. Meeting the requirements of supporting multiple services, including voice and video, on converged IP/MPLS networks, and improvements in quality of service (QoS), resiliency, availability, and scalability mean that IP networks can now deliver the service quality required for the highest-quality video. Consequently, IP is becoming the preferred network convergence technology capable of meeting the requirements of both real-time and file-based video services, both in secondary distribution networks such as IPTV over residential broadband (for example, DSL and cable), and also increasingly for primary distribution and contribution networks. IP/MPLS transport networks are becoming vital to the delivery of next-generation video services.

The increasing importance of real-time video transport is a significant contributor to the evolution of service provider networks, prompting providers to reconsider their core network transport architecture and protocol choices, and to invest in new platforms and technologies. A crucial question for service providers during these upgrade cycles is how to most efficiently deliver the required video quality of experience (QoE) to the end user while controlling network costs and complexity to maintain profit margins. The answer to this question requires understanding network requirements for specific video services and how IP-based technologies behave in the delivery of these services.

This document describes the requirements, challenges, and the relevant IP technologies that are involved in the delivery of video services for secondary distribution service providers, and describes how Cisco is developing these technologies to address market needs.

Challenges for Video Transport

The creation of video content and its distribution is a multistage process that follows a lifecycle from acquisition at the source, through production and packaging of the content, to distribution to viewers. IP video transport service providers can be categorized by the types of services that they provide: contribution, primary distribution, or secondary distribution. In some cases, providers and networks may support more than one of the following services.

- Video contribution providers manage the capture and initial processing of video content and its transport prior to distribution. For example, the content may be transferred from one studio to another or from an event location such as a sporting event to a preproduction facility where the content is reworked before distribution. The video streams used in video contribution may be uncompressed at very high stream rates, with standard-definition (SD) stream rates of 270 Mbps and HD stream rates of 1.485 or 2.970 Gbps. Contribution services can be point-to-point (unicast) or point-to-multipoint (multicast).
- Primary video distribution providers manage the transport of video content from production environments to secondary distribution. Primary distribution services are normally compressed, ranging from MPEG-2, MPEG-4 to JPEG 2000 depending on the quality requirements. Networks supporting primary distribution services may share requirements from video contribution and secondary distribution providers.
- Secondary video distribution providers manage the transport of video content from primary video distribution providers to end consumers. IPTV, cable TV, and video over DOCSIS® are examples of secondary distribution services. Secondary distribution services are normally compressed using MPEG-2 or MPEG-44

standards, with rates of 2-4 Mbps for SD and 8-20 Mbps for HD. Secondary distribution services can be point-to-point (unicast-based VoD) or point-to-multipoint (multicast-based IPTV).

Figure 2 shows the relationships between the provider categories.

Figure 2. Video Service Provider Categories

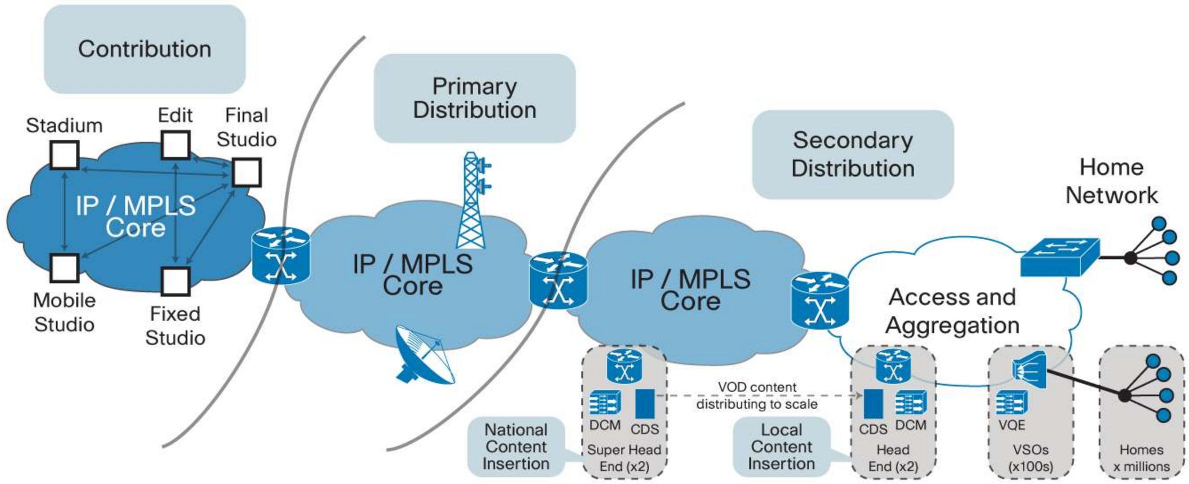

All service provider networks must combine simplicity, versatility, scalability, security, and manageability with cost effectiveness. Additionally, all categories of video transport services share the following challenges:

- What transport technology to use: native IP or IP/MPLS?
- How to meet the required SLAs for premium video services?
- How to manage and monitor the service, to help ensure that the IP network delivers the required SLAs for video, and to identify and isolate problem areas?

Each category of video transport service places its own requirements on IP/MPLS network deployments.

An understanding of the detailed network requirements of each video service offering and an understanding of how IP-based technologies behave in the delivery of these services are the first steps in designing a network optimized for their transport. The rest of this document focuses on network requirements and technologies optimized for real-time streamed secondary distribution video transport services, such as broadcast IPTV and cable TV.

Reference Architecture for Secondary Distribution

Figure 3 shows a reference architecture for secondary distribution video services.

Figure 3. Secondary Distribution Reference Architecture

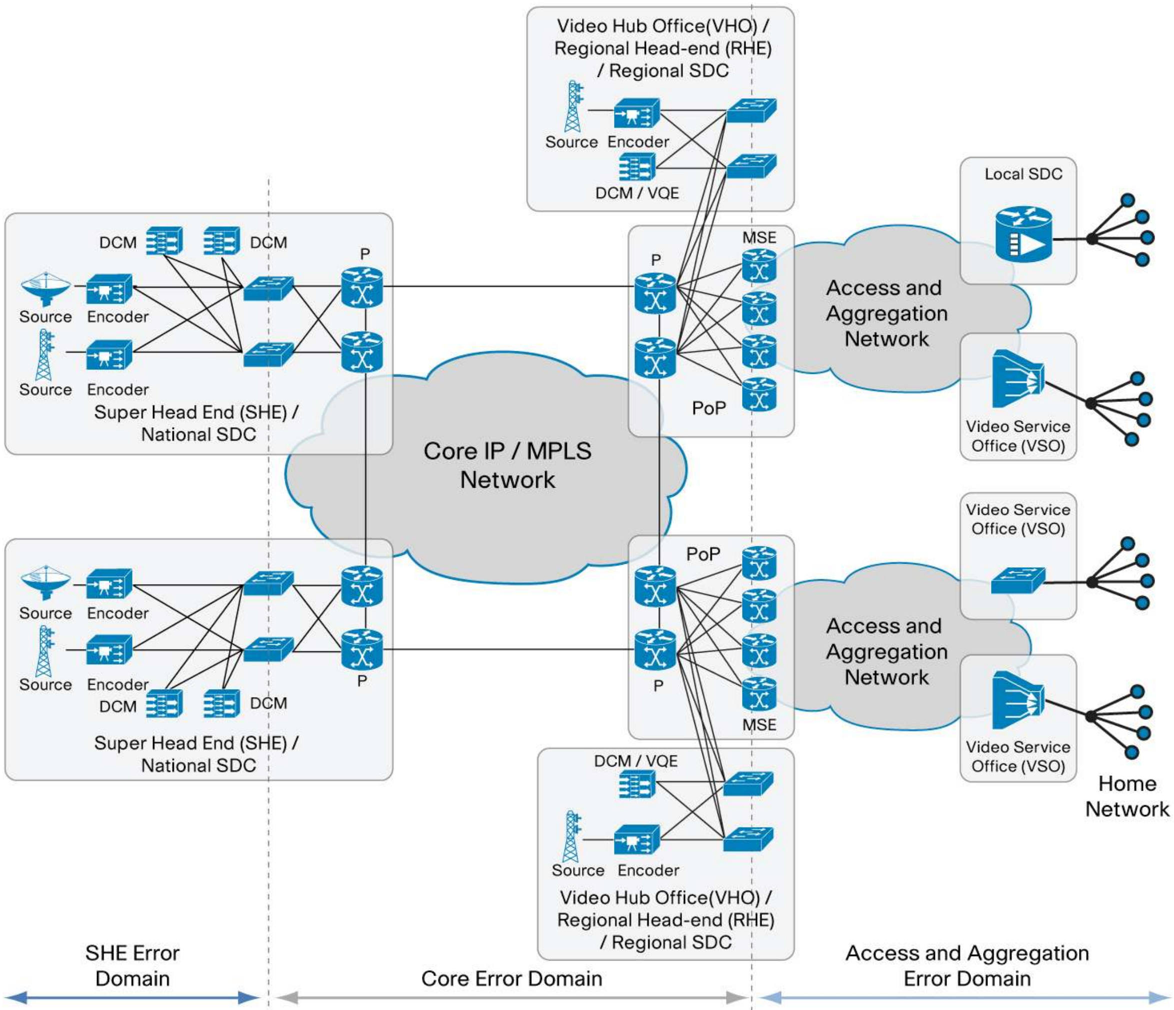

A secondary distribution video transport network may include the following functional locations:

- The super headend (SHE) is the location where common or national channels are acquired. A service outage at the SHE can affect all subscribers; therefore, there are typically two SHEs for resiliency. The two SHE locations can be used in active: active or active+standby configurations. The functions within the SHE include:
  - Encoders that encode the received video signal to MPEG: for example, Cisco D9032 MPEG-2 Encoder (MPEG-2 standard definition), Cisco HD Encoder Model D9050 (MPEG-2 high definition), Cisco Encoder Model D9034 (H.264 standard definition), and Cisco D9054 HDTV Encoder (H.264 high definition)
  - Video multiplexers, which are responsible for sending and recovering the encoded video streams, delivering them in the format required for the common channel lineup: for example, Cisco DCM Model D9900. The Cisco DCM in the SHE represents the demarcation point between the core system and the SHE as defined by the Cisco Video/IPTV over Broadband system.

  The cable TV equivalent to the SHE is the national Service Data Center (SDC).
- The regional headend (RHE) is the location where local channels are acquired and then added to the common channel lineup to create the channel lineup for a given region. The components within the RHE are typically a subset of those within the SHE. If the national channel lineup needs to be customized for each region, video-multiplexing equipment at the RHE can be used to renfigure the national channel lineup. A

video multiplexer (for example, Cisco DCM) can be used at the RHE to provide failover between redundant copies of the national channel lineup, while providing a single good copy (combined with local channels) to the region. The RHE is sometimes located with the video hub office (VHO). In smaller deployments or where there is no local channel acquisition, RHEs may not be necessary.

- Points of presence (POPs) are locations of the termination devices for IP services. Examples of IP termination devices include:
  - Broadband remote-access servers (for example, Cisco 10000 Series Routers or Cisco ASR 1000 Series for residential Internet services)
  - Multiservice edge and provider edge (for example, Cisco ASR 9000 Series routers, Cisco 7600 Series Routers, Cisco XR 12000 Series Routers, or Cisco CRS-1 and CRS-3 for business VPN services)

  Often the POPs also contain distribution and core nodes: for example in the, Cisco ASR 9000 Series, Cisco 7600 Series, Cisco XT 12000 Series. or Cisco CRS-1 and CRS-3. In many deployments, the VHO and POP, and sometimes the RHE, are located in the same facility.
- The video serving office (VSO) is the usual location of the aggregation nodes (and sometimes the access nodes). The VSO is typically the provider-owned facility closest to the subscriber, for example, the central office or local exchange. This may be where the Cisco Visual Quality Experience (VQE) Application is located. The equivalent to the VSO in cable is the local SDC.
- The video data center (VDC) is the location of most of the video application components and associated routing and switching infrastructure. VDCs are often located with the SHE, video headend (VHE), and RHE. Examples of these application components are the servers needed for back-office and middleware support.

From a transport perspective, the end-to-end service can be viewed as three separate error domains, which allows us to analyze the behavior of video streams and the effects of errors on these video streams in different parts of the network.

- The SHE error domain extends from the source to the video multiplexer (for example, DCM) located in the SHE, which is the demarcation point between the SHE and the core error domain.
- The core and edge error domain extends from the DCM located in the SHE to the VHO, which represents the demarcation point between the core and edge network and the access and aggregation network.
- The access and aggregation error domain includes everything on the subscriber side of the VHO.

Different strategies for SLA assurance and error recovery may be applied in each error domain.

Video SLA Requirements

Among secondary distribution video services, real-time streaming applications such as IPTV and cable TV have the most stringent SLA requirements. With video streaming applications, a receiver client requests a video that is stored on a server or produced in real time; the source server streams the video to the receiver, which starts to play the video before all of the video stream data has been received. Video streaming is used both for broadcasting video channels (commonly delivered as IP multicast) and for VoD, which is delivered as IP unicast. For distributed VoD architectures, where VoD caches are located at the edge of the network, the core IP/MPLS network is used for distribution (not in real time) of the content to the VoD caches. This distribution is essentially a file transfer, which has less stringent requirements on the core network than IPTV, for example, which is streamed in real time over the core and access and aggregation networks.

For secondary distribution, IP-based streaming video is commonly transported as a data stream encoded using the standards defined by MPEG and transported using the Real-time Transport Protocol (RTP, defined in RFC 3550), over the User Datagram Protocol (UDP, defined in RFC 768) over IP, or transported directly using UDP over IP. MPEG defines the encoding used for the actual video stream, while RFC 2250, RFC 2343, and RFC 3640 define how real-time audio and video data are formatted for RTP transport. RTP is the transport layer protocol, which manages the delivery of that stream from sender to receivers.

We can define the essential SLA requirements for an IP-based video transport service in terms of delay, jitter, and packet loss. Let's examine the network SLA requirements for a real-time streaming IPTV service.

Network Delay

One-way network delay characterizes the time difference between the receipt of an IP packet at a defined network ingress point and its transmission at a defined network egress point. Network delays include four components: propagation delay along the network path (approximately 5 ms per 1000 km for optical fiber), switching delay and queuing delay at network elements on the path, and serialization delay (the time it takes to transmit the bits of the packet sequentially onto a link). In addition, application response might be subject to network control protocol processing delays (such as multicast processing) and delays due to processing in application end-systems.

In the latest generation of IP/MPLS routers used in high-speed networks, switching delays are typically a few tens of microseconds and serialization delays are negligible. The Differentiated Services (DiffServ) IP QoS architecture (defined in RFC 2475) is used to control queuing delays and help ensure that service providers can meet their network delay SLAs. DiffServ is the baseline technology for IP QoS deployments today; it allows differentiated delay, jitter, and loss commitments to be supported on the same IP network for different types or classes of service. Traffic is classified at the edge of the network and then marked using the differentiated services code point (DSCP) in the IP packet header into a limited number of traffic aggregates or classes. Within the IP/MPLS network, scheduling and queuing control mechanisms are applied to the traffic classes based upon the DSCP marking; all traffic conditioning and dropping is handled intelligently at the network layer using IP DiffServ QoS mechanisms. DiffServ may be used in conjunction with the Integrated Services architecture (IntServ, defined in RFC 1633), which uses the Resource-Reservation Protocol (RSVP; defined in RFC 2205) for explicit admission control and bandwidth reservation.

In practice, with an effective DiffServ implementation and design and effective capacity planning, network delays will generally represent a relatively small proportion of the end-to-end delays. For example, in a national network within Europe, the end-to-end one-way network delays are typically less than 10 ms. Other contributors to the end-to-end delay include de-jitter buffers on the receiver and encoding and decoding delays.

Network Jitter

Network jitter is the variation in network delay caused by factors such as fluctuations in queuing and scheduling delays at network elements. We can generally consider jitter to be a variation of the one-way delay for two consecutive packets. De-jitter buffers are used to remove the delay variation the network causes by turning variable network delays into constant delays at the receiver. If a video de-jitter buffer is appropriately sized to accommodate the maximum value of network jitter possible, jitter will not delay playout beyond the worst-case end-to-end network delay.

Jitter can affect end-user interactivity by affecting the time it takes for the user to change from one TV channel to another (known as the channel-change time). For video-streaming applications, service providers typically target one-way network delays of less than 100 ms to achieve overall delays of 1 to 2 seconds. DiffServ IP QoS mechanisms are used to control network delays, and thus to set the maximum network jitter. In addition, advanced capabilities such as Cisco VQE technology, can enable rapid channel change for improved viewer QoE.

Packet Loss

Packet loss characterizes the packet drops that occur between a defined network ingress point and a defined network egress point. We consider a packet lost if it does not arrive at the specified egress point within a defined time period. Network packet loss has three primary causes.

- Congestion: When congestion occurs, queues build up and the network drops packets. DiffServ IP QoS mechanisms and capacity planning are used to help ensure that congestion does not occur, so no packets are dropped due to congestion. These mechanisms may be augmented with RSVP-based admission control.
- Lower-layer errors: Bit errors, which might occur due to noise or attenuation in the transmission medium, can result in dropped packets. Actual bit error rates vary depending on the underlying Layer 1 or Layer 2 technologies used, which are different for different parts of the network. In practice, fiber-based optical links might support bit-error rates as low as 1 bit in error out of 1013 transmitted bits (1e-13), whereas asymmetric DSL (ADSL) services might have bit-error rates as high as 1e-3. Some link-layer technologies employ reliability mechanisms, such as Forward Error Correction (FEC), to recover from commonly occurring bit-error cases and thus reduce the effective packet loss rate.
- Network element failures: Most networks are built to be resilient; however network element failures, such as link or node failures, can result in losses of network connectivity, which cause packets to be dropped until the network connectivity is restored around the failed network element. The resulting packet loss period depends on the capabilities of the underlying network technologies and implementations; these are described in the next section.

In summary, the DiffServ architecture is used in IP/MPLS networks to help ensure that the required delay, jitter, and loss SLAs can be achieved; this may be combined with RSVP-based admission control. With appropriate network engineering and capacity planning, the only network events that should result in a visual impairment to the video service are packet losses due to lower-layer errors or network element failures, which should be limited.

However, any single unrecovered video packet loss may result in a visual impairment. With compressed video, the effect of even limited packet loss can be significant. Depending upon factors such as the specific encoding and compression scheme, losses of different packet types appear as different types and durations of visual impairment or artefact. With MPEG encoding, a single packet loss causes a slice error at the minimum (Figure 4); multiple packet loss can result in blocking and pixelization or ghosting, or freezing and even the complete loss of the video image.

Figure 4. Artifacts Caused by Packet Loss: Slice Errors, Blocking and Pixelization, and Ghosting (Source material copyright SMPTE; used with permission)

The graph in Figure 5 plots the results of testing that demonstrates the relationships between network outage duration (contiguous packet loss) and the resulting visual impairment duration. This testing was undertaken for MPEG-2 encoded video traffic with high- and low-motion video clips, with both SD and HD. The bounded area in the graph shows the range of the resulting visual impairment, from worst case to best case for each test case and scenario.

Figure 5. MPEG-2: Relation of the Duration of Impairment to the Duration of Packet Loss in Standard-Definition (SD) and High-Definition (HD) Videos with High and Low Motion in Best- and Worst-Case Scenarios

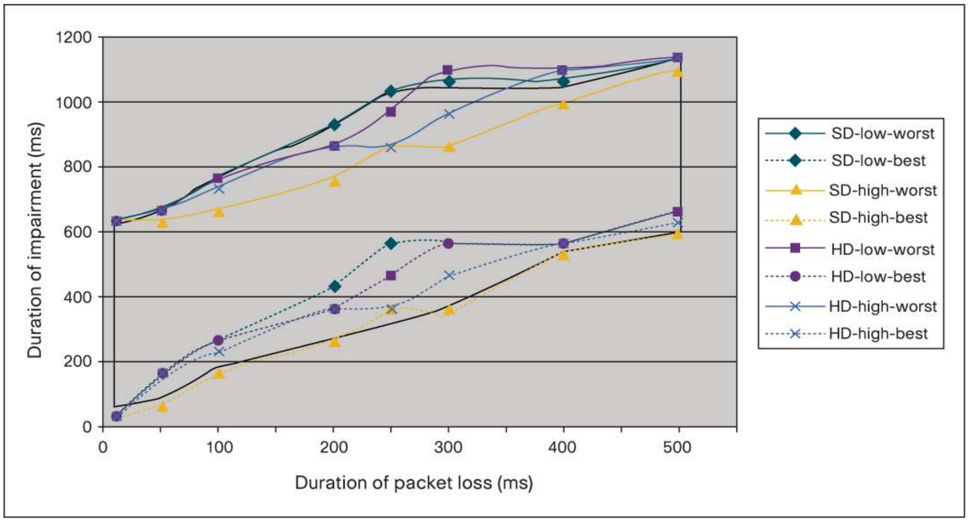

The results show that even 50 ms of packet loss can result in a significant visual impairment, which may last up to about 700 ms. With MPEG-4 encoding, impairments resulting from an equivalent duration of packet loss will generally be greater than for MPEG-2. For more information about this study, see the following documents:

Jason Greengrass, John Evans, Ali C. Begen, "Not All Packets Are Equal, Part I: Streaming Video Coding and SLA Requirements," IEEE Internet Computing, vol. 13, no. 1, January 2009, pp. 70-75
http://www.cisco.com/en/US/solutions/collateral/ns341/ns524/ns610/jg-je-ab-ieee-int-comp-jan09.pdf

- Jason Greengrass, John Evans, Ali C. Begen, "Not All Packets Are Equal, Part II: The Impact of Packet Loss on Streaming Video," IEEE Internet Computing, vol. 13, no. 2, March 2009, pp. 74-84
http://www.cisco.com/en/US/solutions/collateral/ns341/ns524/ns610/jg-je-ab-ieee-int-comp-mar09.pdf

Thus when the highest quality is required, even limited loss may be unacceptable. In these cases, loss recovery may be required, as discussed in the next section.

Availability

Network availability is the fraction of time that network connectivity is available between a network ingress point and a network egress point. For video, however, simply having connectivity is not enough, so service availability is often a more meaningful metric. Service availability is a compound metric, defined as the fraction of time the service is available between a specified ingress point and a specified egress point within the bounds of the other defined SLA metrics for the service, for example, delay, jitter, and loss.

The highest levels of service availability in IP/MPLS networks are achieved by combining the following approaches:

- Selecting carrier-class network elements with resilient processors and power supplies, resulting in high mean time between failures (MTBF) and low mean time to repair (MTTR) and thus high element availability
- Designing the network to be resilient with no single points of failure, employing redundancy in both network elements and links; this requires identifying any shared risks (for example, a situation in which a failure of a duct may affect two links) and constructing the network so that the failure of those shared risks will not unacceptably affect service
- Using IP and MPLS fast-convergence and fast-reroute technologies, augmented by rapid failure-detection capabilities such as IPoDWDM technology, to rapidly switch to alternate paths in the presence of network element failure, minimizing packet loss resulting from the failure

- Employing advanced high-availability techniques to minimize the impairment and packet loss that may result from route processor upgrades or failures
- Using DiffServ QoS, admission control, and capacity planning to meet the SLA requirements for delay, jitter, and loss, if necessary in cases of network element (link and node) failure
- Using transport and application-level approaches to recover from any loss experienced, thus providing lossless transport
- Using a service management solution that is closely coupled to the video transport network, to rapidly isolate and identify faults that affect service when they occur

The use of these techniques and technologies, which are described in more detail in the next section, provide network and service availability of greater than 99.999 percent with IP/MPLS networks.

Video-Enabling Transport Technologies

These combined SLA requirements place exacting requirements on the network. Network infrastructures that were designed for best-effort delivery, with relatively low bandwidths and high latencies, will not be able to meet the requirements of next-generation video services. Networks optimized for premium video services, however, will support other applications. As a result, the latest network infrastructures are heavily influenced by video applications.

Medianet Technologies

To support these requirements, Cisco is developing a suite of networking technologies that span the range of video-optimized products and define the medianet - an intelligent network optimized for rich media. For more details about Cisco medianet, see http://www.cisco.com/web/solutions/medianet/sp.html.

The following medianet technologies can be applied to enhance the SLAs offered for transported video services, allowing broadcasters and service providers to choose the optimal set of capabilities to meet their video service requirements.

Fabric-Based Multicast Replication

Efficient multicast replication is fundamental to supporting point-to-multipoint video services such as broadcast TV. The Cisco video-optimized platforms use innovative service-intelligent switching fabrics built for massive multicast replication and forwarding, which provide superior in-fabric replication compared to products that use inefficient line-card-based replication.

Multi-Priority DiffServ Schedulers

The Cisco DiffServ IP/MPLS QoS implementation is mature and very widely deployed in the industry. Cisco has led the industry in developing multiple priority-scheduling implementations that have been optimized through focused development and years of experience supporting premium services, including video. These implementations provide end-to-end jitter of less than 1 ms in high-speed multiservice IP/MPLS networks, as described in the following document:

- Clarence Filsfils and John Evans, "Deploying Diffserv in IP/MPLS Backbone Networks for Tight SLA Control," IEEE Internet Computing, vol. 9, no. 1, January 2005, pp. 58-65. http://www.cisco.com/en/US/prod/collateral/routers/ps167/prod_white_paper0900aecd802232cd.pdf

The following study demonstrates that these SLAs are achievable in high-speed multiservice IP/MPLS networks today; worst-case jitter was measured at less than 700 microseconds for probes sent at 1 Mbps during a 7-day period between the U.S. East- and West-coast POPs of a Tier-1 IP network service provider.

- Steve Casner, Cengiz Alaettinoglu, Chia-Chee Kuan, “A Fine-Grained View of High-Performance Networking,” Packet Design, NANOG 22, May 20-22, 2001.

Fast IP-Routing Protocol Convergence

In IP networks, convergence is the process by which routing protocols find alternative paths around failed network elements (links or nodes). Cisco leads the industry in the development of implementation and protocol optimizations such as IP prefix prioritization, which can reliably deliver significantly subsecond network convergence times in cases of network element failure for unicast routing protocols such as Open Shortest Path First (OSPF), Intermediate System to Intermediate System (IS-IS), and Border Gateway Protocol (BGP), and multicast routing protocols such as Protocol Independent Multicast (PIM). For details, see the following document:

- Pierre Francois, Clarence Filsfils, John Evans and Olivier Bonaventure, "Achieving subsecond IGP convergence in large IP networks," ACM SIGCOMM Computer Communication Review, Vol. 35, Issue 3 (July 2005), pp. 35-44
  http://www.cisco.com/en/US/solutions/collateral/ns341/ns524/ns610/cf-je-ccr-igp-convergence.pdf

As a specific example, for 400 multicast groups (potentially equivalent to 400 broadcast TV channels), convergence times of about 200 ms are realistically achievable today with optimized implementations. Fast Convergence represents the baseline for all video transport deployments and is available today through Cisco IOS® Software and Cisco IOS-XR Software across the range of video-optimized network products offered by Cisco.

Fast Convergence may be enhanced with other technologies where faster recovery from failures is required, for premium video transport services.

MoFRR

MoFRR is a simple enhancement to PIM Sparse Mode (PIM-SM) processing (defined in RFC 4601), which can provide the capability to instantiate multiple spatially redundant branches of the same multicast tree between a source and receiver. MoFRR can be used to further reduce multicast convergence times to a few hundreds of milliseconds in link and node failure cases, without significantly increasing network complexity. More information about MoFRR is available in the following documents:

- A. Karan, C. Filsfils, D. Farinacci, “Multicast only Fast Re-Route,” IETF draft:
  http://www.ietf.org/internet-drafts/draft-karan-mofrr-00.txt
- Dino Farinacci, Clarence Filsfils, “A Simple and Efficient 0(50msec) Resilience Technology for IPTV,” NANOG 42, February 2008 http://www.nanog.org/meetings/nanog42/presentations/DinoFilsfils_iptv.pdf

MPLS Traffic Engineering with Fast Reroute

MPLS Traffic Engineering (MPLS TE, defined in RFC 3209) uses the resource RSVP (defined in RFC 2205) to signal MPLS TE tunnels, which can be used to explicitly define video-connection paths through the network and provide admission control and bandwidth reservation. MPLS TE can also be used in conjunction with DiffServ for DiffServ-aware (class aware) traffic engineering. The Cisco MPLS TE implementation is the most widely deployed in the industry. Cisco has developed point-to-multipoint (P2M) MPLS traffic engineering (defined in RFC 4875) across its range of video-enabled routers, in support of premium multicast-based video transport applications, such as video contribution environments.

Where MPLS TE is deployed, Fast Reroute (FRR, defined in RFC 4090) can be used in addition to Fast Convergence. FRR allows precalculated backup TE tunnels to be used to protect against link, interface, and line-card failures, so that traffic can be rerouted around the failed element, restoring connectivity within 50 ms of the failure.

IPoDWDM

Cisco introduced IPoDWDM to the industry in December 2005. Since then, providers have used the power of IPoDWDM to combine and distribute video content rapidly and efficiently over all-IP networks. The Cisco IPoDWDM solution allows the reduction of transport elements, reducing CapEx and OpEx. Also, through integration of the transponder, the router line card participates in G.709 signaling and gains information about the link bit error rate (BER). As a result, in some failure cases the router can detect that links are degrading before they have actually failed (before they have started to drop packets), and can proactively switch to an alternative path. This proactive protection capability can be used in conjunction with Fast Convergence and FRR, with the potential to significantly reduce the losses resulting from network failure. The Cisco IPoDWDM solution is available across the Cisco core, edge, and aggregation platforms.

High-Availability Technologies

Cisco high-availability technologies include Nonstop Forwarding (NSF), Stateful Switchover (SSO), and In Service Software Upgrade (ISSU), which reduces the impact and packet loss that may result from upgrades or failures of router control-plane processors.

Transport and Application Approaches to Loss Recovery

The preceding network-level technologies may be further augmented by transport- and application-level approaches to recover from any loss experienced, so as to provide lossless transport for premium video services. Different deployment models for loss management and recovery may be applied in each error domain.

FEC

FEC (Figure 6) adds redundancy to the transmitted data to allow the receiver to detect and correct errors (within some limits) without the need to resend any data. FEC is most useful for recovering from intermittent packet loss due to lower-layer bit errors, for example, errors in the access and aggregation error domain.

Figure 6. Forward Error Correction (FEC)

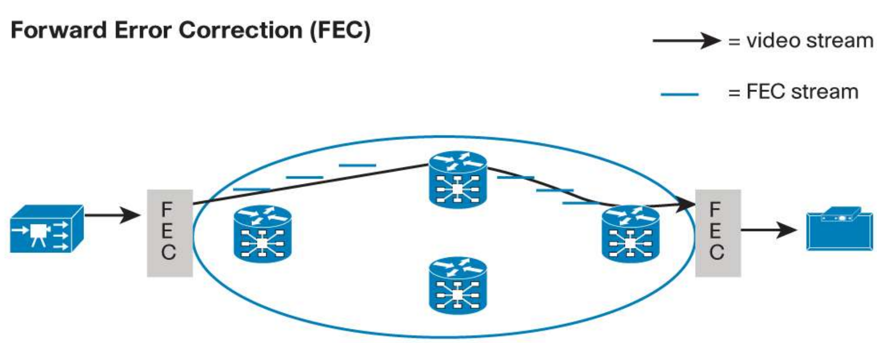

Spatial Redundancy

With spatial redundancy (Figure 7), also known as live-live, two identical streams are sent over diverse paths between the sender and receiver. If a network failure affects one stream, the other will still be received and the playout will be will be uninterrupted. Spatial redundancy is most useful for recovering from burst losses due to network failure events, such as link or node failures (for example, in the core and potentially in the aggregation error domains).

Figure 7. Spatial Redundancy (Live-Live)

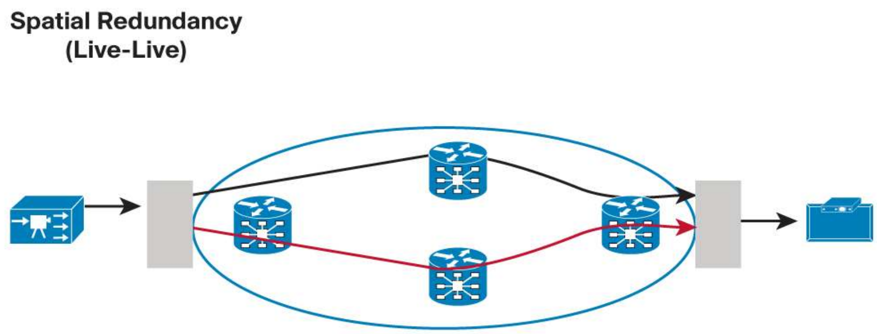

Several IP/MPLS network technologies can be used to help ensure that the streams follow spatially diverse paths, including MoFRR, MPLS TE, and Multi-Topology Routing (MTR). The resulting potential deployment models offer different advantages and disadvantages, and different tradeoffs between cost and complexity. The criteria for analyzing the suitability of these models to a particular deployment and error domain are: the amount of loss and delay each method incurs, the bandwidth each method uses, and the complexity and cost each model produces. An ideal solution has the following characteristics:

- Provides lossless transport through failures of SHE and core components
- Provides negligible overall delay to the transported stream
- Requires only the basic MPEG stream bandwidth to be provisioned over the working and failure case paths
- Does not significantly increase core network costs for design, deployment, and operations

Table 1 summarizes a relative analysis of some of the approaches that can be used by secondary distribution services to minimize loss and therefore improve quality of video transport through the core.

Table 1. Analysis of Models for Spatially Redundant Video Transport

Model	**QoE**	**BW Usage per Stream (Y)** Where X = MPEG stream bw, Z = FEC o/h, and W = FEC + TR overhead, (W < Z)	**Network Complexity**	**Video End System Complexity**	**Delay Impact**
Fast Convergence	Lossy; target < 1 GOP	Working case: Y = X Failure case: Y = X	Low	Low	Zero
MPLS TE FRR	Lossy; < 1 GOP	Working case: Y = X Failure case: Y = >2X	Medium	Low	Zero
MoFRR	Lossy; < 1 GOP	Working case: Y = X Failure case: Y = X	Low	Low	Zero
MoFRR + SR	Lossless	Working case: Y = X Failure case: Y = X	Low	Medium	Low
MPLS TE + SR	Lossless	Working case: Y = X Failure case: Y = X	High	Medium	Low
MTR + SR	Lossless	Working case: Y = X Failure case: Y = X	High	Medium	Low

Real-Time Retransmission

Media streams that use RTP are to some extent resilient, because receivers may use the mechanisms defined for the Real-Time Control Protocol (RTCP) to report packet reception statistics and thus allow a sender to adapt its transmission behavior. Additional techniques are defined within the IETF to extend the basic capabilities of RTCP, to allow faster feedback of packet loss from receivers to senders (defined in RFC 4585), so lost packets may be retransmitted (defined in RFC 4588). Within a defined time period, receivers can detect sequence-number gaps in the received stream indicating lost packets, and report these using RTCP NonAcknowledgements (NACKs) to the sender, which retransmits the lost packets.

Real-time retransmission is most useful for recovering from intermittent packet loss due to lower-layer bit errors, for example, errors in the access and aggregation error domain.

These capabilities are being developed in the Cisco VQE and DCM products. The Cisco combined network and video solution can deliver a lossless video transport solution end-to-end, satisfying the most stringent video SLA requirements for premium contribution and distribution services.

Service Management and Monitoring

The ability to verify that the IP network is delivering the required SLAs for video, and to identify problem areas, is essential to successful IP video deployments. This verification includes both element management and service management. Service management includes service monitoring at the device, system, and application levels. Cisco has developed the following capabilities to provide enhanced video transport service management and monitoring.

Embedded Network Video Monitoring

Embedded network video transport monitoring can detect issues with individual video flows at identified routers. This capability is being developed across the Cisco video-optimized routers to provide metrics including Media Delivery Index (MDI), Media Rate Variation (MRV), Media Loss Rate (MLR), and Media Stop Events (MSE).

These embedded video monitoring capabilities provide pervasive router support and deployment for comprehensive and cost-optimized video-quality problem detection and isolation, with increasing specificity of troubleshooting down to a particular device or link.

Trap and Clone

As a complement to embedded video monitoring, trap and clone provides the capability to tap into a video stream within the network, copy the packets from that stream, and transport them to a location where a video quality analyzer can assess the quality of the video stream. This capability makes efficient use of video quality analyzers for detailed video stream quality monitoring, with reduced CapEx.

Cisco VAMS

The Cisco VAMS is an innovative solution for video services, which provides an end-to-end view of the condition of the video transport network. Cisco VAMS combines information about the network topology and status with an assessment of the quality of the transported video services, through embedded or standalone video quality monitoring, to facilitate the rapid isolation, identification, and impact analysis of faults that may affect the quality of transported video services.

Video-Optimized Network Platforms

Cisco has embedded video intelligence in its core and edge portfolio of router products, offering a wide selection of products that provide medianet-enabled, video-optimized transport and that address the full range of network requirements.

Figure 8. Cisco Network Platforms

The Cisco portfolio of router products supports the strictest SLAs and enhanced video monitoring capabilities demanded for the highest-quality video transport services.

Cisco CRS

The Cisco CRS platform is a carrier-routing system that offers continuous system operation, superior service flexibility, and system longevity. Powered by Cisco IOS XR Software, Cisco CRS can scale system capacity up to 92 Tbps (Cisco CRS-1) and 322 Tbps (Cisco CRS-3).

The Cisco CRS-1 and CRS-3 carrier routing systems can power OC-768c/STM-256c IP and DWDM interfaces while supporting up to 1152 40-Gbps line cards of Packet-over-SONET (PoS), wavelength-division multiplexing (WDM), and Ethernet interfaces. The Cisco CRS routers help simplify today's networks while protecting investments for decades to come.

The Cisco CRS-1 is built on Cisco IOS XR Software, a self-healing operating system for multishelf, multiterabit carrier infrastructures. This microkernel-based operating system provides highly specific process independence, fault containment, and isolation. With these capabilities, the Cisco CRS-1 can be maintained, upgraded, enhanced, and scaled without requiring service interruptions.

The Cisco CRS platform combines:

- The Cisco QuantumFlow Array, a chipset architecture engineered in multiple dimensions of scale, services, and savings
- Cisco Silicon Packet Processor (SPP), a highly sophisticated 40-Gbps application-specific integrated circuit (ASIC)
- Cisco IOS XR Software, with Cisco Service Separation Architecture (SSA) and Cisco Service-Intelligent Switch Fabric to provide excellent service flexibility and capability

With comprehensive service separation and complete line-rate feature flexibility, the Cisco CRS can deliver the capabilities needed by the most advanced converged network services today and tomorrow. For more information about the Cisco CRS, see http://www.cisco.com/en/US/products/ps5763.

Cisco 7600 Series Routers

Cisco 7600 Series Routers offer integrated, high-density Ethernet switching, carrier-class IP/MPLS routing, and 10-Gbps interfaces. Video and multicast admission control capabilities in Cisco 7600 Series Routers provide superior end-to-end admission control for broadcast TV and VoD.

The Cisco 7600 Series supports high-performance IP/MPLS features as well as scalable personalized IP services at the network edge, improving operational efficiency and increasing return on network investment. The Cisco 7600 Series is the industry's first carrier-class edge router to offer integrated, high-density Ethernet switching, carrier-class IP/MPLS routing, and 10-Gbps interfaces, helping service providers to deliver both consumer and business services over a single converged carrier network.

Cisco 7600 Series Routers provide outstanding availability, scalability, security, flexibility, and the features that broadcasters and service providers demand for video transport services. For more information about the Cisco 7600 Series, see http://www.cisco.com/go/7600

Cisco ASR 9000 Series Routers

The Cisco ASR 9000 Series offers up to 400 Gb of capacity per slot, for up to 6.4 Tb of total capacity. With these routers, the edge gains a robust and service-tested operating environment that meets the reliability and performance requirements for a nonstop video experience.

In 10-slot and 6-slot versions, the Cisco ASR 9000 Series routers are power-efficient, space-efficient, and reliable enough to maintain nonstop video and extensive content service delivery from the network edge. Modular power supplies are brought online only as capacity is increased; the Cisco ASR 9000 Series consumes only the power needed to support the configuration demands of the system. Side-to-back airflow makes the 6-slot Cisco ASR 9000 Series capable of further reducing heat dissipation and power consumption in space-constrained facilities.

The Cisco Advanced Video Services Module (AVSM) is a major innovation, providing terabytes of streaming capacity at the aggregation edge while simultaneously offering content caching and ad insertion. The Cisco AVSM also features integration of Cisco VQE capabilities of fast channel change and error correction, all on a single blade that slides into the Cisco ASR 9000 Series routers. For more information about the Cisco ASR 9000 Series, see http://www.cisco.com/go/asr9000.

Cisco DCM Model D9000

The Cisco DCM offers the next generation of MPEG processing equipment, which combines compactness, flexibility, and high performance for video distribution applications. The Cisco DCM can simultaneously process up to 2000 SD or up to 500 HD MPEG video streams from a single platform. This capability will give operators new flexibility in multiplexing architectures, as well as impressive scalability and cost-per-stream points not previously available. The software is upgradable, with industry-leading capabilities that bring operational and economic benefits in MPEG processing applications, including:

- More efficient transrating, grooming, and encryption
- Transrated video quality that is superior to other offerings, supporting more digital programs in the same bandwidth
- Simultaneous support of existing ASI and the latest IP networked devices
- Scalable processing, from 1 to 2000 streams
- Extremely high network reliability
- Readiness for new applications through software upgrades: MPEG-4 transcoding, VBR rate clamping and Open Channel Associated Signaling (CAS)
- Optional built-in digital video broadcasting (DVB) scrambler that allows easy integration with conditional access systems
- Optional card that facilitates Serial Data Input (SDI) video transport, with support for both Conference of the Parties (COP4) FEC and hitless switchover with spatial redundancy.

For more information about the Cisco DCM, see
http://www.cisco.com/en/US/prod/collateral/video/ps9159/ps9195/ps9230/7004373.pdf

Cisco VQE

Cisco VQE includes both VQE-Client (VQE-C) and VQE-Server (VQE-S) features and provides the following capabilities.

- Error repair and reporting: Help ensure that errors can be detected by any set-top box (STB) and retransmitted within milliseconds to maintain a transparent visual experience. Combining FEC and RTP-based real-time retransmission for error recovery for both unicast and multicast video traffic, Cisco VQE helps to increase data rates over a larger percentage of the subscriber access lines and improves video quality over the entire video aggregation and access network. Cisco VQE technology for error recovery improves the effective network video quality while providing all subscribers with a consistent entertainment experience.

  Whenever errors cannot be transparently repaired, an error detection alert is forwarded to troubleshooting systems, such as Cisco VAMS, so they can instantly initiate repair. Postrepair notifications are sent to consumers, thereby reducing call-center escalations.
- Rapid channel change: To reduce the channel change time, Cisco VQE offers rapid channel change, which effectively eliminates or minimizes the main sources of channel change delay. The Cisco VQE rapid channel change solution employs standard RTP and RTCP protocols to perform the appropriate signaling between the STB and the video aggregation network that will optimize the effective channel change time. This solution is highly scalable in terms of the number of concurrent subscribers supported, enabling service providers to control the CapEx and OpEx associated with their video service delivery.
- Quality of experience monitoring: The Cisco VQE technology 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 allow 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.

For more information about the Cisco VQE technology, see

http://www.cisco.com/en/US/solutions/collateral/ns341/ns524/ns610/net_implementation_white_paper0900aecd8057f290.html.

Service Partner

Effective network and capacity planning is crucial to helping ensure that the SLA requirements for video services can be met. To help service providers meet this goal, Cisco has teamed with Cariden Technologies, which creates innovative software tools for network engineers, architects, planners, and operations personnel. The Cariden software product suite, MATE, is emerging as a standard tool for IP/MPLS network and SLA planning, traffic engineering, and network optimization for service provider networks around the world, including AOL, AT&T,

Deutsche Telecom AG, Orange, Nippon Telegraph and Telephone (NTT), Swisscom, and Time Warner Cable. With support for IP/MPLS technologies, including multicast, MPLS TE, MoFRR, and DiffServ, Cariden MATE is well suited to meet the network and SLA planning needs of video-optimized networks. For more information about Cariden products and services, visit http://www.cariden.com.

Conclusion

The move to IP/MPLS-based networks supporting video transport increases the need to address the main challenges facing service providers looking for next-generation video transport solutions. The expertise in carrier-class IP/MPLS networks and video technology gained through the acquisition of Scientific Atlanta makes Cisco exceptionally capable of addressing these challenges. In combination with the Cisco customer-inclusive, systems-based development approach, the result is an outstanding suite of video-optimized products and solutions to address service provider requirements. These products and solutions take full advantage of embedded video intelligence capable of supporting the strictest SLAs, and provide the enhanced video monitoring capabilities demanded for premium video transport. These capabilities allow converged infrastructures to be shared by different services with reduced CapEx and OpEx, accelerating the deployment of premium video services and the optimal delivery of the highest-quality video.

Visit the following websites for examples of how Cisco products have been used to meet the needs of broadcasters and video service providers around the world:

- Lossless Video Architecture: http://www.cisco.com/en/US/netsol/ns577/networking_solutions_solution_category.html
- Neuf Cegetel: http://newsroom.cisco.com/dlls/2007/prod_061107.html
- Radiotelevisione Italiana (RAI): http://newsroom.cisco.com/dlls/2007/prod_091007b.html
- T-Com: http://newsroom.cisco.com/dlls/2006/prod_120406d.html
- Telia Sonera: http://newsroom.cisco.com/dlls/2007/prod_042407c.html

For More Information

For more information about the Cisco video-optimized network architectures, technologies, and products, contact your local Cisco account team and ask for SPSU Video Optimized Transport Systems. For more information about Cisco products and services, visit http://www.cisco.com

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.**

Printed in USA C11-637031-00 01/11