Innovations for Lossless Video Delivery Architecture

The explosive growth of video over IP and Multiprotocol Label Switching (MPLS) networks requires dependable, cost-effective transport and monitoring solutions because of the high bandwidth demands of video and sensitivity of video traffic to packet loss, delay, and jitter. To maintain a consistent, high-quality video experience for customers, service providers need intelligent solutions to effectively and efficiently transport large sets of video streams, to prevent video packet loss during periods of network congestion, and to monitor quality and facilitate troubleshooting. These solutions ideally should be integrated transparently into the network and network management systems to make up the modern “medianet” – an intelligent, all-IP network optimized for rich media.

This paper provides an overview of three Cisco technologies for optimizing medianets, specifically Cisco® Multicast-Only Fast Reroute (MoFRR), which provides a simple and efficient method for transport of reliable video streams in secondary distribution video applications; hitless switchover or Cisco Live-Live, which provides spatial redundancy for video streams and is useful in contribution video applications; and inline video monitoring (Cisco VidMon), which provides a network-based mechanism for monitoring of video quality.

The Growth of IP Video and Its Implications

Video is rapidly becoming one of the fastest-growing types of applications on the Internet, forecast to comprise 91 percent of all global consumer traffic and to be accessed by more than one billion users by 2014, according to the Cisco 2010 Visual Networking Index (VNI) study. Now accounting for more than one-third of all consumer Internet traffic, video has overtaken peer-to-peer traffic in volume. By 2014, video is also forecast to be 66 percent of the world’s mobile data traffic. This tremendous growth in video traffic encompasses an array of IP video applications – IPTV, cable TV (CATV), peer-to-peer TV, video on demand (VoD), three-dimensional (3D) and high-definition (HD) video, and web-based video conferencing – that are fast replacing traditional broadcast formats.

IP/Multiprotocol Label Switching (MPLS) networks have become well suited to delivering a high-quality video experience because of a range of features used to provide converged data, voice, and video services. Technologies that deliver quality of service (QoS), resiliency, availability, and scalability are making the IP network a preferred medium for the transport of video by service providers, studios, and production houses.

This surge in IP and MPLS video applications and video traffic volumes has led to new scrutiny among service providers to ensure that their networks can deliver video to customers with consistent high quality and rapid failover mechanisms. A poor video experience can be costly to service providers because of lost revenue from disgruntled customers, high volumes of help desk calls, and expensive network operations center troubleshooting or in-home or in-office service calls.

IP Video Characteristics

Video traffic in service provider IP networks is most commonly encoded using MPEG-2 historically, or more recently H.264 (MPEG-4 part 10) standards. Compression is used to reduce the bit rate of IP video to accommodate typical access network bandwidth constraints. An MPEG or H.264 encoder converts and compresses video signals into a series of pictures, frames, or group of pictures (GOP) using various specialized techniques, including:

- *Subsampling* to reduce color information that is less-sensitive to the viewer’s eye
- *Spatial compression* or *intracoding* to remove redundant information in frames
- *Temporal compression* or *interframe coding* to remove redundant information between frames

The temporal compression is the most interesting from a networking design perspective, because loss within a frame can affect subsequent frame delivery. Hence loss is nonlinear (for example, a 50-ms network loss could result in a video artifact lasting 1 second), especially if that loss is within the reference frame (I-frame) used for all temporal compression on the GOP.

To transport MPEG-encoded video over IP networks, encoded frame information is encapsulated within MPEG Transport Stream (TS) packets, which are then transported through the User Datagram Protocol (UDP) over IP. In some advanced IP video delivery networks, the MPEG TS is encapsulated in Real-Time Transfer Protocol (RTP) over User Datagram Protocol (UDP)/IP to deliver advanced service-level agreements (SLAs) by using the sequence number tracking of packets enabled by RTP. An MPEG frame can span multiple IP packets and an IP packet can contain information from two consecutive frames. Therefore, when combined with the encoded temporal compression algorithm, the loss of a single IP packet can lead to a serious loss of information and possibly degrade the viewer's experience.

Video Transport Challenges

For any transport (including SONET/SDH, ATM, or IP) of video content, the main service degradation challenges that must be properly mitigated to support SLAs are network delay, network jitter, and packet loss.

- Network delay results from propagation delays along the network path, switching and queuing delays at network elements on the path, and serialization delay (the delay in transmitting bits sequentially on a link). Poor application performance due to network control protocol processing and processing in application end systems may also contribute to network delay. Delays are usually visible to end customers only when they change channels on a set-top box (STB) or click on a video on demand (VoD) on their PCs. For video streaming, service providers typically want one-way network delays of fewer than 100 milliseconds (msec) to achieve channel change times of 1 to 2 seconds. In modern networks, switching, serialization, and queuing delays are minimal, measured in microseconds. Propagation delay is the outlier that affects only transcontinental traffic (approximately 5 msec for each 1000 km). Therefore for broadcast-only applications, end-to-end delays are not usually significant enough to seriously affect service in a well-designed networking environment.
- Network jitter results from the variation over time of packet latency across a network due to fluctuations in queuing and scheduling delays in network elements. Jitter could theoretically cause a video to flicker and introduce clicking or gaps in accompanying audio. The problem is usually remedied with playout buffers in receiving devices such as STBs. Jitter in modern networks is again measured in the microsecond range and is negligible for a well-designed and capacity-managed IP network.
- Packet loss occurs between network ingress and egress points. Packets are considered lost when they do not arrive within a defined period of time because of network congestion, lower-layer errors (for example, bit errors), or network element failures. Packet loss can lead to video anomalies such as pixilation, picture freeze, or complete loss of picture and audio. Reducing packet loss is the main engineering design problem in the delivery of high-quality video over modern networks.

Achieving Near Lossless Video Transport with Multicast-Only Fast Reroute

Packet loss is the primary challenge for high-quality IP video transport. Mechanisms to minimize its occurrence are a top priority for service providers. However, cost must also be considered because a complex design may provide benefits but can result in overly high operational costs. To reduce packet loss, network equipment vendors have focused on solutions for faster convergence during network failures. Three examples are specifically targeted at multicast delivery: faster Interior Gateway Protocol (IGP) convergence, MPLS Point-to-Multipoint (P2MP) Traffic Engineering (TE), and Cisco Multicast-Only Fast Reroute (MoFRR).

- **Faster IGP convergence:** In IP networks with modern link-state protocols (such as Open Shortest Path First [OSPF] and Intermediate System-to-Intermediate System [IS-IS]), convergence is measured at a node by how fast it receives link-state updates from the affected routers in the network and the time it takes to then run the Shortest Path Forwarding (SPF) algorithm on its link-state database (LSDB). Convergence in the early days of the Internet was measured in the second range. Advances by Cisco over 20 years of IGP design and improved hardware have lowered this value to the sub-200-ms range using IGP prefix prioritization of video prefixes. (Note: Multicast convergence times go above the 200-ms range because Protocol Independent Multicast [PIM] has to converge on Reverse Path Forwarding [RPF] interfaces after the unicast IGP has converged.) This benefit comes with no extra complexity because the fast convergence is default-engineered by Cisco when the IGP is configured on the network. If a service provider is transporting temporarily compressed content where loss is nonlinear and the service provider has a highly available network, then fast convergence best fits a low-cost IPTV delivery service mechanism.
- **Fast Reroute (FRR) for MPLS TE:** The FRR for MPLS TE tunnels feature was designed to provide 50-ms convergence for link failures by defining bandwidth-protected traffic-engineered tunnels in primary and backup. This feature has been extended to video delivery through the support of FRR for MPLS P2MP tunnels. Although very fast, the speed from this solution comes at considerable operational complexity because the primary and backup tunnels must be configured for all links traversed beforehand. Also, because of the P2MP nature of video multicast delivery, the backup tunnels could be doubling or tripling the bandwidth on a failure because of the video service already being delivered on the backup link. Therefore, FRR for P2MP is not a good fit for convergence applications. MPLS P2MP TE, however, is still useful in contribution networks for its constraint-based routing capability.
- **Cisco MoFRR:** A more recent solution from Cisco is MoFRR, which can provide sub-50-ms video-stream convergence times while being operationally simpler and contributing to lower operating expenses than other alternatives. MoFRR is a slight alteration to PIM state processing, which is completely compatible with non-MoFRR PIM state machines. It instantiates multiple spatially redundant branches of the same multicast tree between a source and a receiver if there is Equal Cost Multipath (ECMP). This method provides two active-active paths to the receiver, which are impervious to a single failure and allows the receiver to make the path selection based on quality. If the tail-end receiver router has the capability to make path selection based on traffic statistics on the incoming line-card – the solution becomes operationally low-cost and provides convergence times in the 40-ms range.

A summary of benefits of the Cisco MoFRR solution includes:

- Sub-50-msec multicast video convergence time with minimal configuration or operational overhead
- Video path protection from video source to the receiver router during both link and node failure
- Simple operation that does not require any protocol-level interoperation with other routers in the network and works with native IP Multicast

How MoFRR Works

In a standard PIM environment (Figure 1) with ECMP to the source, the router attached to the receiver chooses one of the upstream paths based on PIM RPF tie mechanisms to receive a video stream.

Figure 1. IP PIM Failover Without MoFRR

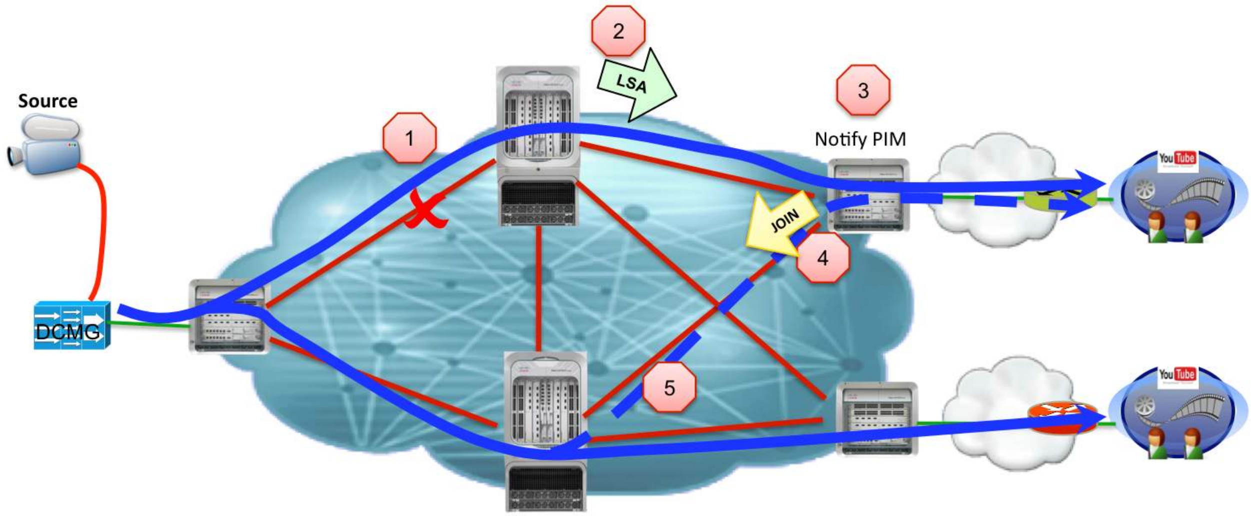

If a failure occurs anywhere on the active path, the tail-end router needs to:

1. Detect the failure…
2. …by receiving link-state updates from neighbors
3. Converge on new unicast routes and then notify PIM
4. Then PIM needs to signal on the new RPF interface
5. Wait for traffic to be restored

As shown in this the traditional multicast model, the amount of time between the failure of the active stream and recovery on an alternate path depends on numerous factors that must happen in a series of processes.

With MoFRR the tail-end routers can avoid steps 2 through 4, because all they have to do is detect a failure in the video stream and switch paths (refer to Figure 2).

This ability to detect a failure in the video stream and switch paths makes MoFRR much faster because it relies on a local tail-end router decision only. Loss of video detection equates to examining traffic statistics for incoming content from S,G routes and switching paths if nothing is received for more than 30 ms.

Figure 2. Cisco MoFRR

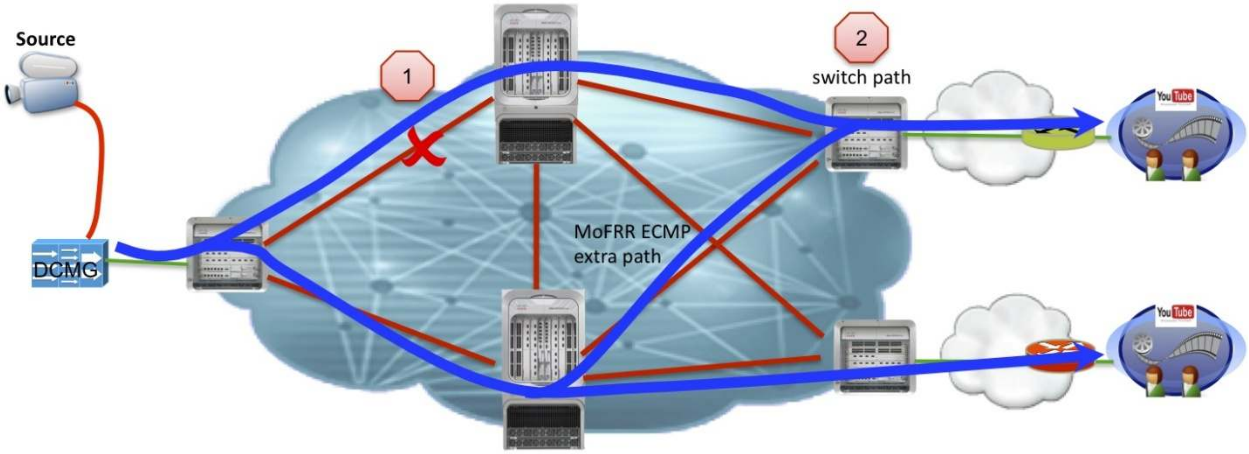

One of the objections sometimes raised about MoFRR is the bandwidth overhead due to the duplicated video stream received on the receiver router. This is not an issue when content is distributed to a dense population of subscribers, and in fact is actually the case for IPTV deployments where video content is dispersed fairly evenly to all points of presence (POPs) and a large percentage of all POPs have to deliver all the standard TV channels to subtended subscribers. With a parallel core, two core tree structures are receiving, transmitting, and replicating multicast to subscriber POPs efficiently. Hence the second ECMP PIM join from the subscriber POPs splices onto the alternative core tree very near to the subscriber POP. Therefore, there is no additional bandwidth overhead in the core network. Additional bandwidth is required only on the last-hop provider-edge router and only on its backup link to the redundant core router. Because this link is a link for resiliency purposes and therefore should be lightly loaded before the failover, providing backup video traffic on this tail-end link is not a problem.

The Simplicity of MoFRR

Simplicity is the big advantage of MoFRR. It provides very fast multicast convergence (sub-50-ms) with minimal configuration. MoFRR needs to be enabled only on the last-hop receiver router and does not require any configuration or protocol-level interaction with other routers. Aside from these operational savings, MoFRR provides 50-msec convergence with native IP PIM, so MPLS TE or FRR are not required.

MoFRR currently operates with native IPv4 Multicast as the video transport option. It is being extended to also work with other video transport technologies such as multicast Label Distribution Protocol (mLDP)-based transport options.

Achieving Lossless Video Transport with Hitless Switchover and Cisco Live-Live

Although MoFRR provides fast convergence, it is a convergence event, so by definition there is a cut in customer connectivity as processes are taking place at the network layer. If customer requirements specify no loss for transporting the content, then it is necessary to move up the software stack to look at application-level methods to recover from loss. Methods to provide lossless transport for video content include:

- Forward Error Correction (FEC)
- Spatial diversity and live merge (Cisco Live-Live)

FEC is useful to repair bit errors but less useful for convergence scenarios where frame loss is sequential and measured in milliseconds or longer (the overhead required by FEC to recover from continuous loss is substantial). Furthermore, the latency that FEC encoding introduces to the content stream is not acceptable for uncompressed contribution content.

Spatial diversity with live merge (Cisco Live-Live) provides lossless transport without any of the deficiencies of FEC (Figure 3). It provides lossless transport by providing dual live content streams across the IP backbone (for spatial diversity) and re-merging them at the receiver edge (an example of live merge) by looking at a sequence tag (using RTP), uniquely identifying the content within the IP transport and dropping duplicates. This process works on the premise that the chances of dual network link failures in a well-engineered and -operated network is close to nil. This premise is backed up by historical operational practice.

Figure 3. Cisco Live-Live

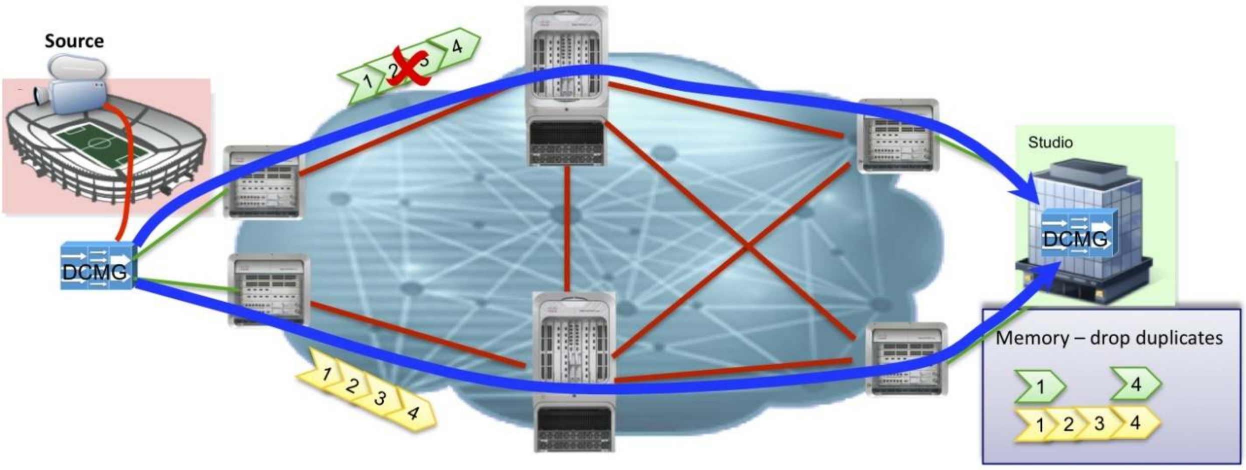

Cisco provides lossless video transport for contribution flows with the IP-video adapter product, the Digital Content Media Gateway. It provides the live merge function on RTP-encapsulated video Serial Digital Interface (SDI) content independent of the particular SDI encoding (from 270 Mbps to 3 Gbps).

Efficient Inline Video Monitoring

Until recently, IP networks were built for applications that can handle 2- to 4-second outages through retransmission of packets. But for video, as previously mentioned, a loss of even a few packets can result in visible video degradation. Existing IP Next-Generation Networks (NGNs) have network quality-monitoring tools such as IP SLA to measure packet loss, jitter, and delay, and these tools can be used to troubleshoot voice and data networks. But video traffic is sufficiently different that a stream that is deemed acceptable by IP SLA might still not deliver an acceptable, high-quality video experience. Therefore video performance monitoring requires other specialized tools.

Three fundamentally different architectural solutions are available to service providers for the monitoring of IP video:

- Video performance-monitoring probes placed at strategic network locations
- Probes deployed in router blades from the regional headend to the network edge
- Inline solutions that are integrated into the IP NGN

Probes are external and often expensive devices that are connected to the network to monitor video streams. Deployed at various critical points in the network, they receive duplicate video streams from the routers for analysis. Alternately, the probes can send Internet Group Management Protocol (IGMP) joins like a normal receiver to obtain video streams. The probes produce video-quality metrics that the NOC collects and analyzes. Probes are not integrated with the routers or the video path, so if problems are detected manual intervention is required to find the exact location of the problem in the path before the quality problem can be fixed – further slowing the process and affecting customer satisfaction. Each probe can monitor only a couple of channels, and networks have hundreds or thousands of channels, requiring a commensurate number of dedicated probes and resulting in high capital expenditures (CapEx).

Blade-based service-monitoring modules require costly blades and accompanying licenses for deployment throughout the network. They are integrated with routers and can handle more channels than individual video probes because of the large internal backplane bandwidth between the service blade and the router itself. Again, reports of video problems are referred back to the NOC, and fixing the problems involves reactive, manual intervention from operations personnel, because the blade-based monitoring is not in the path. The blade-based solution involves high OpEx and CapEx, especially as networks scale to handle more video for more customers.

Inline monitoring solutions, such as Cisco VidMon, provide network-based monitoring through the router transport line card. The Cisco VidMon solution monitors the video stream inline and in real time without sacrificing transport performance or scalability.

The benefits of Cisco VidMon over blade-based video service-monitoring modules and standalone video problems include:

- Tight integration between video-quality monitoring and transport operations: If the Cisco VidMon feature detects a video-quality problem, it can communicate to the router, which can automatically switch over to the backup path automatically and instantly, without any manual intervention.
- High scalability and performance: The router with the Cisco VidMon line card can monitor thousands of video streams in line and in real time and with no performance degradation.
- CapEx and OpEx savings: No additional hardware is required for Cisco VidMon to perform inline video monitoring; the transport hardware handles this monitoring. Operationally, no manual intervention from NOC personnel is required.

Cisco has compared the cost of blade-based monitoring with inline, real-time monitoring with Cisco VidMon (Figure 4).

Figure 4. Cisco VidMon Inline Monitoring vs. Blade-Based Monitoring: CapEx Comparison (14 aggregation nodes each monitor 40 Gbps per chassis)

Implementing blade-based solutions can incur capital costs 22 times higher than an equivalent Cisco VidMon solution. This analysis is based on deployment of 14 aggregation nodes. As for OpEx, with Cisco VidMon each intelligent node on the network is provisioned once. By contrast, multiple-blade solutions require complete system upgrades to manage and upgrade every node, creating ongoing operational challenges for networking staff.

Cisco VidMon Inline Video Monitoring Integrated with Cisco Video Assurance Management Solution

Video-quality metrics gathered by Cisco VidMon (Figure 5) can be integrated into centralized network management solutions such as the Cisco Video Assurance Management Solution (VAMS).

Figure 5. Cisco VidMon Video Metrics

Metric	Applicability
Media Delivery Index (MDI)	Measures MPEG 2/4 headers for loss and delay
Media Discontinuity Counter (MDC)	Measures MPEG 2/4 headers for the number of times loss was detected
Media Rate Variation (MRV)	Measures IP/UDP headers for delivery variations
RTP loss and jitter	Measures RTP loss and delay by examining the RTP header
Media Stop Event (MSE)	Notification of a monitored flow stops receiving traffic

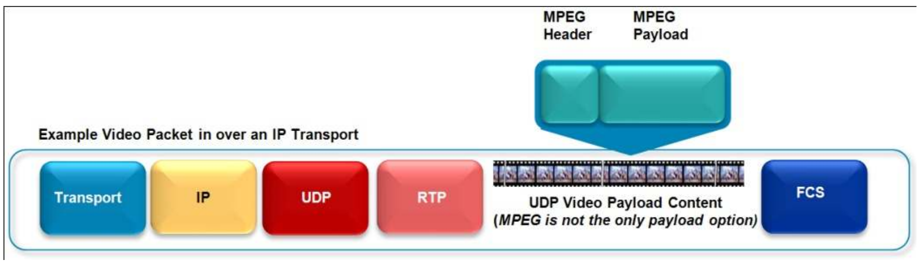

Cisco VAMS is a reference architecture for management application responsible for the monitoring and troubleshooting of video. The VAMS architecture includes mechanisms for polling the video quality metrics from the routers through a Simple Network Management Protocol MIB as well as receiving traps to create an end-to-end picture of the video quality in the network. The VAMS architecture also includes standard video probes that can be used in conjunction with the VidMon inline router metrics to provide additional visibility when needed. Management applications such as the Cisco Multicast Manager (CMM) can collect simultaneous VidMon and Video Probe statistics and overlay these quality metrics to topology map of a given video flow.

With Cisco VidMon, the router platform can monitor thousands of video flows simultaneously. A centralized network monitoring architecture such as Cisco VAMS can isolate the root cause of the poor-quality video flow to a certain part of the network to ease troubleshooting (Figure 6).

Figure 6. Inline Video Monitoring and Cisco VAMS Integration

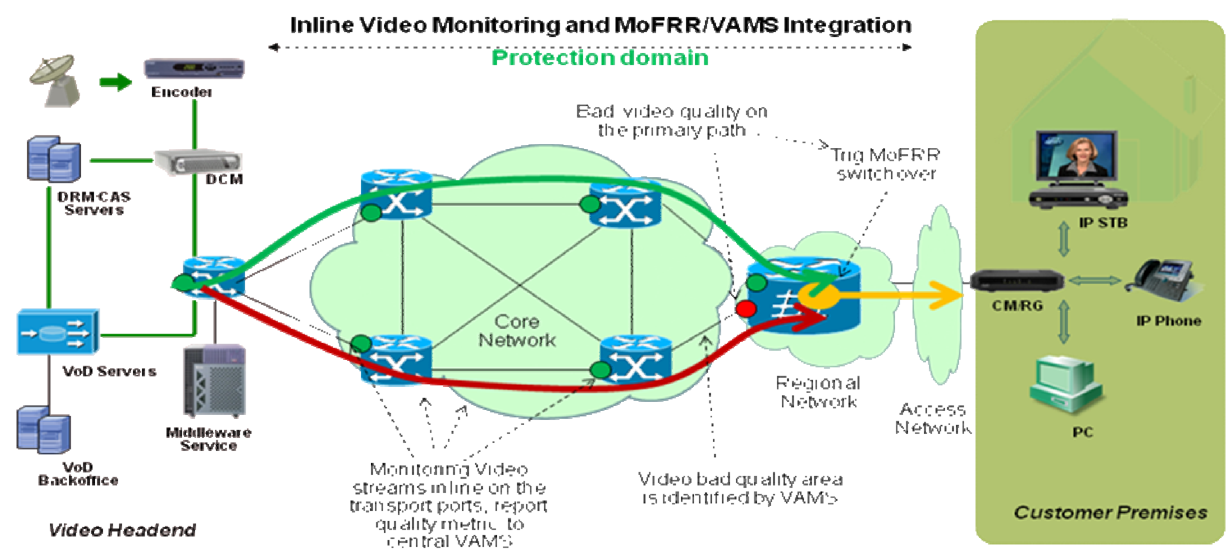

Conclusion

Video quality remains the top priority among service providers as they architect IP video networks. With innovative Cisco solutions such as Cisco MoFRR, hitless switchover (Cisco Live-Live), and Cisco VidMon, network routers can transport video with much higher quality and detect quality problems efficiently and proactively. MoFRR can provide sub-50-msec video path protection in leading Cisco router platforms and software without additional OpEx. Cisco Live-Live can provide lossless transport using Cisco video transcoders, and increasingly Cisco routers; and Cisco VidMon can detect both hard and soft failures, including link-quality problems. Both solutions deliver new benefits that provide a proactive instead of a reactive approach to video monitoring and failover while at the same time lowering the CapEx and OpEx of these solutions when compared to competitive approaches.

For More Information

For more information about the Cisco IP Next-Generation Network Carrier Ethernet System, please visit: http://www.Cisco.com/go/ce.

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 C22-636808-00 11/10