Toward Lossless Video Transport

Temporary outages in IP packet delivery can occur during network reconvergence events, and conventional error-control methods that end points use can't effectively deal with such extended blocks of packet loss. Service providers can employ a number of network technologies to minimize the IP packet loss during such outages, ultimately resulting in lossless video transport.

John Evans, Ali C. Begen, Jason Greengrass, and Clarence Filsfils
Cisco

IP is becoming the dominant network technology for video transport, and, in turn, video is becoming an increasingly significant component of IP network traffic (see up-to-date predictions at www.ciscovnipulse.com). In previous articles for *IEEE Internet Computing*, we described the IP service-level requirements for a video-transport service and considered the network factors that impacted viewers' quality of experience (QoE) for IP-based video streaming services such as IPTV and cable TV (CATV).[1,2] We also highlighted the impact that different durations of IP packet loss had on viewers' QoE, showing that due to the compression employed in video coding, even a single lost packet could result in several hundreds of milliseconds of visible impact to viewers.

Considering the primary causes of packet loss in IP networks – congestion, lower-layer bit errors, and network-element failures – service providers can use the Differentiated Services (www.ietf.org/rfc/rfc2475.txt) and Integrated Services (www.ietf.org/rfc/rfc1633.txt) architectures to engineer an IP or multiprotocol label switching (MPLS) network to help deliver the required delay, jitter, and loss rates for a video-transport service.[3] However, packet loss can still occur due to network failure events or lower-layer errors, which might result in visual impairments to a video service.

Several IP/MPLS network technology approaches can help minimize the duration of packet loss resulting from network-reconvergence events, but packet loss might still occur. Service providers can further augment these network-level technologies with application- or transport-level approaches to recover from any loss experienced due to either reconvergence events or lower-layer errors.

Here, we give an overview of the different network-reconvergence technologies and loss-recovery approaches, and compare various deployment models and their impact on a broadcast IP video service. We define such a service as one in which a source originates broadcast video content, which is distributed across IP or MPLS core and aggregation networks to the last-mile

1089-7801/11/$26.00

access connection and then the subscribers. Feedback from service providers suggests that the mean time between any type of failure for an element within a core and aggregation network is longer than 100 hours.

Comparison Methodology

We compare different deployment models in terms of the following characteristics:

- *QoE.* In comparing the different approaches for video transport, we must consider whether the models provide completely lossless video transport or experience packet loss, which degrades QoE. For lossy approaches, we compare the potential video impairments resulting from the anticipated packet loss. Error-concealment methods can help minimize the visible impact of such impairments. We use our previous findings in our comparisons to estimate the resulting QoE.[2]
- *Capacity provisioned and consumed.* It's standard practice for IP/MPLS transport networks to be built with redundancy in all but the last-mile access connection to the subscriber, so that they're resilient to single-element (such as a link or node) failure cases wherever possible. It's also normal for service providers to provision sufficient capacity on both working- and failure-case paths to support their premium services, such as video. When comparing different approaches, we consider the link bandwidth consumed and hence how much providers must provision per video stream on both working- and failure-case paths, relative to the basic video stream bandwidth *B*.
- *Delay.* Different models have different impacts on the delay of the transported video stream. Delay introduced by some models might have an adverse impact on the interactivity viewers experience.
- *Deployment considerations.* These include requirements imposed on the network to support the approach and the network's overall design, deployment, and operational complexity in support of that model.
- *Applicability.* These factors will affect different deployment models' applicability. Not all models will be suitable everywhere — for example, are they viable for the core network as well as aggregation and access networks?

An idealized video-transport solution would be lossless in the presence of packet loss due to network failure cases and lower-layer errors; would be bandwidth efficient, requiring that only the basic video stream be provisioned on working- and failure-case paths; would add negligible overall delay to the transported video stream; and wouldn't significantly increase network cost and complexity with respect to design, deployment, and operations. Such a solution might use different models in different parts of the network.

Network-Reconvergence Technologies

Networks are commonly built resiliently, so that if any single component (such as a link or router) between the source and the last mile to the subscriber fails, an alternate path will be available. The last mile to the subscriber is normally a single point of failure due to cost constraints. In the rest of the network, network-element failures can cause packets to drop until connectivity is restored around the failure. The resulting loss period depends on which underlying network technologies the provider uses. The various potential choices can result in loss periods of tens of milliseconds or more. In this article, we consider three approaches.

Fast IP Routing Protocol Convergence

Reachability or connectivity in IP networks is established by Interior Gateway (routing) Protocols (IGPs). For IP unicast traffic, after a network-element failure, even if an alternate path exists around the failure, a loss of connectivity will cause packet loss until the routing protocol establishes an alternate path. The process of finding alternate paths is called *convergence*; *convergence time* is the time between a network-reconvergence event — for instance, a link or node failure or recovery that causes a routing protocol recalculation — and when convergence occurs. The convergence time comprises the time to detect the failure, propagate the news of the failure throughout the network, and recalculate new paths at nodes or routers affected by the failure. The network-reconvergence process might need to complete at multiple routers before the end-to-end connectivity is re-established.

Multicast lets a source send traffic to a group of interested receivers. IP multicast is the default technology used for IP-based broadcast

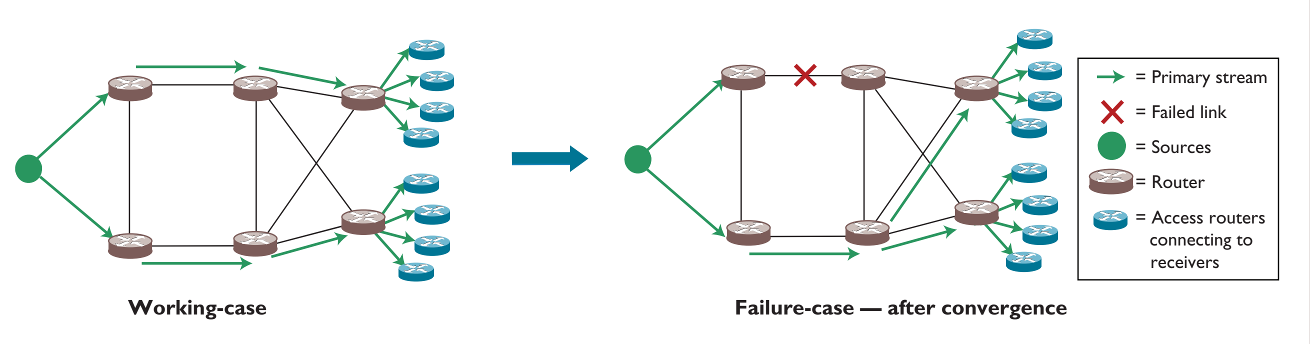

Figure 1. Capacity provisioned and consumed for fast convergence. In IPTV networks, the working-case paths for some receivers are the failure-case paths for others.

video services, such as IPTV and CATV, because of the scalability and bandwidth efficiencies it offers for point-to-multipoint services — that is, the source must send a packet only once, even if it needs to be delivered to millions of receivers. IP multicast connectivity depends on both multicast and unicast routing protocols. Protocol Independent Multicast Source-Specific Multicast (PIM-SSM; www.ietf.org/rfc/rfc4607.txt) is commonly used as the multicast IP routing protocol for point-to-multipoint video services. With PIM-SSM, the network builds separate multicast distribution trees (MDTs) rooted at each multicast source, and clients connected to the tree receive content directly from the source. The MDT is built from the receivers (the leaves) to the source (the root). In doing so, PIM-SSM relies on the unicast routing tables that the IGP establishes to determine which route is preferred and hence the interface toward the source. Multicast traffic on a particular MDT is accepted only if it's received on the preferred interface toward the source. Following a reconvergence event, the preferred path might change, so the interface toward the source might change. Thus, the multicast convergence time depends on the unicast convergence time.

We consider a network deployment in which a broadcast video service uses a video source that's originating a single IP stream per broadcast channel. This stream is distributed across the core, aggregation, and access networks to receivers using IP multicast, with fast convergence used to reduce the loss of connectivity resulting from reconvergence events. In practice, losses due to lower-layer errors might also occur throughout the network, although this is generally more prevalent in access networks and constrained to less than 10 ms. Such a deployment has the following characteristics.

QoE. Following a network-element failure, a loss in connectivity will occur for that stream until the IGP reconverges on the alternate path. In well-designed networks, where the routing protocol implementation is optimized for fast convergence, for 400 multicast groups (potentially equivalent to 400 broadcast TV channels), we can realistically achieve convergence times of approximately 300 ms. Considering the results we presented previously,[2] the resulting video impairments on an impacted video stream might be visible to viewers for up to approximately one second.

Capacity provisioned and consumed. Service providers generally provision sufficient core network bandwidth to be able to support the IPTV load both in working-case and single-element failure-case conditions. For this model, the basic video stream bandwidth B must be provisioned per channel on working- and failure-case paths, as Figure 1 shows. In normal working-case conditions, bandwidth B is consumed only on the working-case paths. Where the receivers are spread around the topology, however, as is common with IPTV, the net effect is that the working-case paths for some receivers are the failure-case paths for other receivers, so in practice, bandwidth B will be consumed on both working- and failure-case paths most of the time.

Delay. Typically, the failure-case path will have a similar delay to the working-case path, in which case this model adds no significant delay to the transported stream.

Deployment considerations. This approach is the simplest in terms of network complexity in that no extra requirements are placed on the network.

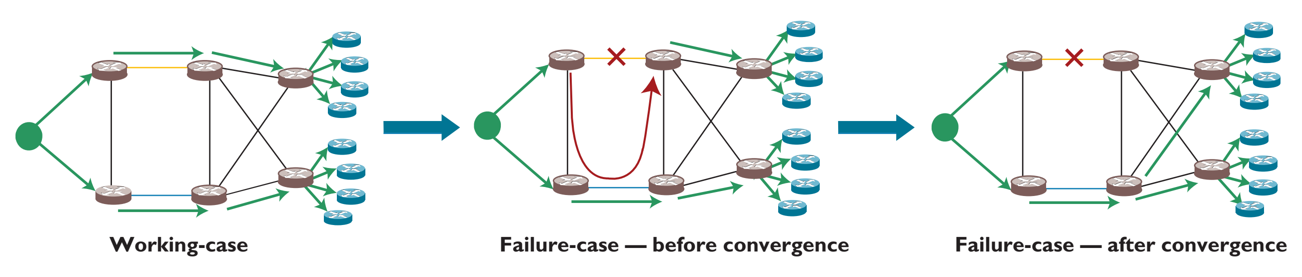

Figure 2. Capacity provisioned and consumed for MPLS Traffic Engineering Fast Reroute (MPLS-TE FRR). To protect from failure of the orange link, the blue link must be provisioned for twice the normal multicast load.

Applicability. We can apply this model in both the core and aggregation networks and consider it as a baseline for broadcast video deployments, against which we can compare other models.

MPLS Traffic Engineering Fast Reroute

MPLS Traffic Engineering (MPLS-TE; www.ietf.org/rfc/rfc3209.txt) uses the Resource Reservation Protocol (RSVP; www.ietf.org/rfc/rfc2205.txt) to signal MPLS-TE tunnels, which can be used to explicitly define which paths through the network video connections take and to provide admission control and bandwidth reservation along the paths. RFC 4875 (www.ietf.org/rfc/rfc4875.txt) defines enhancements to MPLS-TE supporting point-to-multipoint services.

Where MPLS-TE is deployed, providers can use Fast Reroute (FRR; www.ietf.org/rfc/rfc4090.txt) in addition to fast convergence. MPLS-TE FRR is a local protection scheme that enables precalculated backup TE tunnels for protection against link, interface, and linecard failures. When a failure occurs, no delays are associated with the distribution of updated routing information or routing table recalculation prior to connectivity being restored. This is unlike IGP convergence, in which a distributed computation process is triggered after a failure. Consequently, where an alternate path exists, connectivity restoration following network-element failures is always likely to be faster with MPLS-TE FRR than with IGP fast convergence.

We consider a network deployment in which MPLS-TE FRR reduces the loss of connectivity following reconvergence events. Such a deployment has the following characteristics.

QoE. Where service providers use MPLS-TE FRR, the loss of connectivity following reconvergence events is shorter — typically within 50 ms — and more deterministic than with IGP fast convergence. It's also independent of the number of multicast groups. Our earlier results, however, demonstrate that the resulting video impairments might still be visible to viewers for up to approximately 700 ms.[2]

Capacity provisioned and consumed. In working-case conditions, the bandwidth consumed for this model is the same as for the previous one. When providers use point-to-point MPLS-TE FRR backup tunnels to protect multicast services, however, traffic duplication (with link protection) or worse (with node protection) might occur, while the MPLS-TE FRR protection is active following network-element failure cases — that is, before the network has converged around the failure, as Figure 2 shows. This issue is only present as a transient, which lasts from the time of the failure event until the network has converged. However, this is the time period over which the FRR deployment is expected to provide better QoE for an IPTV service than fast convergence. So, service providers must provision bandwidth $\geq 2 \times B$ on failure-case paths for MPLS-TE FRR's faster rerouting capabilities to be effective.

Delay. Typically, the failure-case path will have a similar delay to the working-case path, in which case this model adds no significant delay to the transported stream.

Deployment considerations. This approach implicitly requires deploying MPLS-TE, which incurs an additional level of complexity for network design, deployment, and operations. Furthermore, additional planning is required to make sure that backup MPLS-TE tunnels are diverse from the primary elements they're protecting.

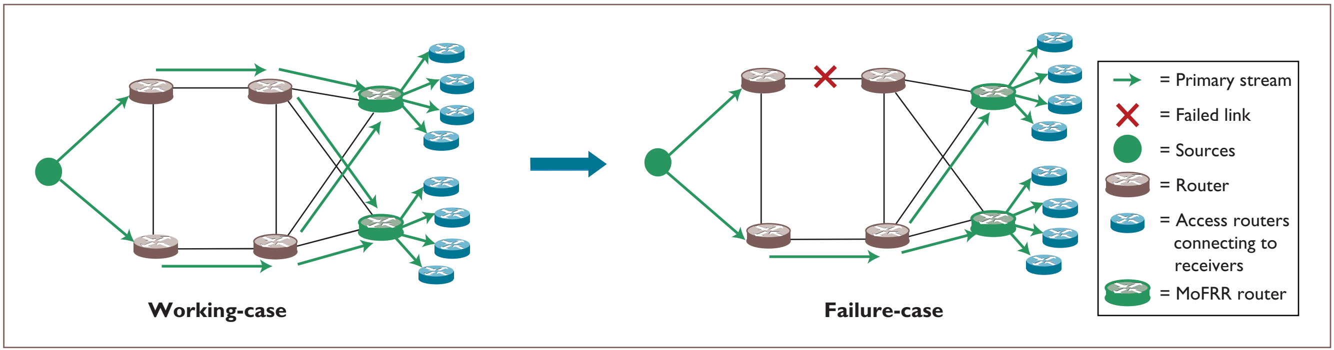

Figure 3. Capacity provisioned and consumed for Multicast Only Fast Reroute (MoFRR). The basic video stream bandwidth B *must be provisioned per channel on working- and failure-case paths, the same as for the IGP fast-convergence case.*

Applicability. The bandwidth requirements of using MPLS-TE with FRR can limit its applicability to core networks, where 2 × B capacity overprovisioning might be acceptable for transient periods.

Multicast Only Fast Reroute

Multicast Only Fast Reroute (MoFRR) is a simple enhancement to PIM-SSM processing that provides the ability to instantiate multiple branches of the same multicast tree between a source and receiver. If the branches are spatially redundant, MoFRR can further reduce the effective multicast convergence time. For instance, when a network-reconvergence event impacts one branch, the device implementing MoFRR simply needs to detect the failure and switch to the working branch.

We consider a network deployment in which a source is originating two resilient multicast streams carrying the same content. A router between the source and the last-mile access connection will receive both streams, from which it will forward on only one. This router uses MoFRR to instantiate the resilient multicast streams. Such a deployment has the following characteristics.

QoE. MoFRR effectively reduces the losses experienced due to network-element failures by removing the dependency on IGP convergence. MoFRR can potentially reduce the effective loss duration to less than 100 ms;[4] correlating with our previous results, we can conclude that the resulting impairments might still be visible to viewers for up to approximately 750 ms.

Capacity provisioned and consumed. A common misconception exists that approaches benefiting from resilient multicast streams consume additional network capacity compared to models using a single stream or forward error correction (FEC). Where receivers are spread around the topology and multiple paths go back to the source, however — as is common in IPTV — the net effect is that the working-case paths for some receivers are the failure-case paths for others. Hence, the resilient multicast streams are generally already established. In practice, this means that the basic video stream bandwidth B must be provisioned per channel on working- and failure-case paths, the same as for the IGP fast-convergence case, as Figure 3 shows.

Delay. Typically, the failure-case path will have a similar delay to the working-case path, in which case this model adds no significant delay to the transported stream.

Deployment considerations. This approach is relatively simple to design, deploy, and operate. The only network requirement is for MoFRR, and it's a simple configuration if one is required.

Applicability. This model can be applied in both the core and aggregation networks.

Loss-Recovery Approaches

If service providers don't employ loss-recovery approaches, the loss of an MPEG frame will generally result in a visible impairment. Where packet loss occurs, we can deal with it using loss-recovery techniques in conjunction with the network-convergence technologies we described in the previous section.

We consider four approaches for loss repair in streaming video. (A detailed overview of these approaches and their variants is available elsewhere.[5,6])

Forward Error Correction

FEC relies on redundant information being sent along with the media stream to allow lost packet recovery without the need for retransmission. Many FEC schemes are available today. Here, we use the parity FEC scheme due to its ease of implementation to demonstrate the FEC bandwidth overhead versus delay trade-off. Other FEC schemes might offer a different trade-off, but the concepts we present still generally apply.

In parity FEC, the sender-side encoder performs the exclusive-OR (XOR) operation on a source block of packets arranged in a matrix of D rows by L columns to generate redundant parity packets. The receiver-side decoder uses the parity packets generated by applying FEC across the packets in the same columns to recover from bursty losses and the parity packets generated by applying FEC across the packets in the same rows to recover from isolated losses.

The structure of the matrix that forms the source block impacts the loss burst size the FEC can protect against and the bandwidth overhead associated with the FEC data. The source block's size and the regime in which the parity packets are sent determine the delay the FEC introduces; hence, careful selection of D and L values is important. (A detailed analysis on the parity FEC scheme is available elsewhere.[7])

We consider a network deployment in which a broadcast video service uses a video source that's originating a single IP stream per broadcast channel, which is distributed across the core, aggregation, and access networks to receivers using IP multicast. This deployment uses either fast convergence or MPLS-TE FRR to reduce connectivity loss following reconvergence events, and it uses FEC end-to-end from the source to the receiver to recover from resulting packet losses or losses due to lower-layer errors. Such a deployment has the following characteristics.

QoE. FEC can potentially recover from losses due to both network-reconvergence events and lower-layer errors, resulting in a lossless video transport solution, assuming that service providers can configure the FEC parameters to match a particular service's requirements and taking into account the underlying network's characteristics. We generally apply FEC to the end-to-end stream, so it can protect against losses occurring across the core and aggregation networks, as well as across access lines.

Capacity provisioned and consumed. FEC incurs an overhead per protected stream. If B is the basic video stream bandwidth per channel, we must provision $(1 + O) \times B$ total bandwidth per channel on working- and failure-case paths to support the FEC streams' bandwidth for each channel FEC protects. O accounts for the FEC overhead, which is often determined based on the FEC code, the FEC latency budget, and the duration of packet loss that the FEC must protect against.

If we were to use FEC in conjunction with point-to-point MPLS-TE FRR, the FEC overhead might be compounded by the bandwidth inefficiencies we previously described for MPLS-TE FRR. As a result, we need to provision greater than $2 \times (1 + O) \times B$ bandwidth on failure-case paths for MPLS-TE's FRR capabilities to be effective.

Delay. The additional delay an FEC scheme incurs depends on the size of the FEC block (that is, the block containing the source as well as the FEC packets) and the time it takes to perform the FEC decoding operations.

Deployment considerations. Because FEC streams and media streams typically follow the same path, this approach places no explicit requirements on the network.

Applicability. To understand FEC's applicability, we consider the overhead (O) and delay impact from using FEC to cover the packet losses experienced for some specific scenarios:

- *Lower-layer errors.* Consider a scenario in which the FEC scheme must recover from losses resulting from lower-layer errors lasting for 20 ms. Assuming an impacted 3.75-Mbps MPEG2 video stream at 350 IP packets per second (pps), a 20-ms outage would result in the loss of seven IP packets. In the parity FEC scheme, we could protect against this much loss using an L value of at least 7. To keep the overhead at 20 percent, we use a D value of 5. Under the bursty transmission scheme,[7] we would need a latency budget of $(D + 1) \times L = 42$-packet duration, which would equal 120 ms. Larger D values would reduce the overhead at the cost of increasing delay.

- *Network-reconvergence events.* To recover from losses resulting from typical network-reconvergence events, the FEC scheme would need to be able to cover a loss of 300 ms. For the video stream we described previously, a 300-ms outage would result in the loss of 105 IP packets. We could protect against this using an L value of at least 105. To keep the delay the FEC introduces at an acceptable level, we would need to choose a small D value. If we choose a D value of 2, the overhead becomes 50 percent. Despite this high overhead, we would still need a latency budget of 315-packet duration, which would equal 900 ms. Clearly, FEC's overhead or delay will increase proportionally as the outage duration increases. Even if we use MPLS-TE FRR, and the packet loss following a network failure is only 50 ms, we would lose 18 IP packets; with $D = 2$ and $L = 18$, the resulting overhead would be 50 percent, and the latency budget would equal more than 100 ms. Hence, due to the resulting bandwidth overhead or latency impact, using FEC to recover from packet losses due to network failure events is unlikely to be viable in practice.

FEC's bandwidth and delay impact are likely to limit its applicability to providing protection for limited durations of contiguous packet loss — for example, spurious packet loss due to lower-layer errors that might be prevalent in the access last mile, rather than protecting against losses due to network-reconvergence events.

Temporal Redundancy

Techniques using temporal redundancy break a stream into blocks; the sender then transmits each block twice, separated in time. The block repetition pattern is such that within a time window greater than the block-separation period, the receiver should receive two copies of each packet, and will select one to use. If packet loss occurs, lasting for a duration less than the block separation period, the receiver should get at least one copy of each packet.

We consider a network deployment in which a broadcast video service uses a video source that's originating a single IP stream per broadcast channel that's distributed across the core, aggregation, and access networks to receivers using IP multicast. The deployment uses either fast convergence or MPLS-TE FRR to reduce the loss of connectivity following core reconvergence events and temporal redundancy to recover from losses experienced due to network-reconvergence events or lower-layer errors. Such a deployment has the following characteristics.

QoE. Temporal redundancy can cope with individual packet losses and losses due to network-element failures, resulting in a lossless video transport solution.

Capacity provisioned and consumed. Temporal redundancy incurs a 100 percent overhead per protected stream — that is, if B is the basic video stream bandwidth per channel, then we must provision $2 \times B$ bandwidth per channel on working- and failure-case paths to support temporally redundant streams for all channels.

If we used temporal redundancy in conjunction with point-to-point MPLS-TE FRR, temporal redundancy's overhead might be compounded by the bandwidth inefficiencies we previously described for MPLS-TE FRR, with the result that we must provision $\geq 4 \times B$ bandwidth on failure-case paths for MPLS-TE's FRR capabilities to be effective.

Delay. To be effective, temporal redundancy must incur a delay on the received stream at least equal to the block-separation period, which in turn must be at least the period of the loss in connectivity that the technique is aiming to protect against. This depends on whether the protection is protecting against losses due to lower-layer errors or network-reconvergence events and, if the latter, whether the service provider is using fast routing protocol convergence or MPLS-TE FRR.

Deployment considerations. Because the temporally redundant streams follow the same path, this approach places no explicit requirements on the network.

Applicability. Due to the 100 percent overhead temporal redundancy incurs, and to limited bandwidth available in the aggregation networks and access lines, this approach is likely to be viable only in core networks.

Spatial Redundancy

Techniques using spatial (physical) redundancy send two streams between the source and destination over different network paths. In normal working-case network conditions, the destination (or an intermediate point on the path to the destination) will effectively receive two copies of each packet, from which it will select one. If a network failure affects a transmitted stream, the destination or the intermediate point will continue receiving the other stream.

We consider a network deployment in which a source is originating two resilient multicast streams carrying the same content. A router or other intermediate system between the source and the last-mile access connection will receive both streams, compare the received packets from those streams, and forward on a single copy of each packet. We can use several techniques to instantiate the resilient streams. Such a deployment has the following characteristics.

QoE. Spatial redundancy can cope with individual packet losses and losses due to network-element failures, resulting in a lossless video transport solution. If a network failure affects one transmitted stream, the other will still be received.

Capacity provisioned and consumed. As we explained earlier, in IPTV deployments, the working-case paths for some receivers are often the failure-case paths for other receivers, so the resilient multicast streams are generally already established. In practice, this again means that the basic video stream bandwidth B must be provisioned per channel on working- and failure-case paths.

Delay. Spatial redundancy only introduces sufficient delay to receive the two streams — that is, the difference between the two paths' transmission delays.

Deployment considerations. For maximum service availability, a model employing spatial redundancy requires that the two video streams follow different network paths. We can ensure this using techniques we previously described such as MoFRR or MPLS-TE. The chosen technique impacts the resulting solution's complexity.

Applicability. Spatial redundancy doesn't normally use resilient access lines, so this approach is most applicable in core and aggregation networks.

Retransmission

Media streams that use the Real-Time Transport Protocol (RTP; www.ietf.org/rfc/rfc3550.txt) are to some extent resilient, in that receivers can use the mechanisms defined for the RTP Control Protocol (RTCP) to report packet reception statistics and thus let a sender adapt its transmission behavior. Additional techniques have been defined within the IETF that extend RTCP's basic capabilities to allow for faster feedback of packet loss (www.ietf.org/rfc/rfc4585.txt) and allow retransmissions (www.ietf.org/rfc/rfc4588.txt).

We consider a network deployment in which a broadcast video service uses a video source that's originating a single IP stream per broadcast channel, distributed across the core, aggregation, and access networks to receivers using IP multicast. The deployment uses either fast convergence or MPLS-TE FRR to reduce the loss of connectivity following reconvergence events and retransmission potentially to recover from losses experienced due to network-reconvergence events or lower-layer errors. Such a deployment has the following characteristics.

QoE. Retransmission can cope with individual packet losses, but losses due to network-element failures can result in substantial retransmissions and delay. When many receivers exist, doing multicast retransmissions might help.

Capacity provisioned and consumed. Real-time retransmission is a reactive scheme that resends only those packets that were lost, so it incurs a minimal bandwidth overhead.

Delay. This approach adds a delay equal to the sum of the loss-detection time and round-trip time (RTT) between the receiver and the retransmission server. To avoid increasing the channel-change times, this delay must remain small, limiting the scope of real-time retransmissions to cases in which the retransmission server is close enough to the receivers.

Table 1. Relative analysis of various deployment models.

Model	QoE	Bandwidth usage*	Network complexity	Delay impact	Potential applicability
Fast convergence (FC)	Lossy; < 1 group of pictures (GoP)	Working-case: B Failure-case: B	Low	Zero	Core and aggregation
MPLS Traffic Engineering Fast Reroute (MPLS-TE FRR)	Lossy; < 1 GoP	Working-case: B; failure-case: $\geq 2 \times B$	Medium	Zero	Core
Multicast Only Fast Reroute (MoFRR)	Lossy; < 1 GoP	Working-case: B; failure-case: B	Low	Zero	Core and aggregation
FC + Forward Error Correction (FEC)	Lossless	Working-case: $(1 + O) \times B$; failure-case: $(1 + O) \times B$	Low	High	To protect against losses due to lower-layer errors only
MPLS-TE FRR + FEC	Lossless	Working-case: $(1 + O) \times B$; failure-case: $\geq 2 \times (1 + O) \times B$	Medium	High	To protect against losses due to lower-layer errors only
FC + temporal redundancy (TR)	Lossless	Working-case: $2 \times B$; failure-case: $2 \times B$	Low	High	Core
MPLS-TE FRR + TR	Lossless	Working-case: $2 \times B$; failure-case: $\geq 2 \times B$	Medium	High	Core
MoFRR + spatial redundancy (SR)	Lossless	Working-case: B; failure-case: B	Low	Low	Core and aggregation
MPLS-TE + SR	Lossless	Working-case: B Failure-case: B	High	Low	Core and aggregation

*B = video stream bandwidth, O = FEC overhead

This drives the need for local retransmission servers closer to the receivers than the actual source.

Deployment considerations. This approach places no explicit requirements on the network.

Applicability. The scaling concerns associated with using retransmission to protect against losses due to network-element failures, and the requirement to constrain the RTT between the receivers and the retransmission server, limit retransmission's applicability to the aggregation and access networks.

The combinations of different network-reconvergence technologies and loss-recovery approaches result in a potentially confusing number of possible deployment models, which have different pros and cons and costs versus complexity trade-offs. Table 1 summarizes a relative analysis of some of the models service providers can use to minimize loss and therefore improve the quality of video transport through the network in the context of IP-based video streaming services such as IPTV and CATV.

Starting in early 2011, we've seen initial deployments of these technologies worldwide, throughout the large and local service provider networks. As the dominance of IP video traffic increases, we'll likely see more pervasive deployment of these technologies.

References

1. J. Greengrass, J. Evans, and A.C. Begen, "Not All Packets Are Equal, Part 1: Streaming Video Coding and SLA Requirements," *IEEE Internet Computing*, vol. 13, no. 1, 2009, pp. 70–75.
2. J. Greengrass, J. Evans, and A.C. Begen, "Not All Packets Are Equal, Part 2: The Impact of Network Packet Loss on Video Quality," *IEEE Internet Computing*, vol. 13, no. 2, 2009, pp. 74–82.
3. C. Filsfils and J. Evans, "Deploying DiffServ in Backbone Networks for Tight SLA Control," *IEEE Internet Computing*, vol. 9, no. 1, 2005, pp. 66–74.
4. D. Farinacci and C. Filsfils, "A Simple and Efficient 0(50 msec) Resilience Technology for IPTV," *Proc. North Am. Network Operators' Group Meeting* (NANOG42), NANOG, 2008; www.nanog.org/meetings/nanog42/presentations/DinoFilsfils_iptv.pdf.

5. Y. Wang and Q-F. Zhu, “Error Control and Concealment for Video Communication: A Review,” *Proc. IEEE*, vol. 86, May 1998, pp. 974–997.
6. D. Wu, Y.T. Hou, and Y.-Q. Zhang, “Transporting Real-Time Video over the Internet: Challenges and Approaches,” *Proc. IEEE*, vol. 88, Dec. 2000, pp. 1855–1877.
7. A.C. Begen, “Error Control for IPTV over xDSL Networks,” *Proc. 5th Ann. IEEE Consumer Communications and Networking Conf.* (CCNC), IEEE Press, 2008, pp. 632–637.

John Evans is a distinguished engineer within Cisco’s Development Organization, where he works on the definition of network architectures for service providers. His work focuses on IP routing and core IP/MPLS technologies, traffic engineering and management, quality of service, and video transport. Evans has an MSc in communications engineering from the University of Manchester. He’s a senior member of IEEE and a fellow of the Institution of Engineering and Technology (IET). Contact him at joevans@cisco.com.

Ali C. Begen is with the Video and Content Platforms Research and Advanced Development Group at Cisco. His interests include networked entertainment, Internet multimedia, transport protocols, and content distribution. Begen has a PhD in electrical and computer engineering from the Georgia Institute of Technology. He’s a member of IEEE and the ACM. Contact him at abegen@cisco.com.

Jason Greengrass is a technical lead engineer in Cisco’s Systems Development Unit. He specializes in the testing and deployment of triple-play network solutions, with a focus on IP-based video quality testing and analysis. He’s a Cisco Certified Internet Expert at Routing and Switching. Contact him at jasong@cisco.com.

Clarence Filsfils is a distinguished engineer at Cisco, where he focuses on IP core routing and capacity-management design. He has an engineering degree in computer science from the University of Liege, Belgium, and a business degree from the Solvay Business School, Brussels. Contact him at cf@cisco.com.

cn *Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.*