Adaptive Motion Control

Intelligent Packet Loss Recovery

LifeSize Communications White Paper

June 2010

Table of Contents

Executive Summary

Packet loss is becoming more common as more applications, services and systems use network resources and bandwidth. Even the slightest interruption of data transfer can detrimentally influence the quality of experience for internet protocol (IP) video conferences. Internet Service Providers (ISPs), infrastructure equipment manufacturers and video endpoint developers try to minimize packet loss with solutions ranging from dedicated network resources to forward error correction (FEC).

LifeSize Communications, a leading developer and manufacturer of high definition (HD) video communication systems, addresses packet loss with Adaptive Motion Control (AMC). AMC applies industry standard techniques such as bitrate downspeeding, packet marking and FEC to overcome network congestion and deliver an unparalleled experience in HD video communication. By answering packet loss intelligently with AMC, LifeSize equipment can provide a high quality of experience without network administrator assistance or unnecessary bandwidth allocation.

AMC provides uninterrupted video conferencing in network environments, where there is less than 5 percent packet loss. The end result is an unmatched video communication experience that leads to increased productivity.

Packet Loss

Packet loss, as the name implies, is the failure of one or more data packets to arrive at their destination. Loss is represented in a percentage value, where 5 percent means that 5 percent of the total transmitted data packets did not arrive at the intended destination.
Causes of packet loss range from internal local area network (LAN) to external wide area network (WAN) sources. Within an organization's LAN, problems can manifest from network equipment, video systems, and/or cabling. Externally, reasons for packet loss include excessive utilization of the organization's internet service provider (ISP), numerous hand-offs between ISPs, and/or conflicting data transmission protocols between the LAN and the WAN.

Packet loss can affect audio, video and/or data components, thereby affecting the entire quality of the video conference experience. Video is the most susceptible to packet loss, since it utilizes the majority of data. Video issues include blocking, distortion, smearing, flickering, audio/video synchronization, and freezing frames, as shown in *Figure 1.*

Figure 1. *Screen capture of video images with and without packet loss*

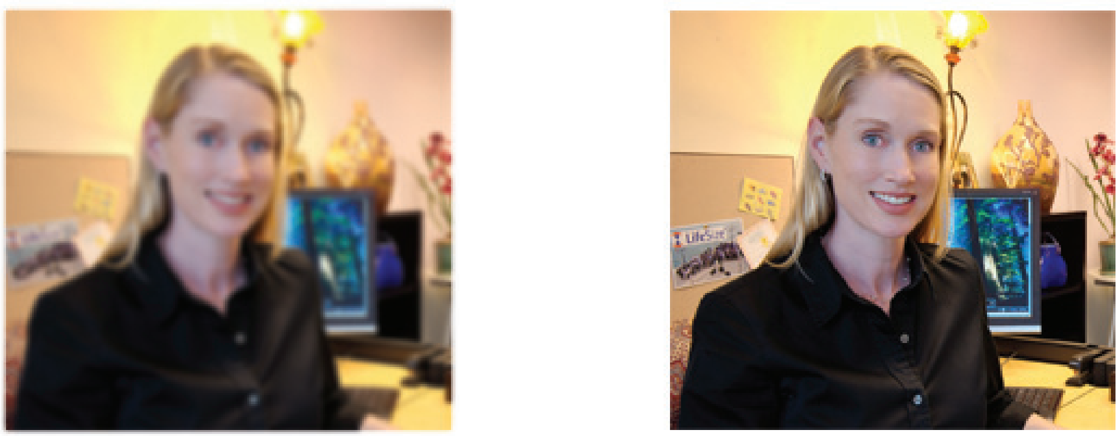

Image With Packet Loss | Image Without Packet Loss

While video issues can degrade the experience of a video call, audio issues are much less tolerable. With too many audio issues, the video conference experience breaks down. Users then typically resort to more traditional telephony devices.

Current Solutions

A general understanding of quality of service (QoS) is necessary in order to evaluate the current options for addressing packet loss. QoS is a suite of tools and methods that prioritize and distinguish data on internal and external networks. This suite is available to ISPs, infrastructure devices and video systems. When used in unison across products and networks to prioritize properly marked packets, QoS can deliver video without packet loss.

Network Provisioning

At the ISP level, organizations can influence QoS by allocating enough resource and bandwidth for all IP communications. The advantage with this approach is a dedicated pipeline for exchange of telephony, data and video between all of the organization's offices. The disadvantages are cost and availability.

Cost of bandwidth is a critical factor, since cost per megabit per second (Mbps) can vary greatly by location and by ISPs. Strict attention to total network usage is required in order to maximize the network benefit at the right cost. Overprovisioning the network does not provide a cost effective solution. For example, if one location needs 45 Mbps to compensate for maximum usage by all forms of IP communication, the location only experiences this maximum amount of traffic four or five days per month, on average. Customers can help minimize this issue by properly provisioning their networks for average utilization and employing other tools, discussed below, to help manage peak periods.

Customers can differentiate video communications by internal versus external traffic and provision their network accordingly. If a large portion is external communication, the customer is more dependent on the ISP and other partner or customer organizations deploying adequate bandwidth. This difference in resource allocation between multiple organizations can lead to packet loss when overutilization occurs, especially when calling into locations that have provisioned low bandwidth.

Reservation Protocol

To minimize bandwidth cost and the possibility of overutilization during high traffic periods, organizations can deploy a reservation protocol approach at the network level. Reservation protocol requires all network equipment to support an allocation provision that will reserve the necessary bandwidth and prioritize certain types of data packets.

In this reservation protocol scenario, the video equipment or a proxy will inform the network infrastructure that it needs a dedicated amount of bandwidth for a video conference. The network infrastructure will choose the best route, prioritize the video packets, and maintain the bandwidth throughout the video call. Upon completion of the conference, the network infrastructure will release the bandwidth "reservation" for usage by other data packets.

By implementing this solution, organizations can adopt a bandwidth strategy that mimics everyday usage better, while providing a high QoS for video conferencing. Deployment and implementation will require years of investment, with no guarantee of migration across multiple service providers. To illustrate, Cisco and Juniper are the only current network providers that support reservation protocols. The likelihood that Cisco or Juniper is available at all locations of business for an organization is very low. Therefore, the cost to an organization and its customers, partners and suppliers for adoption of such a video communication solution could be enormous.

Packet Marking and Queuing

Video endpoints are similar to other network infrastructure products. Within each video endpoint, a method exists to mark data packets emitting from that device. Industry predefined identifiers can label audio, data and video packets. According to Request for Comments (RFC) 4594 from the Internet Engineering Task Force (IETF), video communication systems can label their video data packet with an identifier of 34 for L3 classification.
(see Table 1).

Table 1: RFC 4594 Labeling Recommendation

Application	L3 Classification		IETF
	PHB	DSCP	RFC
Network Control	CS6	48	RFC 2474
VoIP Telephony	EF	46	RFC 3246
Call Signaling	CS5	40	RFC 2474
Multimedia Conferencing	AF41	34	RFC 2597
Real-Time Interactive	CS4	32	RFC 2474
Multimedia Streaming	AF31	26	RFC 2597
Broadcast Video	CS3	24	RFC 2474
Low-Latency Data	AF21	18	RFC 2597
OAM	CS2	16	RFC 2474
High-Throughput Data	AF11	10	RFC 2597
Best Effort	DF	0	RFC 2474
Low-Priority Data	CS1	8	RFC 3662

Table created from RFC 4594 guidelines from IETF website

To support prioritization of video packets within the organization's LAN, network administrators need to program routers, servers and switches to treat packets, marked with identifier 34, with greater urgency. Outside the LAN, the treatment of labeled packets depends upon the ISP. ISPs that support multiprotocol label switching (MPLS) can preserve the identifier of the video packet, prioritizing video over all other data packets. In contrast, ISPs that do not support MPLS will ignore or remove the identifier and treat video packets like other data packets, which can result in packet loss during periods of overutilization.

H.264 Truncating

The majority of current video conferencing solutions uses the H.263 or H.264 video codec. H.264 SVC, a derivative of the H.264 standard video codec, divides the video frame into several layered frames. The core frame contains the basic video image. Each additional frame enhances the video in terms of resolution, frame rate and quality. When the decoder receives all of the transmitted frames from the encoder, it reproduces the original transmitted image.

To illustrate, assume a H.264 SVC encoder divides a HD video image into 20 frames for transmission to a H.264 SVC decoder. Frame 1, the core frame, contains the basic video image. Frames 2 to 20, enhancement frames, contain information to increase the frame rate, resolution and quality of the core frame. In the event that one or more of the enhancement frames disappear as a result of packet loss, it is still possible to reassemble most of the image from the remaining frames. Without the missing enhancement frames, the reassembled image could suffer from decreased frame rates, reduced video resolution, or degraded image quality.

H.264 SVC can work in unison with any of the previously mentioned solutions for packet loss. The control lever for this solution is internal to the video conferencing equipment. The caveat, however, is this solution is proprietary and both endpoints, receiving and sending, need to support the same H.264 SVC codec. If this caveat is not met, then the video systems will not be able to communicate with one another.

Bitrate Downspeeding

Audio and video exhibit different qualities at different bitrates. For WMA, AAC or MP3 audio files, a 128 kilobit per second (Kbps) file will sound better than a 64 Kbps file. With video, certain bitrates are necessary to achieve HD resolutions at 1280x720 and 1920x1080. In addition to the resolution, the bitrate also influences the frame rate, which can range from 30 frames per second (fps) to 60 fps.

Controlling the bitrate of video calls resides within the endpoint. The endpoint detects network congestion or packet loss during the video conference. As packet loss increases, the endpoint starts to decrease the required bitrate in increments of 32 or 64 Kbps.

The advantage of this solution is the gradual minimization of network usage by the video conference. The disadvantage is the reduction in video resolution and/or frame rate of the conference. In some instances, bitrate downspeeding could result in standard definition (SD) quality video at 640x480 resolution or common intermediate format (CIF) quality video at 352x288 resolution. Downspeeding could also reduce frame rates to 15 fps or less.

Video conference equipment manufacturers determine the downspeeding methodology to incorporate into their video endpoints. Some manufacturers implement a downspeeding scheme, which supports a few resolutions with varying frame rates, ranging from 5 to 60 fps, at each resolution. The result is a constant resolution over a range of bitrates, while the frame rate increases incrementally with increasing bitrate. Other manufacturers prefer to implement a methodology that supports numerous resolutions, while maintaining a constant frame rate of 30 or 60fps. The outcome is a superior video conferencing experience, where motion remains constant and resolution correlates to bitrate.

Forward Error Correction (FEC)

Sometimes, the best defensive position is to go on the offensive. Forward Error Correction is a proactive approach to packet loss. FEC assumes that network congestion is always present. In addition, FEC assumes a certain percentage of packet loss is recoverable. Beyond this percentage, packet loss will be too great to recover.

FEC works by sending and receiving extra data packets during the video conference, regardless of network congestion or actual packet loss. These extra data packets can compensate for up to 5 percent packet loss.
FEC is a good solution that can work independently or in unison with other packet loss solutions; however, FEC usage does warrant some considerations. FEC will consume additional bandwidth, which could lead to further network congestion. In addition, the coding of FEC is proprietary to each video endpoint manufacturer, which requires organizations to deploy endpoints from only one supplier.

LifeSize's Adaptive Motion Control (AMC)

As a video endpoint manufacturer, LifeSize can control elements that directly relate to video systems. Furthermore, LifeSize can strongly recommend how organizations provision their networks or develop their network protocol procedures. Therefore, LifeSize has developed Adaptive Motion Control to help address the issue of packet loss.

AMC is LifeSize's solution, that combines standards-based and proprietary tools, to minimize video degradation from packet loss. Most features of AMC work in the background, and do not require extensive network administrator support. This allows local control, where features can be easily enabled or disabled, depending upon the conditions of the network.

Packet Marking

All LifeSize endpoints support packet marking, which allows organizations to configure endpoints to work in a dedicated or provisioned network. *(see Figure 3).*

Figure 3. *Screen capture of LifeSize Packet Marking user interface menu*

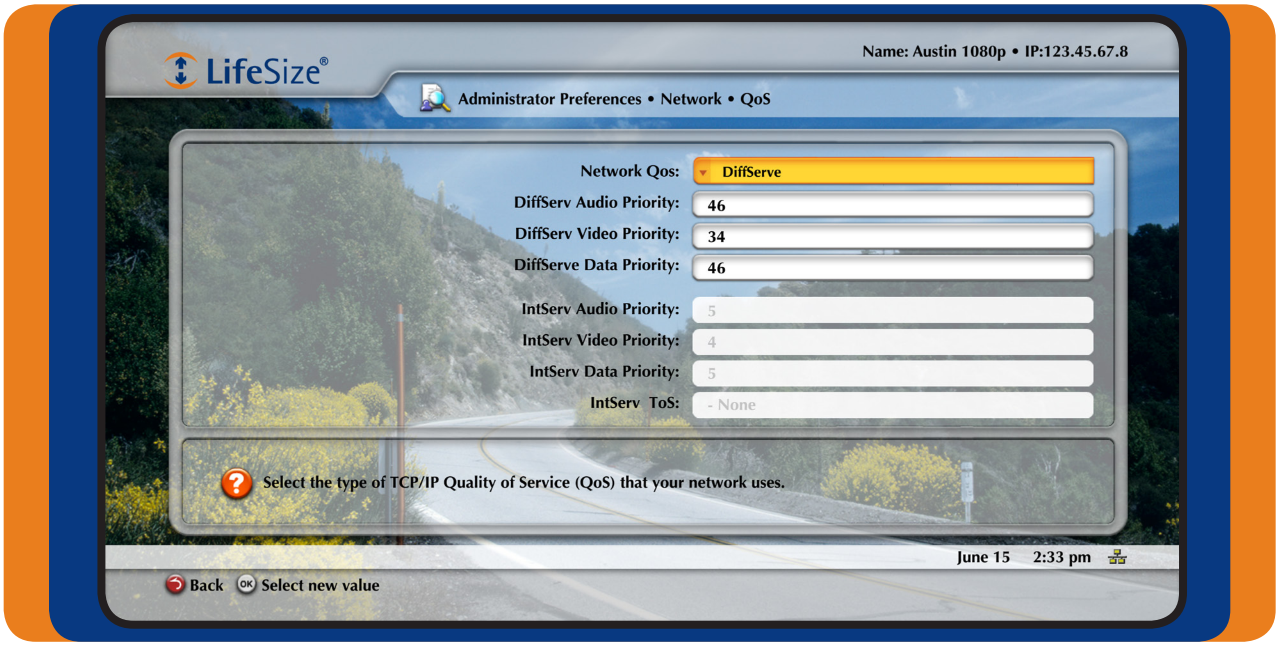

Bitrate Downspeeding

LifeSize supports bitrate downspeeding in all of its video endpoints. However, LifeSize's method of downspeeding is superior to that of other manufacturers. The key components behind LifeSize's bitrate downspeeding are the superior utilization of bandwidth and the total number of available resolutions.

Unlike other video systems, the majority of LifeSize endpoints can deliver HD quality video with a frame rate of 30fps at 768 Kbps. Most competitors require 1.1 Mbps or higher to achieve the same resolution and frame rate. This lower bandwidth utilization for HD video instantly places LifeSize at an advantage when networks experience congestion. *(see Chart 1).*

***Chart 1:** Bandwidth utilization and resolution at 30fps.*

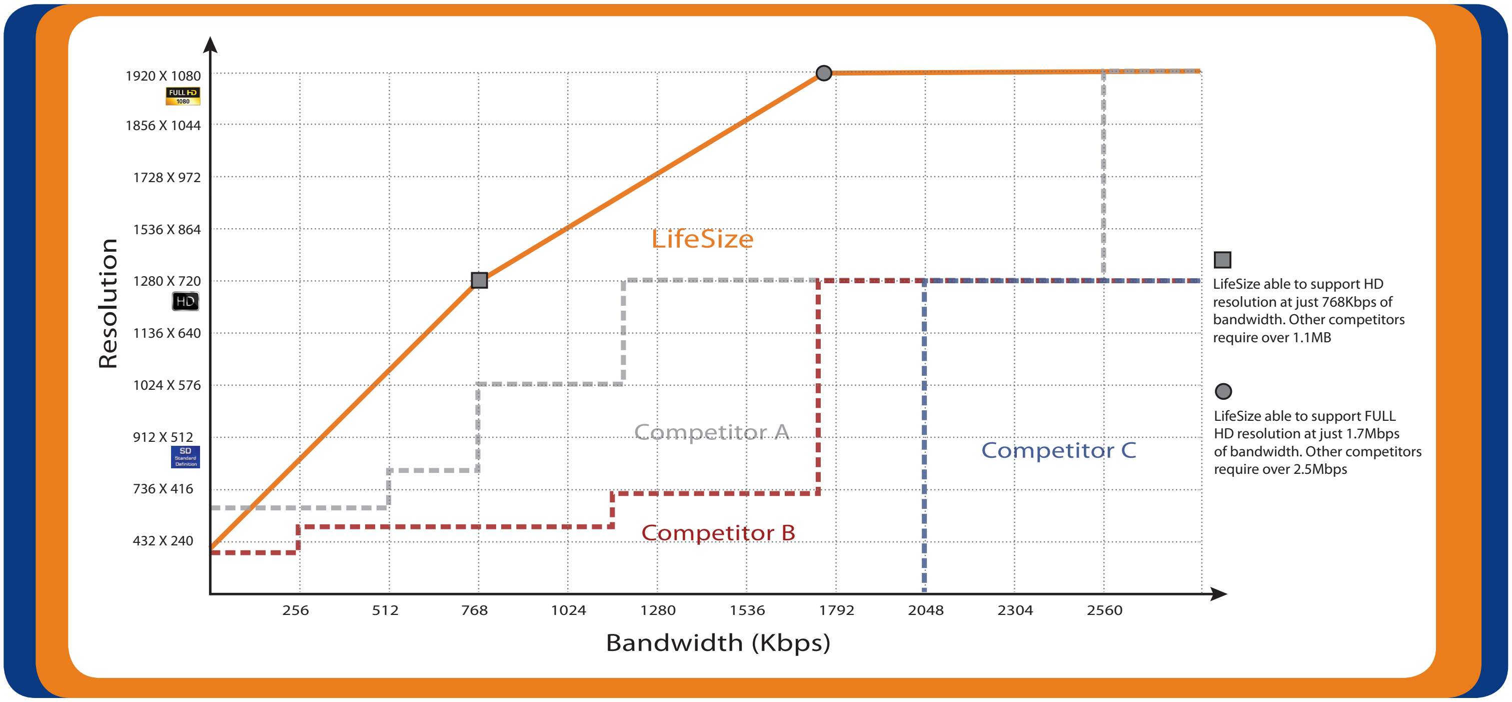

In the event that LifeSize does need to reduce its bitrate to support the video conference, LifeSize's approach to downspeeding always maintains the frame rate at 30 or 60 fps (depends upon selection of 720p60 or 1080p30 mode). This approach contradicts methods taken by LifeSize competitors. Most competitors will try to maintain the resolution, while reducing the frame rate. This method of downspeeding can lead to a stop-motion visual experience, where the HD resolution is moving at 10 or less frames per second.

LifeSize's competitors use this downspeeding technique because they only support 5 to 10 display resolutions. For some competitors, the downspeeding in resolution from HD, while maintaining a 30 fps frame rate, could lead to SD quality video communication. At 384 Kbps bandwidth, LifeSize is able to deliver DVD-quality standard resolution (848x480) at 30 fps, while the competition cannot even deliver SD quality video at 30 fps.

LifeSize makes downspeeding more user friendly by supporting over 200 resolutions, at frame rates of 30 fps or 60 fps.

For example, in a scenario where a HD video call at 1280x720 at 30fps uses 768 Kbps of bandwidth, the available bandwidth could then decrease to 512 Kbps due to network congestion. In response, the LifeSize system will decrease the resolution to only 1024x576 while maintaining 30 fps. In a similar situation with Competitor A, the initial video call using 768 Kbps bandwidth is at 1024x576 resolution (below HD video quality) at 30 fps. If the bandwidth decreases to 512 Kbps due to network congestion, Competitor A's system will decrease the resolution to only 768x448 (below DVD video quality) and maintain 30 fps.

Forward Error Correction

The main difference in LifeSize's approach to implementing FEC is its network intelligence. Most FEC solutions send extra data packets regardless of network condition. LifeSize's FEC is not a continuously "on" solution. LifeSize's FEC intelligently monitors the video conference for packet loss and jitter. When it detects packet loss, FEC will start to send extra data packets to compensate for the packet loss. When packet loss is no longer present, FEC will stop sending extra data packets.

LifeSize's approach to FEC minimizes network bandwidth overhead while maximizing user experience and productivity. The user does not need to think about activating or deactivating the FEC function of the endpoint. The user can simply leave FEC enabled and just make video calls. FEC will automatically initiate when it detects packet loss.

LifeSize Advantage

Network congestion and packet loss are both part of IP communications. Users will increase usage of IP communications and manufacturers will continue development of new networking devices. Business organizations must properly provision network resources for all forms of IP communication. Good planning and provisioning towards appropriate network allocation is the best approach; however, it will require time for implementation, money for infrastructure investment, and ISP agreements for network protocols.

With AMC, businesses can enjoy higher quality and more reliable video conferences and worry less about the network causing poor quality video experiences. The end result is higher user satisfaction and greater productivity for the organization.