

Polycom's Lost Packet Recovery (LPR) Capability

Polycom's Lost Packet Recovery (LPR) Capability

January 2008



Table of Contents

<i>Executive Summary</i>	1
<i>What is Packet Loss?</i>	2
<i>What Causes Packet Loss?</i>	2
<i>What Impact Does Packet Loss have on IP Videoconferencing?</i>	2
<i>What is LPR?</i>	3
<i>What are the Benefits of LPR?</i>	3
<i>How Does LPR Work?</i>	4
Lost Packet Recovery (LPR)	4
Polycom Dynamic Bandwidth Allocation (DBA)	5
Polycom Video Error Concealment (PVEC)	5
<i>How is LPR Different?</i>	5
<i>What is the Downside of Using LPR?</i>	6
<i>What Video Systems Support LPR Today?</i>	6
<i>Is LPR Proprietary?</i>	6
<i>Does LPR Work with Any Codec?</i>	6
<i>How Did Wainhouse Research Evaluate LPR?</i>	7
Part 1 – Standard Video Call Test	8
Part 2 – H.239 / Dual Stream Test	11
<i>Summary</i>	12
<i>About Wainhouse Research</i>	14
About the Author	14
<i>About Polycom</i>	14

List of Figures

Figure 1: Flowchart of the LPR Process	4
Figure 2: The LPR Test Environment	7
Figure 3: Impact of Packet Loss during 1472 kbps Video Call	9
Figure 4: Screenshot of 384 kbps Test Call w/5% Packet Loss - Without and With LPR	10
Figure 5: Screenshot of H.239 Test Call w/5% Packet Loss – Without and With LPR	11

Executive Summary

Even a small amount of network error / packet loss can have a dramatic impact on the user experience during an IP video call. For this reason, the ability to host successful video calls over “lossy” IP networks has been a challenge since the emergence of IP videoconferencing in the mid- to late-1990s. In recent years, the increased use of low-cost, shared, and over-subscribed network lines (e.g. DSL, cable, satellite, high congestion LANs and WANs, general Internet, etc.) and the use of higher call speeds (often required to support higher video resolutions) have made this problem even more critical.

In Q4 2007, Polycom asked Wainhouse Research (WR) to evaluate the performance of its newly released LPR (or Lost Packet Recovery) feature; an algorithm available on Polycom’s HDX video systems. LPR protects IP video calls from the impact of packet loss by temporarily allocating a small portion of the call bandwidth to be used for forward error correction (FEC). The theory behind LPR’s operation is that it’s better to provide a solid video call at a slightly lower call speed than a video call plagued by video and audio issues (blocking, freezing, hiccups, etc.).

As a part of this effort, WR researched and analyzed the methods used by currently available videoconferencing systems to conceal the effect of network errors. In addition, WR set up a number of video systems in our Atlanta test lab, placed a series

LPR’s ability to protect the video call experience from the effects of packet loss is exceptional.

of test calls including varying levels of packet loss, and observed the results. This allowed WR to assess each system’s ability to protect video calls from packet loss in a variety of conditions.

After placing several hundred test calls between various systems, WR concluded that LPR’s ability to protect the video call experience from the effects of packet loss – especially random / bursty packet loss - is exceptional. Although LPR is only available on Polycom HDX systems at this time, Polycom expects to add LPR to its VSX, RMX, and RSS product lines in the near future.

Particularly notable aspects of LPR include:

- 1) LPR’s performance is especially strong when tackling packet loss of 5% or less, making LPR well suited to address the vast majority of network situations.
- 2) While most of the video systems on the market today include some form of packet loss / error concealment capability, Polycom LPR is one of only two error protection schemes available today that uses forward error correction (FEC) to recover lost data.
- 3) One of LPR’s differentiators and strengths is that it protects all parts of the video call, including the audio, video, and content / H.239 channels, from packet loss.

What is Packet Loss?

Packet loss is the failure of one or more transmitted packets [of data] to arrive at their destination¹ and is typically cited as a percentage. For example, a 10% packet loss means that 10% of the packets transmitted over that time interval did not arrive at their destination.

What Causes Packet Loss?

Packet loss can be caused internally (within a company's LAN) or externally (over a company's WAN or within a network provider's core). Specific causes of packet loss include:

- Network congestion (over-utilization, over-subscription, hand-offs between carriers)
- Higher priority traffic blocking lower priority traffic
- Network equipment problems (failed switches, routers, etc.)
- Setting / configuration problems (10/100 Mbps mismatches, use of duplicate IP addresses, etc.)
- Video equipment issues
- Cabling issues (wiring issues, connectorization issues)

What Impact Does Packet Loss have on IP Videoconferencing?

The impact of packet loss on an IP video call varies widely and can include:

- Impact on video quality
 - Video tiling or blocking
 - Localized distortion (unclear areas of the image)
 - Video smearing
 - Frequent screen refreshes / flickering
 - Lip sync problems
 - Decreased frame rate
 - Frozen images
- Impact on audio quality
 - General audio distortion
 - Hiccups or chirping
 - Audio drop-outs
- Impact on content / presentation quality
 - Blurred or distorted slides
 - Slow page flips / screen refreshes
 - Frozen images

In addition, packet loss can also cause excessive latency (delay) and even call disconnections.

¹ Source: Searchnetworking.com Definitions (powered by WhatIs.com)

The impact of packet loss on an IP video call depends upon a number of factors including:

- the percent of packet loss
- the distribution of the packet loss over time
- the capabilities of the endpoints / devices participating in the video call

As one might expect, higher levels of packet loss usually have a more significant impact on the videoconferencing session.

What is LPR?

LPR (or Lost Packet Recovery) is a recently released Polycom algorithm designed to protect IP video calls from the impact of network packet loss. LPR uses forward error correction (FEC), an error recovery method in which the sending system adds redundant data to its outgoing data stream to allow the receiving system to detect and correct errors without having to ask the sending system to re-transmit the missing information. The ability to correct errors without having to wait for network transmissions makes FEC especially well suited for real-time communications such as television broadcasts, voice over IP (VoIP), and IP videoconferencing.

What are the Benefits of LPR?

LPR offers five key benefits:

1) Packet Loss Protection

LPR allows users to conduct high quality video calls over loss-prone IP networks (DSL, cable, satellite, high contention LANs / WANs, etc.) without suffering the effects of packet loss.

2) Compensation for Temporary Network Issues

LPR protects video calls from short-term network issues by temporarily adjusting the bit rate of the call in progress. Once the network trouble subsides, LPR automatically re-adjusts the bit rate to the highest possible level (based on the initial call speed / current bandwidth availability).

3) Decreased Latency / Delay

By allowing the receive system to reconstruct and recover lost data in real-time, LPR eliminates the need for the receive system to wait for possible missing, delayed, or mis-ordered information. This, in turn, reduces the jitter buffer and associated delay.

4) Resource Efficiency

LPR allows an organization to use all available bandwidth for its video calls. Users simply place their calls at the ideal call speed, and LPR will automatically adjust the bit rate to use all available bandwidth. This is especially valuable for calls placed over network links with varying capacity (shared links, lines with significant congestion).

5) Total Media Protection

LPR's protection includes all elements of the videoconference call; voice, video, and content.

How Does LPR Work?

Polycom video endpoints monitor all active calls for incoming packet loss. Once packet loss is detected, the endpoint uses three tools to protect the user experience; Lost Packet Recovery (LPR), Dynamic Bandwidth Allocation (DBA), and if LPR is not supported, Polycom Video Error Concealment (PVEC).

Lost Packet Recovery (LPR)

Unlike most error concealment / avoidance algorithms (including Polycom's own PVEC) that involve only the system receiving the packet loss (referred to as the "receiving" or "receive" system below), LPR involves both video systems in the call.

As shown in the flowchart below, LPR works by temporarily allocating a portion of the call bandwidth into a data channel (which we will refer to as the FEC channel within this document) used for sending FEC data to the receive system. Through an iterative process, LPR increases and decreases the size of the FEC data channel until it finds the minimum bandwidth that must be allocated to the FEC data channel to allow the receive system to recover all lost packets.

While FEC is in use (meaning that the size of the FEC data channel is more than 0 kbps), the systems repeatedly test whether they can decrease the size of the FEC data channel and thereby make a greater portion of the call bandwidth available for the audio, video, and content data. Hence, LPR uses bandwidth only when packet loss is detected. This makes LPR ideally suited for environments with random or bursts of packet loss (like the Internet).

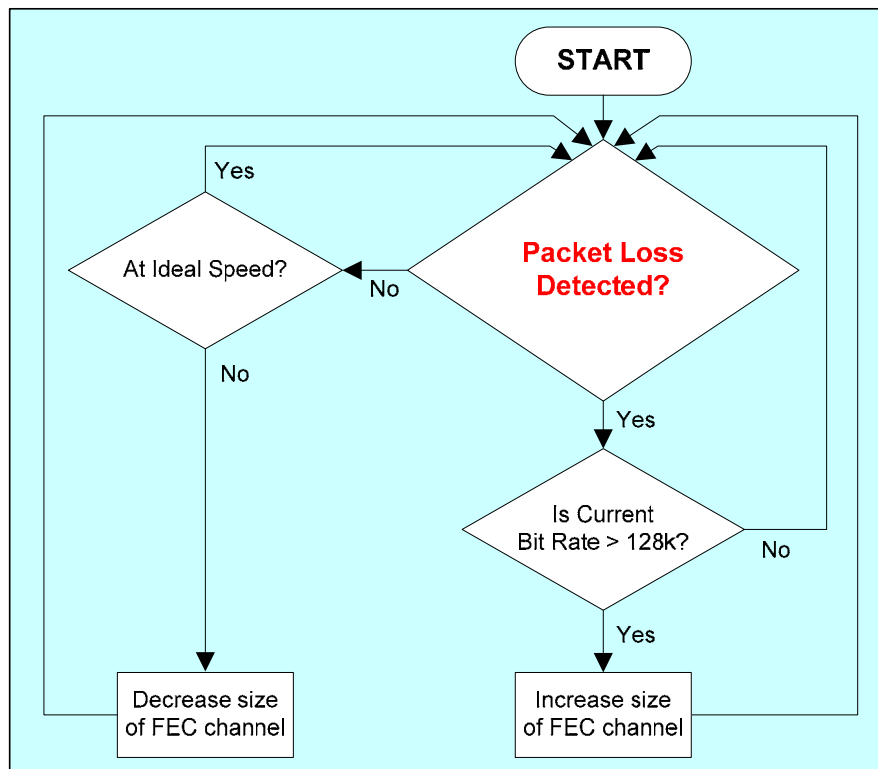


Figure 1: Flowchart of the LPR Process

Polycom Dynamic Bandwidth Allocation (DBA)

For environments or situations where consistent packet loss is noted, Polycom's Dynamic Bandwidth Allocation (DBA) capability is used in conjunction with LPR. DBA is an algorithm that automatically and dynamically adjusts the video bit rate during a video call in order to eliminate / avoid packet loss.

For example, if a consistent 10% packet loss was detected during a 384 kbps video call (320 kbps used for video, 64 kbps used for audio), DBA would decrease the video bit rate by roughly 10%, from 320 kbps to ~ 288 kbps, and then re-sample the signal to see if any packet loss remains. If necessary, DBA would drop the video bit rate further until no packet loss is detected for a few seconds.

If, after downspeeding, DBA determines that the packet loss was only temporary (less than a few seconds in duration), DBA will gradually increase the video bit rate. However, if DBA continues to detect packet loss, it will no longer adjust the video bit rate and will allow the call to proceed. This makes DBA well suited for situations involving packet loss caused by network overutilization or throughput-related issues.

Polycom Video Error Concealment (PVEC)

For calls when both participating systems do not support LPR, Polycom's Video Error Concealment (PVEC) is used. PVEC is an algorithm for IP video quality of service that compensates for the effects of packet loss by using packet information from neighboring macroblocks, prior frames, and future frames to estimate the content of the current video frame. Unlike LPR which recovers lost packets and DBA which avoids packet loss, PVEC works to mask the impact of packet loss.

How is LPR Different?

LPR differs from most packet loss concealment algorithms in three ways:

1) Packet Recovery vs. Packet Loss Masking

Most other solutions either try to hide the impact of the packet loss (error concealment) or avoid packet loss by decreasing the video bit rate (downspeeding). LPR, however, takes things one step further by actually recovering lost data packets.

2) Self-Healing Feature

Many competing solutions decrease the bit rate to avoid / decrease the impact of packet loss, but most do not increase the bit rate once the packet loss subsides. LPR, on the other hand, returns the call bit rate back to its original value once packet loss disappears.

3) Full Coverage

Most competing systems' error concealment schemes protect only the video channel, while LPR covers all three parts (voice, video, and content) of the video call.

What is the Downside of Using LPR?

The idea behind LPR is to make acceptable compromises to keep packet loss from severely impacting the video call experience. While our testing showed that these compromises are minor or not noticeable in most call situations, the tradeoffs include:

1) Bandwidth Compromise

LPR's use of bandwidth for the FEC channel means that less bandwidth is available for the video call itself. For low bandwidth connections (384 kbps or lower), this may cause a temporary decrease in frame rate or video resolution.

2) Latency Increase

When using LPR, the send system must delay the transmission of the outgoing audio, video, and content data until it can complete the calculations for the FEC channel. Similarly, the receive system must process the incoming FEC channel and perform any necessary packet recovery processing. Depending upon the situation, this may add a few milliseconds of latency to the video call. That said, LPR can also decrease call latency in some circumstances since when LPR is in use, the receive system no longer needs to wait for delayed / lost information to arrive.

3) Additional Processing Burden

The processing required for LPR adds some additional burden on each of the participating video systems. Assuming the systems have adequate processing power, which is likely to be the case for any relatively current video system, this is a non-issue.

Considering the alternative (a flawed or even useless video call), WR believes that the compromises associated with LPR – even in the extreme – represent a very favorable trade.

What Video Systems Support LPR Today?

As of this writing, LPR is supported on all Polycom HDX video systems and solutions based on the HDX including Polycom's RealPresence™ Experience (RPX™) and Telepresence Experience (TPX™) solutions. Polycom expects to add LPR support to the entire VSX product line and its RSS 2000 video capture system in Q2 2008, with LPR support on the RMX 2000 expected to follow soon after. Once implemented, LPR will also protect calls between unlike LPR-capable systems (e.g. HDX and VSX).

Is LPR Proprietary?

YES – At this time LPR is a proprietary, Polycom-only capability. However, Polycom plans to submit LPR to the appropriate standards-bodies (IETF and/or ITU) in the future.

Does LPR Work with Any Codec?

YES – Because LPR is implemented at the communications / data packet level, it is codec agnostic. LPR will work with the codecs most commonly used during video calls including H.264, H.263, G.722, G.722.1C / Siren14 and more. In addition, it can be used with codecs that will be introduced in the future such as the H.265 video protocol already being discussed within the industry.

How Did Wainhouse Research Evaluate LPR?

The first step of the LPR evaluation was defining our expectations for network error compensation systems overall (described below).

General Functionality

- Protect the user experience for packet loss levels matching those found in a typical network
- Compensate for both constant packet loss and random / bursty packet loss situations
- Protect all media flows (video, audio, content) within a videoconferencing session
- Prioritize the protection of audio streams over video streams

Methodology / Implementation

- Operate automatically without needing the user to turn the capability on or off.
- Operate transparently without impacting the overall system performance or disabling other features (supported video resolution, media encryption, etc.)
- Select the appropriate compensation method to use for each situation (eg. downspeeding, forward error correction, error concealment, etc.).
- Self-heal whenever possible (for example, when the packet loss disappears) by returning the call to its original state (bit rate used, video resolutions, etc.).
- Recognize when it cannot improve the experience and avoid making the situation worse
- Be codec and system agnostic

As described below, LPR met – and in many cases exceeded – WR’s expectations.

To evaluate LPR’s performance, WR set up a variety of SD and HD videoconferencing systems from numerous vendors in our Atlanta test lab. We then placed various test calls between like systems (e.g. Polycom HDX to Polycom HDX) over our Local Area Network.

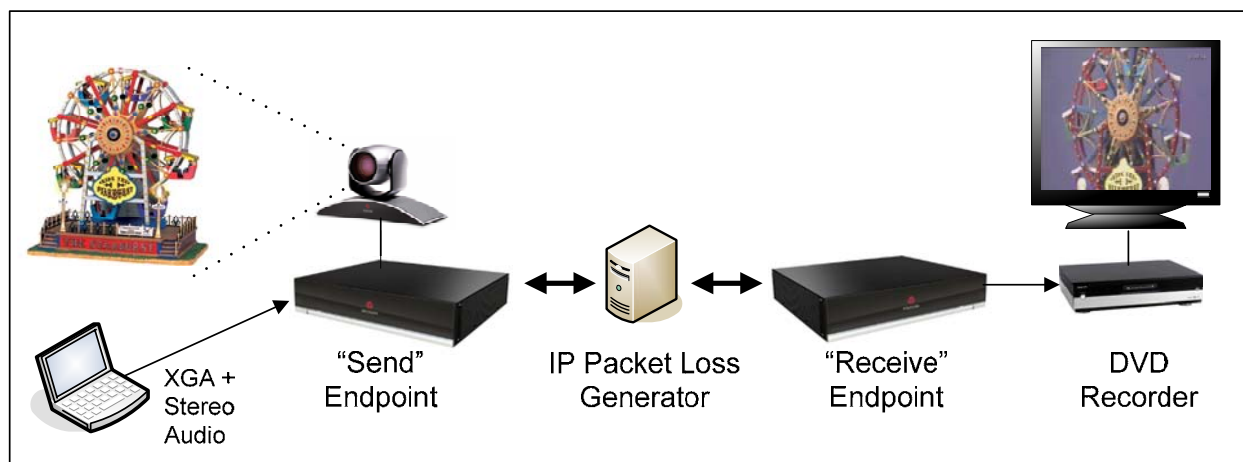


Figure 2: The LPR Test Environment

It is worth noting that the challenges faced by the test endpoints in the WR test lab were more aggressive than those found in the typical enterprise in a variety of ways including:

Packet Loss Type - A software-based packet loss generator (running on a notebook PC) was used to inject packet loss into the test calls. The packet loss used during the testing was constant (e.g. 1%, 2%, etc. of the connection speed), regardless of the call rate / bit rate used.

Packet Loss Percent – For this evaluation, packet loss percentages of 1%, 2%, 5%, and 10% were used. Considering that most high quality networks boast packet loss guarantees of 0.01% or less and that the average packet loss noted on the public Internet is less than 5%², testing at 10% was certainly aggressive.

A software-based packet loss generator (running on a notebook PC) was used to inject packet loss into the test calls. The packet loss used during the testing was constant (e.g. 1%, 2%, etc. of the connection speed), regardless of the call rate / bit rate used.

Motion of the Camera Image - To ensure the consistency of the test environment, all system cameras were adjusted to capture the image of a motorized carousel. This means that ~ 80% (or more) of the image was changing at all times.

H.239 Input Signal – For the H.239 testing, WR utilized an IBM X40 ThinkPad (providing an XGA / 1024x768 resolution signal with a 4:3 aspect ratio) playing a video clip of the trailer for the movie Matrix Revolutions. This video clip clearly included a greater degree of motion than the typical H.239 content.

As with all WR evaluations, unless otherwise stated all systems settings were set for default or auto. For this evaluation, all video systems were set to favor motion over sharpness. As shown in the drawing above, to ensure the consistency of the test environment, the audio signals used during the testing were injected into the audio inputs of each “send” system (in lieu of acoustically capturing the audio using the system microphones).

Part 1 – Standard Video Call Test

For the standard video call testing, WR placed a series of video calls at different speeds (384 kbps, 768 kbps, 1472 kbps, and 4 Mbps) between like systems. During each call, the test team inserted increasing levels of packet loss (1%, 2%, 5%, and finally 10%) into the test call as instructed / cued by a pre-recorded audio clip and documented the results.

² Source: Internet Traffic Report (www.internettrafficreport.com)

Results of the Lost Packet Recovery (LPR) Testing

For this part of the evaluation, WR conducted three rounds of testing as follows:

1) LPR / PVEC disabled and DBA enabled (baseline testing)

As expected, when LPR / PVEC was disabled, the impact of the injected packet loss was obvious and included video artifacts (pixelization, smearing, etc.) and audio issues (chirps, hiccups, distortion). Even at 1% packet loss, the user experience was significantly impacted.

The impact of DBA was noted during this round of testing, but since the packet loss injected was constant (the same percentage regardless of the bit rate / call speed used), DBA's downspeeding did not eliminate the packet loss.

2) LPR / PVEC enabled and DBA disabled

Throughout these test calls, LPR's performance was exceptional in that it totally or substantially eliminated the effect of the injected packet loss on the user experience. Specific results include:

- With 1% and 2% packet loss injected, LPR completely eliminated the effect of the injected packet loss and yielded a "no-packet-loss" call experience.
- At 5% packet loss, the impact of the packet loss was minor (slight audio hiccups and a bit of video smearing during the 4Mbps test call).
- Even at 10% packet loss, LPR provided a user experience that WR believes the typical user would consider acceptable.

Even with 10% packet loss inserted, LPR provided an acceptable user experience.

As expected based on LPR's method of operation, when packet loss was injected, the bit rate of the video channel did not change, but the bit rate used within the video channel decreased. Once the packet loss generator was disabled, the video bit rate used returned to its initial value. This verified the self-healing feature of LPR.

3) LPR / PVEC and DBA enabled

With DBA enabled, LPR's performance was as strong as during the prior round of testing. As expected based on DBA's method of operation, when packet loss was injected, the bit rate of the video channel decreased.

Packet Loss	Video Resolution	Video Rate Available (kbps)	Video Rate Used (kbps)	Frame Rate	Audio Protocol / Bit Rate (kbps)	Comment
0%	4SIF	1408	1050	30	Siren22 @ 64 kbps	Video / audio fine.
1%	4SIF	1024	800	30	Siren22 @ 64 kbps	Video / audio fine.
2%	4SIF	1024	800	30	Siren22 @ 64 kbps	Video / audio fine.
5%	4SIF	1024	800	30	Siren22 @ 64 kbps	Video fine. VERY minor audio hiccups (every 10 – 20 seconds)
10%	4SIF	1024	800	28 – 30	Siren22 @ 64 kbps	Video artifacts (smearing every few secs) and i-frames noted. Minor audio hiccups.

Figure 3: Impact of Packet Loss during 1472 kbps Video Call

The screenshots below show the impact of 5% packet loss during a 384 kbps video call both with and without LPR. Note the smearing of the video image on the left (without LPR / PVEC). As shown on the right, the combination of LPR, PVEC, and DBA totally eliminated the effects of the injected packet loss. Note that the image distortion during the actual video call was more pronounced than represented by the static image below left.




Figure 4: Screenshot of 384 kbps Test Call w/5% Packet Loss - Without and With LPR

Test Results Using Other Leading Video Systems

The error concealment performance of the other leading systems ranged from marginal to poor. Specific results included:

- Even with 1% packet loss inserted, video artifacts were observed
- As packet loss increased, the impact on the video signal became more apparent and included video tiling, smearing, and decreased frame rate.
- With 5% and 10% packet loss injected, the constant video smearing and screen refreshes (due to frequent i-frame requests) made the video images unusable.
- At 5% and 10% packet loss, audio performance was impacted on some systems.

During the testing, WR noted that many systems employed some form of downspeeding (similar to Polycom's DBA) to eliminate packet loss. Since the packet loss injected during our testing was constant, decreasing the video bit rate did not eliminate the packet loss.

Part 2 – H.239 / Dual Stream Test

For this part of the testing, the test team placed a video call at 1472 kbps between like systems, connected the H.239 signal source, activated H.239, and set the receive system to use a side-by-side on-screen layout showing both the incoming camera image and H.239 signal. The test team then inserted varying levels of packet loss into the call as instructed by the audio clip and documented the results.

Results of the Lost Packet Recovery (LPR) Testing

LPR's strong performance continued during the H.239 testing.

- At 1% packet loss, LPR completely protected the audio, primary video, and content video from the packet loss.
- At 2% and 5% packet loss, the only artifact noted was a slight decrease in the frame rate of the primary video channel.
- At 10% packet loss, video artifacts were noted on the primary video channel and pixelization was noted on the content / H.239 channel. However, LPR was able to eliminate / mask enough of the packet loss to provide an acceptable user experience.




Figure 5: Screenshot of H.239 Test Call w/5% Packet Loss – Without and With LPR

Test Results Using Other Leading Video Systems

As with the prior round of testing, the error concealment performance of the other leading systems during the H.239 test ranged from marginal to poor.

- Even with only 1% packet loss inserted, the effects of the of the packet loss on both the primary video and H.239 video channels were immediately apparent
- At packet loss levels above 1%, the video call experience was unacceptable.

Part 3 – One Hour Test

For this part of the testing, the test team placed a video call at 1472 kbps between like systems, injected 10% packet loss, and left the call standing for a one-hour period. At 15 minute intervals, the test team assessed the quality of the video call in progress and documented the call statistics (call rate, actual bit rates, video resolution used, etc.).

Results of the One Hour Test Using Lost Packet Recovery (LPR)

- The combination of LPR and PVEC was able to eliminate / mask most, but not all of the effects caused by the insertion of 10% packet loss.
- Remaining artifacts included video smearing every few seconds and screen refreshes every 5 to 10 seconds.
- Although the user experience would not be perfect, WR believes a talking-head video call could still be successful under this condition.
- Most importantly, the user experience did NOT degrade over time; the performance at the 1 minute mark was the same as that at the 60 minute mark. This is a clear differentiator for LPR.

Results of the One Hour Test Using Other Leading Video Systems

- Immediately after the 10% packet loss was injected, all of the test video calls became unusable due to massive video smearing and frequent screen refreshes.
- As time progressed, the user experience continued to degrade until a lowest possible performance threshold / video bit rate was reached.

For example, one system used 1300 kbps for the video channel at the start of the test call. Immediately after the 10% packet loss was applied, the video bit rate used dropped to 958 kbps. By the 15 minute mark, the video bit rate in use had dropped to only 56 kbps; a level too low to support a full motion video call at any reasonable video resolution.

Overall, none of the non-LPR capable systems were able to provide an acceptable user experience while receiving 10% packet loss.

Summary

Polycom's Lost Packet Recovery algorithm provides a significant benefit for IP videoconferencing users, delivering high quality audio and video performance in lossy network environments. The benefit

afforded by LPR will increase exponentially as the number of video devices (endpoints, bridges, gateways, etc.) supporting LPR grows.

About Wainhouse Research

Wainhouse Research (www.wainhouse.com) is an independent market research firm that focuses on critical issues in rich media communications and conferencing. The company conducts multi-client and custom research studies, consults with end users on key implementation issues, publishes white papers and market statistics, and delivers public and private seminars as well as speaker presentations at industry group meetings. Wainhouse Research publishes Conferencing Markets & Strategies, a three-volume study that details the current market trends and major vendor strategies in the multimedia networking infrastructure, endpoints, and services markets, as well as a variety of segment reports, the free newsletter The Wainhouse Research Bulletin, and the PLATINUM (www.wrplatinum.com) content website.

About the Author

Ira M. Weinstein is a Senior Analyst and Partner at Wainhouse Research, and a 15-year veteran of the conferencing, collaboration and audio-visual industries. Prior to joining Wainhouse Research, Ira was the VP of Marketing and Business Development at IVCi, managed a technology consulting company, and ran the global conferencing department for a Fortune 50 investment bank. Ira's current focus includes IP video conferencing, network service providers, global management systems, scheduling and automation platforms, ROI and technology justification programs, and audio-visual integration. Mr. Weinstein holds a B.S. in Engineering from Lehigh University and can be reached at iweinstein@wainhouse.com.

About Polycom

Polycom, Inc. is the worldwide leader in unified collaborative communications (UCC) that maximize the efficiency and productivity of people and organizations by integrating the broadest array of high definition video, wired and wireless voice, and content solutions to deliver the ultimate collaborative experience. Polycom's high quality, standards-based conferencing and collaboration solutions are easy to deploy and manage, as well as intuitive to use. Supported by an open architecture, they integrate seamlessly with leading telephony, workplace wireless telephony, and presence-based networks. With its market-driving technologies, best-in-class products, alliance partnerships, and world-class service, Polycom is the smart choice for organizations seeking proven solutions and a competitive advantage from on-demand communications and collaboration. For additional information, call 800-POLYCOM or visit the Polycom web site at www.Polycom.com.