

TANDBERG and Packet Loss

TANDBERG

D50165, Rev 2.0

Table of Contents

1.	INTRODUCTION	3
2.	QUALITY OF SERVICE	3
2.1	RSVP (IETF RFC 2205)	3
2.2	DIFFERENTIAL SERVICES	4
2.3	IP PRECEDENCE AND TYPE OF SERVICE	5
3.	VIDEO COMPRESSION	6
4.	INTELLIGENT PACKET LOSS RECOVERY (IPLR^{TF})	6
4.1	PACKET LOSS HANDLING IN THE RECEIVER	7
4.2	PACKET LOSS HANDLING IN THE TRANSMITTER	7
5.	DOWNSPEEDING	7
6.	APPENDIX	8
6.1	POLYCOM VIDEO ERROR CONCEALMENT (PVEC).....	8
6.2	TANDBERG IPLR vs. POLYCOM PVEC.....	8
6.3	CONSIDERATIONS AROUND H.264	9

1. Introduction

This paper is intended to describe the packet loss robustness features available with TANDBERG's products. These features include :

- Intelligent Packet Loss Recovery^{TF} (IPLR^{TF})
- Downspeeding
- Quality of Service (QoS).

The above mentioned features are all available on the following products:

TANDBERG MCU 16+16 / 8+8

TANDBERG GW

TANDBERG 8000

TANDBERG 7000

TANDBERG 6000

TANDBERG 2500

TANDBERG 1000

TANDBERG 880

TANDBERG 800

TANDBERG 550

TANDBERG 500

The IPLR feature is not intended to replace Quality of Service (QoS) techniques which may be employed on IP networks to reserve bandwidth or set precedence on video traffic. However, in the rare event that QoS techniques fail, IPLR is designed to overcome transient packet loss situations and reduce the impact these losses may have on the current video traffic.

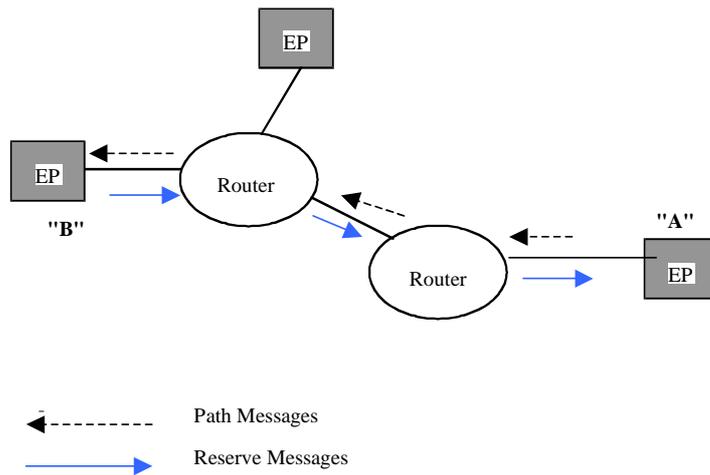
To fully understand what IPLR offers, an understanding of QoS and video compression are recommended. This paper briefly describes each before describing the details of TANDBERG's IPLR.

2. Quality of Service

TANDBERG has implemented a suite of QoS techniques into their endpoint products. Currently, TANDBERG has the most complete QoS offering on the market. These techniques range from reservation protocols such as RSVP to prioritization schemes such as DSCP and IP Precedence.

2.1 RSVP (IETF RFC 2205)

The RSVP standard is used by an endpoint to request certain qualities from the network that will transport the video and audio data. This request is made at each node throughout the network and each node must comply or the reservation will fail.

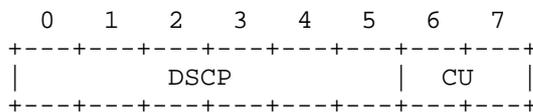


RSVP is a protocol like TCP and UDP. It has no concept of ports and hence there must be a suitable mechanism in the firewall to enable RSVP traffic. To enable RSVP you need to find the RSVP setting and turn it on. Consult your firewall's technical manual to determine if there are separate settings for incoming and outgoing RSVP traffic

2.2 Differential Services

Differential Services is another method of QoS offered by TANDBERG that utilizes 6 bits of the Type of Services Byte. This method is replacing IP Precedence as the preferred method for setting priority of packet traffic.

From the RFC:



DSCP: differentiated services codepoint

CU: currently unused

The DSCP settings have a range of 0-63 and are configurable by the user or hard coded depending on the endpoints being used. The TANDBERG endpoints offer independent configuration of the values for control, video and audio.

2.3 IP Precedence and Type of Service

TANDBERG has implemented IP Precedence and Type of Service (TOS). IP Precedence allows the video terminal to prioritize its video and audio above or below other IP traffic on the same network.

From the RFC:

```
Bits 0-2:  Precedence.
Bit   3:  0 = Normal Delay,      1 = Low Delay.
Bits  4:  0 = Normal Throughput, 1 = High Throughput.
Bits  5:  0 = Normal Reliability, 1 = High Reliability.
Bit  6-7:  Reserved for Future Use.
```

0	1	2	3	4	5	6	7
PRECEDENCE			D	T	R	0	0

Precedence

```
111 - Network Control
110 - Internetwork Control
101 - CRITIC/ECP
100 - Flash Override
011 - Flash
010 - Immediate
001 - Priority
000 - Routine
```

The TOS values available are:

- maximum throughput
- minimum monetary cost
- minimum delay
- maximum reliability
- normal (off).

This TOS feature is used to allow routers to make decisions on how to handle traffic if congestion occurs.

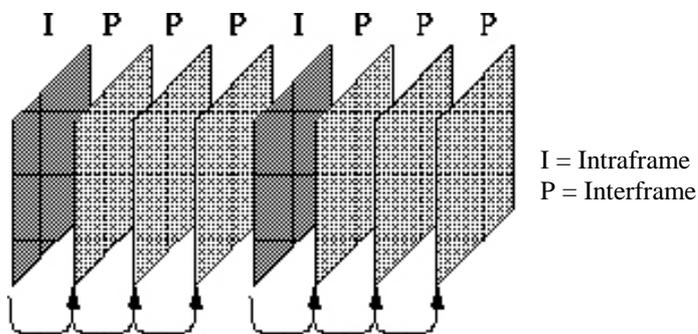
The QoS capabilities of the TANDBERG systems can be used to manage a private LAN or WAN more effectively. IP Precedence and TOS help prioritize IP traffic, control congestion, and allow for a better integration of video conferencing into an existing LAN. In times of high network stress, these features will help the routers to drop lower precedence data to increase router stability and offer higher probability that the video conferencing data will reach its destination. In some instances it may help provide a larger proportion of the link without having to configure policy. In the near future Differential Services may be available from ISP's and ASP's that will use the TOS information to offer better services to the users.

3. Video Compression

The two most common video compression algorithms used in video conferencing today are H.261 and H.263.

Interframe compression is compression applied to a sequence of video frames, rather than a single image. In general, relatively little changes from one video frame to the next. Interframe compression exploits the similarities between successive frames, known as *temporal redundancy*, to reduce the volume of data required to describe the sequence.

Because interframes require information from previous frames, the loss of interframes can be devastating to the process of recompiling the video image. Interframes require less information to transmit than intraframes.



Intraframe compression is compression applied to still images, such as photographs and diagrams, and exploits the redundancy within the image, known as *spatial redundancy*. Intraframe compression techniques can be applied to individual frames of a video sequence.

Because intraframes contain all of the information needed in a single frame and do not require information from previous frames, the loss of an intraframe is not as critical. However, to transmit all frames as intra would greatly reduce the efficiency of the compression as the amount of information needed to send an intraframe is much higher than an interframe.

To further increase flexibility, each frame can be further divided into blocks, where each block can be an intrablock or interblock.

4. Intelligent Packet Loss Recovery (IPLR^{TF})

The IPLR feature makes use of several components that address specific issues associated with packet loss. IPLR has the ability to improve the robustness of the transmitted bit-stream, the received image and to downspeed the call rate if excessive packet loss is detected. All of these components are compliant with existing ITU-T standards.

IPLR is the first packet loss recovery technique that is completely standards based and works with both H.261, H.263, and H.264 video algorithms.

4.1 Packet loss handling in the receiver

In the event packet loss is detected on the received image, the decoder will attempt to estimate what information has been lost due to packet loss. While the information that was lost cannot be reproduced, the effects of the lost information can be minimized. This feature is always working in the TANDBERG codec, continuously trying to improve the image whenever a packet loss occurs.

4.2 Packet loss handling in the transmitter

In the event frequent packet loss is detected (a level of approximately 1-2%), the encoder will enter a 'robustness mode'. In this event the encoder will begin transmitting intrablocks, which do not require previous information to build the image. Because intraframes can dramatically reduce the quality of the visual image and framerate, IPLR prefers to send intrablocks distributed over several frames in such a way that the image is completely rebuilt over several frames. In addition, the codec uses an intelligent routine to decide when to use intrablocks vs. interblocks to ensure the best possible picture. This provides the most efficient trade off between framerate/resolution and packet loss robustness during packet loss situations. These routines are compliant with the ITU-T H.261, H.263, and H.264 standards and are completely interoperable with any manufacturer's H.320 and H.323 based endpoints.

5. Downspeeding

In addition to IPLR, TANDBERG supports Downspeeding on IP. If excessive packet loss is detected, above 10% for 2 seconds, the TANDBERG unit will automatically downspeed the far end using flow control in steps of 64kbps until acceptable packet loss is detected. If the codec successfully downspeeds to 192kbps and excessive packet loss is still detected, the TANDBERG will upspeed back to the original call rate and not attempt to downspeed again for the duration of the call. The assumption being that packet loss in this case is not attributed to the call rate.

6. Appendix

6.1 Polycom Video Error Concealment (PVEC)

To fully understand the power of IPLR, a closer examination of other methods of dealing with packet loss is required. Polycom currently implement a feature called PVEC that attempts to sidestep the effects of packet loss. The developed feature is proprietary in nature and therefore will currently only work on the ViewStation FX, EX and VS4000 codecs.

TANDBERG's IPLR minimizes the effects of lost information, while Polycom attempts to minimize the information lost due to packet loss. Realizing each frame of video is coded and inserted into a IP packet, Polycom approached the problem by dividing the video frame into 4 areas and inserting those 4 areas of information into 4 different packets. If one packet is lost, in theory there should be enough information from the other 3 areas to make up for the lost area. This works well, unless the packet loss problem is dramatic.

If 4 packets are now needed to send the same image previously sent in 1 packet, it therefore follows that to maintain the same framerate, the packet rate must increase by a factor of 4. This means in packet loss situations, Polycom will typically lose up to 4 times as many packets as the traditional video system. In addition, 4 times the number of packets means 4 times the amount of packet overhead. TANDBERG offers a comparable performance, if not better, with less packet overhead and is standards compliant.

6.2 TANDBERG IPLR vs. Polycom PVEC

After reviewing the methods used by the major manufacturers today, the advantages of the TANDBERG IPLR method can be seen. The following table describes some of these advantages when compared to Polycom's PVEC.

	TANDBERG IPLR	Polycom PVEC
Implemented on total range of products	Yes	No, FX/EX/VS4000 only
Standards compliant	Yes	No, proprietary
Works with all vendors' endpoints	Yes	No, both endpoints must be FX/VS4000
Works with all vendors' MCUs	Yes	No MCU supports PVEC today
Works with all vendors' gateways	Yes	No gateway supports PVEC today
Works in dual image mode	Yes, Duo Video	No, Visual Concert

The Polycom PVEC method offers similar performance to IPLR, but it is not a solution based feature as it is limited to the FX/VS4000 products only.

Additionally, TANDBERG will perform better in high jitter and high latency environments which are typically found on non-QoS enabled networks such as the public Internet.

6.3 Considerations around H.264

Parts of the Polycom PVEC algorithm have been incorporated into the new ITU-T H.264 video standard. This means portions of the PVEC technique will be available to all manufacturers to improve the error resiliency of the video. It is likely that the PVEC technique will be 'standard' when Polycom is running in H.264 mode only.

Since H.264 is a standard, the portions of PVEC that are contained in H.264 are available to all manufacturers, including TANDBERG. In addition, this allows TANDBERG to use intelligent decisions in the software to apply either IPLR or some of the PVEC techniques depending on which will perform best at a given time. IPLR is still standards compliant with H.261 and H.263 in addition to H.264, unlike PVEC. TANDBERG is able to offer the same standardized error resiliency in H.264 as Polycom, and can go beyond by providing IPLR during H.264, H.263 and H.261 conferences. This makes the TANDBERG offering the most complete, error resilient solution on the market.