ATIS STANDARD

ATIS-0100035

ATIS Standard on -

TELEPRESENCE QUALITY OF EXPERIENCE AND QUALITY OF SERVICE

As a leading technology and solutions development organization, ATIS brings together the top global ICT companies to advance the industry's most-pressing business priorities. Through ATIS committees and forums, nearly 200 companies address cloud services, device solutions, M2M communications, cyber security, ehealth, network evolution, quality of service, billing support, operations, and more. These priorities follow a fast-track development lifecycle—from design and innovation through solutions that include standards, specifications, requirements, business use cases, software toolkits, and interoperability testing.

ATIS is accredited by the American National Standards Institute (ANSI). ATIS is the North American Organizational Partner for the 3rd Generation Partnership Project (3GPP), a founding Partner of oneM2M, a member and major U.S. contributor to the International Telecommunication Union (ITU) Radio and Telecommunications sectors, and a member of the Inter-American Telecommunication Commission (CITEL). For more information, visit < www.atis.org >.

ATIS-0100035, *Telepresence Quality of Experience and Quality of Service*

Is an ATIS Standard developed by the **ATIS Network Performance, Reliability, and Quality of Service Committee (PRQC)**.

Published by

Alliance for Telecommunications Industry Solutions
1200 G Street, NW, Suite 500
Washington, DC 20005

ATIS-0100035

ATIS Standard

Telepresence Quality of Experience and Quality of Service

Alliance for Telecommunications Industry Solutions

Approved July, 2012

Abstract

This ATIS Standard evaluates and offers recommendations on the QoS and QoE aspects of telepresence, and examines these QoS and QoE aspects for the development of interworking requirements to support desirable levels of telepresence quality.

Foreword

The Alliance for Telecommunication Industry Solutions (ATIS) serves the public through improved understanding between providers, customers, and manufacturers. The Network Performance, Reliability, and Quality of Service Committee (PRQC) develops and recommends standards, requirements, and technical reports related to the performance, reliability, and associated security aspects of communications networks, as well as the processing of voice, audio, data, image, and video signals, and their multimedia integration. PRQC also develops and recommends positions on, and foster consistency with, standards and related subjects under consideration in other North American and international standards bodies.

The mandatory requirements are designated by the word *shall* and recommendations by the word *should*. Where both a mandatory requirement and a recommendation are specified for the same criterion, the recommendation represents a goal currently identifiable as having distinct compatibility or performance advantages. The word *may* denotes a optional capability that could augment the standard. The standard is fully functional without the incorporation of this optional capability.

Suggestions for improvement of this document are welcome. They should be sent to the Alliance for Telecommunications Industry Solutions, PRQC, 1200 G Street NW, Suite 500, Washington, DC 20005.

At the time it approved this document, PRQC, which is responsible for the development of this report, had the following members:

P. Tarapore, PRQC Chair (AT&T)

J. Colombo, PRQC Vice-Chair (Verizon)

P. Tarapore, Technical Editor (AT&T)

C. Dvorak, Technical Editor (AT&T)

M. Geller, Technical Editor (Cisco)

C.A. Underkoffler, ATIS Chief Editor

Table of Contents

Table of Figures

Table of Tables

ATIS Standard on –

Telepresence Quality of Experience and Quality of Service

1 Scope & Purpose

Telepresence enables people at different locations to communicate with each other in an experience that is as close to "face-to-face" as possible. Originally, this service was implemented in an intra-service provider (SP) (single administrative domain) at the business level. This intra-SP success led to telepresence on an inter-SP basis, delivered across multiple domains/multiple SPs via dedicated pipes, bilateral SP agreements, or various networks.

To enable a successful telepresence experience in this expanded environment, many aspects related to telepresence planning need to be considered. Two critical aspects are:

- *Telepresence Quality of Service (QoS)*: End-to-end performance parameters that focus on the IP transport from the underlying networks and enterprise segments impact the service QoS. These parameters are IP Transfer Delay, Delay Variation (Jitter), Packet Loss, and Errored Packets.
- *Telepresence Quality of Experience (QoE)*: Service level parameters such as video frame loss that impact the actual experience of the service viewers.

The purpose of this ATIS Standard is the evaluation and recommendation of the QoS and QoE aspects of telepresence described above.

The scope of this ATIS Standard is the examination of these QoS and QoE aspects, and the development of interworking requirements to support desirable levels of telepresence quality.

2 Normative References

The following standards contain provisions which, through reference in this text, constitute provisions of this Standard. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this Standard are encouraged to investigate the possibility of applying the most recent editions of the standards indicated below.

[ATIS-0200005] ATIS-0200005, *Cloud Framework for Telepresence Service*, February 2012.[1]

[RFC2474] IETF RFC 2474, *Definition of the Differentiated Services Field in the IPV4 and IPv6 Headers*, December 1998.[2]

[RFC4594] IETF RFC 4594, *Configuration Guidelines for Diffserv Service Classes*, August 2006.[2]

[Y.1540] ITU-T Recommendation Y.1540, *Internet Protocol Data Communication Service – IP Packet Transfer and Availability Performance Parameters*, March 2011.[3]

[Y.1541] ITU-T Recommendation Y.1541, *Network performance objectives for IP-based services,* December 2011.[3]

[1] This document is available from the Alliance for Telecommunications Industry Solutions (ATIS), 1200 G Street N.W., Suite 500, Washington, DC 20005. < https://www.atis.org/docstore/default.aspx >

[2] This document is available from the Internet Engineering Task Force (IETF). < http://www.ietf.org >

[3] This document is available from the International Telecommunications Union. < http://www.itu.int/ITU-T/ >

[Y.1542] ITU-T Recommendation Y.1542, *Framework for achieving end-to-end IP performance objectives*, June 2010.[3]

3 Definitions

For a list of common communications terms and definitions, please visit the *ATIS Telecom Glossary*, which is located at < http://www.atis.org/glossary >.

Telepresence: For purposes of this document, *telepresence* is defined as a specific sub-set of video teleconferencing, in which the underlying technologies support an enhanced feeling or appearance that the other parties are present, due to a close approximation of the in-person, face-to-face experience. Factors involved include nearly life-size video images of faces, with the accompanying excellent eye contact; the lack of noticeable delay; excellent spatial resolution and motion handling; and enhanced interaction with the remote location. When all this is achieved, the telepresence experience is often described as being "immersive".

The term *telepresence* can have other, more general, usages that can include things like *telerobotics*, but the definition given above clearly limits the meaning here to communications involving people at near life-size images.

Also, in this document, the term *telepresence* is used generically – meaning that no specific network, system, or premises equipment platform is assumed. The focus is instead on the application, and the QoE perceived while that application is used.

Definitions of other terms used in this document are as follows:

Bandwidth: A characteristic of a communication channel that is the amount of information that can be passed through it in a given amount of time, usually expressed in bits per second [ATIS Glossary].

IP Packet Transfer Delay: The time between the occurrence of two corresponding IP packet reference events, an *ingress event* (IP packet leaves the source) and an *egress event* (IP packet arrives at the destination) [ITU-T Recommendation Y.1540] – see Appendix A.

IP Packet Delay Variation (Jitter): End-to-end 2-point IP packet delay variation (PDV) is defined based on the observations of corresponding IP packet arrivals at the ingress (source) and egress (destination). These observations characterize the variability in the pattern of IP packet arrival events at the egress MP and the pattern of corresponding events at the ingress MP with respect to a reference delay. The PDV for an IP packet between the source and destination is the difference between the absolute IP packet transfer delay of the packet and a defined reference IP packet transfer delay between the source and destination [ITU-T Recommendation Y.1540] – see Appendix A.

IP Packet Error Ratio: IP packet error ratio is the ratio of total errored IP packet outcomes to the total of successful IP packet transfer outcomes plus errored IP packet outcomes in a population of interest – see Appendix A.

IP Packet Loss Ratio: IP packet loss ratio is the ratio of total lost IP packet outcomes to total transmitted IP packets in a population of interest [ITU-T Recommendation Y.1540] – see Appendix A.

Motion Handling: For the purposes of this document, video *motion* is any frame-to-frame change in the spatial image; and *motion handling* is the ability to deliver video motion to users free of commonly occurring temporal artifacts (e.g., "mosquito noise", "ghosting", "jerkiness", and "smearing").

Resolution: A parameter that specifies the ability to distinguish video detail in the spatial dimension or the temporal dimension [ITU-T Recommendation P.10].

Video Frame Jitter: Video Frame Jitter is defined as the difference between the actual arrival time and the expected arrival time of a Video Frame [ATIS-0200005].

4 Acronyms & Abbreviations

AF	Assured Forwarding
ATIS	Alliance of Telecommunications Industry Solutions
CS	Class Selector
DSCP	Diffserv Code Point
EF	Expedited Forwarding
IP	Internet Protocol
IPTD	IP Packet Transfer Delay
IPDV	IP Packet Delay Variation
IPER	IP Packet Error Ratio
IPLR	IP Packet Loss Ratio
Kbps	Kilobits per Second
MFSR	Maximum Frame Serialization Rate
NSE	Network Section Ensemble
PHB	Per Hop Behavior
QoE	Quality of Experience
QoS	Quality of Service
SIP	Session Initiation Protocol
SP	Service Provider
UNI	User-Network Interface

5 Introduction – Telepresence as an Application

Telepresence is a service that will be used much more widely with the achievement of interoperable networks and equipment. Inter-SP telepresence services enable the delivery of communication capabilities that allow users to collaborate easily with other telepresence systems via point-to-point or multi-point sessions.

One important challenge presented in the inter-SP environment is that the calls/sessions must traverse multiple trust and/or administrative domains. Often these domains represent separate companies or different subsidiaries of the same company[4].

There is a similarity here with the delivery of Session Initiation Protocol (SIP) based voice services across multiple SP domains. The experiences of the current federations and “peering” models for delivery of SIP-based services can be expanded for delivery of SIP-based telepresence services.

Another important challenge in inter-SP telepresence solutions comes from the fact that a truly immersive telepresence experience requires QoS-enabled, low-latency, IP transport paths.

Accordingly, the focus of this ATIS Standard is on the support of telepresence with acceptable QoE achieved across multiple SPs with interoperable systems. This requires, amongst other things, adequate QoS across the end-to-end IP path, as depicted in Figure 1.

[4] The combination of these SP domains can be considered to be a telepresence delivery “cloud”.

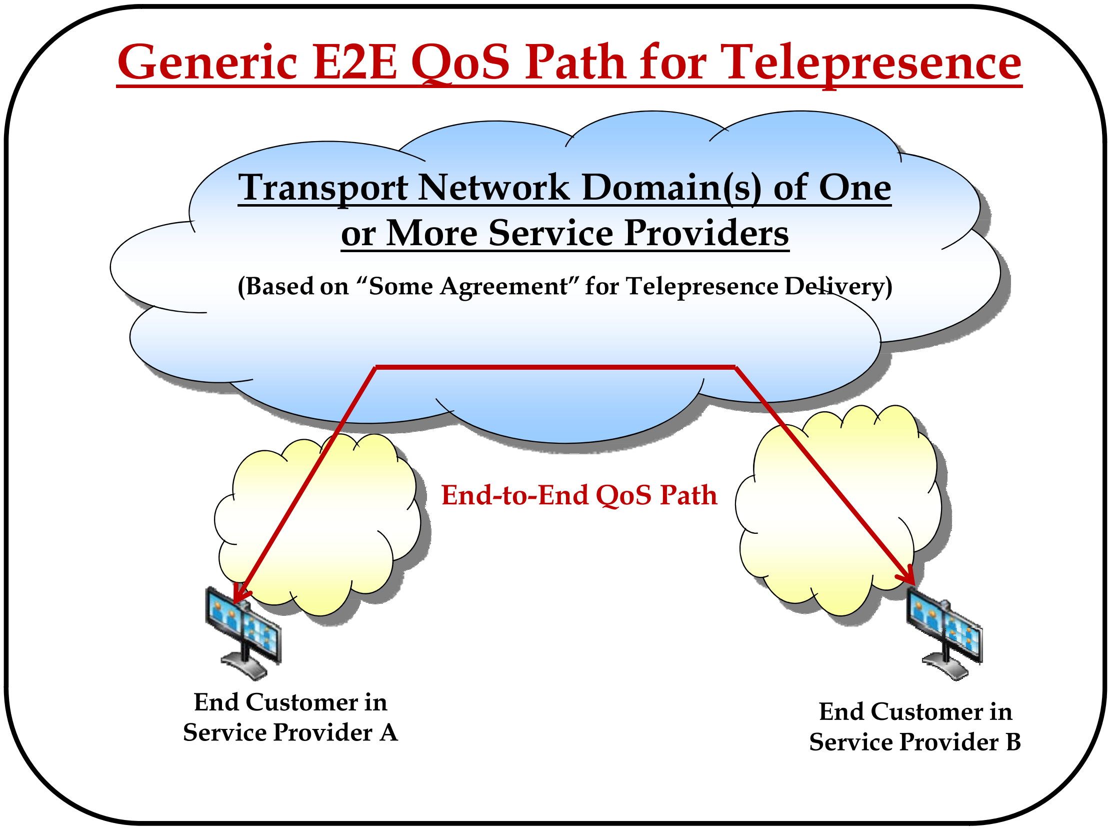

Figure 1 - End to End QoS Network Path for Telepresence Service

6 Telepresence Infrastructure from an End-to-End Perspective

The overall architecture of the Telepresence Service is described as follows.

The telepresence application comprises a supporting infrastructure (referred to as the Telepresence Infrastructure) for the following services:

- Call Services
- Session Services
- Media Services
- Endpoint Services Model

The Telepresence Infrastructure is supported by an underlying Transport Network Infrastructure. Given that the delivery of a satisfactory telepresence experience strongly depends on carrier-grade network quality, it is expected that the Transport Network Infrastructure will be provided by well managed Service Provider network domains[5]. This enables best-practice carrier capabilities to satisfy the stringent QoS requirements such that completely acceptable telepresence QoE can be delivered to end customers.

The Transport Infrastructure provides a number of key aspects required for visibility into network performance, call admission, and service quality. These elements include built-in probe technology to the network elements

[5] In principle, the transport network domain can be provided by the public Internet as well. However, in such a case, there can be no guarantees for delivering a satisfactory telepresence experience.

from end to end, as well as tools to passively examine the traffic as it traverses the network. All of this visibility enables the network operator to manage the services deployed, both proactively for network planning, and reactively for troubleshooting. Once a proper baseline is established, the network can be modeled to control how it behaves to deliver QoE.

The following sections of this report address further details of telepresence quality levels and associated traffic characteristics.

7 Telepresence Quality of Service Requirements

As depicted in Figure 1, the network level QoS is defined by the path of the telepresence service between the two customer end points. The QoS guidance provided here is generic and is applicable to all types of transport network configurations such as:

- *Single Transport Network Domain*: Both customers access telepresence services over a single SP's network.
- *Customers Access Telepresence From Two Different SPs*:
    - The two access networks are interconnected via one transport network operated by an independent SP.
    - The two access networks are interconnected via multiple transport networks operated a group of independent SPs. This is the configuration shown in Figure 1 and is considered to be the most general case.
    - The two access networks are interconnected via the public Internet. For this scenario, the "Best Effort" nature of the public Internet cannot enforce the stringent QoS requirements necessary for an acceptable telepresence experience. Hence, no QoS guarantees can be made here; this scenario is included for completeness.

Regardless of the configuration of the transport network(s) used for the IP path, network level (UNI-to-UNI) performance objectives for delay, jitter, and packet loss must be met if the resulting QoE is to be acceptable. These performance objectives have been defined for IP-based networks by ITU-T Recommendation Y.1541 [Y.1541] in the form of six IP QoS classes (see Appendix B)[6]. Each QoS class consists of network performance objectives (upper bounds) for packet delay, delay variation, and loss, as well as for payload errors within the IP packets. Because only Classes 0 and 1 have an upper bound on packet delay variation, only these two of the six IP QoS classes in Y.1541 can be meaningfully considered for a demanding application like telepresence.

Classes 0 and 1 of Y.1541 differ only in the amount of allowed packet delay (100 and 400 ms, respectively). The allowed delay variation (50 ms), loss (1×10^{-3}) and errors (1×10^{-4}) are the same for both classes.

While it is desirable to meet the 100 ms network delay objective of Class 0, in practice it is often not possible to do so. Accordingly, the telepresence transport infrastructure should satisfy the network performance objectives of Y.1541 Class 1 (where the allowed network delay is 400 ms).

> **QoS Requirement R1**: The performance of the IP network path should satisfy the ITU-T Recommendation Y.1541 Class 1 upper bounds of 400 ms delay, 50 ms delay variation, 1×10^{-3} loss, and 1×10^{-4} errors.
>
> > NOTE: Although the Class 1 delay objective has an upper bound of 400 ms, it is strongly recommended to keep the network delay well below this – preferably below 200 ms – to assure that there are no noticeable delay effects. With careful transport network planning, delay values below 200 ms are often readily achievable.

[6] Note that Appendix VIII in ITU-T Recommendation Y.1541 provides useful information on digital TV.

It must be emphasized that these Y.1541 objective values are for the entire IP network path (UNI-to-UNI), even if network segments in the path are provided by different operators. In addition, statically allocating these "end-to-end" objectives to individual operators is not feasible because of serious practical and technical obstacles. Different approaches to realizing these objectives across multi-operator paths have been assessed in detail by the ITU-T (see Y.1542, "Framework for achieving end-to-end IP performance objectives").

8 Telepresence Traffic Characteristics

The previous section dealt with the IP network performance that is required to support acceptable telepresence QoE. However, to fully describe what affects telepresence QoE, the traffic characteristics of the telepresence video stream must be understood – for two reasons. First, the realized video quality obviously depends on the spatial resolution and motion rendering – each of which ultimately drives the bandwidth required to support the video stream. Second, and perhaps more subtle, is the video frame rate and size generated by the source codec, because it is these video frames that must be serialized as they are transmitted across the underlying IP networks.

This section deals with the telepresence traffic characteristics; the next section deals with the video frame aspects of telepresence QoE.

Table 1 (see below) illustrates the relationships between spatial resolution and motion handling aspects, with the associated bandwidth requirements for different telepresence system configurations [ATIS-0200005]:

Table 1 - Telepresence Traffic Characteristics: Max Bandwidth Consumption (kbps)

Maximum Bandwidth Consumption Kilobits Per Second (kbps)								
Resolution		**1080p**	**1080p**	**1080p**	**720p**	**720p**	**720p**	**720p**
Motion Handling		Best	Better	Good	Best	Better	Good	Lite
Video per Screen (kbps)		4000	3500	3000	2250	1500	1000	936
Audio per Microphone (kbps)		64	64	64	64	64	64	64
Auto Collaborate Video Chan.		500	500	500	500	500	500	100
Auto Collaborate Audio Chan.		64	64	64	64	64	64	64
Single Screen Systems		4628	4128	3628	2878	2128	1628	1164
Total Audio and Video (kbps)		4756	4256	3756	3006	2256	1756	1292
Triple Screen Systems **Total Audio and Video (kbps)**		12756	11256	9756	7606	6256	3766	
+ 20% for Layer 2-4 overhead								
Single Screen Systems *max* bandwidth (kbps)	Tx	5554	4954	4358	3454	2554	1954	1397
Includes layer 2-4 overheard	Rx	5707	6107	4607	3607	2707	2107	1550
Triple screen systems *max* bandwidth (kbps) ***Includes layer 2-4 overheard***		15307	13507	11707	9007	6307	4507	

Optional Ad-on Features (kbps) – ***Not* applicable to 720p Lite**				
30fps Auto Collaborate		4000	+20% for Layer 2-4 overhead	4800
CTRS Recording in CIF		704	+20% for Layer 2-4 overhead	845
SD Interoperability	**Audio**	704	+20% for Layer 2-4 overhead	922
	Video	64		
WebEx On Touch	**Audio**	304	+20% for Layer 2-4 overhead	442
	Video	64		

This table demonstrates bandwidth usage for two levels of resolution (1080p and 720p) each with associated levels of motion handling (“Best”, Better”, and “Good”). The bandwidth requirement to avoid potential QoE impairments during the telepresence session is as follows:

> **QoE Requirement R1**: The bandwidth required across the end-to-end path should take into account all necessary factors for enabling the desired levels of telepresence service for resolution and motion handling.
>
> > NOTE: This requirement has two implications: First, if it is known that a limited amount of bandwidth is available (for example, due to the access network), then the resolution and motion handling combination chosen must not exceed the limited bandwidth. Conversely, if it is known that bandwidth is not constrained by the end-to-end transport path, then the desired combination of resolution and motion handling can be made independently from bandwidth considerations – as long as that bandwidth is guaranteed. Of course cost may then be a factor.

9 Video Frame Jitter Considerations

Now that telepresence traffic characteristics have been described, the QoE-affecting aspect of video frame jitter can be meaningfully examined. Unlike IP network impairments, the effects of not adequately handling video frame jitter directly impacts telepresence QoE. This is because of the direct influence on the video decoder, and thus the associated visible image impairments (gradual screen repair, full screen refreshes, picture freeze, etc.).

The first important consideration in understanding video frame jitter is knowing that the video frame rate is 30 frames/s, so that a video frame is generated by the source codec every 33 ms. However, these video frames can vary greatly in size (from 1 to 65 Kbytes, with the average being 16 Kbytes), depending on the image complexity (e.g., spatial and temporal redundancy).

These varying-sized video frames are packetized and then transmitted across facilities operating at widely varying underlying clock rates (e.g., T1, 10 Mb/s Ethernet, T3, etc.). These clock rates, combined with the video frame size, determine how much time it takes to completely transmit each frame – serialization delay--and clearly this time will vary as the frame size varies. Thus video frame jitter is generated, as the following figure graphically depicts [ATIS-0200005]:

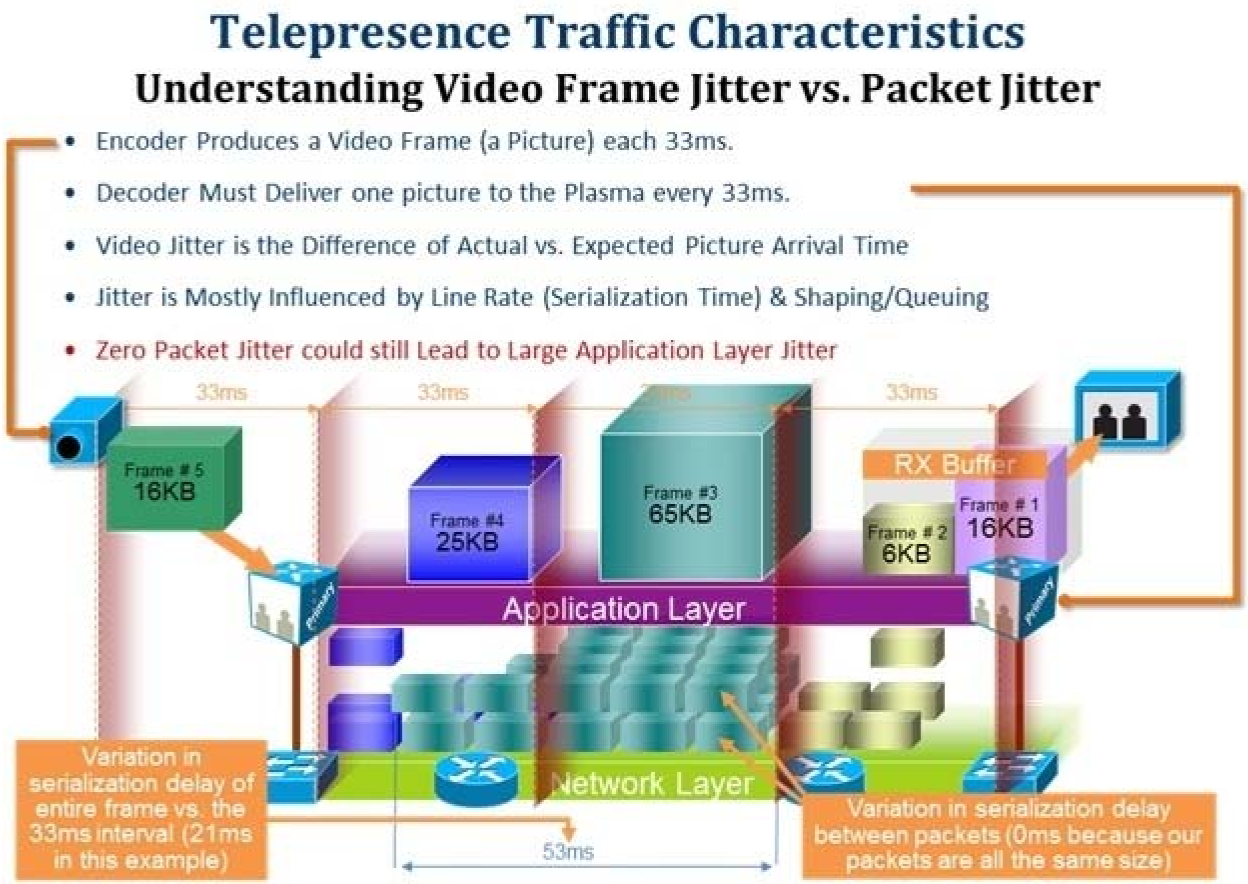

Figure 2 - Understanding Video Frame Jitter versus IP Packet Jitter

The previous figure shows an example of how video frame jitter is generated. It also shows that the frame buffer in the receive telepresence system must have enough storage to readily accommodate this jitter, as the frames will not arrive when “expected” – yet every 33 ms it has to provide the next frame (independent of size) to the decoder.

To appreciate what telepresence systems have to do to accommodate possible ranges of video frame jitter, the following table shows maximum jitter as a function of clock rate and frame size [ATIS-0200005]:

Table 2 - Understanding Video Frame Jitter Versus IP Packet Jitter

Per Screen	T1	E1	4 x T1	4 x E1	10Mbps Ethernet	E3	T3	OC3	OC12
Circuit Bit Rate (kbps)	1, 544	2, 048	6, 176	8,192	10,000	34,368	44,736	155,520	622,080
Max Frame Size	65,000 Bytes								
Average Frame Size	16,000 Bytes								
Min Frame Size	1,000 Bytes								
Max-Frame Serialization Rate (MFSR)	337 ms	265 ms	84 ms	64 ms	52 ms	15 ms	12 ms	4 ms	<1 ms
Average-Frame Serialization Rate	83 ms	63 ms	21 ms	16 ms	13 ms	4 ms	3 ms	<1 ms	<1 ms
Min-Frame Serialization Rate	5 ms	4 ms	1 ms	<1 ms	<1 ms	<1 ms	<1 ms	<1 ms	<1 ms
Max Jitter (MFSR vs. 33ms)	304 ms	221 ms	51 ms	31 ms	19 ms	0 ms	0 ms	0 ms	0 ms

It has not yet been explicitly stated, but it should be clear that the video frame jitter values given in the table above are *completely independent* of the IP packet delay variation described earlier. In other words, even if IP packet delay variation were *zero*, these video frame jitter values would still occur. While packet delay variation (packet jitter) can contribute to video frame jitter, especially in congested IP networks with routers using certain queuing and shaping algorithms, its contribution to video frame jitter is often secondary to that described above.

QoS Requirement R1 (Section 7) defines the target for IP packet delay variation and QoE Requirement R2 (Section 10) identifies packet marking as a means to minimize IP packet delay variation.

10 IP Network Considerations

Because of the potential impacts of video frame jitter on QoE, efforts must be made to avoid it, and so this section addresses IP network considerations – such as the importance of things like marking packets for appropriate network treatment.

Table 3 - DSCP Guidance for Various Components of Telepresence Service

Telepresence Network Design

RFC 4594 Configuration Guidelines for DiffServ Classes

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 / telepresence	CS4	32	RFC 2474
Multimedia Streaming	AF31	26	RFC 2597
Broadcast Video	CS3	24	RFC 2474
Low-Latency / Transactional Data	AF21	18	RFC 2597
Operations / Administration / Management	CS2	16	RFC 2474
High-Throughput / Bulk Data	AF11	10	RFC 2597
Best Effort	DF	0	RFC 2474
Low-Priority / Scavenger Data	CS1	8	RFC 3662

In order to ensure consistent treatment of telepresence IP packets across all domains in the end-to-end IP path, it is necessary to mark the packets via Diffserv Code Point 32 [RFC2474]. Table 3 provides guidance on potential DSCP values for various applications, where it can be seen a DSCP value of 32 is assigned to telepresence. This leads to the next QoE requirement of this document:

> **QoE Requirement R2:** Each telepresence service component should be marked with a DSCP value of 32 by the initiating telepresence domain so that all subsequent domains in the IP path can provide appropriate packet treatment per pre-arranged negotiations between SPs.

11 Conclusions

This ATIS Standard describes QoS and QoE aspects of telepresence applications over IP networks. The importance of appreciating the bandwidth aspect of various telepresence configurations is emphasized, as is the potential impact of video frame jitter.

To achieve acceptable telepresence QoS and QoE, three requirements are provided that are especially important to meet when the end-to-end telepresence service is realized using systems and networks provided by different manufacturers and operators, respectively.

The methods and processes (e.g., use of MPLS tunnels for reserving end-to-end bandwidth, algorithms for minimizing video frame jitter, etc.) that could be implemented for enabling these requirements are for further study.

Appendix A
(informative)

A IP Packet Quality of Service Metrics

The detailed descriptions of these critical performance parameters are taken from ITU-T Recommendation Y.1540 [Y.1540].

A.1 IP Packet Transfer Delay (IPTD)

IP packet transfer delay is defined for all successful and errored packet outcomes across a basic section or an NSE. IPTD is the time, $(t_2 – t_1)$ between the occurrence of two corresponding IP packet reference events, ingress event $IPRE_1$ at time t_1 and egress event $IPRE_2$ at time t_2, where $(t_2 > t_1)$ and $(t_2 – t_1) \leq T_{max}$. If the packet is fragmented within the NSE, t_2 is the time of the final corresponding egress event. The end-to-end IP packet transfer delay is the one-way delay between the MP at the SRC and DST as illustrated in Figure A.1.

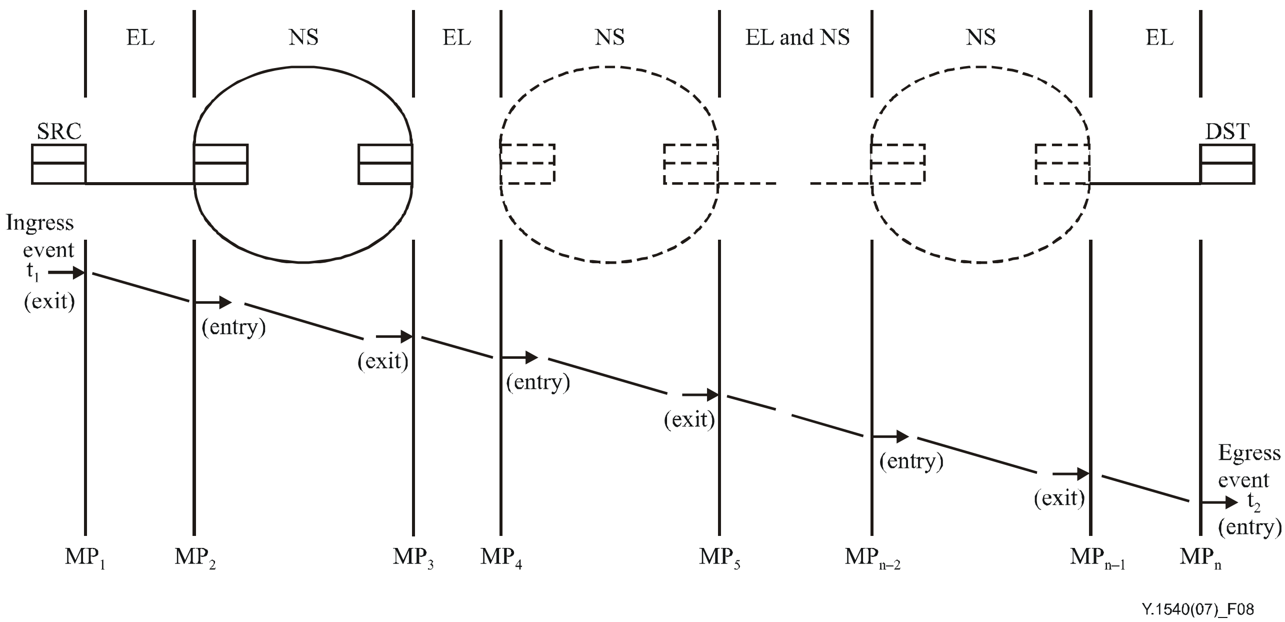

Figure A.1 - IP packet transfer delay events

A.2 End-to-end 2-point IP Packet Delay Variation (IPDV)

The variations in IP packet transfer delay are also important. Streaming applications might use information about the total range of IP delay variation to avoid buffer underflow and overflow. Extreme variations in IP delay will cause TCP retransmission timer thresholds to grow and may also cause packet retransmissions to be delayed or cause packets to be retransmitted unnecessarily.

End-to-end 2-point IP packet delay variation (PDV) is defined based on the observations of corresponding IP packet arrivals at ingress and egress MP (e.g., MP_{DST}, MP_{SRC}). These observations characterize the variability in the pattern of IP packet arrival events at the egress MP and the pattern of corresponding events at the ingress MP with respect to a reference delay.

The 2-point PDV (v_k) for an IP packet k between SRC and DST is the difference between the absolute IP packet transfer delay (x_k) of packet k and a defined reference IP packet transfer delay, $d_{1,2}$, between those same MPs (see Figure A.2): $v_k = x_k – d_{1,2}$.

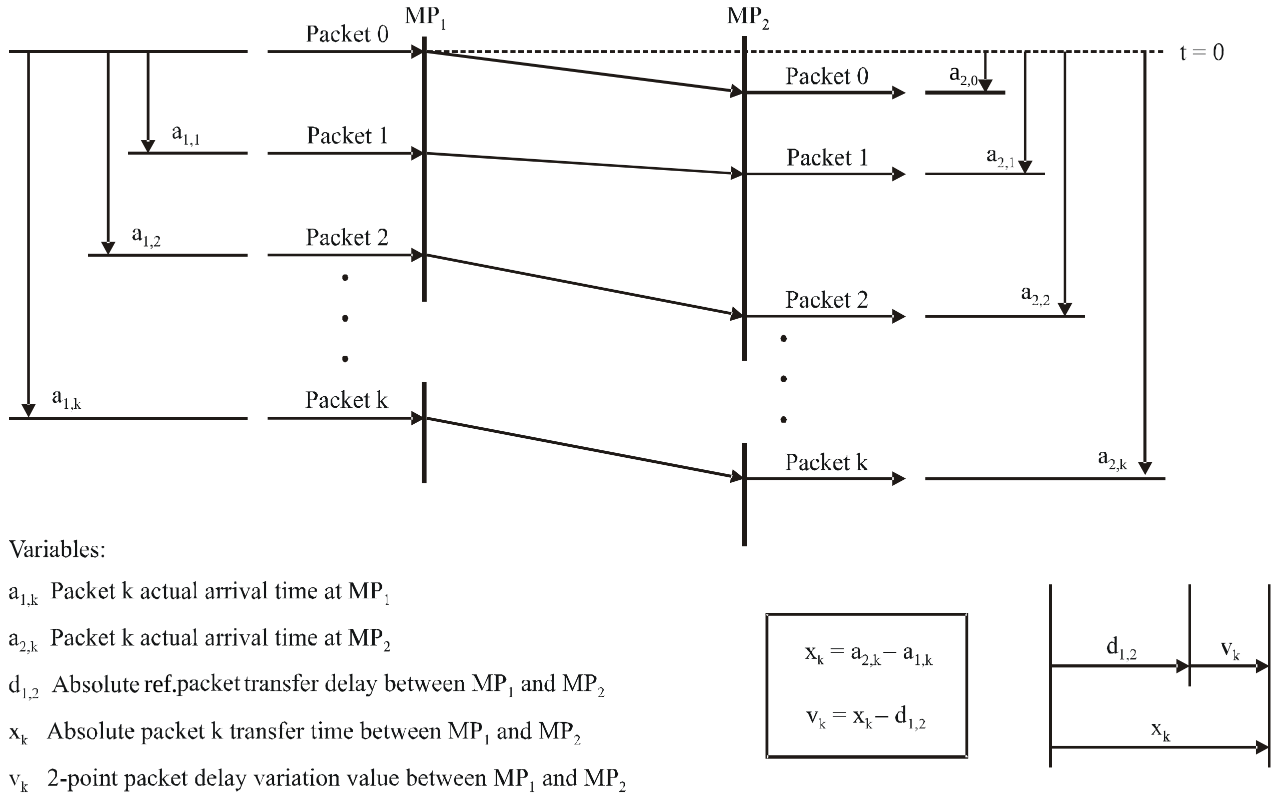

Figure A.2 - 2-point IP packet delay variation

The reference IP packet transfer delay, $d_{1,2}$, between SRC and DST is the absolute IP packet transfer delay experienced by a selected IP packet between those two MPs.

Positive values of 2-point IPDV correspond to IP packet transfer delays greater than those experienced by the reference IP packet; negative values of 2-point PDV correspond to IP packet transfer delays less than those experienced by the reference IP packet. The distribution of 2-point PDVs is identical to the distribution of absolute IP packet transfer delays displaced by a constant value equal to $d_{1,2}$.

A.3 Using Minimum Delay as the Basis for Delay Variation

As illustrated in Figure A.2, the delay variation of an individual packet is naturally defined as the difference between the actual delay experienced by that packet and a nominal or reference delay. The preferred reference (used in ITU-T Y.1541 IPDV objectives) is the minimum delay of the population of interest. This ensures that all variations will be reported as positive values, and simplifies reporting the range of variation (the maximum value of variation is equal to the range). Distributions of delay variation in IP networks often exhibit a bias toward the minimum (e.g., the minimum and the mode are equal). Many more useful capabilities of this form of delay variation – PDV, using the minimum delay as reference – are detailed in [IETF RFC 5481].

Use of the average delay as the delay variation reference is depreciated in this version of this Recommendation.

In previous versions of this Recommendation, there was an alternative to using the minimum packet delay as the nominal delay: to use the average delay of the population of interest as the nominal or reference delay. This has the effect of centering the distribution of delay variation values on zero (when the distribution is symmetrical), and produces both positive and negative variations. However, the average delay of the population may be distinctly different from the delay of any individual packet, creating an artificial reference for variation (e.g., when a bimodal distribution is present).

A.4 Quantile-based Limits on IP Packet Delay Variation

The preferred method (used in ITU-T Y.1541 objectives) for summarizing the delay variation of a population of interest is to select upper and lower quantiles of the delay variation distribution, and then measure the distance between those quantiles. For example, select the $1 – 10^{–3}$ quantile and the 0 quantile (or minimum), make measurements, and observe the difference between the delay variation values at these two quantiles. This example would help application designers determine the de-jitter buffer size for no more than 0.1% total buffer overflow.

An objective for IP packet delay variation could be established by choosing an upper bound for the difference between pre-specified quantiles of the delay variation distribution. For example, "The difference between the 99.9 quantile and the minimum of the packet delay variation should be no more than 50 ms."

A.5 IP Packet Loss Ratio (IPLR)

IP packet loss ratio is the ratio of total lost IP packet outcomes to total transmitted IP packets in a population of interest.

NOTE: Metrics for describing one-way loss patterns may be found in [b-IETF RFC 3357]. Consecutive packet loss is of particular interest to certain non-elastic real-time applications, such as voice and video.

A.6 IP Packet Error Ratio (IPER)

IP packet error ratio is the ratio of total errored IP packet outcomes to the total of successful IP packet transfer outcomes plus errored IP packet outcomes in a population of interest.

Appendix B
(informative)

B Y.1541 Performance Classes

Table B.1 - IP network QoS class definitions and network performance objectives -

Network performance parameter	Nature of network performance objective	QoS Classes					
		Class 0	Class 1	Class 2	Class 3	Class 4	Class 5 Unspecified
IPTD	Upper bound on the mean IPTD (Note 1)	100 ms	400 ms	100 ms	400 ms	1 s	U
IPDV	Upper bound on the $1 - 10^{-3}$ quantile of IPTD minus the minimum IPTD (Note 2)	50 ms (Note 3)	50 ms (Note 3)	U	U	U	U
IPLR	Upper bound on the packet loss probability	1×10^{-3} (Note 4)	1×10^{-3} (Note 4)	1×10^{-3}	1×10^{-3}	1×10^{-3}	U
IPER	Upper bound	1×10^{-4} (Note 5)					U

General Notes:

The objectives apply to public IP Networks. The objectives are believed to be achievable on common IP network implementations. The network providers' commitment to the user is to attempt to deliver packets in a way that achieves each of the applicable objectives. The vast majority of IP paths advertising conformance with ITU-T Rec. Y.1541 should meet those objectives. For some parameters, performance on shorter and/or less complex paths may be significantly better.

An evaluation interval of 1 minute is suggested for IPTD, IPDV, and IPLR and – in all cases – the interval must be recorded with the observed value. Any minute observed should meet these objectives.

Individual network providers may choose to offer performance commitments better than these objectives.

"U" means "unspecified" or "unbounded". When the performance relative to a particular parameter is identified as being "U" the ITU-T establishes no objective for this parameter and any default Y.1541 objective can be ignored. When the objective for a parameter is set to "U", performance with respect to that parameter may, at times, be arbitrarily poor.

NOTE 1 – Very long propagation times will prevent low end-to-end delay objectives from being met. In these and some other circumstances, the IPTD objectives in Classes 0 and 2 will not always be achievable. Every network provider will encounter these circumstances and the range of IPTD objectives in Table 1 provides achievable QoS classes as alternatives. The delay objectives of a class do not preclude a network provider from offering services with shorter delay commitments. According to the definition of IPTD in ITU-T Rec. Y.1540, packet insertion time is included in the IPTD objective. This Recommendation suggests a maximum packet information field of 1500 bytes for evaluating these objectives.

NOTE 2 – The definition of the IPDV objective (specified in ITU-T Rec. Y.1540) is the 2-point IP Packet Delay Variation. See ITU-T Rec. Y.1540 and Appendix II for more details on the nature of this objective. For planning purposes, the bound on the mean IPTD may be taken as an upper bound on the minimum IPTD and, therefore, the bound on the $1 – 10^{-3}$ quantile may be obtained by adding the mean IPTD and the IPDV value (e.g., 150 ms in Class 0).

NOTE 3 –This value is dependent on the capacity of inter-network links. Smaller variations are possible when all capacities are higher than primary rate (T1 or E1), or when competing packet information fields are smaller than 1500 bytes (see Appendix IV).

NOTE 4 – The Class 0 and 1 objectives for IPLR are partly based on studies showing that high quality voice applications and voice codecs will be essentially unaffected by a 10^{-3} IPLR.

NOTE 5 – This value ensures that packet loss is the dominant source of defects presented to upper layers, and is feasible with IP transport on ATM.