ITU-T QoS Standards for IP-Based Networks

Neal Seitz, NTIA/ITS

ABSTRACT

To support IP/PSTN convergence, future IP networks will need to provide reliable differentiated QoS to a diverse set of user applications, including telephony. To achieve end-to-end QoS solutions, IP network providers will need to agree on a common set of IP packet transfer performance parameters and QoS objectives. This article describes two new ITU-T Recommendations, Y.1540 and Y.1541, that document such an agreement.

INTRODUCTION

It has been widely assumed that today's circuit-switched and packet-switched networks will gradually come together in an Internet Protocol (IP)-based infrastructure that carries both public switched telephone network (PSTN) and traditional Internet application traffic. This "convergence" scenario has great appeal: it offers both cost savings, through technology consolidation, and industry growth, through new service creation. However, convergence has been slow to materialize. From a technical viewpoint, the major stumbling block has been quality of service (QoS). Traditional IP networks take a best effort approach to quality, affording users a fair share of the available network resources but not ensuring that any particular performance levels will be met. The best effort paradigm has been spectacularly successful in supporting non-real-time data applications (email, file transfer) and has been extended to near-real-time multimedia applications (audio/video streaming, Web browsing). Given the current abundance of bandwidth on many routes, the best effort paradigm even meets the needs of many users today for interactive voice telephony and other real-time applications. However, it does not reliably provide the quality users expect in interactive voice telephony and other demanding real-time applications when bandwidth limitations appreciably increase latency or packet loss. To realize the full benefits of convergence, future IP-based networks will need to implement new resource sharing paradigms capable of reliably providing differentiated QoS to a large and diverse set of end user applications, including, importantly, voice over IP (VoIP). Such resource sharing will need to be coordinated among independent jurisdictions in a multiprovider environment.

An end-to-end IP QoS solution enabling successful IP/PSTN convergence will likely be realized in three steps:

- Achieving network provider agreement on a common set of IP performance parameters and QoS objectives
- Deploying network mechanisms that can support the specified QoS objectives on a terminal-to-terminal basis[1]
- Embedding the QoS objectives in signaling protocols to enable on-demand creation of QoS-assured IP flows

International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) Study Group 13 has recently completed two international standards (Recommendations) that fulfill the first of these steps. The first Recommendation, Y.1540, defines standard performance parameters for packet transfer in IP-based networks [1]. The second, Y.1541, specifies network-interface-to-network-interface (NI–NI) objectives for the Y.1540 parameters and clusters these numerical objectives in six IP network QoS classes [2]. This article describes the development of these new Recommendations, summarizes their technical content, and identifies additional work that will need to be done to maximize their beneficial use in future QoS-assured IP networks.

RECOMMENDATION Y.1540

Recommendation Y.1540 defines the parameters to be used in specifying and assessing the speed, accuracy, dependability, and availability of IP packet transfer in international data communications. The parameters may be used in characterizing end-to-end IP flows and the individual network portions that support such flows. The Y.1540 definitions address connectionless transport as a distinguishing aspect of IP. Y.1540 applies to IP flows provided using IP Version 4 (IPv4).[2]

The intended users of Recommendation

[1] *The mechanisms deployed by network jurisdictions may differ.*

[2] *The principles and definitions apply generally to IPv6 as well. IPv6 is expected to be addressed in future work.*

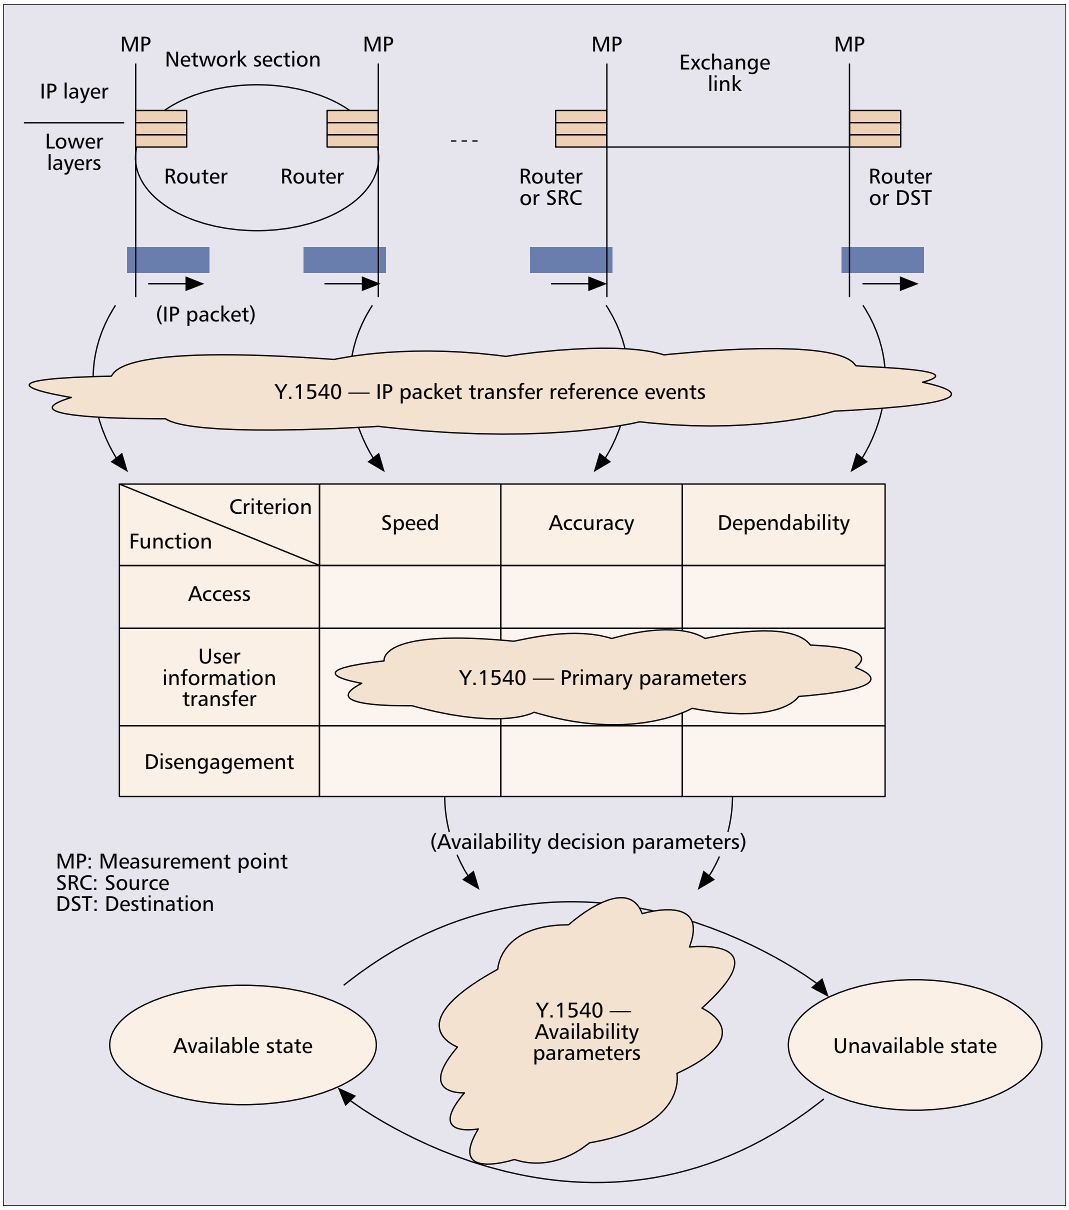

Figure 1. *The reference model and Y.1540 scope.*

Manufacturers will use Y.1540 in developing and marketing equipment conforming to provider specifications. End users will apply the Recommendation in assessing the performance that IP-based networks actually deliver between terminals.

Y.1540 are IP network providers, equipment manufacturers, and end users. Network providers will use Y.1540 in planning, developing, and assessing IP networks in relation to user performance needs. One large network provider is already using Y.1540 parameters in monitoring IP packet transfer performance [3]. Manufacturers will use Y.1540 in developing and marketing equipment conforming to provider specifications. End users will apply the Recommendation in assessing the performance IP-based networks actually deliver between terminals.

PARAMETER DEVELOPMENT PROCESS

Figure 1 illustrates the three-step process ITU-T SG 13 has traditionally used in developing digital network performance parameters, and identifies the scope of Recommendation Y.1540 in that context. The first step is to define the interfaces at which the parameters will be applied, and the specific events that can be observed at those interfaces. The network is modeled as a concatenation of *network sections* and interconnecting *exchange links*. The interfaces between them, called *measurement points* (MPs), are functional boundaries at which standardized communication protocols can be observed. The performance-significant events that can be counted, timed, or compared at MPs are called *reference events* (REs). Specific REs are defined by the interface protocol.

The second step is to define a set of *primary parameters* that collectively characterize the network's performance. The primary parameters relate to particular communication functions and are defined in terms of the REs. A *communication function* defines the expected response of a network (or network portion) to a specified external stimulus; the stimuli and responses are REs. Three generic communication

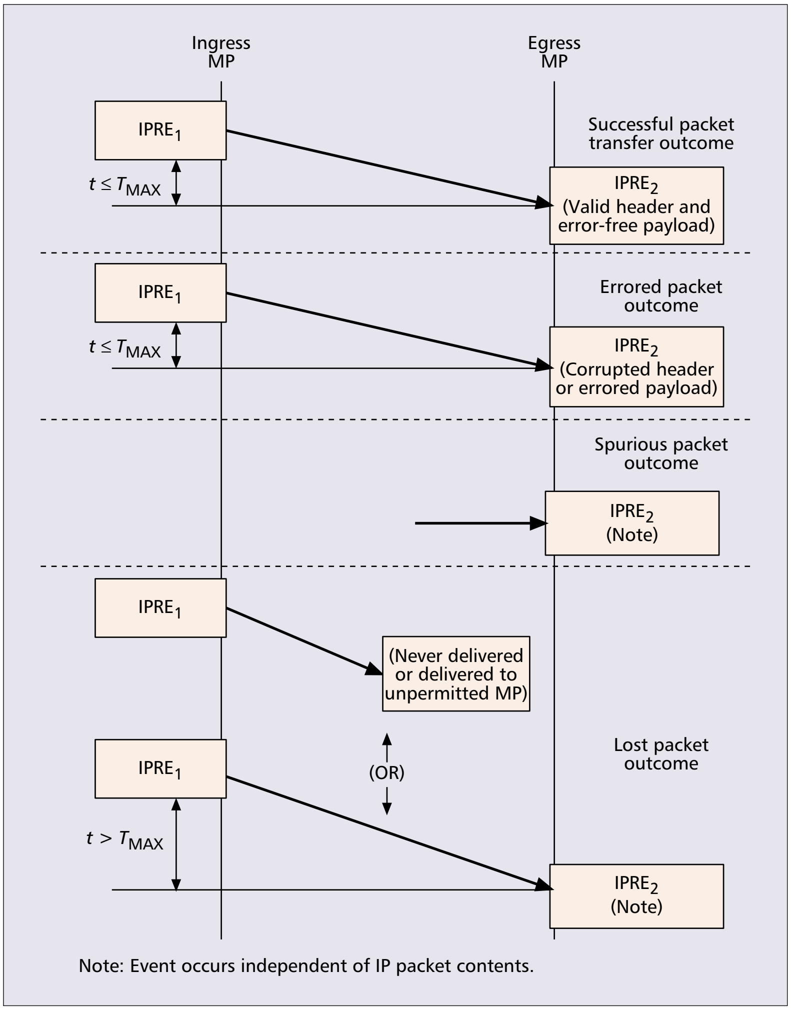

Figure 2. *IP packet transfer outcomes.*

[3] *Although the REs are packet layer events, they need not be observed within the IP protocol stack of the host terminal equipment or router; their times of occurrence can be approximated by observing the IP packets crossing the associated physical interface. The times of occurrence of REs are defined so as to attribute packet insertion (modulation) times to the links.*

[4] *Y.1540 also defines a multipacket* severe loss block outcome *and an associated parameter,* severe loss block ratio*, to identify and limit consecutive or clustered packet loss events. No objective for this parameter is specified in Y.1541.*

functions are commonly used in digital network performance description: *access*, *user information transfer*, and *disengagement*. Statistically, the performance parameters are random variables on a sample space that distinguishes the possible outcomes a function performance attempt may encounter. For any discrete function, three general types of outcomes can be distinguished: successful performance, incorrect performance, and nonperformance. The corresponding user performance concerns (or *criteria*) are *speed*, *accuracy*, and *dependability*. These are related to the three generic communication functions in the familiar 3 × 3 *matrix* [4]. One or more primary parameters are defined to address each function/criterion combination in the matrix. The matrix approach helps ensure that no significant attribute of performance is overlooked.

The third step in parameter development is to define a set of *availability parameters* to characterize performance from a more macroscopic, longer-term viewpoint. The availability parameters are defined on the basis of observed values for a subset of the primary parameters, the *availability decision parameters*. The communication path between a pair (or set) of users is determined to be in either the *available state* or the *unavailable state* by an availability function that compares observed values for the decision parameters with corresponding outage thresholds over successive observation periods. The availability parameters characterize the resulting binary random process in statistical terms.

In developing Y.1540, Study Group 13 agreed that the MPs of interest are the jurisdictional boundaries that separate independently operated IP networks (*autonomous systems*) from end user terminals and each other. The relevant interface protocol is IPv4, and the relevant information units are IP packets. An *IP packet transfer reference event* (IPRE) occurs, for a specified source/destination (SRC/DST) pair, when an IP packet with the defining SRC/DST IP addresses (and a valid header checksum) crosses an MP.[3] The only communication function considered in Recommendation Y.1540 is *IP packet transfer*; the access and disengagement functions are not addressed. This reflects the connectionless nature of today's IP networks. ITU-T SG 13 and other groups are developing performance parameters for IP networks that might include such functions in the future (e.g., setup and release of connection-oriented flows).

Recommendation Y.1540 defines four individual packet transfer outcomes based on IP packet transfer REs at the MPs. These are illustrated in simplified form in Fig. 2. An IP packet input to a section at an ingress MP will encounter one of three outcomes: *successful transfer*, *error*, or *loss*. An IP packet that appears at an egress MP with no corresponding input is said to be *spurious*. The IP packet transfer events and outcomes are defined more formally in Y.1540, taking into account global routing information and the possibility of packet fragmentation. Variable routing is addressed by defining, at a given time and relative to a given end-to-end IP flow and network portion, a set of *permissible* ingress and egress MPs. Packet fragmentation is addressed by considering explicitly, in the packet transfer outcome definitions, cases where an RE at one MP results in several corresponding events at other MPs.[4]

IP Packet Transfer Performance Parameters

Recommendation Y.1540 defines five IP packet transfer performance parameters on the basis of the outcomes shown in Fig. 2.

IP packet transfer delay (IPTD) is the time ($t_2 - t_1$) between the occurrence of two corresponding IP packet transfer reference events: an ingress event RE_1 at time t_1 and an egress event RE_2 at time t_2, where ($t_2 > t_1$) and ($t_2 - t_1$) ≤ T_{max}. IPTD is defined for all successful and errored packet transfer outcomes. If the packet is fragmented, t_2 is the time of the final corresponding egress event. *Mean IP packet transfer delay*, the parameter actually specified in Recommendation Y.1541, is the arithmetic average of IP packet transfer delays for a population of interest.

IP packet delay variation (IPDV) is defined based on observations of corresponding IP pack-

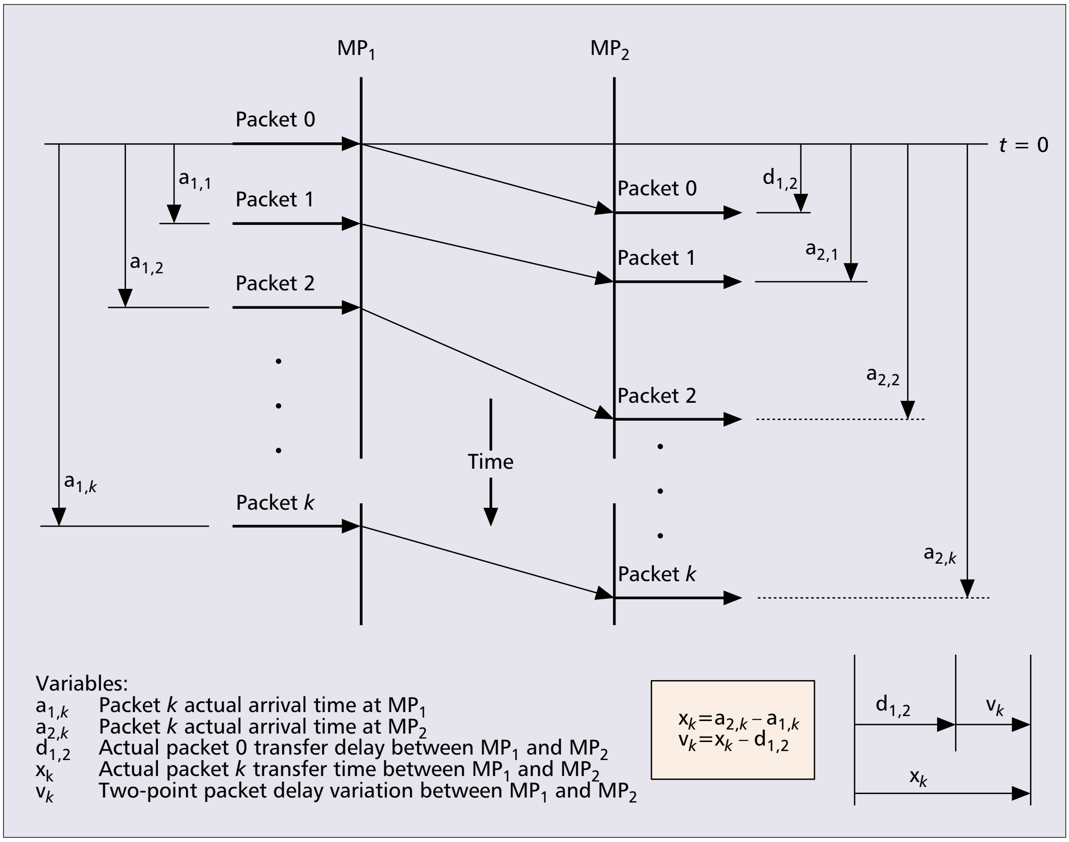

Figure 3. *IP packet delay variation.*

The IP packet transfer events and outcomes are defined more formally in Y.1540, taking into account global routing information and the possibility of packet fragmentation.

et arrivals at ingress and egress MPs (e.g., MP_1, MP_2 in Fig. 3). The packet delay variation (v_k) for an IP packet k between MP_1 and MP_2 is the difference between the absolute IP packet transfer delay (x_k) of the packet and a defined reference IP packet transfer delay, $d_{1,2}$, between those same MPs: $v_k = x_k - d_{1,2}$. The reference IP packet transfer delay, $d_{1,2}$, between SRC and DST is the absolute IP packet transfer delay experienced by the first IP packet between those two MPs.[5]

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

IP packet error ratio (IPER) 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.

Spurious IP packet rate (SIPR) at an egress MP is the total number of spurious IP packets observed at that egress MP during a specified time interval divided by the time interval duration (equivalently, the number of spurious IP packets per second). This parameter is expressed as a time rate rather than a ratio because the mechanisms that cause spurious IP packets have little to do with the number of IP packets transmitted.

Although not exhaustive, these parameters collectively address the major performance concerns of IP network users. IP packet transfer delay describes the average time a network takes to transfer packets between ingress and egress MPs. IPTD limits will be crucial to the successful deployment of VoIP, videoconferencing, and real-time data applications, and will strongly influence customer acceptance of others. IP packet delay variation characterizes jitter in the timing of packet transfer reference events at an egress interface with reference to the corresponding pattern of ingress events. IPDV must be controlled to avoid underflow or overflow in IP routers or terminal buffers. IP packet loss ratio expresses the likelihood that a packet entrusted to a network at an ingress interface is not delivered to the appropriate egress point(s). IPLR must be limited to ensure intelligibility and acceptable image quality in voice and real-time video applications, and to maintain reasonable efficiency in other applications. (Consecutive packet loss is of particular interest to certain nonelastic real-time applications such as voice and video. The parameter severe loss block ratio is one way to characterize such events.) IP packet error ratio and spurious IP packet rate express the likelihood that user data delivered at an egress interface differs from the input data as a result of corruption, duplication, or misrouting in the network.

Absent from the normative set of Y.1540 metrics is any parameter that describes the user data transfer rate, or *throughput*, a network portion provides. Y.1541 notes that throughput and other flow-related issues may be addressed using an IP network *traffic descriptor* defined in a companion Recommendation, Y.1221 [5]. A framework for possible future work in defining bulk transfer capacity metrics is defined in RFC 3148 [6].

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

ITU-T Recommendation Y.1541 specifies numerical values to be achieved, on international IP network paths between end user terminals, for each of the key performance parameters defined in Recommendation Y.1540.

AVAILABILITY PARAMETERS

As defined in Y.1540, availability applies to a unidirectional IP packet flow between a specified pair (or set) of MPs. The Y.1540 availability function specifies IPLR as the sole availability decision parameter. For a specified flow, a network portion is defined to be *available* over an observation period if the observed IPLR value for the flow is less than a specified threshold, c_1. Otherwise, the portion is *unavailable*. Y.1540 specifies a value for c_1 of 0.75, and notes that specifications of expected IPLR performance should exclude all periods of unavailability (i.e., all time intervals during which the observed IPLR exceeds c_1). The Recommendation defines a minimum availability observation period of 5 min.[6] The Y.1540 availability definition is intended to be usable in characterizing network performance for both normal traffic between a source and destination, and "synthetic" traffic generated by test sets or other measurement devices. Y.1540 notes that test traffic should be limited so that it does not cause congestion, which might bias the test results.

Recommendation Y.1540 defines two availability performance parameters. For a given network portion and flow, *percent availability* is the percentage of scheduled available time during which the portion actually maintains the flow in the available state. *Percent unavailability* is its complement, that is, the percentage of scheduled available time during which the flow is in the unavailable state. In any single specification, the two values add to 100 percent. The base period is limited to scheduled available time to exclude any agreed upon unavailable periods (e.g., planned downtime for preventive maintenance).

RECOMMENDATION Y.1541

ITU-T Recommendation Y.1541 specifies numerical values to be achieved, on international IP network paths between end user terminals, for each of the key performance parameters defined in Recommendation Y.1540. The specified values are grouped in a number of distinct QoS classes to establish a practical basis for communication between end users and network providers, and among providers, on the quality to be supported in end-to-end IP paths.

PERFORMANCE SPECIFICATION PROCESS

Defining numerical objectives is the key — and usually the most difficult — step in ITU-T's development of performance specifications for international communications. On one hand, network operators recognize that end-to-end performance is limited by the worst performing network portion, and user application requirements can only be met if each concatenated network is designed and operated with these requirements in mind. On the other hand, operators are aware that performance specifications strongly impact network economics, play a role in competitive product differentiation and marketing, and represent commitments users expect them to meet. It is a testament to ITU-T's effectiveness that these negotiations regularly succeed, enabling technologically diverse, independently operated networks to cooperate in providing reliable high-quality communications worldwide.

In defining end-to-end performance objectives for particular user applications and network technologies, ITU-T Study Group 13 has historically undertaken two complementary and concurrent performance assessment efforts. The first, "top-down" effort translates end user application requirements and quality expectations into numerical values for the ITU-T standardized parameters, *as observed at the user/network interfaces*. Top-down translation is done for each broad category of user application, and must address the variability in terminal functionality and performance. The second, "bottom-up" effort translates technical specifications defining the capabilities and limitations of individual network elements into numerical values for the same ITU-T standardized performance parameters, *observed at the same user/network interfaces*. Bottom-up translation is based on reference configurations that identify typical concatenations of network links and nodes, and worst case values for key variables such as the geographical distances between network interfaces. In shared networks, other variables such as network capacity, offered traffic, and resource management mechanisms must also be considered. Ideally, the top-down and bottom-up efforts produce overlapping ranges of values for the standardized parameters, from which one or more objectives can be specified.

The process of selecting numerical performance objectives was especially challenging in the case of Recommendation Y.1541. From a top-down perspective, the key challenge was to encompass the diverse set of user applications and terminals. Study Group 13 participants were able to bound and segment the application space by systematically considering, for each key application, the functional relationship between user satisfaction and the Y.1540 parameter values. The top-down analysis was greatly facilitated by close liaison with ITU-T Study Group 12, which has specialized for many years in relating the perceived quality and end user acceptance of particular applications and media (e.g., speech, image, text) with transmission network impairments, taking into account the performance of telephone, audiovisual, and interactive voice terminals. One powerful tool for establishing such relationships is the "E-model," described in [7]. Relevant national and regional standards were also considered (e.g., see [8]).

Study Group 13 received strong bottom-up input from network providers in characterizing the performance capabilities and limitations of available IP network elements and defining realistic end-to-end reference configurations. The top-down and bottom-up analyses quickly confirmed that no single set of IP network performance levels could economically support all envisioned applications for future IP-based networks; accordingly, Study Group 13 undertook to define several sets of performance objectives — the Y.1541 QoS classes.

The selection of QoS classes to be included in Y.1541 was debated in Study Group 13 Working Party 4 over several meetings. In early dis-

[6] *Monitoring of lower-layer performance and network element faults may allow identification of impending unavailability in a shorter time, and direct corrective action.*

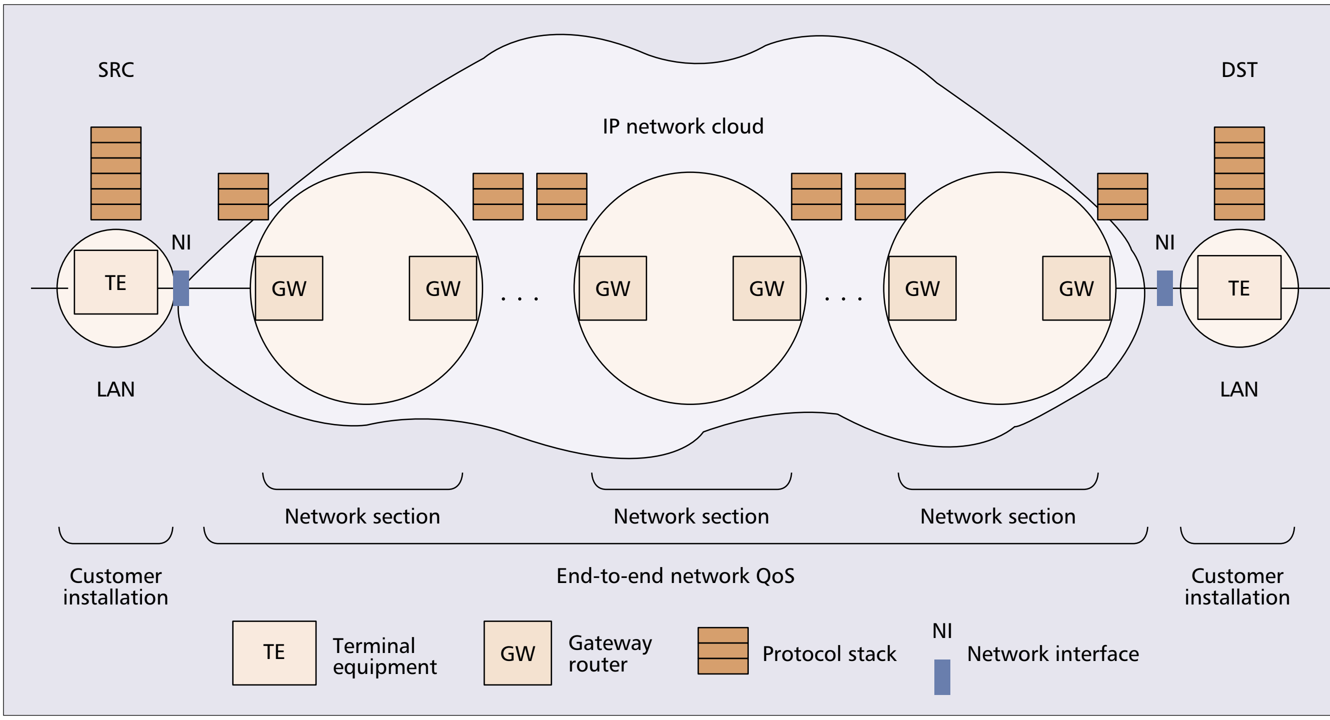

Figure 4. *The end-to-end reference path for QoS objectives.*

cussions, participants considered a per-parameter negotiation approach that would have enabled users to specify values for each parameter independently. It was quickly agreed that allowing continuous freedom of choice would be too complicated to implement. In fact, there was a strong consensus that the number of distinct QoS classes specified in Y.1541 should be strictly minimized to avoid overcomplicating the Recommendation (and, more important, the network technologies required to implement it). To provide widest coverage, the group agreed that the defined classes should collectively encompass a broad set of applications and a high percentage of user needs on emerging IP networks. In addition to the traditional Internet applications, these include point-to-point telephony, multimedia teleconferencing, and interactive data transfer (e.g., signaling). The group concluded that the needs of a few particularly demanding applications (e.g., high-resolution real-time video distribution, high-bandwidth TCP connections) would not be reflected in the standard classes at this time. It was agreed that each QoS class should encompass a group of applications whose performance requirements are similar, but significantly different from those of other classes. One heuristic the group used to limit the complexity of the QoS class structure was to ask, for each pair of proposed QoS classes, whether operators of managed IP networks were likely to do something different in implementing them. QoS classes were distinguished only if the answer to that question was yes.

Y.1541 Reference Path

The end-to-end IP performance objectives defined in Recommendation Y.1541 apply from NI to NI, as shown in Fig. 4. The end-to-end IP network path includes the set of network sections and exchange links that transport IP packets from SRC to DST; the protocols below and including the IP layer within the SRC and DST may also be considered part of the IP network. Network sections correspond with operator domains, and may include IP access network architectures. The customer installation includes all terminal equipment, such as a host and any on-premises router or LAN.

Performance Objectives and QoS Classes

The Recommendation Y.1541 performance objectives and QoS classes are specified in Table 1. Each QoS class creates a specific combination of bounds on a subset of the performance values. The classes and associated performance objectives apply to IP packet flows between the MPs that delimit the end-to-end IP network (i.e., the NIs shown in Fig. 4). An IP *packet flow* is the traffic associated with a given connection or connectionless stream having the same source host (SRC), destination host (DST), class of service, and session identification. Other documents may use the terms microflow or subflow when referring to traffic streams with this degree of classification.

Classes 0 and 1 place upper bounds on packet transfer delay and packet loss. They also limit packet delay variation. Classes 2 and 3 place upper bounds on packet transfer delay and packet loss, but do not limit packet delay variation. Classes 0 and 2 differ from classes 1 and 3 in their packet transfer delay objectives. Class 4 limits packet loss and provides a very soft upper bound on delay. Y.1541 also defines an unspecified class (class 5) that provides no specific performance guarantees. The value for the single packet error ratio objective was chosen to ensure that packet loss is the dominant cause of defects presented to upper layers. The QoS 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	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)	50 ms	50 ms	U	U	U	U
IPLR	Upper bound on the packet loss probability	$1*10^{-3}$	$1*10^{-3}$	$1*10^{-3}$	$1*10^{-3}$	$1*10^{-3}$	U
IPER	Upper bound	$1*10^{-4}$	$1*10^{-4}$	$1*10^{-4}$	$1*10^{-4}$	$1*10^{-4}$	U

Note: U = unspecified.

Table 1. *IP QoS class definitions and network performance objectives.*

are applicable when access link speeds are at the T1 or E1 rate or higher. The IPTD objectives of classes 0 and 2 will not always be achievable on long paths.

Y.1541 assumes that the user and network provider have agreed on a traffic profile that applies to one or more packet flows in a QoS class. At present, the agreeing parties may use whatever capacity specifications they consider appropriate as long as they allow both enforcement and verification. For example, peak bit rate (including lower layer overhead) may be sufficient. When protocols and systems supporting dynamic requests are available, users may negotiate a *traffic contract* that specifies one or several traffic parameters in accordance with Recommendation Y.1221.

Networks offering IP communications in accordance with Y.1541 are expected to support these end-to-end bounds for the lifetime of an established flow as long as the users (and other networks) do not exceed the agreed capacity. Y.1541 stipulates that networks complying with Y.1541 are not required to support agreed QoS values if the specified capacity is exceeded. A network observing such excess flow may discard a number of packets equal to the number of excess packets. Such discarded packets are not counted as lost packets in assessing the network's IPLR performance.

In addition to the performance objectives and QoS classes, Recommendation Y.1541 specifies various ancillary variables (minimum observation periods, test packet lengths, sample sizes, etc.) to facilitate performance estimation and comparison. As an example, a minimum evaluation interval of 10–20 s is recommended to assess VoIP at typical packet rates (50–100 packets/s). The recommended evaluation interval for loss, delay, and IPDV is 1 min, striking a balance between statistical confidence and relevance to user experience.

Table 2 (from Y.1541) provides guidance on the applicability and engineering of the QoS classes. Y.1541 notes that these guidelines are completely discretionary; network providers may use whatever node mechanisms, routing constraints, or other techniques they choose.

Conclusions and Further Work

ITU-T Recommendations Y.1540 and Y.1541 together provide a key piece of the IP QoS puzzle. Y.1540 defines standard performance param-

QoS class	Applications (examples)	Node mechanisms	Network techniques
0	Real-time, jitter-sensitive, high interaction (VoIP, video teleconferencing)	Separate queue with preferential servicing, traffic grooming	Constrained routing and distance
1	Real-time, jitter-sensitive, interactive (VoIP, video teleconferencing).		Less constrained routing and distance
2	Transaction data, highly interactive (e.g., signaling)	Separate queue, drop priority	Constrained routing and distance
3	Transaction data, interactive		Less constrained routing and distance
4	Low loss only (short transactions, bulk data, video streaming)	Long queue, drop priority	Any route/path
5	Traditional applications of default IP networks	Separate queue (lowest priority)	Any route/path

Table 2. *Guidance for IP QoS classes.*

eters for packet transfer in IP-based networks. Y.1541 specifies NI–NI objectives for the Y.1540 parameters and clusters these numerical performance objectives in six distinct IP network QoS classes. As a set, the Y.1541 classes encompass the major IP user application categories. They are relatable to IP network QoS mechanisms that are implementable. The performance values they specify can be achieved in practical networks, and can be verified at jurisdictional boundaries by instrumenting terminal equipment or interworking functions. They document an important agreement among network providers, equipment manufacturers, and end users on the quality levels that will need to be supported to provide assured quality to a wide range of IP applications, including telephony. They can be used as a basis for QoS negotiation among networks. They may also meet the need for a *lingua franca* to support QoS interworking among different technologies.

Although Y.1540/Y.1541 represent a useful step forward, the successful evolution of an IP-based next-generation network supporting a dynamic set of specific QoS classes is not ensured. QoS mechanisms are not widely deployed in IP-based networks today. Although static QoS class agreements could be implemented today by associating packet markings (e.g., TOS or DiffServ code points) with specific QoS classes, work is still needed to define a more flexible QoS architecture [9] and identify how to implement the Y.1541 QoS classes in signaling protocols [10]. Providers will need to define, and probably standardize, a means of apportioning performance objectives among the several independent networks that typically will interoperate in providing QoS-assured IP flows between end user terminals. In short, continuing IP/PSTN convergence will require a convergence of thought and action regarding IP network QoS. ITU-T Study Group 13 and other standards organizations are working to achieve that goal.

References

[1] ITU-T Rec. Y.1540, "IP Packet Transfer and Availability Performance Parameters," Dec. 2002.

[2] ITU-T Rec. Y.1541, "Network Performance Objectives for IP-Based Services," May 2002.

[3] L. Ciavattone, A. Morton, and G. Ramachandran, "Standardized Active Measurements on a Tier 1 IP Backbone," *IEEE Commun. Mag.*, this issue.

[4] ITU-T Rec. I.350, "General Aspects of Quality of Service and Network Performance in Digital Networks, Including ISDNs," Mar. 1993.

[5] ITU-T Rec. Y.1221, "Traffic Control and Congestion Control in IP Based Networks," Mar. 2002.

[6] M. Mathis and M. Allman, "A Framework for Defining Empirical Bulk Transfer Capacity Metrics," IETF RFC 3148, July 2001.

[7] M. Perkins *et al.*, "Speech Transmission Performance Planning in Hybrid IP/SCN Networks," *IEEE Commun. Mag.*, July 1999, pp.126–31.

[8] "Quality of Service for Business Multimedia Conferencing," ANS T1.522, Apr. 2000.

[9] H.-L. Lu and I. Faynberg, "An Architectural Framework for Support of QoS In Packet Networks," *IEEE Commun. Mag.*, this issue.

[10] ITU-T Study Group 11 draft TRQ.IPQoS.SIG.CS1, "Signaling Requirements for IP-QoS," Nov. 2002.

Biography

NEAL SEITZ (neal@its.bldrdoc.gov) is a senior engineer at the Institute for Telecommunication Sciences (ITS), the telecommunications research and engineering arm of the U.S. Department of Commerce's National Telecommunications and Information Administration (NTIA). He has held leadership positions in telecommunication performance standards committees for over 25 years. He currently chairs ITU-T Study Group 13 Working Party 4, which develops ITU-T Recommendations on network performance and resource management for multiprotocol and IP-based networks and their interworking. He has participated in the leadership of ANSI-accredited Standards Committee T1 (Telecommunications) Subcommittee T1A1 (Performance, Reliability, and Signal Processing) since 1984. He has directed and contributed to the development of over a dozen national and international telecommunication performance standards, including user-oriented technology-independent QoS standards for digital services, and provider-oriented technology-specific performance standards for packet switching, frame relay, ISDN, B-ISDN, ATM, and IP-based networks. He holds two U.S. patents for innovations in multichannel speech compression technology.

Although Y.1540/Y.1541 represent a useful step forward, the successful evolution of an IP-based next-generation network supporting a dynamic set of specific QoS classes is not ensured.