

International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Y.1540

(11/2007)

SERIES Y: GLOBAL INFORMATION
INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS
AND NEXT-GENERATION NETWORKS

Internet protocol aspects – Quality of service and network
performance

**Internet protocol data communication service –
IP packet transfer and availability performance
parameters**

ITU-T Recommendation Y.1540



ITU-T Y-SERIES RECOMMENDATIONS
**GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT-
GENERATION NETWORKS**

GLOBAL INFORMATION INFRASTRUCTURE	
General	Y.100–Y.199
Services, applications and middleware	Y.200–Y.299
Network aspects	Y.300–Y.399
Interfaces and protocols	Y.400–Y.499
Numbering, addressing and naming	Y.500–Y.599
Operation, administration and maintenance	Y.600–Y.699
Security	Y.700–Y.799
Performances	Y.800–Y.899
INTERNET PROTOCOL ASPECTS	
General	Y.1000–Y.1099
Services and applications	Y.1100–Y.1199
Architecture, access, network capabilities and resource management	Y.1200–Y.1299
Transport	Y.1300–Y.1399
Interworking	Y.1400–Y.1499
Quality of service and network performance	Y.1500–Y.1599
Signalling	Y.1600–Y.1699
Operation, administration and maintenance	Y.1700–Y.1799
Charging	Y.1800–Y.1899
NEXT GENERATION NETWORKS	
Frameworks and functional architecture models	Y.2000–Y.2099
Quality of Service and performance	Y.2100–Y.2199
Service aspects: Service capabilities and service architecture	Y.2200–Y.2249
Service aspects: Interoperability of services and networks in NGN	Y.2250–Y.2299
Numbering, naming and addressing	Y.2300–Y.2399
Network management	Y.2400–Y.2499
Network control architectures and protocols	Y.2500–Y.2599
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899

For further details, please refer to the list of ITU-T Recommendations.

ITU-T Recommendation Y.1540

Internet protocol data communication service – IP packet transfer and availability performance parameters

Summary

ITU-T Recommendation Y.1540 defines parameters that may be used in specifying and assessing the performance of speed, accuracy, dependability and availability of IP packet transfer of international Internet protocol (IP) data communication service. The defined parameters apply to end-to-end, point-to-point IP service and to the network portions that provide, or contribute to the provision of, such service in accordance with the normative references specified in clause 2. Connectionless transport is a distinguishing aspect of the IP service that is considered in this Recommendation.

Source

ITU-T Recommendation Y.1540 was approved on 13 November 2007 by ITU-T Study Group 12 (2005-2008) under the ITU-T Recommendation A.8 procedure.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure e.g. interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at <http://www.itu.int/ITU-T/ipr/>.

© ITU 2008

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

CONTENTS

	Page
1 Scope	1
2 References.....	3
3 Abbreviations.....	3
4 Layered model of performance for IP service	5
5 Generic IP service performance model.....	6
5.1 Network components.....	6
5.2 Exchange links and network sections.....	7
5.3 Measurement points and measurable sections.....	8
5.4 IP packet transfer reference events (IPREs).....	9
5.5 IP packet transfer outcomes.....	10
6 IP packet transfer performance parameters	15
7 IP service availability	20
7.1 IP service availability function.....	20
7.2 IP service availability parameters.....	22
Appendix I – IP packet routing considerations	23
Appendix II – Secondary terminology for IP packet delay variation	24
II.1 Introduction	24
II.2 Definition of inter-packet delay variation	24
II.3 Definition of 1-point packet delay variation	24
II.4 Guidance on applying the different parameters.....	25
Appendix III – Rate and throughput capacity related parameters	26
III.1 Definition of IP packet rate parameters.....	26
III.2 References for throughput parameters and measurements.....	26
III.3 Open issues.....	26
Appendix IV – Minimal test of IP service availability state and sampling estimation of IP service availability parameters	28
IV.1 Minimal test of IP service availability state (for test methodologies and test sets)	28
IV.2 Sampling estimation of IP service availability.....	28
Appendix V – Material relevant to IP performance measurement methods.....	29
Appendix VI – Background on IP service availability	30
VI.1 Introduction	30
VI.2 Background.....	30
VI.3 Definitions of the regions in Figure VI.1	31
VI.4 Summary.....	31
Bibliography.....	32

ITU-T Recommendation Y.1540

Internet protocol data communication service – IP packet transfer and availability performance parameters

1 Scope

This Recommendation defines parameters that may be used in specifying and assessing the performance of speed, accuracy, dependability and availability of IP packet transfer of international Internet protocol (IP) data communication service. The defined parameters apply to end-to-end, point-to-point IP service and to the network portions that provide, or contribute to the provision of, such service in accordance with the normative references specified in clause 2. Connectionless transport is a distinguishing aspect of the IP service that is considered in this Recommendation.

For the purpose of this Recommendation, end-to-end IP service refers to the transfer of user-generated IP datagrams (referred to in this Recommendation as IP packets) between two end hosts as specified by their complete IP addresses. This differs from the boundaries implied by the phrase "end-to-end" in some other Recommendations. For example, [ITU-T P.10] defines end-to-end quality as related to the performance of a communication system, including all terminal equipment. For voice services, end-to-end is equivalent to mouth-to-ear quality.


NOTE 1 – This Recommendation defines parameters that can be used to characterize IP service provided using IPv4 and IPv6; applicability or extension of this Recommendation to other protocols (e.g., RSVP) is for further study.

NOTE 2 – Recommendations for the performance of point-to-multipoint IP service are currently under development.

The Y.1540 performance parameters are intended to be used in planning and offering international IP service. The intended users of this Recommendation include IP service providers, equipment manufacturers and end users. This Recommendation may be used by service providers in the planning, development and assessment of IP service that meets user performance needs; by equipment manufacturers as performance information that will affect equipment design; and by end users in evaluating IP service performance.

The scope of this Recommendation is summarized in Figure 1. The IP service performance parameters are defined on the basis of IP packet transfer reference events that may be observed at measurement points (MPs) associated with specified functional and jurisdictional boundaries. For comparability and completeness, IP service performance is considered in the context of the 3×3 performance matrix defined in [ITU-T I.350]. Three protocol-independent communication functions are identified in the matrix: access, user information transfer and disengagement. Each function is considered with respect to three general performance concerns (or "performance criteria"): speed, accuracy and dependability. An associated two-state model provides a basis for describing IP service availability.

NOTE 3 – In this Recommendation, the user information transfer function illustrated in Figure 1 refers to the attempted transfer of any IP packet, regardless of its type or contents.



Y.1540(07)_F01

Figure 1 – Scope of this Recommendation

The performance parameters defined in this Recommendation describe the speed, accuracy, dependability and availability of IP packet transfer as provided by IP data communication service. Future ITU-T Recommendations may be developed to provide standard methods of measuring the Y.1540 performance parameters in an international context. The end-to-end performance of international IP services providing access and disengagement functions (e.g., domain name service) and higher-layer transport capabilities (e.g., transmission control protocol) may be addressed in separate Recommendations.

This Recommendation is structured as follows: Clause 1 specifies its scope. Clause 2 specifies its normative references. Clause 3 provides a list of abbreviations. Clause 4 illustrates the layered model that creates the context for IP performance specification. Clause 5 defines the model used for IP performance, including network sections and measurement points, reference events and outcomes. Clause 6 uses this model to define IP packet transfer performance parameters. Clause 7 then defines IP service availability parameters. Appendix I describes IP packet routing considerations and their effects on performance. Appendix II provides secondary terminology for IP packet delay variation. Appendix III describes some possible metrics for IP packet rate and references material for assessing the throughput and throughput capacity of IP service. Appendix IV describes estimation of IP service availability. Appendix V presents considerations for measuring the Y.1540 parameters. Appendix VI gives some background on IP service availability.

NOTE 4 – The Y.1540 parameters may be augmented or modified based upon further study of the requirements of the IP applications (e.g., interactive, block, stream) to be supported.

NOTE 5 – The Y.1540 speed, accuracy and dependability parameters are intended to characterize IP service in the available state.

NOTE 6 – The parameters defined in this Recommendation can apply to a single end-to-end IP service between two end hosts identified by their IP addresses. The parameters can also be applied to those IP packets from a given end-to-end IP service that are offered to a given network or exchange link.

NOTE 7 – The Y.1540 parameters are designed to characterize the performance of service provided by network elements between specified section boundaries. However, users of this Recommendation should be aware that network elements outside the specified boundaries can sometimes influence the measured performance of the elements between the boundaries. Examples are described in Appendix V.

NOTE 8 – The parameters defined in this Recommendation can also be applied to any subset of the IP packets offered to a given set of network equipment. Methods for aggregating performance over a set of network equipment or over an entire network are outside of the scope of this Recommendation.

NOTE 9 – This Recommendation does not provide the tools for explicit characterization of routing stability. However, the effects of route instability can be quantified using the loss, delay and severe loss block parameters defined in this Recommendation.

NOTE 10 – Specification of numerical performance objectives for some of the Y.1540 performance parameters may be found in [ITU-T Y.1541].

NOTE 11 – The word "provisional", as used in this Recommendation, means that there is agreement on the stability of the value referenced, but that the value may change following further study, or on the basis of real network operational experience.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

- [ITU-T I.350] ITU-T Recommendation I.350 (1993), *General aspects of quality of service and network performance in digital networks, including ISDNs*.
- [ITU-T P.10] ITU-T Recommendation P.10/G.100 (2006), *Vocabulary for performance and quality of service*.
- [ITU-T Y.1541] ITU-T Recommendation Y.1541 (2006), *Network performance objectives for IP-based services*.
- [IETF RFC 791] IETF RFC 791 (1981), *Internet Protocol*.
<<http://www.ietf.org/rfc/rfc791.txt>>
- [IETF RFC 2460] IETF RFC 2460 (1998), *Internet Protocol, Version 6 (IPv6) Specification*.
<<http://www.ietf.org/rfc/rfc2460.txt>>
- [IETF RFC 4737] IETF RFC 4737 (2006), *Packet Reordering Metrics*.
<<<http://www.ietf.org/rfc/rfc4737.txt>>>

3 Abbreviations

This Recommendation uses the following abbreviations:

- ATM Asynchronous Transfer Mode
BTC Bulk Transfer Capacity

DSCP	Differentiated Services Code Point
DST	Destination host
EL	Exchange Link
ER	Edge Router
FTP	File Transfer Protocol
HTTP	Hypertext Transfer Protocol
IP	Internet Protocol
IPDV	Internet Protocol packet Delay Variation
IPDR	Internet Protocol packet Duplicate Ratio
IPER	Internet Protocol packet Error Ratio
IPLR	Internet Protocol packet Loss Ratio
IPOR	Octet-based IP packet Rate
IPPR	Internet Protocol Packet rate
IPRE	Internet Protocol packet transfer Reference Event
IPRR	Internet Protocol packet Reordered Ratio
IPSLBR	Internet Protocol packet Severe Loss Block Ratio
IPTD	Internet Protocol packet Transfer Delay
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISP	Internet Service Provider
LL	Lower Layers (protocols and technology supporting the Internet Protocol layer)
M_{av}	The minimum number of packets recommended for assessing the availability state
MP	Measurement Point
MTBISO	Mean Time Between IP Service Outages
MTTISR	Mean Time To Internet Protocol Service Restoral
N	The number of packets in a throughput probe of size N
NS	Network Section
NSE	Network Section Ensemble
NSP	Network Service Provider
PDH	Plesiochronous Digital Hierarchy
PDV	Packet Delay Variation
PIA	Percent Internet Protocol service Availability
PIU	Percent Internet Protocol service Unavailability
pkt	Internet Protocol datagram (IP packet)
QoS	Quality of Service
R	Router
RFC	Request for Comments

RIPR	Replicated Internet Protocol Packet Ratio
RSVP	Resource reSerVation Protocol
RTP	Real-time Transport Protocol
SDH	Synchronous Digital Hierarchy
SRC	Source host
STD	Standard
T_{av}	Minimum length of time of Internet Protocol availability; minimum length of time of Internet Protocol unavailability
TCP	Transmission Control Protocol
T_{max}	Maximum Internet Protocol packet delay beyond which the packet is declared to be lost
ToS	Type of Service
T_s	Length of time defining the block in the severe loss block outcome
TTL	Time To Live
UDP	User Datagram Protocol

4 Layered model of performance for IP service

Figure 2 illustrates the layered nature of the performance of IP service. The performance provided to IP service users depends on the performance of other layers:

- Lower layers that provide (via "links") connection-oriented or connectionless transport supporting the IP layer. Links are terminated at points where IP packets are forwarded (i.e., "routers", "SRC" and "DST") and thus have no end-to-end significance. Links may involve different types of technologies, for example, ATM, frame relay, SDH, PDH, ISDN and leased lines. There may be several layers of protocols and services below the IP layer, and these, in the end, make use of various types of physical media.
- The IP layer that provides connectionless transport of IP datagrams (i.e., IP packets). The IP layer has end-to-end significance for a given pair of source and destination IP addresses. Certain elements in the IP packet headers may be modified by networks, but the IP user data may not be modified at or below the IP layer.
- Higher layers, supported by IP, that further enable end-to-end communications. Upper layers may include, for example, TCP, UDP, FTP, RTP and HTTP. The higher layers will modify and may enhance the end-to-end performance provided at the IP layer.

NOTE 1 – Clause 5 defines an IP service performance model and more precisely defines key terms used in this layered model.

NOTE 2 – Performance interactions among these layers are for further study.




Figure 2 – Layered model of performance for IP service – Example

5 Generic IP service performance model

This clause defines a generic IP service performance model. The model is primarily composed of two types of sections: the exchange link and the network section. These are defined in clause 5.2. They provide the building blocks with which any end-to-end IP service may be represented. Each of the performance parameters defined in this Recommendation can be applied to the unidirectional transfer of IP packets on a section or a concatenated set of sections.

Clause 5.4 specifies the set of IP packet transfer reference events that provide the basis for performance parameter definition. These reference events are derived from and are consistent with relevant IP service and protocol definitions. Clause 5.5 then uses those reference events to enumerate the possible outcomes when a packet is delivered into a section.

NOTE – Incorporation of all or part of the Y.1540 performance model and reference events into [b-ITU-T I.353] is for further study.

5.1 Network components

5.1.1 host: A computer that communicates using the Internet protocols. A host implements routing functions (i.e., it operates at the IP layer) and may implement additional functions including higher layer protocols (e.g., TCP in a source or destination host) and lower layer protocols (e.g., ATM).

5.1.2 router: A host that enables communication between other hosts by forwarding IP packets based on the content of their IP destination address field.


5.1.3 source host (SRC): A host and a complete IP address where end-to-end IP packets originate. In general, a host may have more than one IP address; however, a source host is a unique association with a single IP address. Source hosts also originate higher layer protocols (e.g., TCP) when such protocols are implemented.

5.1.4 destination host (DST): A host and a complete IP address where end-to-end IP packets are terminated. In general, a host may have more than one IP address; however, a destination host is a unique association with a single IP address. Destination hosts also terminate higher layer protocols (e.g., TCP) when such protocols are implemented.

5.1.5 link: A point-to-point (physical or virtual) connection used for transporting IP packets between a pair of hosts. It does not include any parts of the hosts or any other hosts; it operates

below the IP layer. For example, a link could be a leased line or it could be implemented as a logical connection over an Ethernet, a frame relay network, an ATM network, or any other network technology that functions below the IP layer.

Figure 3 illustrates the network components relevant to IP service between a SRC and a DST. Links, which could be dial-up connections, leased lines, rings, or networks are illustrated as lines between hosts. Routers are illustrated as circles and both SRC and DST are illustrated as triangles.



Y.1540(07)_F03

Figure 3 – IP network components

5.2 Exchange links and network sections

5.2.1 exchange link (EL): The link connecting:

- 1) a source or destination host to its adjacent host (e.g., router) possibly in another jurisdiction, sometimes referred to as an access link, ingress link or egress link; or
- 2) a router in one network section with a router in another network section.

Note that the responsibility for an exchange link, its capacity, and its performance, is typically shared between the connected parties.

NOTE – "Exchange link" is roughly equivalent to the term "exchange" as defined in [b-IETF RFC 2330].

5.2.2 network section (NS): A set of hosts together with all of their interconnecting links that together provide a part of the IP service between a SRC and a DST, and are under a single (or collaborative) jurisdictional responsibility. Some network sections consist of a single host with no interconnecting links. Source NS and destination NS are particular cases of network sections. Pairs of network sections are connected by exchange links.

NOTE – "Network section" is roughly equivalent to the term "cloud" as defined in [b-IETF RFC 2330].

Any set of hosts interconnected by links could be considered a network section. However, for the (future) purpose of IP performance allocation, it will be relevant to focus on the set of hosts and links under a single (or collaborative) jurisdictional responsibility (such as an ISP or an NSP). These hosts typically have the same network identifier in their IP addresses. Typically, they have their own rules for internal routing. Global processes and local policies dictate the routing choices to destinations outside of this network section (to other NS via exchange links). These network sections are typically bounded by routers that implement the IP exterior gateway protocols.

5.2.3 source NS: The NS that includes the SRC within its jurisdictional responsibility. In some cases, the SRC is the only host within the source NS.

5.2.4 destination NS: The NS that includes the DST within its jurisdictional responsibility. In some cases, the DST is the only host within the destination NS.

Figure 4 illustrates the network connectivity relevant to IP service between a SRC and a DST. At the edges of each NS, gateway routers receive and send packets across exchange links.




Figure 4 – IP network connectivity

5.3 Measurement points and measurable sections

5.3.1 measurement point (MP): The boundary between a host and an adjacent link at which performance reference events can be observed and measured. Consistent with [b-ITU-T I.353], the standard Internet protocols can be observed at IP measurement points. [b-ITU-T I.353] provides more information about MPs, for digital services.

NOTE – The exact location of the IP service MP within the IP protocol stack is for further study.

A section or a combination of sections is measurable if it is bounded by a set of MPs. In this Recommendation, the following sections are measurable.

5.3.2 basic section: Either an EL, an NS, a SRC or a DST. Basic sections are delimited by MPs.

The performance of any EL or NS is measurable relative to any given unidirectional end-to-end IP service. The *ingress MPs* are the set of MPs crossed by packets from that service as they go into that basic section. The *egress MPs* are the set of MPs crossed by packets from that service as they leave that basic section.

5.3.3 end-to-end IP network: The set of ELs and NSs that provide the transport of IP packets transmitted from SRC to DST. The MPs that bind the end-to-end IP network are the MPs at the SRC and the DST.

The end-to-end IP network performance is measurable relative to any given unidirectional end-to-end IP service. The *ingress MPs* are the MPs crossed by packets from that service as they go into the end-to-end network at the SRC. The *egress MPs* are the MPs crossed by packets from that service as they leave the end-to-end network at the DST.


5.3.4 network section ensemble (NSE): An NSE refers to any connected subset of NSs together with all of the ELs that interconnect them. The term "NSE" can be used to refer to a single NS, two

NSs, or any number of NSs and their connecting ELs. Pairs of distinct NSEs are connected by exchange links. The term "NSE" can also be used to represent the entire end-to-end IP network. NSEs are delimited by MPs.

The performance of any given NSE is measurable relative to any given unidirectional end-to-end IP service. The *ingress MPs* are the set of MPs crossed by packets from that service as they go into that NSE. The *egress MPs* are the set of MPs crossed by packets from that service as they leave that NSE.

5.4 IP packet transfer reference events (IPREs)

In the context of this Recommendation, the following definitions apply on a specified end-to-end IP service. The defined terms are illustrated in Figure 5.



Y.1540(07)_F05

NOTE 1 – IP exit events for packets A and C.

NOTE 2 – IP entry events for packets B and D.

Figure 5 – Example IP packet transfer reference events

An IP packet transfer event occurs when:

- an IP packet crosses a measurement point (MP); and
- standard IP procedures applied to the packet verify that the header checksum is valid; and
- the source and destination address fields within the IP packet header represent the IP addresses of the expected SRC and DST.

NOTE – The IP packet header contains information including type of service (ToS) or differentiated services code point (DSCP). How such information may affect packet transfer performance is for further study.

IP packet transfer reference events are defined without regard to packet fragmentation. They occur for every IP packet crossing any MP regardless of the value contained in the "more-fragments flag".

Four types of IP packet transfer events are defined:

5.4.1 IP packet entry event into a host: An IP packet transfer entry event into a host occurs when an IP packet crosses an MP entering a host (NS router or DST) from the attached EL.

5.4.2 IP packet exit event from a host: An IP packet transfer exit event from a host occurs when an IP packet crosses an MP exiting a host (NS router or SRC) into the attached EL.

5.4.3 IP packet ingress event into a basic section or NSE: An IP packet transfer ingress into a basic section or NSE event occurs when an IP packet crosses an ingress MP into a basic section or an NSE.

5.4.4 IP packet egress event from a basic section or NSE: An IP packet transfer egress event from a basic section or NSE occurs when an IP packet crosses an egress MP out of a basic section or an NSE.


NOTE 1 – IP packet entry and exit events always represent, respectively, entry into and exit from a host. IP packet ingress events and egress events always represent ingress into and egress from a section or an NSE. To illustrate this point, note that an ingress into an EL creates an exit event from the preceding host, while an ingress into an NS is an entry event because, by definition, NSs always have hosts at their edges.

NOTE 2 – For practical measurement purposes, IP packet transfer reference events need not be observed within the IP protocol stack of the host. Instead, the time of occurrence of these reference events can be approximated by observing the IP packets crossing an associated physical interface. This physical interface should, however, be as near as possible to the desired MP. In cases where reference events are monitored at a physical interface, the time of occurrence of an exit event from a host is approximated by the observation of the first bit of the IP packet coming from the host or test equipment. The time of occurrence of an entry event into a host is approximated by the observation of the last bit of the IP packet going to the host or test equipment.

5.5 IP packet transfer outcomes

By considering IP packet transfer reference events, a number of possible IP transfer outcomes may be defined for any packet attempting to cross a basic section or an NSE. A transmitted IP packet is either *successfully transferred*, *errored* or *lost*. A delivered IP packet for which no corresponding IP packet was offered is said to be *spurious*. Figure 6 illustrates the IP packet transfer outcomes.

The definitions of IP packet transfer outcomes are based on the concepts of *permissible ingress MP*, *permissible egress MP* and *corresponding packets*.



NOTE – Outcome occurs independent of IP packet contents.

Figure 6 – IP packet transfer outcomes

5.5.1 Global routing information and permissible output links

In theory, in a connected IP network, a packet can be delivered to any router, NS or NSE, and still arrive at its destination. However, global routing information defines a restricted set of destination addresses that each network (autonomous system) is willing and able to serve on behalf of each of its adjoining NS. It is reasonable to assume that (in the worst case) an NS will completely discard any packets with destination addresses for which that NS has announced an inability (or an unwillingness) to serve. Therefore all IP packets (and fragments of packets) leaving a basic section should only be forwarded to other basic sections as *permitted* by the available global routing information.

For performance purposes, the transport of an IP packet by an NSE will be considered successful only when that NSE forwards all of the packet contents to other basic sections as permitted by the currently available global routing information. If the destination address corresponds to a host attached directly to this NSE, the only permitted output and the only successful IP transport is a forwarding to the destination host.

NOTE 1 – IP procedures include updating of global routing information. An NS that was permissible may no longer be permissible following an update of the routing information shared between NSs. Alternatively, an

NS that was not previously permissible may have become permissible after an update of the global routing information.

NOTE 2 – Routing information can be supplemented by information about the relative suitability of each of the permitted output links. The performance implications of that additional information are for further study.

At a given time, and relative to a given end-to-end IP service and a basic section or NSE:


- an ingress MP is a *permissible ingress MP* if the crossing of this MP into this basic section or NSE is permitted by the global routing information;
- an egress MP is a *permissible egress MP* if the crossing of this MP leads into another basic section that is permitted by the global routing information.

5.5.2 Corresponding events

Performance analysis makes it necessary to associate the packets crossing one MP with the packets that crossed a different MP. Connectionless routing means a packet may leave a basic section on any one of (possibly) several permissible egress MP. Packet fragmentation means that a packet going into a basic section may leave in fragments, possibly into several different other basic sections. Finally, connectionless IP routing may even send a packet or a fragment back into a basic section it has already traversed (possibly due to the updating of routing tables).

An IP egress event is said to *correspond* to an earlier ingress event if they were created by the "same" IP packet. This concept applies whether the packet at the egress MP is the whole packet or just a fragment of the original. Figure 7 illustrates a case where a packet goes into NS C from NS B and is fragmented into two parts in NS C. One of the fragments is sent to NS D and the other to NS F. Both of these egress events *correspond* to the single ingress event. To avoid confusion resulting from packets re-entering the NSE, this concept of *correspondence* also requires that this be the first time (since its ingress) this particular content has departed from the NSE.

The practical determination of whether IP reference events are corresponding is usually *ad hoc* and will often rely on consideration of the IP addresses, the global routing information, the IP packet identification field, other header information and the IP packet contents.



An IP packet from SRC to DST enters NS C, creates an ingress event, is fragmented, and creates two corresponding egress events, b) and c).

Y.1540(07)_F07

Figure 7 – Corresponding events when fragmentation occurs

5.5.3 Notes about the definitions of successful, errored, lost and spurious packet outcomes

Each of the following definitions of individual packet outcomes is based on observing IP reference events at IP measurement points. By selecting the appropriate IP measurement points, each definition can be used to evaluate the performance of a particular EL, a particular NS, a particular NSE, and they can be applied to the performance of end-to-end services.

These outcomes are defined without restriction to a particular packet type (ToS, DSCP, protocol, etc.). IP performance will differ by packet type.

In each definition, the possibility of packet fragmentation is accounted for by including the possibility that a single IP reference event could result in several subsequent events. Note that if any fragment is lost, the whole original packet is considered lost. If no fragments are lost, but some are errored, the entire original packet is considered errored. For the delivery of the original packet to be considered successful, each fragment must be successfully delivered to one of the permissible output ELs.

5.5.4 successful IP packet transfer outcome: A successful packet transfer outcome occurs when a single IP packet reference event at a permissible ingress MP_0 results in one (or more) corresponding reference event(s) at one (or more) egress MP_i , all within a specified time T_{max} of the original ingress event and:

- 1) all egress MP_i where the corresponding reference events occur are permissible; and
- 2) the complete contents of the original packet observed at MP_0 are included in the delivered packet(s); and
- 3) the binary contents of the delivered IP packet information field(s) conform exactly with that of the original packet; and
- 4) the header field(s) of the delivered packet(s) is (are) valid.

NOTE – The value of T_{max} is recommended to be set at 3 seconds for general use. Some global end-to-end paths may require a larger value of T_{max} to ensure that packets with long transfer times have adequate opportunity to arrive. The value of 3 seconds has been used in practice.

5.5.5 errored IP packet outcome: An errored packet outcome occurs when a single IP packet reference event at a permissible ingress MP_0 results in one (or more) corresponding reference event(s) at one (or more) egress MP_i , all within T_{max} time of the original reference event and:

- 1) all egress MP_i where the corresponding reference events occur are permissible; and
- 2) the complete contents of the original packet observed at MP_0 are included in the delivered packet(s); and
- 3) either:
 - the binary contents of the delivered IP packet information field(s) do not conform exactly with that of the original packet; or
 - one or more of the header field(s) of the delivered packet(s) is (are) corrupted.

NOTE – Most packets with errored headers that are not detected by the header checksum at the IP layer will be discarded or redirected by other IP layer procedures (e.g., based on corruption in the address or ToS/DSCP fields). The result is that no reference event is created for the higher layer protocols expecting to receive this packet. Because there is no IP reference event, these packet transfer attempts will be classified as lost packet outcomes. Errored headers that do not result in discarding or misdirecting will be classified as errored packet outcomes.

5.5.6 lost IP packet outcome: The definition of a lost IP packet outcome is predicated on a definition for a *misdirected packet*.

A misdirected packet occurs when a single IP packet reference event at a permissible ingress MP_0 results in one (or more) corresponding reference event(s) at one (or more) egress MP_i , all within a specified T_{max} time of the original reference event and:

- 1) the complete contents of the original packet observed at MP_0 are included in the delivered packet(s); but
- 2) one or more of the egress MP_i where the corresponding reference events occur is (are) not permissible egress $MP(s)$.

A lost packet outcome occurs when a single IP packet reference event at a permissible ingress MP_0 results only in one or more misdirected packet outcomes, or when some or all of the contents of that packet do not result in any IP reference event at any egress MP within the time T_{max} .

5.5.7 spurious IP packet outcome: A spurious IP packet outcome occurs for a basic section, an NSE, on an end-to-end IP service when a single IP packet creates an egress event for which there was no corresponding ingress event.

5.5.8 secondary IP packet outcomes: The following outcomes are based on the fundamental outcomes described above.

5.5.8.1 in-order and reordered IP packet outcomes: The definition of these IP packet outcomes requires some background discussion.

In-order packet delivery is a property of successful packet transfer attempts, where the sending packet order is preserved on arrival at the destination host (or measurement point). Arrival order is determined by position relative to other packets of interest, though the extent to which a given packet has been reordered may be quantified in the units of position, time and payload byte distances. A reordered packet performance parameter is relevant for most applications, especially when assessing network support for real-time media streams, owing to their finite ability to restore order or when the performance implies a lack of that capability. Packets usually contain some unique identifier applied at the SRC, sometimes assumed to be a sequence number, so this number or other information (such as time stamps from the MP_0) is the reference for the original order at the source. The evaluation of arrival order also requires the ability to determine which specific packet is the "next expected" packet, and this is greatly simplified where sequence numbers are consecutive increasing integers.

An in-order packet outcome occurs when a single IP packet reference event at a permissible egress measurement point results in the following:

- The packet has a sequence number greater than or equal to the next expected packet value. The next expected value increases to reflect the arrival of this packet, setting a new value of expectation.

A reordered or out-of-order packet outcome occurs when a single IP packet reference event at a permissible egress measurement point results in the following:

- The packet has a sequence number lower than the next expected packet value and therefore the packet is reordered. The next expected value does not change due to arrival of this packet.

5.5.8.2 IP packet severe loss block outcome: An IP packet severe loss block outcome occurs for a block of packets observed during time interval T_s at ingress MP_0 when the ratio of lost packets at egress MP_i to total packets in the block exceeds s_1 .

The value of time interval T_s is provisionally set at 1 minute. The value of threshold s_1 is provisionally set at 0.2. Evaluation of successive blocks (time intervals) should be non-overlapping.

NOTE – These values are intended to identify IP path changes due to routing updates, which cause significant degradation to most user applications. The values may change following further study and experience. Lower values of s_1 would capture additional network events that may affect the operation of connectivity-sensitive applications. Also, significant degradation to video and audio applications may be well correlated with the IPSLB outcome when using T_s block lengths of approximately 1 second, and use of this value may be important in the future.

The minimum number of packets that should be used in evaluating the severe loss block outcome is M_{lb} , and these packets should be spread throughout a T_s interval. The value of M_{lb} is for further study.

5.5.8.3 duplicate IP packet outcome: A duplicate packet transfer outcome is a subset of successful packet outcomes, and occurs when a single IP packet reference event at a permissible

ingress MP_0 results in two or more corresponding reference event(s) on at least one permissible egress MP_i , and the binary information fields of all the output packets are identical to the original packet. The egress reference event at MP_i for a duplicate packet occurs subsequently to at least one other corresponding egress reference event for the original packet (usually also at MP_i).

Note that in point-to-point communication, there is only one permissible egress MP_i where the destination host is directly attached to the NSE. In point-to-multipoint communication, there may be many permissible egress MP_i for the various destinations.

5.5.8.4 replicated IP packet outcome: A replicated packet transfer outcome occurs when a single IP packet reference event at a permissible ingress MP_0 results in two or more corresponding reference event(s) on at least one permissible egress MP_i , and the binary information fields of all the output packets are identical to the original packet. The egress reference event at MP_i for a replicated packet is the first for the original packet and occurs prior to at least one other egress reference event for a duplicate packet (usually also at MP_i).

6 IP packet transfer performance parameters

This clause defines a set of IP packet information transfer performance parameters using the IP packet transfer outcomes defined in clause 5.5. All of the parameters may be estimated on the basis of observations made at MP that bound the basic section or NSE under test.

NOTE – Definitions of additional IP packet transfer performance parameters (e.g., severely errored IP packet block ratio) are for further study.

6.1 packet qualifications: This clause defines key terminology for qualifying the applicability of performance parameters to sets of packets.

6.1.1 populations of interest: Most of the performance parameters are defined over sets of packets called *populations of interest*. For the *end-to-end case*, the population of interest is usually the total set of packets being sent from SRC to DST. The measurement points in the end-to-end case are the MP at the SRC and DST.

For a basic section or NSE and relative to a particular SRC and DST pair, the population of interest at a particular permissible ingress MP is that set of packets being sent from SRC to DST that are routed into the basic section or NSE across that specific MP. This is called the *specific-ingress case*.

The total population of interest for a basic section or NSE relative to a particular SRC and DST pair is the total set of packets from SRC to DST that are delivered into the section or NSE across any of its permissible ingress MPs. This is called the *ingress-independent case*.

Each of these IP performance parameters are defined without reference to a particular packet type (ToS, DSCP, protocol, etc.) Performance will differ by packet type and any statement about measured performance should include information about which packet type or types were included in the population.

6.1.2 packet flow: A packet flow is the set of packets associated with a given connection or connectionless stream having the same source host address (SRC), destination host address (DST), class of service, and session identification (e.g., port numbers from a higher-layer protocol). Other documents may use the terms microflow or subflow when referring to packet streams with this degree of classification. A packet flow is the most common example of a population of interest.

IPv6 packets have an additional field for the source host to label sequences of packets which should receive some special treatment in IPv6 routers. This field is called the flow label and, in combination with the source address, uniquely defines a packet flow.

6.2 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 IPRE₁ at time t_1 and egress event IPRE₂ 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 8.




Figure 8 – IP packet transfer delay events
(illustrated for the end-to-end transfer of a single IP packet)

6.2.1 mean IP packet transfer delay: Mean IP packet transfer delay is the arithmetic average of IP packet transfer delays for a population of interest.

6.2.2 minimum IP packet transfer delay: Minimum IP packet transfer delay is the smallest value of IP packet transfer delay among all IP packet transfer delays of a population of interest. This includes propagation delay and queuing delays common to all packets. Therefore, this parameter may not represent the theoretical minimum delay of the path between MP.

6.2.3 median IP packet transfer delay: The median IP packet transfer delay is the 50th percentile of the frequency distribution of IP packet transfer delays from a population of interest. The median is the middle value once the transfer delays have been rank-ordered. To obtain this middle value when the population contains an even number of values, then the mean of the two central values is used.

6.2.4 end-to-end 2-point IP packet delay variation: 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 9): $v_k = x_k - d_{1,2}$.




Figure 9 – 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 the first 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 $d_{1,2}$.

6.2.4.1 Using minimum delay as the basis for delay variation: As illustrated in Figure 9, 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 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).

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 centring 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).

6.2.4.2 Quantile-based limits on IP packet delay variation: The preferred method (used in 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."

6.2.4.3 Interval-based limits on IP packet delay variation: An alternative method for summarizing the IP packet delay variation experienced by a population of interest is to pre-specify a delay variation interval, e.g., 50 ms, and then observe the percentage of individual packet delay variations that fall inside and outside of that interval. If the 50 ms interval were used, application with fixed buffer sizes of at or near 50 ms would then know approximately how many packets would cause buffer over- or under-flow.

NOTE – If this method is used for summarizing IP packet delay variation, the delay variant of individual packets should be calculated using the minimum delay as nominal in clause 6.2.4.1, instead of the definition of clause 6.2.4 using the first packet. Using the definition of clause 6.2.4, the pre-selected interval (e.g., the 50 ms) might occasionally be anchored on an unusually large or small value.

An objective for IP packet delay variation could be established by choosing a lower bound for the percentage of individual packet delay variations that fall within a pre-specified interval. For example, "≥99.9% of packet delay variations should be within the interval [0 ms, 50 ms]."

6.2.4.4 Secondary parameters for IP packet delay variation: One or more parameters that capture the effect of IP packet delay variations on different applications may be useful. It may be appropriate to differentiate the (typically small) packet-to-packet delay variations from the potentially larger discontinuities in delay that can result from a change in the IP routing. Appendix II gives several secondary definitions of delay variation and guidance on their use.

6.3 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.

6.4 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.

6.5 spurious IP packet rate: Spurious IP packet rate 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 service-second)¹.

6.6 IP packet reordered ratio (IPRR): An IP packet reordered ratio is the ratio of the total reordered packet outcomes to the total of successful IP packet transfer outcomes in a population of interest.

Figure 10 illustrates an out-of-order packet outcome for packet 2, and a hypothetical tolerance on arrival time with a playout buffer that can restore order.




Figure 10 – Illustration of reordered arrival

If separate reordering events can be distinguished, then an event count may also be reported (along with the event criteria).

It is also possible to assert the degree to which a packet is reordered. Any packet whose sequence number causes the next expected value to increment by more than the standard increment indicates a discontinuity in the arrival order. From this point on, any (reordered) packets with sequence number less than the next expected value can be quantified with a distance with respect to the discontinuity. The distance may be in units of position, time or the sum byte payloads of intervening packets. Referring to Figure 10 for an example, packet 2 can be said to be "late" by δt seconds, or 1 packet in terms of position.

[IETF RFC 4737] should be consulted for additional reordering parameters.

6.7 IP packet severe loss block ratio (IPSLBR): An IP packet severe loss block ratio is the ratio of the IP packet severe loss block outcomes to total blocks in a population of interest.

NOTE – This parameter can identify multiple IP path changes due to routing updates, also known as route flapping, which causes significant degradation to most user applications.

6.8 IP packet duplicate ratio (IPDR): IP packet duplicate ratio is the ratio of total duplicate IP packet outcomes to the total of successful IP packet transfer outcomes minus the duplicate IP packet outcomes in a population of interest.

6.9 replicated IP packet ratio (RIPR): The replicated IP packet ratio is the ratio of total replicated IP packet outcomes to the total of successful IP packet transfer outcomes minus the duplicate IP packet outcomes in a population of interest.

¹ Since the mechanisms that cause spurious IP packets are expected to have little to do with the number of IP packets transmitted across the sections under test, this performance parameter is not expressed as a ratio, only as a rate.

6.10 Flow-related parameters: It is useful to characterize performance in terms of flow or throughput-related parameters that evaluate the ability of IP networks or sections to carry quantities of IP packets. It should be noted that a parameter intended to characterize the throughput of an IP application would not necessarily be an accurate estimate of the amount of resources available to that application; this is because the higher layer protocols over IP (e.g., TCP) also influence the throughput experienced.

In the present version of this Recommendation, it is recommended that all flow- or throughput-related parameters should fulfil the following requirements:

- 1) A parameter characterizing the throughput offered to an IP service should relate the amount of IP packets successfully transported by an IP network or section to the amount of IP packets that were delivered into this network or section.
- 2) The throughput-related parameter should apply to an end-to-end IP network and to the IP transport across an EL, NS or NSE.

Some flow- or throughput-related parameters attempt to characterize the throughput capacity of an IP network, i.e., its ability to sustain a given IP packet transfer rate. It is recommended that any such parameters should fulfil the following additional requirements:

- 1) The traffic pattern offered to the IP network or section should be described, since the ability of the IP network or section to successfully deliver these packets depends on this traffic pattern.
- 2) The rate at which traffic is offered should not exceed the capacity (in bits per second) of the link that connects the sections under test with the destination sections that are not under test.
- 3) In any individual statement about throughput performance, the type of IP packet considered should be declared.

It is also recommended to follow the guidelines for throughput-related parameters and their measurement found in the RFC 3148 framework for bulk transfer capacity metrics. All parameters related to flow and throughput remain under study. Appendix III presents some candidate throughput-related parameters and an experimental method of measurement.

7 IP service availability

IP service availability is applicable to end-to-end IP service, basic sections and NSE.

An availability function (defined in clause 7.1) serves to classify the total scheduled service time for an IP service into available and unavailable periods. On the basis of this classification, both percent IP availability and percent IP unavailability are defined in clause 7.2. Finally, a two-state model of IP service availability serves as the basis for defining related availability parameters in clause 7.2.

NOTE – Unless otherwise noted by an IP service provider, the scheduled service time for IP service is assumed to be 24 hours a day, seven days a week.

7.1 IP service availability function

The basis for the IP service availability function is a threshold on the IPLR performance.

The IP service is available on an end-to-end basis if the IPLR for that end-to-end case is smaller than the threshold c_1 defined in Table 1.

Table 1 – IP service availability function

Outage criterion	Threshold
IPLR > c_1	$c_1 = 0.75$
<p>NOTE – The value of 0.75 for c_1 is considered provisional and is identified as requiring further study. Values of 0.9 and 0.99 have also been suggested for c_1. However, at the time of approval of this Recommendation the majority of causes for unavailability appear to stem from failures where the loss ratio is essentially 100%, and unavailable periods of more than 5 minutes accompany such failures. When IP networks support multiple qualities of service, it may be appropriate to consider different values of c_1 for different services. In this case, c_1 values of between 0.03 and 0.2 (based on resilience of different speech coders) have been suggested for services offering Y.1541 class 0 or class 1, and c_1 of 0.75 for Y.1541 class 5.</p> <p>The threshold c_1 is only to be used for determining when the IP network resources are (temporarily) incapable of supporting a useful IP packet transfer service. The value c_1 should not be considered a statement about IPLR performance nor should it be considered an IPLR objective suitable for any IP application. Performance objectives established for IPLR should exclude all periods of service unavailability, i.e., all time intervals when the $IPLR > c_1$.</p>	

Relative to a particular SRC and DST pair, *a basic section or an NSE is available for the ingress-independent case* if the IPLR for that pair is smaller than the threshold c_1 , as measured across all permissible ingress MPs.

Relative to a particular SRC and DST pair, *a basic section or an NSE is available for the specific-ingress case* if the IPLR for that pair is smaller than the threshold c_1 , as measured from a specific permissible ingress MP.

NOTE 1 – From an operations perspective, it will be possible to measure and/or monitor availability from a specific ingress MP and then use this information to create inferences about the ingress-independent availability.

NOTE 2 – The quantitative relationship between end-to-end IP service availability and the IP service availability of the basic section or NSE remains for further study.

If the outage criteria given by Table 1 is satisfied (i.e., IPLR exceeds its threshold), the IP service is in the unavailable state (experiences an outage). The IP service is in the available state (no outage) if the outage criteria is not satisfied. The minimum number of packets that should be used in evaluating the IP service availability function is M_{av} (the value of M_{av} is for further study. When tests of availability use end-user generated traffic, M_{av} of 1000 packets has been suggested). The minimum duration of an interval of time during which the IP service availability function is to be evaluated is T_{av} . T_{av} is provisionally defined to be five minutes. Study has revealed that this value is consistent with practical limits on IP layer operations. Monitoring of lower layer performance and network element faults may be able to identify impending unavailability in a shorter time, and direct corrective action. Appendix VI gives the rationale for the current IP service availability function definition and values for T_{av} and c_1 .

NOTE 3 – The outage criterion based on the IPLR is expected to satisfactorily characterize IP service availability. However, IP service availability might also take into account severely degraded performance for IPER and/or spurious IP packet rate. The inclusion of additional availability decision parameters and their associated thresholds remains for further study.

NOTE 4 – This unidirectional definition of availability is motivated by the fact that IP packets often traverse very different routes from SRC to DST than they traverse from DST to SRC. If, from an IP network user perspective, a bidirectional availability definition is needed, a bidirectional definition can be easily derived from this unidirectional definition.

It is intended that this definition of IP service availability be applicable to both end-user generated IP traffic (i.e., the normal flow of IP packets between the SRC and the DST) as well as to traffic generated by test sets and test methodologies. In either case, the source of the IP traffic should be

documented when reporting availability findings. Such documentation should include the specific types of packets used in each direction of flow.

Traffic generated specifically to test the availability state should be limited so that it does not cause congestion. This congestion could affect other traffic and/or could significantly increase the probability that the outage criteria will be exceeded.

More information on the determination of the availability state can be found in Appendix IV.

7.2 IP service availability parameters

7.2.1 percent IP service unavailability (PIU): The percentage of total scheduled IP service time (the percentage of T_{av} intervals) that is (are) categorized as unavailable using the IP service availability function.

7.2.2 percent IP service availability (PIA): The percentage of total scheduled IP service time (the percentage of T_{av} intervals) that is (are) categorized as available using the IP service availability function:

$$PIU = 100 - PIA$$

NOTE – Because the IPLR typically increases with increasing offered load from SRC to DST, the likelihood of exceeding the threshold c_1 increases with increasing offered load. Therefore, PIA values are likely to be smaller when the demand for capacity between SRC and DST is higher.

Appendix IV provides information on sampling to determine the PIA and PIU.

Appendix I

IP packet routing considerations

(This appendix does not form an integral part of this Recommendation)

This appendix describes IP packet routing considerations relevant to the characterization of IP service performance.

IP packet routing is determined by each network operator's policies and configurations for routing protocols, and choices of the protocols themselves. For example, operators configure a parameter for the "cost" of traversing each link in their network, and the routing algorithm computes the lowest-cost route to the destination based on its knowledge of the current state of network topology. Clearly, the path a packet takes from source to destination greatly influences the transfer delay it will experience (from both transport and queuing), as well as exposure to other impairments such as loss, errors, duplication and reordering.

Another way in which routing protocols influence packet transfer performance is in their automated response to changes in network topology, such as link or router failures, or maintenance action to take a network element out of service. When the network topology changes due to failure, a recovery process restores the affected connectivity over the remaining network topology, if possible. This process is called "re-routing" or "re-convergence", and typically contains the following steps (each requiring time to execute):

- 1) Failure/event detection.
- 2) Path computation.
- 3) Advertisement.
- 4) Forwarding table update.

Again, options for timers configured by the operator determine the duration of the re-routing process to a great extent. Operators also have the option to set waiting times between executions of the routing algorithm, which conserves processing resources but may lengthen the response to a failure in some cases.

Sub-IP networking technologies, such as SONET rings and MPLS-TE fast re-route, enable sub-second restoration from link or router failures.

Appendix II

Secondary terminology for IP packet delay variation

(This appendix does not form an integral part of this Recommendation)

II.1 Introduction

This Recommendation specifies a single primary/normative definition that assesses the variation in a set of delays with respect to a reference delay. This appendix provides two informative/secondary definitions in the clauses that follow (based on IETF's inter-packet delay variation, and a modification of 1-point cell delay variation). This appendix also gives guidance on when each parameter is most appropriate, and relates the results of observations with the different parameters.

There are two additional approaches to quantifying delay variation:

- 1) A parameter based on [b-IETF RFC 3393] that ascertains the inter-packet delay variation.
- 2) A parameter similar to the 1-point cell delay variation described in [b-ITU-T I.356], which assesses the packet arrival spacing at a single interface with respect to an ideal arrival interval.

Note that [b-ITU-T I.356] included two different variation definitions, both 2-point and 1-point.

The Y.1541 IP performance objectives for PDV are in terms of the normative 2-point packet delay variation parameter in this Recommendation.

II.2 Definition of inter-packet delay variation

[b-IETF RFC 3393] defines delay variation as follows:

- A definition of the IP packet delay variation (IPDV) can be given for packets inside a stream of packets.
- The IPDV of a pair of packets within a stream of packets is defined for a selected pair of packets in the stream going from measurement point MP1 to measurement point MP2.
- The IPDV is the difference between the one-way-delay of the selected packets.

A selection function unambiguously determines the pair of packets used in each calculation of the delay variation metric. Only packets that arrive successfully are used in IPDV calculations.

The first selection function defined is for adjacent packets in the stream. The 1-way delay of the current packet has the 1-way delay of the previous packet subtracted from it to determine the current packet's IPDV. If either of the packets in the pair (or both) are lost, then the IPDV is undefined.

Another important example is the selection function that produces an equivalent delay variation assessment to the 2-point PDV parameter defined in clause 6.2.4. The pair of packets always includes the current packet and the packet with the minimum 1-way delay in the stream. The 2-point PDV for all arriving packets is calculated by subtracting the minimum delay from their 1-way delay values (the reference delay is the minimum delay).

II.3 Definition of 1-point packet delay variation

The fundamental notion of a 1-point delay variation parameter is the comparison between the actual arrival pattern and the intended (usually periodic) arrival pattern. Some variations of this definition include a "skipping clock" adjustment (when cells or packets arrive late/behind their ideal arrival time), as in [b-ITU-T I.356]. The definition below does not implement the skipping clock feature, since there is no clear bias if the reference pattern is established arbitrarily.

The 1-point PDV (y_k) for packet k at an MP is the difference between the packet's reference arrival time (c_k) and actual arrival time (a_k) at the MP: $y_k = c_k - a_k$. The reference arrival time pattern (c_k) is defined as follows:

$$c_0 = a_0 = 0,$$
$$c_{k+1} = c_k + T$$

where T is ideal packet spacing

Positive values of 1-point PDV ("early" packet arrivals) correspond to packet clumping; negative values of 1-point PDV ("late" packet arrivals) correspond to gaps in the packet stream.

II.4 Guidance on applying the different parameters

Guidance that serves the practical side of measurement is as follows:

- When synchronized clocks are not possible (or temporarily unavailable) in measurement devices:
 - 1) 1-point packet delay variation (1-point PDV) is a possible substitute for 1-way delay range/histogram, applicable for measurements on packet streams with periodic sending times (once the reference arrival time is appropriately set).
 - 2) IPPM inter-packet delay variation is applicable to all traffic flow types.
 - 3) When clock error is stable, the Y.1540 2-point PDV can be calculated and used.
- When synchronized clocks are available in measurement devices:
 - 1) The Y.1540 PDV 1-way delay range/histogram calculation is useful for a range of assessment tasks, including assessment of de-jitter buffer size.
 - 2) IPPM inter-packet delay variation adds a parameter with sensitivity to sequential/short-term variation and some immunity to route changes.

The IPPM inter-packet metric IPDV is similar to the calculation of inter-arrival jitter measurement in real-time control protocol (RTCP) reports. Real time protocol (RTP) gives the calculation of inter-arrival jitter in clause 6.3.1 of [b-IETF RFC 3550], with a sample implementation in an appendix. Although there are some differences in method (RTCP inter-arrival jitter uses order of arrival, as opposed to sending sequence with IPDV), there should be a favourable comparison between a "smoothed jitter" computed using IPDV singletons and the RTCP reports of jitter in many circumstances (if many packets were reordered, the results would probably not agree). It would be valuable to have a parameter that can be related to measurements made by user's endpoints. The IPDV metric with adjacent packet pairs is also less susceptible to route changes during a measurement interval, where the effect would only be observed in measurement pairs spanning the route change.

A positive attribute of 1-point PDV is its simplicity. The capability of assessing periodic streams within a single network element is highly advantageous.

A point that must be made clear in all variation parameter specifications is the effect of packet length. Since insertion time is included in transfer delay (first-bit to last-bit), packets with varying size have an inherent delay variation. Network specifications and tests should use packets with a single size to simplify interpretation of the results (and the size must be recorded).

It is worthwhile noting that the IETF IPPM Working Group has work in progress to compare the delay variation metrics in more dimensions and to provide detailed guidance.

Appendix III

Rate and throughput capacity related parameters

(This appendix does not form an integral part of this Recommendation)

This appendix, which is for further study, presents metrics and techniques for assessing the rate and aspects of the throughput capacity of IP networks. The specific proposal for a throughput probe that appeared in previous versions of this Recommendation has been deleted, since some of the assumptions about maximum TCP window size settings and packet sizes are no longer realistic. The open study questions are still valuable, and have been retained.

III.1 Definition of IP packet rate parameters

Two types of rate parameters are currently envisaged. One parameter measures rate in terms of rate of successfully transmitted IP packets; another parameter is octet-based and measures the rate in terms of the octets that have been transmitted in those packets.

III.1.1 IP packet rate (IPPR): For a given population of interest, the IP packet rate at an egress MP is the total number of IP packet transfer reference events observed at that egress MP during a specified time interval divided by the time interval duration (equivalently, the number of IP packet transfer reference events per service-second).

III.1.2 octet-based IP packet rate (IPOR): For a given population of interest, the octet-based IP packet rate at an egress MP is the total number of octets transmitted in IP packets that result in an IP packet transfer reference event at that egress MP during a specified time interval divided by the time interval duration (equivalently, the number of octets in the IP packets resulting in IP packet reference events per service-second).

III.2 References for throughput parameters and measurements

Throughput parameter definitions are considered controversial because the measurements have many dependencies and results from different measurement techniques may not be comparable. [b-IETF RFC 3148] provides the IETF's guidance on the development of metrics of this class. The authors wisely point out that, in order to capture the flow-control aspects of a particular measurement tool, the areas normally left flexible in a protocol must be tightly specified to measure bulk transfer capacity (BTC).

At present, the IETF IPPM working group is developing fundamental definitions for network capacity. It is hoped that measurement methods having relevance to this topic of throughput capacity will follow.

III.3 Open issues

The following questions can be investigated with a directed test programme. Answers to these questions would affirm or contradict the usefulness of throughput probes in assessing network capacity:

- Is IP packet loss really greater for throughput probes than for isolated IP packets?
- Is IP packet loss for throughput probes really larger than the packet loss during a streaming application that sustains an equivalent source rate for long periods of time? Is the upper bound so high as to be useless in predicting long-term performance of streaming applications?

- Is the throughput corruption ratio really an upper bound on corrupted TCP windows? Is the upper bound so high as to be useless in calculating long-term TCP performance?
- Since throughput probes do not have slow start operation, is there any substantial risk to other applications from infrequent testing with throughput probes?

Appendix IV

Minimal test of IP service availability state and sampling estimation of IP service availability parameters

(This appendix does not form an integral part of this Recommendation)

This appendix, which is for further study, describes a minimum test for determining whether an IP service, a basic section or an NSE is in the available state or the unavailable state. In a future version, it will provide methods for sampling estimation of the IP service availability parameters.

IV.1 Minimal test of IP service availability state (for test methodologies and test sets)

Clause 7.1 requires that at least M_{av} packets be used to evaluate the availability state. Test methodologies and test sets should attempt at least M_{av} packets spread throughout a T_{av} interval of time. For end-user generated traffic, successive T_{av} intervals of time might be concatenated until the requirement of at least M_{av} ingress events is fulfilled. This is for further study.

The following describes the minimum amount of effort that is necessary to decide the availability state during a single T_{av} interval of time. Repeated applications of this test are necessary in order to determine the PIA and the PIU. This minimum test of IP service availability is applicable to test methodologies and test sets; some requirements for end-user generated traffic are presented in clause 7.1. Any other test of IP service availability that (statistically) performs at least as well as this test is an acceptable test of IP availability. This test of IP availability is applicable end-to-end or in the specific-ingress case for a basic section or an NSE.

- Step 1: Determine the SRC and the DST.
- Step 2: Position test sets or activate test scripts at the appropriate measurement points.
- Step 3: At a predetermined time, start sending M_{av} IP packets distributed over the time duration T_{av} .
- Step 4: If the number of lost packet outcomes is greater than $c_1 \times M_{av}$ then the IP service is unavailable over the T_{av} interval of time.
- Step 5: If the IP service (basic section or NSE) is not declared unavailable as per the results of step 4, then it is available over this T_{av} interval of time.

IV.2 Sampling estimation of IP service availability

Random samples of the availability state using the minimum test above may be sufficient for estimating PIA and PIU. In order to estimate the duration of contiguous time in an available or an unavailable state, sampling must be much more frequent. [b-ITU-T X.137] provides procedures for X.25/X.75 networks that might also be suitable for IP service.

Appendix V

Material relevant to IP performance measurement methods

(This appendix does not form an integral part of this Recommendation)

This appendix, which is for further study, will describe important issues to consider as IP performance measurement methods are developed. It will describe the effects of conditions external to the sections under test, including traffic considerations, on measured performance.

The following conditions should be specified and controlled during IP performance measurements:

- 1) Exact sections being measured:
 - SRC and DST for end-to-end measurements;
 - MP bounding an NSE being measured.

NOTE – It is not necessary to measure between all MP pairs or all SRC and DST pairs in order to characterize performance.
- 2) Measurement time:
 - how long samples were collected;
 - when the measurement occurred.
- 3) Exact traffic characteristics:
 - rate at which the SRC is offering traffic;
 - SRC traffic pattern;
 - competing traffic at the SRC and DST;
 - IP packet size.
- 4) Type of measurement:
 - in-service or out-of-service;
 - active or passive.
- 5) Summaries of the measured data:
 - means, worst-case, empirical quantiles;
 - summarizing period:
 - short period (e.g., one hour);
 - long period (e.g., one day, one week, one month).

Appendix VI

Background on IP service availability

(This appendix does not form an integral part of this Recommendation)

VI.1 Introduction

This appendix gives the rationale for the current IP service availability function definition in clause 7. The purpose is to provide additional background information and aid the appreciation for this complex and important topic.

VI.2 Background

There are many ways to define availability, and many perspectives that translate into evaluation using a range of sensitivities and time-scales. This Recommendation uses a simple, adequate definition (from a network operator's perspective) that specifies the minimum evaluation conditions. In order to understand why the IP service availability function is sufficient, an understanding of the causes of unavailability is needed.

Figure VI.1 shows a Venn diagram where the universe is all service time. The body of this Recommendation notes that IP service providers may identify maintenance intervals where service availability is not guaranteed. Thus, the service time universe is usually different from the universe of *all* time.




Figure VI.1 – Illustration of service time as a Venn diagram

We indicate that service time is divided in two main categories: available time (on the left) and unavailable time (on the right). Note that the relative sizes are not to scale, since available time is usually much larger than unavailable time.

VI.3 Definitions of the regions in Figure VI.1

Unavailable time is composed of the following regions:

- **Not accessible:** The service user is unable to communicate with the IP network because of failure in the access network transport or network elements. The access link itself or router interface failure are common causes. Packet loss ratio is typically 100%, and this failure will often take much longer than 5 minutes to correct. Maintenance forces should be almost immediately alerted to the failure by fault management systems.
- **Not continuous:** The service user is unable to communicate with the desired destination because of a failure in IP network global routing information. The user may be able to communicate with some destinations, but not the desired destination. Packet loss ratio is typically 100% and this failure will often take much longer than 5 minutes to correct.
- **Not accessible, not continuous:** The service user is unable to communicate while both of the above conditions exist simultaneously.
- **Poor performance:** The service user is unable to communicate reliably with the desired destination. The packet loss ratio is 75% or greater, and the user will deem the service unavailable for communicating with almost any form of IP network application. When congestion is the primary cause for this level of packet loss, end-to-end flow control should be activated to alleviate it (as provided in TCP).

Available time is composed of the following regions:

- **Y.1541 class n-compliant:** The service user is able to communicate with the desired destination and the packet transfer performance is compliant with the objectives of the agreed class. Evaluation of this state is usually conducted in 1-minute intervals. Note that any user application will have specific capacity needs; the ability to support a traffic contract (as defined in [b-ITU-T Y.1221]) must also be considered.
- **Not compliant:** The service user is able to communicate with the desired destination, but the packet transfer performance does not meet one or more of the objectives of the agreed class. Evaluation of this state is usually conducted in 1-minute intervals.
- **IP packet severe loss block (IPSLB):** The service user is able to communicate with the desired destination, but the packet transfer performance does not meet one or more of the objectives of the agreed class. Specifically, the loss ratio is sufficient to determine that an IPSLB has occurred (provisionally defined as more than 20% loss in a 1-minute interval).

VI.4 Summary

We observe that the criteria of the IP service availability function are only important in the poor performance region, and that the unavailable time contributed by this region is small compared to the other causes of unavailability. Therefore, the evaluation of state based on loss alone, and the criteria provisionally agreed for state evaluation (5 minutes, 75% loss), are deemed sufficient.

Bibliography

- [b-ITU-T I.353] ITU-T Recommendation I.353 (1996), *Reference events for defining ISDN and B-ISDN performance parameters.*
- [b-ITU-T I.356] ITU-T Recommendation I.356 (2000), *B-ISDN ATM layer cell transfer performance.*
- [b-ITU-T X.25] ITU-T Recommendation X.25 (1996), *Interface between Data Terminal Equipment (DTE) and Data Circuit-terminating Equipment (DCE) for terminals operating in the packet mode and connected to public data networks by dedicated circuit.*
- [b-ITU-T X.75] ITU-T Recommendation X.75 (1996), *Packet-switched signalling system between public networks providing data transmission services.*
- [b-ITU-T X.137] ITU-T Recommendation X.137 (1997), *Availability performance values for public data networks when providing international packet-switched services.*
- [b-ITU-T Y.1221] ITU-T Recommendation Y.1221 (2002), *Traffic control and congestion control in IP-based networks.*
- [b-IETF RFC 768] IETF RFC 768 (1980), *User Datagram Protocol.*
<<http://www.ietf.org/rfc/rfc768.txt>>
- [b-IETF RFC 792] IETF RFC 792 (1981), *Internet Control Message Protocol.*
<<http://www.ietf.org/rfc/rfc792.txt>>
- [b-IETF RFC 793] IETF RFC 793 (1981), *Transmission Control Protocol.*
<<http://www.ietf.org/rfc/rfc793.txt>>
- [b-IETF RFC 919] IETF RFC 919 (1984), *Broadcasting Internet Datagrams.*
<<http://www.ietf.org/rfc/rfc919.txt>>
- [b-IETF RFC 922] IETF RFC 922 (1984), *Broadcasting Internet datagrams in the presence of subnets.*
<<http://www.ietf.org/rfc/rfc922.txt>>
- [b-IETF RFC 950] IETF RFC 950 (1985), *Internet Standard Subnetting Procedure.*
<<http://www.ietf.org/rfc/rfc950.txt>>
- [b-IETF RFC 959] IETF RFC 959 (1985), *File Transfer Protocol.*
<<http://www.ietf.org/rfc/rfc959.txt>>
- [b-IETF RFC 1305] IETF RFC 1305 (1992), *Network Time Protocol (Version 3) Specification, Implementation and Analysis.*
<<http://www.ietf.org/rfc/rfc1305.txt>>
- [b-IETF RFC 1786] IETF RFC 1786 (1995), *Representation of IP Routing Policies in a Routing Registry (ripe-8I++).*
<<http://www.ietf.org/rfc/rfc1786.txt>>
- [b-IETF RFC 1812] IETF RFC 1812 (1995), *Requirements for IP Version 4 Routers.*
<<http://www.ietf.org/rfc/rfc1812.txt>>
- [b-IETF RFC 2018] IETF RFC 2018 (1996), *TCP Selective Acknowledgment Options.*
<<http://www.ietf.org/rfc/rfc2018.txt>>
- [b-IETF RFC 2330] IETF RFC 2330 (1998), *Framework for IP Performance Metrics.*
<<http://www.ietf.org/rfc/rfc2330.txt>>
- [b-IETF RFC 3148] IETF RFC 3148 (2001), *A Framework for Defining Empirical Bulk Transfer Capacity Metrics.*
<<http://www.ietf.org/rfc/rfc3148.txt>>
- [b-IETF RFC 3357] IETF RFC 3357 (2002), *One-way Loss Pattern Sample Metrics.*
<<http://www.ietf.org/rfc/rfc3357.txt>>

- [b-IETF RFC 3393] IETF RFC 3393 (2002), *IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)*.
<<http://www.ietf.org/rfc/rfc3393.txt>>
- [b-IETF RFC 3432] IETF RFC 3432 (2002), *Network performance measurement with periodic streams*.
<<http://www.ietf.org/rfc/rfc3432.txt>>
- [b-IETF RFC 3550] IETF RFC 3550 (2003), *RTP: A Transport Protocol for Real-Time Applications*.
<<http://www.ietf.org/rfc/rfc3550.txt>>

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects and next-generation networks
Series Z	Languages and general software aspects for telecommunication systems