



INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

I.380

(02/99)

SERIES I: INTEGRATED SERVICES DIGITAL
NETWORK

Overall network aspects and functions – General network
requirements and functions

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

ITU-T Recommendation I.380

(Previously CCITT Recommendation)

ITU-T I-SERIES RECOMMENDATIONS
INTEGRATED SERVICES DIGITAL NETWORK

GENERAL STRUCTURE	
Terminology	I.110–I.119
Description of ISDNs	I.120–I.129
General modelling methods	I.130–I.139
Telecommunication network and service attributes	I.140–I.149
General description of asynchronous transfer mode	I.150–I.199
SERVICE CAPABILITIES	
Scope	I.200–I.209
General aspects of services in ISDN	I.210–I.219
Common aspects of services in the ISDN	I.220–I.229
Bearer services supported by an ISDN	I.230–I.239
Teleservices supported by an ISDN	I.240–I.249
Supplementary services in ISDN	I.250–I.299
OVERALL NETWORK ASPECTS AND FUNCTIONS	
Network functional principles	I.310–I.319
Reference models	I.320–I.329
Numbering, addressing and routing	I.330–I.339
Connection types	I.340–I.349
Performance objectives	I.350–I.359
Protocol layer requirements	I.360–I.369
General network requirements and functions	I.370–I.399
ISDN USER-NETWORK INTERFACES	
Application of I-series Recommendations to ISDN user-network interfaces	I.420–I.429
Layer 1 Recommendations	I.430–I.439
Layer 2 Recommendations	I.440–I.449
Layer 3 Recommendations	I.450–I.459
Multiplexing, rate adaption and support of existing interfaces	I.460–I.469
Aspects of ISDN affecting terminal requirements	I.470–I.499
INTERNETWORK INTERFACES	I.500–I.599
MAINTENANCE PRINCIPLES	I.600–I.699
B-ISDN EQUIPMENT ASPECTS	
ATM equipment	I.730–I.739
Transport functions	I.740–I.749
Management of ATM equipment	I.750–I.799

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

ITU-T RECOMMENDATION I.380

INTERNET PROTOCOL DATA COMMUNICATION SERVICE – IP PACKET TRANSFER AND AVAILABILITY PERFORMANCE PARAMETERS

Summary

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.

Source

ITU-T Recommendation I.380 was prepared by ITU-T Study Group 13 (1997-2000) and was approved under the WTSC Resolution No. 1 procedure on the 26th of February 1999.

FOREWORD

ITU (International Telecommunication Union) is the United Nations Specialized Agency in the field of telecommunications. The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of the ITU. The 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 Conference (WTSC), which meets every four years, establishes the topics for study by the ITU-T Study Groups which, in their turn, produce Recommendations on these topics.

The approval of Recommendations by the Members of the ITU-T is covered by the procedure laid down in WTSC Resolution No. 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 term *recognized operating agency (ROA)* includes any individual, company, corporation or governmental organization that operates a public correspondence service. The terms *Administration*, *ROA* and *public correspondence* are defined in the *Constitution of the ITU (Geneva, 1992)*.

INTELLECTUAL PROPERTY RIGHTS

The ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. The 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, the ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementors are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database.

© ITU 1999

All rights reserved. No part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from the ITU.

CONTENTS

	Page
1 Scope	1
2 Normative references.....	3
3 Abbreviations	3
4 Layered model of performance for IP service.....	6
5 Generic IP service performance model	6
5.1 Network components.....	6
5.2 Circuit sections 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.....	11
5.5.1 Global routing information and permissible output links	12
5.5.2 Corresponding events.....	12
5.5.3 Notes about the definitions of successful, errored, lost and spurious packet outcomes	13
6 IP packet transfer performance parameters	14
6.6 Flow related parameters	16
7 IP service availability	17
7.1 IP service availability function.....	17
7.2 IP service availability parameters.....	18
Appendix I – IP packet routing considerations.....	18
Appendix II – Terminology related to IP packet delay variation.....	19
II.1 End-to-end 2-point IP packet delay variation.....	19
II.2 Using average delay as the basis for delay variation.....	20
II.3 Interval-based limits on IP packet delay variation	20
II.4 Quantile-based limits on IP packet delay variation.....	20
Appendix III – Flow and throughput capacity related parameters.....	20
III.1 Definition of IP throughput parameters.....	20
III.2 Measurements using throughput probes.....	21
III.2.1 Destination limited source	21
III.2.2 Throughput probe.....	21
III.2.3 Probe performance parameters.....	22
III.2.4 Creating lower bounds on capacity currently available to applications.....	22
III.2.5 Open issues	23

	Page
Appendix IV – Minimal test of IP service availability state and sampling estimation of IP service availability parameters	23
IV.1 Minimal test of IP the service availability state (for test methodologies and test sets)	23
IV.2 Sampling estimation of IP service availability	24
Appendix V – Material relevant to IP performance measurement methods.....	24
Appendix VI – Bibliography	25

Recommendation I.380

INTERNET PROTOCOL DATA COMMUNICATION SERVICE – IP PACKET TRANSFER AND AVAILABILITY PERFORMANCE PARAMETERS

(Geneva, 1999)

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.

NOTE 1 – This Recommendation defines parameters that can be used to characterize IP service provided using IPv4; applicability or extension of I.380 to other IP services (e.g. guaranteed service) and other protocols (e.g. IPv6, RSVP) is for further study.

NOTE 2 – Recommendations for the performance of point-to-multipoint IP service are for further study.

The I.380 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 Recommendation 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.

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 I.380 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 preliminary concepts related to IP packet delay variation. Appendix III describes some possible techniques 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 I.380 parameters. Finally, Appendix VI provides a bibliography.

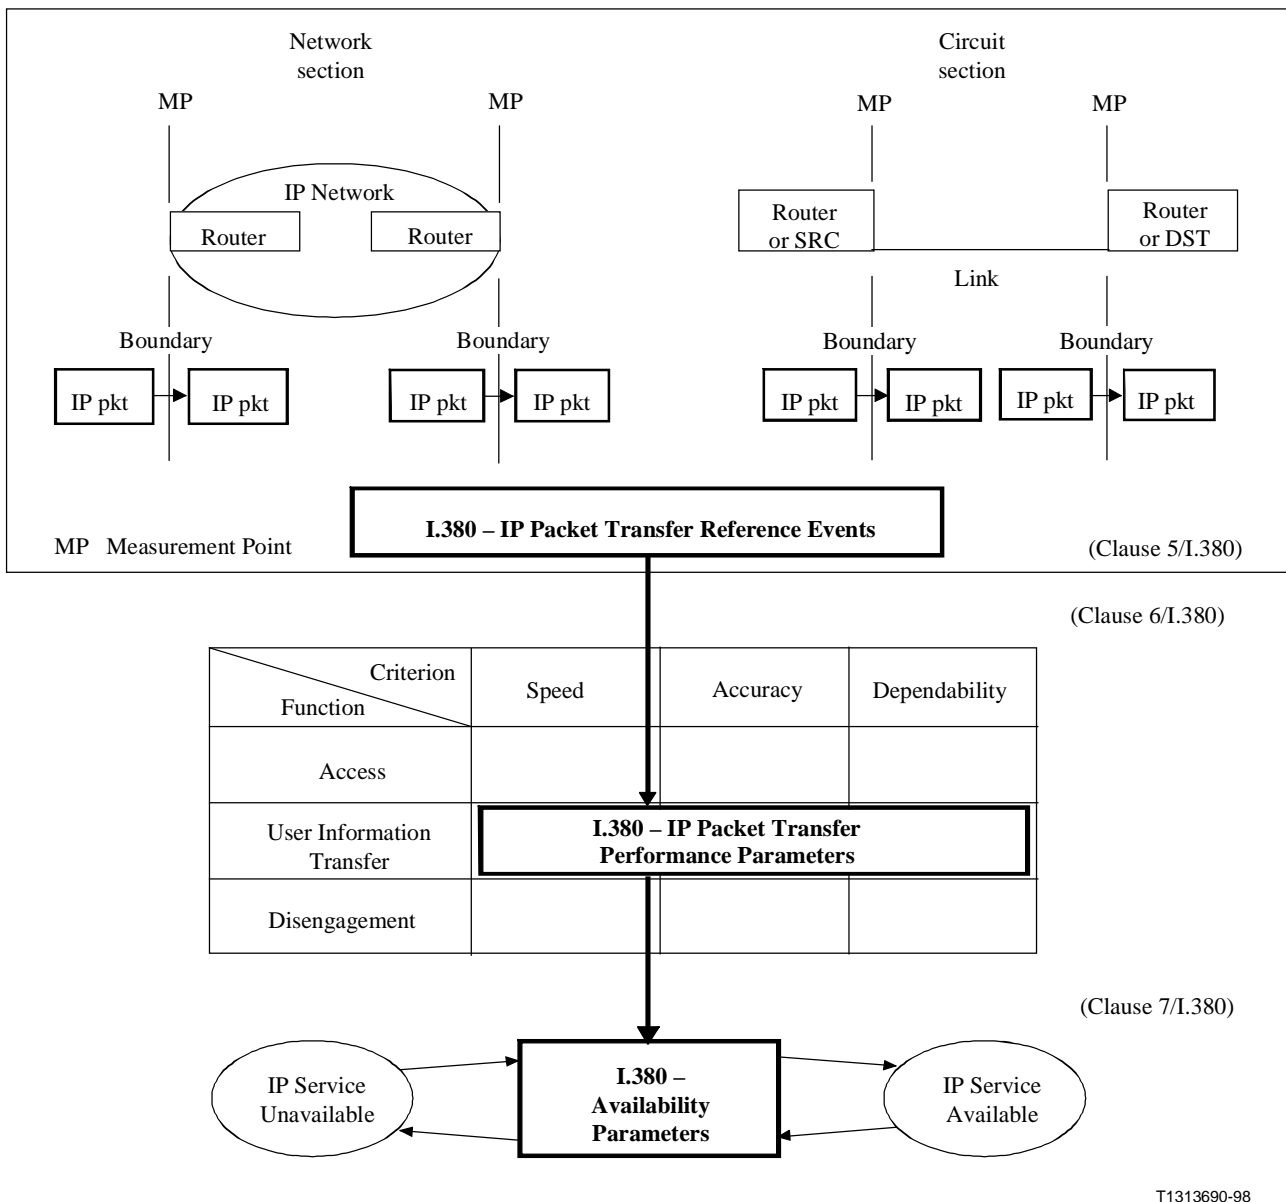


Figure 1/I.380 – Scope of ITU-T I.380

NOTE 4 – The I.380 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 I.380 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 circuit section.

NOTE 7 – The I.380 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 and delay parameters defined in this Recommendation. See Appendix I.

NOTE 10 – Specification of numerical performance objectives for some or all of the I.380 performance parameters is for further study. No objectives are specified in this version of the Recommendation.

2 Normative 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; all 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.


- ITU-T Recommendation I.350 (1993), *General aspects of quality of service and network performance in digital networks, including ISDNs*.
- RFC 791 (STD-5) – *Internet Protocol (IP)*, DARPA Internet program protocol specification, September 1981.

3 Abbreviations

This Recommendation uses the following abbreviations:

ATM	asynchronous transfer mode
CS	circuit section
DST	destination host
FTP	file transfer protocol
gw	gateway router
HTTP	hypertext transfer protocol
IETF	Internet Engineering Task Force
IP	Internet protocol
IPER	IP packet error ratio
IPLR	IP packet loss ratio

IPOT	octet based IP packet throughput
IPPT	IP packet throughput
IPRE	IP packet transfer reference event
IPTD	IP packet transfer delay
ISP	Internet service provider
ITU-T	International Telecommunication Union – Telecommunication Standardization Sector
LL	lower layers, protocols and technology supporting the IP 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 IP 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
PIA	percent IP service availability
PIU	percent IP service unavailability
pkt	IP datagram (IP packet)
QoS	quality of service
R	router
RFC	Request for Comment
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 IP availability; minimum length of time of IP unavailability
TCP	transmission control protocol
T_{max}	maximum IP packet delay beyond which the packet is declared to be lost
TOS	type of service
TTL	time to live
UDP	user datagram protocol



T1313700-98

Figure 2/I.380 – Layered model of performance for IP service – example

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.

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 circuit section and the network section. These are defined in 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.

Subclause 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. Subclause 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 I.380 performance model and reference events into Recommendation 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.




Figure 3/I.380 – IP network components

5.2 Circuit sections and network sections

5.2.1 circuit section (CS): The link connecting:

- 1) a source or destination host to its adjacent host (e.g. router) possibly in another jurisdiction; or
- 2) a router in one network section with a router in another network section.

Note that the responsibility for a circuit section, its capacity, and its performance is typically shared between the connected parties.

NOTE – "Circuit section" is roughly equivalent to the term "exchange" as defined in 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 circuit sections.

NOTE – "Network section" is roughly equivalent to the term "cloud" as defined in 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 circuit sections). 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 circuit sections.




Figure 4/I.380 – 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 Recommendation I.353, the standard Internet protocols can be observed at IP measurement points. Recommendation I.353 provides more information about MP 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 a CS, an NS, a SRC, or a DST. Basic sections are delimited by MP.

The performance of any CS 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 CS and NS 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 CSs that interconnect them. The term NSE can be used to refer to a single NS, two NSs, or any number of NS and their connecting CS. Pairs of distinct NSEs are connected by circuit sections. The term NSE can also be used to represent the entire end-to-end IP network. NSEs are delimited by MP.

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.

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 about higher layer protocols including Type of Service (TOS). 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".



T1313730-98

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

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

Figure 5/I.380 – Example IP packet transfer reference events

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 CS.

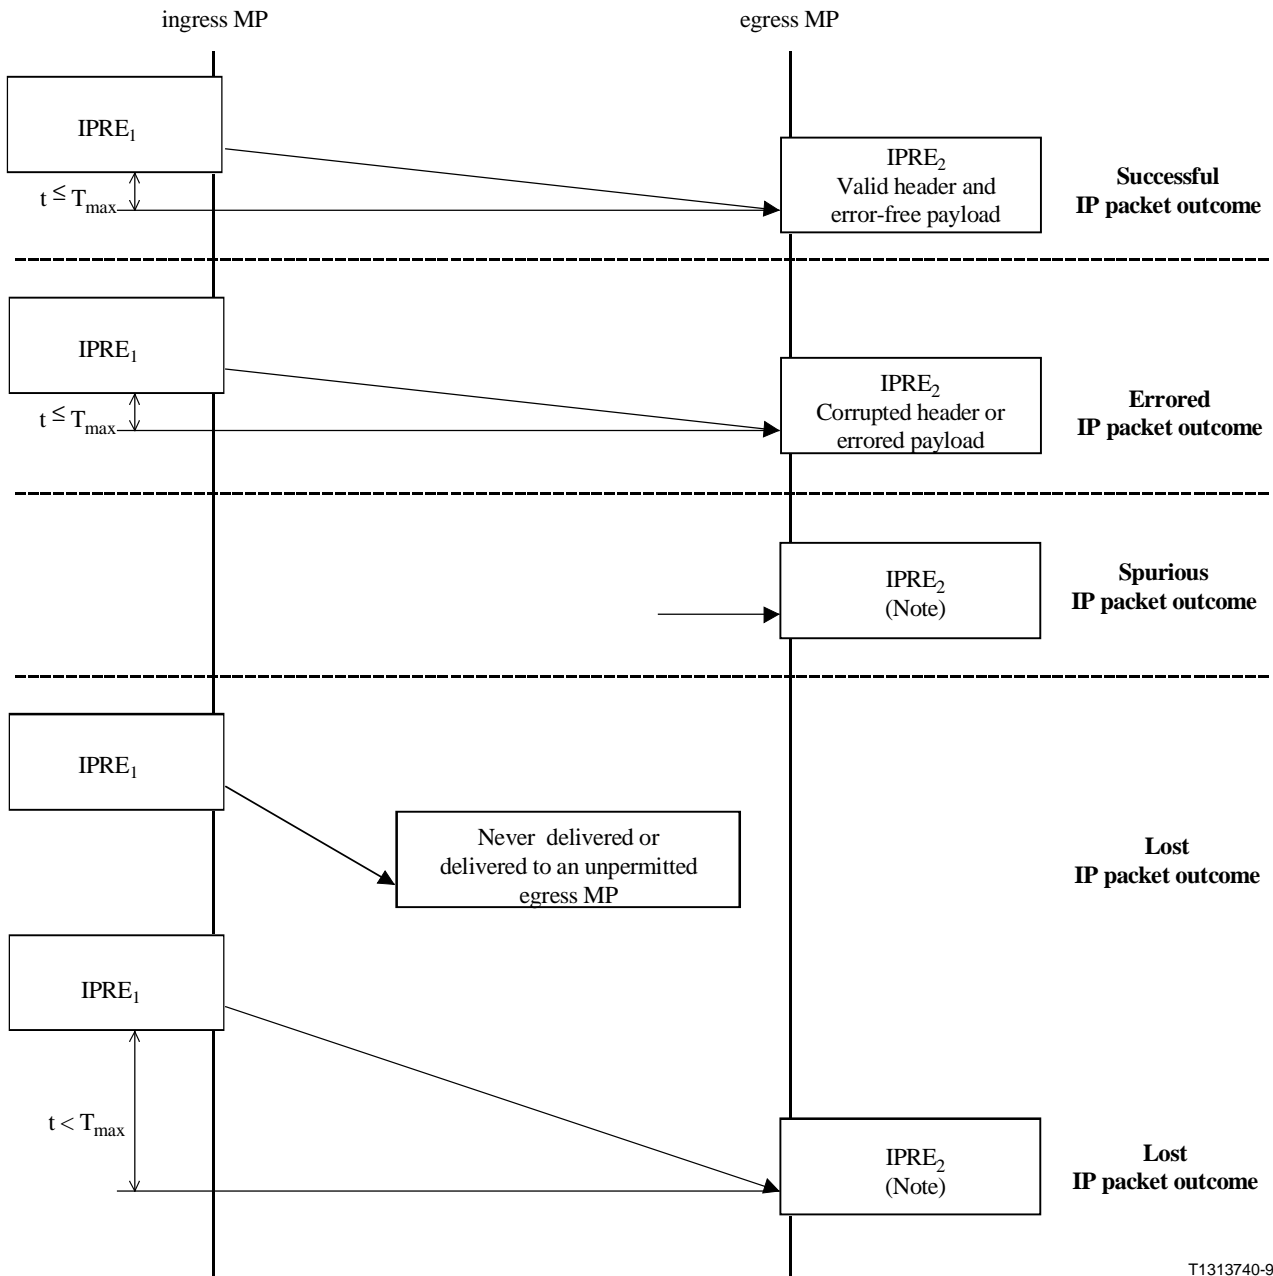
5.4.2 IP packet exit event from a host: An IP packet transfer exit event from a host occurs when an IP packet crosses an MP exiting a host (NS router or SRC) into the attached CS.

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 a 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 a 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 a CS 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.



T1313740-98

Figure 6/I.380 – IP packet transfer outcomes

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.

NOTE – Definition of other IP packet transfer outcomes (e.g. definition of a "Severely Errored IP Packet Block Outcome" based on a time duration or sequence of packets) is for further study.

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

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. A NS that was permissible may no longer be permissible following an update of the routing information shared between NSs. Alternatively, a 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:

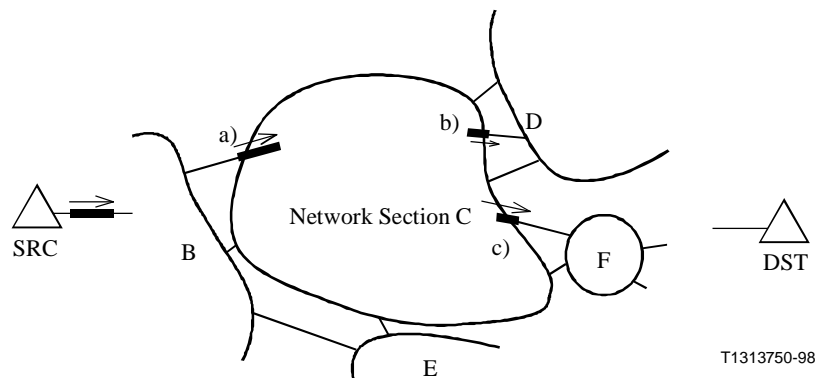
- an ingress MP is a *permissible ingress MP* if the crossing of this MP into this basic section or NSE is permitted by the global routing information;
- an egress MP is a *permissible egress MP* if the crossing of this MP leads into another basic section that is permitted by the global routing information.

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).

Figure 7/I.380 – 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 CS, a particular NS, a particular NSE, and they can be applied to the performance of end-to-end service.

These outcomes are defined without restriction to a particular packet type (TOS, 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 CS.

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 for further study. A value of 255 seconds has been suggested.

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 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 are not permissible egress MP.

A lost packet outcome occurs when a single IP packet reference event at a permissible ingress MP_0 results in a misdirected packet outcome 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 end-to-end when a single IP packet creates an egress event for which there was no corresponding ingress event.

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 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.

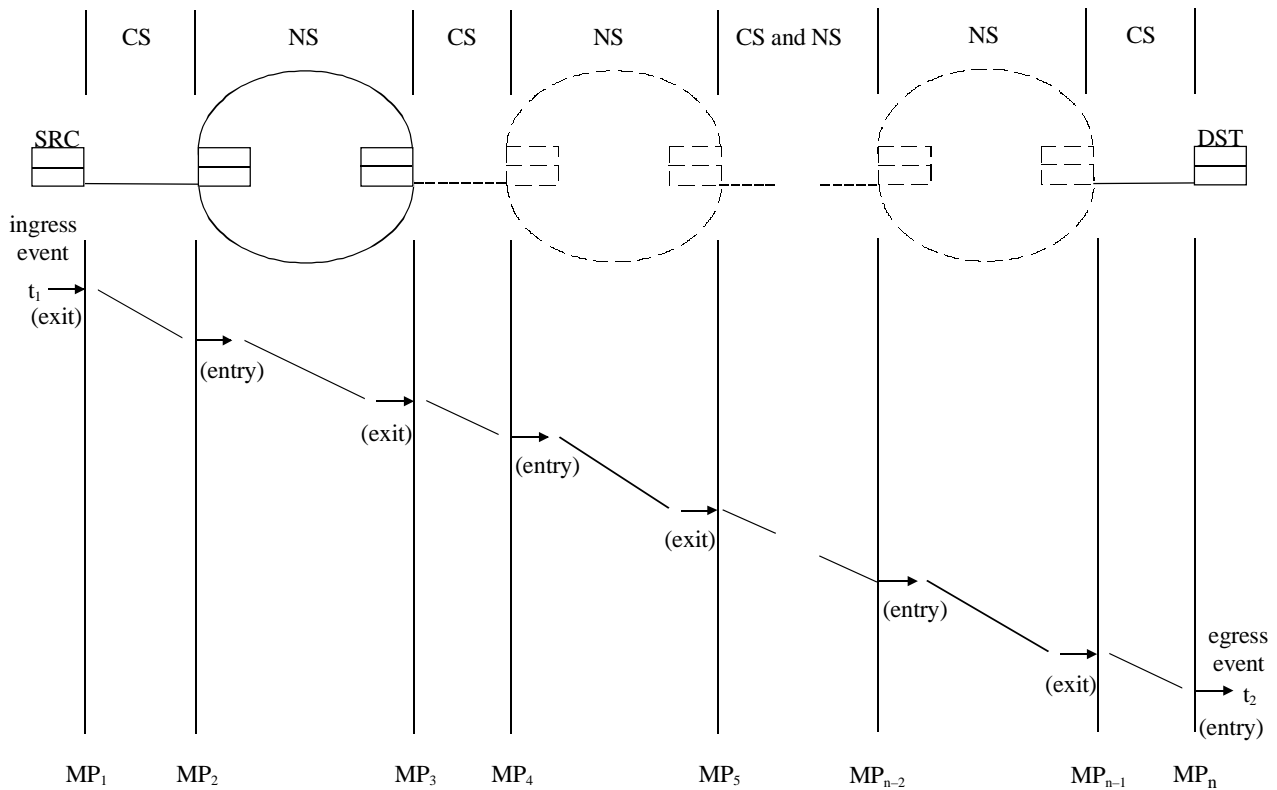
6.1 populations of interest: Most of the performance parameters are defined over sets of packets called *populations of interest*. For the *end-to-end case*, the population of interest is usually the total set of packets being sent from SRC to DST. The measurement points in the end-to-end case are the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

Each of these IP performance parameters are defined without reference to a particular packet type (TOS, 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.

6.2 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 $IPRE_1$ at time t_1 and egress event $IPRE_2$ at time t_2 , where $(t_2 > t_1)$ and $(t_2 - t_1) \leq T_{max}$. If the packet is fragmented within the NSE, t_2 is the time of the final corresponding egress event. The end-to-end IP packet transfer delay is the one-way delay between the MP at the SRC and DST as illustrated in Figure 8.



T1313760-98

Figure 8/I.380 – 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 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. 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. 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 describes some terminology that might be useful in quantifying aspects of IP packet delay variation.

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.

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 Flow related parameters

Currently in IPv4-based networks, the traffic offered on an end-to-end IP service is not checked for its conformance to an agreed traffic pattern. Furthermore, IPv4 networks can limit the rate at which packets are offered by a SRC only by discarding those packets. Finally, today's IP networks usually make no formal commitment to deliver any of the offered traffic.

However, it is useful to characterize the performance delivered by sections 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 that characterizes 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 fulfill 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 a CS, an NS or an 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 fulfill 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.

Appendix III proposes some throughput related parameters that are currently considered for inclusion in this Recommendation. All parameters related to flow and throughput remain under study.

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.

7 IP service availability

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

An availability function (defined in 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 7.2. Finally, a two-state model of IP service availability serves as the basis for defining related availability parameters in 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.

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 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.

Table 1/I.380 – IP service availability function

Outage criterion	Threshold
$IPLR > c_1$	$c_1 = 0.75$

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 . When IP networks support multiple qualities of service, it may be appropriate to consider different values of c_1 for different services.

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 $IPLR > c_1$.

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.) 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.)

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, which is for further study, will describe IP packet routing considerations relevant to the characterization of IP service performance.

APPENDIX II

Terminology related to IP packet delay variation

This appendix, which is for further study, describes terminology that may be helpful in defining useful IP packet delay variation parameters.

II.1 End-to-end 2-point IP packet delay variation

End-to-end 2-point IP packet delay variation 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 reference events at the egress MP with reference to the pattern of corresponding reference events at the ingress MP.

The 2-point packet delay variation (v_k) for an IP packet k between SRC and DST 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 (see Figure II.1): $v_k = x_k - d_{1,2}$.

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.

Positive values of 2-point PDV 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 $d_{1,2}$.

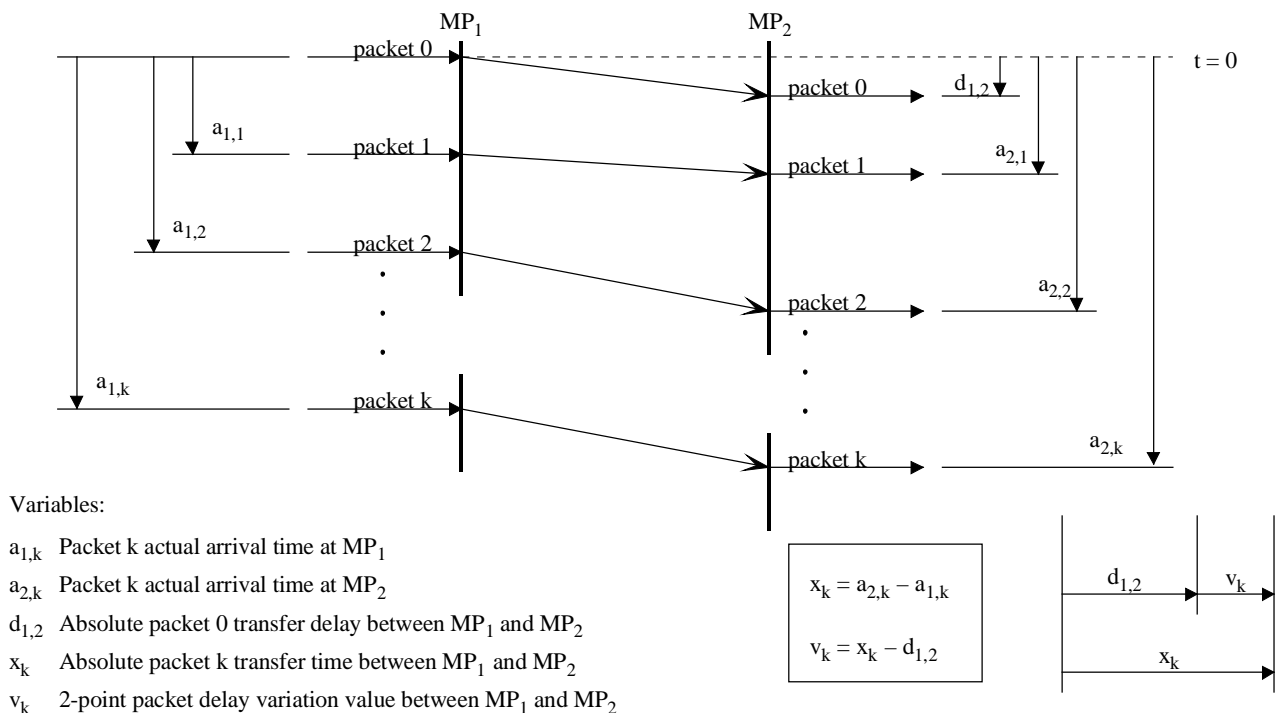


Figure II.1/I.380 – 2-point IP packet delay variation

II.2 Using average delay as the basis for delay variation

As illustrated in II.1, the delay variation of an individual packet is naturally defined as the difference between the actual delay experienced by that packet and a nominal (expected) delay. An alternative to using the first packet delay as the nominal delay is to use the average delay of the population of packets as the nominal delay. This has the effect of centering the distribution of delay variation values on zero.

II.3 Interval-based limits on IP packet delay variation

One method for summarizing the IP packet delay variation experienced by a population of packets is to pre-specify a delay variation interval, e.g. ± 30 milliseconds, and then observe the percentage of individual cell delay variations that fall inside and outside of that interval. If the ± 30 millisecond interval were used, application with fixed buffer sizes of at or near 60 milliseconds 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 definition in II.2, instead of the definition of II.1. Using the definition of II.1, the pre-selected interval (e.g. the ± 30 milliseconds) might occasionally be centered on an unusually large 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, "≥95% of packet delay variations should be within the interval [–30 msec, +30 msec]."

II.4 Quantile-based limits on IP packet delay variation

An alternative for summarizing the delay variation of a population of IP packets is to select upper and lower quantiles of the delay variation distribution and then measure the distance between those quantiles. For example, select the 99.5% ile and then 0.5% ile, make measurements, and observe the difference between the delay variation values at these two quantiles. This example would help application designers decide how to design for no more than 1% total buffer over- and under-flow.

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 99.5% ile and the 0.5% ile of the packet delay variation should be no more than 100 milliseconds."

APPENDIX III

Flow and throughput capacity related parameters

This appendix, which is for further study, presents metrics and techniques currently proposed for assessing the flow and throughput capacity of IP networks.

III.1 Definition of IP throughput parameters

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

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

III.1.2 Octet based IP packet throughput (IPOT): For a given population of interest, the octet based IP packet throughput at an egress MP is the total number of octets transmitted in IP packets that were successfully transmitted at that egress MP during a specified time interval divided by the time interval duration (equivalently, the number of octets in successfully transmitted IP packets per service-second).

III.2 Measurements using throughput probes

Throughput probes might be used to characterize the network's current capability to support additional traffic. By virtue of its brevity, a probe will not contribute in a major way to congestion. Any consequential congestion is further mitigated because the rate at which the throughput probe can be transmitted is bounded (III.2.1). The net effect is that widely scattered sampling using throughput probes will probably not place an excessive burden on the networks under test.

By virtue of their length, throughput probes will at least yield relative information about how much capacity is available for traffic between the SRC and DST. Subclause III.2.4 shows how the performance of the network in delivering throughput probes might be useful in creating lower bounds for the effective throughput performance of live IP applications.

III.2.1 Destination limited source

Let s be the link speed, in bits per second, of the link connecting the NSE under test to the destination host (DST). (If the link is a virtual connection such as a frame relay network, let s be its virtual carrying capacity in bits per second.) Let $\{p_1, p_2, p_3, \dots\}$ be the complete set of packets transmitted by the source host (SRC) to the DST, over its link to the NSE under test. Let t_1 be the instant in time the p_1 is transmitted by SRC. Let b_i be the number of bits in packet p_i including IP headers. Then the source is *destination limited* if for every packet p_j , the transmission of p_j does not

begin before

$$t_j = t_1 + \frac{1}{s} \sum_{i=1}^{j-1} b_i$$

NOTE 1 – If the link speed from the SRC to the NSE under test is equal or lower than s , the source is automatically destination limited.

NOTE 2 – If there is traffic from other sources using the same link from the NSE to DST, this traffic reduces the value of s used in this definition. This case requires further study.

NOTE 3 – It is never possible to sustain higher throughput than can be achieved using a fast destination limited source.

III.2.2 Throughput probe

A throughput probe is a sequence of N (<30), 576-byte IP packets transmitted from a destination limited SRC to a DST. In general, a significant amount time should elapse between the transmission of throughput probes for a given SRC and DST pair. At a minimum, if at least one of the N packets results in a lost packet outcome, another throughput probe should not be initiated until at least T_{\max} seconds after the time when the last of the lost packets was transmitted.

NOTE 1 – N is provisionally bounded by 30 because TCP implementations commonly advertise maximum window sizes that could allow up to 29 packets to be transmitted without acknowledgment (16 000 TCP payload bytes.)

NOTE 2 – The 576-byte packet is chosen because it is the maximum packet size all IP hosts are required to accept.

NOTE 3 – Enforcing the minimum separation between throughput probes helps ensure that one probe does not cause congestion for its successor and helps ensure that pairs of probe results are not correlated.

A *maximized throughput probe* is a throughput probe for which:

$$t_j = t_1 + \frac{1}{s} \sum_{i=1}^{j-1} b_i \quad (\text{allowing for reasonable clock differences}).$$

NOTE 4 – The most stressful tests will be those done with maximized throughput probes, but testing in certain contexts may allow for (or even prefer) testing with probes that are not maximized.

III.2.3 Probe performance parameters

NOTE 1 – If values are ever standardized for throughput probe performance, every value will be associated with its applicable probe size(s). It may be appropriate to use larger values of N for higher speed destination links. These issues are for further study.

NOTE 2 – As with other measures of throughput, when values for probe corruption ratio and probe packet ratio are specified, the competing traffic on the source link and destination link must be limited, controlled and reported. Because loading on networks will vary with time of day, time of day must also be controlled and reported in connection with throughput probe performance specifications.

III.2.3.1 probe corruption ratio: For an ensemble of throughput probes of given probe size, N, the probe corruption ratio is the fraction of those probes that have one or more lost packet outcomes at DST.

III.2.3.2 probe packet ratio: For an ensemble of throughput probes of given probe size, N, the probe packet ratio is the fraction of the packets within those probes that result in a successful or an errored packet outcome at DST.

III.2.4 Creating lower bounds on capacity currently available to applications

Today's dominant applications of IP networks are TCP implementations. These applications respond to congestion by slowing the rate at which they are transmitting (by reducing their window size) when loss is detected. When a new source of traffic is added to a router's burden, that new traffic increases the probability of queue overflow and increases the loss probability for each competing TCP application. That causes TCP applications to back off which in turn creates more room for the new traffic. Therefore, all other things being constant, new traffic will experience higher loss probabilities at the beginning of its transaction than it will experience later. An application running at its top speed will get better throughput (loss) performance after the competing TCP sources have backed off.

Similarly, an isolated throughput probe of size N is expected to experience a higher loss ratio than an application that attempts to sustain high throughput for more than N packets. For this reason, it is felt that throughput probe performance is a basis for constructing lower bounds on application throughput.

If a maximized throughput probe encounters no bottleneck and none of its packets are lost, the indication is that the network can, at least for the near-term, fully support destination limited throughput from SRC to DST. Also, if the throughput probe experienced no loss, it is likely that the throughput probe has not created much loss for its competing applications either. Those competing applications may only experience a temporary increase in IP packet delay during the test.

If a maximized throughput probe encounters a bottleneck and some of its packets are lost, the indication is that the network cannot immediately support the attempted level of throughput from SRC to DST. The near-term sustainable throughput might be lower bounded by the number of probe packets that were delivered. Over a longer time interval, if the destination limited SRC were to continue transmitting, competing TCP traffic would back off and the successful target traffic throughput would increase.

If a throughput probe experiences loss, it is likely that some of the competing connections will also have experienced loss during the test. Any TCP applications that experienced loss will reduce their window size. Since the throughput probe is short, the next TCP window will not compete with the probe, so the window size will immediately start to grow back to its original "equilibrium." This is a more acceptable outcome than would occur with a sustained test of throughput capacity.

III.2.5 Open issues

There is currently no empirical evidence to support many of the basic assertions about throughput probes presented above. The following questions can be investigated with a directed test program. 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, 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 the service availability state (for test methodologies and test sets)

Subclause 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 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. Recommendation 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, 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) the 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

Bibliography

- RFC 768 (STD-6) – *User Datagram Protocol.*
- RFC 792 (STD-5) – *Internet Control Message Protocol.*
- RFC 793 (STD-7) – *Transmission Control Protocol.*
- RFC 919 (STD-5) – *IP Broadcast datagrams.*
- RFC 922 (STD-5) – *Broadcasting Internet datagrams in the presence of subnets.*
- RFC 950 – *Internet Standard Subnetting Procedure (updates RFC 792).*
- RFC 959 (STD-9) – *File Transfer Protocol (FTP).*
- RFC 1305 – *Network Time Protocol (Version 3) Specification, Implementation and Analysis.*
- RFC 1786 – *Representation of IP Routing Policies in a Routing Registry.*
- RFC 1812 – *Requirements for IP Version 4 Routers.*
- RFC 2018 – *TCP Selective Acknowledgment Options.*
- RFC 2330 – *Framework for IP Performance Metrics.*

ITU-T RECOMMENDATIONS SERIES

- Series A Organization of the work of the ITU-T
- Series B Means of expression: definitions, symbols, classification
- Series C General telecommunication statistics
- 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 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 TMN and network maintenance: international transmission systems, telephone circuits, telegraphy, facsimile and leased circuits
- 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 and open system communications
- Series Y Global information infrastructure and Internet protocol aspects
- Series Z Languages and general software aspects for telecommunication systems