

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Y.1540

(07/2016)

SERIES Y: GLOBAL INFORMATION
INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS
AND NEXT-GENERATION NETWORKS, INTERNET OF
THINGS AND SMART CITIES

Internet protocol aspects – Quality of service and network
performance

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

Recommendation ITU-T 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
IPTV over NGN	Y.1900–Y.1999

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
Smart ubiquitous networks	Y.2600–Y.2699
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899
Carrier grade open environment	Y.2900–Y.2999
Future networks	Y.3000–Y.3099

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

Recommendation ITU-T Y.1540

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

Summary

Recommendation ITU-T 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 services. 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.

History

Edition	Recommendation	Approval	Study Group	Unique ID*
1.0	ITU-T I.380	1999-02-26	13	11.1002/1000/4573
1.0	ITU-T Y.1540	1999-02-26	13	11.1002/1000/5302
2.0	ITU-T Y.1540	2002-12-14	13	11.1002/1000/6189
2.1	ITU-T Y.1540 (2002) Amd. 1	2003-08-01	13	11.1002/1000/6975
3.0	ITU-T Y.1540	2007-11-13	12	11.1002/1000/9270
3.1	ITU-T Y.1540 (2007) Amd.1	2009-03-19	12	11.1002/1000/9727
4.0	ITU-T Y.1540	2011-03-01	12	11.1002/1000/11079
4.1	ITU-T Y.1540 (2011) Amd.1	2016-01-21	12	11.1002/1000/12761
5.0	ITU-T Y.1540	2016-07-29	12	11.1002/1000/12975

* To access the Recommendation, type the URL <http://handle.itu.int/> in the address field of your web browser, followed by the Recommendation's unique ID. For example, <http://handle.itu.int/11.1002/1000/11830-en>.

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 2016

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

Table of Contents

		Page
1	Scope.....	1
2	References.....	3
3	Abbreviations and acronyms	4
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.....	9
5.4	IP packet transfer reference events (IPREs).....	9
5.5	IP packet transfer outcomes.....	11
6	IP packet transfer performance parameters	16
6.1	Packet qualifications.....	16
6.2	IP packet transfer delay	17
6.3	IP packet error ratio (IPER).....	20
6.4	IP packet loss ratio (IPLR)	20
6.5	Spurious IP packet rate	21
6.6	IP packet reordered ratio (IPRR).....	21
6.7	IP packet severe loss block ratio (IPSLBR)	21
6.8	IP packet duplicate ratio (IPDR)	22
6.9	Replicated IP packet ratio (RIPR)	22
6.10	Stream repair parameters	22
6.11	Capacity parameters	22
6.12	Flow-related parameters	24
7	IP service availability	25
7.1	IP service availability function.....	25
7.2	IP service availability parameters.....	27
	Appendix I – IP packet routing considerations.....	28
	Appendix II – Secondary terminology for IP packet delay variation	29
	II.1 Introduction	29
	II.2 Definition of inter-packet delay variation	29
	II.3 Definition of 1-point packet delay variation	29
	II.4 Guidance on applying the different parameters.....	30
	Appendix III – Rate and throughput capacity related parameters	31
	III.1 Definition of IP packet rate parameters.....	31
	III.2 References for throughput parameters and measurements.....	31
	III.3 Open issues	31
	Appendix IV – Tests of IP service availability state and sampling estimation of IP service availability parameters.....	33

	Page
IV.1 Minimal test of IP service availability state (for test methodologies and test sets)	33
IV.2 Test of IP service availability state (using sequential probability ratio test)..	33
IV.3 Alternate test of statistical significance to determine IP service availability	35
IV.4 Sampling estimation of IP service availability	36
Appendix V – Material relevant to IP performance measurement methods	37
Appendix VI – Background on IP service availability	38
VI.1 Introduction	38
VI.2 Background.....	38
VI.3 Definitions of the regions in Figure VI.1	39
VI.4 Summary.....	39
Appendix VII – Packet performance parameters for estimation and optimization of stream repair techniques	40
VII.1 Introduction	40
VII.2 Short description of application-layer stream repair techniques	41
VII.3 Simple model of application-layer stream repair techniques	41
VII.4 Example of performance parameters to characterize stream repair variables.....	42
VII.5 Discussion of parameter measurement and usage	42
VII.6 Additional considerations	43
Appendix VIII – IP-layer capacity framework	44
VIII.1 Introduction	44
VIII.2 Terminology and relation to IETF RFC 5136	44
VIII.3 Items for further study	45
Appendix IX – Explanation of TCP-based measurement inadequacy to meet normative requirements	46
IX.1 Introduction	46
IX.2 Comparison with normative requirements	46
Bibliography.....	48

iv

Recommendation ITU-T 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 services. The defined parameters apply to the 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 Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6); applicability or extension of this Recommendation to other protocols (e.g., resource reservation protocol (RSVP)) is for further study.

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

The ITU-T 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 (IPREs) 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.




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 the IP data communication service. Future ITU-T Recommendations may be developed to provide standard methods of measuring the ITU-T 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 reference 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 ITU-T Y.1540 parameters. Appendix VI gives some background on IP service availability. Appendix VII offers background information on the stream repair parameters, and Appendix VIII adds information on capacity parameters (including a mapping to prior IETF metrics and items for further study).

NOTE 4 – The ITU-T 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 ITU-T 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 (EL).

NOTE 7 – The ITU-T 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 ITU-T 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] Recommendation ITU-T I.350 (1993), *General aspects of quality of service and network performance in digital networks, including ISDNs*.
- [ITU-T P.10] Recommendation ITU-T P.10/G.100 (2006), *Vocabulary for performance and quality of service*.
- [ITU-T Y.1541] Recommendation ITU-T 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>>
- [IETF RFC 5136] IETF RFC 5136 (2008), *Defining Network Capacity*.
<<http://www.ietf.org/rfc/rfc5136.txt>>
- [IETF RFC 5481] IETF RFC 5481 (2009), *Packet Delay Variation Applicability Statement*.
<<http://www.ietf.org/rfc/rfc5481.txt>>

3 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

ARQ	Automatic Repeat-request
ATM	Asynchronous Transfer Mode
BTC	Bulk Transfer Capacity
DSCP	Differentiated Services Code Point
DST	Destination host
EL	Exchange Link
ER	Edge Router
FEC	Forward Error Correction
FTP	File Transfer Protocol
HTTP	Hypertext Transfer Protocol
IP	Internet Protocol
IPDR	Internet Protocol packet Duplicate Ratio
IPDV	Internet Protocol packet Delay Variation
IPER	Internet Protocol packet Error Ratio
IPIBR	Internet Protocol packet Impaired Block Ratio
IPIIR	Internet Protocol packet Impaired Interval Ratio
IPLR	Internet Protocol packet Loss Ratio
IPOR	Octet-based IP packet Rate
IPPM	IP Performance Metrics
IPPR	Internet Protocol Packet Rate
IPRE	Internet Protocol packet transfer Reference Event
IPRR	Internet Protocol packet Reordered Ratio
IPSLB	Internet Protocol packet Severe Loss Block outcome
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
QoS	Quality of Service
R	Router
RIPR	Replicated Internet protocol Packet Ratio
RSVP	Resource reservation Protocol
RTCP	Real-Time Control Protocol
RTO	Retransmission Time Out
RTP	Real-time Transport Protocol
RTT	Round-Trip-Time
SDH	Synchronous Digital Hierarchy
SPRT	Sequential Probability Ratio Test
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 (LL) 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, asynchronous transfer mode (ATM), frame relay, synchronous digital hierarchy (SDH), plesiochronous digital hierarchy (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, transmission control protocol (TCP), user datagram protocol (UDP), file transfer protocol (FTP), real-time transport protocol (RTP) and hypertext transfer protocol (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 (NS). 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 ITU-T 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 (DST)) 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 (SRC) 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 – 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 Internet service provider (ISP) or a network service provider (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 (MPs). [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)

A network section ensemble (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.



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 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*.




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 the entire 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.