INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Y.1561

(05/2004)

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

Internet protocol aspects – Quality of service and network performance

Performance and availability parameters for MPLS networks

ITU-T Recommendation Y.1561

ITU-T Y-SERIES RECOMMENDATIONS

GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT GENERATION NETWORKS

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

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

ITU-T Recommendation Y.1561

Performance and availability parameters for MPLS networks

Summary

This Recommendation defines parameters that may be used in specifying and assessing the performance of speed, accuracy, dependability, and availability of packet transfer over a Label Switched Path on a Multi-Protocol Label Switching (MPLS) network. The defined parameters apply to end-to-end, point-to-point and multipoint-to-point LSP and to any MPLS domain that provides, or contributes to the provision of, packet transfer services.

Two categories of MPLS networks are considered:

1) TE-LSP: Traffic Engineering Label Switched Path, or configured LSP. These are point-point paths.

2) LDP-based LSP: This includes point-to-point and multipoint to point LSPs.

Source

ITU-T Recommendation Y.1561 was approved on 7 May 2004 by ITU-T Study Group 13 (2001-2004) under the ITU-T Recommendation A.8 procedure.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications. 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, implementors are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database.

CONTENTS

ITU-T Recommendation Y.1561

Performance and availability parameters for MPLS networks

1 Scope

This Recommendation defines parameters that may be used in specifying and assessing the performance of speed, accuracy, dependability, and availability of packet (labelled or not – in the latter case, Penultimate Hop Popping (PHP) may cause loss of Label Switched Paths (LSP) identity at the network edges) transfer over an LSP on a Multi-Protocol Label Switching (MPLS) network. The defined parameters apply to end-to-end, point-to-point and multipoint-to-point LSP and to any MPLS domain that provides, or contributes to the provision of, packet transfer services in accordance with the normative references specified in clause 2.

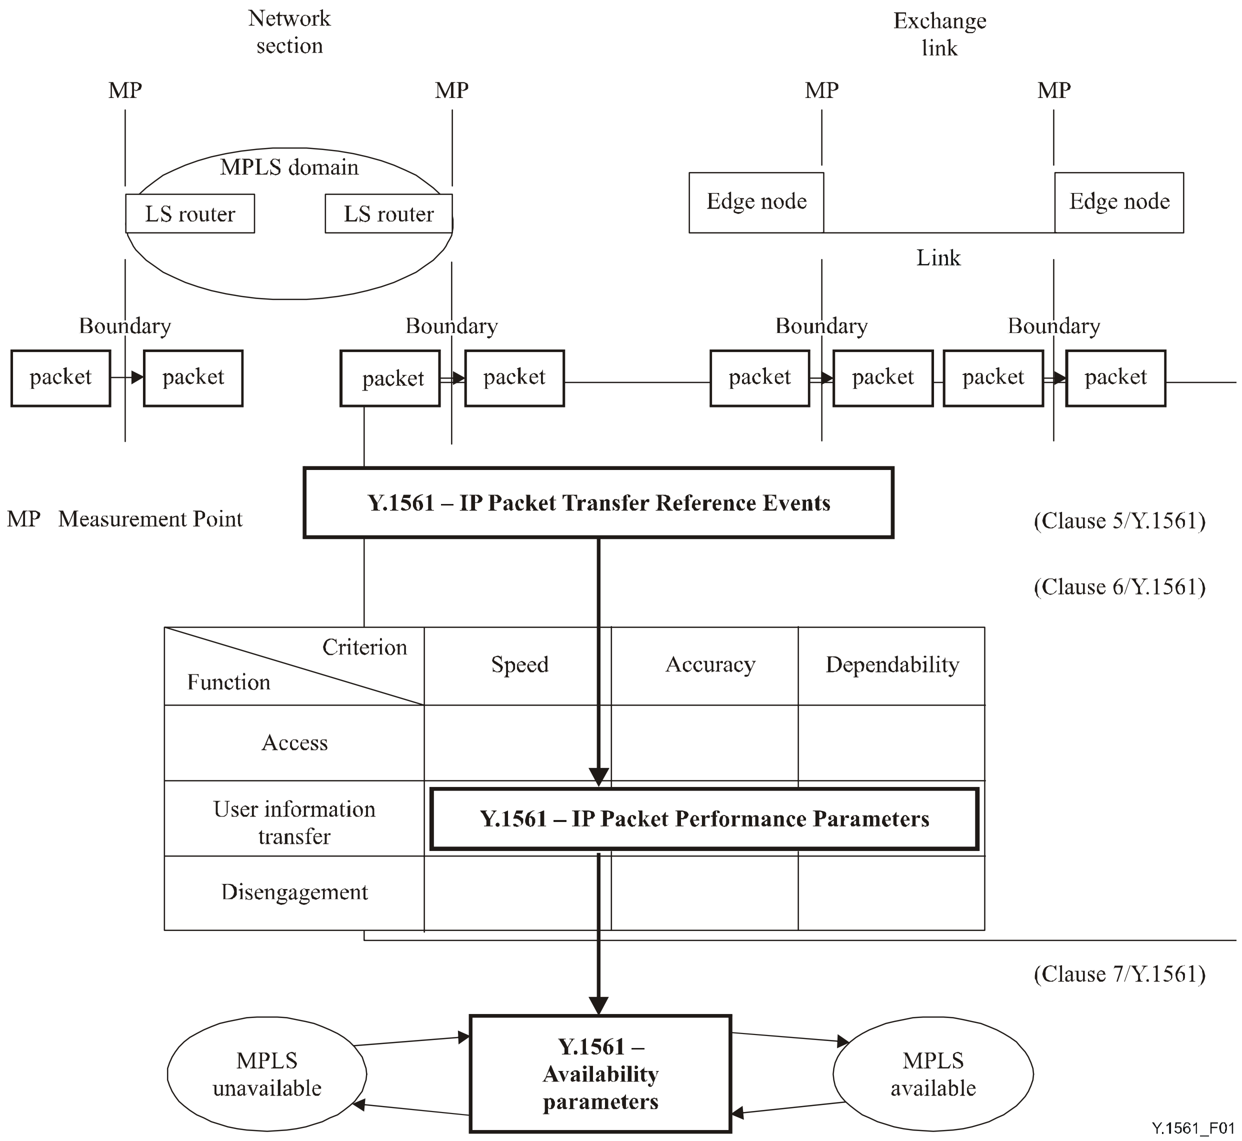

Criterion / Function	Speed	Accuracy	Dependability
Access			
User information transfer	Y.1561 – IP Packet Performance Parameters		
Disengagement			

Figure 1/Y.1561 – Scope of this Recommendation

The scope of this Recommendation is summarized in Figure 1. The MPLS network performance parameters are defined on the basis of packet transfer reference events that may be observed at measurement points (MPs) associated with specified functional and jurisdictional boundaries. The measurement points may be at the ends of LSPs. For comparability and completeness, MPLS network performance is considered in the context of the 3 × 3 performance matrix defined in ITU-T Rec. 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 MPLS network availability.

The performance of MPLS networks providing access and disengagement functions (e.g., Resource ReserVation Protocol – Traffic Engineering, RSVP-TE) and supporting capabilities (e.g., Label Distribution Protocol, LDP, as per RFC 3036) may be addressed in separate Recommendations.

Two categories of MPLS networks will be considered here:

1) TE-LSP: Traffic Engineering Label Switched Path, or configured LSP. These are point-point paths. The paths are connection-oriented, explicitly routed, and fixed.

2) LDP-based LSP: This includes point-to-point and multipoint to point LSPs. The paths behave more like IP, using the Interior Gateway Protocol (IGP), such as Open Shortest Path First (OSPF) or Intermediate System to Intermediate System (IS-IS) for routing. The LSPs are connectionless.

The point-to-point and multipoint-to-point topologies are accommodated through the concept of *populations of interest*, defined in 6.1. The case of PHP causing loss of LSP identity is specifically addressed in 5.4, where there are three optional criteria for packet transfer reference events. Correspondence between ingress and egress reference events has been dealt with here to the same degree as in different Recommendations (e.g., ITU-T Rec. Y.1540), in 5.5.2.

In this Recommendation, the general term *packet* refers to an IP packet with header and information field, or to other protocols with combinations of header and information fields, so long as there is a standard that describes the encapsulation as an MPLS packet.

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 Recommendation I.350 (1993), *General aspects of quality of service and network performance in digital networks, including ISDNs*.

– ITU-T Recommendation I.353 (1996), *Reference events for defining ISDN and B-ISDN performance parameters*.

– ITU-T Recommendation Y.1540 (2002), *Internet protocol data communication service – IP packet transfer and availability performance parameters*.

– ITU-T Recommendation Y.1711 (2004), *Operation and maintenance mechanism for MPLS networks*.

– IETF RFC 3031 (2001), *Multiprotocol Label Switching Architecture*.

– IETF RFC 3032 (2001), *MPLS Label Stack Encoding*.

– IETF RFC 3036 (2001), *LDP Specification*.

– IETF RFC 3107 (2001), *Carrying Label Information in BGP-4*.

– IETF RFC 3429 (2002), *Assignment of the 'OAM Alert Label' for Multiprotocol Label Switching Architecture (MPLS) Operation and Maintenance (OAM) Functions*.

3 Definitions

This Recommendation defines the following terms:

3.1 Forwarding Equivalence Class (FEC): A category of IP packets that receive the same forwarding treatment.

3.2 Label Switched Path (LSP): The path through one or more LSRs at one level of the hierarchy followed by a packet in a particular FEC.

3.3 Penultimate Hop Popping (PHP): An optional feature in MPLS, where the label stack may be popped (removed) at the penultimate Label Switching Router of the LSP, rather than at the LSP Egress.

4 Abbreviations

This Recommendation uses the following abbreviations:

CE	Customer Edge Router
CR-LDP	Constraint-based Routing – Label Distribution Protocol
DSCP	Differentiated Services Code Point
DST	Destination
EL	Exchange Link
EXP	Experimental
FEC	Forwarding Equivalence Class
IGP	Interior Gateway Protocol
IP	Internet Protocol
IS-IS	Intermediate System to Intermediate System
LDP	Label Distribution Protocol
LSP	Label Switched Path
LSR	Label Switching Router
MP	Measurement Point
MPLS	Multi-Protocol Label Switching
NS	Network Section
NSE	Network Section Ensemble
OSPF	Open Shortest Path First
PDV	Packet Delay Variation
PE	Provider Edge Label Switching Router
PER	Packet Error Ratio
PHP	Penultimate Hop Popping
PIA	Percent service availability
PIU	Percent service unavailability

PLR	Packet Loss Ratio
PRE	Packet Transfer Reference Event
PSLBR	Packet Severe Loss Block Ratio
PTD	Packet Transfer Delay
RSVP-TE	Resource ReserVation Protocol – Traffic Engineering
RTPTD	Round Trip Packet Transfer Delay
SLB	Severe Loss Block
SPR	Spurious Packet Rate
SRC	Source
TLV	Type-Length-Value-tuple
ToS	Type of Service
TTL	Time To Live
UDP	User Datagram Protocol
UNI	User Network Interface

5 Layered protocol reference model and performance model for MPLS

Figure 2 illustrates the layered nature of MPLS transport service. The performance provided to layers above MPLS depends both on the MPLS layer performance and the layers below MPLS:

- Lower layers that provide connection-oriented or connection transport supporting the MPLS layer.
- The MPLS layer that transports packets. This layer has significance across the MPLS Domain(s), and provides the Label Switched Path (LSP). In the case where Penultimate Hop Popping is employed, the label stack is returned to the depth on ingress in the penultimate node.
- Higher Layers, including the IP layer, that further enable end-to-end communication.

This clause defines a generic MPLS transport network performance model, composed of network sections and exchange links that interconnect network sections. The performance parameters defined here may be applied to the unidirectional transfer of packets on a network section, or across a single MPLS Domain, as defined below (this is the scope of OA&M measurements, such as those defined in ITU-T Rec. Y.1711). The parameters may also be applied to combinations of network sections and exchange links, when the mapping between labels and routes has been distributed between AS according to standardized protocols, such as those defined in RFC 3107. The performance parameters are based on reference events and packet transfer outcomes, also defined below.

In particular, Figure 2 illustrates one of the challenges of MPLS architectures, where the MPLS layer may not exist across the entire measurement path.

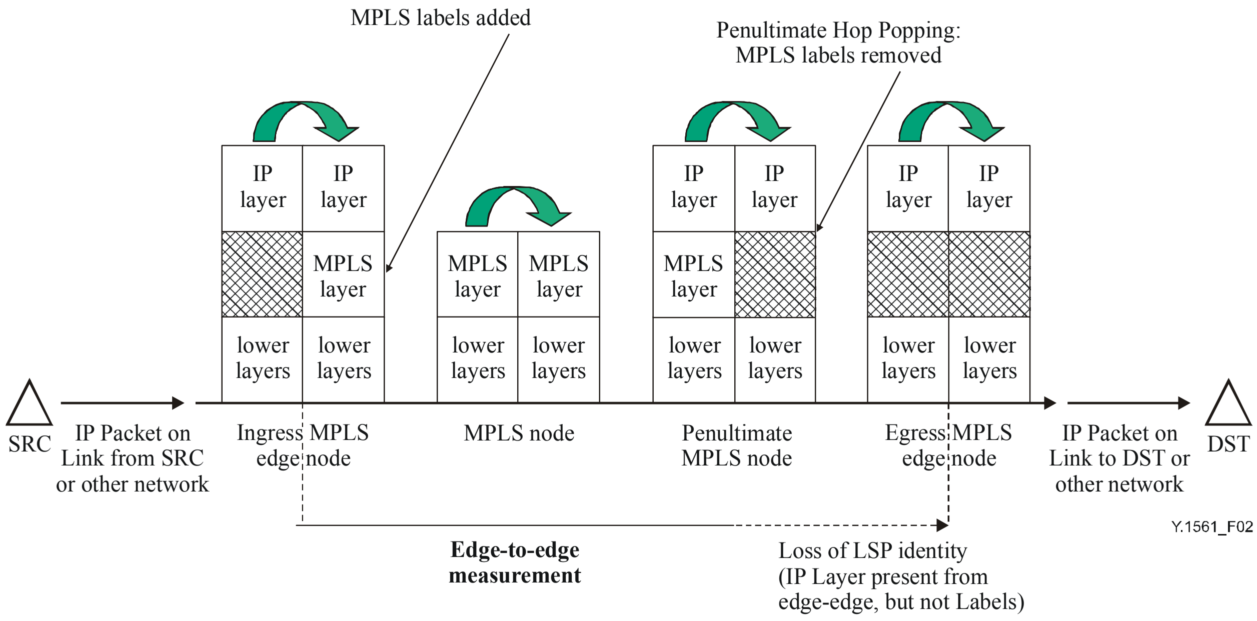

Figure 2/Y.1561 – Layered model of performance for MPLS

As noted earlier, packets using protocols other than IP may be encapsulated with MPLS labels and transported over MPLS networks, but the PHP option would not be used and labels will be present from network edge-to-edge. In some cases, MPLS networks may provide transport between User-Network Interfaces (UNI), and offer an end-to-end transport service to higher-layer protocols.

5.1 Network components

Fundamental components of IP networks are defined in ITU-T Rec. Y.1540.

The following network components have been defined in RFC 3031:

5.1.1 LSR (label switching router): An MPLS node which is capable of forwarding native L3 packets.

NOTE – As used here and in the definitions that follow, L3 refers to the IP layer.

5.1.2 MPLS domain: A contiguous set of nodes which operate MPLS routing and forwarding and which are also in one Routing or Administrative Domain.

5.1.3 MPLS edge node: An MPLS node that connects an MPLS domain with a node which is outside of the domain, either because it does not run MPLS, and/or because it is in a different domain. Note that if an LSR has a neighbouring host which is not running MPLS, that LSR is an MPLS edge node.

5.1.4 MPLS egress node: An MPLS edge node in its role in handling traffic as it leaves an MPLS domain.

5.1.5 MPLS ingress node: An MPLS edge node in its role in handling traffic as it enters an MPLS domain.

5.1.6 MPLS node: A node which is running MPLS. An MPLS node will be aware of MPLS control protocols, will operate one or more L3 routing protocols, and will be capable of forwarding packets based on labels. An MPLS node may optionally be also capable of forwarding native L3 packets.

This Recommendation also defines:

5.1.7 MPLS network: A network that consists of one or more MPLS domains, having one or more LSPs from network ingress node to network egress node.

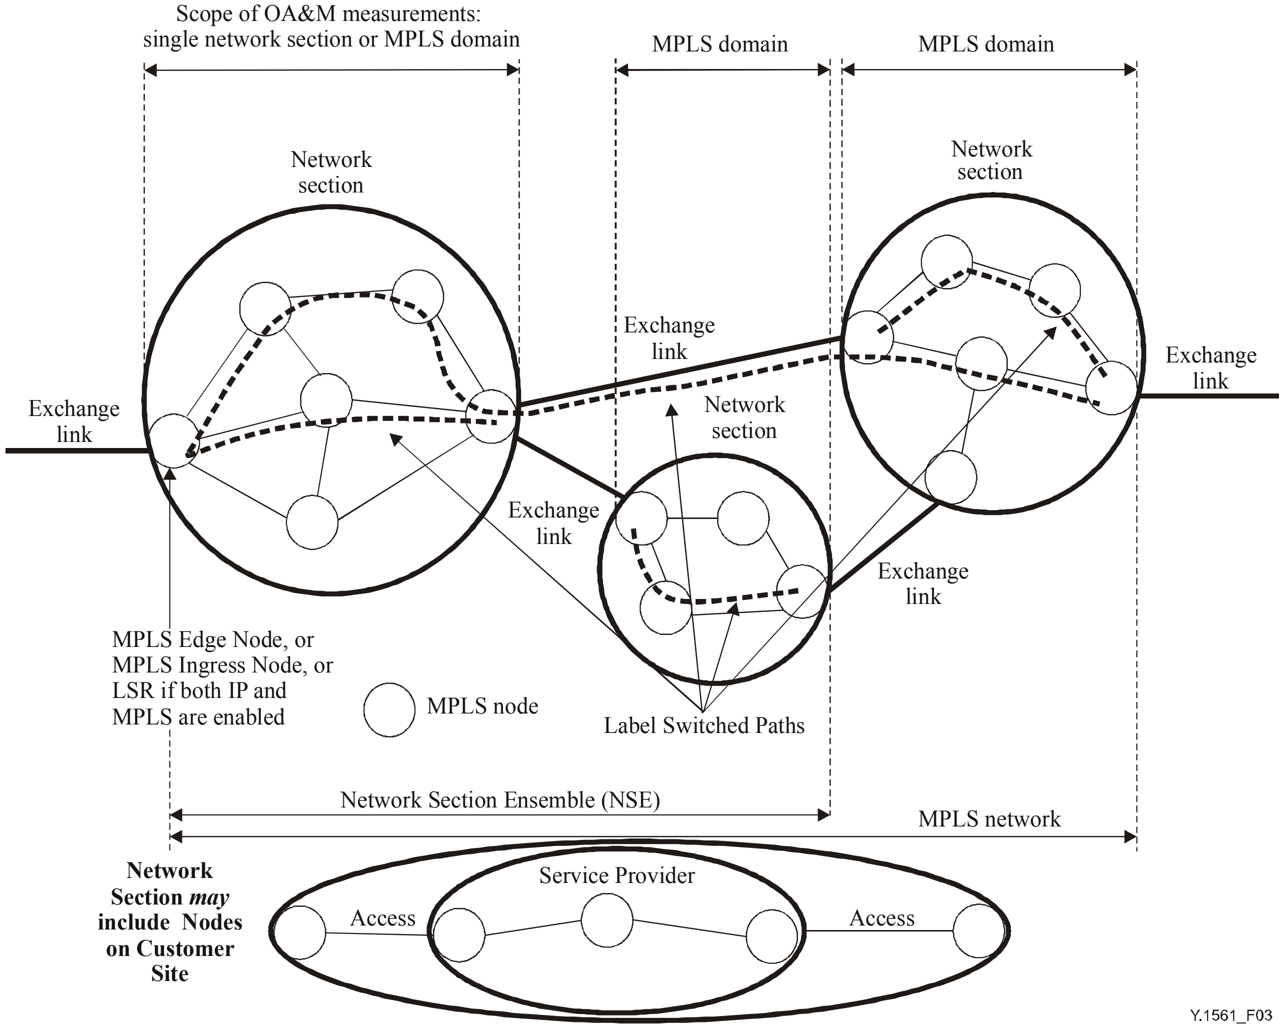

Figure 3/Y.1561 – MPLS network connectivity

Table 1 gives the hierarchical relationship of the terminology used here.

Table 1/Y.1561 – Node terminology

MPLS node							
Edge node				**Interior node**			
Ingress		Egress					
LSR(IP)	Non-LSR	LSR(IP)	Non-LSR				

5.2 Exchange links and network sections

5.2.1 Exchange Link (EL)

The link connecting:

1) a source or destination host (or router) to its adjacent 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 equivalent to the term "label switched hop" as defined in RFC 3031.

5.2.2 Network Section (NS)

A set of MPLS nodes together with all of their interconnecting links that together provide all or part of the MPLS network between an ingress node and an egress node, 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 synonymous with the term "MPLS Domain" as defined in RFC 3031.

5.3 Measurement points and measurable sections

5.3.1 Measurement Point (MP)

The boundary between a host or MPLS Edge Node and an adjacent link at which performance reference events can be observed and measured. Consistent with ITU-T Rec. I.353, any of the standard Internet protocols can be observed at measurement points.

NOTE – The exact location of the MPLS MP within the 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

An EL, NS, a SRC, or a DST. Basic sections are delimited by MP.

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

5.3.3 End-to-end MPLS transport on a label switched path

The set of EL and NS that provide the transport of packets transmitted from MPLS edge node to MPLS edge node on an MPLS network. The MPs that bind the end-to-end MPLS network are the MPs at the ingress node of the first MPLS domain and at the egress node of the last MPLS domain that form the label switched path (LSP).

The end-to-end MPLS network performance is measurable relative to any given unidirectional label switched path. The *ingress MPs* are the MPs crossed by packets from a FEC as they enter the LSP. The *egress MPs* are the MPs crossed by packets from that FEC as they leave that LSP.

5.3.4 Network section ensemble (NSE)

An NSE refers to any connected subset of NSs together with all of the ELs that interconnect them. The term NSE can be used to refer to a single NS, two NSs, or any number of NS and their connecting EL. Pairs of distinct NSEs are connected by exchange links. The term NSE can also be used to represent the entire end-to-end MPLS transport. NSEs are delimited by MP.

The performance of any given NSE is measurable relative to any given unidirectional label switched path provided by the NSE. The *ingress MPs* are the set of MPs crossed by packets from a service as they go into an NSE. The *egress MPs* are the set of MPs crossed by packets from that service as they leave that NSE.

5.4 Packet transfer reference events (PREs)

In the context of this Recommendation, the following definitions apply on a specified end-to-end MPLS Network. The defined terms are illustrated in Figure 4.

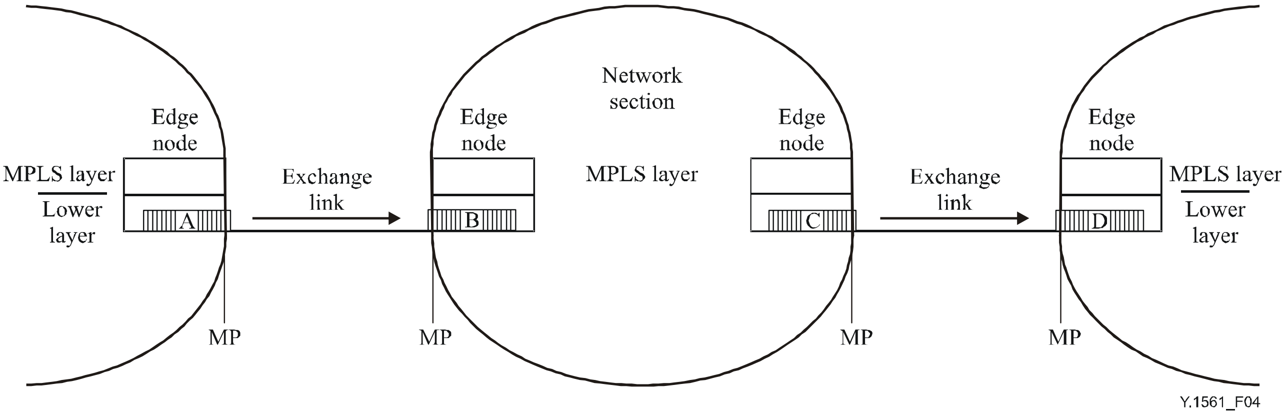

NOTE 1 – Packet exit events for packets A and C.
NOTE 2 – Packet entry events for packets B and D.

Figure 4/Y.1561 – Example packet transfer reference events

A packet transfer event occurs when:

– a packet crosses a measurement point (MP);

– standard procedures confirm that the packet header is valid, e.g., MPLS procedures are applied to validate the label(s), or other header procedures as appropriate;

– the packet is a member of the FEC of interest, as determined by any of the following:

  - the label value within the label contains the expected value and the TTL is non-zero;

    or (the cases below are applicable when PHP removes the label with LSP identity);

  - in the Y.1711 OAM Connectivity Verification (CV) flow packet case, the packet payload contains the OAM function type codepoint, the OAM Payload consistent with the function type, and the Trail Termination Source Identifier field contains the expected LSP ID and the IP address of the expected SRC; or

  - the source and destination address fields within the IP packet header represent the IP addresses of the expected SRC and DST (within the FEC). Information in the packet payload (e.g., inserted by a measurement system) may supplement the header information; or

  - in the LSP-PING packet case, the MPLS echo request packet must be well-formed (valid) at all supporting layers, including the UDP layer and the request format in the UDP payload with the required FEC Stack Type-Length-Value-tuple (TLV).

NOTE 1 – The applicability of Y.1711 OAM messages with PHP is given in RFC 3429. In summary, the ultimate node receiving the OAM packet must be an MPLS LSR to interpret the Y.1711 label and payload correctly. If the ultimate node has no MPLS label look-up or processing, then Y.1711 is not applicable.

NOTE 2 – The MPLS label contains 3 EXP bits and the IP packet header contains information including Type of Service (ToS) or Differentiated Services Code Point (DSCP). This information may affect packet transfer performance, and must be specified if used (set to non-default value).

Packet transfer reference events are defined without regard to packet fragmentation. They occur for every packet crossing any MP regardless of the value contained in the "more-fragments flag". If fragmentation is necessary, an LSR may silently discard the packet (as per RFC 3032).

Four types of packet transfer events are defined:

5.4.1 Packet entry event into a node

A packet transfer entry event into a node occurs when a packet crosses a MP entering a node (LSR or MPLS edge node) from the attached EL.

5.4.2 Packet exit event from a node

A packet transfer exit event from a node occurs when a packet crosses a MP exiting a node (LSR or MPLS edge node) into the attached EL.

5.4.3 Packet ingress event into a basic section or NSE

A packet transfer ingress into a basic section or NSE event occurs when a packet crosses an ingress MP into a basic section or a NSE.

5.4.4 Packet egress event from a basic section or NSE

A packet transfer egress event from a basic section or NSE occurs when a packet crosses an egress MP out of a basic section or a NSE.

NOTE 1 – Packet entry and exit events always represent, respectively, entry into and exit from a node. 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 node, while an ingress into an NS is an entry event because, by definition, NSs always have nodes at their edges.

NOTE 2 – For practical measurement purposes, Packet transfer reference events need not be observed within the protocol stack of the node. Instead, the time of occurrence of these reference events can be approximated by observing the 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 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 packet going to the host or test equipment.

5.5 Packet transfer outcomes

By considering packet transfer reference events, a number of possible transfer outcomes may be defined for any packet attempting to cross a basic section or an NSE. A transmitted packet is either *successfully transferred, errored or lost*. A delivered packet for which no corresponding packet was offered is said to be *spurious*. Figure 5 illustrates the packet transfer outcomes.

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

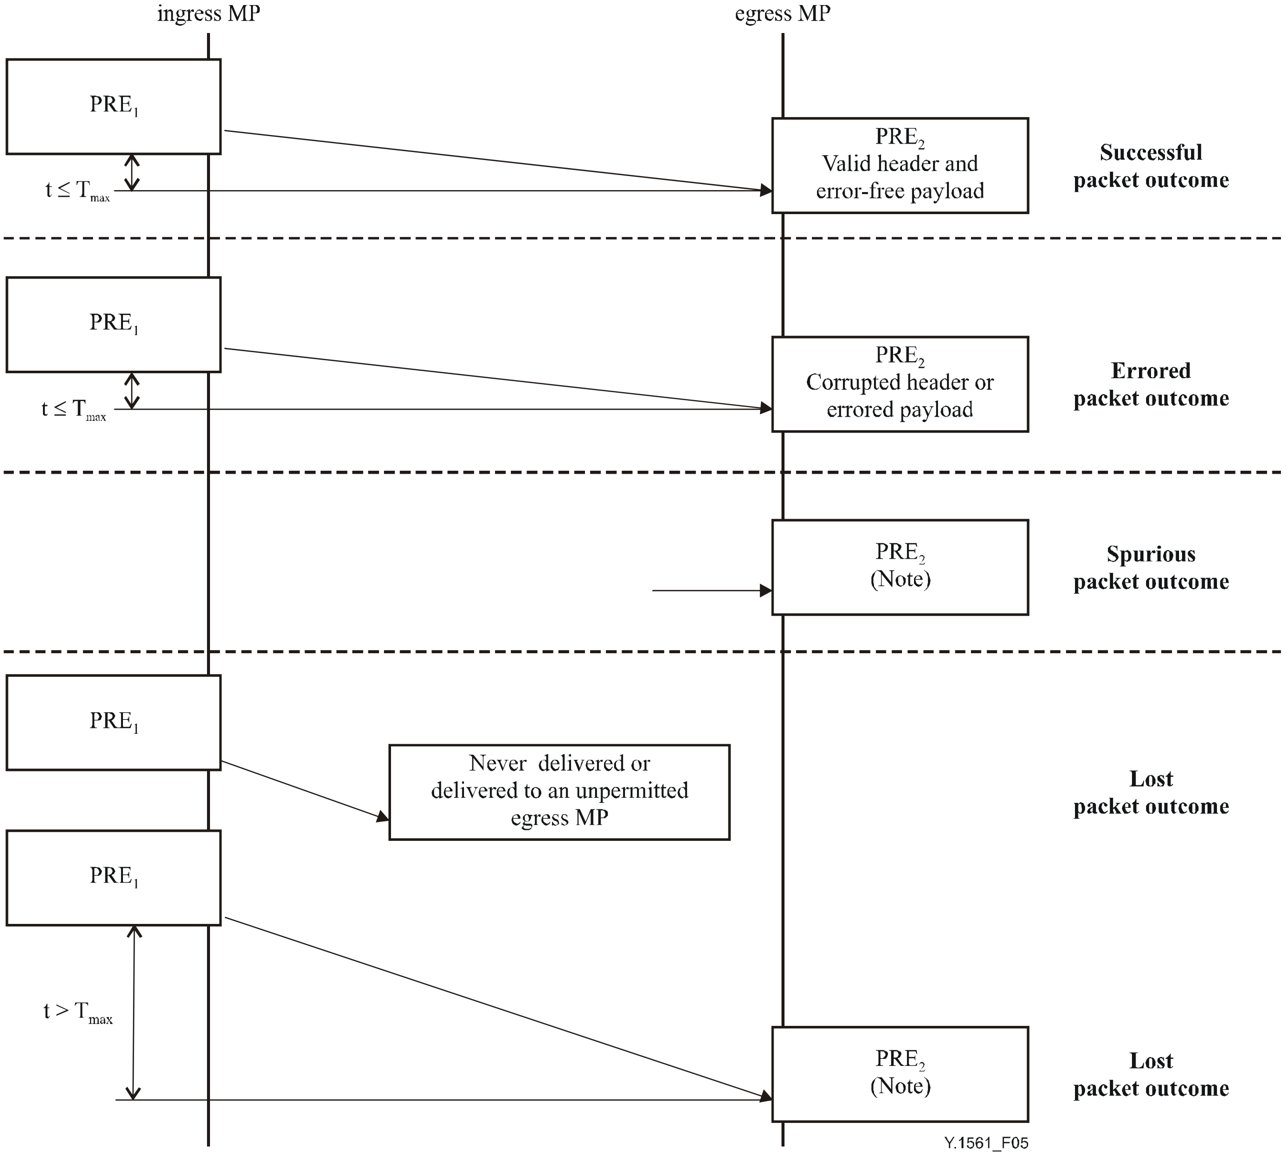

NOTE – Outcome occurs independent of packet contents.

Figure 5/Y.1561 – Packet transfer outcomes

5.5.1 Global routing information and permissible output links

All 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 successful output is to forward the packet to the destination host.

NOTE – Routing protocol 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.

At a given time, and relative to a given end-to-end MPLS network 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 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 egress event is said to *correspond* to an earlier ingress event if they were created by the "same" packet. This concept applies whether the packet at the egress MP is the whole packet or just a fragment of the original. Figure 6 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 reference events are corresponding is usually *ad hoc* and will often rely on consideration of the addresses, the global routing information, the packet identification field, other header information and the packet contents (e.g., the LSP-PING UDP layer and the request format in the UDP payload with the required FEC Stack TLV, or Y.1711 CV flow payload).

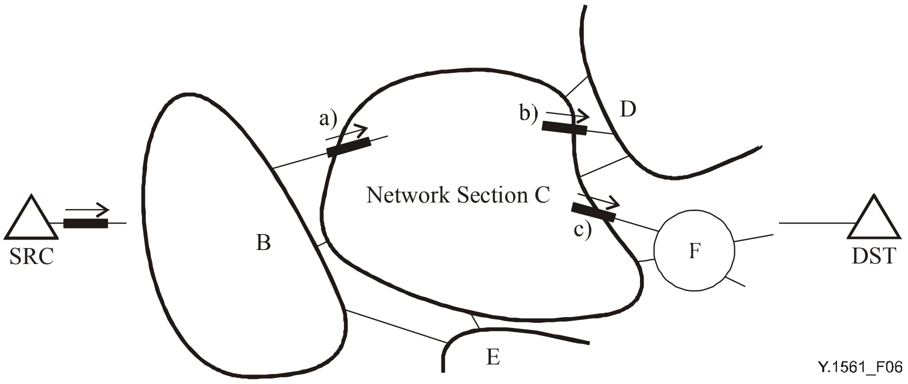

A packet from SRC to DST enters NS C, creates an ingress event, is fragmented, and creates two corresponding egress events, b) and c).

Figure 6/Y.1561 – 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 packet reference events at MPLS measurement points. By selecting the appropriate measurement points, each definition can be used to evaluate the performance of a particular EL, a particular NS, a particular NSE, and they can be applied to the performance of end-to-end networks.

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

In each definition, the possibility of packet fragmentation is accounted for by including the possibility that a single packet 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 EL.

5.5.4 Successful packet transfer outcome

A successful packet transfer outcome occurs when a single 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 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 provisionally set at 3 seconds. Some global end-end paths may require a larger value of T_{max}. The value of 3 seconds has been used in practice.

5.5.5 Errored packet outcome

An errored packet outcome occurs when a single 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 packet information field(s) do not conform exactly with that of the original packet; or
   – one or more of the label or header field(s) of the delivered packet(s) is (are) corrupted.

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

5.5.6 Lost packet outcome

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

A misdirected packet occurs when a single 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 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 packet reference event at any egress MP within the time T_{max}.

5.5.7 Spurious packet outcome

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

5.5.8 Packet severe loss block outcome

A severe loss block (SLB) outcome occurs for a block of packets observed during time interval T_{lb} at ingress MP_0 when the ratio of lost packets at egress MP_i to total packets in the block exceeds s1.

The value of time interval T_{lb} is provisionally set at 1 second. The value of threshold s1 is provisionally set at 0.15. Evaluation of successive blocks (time intervals) should be non-overlapping.

NOTE – The values may change following further study and experience. Current values of T_{lb} and s1 capture network events that may affect the operation of connectivity-sensitive applications. For example, degradation to video and audio applications may be well correlated with the SLB outcome as defined here.

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

5.5.9 Consecutive SLB

When the conditions required for a severe loss block (SLB) outcome occur in successive (non-overlapping) time intervals T_{lb} at ingress MP_0, then a Consecutive SLB outcome occurs.

6 Packet transfer performance parameters

This clause defines a set of information transfer performance parameters using the 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.

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 traversing the Label Switched Path. The measurement points in the end-to-end case are the MP at the MPLS ingress node(s) where packets enter the LSP and at the MPLS egress node where packets exit the LSP.

For a basic section or NSE and relative to a particular LSP, the population of interest at a particular permissible ingress MP is that set of packets traversing the LSP that are routed into the basic section or NSE across that specific MP. This is called the *specific-ingress case*, and applies to any point-to-point LSP, such as those created to provide MPLS-TE/RSVP-TE LSPs.

The total population of interest for a basic section or NSE relative to a particular LSP is the total set of packets traversing the LSP that are delivered into the section or NSE across any of its permissible ingress MP. This is called the *ingress-independent case*, and does not apply with point-to-point MPLS-TE. This population is better suited to characterization of the multipoint-to-point LSP topology.

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

6.2 Packet Transfer Delay (PTD)

Packet transfer delay is defined for all successful and errored packet outcomes across a basic section or an NSE. PTD is the time, $(t_2 - t_1)$ between the occurrence of two corresponding packet reference events, ingress event PRE_1 at time t_1 and egress event PRE_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 packet transfer delay is the one-way delay between the MP at the opposite ends of the LSP as illustrated in Figure 7.

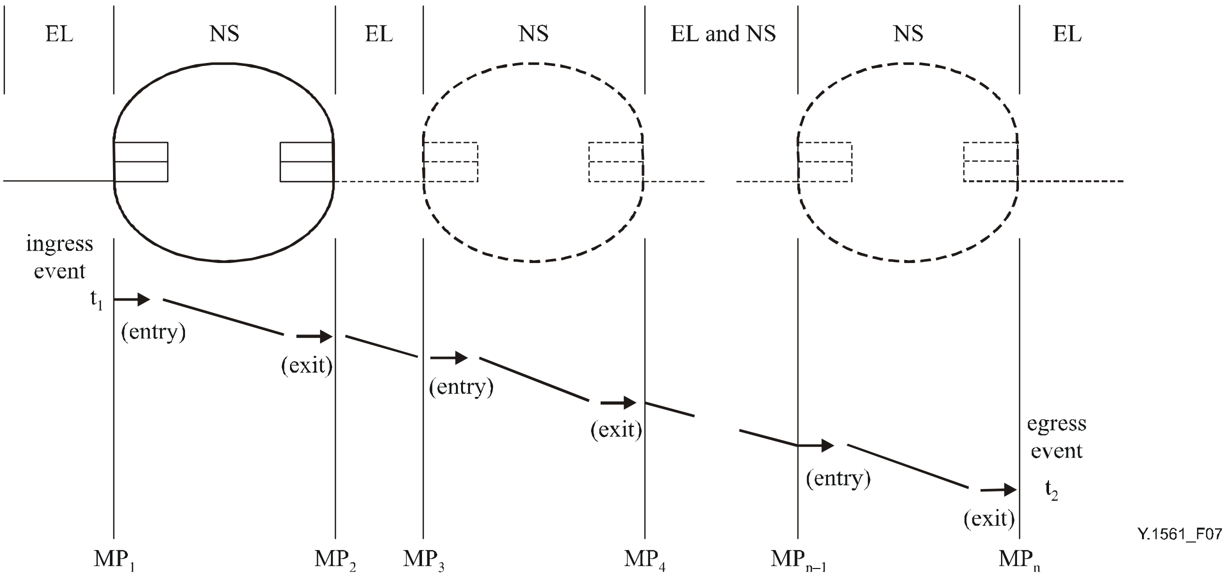

Figure 7/Y.1561 – Packet transfer delay events (illustrated for the end-to-end LSP transfer of a single packet)

6.2.1 Mean packet transfer delay

Mean packet transfer delay is the arithmetic average of packet transfer delays for a population of interest.

6.2.2 End-to-end 2-point Packet Delay Variation (PDV)

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

End-to-end 2-point packet delay variation is defined based on the observations of corresponding packet arrivals at ingress and egress MP (e.g., MP_{DST}, MP_{SRC}). These observations characterize the variability in the pattern of 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 a packet k between SRC and DST is the difference between the absolute packet transfer delay (x_k) of the packet and a defined reference packet transfer delay, $d_{1,2}$, between those same MPs (see Figure 8): $v_k = x_k - d_{1,2}$.

The reference packet transfer delay, $d_{1,2}$, is the absolute packet transfer delay experienced by the first packet between those two MPs (in this example, other reference delays are allowed).

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

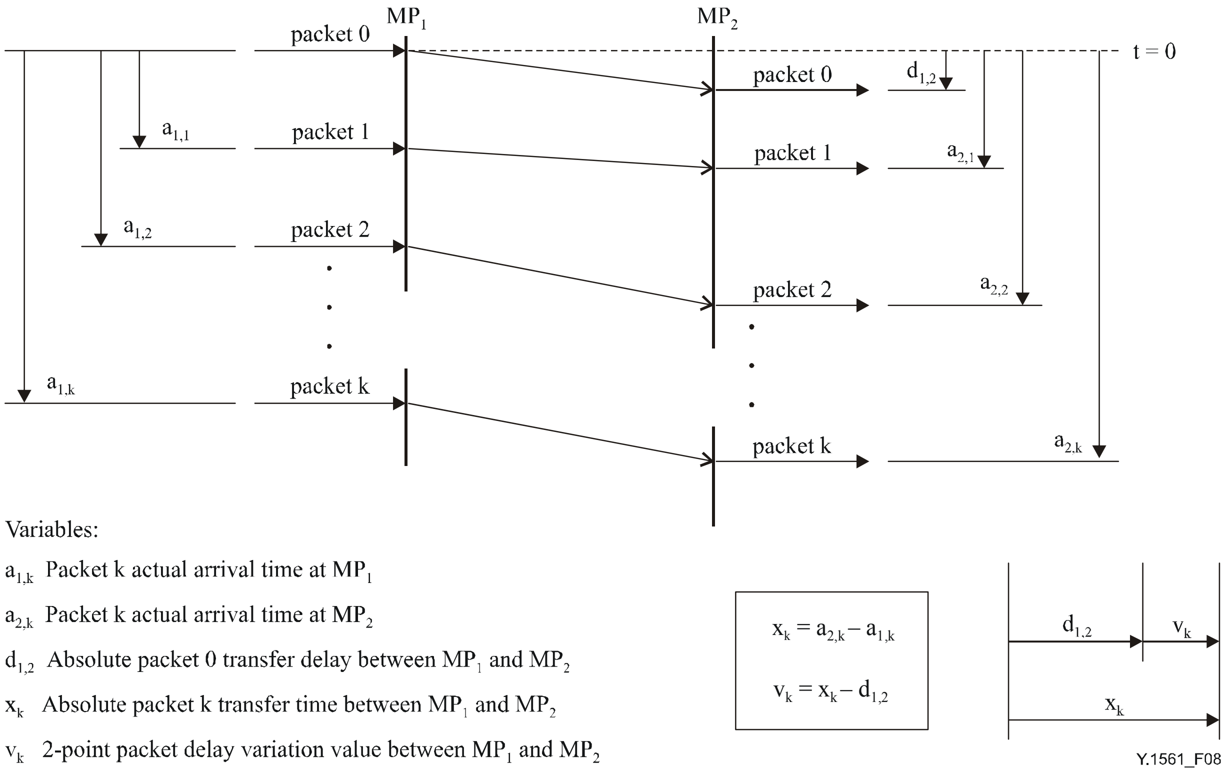

MP$_1$
MP$_2$
t = 0
packet 0
packet 1
packet 2
packet k
$a_{1,1}$
$a_{1,2}$
$a_{1,k}$
$d_{1,2}$
$a_{2,1}$
$a_{2,2}$
$a_{2,k}$
Variables:
$a_{1,k}$ Packet k actual arrival time at MP$_1$
$a_{2,k}$ Packet k actual arrival time at MP$_2$
$d_{1,2}$ Absolute packet 0 transfer delay between MP$_1$ and MP$_2$
x_k Absolute packet k transfer time between MP$_1$ and MP$_2$
v_k 2-point packet delay variation value between MP$_1$ and MP$_2$
$x_k = a_{2,k} - a_{1,k}$
$v_k = x_k - d_{1,2}$
$d_{1,2}$
v_k
x_k
Y.1561_F08

Figure 8/Y.1561 – 2-point packet delay variation

6.2.2.1 Using minimum delay or average delay as the basis for delay variation

As illustrated in Figure 8, 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 centring the distribution of delay variation values on zero (when the distribution is symmetrical).

It simplifies the analysis of delay variation range to use the packet with the minimum delay as the reference delay, and this is a recognized alternative.

6.2.2.2 Interval-based limits on packet delay variation

One method for summarizing the 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 packet 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 underflow.

NOTE – If this method is used for summarizing packet delay variation, the delay variant of individual packets should be calculated using the definition (using the average delay as nominal) in 6.2.2.1, instead of the definition of 6.2.2. Using the definition of 6.2.2, the pre-selected interval (e.g., the ±30 milliseconds) might occasionally be centred on an unusually large or small value.

An objective for 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 using the minimum delay as nominal. For example, "≥ 95% of packet delay variations should be within the interval [0 ms, +30 ms]".

6.2.2.3 Quantile-based limits on packet delay variation

An alternative for summarizing the delay variation of a population of 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 99.9% ile and then 0.1% 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 underflow.

An objective for 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 99.1% ile and the 0.1% ile of the packet delay variation should be no more than 100 milliseconds".

6.2.2.4 Secondary Parameters for packet delay variation

One or more parameters that capture the effect of 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 routing. Appendix II/Y.1540 describes additional delay variation parameters.

6.2.3 Round-Trip Packet Transfer Delay

Round Trip Packet Transfer Delay (RTPTD) is defined as the sum of the one-way delays (PTD) for two LSPs. The pair of LSPs must exist between two MPs at the opposite ends of a basic section or NSE.

Since PTD is the time, ($t_2 - t_1$) between the occurrence of two corresponding packet reference events, the RTPTD only includes the packet transfer time in each direction. The time required to generate or re-generate a packet must not be included. In practice, this issue has been addressed by adding multiple timestamps in test packets (for example, see the Timestamp Request/Reply format in RFC 792).

6.3 Packet Error Ratio (PER)

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

6.4 Packet Loss Ratio (PLR)

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

6.5 Spurious Packet Rate (SPR)

Spurious packet rate at an egress MP is the total number of spurious packets observed at that egress MP during a specified time interval divided by the time interval duration (equivalently, the number of spurious packets per service-second).[1]

6.6 Packet Severe Loss Block Ratio (PSLBR)

A packet severe loss block ratio is the ratio of the packet severe loss block outcomes to total blocks in a population of interest.

NOTE – This parameter can identify path changes due to failures routing updates, and may cause degradation to user applications.

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

6.7 Recovery Time

The count of successive T_{lb} that form a Consecutive SLB outcome at ingress MP_0 is defined as the Recovery Time.

NOTE – As implied by its name, this parameter attempts to capture any form of transient event that interrupts the packet transfer on an LSP for more than one second. Such events may occur when "fast" recovery mechanisms do not restore connectivity with sufficiently small loss ratio.

7 Availability

MPLS service availability is applicable to edge-to-edge service, basic sections and NSE.

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

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

This service function evaluates availability for the following uses:

- connection-oriented packet transfer services;
- continuous stream real-time applications, such as voice and video;
- high-volume interactive packet services, where suspension of packet transfer may cause customer equipment to attempt restoration using alternate networks.

We note that even on connectionless packet transfer networks some fraction of the total may be connection-oriented traffic. Thus, we define a single availability function.

7.1 Availability service function for connection-oriented services

Connection-oriented services require a more continuous packet transfer than other packet services. We have defined a Severe Loss Block (SLB) outcome where the value of time interval T_{lb} is (provisionally) set at 1 second, and the value of the loss threshold s1 is provisionally set at 0.15. Evaluation of successive blocks (time intervals) should be non-overlapping.

Relative to a particular MPLS Ingress Node and Egress Node pair, the availability for *a basic section or an NSE in the specific-ingress case*, is evaluated as follows:

The onset of unavailability begins with the occurrence of ten consecutive SLBs. These ten seconds are part of unavailable time. A period of unavailability ends with the occurrence of ten consecutive seconds, none of which are SLB. These ten seconds are part of available time. The ten-second criteria are supported using a sliding window with one-second granularity.

7.2 Availability parameters

7.2.1 Percent MPLS service unavailability (PIU)

The percentage of total scheduled service time that is categorized as unavailable using the MPLS service availability function.

7.2.2 Percent MPLS service availability (PIA)

The percentage of total scheduled service time that is categorized as available using the MPLS service availability function.

PIU = 100 – PIA

NOTE – Because the PLR typically increases with increasing offered load from SRC to DST, the likelihood of exceeding the threshold s^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.

8 Security

This Recommendation does not specify a protocol. Hence, there are a few areas where security issues may arise, and all are associated with implementation of the performance parameters in measurement systems.

Measurement systems that assess the performance of networks according to the parameter definitions defined in this Recommendation should limit the measurement traffic to appropriate levels to avoid abuse (e.g., Denial of Service Attack). Administrations or Operators should agree on acceptable levels of measurement traffic in advance.

Systems that monitor user traffic for the purpose of measurement must maintain the confidentiality of user information.

Systems that attempt to make measurements may employ techniques (e.g., cryptographic hash) to determine if additional traffic has been inserted by an attacker appearing to be part of the population of interest.

BIBLIOGRAPHY

– IETF RFC 792: Internet Control Message Protocol, *J. Postel*, September 1981.

– Detecting MPLS Data Plane Failures, *IETF work in progress*.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of 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	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	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, Internet protocol aspects and Next Generation Networks**
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2004