International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Y.1563

(01/2009)

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

Internet protocol aspects – Quality of service and network performance

---

Ethernet frame transfer and availability performance

Recommendation ITU-T Y.1563

ITU-T

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.

Recommendation ITU-T Y.1563

Ethernet frame transfer and availability performance

Summary

Recommendation ITU-T Y.1563 defines parameters that may be used in specifying and assessing the performance of speed, accuracy, dependability and availability of Ethernet frame transfer of an Ethernet communication service. The defined parameters apply to end-to-end, point-to-point connections and multipoint connectivity in the Ethernet layer 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.

Source

Recommendation ITU-T Y.1563 was approved on 13 January 2009 by ITU-T Study Group 12 (2009-2012) under Recommendation ITU-T A.8 procedures.

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/.

CONTENTS

Page

Recommendation ITU-T Y.1563

Ethernet frame transfer and availability performance

1 Scope

This Recommendation defines parameters that may be used in specifying and assessing the performance of speed, accuracy, dependability and availability of Ethernet frame transfer on an Ethernet network. It provides a simple set of constructs (terms, definitions and illustrations) and the semantics that allow the performance parameters to be applied to Ethernet networks in many useful ways.

This Recommendation describes Ethernet networks in terms of basic sections, such as networks and links, and measurement points at the boundaries between them. These descriptive terms and concepts have been developed and refined over a series of Recommendations addressing various networking technologies (e.g., ATM, IP, MPLS), and applied to Ethernet technology in this Recommendation. Ethernet networks described using these basic sections (and a few well-known concepts in the networking industry) can be related to Ethernet layer networks defined in [ITU-T G.8010] (architecture of Ethernet layer networks) and Ethernet services defined in [b-ITU-T G.8011]. [ITU-T G.8010] provides a functional specification of Ethernet bridges as defined in [b-IEEE 802.1D] and [b-IEEE 802.1Q]. Appendix VI provides the relationship.

The defined parameters apply to networks that provide, or contribute to the provision of, Ethernet services in accordance with the normative references specified in clause 2.

It applies to point-to-point Ethernet topologies and to point-to-multipoint Ethernet topologies. Multipoint-to-multipoint topologies which add a greater level of complexity are left for further study.

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

Access and disengagement phases are beyond the scope of this Recommendation. In particular, additional Ethernet frames associated to MAC learning procedures are not considered in this Recommendation, since they are related to the access phase.

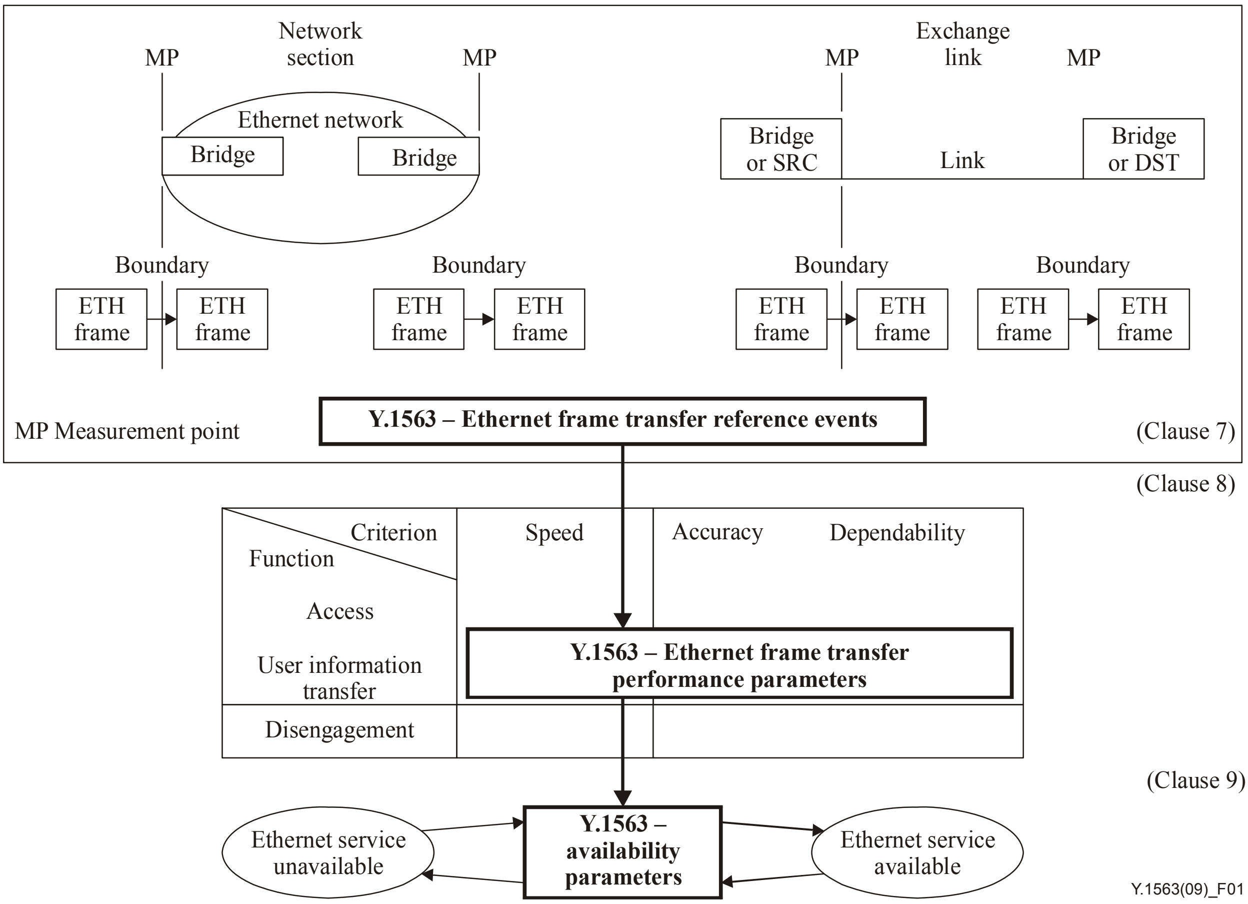

Figure 1 – Scope of this Recommendation

The objectives associated with the Ethernet user information transfer performance parameters are for further study. Another Recommendation will specify these performance objectives.

The end-to-end performance of Ethernet networks providing access and disengagement functions and higher-layer transport capabilities (e.g., Internet 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 lists the terms used and/or defined in this Recommendation. Clause 4 provides a list of abbreviations. Clause 6 illustrates the layered model that creates the context for Ethernet performance specification. Clause 7 defines the model used for Ethernet performance, including network sections and measurement points, reference events and outcomes. Clause 8 uses this model to define Ethernet frame transfer performance parameters. Clause 9 then defines Ethernet service availability parameters. Appendix I describes the implications of sampling large multicast groups on performance parameters definitions and possibly the associated objectives. Appendix II provides alternative parameters which use the framework of point-to-point parameters at registered destinations of a multicast group defined in this Recommendation. Appendix III proposes to harmonize terminologies for frame delay and frame delay variation metrics developed by ITU-T and MEF. Appendix IV describes different networking opportunities making use of Ethernet technology. Finally, Appendix V presents OAM measurement considerations, and Appendix VI addresses the mapping to alternate descriptive terms for Ethernet architectures in [ITU-T G.8010].

NOTE 1 – The speed, accuracy and dependability parameters of this Recommendation are intended to characterize Ethernet service in the available state.

However, no measurements of these parameters made during periods of unavailable time are ever used for determining long-term frame transfer performance, and mechanisms must be established to exclude all performance measurement results collected during unavailable periods from any estimations of long-term Ethernet frame performance parameters.

NOTE 2 – The parameters of this Recommendation 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.

NOTE 3 – The parameters defined in this Recommendation can also be applied to any subset of the Ethernet frames 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 4 – 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 G.805] Recommendation ITU-T G.805 (2000), *Generic functional architecture of transport networks.*

[ITU-T G.809] Recommendation ITU-T G.809 (2003), *Functional architecture of connectionless layer networks.*

[ITU-T G.8010] Recommendation ITU-T G.8010/Y.1306 (2004), *Architecture of Ethernet layer networks*.

[IEEE 802.1ad] IEEE 802.1ad (2005), *IEEE Standard for Local and metropolitan area networks – Virtual Bridged Local Area Networks – Amendment 4: Provider Bridges*.

[IEEE 802.1ag] IEEE 802.1ag (2007), *IEEE Standard for Local and metropolitan area networks – Virtual Bridged Local Area Networks – Amendment 5: Connectivity Fault Management*.

[IEEE 802.1ah] IEEE 802.1ah (2008), *IEEE Standard for Local and metropolitan area networks – Virtual Bridged Local Area Networks – Amendment 6: Provider Backbone Bridges*.

[IEEE 802.3] IEEE 802.3 (2002), *IEEE Standard for Information technology – Telecommunications and Information Exchange Between Systems – Local and Metropolitan Area Networks – Specific Requirements – Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications.*

[IETF RFC 3393] IETF RFC 3393 (2002), *IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)*.

[MEF 10.1] MEF 10.1 (2006), *Ethernet Services Attributes Phase 2*.

3 Definitions

3.1 Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

3.1.1 connection: [ITU-T G.805].

3.1.2 connection point: [ITU-T G.805].

3.1.3 Ethernet frame (equivalent to MAC frame): [IEEE 802.3].

3.1.4 flow: [ITU-T G.809].

3.1.5 link: [ITU-T G.805].

3.2 Terms defined in this Recommendation

This Recommendation defines the following terms in clauses 7, 8 and 9:

3.2.1 comparative group delivery ratio

3.2.2 comparative group loss ratio

3.2.3 delay variation range over group

3.2.4 frame delay variation

3.2.5 frame error ratio

3.2.6 frame loss ratio

3.2.7 frame reference event

3.2.8 frame transfer delay

3.2.9 global mean transfer delay

3.2.10 group loss ratio

3.2.11 group mean transfer delay

3.2.12 loss ratio range over group

3.2.13 mean Ethernet frame transfer delay

3.2.14 mean transfer delay range over group

3.2.15 median Ethernet frame transfer delay

3.2.16 minimum Ethernet frame transfer delay

3.2.17 maximum Ethernet frame transfer delay

3.2.18 multicast group: A multicast group is a set of bridges or end points which have been configured to receive frames sent to a given Ethernet multicast address, or which belong to the same VLAN. This set of bridges and their associated egress measurement points provide a key qualification to the point-to-multipoint population of interest.

3.2.19 percent Ethernet service availability

3.2.20 percent Ethernet service unavailability

3.2.21 severe errored second

3.2.22 spurious Ethernet frame rate

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

ATM Asynchronous Transfer Mode

CCM Continuity Check Message

CE Customer Edge

DST Destination customer edge

EL Exchange Link

EPL Ethernet Private Line

EPLAN Ethernet Private Local Area Network

ETY Ethernet physical layer network

EVC Ethernet Virtual Connection

EVPL Ethernet Virtual Private Line

EVPLAN Ethernet Virtual Private Local Area Network

FCS Frame Check Sequence

FDV Frame Delay Variation

FER Frame Error Ratio

FLR Frame Loss Ratio

FRE Frame Reference Event

FTD Frame Transfer Delay

GFP Generic Framing Procedure

IFDV Inter-Frame Delay Variation

IPDV Internet Protocol Packet Delay Variation

IP Internet Protocol

IPTV Internet Protocol Television

LAN Local Area Network

LOC Loss of Continuity

MAC Media Access Control

MEG Maintenance Entity Group

MEP Maintenance End Point

MMG Mismerge (defect)

MP Measurement Point

MPLS MultiProtocol Label Switching

NNI Network-to-Network Interface

NS Network Section

NSE Network Section Ensemble

OAM Operation, Administration and Maintenance

OTN Optical Transport Network

PDH	Plesiochronous Digital Hierarchy
PE	Provider Edge
PEA	Percent Ethernet service Availability
PEU	Percent Ethernet service Unavailability
R_{max}	Maximum number of successful frames at all destinations in a group, $max(R_n)$
R_n	Number of successful frames at destination n, with respect to a population of interest
RSTP	Rapid Spanning Tree Protocol
S	the number of frames a Source transmits that constitute the population of interest
SDH	Synchronous Digital Hierarchy
SLS	Service Level Specification
SRC	Source Customer edge
STP	Spanning Tree Protocol
TCP	Transmission Control Protocol
T_{max}	Maximum Ethernet frame delay beyond which the frame is declared to be lost
T_{POI}	Interval of time corresponding to transmission of the Population Of Interest
$T_{R,n}$	Interval corresponding to reception of the population of interest at destination n
UNI	User-to-Network Interface
VID	Virtual local area network Identifier
VLAN	Virtual Local Area Network

5 Conventions

None.

6 Layered model of performance for Ethernet service

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

– Lower layers that provide (via "links") connection-oriented or connectionless transport supporting the Ethernet layer. Links are terminated at points where Ethernet frames are forwarded (i.e., "bridges", "SRC" and "DST") and thus have no end-to-end significance. Links may involve different types of technologies, for example, SDH, OTN, PDH, MPLS, ATM and ETY.
– The Ethernet layer that provides transport of Ethernet frames. The Ethernet layer has end-to-end significance for a given pair of source and destination MAC addresses, possibly within the scope of an Ethernet virtual local area network (VLAN).
– Higher layers, supported by Ethernet, that further enable end-to-end communications. Upper layers may include, for example, IP, MPLS and Ethernet (different networking opportunities using Ethernet in higher layers are described in Appendix IV). The higher layers will modify and may enhance the end-to-end performance compared to that provided at the Ethernet layer (for example, frame loss may result in retransmission and delayed delivery to a higher layer).

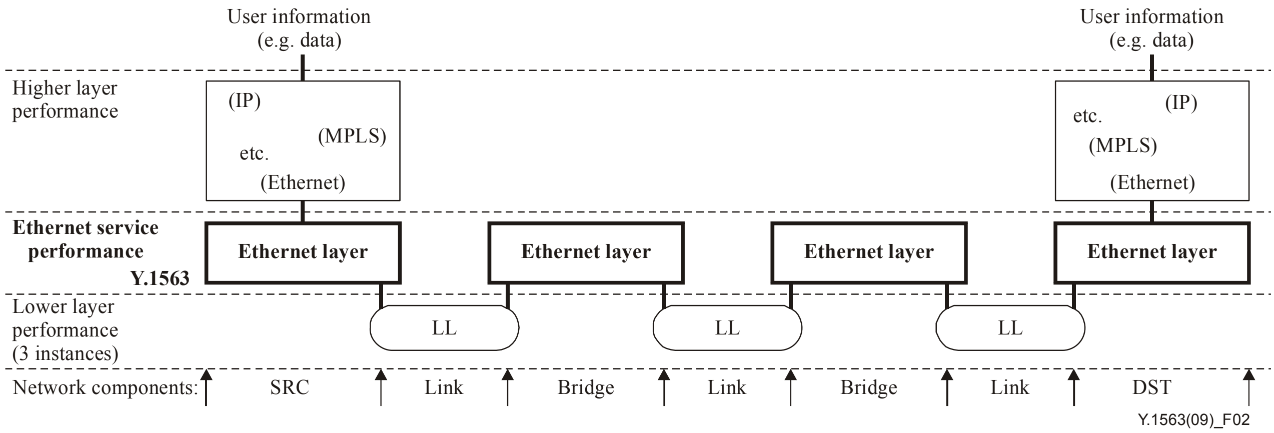

Note – The Ethernet layer is not terminated between the SRC and DST.

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

7 Ethernet service performance model, reference events and outcomes

This clause defines a generic Ethernet service performance model. The model is primarily composed of two types of sections: the exchange link and the network section. These are defined in clause 7.2. They provide the building blocks with which any end-to-end Ethernet service may be represented. Each of the performance parameters defined in this Recommendation can be applied to the unidirectional transfer of Ethernet frames on a section or a concatenated set of sections.

Clause 7.3 defines the measurement points and measurable sections. Clause 7.4 specifies the set of Ethernet frame transfer reference events that provide the basis for performance parameter definition. These reference events are derived from and are consistent with relevant Ethernet service and protocol definitions. Clause 7.5 then uses those reference events to enumerate the possible outcomes when a frame is delivered into a section. Clause 7.6 defines the populations of interest.

7.1 Network components

The ETH topological components are defined in [ITU-T G.8010].

The following groupings of network components are defined.

7.2 Links and network sections

7.2.1 Exchange link (EL)

The link connecting:

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

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

7.2.2 Network section (NS)

A set of bridges together with all of their interconnecting links that together provide a part of the Ethernet service between an SRC and a DST, and are under a single (or collaborative) jurisdictional responsibility. Some network sections consist of a single bridge 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 "transport operator network" as defined in [b-ITU-T G.8011]. It is also possible that a transport operator network contains more than one network section.

Any set of bridges interconnected by links could be considered a network section. However, for the (future) purpose of Ethernet performance allocation, it will be relevant to focus on the set of bridges and links under a single (or collaborative) jurisdictional responsibility. Typically, these bridges belong to the same domain. The destination address, VLAN ID, and the configuration of the Ethernet network (e.g., forwarding rules) dictate the permissible egress port(s) out of this network section (to other NS via exchange links).

7.3 Measurement points and measurable sections

This clause defines the key elements of this Recommendation's performance measurement and specification framework. See Appendix VI for additional information on measurement points and measurable sections in the ITU-T G.8010 architecture.

7.3.1 Measurement point (MP)

A measurement point is the boundary between a bridge and an adjacent link at which performance reference events can be observed and measured.

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.

7.3.2 Basic section

A basic section is either an EL or a NS. Basic sections are delimited by MPs.

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

7.3.3 End-to-end Ethernet network

An end-to-end Ethernet network is the set of EL and NS that provide the transport of Ethernet frames transmitted from SRC to DST. The MPs that bind the end-to-end Ethernet network are the MPs at the SRC and the DST.

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

7.3.4 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 Ethernet network. NSEs are delimited by MPs.

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

7.4 Ethernet frame transfer reference events (FREs)

In the context of this Recommendation, the following definitions apply on a specified point-to-point Ethernet connection or point-to-multipoint Ethernet connection. The defined terms are illustrated in Figure 3.

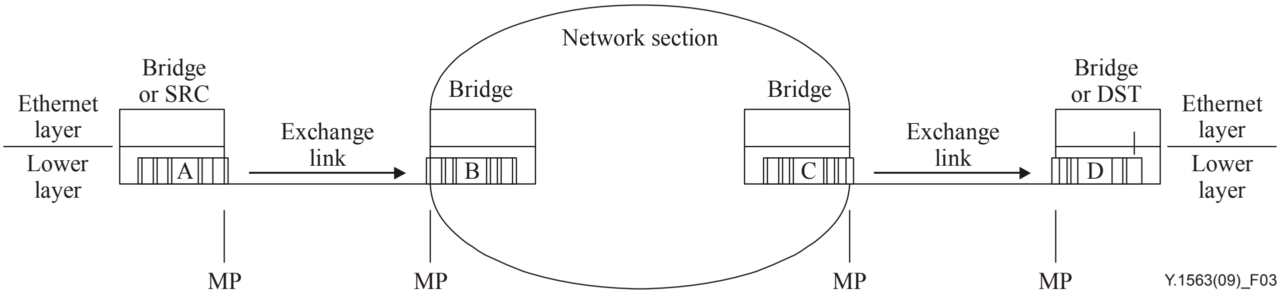

NOTE 1 – Ethernet exit events for frames A and C.
NOTE 2 – Ethernet entry events for frames B and D.

Figure 3 – Example Ethernet frame transfer reference events

An Ethernet frame transfer event occurs when:

– an Ethernet frame crosses a measurement point (MP); and

– standard Ethernet procedures applied to the frame verify that the FCS is valid; and

– the source and destination address fields within the Ethernet frame header contain the expected MAC addresses.

Four types of Ethernet frame transfer events are defined below.

7.4.1 Ethernet frame entry event into a CE

An Ethernet frame transfer entry event into a bridge occurs when an Ethernet frame crosses an MP entering a bridge (PE or DST) from the attached EL.

7.4.2 Ethernet frame exit event from a CE

An Ethernet frame transfer exit event from a bridge occurs when an Ethernet frame crosses an MP exiting a bridge (PE or SRC) into the attached EL.

7.4.3 Ethernet frame ingress event into a basic section or NSE

An Ethernet frame transfer ingress into a basic section event occurs when an Ethernet frame crosses an ingress MP into a basic section.

7.4.4 Ethernet frame egress event from a basic section or NSE

An Ethernet frame transfer egress event from a basic section occurs when an Ethernet frame crosses an egress MP out of a basic section.

NOTE 1 – Ethernet frame entry and exit events always represent, respectively, entry into and exit from an end station. Ethernet frame 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 bridge, while an ingress into an NS is an entry event because, by definition, NSs always have end stations at their edges.

NOTE 2 – For practical measurement purposes, Ethernet frame transfer reference events need not be observed within the MAC layer of the end station. Instead, the time of occurrence of these reference events can be approximated by observing the Ethernet frames 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 an end station is approximated by the observation of the first bit of the Ethernet frame coming from the end station or test equipment. The time of occurrence of an entry event into an end station is approximated by the observation of the last bit of the Ethernet frame going to the end station or test equipment.

7.5 Ethernet frame transfer outcomes

By considering Ethernet frame transfer reference events associated with a particular source-destination pair, a number of elementary Ethernet frame transfer outcomes may be defined for any frame attempting to cross a basic section or an NSE. A transmitted Ethernet frame is either *successfully transferred, errored, misdirected,* or *lost*. A delivered Ethernet frame for which no corresponding Ethernet frame was offered is said to be *spurious*.

Figure 4 illustrates the Ethernet frame transfer outcomes.

The definitions of Ethernet frame transfer outcomes are based on the concepts of *permissible ingress MP*, *permissible egress MP* and *corresponding frames*.

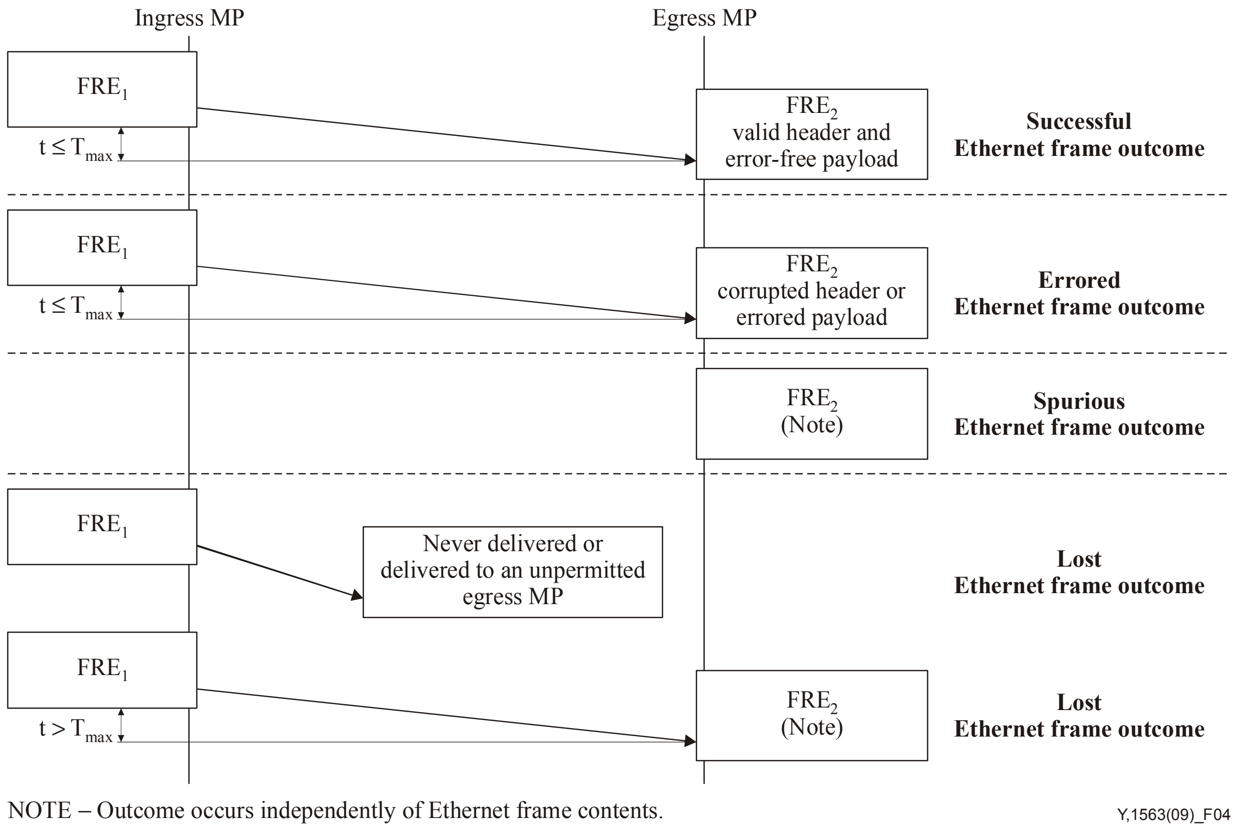

Figure 4 – Ethernet frame transfer outcomes

7.5.1 Global forwarding information and permissible output links

In theory, once connectivity has been established in an Ethernet network, a frame can be delivered to any switch, NS or NSE within the network, and still arrive at its destination. However, global forwarding information for an NS may define a restricted set of destination addresses (MAC address, or a combination of MAC address and VLAN ID) that each network is willing and able to serve on behalf of each of its adjoining NSs. The global forwarding information for an NS is the consolidation of all the local bridging information in that NS. It is reasonable to assume that (in the worst case) an NS will completely discard any frames with destination MAC addresses for which that NS is unable or not configured to serve. Therefore, all Ethernet frames leaving a basic section should only be forwarded to other basic sections as *permitted* by the available global forwarding information.

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

Address learning procedures allow Ethernet switches to learn the MAC addresses of unknown bridges in order to identify on which port the traffic should be sent and establish connectivity. This operation may be considered as part of the access phase of the communication, and as such only the output links (or egress MP) identified in the bridge table for a given end-to-end Ethernet service will be considered permissible.

NOTE – Ethernet procedures and protocols (STP, RSTP) include updating of global forwarding information. An NS that was permissible may no longer be permissible following an update of the forwarding information shared between NSs, and this may include information from other networking technologies that support the Ethernet network. Alternatively, a NS that was not previously permissible may have become permissible after an update of the global forwarding information. The notion of permissible basic sections should be studied further in the context of Ethernet networks.

At a given time, and relative to a given end-to-end Ethernet 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 forwarding 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 forwarding information.

7.5.2 Corresponding events

Performance analysis makes it necessary to associate the frames crossing one MP with the frames that crossed a different MP. A frame should leave a basic section on more than one permissible egress MP when the section performs replication in a point-to-multipoint case or when flooding during MAC learning. Successive frames (intended for a given destination) may use a different egress MP over time due to network topology changes.

An Ethernet egress event is said to *correspond* to an earlier ingress event if they were created by the "same" Ethernet frame. In the point-to-multipoint case, several egress events usually *correspond* to a single ingress event (e.g., MAC learning frames, broadcast frames).

The practical determination of whether Ethernet reference events are corresponding is usually ad hoc and will often rely on consideration of the MAC addresses, VLAN identifiers, the global bridging information, other header information and the Ethernet frame contents.

When maintenance end points (MEPs) are used ([b-ITU-T Y.1731], [IEEE 802.1ag]), accurate frame counting requires access to information about the MEP that corresponds with the source address, and this information can come from provider bridge [IEEE 802.1ad] and provider backbone bridge [IEEE 802.1ah] encapsulations. In the absence of this information, customer MAC addresses need to be maintained in MEPs possibly with a large overhead.

7.5.3 Successful Ethernet frame transfer outcome

A successful frame transfer outcome occurs when a single Ethernet frame 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 frame observed at MP_0 are included in the delivered frame(s); and

3) the binary contents of the delivered Ethernet frame information field(s) conform exactly with that of the original frame; and

4) the header field(s) of the delivered frame(s) is (are) valid.

NOTE – The value of T_{max} is for further study.

7.5.4 Errored Ethernet frame outcome

An errored frame outcome occurs when a single Ethernet frame 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 frame observed at MP_0 are included in the delivered frame(s); and
3) either:
   – the binary contents of the delivered Ethernet frame information field(s) do not conform exactly with that of the original frame; or
   – one or more of the header field(s) of the delivered frame(s) is (are) corrupted.

NOTE – Most frames with errored headers that are not detected by the FCS at the Ethernet layer will be discarded. The result is that no reference event is created for the higher layer protocols expecting to receive this frame. Because there is no Ethernet reference event, these frame transfer attempts will be classified as lost frame outcomes. Errored headers that do not result in discarding or misdirecting will be classified as errored frame outcomes.

7.5.5 Lost Ethernet frame outcome

A lost frame outcome occurs when there is a single Ethernet frame reference event at a permissible ingress MP_0, and when some or all of the contents corresponding to that ingress frame do not result in an Ethernet reference event at a permissible egress MP_i within the time T_{max}.

A lost Ethernet frame outcome may in fact be one or more *misdirected frame* outcomes (which were not observed), as defined below.

A misdirected frame occurs when a single Ethernet frame 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 frame observed at MP_0 are included in the delivered frame(s); but
2) one or more of the egress MP_i where the corresponding reference events occur are not permissible egress MP(s).

7.5.6 Spurious Ethernet frame outcome

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

A spurious Ethernet frame may have diverse origins. As an example, it may result from bit errors modifying the Ethernet frame MAC destination address, but not detected by the FCS.

NOTE – Alternatively, the cause could be an error in a forwarding table. In this case, this should go in the misdelivered outcome.

7.5.7 Secondary Ethernet frame outcomes

The following outcomes are based on the fundamental outcomes described above.

7.5.7.1 In-order and reordered Ethernet frame outcomes

The definition of these Ethernet frame outcomes requires some background discussion.

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

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

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

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

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

7.5.7.2 Duplicate Ethernet frame outcome

A duplicate Ethernet frame outcome is a subset of successful frame outcomes, and occurs when a single Ethernet frame reference event at a permissible ingress MP_0 results in two or more corresponding reference event(s) on at least one permissible egress $MP_{i,}$ and the binary information fields of all the output frames are identical to the original frame. The egress reference event at MP_i for a duplicate frame occurs subsequently to at least one other corresponding egress reference event for the original frame (usually also at MP_i).

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

7.5.7.3 Replicated Ethernet frame outcome

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

7.6 Frame qualifications

Most of the performance parameters are defined over sets of frames called *populations of interest*.

Descriptions of the population of interest must include:

1) The interval of time from the first to the last ingress reference events (at the source MP and destination MP), T_{POI}. The corresponding time interval at a particular egress MP n is $T_{R,n}$.
2) The number of frames in the population (all such frames must correspond to ingress reference events).
3) The set of permissible ingress and egress MPs during T_{POI}.
4) Other qualifying aspects from the frame header, such as the source and group MAC addresses, etc.

7.6.1 Population of interest for the point-to point case

For the *end-to-end case*, the population of interest is usually the total set of frames being sent from SRC to DST, over a VLAN where appropriate. 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 frames being sent from SRC to DST that are forwarded 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 frames 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*.

7.6.2 Ethernet frame flow

An Ethernet frame flow is the set of frames associated with a given connection or connectionless stream having the same source MAC address (SRC), the same destination MAC address (DST) in case of point-to-point connections or the same set of destinations in case of point-to-multipoint connections, the same Ethernet VLAN tag, or the same session identification (e.g., IP addresses or port numbers from a higher-layer protocol). The VLAN tag includes both a VLAN ID and the user priority bits, commonly referred to as the p-bits.

An Ethernet frame flow is the most common example of a population of interest. Some in-service measurements (or monitoring) may be conducted on the aggregate of many frame flows for practical reasons.

7.6.3 Population of interest for the point-to-multipoint case

For the *point-to-multipoint case*, the population of interest is usually the total set of frames being sent from source to a set of destinations that are members of a specific multipoint VLAN (a VLAN with two or more permissible egress MP). The measurement points in the typical NSE "end-to-end" case are the MPs at the source and destinations.

It is important to note that the set of permissible ingress and egress MPs may change during T_{POI}. In some forms of point-to-multipoint communication, the source transmits to the multicast group continuously, and destinations may join or leave the group whenever they wish (this would correspond to a user viewing the live television channels offered in an IPTV system). A destination's group membership activity determines the portion of the population of interest that is relevant to the calculation of its point-to-point parameters.

Thus, when a source has transmitted S frames during T_{POI}, and a specific destination n joins the group while T_{POI} is in progress, then the number of frames relevant to calculating that destination's point-to-point parameters is S_n. The first frame considered to count toward S_n is the first frame observed that is part of a particular multicast group. If a destination leaves the group during T_{POI}, then the last frame considered to count toward S_n is the last frame observed that is part of a particular multicast group. Likewise, the number of frames successfully delivered to a particular destination is R_n.

A set of n destinations is designated by $D = \{D_1, D_2, D_3, D_4, D_5, ..., D_n\}$, registered for the multicast group included in the population of interest.

8 Ethernet frame transfer performance parameters

This clause defines a set of Ethernet frame information transfer performance parameters using the Ethernet frame transfer outcomes defined in clause 7.5. All of the parameters may be estimated on the basis of observations made at MP that bound the basic section or NSE under test.

The performance parameters in packet networks are usually defined in terms of packet loss and packet delay. Ethernet services are no different and their performance can be identified based on frame transfer delay and frame loss. Performance parameters are useful for monitoring the quality of a connection and are usually a part of the overall network management and OAM. The performance of point-to-multipoint frame distribution to a set of destinations can be considered a set of point-to-point frame transfers, and characterized using any or all of the point-to-point parameters. Parameters that are specific to the point-to-multipoint case are also defined below.

The names of the point-to-multipoint parameters employ two adjectives with the following meanings:

- **Global**: equal weighting given to all *frames* in the population of interest.
- **Group**: equal weighting given to the point-to-point parameters calculated for each *destination* that is member of the group.

Both types of parameters take the possibility of group membership changes into account.

Some parameters can be taken as more primary than others. For example, a global frame loss ratio result of zero indicates that all frames have been delivered to all destinations, and there is no need to investigate any per-destination point-to-point results for that population of interest.

8.1 Point-to-point Ethernet frame transfer delay (FTD)

Point-to-point frame transfer delay is the time required to transfer a frame from its source to its intended destination. Using round-trip delays may be worthwhile inclusions in this Recommendation because of their simplicity. These inclusions are for further study.

Ethernet frame transfer delay is defined for all successful and errored frames outcomes across a transport operator network. FTD is the time, $(t_2 - t_1)$ between the occurrence of two corresponding Ethernet frame reference events, ingress event FRE_1 at time t_1 and egress event FRE_2 at time t_2, where $(t_2 > t_1)$ and $(t_2 - t_1) \leq T_{max}$. The end-to-end Ethernet frame transfer delay is the one-way delay between the MP at the SRC and DST as illustrated in Figure 5.

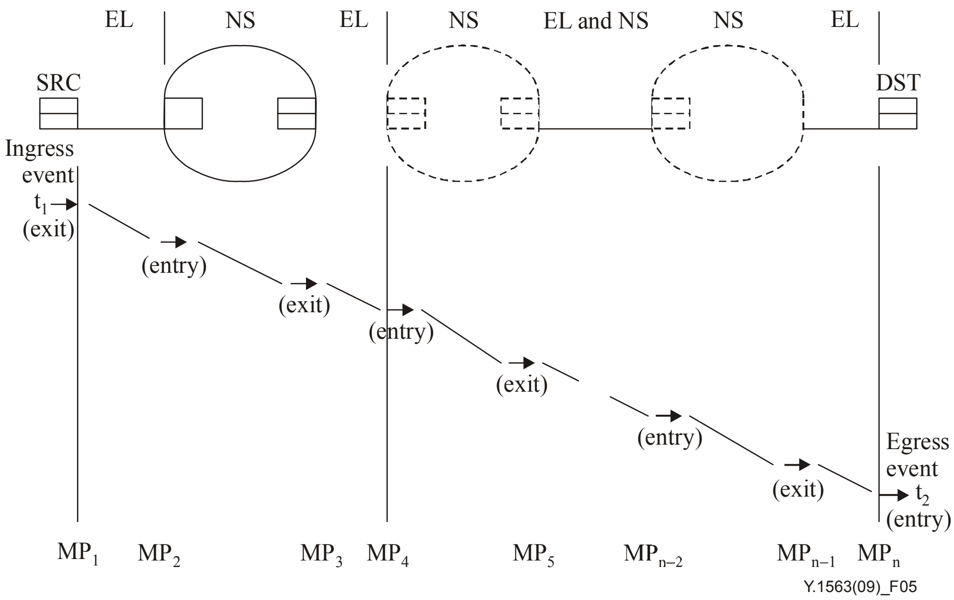

Figure 5 – Ethernet frame transfer delay events (illustrated for the end-to-end transfer of a single Ethernet frame)

Due to the nodal queuing delay, frame transfer delay is a random quantity that is usually characterized by its mean and variations.

8.1.1 Mean Ethernet frame transfer delay

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

8.1.2 Minimum Ethernet frame transfer delay

Minimum Ethernet frame transfer delay is the smallest value of Ethernet frame transfer delay among all Ethernet frame transfer delays of a population of interest. This includes propagation delay and queuing delays common to all frames. Therefore, this parameter may not represent the theoretical minimum delay of the path between MPs.

8.1.3 Maximum Ethernet frame transfer delay

Maximum Ethernet frame transfer delay is the largest value of Ethernet frame transfer delay among all Ethernet frame transfer delays of a population of interest.

8.1.4 Median Ethernet frame transfer delay

The median Ethernet frame transfer delay is the 50th percentile of the frequency distribution of Ethernet frame transfer delays from a population of interest. The median is the middle value once the transfer delays have been rank-ordered. To obtain this value if the population contains an even number of values, the mean of the two central values will be used.

8.1.5 End-to-end 2-point Ethernet frame delay variation

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

End-to-end 2-point Ethernet frame delay variation is defined based on the observations of corresponding Ethernet frame arrivals at ingress and egress MP (e.g., MP_{DST}, MP_{SRC}). These observations characterize the variability in the pattern of Ethernet frame arrival events at the egress MP and the pattern of corresponding events at the ingress MP with respect to a reference delay.

The 2-point frame delay variation (v_k) for an Ethernet frame k between SRC and DST is the difference between the absolute Ethernet frame transfer delay (x_k) of frame k and a defined reference Ethernet frame transfer delay, $d_{1,2}$, between those same MPs (see Figure 6): $v_k = x_k - d_{1,2}$.

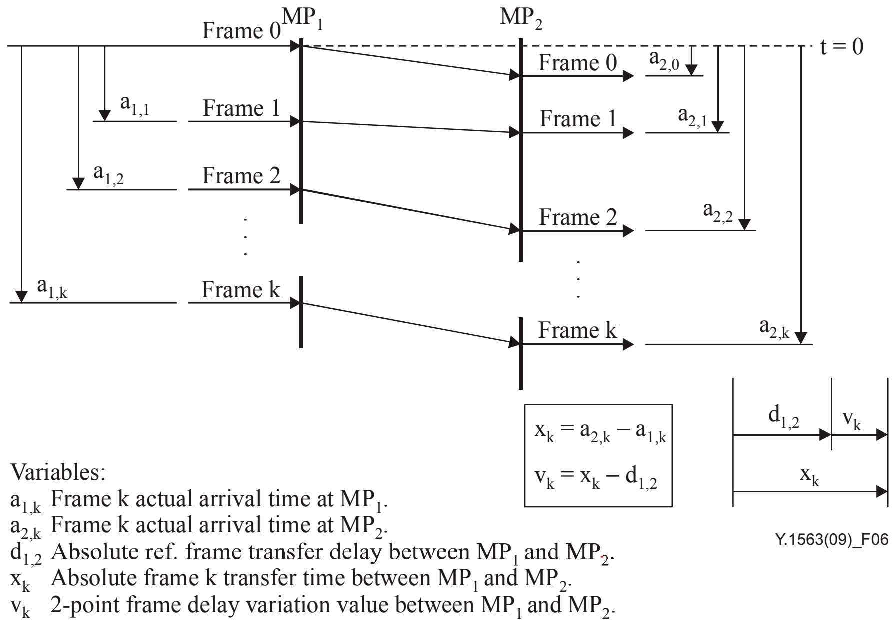

Figure 6 – 2-point Ethernet frame delay variation

The reference Ethernet frame transfer delay, $d_{1,2}$, is the absolute Ethernet frame transfer delay experienced by a selected Ethernet frame between MP_1 and MP_2 .

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

This Recommendation terminology for frame delay and frame delay variation differs somewhat from the terminology used presently by the Metro Ethernet Forum. Appendix III provides the mapping between similar metrics developed by these two standards bodies.

8.1.5.1 Using minimum delay as the basis for delay variation

As illustrated in Figure 6, the delay variation of an individual frame is naturally defined as the difference between the actual delay experienced by that frame and a nominal or reference delay. The preferred reference is the minimum delay of the population of interest. This ensures that all variations will be reported as positive values, and simplifies reporting the range of variation (the maximum value of variation is equal to the range).

There is an alternative to using the minimum frame delay as the nominal delay: to use the average delay of the population of interest as the nominal or reference delay. This has the effect of centring the distribution of delay variation values on zero (when the distribution is symmetrical), and produces both positive and negative variations. However, the average delay of the population may not match the delay of any individual frame, creating an inappropriate reference for variation (e.g., when a bimodal distribution is present).

8.1.5.2 Quantile-based limits on Ethernet frame delay variation

The preferred method for summarizing the delay variation of a population of interest is to select upper and lower quantiles of the delay variation distribution and then measure the distance between those quantiles. For example, select the $1\text{-}10^{-3}$ quantile and the 0 quantile (or minimum), make measurements, and observe the difference between the delay variation values at these two quantiles.

This example would help application designers determine the de-jitter buffer size for no more than 0.1% total buffer overflow.

An objective for Ethernet frame delay variation could be established by choosing an upper bound for the difference between pre-specified quantiles of the delay variation distribution. For example, "the difference between the 99.9 quantile and the minimum of the frame delay variation should be no more than 50 ms."

8.1.5.3 Interval-based limits on Ethernet frame delay variation

An alternative method for summarizing the Ethernet frame delay variation experienced by a population of interest is to pre-specify a delay variation interval, e.g., 50 ms, and then observe the percentage of individual frame delay, variations that fall inside and outside of that interval. If the 50 ms interval were used, applications with fixed buffer sizes of at or near 50 ms would then know approximately how many frames would cause buffer overflow or underflow.

NOTE – If this method is used for summarizing Ethernet frame delay variation, the delay variant of individual frames should be calculated using the minimum delay as nominal in clause 8.1.5.1., instead of the definition of clause 8.1.5 using the first frame. Using the definition of clause 8.1.5, the pre-selected interval (e.g., the 50 ms) might occasionally be anchored on an unusually large or small value.

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

8.1.5.4 Secondary parameters for Ethernet frame delay variation

Annex A gives secondary methods for assessing delay variation. Some aspects of these definitions and guidance on their use is for further study.

8.2 Point-to-multipoint Ethernet frame transfer delay

For the *point-to-multipoint case*, the population of interest (usually the total set of frames that have been sent from the source to a set of destinations that have registered as members of a specific multicast group) forms the matrix displaying frame transfer delays as illustrated in Figure 7. The measurement points in the typical NSE case are the MP at the source and destinations.

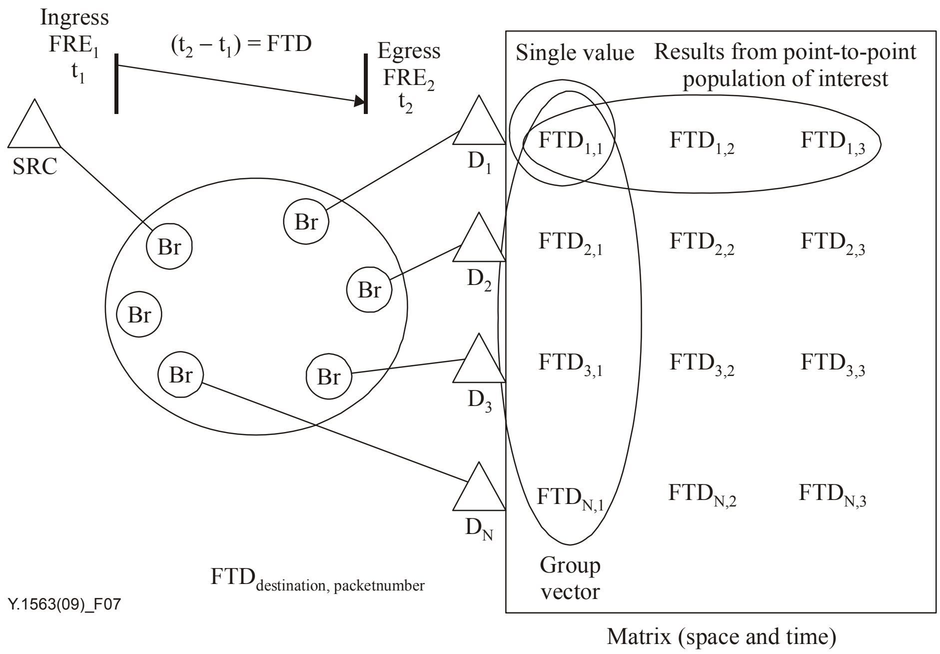

Figure 7 – Illustration of point-to-multipoint terms

In addition to S_n and R_n, which are defined in clause 7.6.3, the other point-to-point outcomes can be represented in a similar way. Lost frame outcomes for destination n are counted as L_n, errored frame outcomes are E_n, spurious frame outcomes as F_n, and so forth. Thus, for each point-to-multipoint population of interest, there are sets of counts as follows:

$$\overline{S} = \{S_1, S_2, ..., S_N\} \quad \overline{R} = \{R_1, R_2, ..., R_N\} \quad \overline{L} = \{L_1, L_2, ..., L_N\}$$

and sets of point-to-point parameters, such as:

$$\overline{FLR} = \{FLR_1, FLR_2, ..., FLR_N\} \quad \overline{FDV} = \{FDV_1, FDV_2, ..., FDV_N\}$$

where the indices are for destinations (these are vectors of point-to-point parameter results).

On the other hand, the set of permissible MPs for an NS may be revised due to forwarding adaptation to equipment failures during T_{POI}. This category of changes to the permissible set is expected to be infrequent.

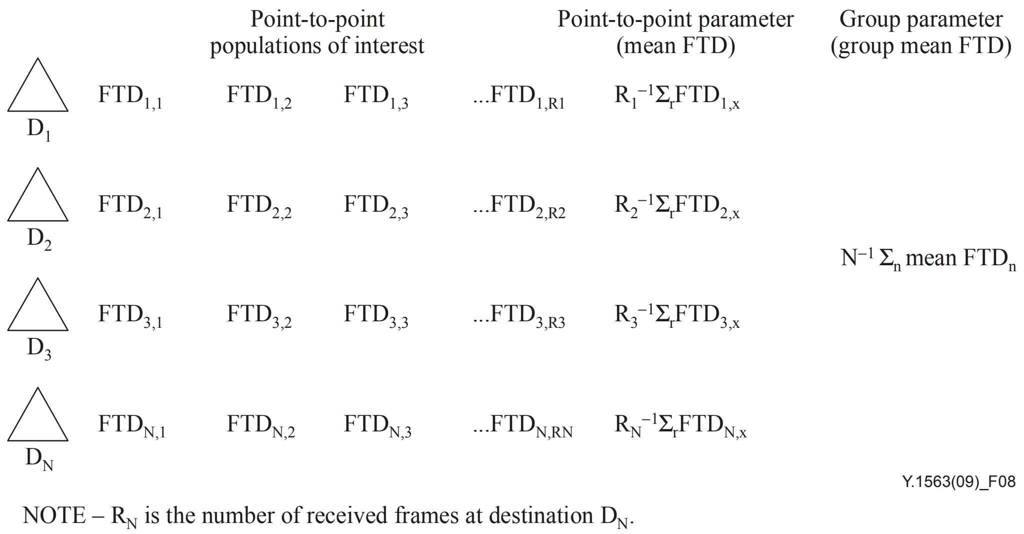

Figure 8 – Illustration of a point-to-multipoint group parameter calculation

Note that with group parameters, *any* statistic can be applied to the point-to-point populations of interest, and then *any* statistic can be applied to the point-to-point parameters, not just the mean as illustrated in Figure 8. With global parameters, *any* statistic can be applied to the complete matrix of point-to-point results. Appendix II discusses the possibilities in more detail.

8.2.1 Global mean transfer delay

The overall mean frame transfer delay for all individual destinations, calculated as the sum of frame transfer delays for all successful Ethernet frame transfer outcomes divided by the total successful Ethernet frame transfer outcomes at all individual destinations.

Using the concepts and symbols introduced above, the mathematical representation of this parameter is:

$$\frac{\sum_{n,r}^{N,R_n} FTD_{n,r}}{\sum_{n=1}^{N} R_n}$$

8.2.2 Group mean transfer delay

The overall mean frame transfer delay for all individual destinations, calculated as the sum of point-to-point mean frame transfer delays for all destinations divided by the number of individual destinations (N). This parameter is illustrated in Figure 8.

The mathematical representation of this parameter is:

$$\frac{\sum_{n=1}^{N} mean(FTD_n)}{N} \text{ where } mean(FTD_n) = R_n^{-1} \sum_{r}^{R_n} FTD_{n,r}$$

8.2.3 Mean transfer delay range over group

This range is determined from the minimum and maximum values of the point-to-point mean frame transfer delays for the set of destinations in the group and a population of interest. Both the minimum and the maximum are recorded, and the range is the difference between the maximum and the minimum.

The mathematical representation of this parameter is:

$$\max(\overline{mean(FTD)}), \min(\overline{mean(FTD)})$$

8.2.4 End-to-end 2-point delay variation range over group

This range is determined from the minimum and maximum values of the point-to-point Ethernet frame delay variation for the set of destinations in the group and a population of interest, where the delay variation is usually expressed as the $1\text{-}10^{-3}$ quantile of one-way delay minus the minimum one-way delay. If a more demanding service is considered, one alternative is to use the $1\text{-}10^{-5}$ quantile and, in either case, the quantile used should be recorded with the results. Both the minimum and the maximum are recorded, and both values are given to indicate the range.

The mathematical representation of this parameter is:

$$\max(\overline{FDV}), \min(\overline{FDV})$$

8.3 Ethernet frame error ratio (FER)

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

8.4 Ethernet frame loss ratio (FLR)

8.4.1 Ethernet frame loss ratio

The ratio of total lost Ethernet frame outcomes to total transmitted Ethernet frames in a population of interest.

In point-to-multipoint configurations, it can also be useful to compare the successful frame transfers among destinations using the destination with the largest number of successful frame transfers as the reference.

8.4.2 Global loss ratio

The overall loss ratio for all individual destinations and a frame population of interest, calculated as the sum of all lost frame outcomes divided by the sum of frames transmitted to each destination while a member of the specified group.

Using the concepts and symbols introduced above, the mathematical representation of this parameter is:

$$\frac{\sum_{n=1}^{N} L_n}{\sum_{n=1}^{N} S_n}$$

8.4.3 Mean group loss ratio

The mean group loss ratio for all registered destinations and a frame population of interest is calculated as the sum of all point-to-point Ethernet frame loss ratios (FLRs) divided by the number of registered destinations that were members of the specified group during T_{POI}.

The mathematical representation of this parameter is:

$$\frac{\sum_{n=1}^{N} FLR_n}{N} \text{ where } FLR_n = \frac{L_n}{S_n}$$

8.4.4 Loss ratio range over group

The loss ratio range is determined from the minimum and maximum values of the point-to-point Ethernet frame loss ratios for the set of destinations in the group and a population of interest. Both the maximum and the minimum are recorded, and both values are given to indicate the range.

The mathematical representation of this parameter is:

$$\max(\overline{FLR}), \min(\overline{FLR})$$

This parameter may be based on a low percentile and a high percentile, such as 1% and 99% or other values, rather than minimum and maximum, because these cannot be known with certainty when sub-sampling the population. In any case, the basis of the range must be included with the results. See Appendix I for further discussion.

8.4.5 Comparative group delivery ratio

The ratio between the number of successful Ethernet frame transfer outcomes R_n, for a particular individual destination D_n, and the largest number of successful Ethernet frame transfer outcomes at another individual destination, designated R_{max}, for the population of interest, where both destinations were registered group members throughout T_{POI}. The mathematical representation of this parameter is:

$$R_n / R_{\max}$$

Note that the use of R_{max} enables a destination-only assessment, but R_{max} may not equal the transmitted frame count, S for the population of interest. Also, note that the one's-complement of this parameter would be the comparative group loss ratio

8.5 Spurious Ethernet frame rate

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

8.6 Frame rate-related parameters

Two types of rate-related parameters may currently be defined. One type of parameter measures the rate in terms of the rate of successfully transmitted Ethernet frames; another type of parameter is octet-based and measures the rate in terms of the octets that have been transmitted in those frames.

Point-to-point Ethernet frame rate (FR): For a given population of interest, the Ethernet frame rate at an egress MP is the total number of Ethernet frame transfer reference events observed at that egress MP during a specified time interval divided by the time interval duration, T_R (equivalently, the number of Ethernet frame transfer reference events per service-second).

Accounting for the possibility that destinations may join or leave a group during T_{POI}, the mathematical representation of this parameter is:

$$R_n / T_{R,n}$$

where $T_{R,n}$ is the time interval corresponding to reception of the population of interest at destination n.

Point-to-point octet-based Ethernet frame rate (FOR): For a given population of interest, the octet-based Ethernet frame rate at an egress MP is the total number of octets transmitted in **Ethernet frame payloads and headers** (including FCS) that result in an Ethernet frame transfer reference event at that egress MP during a specified time interval divided by the time interval duration, T_{POI} (equivalently, the number of octets in the Ethernet frames resulting in Ethernet frame reference events per service-second).

Group mean Ethernet frame rate: The overall mean Ethernet frame rate for all registered destinations, calculated as the sum of FR for all destinations divided by the number of registered destinations (N).

Group mean octet-based Ethernet frame rate: The overall mean Ethernet frame rate for all registered destinations, calculated as the sum of FOR for all destinations divided by the number of registered destinations (N).

The mathematical representation of these parameters is:

$$mean(\overline{FR}) = \frac{\sum_{n=1}^{N}(FR_n)}{N} \quad \text{and} \quad mean(\overline{FOR}) = \frac{\sum_{n=1}^{N}(FOR_n)}{N}$$

One-way Ethernet frame rate range over group: This range is determined from the minimum and maximum values of the point-to-point mean one-way Ethernet frame rate for the set of Destinations in the group and a population of interest. Both the minimum and the maximum are recorded, and the range is the difference between the maximum and the minimum.

The mathematical representation of this parameter is:

$$\max(\overline{FR}), \min(\overline{FR})$$

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

(and similar for octet-based rate).

8.7 Comparison with objectives (general calculation for all point-to-multipoint parameters)

Typically, the users of performance parameters need to make comparisons with objectives. This clause treats the point-to-multipoint parameters as a general case. Results collected for a population of interest and a set of registered destinations should be compared with an objective, O, as follows:

Percent meeting objective (PMO): The percentage of total destinations with point-to-point performance that is categorized as meeting the stated objective for a specific population of interest.

The objectives are evaluated over sets of point-to-point parameters, such as the following for FLR:

$$\frac{\text{Count}(\overline{FLR} | FLR_n \leq O_{IPLR})}{\text{Count}(\overline{FLR})} \times 100$$

where the Count( ) function determines the number of elements in the set that meet the stated condition.

8.8 Organization of parameters according to use case

This clause categorizes the performance parameters according to the audience most likely to benefit from the results expressed in those terms. In Table 1, customer representatives are the persons responsible for a large community of users, and may act on the user's behalf for contract negotiations or bill-paying. Network operators and individual users are also included.

Table 1 – Point-to-multipoint performance parameters organized by use case

	Customer representative-oriented parameters	Network operator-oriented parameters
Throughput (rates)	Min and max rates over group	Mean, min and max rates
Loss	Group loss ratio	Loss ratio range over group
Delay	Global mean	Delay range over group
Delay variation	Range over group	Range over group

The point-point metrics are best suited for the needs of individual users.

Availability is important to all categories, and parameters are defined in clause 9.

9 Ethernet service availability

The Ethernet service availability definition is based on a model which uses two states corresponding to the ability or inability of the network to sustain the service in the available state. Transitions between the states of the model are governed by the occurrence of patterns of severe errored seconds in the Ethernet layer (SES_{ETH}). This Recommendation views availability from the network perspective, where availability performance is characterized independently of user behaviour.

9.1 Severe errored second (SES_{ETH})

A severe errored second (SES_{ETH}) outcome occurs for a block of frames observed during a one-second interval at ingress MP_0 when the corresponding FLR (i.e., the ratio of lost frames to total frames in the block) at egress MP_i exceeds s_1.

A provisional value s_1 of 0.5 is proposed, and different values may also be chosen depending on the class of service (CoS).

Evaluation of successive one-second intervals is non-overlapping.

NOTE – Measurement methodologies for the service availability definition are beyond the scope of this Recommendation, and are for further study. The minimum number of frames that should be used in evaluating the SES_{ETH} outcome, M, is one aspect that requires further consideration, along with the state designation when the number of frames present is less than M (when using non-dedicated frames, or live traffic, as a basis for measurement).

Note that provisional values are subject to change following additional study and real network experience.

9.2 Criteria for entry to and exit from unavailable state

A period of unavailable time begins at the onset of 10 consecutive SES_{ETH} outcomes. The corresponding period of time is considered to be part of unavailable time. During the unavailable time period, the Ethernet network is in unavailable state. A new period of available time begins at the onset of 10 consecutive non-SES_{ETH} outcomes. The corresponding period of time is considered to be part of available time. During the available time period, the Ethernet network is in available state. Figure 9 illustrates the definition of criteria for transition to/from the unavailable state.

This definition of availability has been chosen to allow comparison with other link layer techniques.

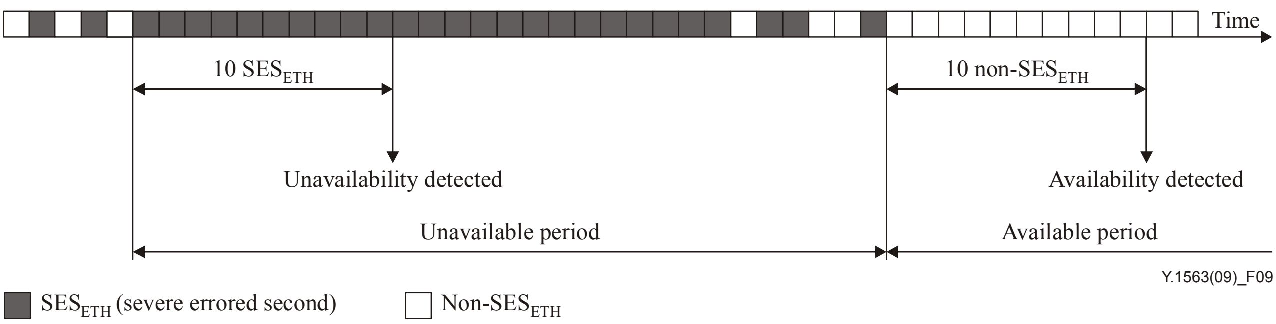

Figure 9 – Example of unavailability determination

All 10-second intervals evaluated for state determination must be entirely composed of scheduled service time (the time interval when the service is intended to be operational, and is usually specified in a service agreement). This means that all seconds of scheduled service time are evaluated at least once for state determination purposes.

Because an Ethernet service is bidirectional, an Ethernet network is in the unavailable state if either one, or both directions, are in the unavailable state. The unidirectional availability can be measured by the criteria mentioned above.

9.3 Ethernet service availability parameters

9.3.1 Percent Ethernet service unavailability (PEU)

The percentage of total scheduled Ethernet service time (the percentage of one-second intervals) that is (are) categorized as unavailable using the Ethernet service availability function.

9.3.2 Percent Ethernet service availability (PEA)

The percentage of total scheduled Ethernet service time (the percentage of one-second intervals) that is (are) categorized as available using the Ethernet service availability function:

$$PEU = 100 - PEA$$

Annex A

Secondary terminology for assessing delay variation

(This annex forms an integral part of this Recommendation)

A.1 Introduction

The body of this Recommendation specifies a primary/normative definition for frame delay variation which will be the basis for future numerical objectives on performance of Ethernet services. This annex describes a secondary method based on IETF's inter-packet delay variation as used in the Metro Ethernet Forum (MEF), and briefly lists areas intended for further study.

A.2 Definition of IP packet delay variation

[IETF RFC 3393] defines delay variation as follows:

- A definition of the IP packet delay variation (IPDV) can be given for packets inside a stream of packets.
- The IPDV of a pair of packets within a stream of packets is defined for a selected pair of packets in the stream going from measurement point MP_1 to measurement point MP_2.
- The IPDV is the difference between the one-way delay of the selected packets.

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

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

The IETF definition has considerable flexibility. Either the ITU-T Y.1563 2-point frame delay variation parameter (clause 8.1.5) or the MEF 10.1 method described below can be derived by specifying the appropriate selection functions.

A.3 Definition of inter-frame delay variation (IFDV) for Ethernet frame transfer

The IFDV $\left(v_k^{'}\right)$ is the difference between the frame reference arrival (c_k) time and the actual arrival time (d_k) at the point of interest, e.g., network egress.

$$v_k^{'} = d_k - c_k$$

Where $c_k = d_j + \tau_{kj}$, j < k and τ_{kj} is the inter-arrival time between the j-th and the k-th frames as shown in Figure A.1. The IFDV parameter may take positive and negative values. Negative values are an indication of the network tendency to clump frames together while positive values are an indication of the network tendency to spread frames apart. The relationship between the various delay values is depicted in Figure A.1. Figure A.1 shows a scenario where a frame arrives at the egress before its reference arrival time.

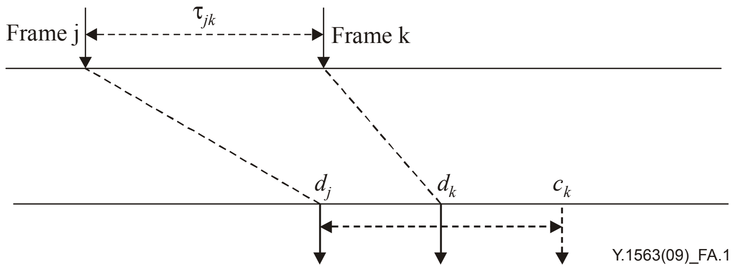

Figure A.1 – Inter-frame delay variation

Since the inter-frame time τ_{ij} is generally a variable, one can choose the two frames *j* and *k* such that there is a predefined fixed separation between them, Δt ($\tau_{ij} = \Delta t$, for all *i* and *j*). In this case, the reference arrival for any frame k is equal to:

$$c_k = d_j + \Delta t$$

where d_j is the arrival time of a frame that is separated by time Δt from frame k at ingress.

A.4 Summarizing IFDV performance

There could be a number of ways for summarizing the IFDV performance. One method that has seen implementation and is used for certification of Ethernet products is that specified in [MEF 10.1]. With this method, the absolute values of the IFDV $\left|v'_k\right|$ are taken, and the IFDV value that achieves a predetermined percentile, P, is determined.

A.5 Areas for further study

When considering the IFDV method of measurement, the following aspects should be studied further:

- What value(s) should be recommended for the ingress packet spacing, Δt?
- How many values of Δt are necessary to provide a sufficient characterization of network performance?
- Must each frame pair be separated by Δt? Or, should the frame pairs overlap in time?
- What are the measurement applications where this method is suitable?

Appendix I

Implications of sampling large groups

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

I.1 Introduction

There are always limitations on the practical size of the group measurements. Specifically, the scale of large multicast groups tends to prevent measurements at every member destination.

I.2 Performance at sample destinations

It should be permissible to measure at a sub-set of the destinations in a group and report the measured range of loss ratio variation, or other parameters.

If the sub-set can be selected with knowledge of the multicast tree structure, then all of the tree but the final links to the un-sampled destinations can be assessed.

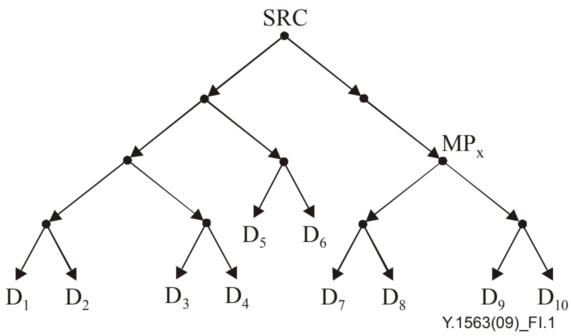

Figure I.1 – Sample of destinations

For example, if the performance of all odd numbered or all even numbered destinations is sampled, then all shared branches and nodes of the tree are assessed.

Note that the population of interest could be defined as described in the body of this Recommendation, but with the additional qualification "that are successfully delivered to the intermediate node with MP_x". This flexibility allows assessment of a partial multicast tree, where R_x becomes the effective S for this population of interest, and frame loss parameter results for destinations 7 through 10 may assist in the isolation of a cause of loss within the tree.

If objectives are developed for a multicast tree, then the measurement and sampling methodology is critical in determining how the objectives are constructed, and possibly even the numerical values set for the objectives.

The body of this Recommendation specifies several parameters that emphasize the performance range across all destinations in a group. When a sub-set of destinations is measured, it is possible to report the sample range, or to report the range between several percentiles (1% and 99%, for example) because the true performance range of the group cannot be known with certainty.

Appendix II

Alternative parameters using the framework of this Recommendation

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

II.1 Introduction

It is possible to calculate many different performance parameters using the framework of point-to-point parameters at registered destinations of a multicast group. At the same time, it is impossible to anticipate all the ways that point-to-multipoint parameters might be used, and therefore only the statistics that are anticipated to be relevant for many uses have been specified in the body of this Recommendation. This appendix is intended to give a view of the alternate possibilities.

II.2 Point-to-multipoint performance parameters

The lists that follow give examples of the point-to-multipoint parameters that can be created from frame transfer outcomes and point-to-point parameters. Parameters expressed as ratios are treated separately from continuous-value parameters.

For the point-to-multipoint case, the atomic measurement results (the point-to-point measurements) can be organized into a matrix where the horizontal axis represents the time and the vertical axis represents the destinations, as in Figure 8. This matrix can be processed by rows (creating point-to-point parameters), by columns (creating vectors describing the performance at all destinations for a given frame transfer attempt) or the entire matrix can be processed. This post-processing could be a minimum/maximum, mean/median or range calculation as well. Then these post-processed parameters can be further post-processed (e.g., a column-maximum can be averaged over time, etc.).

As an example, the following interpretation is presented:

Loss (and other outcomes that use a ratio in the parameter):

- Point-to-point
  - Single outcome
    - i) a frame is lost/errored/...
  - Calculated values for a population of interest
    - i) ratio (e.g., Ethernet frame loss ratio in clause 8.4.1)
- Point-to-multipoint
  - Ratio
    - i) singles (→ group vector)
    - ii) populations of singles (→ whole matrix)
    - iii) of ratios
  - Mean
    - i) of ratios (e.g., mean group loss ratio in clause 8.4.3)

Delay (and others that produce results in a continuous range):

- Point-to-point
  - Single value
    - i) delay of a frame
    - ii) delay variation of two frames

- Calculated values for a population of interest
    - i) minimum/maximum
    - ii) mean/median
    - iii) range

• Point-to-multipoint
  - Minimum/maximum
    - i) singles (→ group vector)
    - ii) populations of singles (→ whole matrix)
    - iii) min/mean/range
  - Mean/median
    - i) singles
    - ii) populations of singles (e.g., global mean transfer delay in clause 8.2.1)
    - iii) min/mean/range
  - Range
    - i) singles
    - ii) population of singles
    - iii) min/mean/range (e.g., mean transfer delay range over group in clause 8.2.3)

II.3 Relationship to other standards

MEF defines two types of multipoint services, a tree and a LAN. Neither of these is the same as the point-to-multipoint connection discussed here. However, for comparison purposes and with simplifications, the MEF multipoint delay/delay variation/loss ratio performance can be handled as a maximum of the vector point-to-point performance parameters (at each destination), and the availability performance as the minimum of the vector. Then the performance objectives are defined for these maxima/minima in [MEF 10.1].

Appendix III

Harmonizing terminology for frame delay and frame delay variation metrics

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

III.1 Introduction

The terminology for frame delay and frame delay variation used in this Recommendation differs somewhat from the terminology used presently by the Metro Ethernet Forum (MEF). For example, this Recommendation specifies a single primary/normative definition for frame delay variation, 2-point frame delay variation, which assesses the variation in a set of one-way delays with respect to a single reference delay (using the minimum delay is strongly recommended as the reference delay). This appendix provides a mapping between similar metrics developed by these two standards bodies. It also discusses other forms of delay variation metrics.

III.2 Frame delay and frame delay variation

In their technical specification [MEF 10.1], the MEF defined its service performance attributes (or metrics) in section 6.9. There are different performance attributes for point-to-point Ethernet virtual connections (EVCs) and multipoint-to-multipoint EVCs. Point-to-point EVCs have exactly two user-to-network interfaces (UNIs), supporting a bidirectional frame transfer service between the designated UNI pair.

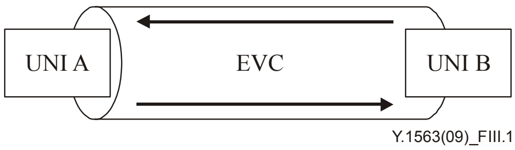

Figure III.1 – Point-to-point Ethernet virtual connections (EVCs)

In most ways, the definition of point-to-point EVC frame delay in [MEF 10.1] is similar to the definition in this Recommendation. One-way delay uses the first bit on ingress and the last bit on egress. However, the *population of interest* in [MEF 10.1] includes successfully transferred frames in both directions on the EVC. This is unusual, and a follow-on effort is in progress to separate the populations by direction (e.g., the population of interest would include frames transferred on the <A, B> direction). This follow-on work will be captured in an incremental update to [MEF 10.1], provisionally identified as [b-MEF 10.1.1].

A service level specification (SLS) for a point-to-point EVC frame delay would be based on a specified percentile of the one-way delay of subscriber traffic submitted during an evaluation time interval, T (the traffic must meet several other qualifications, see [MEF 10.1]). The MEF follow-on work will allow for the calculation of multiple percentiles of the delay distribution. Also, work on [b-MEF 10.1.1] proposes to add a new SLS objective on the difference between two percentiles of the one-way delay of subscriber traffic. This is the foundation for a direct comparison between the MEF performance attributes and the 2-point frame delay variation parameter specified in this Recommendation.

When the objective for the difference between percentiles (or "one-way frame delay range") is based on the same values used in the 2-point frame delay variation parameter (e.g., the 99.9 percentile and the minimum), then with the other modifications described above, the MEF and ITU-T will be able to specify equivalent metrics.

The MEF follow-on effort also intends to add a mean one-way frame delay objective, consistent with widespread usage in SLS and this Recommendation.

III.3 Inter-frame delay variation

[b-MEF 10.1.1] clarified that the form of delay variation specified in [MEF 10.1] is more accurately referred to as inter-frame delay variation (IFDV). This attribute is loosely based on [IETF RFC 3393], where the variation is assessed between packets sent with constant time spacing. Thus, the reference delay for each value of delay variation is the delay of a packet following the current packet by a constant Δt. Discussions in the follow-on activity have revealed that network characterization requires use of several values of Δt (on the order of milliseconds, seconds and tens of seconds).

IFDV is recognized as another form of delay variation assessment in this appendix and in Annex A. However, the 2-point frame delay variation specified in body of this Recommendation is preferred.

For one reason, there are tested and verified methods for combining segment measurements to estimate the UNI-to-UNI performance based on the 2-point FDV form (see clause 8 of [b-ITU-T Y.1541] for one such method).

The ITU-T Y.1541 IP performance objectives for packet delay variation (PDV) are specified in terms of the normative 2-point packet delay variation parameter, which is identical to the 2-point FDV in this Recommendation. This is because objectives using the 2-point PDV have been more straightforward to rationalize (or to relate to real-world networking aspects) than IFDV and similar forms.

It is worthwhile noting that the IETF IPPM Working Group has completed [b-IETF RFC 5481] comparing the two delay variation metrics in additional dimensions and to provide detailed guidance. [b-IETF RFC 5481] analyses the two forms in order to determine the tasks and circumstances that are most appropriate for each form.

III.4 Concatenation of inter-frame delay variation

The proposed method is an extension of the approach for combining 2-point delay variation measurements that appear in Appendix IV of [b-ITU-T Y.1541].

Both these methods employ approximations that produce computationally simple procedures while retaining an accuracy that should be sufficient for many planning purposes.

Suppose Ethernet frames flow through n network segments, which could be routers or networks themselves. Let T_k denote the 1-point frame delay variation through the k-th segment. The following procedure can be used to estimate the corresponding 1-point frame delay variation T of the flow measured upon output from the n-th (last) segment. Suppose we are given a set of measurements $m_1, m_1, \cdots, m_r$ of $|T_k|$ for each $k = 1, 2, \cdots, n$ and a fixed probability p. First define x to be the value satisfying $\Phi(x) = \frac{p+1}{2}$ where Φ is the standard normal (mean 0, variance 1) distribution function. Then:

- For each k compute the p-th quantile t_k satisfying $\Pr(|T_k| < t_k) = p$.

- For each k compute $\kappa_k = \frac{1}{r}\sum_{i=1}^{r} m_i^4 - 3 \cdot \sigma_k^4$ where $\sigma_k^2 = \frac{1}{r}\sum_{i=1}^{r} m_i^2$.

  NOTE – κ_k and σ_k^2 are in fact the 4th cumulant and variance of T_k . Alternatively, the cumulants can be estimated from the expression:

$$\kappa_k = 24 \cdot \frac{\sigma_k^4}{x} \cdot \frac{t_k / \sigma_k - x}{x^2 - 3}.$$

- Compute the end-to-end variance σ^2 and 4th cumulant of $T = \sum_{k=1}^{n} T_k$ as:

$$\sigma^2 = \sum_{k=1}^{n} \sigma_k^2 \qquad \kappa = \sum_{k=1}^{n} \kappa_k$$

- The p-th quantile t satisfying $\Pr(|T| < t) = p$ is finally computed as:

$$t = \sigma \cdot x \left\{ 1 + \frac{\kappa}{24 \cdot \sigma^4} (x^2 - 3) \right\}.$$

Appendix IV

Networking technologies making use of Ethernet

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

Different networking opportunities can make use of Ethernet. This appendix, which requires further study, will summarize the raw definitions of VLAN, PB (provider bridging) and PBB (provider backbone bridging). These technologies correspond to different encapsulations of Ethernet frames as shown in Figure IV.1.

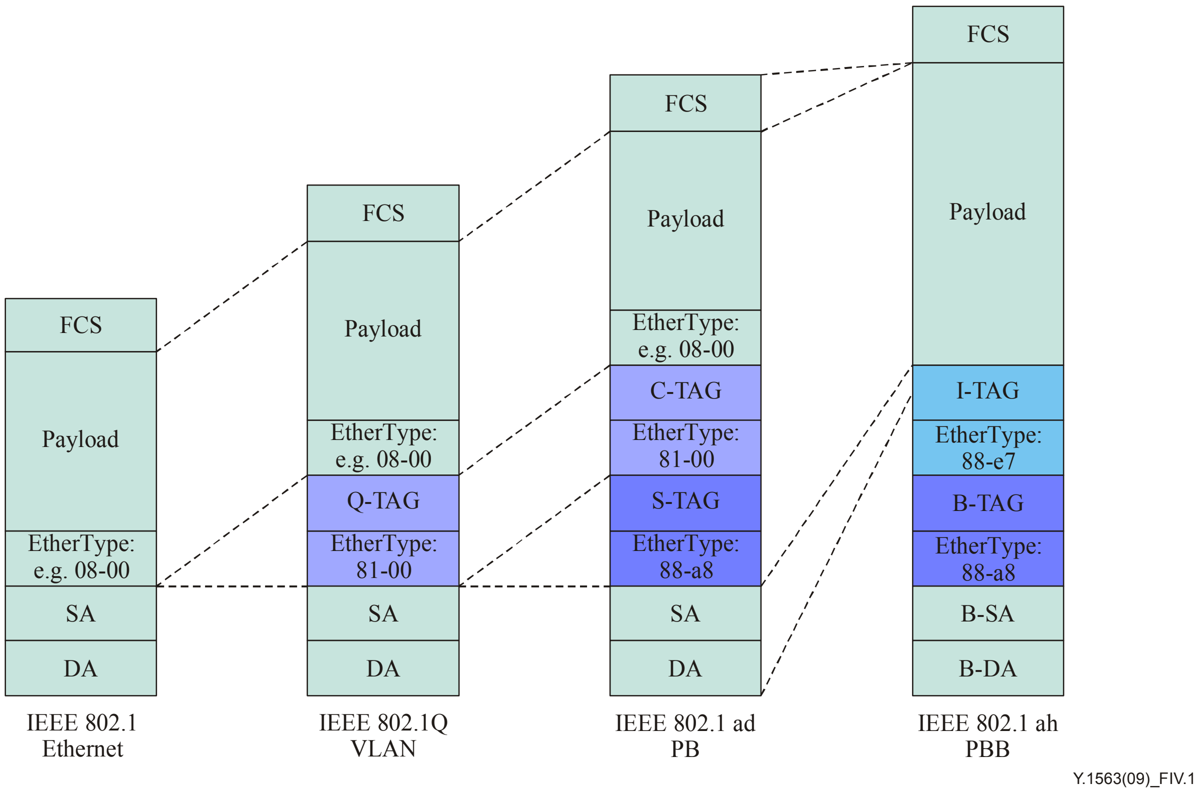

Figure IV.1 – Different Ethernet technologies

Appendix V

OAM measurement considerations

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

V.1 Introduction

Operation, administration and management (OAM) systems can introduce various types of maintenance frames into the subscriber traffic. This appendix examines the differences between OAM methods of measurement and the corresponding parameters defined in this Recommendation.

One type of OAM frame conveys frame counts at a particular MP, an ingress UNI MP, for example. Assuming that the MP counter only counts frames intended for some egress UNI MP, then the ingress count can be compared with the egress count to determine whether any frames were lost.

This Recommendation defines loss by tracking the transfer of *corresponding frames* between MPs, and uses a waiting time threshold to distinguish between loss and unlikely long delay. In practice, all frames have a useful lifetime. Thus, the waiting time, T_{max}, can be set fairly long (on the order of seconds). If the apparent delay of a frame exceeds T_{max}, then the frame is declared lost.

The OAM counter system does not use a rigorous time-out to determine frame loss. This appendix examines the ramifications of the OAM approach.

V.2 Loss performance based on counters

When OAM counter packets are introduced into a stream on a periodic basis, the delay variation influences the ingress frame counts. The ingress count can be quite accurate. However, the following factors influence the accuracy of the egress counts:

1) Frames may be counted at one MP which did not pass the other MP.
2) Assuming that frame order cannot be changed by the network, the arrival of an OAM frame indicates that no more frames applying to the count in that OAM packet can possibly arrive.
3) Where the order of frame delivery can change between MPs, then the frame counts carried in OAM frames may not be valid. This is especially true when the OAM frame itself is delivered out-of-order, or when a frame re-ordering event takes place in the vicinity of an OAM frame.

The interval between OAM counter frames, T_{OAM}, plus the one-way delay of the OAM frame, determines the maximum waiting time for frame arrival at a given counter location. T_{OAM} may not be equal to T_{max}, and thus the waiting time applied in an OAM measurement may differ from the specified T_{max}.

Appendix VI

Mapping to Recommendation ITU-T G.8010 terms and conventions (and to those of other associated Recommendations)

VI.1 Introduction

The body of this Recommendation uses the terminology and architectural conventions for connectionless networks that were developed beginning with [ITU-T I.380] ([ITU-T I.380] was renumbered [ITU-T Y.1540] a year later), for IP-based network service performance. This terminology and architecture has been adapted to the specific case of Ethernet connectionless networks, as it was for [b-ITU-T Y.1561] (MPLS) in 2004.

A telecommunication network is a complex network which can be described in a number of different ways, depending on the purpose of the description, and therefore other Recommendations may use alternate terminology and conventions. For example, [ITU-T G.8010] relies on other Recommendations developed to serve the needs of Study Group 15. This appendix provides a mapping between key terms of [ITU-T G.8010] and this Recommendation. Annexes A and B of [b-ITU-T G.8011] also provide a mapping between Ethernet service attributes/interfaces and the architecture of [ITU-T G.8010].

VI.2 Mapping between Ethernet network performance and architecture Recommendations

Table VI.1 – Mapping between sets of terminology

ITU-T Y.1563	ITU-T G.8010 ETH
Source (SRC)	Flow termination source.
Destination (DST)	Flow termination sink.
Measurement point (MP)	Possibly a maintenance entity associated with an access point, flow point, or a termination flow point.
Point-to-point topology (2 MP)	Possibly Ethernet access groups with a point-to-point relationship association, and the Ethernet flow domains and Ethernet flow point pool links between them.
Point-to-multipoint topology (n > 2 MP)	Possibly Ethernet access groups with a point-to-multipoint relationship association, and the Ethernet flow domains and Ethernet flow point pool links between them (also known as a rooted multipoint Ethernet VLAN).
Bridge*	Special case of a flow domain.
Exchange link	Flow point pool link, link flow domain fragment, or link connection when used to connect a user's flow domain to a network provider's flow domain (access link), or when used to connect between provider's flow domains.
Network section	Possibly a flow domain, (sub)network flow domain fragment which is under a single jurisdictional boundary.
Network section ensemble	One or more flow domains or (sub)network flow domain fragments (connected by flow point pool links), or an entire layer network.
Virtual LAN (VLAN)*	Possibly a flow point pool component link and corresponding flow domain fragment(s).

Table VI.1 – Mapping between sets of terminology

ITU-T Y.1563	ITU-T G.8010 ETH
End-to-end Ethernet network	Possibly a layer network, (sub)network flow domain fragment, or connectionless trail.
Population of interest	Possibly a flow domain flow, or a flow domain fragment.
* New term in this Recommendation, not used in earlier Recommendations on performance.	

Measurement points (MPs) may be assigned in the architecture wherever needed (and practical). For example, measurement points may be located at the ETH access point between the ETH/client adaptation function and the ETH flow termination function, or between the ETH/ETH adaptation function and the ETH flow termination function. ETH service performance may be monitored between MP at ETH layer access points (NC-MEG) or ETH sublayer access points (TC-MEG, LC-MEG). MP at ETH (sub)layer access points may be located in UNI-C ports, UNI-N ports and E-NNI ports. Note that one port may contain more than one measurement point for an ETH flow domain fragment.

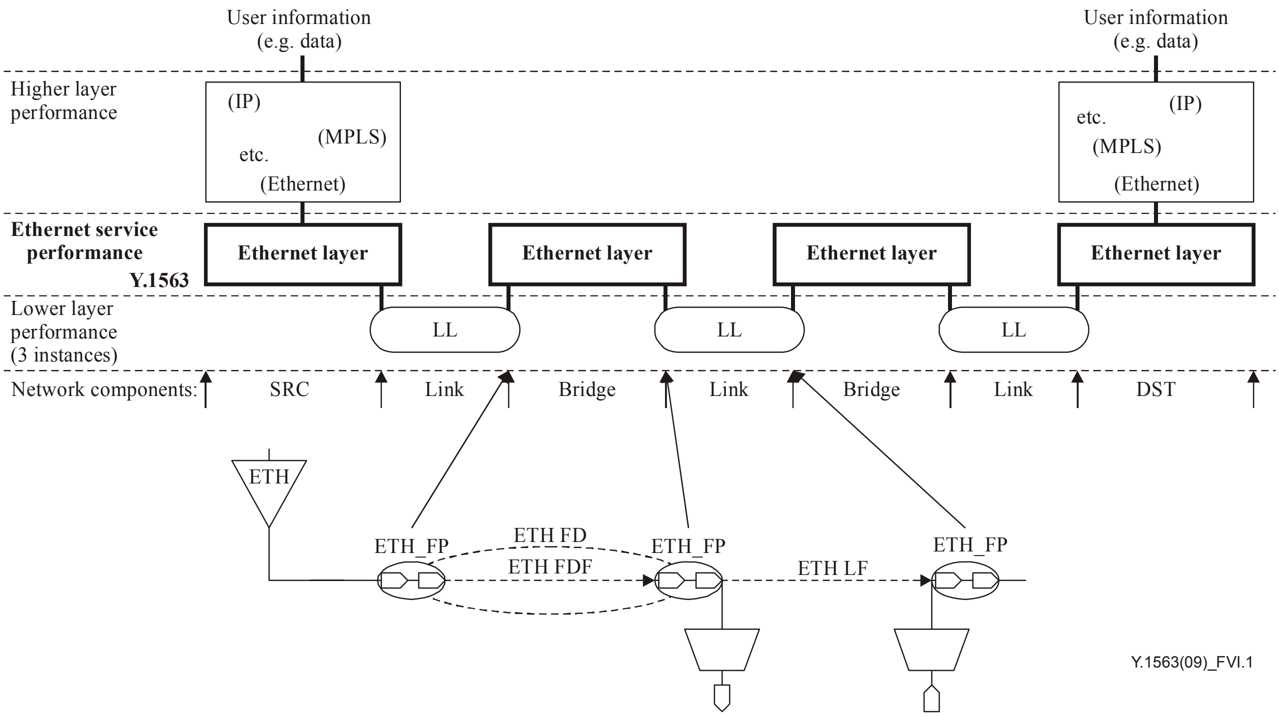

Figure VI.1 – Example mapping between Figure 2 of this Recommendation and Figure 4 of [ITU-T G.8010]

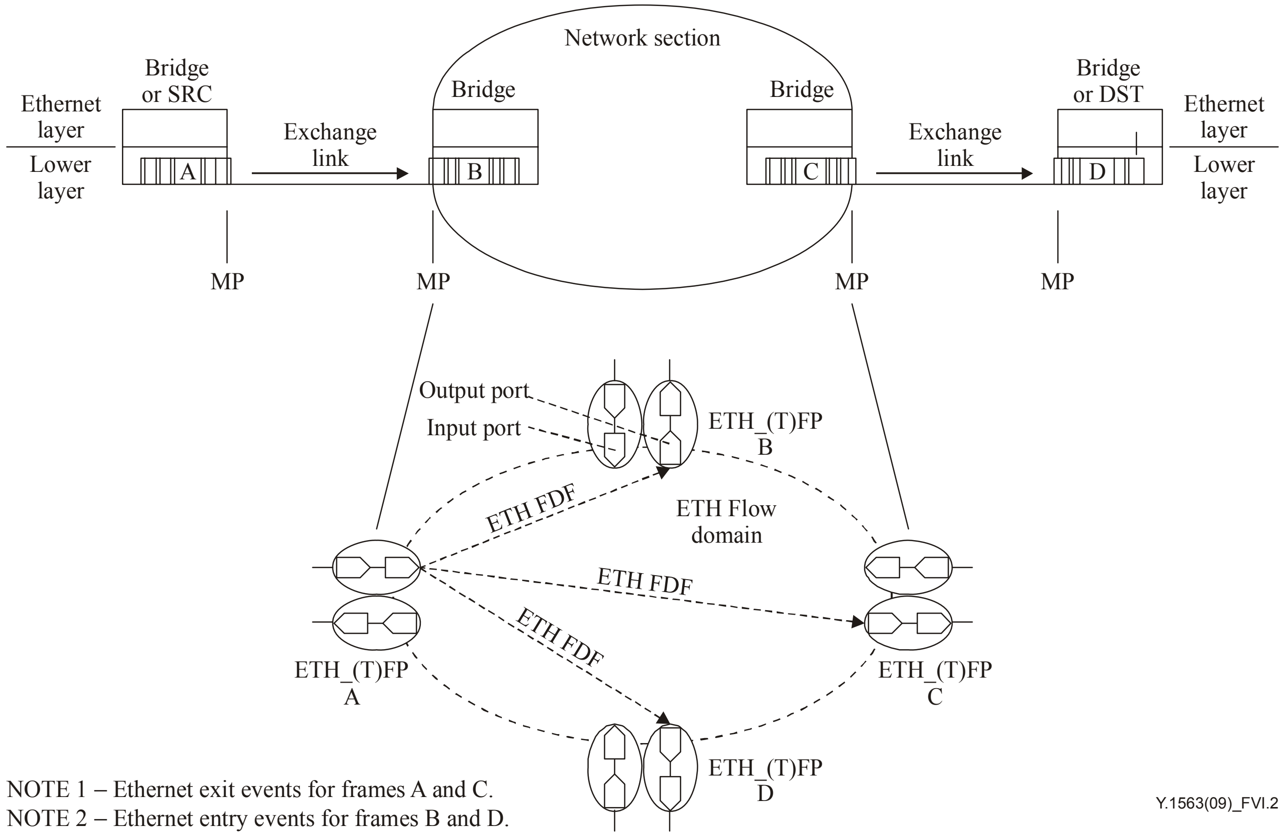

Figure VI.2 – Example mapping between Figure 3 of this Recommendation and Figure 6 of [ITU-T G.8010]

VI.3 In-service measurement and applicability of performance and availability parameters

Frame transfer delay (clause 8.1) and frame transfer delay variation (clause 8.1.5) are applicable between ingress and egress points of a maintenance entity group (MEG) level within an ETH flow domain fragment.

For the in-service monitoring ETH-CCM with its CC, LM and RDI, functions may be used to monitor the performance of an ETH flow domain fragment. An ETH flow domain fragment may have three types of maintenance entity groups (MEGs); one network connection MEG, up to six tandem connection MEGs and up to one link connection MEG. The performance of each MEG level in an ETH flow domain fragment is monitored in-service by means of an MEG level-specific ETH-CCM. The presence of lost or spurious frames is determined by counting the transmitted frames with a specific priority in the ETH flow termination source function, transmitting this value in the CCM frame, counting the received frames with this priority in the ETH flow termination sink function, extracting the transmitted frame count values from the incoming CCM frames and processing this information. Note that the CCM bit rate (100 ms period) is 7.2 kbit/s. [b-ITU-T Y.1731], [b-ITU-T G.8021], [b-ITU-T G.8051] and [b-ITU-T G.7710] should be consulted for additional specifications on in-service monitoring.

The ETH-CC function (supported in the ETH-CCM OAM) is one in-service monitoring function which determines if there is continuity and connectivity. Loss of continuity and/or incorrect connectivity will result in LOC and MMG defects. Those defects contribute to the defect second performance parameter and are the first contribution to the error performance primitives of the ETH flow domain fragment (indicate 100% frame loss). In addition, ETH-LM function (supported in the ETH-CCM OAM) is available for 2-port ETH flow domain fragments to determine the frame loss during periods without loss of continuity or wrong connectivity conditions. ETH-CCM is very important, but it is only one option for performance monitoring (and not ideal for some aspects of the task).

Bibliography

[b-ITU-T G.707] Recommendation ITU-T G.707/Y.1322 (2003), *Network node interface for the synchronous digital hierarchy (SDH).*

[b-ITU-T G.709] Recommendation ITU-T G.709/Y.1331 (2001), *Interfaces for the Optical Transport Network (OTN).*

[b-ITU-T G.827] Recommendation ITU-T G.827 (2003), *Availability performance parameters and objectives for end-to-end international constant bit-rate digital paths*.

[b-ITU-T G.7041] Recommendation ITU-T G.7041/Y.1303 (2001), *Generic framing procedure (GFP).*

[b-ITU-T G.7042] Recommendation ITU-T G.7042/Y.1303 (2001), *Link capacity adjustment scheme (LCAS) for virtual concatenated signals.*

[b-ITU-T G.7710] Recommendation ITU-T G.7710/Y.1701 (2007), *Common equipment management function requirements*.

[b-ITU-T G.8011] Recommendation ITU-T G.8011/Y.1307 (2004), *Ethernet over Transport – Ethernet services framework*.

[b-ITU-T G.8011.1] Recommendation ITU-T G.8011.1/Y.1307.1 (in force), *Ethernet private line service*.

[b-ITU-T G.8011.2] Recommendation ITU-T G.8011.2/Y.1307.2 (2005), *Ethernet virtual private line service*.

[b-ITU-T G.8012] Recommendation ITU-T G.8012/Y.1308 (2004), *Ethernet UNI and Ethernet NNI*.

[b-ITU-T G.8021] Recommendation ITU-T G.8021/Y.1341 (2007), *Characteristics of Ethernet transport network equipment functional blocks*.

[b-ITU-T G.8051] Recommendation ITU-T G.8051/Y.1345 (2007), *Management aspects of the Ethernet-over-Transport (EoT) capable network element*.

[b-ITU-T I.350] Recommendation I.350 (1993), *General aspects of quality of service and network performance in digital networks, including ISDNs*.

[b-ITU-T I.380] Recommendation ITU-T I.380 (1999), *Internet protocol data communication service – IP packet transfer and availability performance parameters*.

[b-ITU-T Y.1540] Recommendation ITU-T Y.1540 (2007), *Internet protocol data communication service – IP packet transfer and availability performance parameters*.

[b-ITU-T Y.1541] Recommendation ITU-T Y.1541 (2006), *Network performance objectives for IP-based services*.

[b-ITU-T Y.1561] Recommendation ITU-T Y.1561 (2004), *Performance and availability parameters for MPLS networks*.

[b-ITU-T Y.1730] Recommendation ITU-T Y.1730 (2004), *Requirements for OAM functions in Ethernet-based networks and Ethernet services*.

[b-ITU-T Y.1731] Recommendation ITU-T Y.1731 (2008), *OAM functions and mechanisms for Ethernet based networks*.

[b-IEEE 802] IEEE 802 (2001), *IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture*.

[b-IEEE 802.1D] IEEE 802.1D (2004), *IEEE Standard for Local and metropolitan area networks: Media Access Control (MAC) Bridge*s.

[b-IEEE 802.1Q] IEEE 802.1Q (2005), *IEEE standard for Local and metropolitan area networks: Virtual bridged local area network*s.

[b-IETF RFC 5481] IETF RFC 5481 (2009), *Packet Delay Variation Applicability Statement*.

[b-MEF 10.1.1] MEF 10.1.1 (2007), *Amendment to Ethernet Services Attributes Phase 2*.

[b-MEF 6] MEF 6 (2004), *Metro Ethernet Services Definitions Phase I.*

SERIES OF ITU-T RECOMMENDATIONS

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

Printed in Switzerland
Geneva, 2009