

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

G.8013/Y.1731

Corrigendum 2
(08/2019)

SERIES G: TRANSMISSION SYSTEMS AND MEDIA,
DIGITAL SYSTEMS AND NETWORKS

Packet over Transport aspects – Ethernet over Transport
aspects

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

Internet protocol aspects – Operation, administration and
maintenance

Operation, administration and maintenance (OAM)
functions and mechanisms for Ethernet-based
networks

Corrigendum 2

Recommendation ITU-T G.8013/Y.1731 (2015) –
Corrigendum 2

ITU-T G-SERIES RECOMMENDATIONS

TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

INTERNATIONAL TELEPHONE CONNECTIONS AND CIRCUITS	G.100–G.199
GENERAL CHARACTERISTICS COMMON TO ALL ANALOGUE CARRIER-TRANSMISSION SYSTEMS	G.200–G.299
INDIVIDUAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON METALLIC LINES	G.300–G.399
GENERAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON RADIO-RELAY OR SATELLITE LINKS AND INTERCONNECTION WITH METALLIC LINES	G.400–G.449
COORDINATION OF RADIOTELEPHONY AND LINE TELEPHONY	G.450–G.499
TRANSMISSION MEDIA AND OPTICAL SYSTEMS CHARACTERISTICS	G.600–G.699
DIGITAL TERMINAL EQUIPMENTS	G.700–G.799
DIGITAL NETWORKS	G.800–G.899
DIGITAL SECTIONS AND DIGITAL LINE SYSTEM	G.900–G.999
MULTIMEDIA QUALITY OF SERVICE AND PERFORMANCE – GENERIC AND USER-RELATED ASPECTS	G.1000–G.1999
TRANSMISSION MEDIA CHARACTERISTICS	G.6000–G.6999
DATA OVER TRANSPORT – GENERIC ASPECTS	G.7000–G.7999
PACKET OVER TRANSPORT ASPECTS	G.8000–G.8999
Ethernet over Transport aspects	G.8000–G.8099
MPLS over Transport aspects	G.8100–G.8199
Synchronization, quality and availability targets	G.8200–G.8299
Service Management	G.8600–G.8699
ACCESS NETWORKS	G.9000–G.9999

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

Recommendation ITU-T G.8013/Y.1731

Operation, administration and maintenance (OAM) functions and mechanisms for Ethernet-based networks

Corrigendum 2

Summary

Recommendation ITU-T G.8013/Y.1731 provides mechanisms for user-plane OAM functionality in Ethernet networks according to the requirements and principles given in Recommendation ITU-T Y.1730. This Recommendation is designed specifically to support point-to-point connections and multipoint connectivity in the ETH layer as identified in Recommendation ITU-T G.8010/Y.1306.

The OAM mechanisms defined in this Recommendation offer capabilities to operate and maintain network and service aspects of the ETH layer.

Corrigendum 1 provides:

- Editorial modifications in clauses 7.3.2, 9.3.2, 9.4.2, 9.6.2, 9.14.2, 9.15.2, 9.16.2, 9.22.2, 9.23.2 and 9.24.2;
- A correction in Table V.1;
- Updates of the publication date for some references.

Amendment 1 provides:

- Modification to "Frame delay" in clause 8.
- Updates to clauses 2 and 3, and to the Bibliography.

Corrigendum 2 provides:

- An editorial modification to clause 7.13.1.

History

Edition	Recommendation	Approval	Study Group	Unique ID*
1.0	ITU-T Y.1731	2006-05-22	13	11.1002/1000/7192
2.0	ITU-T Y.1731	2008-02-29	13	11.1002/1000/9347
2.1	ITU-T Y.1731 (2008) Amd. 1	2010-07-29	15	11.1002/1000/10925
3.0	ITU-T G.8013/Y.1731	2011-07-22	15	11.1002/1000/11136
3.1	ITU-T G.8013/Y.1731 (2011) Cor. 1	2011-10-29	15	11.1002/1000/11418
3.2	ITU-T G.8013/Y.1731 (2011) Amd. 1	2012-05-07	15	11.1002/1000/11511
4.0	ITU-T G.8013/Y.1731	2013-11-06	15	11.1002/1000/12029
4.1	ITU-T G.8013/Y.1731 (2013) Amd. 1	2015-02-22	15	11.1002/1000/12381
5.0	ITU-T G.8013/Y.1731	2015-08-13	15	11.1002/1000/12552
5.1	ITU-T G.8013/Y.1731 (2015) Cor. 1	2018-03-16	15	11.1002/1000/13542
5.2	ITU-T G.8013/Y.1731 (2015) Amd. 1	2018-11-29	15	11.1002/1000/13757
5.3	ITU-T G.8013/Y.1731 (2015) Cor. 2	2019-08-29	15	11.1002/1000/14007

Keywords

Ethernet OAM, fault management, IEEE 802.1Q, OAM frame format, OAM mechanism, performance monitoring.

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

© ITU 2019

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

Table of Contents

	Page
Corrigendum 2	1
1 Scope.....	1
2 References.....	1
3 Definitions	3
3.1 Terms defined elsewhere	3
3.2 Terms defined in this Recommendation.....	4
4 Abbreviations and acronyms	5
5 Conventions	7
5.1 ME group (MEG)	7
5.2 Traffic conditioning point (TrCP)	8
5.3 MEG level	8
5.4 OAM transparency	8
5.5 Representation of octets	9
6 OAM relationships.....	9
6.1 MEs, MEPs, MIPs and TrCPs relationship	9
6.2 MEs, MEGs and MEG level relationship.....	9
6.3 MEPs and MIPs configuration	11
7 OAM functions for fault management.....	11
7.1 Ethernet continuity check (ETH-CC).....	11
7.2 Ethernet loopback (ETH-LB)	13
7.3 Ethernet link trace (ETH-LT)	16
7.4 Ethernet alarm indication signal (ETH-AIS).....	18
7.5 Ethernet remote defect indication (ETH-RDI)	20
7.6 Ethernet locked signal (ETH-LCK).....	21
7.7 Ethernet test signal (ETH-Test).....	22
7.8 Ethernet automatic protection switching (ETH-APS).....	23
7.9 Ethernet maintenance communication channel (ETH-MCC)	23
7.10 Ethernet experimental OAM (ETH-EXP)	24
7.11 Ethernet vendor-specific OAM (ETH-VSP)	24
7.12 Ethernet client signal fail (ETH-CSF)	24
7.13 Ethernet bandwidth notification (ETH-BN)	26
7.14 Ethernet expected defect function (ETH-ED)	28
8 OAM functions for performance monitoring	29
8.1 Frame loss measurement (ETH-LM).....	30
8.2 Frame delay measurement (ETH-DM).....	33
8.3 Throughput measurement.....	36

	Page
8.4 Synthetic loss measurement (ETH-SLM)	36
9 OAM PDU types.....	39
9.1 Common OAM information elements	39
9.2 CCM PDU	42
9.3 LBM PDU	44
9.4 LBR PDU	46
9.5 LTM PDU	47
9.6 LTR PDU.....	49
9.7 AIS PDU	51
9.8 LCK frame.....	52
9.9 TST PDU	53
9.10 APS PDU.....	53
9.11 MCC PDU	54
9.12 LMM PDU.....	55
9.13 LMR PDU	56
9.14 IDM PDU.....	57
9.15 DMM PDU	58
9.16 DMR PDU	59
9.17 EXM PDU	60
9.18 EXR PDU	61
9.19 VSM PDU	62
9.20 VSR PDU	63
9.21 Client signal fail (CSF).....	64
9.22 SLM PDU	65
9.23 SLR PDU.....	66
9.24 1SL PDU	67
9.25 BNM PDU	68
9.26 EDM PDU	69
10 OAM frame addresses	70
10.1 Multicast destination addresses	70
10.2 CCM	71
10.3 LBM	71
10.4 LBR	71
10.5 LTM.....	71
10.6 LTR.....	71
10.7 AIS.....	71
10.8 LCK	71
10.9 TST	72
10.10 APS.....	72

	Page
10.11 MCC	72
10.12 LMM.....	72
10.13 LMR	72
10.14 IDM.....	72
10.15 DMM	72
10.16 DMR.....	72
10.17 EXM	72
10.18 EXR	72
10.19 VSM	72
10.20 VSR	72
10.21 CSF.....	72
10.22 SLM.....	73
10.23 SLR.....	73
10.24 1SL	73
10.25 BNM.....	73
10.26 EDM	73
11 OAM PDU validation and versioning	74
11.1 OAM PDU transmission.....	74
11.2 OAM PDU validation in reception.....	74
11.3 OAM PDU reception after validation.....	75
Annex A – MEG ID format	76
A.1 ICC based MEG_ID format.....	77
A.2 Global MEG ID format based on CC and ICC.....	78
Annex B – Ethernet link trace (ETH-LT) of ITU-T Y.1731 interoperability considerations..	79
B.1 Ethernet link trace (ETH-LT) as defined in ITU-T Y.1731	79
B.2 Interworking with ITU-T Y.1731	79
Appendix I – Ethernet network scenarios	81
I.1 Shared MEG levels example.....	81
I.2 Independent MEG levels example	82
Appendix II – Frame loss measurement	83
II.1 Simplified calculation for frame loss	84
II.2 Frame counter wrapping periodicity	85
Appendix III – Network OAM interworking.....	86
Appendix IV – Mismatch detection limitation	87
Appendix V – Terminology alignment with IEEE 802.1Q	88
Appendix VI – Examples showing accuracy for ETH-SLM measurement.....	89
Appendix VII – ETH-LM and link aggregation	90
Bibliography.....	92

Introduction

ITU-T has prepared Recommendation ITU-T G.8013/Y.1731 in cooperation with the IEEE Project 802.1ag (Connectivity fault management). Since the IEEE work is now complete, this Recommendation contains amendments to fully align the final results and includes the appropriate normative references to IEEE documents. Moreover, further detailed work on the implementation details (i.e., the specification of the equipment functions) has been undertaken by ITU-T.

Recommendation ITU-T G.8013/Y.1731

Operation, administration and maintenance (OAM) functions and mechanisms for Ethernet-based networks

Corrigendum 2

Editorial note – This is a complete-text publication. Modifications introduced by this corrigendum are shown in revision marks relative to Recommendation ITU-T G.8013/Y.1731 (2015) plus Corrigendum 1 (2018) and Amendment 1 (2018).

1 Scope

This Recommendation specifies mechanisms required to operate and maintain the network and service aspects of the ETH layer. It also specifies the Ethernet OAM frame formats and syntax and semantics of OAM frame fields. The OAM mechanisms as described in this Recommendation apply to both point-to-point ETH connections and multipoint ETH connectivity including both multipoint-to-multipoint and rooted-multipoint connections. The OAM mechanisms as described in this Recommendation are applicable to any environment independently of how the ETH layer is managed (e.g., using network management systems or operational support systems).

The architectural basis for this Recommendation is the Ethernet specification [ITU-T G.8010] which also accounts for [IEEE 802.1Q] and [IEEE 802.3]. The OAM functions of the server layer networks used by the Ethernet network are not within the scope of this Recommendation. The OAM functions of the layers above the ETH layer are not within the scope of this Recommendation either.

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.806] Recommendation ITU-T G.806 (2012), *Characteristics of transport equipment – Description methodology and generic functionality*.
- [ITU-T G.809] Recommendation ITU-T G.809 (2003), *Functional architecture of connectionless layer networks*.
- [ITU-T G.826] Recommendation ITU-T G.826 (2002), *End-to-end error performance parameters and objectives for international, constant bit-rate digital paths and connections*.
- [ITU-T G.7710] Recommendation ITU-T G.7710/Y.1701 (2012), *Common equipment management function requirements*.
- [ITU-T G.8010] Recommendation ITU-T G.8010/Y.1306 (2004), *Architecture of Ethernet layer networks*.

- [ITU-T G.8021] Recommendation ITU-T G.8021/Y.1341 (2018), *Characteristics of Ethernet transport network equipment functional blocks.*
- [ITU-T G.8031] Recommendation ITU-T G.8031/Y.1342 (2015), *Ethernet linear protection switching.*
- [ITU-T G.8032] Recommendation ITU-T G.8032/Y.1344 (2015), *Ethernet ring protection switching.*
- [ITU-T G.8113.1] Recommendation ITU-T G.8113.1/Y.1372.1 (2016), *Operations, administration and maintenance mechanism for MPLS-TP in packet transport networks.*
- [ITU-T M.1400] Recommendation ITU-T M.1400 (2013), *Designations for interconnections among operators' networks.*
- [ITU-T O.150] Recommendation ITU-T O.150 (1996), *General requirements for instrumentation for performance measurements on digital transmission equipment.*
- [ITU-T T.50] Recommendation ITU-T T.50 (1992), *International Reference Alphabet (IRA) (Formerly International Alphabet No. 5 or IA5) – Information technology – 7-bit coded character set for information interchange.*
- [ITU-T Y.1563] Recommendation ITU-T Y.1563 (2009), *Ethernet frame transfer and availability performance.*
- [ITU-T Y.1564] Recommendation ITU-T Y.1564 (2011), *Ethernet service activation test methodology.*
- [ITU-T Y.1730] Recommendation ITU-T Y.1730 (2004), *Requirements for OAM functions in Ethernet-based networks and Ethernet services.*
- [ITU-T Y.1731] Recommendation ITU-T Y.1731 (2006), *OAM functions and mechanisms for Ethernet based networks.*
- [IEC 61588] IEC 61588 (2009), *Precision clock synchronization protocol for networked measurement and control systems.*
<<https://webstore.iec.ch/publication/5639>>
- [IEEE 1588] IEEE 1588-2002, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems.*
<<http://standards.ieee.org/findstds/standard/1588-2002.html>>
- [IEEE 802] IEEE 802-2014, *IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture.*
<<http://standards.ieee.org/findstds/standard/802-2014.html>>
- [IEEE 802.1Q] IEEE 802.1Q-2018, *IEEE Standard for Local and metropolitan area networks – Bridges and Bridged Networks.*
- [IEEE 802.3] IEEE 802.3-2018, *IEEE Standard for Ethernet.*
<<http://standards.ieee.org/findstds/standard/802.3-2018.html>>
- [ISO 3166-1] ISO 3166-1 (2013), *Codes for the representation of names of countries and their subdivisions – Part 1: Country codes.*
- [MEF 10.3] MEF 10.3 (2013), *Ethernet Services Attributes Phase 3.*
<https://www.mef.net/Assets/Technical_Specifications/PDF/MEF_10.3.pdf>

3 Definitions

3.1 Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

- 3.1.1 **adaptation:** [ITU-T G.809].
- 3.1.2 **adapted information:** [ITU-T G.809].
- 3.1.3 **client/server relationship:** [ITU-T G.809].
- 3.1.4 **connection point:** [ITU-T G.805].
- 3.1.5 **connectionless trail:** [ITU-T G.809].
- 3.1.6 **defect:** [ITU-T G.806].
- 3.1.7 **failure:** [ITU-T G.806].
- 3.1.8 **flow:** [ITU-T G.809].
- 3.1.9 **flow domain:** [ITU-T G.809].
- 3.1.10 **flow domain flow:** [ITU-T G.809].
- 3.1.11 **flow point:** [ITU-T G.809].
- 3.1.12 **flow point pool:** [ITU-T G.809].
- 3.1.13 **flow point pool link:** [ITU-T G.809].
- 3.1.14 **flow termination:** [ITU-T G.809].
- 3.1.15 **flow termination sink:** [ITU-T G.809].
- 3.1.16 **flow termination source:** [ITU-T G.809].
- 3.1.17 **layer network:** [ITU-T G.809].
- 3.1.18 **link:** [ITU-T G.805].
- 3.1.19 **link connection:** [ITU-T G.805].
- 3.1.20 **link flow:** [ITU-T G.809].
- 3.1.21 **network:** [ITU-T G.809].
- 3.1.22 **network connection:** [ITU-T G.805].
- 3.1.23 **organizationally unique identifier:** [IEEE 802].
- 3.1.24 **port:** [ITU-T G.809].
- 3.1.25 **reference point:** [ITU-T G.809].
- 3.1.26 **termination connection point:** [ITU-T G.805].
- 3.1.27 **termination flow point:** [ITU-T G.809].
- 3.1.28 **traffic unit:** [ITU-T G.809].
- 3.1.29 **trail:** [ITU-T G.805].
- 3.1.30 **trail termination:** [ITU-T G.805].
- 3.1.31 **transport:** [ITU-T G.809].
- 3.1.32 **transport entity:** [ITU-T G.809].
- 3.1.33 **transport processing function:** [ITU-T G.809].

3.2 Terms defined in this Recommendation

This Recommendation defines the following terms:

3.2.1 dual-ended: A type of protocol messaging whereby an initiating MEP sends a request PDU to a receiving peer MEP, which does not send any response. In performance monitoring, the receiving MEP performs measurement calculations.

3.2.2 far-end: In single-ended or dual-ended messaging, the direction and information relating to a one-way measurement from an initiating MEP to a receiving or responding peer MEP.

3.2.3 initiating MEP: An initiating MEP initiates measurements by sending request PDUs, and in the case of single-ended messaging, receiving response PDUs.

3.2.4 in-profile: For frames that belong to the same instance of an Ethernet service and priority, in-profile frames are defined as frames for which priority corresponds to a common value <X> and the network has determined that drop-eligibility corresponds to <false>. Procedures to determine drop-eligibility for a frame and procedures to convey priority and drop-eligibility on a frame are network and service specific.

3.2.5 in-service OAM: OAM actions that are carried out while the data traffic is not interrupted, with an expectation that data traffic remains transparent to OAM actions.

3.2.6 maintenance entity: The entity between two of the flow/connection points in a maintenance entity group.

3.2.7 maintenance entity group (MEG): A group defined, for the purpose of fragment or connection monitoring, between a set of flow or connection points within a fragment/connection. This set of flow or connection points may be located at the boundary of one administrative domain or a protection domain, or at the boundaries of two adjacent administrative domains. The maintenance entity group consists of one or more maintenance entities.

3.2.8 MEG end point (MEP): A point marking the end of an ETH MEG that is capable of initiating and terminating OAM frames for fault management and performance monitoring. A MEP does not add a new forwarding identifier to the transit ETH flows. A MEP does not terminate the transit ETH flows, though it can observe these flows (e.g., count frames).

3.2.9 MEG intermediate point (MIP): An intermediate point in a MEG that is capable of reacting to some OAM frames. A MIP does not initiate OAM frames. A MIP takes no action on the transit ETH flows.

3.2.10 near-end: In single-ended messaging, the direction and information relating to a one-way measurement from a responding peer MEP to an initiating MEP.

3.2.11 on-demand OAM: OAM actions that are initiated via manual intervention for a limited time to carry out diagnostics. On-demand OAM can result in singular or periodic OAM actions during the diagnostics time interval.

3.2.12 one-way: A measurement of the performance of frames that is achieved in one direction from an initiating MEP to a peer MEP, or vice versa; i.e., a unidirectional measurement.

3.2.13 out-of-service OAM: OAM actions that are carried out while the data traffic is interrupted.

3.2.14 peer MEP: In a given MEP, all of the other MEPs in the same MEG.

3.2.15 proactive OAM: OAM actions that are carried out continuously to permit proactive reporting of fault or performance results.

3.2.16 responding MEP: In single-ended messaging, an MEP that receives request PDUs from an initiating MEP and transmits corresponding response PDUs.

3.2.17 receiving MEP: In dual-ended messaging, an MEP that receives request PDUs from an initiating MEP, and does not respond to them.

3.2.18 server MEP: A server MEP represents the compound function of the server layer termination function and server/ETH adaptation function, which is used to notify the ETH layer MEPs of conditions of the server layer that may be relevant to these ETH MEPs (e.g., failure or degradation detection by the server layer termination function or server/ETH adaptation function), where the server layer termination function is expected to run OAM mechanisms specific to that server layer.

3.2.19 single-ended: A type of protocol messaging whereby an initiating MEP sends a request PDU to a responding peer MEP, and the peer MEP responds by sending a response PDU that contains the original request data plus any additional data added by the responder. In performance monitoring, the initiating MEP performs measurement calculations.

3.2.20 two-way: A measurement of the performance of frames that is achieved from an initiating MEP to responding peer MEP, and back to the initiating MEP; i.e., a bidirectional or round-trip measurement.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

1DM	One-way Delay Measurement
1SL	One-way Synthetic Loss measurement
AIS	Alarm Indication Signal
APS	Automatic Protection Switching
BNM	Bandwidth Notification Message
CCM	Continuity Check Message
CoS	Class of Service
CP	Connection Point
CSF	Client Signal Fail
DA	Destination MAC Address
DEI	Drop Eligible Indicator
DMM	Delay Measurement Message
DMR	Delay Measurement Reply
EDM	Expected Defect Message
ETH	Ethernet MAC layer network
ETH-AIS	Ethernet Alarm Indication Signal function
ETH-APS	Ethernet Automatic Protection Switching function
ETH-BN	Ethernet Bandwidth Notification function
ETH-CC	Ethernet Continuity Check function
ETH-CSF	Ethernet Client Signal Fail function
ETH-DM	Ethernet Delay Measurement function
ETH-ED	Ethernet Expected Defect function
ETH-EXP	Ethernet Experimental OAM function

ETH-LB	Ethernet Loopback function
ETH-LCK	Ethernet Lock signal function
ETH-LM	Ethernet Loss Measurement function
ETH-LT	Ethernet Link Trace function
ETH-MCC	Ethernet Maintenance Communication Channel function
ETH-RDI	Ethernet Remote Defect Indication function
ETH-SLM	Ethernet Synthetic Loss Measurement function
ETH-Test	Ethernet Test function
ETH-TFP	Ethernet Termination Flow Point
ETH-VSP	Ethernet Vendor-Specific OAM function
ETY	Ethernet PHY layer network
EXM	Experimental OAM Message
EXR	Experimental OAM Reply
FLR	Frame Loss Ratio
FT	Flow Termination
GNM	Generic Notification Message
ICC	ITU Carrier Code
LBM	Loopback Message
LBR	Loopback Reply
LCK	Locked
LMM	Loss Measurement Message
LMR	Loss Measurement Reply
LOC	Loss Of Continuity
LTM	Link Trace Message
LTR	Link Trace Reply
MAC	Media Access Control
MCC	Maintenance Communication Channel
ME	Maintenance Entity
MEG	ME Group
MEL	MEG Level
MEP	MEG End Point
MIP	MEG Intermediate Point
NMS	Network Management System
NNI	Network Node Interface
NT	Network Termination
OAM	Operation, Administration and Maintenance
OSS	Operations Support System

OTN	Optical Transport Network
OUI	Organizationally Unique Identifier
PCP	Priority Code Point
PDU	Protocol Data Unit
PE	Provider Edge
PHY	Ethernet Physical layer entity consisting of the PCS, the PMA, and, if present, the PMD sublayers
PRBS	Pseudo-Random Bit Sequence
RDI	Remote Defect Indication
SA	Source MAC Address
SES	Severely Errored Seconds
SLA	Service Level Agreement
SLM	Synthetic Loss Message
SLR	Synthetic Loss Reply
SRV	Server
STP	Spanning Tree Protocol
TCI	Tag Control Information
TLV	Type, Length and Value
TrCP	Traffic Conditioning Point
TST	Test PDU
TTL	Time To Live
UMC	Unique MEG ID Code
UNI	User Network Interface
UNI-C	Customer side of UNI
UNI-N	Network side of UNI
VLAN	Virtual LAN
VSM	Vendor-Specific OAM Message
VSR	Vendor-Specific OAM Reply

5 Conventions

The diagrammatic conventions for connection-oriented and connectionless layer networks described in this Recommendation are those of [ITU-T G.805], [ITU-T G.809] and [ITU-T G.8010].

For the purposes of this Recommendation, the following OAM terms and diagrammatic conventions are also defined.

5.1 ME group (MEG)

An ME group (MEG) includes different MEs that satisfy the following conditions:

- All MEs in a MEG exist in one same administrative boundary.
- All MEs in a MEG have the same MEG level (see clause 5.3).

- All MEs in a MEG belong to the same point-to-point ETH connection or multipoint ETH connection.

For a point-to-point ETH connection, a MEG contains a single ME.

For a multipoint ETH connection containing n end-points, a MEG contains $n*(n - 1)/2$ MEs.

For a rooted-multipoint ETH connection containing k root and m leaf end-points, it is possible, but not required, for a MEG to contain MEs between leaf end-points; if it does not, the MEG contains $k \times (k - 1)/2 + k \times m$ MEs.

5.2 Traffic conditioning point (TrCP)

A traffic conditioning point (TrCP) is an ETH flow point which is capable of an ETH traffic conditioning function, as specified in [ITU-T G.8010].

5.3 MEG level

In the case where MEGs are nested, the OAM flow of each MEG has to be clearly identifiable and separable from the OAM flows of the other MEGs. In cases where the OAM flows are not distinguishable by the ETH layer encapsulation itself, the MEG level in the OAM frame distinguishes between the OAM flows of nested MEGs.

Eight MEG levels are available to accommodate different network deployment scenarios.

When customer, provider and operator data path flows are not distinguishable based on means of the ETH layer encapsulations, the eight MEG levels can be shared amongst them to distinguish between OAM frames belonging to nested MEGs of customers, providers and operators. The default MEG level assignment amongst customer, provider and operator roles is:

- The customer role is assigned three MEG levels: 7, 6 and 5.
- The provider role is assigned two MEG levels: 4 and 3.
- The operator role is assigned three MEG levels: 2, 1 and 0.

The default MEG level assignment can be changed via a mutual agreement among customer, provider and/or operator roles.

Though eight MEG levels are available, not all MEG levels may be used. When not all eight MEG levels are used, there is no constraint on the continuity of MEG levels (e.g., MEG levels 7, 5, 2 and 0 may be used). The number of MEG levels used depends on the number of nested MEs for which the OAM flows are not distinguishable based on the means of the ETH layer encapsulation.

The specific assignment of MEG levels across different roles in specific deployments is outside of the scope of this Recommendation. Refer to [ITU-T G.8010] for some examples.

5.4 OAM transparency

OAM transparency refers to the ability to allow the transparent carrying of OAM frames belonging to higher-level MEGs across other lower-level MEGs when the MEGs are nested.

OAM frames belonging to an administrative domain originate and terminate in MEPs present at the boundary of that administrative domain. A MEP prevents OAM frames corresponding to a MEG in the administrative domain, from leaking outside that administrative domain. However, when a MEP is not present or is faulty, the associated OAM frames could leave the administrative domain.

Similarly, a MEP present at the boundary of an administrative domain protects the administrative domain from OAM frames belonging to other administrative domains. The MEP allows OAM frames from outside administrative domains belonging to higher-level MEs to pass transparently, while it blocks OAM frames from outside administrative domains belonging to same or lower-level MEs.

The customer role can use any of the eight MEG levels when the MEG levels are not shared with provider and operator roles, as mentioned in clause 5.3. However, if MEG levels are shared with provider and operator roles, transparency of customer's OAM frames across provider's and/or operator's administrative domains will only be guaranteed for mutually agreed MEG levels, e.g., default MEG levels 7, 6 and 5. Similarly, transparency of a provider's OAM frames across an operator's administrative domain when MEG levels are shared will be guaranteed for mutually agreed MEG levels, e.g., default MEG levels 4 and 3, while the operator role can use default MEG levels 2, 1, and 0.

OAM frames can be prevented from leaking by implementing an OAM filtering process in the MEP atomic functions.

5.5 Representation of octets

In this Recommendation, octets are represented as defined in [IEEE 802.1Q].

When consecutive octets are used to represent a binary number, the lower octet number has the most significant value. As an example, if Octet1 and Octet2 in Figure 5.5-1 represent a binary number, Octet1 has the most significant value.

The bits in an octet are numbered from 1 to 8, where bit 1 is the least significant bit (LSB) and bit 8 is the most significant bit (MSB).

	1								2								3								4							
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	Octet1								Octet2								Octet3								Octet4							
5	Octet5								Octet6								Octet7								Octet8							
9	Octet9								Octet10								Octet11								Octet12							
:																																

Figure 5.5-1 – Example PDU format

6 OAM relationships

6.1 MEs, MEPs, MIPs and TrCPs relationship

Appendix I provides different network scenarios to show how MEGs, MEPs and MIPs at different MEG levels can be deployed, and where TrCPs are likely to be placed.

NOTE – Not all MEGs and corresponding MEPs and MIPs may be used or provided in the example network scenarios in Appendix I. For example, providers may not provide customer MIPs.

6.2 MEs, MEGs and MEG level relationship

The MEPs associated with an administrative domain operate at an assigned MEG level. Inter-domain MEPs, associated with MEGs between two administrative domains, can operate at a MEG level agreeable between the two administrative domains, so that associated inter-domain OAM flows are prevented from leaking into either administrative domain. The default MEG level for inter-domain OAM flows is 0.

MEs in Ethernet networks are indicated in Figures 23 and 24 of [ITU-T G.8010] and Ethernet MEs are defined in clause 9 of [ITU-T Y.1730]. MEs can nest but not overlap. Figure 6.2-1 illustrates an example of MEs associated with a point-to-point connection administrative domain.

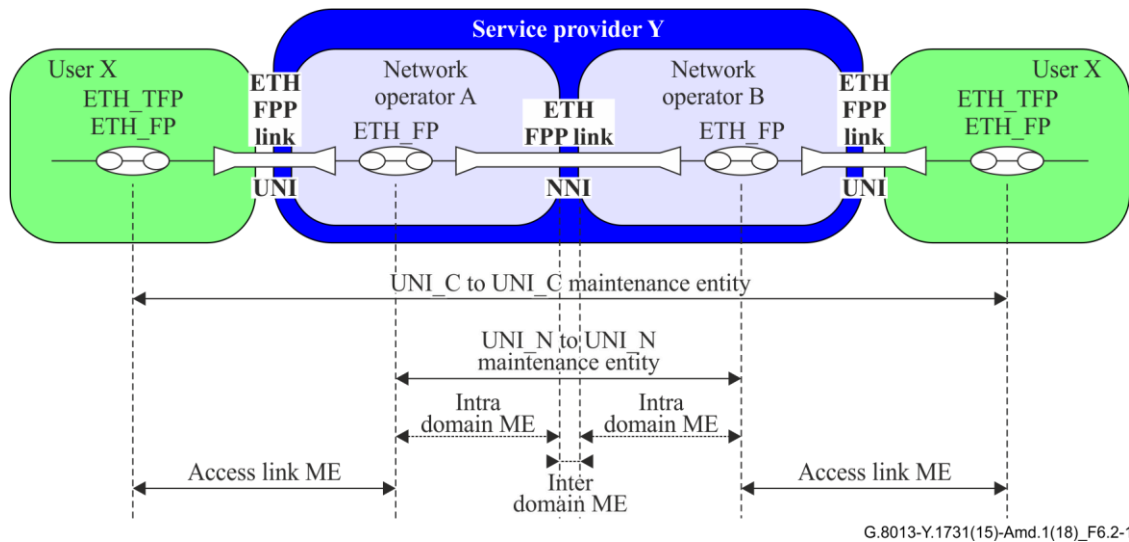


Figure 6.2-1 – Example of MEs associated with a point-to-point connection administrative domain shown in Figure 23 of [ITU-T G.8010]

Table 6-1 highlights possible MEG level assignments for MEGs within the context of customer, provider and operator administrative domains that share the MEG levels, as mapped to [ITU-T G.8010] and [ITU-T Y.1730].

Table 6-1 – Example MEG level assignments for shared MEG levels

ITU-T G.8010 MEG	ITU-T Y.1730 ME	MEG level(s)
UNI-C to UNI-C ME	UNI-UNI (Customer)	7, 6, or 5
UNI-N to UNI-N ME	UNI-UNI (Provider)	4, or 3
Intra-domain ME	Segment (PE-PE) intra-provider	4, or 3
Inter-domain ME	Segment (PE-PE) inter-provider (Provider – Provider)	0 (default)
Access link ME	ETY link OAM – UNI (Customer – Provider)	0 (default)
Inter-domain ME	ETY link OAM – NNI (Operator – Operator)	0 (default)

As mentioned in clause 5.3, MEG levels are shared when the OAM flows of nested MEGs of customer, provider and operator cannot be distinguished based on ETH layer encapsulation. However, when OAM flows of nested MEGs of customer, provider and operator can be distinguished based on ETH layer encapsulation, MEG levels are not shared except for inter-domain MEGs (e.g., MEGs between customer and provider, MEGs between provider and operator, MEGs between operators, MEs between providers, etc.).

Table 6-2 highlights possible MEG level assignments for MEs within the context of customer, provider and operator administrative domains that do not share the MEG levels but require inter-domain MEs.

Table 6-2 – Example MEG level assignments for independent MEG levels

ITU-T G.8010 MEG	ITU-T Y.1730 ME	MEG level(s)
UNI-C to UNI-C ME	UNI-UNI (Customer)	7to1
UNI-N to UNI-N ME	UNI-UNI (Provider)	7to1
Intra-domain ME	Segment (PE-PE) intra-provider	7to1
Inter-domain ME	Segment (PE-PE) inter-provider (Provider – Provider)	0 (default)
Access link ME	ETY link OAM – UNI (Customer – Provider)	0 (default)
Inter-domain ME	ETY link OAM – NNI (Operator – Operator)	0 (default)

Furthermore, if inter-domain MEs are not required, each customer, provider and operator can use any of the eight MEG levels. However, as already stated in clause 5.3, not all MEG levels may be used.

6.3 MEPs and MIPs configuration

MEG end-points (MEPs) and MEG intermediate points (MIPs) are configured via the management plane and/or control plane. The management plane configurations may be carried out through manual local administration of each device or via network management systems (NMS).

This configuration is outside the scope of this Recommendation.

7 OAM functions for fault management

OAM functions for fault management allow detection, verification, localization and notification of different defect conditions.

7.1 Ethernet continuity check (ETH-CC)

The Ethernet continuity check function (ETH-CC) is used for proactive OAM. It is used to detect loss of continuity (LOC) between any pair of MEPs in a MEG. ETH-CC also allows detection of unintended connectivity between two MEGs (mismerge), unintended connectivity within the MEG with an unexpected MEP (unexpected MEP) and other defect conditions (e.g., unexpected MEG level, unexpected period, etc.). ETH-CC is applicable for fault management, performance monitoring, or protection switching applications.

A MEP must always report reception of a frame with unexpected ETH-CC information. ETH-CC transmission may be enabled or disabled in a MEG. When ETH-CC transmission is enabled in a MEG, all MEPs are enabled to periodically transmit frames with ETH-CC information to their peer MEPs in the MEG. The ETH-CC transmission period is the same for all MEPs in the MEG. When a MEP is enabled to generate frames with ETH-CC information, it also expects to receive frames with ETH-CC information from its peer MEPs in the MEG.

When ETH-CC transmission is disabled in a MEG, all MEPs are disabled to transmit frames with ETH-CC information.

The specific configuration information required by each MEP to support ETH-CC is as follows:

- MEG ID – identifies the MEG to which the MEP belongs.
- MEP ID – MEP's own identity in the MEG.
- List of peer MEP IDs – list of peer MEPs in the MEG. For a point-to-point MEG with a single ME, the list would consist of a single MEP ID for the peer.
- MEG level – MEG level at which the MEP exists.
- ETH-CC transmission period – This is application dependent. ETH-CC has three different applications (for each application, a default transmission period is specified):

- Fault management: default transmission period is 1 s (i.e., transmission rate of 1 frame/second);
 - Performance monitoring: default transmission period is 100 ms (i.e., transmission rate of 10 frames/second);
 - Protection switching: default transmission period is 3.33 ms (i.e., transmission rate of 300 frames/second).
- Priority – identifies the priority of frame with ETH-CC information. By default, the frame with ETH-CC information is transmitted with the highest priority available to the data traffic. This information is configurable per operation.
 - Drop eligibility – Frames with ETH-CC information are always marked as drop ineligible. This information is not necessarily configured.

A MIP is transparent to the ETH-CC information and therefore does not require any configuration information to support ETH-CC.

When a MEP does not receive ETH-CC information from a peer MEP, in the list of peer MEPs, within an interval of 3.5 times the ETH-CC transmission period, it detects loss of continuity to that peer MEP. The interval corresponds to a loss of three consecutive frames carrying ETH-CC information from the peer MEP. ETH-CC also allows detection of other defect conditions as described in clause 7.1.2.

The OAM PDU used for ETH-CC information is CCM, as described in clause 9.2. Frames which carry the CCM PDU are called CCM frames.

7.1.1 CCM (with ETH-CC information) transmission

When ETH-CC is enabled, a MEP periodically transmits CCM frames as often as the configured transmission period. The transmission period can be one of the following seven values:

- 3.33 ms: default transmission period for protection switching application (transmission rate of 300 frames/second);
- 10 ms: (transmission rate is 100 frames/second);
- 100 ms: default transmission period for performance monitoring application (transmission rate of 10 frames/second);
- 1 s: default transmission period for fault management application (transmission rate of 1 frame/second);
- 10 s: (transmission rate of 6 frames/minute);
- 1 min: (transmission rate of 1 frame/minute);
- 10 min: (transmission rate of 6 frames/hour);

NOTE – Even though seven different values are specified for the transmission period, the default values are recommended based on the application area for which ETH-CC is being used. When a transmission period other than the default value for an application area is used, the behaviour of the intended application is not guaranteed.

The Period field in CCM is transmitted with a value for the transmission period configured at the transmitting MEP, so that a receiving MEP can detect an unexpected period, if the transmission period is not the same across the transmitting and receiving MEPs.

7.1.2 CCM (with ETH-CC information) reception

When a MEP receives a CCM frame, it examines it to ensure that its MEG ID matches the configured MEG ID in the receiving MEP, and that the MEP ID in the CCM frame is one from the configured list of peer MEP IDs. The information in the CCM frame is catalogued in the receiving MEP.

CCM frames allow the detection of different defect conditions, which include:

- If no CCM frames from a peer MEP are received within the interval equal to 3.5 times, the receiving MEP's CCM transmission period, loss of continuity with peer MEP is detected.
- If a CCM frame with a MEG level lower than the receiving MEP's MEG level is received, unexpected MEG level is detected.
- If a CCM frame with the same MEG level but a MEG ID different than the receiving MEP's own MEG ID is received, mismerge is detected.
- If a CCM frame with the same MEG level and a correct MEG ID but with an incorrect MEP ID, including the receiving MEP's own MEP ID, is received, unexpected MEP is detected.
- If a CCM frame is received with a correct MEG level, a correct MEG ID, a correct MEP ID, but with a period field value different to the receiving MEP's own CCM transmission period, unexpected period is detected.

A receiving MEP must notify the equipment fault management process when it detects the above defect conditions.

7.2 Ethernet loopback (ETH-LB)

The Ethernet loopback function (ETH-LB) is used to verify connectivity of a MEP with a MIP or peer MEP(s). There are two ETH-LB types:

- Unicast ETH-LB.
- Multicast ETH-LB.

7.2.1 Unicast ETH-LB

Unicast ETH-LB is an on-demand OAM function that can be used for the following applications:

- to verify bidirectional connectivity of a MEP with a MIP or a peer MEP;
- to perform a bidirectional in-service or out-of-service diagnostics test between a pair of peer MEPs. This includes verifying bandwidth throughput, detecting bit errors, etc.

Frames with unicast ETH-LB information can be transmitted in several ways for different on-demand command types, e.g., single transmission, repetitive transmission, etc. The specific on-demand command types are outside the scope of this Recommendation.

When used to verify bidirectional connectivity, a MEP sends a unicast frame with ETH-LB request information and expects to receive a unicast frame with ETH-LB reply information from a MIP or peer MEP within a specified period of time. The MIP or peer MEP is identified by its MAC address. This MAC address is encoded in the DA of the unicast request frame. If the MEP does not receive the unicast frame with ETH-LB reply information within the specified period of time, loss of connectivity with the MIP or peer MEP can be inferred. Unicast ETH-LB can also be used to test the bidirectional connectivity with different frame sizes between a MEP and a MIP or peer MEP.

When used for performing bidirectional diagnostics tests, a MEP sends unicast frames with ETH-LB request information to a peer MEP. This ETH-LB request information includes test patterns. When out-of-service diagnostic tests are performed, data traffic is not delivered on either side of the diagnosed ME. Instead the MEPs are configured to send frames with ETH-LCK information, as described in clause 7.6, for the immediate client MEG level on either side of the ME.

NOTE 1 – Unicast ETH-LB can be used to perform only one of the two applications at any time. It must finish the pending on-demand command related to one application (either connectivity verification or diagnostic test) before it can act on a new on-demand command for the other application.

NOTE 2 – The maximum rate at which frames with unicast ETH-LB information can be sent without adversely impacting the data traffic for in-service bidirectional connectivity verification or in-service bidirectional diagnostic tests is outside the scope of this Recommendation. It may be mutually agreed between the user of unicast ETH-LB and the user of the service.

Specific configuration information required by a MEP to support unicast ETH-LB is as follows:

- MEG level – MEG level at which the MEP exists.
- Unicast MAC address of remote MIP or MEP to which ETH-LB is intended. This information is configurable per operation.
- Data – optional element whose length and contents are configurable at the MEP. The contents can be a test pattern and an optional checksum. Examples of test patterns include pseudo-random bit sequence (PRBS) ($2^{31}-1$) as specified in clause 5.8 of [ITU-T O.150], all '0' pattern, etc. For bidirectional diagnostic test applications, configuration is required for a test signal generator and a test signal detector associated with the MEP.
- Priority – identifies the priority of frames with unicast ETH-LB information.
- Drop eligibility – identifies the eligibility of frames with unicast ETH-LB information to be discarded when congestion conditions are encountered.

NOTE 3 – Additional configuration information elements may be needed for repetitive transmission, e.g., repetition rate, total interval of repetition, etc. These additional configuration information elements are outside the scope of this Recommendation.

A remote MIP or MEP, upon receiving the unicast frame with ETH-LB request information which is addressed to the MIP or MEP, responds with a unicast frame with ETH-LB reply information.

Specific configuration information required by a MIP to support unicast ETH-LB is as follows:

- MEG level – MEG level at which the MIP exists.

The OAM PDU used for unicast-LB request information is LBM, as described in clause 9.3. The OAM PDU used for unicast-LB reply information is LBR, as described in clause 9.4. Unicast frames carrying the LBM PDU are called unicast LBM frames. Unicast frames carrying the LBR PDU are called unicast LBR frames.

7.2.1.1 Unicast LBM transmission

Unicast LBM frames are transmitted by a MEP on an on-demand basis.

When used for bidirectional connectivity verification, a MEP transmits a unicast LBM frame addressed to a MIP or peer MEP with a specific transaction ID inserted in the Transaction ID/Sequence Number field. After unicast LBM frame transmission, a MEP expects to receive a unicast LBR frame within 5 seconds. The transmitted transaction ID is therefore retained by the MEP for at least 5 seconds after the unicast LBM frame is transmitted. A different transaction ID must be used for every unicast LBM frame and no transaction ID from the same MEP may be repeated within one minute.

A MEP can optionally use a Data TLV or Test TLV. When configured for checking the successful transmission of different frame sizes, the MEP uses a Data TLV. However, when used for diagnostic tests, a MEP transmits a unicast LBM frame addressed to the peer MEP with a test TLV. The test TLV is used to carry the test pattern generated by a test signal generator associated with the MEP. When the MEP is configured for an out-of-service diagnostic test, the MEP also generates LCK frames, as described in clause 7.6, at the client MEG level.

7.2.1.2 Unicast LBM reception and LBR transmission

Whenever a valid unicast LBM frame is received by a MIP or MEP, an LBR frame is generated and transmitted to the initiating MEP. A unicast LBM frame with a valid MEG level and a destination MAC address equal to the MAC address of responding MIP or MEP is considered to be a valid unicast LBM frame. Every field in the unicast LBM frame is copied to the LBR frame with the following exceptions:

- The source and destination MAC addresses are swapped.
- The OpCode field is changed from LBM to LBR.

Furthermore, when a responding MEP is configured for an out-of-service diagnostic test, it also generates LCK frames, as described in clause 7.6, at the client MEG level.

7.2.1.3 LBR reception

When a MEP configured for connectivity verification receives an LBR frame addressed to it with the same MEG level as its own MEG level, and with an expected transaction ID and within 5 seconds after transmitting the unicast LBM frame, the LBR frame is valid. Otherwise the LBR frame addressed to it is invalid and is discarded.

When a MEP configured for a diagnostics test receives an LBR frame addressed to it with the same MEG level as its own MEG level, the LBR frame is valid. The test signal receiver associated with MEP may also validate the received sequence number against expected sequence numbers.

If a MIP receives an LBR frame addressed to it, such an LBR frame is invalid and the MIP should discard it.

7.2.2 Multicast ETH-LB

The multicast ETH-LB function is used to verify bidirectional connectivity of a MEP with its peer MEPs. Multicast ETH-LB is an on-demand OAM function. When a multicast ETH-LB function is invoked on a MEP, the MEP returns to the initiator of multicast ETH-LB a list of its peer MEPs with whom the bidirectional connectivity is detected.

When multicast-LB is invoked on a MEP, a multicast frame with ETH-LB request information is sent from a MEP to its peer MEPs. The MEP expects to receive a unicast frame with ETH-LB reply information from its peer MEPs within a specified period of time. Upon reception of a multicast frame with ETH-LB request information, the receiving MEPs validate the multicast frame with ETH-LB request information and transmit a unicast frame with ETH-LB reply information after a randomized delay in the range of 0 to 1 second.

Specific configuration information required by each MEP to support multicast ETH-LB is as follows:

- MEG level – MEG level at which the MEP exists.
- Priority – identifies the priority of multicast frames with ETH-LB request information. This information is configurable per operation.
- Drop eligibility – Multicast frames with ETH-LB request information are always marked as drop ineligible. This information is not necessarily configured.

A MIP is transparent to the multicast frames with ETH-LB request information and therefore does not require any information to support multicast ETH-LB.

The OAM PDU used for multicast ETH-LB request information is LBM, as described in clause 9.3. The OAM PDU used for ETH-LB reply is LBR, as described in clause 9.4. Multicast frames carrying the LBM PDU are called multicast LBM frames.

7.2.2.1 Multicast LBM transmission

Multicast LBM frames are transmitted by a MEP on an on-demand basis. After transmitting the multicast LBM frame with a specific transaction ID, the MEP expects to receive LBR frames within 5 seconds. The transmitted transaction ID is therefore retained for at least 5 seconds after the multicast LBM frame has been transmitted. A different transaction ID must be used for every multicast LBM frame, and no transaction ID from the same MEP may be repeated within one minute.

7.2.2.2 Multicast LBM reception and LBR transmission

Whenever a valid multicast LBM frame is received by a MEP, an LBR frame is generated and transmitted to the initiating MEP after a randomized delay in the range of 0 to 1 second. The validity of the multicast LBM frame is determined based on the correct MEG level.

Every field in the multicast LBM frame is copied to the LBR frame with the following exceptions:

- The source MAC address in the LBR frame is the unicast MAC address of the responding MEP. The destination MAC address in the LBR frame is copied from the source MAC address of the multicast LBM frame which should be a unicast address.
- The OpCode field is changed from LBM to LBR.

7.2.2.3 LBR reception

When an LBR frame is received by a MEP with an expected transaction ID and within 5 seconds of transmitting the multicast LBM frame, the LBR frame is valid. If a MEP receives an LBR frame with a transaction ID that is not in the list of transmitted transaction IDs maintained by the MEP, the LBR frame is invalid and is discarded.

If a MIP receives an LBR frame addressed to it, such an LBR frame is invalid and the MIP should discard it.

7.3 Ethernet link trace (ETH-LT)

The Ethernet link trace function (ETH-LT) is an on-demand OAM function that can be used for the following two purposes:

- Adjacent relation retrieval – The ETH-LT function can be used to retrieve an adjacency relationship between a MEP and a peer MEP or MIP. The result of running ETH-LT function is a sequence of MIPs from the initiating MEP until the target MIP or MEP. Each MIP and/or MEP is identified by its MAC address.
- Fault localization – The ETH-LT function can be used for fault localization. When a fault (e.g., a link and/or a device failure) or a forwarding plane loop occurs, the sequence of MIPs and/or MEP will likely be different from the expected one. Difference in the sequences provides information about the fault location.

ETH-LT request information is initiated in a MEP on an on-demand basis. After transmitting a frame with ETH-LT request information, the MEP expects to receive frames with ETH-LT reply information within a specified period of time. Network elements containing MIPs or MEPs and receiving the frame with ETH-LT request information respond selectively with frames containing ETH-LT reply information.

A network element containing MIP or MEP responds with a frame with ETH-LT reply information upon receiving a valid frame with ETH-LT request information only if:

- the network element where the MIP or MEP resides is aware of the TargetMAC address in the ETH-LT request information and associates it to a single egress port, where the egress port is not the same as the port on which the frame with ETH-LT request information was received; or
- the TargetMAC address is the same as the MIP's or MEP's own MAC address.

A network element containing MIPs may also relay the frame with ETH-LT request information, as described in clause 7.3.2.

The specific configuration information required by a MEP to support ETH-LT is as follows:

- MEG level – MEG level at which the MEP exists.
- Priority – identifies the priority of the frames with ETH-LT request information. This information is configurable per operation.
- Drop eligibility – Frames with ETH-LT information are always marked as drop ineligible. This information is not necessarily configured.
- Target MAC address (usually of MIPs or MEPs of the MEG, but not limited to that) for which ETH-LT is intended. This information is configurable per operation.

- TTL – allows the receiver to determine if frames with ETH-LT request information can be terminated. TTL is decremented every time frames with ETH-LT request information are relayed. Frames with ETH-LT request information with $TTL \leq 1$ are not relayed.

The specific configuration information required by a MIP to support ETH-LT is as follows:

- MEG level – MEG level at which the MIP exists.

The PDU used for ETH-LT request information is LTM, as described in clause 9.5. The PDU used for ETH-LT reply information is LTR, as described in clause 9.6. Frames carrying the LTM PDU are called LTM frames. Frames carrying the LTR PDU are called LTR frames.

NOTE 1 – As each network element, containing the MIPs or MEP, needs to be aware of the TargetMAC address in the received LTM frame and associates it to a single egress port, in order that the MIP or MEP can process the received LTM frames, a unicast ETH-LB to the TargetMAC address could be performed by a MEP before transmitting the LTM frame. This would ensure that the network elements along the path to the TargetMAC address would have information about the route to the TargetMAC address if the TargetMAC address is reachable in the same MEG.

NOTE 2 – During a failure condition the information about the route to the TargetMAC address may age out after a certain time. The ETH-LT function has to be performed before the age out occurs in order to provide information about the route.

7.3.1 LTM transmission

An LTM frame is transmitted by a MEP on an on-demand basis. If the MEP resides at an ingress port, the LTM frame is forwarded towards the network element's own ETH-LT responder. However, if the MEP resides on an egress port, the LTM frame is transmitted out of that egress port. The LTM frame contains an LTM egress identifier TLV which identifies the network element initiating the LTM frame.

NOTE – ETH-LT responder is not defined in [ITU-T Y.1731], only the MEP and MIP of ingress and egress ports are defined. And LTM egress identifier TLV is regarded as optional in [ITU-T Y.1731].

After transmitting the LTM frame with a specific transaction number, the MEP expects to receive LTR frames within 5 seconds. The transaction number of each LTM frame transmitted is therefore retained for at least 5 seconds after the LTM frame is transmitted. A different transaction number must be used for every LTM frame, and no transaction number from the same MEP may be repeated within one minute.

7.3.2 LTM reception and forwarding, and LTR transmission

If an LTM frame is received by a MEP or MIP, it forwards the LTM frame to the network element's ETH-LT responder, which performs the following validation:

- Only LTM frames with the same MEG level as the receiving MEP's or MIP's own MEG level are validated.
- Thereafter, the TTL field value of the LTM frame is checked. If the TTL field value is 0, the LTM frame is discarded (a TTL field value of 0 is an invalid value).
- Thereafter, the LTM frame is checked to see if LTM egress identifier TLV is present. The LTM frame is discarded if it does not contain LTM egress identifier TLV. It is noted that the LTM frame generated by [ITU-T Y.1731] may not contain the LTM egress identifier TLV. See Annex B for keeping the compatibility, i.e., LTM frame TLV may be processed at the MIP or MEP even if the LTM egress identifier TLV is absent.

If the LTM frame is valid, the ETH-LT responder does the following:

- It determines the destination address for the LTR frame from the OriginMAC address in the received LTM frame.
- If the network element is aware of the TargetMAC address in the LTM frame and associates it with a single egress port, where the egress port is not the same as the ingress port, or the LTM frame terminates at the MIP or MEP (when the TargetMAC address is the MIP's or

MEP's own MAC address), an LTR frame is sent backwards to the initiating MEP after a random time interval in the range of 0 to 1 second.

- Furthermore, if the above condition applies and the LTM frame does not terminate at the MIP (when the TargetMAC address is not the same as the MIP's own address, if received by a MIP) and the TTL field in the LTM frame is greater than 1, the LTM frame is forwarded towards the single egress port. All the fields of the relayed LTM frame are the same as the original LTM frame except for TTL which is decremented by 1, the source address which becomes the MIP's own MAC address, and LTM egress identifier TLV which identifies the network element relaying the modified LTM frame. It is noted that MIPs supporting [ITU-T Y.1731] may forward the LTM egress identifier TLV as it is. See Annex B for keeping the compatibility.
- Furthermore, when the TargetMAC address is not the same as the MEP's own address, if received by a MEP, the LTM frames are always terminated at the MEP and the MEP does not send back the LTR frames.

The LTR frame contains LTR egress identifier TLV which identifies the source and destination of the LTM that triggered the transmission of this LTR. LTR egress identifier TLV contains the Last Egress Identifier field which identifies the network element that originated or forwarded the LTM frame for which this LTR frame is the response. This field takes the same value as the LTM egress identifier TLV of that LTM frame. The LTR egress identifier TLV also contains the Next Egress Identifier field which identifies the network element that transmitted this LTR frame, and can relay a modified LTM frame to the next hop. This field takes the same value as the LTM egress identifier TLV of the relayed modified LTM frame, if any. If no modified LTM frame is relayed, the FwdYes bit of the Flags field in the LTM frame is clear and the contents of next egress identifier are undefined, and must be ignored by the LTR frame receiver.

Additionally, if the LTM frame was received by a MIP or MEP at an ingress port, the LTR frame includes a reply ingress TLV which describes the MIP or MEP at the ingress port.

Similarly, if the LTM frame was not received by a MEP at the ingress port, and if the egress port has a MIP or MEP, the LTR frame includes a reply egress TLV which describes the MIP or MEP at the egress port.

It is noted that both included Reply Ingress TLV and Reply Egress TLV are documented as optional in [ITU-T Y.1731] so that they may not be included in the LTR frame of that version. See Annex B for keeping the compatibility.

7.3.3 LTR reception

When an LTR frame is received by a MEP with an expected transaction number and within 5 seconds of transmitting the LTM frame, the LTR frame is valid. If a MEP receives an LTR frame with a transaction number that is not in the list of transmitted transaction numbers maintained by the MEP, the LTR frame is invalid.

If a MIP receives an LTR frame addressed to it, such an LTR frame is invalid and the MIP should discard it.

7.4 Ethernet alarm indication signal (ETH-AIS)

The Ethernet alarm indication signal function (ETH-AIS) is used to suppress alarms following detection of defect conditions at the server (sub) layer. Due to independent restoration capabilities provided within the Spanning Tree Protocol (STP) environments, ETH-AIS is not expected to be applied in the STP environments.

Transmission of frames with ETH-AIS information can be enabled or disabled on a MEP (or on a server MEP).

Frames with ETH-AIS information can be issued at the client MEG level by a MEP, including a server MEP, upon detecting defect conditions. For example, the defect conditions may include:

- signal fail conditions in the case that ETH-CC is enabled;
- AIS condition or LCK condition in the case that ETH-CC is disabled.

NOTE – Since a server MEP does not run ETH-CC, a server MEP can transmit frames with ETH-AIS information upon detection of any signal fail condition.

For multipoint ETH connectivity, a MEP cannot determine the specific server (sub) layer entity that has encountered defect conditions upon receiving a frame with ETH-AIS information. More importantly, it cannot determine the associated subset of its peer MEPs for which it should suppress alarms since the received ETH-AIS information does not contain that information. Therefore, upon reception of a frame with ETH-AIS information, the MEP will suppress alarms for all peer MEPs whether there is still connectivity or not.

However, for a point-to-point ETH connection, a MEP has only a single peer MEP. Therefore, there is no ambiguity regarding the peer MEP for which it should suppress alarms when it receives the ETH-AIS information.

Only a MEP, including a server MEP, is configured to issue frames with ETH-AIS information. Upon detecting a defect condition the MEP can immediately start transmitting periodic frames with ETH-AIS information at a configured client MEG level. A MEP continues to transmit periodic frames with ETH-AIS information until the defect condition is removed. Upon receiving a frame with ETH-AIS information a MEP detects the AIS condition and suppresses loss of continuity alarms associated with all its peer MEPs. A MEP resumes loss of continuity alarm generation upon detecting loss of continuity defect conditions in the absence of an AIS condition.

The specific configuration information required by a MEP to support ETH-AIS transmission is as follows:

- Client MEG level – MEG level at which the most immediate client layer MIPs and MEPs exist.
- ETH-AIS transmission period – determines transmission periodicity of frames with ETH-AIS information.
- Priority – identifies the priority of frames with ETH-AIS information.
- Drop eligibility – Frames with ETH-AIS information are always marked as drop ineligible. This information is not necessarily configured.

Specific configuration information required by a MEP to support ETH-AIS reception is as follows:

- Local MEG level – MEG level at which the MEP operates.

A MIP is transparent to frames with ETH-AIS information and therefore does not require any information to support ETH-AIS functionality.

The PDU used for ETH-AIS information is AIS, as described in clause 9.7. Frames carrying the AIS PDU are called AIS frames.

7.4.1 AIS transmission

A MEP, upon detecting a defect condition, can transmit AIS frames in a direction opposite to its peer MEP(s). The periodicity of AIS frames transmission is based on the AIS transmission period. An AIS transmission period of 1 second is recommended. The first AIS frame must always be transmitted immediately following the detection of a defect condition.

The client (sub) layer may consist of multiple MEGs that should be notified to suppress alarms resulting from defect conditions detected by the server (sub) layer MEP. The server (sub) layer MEP, upon detecting the signal fail condition, needs to send AIS frames to each of these client (sub) layer

MEGs. In such cases, the first AIS frame for all client (sub) layer MEGs must be transmitted within 1 second of defect condition.

NOTE – To support ETH-AIS across current equipment, which may be stressed when issuing AIS frames every 1 second potentially across all 4094 VLANs, another AIS transmission period of 1 minute is also supported. An AIS frame communicates the used AIS transmission period via the Period field.

7.4.2 AIS reception

Upon receiving an AIS frame, a MEP examines it to ensure that its MEG level corresponds to its own MEG level. The Period field indicates the period at which the AIS frames can be expected. Upon receiving an AIS frame, the MEP detects the AIS defect condition. Following the detection of the AIS defect condition, if no AIS frames are received within an interval of 3.5 times the AIS transmission period indicated in the AIS frames received before, the MEP clears the AIS defect condition.

7.5 Ethernet remote defect indication (ETH-RDI)

The Ethernet remote defect indication function (ETH-RDI) can be used by a MEP to communicate to its peer MEPs that a defect condition has been encountered. ETH-RDI is used only when ETH-CC transmission is enabled.

ETH-RDI has the following two applications:

- Single-ended fault management: The receiving MEP detects an RDI defect condition, which gets correlated with other defect conditions in this MEP and may become a fault cause. The absence of received ETH-RDI information in a single MEP indicates the absence of defects in the entire MEG.
- Contribution to far-end performance monitoring: It reflects that there was a defect condition in the far-end which is used as an input to the performance monitoring process.

A MEP that is in a defect condition transmits frames with ETH-RDI information. A MEP, upon receiving frames with ETH-RDI information, determines that its peer MEP has encountered a defect condition. However, for multipoint ETH connectivity, a MEP, upon receiving frames with ETH-RDI information, cannot determine the associated subset of its peer MEPs with which the MEP transmitting RDI information encounters defect conditions, as the transmitting MEP itself does not always have that information.

The specific configuration information required by a MEP to support the ETH-RDI function is as follows:

- MEG level – MEG level at which the MEP exists.
- ETH-RDI transmission period – application dependent and is configured to be the same as ETH-CC transmission period.
- Priority – identifies the priority of frames with ETH-RDI information. The priority is the same as ETH-CC priority.
- Drop eligibility – Frames with ETH-RDI information are always marked as drop ineligible. This information is not necessarily configured.

A MIP is transparent to frames with ETH-RDI information and therefore does not require any configuration information to support ETH-RDI functionality.

The PDU used to carry ETH-RDI information is CCM, as described in clause 9.2.

7.5.1 CCM with ETH-RDI transmission

A MEP, upon detecting a defect condition with its peer MEP, sets the RDI field in the CCM frames for the duration of the defect condition. CCM frames, as described in clause 7.1.1, are transmitted periodically based on the CCM transmission period, when the MEP is enabled for CCM frames

transmission. When the defect condition clears, the MEP clears the RDI field in the CCM frames in subsequent transmissions.

7.5.2 CCM with ETH-RDI reception

Upon receiving a CCM frame, a MEP examines it to ensure that its MEG level corresponds to its configured MEG level and detects RDI condition if the RDI field is set. For a point-to-point ETH connection, a MEP can clear the RDI condition when it receives the first CCM frame from its peer MEP with the RDI field cleared. For multipoint ETH connectivity, a MEP can clear the RDI condition when it receives the CCM frames from its entire list of peer MEP with the RDI field cleared.

7.6 Ethernet locked signal (ETH-LCK)

The Ethernet locked signal function (ETH-LCK) is used to communicate the administrative locking of a server (sub) layer MEP and consequential interruption of data traffic forwarding towards the MEP expecting this traffic. It allows a MEP receiving frames with ETH-LCK information to differentiate between a defect condition and an administrative locking action at the server (sub) layer MEP. An example of an application that would require administrative locking of a MEP is the out-of-service ETH-Test, as described in clause 7.7.

A MEP continues to transmit periodic frames with ETH-LCK information at the configured client MEG level until the administrative/diagnostic condition is removed.

A MEP extracts frames with ETH-LCK information at its own MEG level and detects a LCK condition, which contributes to the signal fail condition of the MEP. The signal fail condition may result in the transmission of AIS frames to its client MEPs.

Specific configuration information required by a MEP to support ETH-LCK transmission is as follows:

- Client MEG level – MEG level at which the most immediate client layer MIPs and MEPs exist.
- ETH-LCK transmission period – determines transmission periodicity of frames with ETH-LCK information.
- Priority – identifies the priority of frames with ETH-LCK information.
- Drop eligibility – Frames with ETH-LCK information are always marked as drop ineligible. This information is not necessarily configured.

The specific configuration information required by a MEP to support ETH-LCK reception is as follows:

- Local MEG level – MEG level at which the MEP operates.

A MIP is transparent to the frames with ETH-LCK information and therefore does not require any information to support ETH-LCK functionality.

The PDU used for ETH-LCK information is LCK, as described in clause 9.8. Frames carrying the LCK PDU are called LCK frames.

7.6.1 LCK transmission

A (server) MEP, when administratively locked, transmits LCK frames to each of its client (sub-) layer MEGs as shown in Figure 7.6-1.




Figure 7.6-1 – Example of ETH-LCK transmission

The periodicity of LCK frames transmission is based on the LCK transmission period. The LCK transmission period is the same as the AIS transmission period. The first LCK frame must always be transmitted immediately following the administrative/ diagnostic action.

The client (sub) layer may consist of multiple MEGs that should be notified to suppress alarms resulting from intentional maintenance/diagnostic related configuration at the server (sub) layer MEP. The server (sub) layer MEP, upon being administratively locked, needs to send LCK frames to each of its client (sub) layer MEGs. In such cases, the first LCK frame for all client (sub) layer MEGs must be transmitted within 1 second of defect condition.

7.6.2 LCK reception

Upon receiving an LCK frame, a MEP examines it to ensure that its MEG level corresponds to its configured MEG level. The Period field indicates the periodicity at which the LCK frames can be expected. Upon receiving an LCK frame, the MEP detects an LCK condition. Following detection of an LCK condition, if no LCK frames are received within an interval of 3.5 times the LCK transmission period indicated in the LCK frames received before, the MEP clears the LCK condition.

7.7 Ethernet test signal (ETH-Test)

The Ethernet test signal function (ETH-Test) is used to perform one-way on-demand in-service or out-of-service diagnostics tests. This includes verifying bandwidth throughput, frame loss, bit errors, etc.

When configured to perform such tests, a MEP inserts frames with ETH-Test information with specified throughput, frame size and transmission patterns.

When the out-of-service ETH-Test function is performed, client data traffic is disrupted in the diagnosed entity. The MEP configured for the out-of-service test transmits LCK frames, as described in clause 7.6, for the immediate client ETH (sub) layer.

When an in-service ETH-Test function is performed, data traffic is not disrupted and the frames with ETH-Test information are transmitted in such a manner that a limited part of the service bandwidth is utilized. This rate of transmission for frames with ETH-Test information is pre-determined for in-service ETH-Test function.

NOTE 1 – The maximum rate at which frames with ETH-Test information can be sent without adversely impacting the data traffic for an in-service ETH-Test is outside the scope of this Recommendation. It may be mutually agreed between the user of ETH-Test and the user of the service.

The specific configuration information required by a MEP to support ETH-Test is as follows:

- MEG level – MEG level at which the MEP exists.
- Unicast MAC address of the peer MEP for which ETH-Test is intended; this information is configurable per operation.
- Data – optional element whose length and contents are configurable at the MEP. The contents can be a test pattern and an optional checksum. Examples of test patterns include pseudo-random bit sequence (PRBS) ($2^{31}-1$) as specified in clause 5.8 of [ITU-T O.150], all '0' pattern, etc. At the initiating MEP, configuration is required for a test signal generator

which is associated with the MEP. At a receiving MEP, configuration is required for a test signal detector which is associated with the MEP.

- Priority – identifies the priority of frames with ETH-Test information. This information is configurable per operation.
- Drop eligibility – Identifies the eligibility of frames with ETH-Test information to be dropped when congestion conditions are encountered.

NOTE 2 – Additional configuration information elements may be needed, such as the transmission rate of ETH-Test information, the total interval of ETH-Test, etc. These additional configuration information elements are outside the scope of this Recommendation.

A MIP is transparent to the frames with ETH-Test information and therefore does not require any configuration information to support ETH-Test functionality.

A MEP inserts frames with ETH-Test information towards a targeted peer MEP. The receiving MEP detects these frames with ETH-Test information and makes the intended measurements.

The PDU used for ETH-Test information is TST, as described in clause 9.9. Frames carrying the TST PDU are called TST frames.

7.7.1 TST transmission

A test signal generator associated with a MEP can transmit TST frames as often as the test signal generator configuration. Each TST frame is transmitted with a specific sequence number. A different sequence number must be used for every TST frame, and no sequence number from the same MEP may be repeated within one minute.

When a MEP is configured for an out-of-service test, the MEP also generates LCK frames for the immediate client MEG level.

7.7.2 TST reception

When a MEP receives TST frames, it examines them to ensure that the MEG level corresponds to its own configured MEG level. If the receiving MEP is configured for ETH-TST function, the test signal detector associated with the MEP detects bit errors from the pseudo-random bit sequence of the received TST frames and reports such errors. Furthermore, when the receiving MEP is configured for an out-of-service test, it also generates LCK frames for the client MEG level.

7.8 Ethernet automatic protection switching (ETH-APS)

The Ethernet automatic protection switching function (ETH-APS) is used to control protection switching operations to enhance reliability. The specific details of protection switching operations are outside the scope of this Recommendation.

The OAM frame type used for ETH-APS is APS frame, as described in clause 9.10.

Applications of ETH-APS mechanisms are defined in [ITU-T G.8031] and [ITU-T G.8032].

7.9 Ethernet maintenance communication channel (ETH-MCC)

The Ethernet maintenance communication channel function (ETH-MCC) provides a maintenance communication channel between a pair of MEPs. ETH-MCC can be used to perform remote management. The specific use of ETH-MCC with an OUI other than the ITU-T OUI (00-19-A7) is outside the scope of this Recommendation.

A MEP can send a frame with ETH-MCC information to its peer MEP with remote maintenance request, remote maintenance reply, notification, etc.

Specific configuration information required by a MEP to support ETH-MCC is as follows:

- MEG level – MEG level at which the MEP exists.

- Unicast MAC address of the remote MEP for which ETH-MCC is intended.
- OUI – organizationally unique identifier (OUI) used to identify the organization defining a specific format and meaning of ETH-MCC.
- Data – additional information that may be needed and is dependent on the specific application of ETH-MCC. Application specific information is outside the scope of this Recommendation.
- Priority – identifies the priority of frames with ETH-MCC information. This information is configurable per operation.
- Drop eligibility – Frames with ETH-MCC information are always marked as drop ineligible. This information is not necessarily configured.

A peer MEP, upon receiving a frame with ETH-MCC information and with a correct MEG level, passes the ETH-MCC information to the management agent which may additionally respond.

A MIP is transparent to the frames with ETH-MCC information and therefore does not require any configuration information to support ETH-MCC functionality.

The PDU used for ETH-MCC information is MCC, as described in clause 9.11. Frames carrying the MCC PDU are called MCC frames.

7.10 Ethernet experimental OAM (ETH-EXP)

Ethernet experimental OAM (ETH-EXP) is used for the experimental OAM functionality that can be provided within an administrative domain on a temporary basis. Interoperability of the experimental OAM functionality and hence use of ETH-EXP containing a given OUI, is not expected across different administrative domains.

NOTE – Use for other different purposes, requiring for example processing of an embedded SDO-specific OUI, is undesirable and not recommended.

The specific application of ETH-EXP is outside the scope of this Recommendation.

EXM PDU, as described in clause 9.17, and EXR PDU, as described in clause 9.18, can be used for experimental OAM. Details of experimental OAM mechanisms are outside the scope of this Recommendation.

7.11 Ethernet vendor-specific OAM (ETH-VSP)

Ethernet vendor-specific OAM (ETH-VSP) is used for vendor-specific OAM functionality that may be provided by a vendor across its own equipment. Interoperability of vendor-specific OAM functionality and hence use of ETH-VSP containing a given OUI, is not expected across different vendors' equipment.

NOTE – Use for other different purposes, requiring for example processing of an embedded SDO-specific OUI, is undesirable and not recommended.

The specific application of ETH-VSP is outside the scope of this Recommendation.

VSM PDU, as described in clause 9.19, and VSR PDU, as described in clause 9.20, can be used for vendor-specific OAM. Details of vendor-specific OAM mechanisms are outside the scope of this Recommendation.

7.12 Ethernet client signal fail (ETH-CSF)

The Ethernet client signal fail function (ETH-CSF) is used by a MEP to propagate to a peer MEP the detection of a failure or defect event in an Ethernet client signal when the client itself does not support appropriate fault or defect detection or propagation mechanisms such as ETH-CC or ETH-AIS. The ETH-CSF messages propagate in the direction from the Ethernet MEP, associated with the ingress client port detecting the failure or defect event, to the Ethernet peer MEP.

ETH-CSF is only applicable to point-to-point Ethernet transport applications. In particular, the use of ETH-CSF with [IEEE 802.1Q] or other Ethernet Spanning Tree Protocol (STP)-based networking environments is strictly restricted to point-to-point segments of the Ethernet flow. The use of client signal fail indications to support client failure applications is described in Appendix VIII of [ITU-T G.806].

Specific configuration information required by a MEP to support ETH-CSF transmission is:

- Local MEG level – MEG level at which the initiating MEP operates.
- ETH-CSF transmission period – determines transmission periodicity of frames with ETH-CSF information.
- Priority – identifies the priority of frames with ETH-CSF information.
- Drop eligibility – Frames with ETH-CSF information are always marked as drop ineligible.

Specific configuration information required by a MEP to support ETH-CSF reception is:

- Local MEG level – MEG level at which the receiving MEP operates.

A MIP is transparent to frames with ETH-CSF information and therefore does not require any information to support ETH-CSF functionality.

The ETH-CSF message indicates also the type of defect. Three CSF defect types are currently defined:

- Client loss of signal (C-LOS);
- Client forward defect indication (C-FDI);
- Client reverse defect indication (C-RDI).

The PDU used to convey ETH-CSF information is referred to as CSF PDU, as described in clause 9.21. Frames carrying the ETH-CSF indications are also referred to as CSF frames.

7.12.1 CSF transmission

Frames with ETH-CSF information can be issued by a MEP, upon notification of an Ethernet CSF event from the corresponding ingress client port. Detection rules for Ethernet CSF events are Ethernet client and application specific.

Transmission of packets with CSF information can be enabled or disabled on a MEP.

Upon receiving an Ethernet CSF notification from the ingress client port the associated MEP can immediately start periodic transmission of frames with ETH-CSF information. A MEP continues periodic transmission of frames with ETH-CSF information until the Ethernet CSF indication is removed by the source adaptation function.

Clearing an Ethernet CSF condition is Ethernet client and application specific. The clearance of the Ethernet CSF condition by the source adaptation function is communicated to the peer MEP via:

- the non-sending ETH-CSF, or
- the forwarding of a ETH-CSF PDU with client defect clear indication (C-DCI) information.

7.12.2 CSF reception

Upon receiving a CSF frame with ETH-CSF information a MEP declares the beginning or end of an Ethernet remote CSF condition, depending on the received ETH-CSF information as described in [ITU-T G.8021] and propagates this Ethernet client defect condition towards the corresponding egress client port. An Ethernet MEP detects an Ethernet remote CSF condition when an ETH-CSF PDU with no C-DCI information is received.

The clearance of the Ethernet remote CSF condition by the Ethernet client is detected when:

- no ETH-CSF frame is received in within an interval of N times the CSF transmission period ms (suggested value of N is 3.5), or
- an ETH-CSF PDU with client defect clear indication (C-DCI) information is received.

Note that consequent actions by the sink adaptation function associated with the MEP to propagate the received ETH-CSF information to the Ethernet client are by definition Ethernet client and application specific.

7.13 Ethernet bandwidth notification (ETH-BN)

The Ethernet Bandwidth Notification function (ETH-BN) is used by a server MEP to signal the server layer link bandwidth in the transmit direction to a MEP at the client layer, for example when the server layer runs over a microwave link which has the capability to adapt its bandwidth according to the prevailing atmospheric conditions. Frames with ETH-BN information carry the current and nominal bandwidth of the server layer link. On receiving frames with ETH-BN information, the client layer MEP can use bandwidth information to adjust service policies, e.g., to reduce the rate of traffic being directed towards the degraded link.

Transmission of frames with ETH-BN information can be enabled or disabled on a server MEP. Only a server MEP can transmit frames with ETH-BN information.

When enabled, frames with ETH-BN information are transmitted at the client MEG level by a server MEP, upon detecting bandwidth degradation conditions. A server MEP continues to transmit periodic frames with ETH-BN information until the full bandwidth is restored. In addition, periodic frames with ETH-BN information may optionally be sent when there is no degradation or when the bandwidth degrades to 0.

In a multipoint client MEG, frames with ETH-BN information may need to include a Port Identification (Port ID), to identify which port is associated with the ETH-BN information. This is required if server MEPs for different links transmit frames using the same source MAC address.

Upon receiving a frame with ETH-BN information, a MEP passes the received information to the management system. The management system may take further action to reduce the rate of traffic being directed towards the degraded link or otherwise adjust the service policy for the link.

NOTE – Use of ETH-BN for protection switching is for further study.

The specific configuration information required by a server MEP to support ETH-BN transmission is as follows:

- Client MEG level – MEG level at which the most immediate client layer MIPs and MEPs exist.
- ETH-BN transmission period – determines transmission periodicity of frames with ETH-BN information.
- Hold time – determines the time between detecting degradation, and transmission of the first frame with BNM information that indicates degradation (up to 10 s).
- Priority – identifies the priority of frames with ETH-BN information.
- Drop eligibility – Frames with ETH-BN information are always marked as drop ineligible. This information is not necessarily configured.
- Port ID – A 32-bit unique identifier for the port; this is needed in multipoint MEGs if frames with ETH-BN information about different ports would otherwise be identical. It is optional otherwise. The value must be unique over all server links within the client MEG.

Specific configuration information required by a MEP to support ETH-BN reception is as follows:

- Local MEG level – MEG level at which the MEP operates.

A MIP is transparent to frames with ETH-BN information and therefore does not require any information to support ETH-BN functionality.

The PDU used for ETH-BN information is BNM, as described in clause 9.25. Frames carrying the BNM PDU are called BNM frames.

7.13.1 BNM transmission

A server MEP, upon detecting a transmission bandwidth degrade condition, can transmit periodic BNM frames in a direction opposite to its peer server MEP, indicating that the current bandwidth is less than the nominal bandwidth. Transmission of BNM frames is shown in Figure 7.13-1.

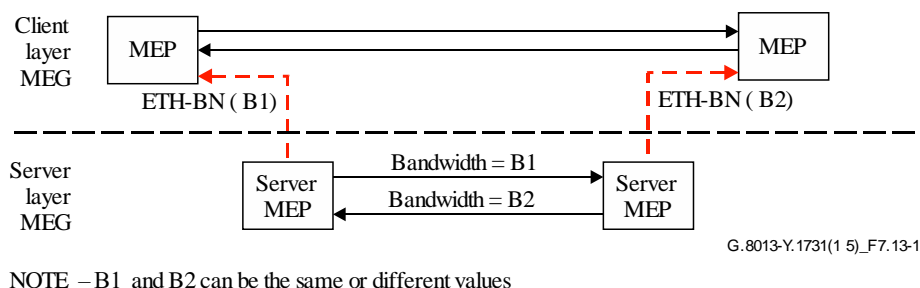


Figure 7.13-1 – Example of ETH-BN transmission

A server MEP may also transmit periodic BNM frames when there is no degradation, indicating that the current and nominal bandwidth are the same, or when the port detects failure, to indicate that the current bandwidth is 0.

NOTE 1 – When the port detects failure, AIS frames are also transmitted by the peer server MEP.

Upon detecting a change in the transmission bandwidth, the first BNM frame that indicates the new transmission bandwidth must be transmitted after the hold time (up to 10s) after detection of the transmission bandwidth change provided the condition has persisted for that time. If the change lasts for less than the hold time, no BNM frame indicating the change of transmission bandwidth is transmitted.

NOTE 2 – BNM notifications are expected to be used where the server layer is a microwave link that uses adaptive bandwidth modulation. A hold time is used to prevent notifications if the degradation is very short, such as might be caused by an object passing through the line of sight of the microwave. The applicability of BNM notifications to other technologies is for further study.

The first BNM frames, indicating the change of transmission bandwidth, are transmitted in quick succession so that reliable and fast actions at the receiver MEP are possible even if some BNM frames are lost or corrupted. The interval and the number of the first BNM frames are implementation specific.

~~The periodicity of BNM frame transmission is based on the configured value, and this is also communicated via the Period field in all the BNM frames. When detecting transmission bandwidth degrade, after the transmission of the first BNM frames, the server MEP starts transmission of periodic BNM frames; otherwise, w~~When detecting the full bandwidth recovery or a link failure, after the transmission of first BNM frames, the server MEP may cease transmission of ~~periodic~~ periodic BNM frames. Alternatively, t

~~The periodic BNM frames may be transmitted by configuration even when there is no degradation or full bandwidth recovery. The periodicity is based on the same configured value for periods of degradation. The interval of periodic BNM frame transmission is based on the configured value, and this is also communicated via the Period field in all the BNM frames.~~

7.13.2 BNM reception

Upon receiving a BNM frame, a MEP examines it to ensure that its MEG level corresponds to its own MEG level. The Period field indicates the period at which the BNM frames can be expected. The source MAC, Port ID and bandwidth information are extracted and passed to the management system. Subsequently, if no BNM frames are received within an interval of 3.5 times the BNM transmission period indicated in the last BNM frame received, the MEP signals to the management system that it no longer has any bandwidth information (e.g., because the full bandwidth has been restored).

As described in clause 7.13.1, the first BNM frames are transmitted in quick succession upon detecting a change in the transmission bandwidth. In this case, BNM frames are also received in quick succession to detect the change of bandwidth.

7.14 Ethernet expected defect function (ETH-ED)

The Ethernet expected defect function (ETH-ED) is used by a MEP to signal to its peer MEPs that transmission of CCM frames is expected to be interrupted, without any interruption to data frames, and that the consequent loss of continuity defects at the peer MEPs should therefore be suppressed. Frames with ETH-ED information carry the MEP ID of the MEP and the expected duration of the interruption.

Frames with ETH-ED information are transmitted by a MEP shortly in advance of an expected interruption of CCM frame transmission, if no interruption is expected in the forwarding of data frames. Examples of this are when an in-service software or firmware upgrade is performed, or when a new MEP is added to an existing MEG.

Upon receiving a frame with ETH-ED information, a MEP passes the received information to the element management function (EMF). If enabled by the management system, the EMF can take action to disable the reception of CCMs, and hence avoid any loss of continuity defects that would otherwise be triggered.

NOTE – Further details on how expected defect notifications can be used, and considerations on handling received notifications in the EMF at the peer MEP can be found in Appendix IX of [ITU-T G.8021].

The specific configuration information required by a MEP to support ETH-ED transmission is as follows:

- MEG level – MEG level at which the MEP exists.
- MEP ID – The MEP's identity within the MEG.
- Expected defect duration – duration for which the peer MEPs are requested to suppress loss of continuity alarms.
- ETH-ED transmission period – determines transmission periodicity of frames with ETH-ED information.
- Priority – identifies the priority of frames with ETH-ED information.
- Drop eligibility – Frames with ETH-ED information are always marked as drop ineligible. This information is not necessarily configured.

Specific configuration information required by a MEP to support ETH-ED reception is as follows:

- Local MEG level – MEG level at which the MEP operates.

A MIP is transparent to frames with ETH-ED information and therefore does not require any information to support ETH-ED functionality.

The PDU used for ETH-ED information is EDM, as described in clause 9.26. Frames carrying the EDM PDU are called EDM frames.

7.14.1 EDM transmission

A MEP can transmit one or more periodic EDM frames shortly in advance of an expected interruption of CCM frame transmission, or when CCM frame transmission has not yet commenced. Transmission of EDM frames ceases once the interruption occurs or when normal CCM transmission is (re)started.

7.14.2 EDM reception

Upon receiving an EDM frame, a MEP examines it to ensure that its MEG level corresponds to its own MEG level. The source MEP ID and the expected duration are extracted and passed to the management system.

8 OAM functions for performance monitoring

OAM functions for performance monitoring allow measurement of different performance parameters. The functions and measurement methods for point-to-point ETH connections and multipoint ETH connectivity are defined.

This Recommendation covers the following performance parameters which are based on [MEF 10.3].

- **Frame loss ratio**

Frame loss ratio is defined as a ratio, expressed as a percentage, of the number of frames not delivered divided by the total number of frames during time interval T, where the number of frames not delivered is the difference between the number of frames arriving at the ingress ETH flow point to be delivered to the egress ETH flow point and the number of frames delivered at the egress ETH flow point in a point-to-point ETH connection or multipoint ETH connectivity. The frame loss ratio may be measured using either service frames or synthetic frames, belonging to single CoS. The use of synthetic frames can also be applicable for multipoint ETH connectivity. The use of service frames is applicable only to point-to-point ETH connection where all the frames arriving at the ingress ETH flow point are to be delivered to the egress ETH flow point.

- **Frame delay**

Frame delay can be expressed as a one-way delay for a frame, where one-way frame delay is defined as the time elapsed from the start of transmission of the first bit of the frame at a source reference point until the reception of the last bit of the same frame at the destination reference point. When a two-way delay is measured, a loopback is performed at the frame's destination reference point and the frame is received at the original source reference point. In the round-trip case, there are four timestamps available which enable both one-way and two-way delay calculations. Ideally, the mean one-way frame delay should be accessible for a set of frames. Mean one-way frame delay is defined in [ITU-T Y.1563]. The service frames belong to the same CoS instance on a point-to-point ETH connection or multipoint ETH connectivity.

- **Frame delay variation**

Frame delay variation is a measure of the variations in the frame delay between a pair of service frames. The service frames belong to the same CoS instance on a point-to-point ETH connection or multipoint ETH connectivity.

- **Availability**

The Ethernet service definition is defined in [ITU-T Y.1563]. Although the mechanisms defined in this Recommendation can contribute to availability-related measurements, the details of measurement methods in this Recommendation are for further study.

Frame performance parameters are applicable to frames that conform to an agreed-upon priority level X and that are deemed by the network as not drop-eligible (i.e., so called "green" frames) of bandwidth profile conformance. Such "green" frames are also called in-profile (see [ITU-T G.8021]).

Service frames are admitted at the ingress ETH flow point of a point-to-point ETH network, tandem or link connection and should be delivered to the egress ETH flow point.

In addition, another performance parameter is identified as per [b-IETF RFC 2544]:

– **Throughput**

Throughput is the average rate of successful traffic delivery over a communication channel. This is typically measured under test conditions, i.e., an out-of-service test, where there is no service traffic for the Ethernet service under test. [ITU-T Y.1564] defines a methodology to test Ethernet-based services at the service activation stage, using an out-of-service test. The Recommendation describes service configuration tests to verify bandwidth profiles and other Ethernet service attributes. Procedures used for out-of-service testing other than for Ethernet service activation can also be found in [b-IETF RFC 2544]. The procedure for in-service testing is for further study.

8.1 Frame loss measurement (ETH-LM)

Ethernet loss measurement function (ETH-LM) is used to collect counter values applicable for ingress and egress service frames where the counters maintain a count of transmitted and received data frames between a pair of MEPs.

ETH-LM is performed by sending frames with ETH-LM information to a peer MEP and similarly receiving frames with ETH-LM information from the peer MEP. Each MEP performs frame loss measurements which contribute to unavailable time. Since a bidirectional service is defined as unavailable if either of the two directions is declared unavailable, ETH-LM must facilitate each MEP to perform near-end and far-end frame loss measurements.

For a MEP, near-end frame loss refers to frame loss associated with ingress data frames while far-end frame loss refers to frame loss associated with egress data frames. Both near-end and far-end frame loss measurements contribute to near-end severely errored seconds (near-end SES) and far-end severely errored seconds (far-end SES), respectively, which together contribute to unavailable time, in a manner similar to [ITU-T G.826] and [ITU-T G.7710].

A MEP maintains the following two local counters for each peer MEP and for each priority class being monitored in a point-to-point ME for which loss measurements are to be performed:

- TxFCI: counter for in-profile data frames transmitted towards the peer MEP.
- RxFCI: counter for in-profile data frames received from the peer MEP.

TxFCI and RxFCI counters do not count the OAM frames transmitted or received by the MEP at the MEP's MEG level in some conditions (see Notes). However, the counters do count OAM frames from the higher MEG levels that pass through the MEPs in a manner similar to the data frames.

NOTE 1 – Both proactive and on-demand ETH-LM count OAM frames as follows:

For single-ended ETH-LM, OAM frames that are only used for proactive functions used by termination functions (e.g., those for ETH-CC) are counted.

For dual-ended ETH-LM, OAM frames for proactive functions used by termination functions are NOT counted.

In both cases:

Proactive OAM frames used by adaptation functions (e.g., those for ETH-APS and ETH-CSF) are counted.

OAM frames that can be used for on-demand functions (e.g., those for ETH-LB, ETH-LT and on-demand ETH-LM, ETH-DM and ETH-SLM) are NOT counted.

NOTE 2 – As OAM frames for ETH-AIS and ETH-LCK are only sent in the defect conditions where the result of loss measurements is invalid, it is unnecessary to count these frames.

The method of loss measurement involving pairs of consecutive frames with ETH-LM information, as shown in clauses 8.1.1.2 and 8.1.2.3, alleviates lack of synchronization across the initial counter values at the initiating and receiving MEPs. Furthermore, when a MEP detects a loss-of-continuity defect condition, it ignores loss measurements during the defect condition and assumes 100% losses.

NOTE 3 – The level of accuracy in the loss measurements is dependent on how frames with ETH-LM information are added to the data stream after the counter values are copied in the ETH-LM information. For example, if additional data frames get transmitted and/or received between the time of reading the counter values and adding the frame with ETH-LM information to the data stream, the counter values copied in ETH-LM information become inaccurate. However, a hardware-based implementation which is able to add frames with ETH-LM information to the data stream immediately after reading the counter values, provides enhanced accuracy.

NOTE 4 – Details on the processing of counters used for the transmitted and received data frames are described in [ITU-T G.8021].

NOTE 5 – In profile frames are so called 'green' frames where drop-eligibility is 'false'. Network operators or administrators can configure the encoding method to identify green frames. For example, green frames are those where DEI field is 'false', and yellow frames are those where such field is 'true'. PCP or PCP/DEI can be used for this identification.

Specific configuration information required by a MEP to support ETH-LM is as follows:

- MEG level – MEG level at which the MEP exists.
- Unicast MAC address of a peer MEP to which ETH-LM is intended. Multicast Class 1 MAC address is also allowed.
- ETH-LM transmission period – The default transmission period is 100 ms (i.e., a transmission rate of 10 frames/second). The ETH-LM transmission period should be such that the frame and/or octet counters whose values are carried in ETH-LM information should not wrap around to the same value even if one or more ETH-LM frames are lost. This is primarily a concern for frame loss measurements at lower priority levels. Refer to clause II.2 for examples of frame counter wrapping periods.
- Priority – identifies the priority of the frames with ETH-LM information. This information is configurable per operation.
- Drop eligibility – Frames with ETH-LM information are always marked as drop ineligible. This information is not necessarily configured.

A MIP is transparent to frames with ETH-LM information and therefore does not require any information to support ETH-LM functionality.

ETH-LM can be performed in two ways:

- dual-ended ETH-LM (see clause 8.1.1);
- single-ended ETH-LM (see clause 8.1.2).

8.1.1 Dual-ended ETH-LM

Dual-ended ETH-LM is used as proactive OAM for performance monitoring and is applicable to fault management. In this case, each MEP sends periodic dual-ended frames with ETH-LM information to its peer MEP in a point-to-point ME to facilitate frame loss measurements at the peer MEP. Each MEP terminates the dual-ended frames with ETH-LM information and makes the near-end and far-end loss measurements. This function is used for performance monitoring at the same priority level as used for ETH-CC.

The PDU used for dual-ended ETH-LM information is CCM, as described in clause 9.2.

8.1.1.1 CCM with dual-ended ETH-LM transmission

When configured for proactive loss measurement, a MEP periodically transmits CCM frames with the following information elements:

- TxFCf: TxFCf is the value of the local counter TxFCI at the time of transmission of the CCM frame.
- RxFCb: RxFCb is the value of the local counter RxFCI at the time of reception of the last CCM frame from the peer MEP.
- TxFCb: TxFCb is the value of TxFCf in the last received CCM frame from the peer MEP.

The CCM PDU is transmitted with a period value equal to the CCM transmission period configured for performance monitoring application at the transmitting MEP. The receiving MEP detects an unexpected period defect condition if the CCM transmission period is not the same as the configured value.

8.1.1.2 CCM with dual-ended ETH-LM frame reception

When configured for proactive loss measurement, a MEP, upon receiving a CCM frame, uses the following values to make near-end and far-end loss measurements:

- Received CCM frame's TxFCf, RxFCb, and TxFCb values and local counter RxFCI value at time this CCM frame was received. These values are represented as TxFCf[t_c], RxFCb[t_c], TxFCb[t_c], and RxFCI[t_c], where t_c is the reception time of the current frame.
 - Previous CCM frame's TxFCf, RxFCb, and TxFCb values and local counter RxFCI value at time the previous CCM frame was received. These values are represented as TxFCf[t_p], RxFCb[t_p], TxFCb[t_p], and RxFCI[t_p], where t_p is the reception time of the previous frame.
- $$\text{Frame Loss}_{\text{far-end}} = |\text{TxFCb}[t_c] - \text{TxFCb}[t_p]| - |\text{RxFCb}[t_c] - \text{RxFCb}[t_p]|$$
- $$\text{Frame Loss}_{\text{near-end}} = |\text{TxFCf}[t_c] - \text{TxFCf}[t_p]| - |\text{RxFCI}[t_c] - \text{RxFCI}[t_p]|$$

If the Period field value in the received CCM frame is different than the MEP's own configured CCM transmission period, the MEP detects an unexpected period defect condition.

8.1.2 Single-ended ETH-LM

Single-ended ETH-LM is used for on-demand and proactive OAM. In this case, a MEP sends frames with ETH-LM request information to its peer MEP and receives frames with ETH-LM reply information from its peer MEP to carry out loss measurements.

The PDU used for single-ended ETH-LM request is LMM, as described in clause 9.12. The PDU used for single-ended ETH-LM reply is LMR, as described in clause 9.13. Frames which carry the LMM PDU are called LMM frames. Frames which carry the LMR PDU are called LMR frames. The same LMM and LMR frame formats can be used for proactive and on-demand single-ended ETH-LM. The distinction of proactive LMM/LMR frames from on-demand LMM/LMR frames is by the value of a flag field in the LMM/LMR frames.

8.1.2.1 LMM transmission

When configured for single-ended loss measurement, a MEP periodically transmits LMM frames with the following information element:

- TxFCf: Value of the local counter TxFCI at the time of LMM frame transmission.

8.1.2.2 LMM reception and LMR transmission

Whenever a valid LMM frame is received by a MEP, an LMR frame is generated and transmitted to the initiating MEP. An LMM frame with a valid MEG level and a destination MAC address equal to the receiving MEP's MAC address is considered to be a valid LMM frame. An LMR frame contains the following values:

- TxFCf: Value of TxFCf copied from the LMM frame.
- RxFCf: Value of local counter RxFCI at the time of LMM frame reception.
- TxFCb: Value of local counter TxFCI at the time of LMR frame transmission.

8.1.2.3 LMR reception

Upon receiving an LMR frame, a MEP uses the following values to make near-end and far-end loss measurements:

- Received LMR frame's TxFCf, RxFCf, and TxFCb values and local counter RxFCI value at the time this LMR frame was received. These values are represented as TxFCf[t_c], RxFCf[t_c], TxFCb[t_c], and RxFCI[t_c], where t_c is the reception time of the current reply frame.
- Previous LMR frame's TxFCf, RxFCf, and TxFCb values and local counter RxFCI value at the time the previous LMR frame was received. These values are represented as TxFCf[t_p], RxFCf[t_p], TxFCb[t_p], and RxFCI[t_p], where t_p is the reception time of the previous reply frame.

$$\text{Frame Loss}_{\text{far-end}} = |\text{TxFCf}[t_c] - \text{TxFCf}[t_p]| - |\text{RxFCf}[t_c] - \text{RxFCf}[t_p]|$$

$$\text{Frame Loss}_{\text{near-end}} = |\text{TxFCb}[t_c] - \text{TxFCb}[t_p]| - |\text{RxFCI}[t_c] - \text{RxFCI}[t_p]|$$

8.2 Frame delay measurement (ETH-DM)

Frame delay measurement (ETH-DM) can be used for on-demand or proactive OAM to measure frame delay and frame delay variation. Frame delay and frame delay variation measurements are performed by sending periodic frames with ETH-DM information to the peer MEP and receiving frames with ETH-DM information from the peer MEP during a proactive measurement session and/or diagnostic interval. Each MEP may perform frame delay and frame delay variation measurement.

When a MEP is enabled to generate frames with ETH-DM information, it periodically sends frames with ETH-DM information to its peer MEP in the same ME. When a MEP is enabled to generate frames with ETH-DM information, it also expects to receive frames with ETH-DM information from its peer MEP in the same ME.

Specific configuration information required by a MEP to support ETH-DM is as follows:

- MEG level – MEG level at which the MEP exists.
- Unicast MAC address of a peer MEP to which ETH-DM is intended. Multicast MAC address is also allowed for multipoint ETH connectivity. In case of a multipoint ETH connectivity, a MEP may activate multiple monitoring towards different peer MEPs simultaneously. In this case, each MEP needs to manage the result of monitoring on a per peer-MEP basis.
- DM application – identifies the application, i.e., proactive versus on-demand delay measurement. This information is configurable per operation. A MEP may activate proactive and on-demand monitoring simultaneously on the same CoS level and towards the same peer MEP. In this case, each MEP needs to manage the result of monitoring on a per peer-MEP basis.
- Data – optional data element whose length is configurable at the MEP. The inclusion of the optional data element in the DM frame is to support configurable DM frame size.
- Priority – identifies the priority of the frames with ETH-DM information. This information is configurable per operation. A MEP may activate the multiple monitoring on different CoS levels simultaneously. In this case, each MEP needs to manage the result of monitoring per CoS level.
- Drop eligibility – Frames with ETH-DM information are always marked as drop ineligible. This information is not necessarily configured.
- Test ID – can optionally be used to distinguish each DM measurement if multiple measurements are simultaneously activated. It must be unique at least within the context of any DM measurement type (single-ended/dual-ended and on-demand/proactive) for the MEG and initiating MEP.

NOTE 1 – Additional configuration information elements may be needed, such as the transmission rate of ETH-DM information, the total interval of ETH-DM, etc. These additional configuration information elements are outside the scope of this Recommendation.

A MIP is transparent to the frames with ETH-DM information and therefore does not require any information to support ETH-DM functionality.

A MEP transmits frames with ETH-DM information with the following information element:

- TxTimeStampf: Timestamp at the transmission time of ETH-DM frame

The receiving MEP can compare this value with the RxTimeef, the time at the reception of ETH-DM frame and calculate the one-way frame delay as:

$$\text{Frame Delay} = \text{RxTimeef} - \text{TxTimeStampf}$$

However, one-way frame delay measurement requires that the time and phase at the initiating MEP and the receiving MEPs are synchronized. For the purposes of frame delay variation measurement, which is based on the difference between subsequent frame delay measurements, the requirement for the time and phase synchronizations can be relaxed since the out-of-phase period can be eliminated in the difference of subsequent frame delay measurements.

If it is not practical for the clocks to be synchronized, which is expected to be the most common scenario, the frame delay measurement can be made only for two-way measurements, where the MEP transmits a frame with ETH-DM request information with the TxTimeStampf, and the receiving MEP responds with a frame with ETH-DM reply information with TxTimeStampf copied from the ETH-DM request information. The MEP receiving the frame with ETH-DM reply information compares the TxTimeStampf with the RxTimeb, which is the time at the reception of frame with ETH-DM reply information and calculates the two-way frame delay as:

$$\text{Frame Delay} = \text{RxTimeb} - \text{TxTimeStampf}$$

The MEP can also make two-way frame delay variation measurements based on its ability to calculate the difference between two subsequent two-way frame delay measurements.

NOTE 2 – To allow a more precise two-way frame delay measurement, the MEP replying to frame with ETH-DM request information can also include two additional timestamps in the ETH-DM reply information: RxTimeStampf (Timestamp at the time of receiving frame with ETH-DM request information), and TxTimeStampb (Timestamp at the time of transmitting frame with ETH-DM reply information).

ETH-DM can be performed in two ways:

- Dual-ended ETH-DM (see clause 8.2.1).
- Single-ended ETH-DM (see clause 8.2.2).

NOTE 3 – In previous revisions of this recommendation, dual-ended ETH-DM and single-ended ETH-DM were known as one-way ETH-DM and two-way ETH-DM respectively.

8.2.1 Dual-ended ETH-DM

In this case, each MEP sends frames with dual-ended ETH-DM information to its peer MEP to facilitate one-way frame delay and/or one-way frame delay variation measurements at the peer MEP.

NOTE – If the clocks between the two MEPs are synchronized, one-way frame delay measurement can be carried out; otherwise, only one-way frame delay variation measurement can be performed.

The PDU used for dual-ended ETH-DM is 1DM, as described in clause 9.14. Frames which carry the 1DM PDU are called as 1DM frames. The same 1DM frame format can be used for proactive and on-demand dual-ended ETH-DM. The distinction of a proactive 1DM frame from an on-demand 1DM frame is by the value of a flag field in the 1DM frame.

NOTE – In previous revisions of this recommendation, dual-ended ETH-DM was known as one-way ETH-DM.

8.2.1.1 1DM transmission

When configured for dual-ended ETH-DM, a MEP periodically transmits 1DM frames with the TxTimeStampf value. A MEP can optionally use a Test ID TLV and/or a Data TLV. The MEP uses a Test ID TLV, containing a Test ID that is used to run multiple tests simultaneously, when configured. The MEP uses a Data TLV when configured for measuring delay and delay variation for different frame sizes.

8.2.1.2 1DM reception

When configured for dual-ended ETH-DM, a MEP, upon receiving a valid 1DM frame, uses the following values to make one-way frame delay measurement. A 1DM frame with a valid MEG level and a destination MAC address equal to the receiving MEP's MAC address or multicast Class 1 MAC Address is considered to be a valid 1DM frame. These values serve as input to the one-way frame delay variation measurement:

- 1DM frame's TxTimeStampf value
 - RxTimeef, which is the time at reception of the 1DM frame
- $$\text{Frame Delay}_{\text{one-way}} = \text{RxTimeef} - \text{TxTimeStampf}$$

8.2.2 Single-ended ETH-DM

A MEP sends frames with ETH-DM request information to its peer MEP and receives frames with ETH-DM reply information from its peer MEP to carry out two-way frame delay and two-way frame delay variation measurements. If two optional timestamps of RxTimeStampf and TxTimeStampb are supported on its peer MEP, the results of one-way frame delay and one-way frame delay variation measurements can be also calculated by the same ETH-DM request/reply information.

NOTE – Regarding the one-way measurements, if the clocks between the two MEPs are synchronized, one-way frame delay measurement can be carried out. Otherwise, only one-way frame delay variation measurement can be performed.

The PDU used for ETH-DM request is DMM, as described in clause 9.15. The PDU used for ETH-DM reply is DMR, as described in clause 9.16. Frames which carry the DMM PDU are called as DMM frames. Frames which carry the DMR PDU are called as DMR frames. The same DMM and DMR frame formats can be used for proactive and on-demand single-ended ETH-DM. The distinction of proactive DMM/DMR frames from on-demand DMM/DMR frames is by the value of a flag field in the DMM/DMR frames.

NOTE – In previous revisions of this recommendation, single-ended ETH-DM was known as two-way ETH-DM.

8.2.2.1 DMM transmission

When configured for single-ended ETH-DM, a MEP periodically transmits DMM frames with the TxTimeStampf value. A MEP can optionally use a Test ID TLV and/or a Data TLV. The MEP uses a Test ID TLV, containing a Test ID that is used to run multiple tests simultaneously, when configured. The MEP uses a Data TLV when configured for measuring delay and delay variation for different frame sizes.

8.2.2.2 DMM reception and DMR transmission

Whenever a valid DMM frame is received by a MEP, a DMR frame is generated and transmitted to the initiating MEP. A DMM frame with a valid MEG level and a destination MAC address equal to the responding MEP's MAC or multicast Class 1 MAC Address is considered to be a valid DMM frame. Every field in the DMM frame is copied to the DMR frame with the following exceptions:

- The source MAC address is copied to the destination MAC address and the source MAC address is filled with the MEP MAC address.
- The OpCode field is changed from DMM to DMR.

NOTE – As an option, two additional timestamps may be used in the DMR frame to take into account the processing time at the responding MEP: RxTimeStampf (timestamp at the time of receiving the DMM frame) and TxTimeStampb (timestamp at the time of transmitting the DMR frame).

8.2.2.3 DMR reception

Upon receiving a DMR frame, a MEP uses the following values to calculate two-way frame delay. This value serves as input for two-way frame delay variation measurement:

- DMR frame's TxTimeStampf value
 - RxTimeb – reception time of the DMR frame
- $$\text{Frame Delay}_{\text{two-way}} = \text{RxTimeb} - \text{TxTimeStampf}$$

If the additional timestamps are carried in the DMR frame, which is determined by non-zero values of the RxTimeStampf and TxTimeStampb fields, the frame delay for one-way and two-way can be calculated to be:

$$\text{Frame Delay}_{\text{two-way}} = (\text{RxTimeb} - \text{TxTimeStampf}) - (\text{TxTimeStampb} - \text{RxTimeStampf})$$

$$\text{Frame Delay}_{\text{one-way_far}} = \text{RxTimeStampf} - \text{TxTimeStampf}$$

$$\text{Frame Delay}_{\text{one-way_near}} = \text{RxTimeb} - \text{TxTimeStampb}$$

8.3 Throughput measurement

[b-IETF RFC 2544] specifies measuring the throughput by sending frames at an increasing rate (up to the theoretical maximum), graphing the percentage of frames received and reporting the rate at which frames start being dropped. In general this rate is dependent on the frame size.

The mechanisms specified in this Recommendation, e.g., unicast ETH-LB (e.g., LBM and LBR frames with the data field) and ETH-Test (e.g., TST frames with the data field) can be used for performing the throughput measurements. A MEP can insert TST frames or LBM frames with configured size, pattern, etc., at a rate to exercise the throughput and make one-way or two-way measurements.

8.4 Synthetic loss measurement (ETH-SLM)

Synthetic loss measurement is a mechanism to measure frame loss using synthetic frames, rather than data traffic. A number of synthetic frames are sent and received, and the number of those that are lost is hence calculated. This can be treated as a statistical sample, and used to approximate the frame loss ratio of data traffic.

ETH-SLM collects counters so as to maintain a count of transmitted and received synthetic frames between a set of MEPs.

ETH-SLM is used to perform on-demand or proactive tests by sending a finite number of frames with ETH-SLM information to one or multiple peer MEPs and similarly receiving frames with ETH-SLM information from the peer MEPs. Each MEP then performs frame loss measurements which contribute to unavailable time. Since a bidirectional service is defined as unavailable if either of the two directions is declared unavailable, ETH-SLM must facilitate each MEP to perform near-end and far-end synthetic frame loss measurements.

A MEP maintains the following local counters for each Test ID and for each peer MEP being monitored in a ME for which loss measurements are to be performed:

- TxFCI: number of synthetic frames transmitted towards the peer MEP, part of a given Test ID. An initiating MEP increments this number for successive transmission of synthetic frames with ETH-SLM request information while a responding MEP increments it for successive transmission of synthetic frames with ETH-SLM reply information.
- RxFCI: number of synthetic frames received from the peer MEP and part of a given Test ID. An initiating MEP increments this number for successive reception of synthetic frames with

ETH-SLM reply information while a responding MEP increments it for successive reception of synthetic frames with ETH-SLM request information.

The method of loss measurement involves series of frames with increasing values of TxFCI with ETH-SLM information, as shown in clauses 8.4.1 and 8.4.2.

NOTE 1 – No synchronization is required of the Test ID value between initiating and responding MEPs, as the Test ID is configured at the initiating MEP, and the responding MEP uses the Test ID it receives from the initiating MEP. The allocation and release of local counter resources for each Test ID at the responding MEP is outside the scope of this Recommendation.

The specific configuration information required by a MEP to support ETH-SLM is as follows:

- MEG level – MEG level at which the MEP exists.
- Data – optional data element whose length is configurable at the MEP. The inclusion of the optional data element in the SLM frame is to support configurable SLM frame size.
- Destination MAC address – identifies the target peer MEP.
- Test ID – used to distinguish each SL measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SL measurement for the MEG and initiating MEP.
- Priority – identifies the priority of the frames with ETH-SLM information. This information is configurable per operation.
- Drop eligibility – Frames with ETH-SLM information are always marked as drop ineligible. This information is not necessarily configured.

A MIP is transparent to frames with ETH-SLM information and therefore does not require any information to support ETH-SLM functionality.

NOTE 2 – As ETH-SLM is a sampling technique, it is inevitably less accurate than counting the service frames. Furthermore, the accuracy depends on the number of SLM frames used or the period for transmitting SLM frames. The number of SLM frames or the period for SLM frames is outside the scope of this Recommendation, but some examples of accuracy are included for information in Appendix VI.

8.4.1 Single-ended ETH-SLM

Single-ended ETH-SLM is used for proactive or on-demand OAM. It carries out synthetic loss measurements applicable to both point-to-point ETH connection and multipoint ETH connectivity. It allows a MEP to initiate and report far-end and near-end loss measurements associated with one or a set of peer MEPs part of the same MEG.

The selection of on-demand or proactive is performed by the management function that initiates the test, however this is local information and does not need to be conveyed in the PDU.

For single-ended operation, a MEP sends frames with ETH-SLM request information to its peer MEP(s) and receives frames with ETH-SLM reply information from its peer MEP(s) to carry out synthetic loss measurements.

The PDU used for single-ended ETH-SLM request is SLM, as described in clause 9.22. The PDU used for single-ended ETH-SLM reply is SLR, as described in clause 9.23. Frames which carry the SLM PDU are called SLM frames. Frames which carry the SLR PDU are called SLR frames.

8.4.1.1 SLM transmission

A MEP periodically transmits SLM frames with the following information elements included:

- Test ID: Test ID is a value containing a number configured by the MEP and then used to run multiple tests simultaneously.
- Source MEP ID: Source MEP ID is the MEP's own identity in the MEG.
- TxFCf: TxFCf is the value of the local counter TxFCI at the time of SLM frame transmission.

- TxFCb: TxFCb is set always to zero. Reserved for SLR transmission.

8.4.1.2 SLM reception and SLR transmission

Whenever a valid SLM frame is received by a MEP, an SLR frame is generated and transmitted to the initiating MEP. An SLM frame with a valid MEG level and a destination MAC address equal to the responding MEP's MAC address or a multicast Class 1 MAC address is considered to be a valid SLM frame. Every field in the SLM frame is copied to the SLR frame with the following exceptions:

- The source MAC address is copied to the destination MAC address and the source MAC address is filled with the MEP MAC address.
- The OpCode is changed from SLM to SLR.
- Responder MEP ID: MEP's own identity in the MEG.
- TxFCb: value of the local counter RxFCI at the time of SLR frame transmission.

Note that as an SLR frame is generated every time an SLM frame is received, RxFCI in the responder is equal to the number of SLM frames received and also equal to the number of SLR frames sent. In other words, in the responder, RxFCI = TxFCI.

8.4.1.3 SLR reception

After transmission of a SLM frame (with a given TxFCf value), a MEP will expect to receive a corresponding SLR frame (carrying same TxTCf value) from its peer MEP(s). In on-demand mode, SLR frames received more than 5s after the command that terminates SL measurement must be discarded, as specified in [ITU-T G.8021].

With the information contained in SLR frames, a MEP determines frame loss for given measurement periods. The measurement period is a time interval during which the number of SLM frames transmitted is statistically adequate to make a measurement at a given accuracy. (See Appendix VI.) A MEP uses the following values to determine near-end and far-end frame loss in the measurement period:

- Last received SLR frame's TxFCf and TxFCb values and local counter RxFCI at the end of the measurement period. These values are represented as TxFCf[t_c], TxFCb[t_c] and RxFCI[t_c], where t_c is the end time of the measurement period.
- SLR frame's TxFCf and TxFCb values of the first received SLR frame after the test starts and local counter RxFCI at the beginning of the measurement period. These values are represented as TxFCf[t_p], TxFCb[t_p] and RxFCI[t_p], where t_p is the start time of the measurement period.

$$\text{Frame loss}_{\text{far-end}} = | \text{TxFCf}[t_c] - \text{TxFCf}[t_p] | - | \text{TxFCb}[t_c] - \text{TxFCb}[t_p] |$$

$$\text{Frame loss}_{\text{near-end}} = | \text{TxFCb}[t_c] - \text{TxFCb}[t_p] | - | \text{RxFCI}[t_c] - \text{RxFCI}[t_p] |$$

NOTE – If there are SLMs at the end of the measurement period for which no corresponding SLRs have been received within the timeout period (i.e., SLMs with sequence numbers after the sequence number of the last SLR received), it is not possible to determine whether they were lost in the near-end or far-end direction.

8.4.2 Dual-ended ETH-SLM

Dual-ended ETH-SLM can be used for on-demand and proactive OAM. It carries out loss measurements applicable to both point-to-point ETH connection or multipoint ETH connectivity. It allows a MEP in a MEG to send periodic dual-ended frames with ETH-SLM information to its peer MEP(s) to facilitate frame loss measurement at the peer MEP. The receiving MEP terminates the dual-ended frames and makes the near-end loss measurements.

The selection of on-demand or proactive is performed by the management function that initiates the test; however this is local information and does not need to be conveyed in the PDU.

Dual-ended ETH-SLM is suitable where it is required and practical to measure unidirectional FLR from every MEP to all its peer MEPs (e.g., any-to-any measurements).

The PDU used for dual-ended ETH-SLM information is 1SL, as described in clause 9.24. Frames which carry the 1SL PDU are called 1SL frames.

8.4.2.1 1SL transmission

When configured for dual-ended operation, a MEP periodically transmits 1SL frames with the following information elements:

- Test ID: Test ID is the value containing a number configured by the MEP and then used to run multiple tests simultaneously.
- Source MEP ID: Source MEP ID is the MEP's own identity in the MEG.
- TxFCf: TxFCf is the value of the local counter TxFCI at the time of 1SL frame transmission.

The 1SL PDU is transmitted with a period value equal to the 1SL transmission period configured for performance monitoring application at the transmitting MEP.

8.4.2.2 1SL reception

When configured for one-way synthetic loss measurements, a MEP, upon receiving a valid 1SL frame, uses the following values to make one-way frame loss measurement. A 1SL frame with a valid MEG level and a destination MAC address equal to the receiving MEP's MAC address or multicast Class 1 MAC address is considered to be a valid 1SL frame.

Whenever a valid 1SL frame is received by a MEP with a given **TxFCf** value, the MEP will expect to receive a subsequent 1SL frame (**TxFCf** value incremented by one).

For a given measurement period, a MEP uses the following values to determine near-end frame loss in the period:

- Last received 1SL frame's **TxFCf** value and local counter **RxFCI** at the end of the measurement period. These values are represented as **TxFCf[t_c]** and **RxFCI[t_c]**, where t_c is the end time of the measurement period.
- 1SL frame's **TxFCf** value of the first received 1SL after the test starts and local counter **RxFCI** at the beginning of the measurement period. These values are represented as **TxFCf[t_p]** and **RxFCI[t_p]**, where t_p is the start time of the measurement period.

$$\text{Frame loss}_{\text{near-end}} = | \text{TxFCf}[t_c] - \text{TxFCf}[t_p] | - | \text{RxFCI}[t_c] - \text{RxFCI}[t_p] |$$

9 OAM PDU types

This clause describes the information elements and formats for different OAM PDU types used to meet the requirements of OAM functions described in clauses 7 and 8.

NOTE – When the values of OAM PDU fields are fixed, they are shown in parentheses in the OAM PDU formats in the following clauses.

9.1 Common OAM information elements

Some information elements are common across the OAM PDUs that are identified in this Recommendation. These information elements are:

- MEG Level: MEG Level is a 3-bit field. It contains an integer value that identifies the MEG level of OAM PDU. Value ranges from 0 to 7.
- Version: Version is a 5-bit field. It contains an integer value that identifies the OAM protocol version. Clause 11 discusses the specifics of OAM PDU validation and versioning with respect to this field.

- **OpCode:** OpCode is a 1-octet field. It contains an OpCode that identifies an OAM PDU type. OpCode is used to identify the remaining content of an OAM PDU. The values of this information field are shown in Table 9-1.
- **Flags:** Flags is an 8-bit field. Use of the bits in this field is dependent on the OAM PDU type.
- **TLV Offset:** TLV Offset is a 1-octet field. It contains the offset to the first TLV in an OAM PDU relative to the TLV Offset field. The value of this field is associated with an OAM PDU type. When the TLV Offset is 0, it points to the first octet following the TLV Offset field.

Other information elements which are not present in OAM PDUs but are conveyed in frames carrying OAM PDUs include:

- **Priority:** Priority identifies the priority of a specific OAM frame.
- **Drop eligibility:** Drop eligibility identifies the drop eligibility of a specific OAM frame.

Table 9-1 – OpCode values

OpCode value	OAM PDU type	OpCode relevance for MEPs/MIPs
OpCodes common with IEEE 802.1		
1	CCM	MEPs
3	LBM	MEPs and MIPs (connectivity verification)
2	LBR	MEPs and MIPs (connectivity verification)
5	LTM	MEPs and MIPs
4	LTR	MEPs and MIPs
0, 6-31, 64-255	Reserved (Note 1)	
OpCodes specific to this Recommendation		
32	GNM (Note 4)	MEPs
33	AIS	MEPs
35	LCK	MEPs
37	TST	MEPs
39	Linear APS	Refer to [ITU-T G.8031]
40	Ring APS	Refer to [ITU-T G.8032]
41	MCC	MEPs
43	LMM	MEPs
42	LMR	MEPs
45	1DM	MEPs
47	DMM	MEPs
46	DMR	MEPs
49	EXM	Outside the scope of this Recommendation
48	EXR	Outside the scope of this Recommendation
51	VSM	Outside the scope of this Recommendation
50	VSR	Outside the scope of this Recommendation
52	CSF	MEPs
53	1SL	MEPs
55	SLM	MEPs
54	SLR	MEPs

Table 9-1 – OpCode values

OpCode value	OAM PDU type	OpCode relevance for MEPs/MIPs
34, 36, 38, 44, 60-63	Reserved (Note 2)	
56-59	Reserved (Note 3)	
NOTE 1 – Reserved for definition by IEEE 802.1.		
NOTE 2 – Reserved for future standardization by ITU-T.		
NOTE 3 – Reserved for definition by MEF. The definition is outside the scope of this Recommendation.		
NOTE 4 – The Generic Notification Message (GNM) PDU type is used to carry other OAM PDUs using the Sub-Opcodes in Table 9-1a		

9.1.1 Common OAM PDU format

The common format used in all OAM PDUs is shown in Figure 9.1-1.




Figure 9.1-1 – Common OAM PDU format

When OpCode 32 (GNM) is used, there is an additional one-octet Sub-OpCode field following the TLV Offset field. Sub-OpCode values are shown in Table 9-1a.

Table 9-1a – Sub-OpCode values

Sub-OpCode value	OAM PDU type
1	BNM
0, 2-255	Reserved (Note)
NOTE – Reserved for future standardization by ITU-T.	

The general format of TLVs is shown in Figure 9.1-2. Type values are specified in Table 9-2.




Figure 9.1-2 – Generic TLV format

NOTE – In an End TLV, Type = 0, and both Length and Value fields are not used.

Table 9-2 – Type values

Type value	TLV name
Types common with IEEE 802.1	
0	End TLV
3	Data TLV
5	Reply ingress TLV
6	Reply egress TLV
7	LTM egress identifier TLV
8	LTR egress identifier TLV
2, 4, 9-31, 64-255	Reserved (Note 1)
Types specific to this Recommendation	
32	Test TLV
33-35	Reserved (Note 2)
36	Test ID TLV
37, 38	Reserved (Note 3)
39-63	Reserved (Note 4)
NOTE 1 – Reserved for definition by IEEE 802.1. NOTE 2 – Reserved for definition by [ITU-T G.8113.1]. NOTE 3 – Reserved for definition by MEF. The definition is outside the scope of this Recommendation. NOTE 4 – Reserved for future standardization by ITU-T.	

9.2 CCM PDU

CCM is used to support ETH-CC function, as described in clause 7.1, ETH-RDI function, as described in clause 7.5, and dual-ended ETH-LM function, as described in clause 8.1.1.

9.2.1 CCM information elements

Information elements carried in CCM to support ETH-CC are:

- Period: Period is a 3-bit information element carried in the three least significant bits of the Flags field. Period contains the value of the CCM transmission period configured at the CCM source. CCM Period values are specified in Table 9-3.
- MEG ID: MEG ID is a 48-octet field which contains the MEG ID of the MEG to which the MEP transmitting the CCM frame belongs. See Annex A.
- MEP ID: MEP ID is a 2-octet field where the 13 least significant bits are used to identify the MEP transmitting the CCM frame. MEP ID is unique within the MEG.

Information element carried in CCM to support ETH-RDI is:

- RDI: RDI is a 1-bit information element carried in most significant bit of Flags field. When the RDI bit is 1, detection of a defect is indicated by the transmitting MEP. When the RDI bit is 0, no defect indication is communicated by the transmitting MEP.

The information elements carried in CCM to support dual-ended ETH-LM are:

- TxFCf: TxFCf is a 4-octet field which carries the value of the counter of in-profile data frames transmitted by the MEP towards its peer MEP, at the time of CCM frame transmission.
- RxFCb: RxFCb is a 4-octet field which carries the value of the counter of in-profile data frames received by the MEP from its peer MEP, at the time of receiving the last CCM frame from that peer MEP.

- TxFCb: TxFCb is a 4-octet field which carries the value of the TxFCf field in the last CCM frame received by the MEP from its peer MEP.

9.2.2 CCM PDU format

The CCM PDU format used by a MEP to transmit CCM information is shown in Figure 9.2-1.

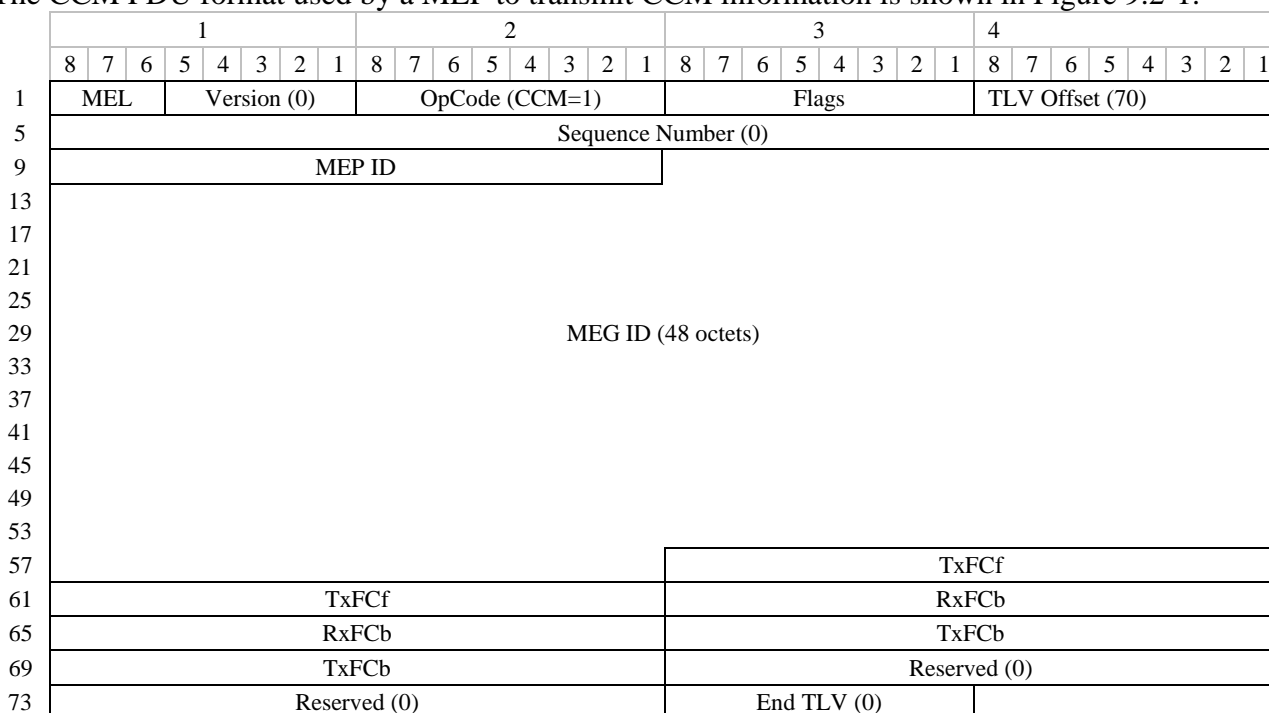


Figure 9.2-1 – CCM PDU format

The fields of the CCM PDU format are as follows:

- MEG Level: Refer to clause 9.1.
- Version: Refer to clause 9.1; value is 0 in the current version of this Recommendation.
- OpCode: value for this PDU type is CCM (1).
- Flags: two information elements in the Flags field for CCM PDU: RDI, and Period as follows:




Figure 9.2-2 – Flags format in CCM PDU

- RDI: Bit 8 is set to 1 to indicate RDI, otherwise it is set to 0.
- Period: Bits 3 to 1 indicate the transmission period with the encoding shown in Table 9-3.

Table 9-3 – CCM period values

Flags[3:1]	Period value	Comments
000	Invalid value	Invalid value for CCM PDUs
001	3.33 ms	300 frames per second
010	10 ms	100 frames per second
011	100 ms	10 frames per second
100	1 s	1 frame per second
101	10 s	6 frames per minute
110	1 min	1 frame per minute
111	10 min	6 frame per hour

- TLV Offset: Set to 70.
- Sequence Number: This field is set to all-ZEROes for this Recommendation.
- MEP ID: a 13-bit integer value identifying the transmitting MEP within the MEG. The three MSBs of the first octet are not used and set to ZERO:

MSB								LSB							
octet 9								octet 10							
8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
0	0	0	MEP ID												

Figure 9.2-3 – MEP ID format in CCM PDU

- MEG ID: 48-octet field. Refer to Annex A for the format used for the MEG ID field.
- TxFCf, TxFCb, RxFCb: 4-octet integer values with samples of the wrap-around frame counters, as specified in clause 9.2.1. These fields are set to all-ZEROes when not used.
- Reserved: Reserved fields are set to all ZEROes.
- End TLV: an all-ZEROes octet value.

9.3 LBM PDU

LBM is used to support ETH-LB request, as described in clause 7.2.

9.3.1 LBM information elements

Information elements carried in LBM include:

- Transaction ID/Sequence Number: Transaction ID/Sequence Number is a 4-octet field that contains the transaction ID/sequence number for the LBM. The receiver is expected to copy the Transaction ID /Sequence Number in the LBR PDU, as described in clause 9.4.
- Data/Test Pattern: Data is an optional field whose length and contents are determined at the transmitting MEP. The contents of Data field can be a test pattern with an additional, optional checksum. The test pattern can be a pseudo-random bit sequence (PRBS) ($2^{31}-1$) as specified in clause 5.8 of [ITU-T O.150], an all '0' pattern, etc.

9.3.2 LBM PDU format

The LBM PDU format used by a MEP to transmit LBM information is shown in Figure 9.3-1.

	1							2							3							4										
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	MEL			Version (0)				OpCode (LBM=3)							Flags (0)							TLV Offset (4)										
5	Transaction ID/Sequence Number																															
9	<i>[Optional TLV starts here; otherwise End TLV]</i>																															
13																																
17																																
:																																
last																									End TLV (0)							

Figure 9.3-1 – LBM PDU format

The fields of the LBM PDU format are as follows:

- MEG Level: Refer to clause 9.1.
- Version: Refer to clause 9.1; value is 0 in the current version of this Recommendation.
- OpCode: value for this PDU type is LBM (3).
- Flags: Set to all-ZEROes.

MSB				LSB			
8	7	6	5	4	3	2	1
Reserved (0)							

Figure 9.3-2 – Flags format in LBM PDU

- TLV Offset: Set to 4.
- Transaction ID/Sequence Number: A 4-octet value containing either the transaction number for the LBM PDU without test pattern or a Sequence number incremented for successive LBM PDUs with a test pattern.
- Optional TLV: A Data TLV or Test TLV as specified in Figure 9.3-3 or Figure 9.3-4 respectively may be included in any LBM transmitted.
- End TLV: all-ZEROes octet value.

	1								2								3								4							
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	Type (3)								Length																Data Pattern							
:																																
:																																
:																																
:																																

Figure 9.3-3 – Data TLV format

The fields of the Data TLV format are as follows:

- Type: identifies TLV type; value for this TLV type is Data (3).
- Length: identifies size, in octets, of the Value field containing the Data Pattern. In a frame where the PDU is limited to 1492 octets, the maximum length value is 1480 (since 12 octets are required for 8 octets of LBM PDU overhead, 3 octets of Data TLV overhead, and 1 octet of End TLV). Any other TLVs, if present in LBM, will furthermore detract from the maximum length value of 1480.
- Data Pattern: an n-octet (n = Length) arbitrary bit pattern. The receiver should ignore it.

1 : : : : :	1								2								3								4							
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
	Type (32)								Length																Pattern Type							
	Test Pattern (NULL, PRBS)																															
	CRC-32 (optional)																															

Figure 9.3-4 – Test TLV format

The fields of the Test TLV format are as follows:

- Type: identifies TLV type; value for this TLV type is Test (32).
- Length: identifies size, in octets, of the Value field containing the pattern type, test pattern and CRC-32. In a frame where the PDU is limited to 1492 octets, the maximum length value is 1480 octets (since 12 octets are required for 8 octets of LBM PDU overhead, 3 octets of Test TLV overhead, and 1 octet of End TLV). Any other TLVs, if present in LBM, will furthermore detract from the maximum length value of 1480. (As one octet is used for pattern type, 1479 octets are available for the test pattern.)
- Pattern Type: identifies test pattern type; values are:
 - 0 'Null signal without CRC-32'
 - 1 'Null signal with CRC-32'
 - 2 'PRBS $2^{-31} - 1$ without CRC-32'
 - 3 'PRBS $2^{-31} - 1$ with CRC-32'
 - 4-255 Reserved for future standardization
- Test Pattern: an n-octet ($n \leq \text{Length}$) test pattern: PRBS $2^{-31} - 1$ or Null (all-zeroes) pattern.
- CRC-32: covers all fields (from Type to last octet before CRC-32).

9.4 LBR PDU

LBR is used to support ETH-LB reply, as described in clause 7.2.

9.4.1 LBR information element

The information elements carried in LBR include:

- Transaction ID/Sequence Number: Transaction ID/Sequence Number is a 4-octet field that is copied from the Transaction ID/Sequence Number field in LBM.
- Data: Data is a field that is copied from the Data field in LBM.

9.5.2 LTM PDU format

The LTM PDU format used by a MEP or MIP to transmit LTM information is shown in Figure 9.5-1.

NOTE – MIPs only transmit LTM information in response to received LTM information.

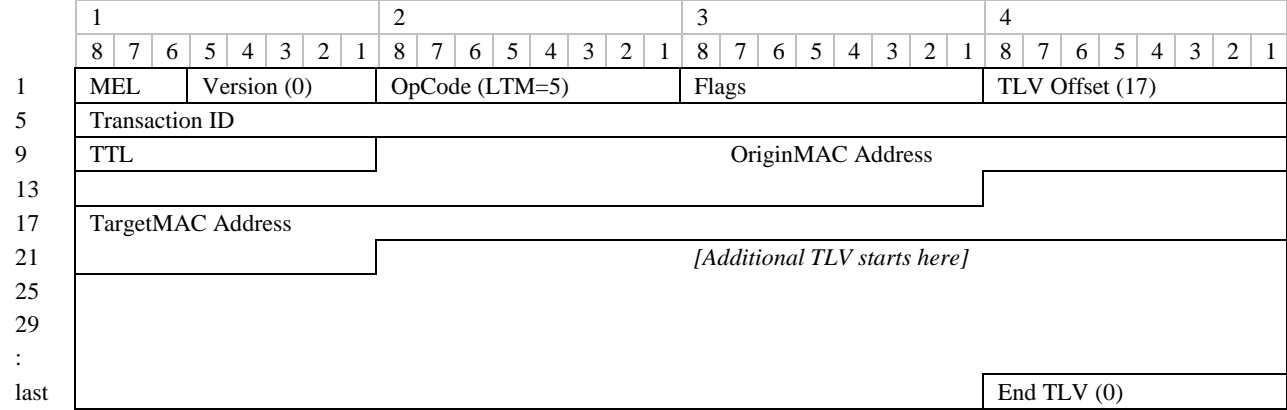


Figure 9.5-1 – LTM PDU format

The fields of the LTM PDU format are as follows:

- MEG Level: Refer to clause 9.1.
- Version: Refer to clause 9.1; value is 0 in the current version of this Recommendation.
- OpCode: The value of this PDU type is LTM (5).
- Flags: The format is as shown in Figure 9.5-2.




Figure 9.5-2 – Flags format in LTM PDU

- HWonly: Bit 8 set to 1. Value 1 indicates that only MAC addresses learned in a bridge's active data forwarding tables is to be used to forward the LTM to the next hop. When forwarding a received LTM, HWonly is copied from incoming LTM value.
- TLV Offset: Set to 17.
- Transaction ID: a 4-octet value containing the transaction ID for the LTM PDU.
- TTL: 1-octet field used to carry a TTL value as specified in clause 9.5.1.
- OriginMAC Address: a 6-octet OriginMAC as specified in clause 9.5.1.
- TargetMAC Address: a 6-octet TargetMAC as specified in clause 9.5.1.
- Additional TLV: LTM Egress Identifier TLV as specified in Figure 9.5-3.
- End TLV: all-ZEROes octet value.




Figure 9.5-3 – LTM egress identifier TLV format

- Type: identifies TLV type; value for this TLV type is LTM egress identifier (7).
- Length: identifies size, in octets, of the Value field containing the egress identifier. This is set to 8.
- Egress Identifier: identifies the MEP initiating the LTM frame or ETH-LT responder relaying the modified LTM frame. Octets 4 and 5 are ZEROes while remaining six Octets 6-11 contain a 48-bit IEEE MAC address unique to network element where the MEP or ETH-LT responder resides.

LTR is used to support ETH-LT reply, as described in clause 7.3.

The information elements carried in LTR include:

- ### 9.6.2 LTR PDU format

	1							2							3							4										
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	MEL			Version (0)				OpCode (LTR=4)							Flags							TLV Offset (6)										
5	Transaction ID																															
9	TTL							Relay Action							<i>[Additional TLVs start here]</i>																	
17																																
21																																
:																																
last																													End TLV (0)			

Figure 9.6-1 – LTR PDU format

- MEG Level: a 3-bit field the value of which is copied from the received LTM PDU.
- Version: Refer to clause 9.1; value is 0 in the current version of this Recommendation.
- OpCode: value of this PDU type is LTR (4).
- Flags: The format is as shown in Figure 9.6-2.

Figure 9.6-2 – Flags format in LTR PDU

- HWonlY: Bit 8 (HWonlY) is copied from incoming LTM value.
- FwdYes: Bit 7 is set to 1 if modified LTM frame was relayed, or set to 0 if no LTM frame was relayed.
- TerminalMEP: Bit 6 is set to 1 if reply egress TLV (or reply ingress TLV, if the reply egress TLV is not present) is a MEP, or set to 0 otherwise.

- TLV Offset: Set to 6.
 - Transaction ID: a 4-octet field the value of which is copied from the LTM PDU.
 - TTL: a 1-octet field the value of which is copied from the LTM PDU after decrementing it by one.
 - Relay Action: a 1-octet field that reports how the data frame targeted by the LTM would be passed through the MAC relay entity to the egress bridge port as described in clause 21.9.5 in [IEEE 802.1Q]. The value is defined in Table 21-27 of [IEEE 802.1Q].
 - Additional TLVs: LTR egress identifier TLV, reply ingress TLV and/or reply egress TLV as specified in Figures 9.6-3, 9.6-4 and 9.6-5 respectively.
- NOTE – Per [IEEE 802.1Q], LTR Egress Identifier TLV, Reply Ingress TLV or Reply Egress TLV are not the only type of TLVs that could be present in a LTM PDU, hence a LTR PDU could include additional TLVs.
- End TLV: all-ZEROes octet value.




Figure 9.6-3 – LTR egress identifier TLV format

The fields of the LTR egress identifier TLV format are as follows:

- Type: identifies TLV type; value for this TLV type is LTR egress identifier (8).
- Length: identifies size, in octets, of the Value field containing the last egress identifier and next egress identifier. This is set to 16.
- Last Egress Identifier: identifies the MEP that initiated, or ETH-LT responder that relayed the LTM frame to which this LTR frame is the response. This field is the same as the egress identifier in the LTM egress identifier TLV of the incoming LTM frame.
- Next Egress Identifier: identifies the ETH-LT responder that transmitted this LTR frame, and which can relay a modified LTM frame to the next hop. If the FwdYes bit of Flags field is 0, the contents of this field are undefined, and ignored by the LTR frame receiver. When not undefined, Octets 12 and 13 are ZEROes while the remaining six octets 14-19 contain a 48-bit IEEE MAC address unique to network element where the ETH-LT responder resides.




Figure 9.6-4 – Reply ingress TLV format

The fields of the reply ingress TLV format are as follows:

- Type: identifies TLV type; value for this TLV type is ingress reply (5).
- Length: identifies size, in octets, of the Value field. This is set to 7.
- Ingress Action: a 1-octet field which is reserved for definition by [IEEE 802.1].
- Ingress MAC Address: a 6-octet field which is reserved for definition by [IEEE 802.1].

Table 9-4 – AIS/LCK period values

Flags[3:1]	Period value	Comments
000-011	Invalid value	Invalid value for AIS/LCK PDUs
100	1s	1 frame per second
101	Invalid value	Invalid value for AIS/LCK PDUs
110	1 min	1 frame per minute
111	Invalid value	Invalid value for AIS/LCK PDUs

- TLV Offset: Set to 0.
- End TLV: all-ZEROes octet value.

9.8 LCK frame

LCK PDU is used to support ETH-LCK function, as described in clause 7.6.

9.8.1 LCK information elements

The information element carried in LCK is:

- Period: Period is a 3-bit information element carried in the three least significant bits of the Flags field. Period contains the value of LCK transmission periodicity. LCK period values are specified in Table 9-4.

9.8.2 LCK PDU format

The LCK PDU format used by a MEP to transmit LCK information is shown in Figure 9.8-1.




Figure 9.8-1 – LCK PDU format

The fields of the LCK PDU format are as follows:

- MEG Level: a 3-bit field that is used to carry the MEG level of the client MEG.
- Version: Refer to clause 9.1; value is 0 in the current version of this Recommendation.
- OpCode: value for this PDU type is LCK (35).
- Flags: one information element in the Flags field for the LCK PDU, Period, as follows:




Figure 9.8-2 – Flags format in LCK PDU

- Period: Bits 3 to 1 indicate transmission period with the encoding in Table 9-4.
- TLV Offset: Set to 0.
- End TLV: all-ZEROes octet value.

9.9 TST PDU

TST PDU is used to support the unidirectional ETH-Test function, as described in clause 7.7.

9.9.1 TST information elements

The information elements carried in TST are:

- Sequence Number: Sequence Number is a 4-octet field that contains the sequence number for the TST frames.
- Test: Test is an optional field whose length and contents are determined at the transmitting MEP. The contents of Test field indicate a test pattern and also carry an optional checksum. The test pattern can be a pseudo-random bit sequence (PRBS) ($2^{31} - 1$) as specified in clause 5.8 of [ITU-T O.150], an all '0' pattern, etc.

9.9.2 TST PDU format

The TST PDU format used by a MEP to transmit TST information is shown in Figure 9.9-1.

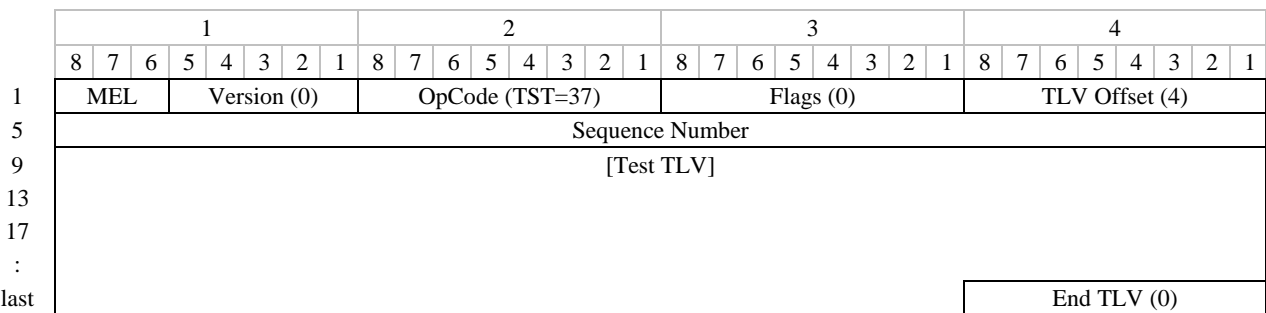


Figure 9.9-1 – TST PDU format

The fields of the TST PDU format are as follows:

- MEG Level: Refer to clause 9.1.
- Version: Refer to clause 9.1; value is 0 in the current version of this Recommendation.
- OpCode: value for this PDU type is TST (37).
- Flags: Set to all-ZEROes.




Figure 9.9-2 – Flags format in TST PDU

- TLV Offset: Set to 4.
- Sequence Number: a 4-octet value containing the sequence number incremented for successive TST PDUs.
- Test TLV: Test TLV as specified in Figure 9.3-4.
- End TLV: all-ZEROes octet value.

9.10 APS PDU

APS is used to support ETH-APS function, as described in clause 7.8.

9.10.1 APS information elements

Information elements carried in APS are outside the scope of this Recommendation.

The fields of the MCC PDU format are as follows:

- MEG Level: Refer to clause 9.1.
- Version: ETH-ED uses this field as described in clause 9.26. Other uses of this field are outside the scope of this Recommendation but they must conform to clause 9.1.
- OpCode: Value for this PDU type is MCC (41).
- Flags: ETH-ED uses this field as described in clause 9.26. Other uses of this field are outside the scope of this Recommendation; however, if not otherwise specified, it shall be set to all-ZEROes.

MSB				LSB			
8	7	6	5	4	3	2	1
Reserved (0)							

Figure 9.11-2 – Flags format in MCC PDU

- TLV Offset: 1-octet field. ETH-ED uses this field as described in clause 9.26. Other uses of this field are outside the scope of this Recommendation but they must conform to clause 9.1.
- OUI: 3-octet field that contains the organizationally unique identifier of the organization defining the format of MCC Data and values SubOpCode.
- SubOpCode: 1-octet field. When the OUI field contains the ITU-T OUI (00-19-A7), ETH-ED uses SubOpCode (1) as described in clause 9.26 and other values are reserved. When a different OUI is used, the values of the SubOpCode are outside the scope of this Recommendation.
- MCC Data: ETH-ED uses this field as described in clause 9.26. Other uses of this field are outside the scope of this Recommendation.
- End TLV: all-ZEROes octet value.

9.12 LMM PDU

LMM is used to support a single-ended proactive and on-demand ETH-LM request, as described in clause 8.1.2.

9.12.1 LMM information elements

Information elements carried in LMM are:

- TxFCf: TxFCf is a 4-octet field which carries the value of counter responsible for counting in-profile data frames transmitted by the MEP towards its peer MEP, at the time of LMM frame transmission.

9.12.2 LMM PDU format

The LMM PDU format used by a MEP to transmit LMM information is shown in Figure 9.12-1.

	1								2								3								4							
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	MEL				Version (1)				OpCode (LMM=43)								Flags								TLV Offset (12)							
5	TxFCf																															
9	Reserved for RxFCf in LMR																															
13	Reserved for TxFCb in LMR																															
17	End TLV (0)																															

Figure 9.12-1 – LMM PDU format

The fields of the LMM PDU format are as follows:

- MEG Level: Refer to clause 9.1.
- Version: Refer to clause 9.1; value for the LMM PDU on this version is set to 1.
- OpCode: value for this PDU type is LMM (43).
- Flags: One information elements in the Flags field, the LSB bit (Type), is used to indicate the type of the LMM operation as follows:

MSB							LSB
8	7	6	5	4	3	2	1
Reserved (0)							Type

Figure 9.12-2 – Flags format in LMM PDU

- Type: Bit 1 is set to 1 if it is the proactive operation, or set to 0 if it is the on-demand operation.
- TLV Offset: Set to 12.
- TxFCf: 4-octet integer values with samples of the frame counters, as specified in clause 9.12.1.
- Reserved: Reserved fields are set to all ZEROes.
- End TLV: an all-ZEROes octet value.

9.13 LMR PDU

LMR PDU is used to support a single-ended proactive and on-demand ETH-LM reply, as described in clause 8.1.2.

9.13.1 LMR information elements

Information elements carried in LMR are:

- TxFCf: TxFCf is a 4-octet field which carries the value of the TxFCf field in the last LMM PDU received by the MEP from its peer MEP.
- TxFCb: TxFCb is a 4-octet field which carries the value of the counter of in-profile data frames transmitted by the MEP towards its peer MEP at the time of LMR frame transmission.
- RxFCf: RxFCf is a 4-octet field which carries the value of the counter of in-profile data frames received by the MEP from its peer MEP, at the time of receiving last LMM frame from that peer MEP.

9.13.2 LMR PDU format

The LMR PDU format used by a MEP to transmit LMR information is shown in Figure 9.13-1.

	1								2								3								4							
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	MEL			Version					OpCode (LMR=42)								Flags								TLV Offset							
5	TxFCf																															
9	RxFCf																															
13	TxFCb																															
17	End TLV (0)																															

Figure 9.13-1 – LMR PDU format

The fields of the LMR PDU format are as follows:

- MEG Level: a 3-bit field the value of which is copied from the last received LMM PDU.

- Version: a 5-bit field the value of which is copied from the last received LMM PDU.
- OpCode: value for this PDU type is LMR (42).
- Flags: a 1-octet field the value of which is copied from the last received LMM PDU.
- TLV Offset: a 1-octet field the value of which is copied from the last received LMM PDU.
- TxFCf: 4-octet field the value of which is copied from the last received LMM PDU.
- RxFCf: 4-octet integer values with samples of the frame counters, as specified in clause 9.13.1.
- TxFCb: 4-octet integer values with samples of the frame counters, as specified in clause 9.13.1.
- End TLV: a 1-octet field the value of which is copied from the LMM PDU.

9.14 1DM PDU

The 1DM PDU is used to support proactive and on-demand dual-ended ETH-DM, as described in clause 8.2.1.

9.14.1 1DM information element

The information element carried in 1DM is:

- TxTimeStamp: TxTimeStamp is an 8-octet field that contains the timestamp of 1DM transmission. The format of TxTimeStamp is equal to the TimeRepresentation format in [IEEE 1588].

9.14.2 1DM PDU format

The 1DM PDU format used by a MEP to transmit 1DM information is shown in Figure 9.14-1.

1								2								3								4								
8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	
1	MEL		Version (1)					OpCode (1DM=45)								Flags								TLV Offset (16)								
5	TxTimeStampf																															
9																																
13																																
17	Reserved for 1DM receiving equipment (0) (for RxTimeStampf)																															
21	[Optional TLVs start here; otherwise End TLV]																															
25																																
29																																
:																																
last																							End TLV (0)									

Figure 9.14-1 – 1DM PDU format

The fields of the 1DM PDU format are as follows:

- MEG Level: Refer to clause 9.1.
- Version: Refer to clause 9.1; value for the 1DM PDU on this version is set to 1.
- OpCode: value for this PDU type is 1DM (45).
- Flags: One information element in the Flags field, the LSB bit (Type), is used to indicate the type of the 1DM operation as follows:

MSB				LSB			
8	7	6	5	4	3	2	1
Reserved (0)						Type	

Figure 9.14-2 – Flags format in 1DM PDU

- Type: Bit 1 is set to 1 if it is the proactive operation, or set to 0 if it is the on-demand operation.
- TLV Offset: Set to 16.
- TxTimeStamp: an 8-octet transmit time stamp field as described in clause 9.14.1.
- Reserved: an 8-octet reserved field set to all-ZEROes.
- Optional TLV(s): A Test ID TLV as specified in Figure 9.14-3 and/or a Data TLV as specified in Figure 9.3-3, with configurable size, in octets may be included in any 1DM transmitted. When a Test ID TLV is included in this field, it is recommended to put Test ID TLV first (prior to Data TLV). For the purpose of ETH-DM, the value part of Data TLV is unspecified.
- End TLV: all-ZEROes octet value.

	1								2								3								4							
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	Type (36)								Length																Test ID							
5	Test ID																															

Figure 9.14-3 – Test ID TLV format

The fields of the Test ID TLV format are as follows:

- Type: identifies TLV type; value for this TLV type is Test ID (36).
- Length: identifies size. Must be 32.
- Test ID: Test ID is a 4-octet field set by the transmitting MEP when used to run multiple tests simultaneously between MEPs.

9.15 DMM PDU

DMM is used to support proactive or on-demand Single-Ended ETH-DM request, as described in clause 8.2.2.

9.15.1 DMM information elements

The information elements carried in DMM are:

- TxTimeStamp: TxTimeStamp is an 8-octet field that contains the timestamp of DMM transmission. The format of TxTimeStamp is equal to the TimeRepresentation format in [IEEE 1588].

9.15.2 DMM PDU format

The DMM PDU format used by a MEP to transmit DMM information is shown in Figure 9.15-1.

	1								2								3								4							
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	MEL			Version (1)					OpCode (DMM=47)								Flags								TLV Offset (32)							
5	TxTimeStampf																															
9																																
13																																
17	Reserved for DMM receiving equipment (0) <i>(for RxTimeStampf)</i>																															
21	Reserved for DMR (0) <i>(for TxTimeStampb)</i>																															
25	Reserved for DMR receiving equipment (0)																															
29																																
33																																
37	<i>[Optional TLV(s) start(s) here; otherwise End TLV]</i>																															
41																																
45																																
:																																
last																									End TLV (0)							

Figure 9.15-1 – DMM PDU format

The fields of the DMM PDU format are as follows:

- MEL Level: Refer to clause 9.1.
- Version: Refer to clause 9.1; value for the DMM PDU is set to 1.
- OpCode: value for this PDU type is DMM (47).
- Flags: Set to all-ZEROes. One information elements in the Flags field, the LSB bit (Type), is used to indicate the type of the DMM operation as follows:

MSB								LSB	
8	7	6	5	4	3	2	1		
Reserved (0)								Type	

Figure 9.15-2 – Flags format in DMM PDU

- Type: Bit 1 is set to 1 if it is the proactive operation, or set to 0 if it is the on-demand operation.
- TLV Offset: Set to 32.
- TxTimeStampf: an 8-octet transmit time stamp field as described in clause 9.15.1.
- Reserved: 24-octet reserved fields are set to all-ZEROes.
- Optional TLV: A Test ID TLV as specified in Figure 9.14-3 and/or a Data TLV as specified in Figure 9.3-3, with configurable size, in octets may be included in any DMM transmitted. When a Test ID TLV is included in this field, it is recommended to put Test ID TLV first (prior to Data TLV). For the purpose of ETH-DM, the value part of Data TLV is unspecified.
- End TLV: an all-ZEROes octet value.

9.16 DMR PDU

DMR is used to support a single-ended ETH-DM reply, as described in clause 8.2.2.

9.16.1 DMR information elements

The information elements carried in DMR are:

- TxTimeStampf: TxTimeStampf is an 8-octet field that contains the copy of TxTimeStampf field in received DMM.

9.17.1 EXM PDU format

The information elements carried in EXM include:

- OUI: OUI is a 3-octet field that contains the organizationally unique identifier of the organization using the EXM.
- SubOpCode: SubOpCode is a 1-octet field that is used to interpret the remaining fields in the EXM frame.
- EXM Data: Depending on the functionality indicated by the OUI and organizationally specific SubOpCode, EXM may carry one or more TLVs. EXM Data is outside the scope of this Recommendation.

9.17.2 EXM PDU format

The EXM PDU format used by a MEP to transmit EXM information is shown in Figure 9.17-1.

	1								2								3								4							
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	MEL			Version (0)					OpCode (EXM=49)								Flags								TLV Offset							
5	OUI																								SubOpCode							
9	[optional EXM data; else End TLV]																															
:																																
:																																
:																									End TLV (0)							

Figure 9.17-1 – EXM PDU format

The fields of the EXM PDU format are as follows:

- MEG Level: Refer to clause 9.1.
- Version: Its specific value for EXM is outside the scope of this Recommendation but it must conform to clause 9.1.
- OpCode: value for this PDU type is EXM (49).
- Flags: outside the scope of this Recommendation.
- TLV Offset: 1-octet field. Its specific value for EXM is outside the scope of this Recommendation but it must conform to clause 9.1.
- OUI: 3-octet field the values of which are outside the scope of this Recommendation.
- SubOpCode: 1-octet field the values of which are outside the scope of this Recommendation.
- EXM Data: format and length of this field are outside the scope of this Recommendation.
- End TLV: all-ZEROes octet value.

9.18 EXR PDU

EXR is used as an experimental OAM reply PDU.

9.18.1 EXR information elements

The information elements carried in EXR include:

- OUI: OUI is a 3-octet field that contains the organizationally unique identifier of the organization using the EXR.
- SubOpCode: SubOpCode is a 1-octet field that is used to interpret the remaining fields in the EXR frame.

- EXR Data: Depending on the functionality indicated by the OUI and organizationally specific SubOpCode, EXR Data may carry one or more TLVs. EXR Data is outside the scope of this Recommendation.

9.18.2 EXR PDU format

The EXR PDU format used to transmit EXR information is shown in Figure 9.18-1.

	1								2								3								4							
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	MEL			Version					OpCode (EXR=48)								Flags								TLV Offset							
5	OUI																								SubOpCode							
9	[Optional EXR Data; else End TLV]																															
:																																
:																																
:																									End TLV (0)							

Figure 9.18-1 – EXR PDU format

The fields of the EXR PDU format are as follows:

- MEL Level: a 3-bit field the value of which is copied from the last received EXM PDU.
- Version: a 5-bit field the value of which is copied from the last received EXM PDU.
- OpCode: value for this PDU type is EXR (48).
- Flags: outside the scope of this Recommendation.
- TLV Offset: 1-octet field. Its specific value for EXR is outside the scope of this Recommendation but it must conform to clause 9.1.
- OUI: 3-octet field the value of which is copied from the last received EXM PDU.
- SubOpCode: 1-octet field the values of which are outside the scope of this Recommendation.
- EXR Data: format and length of this field are outside the scope of this Recommendation.
- End TLV: All-ZEROes octet value.

9.19 VSM PDU

VSM is used as a vendor-specific OAM request PDU.

9.19.1 VSM PDU format

The information elements carried in VSM include:

- OUI: OUI is a 3-octet field that contains the organizationally unique identifier of the organization using the VSM.
- SubOpCode: SubOpCode is a 1-octet field that is used to interpret the remaining fields in the VSM frame.
- VSM Data: Depending on the functionality indicated by the OUI and organizationally specific SubOpCode, VSM Data may carry one or more TLVs. VSM Data is outside the scope of this Recommendation.

9.19.2 VSM PDU format

The VSM PDU format used by a MEP to transmit VSM information is shown in Figure 9.19-1.

	1								2								3								4							
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	MEL			Version (0)					OpCode (VSM=51)								Flags								TLV Offset							
5	OUI																								SubOpCode							
9	[Optional VSM Data; else End TLV]																															
:																																
:																																
:																									End TLV (0)							

Figure 9.19-1 – VSM PDU format

The fields of the VSM PDU format are as follows:

- MEG Level: Refer to clause 9.1.
- Version: Its specific value for VSM is outside the scope of this Recommendation but it must conform to clause 9.1.
- OpCode: value for this PDU type is VSM (51).
- Flags: outside the scope of this Recommendation.
- TLV Offset: 1-octet field. Its specific value for VSM is outside the scope of this Recommendation but it must conform to clause 9.1.
- OUI: 3-octet field the values of which are outside the scope of this Recommendation.
- SubOpCode: 1-octet field the values of which are outside the scope of this Recommendation.
- VSM Data: format and length of this field are outside the scope of this Recommendation.
- End TLV: all-ZEROes octet value.

9.20 VSR PDU

VSR is used as a vendor-specific OAM reply PDU.

9.20.1 VSR information elements

The information elements carried in VSR include:

- OUI: OUI is a 3-octet field that contains the organizationally unique identifier of the organization using the VSR.
- SubOpCode: SubOpCode is a 1-octet field that is used to interpret the remaining fields in the VSR frame.
- VSR Data: Depending on the functionality indicated by the OUI and organizationally specific SubOpCode, VSR Data may carry one or more TLVs. VSR Data is outside the scope of this Recommendation.

9.20.2 VSR PDU format

The VSR PDU format used to transmit VSR information is shown in Figure 9.20-1.

1								2								3								4								
8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	
1	MEL		Version					OpCode (VSR=50)							Flags							TLV Offset										
5	OUI																							SubOpCode								
9	[Optional VSR Data; else End TLV]																															
:																																
:																																
:																																
:																								End TLV (0)								

Figure 9.20-1 – VSR PDU format

The fields of the VSR PDU format are as follows:

- MEG Level: a 3-bit field the value of which is copied from the last received VSM PDU.
- Version: a 5-bit field the value of which is copied from the last received VSM PDU.
- OpCode: value for this PDU type is VSR (50).
- Flags: outside the scope of this Recommendation.
- TLV Offset: 1-octet field. Its specific value for EXR is outside the scope of this Recommendation but it must conform to clause 9.1.
- OUI: 3-octet field the value of which is copied from the last received VSM PDU.
- SubOpCode: 1-octet field the values of which are outside the scope of this Recommendation.
- VSR Data: format and length of this field are outside the scope of this Recommendation.
- End TLV: all-ZEROes octet value.

9.21 Client signal fail (CSF)

The CSF PDU is used to support the ETH-CSF function, as described in clause 7.12.

The CSF PDU format is shown in Figure 9.21-1.




Figure 9.21-1 – CSF PDU format

The fields of the CSF PDU format are as follows:

- MEG Level: a 3-bit field that is used to carry the local MEG level.
- Version: Refer to clause 9.1; value is 0 in the current version of this Recommendation.
- OpCode: value for this PDU type is CSF (52).
- Flags: one information element in the Flags field for CSF PDU. It consists of a 3-bit Type sub-element and a 3-bit Period sub-element formatted as follows:




Figure 9.21-2 – Flags format in CSF PDU

- Type: Bits 6 to 4 indicate the CSF type with the encoding in Table 9-5.

Table 9-5 – CSF type values

Flags[6:4]	Type	Comments
000	LOS	Client loss of signal
001	FDI/AIS	Client forward defect indication
010	RDI	Client reverse defect indication
011	DCI	Client defect clear indication

- Period: Bits 3 to 1 indicate transmission period with the encoding Table 9-6.

Table 9-6 – CSF period values

Flags[3:1]	Period value	Comments
000	Invalid value	Invalid value for CSF PDUs
001	For further study	For further study
010	For further study	For further study
011	For further study	For further study
100	1s	1 frame per second
101	For further study	For further study
110	1 min	1 frame per minute
111	For further study	For further study

- TLV Offset: Set to 0.
- End TLV: all-ZEROes octet value.

9.22 SLM PDU

SLM is used to support single-ended ETH-SLM requests, as described in clause 8.4.1.

9.22.1 SLM information elements

Information elements carried in SLM include:

- Source MEP ID: Source MEP ID is a 2-octet field where the last 13 least significant bits are used to identify the MEP transmitting the SLM frame. MEP ID is unique within the MEG.
- Test ID: Test ID is a 4-octet field set by the transmitting MEP and used to identify a test when multiple tests run simultaneously between MEPs including on concurrent on-demand and proactive tests.
- TxFCf: TxFCf is a 4-octet field which carries the number of SLM frames transmitted by the MEP towards its peer MEP.

9.22.2 SLM PDU format

The SLM PDU format used by a MEP to transmit SLM information is shown in Figure 9.22-1.

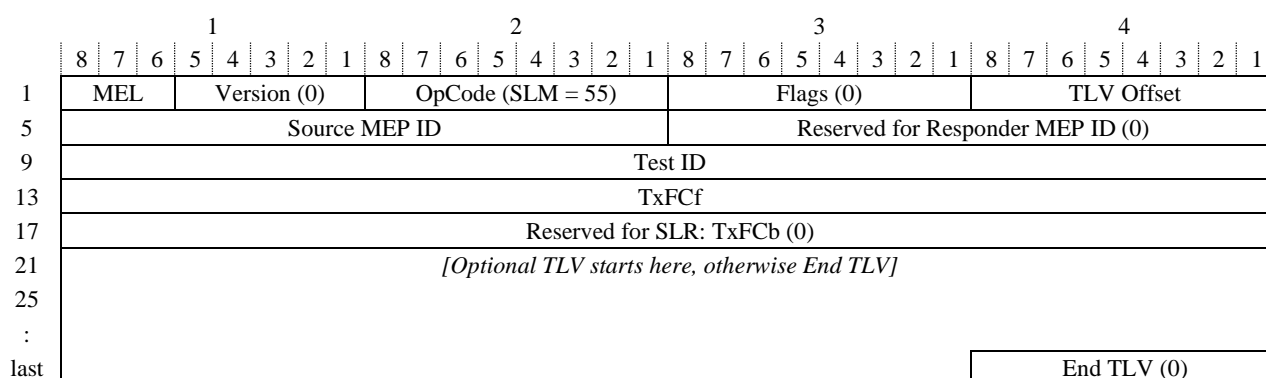


Figure 9.22-1 – SLM PDU format

The fields of the SLM PDU format are as follows:

- MEG Level: Refer to clause 9.1.
- Version: Refer to clause 9.1; value is 0 in the current version of this Recommendation.
- OpCode: value for this PDU type is SLM (55).

- Flags: Set to all-ZEROes.
- TLV Offset: Set to 16.
- Reserved: Reserved fields are set to all ZEROes.
- Source MEP ID: a 2-octet field used to identify the MEP transmitting the SLM frame, as specified in clause 9.22.1.
- Test ID: a 4-octet field used to identify a unique test among MEPs, as specified in clause 9.22.1.
- TxFCf: a 4-octet integer value representing the number of SLM frames transmitted, as specified in clause 9.22.1.
- Optional TLV: A Data TLV as specified in Figure 9.3-3 may be included in any SLM transmitted. For the purpose of ETH-SLM, the value part of Data TLV is unspecified.
- End TLV: an all-ZEROes octet value.

SLR is used to support a single-ended ETH-SLM reply, as described in clause 8.4.1.

The information elements carried in SLR include:

9.23.2 SLR PDU format

	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	MEL				Version				OpCode (SLR = 54)								Flags								TLV Offset(16)							
5	Source MEP ID																Responder MEP ID															
9	Test ID																															
13	TxFCf																															
17	TxFCb																															
21	[Optional TLV starts here, otherwise End TLV]																															
:																																
last																									End TLV (0)							

The fields of the SLR PDU format are as follows:

- Version: a 5-bit field the value of which is copied from the last received SLM PDU.
- OpCode: value for this PDU type is SLR (54).
- Flags: a 1-octet field the value of which is copied from the SLM PDU.
- TLV Offset: a 1-octet field the value of which is copied from the SLM PDU.
- Reserved: Reserved fields are set to all ZEROes.
- Source MEP ID: a 2-octet field the value of which is copied from the SLM PDU.
- Responder MEP ID: a 2-octet field used to identify the MEP transmitting the SLR frame, as specified in clause 9.22.1.
- Test ID: a 4-octet field the value of which is copied from the SLM PDU.
- TxFCf: a 4-octet field the value of which is copied from the SLM PDU.
- TxFCb: a 4-octet integer value representing the number of SLR frames transmitted, as specified in clause 9.22.1.
- Optional TLV: if present in SLM PDU, is copied from the SLM PDU.
- End TLV: a 1-octet field the value of which is copied from the SLM PDU.

9.24 1SL PDU

1SL is used to support proactive and on-demand dual-ended ETH-SLM, as described in clause 8.4.2.

9.24.1 1SL information elements

Information elements carried in 1SL include:

- Source MEP ID: Source MEP ID is a 2-octet field where the last 13 least significant bits are used to identify the MEP transmitting the 1SL frame. MEP ID is unique within the MEG.
- Test ID: Test ID is a 4-octet field set by the transmitting MEP and used to identify when multiple tests run simultaneously towards different MEPs including concurrent on-demand and proactive tests.
- TxTCf: TxTCf is a 4-octet field which carries the number of 1SL frame transmitted by the MEP towards its peer MEPs.

9.24.2 1SL PDU format

The 1SL PDU format used by a MEP to transmit 1SL information is shown in Figure 9.24-1.

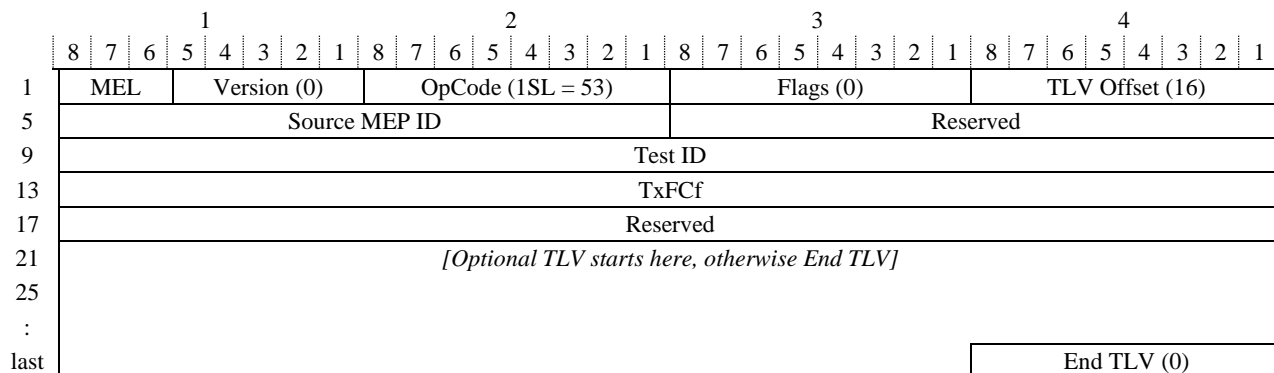


Figure 9.24-1 – 1SL PDU format

The fields of the 1SL PDU format are as follows:

- MEG Level: Refer to clause 9.1.
- Version: Refer to clause 9.1; value is 0 in the current version of this Recommendation.

- OpCode: value for this PDU type is 1SL (53).
- Flags: Set to all-ZEROes.
- TLV Offset: Set to 16.
- Reserved: Reserved fields are set to all ZEROes.
- Source MEP ID: a 2-octet field used to identify the MEP transmitting the 1SL frame, as specified in clause 9.24.1.
- Test ID: a 4-octet field used to identify a unique test among MEPs, as specified in clause 9.24.1.
- TxFCf: a 4-octet integer value representing the number of 1SL frames transmitted, as specified in clause 9.24.1.
- Optional TLV: A Data TLV as specified in Figure 9.3-3 may be included in any 1SL transmitted. For the purpose of ETH-SLM, the value part of Data TLV is unspecified.
- End TLV: an all-ZEROes octet value.

9.25 BNM PDU

The BNM PDU is used to support the ETH-BNM function, as described in clause 7.13.

9.25.1 BNM information elements

The information element carried in BNM is:

- Period: Period is a 3-bit information element carried in the three least significant bits of the Flags field. Period contains the value of BNM transmission periodicity. BNM period values are specified in Table 9-7.
- Nominal Bandwidth: Nominal Bandwidth is nominal full bandwidth of the link, expressed in integer Mb/s.
- Current Bandwidth: Current Bandwidth is current bandwidth of the link, expressed in integer Mb/s.
- Port ID: Port ID is either a non-zero unique identifier for the port or zero if this identifier is not used.

The nominal full bandwidth and the current bandwidth values represent the available bandwidth of the server layer.

9.25.2 BNM PDU format

The BNM PDU format used by a server MEP to transmit BNM information is shown in Figure 9.25-1.

	1								2								3								4							
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	MEL			Version (0)					OpCode (GNM=32)								Flags								TLV Offset (13)							
5	Sub-OpCode (BNM=1)								Nominal Bandwidth																							
9	Nominal Bandwidth (cont)								Current Bandwidth																							
13	Current Bandwidth (Cont)								Port ID																							
17	Port ID (Cont)								End TLV (0)																							

Figure 9.25-1 – BNM PDU format

The fields of the BNM PDU format are as follows:

- MEG Level: a 3-bit field that is used to carry the MEG Level of the client MEG.
- Version: Refer to clause 9.1; value is 0 in the current version of this Recommendation.
- OpCode: value for this PDU type is GNM (32).

- Flags: one information element in the Flags field for the BNM PDU, Period, as follows:

MSB				LSB			
8	7	6	5	4	3	2	1
Reserved (0)					Period		

Figure 9.25-2 – Flags format in BNM PDU

- Period: Bits 3 to 1 indicate a transmission period with the encoding in Table 9-7.

Table 9-7 – BNM period values

Flags[3:1]	Period value	Comments
000	Invalid value	Invalid value for BNM PDUs
001	For further study	For further study
010	For further study	For further study
011	For further study	For further study
100	1s	1 frame per second
101	10s	1 frame per 10 seconds
110	1 min	1 frame per minute
111	Invalid value	Invalid value for BNM PDUs

- TLV Offset: Set to 13.
- Sub-OpCode: value for this PDU type is BNM (1).
- Nominal bandwidth: the nominal full bandwidth of the link, expressed as integer Mb/s.
- Current bandwidth: the current bandwidth of the link, expressed as integer Mb/s.
- Port ID: an optionally used non-zero 32-bit identifier for the port to which the bandwidth information pertains. The value must be unique over all server links within the client MEG. If this identifier is not used, the value should be zero.
- End TLV: all-ZEROes octet value.

9.26 EDM PDU

The EDM PDU is used to support the ETH-ED function, as described in clause 7.14.

9.26.1 EDM information elements

The information element carried in EDM is:

- MEP ID: MEP ID is a 2-octet field where the 13 least significant bits are used to identify the MEP transmitting the EDM frame. MEP ID is unique within the MEG.
- Expected Duration: Expected duration for which the peer MEP is requested to suppress Loss of Continuity defects.

9.26.2 EDM PDU format

The EDM PDU format used by a MEP to transmit EDM information is shown in Figure 9.26-1.

	1								2								3								4							
	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
1	MEL			Version (0)					OpCode (MCC=41)								Flags (0)								TLV Offset (10)							
5	OUI																SubOpCode(EDM=1)															
9	MEP ID																Expected Duration															
13	Expected Duration (cont)																End TLV (0)															

Figure 9.26 – EDM PDU format

The fields of the EDM PDU format are as follows:

- MEG Level: a 3-bit field that is used to carry the MEG Level of the client MEG.
- Version: Refer to clause 9.1; value is 0 in the current version of this Recommendation.
- OpCode: value for this PDU type is MCC (41).
- Flags: Set to all ZEROes.
- TLV Offset: Set to 10.
- OUI: Set to the ITU-T OUI, 00-19-A7.
- Sub-OpCode: value for this PDU type is EDM (1).
- MEP ID: a 13-bit integer value identifying the transmitting MEP within the MEG. The three MSBs of the first octet are not used and set to ZERO.
- Expected Duration: The duration, in seconds, for which the loss of continuity is expected to last (starting from when the first EDM is transmitted).
- End TLV: all-ZEROes octet value.

10 OAM frame addresses

OAM frames are identified by a unique EtherType the value of which is 0x8902. OAM frames processing and filtering at a MEP is based on the OAM EtherType and MEG Level fields for both unicast and multicast DA.

As indicated in clauses 7 and 8, the DA in an OAM frame could be multicast or unicast depending on the specific OAM functionality. The SA in an OAM frame is always unicast.

This clause provides further discussion on the choice of DA in specific OAM functions. Table 10-1 provides a summary of the DAs that are applicable for different OAM types.

NOTE – The choice of MAC DA for Ethernet OAM frames is application-specific. Implementations are not required to support all of the addresses specified in this Recommendation; however, support for those addresses specified in [ITU-T G.8021] is required.

10.1 Multicast destination addresses

The following types of multicast addresses are required depending on the type of OAM function:

- Multicast DA Class 1: OAM frames that are addressed to all peer MEPs in a MEG (e.g., CCM, Multicast LBM, AIS, etc.).
- Multicast DA Class 2: OAM frames that are addressed to all MIPs and peer MEPs in a MEG (e.g., LTM).
- Multicast DA for Ring APS: OAM frames used for Ethernet ring protection.

Normally, a single multicast DA Class 1 address and a single multicast DA Class 2 address would be sufficient. However, for a short-term deployment of Ethernet OAM across the current Ethernet equipment, a multicast DA could also implicitly carry the MEG level. This would require 8 distinct addresses for each of the multicast DA Classes 1 and 2 for the 8 different MEG levels.

Specific values for 8 multicast addresses for Class 1 and 8 Multicast addresses for Class 2 are 01-80-C2-00-00-3x and 01-80-C2-00-00-3y respectively; x represents the MEG level, with x being a value in the range 0-7, and y represents the MEG level, with y being a value in the range 8-F.

In addition, a specific range of multicast DA with ITU OUI (01-19-A7) is used for Ring APS frames. See more details in [ITU-T G.8032].

10.2 CCM

CCM frames are generated with multicast Class 1 DA in a multipoint MEG, and are typically generated with a multicast DA in a point-to-point MEG except as described below.

Using a multicast DA, CCM frames allow discovery of MAC addresses associated with peer MEPs of the receiving MEP. The use of a multicast DA also allows for the detection of misconnections among flow domain fragments. Detection of misconnections is described in clause 7.1.

When detection of the above conditions is important, a multicast DA must be used for CCM frames. When the above conditions are not expected or are not required to be detected and the data frames in different services instances are distinguished using unicast DAs (as in provisioned environments for point-to-point connections), CCM frames are generated with the unicast DA of the peer MEP.

10.3 LBM

LBM frames can be generated with unicast or multicast Class 1 DAs, as per the unicast ETH-LB or multicast ETH-LB functions respectively.

10.4 LBR

LBR frames are always generated with unicast DAs.

10.5 LTM

LTM frames are generated with a multicast Class 2 DA.

A multicast DA is used instead of a unicast DA for LTM frames since in current bridges, the MIPs would not be able to intercept a frame with a unicast DA which was not their own address. Therefore, the MIPs would not be able to reply and would simply forward the LTM frame with the unicast DA. The limitation is that current ports do not look at the EtherType before looking at the DA.

10.6 LTR

LTR frames are always generated with unicast DAs.

10.7 AIS

AIS frames are generated with a multicast Class 1 DA in a multipoint MEG, and are typically generated with a multicast Class 1 DA in a point-to-point MEG except as described below.

In provisioned environments for point-to-point connections where the data frames in different services instances are distinguished using unicast DAs, AIS frames are generated with the unicast DA of the downstream MEP.

10.8 LCK

LCK frames are generated with multicast Class 1 DA in a multipoint MEG, and are typically generated with a multicast Class 1 DA in a point-to-point MEG except as described below.

In provisioned environments for point-to-point connections where the data frames in different services instances are distinguished using unicast DAs, AIS frames are generated with the unicast DA of the downstream MEP.

10.9 TST

TST frames are generated with unicast DAs. TST frames may be generated with a multicast Class 1 DA if multipoint diagnostics are desired.

10.10 APS

For linear APS, refer to [ITU-T G.8031]. For ring APS, refer to [ITU-T G.8032].

10.11 MCC

MCC frames are generated with unicast DAs. For the case when a point-to-point VLAN is being used, a multicast Class 1 DA may be used.

10.12 LMM

LMM frames are generated with unicast DAs. LMM frames may be generated with a multicast Class 1 DA if multipoint measurements are desired.

10.13 LMR

LMR frames are always generated with unicast DAs.

10.14 1DM

1DM frames are generated with unicast DAs. 1DM frames may be generated with a multicast Class 1 DA if multipoint measurements are desired.

10.15 DMM

DMM frames are generated with unicast DAs. DMM frames may be generated with a multicast Class 1 DA if multipoint measurements are desired.

10.16 DMR

DMR frames are always generated with unicast DAs.

10.17 EXM

EXM frames DA is outside the scope of this Recommendation.

10.18 EXR

EXR frames DA is outside the scope of this Recommendation.

10.19 VSM

VSM frames DA is outside the scope of this Recommendation.

10.20 VSR

VSR frames DA is outside the scope of this Recommendation.

10.21 CSF

CSF frames are generated with a multicast Class 1 DA in a multipoint MEG, and are typically generated with a multicast class 1 DA in a point-to-point MEG except as described below.

In provisioned environments for point-to-point connections where the data frames in different services instances are distinguished using unicast DA, CSF frames are generated with the unicast DA of the downstream MEP.

10.22 SLM

SLM frames are generated with unicast DAs. SLM frames may be generated with a multicast Class 1 DA if multipoint measurements are desired.

10.23 SLR

SLR frames are always generated with unicast DAs.

10.24 1SL

1SL frames are generated with unicast DAs. 1SL frames may be generated with a multicast Class 1 DA if multipoint measurements are desired.

10.25 BNM

BNM frames are generated with a multicast Class 1 DA in a multipoint MEG, and are typically generated with a multicast class 1 DA in a point-to-point MEG except as described below.

In provisioned environments for point-to-point connections where the data frames in different services instances are distinguished using unicast DAs, BNM frames are generated with the unicast DA of the downstream MEP.

10.26 EDM

EDM frames are generated with a multicast Class 1 DA in a multipoint MEG, and are typically generated with a multicast class 1 DA in a point-to-point MEG except as described below.

In provisioned environments for point-to-point connections where the data frames in different services instances are distinguished using unicast DAs, EDM frames are generated with the unicast DA of the downstream MEP.

Table 10-1 – OAM frame DA

OAM type	DAs for frames with OAM PDU
CCM	Multicast Class 1 DA or unicast DA
LBM	Unicast DA or multicast Class 1 DA
LBR	Unicast DA
LTM	Multicast Class 2 DA
LTR	Unicast DA
AIS	Multicast Class 1 DA or unicast DA
LCK	Multicast Class 1 DA or unicast DA
TST	Unicast DA or multicast Class 1 DA
Linear APS	Refer to [ITU-T G.8031]
Ring APS	Refer to [ITU-T G.8032]
MCC	Unicast DA or multicast Class 1 DA
LMM	Unicast DA or multicast Class 1 DA
LMR	Unicast DA
IDM	Unicast DA or multicast Class 1 DA
DMM	Unicast DA or multicast Class 1 DA
DMR	Unicast DA
EXM, EXR, VSM, VSR	Outside the scope of this Recommendation
CSF	Multicast Class 1 DA or unicast DA

Table 10-1 – OAM frame DA

OAM type	DAs for frames with OAM PDU
SLM	Unicast DA or multicast Class 1 DA
SLR	Unicast DA
1SL	Unicast DA or Multicast Class 1 DA
BNM	Multicast Class 1 DA or unicast DA
EDM	Multicast Class 1 DA or unicast DA

11 OAM PDU validation and versioning

This clause describes rules for validation and versioning of OAM PDUs, which are designed to ensure that implementations of this Recommendation will interoperate with implementations of future versions of this Recommendation. In addition, these rules allow implementations to provide proprietary, non-standard extensions to the protocol in a way which does not jeopardize interoperability with future versions of this Recommendation or restrict the ability of future versions of this Recommendation to extend the Recommendation functionality.

NOTE 1 – The change to the LTM format between the 2006 and 2008 versions of this Recommendation did not change the version number; however, future revisions to this Recommendation must align with these rules.

NOTE 2 – The rules described here only apply to how PDUs with different versions are interpreted. Further details regarding how PDUs are subsequently processed, where applicable, may be found in the atomic function definitions in [ITU-T G.8021] and [ITU-T G.8032].

NOTE 3 – These rules do not apply to parts of PDUs that are not specified in ITU-T Recommendations, for example, the data fields of VSM, VSR, EXM and EXR PDUs.

11.1 OAM PDU transmission

OAM PDU transmission is required to meet the following requirements:

- The fixed header fields shall be transmitted exactly as specified in this Recommendation.
- All bits defined as "reserved" in this Recommendation shall be transmitted as 0.
- Additional fields shall not be added to the fixed header specified in this Recommendation.
- Code points reserved in this Recommendation or [IEEE 802.1] shall not be transmitted in any OAM PDU; for example, reserved values of the OpCode field (Table 9-1), the TLV Type field (Table 9-2), or the MEG ID format field (Table A.1).
- Additional fields shall not be added to any TLV specified in this Recommendation.

11.2 OAM PDU validation in reception

Received OAM PDUs are subject to a number of validation tests and are discarded without further processing if they fail these tests. This clause does not provide an exhaustive list of such tests, it covers only those aspects that are most important to future compatibility. In addition to the tests specified here, it may be assumed that if an OAM PDU with a particular OpCode does not conform to the corresponding description in clause 9, it fails the tests. The initial validation test is to ensure that the OAM PDU is sufficiently long to contain the MEG level and version fields. OAM PDUs that fail this test are discarded.

The OAM PDU is subsequently processed in accordance with whichever is numerically lower of 1) the Version field in the OAM PDU and 2) the highest version number known to the receiving implementation. That is, a version 1 implementation receiving a version 0 OAM PDU processes it according to version 0, and it processes a version 1 OAM PDU according to version 1. It is noted that the imposition on future versions of this Recommendation that all earlier version implementations

can process received OAM PDUs correctly, that is, that OAM PDUs specified by later versions of this Recommendation must remain valid when processed according to version 0.

The following validation tests are used, according to the version selected as described above:

- The fixed header length, as determined by the TLV Offset field, is not shorter than the length specified by the selected version.
- The OAM PDU is sufficiently long to contain a fixed header of the length specified by the selected version.

If the OAM PDU contains a TLV that needs to be processed, the following validation tests are used, according to the version selected as described above:

- The OAM PDU is sufficiently long to contain a TLV Value field whose length is specified by the TLV Length field.
- A TLV Length field does not indicate a length that is shorter than the minimum length for that TLV as specified by the selected version.

The following criteria shall not be used to validate a received OAM PDU:

- The fixed header can be longer than the length specified by the selected version.
- Bits can be set in the reserved bits of the Flags field.
- A TLV can have a Type field not specified by the selected version of the standard.
- A TLV's Length field can be larger than the value (if any) specified in the selected version of the standard.
- Either the TLV Offset field, or the Length field of the last TLV in the OAM PDU, can indicate a position for the first (next) TLV that coincides with the end of the OAM PDU. That is, the end TLV can be missing from the OAM PDU.
- TLVs may occur in any order in the OAM PDU, unless the descriptions in clause 9 specify otherwise.

NOTE – The selection of the version to be used for processing a received OAM PDU does not impact the version copying requirement if an OAM PDU reply needs to be generated. This means that a version 0 implementation receiving a version 1 OAM PDU request interprets it according to version 0, but replies depending on the replying rules, unless this rule has version dependency. In this case, the reception of a version 1 OAM PDU reply cannot be used as an indication that the OAM PDU request has been processed according to version 1.

11.3 OAM PDU reception after validation

Received OAM PDUs that pass the validation tests described above must be processed in accordance with the following rules, and in accordance with the same version selected for the validation tests (that is, the numerically lower one of the Version field in the OAM PDU and the highest version number known to the receiving implementation).

- Only those fields in the fixed header portion of the OAM PDU that are defined in the selected version are processed; any extra octets in the fixed header, if it is longer than the length specified by the selected version, are ignored.
- Any TLV with a Type field not specified by the selected version is ignored, except if the OAM PDU is forwarded or retransmitted (with or without modification), or if a new OAM PDU is sent in response to the received OAM PDU, the TLV is copied without modification into the forwarded or retransmitted PDU or into the response PDU.
- Any part of the OAM PDU following the end TLV is ignored (the lack of an end TLV is not an error).
- If any TLV's Length field is larger than the value (if any) specified by the selected version, then any octets following those specified by the selected version are ignored.
- All bits undefined in this Recommendation, e.g., reserved bits in the Flags field, are ignored.

Annex A

MEG ID format

(This annex forms an integral part of this Recommendation.)

The features of maintenance entity group identifiers (MEG IDs) are:

- Each MEG ID must be globally unique.
- Where it may be expected that the MEG may be required for path set-up across an inter-operator boundary, the MEG ID must be available to other network operators.
- The MEG ID should not change while the MEG remains in existence.
- The MEG ID should be able to identify the network operator which is responsible for the MEG.
- The generic format of MEG IDs specific to this Recommendation is shown in Figure A.1.

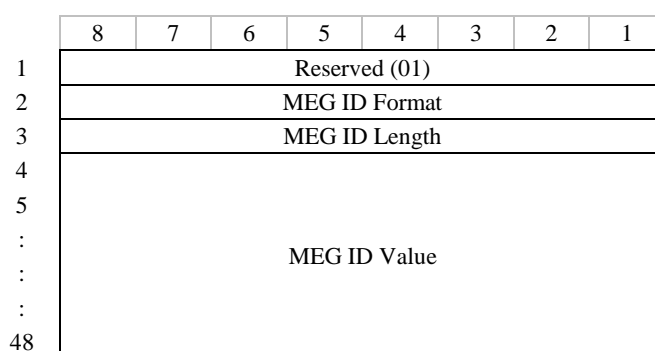


Figure A.1 – Generic MEG ID format

The MEG ID format type is identified by the MEG ID Format field. Specific values of MEG ID format type are defined in Table A.1 and described in clauses A.1 and A.2 below.

Table A.1 – MEG ID format type

MEG ID format type value	TLV name
00, 5-31, 64-255	Reserved (Note 1)
1-4	See below (Note 2)
Types specific to this Recommendation	
32	ICC-based format
33	ICC and CC based Format
34-63	Reserved (Note 3)
NOTE 1 – Reserved for definition by [IEEE 802.1].	
NOTE 2 – Use values as defined in Table 21-20 of [IEEE 802.1Q].	
NOTE 3 – Reserved for future standardization by ITU-T.	

A.1 ICC based MEG_ID format

Figure A.2 shows the format that uses the ITU carrier code (ICC). ICC is a code assigned to a network operator/service provider, maintained by the ITU Telecommunication Standardization Bureau (TSB) as per [ITU-T M.1400].

	8	7	6	5	4	3	2	1
1	Reserved (01)							
2	MEG ID Format (32)							
3	MEG ID Length (13)							
4	0	MEG ID Value[1]						
5	0	MEG ID Value[2]						
15	0	MEG ID Value[12]						
16	0	MEG ID Value[13]						
19	Unused (= all-ZEROes)							
20								
47								
48								

Figure A.2 – ICC-based MEG ID format

The MEG ID value identified by Type 32 consists of 13 characters coded according to [ITU-T T.50] (International Reference Alphabet – 7-bit coded character set for information exchange).

Note that the MEG_ID Type 32 may not be globally unique because, as described in [ITU-T M.1400], the same ICC can exist in different countries. Therefore, the MEG ID Type 32 provides uniqueness only within a country.

Figure A.3 shows the structure of an ICC-based MEG ID Value.

1	2	3	4	5	6	7	8	9	10	11	12	13
ICC		UMC										
ICC			UMC									
ICC				UMC								
ICC					UMC							
ICC						UMC						
ICC							UMC					

Figure A.3 – Structure of ICC- based MEG ID Value

It consists of two subfields: the ITU carrier code (ICC) followed by a unique MEG ID code (UMC). The ITU carrier code consists of 1-6 alphabetic (i.e., A-Z) and or numeric (i.e., 0-9), left-justified characters. The UMC code immediately follows the ICC and shall consist of 7-12 characters, with trailing NULLs, completing the 13-character MEG ID value. The UMC shall be a matter for the organization to which the ICC has been assigned, provided that uniqueness is guaranteed within a country.

A.2 Global MEG ID format based on CC and ICC

Figure A.4 shows the format that uses the ITU carrier code (ICC) with country code (CC). The MEG ID Value is identified by Type 33 and consists of 15 characters coded according to [ITU-T T.50].

Figure A.5 shows the MEG ID Value structure identified by CC and ICC. It consists of three subfields: the Country Code (CC), the ITU carrier code (ICC), followed by a unique MEG ID code (UMC). The country code (alpha-2) is a string of 2 alphabetic characters represented by upper case letters (i.e., A-Z). The country code format is defined in [ISO 3166-1]. The ITU carrier code consists of 1-6 alphabetic, (i.e., A-Z) and or numeric (i.e., 0-9), left-justified characters.

The UMC code immediately follows the ICC and shall consist of 7-12 characters, with trailing NULLs, completing the 15-character MEG ID Value. The UMC shall start with the character "/" if the ICC is less than 6 characters (as illustrated in Figure A.5) and be unique within the context of the organization to which the ITU carrier codes have been assigned.

	8	7	6	5	4	3	2	1
1	Reserved (01)							
2	MEG ID Format (33)							
3	MEG ID Length (15)							
4	0	MEG ID Value[1]						
5	0	MEG ID Value[2]						
17	0	MEG ID Value[14]						
18	0	MEG ID Value[15]						
19	Unused (= all-ZEROes)							
20								
47								
48								

Figure A.4 – CC and ICC based global MEG ID format

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
CC		ICC		/	UMC									
CC		ICC			/	UMC								
CC		ICC				/	UMC							
CC		ICC					/	UMC						
CC		ICC						/	UMC					
CC		ICC							UMC					

Figure A.5 – Structure of CC and ICC based global MEG ID Value

Annex B

Ethernet link trace (ETH-LT) of ITU-T Y.1731 interoperability considerations

(This annex forms an integral part of this Recommendation.)

This annex describes the interworking of Ethernet MEPs and MIPs, supporting different types of Ethernet link trace (ETH-LT) (i.e., ETH-LT as defined in [ITU-T Y.1731] and that are specified in this Recommendation) and identifies the basic requirements to support interworking under the ME where two types of MEPs or MIPs exist.

B.1 Ethernet link trace (ETH-LT) as defined in ITU-T Y.1731

The ETH-LT defined in [ITU-T Y.1731] differs from the one defined in this Recommendation in the following:

- LTM transmission and its PDU, as given in clauses 7.3.1 and 9.5 of [ITU-T Y.1731] do not define the LTM egress identifier TLV and its format, whereas they are defined as mandatory in this Recommendation.
- LTR transmission and its PDU, as given in clauses 7.3.2 and 9.6 of [ITU-T Y.1731] do not define the LTR egress identifier TLV and its format, whereas they are defined as mandatory in this Recommendation. Also, reply ingress TLV and reply egress TLV were optional in [ITU-T Y.1731], whereas they are defined as mandatory in this Recommendation.
- FwdYes and TerminalMEP were defined in bit 7 and bit 6 of the description of fields of the LTR PDU format in clause 9.6.2 of this Recommendation, whereas they were not defined in [ITU-T Y.1731].
- At a MIP, ETH-LT responder was not defined, and both ingress and egress ports could be set as MIP in a v2006 equipment, whereas in this Recommendation ETH-LT responder is defined so that there can only be one MIP per equipment.

B.2 Interworking with ITU-T Y.1731

In the case of an ME consisting of a v2006 MEP that transmits ETH-LTM and some v2008 MIPs, or the case of an ME consisting of a v2006 MEP that transmits ETH-LTM and the case of a v2008 MEP that receives ETH-LTM and transmits ETH-LTR, the v2008 MIP or v2008 MEP may discard ETH-LTM from the v2006 MEP due to the absence of LTM egress identifier TLV. In this case, to maintain interoperability, the v2008 MIP may forward ETH-LTM and transmit ETH-LTR by recognizing that the ETH-LTM does not have the TLV and by behaving as a v2006 MIP. Similarly, the v2008 MEP may transmit ETH-LTR by recognizing that the ETH-LTM does not have the TLV and by behaving as a v2006 MEP. See Figure B.1.




Figure B.1 – Interoperability case 1

In the case of an ME consisting of a v2008 MEP that transmits ETH-LTM and some v2006 MIPs and/or the case of a v2008 MEP that receives ETH-LTM and transmits ETH-LTR, the v2008 MEP receives ETH-LTR without LTR egress identifier TLV and without reply ingress TLV or reply egress TLV generated by v2006 MIPs and/or MEP. The absence of these TLVs in ETH-LTR is considered invalid in the v2008 version. In order to maintain interoperability, the v2008 version may be configured to identify this ETH-LTR as valid. See Figure B.2.




Figure B.2 – Interoperability case 2

In the case of an ME consisting of a v2008 MEP that transmits ETH-LTM and some v2006 MIPs located in both the ingress and egress ports of equipment, the equipment may transmit two ETH-LTRs to the v2008 MEP. In receiving the ETH-LTRs at the v2008 MEP, the behaviour is the same as in the case mentioned above (see Figure B.3). It is noted that this behaviour is compatible with the LTR analysis according to Annex J.5 of [IEEE 802.1Q], as long as each of the MPs that decrement the LTM's TTL field also return an LTR.




Figure B.3 – Interoperability case 3

Appendix I

Ethernet network scenarios

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

I.1 Shared MEG levels example

Figure I.1 provides an example scenario with the default assignment of MEG levels, where the customer, provider and operator roles share the MEG levels. In the figure, triangles represent MEPs, circles represent MIPs, and diamonds represent TrCPs.

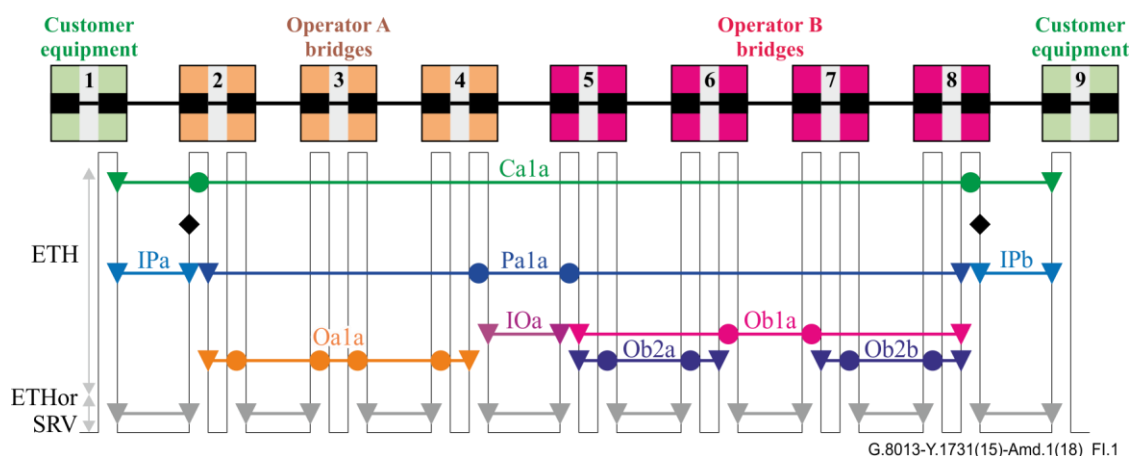


Figure I.1 – Example MEG level assignment for shared MEG levels

- Customer ME (Ca1a) can be assigned a customer MEG level 5. This allows for more customer MEs to be created at higher MEG levels, i.e., 6 and 7, if these customer MEs at additional customer MEG levels are needed.
- Provider ME (Pa1a) can be assigned a provider MEG level 4. This allows for more provider MEs to be created at a lower MEG level, i.e., 3, if additional MEs at a lower provider MEG level are needed.
- End-to-end operator MEs (Oa1a and Ob1a) can be assigned an operator MEG level 2. This allows for more operator MEs to be created at lower MEG levels, i.e., 1 and 0, if these operator MEs at additional operator MEG levels are needed in each operator network.
- Segment operator MEs in Operator B's network (Ob2a and Ob2b) can be now assigned a lower MEG level, for example 1, if Operator B needs such MEs.
- MEs between the customer and provider (IPa and IPb) can be assigned a MEG level 0. This allows providers to filter such OAM frames at UNI_N since the provider is required to provide transparency only to customer MEG levels 7, 6, and 5.
- Inter-operator ME (IOa) can be assigned a MEG level 0. This allows the operator to filter such OAM frames since the operator is required to provide transparency only to customer and provider MEG levels.

I.2 Independent MEG levels example

Figure I.2 provides an example scenario where the customer and service provider do not share the MEG levels. However, the service provider and operator share the MEG levels. In the figure, triangles represent MEPs, circles represent MIPs, and diamonds represent TrCPs.




Figure I.2 – Example MEG level assignment for independent MEG levels

- In the above example, four customer VLANs (11, 12, 21 and 22) and the corresponding customer MEGs (C-TAG, in the figure) are completely independent of the two service provider VLANs (20 and 10) and corresponding service provider MEGs (S-TAG in the Figure).
- As a consequence, the customer and service provider can independently use all eight MEG levels.
- The service provider and operator however share the MEG level space, in a manner similar to the one of Figure I.1. In this case, the eight MEG levels can be agreed mutually between the service provider and the operator.
- In the above example, the customer must send OAM frames as VLAN-tagged or priority-tagged frames to utilize all eight MEG levels independently. However, if the customer uses untagged OAM frames, the MEG levels may not be independent anymore and the customer and provider MEG levels need to be mutually agreed between the customer and the service provider.

Appendix II

Frame loss measurement

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

The four cases below should be taken into account for frame loss calculation.

- No wraparound for either the transmit or the receive counter.
- Only the transmit counter wraps around.
- Only the receive counter wraps around.
- Both transmit and receive counters wrap around.

For each case, the frame loss can be calculated as follows:

- No wrapping around for both transmit and receive counters.




Figure II.1 – No wraparound

For this case, the frame loss can be calculated by the simple calculation:

$$\text{Frame loss} = (CT2 - CT1) - (CR2 - CR1)$$

- Only transmit counter wraps around.




Figure II.2 – Transmit counter wraps around

In this case, frame loss can be calculated by the following equation, as described in the previous clause

$$\begin{aligned} \text{Frame loss} &= ((CTMAX - CT1) + CT2 + 1) - (CR2 - CR1) \\ &= (CT2 - CT1) - (CR2 - CR1) + (CTMAX + 1) \end{aligned}$$

- c) Only the receive counter wraps around.




Figure II.3 – Receive counter wraps around

$$\begin{aligned}\text{Frame loss} &= (CT2 - CT1) - ((CRMAX - CR1) + CR2 + 1) \\ &= (CT2 - CT1) - (CR2 - CR1) - (CRMAX + 1)\end{aligned}$$

- d) Both transmit and receive counters wrap around.




Figure II.4 – Both counters wrap around

$$\begin{aligned}\text{Frame loss} &= ((CTMAX - CT1) + CT2 + 1) - ((CRMAX - CR1) + CR2 + 1) \\ &= (CT2 - CT1) - (CR2 - CR1) + (CTMAX + 1) - (CRMAX + 1)\end{aligned}$$

II.1 Simplified calculation for frame loss

If the calculation is processed in an unsigned value schema, the calculation formula for the frame loss can be greatly simplified by the following characteristics:

$$N + (MAX + 1) \equiv N \text{ mod}(MAX + 1)$$

$$N - (MAX + 1) \equiv N \text{ mod}(MAX + 1)$$

Therefore, the formulas for frame loss (described in the clauses 8.1.1 and 8.1.2) can be transformed as follows.

- a) Frame loss = $\underline{(CT2 - CT1) - (CR2 - CR1)}$
- b) Frame loss = $(CT2 - CT1) - (CR2 - CR1) + CTMAX + 1$
 $= ((CT2 + (CTMAX + 1)) - CT1) - (CR2 - CR1)$
 $= \underline{(CT2 - CT1) - (CR2 - CR1)}$
- c) Frame loss = $(CT2 - CT1) - (CR2 - CR1) - (CRMAX + 1)$
 $= (CT2 - CT1) - ((CR2 + CRMAX + 1) - CR1)$
 $= \underline{(CT2 - CT1) - (CR2 - CR1)}$
- d) Frame loss = $(CT2 - CT1) - (CR2 - CR1) + (CTMAX + 1) - (CRMAX + 1)$
 $= ((CT2 + (CTMAX + 1)) - CT1) - ((CR2 + (CRMAX + 1)) - CR1)$
 $= \underline{(CT2 - CT1) - (CR2 - CR1)}$

As described above, frame loss can be calculated by the single calculation formula for any case if it is calculated in unsigned value schema.

II.2 Frame counter wrapping periodicity

This clause provides a view of wrapping periodicity of 4-octet frame counters for different interface rates and different frame sizes. The interface rates considered are 1 Gbit/s, 10 Gbit/s, and 100 Gbit/s. The frame sizes considered are 64-octet (minimum Ethernet frame size) and 1522-octet (maximum Ethernet frame size).

Table II.1 – Frame counter wrapping period

Interface rate	Frame size	4-octet frame counter wrapping period
1 Gbit/s	64-octet	$(2^{32})/((10^9)/((64+12)*8)) = 2611 \text{ seconds}$
1 Gbit/s	1522-octet	$(2^{32})/((10^9)/((1522+12)*8)) = 52707 \text{ seconds}$
10 Gbit/s	64-octet	$(2^{32})/(((10*(10^9))/((64+12)*8)) = 261 \text{ seconds}$
10 Gbit/s	1522-octet	$(2^{32})/(((10*(10^9))/((1522+12)*8)) = 5270 \text{ seconds}$
100 Gbit/s	64-octet	$(2^{32})/(((100*(10^9))/((64+12)*8)) = 26 \text{ seconds}$
100 Gbit/s	1522-octet	$(2^{32})/(((100*(10^9))/((1522+12)*8)) = 527 \text{ seconds}$

Appendix III

Network OAM interworking

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

The requirements for interworking between layered networks are the following:

- Upon detection of a defect condition in the server layer, the adaptation function between the server and client layer should be able to insert AIS in the client layer.
- The format of AIS inserted is specific to the client layer.

As an example, when the client layer is Ethernet, a server MEP is used.

Appendix IV

Mismerge detection limitation


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

MEPs consider only CCM frames with their own or lower MEG level for defect detection. CCM frames with higher MEG levels are passed through in order to provide OAM transparency as defined in clause 5.4. This behaviour leads to a limitation in the Mismerge detection as shown in Figure IV.1 below.

In case of a mismerge between MEGs with different MEG levels, the MEPs of the MEG with the lower MEG level will not detect any defect as the CCM frames coming from the MEG with the higher MEG level are passed through transparently by the MEPs. The MEPs of the MEG with the higher MEG level will detect Unexpected MEGLevel.


In case of a unidirectional mismerge from the MEG with the higher MEG level to the MEG with the lower MEG level, no defect will be detected.

No defects detected by MEPS of MEG A as only MEG Levels 3 and lower are considered!



a) Bidirectional mismerge

No defects detected by MEPS of MEG A , as only MEG levels 3 and lower are considered!



b) Unidirectional mismerge

G.8013-Y.1731(15)_FV-1

Figure IV.1 – Mismerge detection limitation

Appendix V

Terminology alignment with IEEE 802.1Q

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

The relationship of the terminology used in this Recommendation and [IEEE 802.1Q] is captured below.

Table V.1 – Terminology mapping

ITU-T G.8013/Y.1731 term	IEEE 802.1Q term	Comments
MEG	MA	
MEG ID	MAID (Domain Name + Short MA Name)	Unlike in [IEEE 802.1Q], the MEG ID does not imply a split between domain name and a short MEG name in this Recommendation.
MEG level	MD Level	

Appendix VI

Examples showing accuracy for ETH-SLM measurement

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

Synthetic loss measurement is a sampling technique for measuring frame loss; therefore, the measured FLR will be distributed around the actual loss value according to a binomial distribution. The mean measured FLR will always be equal to the actual FLR, while the standard deviation depends on the number of samples. The standard deviation can therefore be used to illustrate the accuracy of the measured FLR result. Table VI.1 shows the standard deviation for various real loss values and numbers of samples (i.e., number of SLM frames sent). When ETH-SLM is used, the number of samples should be chosen so that the standard deviation is low, when compared to any FLR threshold that is being used to trigger an action. This ensures the chance of false positives is low.

Table VI.1 – Standard deviation for various real loss values and number of samples

Actual FLR	Number of samples	Transmission interval	Std. Dev. (FLR % points)
50%	10	100 ms	15.81%
50%	100	10 ms	5.00%
50%	1000	1 ms	1.58%
10%	10	100 ms	9.49%
10%	100	10 ms	3.00%
10%	1000	1 ms	0.95%
1%	10	100 ms	3.15%
1%	100	10 ms	0.99%
1%	1000	1 ms	0.31%
0.1%	10	100 ms	1.00%
0.1%	100	10 ms	0.31%
0.1%	1000	1 ms	0.1%

Note that if the number of samples is increased by a factor of n , the standard deviation is reduced by a factor of \sqrt{n} .

Appendix VII

ETH-LM and link aggregation

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

Link aggregation (LAG), as specified in [b-IEEE 802.1AX], may impact the effectiveness of OAM mechanisms specified in this Recommendation and in [ITU-T G.8021] and [IEEE 802.1Q]. These OAM mechanisms based on service frames such as ETH-LM require frame ordering preservation while those based on synthetic frames such as ETH-DM and ETH-SLM (and ETH-CC) presume suitable sampling of all feasible transmission links/paths. Though this appendix focuses on ETH-LM, other OAM mechanisms may encounter similar issues when they monitor a fraction of the intended flow. These issues can be avoided for example if LAG is used for protection switching (i.e., a LAG with two aggregated links in which all traffic is forwarded onto the active transport entity) or if LAG is used with flow-aware hashing (i.e., all traffic in a given flow is placed on the same aggregated link).

Specifically considering Ethernet frame loss measurement, the ETH-LM mechanism is in principle capable of accurately detecting single frame loss events on a point-to-point ETH connection between two terminating MEPs (e.g., MEPs A and Z in Figure VII.1 illustrating the scenario discussed hereafter). However, this accuracy may be affected by frame reordering on the ETH connection. The ordering of ETH-LM PDUs relative to the frames that are counted is important.

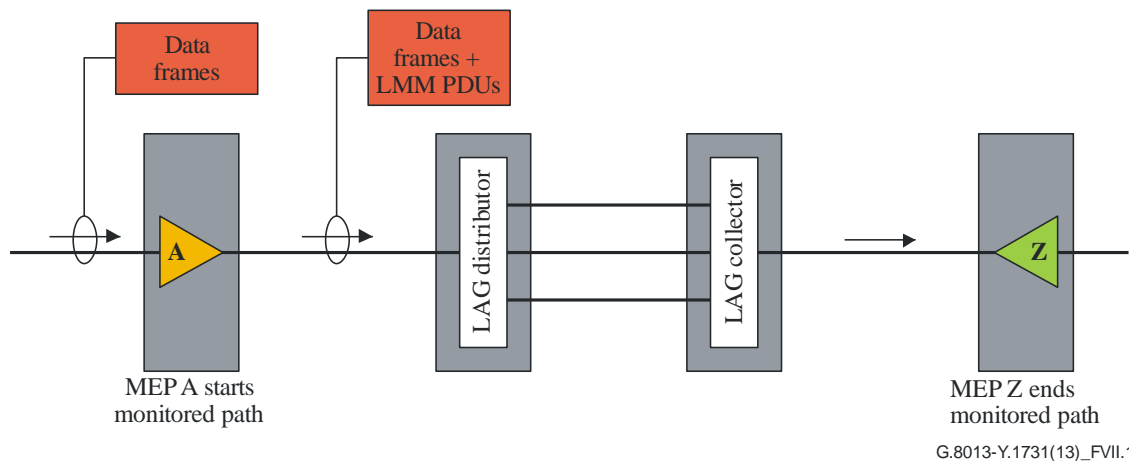


Figure VII.1 – Path monitored for frame loss between two terminating MEPs

The ETH-LM measurement method is based on the assumption that the position of ETH-LM PDUs in the flow of counted frames stays the same between source and sink MEPs. This provides the required synchronization between the counters at both ends of the link. It is a characteristic property of the forwarding in Ethernet bridges to preserve the frame ordering of the MAC service. Some implementations of link aggregation (LAG) however may not guarantee frame ordering preservation over the full aggregated bandwidth. LAG avoids frame reordering by assigning all frames in a given "conversation" to a single aggregated link. This ensures that frame ordering is maintained within each "conversation", but not necessarily between "conversations".

Common implementations of the LAG frame distributor function (the "Distributor") operate largely autonomously and detect "conversations" by hashing not necessarily only on the VLAN identifier (VID) and priority of the exchanged ETH-LM PDUs and counted frames, but rather for example on source and destination MAC and/or IP addresses. The set of ETH-LM PDUs and the frames they are supposed to count will generally contain a variety of values in the fields on which the Distributor bases the assigned hash-value and hence the aggregated link assignment.

Assuming a LAG section/aggregation is to be traversed somewhere along the path between the two terminating MEPs, ETH-LM PDUs and counted frames may be forwarded onto different aggregated links. This takes place even if they are all transmitted in the same VID and at the same priority because the Distributor may consider more frame fields to decide aggregated link assignment. The counted frames themselves may be spread over different aggregated links if they belong to different "conversations". Reordering may also depend on other factors such as the amount of traffic on the LAG section, the variety of frame lengths, or the number of "conversations" that the Distributor can detect.

The LAG frame collector function (the "Collector") is relatively simple compared to the Distributor as it relies on the latter for frame ordering (within "conversations"). As such, it simply passes frames received from aggregated links in the order that it receives them. Because of this, frames with the same VID and at the same priority that were forwarded onto different aggregated links by the Distributor are not reordered by the Collector and are likely to be in a different order before and after traversing the LAG section.

The sink MEP reads its local counter exactly when an LMM PDU is received and compares this reading with the counter in the LMM PDU itself, providing the equivalent count from the source MEP. As illustrated in Figure VII.2, if the LMM PDU shifts position relative to the frames that surround it, i.e., the ones subject to counting, such a comparison will indicate artificial frame loss (or gain), even if there is no actual frame loss (or gain) in reality. This puts a limit on the accuracy that can be achieved with this method of measuring frame loss.

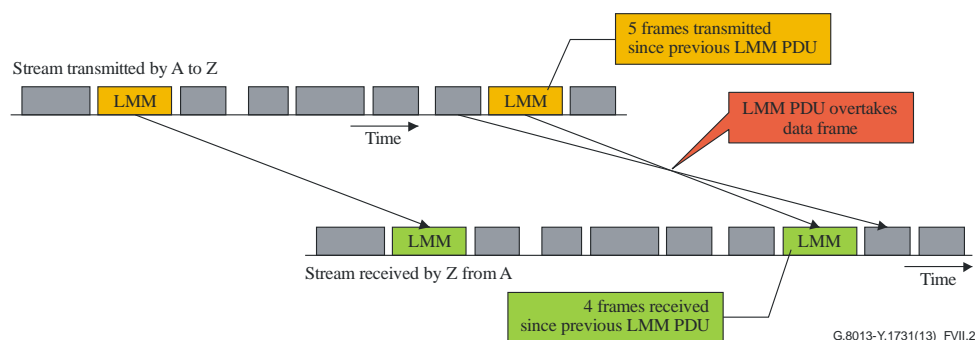


Figure VII.2 – LMM PDU overtakes data frame causing artificial frame loss (or gain)

Due to the many factors affecting ordering on a LAG section, it is difficult to predict how often such errors occur. The error is likely to be plus or minus a few frames. Since ETH-LM PDUs are short they tend to overtake longer frames on a LAG section. Hence an artificial frame loss may be measured before the compensating artificial frame gain is measured. Also, there may be measurement intervals in which there is very little end-user traffic (e.g., on standby connections). In such intervals the (relative) error in the reported frame loss rate caused by reordering can increase significantly. Note that a LAG section typically handles a lot more traffic than just the flow measured by ETH-LM, so the probability of reordering may not depend much on the amount of traffic in the ETH-LM monitored flow itself.

In practice, because the service-frame counters are continuously running, the artificial frame loss or gain cancels out at the next LMR but may be replaced by a new error if misordering continues. If the last LMM and LMR PDUs that are used in a given measurement interval (typically 15-minute or 24-hour long) are both not subject to reordering, any errors made up to that point are compensated. When the measurement interval is long, the errors may be small compared to the number of service frames. However, there may only be a few service frames in the small time intervals used to assess availability, so that the reordering error may be sufficient to cause false or missed unavailability FLR threshold crossings, leading to incorrect unavailability time.

Bibliography

- [b-IEEE 802.1AX] IEEE 802.1AX (2014), *IEEE Standard for Local and metropolitan area networks: Link Aggregation*.
- [b-IETF RFC 2544] IETF RFC 2544 (1999), *Benchmarking Methodology for Network Interconnect Devices*.
<http://www.ietf.org/rfc/rfc2544.txt>

ITU-T Y-SERIES RECOMMENDATIONS

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

GLOBAL INFORMATION INFRASTRUCTURE	
General	Y.100–Y.199
Services, applications and middleware	Y.200–Y.299
Network aspects	Y.300–Y.399
Interfaces and protocols	Y.400–Y.499
Numbering, addressing and naming	Y.500–Y.599
Operation, administration and maintenance	Y.600–Y.699
Security	Y.700–Y.799
Performances	Y.800–Y.899
INTERNET PROTOCOL ASPECTS	
General	Y.1000–Y.1099
Services and applications	Y.1100–Y.1199
Architecture, access, network capabilities and resource management	Y.1200–Y.1299
Transport	Y.1300–Y.1399
Interworking	Y.1400–Y.1499
Quality of service and network performance	Y.1500–Y.1599
Signalling	Y.1600–Y.1699
Operation, administration and maintenance	Y.1700–Y.1799
Charging	Y.1800–Y.1899
IPTV over NGN	Y.1900–Y.1999
NEXT GENERATION NETWORKS	
Frameworks and functional architecture models	Y.2000–Y.2099
Quality of Service and performance	Y.2100–Y.2199
Service aspects: Service capabilities and service architecture	Y.2200–Y.2249
Service aspects: Interoperability of services and networks in NGN	Y.2250–Y.2299
Enhancements to NGN	Y.2300–Y.2399
Network management	Y.2400–Y.2499
Network control architectures and protocols	Y.2500–Y.2599
Packet-based Networks	Y.2600–Y.2699
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899
Carrier grade open environment	Y.2900–Y.2999
FUTURE NETWORKS	Y.3000–Y.3499
CLOUD COMPUTING	Y.3500–Y.3999
INTERNET OF THINGS AND SMART CITIES AND COMMUNITIES	
General	Y.4000–Y.4049
Definitions and terminologies	Y.4050–Y.4099
Requirements and use cases	Y.4100–Y.4249
Infrastructure, connectivity and networks	Y.4250–Y.4399
Frameworks, architectures and protocols	Y.4400–Y.4549
Services, applications, computation and data processing	Y.4550–Y.4699
Management, control and performance	Y.4700–Y.4799
Identification and security	Y.4800–Y.4899
Evaluation and assessment	Y.4900–Y.4999

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

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	Tariff and accounting principles and international telecommunication/ICT economic and policy issues
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	Environment and ICTs, climate change, e-waste, energy efficiency; 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, and associated measurements and tests
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, next-generation networks, Internet of Things and smart cities
Series Z	Languages and general software aspects for telecommunication systems