

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

G.8013/Y.1731

(08/2015)

**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
AND NEXT-GENERATION NETWORKS**

**Internet protocol aspects – Operation, administration and
maintenance**

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

Recommendation ITU-T G.8013/Y.1731

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

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.

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

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

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure, e.g., interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

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

© ITU 2016

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

Table of Contents

		Page
1	Scope.....	1
2	References.....	1
3	Definitions	2
	3.1 Terms defined elsewhere	2
	3.2 Terms defined in this Recommendation.....	4
4	Abbreviations and acronyms	4
5	Conventions	6
	5.1 ME group (MEG)	7
	5.2 Traffic conditioning point (TrCP)	7
	5.3 MEG level	7
	5.4 OAM transparency	7
	5.5 Representation of octets	8
6	OAM relationships.....	8
	6.1 MEs, MEPs, MIPs and TrCPs relationship	8
	6.2 MEs, MEGs and MEG level relationship.....	8
	6.3 MEPs and MIPs configuration	10
7	OAM functions for fault management.....	10
	7.1 Ethernet continuity check (ETH-CC).....	10
	7.2 Ethernet loopback (ETH-LB)	12
	7.3 Ethernet link trace (ETH-LT)	15
	7.4 Ethernet alarm indication signal (ETH-AIS).....	17
	7.5 Ethernet remote defect indication (ETH-RDI)	19
	7.6 Ethernet locked signal (ETH-LCK).....	20
	7.7 Ethernet test signal (ETH-Test).....	21
	7.8 Ethernet automatic protection switching (ETH-APS).....	22
	7.9 Ethernet maintenance communication channel (ETH-MCC)	22
	7.10 Ethernet experimental OAM (ETH-EXP)	23
	7.11 Ethernet vendor-specific OAM (ETH-VSP)	23
	7.12 Ethernet client signal fail (ETH-CSF).....	24
	7.13 Ethernet bandwidth notification (ETH-BN).....	25
	7.14 Ethernet expected defect function (ETH-ED)	27
8	OAM functions for performance monitoring	28
	8.1 Frame loss measurement (ETH-LM).....	29
	8.2 Frame delay measurement (ETH-DM).....	32
	8.3 Throughput measurement	35
	8.4 Synthetic loss measurement (ETH-SLM)	35
9	OAM PDU types.....	38
	9.1 Common OAM information elements.....	39

	Page
9.2 CCM PDU	41
9.3 LBM PDU	43
9.4 LBR PDU	45
9.5 LTM PDU.....	46
9.6 LTR PDU.....	48
9.7 AIS PDU.....	50
9.8 LCK frame.....	51
9.9 TST PDU	51
9.10 APS PDU.....	52
9.11 MCC PDU	53
9.12 LMM PDU.....	54
9.13 LMR PDU	55
9.14 1DM PDU.....	56
9.15 DMM PDU	57
9.16 DMR PDU	58
9.17 EXM PDU	59
9.18 EXR PDU	60
9.19 VSM PDU	61
9.20 VSR PDU	62
9.21 Client signal fail (CSF).....	63
9.22 SLM PDU	64
9.23 SLR PDU.....	65
9.24 1SL PDU	66
9.25 BNM PDU	67
9.26 EDM PDU	69
10 OAM frame addresses	69
10.1 Multicast destination addresses	70
10.2 CCM	70
10.3 LBM	70
10.4 LBR	70
10.5 LTM.....	70
10.6 LTR.....	71
10.7 AIS.....	71
10.8 LCK	71
10.9 TST	71
10.10 APS.....	71
10.11 MCC	71
10.12 LMM.....	71
10.13 LMR	71
10.14 1DM.....	71

	Page
10.15 DMM	71
10.16 DMR	71
10.17 EXM	71
10.18 EXR	72
10.19 VSM	72
10.20 VSR	72
10.21 CSF	72
10.22 SLM	72
10.23 SLR	72
10.24 1SL	72
10.25 BNM	72
10.26 EDM	72
11 OAM PDU validation and versioning	73
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	76
A.2 Global MEG ID format based on CC and ICC	77
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	81
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 include 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.

vi

Recommendation ITU-T G.8013/Y.1731

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

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.8001] Recommendation ITU-T G.8001/Y.1354 (2013), *Terms and definitions for Ethernet frames over transport*.
- [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 (2015), *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 (2012), *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-2014, *IEEE Standard for Local and metropolitan area networks – Bridges and Bridged Networks*
<<http://standards.ieee.org/findstds/standard/802.1Q-2014.html>>
- [IEEE 802.3] IEEE 802.3-2012, *IEEE Standard for Ethernet*.
<<http://standards.ieee.org/findstds/standard/802.3-2012.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*.
<http://www.metroethernetforum.org/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 **dual-ended:** [ITU-T G.8001].
- 3.1.8 **failure:** [ITU-T G.806].
- 3.1.9 **far-end:** [ITU-T G.8001].
- 3.1.10 **flow:** [ITU-T G.809].
- 3.1.11 **flow domain:** [ITU-T G.809].
- 3.1.12 **flow domain flow:** [ITU-T G.809].
- 3.1.13 **flow point:** [ITU-T G.809].
- 3.1.14 **flow point pool:** [ITU-T G.809].
- 3.1.15 **flow point pool link:** [ITU-T G.809].
- 3.1.16 **flow termination:** [ITU-T G.809].
- 3.1.17 **flow termination sink:** [ITU-T G.809].
- 3.1.18 **flow termination source:** [ITU-T G.809].
- 3.1.19 **initiating MEP:** [ITU-T G.8001].
- 3.1.20 **in-profile:** [ITU-T G.8001].
- 3.1.21 **in-service OAM:** [ITU-T G.8001].
- 3.1.22 **layer network:** [ITU-T G.809].
- 3.1.23 **link:** [ITU-T G.805].
- 3.1.24 **link connection:** [ITU-T G.805].
- 3.1.25 **link flow:** [ITU-T G.809].
- 3.1.26 **maintenance entity:** [ITU-T G.8001].
- 3.1.27 **maintenance entity group:** [ITU-T G.8001].
- 3.1.28 **MEG end-point (MEP):** [ITU-T G.8001].
- 3.1.29 **MEG intermediate point (MIP):** [ITU-T G.8001].
- 3.1.30 **near-end:** [ITU-T G.8001].
- 3.1.31 **network:** [ITU-T G.809].
- 3.1.32 **network connection:** [ITU-T G.805].
- 3.1.33 **on-demand OAM:** [ITU-T G.8001].
- 3.1.34 **one-way:** [ITU-T G.8001].
- 3.1.35 **organizationally unique identifier:** [IEEE 802].
- 3.1.36 **out-of-service OAM:** [ITU-T G.8001].
- 3.1.37 **peer MEP:** [ITU-T G.8001].
- 3.1.38 **port:** [ITU-T G.809].



- 3.1.39 proactive OAM:** [ITU-T G.8001].
- 3.1.40 receiving MEP:** [ITU-T G.8001].
- 3.1.41 reference point:** [ITU-T G.809].
- 3.1.42 responding MEP:** [ITU-T G.8001].
- 3.1.43 server MEP:** [ITU-T G.8001].
- 3.1.44 single-ended:** [ITU-T G.8001].
- 3.1.45 termination connection point:** [ITU-T G.805].
- 3.1.46 termination flow point:** [ITU-T G.809].
- 3.1.47 traffic unit:** [ITU-T G.809].
- 3.1.48 trail:** [ITU-T G.805].
- 3.1.49 trail termination:** [ITU-T G.805].
- 3.1.50 transport:** [ITU-T G.809].
- 3.1.51 transport entity:** [ITU-T G.809].
- 3.1.52 transport processing function:** [ITU-T G.809].
- 3.1.53 two-way** [ITU-T G.8001].

3.2 Terms defined in this Recommendation

None.

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 Figure 23 and Figure 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 including 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.

- 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.

.....