

International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Y.1731

(02/2008)

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

Internet protocol aspects – Operation, administration and
maintenance

**OAM functions and mechanisms for Ethernet
based networks**

Recommendation ITU-T Y.1731



ITU-T Y-SERIES RECOMMENDATIONS
**GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT-
GENERATION NETWORKS**

GLOBAL INFORMATION INFRASTRUCTURE	
GENERAL	Y.100–Y.199
SERVICES, APPLICATIONS AND MIDDLEWARE	Y.200–Y.299
NETWORK ASPECTS	Y.300–Y.399
INTERFACES AND PROTOCOLS	Y.400–Y.499
NUMBERING, ADDRESSING AND NAMING	Y.500–Y.599
OPERATION, ADMINISTRATION AND MAINTENANCE	Y.600–Y.699
SECURITY	Y.700–Y.799
PERFORMANCES	Y.800–Y.899
INTERNET PROTOCOL ASPECTS	
GENERAL	Y.1000–Y.1099
SERVICES AND APPLICATIONS	Y.1100–Y.1199
ARCHITECTURE, ACCESS, NETWORK CAPABILITIES AND RESOURCE MANAGEMENT	Y.1200–Y.1299
TRANSPORT	Y.1300–Y.1399
INTERWORKING	Y.1400–Y.1499
QUALITY OF SERVICE AND NETWORK PERFORMANCE	Y.1500–Y.1599
SIGNALLING	Y.1600–Y.1699
OPERATION, ADMINISTRATION AND MAINTENANCE	Y.1700–Y.1799
CHARGING	Y.1800–Y.1899
NEXT GENERATION NETWORKS	
FRAMEWORKS AND FUNCTIONAL ARCHITECTURE MODELS	Y.2000–Y.2099
QUALITY OF SERVICE AND PERFORMANCE	Y.2100–Y.2199
SERVICE ASPECTS: SERVICE CAPABILITIES AND SERVICE ARCHITECTURE	Y.2200–Y.2249
SERVICE ASPECTS: INTEROPERABILITY OF SERVICES AND NETWORKS IN NGN	Y.2250–Y.2299
NUMBERING, NAMING AND ADDRESSING	Y.2300–Y.2399
NETWORK MANAGEMENT	Y.2400–Y.2499
NETWORK CONTROL ARCHITECTURES AND PROTOCOLS	Y.2500–Y.2599
SECURITY	Y.2700–Y.2799
GENERALIZED MOBILITY	Y.2800–Y.2899

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

Recommendation ITU-T Y.1731

OAM functions and mechanisms for Ethernet based networks

Summary

Recommendation ITU-T 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 operation, administration and maintenance mechanisms defined in this Recommendation offer capabilities to operate and maintain the network and service aspects of the ETH layer.

Source

Recommendation ITU-T Y.1731 was approved on 29 February 2008 by ITU-T Study Group 13 (2005-2008) under Recommendation ITU-T A.8 procedure.

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 2009

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

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 Maintenance entity (ME).....	6
5.2 ME group (MEG)	7
5.3 MEG end point (MEP)	7
5.4 MEG intermediate point (MIP)	8
5.5 Traffic conditioning point (TrCP)	8
5.6 MEG level	8
5.7 OAM transparency	9
5.8 Representation of octets	9
6 OAM relationships	9
6.1 MEs, MEPs, MIPs and TrCPs relationship	9
6.2 MEGs and MEG level relationship	10
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).....	20
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).....	23
7.11 Ethernet vendor-specific OAM (ETH-VSP)	24
8 OAM functions for performance monitoring	24
8.1 Frame loss measurement (ETH-LM).....	25
8.2 Frame delay measurement (ETH-DM).....	27
8.3 Throughput measurement.....	30
9 OAM PDU types	30
9.1 Common OAM information elements.....	30
9.2 CCM PDU	32

	Page
9.3	LBM PDU 34
9.4	LBR PDU 36
9.5	LTM PDU 37
9.6	LTR PDU 39
9.7	AIS PDU 41
9.8	LCK PDU 42
9.9	TST PDU 43
9.10	APS PDU 44
9.11	MCC PDU 45
9.12	LMM PDU 46
9.13	LMR PDU 46
9.14	1DM PDU 47
9.15	DMM PDU 48
9.16	DMR PDU 49
9.17	EXM PDU 50
9.18	EXR PDU 51
9.19	VSM PDU 52
9.20	VSR PDU 53
10	OAM frame addresses 54
10.1	Multicast destination addresses 55
10.2	CCM 55
10.3	LBM 56
10.4	LBR 56
10.5	LTM 56
10.6	LTR 56
10.7	AIS 56
10.8	LCK 56
10.9	TST 56
10.10	APS 56
10.11	MCC 56
10.12	LMM 56
10.13	LMR 56
10.14	1DM 57
10.15	DMM 57
10.16	DMR 57
10.17	EXM 57
10.18	EXR 57
10.19	VSM 57
10.20	VSR 57

	Page
Annex A – MEG ID format	58
Appendix I – Defect conditions	60
I.1 Loss of continuity (LOC) condition	60
I.2 Mismerge condition.....	60
I.3 Unexpected MEP (UnexpectedMEP) condition.....	61
I.4 Unexpected MEG level (UnexpectedMEGLevel) condition	61
I.5 Unexpected period (UnexpectedPeriod) condition	61
I.6 Signal fail (SignalFail) condition	62
I.7 AIS condition.....	62
I.8 RDI condition	63
I.9 LCK condition.....	63
Appendix II – Ethernet network scenarios.....	64
II.1 Shared MEG levels example	64
II.2 Independent MEG levels example.....	65
Appendix III – Frame loss measurement	66
III.1 Frame loss calculations.....	66
III.2 Frame counter wrapping periodicity	68
Appendix IV – Network OAM interworking.....	69
Appendix V – Mismerge detection limitation.....	70
Appendix VI – Terminology alignment with IEEE 802.1ag	71
Bibliography.....	72

Introduction

Study Group 13 has prepared this Recommendation 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 Study Group 15.

Recommendation ITU-T Y.1731

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. 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 and/or operational support systems).

The architectural basis for this Recommendation is the Ethernet specification [ITU-T G.8010] which also accounts for [IEEE 802.1D], [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 (2006), *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 (2007), *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 (2007), *Characteristics of Ethernet transport network equipment functional blocks*.
- [ITU-T G.8031] Recommendation ITU-T G.8031/Y.1342 (2006), *Ethernet linear protection switching*.
- [ITU-T M.1400] Recommendation ITU-T M.1400 (2006), *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.1730] Recommendation ITU-T Y.1730 (2004), *Requirements for OAM functions in Ethernet-based networks and Ethernet services.*
- [IEC 61588] IEC 61588 (2004), *Precision clock synchronization protocol for networked measurement and control systems.*
<<http://webstore.iec.ch/webstore/webstore.nsf/artnum/033151>>
- [IEEE 802] IEEE 802-2001, *IEEE Standard for Local and Metropolitan area Networks: Overview and Architecture.*
<<http://ieeexplore.ieee.org/servlet/opac?punumber=7732>>
- [IEEE 802.1ag] IEEE 802.1ag-2007, *IEEE Standard for Local and Metropolitan area Networks: Virtual Bridged Local Area Networks Amendment 5: Connectivity Fault Management.*
<<http://ieeexplore.ieee.org/servlet/opac?punumber=4431834>>
- [IEEE 802.1D] IEEE 802.1D-2004, *IEEE Standard for Local and Metropolitan area Networks: Media Access Control (MAC) Bridges.*
<<http://ieeexplore.ieee.org/servlet/opac?punumber=9155>>
- [IEEE 802.1Q] IEEE 802.1Q-2005, *IEEE Standard for Local and Metropolitan area Networks: Virtual Bridged Local Area Networks.*
<<http://ieeexplore.ieee.org/servlet/opac?punumber=10905>>
- [IEEE 802.3] IEEE 802.3-2002, *Information Technology – Telecommunications and Information Exchange Between Systems – LAN/MAN – Specific Requirements – Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications.*
<<http://ieeexplore.ieee.org/servlet/opac?punumber=10531>>
- [MEF 10] MEF 10 (2004), *Ethernet Services Attributes: Phase 1.*
<http://metroethernetforum.org/MSWord_Documents/MEF10.doc>

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 ETH link:** [ITU-T G.8010]
- 3.1.8 ETH trail:** [ITU-T G.8010]

- 3.1.9 **failure:** [ITU-T G.806]
- 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 **layer network:** [ITU-T G.809]
- 3.1.20 **link:** [ITU-T G.805]
- 3.1.21 **link connection:** [ITU-T G.805]
- 3.1.22 **link flow:** [ITU-T G.809]
- 3.1.23 **multipoint Ethernet connection:** [ITU-T G.8010]
- 3.1.24 **multipoint Ethernet connectivity:** [ITU-T G.8010]
- 3.1.25 **network:** [ITU-T G.809]
- 3.1.26 **network connection:** [ITU-T G.805]
- 3.1.27 **network operator:** [ITU-T G.805]
- 3.1.28 **organizationally unique identifier:** [IEEE 802]
- 3.1.29 **point-to-point Ethernet connection:** [ITU-T G.8010]
- 3.1.30 **port:** [ITU-T G.809]
- 3.1.31 **reference point:** [ITU-T G.809]
- 3.1.32 **service provider:** [ITU-T G.805]
- 3.1.33 **termination connection point:** [ITU-T G.805]
- 3.1.34 **termination flow point:** [ITU-T G.809]
- 3.1.35 **termination flow point pool:** [ITU-T G.809]
- 3.1.36 **traffic unit:** [ITU-T G.809]
- 3.1.37 **transport:** [ITU-T G.809]
- 3.1.38 **transport entity:** [ITU-T G.809]
- 3.1.39 **transport processing function:** [ITU-T G.809]
- 3.1.40 **trail:** [ITU-T G.805]
- 3.1.41 **trail termination:** [ITU-T G.805]

3.2 Terms defined in this Recommendation

This Recommendation defines the following terms:

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

3.2.2 on-demand OAM: On-demand OAM refers to OAM actions which 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.

3.2.3 out-of-service OAM: Out-of-service OAM refers to OAM actions which are carried out while the data traffic is interrupted.

3.2.4 proactive OAM: Proactive OAM refers to OAM actions which are carried out continuously to permit proactive reporting of fault and/or performance results.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

IDM	One-way Delay Measurement
AIS	Alarm Indication Signal
AP	Access Point
APS	Automatic Protection Switching
CCM	Continuity Check Message
CE	Customer Edge
CoS	Class of Service
CP	Connection Point
DA	Destination MAC Address
DMM	Delay Measurement Message
DMR	Delay Measurement Reply
ETH	Ethernet MAC layer network
ETH-AIS	Ethernet Alarm Indication Signal function
ETH-APS	Ethernet Automatic Protection Switching function
ETH-CC	Ethernet Continuity Check function
ETH-DM	Ethernet Delay Measurement function
ETH-EXP	Ethernet Experimental OAM function
ETH_FP	Ethernet Flow Point
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-Test	Ethernet Test function
ETH-TFP	Ethernet Termination Flow Point
ETH-VSP	Ethernet Vendor-Specific OAM function
ETY	Ethernet PHY layer network
EXM	Experimental operations, administration and management Message
EXR	Experimental operations, administration and management Reply
FD	Flow Domain
FP	Flow Point
FPP	Flow Point Pool
FT	Flow Termination
ICC	ITU Carrier Code
LBM	LoopBack Message
LBR	LoopBack Reply
LCK	Locked
LMI	Local Management Interface
LMM	Loss Measurement Message
LMR	Loss Measurement Reply
LOC	Loss of Continuity
LTM	Link Trace Message
LTR	Link Trace Reply
MAC	Media Access Control
MC	Media Converter
MCC	Maintenance Communication Channel
ME	Maintenance Entity
MEG	Maintenance Entity Group
MEL	MEG Level
MEP	MEG End Point
MIB	Management Information Base
MIP	MEG Intermediate Point
NMS	Network Management System
NNI	Network Node Interface
NT	Network Termination
OAM	Operations, Administration and Maintenance
OSS	Operations Support System
OTN	Optical Transport Network
OUI	Organizationally Unique Identifier
PDU	Protocol Data Unit

PE	Provider Edge
PHY	Ethernet PHYSical layer entity
PRBS	Pseudo-Random Bit Sequence
RDI	Remote Defect Indication
SA	Source MAC Address
SES	Severely Errored Seconds
SLA	Service level Agreement
SRV	Server
STP	Spanning Tree Protocol
TC	Traffic Conditioning
TCI	Tag Control Information
TFP	Termination Flow Point
TFPP	Termination Flow Point Pool
TLV	Type, Length, 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
VID	VLAN Identifier
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 Maintenance entity (ME)

ME represents an entity that requires management and is a relationship between two maintenance entity group end points (see clause 5.3). MEs in Ethernet networks have been identified in Figure 23 of [ITU-T G.8010] (see Figure 5-1), Figure 24 of [ITU-T G.8010] and clause 9 of [ITU-T Y.1730]. MEs can nest but not overlap.

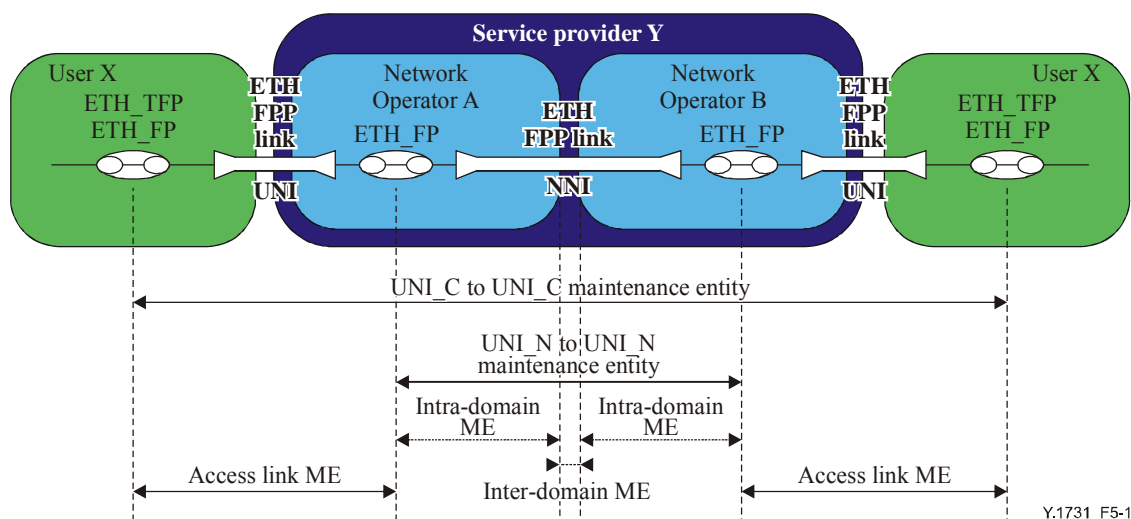


Figure 5-1 – Example of MEs associated with a point-to-point connection administrative domain shown in Figure 23 of [ITU-T G.8010]

The mapping of the MEs as defined in [ITU-T G.8010] and [ITU-T Y.1730] is shown in Table 5-1.

Table 5-1 – MEs as defined in [ITU-T G.8010] and [ITU-T Y.1730]

G.8010 ME	Y.1730 ME
UNI_C to UNI_C ME	UNI-UNI (customer)
UNI_N to UNI_N ME	UNI-UNI (provider)
Intra-domain ME	Segment (PE-PE) intra-provider
Inter-domain ME	Segment (PE-PE) inter-provider (provider-provider)
Access-link ME	ETY link OAM-UNI (customer-provider)
Inter-domain ME	ETY link OAM-NNI (operator-operator)

5.2 ME group (MEG)

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

- MEs in a MEG exist in the same administrative boundary; and
- MEs in a MEG have the same MEG level (see clause 5.6); and
- MEs in a MEG belong to the same point-to-point ETH connection or multipoint ETH connectivity.

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

5.3 MEG end point (MEP)

MEG end point (MEP) marks the end point of an ETH MEG which is capable of initiating and terminating OAM frames for fault management and performance monitoring. The OAM frames are distinct from the transit ETH flows. The OAM frames are added to the aggregate of transit ETH flows and it is assumed that they are subject to the same forwarding treatment as the transit ETH flows being monitored. 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).

A MEP may be described using atomic functions, as per [ITU-T G.8021], that are outside the scope of this Recommendation.

5.3.1 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 upon failure 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 the Server layer.

NOTE – A Server MEP needs to support ETH-AIS function, as described in clause 7.4, where the Server/ETH adaptation function is required to issue frames with ETH-AIS information upon detection of defect at the Server layer by the Server layer termination and/or adaptation function.

A Server MEP may be described using atomic functions, as per [ITU-T G.8021], that are outside the scope of this Recommendation.

5.4 MEG intermediate point (MIP)

MEG intermediate point (MIP) is an intermediate point in a MEG which 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.

A MIP may be described using atomic functions, as per [ITU-T G.8021], that are outside the scope of this Recommendation.

5.5 Traffic conditioning point (TrCP)

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].

5.6 MEG level

In case 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:

- Customer role is assigned three MEG levels: 7, 6 and 5.
- Provider role is assigned two MEG levels: 4 and 3.
- 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 MEGs for which the OAM flows are not distinguishable based on means of the ETH layer encapsulation.

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

5.7 OAM transparency

OAM transparency refers to the ability to allow 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 MEGs to pass transparently; while it blocks OAM frames from outside administrative domains belonging to same or lower level MEGs.

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 5.6. 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 provider's OAM frames across 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.

5.8 Representation of octets

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

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 5-2 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).




Figure 5-2 – Example PDU format

6 OAM relationships

6.1 MEGs, MEPs, MIPs and TrCPs relationship

Appendix II 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 II. For example, providers may not provide customer MIPs.

6.2 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, such 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.

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

G.8010 MEG	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 5.6, 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

G.8010 MEG	Y.1730 ME	MEG level(s)
UNI_C to UNI_C ME	UNI-UNI (customer)	7-1
UNI_N to UNI_N ME	UNI-UNI (provider)	7-1
Intra-domain ME	Segment (PE-PE) intra-provider	7-1
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)

Further, if inter-domain MEs are not required, each customer, provider and operator can use all eight MEG levels. However, as already stated in clause 5.6, not all MEG levels may be used.

6.3 MEPs and MIPs configuration

MEPs and 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 (NMSs).

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.

Appendix I provides an overview of different defect conditions that can be detected using OAM functions.

7.1 Ethernet continuity check (ETH-CC)

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 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 all other 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.

Specific configuration information required by each MEP to support ETH-CC is the following:

- 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 3 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 frames 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 3 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 7 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 of transmission period configured at the transmitting MEP, so that a receiving MEP can detect 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 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 same MEG level but with a MEG ID different from the receiving MEP's own MEG ID is received, Mismatch is detected.
- If a CCM frame with the same MEG level and a correct MEG ID but with an incorrect MEP ID, including 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 than 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)

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, at 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 the following:

- 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 application, 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 the following:

- 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 the remote MIP or remote 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 remote 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 in a direction opposite to the direction where LBM frames are issued.

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 requesting MEP. A Unicast LBM frame with a valid MEG level and a destination MAC address equal to the MAC address of receiving 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.

Further, when a receiving 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 in a direction opposite to the direction where LBR frames are issued.

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

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 ETH-LB is invoked on a MEP, a Multicast frame with ETH-LB request information is sent from a MEP to other peer MEPs in the same MEG. The MEP expects to receive 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 the following:

- 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 as 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 is 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 requesting 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 replying 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)

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

- Adjacent relation Retrieval – ETH-LT function can be used to retrieve adjacency relationship between a MEP and a remote MEP or MIP. The result of running ETH-LT function is a sequence of MIPs from the source MEP until the target MIP or MEP. Each MIP and/or MEP is identified by its MAC address.
- Fault localization – 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.

Specific configuration information required by a MEP to support ETH-LT is the following:

- 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.

Specific configuration information required by a MIP to support ETH-LT is the following:

- 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 for the MIP or MEP to reply, 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

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. LTM frame contains an LTM egress identifier TLV which identifies the network element initiating the LTM frame.

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, forwarding and LTR transmission

If an LTM frame is received by a MEP or MIP, it forwards the LTM frame to 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.

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 originating MEP after a random time interval in the range of 0 to 1 second.
- Further, if the above condition applies and LTM frame does not terminate at the MIP or MEP (when the TargetMAC address is not the same as the MIP's own address, if received by a MIP, or when received by a MEP) and the TTL field in 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.

The LTR frame contains LTR egress identifier TLV which identifies the source and destination of the LTM that triggered transmission of this LTR. LTR egress identifier TLV contains 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. LTR egress identifier TLV also contains 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 if 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.

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)

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 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 AIS condition.

Specific configuration information required by a MEP to support ETH-AIS transmission is the following:

- 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 the following:

- 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 AIS defect condition. Following detection of 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 AIS defect condition.

7.5 Ethernet remote defect indication (ETH-RDI)

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.

Specific configuration information required by a MEP to support ETH-RDI function is the following:

- 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 MEPs with the RDI field cleared.

7.6 Ethernet locked signal (ETH-LCK)

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 the following:

- 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.

Specific configuration information required by a MEP to support ETH-LCK reception is the following:

- 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 MEP, when administratively locked, transmits LCK frames in a direction opposite to its peer MEP(s). 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)

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 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, in the immediate client (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.

Specific configuration information required by a MEP to support ETH-Test is the following:

- 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 transmitting 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 at the immediate client MEG level in the same direction where TST frames are transmitted.

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. Further, when the receiving MEP is configured for an out-of-service test, it also generates LCK frames at the client MEG level in the direction in which the TST frames are received.

7.8 Ethernet automatic protection switching (ETH-APS)

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].

7.9 Ethernet maintenance communication channel (ETH-MCC)

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 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 the following:

- 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 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 remote 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)

ETH-EXP is used for the experimental OAM functionality which can be used within an administrative domain on a temporary basis. Interoperability of the experimental OAM functionality is not expected across different administrative domains.

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)

ETH-VSP is used for vendor-specific OAM functionality which may be used by a vendor across its equipment. Interoperability of vendor-specific OAM functionality is not expected across different vendors' equipment.

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 9.20, can be used for vendor-specific OAM. Details of vendor-specific OAM mechanisms are outside the scope of this Recommendation.

8 OAM functions for performance monitoring

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

This Recommendation covers the following performance parameters which are based on [MEF 10].

- **Frame loss ratio**

Frame loss ratio is defined as a ratio, expressed as a percentage, of the number of service frames not delivered divided by the total number of service frames during time interval T, where the number of service frames not delivered is the difference between the number of service frames arriving at the ingress ETH flow point and the number of service frames delivered at the egress ETH flow point in a point-to-point ETH connection.

- **Frame delay**

Frame delay can be specified as round-trip delay for a frame, where frame delay is defined as the time elapsed since the start of transmission of the first bit of the frame by a source node until the reception of the last bit of the loopbacked frame by the same source node, when the loopback is performed at the frame's destination node.

- **Frame delay variation**

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

Performance parameters are applicable to service frames. Service frames are those frames that conform to an agreed-upon level of bandwidth profile conformance. Service frames are admitted at the ingress ETH flow point of a point-to-point ETH connection and should be delivered to the egress ETH flow point. Specification of bandwidth profile conformance is outside the scope of this Recommendation.

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

- **Throughput**

Throughput is defined as the maximum rate at which no frame is dropped. This is typically measured under test conditions.

NOTE – The definition of availability is outside the scope of this Recommendation. However, the mechanisms defined in this Recommendation can contribute to availability-related measurements.

8.1 Frame loss measurement (ETH-LM)

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:

- **TxFCl**: Counter for in-profile data frames transmitted towards the peer MEP.
- **RxFCl**: Counter for in-profile data frames received from the peer MEP.

TxFCl and RxFCl 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 – For single-ended ETH-LM, both counters do not need to count on-demand OAM frames for ETH-LB, ETH-LT, ETH-LM, ETH-DM and ETH-Test. Instead, proactive OAM frames for ETH-CC and ETH-APS need to be counted. For dual-ended ETH-LM, the counters do not need to count on-demand OAM frames for ETH-LB, ETH-LT, ETH-LM, ETH-DM, and ETH-Test and proactive OAM frames for ETH-CC. However, proactive OAM frames for ETH-APS need to be 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 transmitting and receiving MEPs. Further, 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 outside the scope of this Recommendation.

Specific configuration information required by a MEP to support ETH-LM is the following:

- MEG level – MEG level at which the MEP exists.
- ETH-LM transmission period – 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 III.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.
- Single-ended ETH-LM.

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**: Value of the local counter TxFCI at the time of transmission of the CCM frame.
- **RxFcB**: Value of the local counter RxFCI at the time of reception of the last CCM frame from the peer MEP.
- **TxFcB**: 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 the 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 the 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, in which case the frame loss measurements are not carried out.

8.1.2 Single-ended ETH-LM

Single-ended ETH-LM is used for on-demand 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.

8.1.2.1 LMM transmission

For an on-demand 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 requesting 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)

ETH-DM can be used for on-demand 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 the 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 the following:

- MEG level – MEG level at which the MEP exists.
- Priority – Identifies the priority of the frames with ETH-DM information. This information is configurable per operation.
- Drop Eligibility – Frames with ETH-DM information are always marked as drop ineligible. This information is not necessarily configured.

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 RxTimef, the time at the reception of ETH-DM frame and calculate the one-way frame delay as:

$$\text{Frame delay} = \text{RxTimef} - \text{TxTimeStampf}$$

However, one-way frame delay measurement requires that the clocks at the transmitting 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 clock 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 a 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 a frame with ETH-DM request information) and TxTimeStampb (timestamp at the time of transmitting a frame with ETH-DM reply information).

ETH-DM can be performed in two ways:

- One-way ETH-DM
- Two-way ETH-DM

8.2.1 One-way ETH-DM

In this case, each MEP sends a frame with one-way ETH-DM information to its peer MEP in a point-to-point ME 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 one-way ETH-DM is 1DM, as described in clause 9.14. Frames which carry the 1DM PDU are called as 1DM frames.

8.2.1.1 1DM transmission

When configured for one-way delay measurements, a MEP periodically transmits 1DM frames with the TxTimeStampf value.

8.2.1.2 1DM Reception

When configured for one-way delay measurements, a MEP, upon receiving a 1DM frame, uses the following values to make one-way frame delay measurement. This value serves as input to the one-way frame delay variation measurement:

- 1DM frame's TxTimeStampf value.
- RxTimef, which is the time at reception of the 1DM frame.

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

8.2.2 Two-way 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.

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.

8.2.2.1 DMM transmission

When configured for two-way delay measurements, a MEP periodically transmits DMM frames with the TxTimeStampf value.

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 requesting MEP. A DMM 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 DMM frame. Every field in the DMM frame is copied to the DMR frame with the following exceptions:

- The source and destination MAC addresses are swapped.
- 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 remote 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.

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

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 is calculated to be:

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

8.3 Throughput measurement

[b-RFC 2544] specifies measuring the throughput by sending frames at 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.

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 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. To support the OAM functions specified in this Recommendation, version is always 0.
- **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

Table 9-1 – OpCode values

OpCode value	OAM PDU type	OpCode relevance for MEPs/MIPs
4	LTR	MEPs and MIPs
0, 6-31, 64-255	Reserved (Note 1)	
OpCodes specific to this Recommendation		
33	AIS	MEPs
35	LCK	MEPs
37	TST	MEPs
39	Linear APS	MEPs
40	Ring APS	MEPs
41	MCC	MEPs
43	LMM	MEPs
42	LMR	MEPs
45	IDM	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
32, 34, 36, 38, 44, 52-63	Reserved (Note 2)	
NOTE 1 – Reserved for definition by IEEE 802.1.		
NOTE 2 – Reserved for future standardization by ITU-T.		

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

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
33-63	Reserved (Note 2)
NOTE 1 – Reserved for definition by IEEE 802.1.	
NOTE 2 – 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.
- **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 the most significant bit of the 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.

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

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

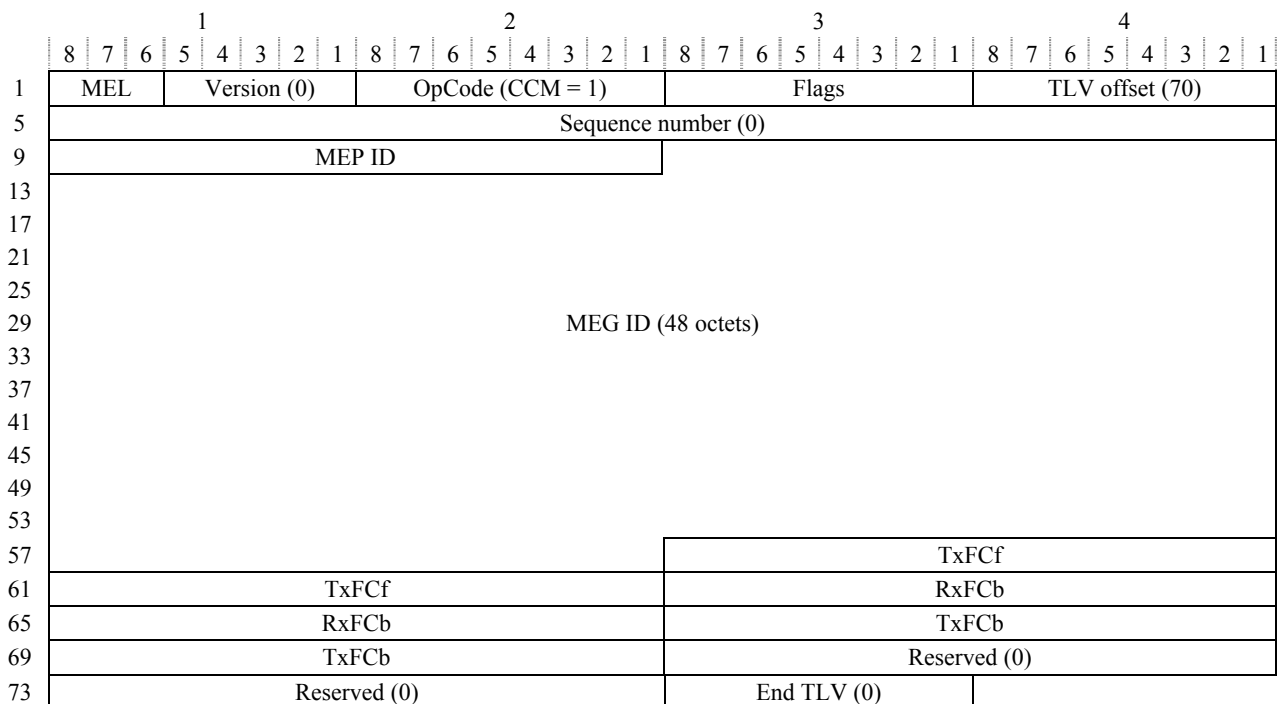


Figure 9.2-1 – CCM PDU format

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 always 0.
- **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 frames 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:



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

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

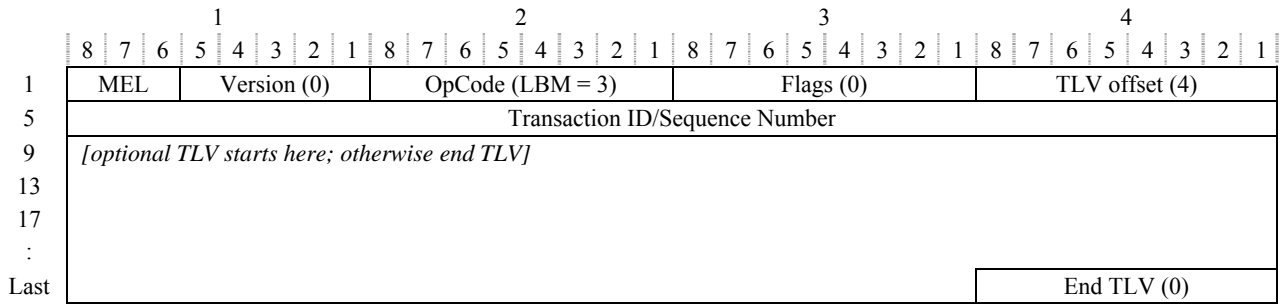


Figure 9.3-1 – LBM PDU format

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 always 0.
- **OpCode:** Value for this PDU type is LBM (3).
- **Flags:** Set to all-ZEROes.




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:** If present, a Data TLV or Test TLV as specified in Figure 9.3-3 or Figure 9.3-4, respectively.
- **End TLV:** All-ZEROes octet value.




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 bytes

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.

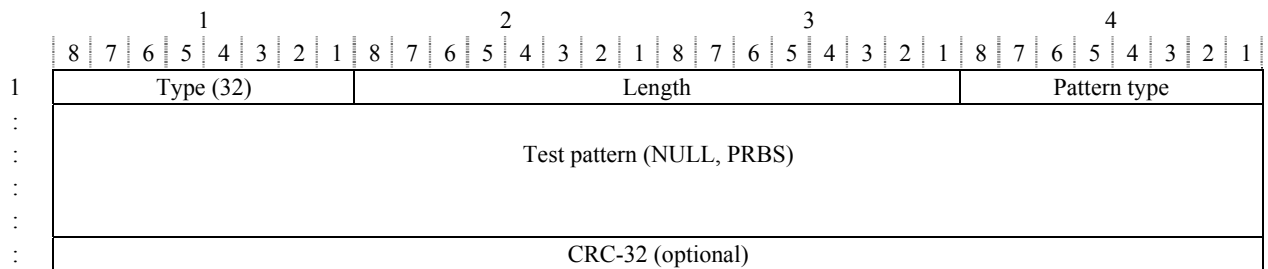


Figure 9.3-4 – Test TLV format

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 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 bytes 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 byte is used for pattern type, 1479 bytes 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.

9.4.2 LBR PDU format

LBR PDU format used by a MEP to transmit LBR information is shown in Figure 9.4-1.

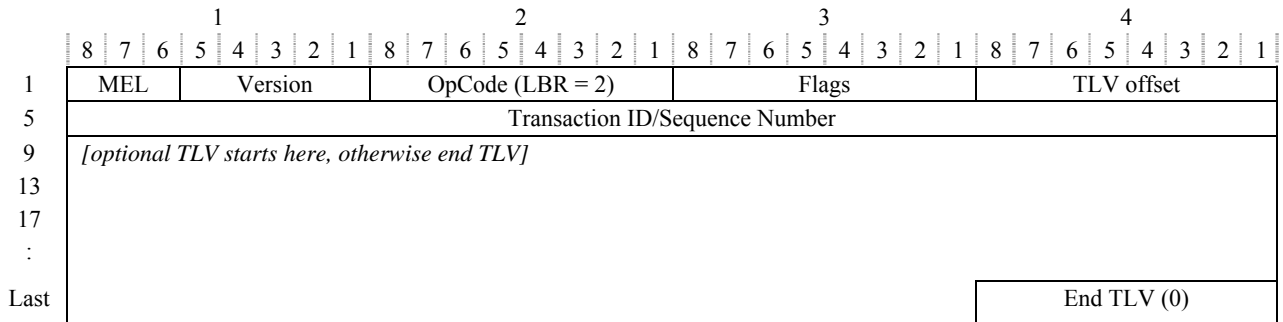


Figure 9.4-1 – LBR PDU format

The fields for the LBR PDU format are as follows:

- **MEG level:** A 3-bit field the value of which is copied from the received LBM PDU.
- **Version:** A 5-bit field the value of which is copied from the LBM PDU.
- **OpCode:** Value for this PDU type is LBR (2).
- **Flags:** A 1-octet field the value of which is copied from the LBM PDU.
- **TLV offset:** A 1-octet field the value of which is copied from the LBM PDU.
- **Transaction ID/Sequence Number:** A 4-octet field the value of which is copied from the LBM PDU.
- **Optional TLV:** If present in LBM PDU, copied from the LBM PDU.
- **End TLV:** A 1-octet field the value of which is copied from the LBM PDU.

9.5 LTM PDU

LTM is used to support ETH-LT request, as described in clause 7.3.

9.5.1 LTM information elements

The information elements carried in LTM include:

- **Transaction ID:** Transaction ID is a 4-octet field that contains the transaction number for the LTM. The receiver is expected to copy transaction number in the LTR PDU, as described in clause 9.6.
- **TTL:** TTL is a 1-octet field used to indicate whether a LTM should be terminated or not by the receiver. When a MIP receives LTM with TTL = 1, the LTM is not relayed. A network element receiving LTM decrements the received TTL value by one and copies it into the TTL field of LTR PDU, as described in clause 9.6, as well as into the LTM that it forwards towards the next hop.
- **TargetMAC:** TargetMAC is a 6-octet field used to carry MAC address of the targeted end-point. An intermediate MIP copies this field into the LTM that it forwards towards the next hop.
- **OriginMAC:** OriginMAC is a 6-octet field used to carry the MAC address of the originating MEP. An intermediate MIP copies this field into the LTM that it forwards towards the next hop.

9.5.2 LTM PDU format

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

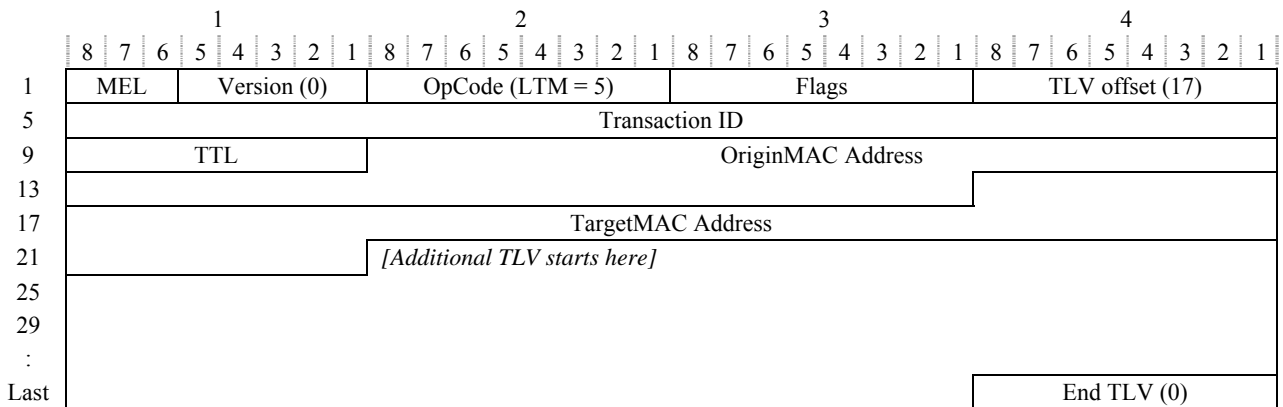


Figure 9.5-1 – LTM PDU format

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 always 0.
- **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

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

- **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 MEP initiating LTM frame or ETH-LT responder relaying modified LTM frame. Octets 4 and 5 are ZEROs 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.

9.6 LTR PDU

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

9.6.1 LTR information elements

The information elements carried in LTR include:

- **Transaction ID:** Transaction ID is a 4-octet field that is copied from the Transaction ID field in LTM.
- **TTL:** TTL is a 1-octet field that contains the TTL field value decremented by 1 from the LTM for which LTR is being sent.

9.6.2 LTR PDU format

LTR PDU format used by a MEP or MIP to transmit LTR information is shown in Figure 9.6-1.

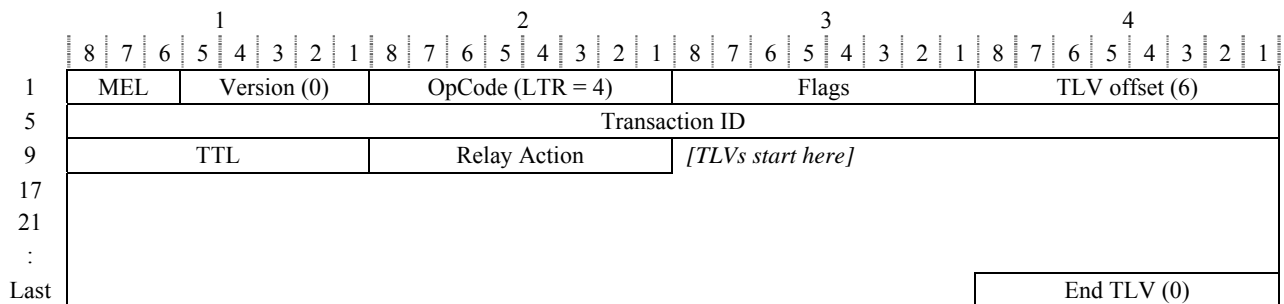


Figure 9.6-1 – LTR PDU format

The fields of the LTR PDU format are as follows:

- **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 always 0.
- **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 which is reserved for use by IEEE 802.1.
- **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.
- **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 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 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 ZEROs while remaining six octets 14-19 contain a 48-bit IEEE MAC address unique to the 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.




Figure 9.6-5 – Reply egress TLV format

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

- **Type:** Identifies TLV type; value for this TLV type is egress reply (6).
- **Length:** Identifies size, in octets, of the value field. This is set to 7.
- **Egress action:** A 1-octet field which is reserved for definition by IEEE 802.1.
- **Egress MAC address:** A 6-octet field which is reserved for definition by IEEE 802.1.

9.7 AIS PDU

AIS PDU is used to support ETH-AIS function, as described in clause 7.4.

9.7.1 AIS information elements

Information element carried in AIS 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 AIS transmission periodicity. AIS period values are specified in Table 9-4.

9.7.2 AIS PDU format

AIS PDU format used by a MEP to transmit AIS information is shown in Figure 9.7-1.




Figure 9.7-1 – AIS PDU format

The fields of the AIS 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 always 0.
- **OpCode:** Value for this PDU type is AIS (33).
- **Flags:** One information element in the Flags Field for AIS PDU: Period as follows:




Figure 9.7-2 – Flags format in AIS PDU

- **Period:** Bits 3 to 1 indicate transmission period with the encoding in Table 9-4.

Table 9-4 – AIS/LCK period values

Flags[3:1]	Period value	Comments
000-011	Invalid value	Invalid value for AIS/LCK PDUs
100	1 s	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 PDU

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

9.8.1 LCK information elements

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

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 always 0.
- **OpCode:** Value for this PDU type is LCK (35).
- **Flags:** One information element in the Flags field for 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 unidirectional ETH-Test function, as described in clause 7.7.

9.9.1 TST information elements

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

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

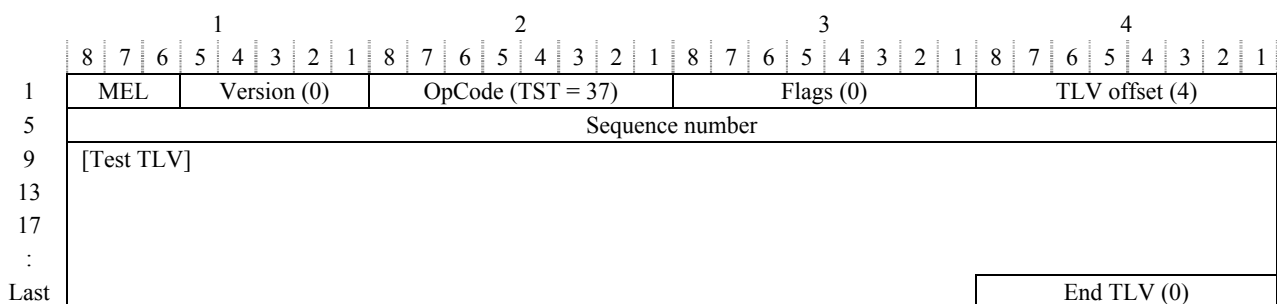


Figure 9.9-1 – TST PDU format

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 always 0.
- **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.

9.10.2 APS PDU format

APS PDU format used by a MEP to transmit APS information is shown in Figure 9.10-1.




Figure 9.10-1 – APS PDU format

The fields of the APS PDU format are as follows:

- **MEG level:** Refer to clause 9.1.
- **Version:** Refer to clause 9.1, value is always 0.
- **OpCode:** Value for this PDU type is (39) for linear APS and (40) for ring APS.
- **Flags:** Set to all-ZEROes.




Figure 9.10-2 – Flags format in APS PDU

- **TLV offset:** 1-octet field. Its specific value for APS is outside the scope of this Recommendation.
- **APS Data:** Format and length of this field are outside the scope of this Recommendation.
- **End TLV:** All-ZEROes octet value.

9.11 MCC PDU

MCC PDU is used to support ETH-MCC, as described in clause 7.9.

9.11.1 MCC information elements

Information elements carried in MCC include:

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

9.11.2 MCC PDU format

MCC PDU format used by a MEP to transmit MCC information is shown in Figure 9.11-1.

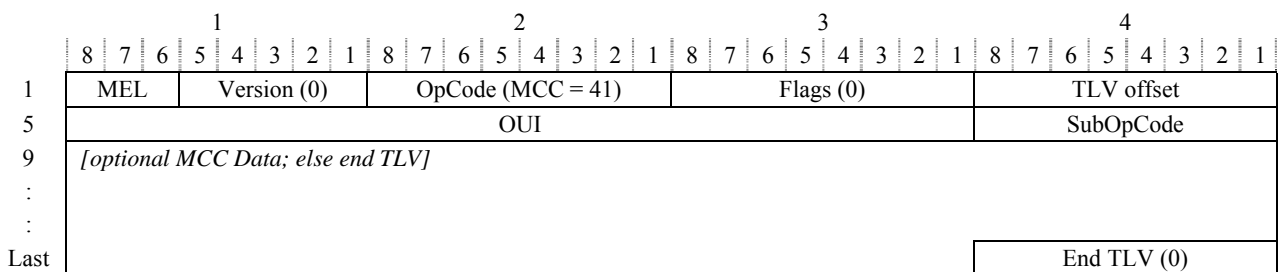


Figure 9.11-1 – MCC PDU format

The fields of the MCC PDU format are as follows:

- **MEG level:** Refer to clause 9.1.
- **Version:** Refer to clause 9.1, value is always 0.
- **OpCode:** Value for this PDU type is MCC (41).
- **Flags:** Set to all-ZEROes.




Figure 9.11-2 – Flags format in MCC PDU

- **TLV offset:** 1-byte field. Its specific value for MCC is outside the scope of this Recommendation.
- **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.
- **MCC data:** Format and length 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 single-ended ETH-LM request, as described in clause 8.1.2.

9.12.1 LMM information elements

Information element carried in LMM is:

- **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

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




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 is always 0.
- **OpCode:** Value for this PDU type is LMM (43).
- **Flags:** Set to all-ZEROes.




Figure 9.12-2 – Flags format in LMM PDU

- **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 single-ended 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

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




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

1DM PDU is used to support one-way ETH-DM, as described in clause 8.2.1.

9.14.1 1DM information element

Information element carried in 1DM is:

- **TxTimeStampf:** TxTimeStampf is an 8-octet field that contains the timestamp of 1DM transmission. The format of TxTimeStampf is equal to the TimeRepresentation format in [IEC 61588].

9.14.2 1DM PDU format

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

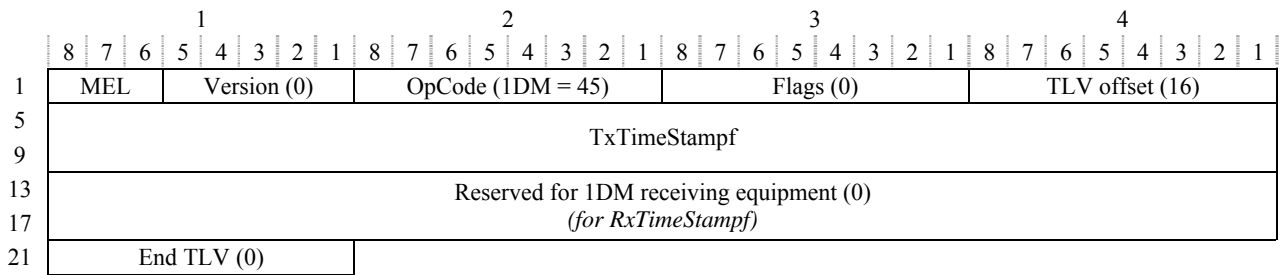


Figure 9.14-1 – 1DM PDU format

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

- **MEG level:** Refer to clause 9.1.
- **Version:** Refer to clause 9.1, value is always 0.
- **OpCode:** Value for this PDU type is 1DM (45).
- **Flags:** Set to all-ZEROes.




Figure 9.14-2 – Flags format in 1DM PDU

- **TLV offset:** Set to 16.
- **TxTimeStampf:** An 8-octet transmit timestamp field as described in clause 9.14.1.
- **Reserved:** Reserved fields are set to all-ZEROes.
- **End TLV:** All-ZEROes octet value.

9.15 DMM PDU

DMM is used to support two-way ETH-DM request, as described in clause 8.2.2.

9.15.1 DMM information elements

Information element carried in DMM is:

- **TxTimeStampf:** TxTimeStampf is an 8-octet field that contains the timestamp of DMM transmission. The format of TxTimeStampf is equal to the TimeRepresentation format in [IEC 61588].

9.15.2 DMM PDU format

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

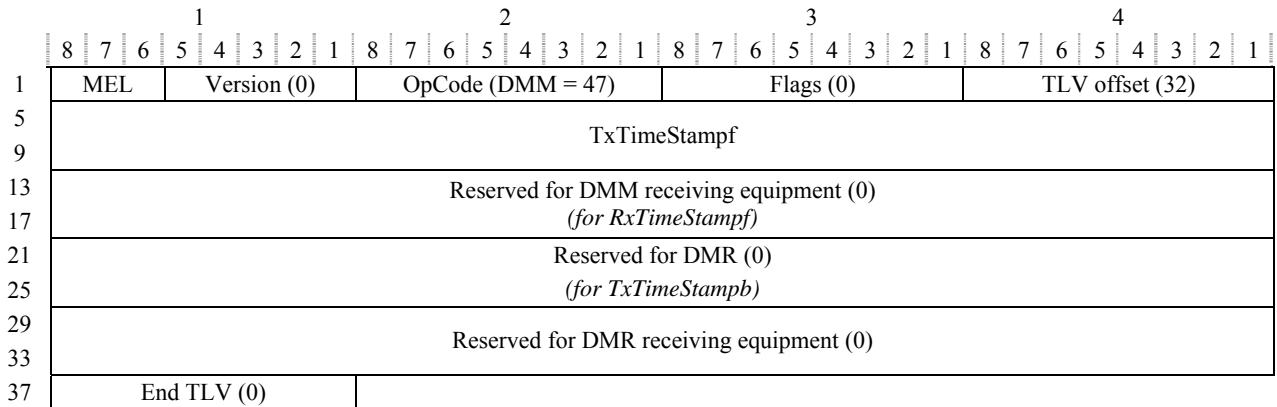


Figure 9.15-1 – DMM PDU format

The fields of the DMM PDU format are as follows:

- **MEG level:** Refer to clause 9.1.
- **Version:** Refer to clause 9.1, value is always 0.
- **OpCode:** Value for this PDU type is DMM (47).
- **Flags:** Set to all-ZEROes.




Figure 9.15-2 – Flags format in DMM PDU

- **TLV offset:** Set to 32.
- **TxTimeStampf:** An 8-octet transmit timestamp field as described in clause 9.15.1.
- **Reserved:** Reserved fields are set to all-ZEROes.
- **End TLV:** An all-ZEROes octet value.

9.16 DMR PDU

DMR is used to support two-way ETH-DM reply, as described in clause 8.2.2.

9.16.1 DMR information elements

Information elements carried in DMR are:

- **TxTimeStampf:** TxTimeStampf is an 8-octet field that contains the copy of TxTimeStampf field in received DMM.
- **RxTimeStampf:** RxTimeStampf is an optional 8-octet field that contains the timestamp of DMM reception. The format of RxTimeStampf is equal to the TimeRepresentation format in [IEC 61588]. When not used, a value of all 0 is used.
- **TxTimeStampb:** TxTimeStampb is an optional 8-octet field that contains the timestamp of DMR transmission. The format of TxTimeStampb is equal to the TimeRepresentation format in [IEC 61588]. When not used, a value of all 0 is used.

9.16.2 DMR PDU format

DMR PDU format used by a MEP to transmit DMR information is shown in Figure 9.16-1.

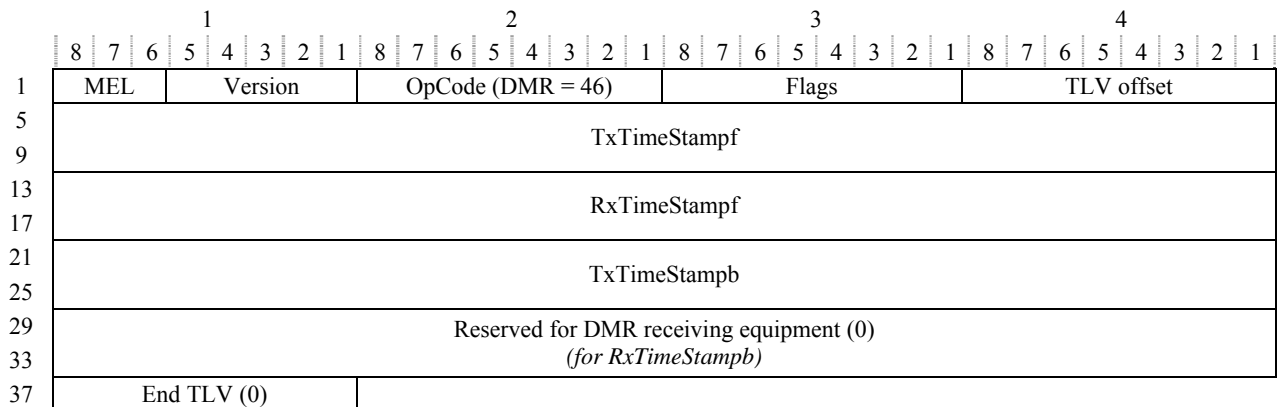


Figure 9.16-1 – DMR PDU format

The fields of the DMR PDU format are as follows:

- **MEG level:** A 3-bit field the value of which is copied from the last received DMM PDU.
- **Version:** A 5-bit field the value of which is copied from the last received DMM PDU.
- **OpCode:** Value for this PDU type is DMR (46).
- **Flags:** A 1-octet field the value of which is copied from the last received DMM PDU.
- **TLV offset:** A 1-octet field the value of which is copied from the last received DMM PDU.
- **TxTimeStampf:** An 8-octet field the value of which is copied from last received DMM PDU.
- **RxTimeStampf:** An 8-octet transmit timestamp field as described in clause 9.16.1.
- **TxTimeStampt:** An 8-octet transmit timestamp field as described in clause 9.16.1.
- **Reserved:** Reserved fields are set to all ZEROes.
- **End TLV:** A 1-octet field the value of which is copied from the DMM PDU.

9.17 EXM PDU

EXM is used as experimental OAM request PDU.

9.17.1 EXM PDU information elements

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

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

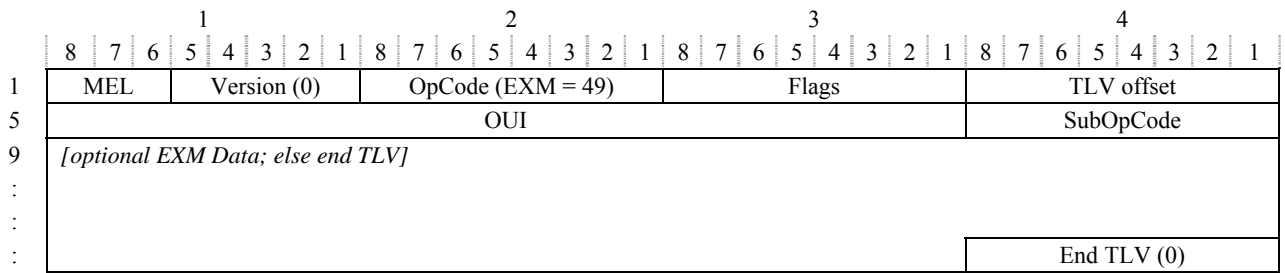


Figure 9.17-1 – EXM PDU format

The fields of the EXM PDU format are as follows:

- **MEG level:** Refer to clause 9.1.
- **Version:** Refer to clause 9.1, value is always 0.
- **OpCode:** Value for this PDU type is EXM (49).
- **Flags:** Outside the scope of this Recommendation.
- **TLV offset:** 1-byte field. Its specific value for EXM is outside the scope of this Recommendation, but 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 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 may carry one or more TLVs. EXR Data is outside the scope of this Recommendation.

9.18.2 EXR PDU format

EXR PDU format used by a MEP to transmit EXR information is shown in Figure 9.18-1.

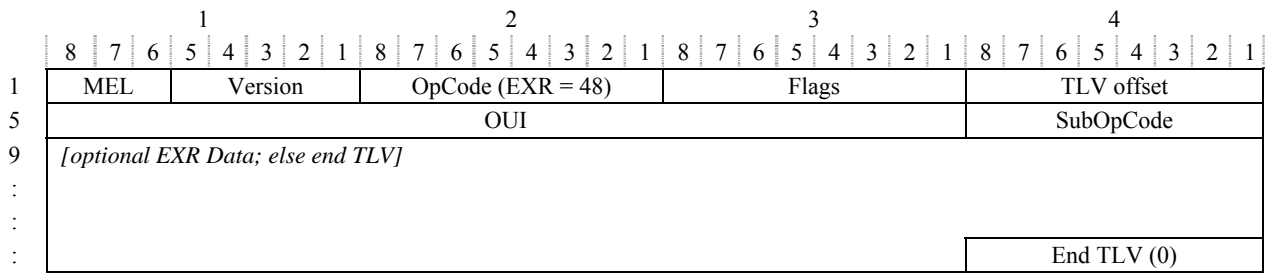


Figure 9.18-1 – EXR PDU format

The fields of the EXR PDU format are as follows:

- **MEG 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-byte field. Its specific value for EXR is outside the scope of this Recommendation, but 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 vendor-specific OAM request PDU.

9.19.1 VSM PDU information elements

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 may carry one or more TLVs. VSM Data is outside the scope of this Recommendation.

9.19.2 VSM PDU information elements

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

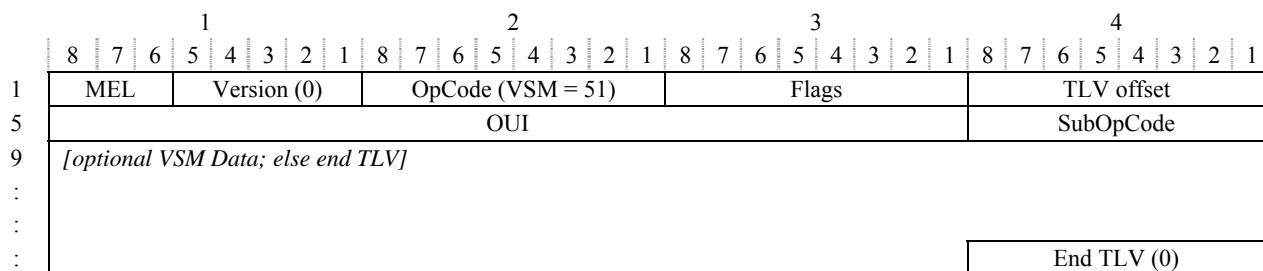


Figure 9.19-1 – VSM PDU format

The fields of the VSM PDU format are as follows:

- **MEG level:** Refer to clause 9.1.
- **Version:** Refer to clause 9.1, value is always 0.
- **OpCode:** Value for this PDU type is VSM (51).
- **Flags:** Outside the scope of this Recommendation.
- **TLV offset:** 1-byte field. Its specific value for VSM is outside the scope of this Recommendation, but 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 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 may carry one or more TLVs. VSR Data is outside the scope of this Recommendation.

9.20.2 VSR PDU format

VSR PDU format used by a MEP to transmit VSR information is shown in Figure 9.20-1.




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-byte field. Its specific value for VSR is outside the scope of this Recommendation, but 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.

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 DAs.

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.

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	Multicast class 1 DA or Unicast DA
Ring APS	Multicast class 1 DA or Unicast DA
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

10.1 Multicast destination addresses

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

- Multicast DA class 1: OAM frames that are addressed to all MEPs in a MEG (e.g., CCM, multicast LBM, AIS, etc.).
- Multicast DA class 2: OAM frames that are addressed to all MIPs and MEPs associated with a MEG (e.g., LTM).

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 MEG level with x being a value in the range 0-7 and y represents MEG level with y being a value in the range 8-F.

10.2 CCM

CCM frames can be generated with a specific Multicast class 1 DA or Unicast DA.

When a Multicast DA is used, CCM frames allow discovery of MAC addresses associated with MEPs. The use of a Multicast DA also allows for 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 service instances are distinguished using the Unicast DAs, CCM frames can be provisioned to use Unicast DAs.

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 frame is generated with a Multicast class 2 DA.

A Multicast DA is used instead of 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 frame can be generated with a Multicast class 1 DA, especially in a multipoint MEG.

Unicast DA may be used in provisioned environments for point-to-point connections. However, it requires that the Unicast DA of the downstream MEP must have been configured on the MEP transmitting AIS.

10.8 LCK

LCK frame can be generated with a Multicast class 1 DA, especially in a multipoint MEG.

Unicast DA may be used in provisioned environments for point-to-point connections. However, it requires that the Unicast DA of the downstream MEP must have been configured on the MEP transmitting LCK.

10.9 TST

TST frames are generated with Unicast DAs. TST frames may be generated with Multicast class 1 DA if multipoint diagnostics are desired.

10.10 APS

APS frames can be generated with a specific Multicast class 1 DA or Unicast DA.

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 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 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 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.

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.

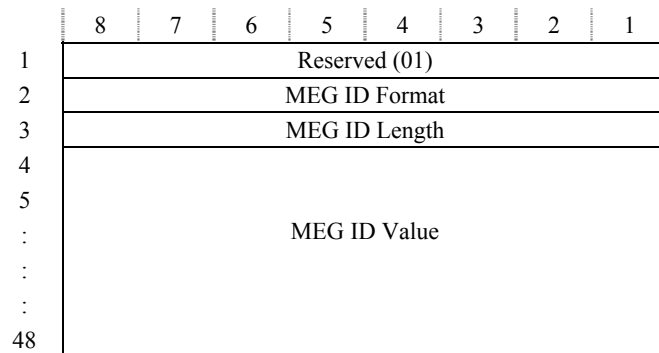


Figure A.1 – Generic MEG ID format

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.

Table A.1 – MEG ID Format type

MEG ID Format type value	TLV name
00-31, 64-255	Reserved (Note 1)
Types specific to this Recommendation	
32	ICC-based format
33-63	Reserved (Note 2)
NOTE 1 – Reserved for definition by IEEE 802.1.	
NOTE 2 – Reserved for future standardization by ITU-T.	

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-T Telecommunication Standardization Bureau (TSB) as per [ITU-T M.1400].

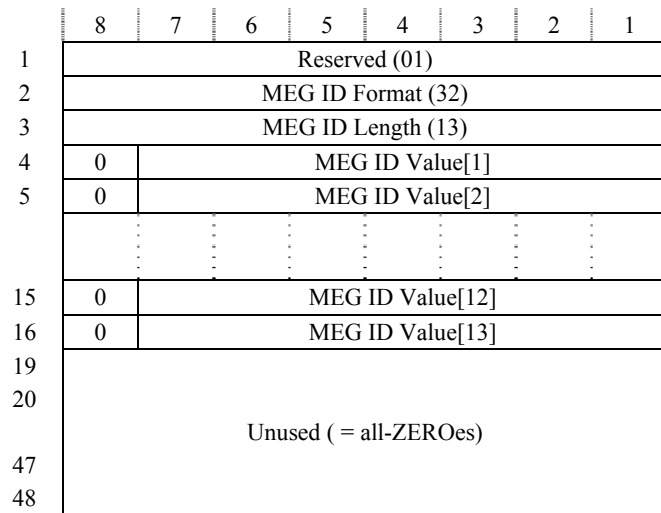


Figure A.2 – ICC-based MEG ID Format

The MEG ID Value identified in Figure A.2 consists of 13 characters coded according to [ITU-T T.50].

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 left-justified characters, alphabetic, or leading alphabetic with trailing numeric. 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.

Appendix I

Defect conditions

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

I.1 Loss of continuity (LOC) condition

A MEP detects LOC with a peer MEP when it stops receiving CCM frames from that peer ME. Such a defect condition can be caused by hard failures (e.g., link failure, device failure, etc.) or soft failures (e.g., memory corruption, mis-configurations, etc.).

Table I.1-1 – LOC entry/exit criteria

LOC(i)	
Entry criteria	A MEP receives no CCM frames from a peer MEP (MEP ID = i) during an interval equal to 3.5 times the CCM transmission period configured at the MEP.
Exit criteria	During an interval equal to 3.5 times the CCM transmission period configured at the MEP, the MEP receives n CCM frames from that peer MEP (MEP ID = i), where $3 \leq n$.

I.2 Mismatch condition

A MEP detects Mismatch when it receives a CCM frame with a correct MEG level (i.e., MEG level same as MEP's own MEG level) but incorrect MEG ID (indicating that frames from a different service instance are merged with the service instance represented by the MEP's own MEG ID). Such a defect condition is most likely caused by mis-configuration, but could also be caused by a hardware/software failure in the network.

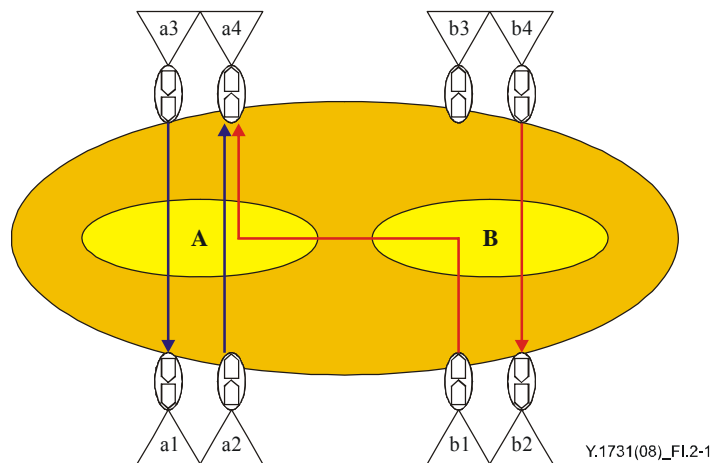


Figure I.2-1 – Mismatch condition

Table I.2-1 – Mismatch entry/exit criteria

Mismatch	
Entry criteria	A MEP receives a CCM frame with correct MEG level but incorrect MEG ID.
Exit criteria	During an interval equal to 3.5 times the CCM transmission period configured at the MEP, the MEP does not receive CCM frames with incorrect MEG ID.

When Mismatch detection is accompanied with LOC detection at a MEP, it points to the mismatch condition where a valid MEP may get swapped with an invalid MEP (belonging to a different MEG) in terms of connectivity in the network.

For limitations of Mismatch detection, see Appendix V.

I.3 Unexpected MEP (UnexpectedMEP) condition

A MEP detects UnexpectedMEP when it receives a CCM frame with a correct MEG level (i.e., MEG level equal to the MEP's own MEG level), a correct MEG ID but an Unexpected MEP ID which includes the MEP's own MEP ID. Determination of Unexpected MEP ID is possible when the MEP maintains a list of its peer MEP IDs. A list of peer MEP IDs must be configured on each MEP during provisioning. This defect condition is most likely caused by mis-configuration.

Table I.3-1 – UnexpectedMEP entry/exit criteria

UnexpectedMEP	
Entry criteria	A MEP receives a CCM frame with correct MEG level, correct MEG ID but with Unexpected MEP ID.
Exit criteria	During an interval equal to 3.5 times the CCM transmission period configured at the MEP, the MEP does not receive CCM frames with an Unexpected MEP ID.

I.4 Unexpected MEG level (UnexpectedMEGLevel) condition

A MEP detects UnexpectedMEGLevel when it receives a CCM frame with incorrect MEG level. Incorrect MEG level is less than the MEP's own MEG level. Such a defect condition is most likely caused by mis-configurations, e.g., improper MEG level configuration, missing MEPs, etc.

Table I.4-1 – UnexpectedMEGLevel entry/exit criteria

UnexpectedMEGLevel	
Entry criteria	A MEP receives a CCM frame with incorrect MEG level.
Exit criteria	During an interval equal to 3.5 times the CCM transmission period configured at the MEP, the MEP does not receive CCM frames with incorrect MEG level.

I.5 Unexpected period (UnexpectedPeriod) condition

A MEP detects UnexpectedPeriod when it receives a CCM frame with a correct MEG level (i.e., MEG level equal to the MEP's own MEG level), a correct MEG ID, a correct MEP ID, but with period field value different than the MEP's own CCM transmission period. Such a defect is most likely caused by mis-configuration.

Table I.5-1 – UnexpectedPeriod entry/exit criteria

UnexpectedPeriod	
Entry criteria	A MEP receives a CCM frame with correct MEG level, correct MEG ID, correct MEP ID, but with a period field value different than its own CCM transmission period.
Exit criteria	During an interval equal to 3.5 times CCM transmission period configured at the MEP, the MEP does not receive CCM frames with incorrect period field value.

I.6 Signal fail (SignalFail) condition

A MEP declares signal fail condition upon detection of defect conditions including loss of continuity, Mismatch, Unexpected MEP, Unexpected MEG level, etc.

Signal fail condition may also be declared by the Server layer termination function to notify the Server/ETH adaptation function, e.g., the Server MEP, about a defect condition in the Server layer.

I.7 AIS condition

A MEP detects AIS when it receives an AIS frame. Such a defect is caused by detection of signal fail condition at a Server layer or reception of AIS at a server (sub-) layer MEP in the case that the MEP does not use ETH-CC function.

Table I.7-1 – AIS entry/exit criteria

AIS	
Entry criteria	A MEP receives an AIS frame.
Exit criteria	During an interval equal to 3.5 times the AIS transmission period indicated in the AIS frames received earlier, the MEP does not receive AIS frames or, when ETH-CC is used, upon clearing of LOC defect at MEP.

A MEP can decide whether it blocks data frames when it detects dAIS. The principle requirement that influences this decision is that data frames ought to be forwarded as much as possible, without the possibility of forwarding wrong data frames downstream. Table I.7-2 provides some examples. A detailed description of data blocking is not provided in this Recommendation.

Table I.7-2 – Examples of data blocking at AIS

AIS generation condition	Data frames blocking decision
LOC, UnexpectedPeriod detection	Do not block
AIS frame reception	Do not block
Mismatch, UnexpectedMEP detection	Block
UnexpectedMEGLevel detection	Block

I.8 RDI condition

A MEP detects RDI when it receives a CCM frame with the RDI field set.

Table I.8-1 – RDI entry/exit criteria

RDI	
Entry criteria	A MEP receives a CCM frame with the RDI field set.
Exit criteria	A MEP receives a CCM frame with the RDI field clear.

I.9 LCK condition

A MEP detects LCK when it receives a LCK frame. Such a defect is caused by an intentional administrative/diagnostic action at a server (sub-) layer MEP that results in disruption of the client data traffic.

Table I.9-1 – LCK entry/exit criteria

LCK	
Entry criteria	A MEP receives a LCK frame.
Exit criteria	During an interval equal to 3.5 times the LCK transmission period indicated in the LCK frames received earlier, the MEP does not receive LCK frames.

Appendix II

Ethernet network scenarios

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

II.1 Shared MEG levels example

Figure II.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.

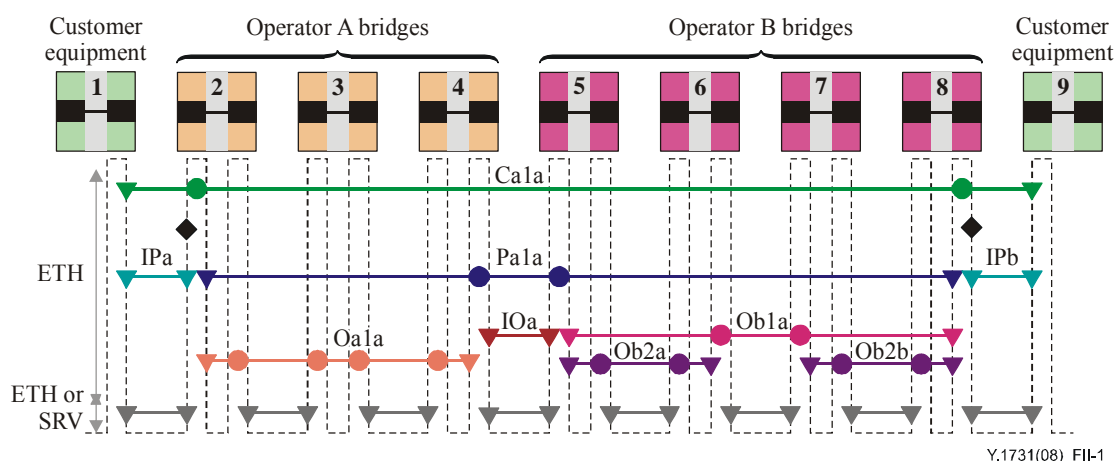


Figure II.1 – Example MEG level assignment for shared MEG levels

- UNI_C to UNI_C 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.
- UNI_N to UNI_N 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.
- UNI_C to UNI_N MEs (IPa and IPb) between the customer and provider 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.

II.2 Independent MEG levels example

Figure II.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 II.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 (represented in green) are completely independent of the two service provider VLANs (20 and 10) and corresponding service provider MEGs (represented in blue).
- 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 in Figure II.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 III

Frame loss measurement

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

III.1 Frame loss calculations

For the frame loss calculation, the four cases below should be taken into account.

- a) No wrapping around for both transmit and receive counters.
- b) Only transmit counter wraps around.
- c) Only receive counter wraps around.
- d) Both transmit and receive counters wrap around.

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

- a) No wrapping around for both transmit and receive counters (see Figure III.1):




Figure III.1 – No wrapping around

For this case, the frame loss can be calculated by the simple formula.

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

- b) Only transmit counter wraps around (see Figure III.2):

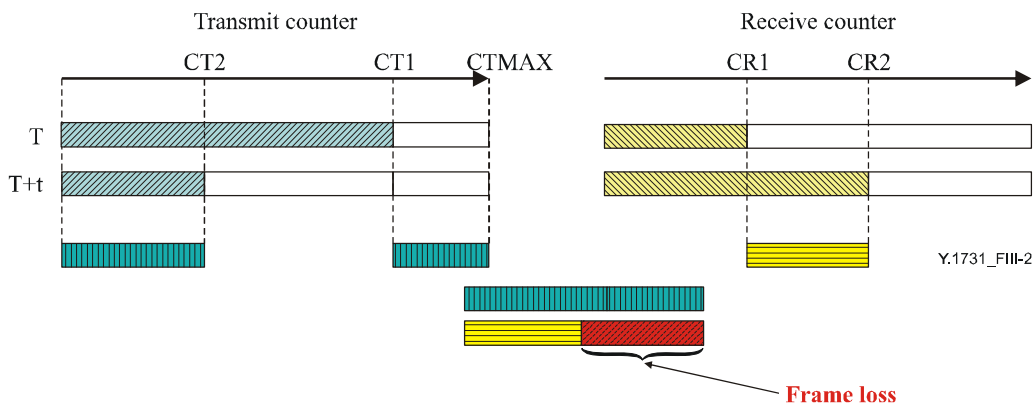


Figure III.2 – Transmit counter wraps around

In this case, frame loss can be calculated by the following equation as described in scenario a):

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

c) Only receive counter wraps around (see Figure III.3):




Figure III.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 (see Figure III.4):

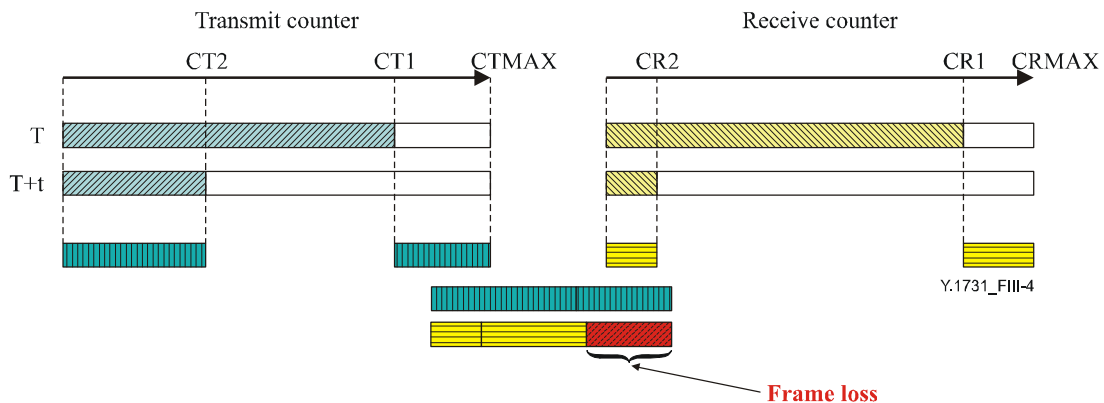


Figure III.4 – Both counters wrap around

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

III.1.1 Simplified calculation for frame loss

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

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

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

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

- a) **Frame loss** = (CT2 – CT1) – (CR2 – CR1)
- b) **Frame loss** = (CT2 – CT1) – (CR2 – CR1) + CTMAX+1
 = ((CT2 + (CTMAX+1)) – CT1) – (CR2 – CR1)
 = (CT2 – CT1) – (CR2 – CR1)
- c) **Frame loss** = (CT2 – CT1) – (CR2 – CR1) – (CRMAX+1)
 = (CT2 – CT1) – ((CR2 + CRMAX+1) – CR1)
 = (CT2 – CT1) – (CR2 – CR1)
- d) **Frame loss** = (CT2 – CT1) – (CR2 – CR1) + (CTMAX+1) – (CRMAX+1)
 = ((CT2 + (CTMAX+1)) – CT1) – ((CR2 + (CRMAX+1)) – CR1)
 = (CT2 – CT1) – (CR2 – CR1)

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

III.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. Frames sizes considered are 64-octet (minimum Ethernet frame size) and 1522-octet (maximum Ethernet frame size).

Table III.1 – Frame counter wrapping period

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

Appendix IV

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, introduced in clause 5.3.1, is used.

Appendix V

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 5.7. This behaviour leads to a limitation in the Mismerge detection as shown in Figure V.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 UnexpectedMEGLevel as defined in clause I.4.

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.

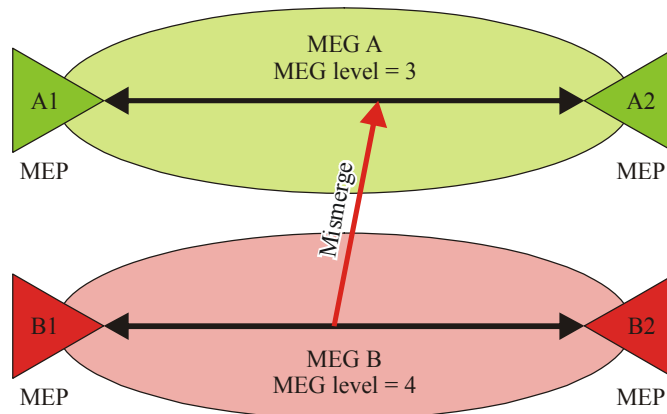
No defects detected by MEPs of MEG A as only MEG levels 3 and lower are considered.



Unexpected MEG level detected by MEPs of MEG B

a) Bidirectional mismerge

No defects detected by MEPs of MEG A as only MEG levels 3 and lower are considered.



No defects detected as no mismerge exists to MEG B

b) Unidirectional mismerge

Y.1731(08_FV-1

Figure V.1 – Mismerge detection limitation

Appendix VI

Terminology alignment with IEEE 802.1ag

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

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

Table VI.1 – Terminology mapping

Y.1731 term	802.1ag term	Comments
MEG	MA	
MEG ID	MAID (domain name + short MA name)	Unlike 802.1ag, the MEG ID does not imply a split between domain name and a short MEG name in this Recommendation.
MEG level	MA level	

Bibliography

- [b-RFC 2544] IETF RFC 2544 (1999), *Benchmarking Methodology for Network Interconnect Devices*.
<<http://www.ietf.org/rfc/rfc2544.txt>>

SERIES OF ITU-T RECOMMENDATIONS

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