INTERNATIONAL TELECOMMUNICATION UNION

TELECOMMUNICATION STANDARDIZATION SECTOR

STUDY PERIOD 2005-2008

STUDY GROUP 13

TD 94 (GEN)

English only

Original: English

Question(s): 5/13 Geneva, 7-17 December 2004

TEMPORARY DOCUMENT

Source: Editor, Y.17ethoam

Title: Draft Recommendation Y.17ethoam - OAM Functions and Mechanisms for Ethernet based networks

This document contains an updated version of draft Recommendation Y.17ethoam "OAM Functions and Mechanisms for Ethernet based networks". This version was drafted during the plenary meeting of SG13 (Geneva, 07-17 December 2004).

Contact: Dinesh Mohan
Nortel Networks
Canada

Tel: +1 613 763 4794
Fax: +1 613 763 2697
Email: mohand@nortelnetworks.com

TABLE OF CONTENTS

OAM Functions and Mechanisms for Ethernet based networks

1 Scope

The scope of this Recommendation is to specify mechanisms required to operate and maintain the network and service aspects of ETH layer. This Recommendation 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 also applicable to environments where ETH layer is managed using network management systems and/or operational support systems.

The architectural basis for this Recommendation is the Ethernet specification G.8010 which also accounts for IEEE 802.1D, 802.1Q, 802.3 and developments of IEEE P802.1ad, P802.1ah provider bridged networks. Furthermore the Connectivity Fault Management currently being defined in IEEE P802.1ag task force is taken into account.

The details of the atomic functions are not within the scope of this Recommendation which are expected to be specified in G.8021. 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 also not within the scope of this Recommendation.

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

NOTE: The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

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

[2] ITU-T Recommendation I.610 (1999), *B-ISDN operation and maintenance principles and functions*.

[3] CCITT Recommendation M.20 (1992), *Maintenance philosophy for telecommunications networks*.

[4] ITU-T Recommendation G.805 (2000), *Generic functional architecture of transport networks*.

[5] ITU-T Recommendation G.8010 (2003), *Architecture of Ethernet Layer Networks*.

[6] ITU-T Recommendation G.8041 (2001) *Generic Framing Procedure (GFP)*.

[EDITOR'S NOTE-Dec2004] TO BE COMPLETED, e.g. G.806

3 Definitions

This Recommendation uses terms defined in ITU-T G.805:

3.1 connection point

3.2 link

3.3 link connection

3.4 network connection

3.5 network operator

3.6 service provider

3.7 termination connection point

3.8 trail

3.9 trail termination

This Recommendation uses terms defined in ITU-T G.809:

3.10 adaptation

3.11 adapted information

3.12 client/server relationship

3.13 connectionless trail

3.14 flow

3.15 flow domain

3.16 flow domain flow

3.17 flow point

3.18 flow point pool

3.19 flow point pool link

3.20 flow termination

3.21 flow termination sink

3.22 flow termination source

3.23 layer network

3.24 link flow

3.25 network

3.26 port

3.27 reference point

3.28 traffic unit

3.29 transport

3.30 transport entity

3.31 transport processing function

3.32 termination flow point

3.33 termination flow point pool

This Recommendation uses terms defined in ITU-T M.20:

3.34 link

3.35 trail

This Recommendation uses terms defined in ITU-T G.806:

3.36 defect

3.37 failure

This Recommendation uses terms defined in ITU-T G.8010:

3.38 ETH trail

3.39 ETH link

3.40 Point-to-point Ethernet connection

3.41 Multipoint Ethernet connectivity

3.42 Multipoint Ethernet connection

This Recommendation defines the following terms:

3.43 xxx

[EDITOR'S NOTE-Dec2004] TO BE COMPLETED

4 Abbreviations

This Recommendation uses the following abbreviations:

AP	Access Point
CE	Customer Edge
CP	Connection Point
DoS	Denial of Service
ETH	Ethernet MAC layer network
ETH-AIS	Ethernet Alarm Indication Signal
ETH-CC	Ethernet Continuity Check
ETH-LB	Ethernet Loopback
ETH-LT	Ethernet Link Trace
ETH-RDI	Ethernet Reverse Defect Indication
ETHS	ETH Segment
ETY	Ethernet PHY layer network
ETYn	Ethernet PHY layer network of order n
FD	Flow Domain
FDF	Flow Domain Flow
FDFr	Flow Domain Fragment
FP	Flow Point
FPP	Flow Point Pool
FT	Flow Termination
MAC	Media Access Control
ME	Maintenance Entity

MEG	ME Group
MEP	MEG End Point
MIP	MEG Intermediate Point
NMS	Network Management System
NNI	Network Node Interface
OAM	Operation, Administration and Maintenance
OTN	Optical Transport Network
PE	Provider Edge
PHY	Ethernet Physical Layer entity consisting of the PCS, the PMA, and, if present, the PMD sublayers
SLA	Service Level Agreement
TC	Traffic Conditioning
TCP	Traffic Conditioning Point
TFP	Termination Flow Point
TFPP	Termination Flow Point Pool
UNI	User Network Interface
UNI-C	Customer side of UNI
UNI-N	Network side of UNI
VID	VLAN Identifier
VLAN	Virtual LAN

[EDITOR'S NOTE-Dec2004] TO BE COMPLETED

5 OAM Conventions

The diagrammatic convention for connection-oriented and connectionless layer networks described in this Recommendation are that of Recommendation G.805, G.809, and G.8010.

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

5.1 Maintenance Entity (ME)

ME is an entity that requires management. MEs in Ethernet networks are defined in Figures 23 and 24 of G.8010 [4] and in section 9 of Y.1730 [1]. MEs can nest but not overlap. The mapping of the MEs as defined in both Recommendations are shown in Table 5-1.

Y.1730 ME	G.8010 ME
UNI-UNI (Customer)	UNI_C to UNI-C ME
UNI-UNI (provider)	UNI_N to UNI_N ME
Segment (PE-PE) intra-provider	Intra Domain ME
Segment (PE-PE) inter-provider (provider – provider)	Inter Domain ME
ETY Link OAM – UNI (customer – provider)	Access Link ME
ETY Link OAM – NNI (operator – operator)	Inter Domain ME

Table 5-1: MEs and OAM Flows

5.2 ME Group (MEG)

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

1. MEs in a MEG exist in the same administrative boundary; and
2. MEs in a MEG have the same ME Level (Section 5.6), and
3. MEs in a MEG belong to the same point-to-point ETH connection or multipoint ETH connection

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

[EDITOR'S NOTE] Reference G.809/G.805 for administrative domain.

[EDITOR'S NOTE] Bind MEG to a protection domain; contributions are invited.

5.3 MEG End Point (MEP)

MEG End Point (MEP) is a short name for an expanded ETH flow point that includes a compound ETH Segment flow termination function (ETHS), which marks the end point of an ETH ME, and a compound ETH Diagnostic flow termination function (ETHD). The ETHS is capable to initiate and terminate proactive OAM signals. MEP's ETHD is capable to initiate and react to diagnostic OAM signals. A MEP is represented by a triangle symbol as shown in Figure 5-1.

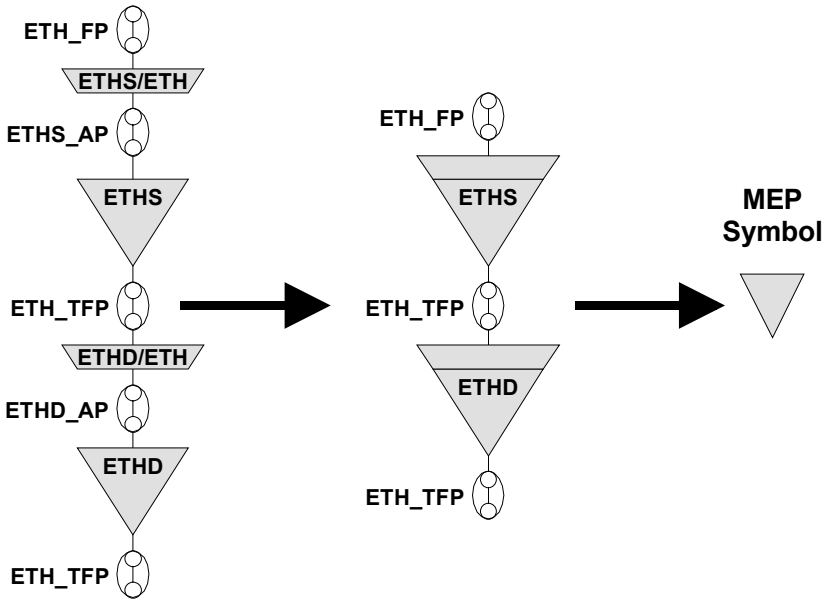

Figure 5-1: MEG End Point (MEP) symbol

5.4 MEG Intermediate Point (MIP)

MEG Intermediate Point (MIP) is a short name for an expanded ETH flow point that includes two compound ETH Diagnostic flow termination functions (ETHD). MIP's ETHD are capable to react to diagnostic OAM signals and do not initiate diagnostic OAM signals. MIP is represented by a circle symbol as shown in Figure 5-2.

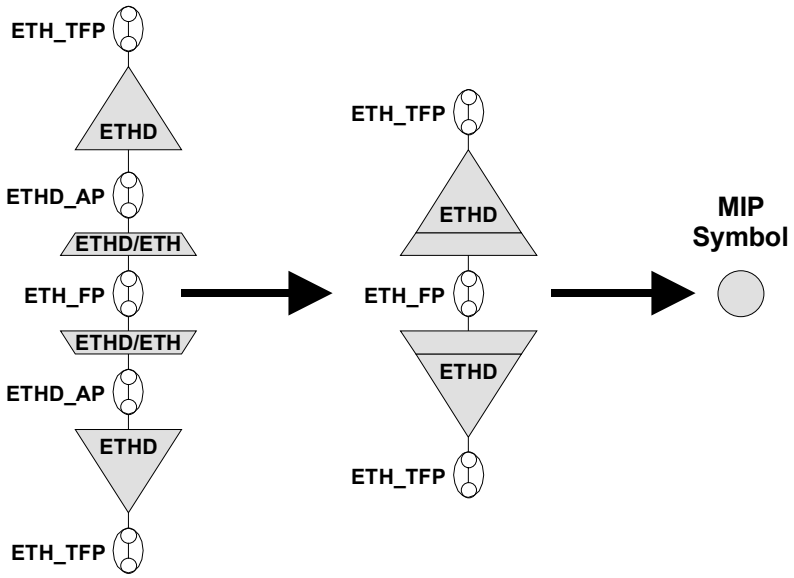

Figure 5- 2 – MEG Intermediate Point (MIP) symbol

5.5 Traffic Conditioning Point (TCP)

Traffic Conditioning Point (TCP) is a short name for an expanded ETH flow point that includes an ETH traffic conditioning function, as specified in Recommendation G.8010. A TCP is represented by a diamond symbol as shown in Figure 5-3.

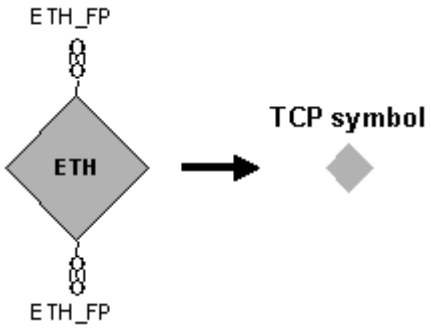

Figure 5-3 – Traffic Conditioning Point (TCP) symbol

5.6 ME Level

At any point in a network, ME Level is used to distinguish between OAM signals belonging to different nested MEs.

Eight ME Levels are available to accommodate different network deployment scenarios. The eight ME Levels are mutually agreed amongst customer, provider and operator entities involved in ETH connections. Default ME Levels assignment amongst customer, provider, and operator entities are defined in the following manner:

- Customers are assigned 3 ME Levels: 0, 1, and 2
- Providers are assigned 2 ME Levels: 3 and 4
- Operators are assigned 3 ME Levels: 5, 6, and 7

Note: Initial ME Level assignment agreement had been {0, 1, and 2} for Operators, {3, and 4} for Providers, and {5, 6, and 7} for Customers. However, the ME Level assignments have been reversed in consideration of aligning with OTN OAM Level assignments across Path, Segment 1-6, and Section.

[EDITOR'S NOTE-Dec2004] Communication with Q.14/15 may be needed for configurations and defaults for EMF – Ethernet Management Functions.

5.7 OAM Transparency

OAM Transparency refers to the ability to allow transparent carrying of OAM signals belonging to higher level MEs across other lower level MEs when these MEs are nested.

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

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

Customer can use any of the eight ME Levels, mentioned in Section 5.6, however, transparency of customer's OAM signals across provider and operator administrative domains will only be

guaranteed for mutually agreed ME Levels e.g. default ME Levels 0,1 and 2. Providers and operators should use only mutually agreed ME Levels, e.g. default ME Levels 3 and 4 for providers and 5,6, and 7 for operators.

OAM signals can be prevented from leaking by implementing an OAM filtering process in the MEP atomic functions. Refer to Appendix VI and VII for further details.

[EDITOR'S NOTE-Dec2004] Communication with Q.9/15 is needed for initiating details of the atomic functions and their processes.

6 OAM Relationships

6.1 MEs and VLANs Relationship

Figure 8 of G.8010 highlights the ETH Flow Domain Fragments (FDFr) which provide connectivity between the (termination) flow points in the fragment. IEEE 802.1Q implementation provides one means of realizing ETH FDFr, where VLAN ID(s) can be used to identify ETH FDFr(s).

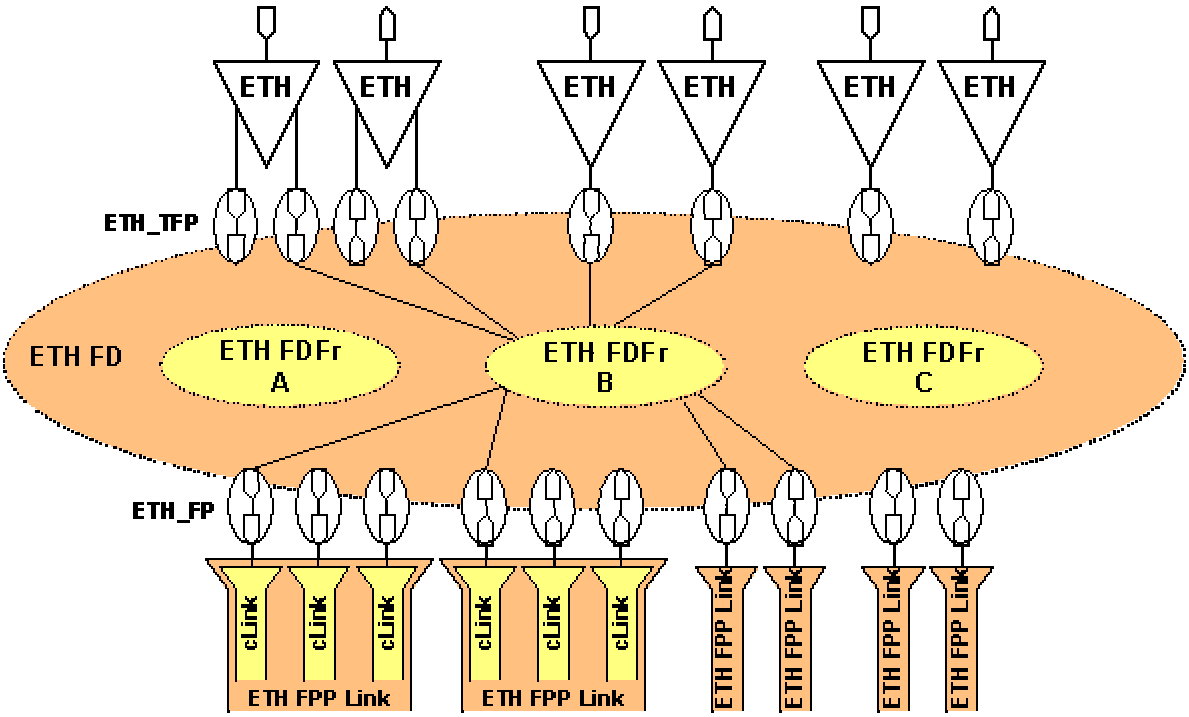

Figure 6-1 – G.8010/Y.1306 (Figure 8) – ETH Flow Domain Fragments

When the provider equipment consists of a dual-relay bridge, segregation of customer service flows, identified by Customer VLANs (C-VLAN), can be achieved by supporting each service instance with a separate Service VLAN (S-VLAN), which can be applied by the provider to customer service frames.

C-VLAN and S-VLAN identify different ETH FDFr(s) and belong to different VLAN spaces. C-VLANs and S-VLANs can therefore be used to segregate MEs belonging to customer and providers within respective Ethernet Flow Domains. The relationship between C-VLAN, S-VLAN and MEs is shown in Figure 6-2. Ethernet OAM flows belonging to C-VLAN and S-VLAN spaces are invisible each other.

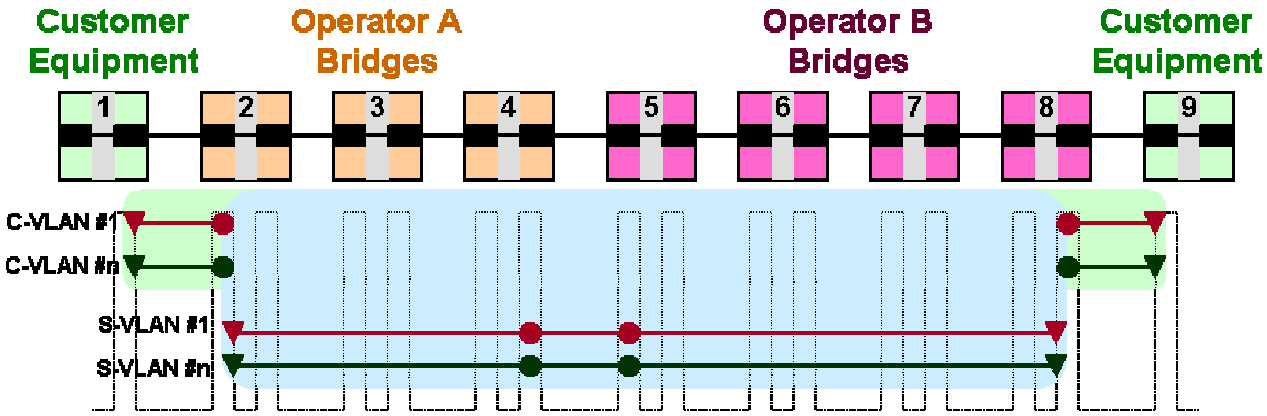

Figure 6-2 – Relationship between MEs and VLANs

As a result, the same OAM mechanisms can be applied independently to C-VLANs and S-VLANs.

Annex B illustrate the relationships of MEs associated with different dual-relay modelled Ethernet network scenarios which use C-VLANs and S-VLANs.

6.2 MEPs/MIPs and Port Status Relationship

As shown in Figure 6-3, a number of MEPs can be associated with a given device port. Left port of Operator A Bridge 2 is associated with 3 different MEPs numbered 1, 2, and 3. MEP 1 is facing "out of device" while MEPs 2 and 3 are facing "in the device". Each of these MEPs is associated with a unique ME Level. Only one MIP, i.e. MIP 4, is shown to be associated with this port.

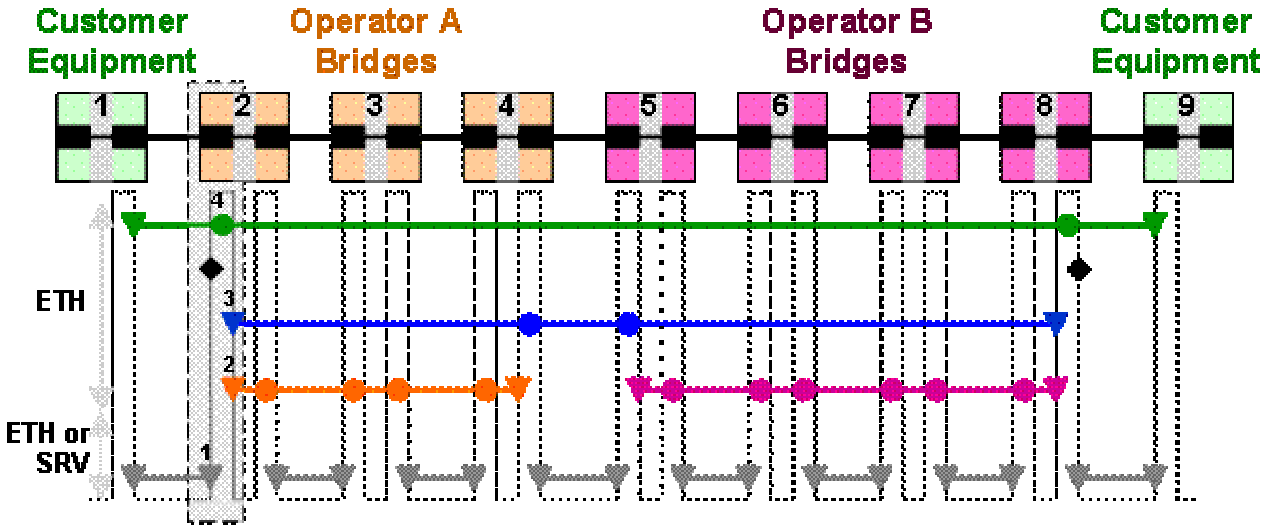

Figure 6-3 – Relationship between Port Status and Maintenance Points

Certain OAM signals can be generated and inserted in MEP's ETH Segment flow termination function i.e. ETHS_FT_So atomic functions. These OAM signals can be extracted and processed in the ETHS_FT_Sk atomic functions. Similarly MEP's ETH Diagnostic flow termination functions are capable to initiate and react to diagnostic OAM signals.

The device port can possibly have different states. Possible port states include:

- Operationally Up

- Operationally Down
- Operationally Blocked
- Administratively Locked
- Administratively Enabled
- Administratively Disabled
- Administratively in Test/Diagnostics

The device port state determines the OAM capabilities of the MEPs and MIPs associated with the port. For example, certain OAM signals may not be inserted or processed unless device port is in "administrative test/diagnostic" state. Similarly, certain OAM signals may not be inserted or processed when the device port is in "operationally blocked" state. In this port state, the MEPs facing "out of device" continue to function normally while MEPs facing "in the device" do not exist. Also, in this port state, the MIPs do not exist.

When a MEP is able to function normally, it is called to be an "active state" MEP. Generally, an "active state" MEP is associated with a port in both "operationally up" and "administratively enabled" states. When a MEP is associated with a port in "administrative test/diagnostic" state, it is called to be a "diagnostic state" MEP.

[EDITOR'S NOTE – Dec2004] Provide standard reference for Port Status e.g. ITU-T (e.g. X.731) and/or IEEE (e.g. IEEE 802.3) references and possible mappings between the ITU-T and IEEE port states. Contributions are invited.

6.3 MEs, MEPs, MIPs and TCPs Relationship

Annex B provides different network scenarios to show how MEs, MEPs and MIPs at different ME Levels can be deployed, and where TCPs are likely to be placed.

Note: Not all MEs and corresponding MEPs and MIPs may be used or allowed in the example network scenarios shown in Annex B. For example, providers may disallow their customers to create MIPs on provider devices.

6.4 MEs and ME Levels Relationship

The MEPs associated with an administrative domain operate at the assigned ME Level. Inter-domain MEPs, associated with MEs between two administrative domains, can operate at a ME Level agreeable between the two administrative domains, such that associated inter-domain OAM flows are prevented from leaking into either administrative domain. The default ME Level for inter-domain OAM flows is one just below the ME Level at which the administrative domain is providing transparency.

Table 6-1 highlights possible ME Level assignments for MEs within the context of Customer, Provider and Operator administrative domains, as mapped to Y.1730 and G.8010.

Y.1730 ME	G.8010 ME	ME Level
UNI-UNI (Customer)	UNI_C to UNI-C ME	0,1, or 2
UNI-UNI (provider)	UNI_N to UNI_N ME	3, or 4
Segment (PE-PE) intra-provider	Intra Domain ME	3, or 4
Segment (PE-PE) inter-provider (provider – provider)	Inter Domain ME	7 (default)

ETY Link OAM – UNI (customer – provider)	Access Link ME	7 (default)
ETY Link OAM – NNI (operator – operator)	Inter Domain ME	7 (default)

Table 6-1: Relationship between OAM flows and MEs

Figure 6-4 provides an example scenario with the default assignment of ME Levels.

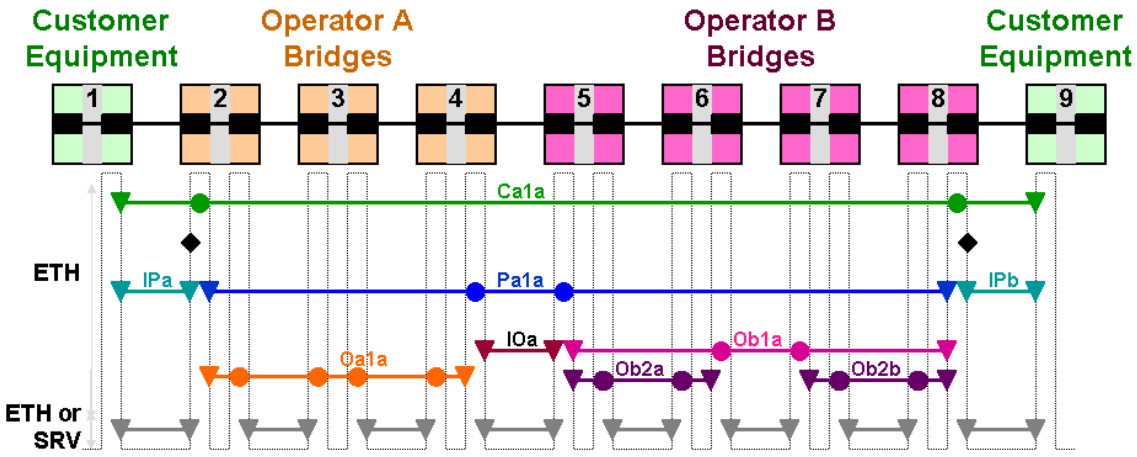

Figure 6-4: Example of default ME Level assignment

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

6.5 MEPs and MIPs Configurations

MEPs and MIPs can be configured along with their ME Levels either manually or automatically. Manual configurations may be performed either through manual local administration of each device or via Network Management Systems (NMS) as indicated in operational scenarios in Annex C. Automatic configurations are also possible via control plane mechanisms and data plane mechanisms using OAM signals. These mechanisms are FFS.

[EDITOR'S NOTE-Dec2004] Contributions are invited to provide solutions for automatic MEP and MIP configurations.

7 OAM Functions for Fault Management

[EDITOR'S NOTE-Dec2004] Use of term OAM signal vs. OAM flows and/or OAM frames needs to be clarified. Contributions invited.

7.1 Ethernet Continuity Check (ETH-CC)

Ethernet Continuity Check (ETH-CC) can be used to detect continuity failures across MEs between a given pair of flow termination functions. ETH-CC can also be used to discover the MAC addresses associated with MEPs. Continuity failures could result due to hard failures (e.g. link failure, device failure, etc.) or soft failures (e.g. software failure, memory corruption, mis-configurations, etc.). Although ETH-CC can be used to detect continuity failures across any given pair of flow termination functions, bound to a pair of flow points, it is particularly useful to detect continuity failures across a given pair of edge flow points.

As shown in Figure 7-1.1, to detect continuity failures with either a given set of flow point or all flow points meeting certain condition(s) within an administrative domain, ETH-CC signal is generated and inserted in the ETHS_FT_So atomic function associated with the sending MEP. The ETH-CC signal is extracted and processed in the ETHS_FT_Sk atomic function associated with receiving MEP. Generation and insertion of ETH-CC can be enabled or disabled in the ETHS_FT_So atomic functions. Processing of ETH-CC can be enabled or disabled in the ETHS_FT_Sk atomic functions.

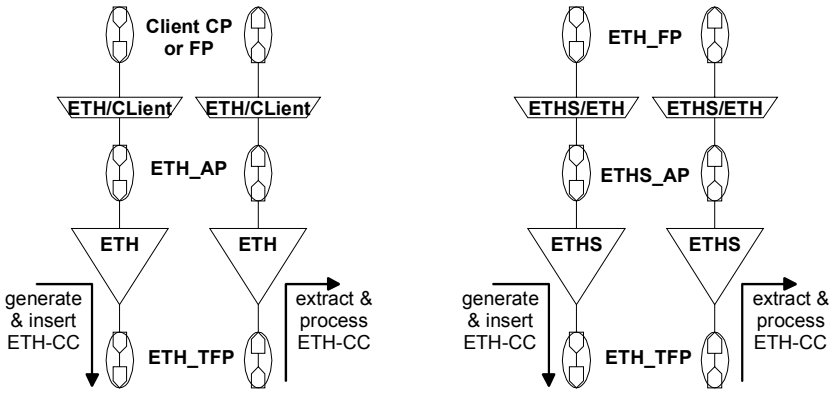

Figure 7-1.1 – Insertion/extraction & processing locations of ETH-CC Signals

[EDITOR'S NOTE-Dec2004] Figure 7-1.1 needs to be revisited to determine the application of 2 parts (Ethernet PATH vs Ethernet SEGMENT0).

The ETH-CC signal is generated with a specific multicast Destination Address (mDA). All MEPs in a MEG should receive this ETH-CC signal. Upon reception of the first ETH-CC signal from a sending MEP, the receiving MEP detects continuity with sending MEP and expects to receive further periodic ETH-CC signals. Once the receiving MEP stops receiving periodic ETH-CC signals from sending MEP, it detects continuity failure. Following detection of continuity failure, the detecting MEP may notify the operator, optionally initiate fault verification followed by optional fault isolation step.

When a MEP starts participating in a network or within a network in a particular service instance for the first time, it may be configured with a list of peer MEP Identifiers. When a MEP does not receive ETH-CC signals from all MEPs in its configured list of peer MEPs, it can detect the continuity failures those MEPs whose ETH-CC signals are missing. Periodicity of ETH-CC signals can be configured at the sending MEP. ETH-CC signal can communicate this periodicity to the receiving MEPs. ETH-CC signal can also communicate the sending MEP identifier to the receiving MEPs.

Figure VIII-1.1 in Appendix VIII shows a multipoint connection with N endpoints with N-1 ETH MEs terminated by each ETHS_FT function. Each of these ETH MEs is to be monitored for continuity. An ETHS_FT_Sk function terminating those N-1 ETH MEs should therefore expect to receive ETH-CC signals from N-1 ETHS_FT_So functions. If less than N-1 ETH-CC signals are received, the ETHS_FT_Sk should be able to state from which of the N-1 ETHS_FT_So functions it is not receiving the ETH-CC signals. If it receives more than expected distinct MEP identifiers, it can determine anomalies (about unexpected entities presence and/or misconnections).

7.1.1 ETH-CC Operations

7.1.1.1 ETH-CC Transmission

Every “active state” MEP can transmit an ETH-CC signal as often as the configured transmission interval. Configured transmission intervals may range from 0.01 seconds to 655.35 seconds. A Lifetime TLV must be transmitted with a value of 3.5 times the configured transmission interval, so that a receiving MEP can lose two ETH-CC signals without declaring a continuity failure. A Lifetime TLV value can range from 1 to 65535 where value 1 corresponds to .035 seconds and value 65535 corresponds to 2293.725 seconds.

7.1.1.2 ETH-CC Reception

Every “active state” MEP or MIP that receives the ETH-CC signal, catalogues ETH-CC signal. Every ETH-CC signal is examined to ensure that its Service Instance Identifier matches that configured in the receiving MEP, and that the MEP Identifier in the ETH-CC signal does not match that of the receiving MEP. The information in the ETH-CC signal is catalogued in the receiving MEP and MIP, indexed by the received MEP Identifier. Information saved includes the Lifetime TLV, so that the information can be timed out, the source MAC address and data path service identifier of the received ETH-CC signal, and the Bridge Port on which it was received. If the value of Lifetime TLV is 0, the catalogued information for the received MEP ID, if any, is discarded.

When an ETH-CC signal is received at either a MEP or a MIP, the source MAC address, data path service identifier, ME Level, and ingress Bridge Port are recorded, indexed by MAC address, , data path service identifier and ME Level, in the Provider Bridge’s ETH-CC Database.

If no ETH-CC signals are received within the time associated with the received ETH-CC’s Lifetime TLV value, loss of continuity with the transmitting MEP is declared.

[EDITOR’S NOTE-Dec2004]Snooping of ETH-CC at MIPs and maintenance of ETH-CC Database at MIPs needs to be made optional to address concerns of heavy-weight ETH-CC and optionality of other OAM functions like ETH-AIS.

[EDITOR’S NOTE-Dec2004]Other OAM functions should not depend on the use of ETH-CC. Information needed for the proper functioning of other OAM functions should be available by other mechanisms.

7.2 Ethernet Loopback (ETH-LB)

ETH-LB can be used to verify connectivity with remote flow point (s). ETH-LB is performed by sending a request ETH-LB signal to remote flow point (s) and expecting an ETH-LB reply signal back which verifies connectivity. When the insertion rate of ETH-LB signals is much slower compared to data rate between the flow points, ETH-LB is suitable to perform in-service connectivity verification.

Though ETH-LB may be initiated at any time, it is particularly useful when verifying connectivity once a continuity failure is detected. A ETH-LB request may be generated either:

- Automatically following detection of continuity failure, where continuity failure can be detected using ETH-CC defined in 7.1, or
- On-demand via an operator initiated command, or
- Periodically.

ETH-LB may be used for fault detection when used on a periodic basis. However, unlike ETH-CC, ETH-LB requires a reply for each request. Reply generation and reply handling by requestor require more processing in ETH-LB than in ETH-CC. While an ETH-CC is suitable for detecting unidirectional continuity failures, ETH-LB can be used to detect bidirectional connectivity failures. ETH-LB signal is generated and inserted in the MEP's ETHD_FT_So functions. It is extracted and processed in the ETHD_FT_Sk functions. Refer to Figure 7.2-1.

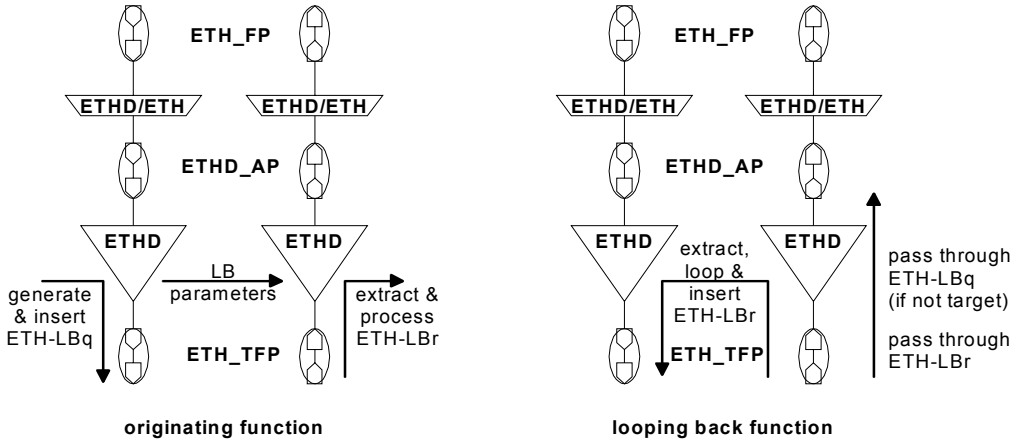

Figure 7.2-1 – Insertion/extraction & processing locations of ETH-LB

ETH-LB can be of two types:

- Unicast ETH-LB
- Multicast ETH-LB

7.2.1 Unicast ETH-LB

Unicast ETH-LB request signal is sent from a MEP to a specific MEP or MIP (with DA = Unicast MAC address of destination flow point). Upon reception of this request signal, the MEP or MIP responds back with Unicast ETH-LB reply signal (with DA = Unicast MAC address of requesting flow point, learnt from request signal). Other flow points that receive this request and/or reply Unicast ETH-LB signal forward these without processing.

Figure VIII-2.1 in Appendix VIII shows a multi-point connection with N endpoints. There are N-1 ETH MEs terminated by an ETHS_FT function. Each of these ETH MEs can be verified for

continuity failures when continuity failures are detected using ETH-CC. An ETHD_FT_Sk function associated with ETHS_FT function can expect to receive Unicast ETH-LB signals from N-1 ETHD_FT_So functions associated with N-1 ETHS_FT functions. For Unicast ETH-LB, ETHD_FT_Sk should receive ETH-LB addressed to itself.

7.2.1.1 Unicast ETH-LB Operations

7.2.1.1.1 Unicast ETH-LB Transmission

Unicast ETH-LB request signal can be transmitted by a MEP either automatically (either periodically or following continuity failure detection) or by operator initiated command (via the CLI or EMS/NMS management interfaces, e.g SNMP MIBs). The Transaction identifier transmitted is retained for at least 5 seconds after the Unicast ETH-LB signal is transmitted. The Transaction Identifier must be changed for every Unicast ETH-LB signal, and no Transaction Identifier from the same MEP may be repeated within one minute.

7.2.1.1.2 Unicast ETH-LB Reception and Reply Transmission

Whenever a valid Unicast ETH-LB request signal is received by a MIP or MEP diagnostic flow termination function, Unicast ETH-LB reply signal is generated and transmitted to the requesting MEP. Every field in the Unicast ETH-LB request signal is copied to the Unicast ETH-LB reply signal with the following exceptions:

1. The source and destination MAC addresses are swapped.
2. The OpCode field is changed from ETH-LB Request to ETH-LB Reply.
3. The Checksum TLV is recalculated to reflect any changes to the message, such as the OpCode field.

7.2.1.1.3 Unicast ETH-LB Reply Reception

When an Unicast ETH-LB reply signal is received by a MIP diagnostic flow termination function, or if the received Transaction ID is not in the list of transmitted Transaction IDs maintained by the MEP, the Unicast ETH-LB reply signal is invalid. The MEP diagnostic flow termination function may examine the TLVs returned in the Unicast ETH-LB reply signal, and declare the signal invalid if the TLVs do not match those sent in the corresponding Unicast ETH-LB request signal.

7.2.2 Multicast ETH-LB

Multicast ETH-LB request signal is sent from a MEP to all other MEPs in a MEG within a OAM boundary (with DA = Multicast DA). Upon reception of this request signal, the receiving MEPs reply back with a Unicast ETH-LB reply signal (with DA = Unicast MAC address of requesting network element, learnt from request signal) after some randomized delay. Finally Unicast ETH-LB reply signal is received and terminated in a source MEP addresses. Refer to Figure VIII-2.3 and VIII-2.4 in Appendix VIII.

The application of Multicast ETH-LB is expected to be limited to diagnostics that may need to be performed before turning on the services. Due to the disruptive nature of Multicast ETH-LB, where a single request signal can result in many reply signals, Multicast ETH-LB is intended for “out-of-service” testing. This can be dictated by requiring destination MEPs or MIPs to respond only if the destination port state is “administratively diagnostics”.

7.2.2.1 Multicast ETH-LB Operations

7.2.2.1.1 Multicast ETH-LB Transmission

Multicast ETH-LB request signal can be transmitted by a MEP via operator initiated command (via the CLI or EMS/NMS management interfaces, e.g. SNMP MIBs). The Transaction identifier transmitted is retained for at least 5 seconds after the Multicast ETH-LB request signal is transmitted. The Transaction Identifier must be changed for every Multicast ETH-LB request signal, and no Transaction Identifier from the same MEP may be repeated within one minute.

7.2.2.1.2 Multicast ETH-LB Reception and Reply Transmission

Whenever a valid Multicast ETH-LB request signal is received by a MEP receive function, Unicast ETH-LB reply signal is generated and transmitted to the requesting MEP following a randomized delay in the range of (TBD). Every field in the Unicast ETH-LB request signal is copied to the Unicast ETH-LB reply signal with the following exceptions:

1. The source MAC address should be the Unicast MAC address of the replying MEP. The destination MAC address should be the unicast MAC address of requesting MEP, learnt from request OAM frame.
2. The OpCode field is changed from ETH-LB Request to ETH-LB Reply.
3. The Checksum TLV is recalculated to reflect any changes to the message, such as the the OpCode field.

7.2.2.1.3 Multicast ETH-LB Reply Transmission

When a Unicast ETH-LB reply signal is received by a MIP diagnostic flow termination function, or if the received Transaction ID is not in the list of transmitted Transaction Ids maintained by the MEP, the Unicast ETH-LB reply signal is invalid. The MEP diagnostic flow termination function may examine the TLVs returned in the Unicast ETH-LB reply signal, and declare the signal invalid if the TLVs do not match those sent in the corresponding Multicast ETH-LB request signal.

[EDITOR'S NOTE-Dec2004] ETH-LB should not depend on the use of CC. Information needed for the proper functioning of other OAM functions should be available by other mechanisms.

7. 3 Ethernet Link Trace (ETH-LT)

The main objectives of an Ethernet Link Trace (ETH-LT) are the following:

- Adjacent Relation Retrieval
- Fault Localization

7.3.1 ETH-LT for Adjacent Relation Retrieval

ETH-LT can be used to identify adjacent MIP and/or MEP. In order to find the relationships, the identification of MIP and/or MEP such as MEP/MIP ID or MAC address, is required. In addition, the sequence of MIPs and/or MEPs in the tested links should be identified

7.3.2 ETH-LT for Fault Localization

ETH-LT can be applied for fault localization. When a fault (eg.in a link and/or bridge failure, etc) and/or a forwarding plane loop occurs, the sequence of MIPs and MEPs in the tested links will likely be different than the expected one. The difference of both sequences provides information of the fault location and faulty element.

ETH-LT request signal is generated and inserted in the ETHD_FT_So functions. It is extracted and processed in the ETHD_FT_Sk functions. Refer to Figure 7.3-1.

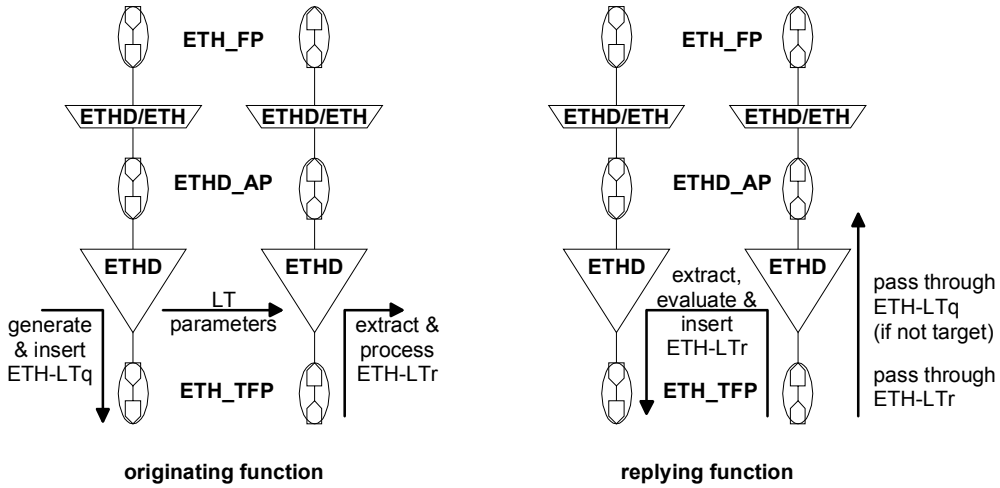

Figure 7.3-1 – Insertion/extraction & processing locations of Non-intrusive ETH-LT OAM

ETH-LT request signal is sent to a target MEP (with DA = Multicast MAC address associated with a ME Level below ETH-CC's Multicast MAC Address) with the target MEP identified in Target Address TLV. Upon reception of this request signal, the MIPs that have a knowledge about the Target Address responds back with Unicast ETH-LT reply signal (with DA = Unicast MAC address of requesting network element, learnt from request signal). These MIPs, which have knowledge about the Target Address also relay ETH-LT request signal out on ports which are associated with the Target Address. Upon reception of this request signal, the MEPs that has the same address as the Target Address responds back with Unicast ETH-LT reply signal (with DA = Unicast MAC address of requesting network element, learnt from request signal) and does not need to relay the request signal any further.

Figure VIII-3.1 in Appendix VIII shows a multi-point connection with N endpoints. There are N-1 ETH MEs terminated by an ETHS_FT function. Each of these ETH MEs can be subjected to Adjacent Relation Retrieval or Fault Localization for continuity failures when continuity failures are detected using ETH-CC. An ETHD_FT_Sk function associated with ETHS_FT function can therefore expect to receive Non-intrusive ETH-LT signals from N-1 ETHD_FT_So functions associated with N-1 ETHS_FT functions.

7.3.3 ETH-LT Operation

7.3.3.1 ETH-LT Origination

ETH-LT request signal can be transmitted by a MEP either automatically (either periodic or following continuity failure detection) or by operator initiated command (via the CLI or EMS/NMS management interfaces e.g. SNMP MIBs) The Transaction Identifier of each ETH-LT request signal transmitted is retained for at least 5 seconds after the ETH-LT signal is transmitted.

7.3.3.2 ETH-LT Reception Link Trace Reception, Forwarding, and Replying

If a ETH-LT request signal is received by a MEP or MIP, and if the data frame targeted by Target Address TLV of ETH-LT signal would pass through the MEP or MIP on ingress or egress through the device, then the device must:

1. Examine the ETH-LT Message's TTL value in TTL TLV, and if 0, discard theETH-LT signal; else
2. Determine the information required to generate a n ETH-LT reply signal to destination address contained in the Source Address TLV of received ETH-LT request signal.;
3. If the data frame targeted by the Target Address TLV of ETH-LT request signal would pass through the device and out a single egress device port , and if ETH-LT signal's's TTL value was greater than 1 when received, then ETH-LT request signal must be relayed on the selected egress port. If the ETH-LT signal's TTL value equals 1 when received, the ETH-LT request signal is not relayed anymore. All fields and TLVs are transmitted exactly as received, except for the source MAC address and TTL value which is decremented by 1.
4. After a random time interval in the range 0-1 second, transmit an ETH-LT reply signal to the originating MEP.

If the data frame targeted by Target Address TLV in ETH-LT request signal would not pass through a MEP or MIP on ingress or egress through the device, then the MIP must pass the ETH-LT request signal through as normal data while the MEP must terminate the ETH-LT request signal.

7.3.3.3 ETH-LT Reply Reception

When an ETH-LT reply signal is received by a MIP receive function, or if the received Transaction ID is not in the list of transmitted Transaction Ids maintained by the MEP, the ETH-LT reply signal is invalid.

Figure 7.3-1 and Table 7.3-1 below show the packet handling and the packet information table by hop respectively

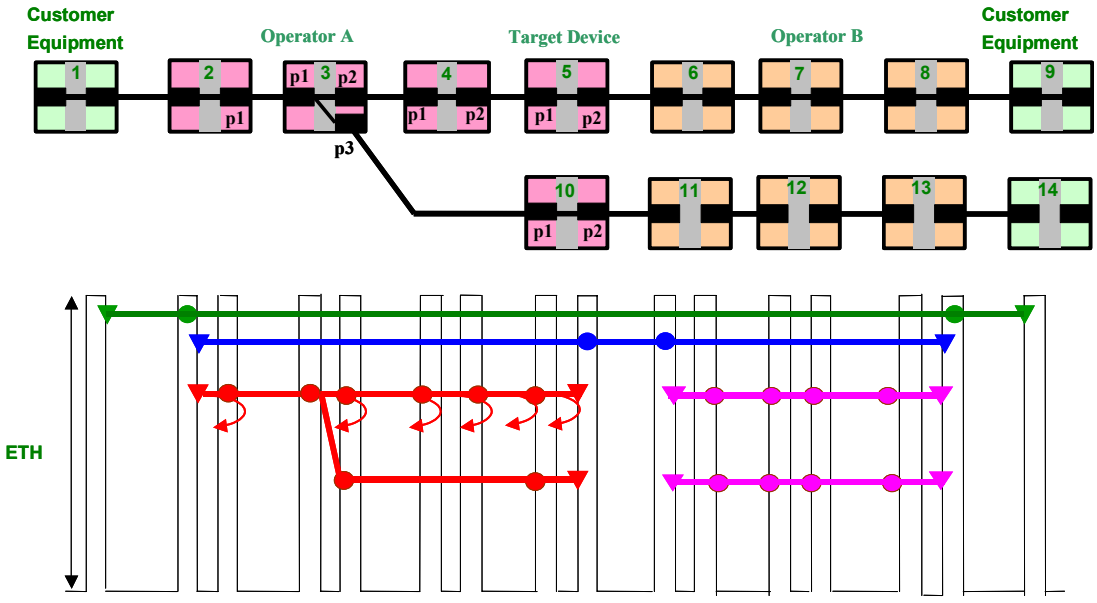

Figure 7.3-1: ETH-LT packet handling

Number of hops	Sent ME ID	Received ME ID
1	2/p1	3/p1
2	3/p1	3/p2

3	3/p2	4/p1
4	4/p1	4/p2
5	4/p2	5/p1
6	5/p1	5/p2

Table 7.3-1: Sorted ETH-LT packet information table by number of hops

[EDITOR'S NOTE-Dec2004] ETH-LT should not depend on the use of CC. Information needed for the proper functioning of ETH-LT should be available by other mechanisms.

7.4 ETH-AIS

ETH layer Alarm Indication Signal (ETH-AIS) can be used to notify client layers about faults detected at server layers such that the ETH-AIS can be used to suppress declaration of same fault at client layers. This allows the fault to be reported to OSS (Operations Support Systems) or NMS (Network Management Systems) by a single layer (at which the fault occurs and is detected) and not by all other higher layers.

Note: The current version describes applicability of ETH-AIS for point to point services offered across infrastructure where automatic reconfiguration mechanisms like STP are not used. Appendix III highlights some scenarios and issues associated with the multipoint services including when a service has only 2 endpoints.

Figure VIII-4.1 in Appendix VIII shows, as an example, how a fault at the ETY layer can be notified via ETH-AIS to higher level MEs.

A Server layer or ETH sublayer MEP Sink function that detects a signal fail condition will insert ETH-AIS in its Srv/ETH_A_Sk[1] or ETHS/ETH_A_Sk function. A ETH sublayer MEP Sink function that detects ETH-AIS at its ME Level will terminate the signal in its ETHS_FT_Sk function, detects dAIS, declares a signal fail condition and inserts in its ETHS/ETH_A_Sk function ETH-AIS (at the higher level).

The termination and re-generation of ETH-AIS within an ETH sublayer MEP Sink function provides some security by preventing internal MEP MAC addresses to be exposed outside a ME domain. Note that a MIP function is transparent to ETH-AIS.

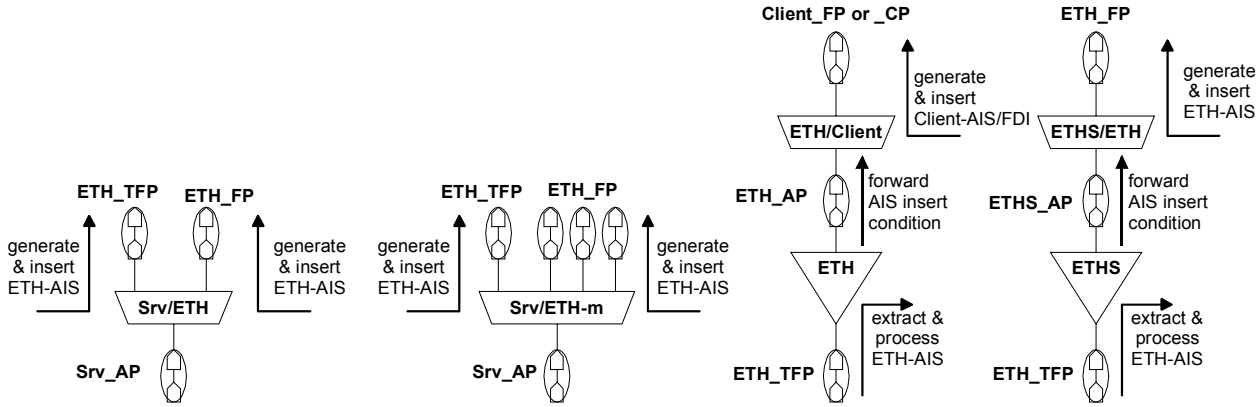

Figure 7.4-2: Insertion/extraction & processing locations of ETH-AIS

[1] This Srv/ETH_A_Sk function will be part of a server layer's MEP.

[EDITOR'S NOTE-Dec2004] ETH-AIS should not depend on the use of CC. Information needed for the proper functioning of ETH-AIS should be available by other mechanisms.

[EDITOR'S NOTE-Dec2004] Refer to Appendix VIII section VIII-4 for more discussion on AIS insertion and extraction points and Appendix III for more discussion on AIS behaviour and issues. Further contributions are invited.

7.5 ETH-RDI

The application of ETH layer Remote Defect Indication (ETH-RDI) is for further study.

Note:

ETH layer is dependent upon an operational ETH Link, where both transmit and receive directions are up. When either transmit or receive direction is physically down at a port of an ETH link, entire port and associated link is marked as operationally down when the auto negotiation function is operated.

However the auto negotiation function is optional, some carriers do not use this function. Even though auto negotiation is operated, it is probable that the connectivity failure occurs for the software failure without physical failure. The auto negotiation function will not operate in this case. Therefore the unidirectional down is not necessarily a rare case, and ETH-RDI is applicable for some cases mainly for point to point case. A following figure shows the ETH-RDI Flow in terms of MIP/MEP model. Furthermore ETH-RDI may be applicable for performance management. Another possible application is to differentiate between administrative shutdown and failure shutdown.

The application for point to multipoint case is F.F.S.

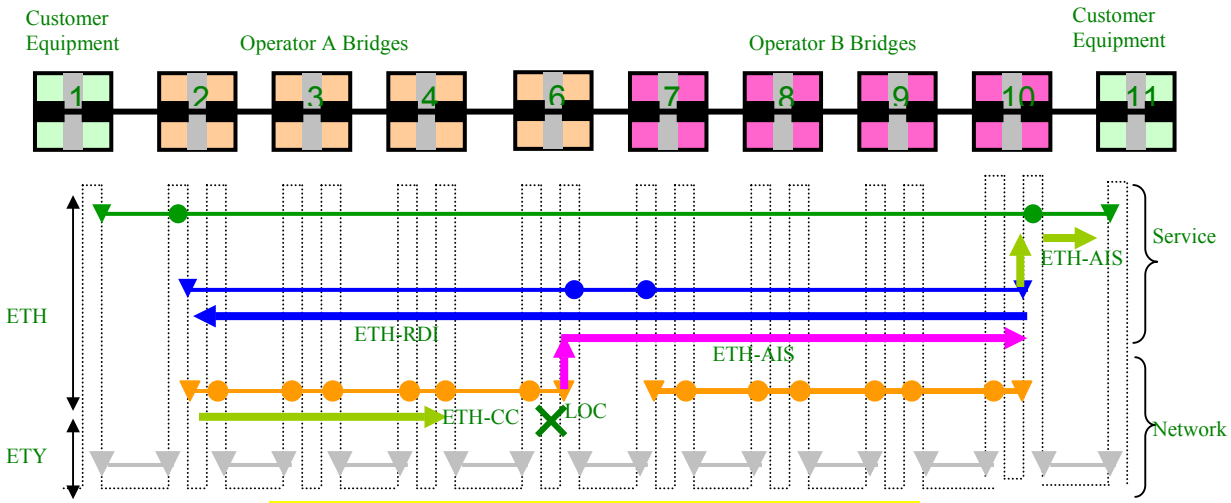

Fig. ETH-RDI Flow in terms of MIP/MEP model.

7.6 Test Signal Generation/Detection Function

The test signal generation function in a MEP/MIP generates test signal with specified throughput (bandwidth), frame size and frame transmission pattern. The detection function in a MEP/MIP detects throughput (bandwidth), frame loss, frame disorders, bit errors, delay and delay variation.

7.6.1 Test modes

This test function can be used in-service and out-of-service.

When out-of-service is conducted, service cannot be offered to the user. For example, this type of test can be used for pre-service test.

Service can be offered to the user when in-service test is conducted. However, since this test uses some of the bandwidth of the service, agreement needs to be made between the user and the network operator on the bandwidth usage.

7.6.2 Frame format

Since test function needs to be done out-of-service and in-service, Ethernet OAM frame format needs to be used so that test frames can be distinguished from normal user data frame.

7.6.3 OAM data

In order to measure frame loss and bit error performance, 32bit sequence number and pseudo-random test sequence (2^31-1) as specified in 5.8/O.150 are included.

7.6.4 Maintenance scenarios

This section shows some examples of maintenance scenarios for point-to-point, in-service and out-of-service case. NOTE: Multipoint application is FFS.

7.6.1.1 Unidirectional Measurement

A MEP at a n edge bridge generates a test signal and another MEP in an edge bridge receives the test signal and measures the performance between these two edge bridges MEPs (Fig. 7.6-1). This scenario is applicable both to in-service test and out-of-service test.

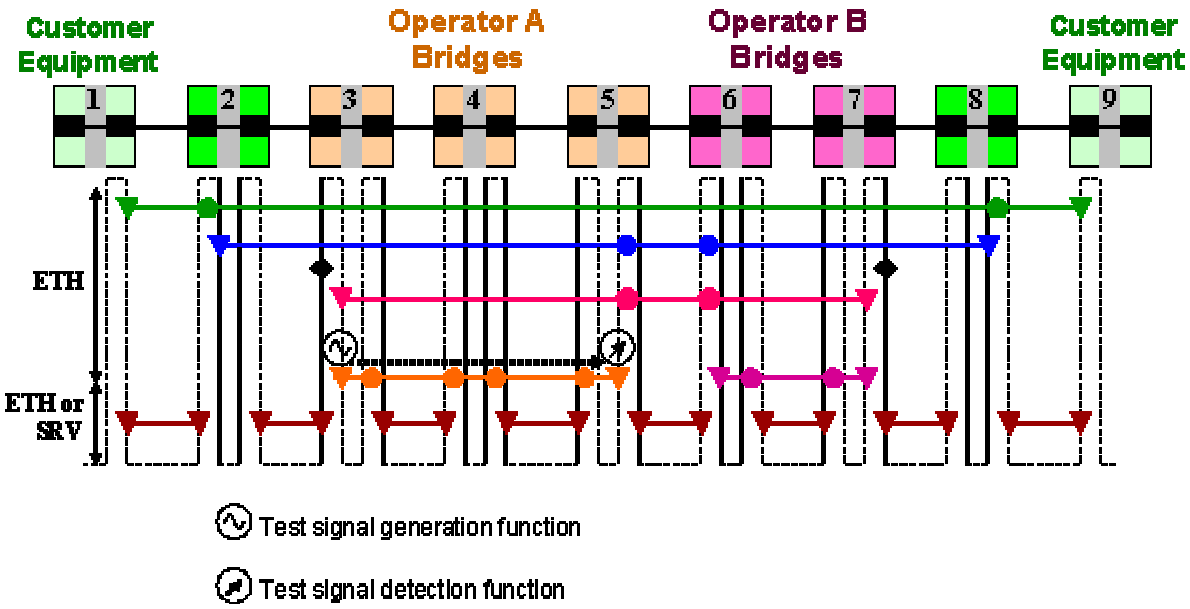

Figure 7.6-1 Unidirectional test.

7.6.1.2 Bidirectional Measurement

A MEP/MIP in a core bridge or an edge bridge generates a test signal. Another MEP/MIP in a edge bridge, core bridge or a access provider device (ex device 8) is put into an intrusive loopback mode. The MEP/MIP generating the test signal sends the test signal towards the MEP/MIP in the intrusive loopback mode and receives the loopbacked test signal. Bidirectional (round trip) performance between these MEPs/MIPs is measured with this (Figs. 7.6-2 and 7.6-3). These scenarios are applicable only to out-of-service test.

EDITER'S NOTE: REALIZING INTRUSIVE LOOPBACK AT HIGH BIT RATE MIGHT CAUSE EXCESSIVE COMPLEXITY. DETAIL MECHANISM TO REALIZE THIS FUNCTION IS FFS. REFER TO LOOPBACK SECTION FOR CURRENT DISCUTION.

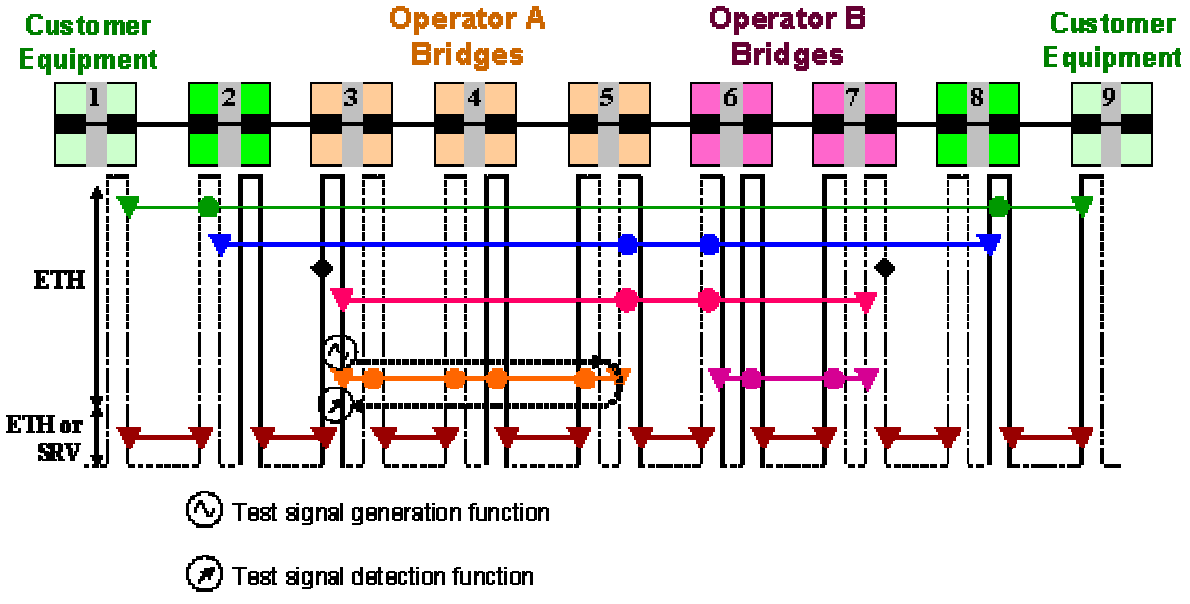

Figure 7.6-2 Bidirectional test (1).

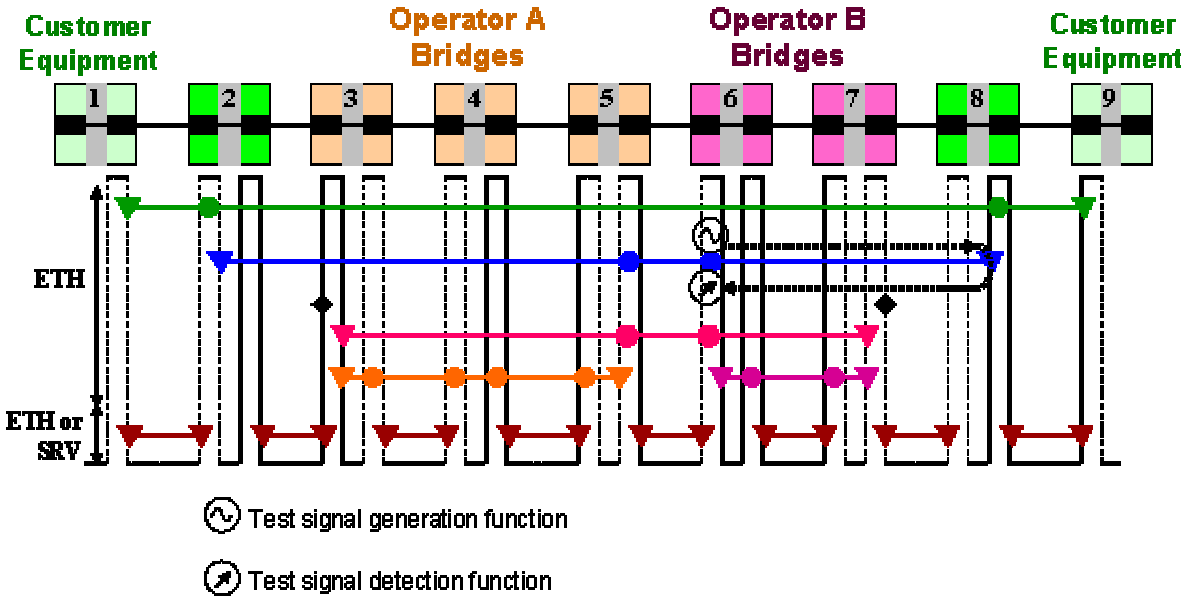

Figure 7.6-3 Bidirectional test (2).

7.7 Ethernet Loopback State Request (ETH-LS)

ETH-LS OAM is used for a request for the local node to be on the loopback state. This is mainly used for the test function. This function is only applicable to the point-to-point case.

There are 4 messages for ETH-LS OAM:

- ETH-LS Set Request (ETH-LS SReq)
- ETH-LS Get Request (ETH-LS GReq)
- ETH-LS Set Response (ETH-LS SRes)
- ETH-LS Get Response (ETH-LS GRes)

The procedure is as following:

1) Set Loopback state

- One MIP/MEP sends an ETH-LS SReq frame.

- The encountered node becomes loop state if it receives ETH-LS SReq frame.

- The encountered node replies ETH-LS SRes that shows whether successful or not.

- The former MIP/MEP judges successful only if it receives reply OK. Other cases, such as NG reply or frame missing, should be not successful.

2) Get Loopback state

- One MIP/MEP sends an ETH-LS GReq frame.

- The encountered node replies ETH-LS GRes that shows the loopback state.

- The former MIP/MEP receives ETH-LS GRes frame.

There are following options of loopback state that is requested via the ETH-LS OAM:

1) Looping back all received frames without any modifications

2) Looping back all received frames where DA and SA are replaced

3) Looping back only ETH-OAM frames where DA and SA are replaced

4) Looping back only for OAM with exact MAC address where DA and SA are replaced

One of the diagnostic states for the ports can be "Loopback".

The details for ETH-LS are F.F.S.

[EDITOR'S NOTE-Dec2004] Based on D-43, Section 7.7 created;

The Loopback functionality that is able to operate by sending test-frames at wire speed or the maximum committed rate up to the customer interface and/or intermediate points is used as a diagnostic tool to detect service degradation or loss of service from a remote site. The Loopback functionality can be performed in three different ways:

1) Looping back all received frames, except for ETH-OAM frames with specific OAM type (Intrusive loopback)
2) Looping back only ETH-OAM frames with a specific OAM type
3) Looping back only ETH-LB frames (with the exact MAC address)

EDITOR'S NOTE: CONTRIBUTIONS ARE NECESSARY TO SHOW DIFFERENT FORMS (THE THREE ABOVE) OF LOOPBACK AND THEIR ADVANTAGES AND DISADVANTAGES

EDITOR'S NOTE: Above three options should be categorize as below:

a. Slow non-intrusive loopback (e.g., less than 1 frame/sec)

b. Fast non-intrusive loopback (option 2 and 3) (e.g., wire speed)

c. Intrusive loopback (option 1)

7.7.1 Intrusive ETH-LB

An Intrusive ETH-LB loopbacks all received frames with the exception of ETH-OAM frames with a specific OAM type. It is intrusive because it affects traffic so it should only be used for out-of-service testing. This function is only applicable to the pont-to-point case.

NOTE: THERE IS A POTENTIAL DANGER OF HAVING TWO LB POINTS WHICH COULD LEAD TO A LOOP IN THE NETWORK.

EDITOR'S NOTE: CLARIFY HOW THE POTENTIAL DANGER COULD BECOME REAL AND ALSO CLARIFY WHAT MEANS WIRE SPEED LOOBPBACK COMPARED TO OTHER LOOPBACKS FRAMES

[EDITOR'S NOTE-Dec2004]Other OAM functions should not depend on the use of CC. Information needed for the proper functioning of other OAM functions should be available by other mechanisms.

8 OAM functions for Performance Management

8.1 Performance Parameters

Current discussions in Metro Ethernet Forum on performance parameters for Ethernet networks and services have focused on the following parameters, which are captured in an MEF working draft (January 2003)

- **Frame Loss (FL)**
  Difference between the number of service frames sent to ingress UNI and the number of service frames received at egress UNI. This is applied to Ethernet Virtual Connection (EVC) which corresponds to UNI_N to UNI_N ME.
- **Frame Delay (FD)**
  Frame delay can be specified in terms of round-trip delay, which is defined as the time elapsed since start of the transmission of the first bit of a frame by the source node until the reception of the last bit of the loop backed frame by the same source node, when the loop back is performed by the frame's destination node.
- **Frame Delay Variation (FDV)**
  Measure of the variations in the frame arrival pattern belonging to the same CoS instance compared to the arrival pattern at the ingress of the MEN.
- **Availability**
  Function of time the ME (associating service UNIs) is in available state. It is specified as a ratio of:

  Total Time ME is in Available State / Total Service Time

  where, **Total Service Time** is viewed as number of time intervals and **Available State** is viewed as interval when service meets FL, FD and FDV bounds. Unavailable state is encountered when at least one of the FL, FD or FDV measures exceed their bounds/thresholds during a time interval. These bounds/thresholds are determined by the class of service (CoS).

  Note 1:
  The definition of Availability should be aligned with Y.1711 and/or Y.MPLSperf . The details of Availability are expected to be defined in a separate Recommendation being developed by Q.6/13 – Y.17EthPerf.

  Note 2:

  For sub rate or virtual services, the frame loss can be associated with both in-profile and out-of-profile service frames.

  Additional performance parameters that may be taken into consideration include:
- **Errored Frame Seconds**
  Indicates if an error (e.g., frame error due to FCS or 8B/10B coding violation) has occurred within the second. This does not take into consideration errors when frames are received error free but are not delivered.
- **Service Status**
  Indicates if the service is in-service or out-of-service. In-service or out-of-service state can be based on **Available state** defined earlier.

- **Frame Throughput**
  Number of frames and/or bytes transmitted to network interface relative to CIR
- **Frame Tx**
  Number of frames transmitted out of the customer facing interface within the (previous) time interval (e.g. 1 second).
- **Frame Rx**
  Number of frames received from the customer facing interface within the (previous) time interval (e.g. 1 second).
- **Frame Drop**
  Number of frames dropped at the customer facing interface within the (previous) time interval (e.g. 1 second).
- **Loopback Status**
  Indicates whether the customer facing interface is in an intrusive Loopback state (potentially due to OAM interactions across Access Link ME).
- **Client Signal Fail**
  Indicates state of Access Link ME.
- **Unavailable Time**
  Number of time intervals (e.g. 1 second) when the service status is unavailable.

8.2 Measurement Mechanisms

Different measurement mechanisms are possible to perform performance measurements. One significant difference across these mechanisms is the level of accuracy of measurements. These mechanisms include:

- **Management plane statistical methods**
  Statistical methods use OAM frames to estimate data path behavior. Such methods are least accurate since they apply approximation to emulate data frames.
  The limitation lies in that the behavior of actual data frames may be quite different due to different addressing, processing, transient congestion conditions etc. Also, error conditions in networks typically happens in bursts thus statistical methods can likely miss those bursts and represent different results.
- **Management plane managed objects**
  Here OAM frames use data path managed objects to calculate performance parameters and are inserted and/or extracted via management plane. These methods are fairly accurate since they use data path statistics to measure data path performance.
  Their limitation lies in that since the insertion and extraction of these OAM frames is done via management plane, in-flight frames need to be accounted for. On the egress side of OAM frame, in-flight frames refers to data frames between accessing egress data path managed objects and actual transmission of OAM frame. On ingress side of OAM frame, in-flight frames refer to data frames between reception of OAM frame and subsequent accessing of ingress data plane managed objects. However, this limitation can be addressed by averaging such measurements across multiple time intervals.
- **Data path OAM frames**
  OAM frames use data path managed objects and are inserted and/or extracted via data plane. This method tends to be most accurate since it does not have the limitation associated with the in-flight frames.

However, the current data path hardware/chips do not support the implementation of such methods since this requires Ethernet data path processing to include automatic insertion and/or extraction of OAM frames with data plane managed object values. Moreover, it would also require changes in hardware/chips to allow ingress and egress filtering rules across OAM frames to protect service provider administrative domains from unintended OAM frames.

Of the three methods mentioned to measure performance the use of management plane managed objects mechanism seems to be the most suitable. The advantage of these mechanisms is that these require no changes in the existing hardware/chips and only require change in OAM client software that needs to be implemented. The steps involved in such measurement mechanism include:

- Collection of managed object (s) information
- Calculation of performance parameter (s)

8.2.1 Performance Management Collection Method

To collect managed object information, general or specific methods can be used. When a generic method is used, it can be applied to collect information across different managed objects e.g. using TLVs as information elements instead of specific information elements. However, when specific method with specific information elements is used, a separate method is needed per managed object or per set of managed objects.

Similarly, it is possible to use either a solicited or unsolicited collection method, where solicited method requires a response after an OAM request frame is sent while unsolicited methods does not require a response to an OAM frame. Some current examples of solicited and unsolicited methods include Loopback and Continuity Check respectively, though these are currently not used as performance management collection methods.

A generic method to send/receive data path managed object information can be used. This is similar to the variable request/response method used in IEEE 802.3ah [section 57.4.3.3/.4]. Also both solicited and unsolicited methods can be used and optionally extend the currently defined Loopback [section 7.2] and Continuity Check [section 7.1]. Note that this extension for PM will require additional processing and therefore should not be used for the measurement of delay.

EDITOR'S NOTE: INPUT ON ATOMIC FUNCTION MODEL FIGURE FROM MAARTEN EITHER FOR HERE OR FOR APPENDIX

8.2.2 Frame Loss Measurement

MEs which can support Frame Loss include:

- Service MEs for point-to-point service with dedicated UNIs
    - UNI_C to UNI_C
    - UNI_N to UNI_N
    - Access Link (UNI)
    - Inter-domain (NNI)
- Network MEs
    - Intra-domain

- Inter-domain

8.2.3 Unsolicited Method

When applied across UNI_N to UNI_N ME, OAM frame is sent every N seconds (e.g. N=1) with **FramesTransmittedOK** value at ingress service UNI. Upon receiving this OAM frame, **FramesTransmittedOK** value is compared with **FramesReceivedOK** value at egress service UNI. Between two such consecutive OAM frames, the FL can be measured as:

Frame Loss = |CT2-CT1| - |CR2-CR1|,

where CT and CR are **FramesTransmittedOK** and **FramesReceivedOK** counts. Consecutive messages help in reducing error introduced by in-flight frames and lack of timing synchronization between sender and receiver. Within a measurement time interval, the Frame loss count can be averaged to improve the accuracy of this measurement.

NOTE: For measurement considerations with possible wrapping of CT/CR, refer to Appendix V

8.2.4 Solicited Method

Requestor sends OAM request frame to receiver every N seconds (e.g. N=1) with its managed objects (MOs) information and expects an OAM response frame with receiver's MOs information.

When applied across UNI_C to UNI_C ME, requestor sends **FramesTransmittedOK** value at egress service UNI and requests **FramesReceivedOK** value from receiver's ingress service UNI. Similarly, when applied across UNI_N to UNI_N ME, requestor sends **FramesReceivedOK** value at ingress service UNI and requests **FramesTransmittedOK** value from receiver's egress service UNI

Upon receiving the OAM request frame, receiver compares received MO information with its corresponding MO information and sends a response OAM frame back to requestor with requested MO information. When applied across UNI_C to UNI_C ME, receiver compares received **FramesTransmittedOK** value with **FramesReceivedOK** value and responds with its **FramesTransmittedOK** value. Similarly, when applied across UNI_N to UNI_N ME, receiver compares received **FramesReceivedOK** value with its **FramesTransmittedOK** value and responds with its **FramesTransmittedOK** value.

Upon receiving OAM response frame, requestor compares original sent value with received values, similar to receiver. It is possible that receiver returns the results of frame loss instead of MO information in response, however, if the MO information is returned, the performance collection method remains generic.

Between two such consecutive OAM frames, the FL can be measured as:

Frame Loss = |CT2-CT1| - |CR2-CR1|,

where CT and CR are **FramesTransmittedOK** and **FramesReceivedOK** counts. Consecutive messages help in reducing error introduced by in-flight frames and lack of timing synchronization between sender and receiver. Within a measurement time interval, the Frame loss count can be averaged to improve the accuracy of this measurement.

NOTE: For measurement considerations with possible wrapping of CT/CR, refer to Appendix V

The above method can be applied for measuring network level Frame Loss. The network level frame loss can be measured within the network independent of the services.

For non-dedicated point-to-point service types with multiplexed service UNI, where a UNI carries more than one service flow, it is possible to measure FL when data path MOs per service instance are supported.

8.2.5 Statistical Method

For multipoint-to-multipoint service type, statistical method across a pair of UNIs can be applied to estimate frame loss.

The requestor sends N OAM request frames to the recipient and receives M response frames back from the recipient such that M <= N. The data path frame loss can be estimated as:

Frame Loss = (N – M) per measurement time interval

As noted earlier, statistical methods are less accurate than proposed method in this contribution.

8.3 Frame Delay Measurement

Services supported include point-to-point and multipoint-to-multipoint between a given pair of UNIs.

MEs across which the frame delay can be measured are:

- Service MEs
    - UNI_C to UNI_C
    - UNI_N to UNI_N

Solicited Method – Loopback

This method measures round-trip or two-way frame delay. Requestor sends OAM request message with its timestamp to the receiver. Receiver replies copying the requestor's timestamp. At the requestor, the difference between the timestamps at the time of receiving the OAM response frame and original timestamp in the OAM response frame results in round trip frame delay.

8.4 Frame Delay Variation Measurement

Services supported include point-to-point and multipoint-to-multipoint between a given pair of UNIs.

MEs across which the frame delay variation can be measured are:

- Service MEs
  - UNI_C to UNI_C
  - UNI_N to UNI_N

Solicited Method – Loopback

This method measures round-trip or two-way frame delay per request and response frame. Within the period of observation, requestor keeps track of maximum frame delay (FD_{max}) and minimum frame delay (FD_{min}). Frame delay variation is calculated as:

Frame Delay Variation or Jitter = FD_{max} - FD_{min}

Information elements for FDV method in OAM Data mentioned in Y.17ethoam [3 - section 15.1] include:

- Sequence number
- Request Timestamp
  - FDV or Jitter = {FD (max) – FD (min)} per measurement time interval
  - Information elements for FD method in OAM Data
    - Sequence number
    - Request Timestamp

8.5 Availability Measurement

Services supported include point-to-point with at least dedicated UNIs.

MEs across which the frame delay variation can be measured are:

- Service MEs
  - UNI_C to UNI_C
  - UNI_N to UNI_N

Measurement Method

Measurement is based on FL, FD and FDV methods. Availability time interval (e.g. 24hr) can be divided into measurement time intervals (e.g. 1 minute). FL, FD and FDV are measured per measurement time interval. If any of the three measures crosses its corresponding thresholds, which are dependent on the service type, the measurement time interval is considered to be unavailable else it is considered to be available.

Availability = (# of available measurement time intervals)/(# of total measurement time intervals)x100%

NOTE: Mechanisms that can be used to measure availability are being proposed here but they will depend on the definition of availability and further details expected to be specified by Ethernet Traffic Management activities (q4/13 and q6/13).

8.6 Other Measurements

As per the unsolicited method explained before, the following parameters can be sent every time interval (e.g. 1 second) to the peer.

8.6.1 Errored Frame Seconds

ME: Access Link ME

Within 1 second, check if any increments in (aFrameCheckSequenceErrors, aAlignmentErrors, aFramesAbortedDueToXsColls, aFramesLostDueToIntMACXmitError, aCarrierSenseErrors, aFrameLostDueToIntMACRcvError)

If yes, declare that 1 second as errored frame second

8.6.2 Service Status

ME: UNI_C to UNI_C ME or UNI_N to UNI_N ME

Within the measurement time interval (e.g. 1 min), declare whether the service is up or down as per availability measurement, explained earlier

8.6.3 Frame Throughput

ME: UNI_N to UNI_N

Within the measurement time interval, aFramesTransmittedOK at egress UNI_N relative to CIR

8.6.4 Frame Tx

ME: Access Link ME

Within 1 second, aFramesTransmittedOK at egress UNI_N

8.6.5 Frame Rx

ME: Access Link ME

Within 1 second, aFramesReceivedOK at ingress UNI_N

8.6.6 Frame Drop

ME: Access Link ME

Within 1 second, ifInDiscards at ingress UNI_N and ifOutDiscards at egress UNI_N.

8.6.7 Loopback Status

ME: Access Link ME

aLoopbackStatus at UNI_N.

8.6.8 Client Signal Fail

ME: Access Link ME

aLinkStatus at UNI_N.

8.6.9 Unavailable Time

ME: UNI_N to UNI_N

This is related to availability definition with the unavailable time intervals being counted within the observation period.

[EDITOR'S NOTE-Dec2004] PM should not depend on the use of CC. Information needed for the proper functioning of PM should be available by other mechanisms.

EDITOR'S NOTE: THE DISCOVERY MATERIAL WAS NOT CONTRIBUTED TO OR DISCUSSED DURING THE LAST THREE MEETINGS SO CONSULTED WITH RAPPORTEUR , IT IS PROPOSED TO DELETE MATERIAL AS BELOW

9 Information Elements

EDITOR'S NOTE: THE COMMON INFORMATION ELEMENTS AND THE ONES EXCLUSIVE FOR CC HAVE BEEN UPDATED AT THE MEETING OF FEBRUARY 2004

9.1 Common Information Elements

- Addressing (DA, SA MAC)
- VLAN ID
- ME Level
- OAM EtherType
- Version
- OAM OpCode
- MPID
- ServiceID
- TransactionID

The Maintenance Point ID (MPID) is necessary as a change of hardware (e.g. an IF card or a bridge is removed/replaced) will imply a change in the MAC address. Each port has a corresponding MAC address.
The Service ID is necessary to identify a service instance. It is globally unique

Note that an MP represents either a MEP or a MIP.

Both the MPID and the Service ID may be TLVs.

Their corresponding position within the generic frame format is FFS.

The MPID is unique within a service instance.

The VLAN ID represents the data plane service instance identifier, when used.

EDITOR'S NOTE: SHOW SOME PROVISIONING MODEL IF POSSIBLE TO ADD INTO AN APPENDIX

EDITOR'S NOTE: COMPARE FUNCTION SPECIFIC INFORMATION ELEMENTS AND DELETE THE ONES THAT ARE COMMON TO ALL AS THEY ARE CAPTURED IN SUBSECTION CALLED COMMON INFORMATION ELEMENTS ABOVE

9.2 Specific Information Elements for Connectiviy Check

- Required CC Information Elements
  - Source MEP ID

Other Functionality/Information Elements are for further study:
e.g.
  - CC Expiry Indication
  - MEP Status - CC Activation/Deactivation Indicator

EDITOR'S NOTE: THE FOLLOWING TEXT REGARDING CC HAS BEEN DELETED AS THERE IS UPDATED (JUNE 2004) MATERIAL IN SECTION 7.1

9.3 Specific Information Elements for Non-intrusive Loopback

EDITOR'S NOTE: CHECK THE FOLLOWING INFORMATION ELEMENTS AS THEY WERE NOT REVISED FOR THE LAST TWO MEETINGS (JUNE 2004)

- Fault Detection
  - OAM Frame Identifier – (Detection of Loops, correlation)
  - Source Port Number – (Identification of specific source port, handle)
  - Destination Port Number – (Identification of specific target port, handle)?
  - Arbitrary data part – (Stress Test: could be used to test different MTUs, pad packet to full size, put worst case patterns)?
  - Note: Addressing used for this case is option (D) i.e. Unicast Destination MAC Address. Needs validation.
- Fault Localization
  - OAM Frame Identifier – (Detection of Loops, correlation)
  - Source Port Number – (Identification of specific source port, handle)
  - Destination Port Number – (Identification of specific target port, handle)?
  - Arbitrary data part – (Stress Test: could be used to test different MTUs, pad packet to full size, put worst case patterns)?
  - Note: Addressing used for this case is option (D) i.e. Unicast Destination MAC Address. Needs validation.
- Performance – for round-trip Delay and Jitter
  - OAM Frame Identifier – (Detection of Loops, correlation)
  - Source Port Number – (Identification of specific source port, handle)?
  - Destination Port Number – (Identification of specific target port, handle)?
  - Source Timestamp
  - Note: Addressing used for this case is option (D) i.e. Unicast Destination MAC Address. Needs validation.

Extra Elements:
  - Response MAC Address
  - Randomized Delay

NOTE: loopback will not be used for discovery purposes because of potential storms in DOS scenarios due to number of replies.

9. 4 Specific Information Elements for Link-Trace (Body)

- OAM Frame Identifier –
- Source Port Number – (Identification of specific source port, handle)
- Destination Port Number – (Identification of specific target port, handle)
- TLV for Checksum – (checksum for part that cannot be changed)
- Target MAC Address
- Source MAC Address
- Hop Count
- Others
  - Periodicity of Loopback – (when used proactively)
  - Randomized Delay - ?

9.5 Performance Monitoring Information Elements

EDITOR'S NOTE: THE FOLLOWING THREE SUBSECTIONS SHOULD BE CONSOLIDATED IN ONE AS THE PERFORMANCE MONITORING OAM FUNCTION SHOULD HAVE INFORMATION ELEMENTS THAT SERVE ALL OF THE PURPOSES

9.5.1 Information elements that can be applied to OAM Data for the Unsolicited Method

- Sequence number
- # of TLVs
- TLVs (Managed Object variable: **FramesTransmittedOK**, value length, value)

9.5.2 Information elements that can be applied to OAM Data for the Solicited Method

- Sequence number
- # of Transmit TLVs (value filled in by requestor, recipient simply copies it back in response)
- # of Request TLVs (value is filled in by recipient and sent back in response)
- TLVs (Managed Object variable: **FramesTransmittedOK & FramesReceivedOK**, value length, value)

9.5.3 Information elements for Frame Delay method in OAM Data

- Sequence number
- Request Timestamp

10 OAM Frame Formats

10.1 Generic OAM Frame Format

A single generic format is defined for all Ethernet OAM frames as shownin Figure10-1 VLAN (VLAN Ether Type + VLAN tag) is optional and if it is present, it indicates the service instance corresponding to the OAM frame. The OAM Ethernet Type is TBD and it identifies the frame as an OAM frame. It should be noted that all the OAM frames carry the same OAM Ethernet Type.

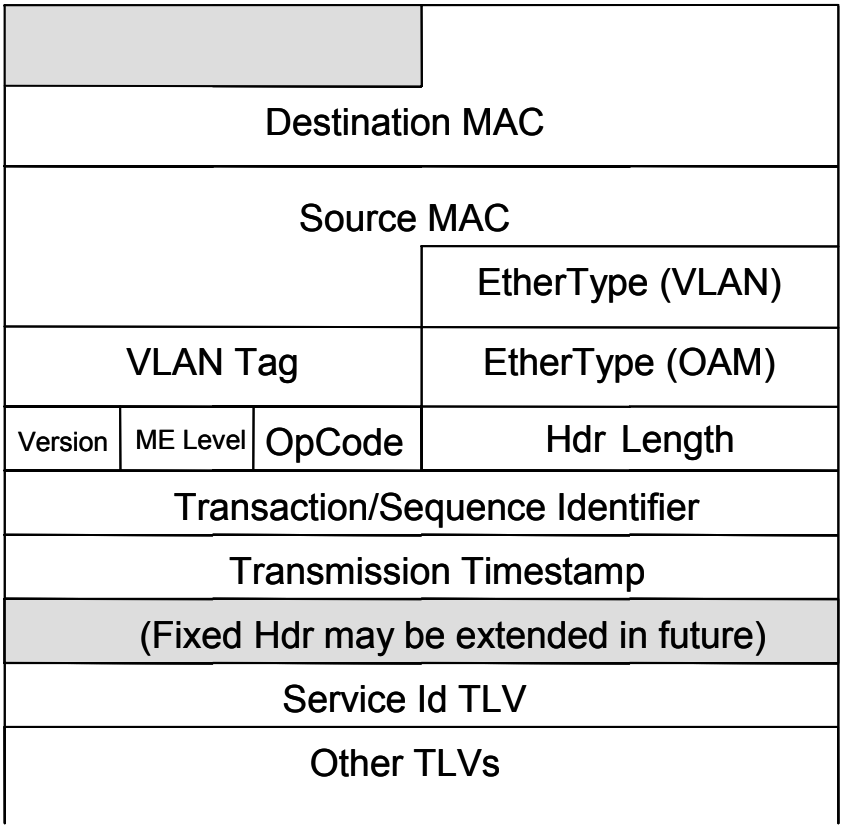

Figure 10-1: Generic OAM Message Format

EDITOR'S NOTE: Following items are under study,

- Bit allocation of Version and ME Level fields;
- Whether timestamp field should be in the common part (non-TLV) or in TLV;
- Whether indication of timestamp is used or not done by Opcode or by a reserved value of timestamp field, or a flag (a bit in the timestamp field).

The fields for the generic OAM frame format are defined as follows:

- **Destination MAC Address**: The destination MAC address can be one of one of the unique Multicast addresses (TBD, and used by ETH-CC and ETH-LT) or a Unicast address of a MEP
- **Source MAC Address**: This is typically the Unicast MAC address of the source MEP. For ETH-LT, it identifies the Unicast address of a MIP or a MEP.
- **VLAN Ether Type and VLAN Tag**: This is an optional field and is used as data plane service identifier at ETH layer.
- **OAM Ethernet Type**: This is a unique Ethernet Type that identifies OAM frames.
- **Version**: The Version field identifies the OAM protocol version. Value for current version is 0x00
- **ME Level**: ME Level identifies the administrative domain of the OAM frame. The value ranges from 0x00 to 0x07. Values 0x00-0x002 identify an operator domain, 0x03-0x04 identify provider domain, and 0x05-0x07 identify a customer domain.

- **OpCode**: The **OpCode** defines the type of OAM frame. OAM frames with unexpected unknown op-codes MUST be silently discarded. The OAM frame types that are defined in this recommendation are:
    - **ETH-CC (0x00)**
    - **Intrusive ETH-LB Reques (0x01)**
    - **Non-intrusive ETH-LB Request (0x02)**
    - **ETH-LB Reply (0x03)**
    - **ETH-LT Request (0x04)**
    - **ETH-LT Reply (0x05)**
    - **ETH-AIS (0x06)**
    - **ETH-RDI (0x07)**
    - **Performance Monitoring Request (0x08)**
    - **Performance Monitoring Reply (0x09)**
    - **Vendor Specific (0xFF)**. The vendor specific op-code is provided to allow vendors or other organizations to extend OAM functions in proprietary ways.
    - **Others:** Other OpCodes may be defined in future.
- **Hdr Length**: The number of bytes in the fixed-length header, starting with the Version field.
- **Transmission/Sequence Identifier**: Supplied by the originator of OAM request and copied in the OAM reply. Semantics of this field are dependent on the OpCode.
- **Transmission Timestamp**: Time at which the OAM frame was transmitted from originating MEP
- **Service ID**: The first TLV that identifies the service instance
- **Other TLVs**: These TLVs correspond toing OpCodeOAM type and sub-type and its format is dependent on the OAM type/sub-type fields.

Note: For further protocol details related to Connectivity Fault Management Related OpCode (e.g. ETH-CC, ETH-LB, ETH-LT, ETH-AIS and ETH-RDI), refer to IEEE 802.1ag.

Annex A

EDITOR'S NOTE: THE MATERIAL OF THIS ANNEX IS TAKEN FROM WD 10 OF INTERIM MEETING OF Q.3/13 (NOVEMBER 2003) AND IT WAS AGREED TO INCLUDE IT IN THIS ANNEX FOR FURTHER CONSIDERATION

AIS/RDI MECHANISM FOR AN ETHERNET POINT-TO-POINT CONNECTION OVER A SINGLE SERVER LAYER (i.e. SDH or OTN)

G.8010 [4] Figure 18 (reproduced below) illustrates the architecture of an Ethernet point-to-point connection.

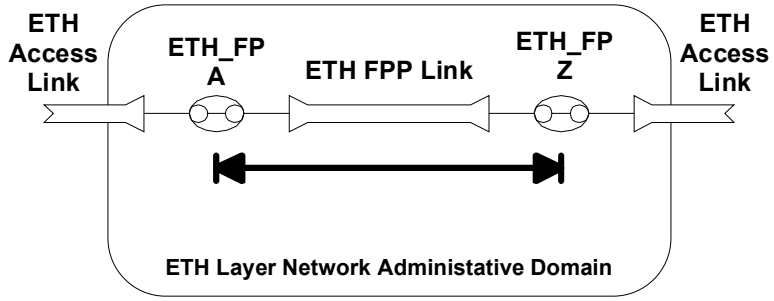

Figure 18/G.8010 – Point-to-point ETH connection (single link)

The representation of the corresponding maintenance areas is illustrated in G.8010 Figure 23 top right (reproduced below).

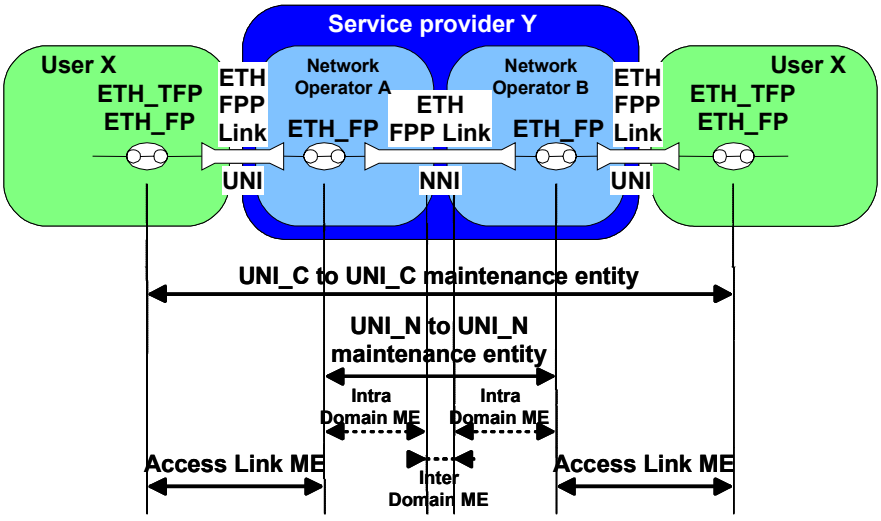

Figure 23 (top right)/G.8010 – Point-to-point ETH connection administrative domain associated MEs

Of note is that there are no ETH flow points at the handoff between the two networks. So for the general case when there are multiple network operators and a single server layer, maintenance of the ETH layer is not possible within those operators' networks, and can only be performed via the server layer. That category of maintenance is called inherent monitoring, also discussed in [4].

To address the lack of AIS and network RDI functionality in EFM OAM, the issues are then:

a) How to convey an access link fault from one side of the network to the other.

b) How to convey a server layer fault to the access links.

In SG 15/Q.12, work has been progressing to define a Service Management Channel (SMC) to facilitate the provider edge NE-to-edge NE exchange of OAM information, as well as support for an intermediate provider NE to query OAM information from, and send test-related commands to, a provider edge NE.

Currently, the direction being taken in Q.12/15 proposes G.7041 [5] GFP-F Client Management Frames (CMF) for conveying the provider edge NE-to-edge NE OAM information, and a Path OH byte for the intermediate provider NE communications channel to a provider edge NE. G.7041 defines CMFs for conveying information about the client signal from an ingress edge NE to the egress edge NE. One of the defined CMF indications is Client Signal Fail (CSF). The figure below illustrates the GFP-F frame format for a CMF with a CSF indication.

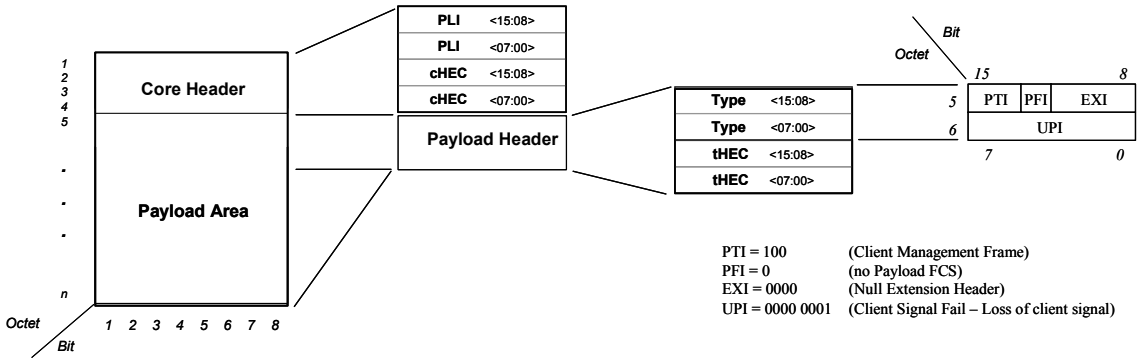

Figure A-1: GFP-F CMF CSF Frame Format

By using the EFM OAM Link Fault flag in conjunction with the GFP-F CMF CSF indication, the necessary AIS and network RDI mechanisms can be provided for an Ethernet point-to-point connection single server application.

A simplifying assumption can be made regarding the conditioning of the Ethernet access links on either side of the SDH/OTN transport network. For a dedicated point-to-point application, the access link is specific to a single service, and since an Ethernet service is bidirectional, a fault in either direction should result in the access link being conditioned as ‘failed'.

The following fault scenarios and accompanying figures illustrate the proposed interworking of the EFM OAM Link Fault flag with the GFP-F CMF CSF indication to appropriately condition the Ethernet access links. Only uni-directional faults are considered, the scenarios can be combined per the superposition principle to describe bi-directional faults.

Scenario 1

In the figure below a uni-directional fault occurs on the east access link on ingress to the carrier network.

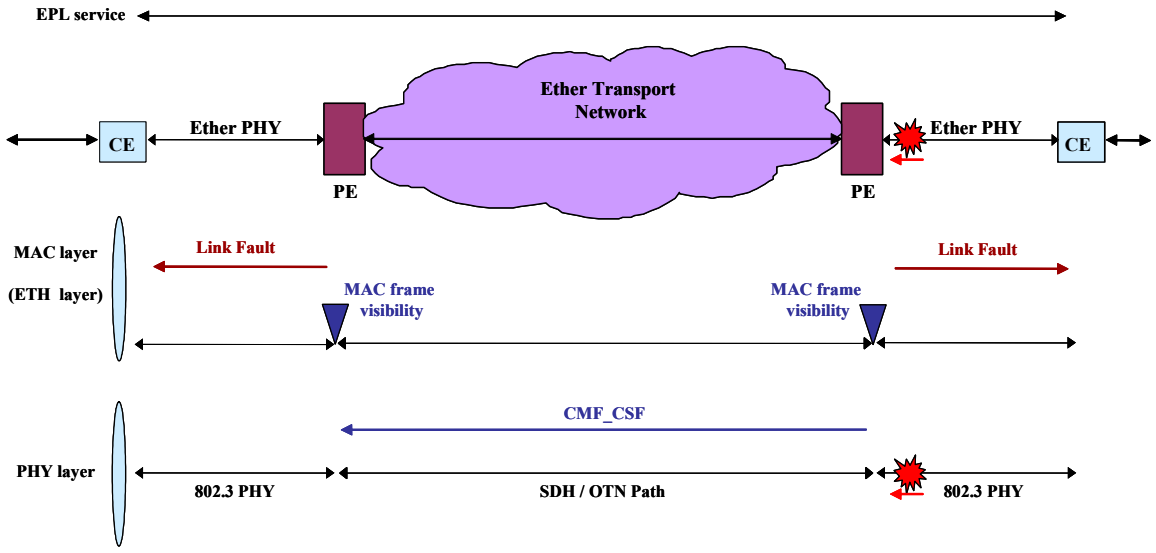

Figure A-2: Fault on Ingress

- The east PE detects the failure:
  - an 802.3ah OAM function sends Link Fault upstream, interspersed with Idles
  - a new function sends a GFP-F CMF CSF indication into the network
- The east CE detects Link Fault:
  - Idles are sent towards the network (and towards the enterprise)
- The west PE detects the GFP-F CMF CSF indication:
  - a new function translates it to 802.3ah Link Fault downstream, interspersed with Idles
- The west CE detects Link Fault:
  - Idles are sent towards the network (and towards the enterprise)

Scenario 2

In the figure below a uni-directional fault occurs westbound within the carrier network.

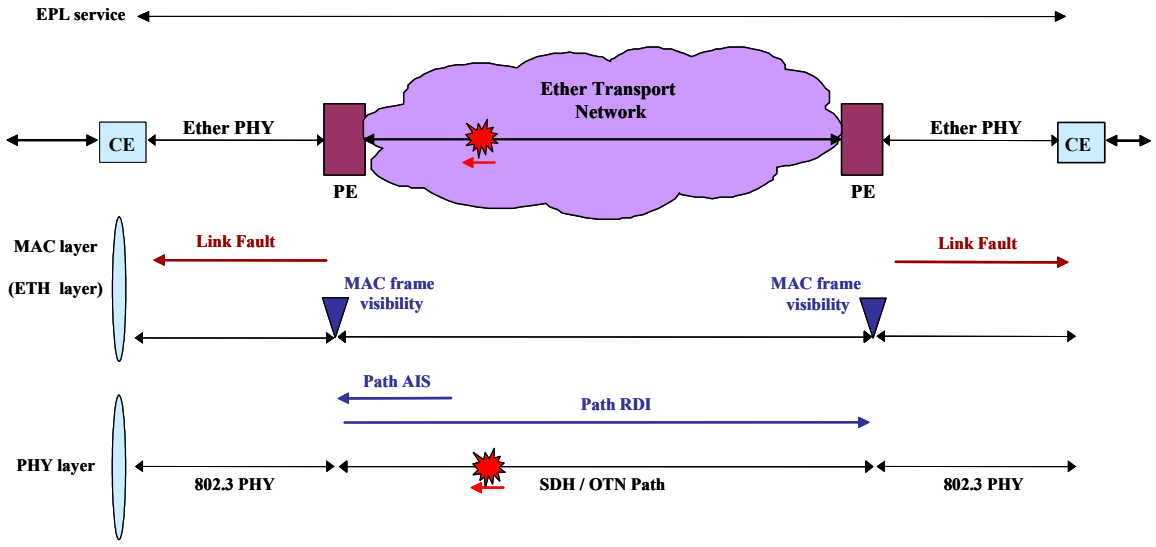

Figure A-3: Fault within Carrier Network

- An NE (or the west PE) in the carrier network detects the failure:
  - SDH Path AIS is generated downstream
- The west PE detects SDH Path AIS (or the fault directly):
  - a new function translates it to 802.3ah Link Fault downstream, interspersed with Idles
  - SDH Path RDI is generated back into the network
- The west CE detects Link Fault:
  - Idles are sent towards the network (and towards the enterprise)
- The east PE detects SDH Path RDI:
  - a new function translates it to 802.3ah Link Fault downstream, interspersed with Idles
- The east CE detects Link Fault:
  - Idles are sent towards the network (and towards the enterprise)

Scenario 3

In the figure below a uni-directional fault occurs on the west access link towards the enterprise network.

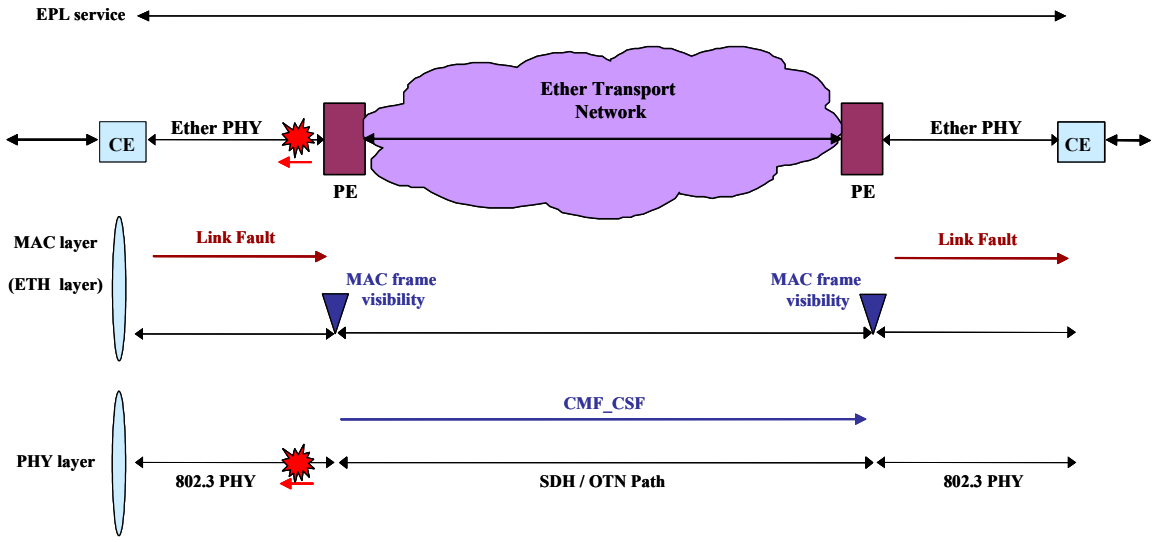

Figure A-4: Fault on Egress

- The west CE detects the failure:
  - an 802.3ah OAM function sends Link Fault upstream, interspersed with Idles
  - Idles are sent towards the enterprise
- The west PE detects Link Fault:
  - a new function translates it to a GFP-F CMF CSF indication into the network
  - Idles are sent towards the CE
- The east PE detects the GFP-F CMF CSF indication:
  - a new function translates it to 802.3ah Link Fault downstream, interspersed with Idles
- The east CE detects Link Fault:
  - Idles are sent towards the network (and towards the enterprise)

Summary

As a result, two maintenance signal translation functions, i.e. EFM OAM Link Fault flag and GFP-F CMF CSF indication can be used by the provider edge of a point-to-point single server application (i.e. SDH or OTN) in order to address network AIS and RDI mechanisms.

[EDITOR'S NOTE-Dec2004] Specific content in Annex A applicable for this Recommendation needs to be identified. Contributions are invited.

Annex B: Ethernet Network Scenarios

B.1 ME, MEP, MIP, and TCP Examples

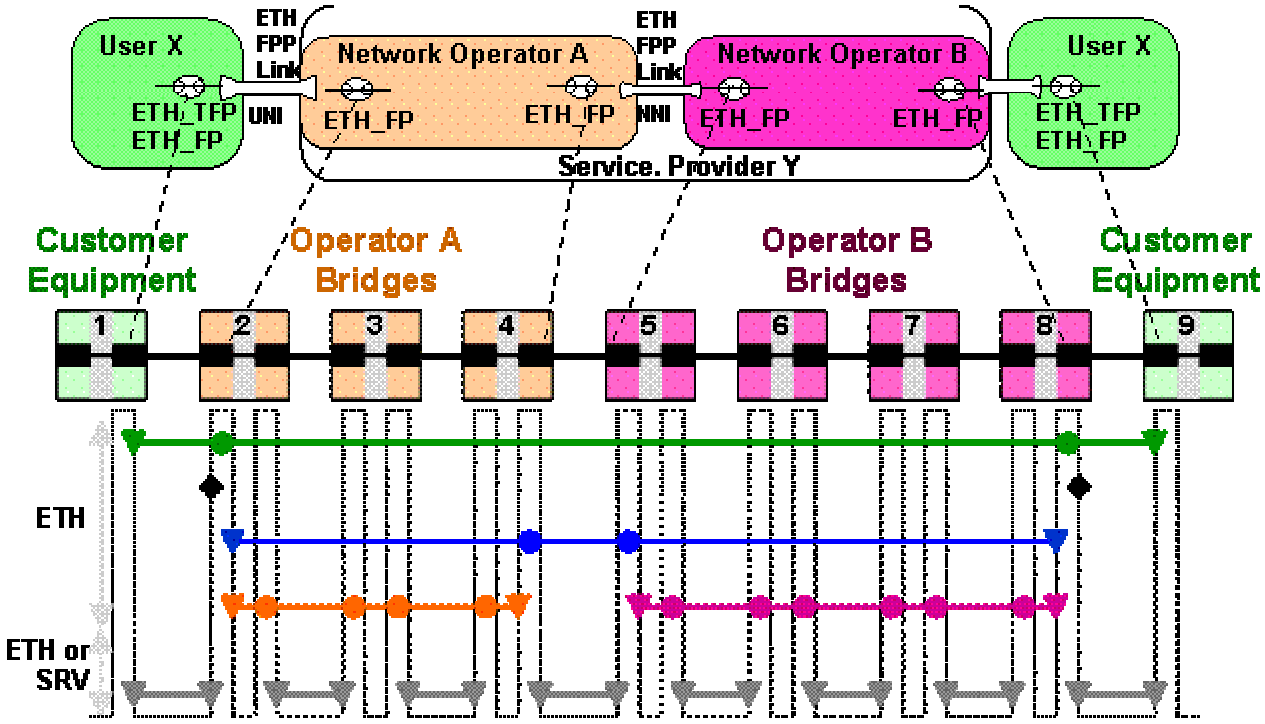

Figure B-1 – Example of ETH MEs with MEPs, MIPs and TCP

- p2p ETH connection between customer equipment 1 and 9 supported by a service provider and two network operators A and B
- green indicates a UNI-C to UNI-C ETH ME (customer) with MEPs in the interface ports facing the network in CE1 and CE9 and MIPs in the network interface ports facing the CEs (B2 and B8)
- blue indicates a UNI-N to UNI-N ETH ME (service provider) with MEPs at the edge of the network (B2,B8) and MIPs at the boundary of the two network operator domains (B4,B5)
- orange and mangenta indicate UNI-N to NNI ETH MEs (network operator) with MEPs at the edge of the operator networks (B2,B4 and B5,B8) and MIPs at each of the other interface ports
- brown indicates ETH link related MEs either realised as ETH MEs (sublayer monitoring) or as server layer MEs (inherent monitoring)
- black indicates location of unidirectional ETH TCPs; left TCP for direction CE1 to CE9 and right TCP for direction CE9 to CE1

NOTE: The black TCPs should be moved to the bottom of the figure if link is sublayer (ETH ME) monitored

Figure B-1-1 – Illustrating the order of MEPs, MIPs and TCP

- same p2p ETH connection as in previous figure, now represented in linear order; ETH_CI traffic units will pass through these MEPs, MIPs, TCPs in the presented order
- note: if link ME is an ETH ME (sublayer monitoring), then the TCPs must flip position with brown MEPs

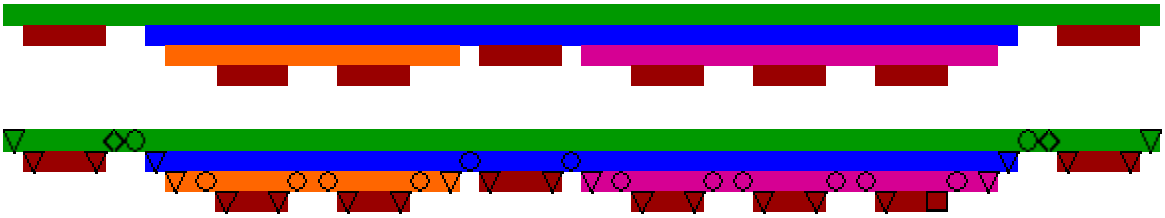

Figure B-1-2 – Illustrating the stacking of ETH MEs or ETH and SRV MEs

- same p2p ETH connection as in previous two figures, now represented as a stack

B.2 ME, MEP, MIP, and TCP in Dual Relay Model: P2P Connection

B.2.1 Dual Relay Model as Single Integrated Provider Device

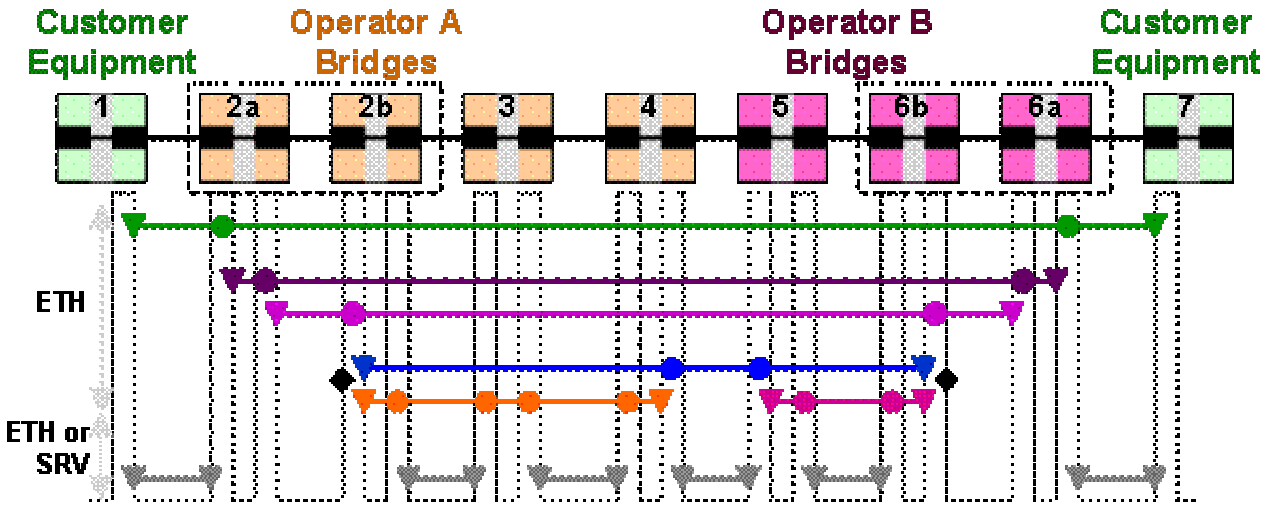

Figure B-2-1: MEs, MEPs, MIPs and TCPs in Dual-relay Provider Devices: one p2p connection

- Provider Device is represented as a dual relay model implemented with both relays. The first relay allows peering of customer L2CP protocols + multiplexing of multiple customer flows onto a single access link between the customer equipment 1 and provider bridge 2 (shown here as 2a and 2b).

- Due to the dual relay model, additional ME are introduced shown here in purple and pink between 2a and 6a. The Purple ME is associated with per customer VLAN at the provider equipment. The Pink ME is associated with per service instance (Service VLAN) that the provider applies to customer service frames. It may be noted that the additional MEs between 2a and 6a for per customer VLAN may be used for diagnostic purposes only since per Service VLAN MEs are sufficient for provider domain.
- Between the dual relays, there are pseudo interfaces which correspond 1-to-1 with the Service VLAN or Provider Tag, which is expected to be inserted at second relay e.g. 2b.
- FFS: The positioning of the TCPs is for further study since TCPs can be positioned at customer access link level, per customer VLAN level and per provider VLAN (aka service) level.

B.2.2 Dual Relay Model with Single Relay as Provider Device

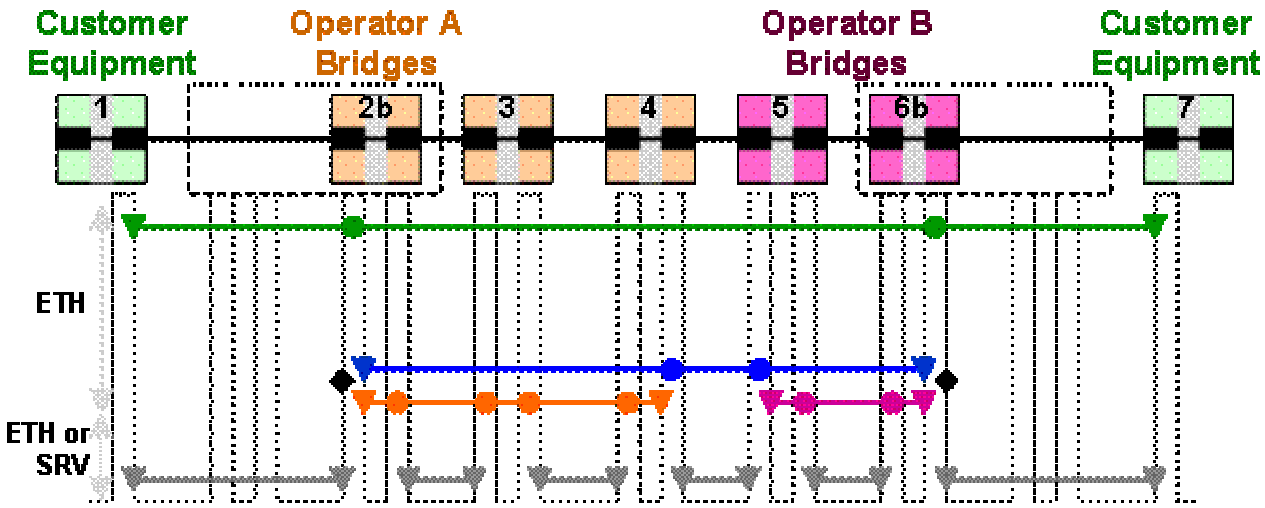

Figure B-2-2: MEs, MEPs, MIPs and TCPs in Dual-relay Provider Devices modelled as a single relay: one p2p connection

- Provider Device is represented as a dual relay model implemented with single relay (shown as 2b). In this case, the second relay does not allow peering of customer L2CP protocols and requires a single link for every service it supports across the customer device 1.
- Customer device 1 is responsible for multiplexing multiple customer flows onto a single service link.
- Due to the provider using a single relay of the dual relay model, the additional MEs that were introduced in Figure 6-5, are expected to be present at the customer device and are not shown here since the customer is expected to manage those. Customer's relationship with the Service Provider is limited to a single service instance ME shown here by the Green ME between 1 and 7.
- The positioning of the TCP is clearer in this case and is shown in above figure. Customer is responsible for per customer VLAN level conditioning

B.3 ME, MEP, MIP, and TCP in Dual Relay Model: Bundling

B.3.1 Dual Relay Model as Single Integrated Provider Device

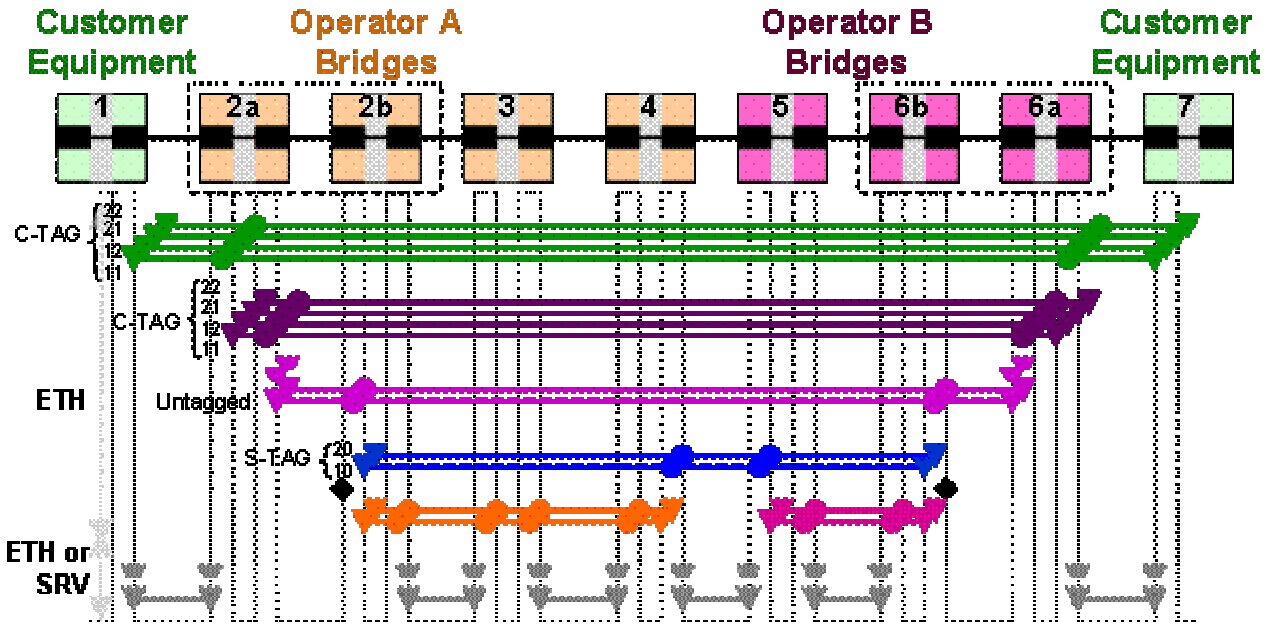

Figure B-3-1: MEs, MEPs, MIPs and TCPs in Dual-relay Provider Devices: two p2p connections with bundling

- Customer is shown using 4 customer VLANs (11, 12, 21, 22). It is also indicated that the customer signs up for two p2p connection services which the provider carries across the provider network using two provider VLANs (10 and 20). It is assumed that 2 customer VLANs (11 and 12) are mapped to provider VLAN 10 and the other two customer VLANs (21 and 22) are mapped to provider VLAN 20.
- Additional MEs introduced between 2a and 6a are replicated per customer VLAN and provider VLAN. It may be noted that the additional MEs between 2a and 6a for per customer VLAN may be used for diagnostic purposes only since per Service VLAN MEs are sufficient for provider domain.
- Additionally the trapezoid entities shown in Figure 6-7 represent the “AIS adaptation” associated with the MEPs where a server level ME multiplexes one or more client level ME.
- MEs corresponding to the dual bridge pseudo interfaces which correspond 1-to-1 with the provider VLANs (10 and 20) are shown as untagged since frames from the first relay (e.g. 2a) are expected to have no provider tag as they arrive at the second relay (e.g. 2b).
- FFS: The positioning of the TCPs is for further study since TCPs can be positioned at customer access link level, per customer VLAN level and per provider VLAN (aka service) level.

B.3.2 Dual Relay Model with Single Relay as Provider Device

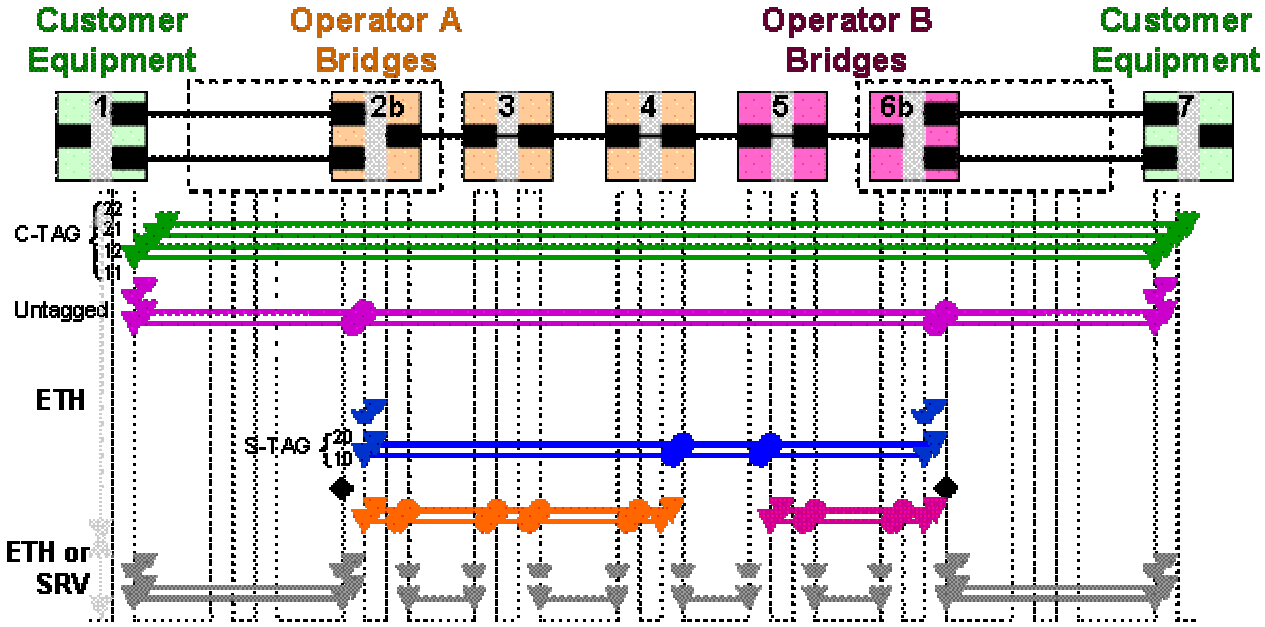

Figure B-3-2: MEs, MEPs, MIPs and TCPs in Dual-relay Provider Devices modelled as single relay: two p2p connections with bundling

- Due to the provider using a single relay of the dual relay model, bundling is realized across the customer device 1 and 7. Two links connect devices 1 and 2b and devices 6b and 7, where each link corresponds to a customer service instance.
- Additional ME is introduced at customer devices to highlight the responsibility of the customer for ME corresponding to per customer VLAN (shown here by 4 different green MEs between customer devices 1 and 7 for customer VLANs 11, 12, 21, and 22) and per service (shown here by 2 different purple MEs between customer devices 1 and 7) . It may be noted that the additional MEs between 1 and 7 for per customer VLAN may be used for diagnostic purposes only since per Service untagged MEs are sufficient for customer domain.
- Additionally the trapezoid entities shown in Figure 6-7 represent the “AIS adaptation” associated with the MEPs where a server level ME multiplexes one or more client level ME.
- The positioning of the TCP is clearer in this case and is shown in above figure. Customer is responsible for per customer VLAN level conditioning.

B.4 ME, MEP, MIP, and TCP in Dual Relay Model: All-to-one Bundling

B.4.1 Dual Relay Model with Single Relay as Provider Device

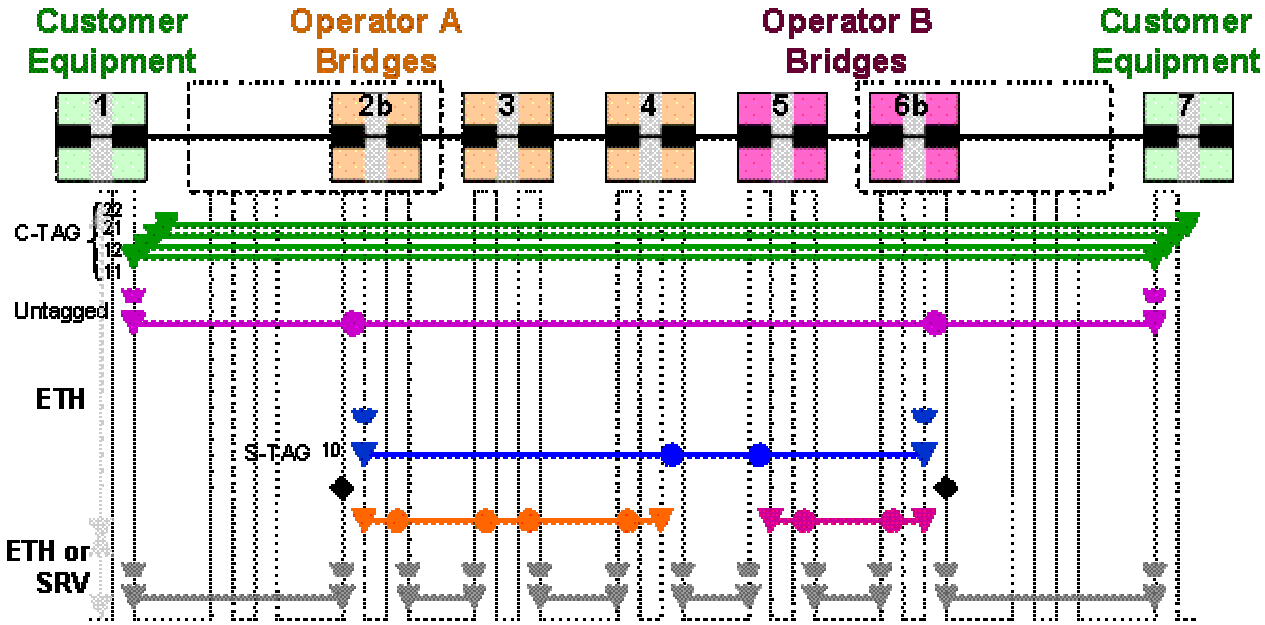

Figure B-4-1: MEs, MEPs, MIPs and TCPs in Dual-relay Provider Devices modelled as single relay: one p2p connection with all-to-one bundling

B.5 ME, MEP, MIP, and TCP in Access Maintenance Scenarios

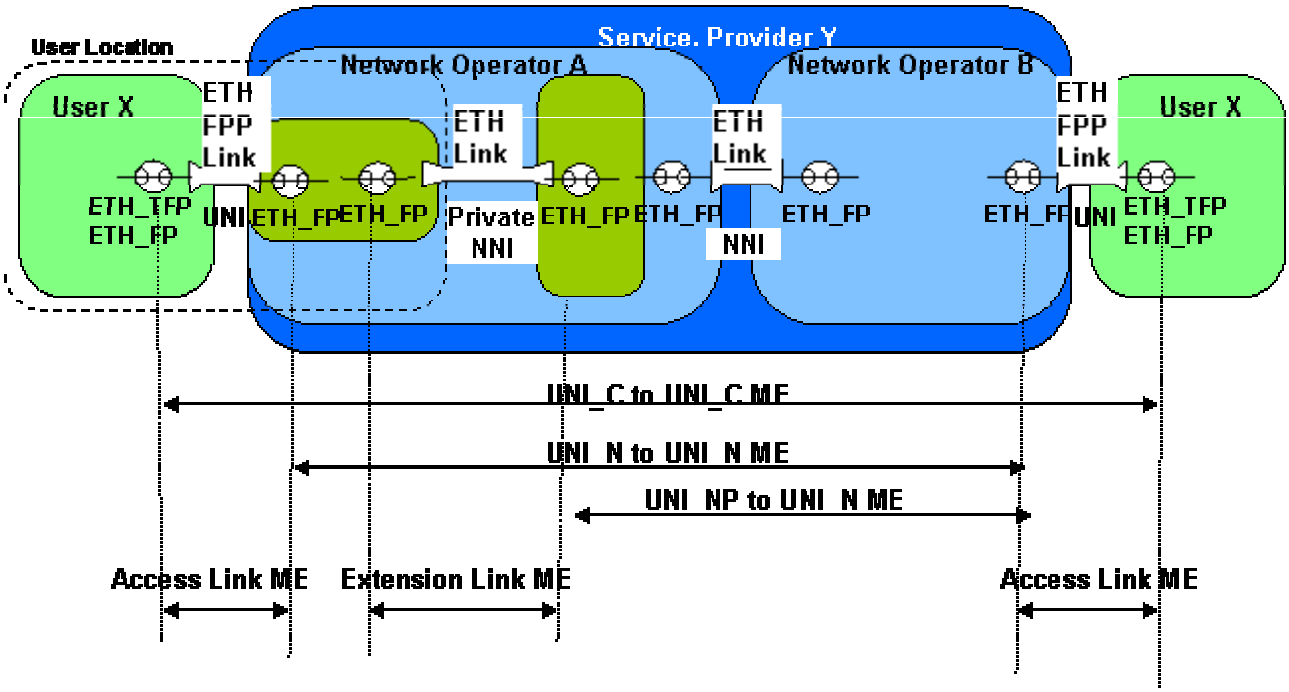

Figure B-5-1: MEs in Access Scenario with a Network Device between User and Provider

Deployment of a network device between the provider and a user introduces the Extension Link ME (for the Private NNI ETH link) and the UNI_NP to UNI_N ME, which together form a subset of the previously defined UNI_N to UNI_N ME.

Figure 6-11 identifies the associated MEPs and MIPs for this access scenario.

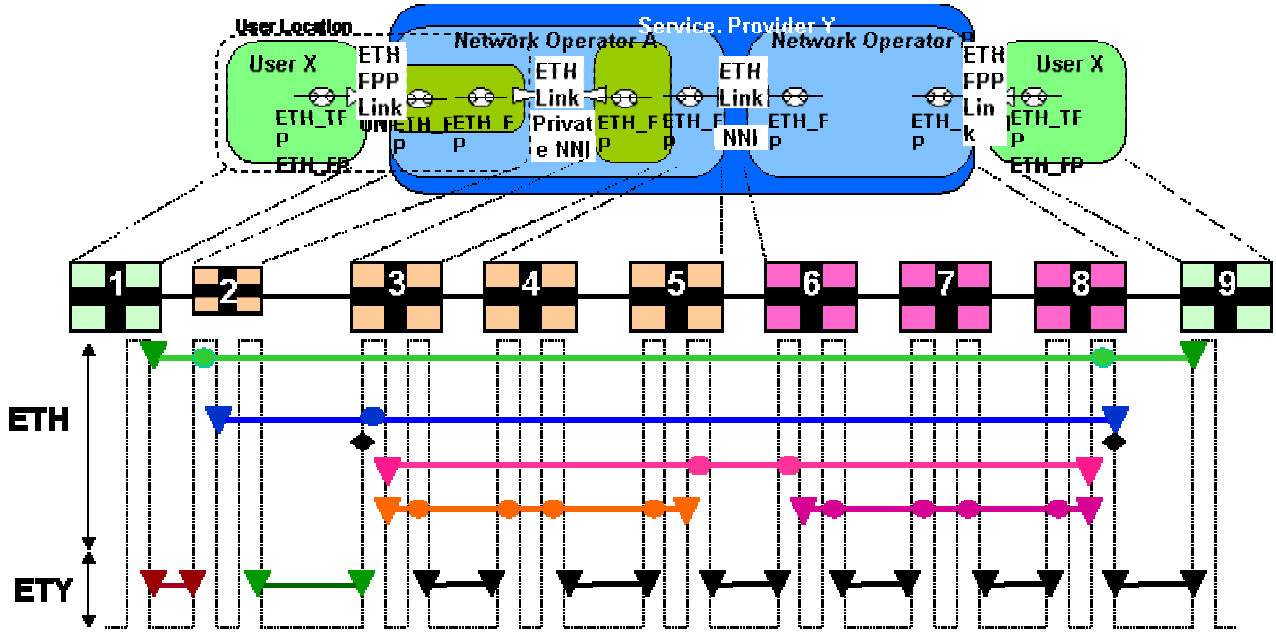

Figure B-5-2: MEs, MEPs, MIPs and TCPs in Access Scenario with a Network Device between User and Provider

- A point-to-point ETH connection between customer equipment 1 and 9 is supported by service provider Y, consisting of a customer premise-located device (2) connected to a metro transport network.
- Light green indicates a UNI-C to UNI-C ETH ME (customer) with MEPs in the interface ports facing the network in CE1 and CE9, and MIPs in the network interface ports facing the CEs (B2, B3 and B8).
- Dark green indicates at ETY layer indicates the Extension Link ETY ME. Brown represents an access link ME.
- Magenta indicates a UNI-NP to UNI-N ETH ME (metro network) with MEPs at the edge of the metro network (B3 and B8) and MIPs at each of the other interface ports in between.
- Blue indicates the UNI-N to UNI-N ETH ME.
- Alternatively the UNI-N to UNI-N ETH ME can be realized by combination of Extension Link ETY ME and UNI-NP to UNI-N ETH ME. In this case, blue ME does not need to be defined.
- Orange and pink indicate intra-domain ETH MEs with MEPs at the edge of network operator A and B respectively (B3-B5 and B6-B8 respectively) and MIPs at each of the other interface ports in between.
- Black indicates ETH link MEs either realized as ETH MEs (sublayer monitoring) or as server layer MEs (inherent monitoring).

EDITOR'S NOTE: SERVICE AWARENESS OF EACH ME MAY BE ADDED. IN THAT CASE, DEFINITION OF SERVICE AWARENESS IS NEEDED.

Other access scenarios are possible where the device 2 could simply be a media converter (MC) with single flow-point. Figure 6.12 represents the scenario where device 2 is a Media Converter

(MC) device where the Network Termination (NT) functionality is present in MC and Line Termination (LT) functionality is present in edge of provider domain.

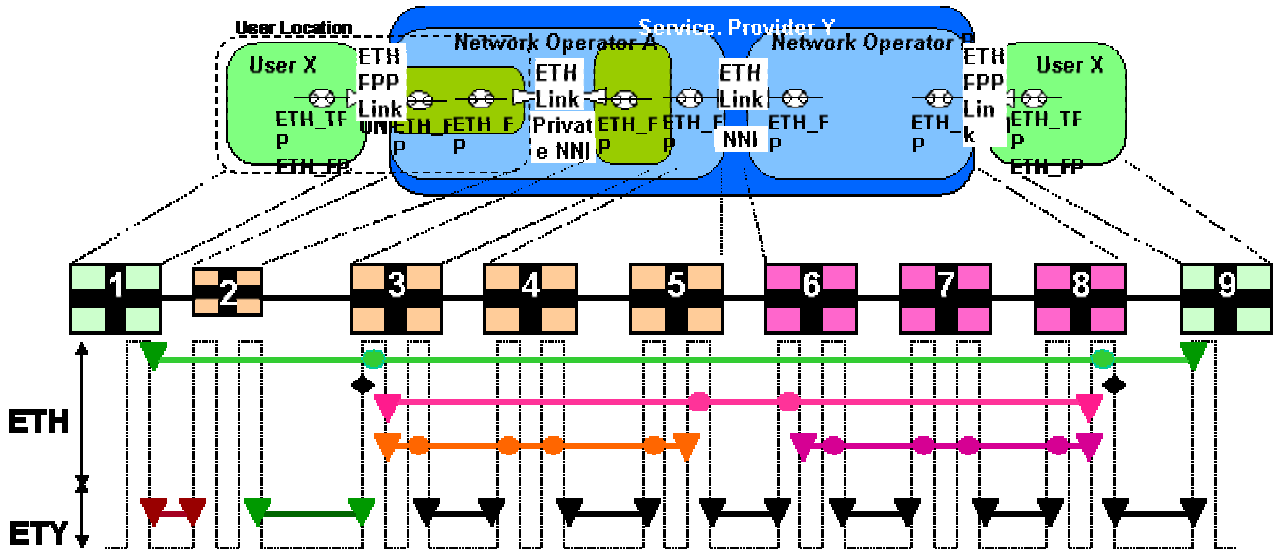

Figure B-5-3: MEs, MEPs, MIPs and TCPs in Access Scenario with a NT Network Device between User and Provider

- A point-to-point ETH connection between customer equipment 1 and 9 is supported by service provider Y. Device 2 here functions as a Media Converter (MC) where the MC realized a Network Termination (NT) device while the Line Termination (LT) functionality is integrated in the edge of network operator (B3).
- Light green indicates a UNI-C to UNI-C ETH ME (customer) with MEPs in the interface ports facing the network in CE1 and CE9, and MIPs in the network interface ports facing the CEs (B3 and B8).
- Dark green indicates at ETY layer indicates the Extension Link ETY ME. Brown represents an access link ME. This scenario requires some stitching between access link ME and Extension link ETY ME.
- Magenta indicates a UNI-NP to UNI-N ETH ME (metro network) with MEPs at the edge of the metro network (B3 and B8) and MIPs at each of the other interface ports in between.
- The UNI-N to UNI-N ETH ME (not shown in the figure) is realized by combination of Extension Link ETY ME and UNI-NP to UNI-N ETH ME.
- Orange and pink indicate intra-domain ETH MEs with MEPs at the edge of network operator A and B respectively (B3-B5 and B6-B8 respectively) and MIPs at each of the other interface ports in between.
- Black indicates ETH link MEs either realized as ETH MEs (sublayer monitoring) or as server layer MEs (inherent monitoring).

[EDITOR'S NOTE-Dec2004] More scenarios to be added e.g. EPL case, Access Networks e.g. Internet Access with 2 VLANs and corresponding MEs. Some of these are expected to come from wd03.

Annex C: OAM Operational Scenarios

[EDITOR'S NOTE-Dec2004] NMS related text from D-53 to be added in this Annex

C.1 Provisioning Example

1. Operator Start-up: The operator-A and Operator-B would each deploy equipment in their domains.
    a. Operator device with Ethernet functionality boots up
    b. By default, each port has a MEP at the ETY or SRV layer. This layer and MEPs are distinct from the Ethernet OAM. E.g. an ETY MEP could correspond to EFM OAM [2].
    c. By default, at ETH layer, each port has a MIP associated with the lowest PHY-ward ME Level (currently 0)
    d. Operator defines its administrative boundary by identifying and configuring ETH layer.
    e. For each configured port, an ETH MEP associated with one of the operator ME Level is created. However, the proactive OAM capabilities like CC may still not be desirable until the operator performs some start-up diagnostics. Towards this objective, the initial configured state of these configured ports can be "Administrative - Diagnostic state"
    f. After all operator devices are deployed, as mentioned in the above steps, the operator may want to run some start-up diagnostics e.g. Multicast ETH-LB to detect any misconnections, or Intrusive ETH-LB to validate connectivity parameters.
    g. After performing the start-up diagnostics, the state of the ETH ports can be restored to "Administratively – Up state"
2. Operator Connections: Subsequent to the Operator Start-up phase, each operator eventually sets up connections based on contracts between them and providers.
3. Provider Start-up: The provider can be either facility based or non-facility based.
    a. Non-facility based Providers may rely upon the Operators to set up ETH MEP at the Provider ME Level; the start-up diagnostics may be limited in this case.
    b. Facility based Providers do not rely upon the Operators to set up their ETH MEP; Rather the provider could follow the same start-up sequence as in 1 with limited start-up diagnostics as compared to Operators.
4. Provider Connections: Subsequent to the Provider Start-up phase, the provider eventually set up connections based on contracts between them and customers.

There are some start-up scenarios, as presented above in steps 1 and 3, which do not necessarily require pro-active OAM capabilities e.g. ETH-CC. However, once the start-up diagnostics are completed, the proactive OAM can be turned on which offers the complete set of OAM functions.

The figures below are example sequences of ME provisioning.

To simplify the sequences, there are several assumptions like below.

- Absolute Assignment Mechanism is used for the ME Level.
    - Assignment/Allocation of each ME Level value is FFS.

- Based on the basic model of ME, there are 1 customer, 1 service provider and 2 network operators.
- All bridges used by the provider to provide the service to the customer can be directly managed by the operator. Therefore provisioning operations are permitted only to an operator of the network operator organization.
  - ➢ The case where the provider uses its own bridges is FFS.
- ServiceIDs and MPIDs are FFS.

EDITOR'S NOTE: THE PROVISIONING SEQUENCES BELOW SHOULD BE DISCUSSED FURTHER MORE CONSIDERING WITH THE PORT/MEP/MIP STATES. CONTRIBUTIONS ARE INVITED.

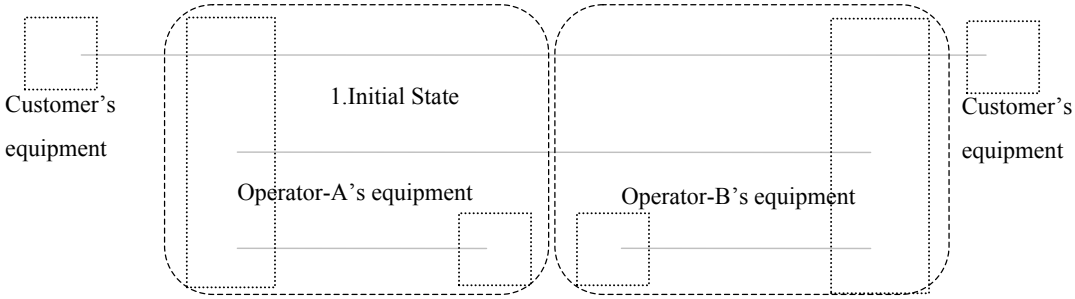

Figure C.1 Initial state

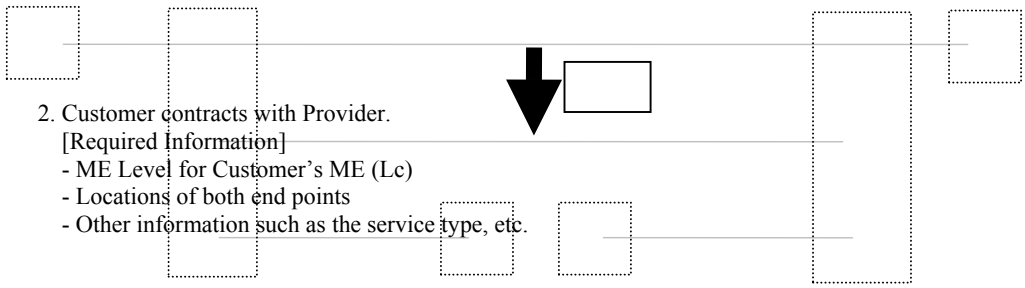

Figure C.2 Contract between Customer and Provider

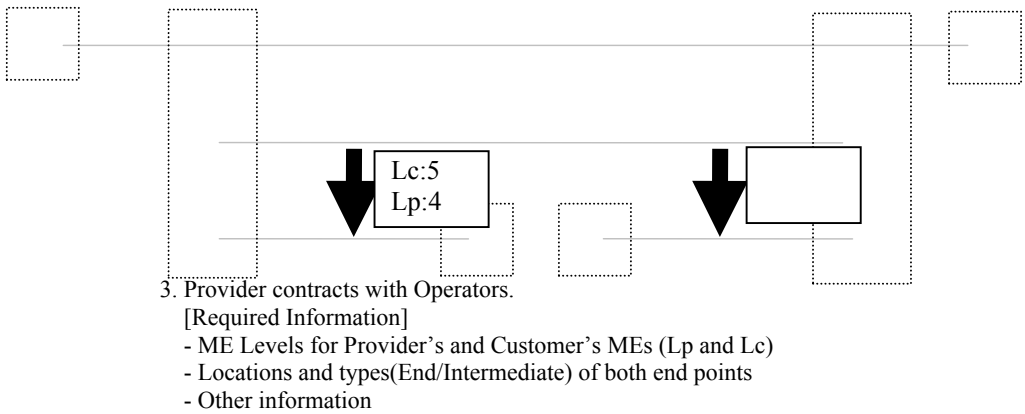

Figure C.3 Contract between Provider and Operators

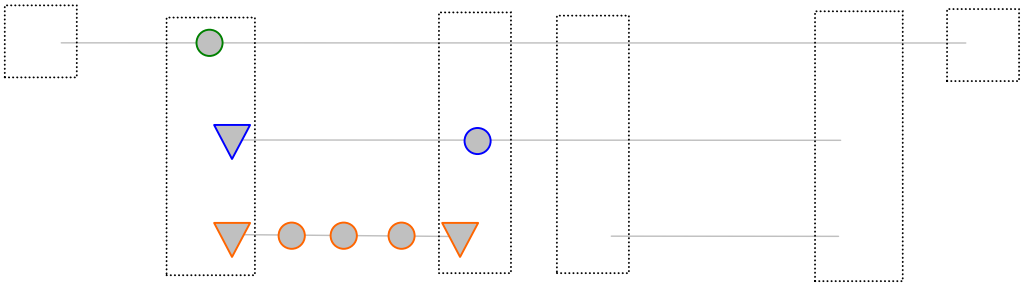

4a. Operator-A creates its own MPs, Provider's MPs and Customer's MIP on its equipments.

Figure C.4a Operator-A creates MPs on its equipments

(Discussion 1) In this example, all MPs are created by the operator. Is this feasible?

There are several options for this.

a) Only Operator can create MPs. (like this example)

b) Each operator in each level can create MPs for that level.

c) Any other case?

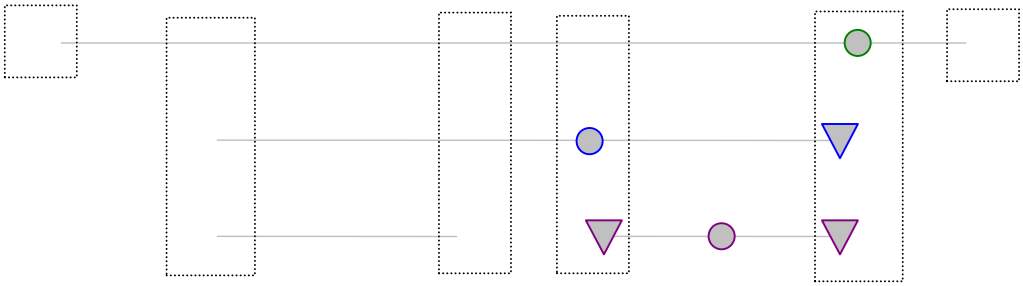

4b. Operator-B creates its MPs, Provider's MPs and Customer's MIP on its equipments.

Figure C.4b Operator-B creates MPs on its equipments

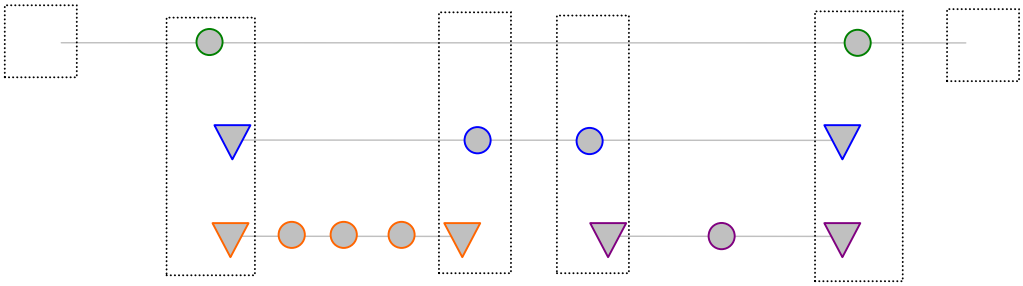

4. after 4a and 4b, all MPs except for the customer's MEPs are created.

Figure C.4 All MPs are created except for the Customer's MEPs

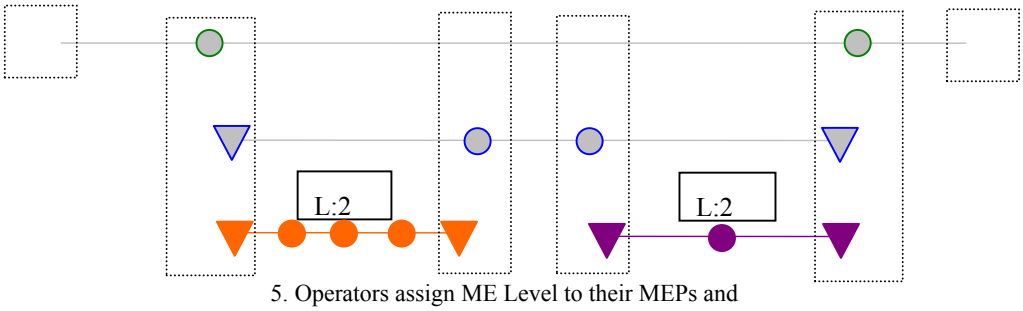

5. Operators assign ME Level to their MEPs and MIPs, then check the connectivity of each ME.

Figure C.5 Operators activate their MPs

(Discussion 2) In this example, all MPs are at first created and then activated. Therefore MEPs and MIPs must have states, such like “disabled”, “activated” and so on. Do we really have to define these states?

(Discussion 3) How to verify the connectivity of ME? For example, at first MEP to MEP LB will be done, then Link Trace will be used to ensure the route, and finally CC will be started, etc. It is better to clarify these basic sequences.

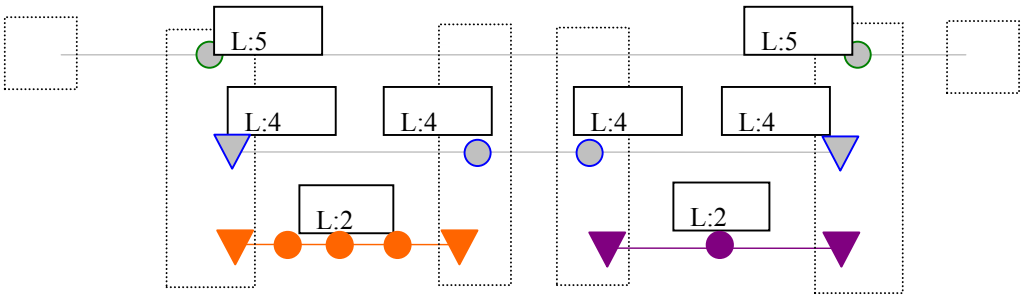

6. Operators assign ME Level to Provider’s MPs and Customer’s MIPs.

Figure C.6 ME Levels are assigned to Provider’s and Customer’s MPs

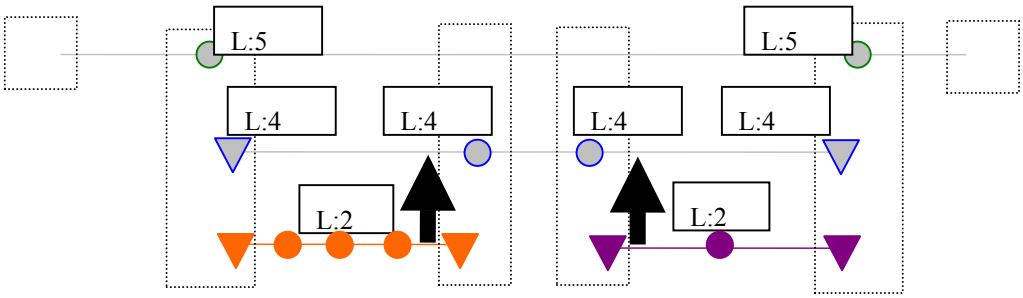

7. Operators inform the information about MPs to Provider.

Figure C.7 Provider receives the necessary information from Operators

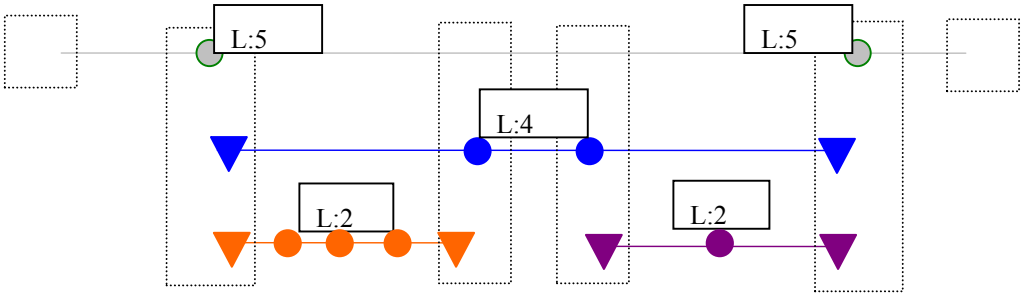

8. Provider verifies the connectivity of its ME.

Figure C.8 Provider activates its MPs

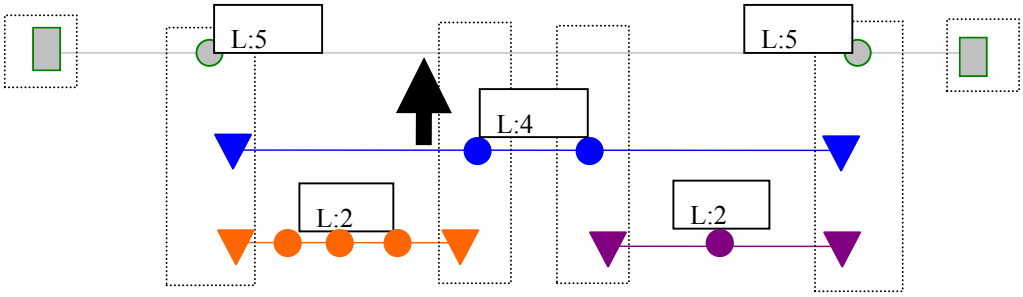

9. Customer creates its MEPs.

Figure C.9 Customer receives the necessary information and creates its MEPs

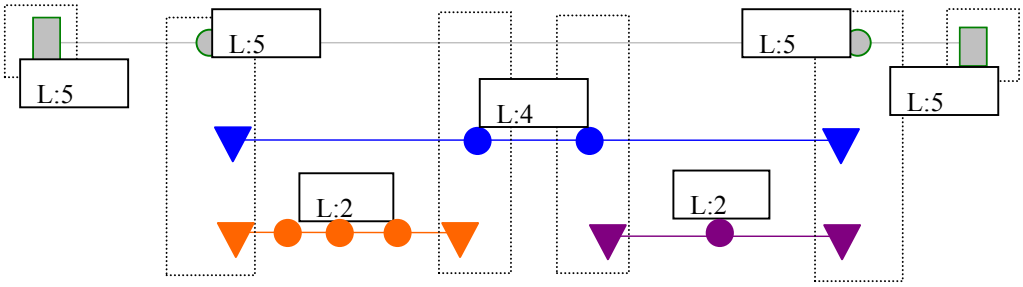

10. Customer assigns ME Levels to its MEPs

Figure C.10 Customer assigns ME Level to its MEPs

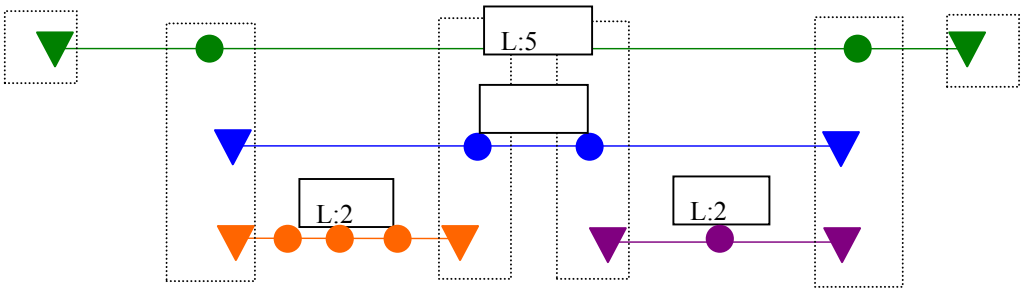

11. Customer verifies the connectivity of its ME.

Figure C.11 Customer activates their ME

C.2 Provisioning Example via Network Management System (NMS)

A number of network operators deploy Ethernet (layer) networks without the in-band Ethernet control plane. Instead those networks deploy network management systems to provide ETH connection management functionality. In this way loop-free p2p and mp ETH connections can be set up in a single step, of which the result can be compared with a "per VLAN spanning tree".

[EDITOR'S NOTE-Dec2004] Contributions highlighting differences between Residential and Business OAM deployments are invited. This is in context of Access Scenarios e.g. DSLAM where optimizing number of CC messages looks attractive but other assertions made indicating CC may not be used in such environments.

Appendix I: OAM Domains and OAM Flows

I.1 OAM Domains

Each provider can be associated with an administrative boundary, called OAM domain. A service may be carried across a single or multiple OAM domains.

As identified in Y.1730, network elements placed at the boundary of provider network serve as edge network elements and are associated with the ingress and egress of a network flow. When an edge network element of a provider performs hand-off of an ETH layer flow, while interacting with edge network element of another provider, that network element serves as an edge hand-off network element. Those network elements that are not associated with the ingress, egress or hand-off of a network flow serve as interior network elements.

It is also possible that a single provider network may have further administrative boundaries. Example is when a provider network consists of different operator networks. If this is the case, one could still identify edge, edge hand-off, and interior network elements within each such administrative boundary.

Ports on a network element in an OAM domain can be classified as interior or exterior to that OAM domain. Interior ports are those on which OAM frames, belonging to an OAM flow, are recognized and processed. Processing may result in either termination of OAM flow or relaying across other ports on the network element. Exterior ports are those on which OAM frames are not recognized and filtered. An edge network element has both interior and exterior ports to an OAM domain, while an interior network element has all its ports marked as interior ports to that OAM domain.

Within an OAM domain, OAM flows may be applicable between edge network elements only (edge hand-off network element is also an edge network element) or across all network elements (i.e. including all interior network element and edge network elements). OAM frames can be Unicast or Multicast frames. The difference between the two is based on the destination MAC address (DA). A Unicast OAM frame has a Unicast DA while a Multicast OAM frame has a Multicast DA. A Multicast OAM frame can associate itself to all edge networks elements or all network elements inside a domain based on its Multicast DA.

NOTE: Refer to G.8010 and G.805.

I.2 OAM Flows

Different OAM flows, as discussed in Section 6.1, can be identified by using an OAM flow identifier within the OAM frame. OAM flow identifier can assume the following values:

- $\text{UNI-UNI}_{\text{Customer}}$
  Customer UNI-UNI flow between reference points on the customer side of the UNI.
- $\text{UNI-UNI}_{\text{Provider}}$
  Provider UNI-UNI flow between reference points on the provider side of the UNI.
- $\text{Segment}_{\text{intra-provider}}$
  Segment OAM flow between flow points within the boundary of a provider network. This may include OAM flow between flow points on the boundary of a provider network or between any flow points within a provider network as required.
- $\text{Segment}_{\text{inter-provider}}$
  Segment OAM flow between flow points inside the boundaries of two or more provider networks. This may include OAM flow between flow points on the boundaries of two or more adjacent provider networks or between any flow points inside the boundaries of two or more provider networks, as required. Note: Under special cases, $\text{Segment}_{\text{inter-provider}}$ may be same as $\text{UNI-UNI}_{\text{Provider}}$.

- $UNI_{Segment}$
  OAM flow between reference points (i.e. TFP and FP) on the customer side and provider side of the UNI.
- $NNI_{Segment}$
  OAM flow between flow points on two edge hand-off network elements connected to each other. Each edge hand-off network element belongs to a different provider network.
- UNI_{Link}
  If the UNI is realized using a single ETY link, this OAM flow can be used for ETY link between customer and provider network.
- $Transit_{Link}$
  Any intermediate ETY link between network elements, this OAM flow can be used.

NOTE: Both UNI_{Link} and $Transit_{Link}$ can be based on IEEE 802.3ah. [However, the reference to IEEE 802.3ah may not be possible, until it becomes a standard, though it is close to being one]

NOTE: It is worth noting that though different OAM flows have been identified, not all will be applicable for all services and/or business models; especially, there may be some limitations within multiple provider scenarios.

An example of MEs and ME Levels can be seen in table I.1

Table I.1: Relationship between OAM flows and MEs

Y.1730	**G.8010**		**Examples**
ME	**ME**	**OAM flows**	**ME**
UNI-UNI (Customer)	UNI_C to UNI-C ME	UNI-UNI Flow	UNI-UNI (Customer)
UNI-UNI (provider)	UNI_N to UNI_N ME	Transit Flow	UNI-UNI (provider)
Segment (PE-PE) intra-provider	Intra Domain ME	Transit Flow	Segment (PE-PE) intra-provider
Segment (PE-PE) inter-provider	Inter Domain ME	Transit Flow Transit Link Flow	Segment (PE-PE) inter-provider
Segment (any to any)		Transit Flow Transit Link Flow	Segment (any to any)
ETY Link OAM - UNI	Access Link ME	UNI Link Flow	ETY Link OAM - UNI
ETY Link OAM - NNI	Inter Domain ME	Transit Link Flow	ETY Link OAM - NNI

Since the OAM Flows have a one-to-one correspondence with the MEs, the ME Levels can be represented by octet values assigned to OAM Flow identifiers as follows:

(It is conceivable that value of OAM Flow Identifiers can be such that filtering can be done based on whether the OAM frame entering or exiting a domain have OAM Flow Identifier value smaller than minimum OAM Flow Identifier configured on the interior and or exterior port of a domain.)

- $UNI\text{-}UNI_{Customer}$ = 255 (0xFF)
- $UNI\text{-}UNI_{Provider}$ = 253 (0xFD)
- $Segment_{inter\text{-}provider}$ = 251 (0xFB)
- $NNI_{Segment}$ = 249 (0xF9)
- $UNI_{Segment}$ = 247 (0xF7)
- $Segment_{intra\text{-}provider}$ = 245 (0xF5)

- UNI_{Link} =
- $Transit_{Link}$ =

And, if the following minimum OAM flow Identifier values are configured across different ports:

- NNI port = 249 (0xF9)
- UNI port = 247 (0xF7)
- Interior port = 245 (0xF5)

Filtering at edge network elements can be achieved such that OAM frames with OAM Flow identifier smaller than minimum OAM Flow identifier are not allowed into or out of OAM domain.

I.3 Fault Types

Two fault types are recognized in relationship with ETH OAM:

1) ETH Discontinuity
2) ETH Misconnection
3) ETH Link Faults

The two first fault types are shown in the following Figure I.3-1:

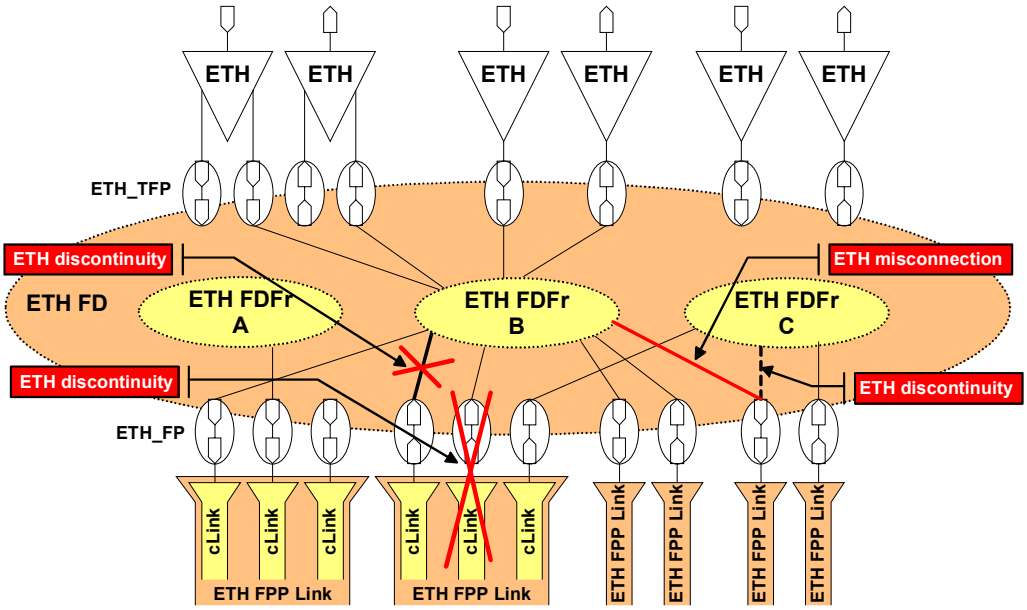

Figure I.3-1

ETH Discontinuity

Causes can be:

– Physical fault (i.e. fibre cut)
– Failure of a bridge
– Looping (customer or provider loops or due to the use of a wrong topology)
– Misconfiguration

ETH Misconnection

Can be caused by:

– Misconfiguration

Appendix II: MEPs and MIPs mapped to IEEE Constructs

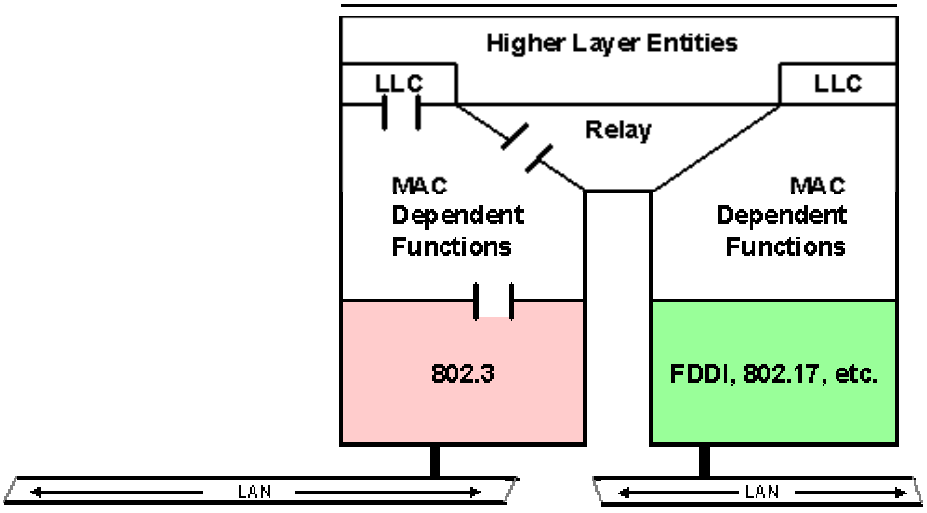

Figure II-1 The "Baggy Pants" diagram: IEEE Std. 802.1D-2003, Fig. 7-3

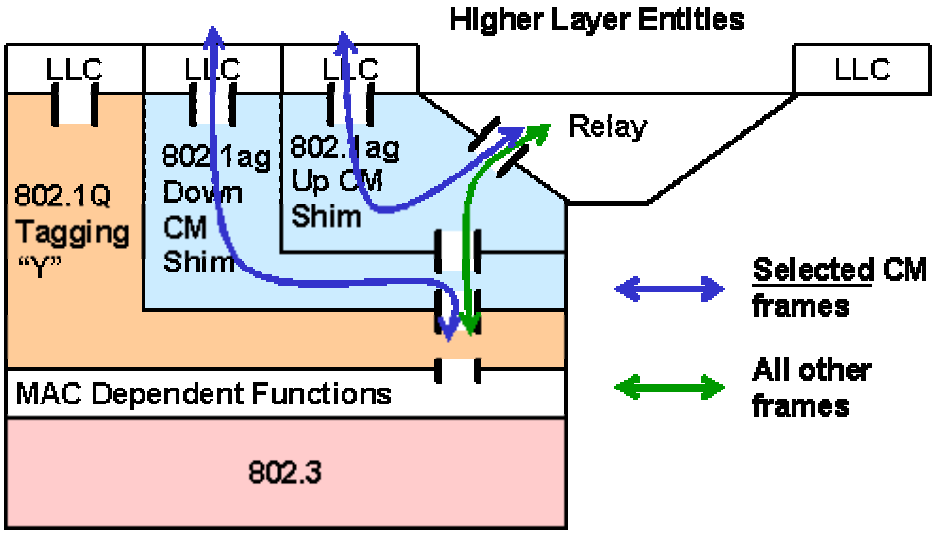

Figure II-2: The MAC stack: 802.1ag Connectivity Management Shim

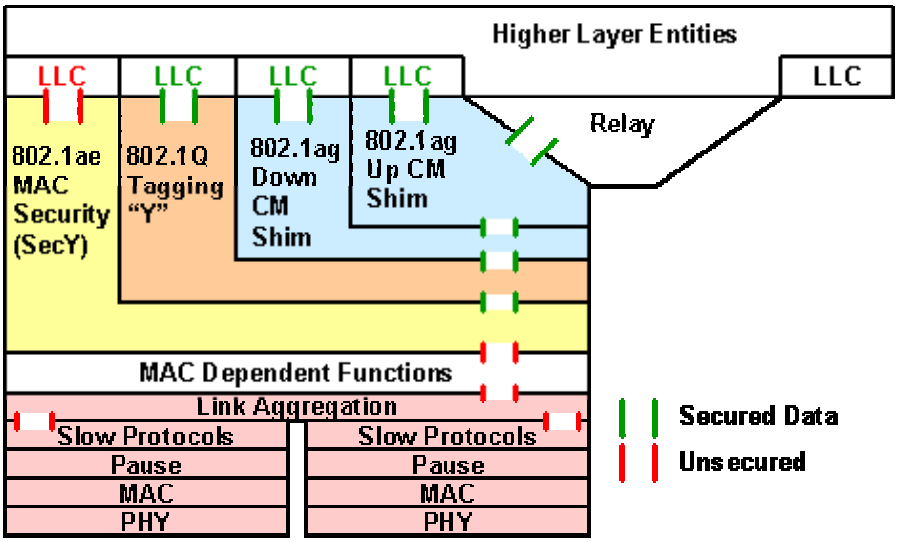

Figure II-3: The MAC stack: CM + Security + 802.3

The examples below relate to Annex B of this Recommendation

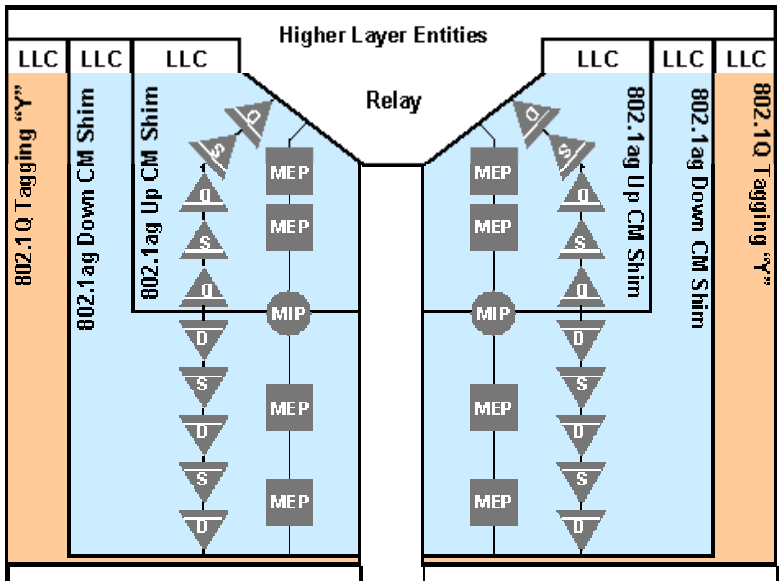

Figure II-4 – Illustrating the location of MEPs and MIPs in the 802.1ag Up and Down CM Shims in the IEEE 802.1 "baggy pants" model

- two representations are shown in parallel: the MEP and MIP shorthand symbolic representation and their associated atomic function representation
- both the up and down shims in an interface port to a bridge may support multiple MEPs; in this example each shim supports two MEPs
- the MIP is located at the junction of up and down shim as it has an ETH Diagnostic function in both the up and the down shim.
- the MEPs and MIP in an interface port have no pre-assigned knowledge of the ME Level they will be operating at; this is represented by making them all colourless (i.e. grey)
- further baggy pants figures will present only MEPs and MIPs, not longer the associated atomic functions

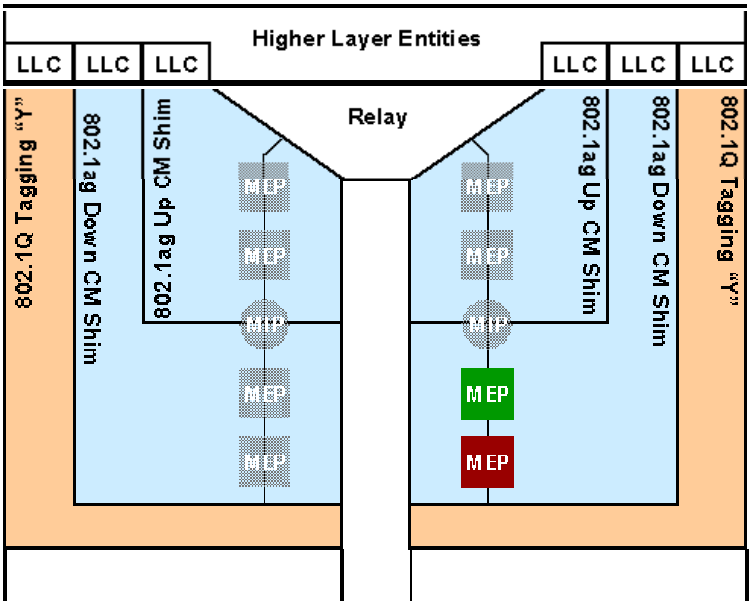

Figure II-5 – Illustrating the location of MEPs in the "baggy pants" model for CE1

- customer equipment number 1 has for the ETH connection of figure 6-2 two MEPs activated in the down shim to monitor the UNI-C to UNI-C connection and the CE1 to B2 link
- the other MEPs and MIPs are made transparent to represent that those are inactive for this connection

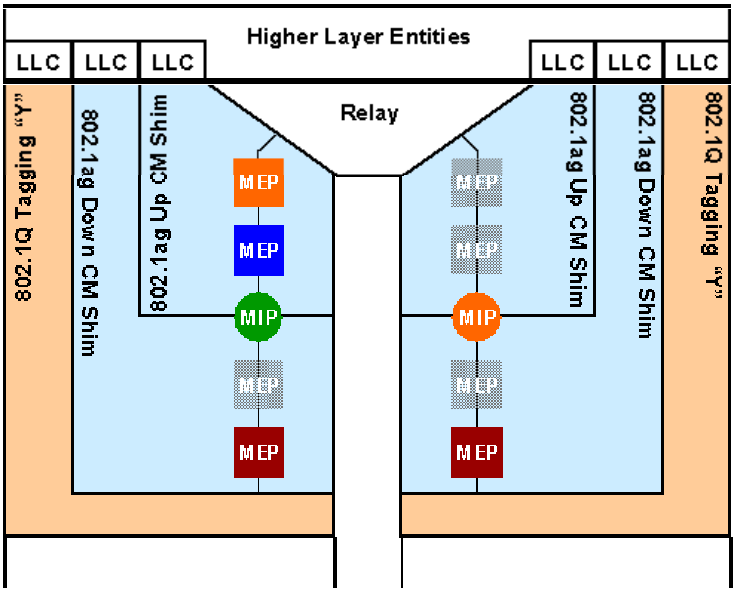

Figure II-6– Illustrating the location of MEPs in the "baggy pants" model for B2

- operator A bridge number 2 has at its customer equipment facing port three MEPs active; one to terminate the link ME, one to terminate the service provider's ETH ME (blue) and one to terminate the operator A's ETH ME
- this port has also an active MIP for use by the customer's UNI-C to UNI-C ETH ME
- the second interface port on this bridge has an active MIP (operator A's ETH ME) and an active MEP (B2 to B3 link)

- note that the choice of the active MEP in the two down shims is arbitrary; the other MEP could have been choosen as well. Equipment should be capable to change the MEP location hitless in order to support the addition of an ME Level above or below an existing ME Level

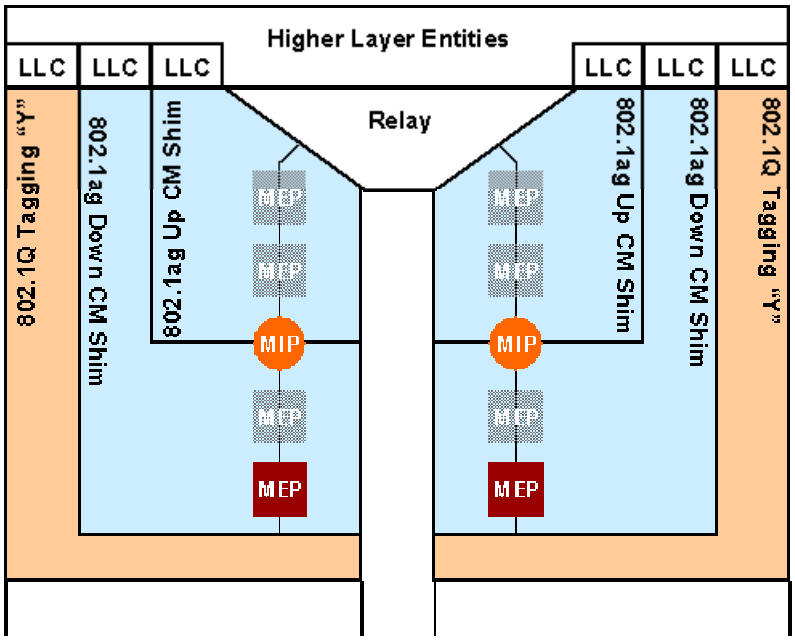

Figure II-7 – Illustrating the location of MEPs in the "baggy pants" model for B3

- operator A bridge number 3 has two ETH link related MEPs and two operator A ETH ME related MIPs active

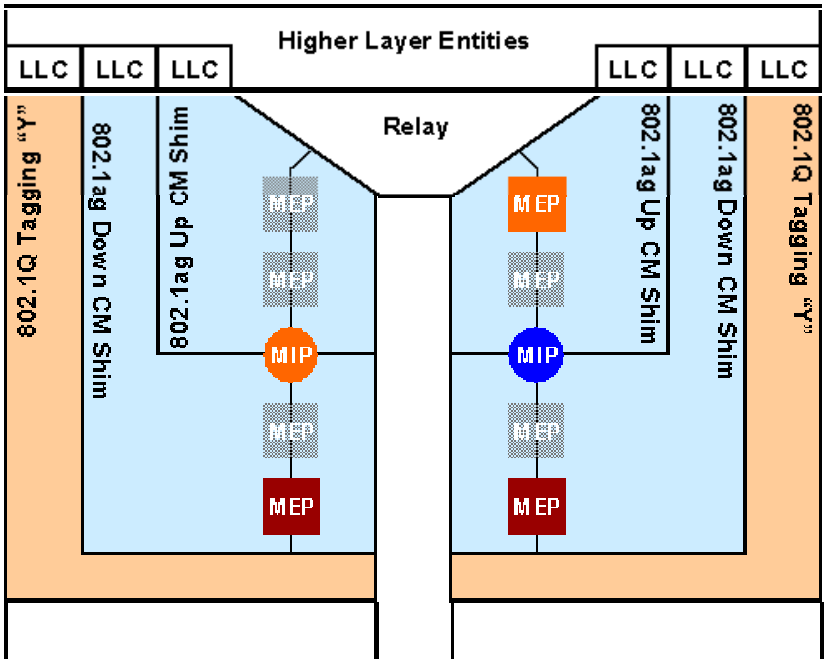

Figure II-8 – Illustrating the location of MEPs in the "baggy pants" model for B4

- operator A bridge number 4 has two ETH link related MEPs active, one in each interface port

- furthermore the operator A domain facing interface port (left) has its MIP active for use in the opertor A ETH ME Level
- the operator B facing interface port (right) terminates operator A's ETH ME and has for that purpose a MEP in the up shim active
- as this right interface port is at a domain boundary, it has to support a MIP for the next higher ETH ME (service provider), to allow fault localization by this service provider (inside network of operator A, inside network of operator B or in the link between A and B)

The reader is assumed to be able to draw the MEP/MIP configurations for the other bridges at this point. Those are not shown therefore.

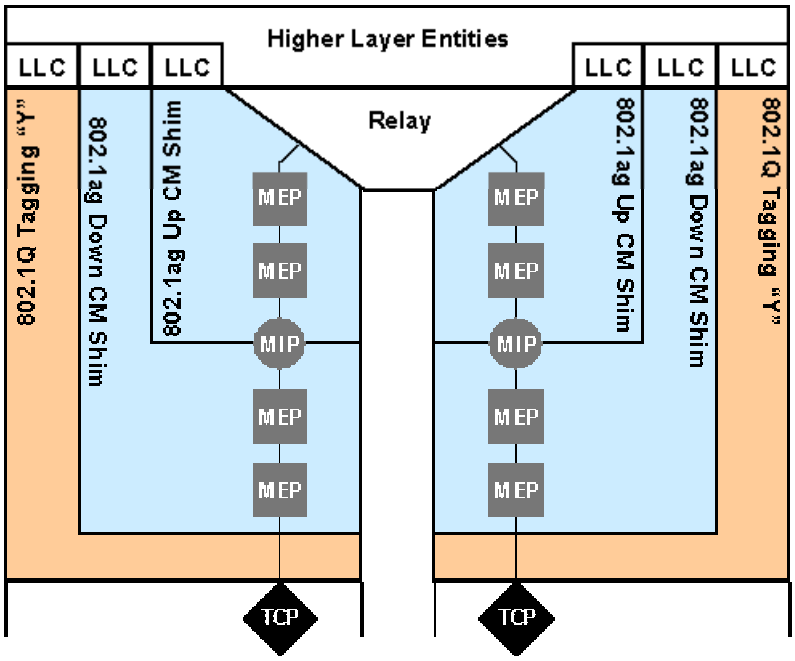

Figure II-9 – Illustrating the location of TCPs in the "baggy pants" model

- the interface port at the network side of a UNI will/may have a TCP that is located below the Down CM Shim in the baggy pants model. In this way the MEPs in the down shim will be able to register the effect (discarding) of the traffic conditioning and report this to the customer and service provider who share the responsibility for this UNI link

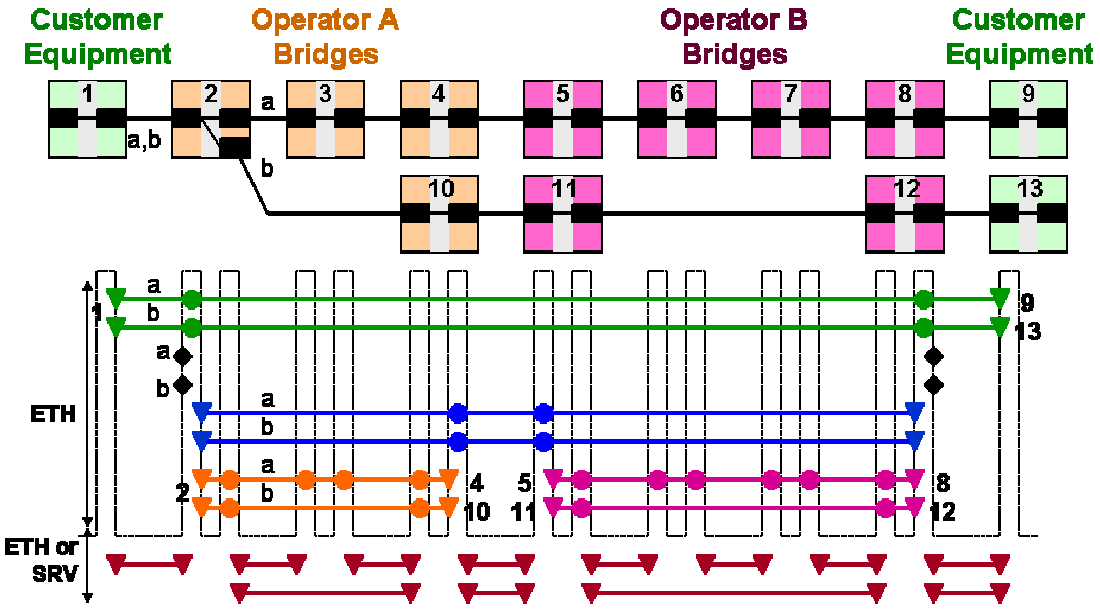

Figure II-10 – Illustrating the location of MEPs, MIPs and TCPs in the "baggy pants" model for the case of a two p2p connections with multiplexed access

- two p2p connections (a,b) with associated CM are depicted
- bridge 2 has two parallel sets of MEPs/MIP/TCP in the UNI facing port
- note: the figure depicts a single ME between CE1 and B2. This implies that this ME is a SRV ME. If it would have been an ME at ETH layer, then there should have been two ETH MEs, one for each p2p connection

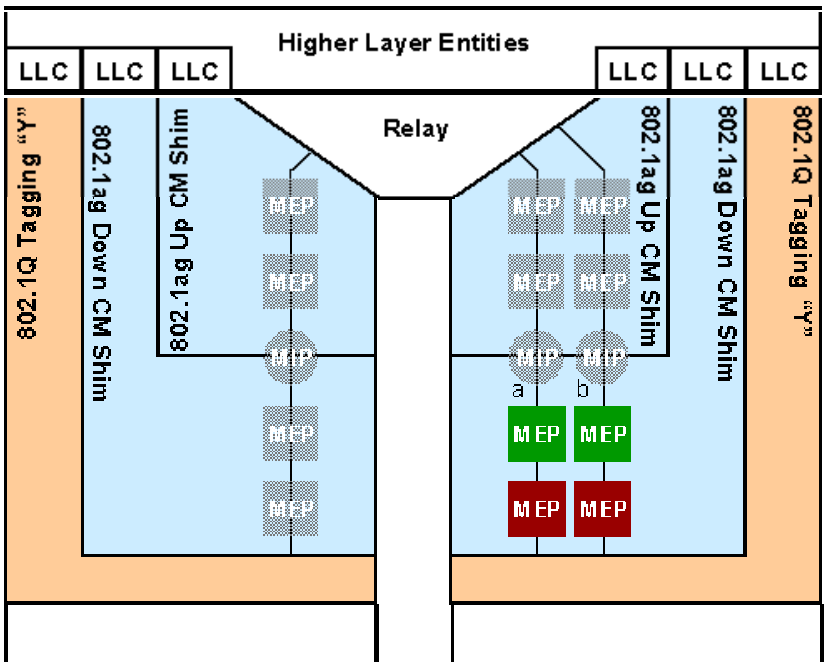

Figure II-11 – Illustrating the location of MEPs, MIPs and TCPs in the "baggy pants" model for the case of a two p2p connections with multiplexed access

- customer equipment number 1 has for the two ETH connections of figure II-10 two MEPs activated in the down shim to monitor the two UNI-C to UNI-C connections and the CE1 to B2 links

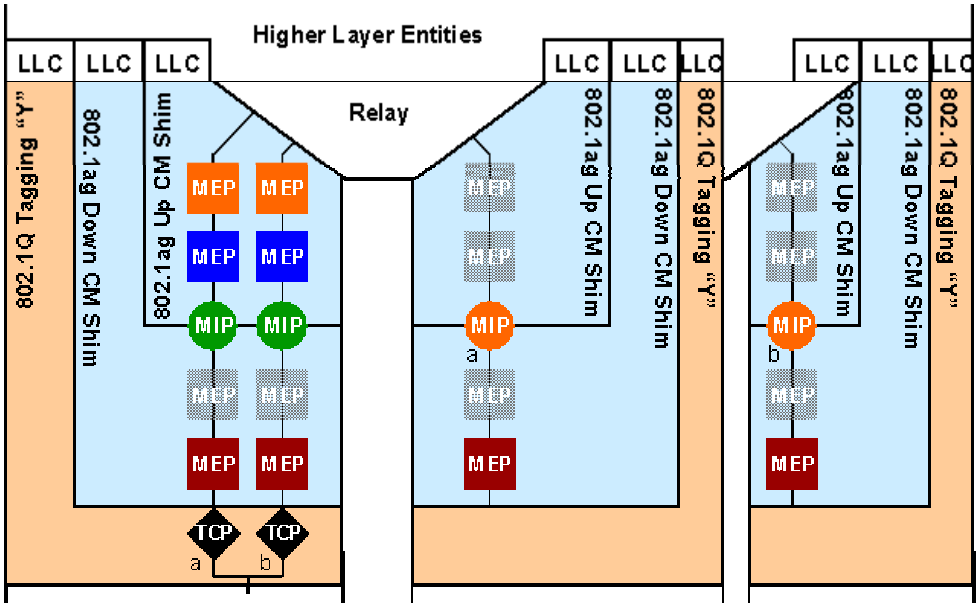

Figure II-12– Illustrating the location of MEPs, MIPs and TCPs in the "baggy pants" model for the case of a two p2p connections with multiplexed access

- operator A bridge number 2 has at its customer equipment facing port for each of the two p2p connections three MEPs active; one to terminate the link ME, one to terminate the service provider's ETH ME (blue) and one to terminate the operator A's ETH ME
- this port has also an active MIPfor each of the two p2p connections for use by the customer's UNI-C to UNI-C ETH MEs
- there are two interface ports facing the network, one for each of the two p2p connections

MEP, MIP, TCP for Dual Relay Model & Bundling MEP, MIP, TCP for Dual Relay Model & Bundling

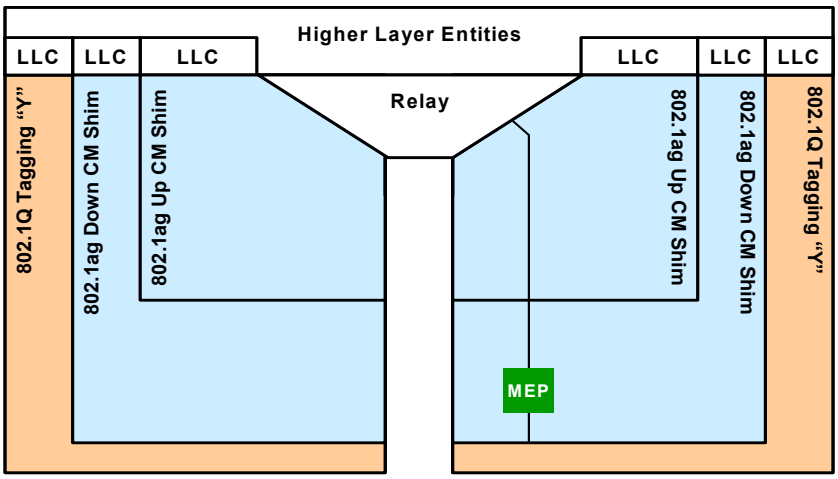

Figure II-13: Customer Bridge 1, example without ETH link ME

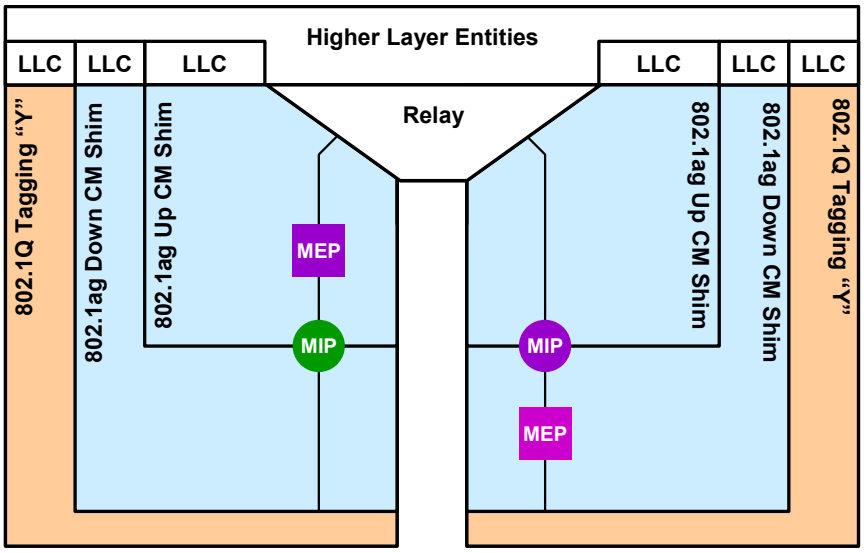

Figure II-14: Provider Bridge 2a, Example without ETH link ME

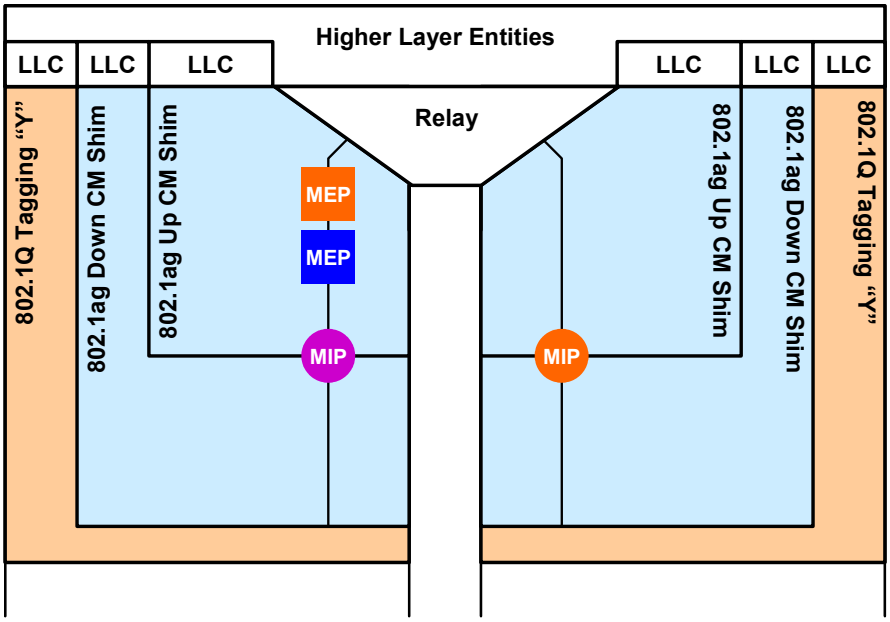

Figure II-15: Provider Bridge 2b, example without ETH link ME
Dual Relay Model with Single Relay as Provider Device

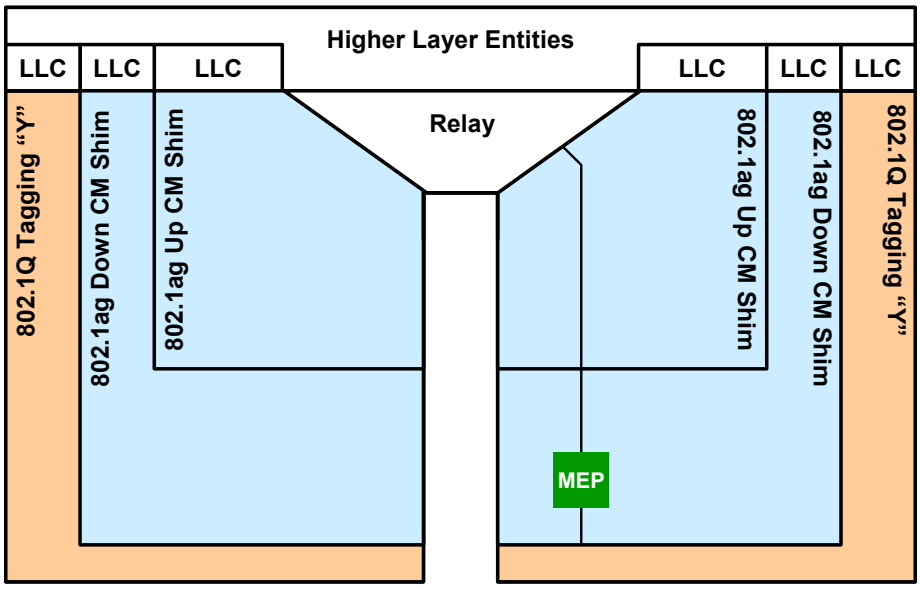

Figure II-16: Customer Bridge 1, example without ETH link ME

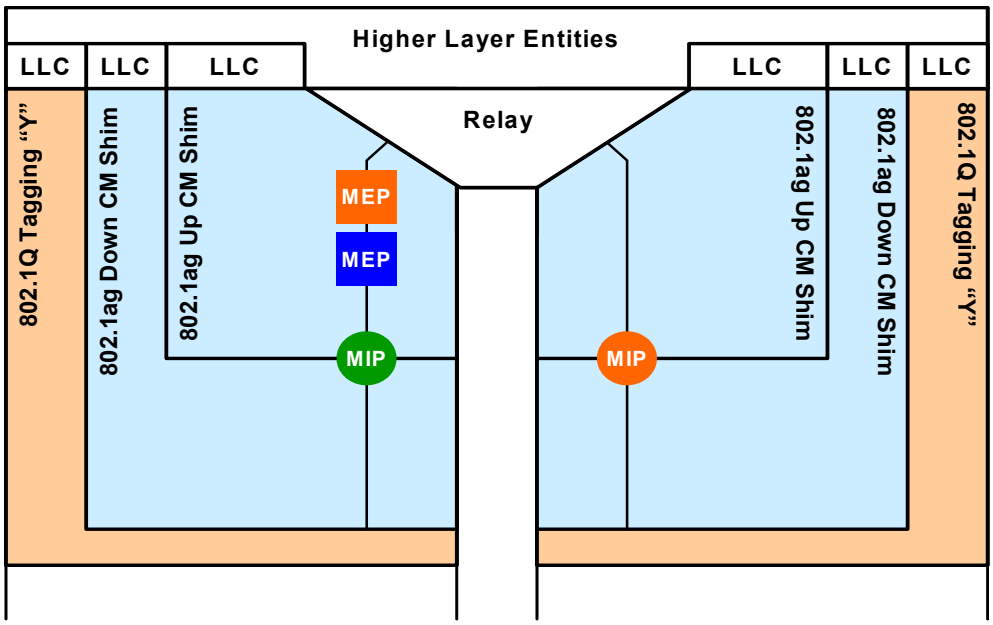

Figure II-17: Provider Bridge 2b, example without ETH link ME
Dual Relay Model with bundling for Single Integrated Provider Device

Appendix III: AIS Considerations & Issues

EDITOR'S NOTE: THE FIRST PART OF THIS APPENDIX (UP TO FIGURE III-7) WAS TAKEN FROM A WD BROUGHT IN NOVEMBER 2003, IT IS NECESSARY TO REVISE IT AND DECIDE WHICH PART OR IF ALL SHOULD BE DELETED OR REPLACED. IT NEEDS ALSO TO BE CONSOLIDATED WITH THE SECOND PART OF THIS APPENDIX (AFTER FIGURE III-7). NECESSARY INPUT FROM DINESH AND MAARTEN AS BOTH WERE THE CONTRIBUTORS

III-1 ETH alarm suppression OAM considerations (ETH-AS considerations)

WD27 introduces a multipoint ETH connection example in Figures 3 and 4/WD27. WD28 illustrates the ETH-AS insertion points and the ETH MEs present on the ETH links. WD28 also introduces three alternatives to identify the ME Level. Two of these alternatives (MELI ID, STID) are being used in this contribution to analyse the ETH-AS behaviour.

Figure III-1 illustrates the MEs present on some of the links in a multipoint ETH connection (see also WD28).

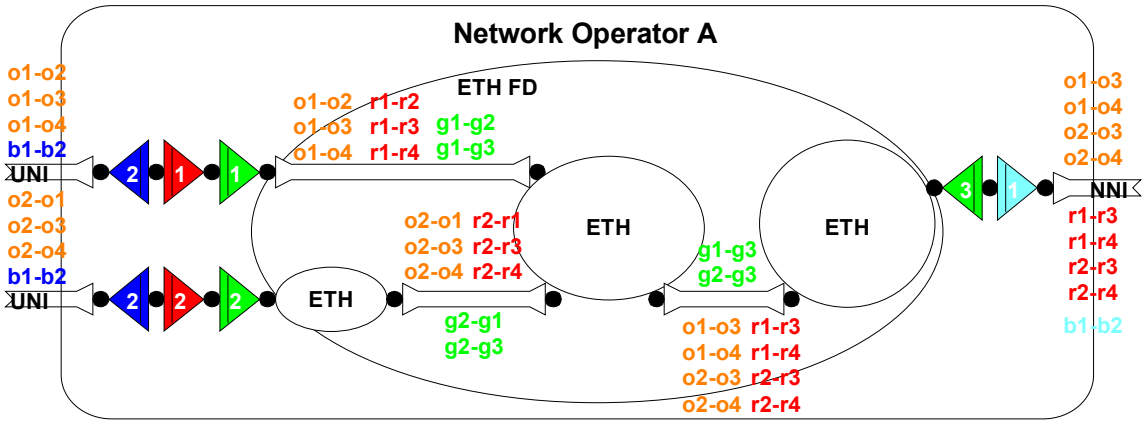

Figure III-1 – ETH MEs on ETH links

III-2 ETH-AS when deploying MELI ID in ETH-CC

When deploying an ETH ME Level instance ID (MELI ID) in ETH-CC OAM frames to identify the ME Level the CC frame belongs to, this MELI ID information can be used at an ETH link end (and an ETH segment end) to learn the set of ETH ME Levels passing through the ETH link and ETH segment. From the port identifier information present in the ETH-CC frames an ETH link end (and an ETH segment end) is able to learn the set of upstream ports that connect through the link or segment. Figure III-2 illustrates this learning at ETH link ends (Srv/ETH(-m)_A_Sk) and ETH segment ends (ETHS/ETH_A_Sk).

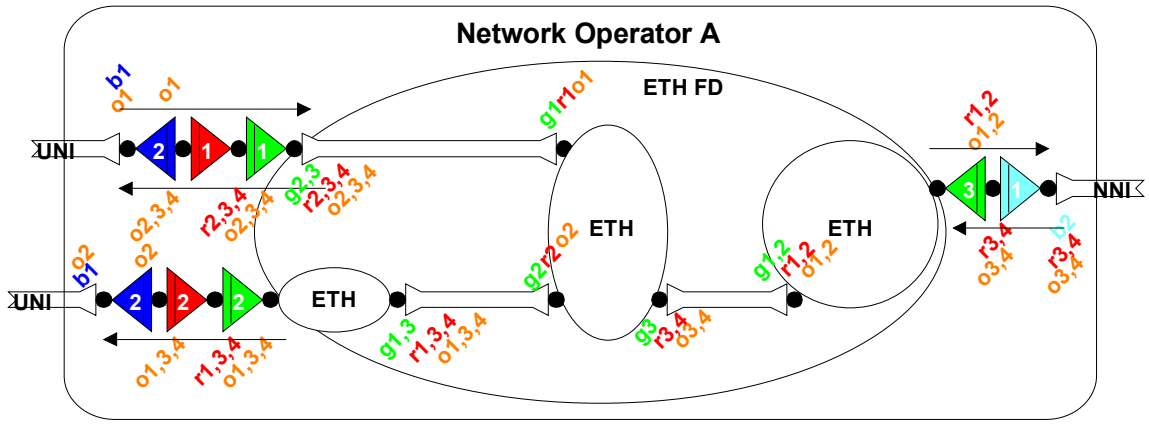

Figure III- 2 – ETH ME Levels and upstream ports learned

Consider a fault occurring in an ETH link in one direction (Figure III-3), a set of MEs (at multiple levels) is impacted. ETH-AS signal generation would in such case use the learned ME Level and upstream port number information and generates ETH-AS frames per ME Level instance, including the set of upstream port numbers.

For the case STP is present, an ETH link fault in a single direction will disable the use of the other direction of the link (for the traffic frames). At this point it is assumed that we will specify a kind of "Reverse Direction Link Down (RDLD)" maintenance signal[2] that runs between a Srv/ETH(-m)_A_So and a Srv/ETH(-m)_A_Sk function to inform the far end of the link that it is down. This signal should then result in ETH-AS signal generation at the far end of the link as well (aAIS = aSSF or dRDLD).

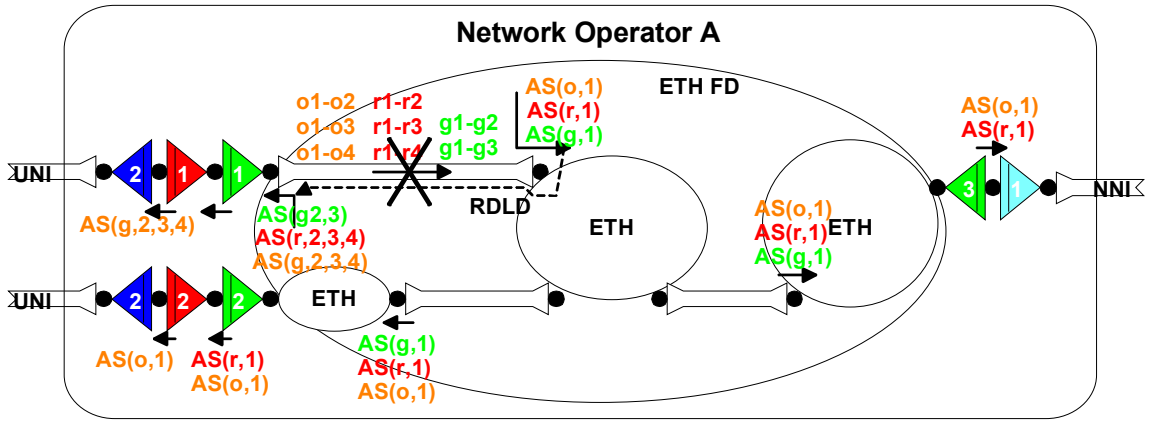

Figure III-3 – ETH-AS insertion example I

---

[2] As a first approximation (and perhaps already sufficient), RDI/BDI signals can be used as RDLD signal. The parameter controlling the port state MAC_Operational = CI_SSF or dRDLD. As a first and perhaps sufficient approximation MAC_Operational = AI_TSF or dRDI (from e.g. Sn_TT_Sk or Sn-X-L_TT_Sk).

The different ETH-AS signals are forwarded[3] by the ETH flow domains and each ETHS_FT_Sk function extracts the ETH-AS signals of its ME Level and processes the included information (upstream port numbers that are disconnected due to fault). It will use this information to suppress the associated loss of continuity fault causes that will be detected as a consequence of the link fault.

The ETH-AS signals for other ME Levels are simply passed through these ETHS_TT_Sk functions.

Figure III-4 present a second example with a bi-directional ETH link fault. Figure III-5 assumes an alternative link being available in the topology, which is initially blocked by spanning tree (or network management, or …). After ETH link fault is detected e.g. STP will restore the ETH connection by taking the black link part of the active topology. At the same time it will block traffic (including ETH-AS OAM) incoming to the ETH-FDs at the end of the failed link. A blocked port will have to flush their learned set of ME Level instances and upstream port numbers.

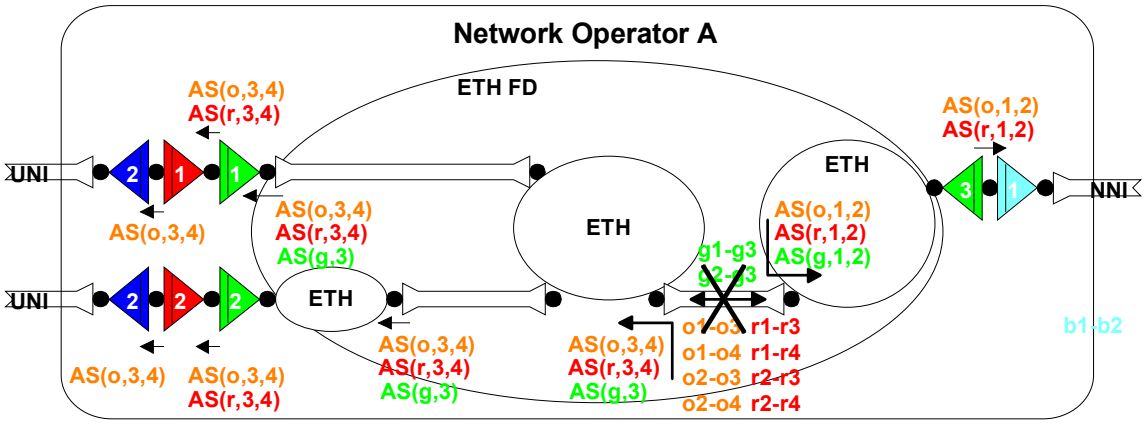

Figure III-4 – ETH-AS insertion example II

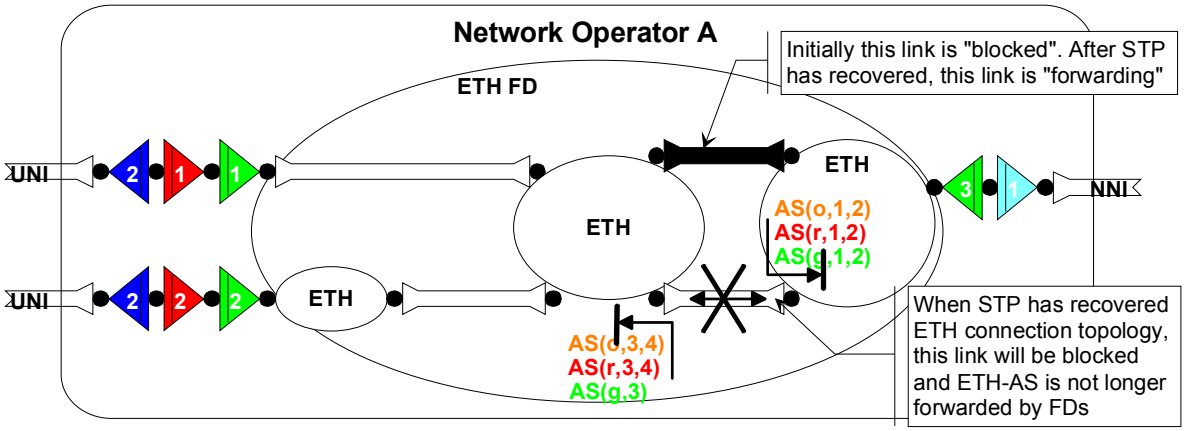

Figure III- 5 – ETH-AS insertion example II with restoration capability

ISSUE: what if the topology only can be partially recovered…

---

[3] On a link fault, the port state changes as far as I understand… will this have any impact on the forwarding of these generated and inserted ETH-AS signals?

NOTE – if instead of bridges an MPLS (VPLS) network would be used that would run Y.1711 OAM, there would be a look alike, feel alike management behaviour; the ETH MEs are now replaced by MPLS MEs…

III-3 ETH-AS when deploying STID in ETH-CC

Figure III-6 illustrates the port identifiers of the ME at the top of the stack within a Srv/ETH adaptation sink function (link end) or ETHS/ETH adaptation sink function (segment end) in a multipoint ETH connection. Much less learning is required in this situation, and that is what is attractive… it also has a price…

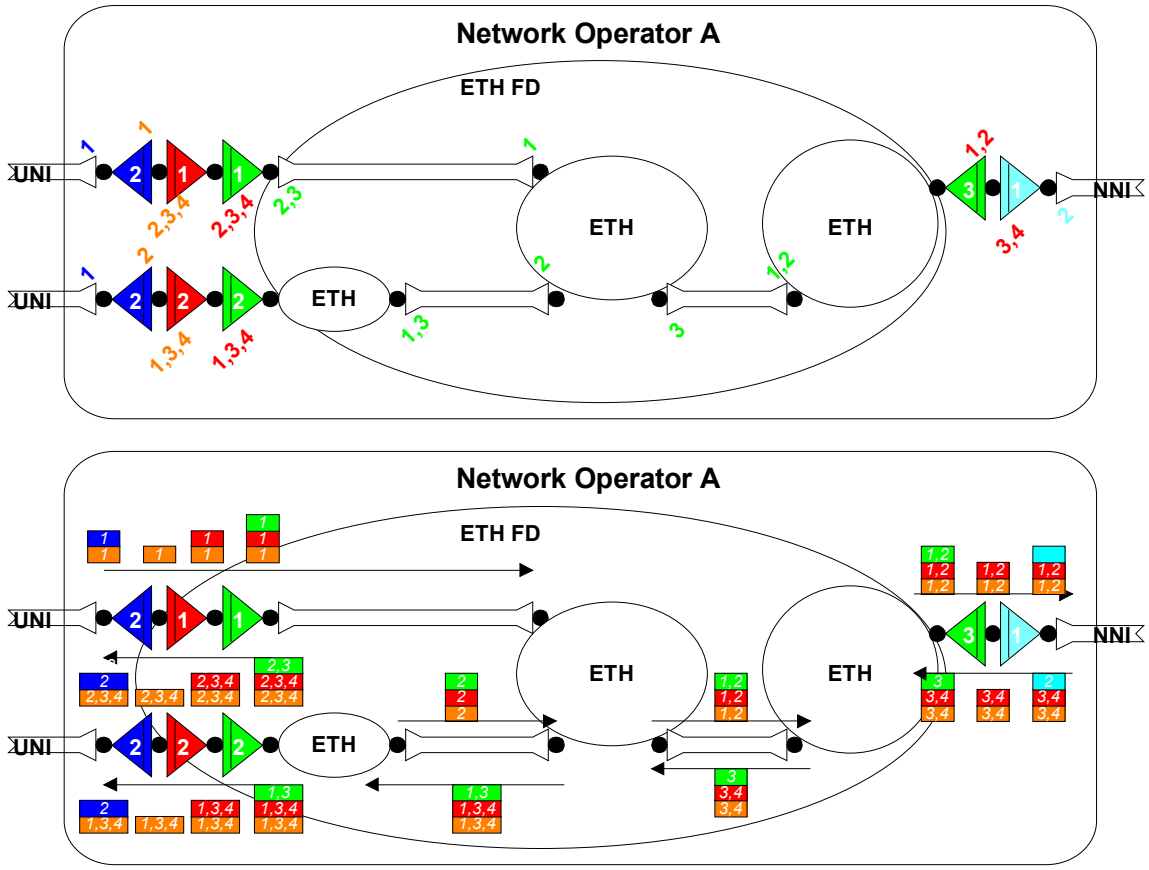

Figure III-6 – ETH ME port identifiers at the top of the stack (top) and full stack (bottom)

A link fault (Figure III-7) will now generate a single ETH-AS frame with upstream port numbers from the ETH link ends for the top level ME. Then at the first segment endpoints (green) these ETH-AS signals are extracted and processed. The signal fail status is forwarded to the adaptation sink function in the segment endpoint, where it has to trigger insertion of ETH-AS for the interrupted top level (red) ME. Unfortunately there is insufficient information at these points to generate ETH-AS frames with specific upstream port number list.

So, should we generate non-specific ETH-AS frames (then also at link ends)? The consequence is that it also will suppress the reporting of a true ETH layer continuity or connectivity fault located elsewhere in the ETH connection… should our ETH OAM be able to detect and report a dual fault condition?

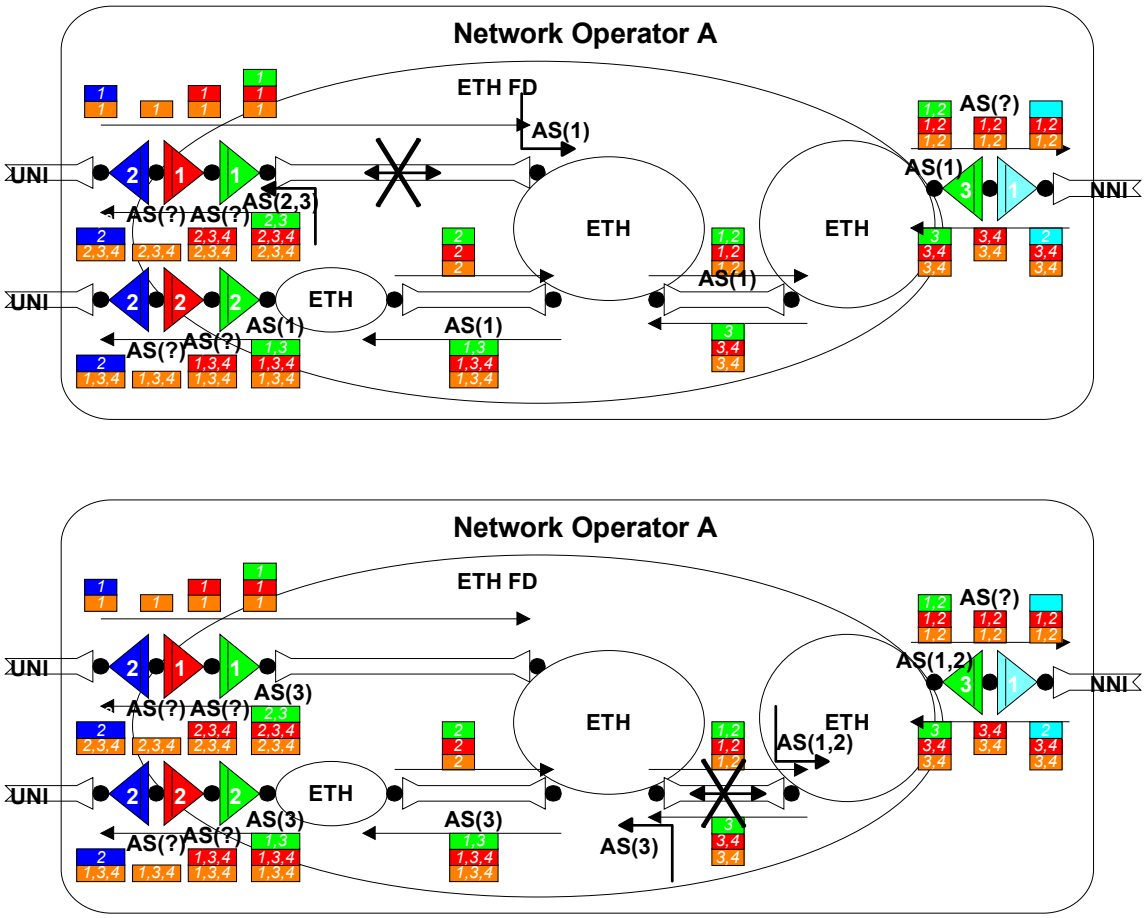
Network Operator A
ETH FD
AS(1)
UNI
UNI
AS(?) AS(?)
AS(2,3)
ETH
ETH
ETH
AS(1)
AS(1)
AS(?)
AS(?)
AS(1)
AS(?)
AS(?)
NNI
Network Operator A
ETH FD
UNI
UNI
AS(3)
AS(?) AS(?)
ETH
ETH
ETH
AS(1,2)
AS(1,2)
AS(?)
AS(?)
AS(3)
AS(?)
AS(3)
AS(3)
NNI

Figure III-7

EDITOR'S NOTE: THE FOLLOWING MATERIAL WAS ADDED FROM WD 12 TO JUNE 7-11 MEETING. IT SHOULD BE CONSOLIDATED WITH THE PREVIOUS MATERIAL: INPUT REQUESTED FROM DINESH AND MAARTEN

Figure III-8 represents a reference network where 3 bidirectional point-to-point services are assumed i.e. S12 (CE1-CE2), S13 (CE1-CE3), and S14 (CE1-CE4). Nodes PE1, PE2, PE3, and PE4 represent the provider edge nodes, while nodes P1, P2, and P3 represent the provider core nodes. The distinction between the core and edge provider nodes is simply that core nodes are not connected to any CE nodes, as per the reference network in Figure III-8.

Since redundancy is shown to exist in the network, links $P2_2$-$P3_3$ and $P1_3$-$P2_1$ may get blocked, either by Spanning Tree Protocol (STP) [3] or manual provisioning. The callouts in Figure 1 represent a view of Forwarding Information Base (FIB).

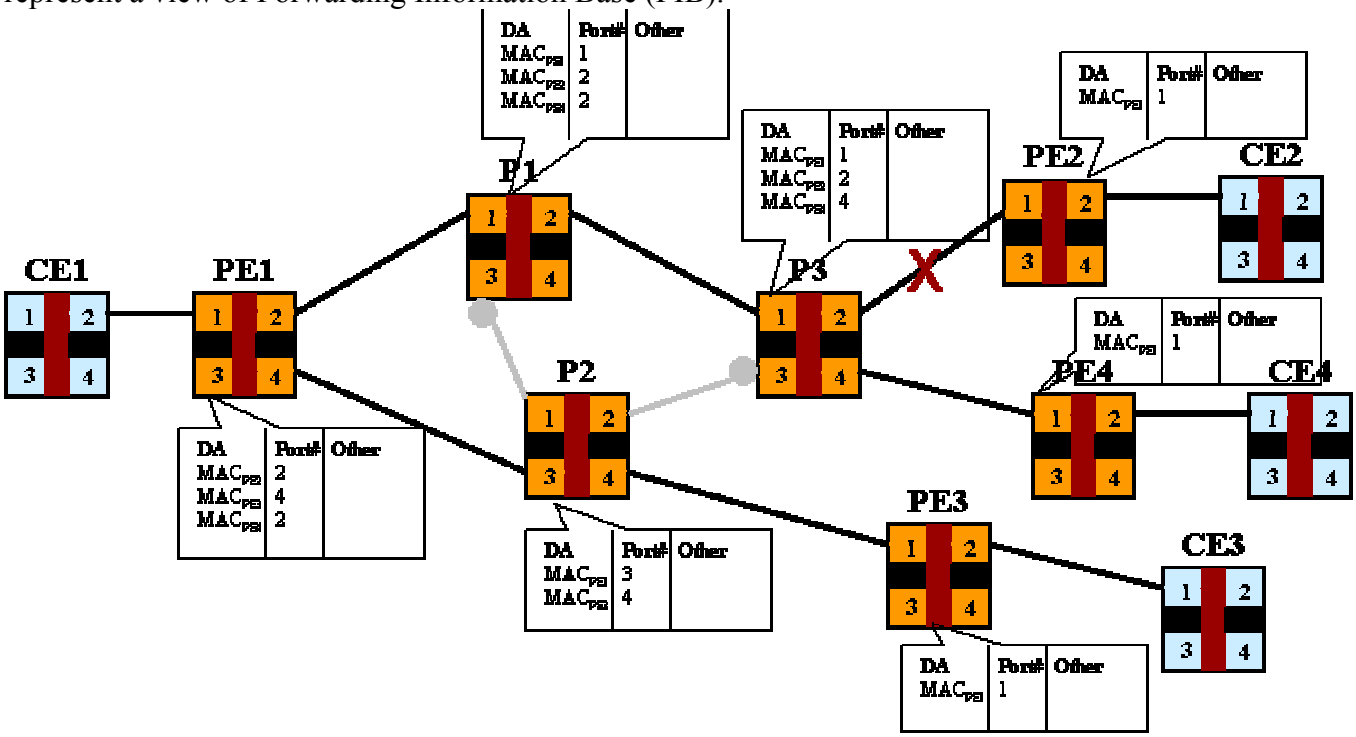

Figure III-8: Connectionless reference network for AIS with link failure scenario #1

A) Link Failure Scenario 1

When a link failure is considered, e.g. link $P3_2PE2_1$, service S12 is affected. Assuming that the link failure is detected on either end of the link, port $P3_2$ and port $PE2_1$ detect this failure. Now the possible options for node P3, if it supports AIS capability, are:

i. Send AIS across all other ports

ii. Send AIS selectively across selective ports

iii. Not send AIS at all

When considering option (i), sending AIS to all ports is not very useful, e.g. PE3 does not have any use for this AIS as the service instance S13 supported by PE3 is not effected by link $P3_2PE2_1$ failure

Option (ii) seems viable as the determination to forward AIS can be made on the basis of service instances e.g. P3 could determine that port $P3_2$ belongs to say VLAN 20, which is also associated with port $P3_1$ for the same point-to-point service instance S12. When sent out across port $P3_1$, the AIS is now received by node P1 across port $P1_2$. Since at P1, only other port associated with same service instance is port $P1_1$, AIS is forwarded to port $P1_1$. Such hop-by-hop forwarding of Ethernet AIS seems pragmatic.

However, one issue may arise when STP or its variants are used which result in flushing of FIBs due to Topology Change Notification (TCN) BPDUs. Under such circumstances hop-by-hop forwarding of AIS is not feasible, as the association of VLANs and corresponding ports on each node is lost due to TCN related flushing.

B) Link Failure Scenario 2

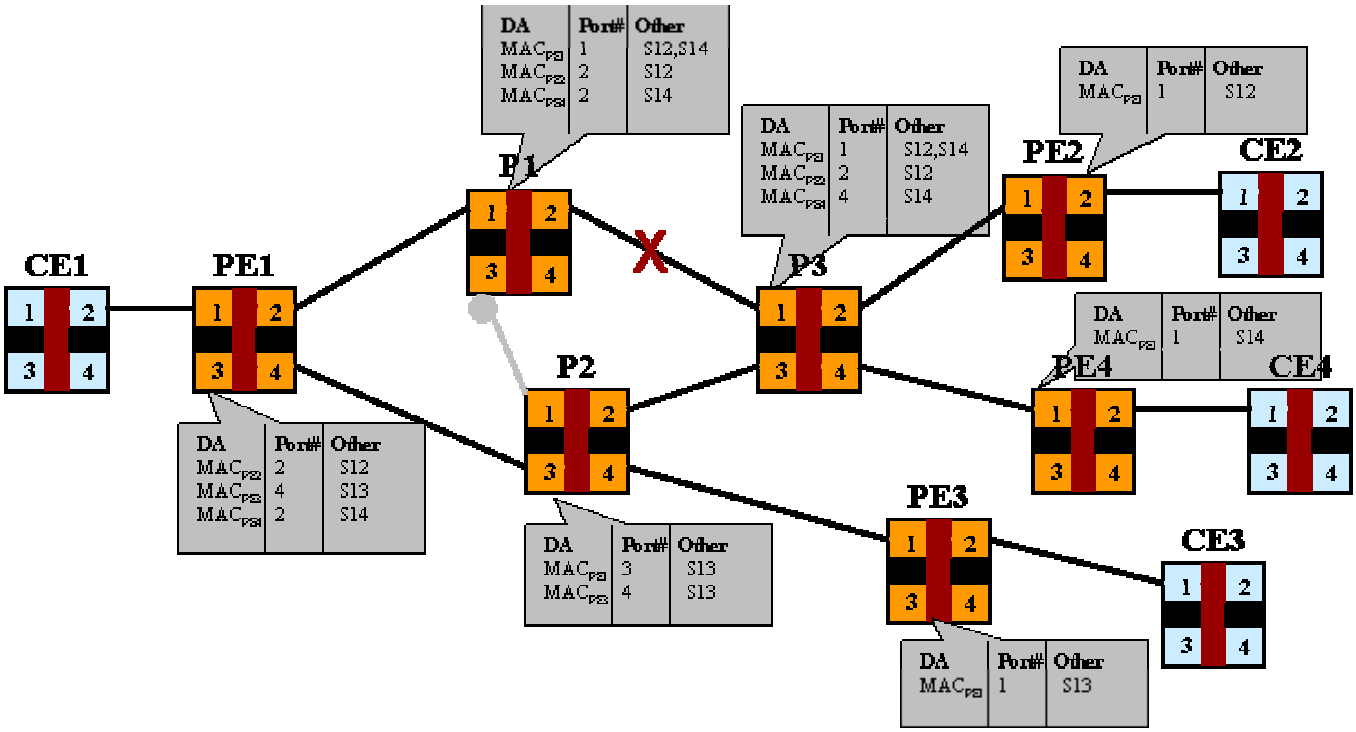

Figure III-9: Connectionless reference network for AIS with link failure scenario #2

When a link failure is considered, e.g. link $P1_2P3_1$, service S12 and S14 are initially affected since link $P2_2P3_3$ is initially blocked, either by STP or its variants or by manual provisioning. However, since this link failure is not a network isolating failure, e.g. link $P2_2P3_3$ is unblocked eventually, and no permanent loss of connectivity is experienced between PE1and PE2 or PE4.

Assuming that the link $P1_2P3_1$ failure is detected on either end of the link, port $P1_2$ and port $P3_1$ detect this failure. Now the possible options for node P3, if it supports AIS capability, are:

i. Send AIS across all other ports

ii. Send AIS selectively across selective ports

iii. Not send AIS at all

Similar to discussions in A), option (ii) is desirable when AIS functionality is supported and required.

However, one issue may arise when link $P2_2P3_3$ is unblocked and port $P3_3$ on node P3 now joins the same service instance as port $P3_1$. Following questions arise:

a) Whether node P3 should forward AIS along ports $P3_2$, $P3_3$, and $P3_4$ or not generate AIS at all i.e. option (iii)?

b) Under what circumstances does the node sending AIS stop sending AIS?

c) If node P3 does send the AIS, what does these AIS mean to node PE2 or PE4 or PE1 since the service is already restored?

d) If node P3 should not send AIS or should stop sending AIS after link $P2_2P3_3$ is unblocked, how does node P3 establish association between the failure and restoration events?

Similarly, when it is assumed from above discussion that node P3 does forward AIS along ports $P3_2$, $P3_3$, and $P3_4$, node P2 is likely to receive both AIS and service frames and other OAM frames (e.g. CC) for the same service instance across the same port. Question arises:

e) Whether node P2 should forward AIS along ports $P2_3$ or should ignore AIS and not forward it?

Further, if now another service instance S23 is created between nodes CE2 and CE3, ports $P3_3$ and $P2_2$ and $P2_4$ are now also associated with S23 service instance. This reflects the need for per service level AIS since otherwise AIS related to link $P1_2P3_1$ failure would get forwarded to node PE3 since port $P3_3$ is now associated with different service instances including S23 and port $P2_2$ is associated with different service instances including S23 as well.

C) Other Issues

Based on the above discussions, it is also important to consider following additional issues:

f) If AIS is required to be generated per service basis, given a single facility could carry thousands of services, the amount of AIS related traffic can be significant, especially around the time when the network has just experienced a fault condition!

g) The above situation is further problematic when the AIS is required to be forwarded along each higher level ME within the network operator, service provider and/or customer domains.

h) Is it always desirable to suppress service level alarms, if the facility level alarms have been detected, OR it is possible that service level alarms are still required independent of network level alarms since the OSS/NMS systems might be set up such.

APPENDIX IV: Reference Managed Objects

Some existing Management Objects (MOs) that can be used for the performance management mechanisms mentioned in Section 8 include:

- **IEEE 802.3-2002**
  - aFramesTransmittedOK [5 – section 5.2.2.1.2]
  - aFramesReceivedOK [5 - 5.2.2.1.5]
- **IEEE 802.1Q-2003**
  - Frames Received [6 - 12.6.1.1.3]
  - Frames Outbound [6 - 12.6.1.1.3]
- **RFC 3635 - Ethernet-like interface MIB (Obsoletes 2665)**
  - IF-MIB
    - ifOutUCastPkts
    - ifOutMulticastPkts
    - ifOutBroadcastPkts
    - ifOutErrors
    - ifOutDiscards
    - ifInUCastPkts
    - ifInMulticastPkts
    - ifInBroadcastPkts
    - ifInErrors
    - ifInDiscards
  - aFramesTransmittedOK = ifOutUCastPkts + ifOutMulticastPkts + ifOutBroadcastPkts – (ifOutErrors + ifOutDiscards)
  - aFramesReceivedOK = ifInUCastPkts + ifInMulticastPkts + ifInBroadcastPkts + (ifInErrors + ifInDiscards)
- **RFC 2674 – VLAN Bridge MIB**
  - dot1qPortVlanStatisticsTable
    - dot1qTpVlanPortInFrames
    - dot1qTpVlanPortOutFrames

Note: It may be noted that these managed objects values eventually wrap. This can lead to inaccurate results when such an event occurs. However, if the time interval of observation is small, the inaccuracy can be avoided. Averaging of the results over the period of observation can alleviate the in flight frames issue.

APPENDIX V: Frame Loss Calculations

V-1 Frame Loss Calculations

For the frame loss calculation, the four cases below should be taken into account when counters with finite digits (bits) are used.

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 following.

A) No wrapping around for both Transmit and Receive Counters

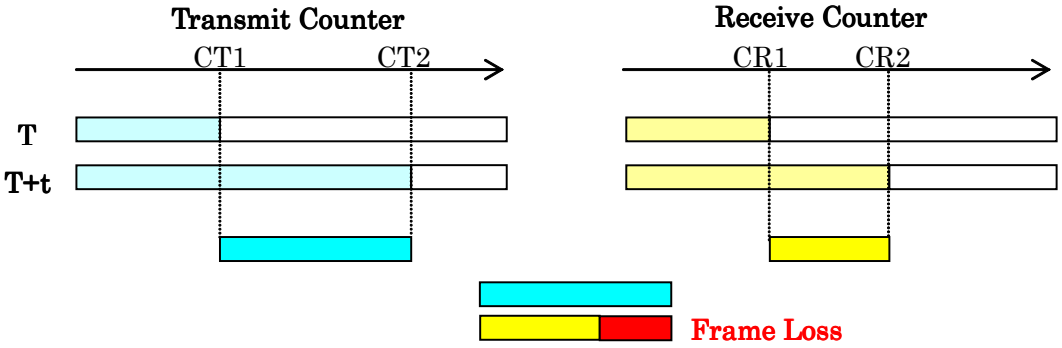

Figure V-1: A) No wrapping around

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

Frame Loss = (CT2 – CT1) – (CR2 – CR1)

B) Only Transmit Counter wraps around

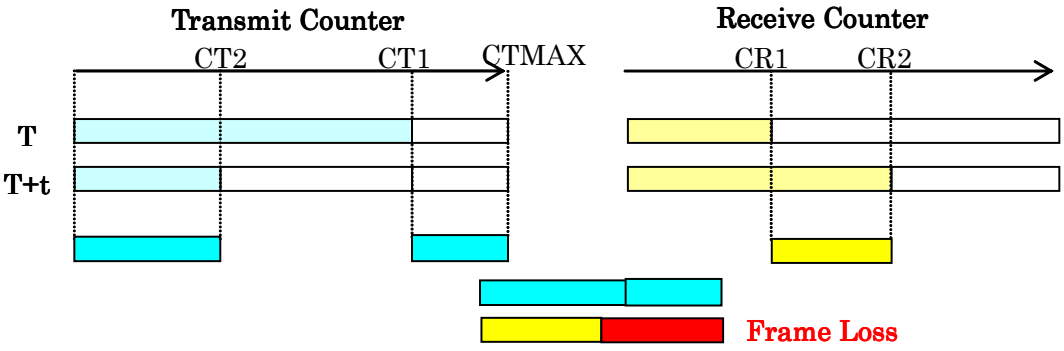

Figure V-2: B) Transmit Counter wraps around

In this case, it can be calculated by the following calculation as is described in the previous section

Frame Loss = ((CTMAX – CT1) + CT2+1) – (CR2 – CR1)

= (CT2 – CT1) – (CR2 – CR1) + (CTMAX+1)

C) Only Receive Counter wraps around

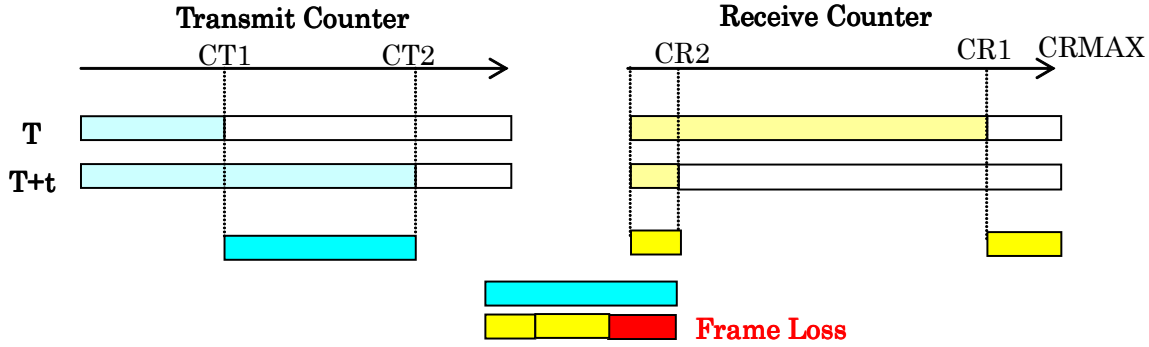

Figure V-3: C) Receive Counter wraps around

Frame Loss = (CT2 – CT1) – ((CRMAX – CR1) + CR2+1)

= (CT2 – CT1) – (CR2 – CR1) - (CRMAX+1)

D) Both Transmit and Receive Counters wrap around

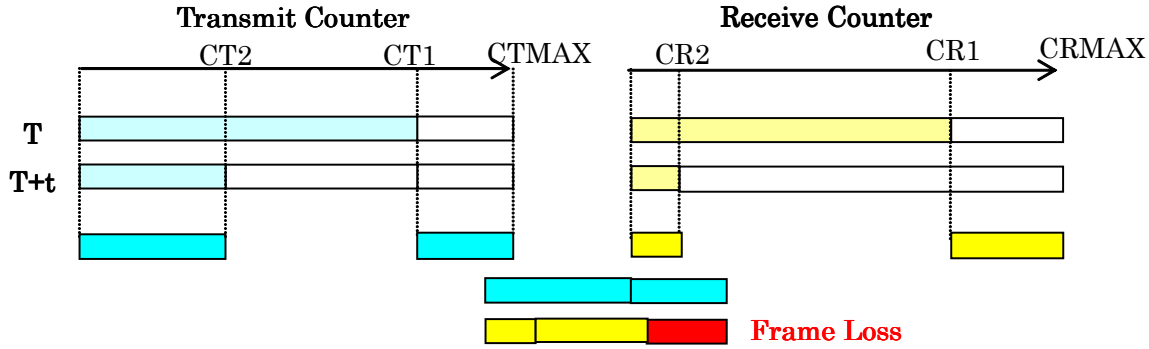

Figure V-4: D) Both Counters wrap around

Frame Loss = ((CTMAX – CT1)+CT2+1) – ((CRMAX – CR1) + CR2+1)

= (CT2 – CT1) – (CR2 – CR1) + (CTMAX+1) - (CRMAX+1)

V-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 greatly simplified by the following characteristics.

N+(MAX+1) ≡N mod(MAX+1)

N-(MAX+1) ≡ N mod(MAX+1)

Therefore each calculation formulas for frame loss which are described in the section 8.2.3 and 8.2.4 can be transformed as below.

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.

If wrapping around of counters happen more than twice, the counters for the wrapping around are required to calculate the frame loss correctly.

APPENDIX VI: OAM Filtering Function

VI-1 The OAM Filtering Functional Block

The OAM filtering function is defined in section 6.5 .

This filtering function should be performed at the MEP. It can be define as part of the ETHS function. The ETHS function consists of the ETHS_A function and the ETHS_FT function. The ETHS_FT is the flow termination function that injects and terminates or processes the OAM flow. The ETHS_A is the adaptation function that adapts OAM flow to ETH flow domain or upper ME Level. The OAM filtering functional block is defined within the ETHS_A functional block.

Fig. VI-1 shows the ETHS/ETH_A_So Function and Fig. VI-2 shows ETHS/ETH_A_Sk Function. The OAM filtering function is defined in ETHS/ETH_A_So function.

Fig. VI-3 shows the ETHS_So and the ETHD_So function. Fig. VI- 4 shows the ETHS_Sk and the ETHD_Sk function. The OAM generation function described in the ETHS_FT_So injects OAM flows. The OAM termination and processing function described in the ETHS_FT_Sk terminates or properly processes the OAM flow from ETHD function. This function terminates or processes OAM flows in case where the OAM level of ME coincides with the OAM level of flows.

The logics of OAM filtering function are

as follows:

(1) ME Level of Flow <= ME Level of MEP: Discard

(2) ME Level of Flow > ME Level of MEP: Forward

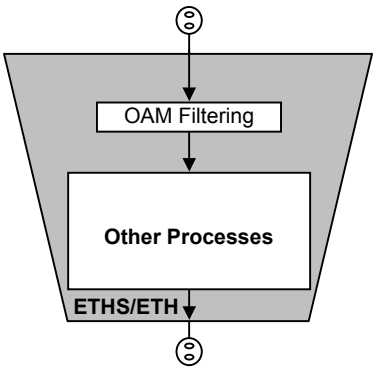

Fig. VI- 1. ETHS/ETH_A_So Function

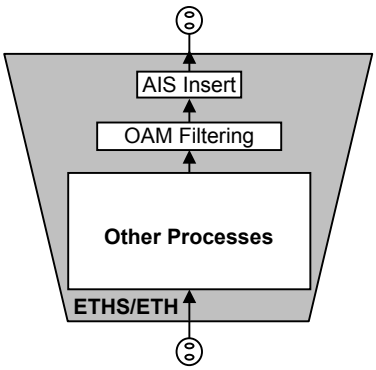

Fig. VI-2. ETHS/ETH_A_Sk Function

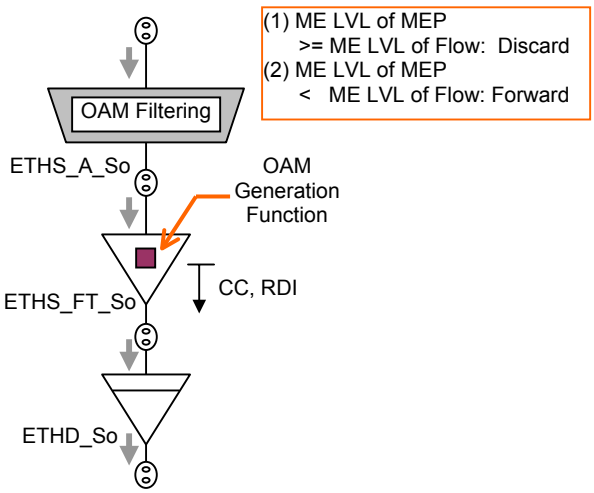

Fig. VI-3. ETHS_So and ETHD_So Function

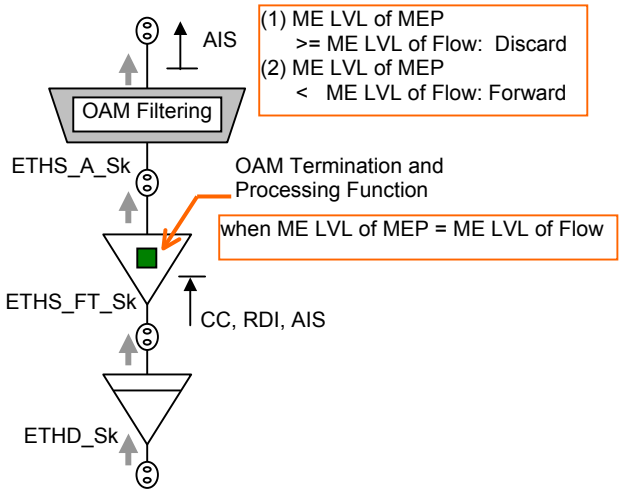

Fig. VI-4.ETHS_Sk and ETHD_Sk Function

APPENDIX VII: ETHS/ETH_A Functional Block

VII-1 The ETHS/ETH_A Functional Block

The common information element is processed by the following orderin the ETHS/ETH_A_So function.

i) DA/SA

ii) VLAN Tag

iii) OAM Type

iv) Version

v) ME Level

vi) OpCode

vii) Transaction/ Service ID

viii) Individual area

According to this order, the ETHS/ETH_A_So function block is derived as shown in Fig. VII- 1. The process of DA/SA and VLAN tag is not done in this function but in the Sev/ETH_A functional block. The detail mechanism of other function block is F.F.S. And the ETHS/ETH_A_Sk functional block and ETHS/ETH_FT_Sk/So functional block is still open.

This material will be included in G.8021.

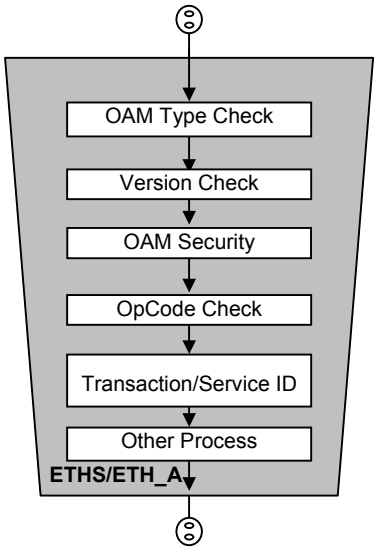

Fig. VII- 1. ETHS/ETH_A_So Function

APPENDIX VIII: OAM Functional Details

VIII-1 ETH-CC

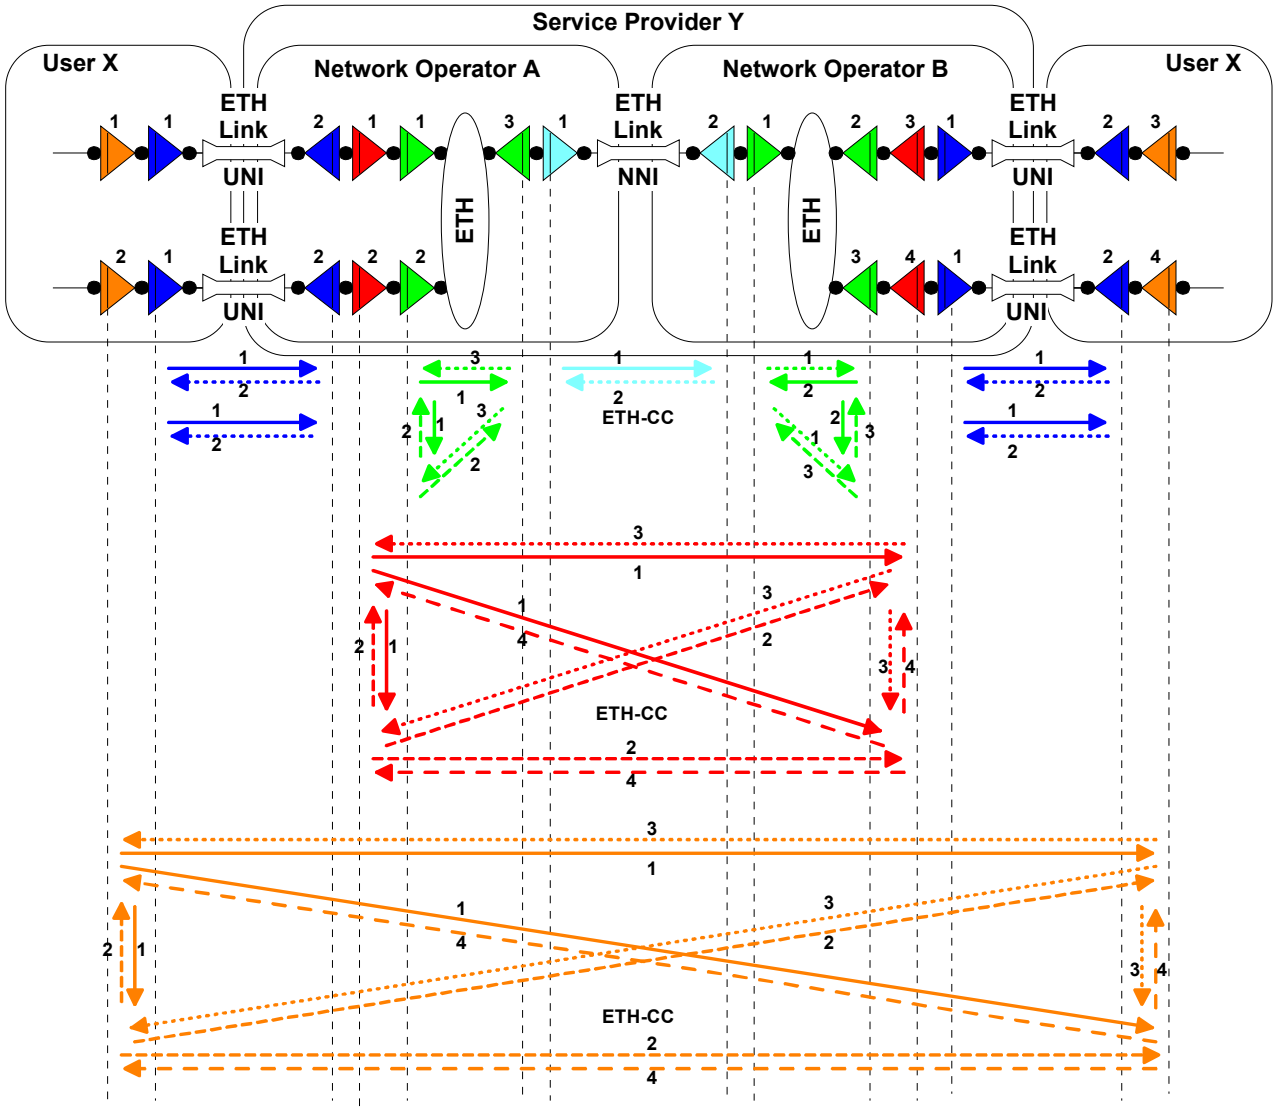

Figure VIII-1.1 – ETH-CC in multi-operator ETH multipoint connection

[EDITOR'S NOTE-Dec2004] Figure 7-1.2 needs to clarify that ETH-CC signals start with multicast DA.

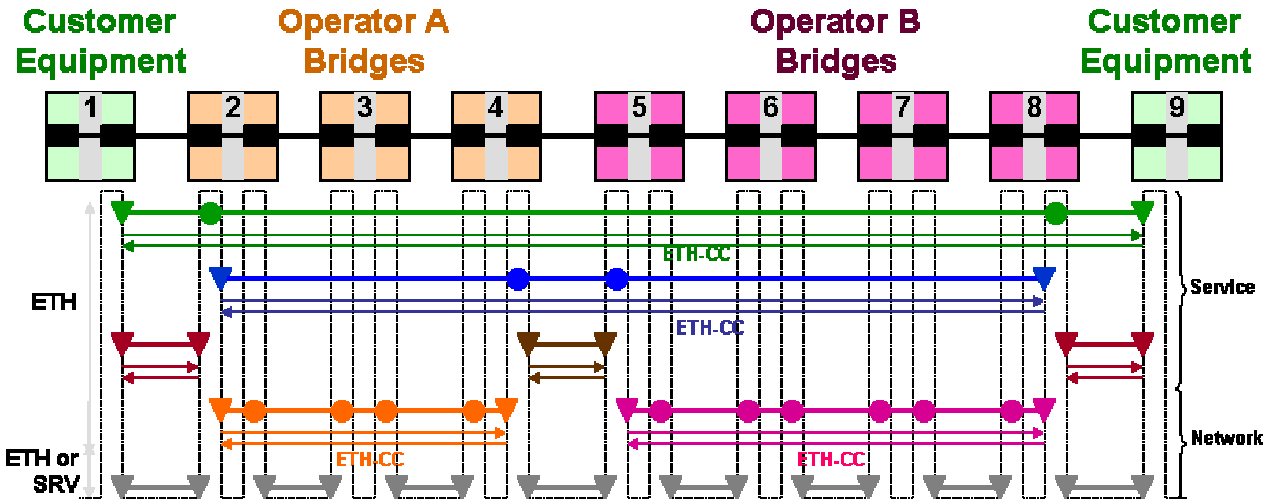

Figure VIII-1.2– ETH-CC in multi-operator ETH point-point connection

VIII-2 ETH-LB

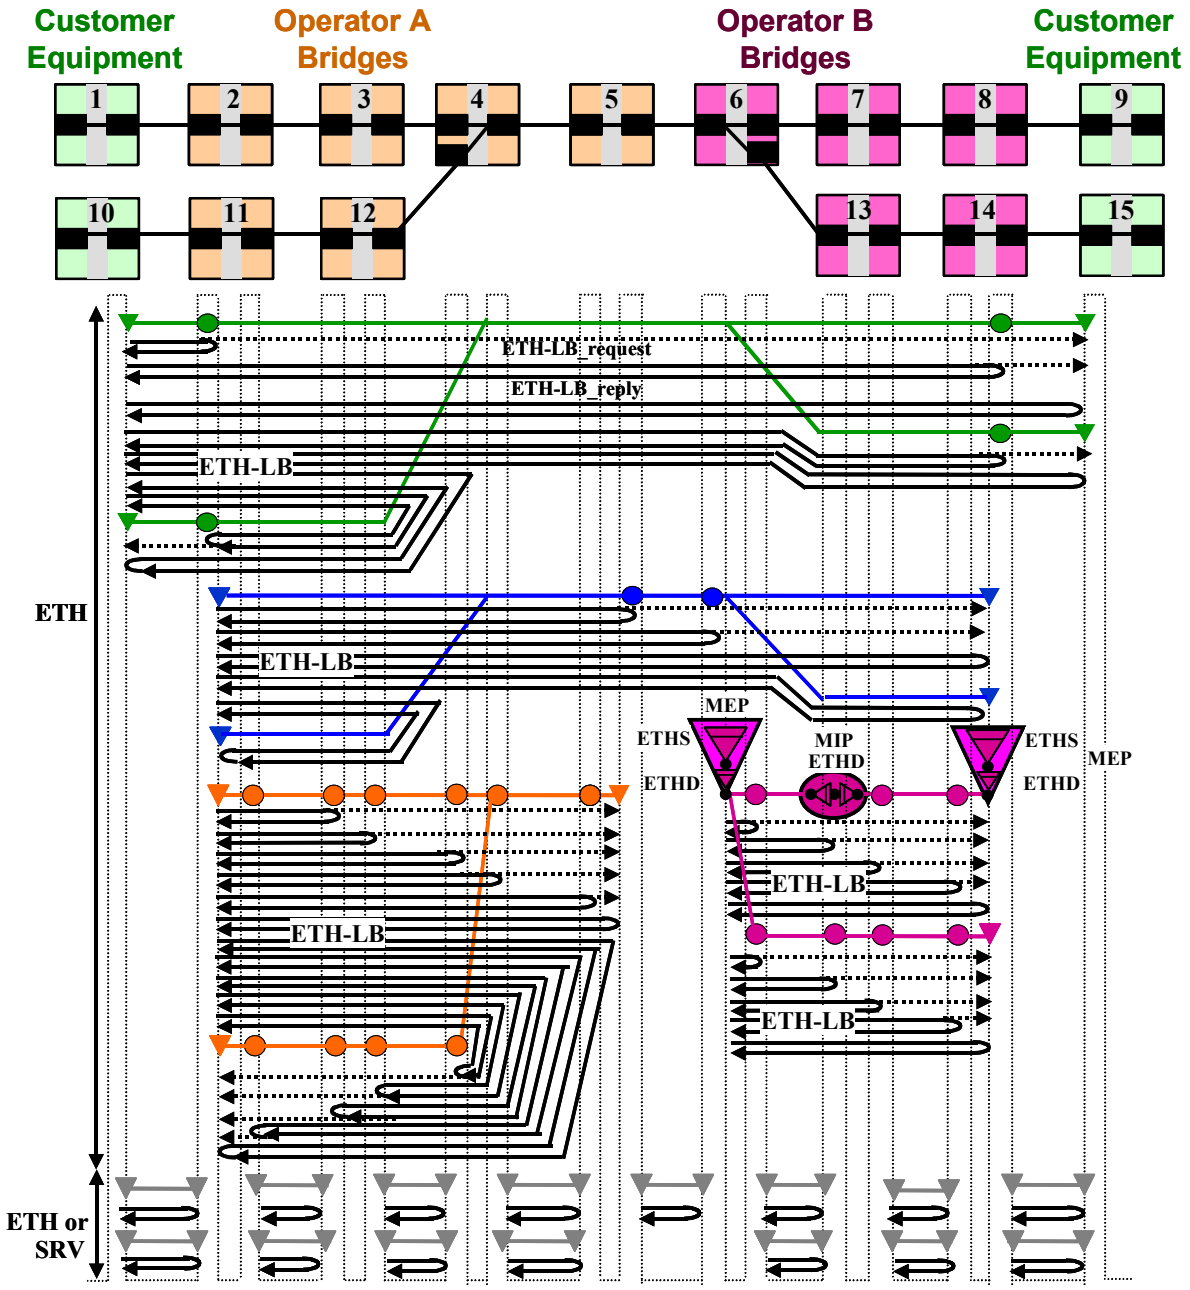

Fig VIII-2.1 – Unicast Non-Intrusive Loopback in ETH multipoint connection

(Dash lines illustrates an ME domain "shield" that stops ETH_LB frames from leaking)

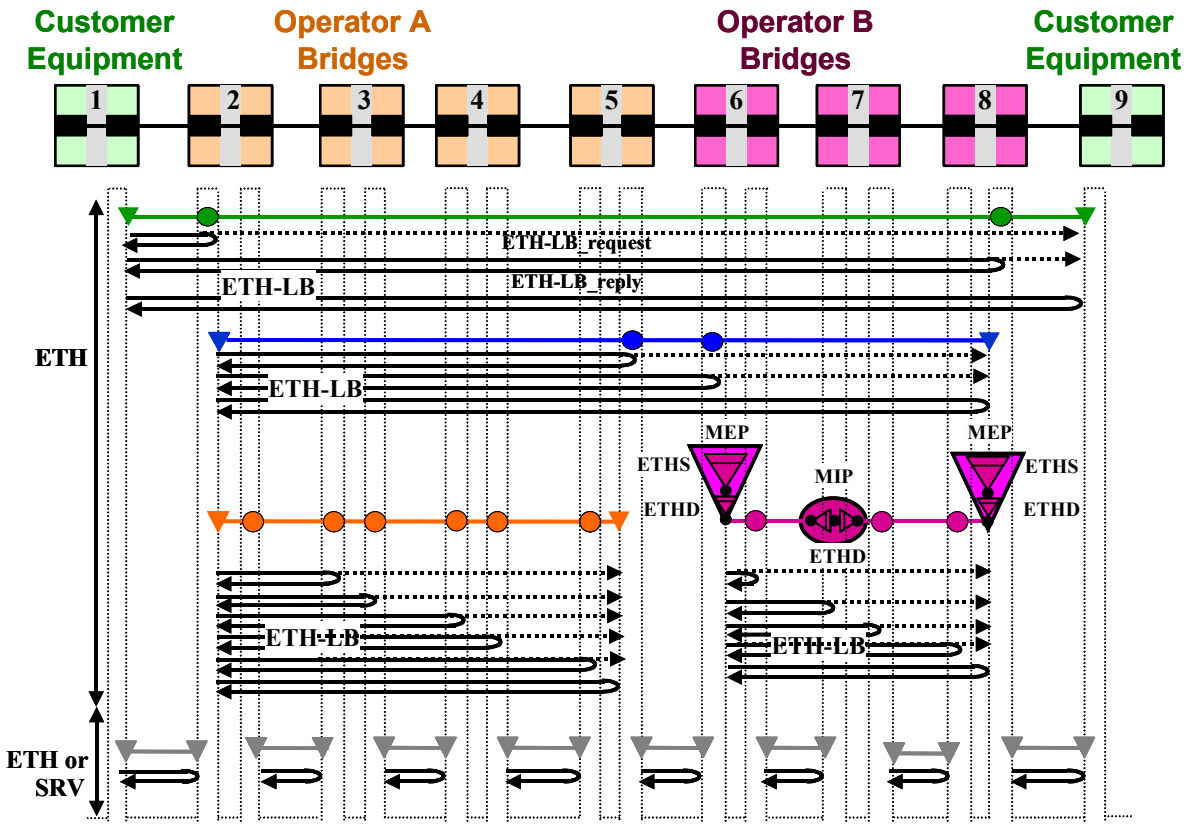

Figure VIII-2.2 Non-intrusive ETH-LB in ETH point-to-point connection

(Dash lines illustrates an ME domain "shield" that stops ETH_LB frames from leaking)

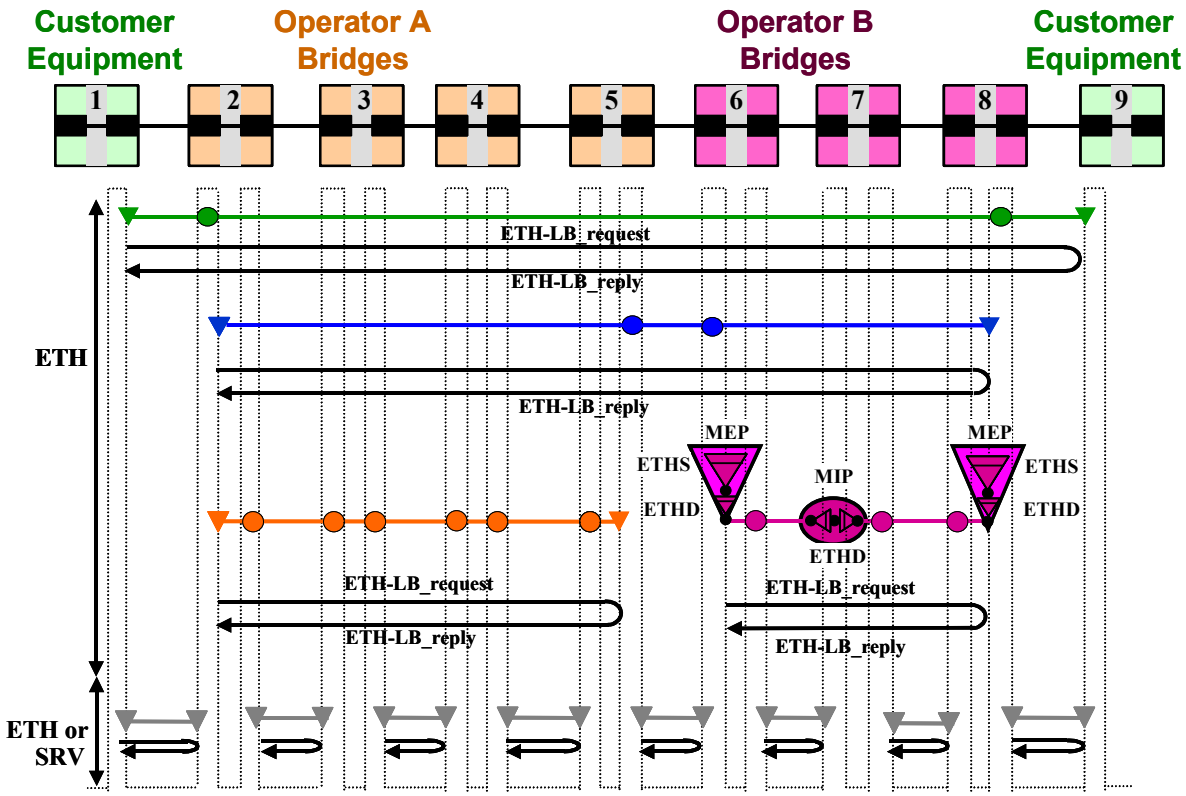

Figure VIII-2.3 Multicast Non-Intrusive Loopback in ETH point-to-point connection

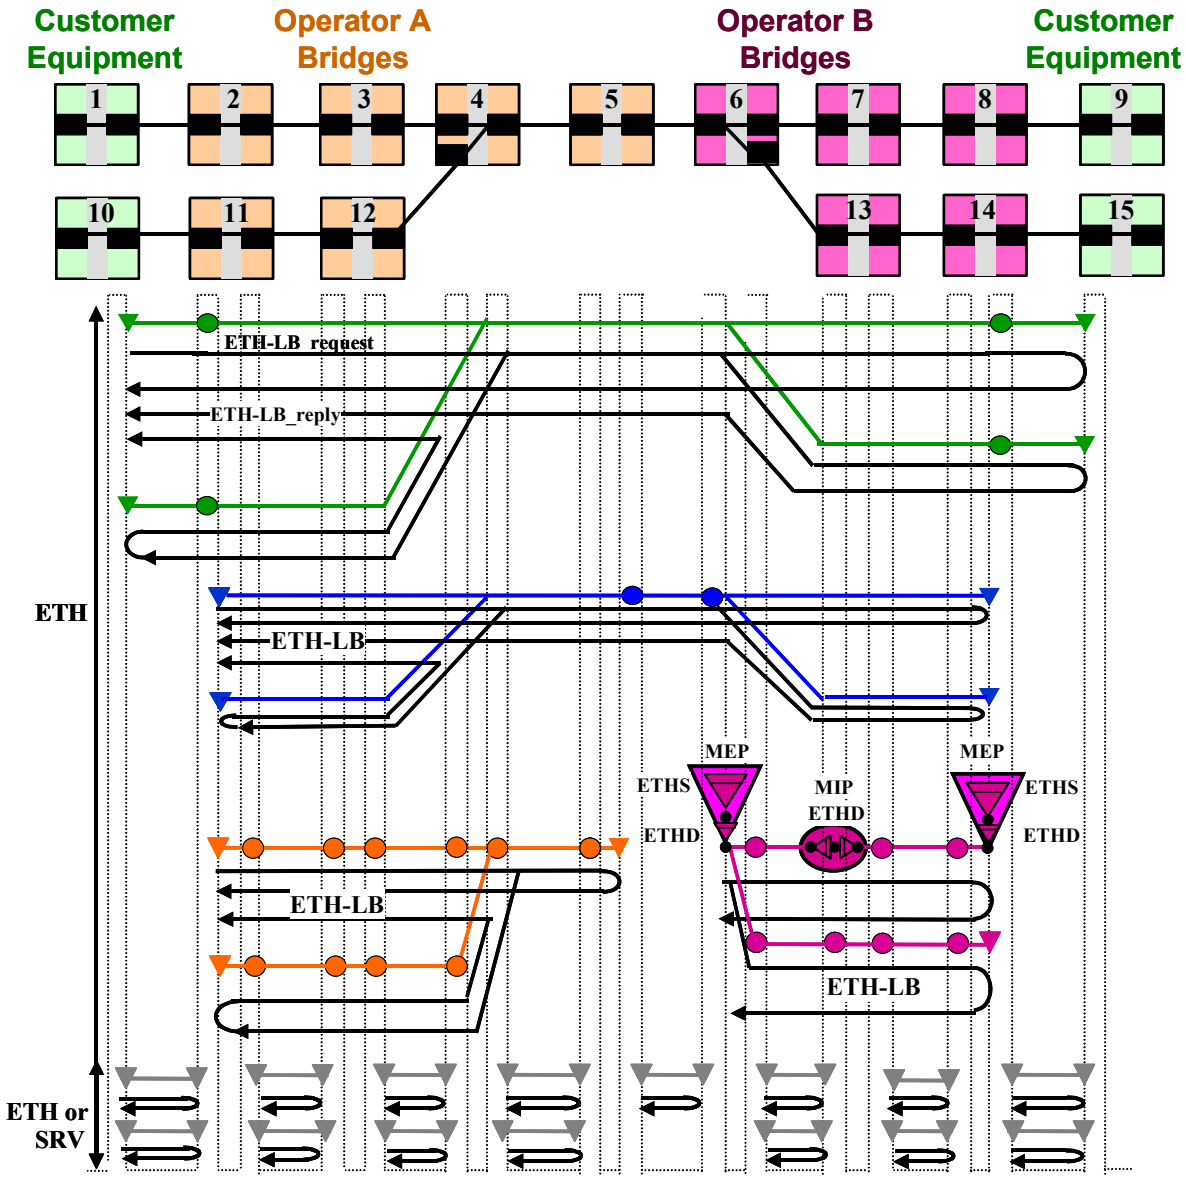

Figure VIII-2.4 Multicast Non-Intrusive loopback in ETH multipoint connection

VIII-3 ETH-LT

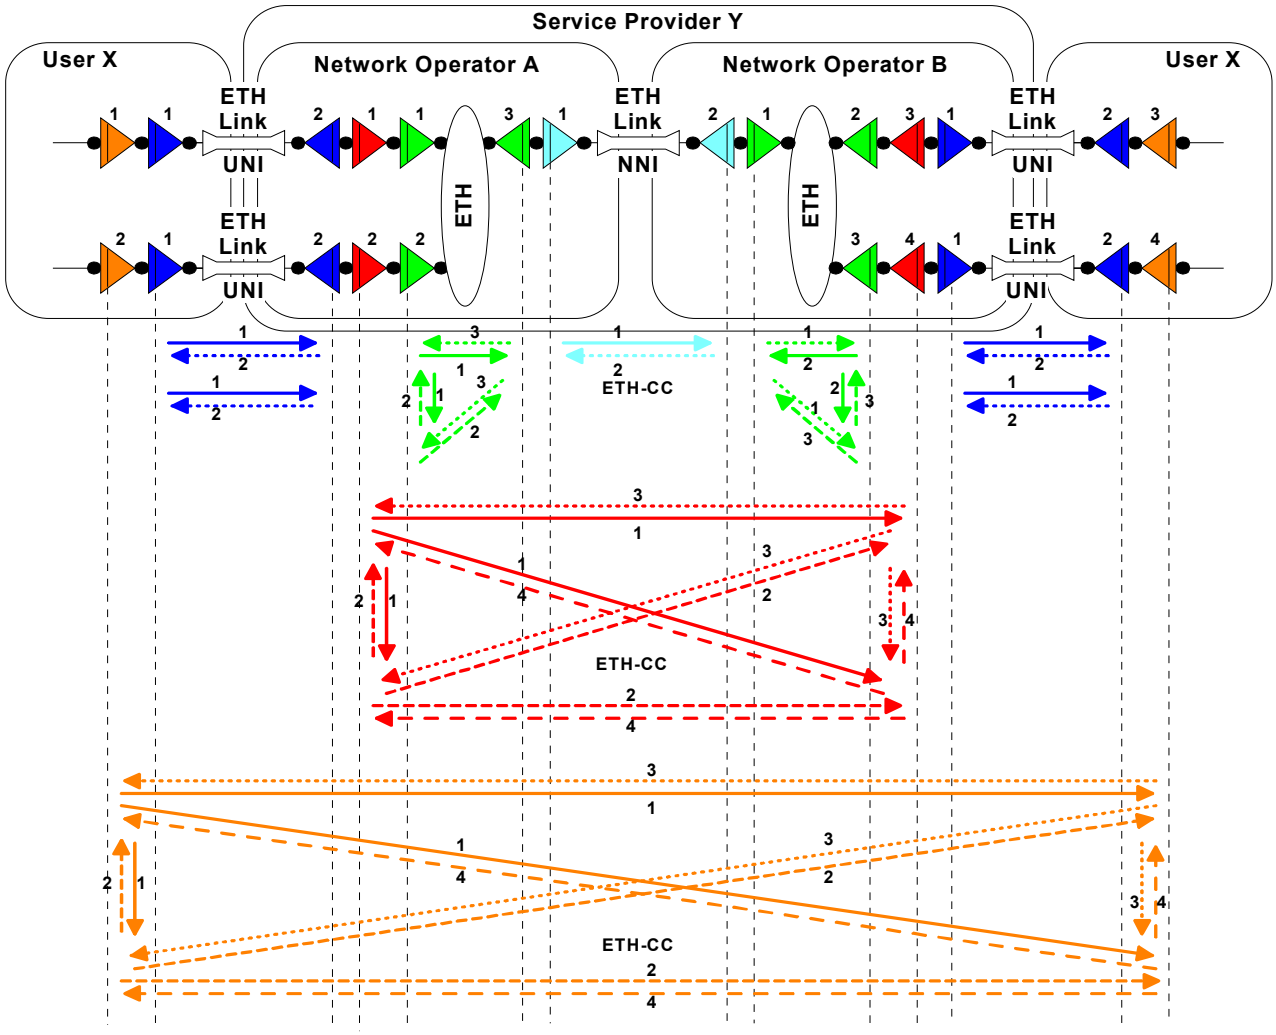

Figure VIII-3.1 –ETH-LT in multi-operator ETH multipoint connection

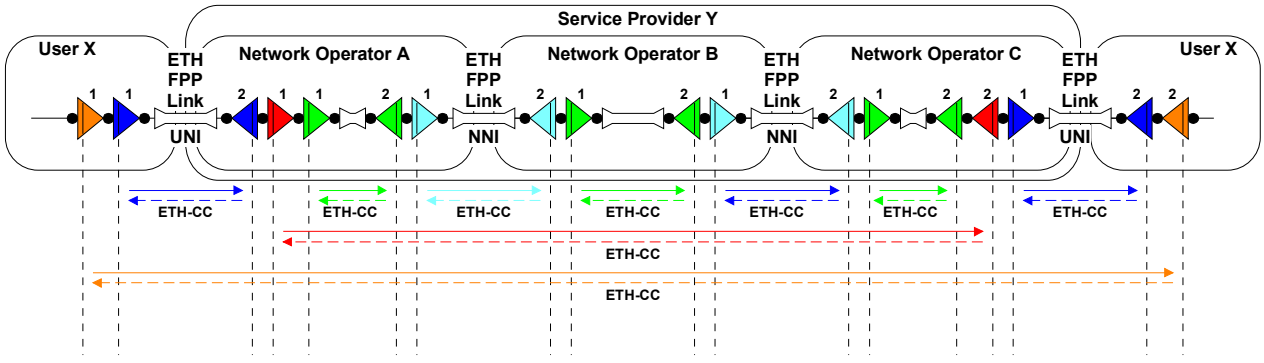

Figure VIII-3.2 - ETH-LT in multi-operator ETH point-to-point connection

NOTE: REPLACE ETH-CC WITH ETH-LT IN FIGURES 7.3-2 and 7.3-3

VIII-4 ETH-AIS

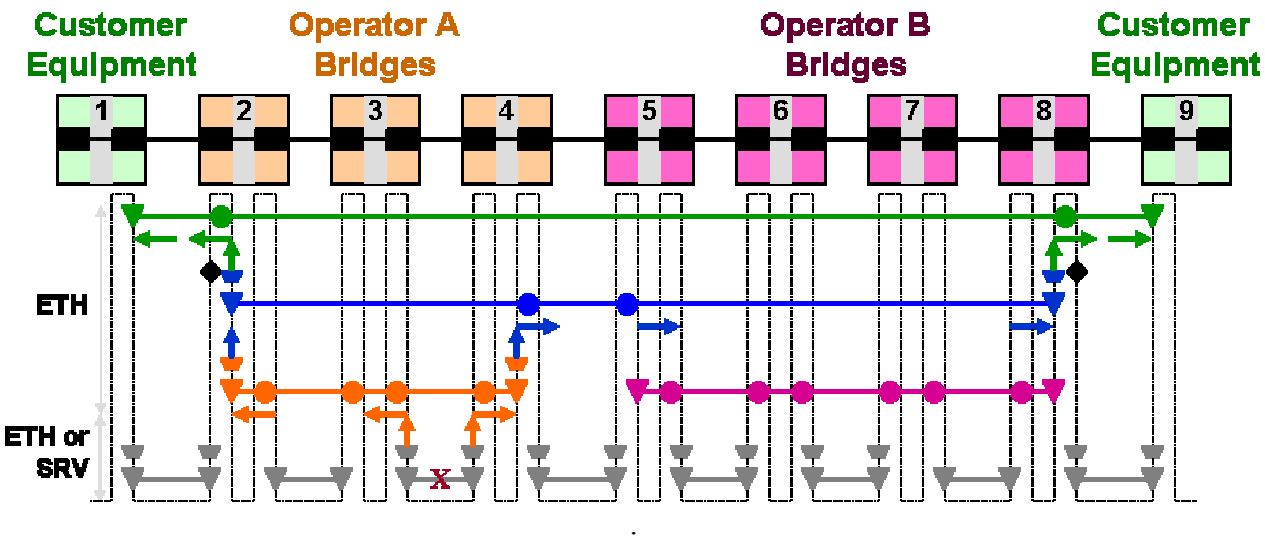

Figure VIII-4.1: ETH-AIS on p2p connection (failure in operator A domain)

- Black indicates ETH link MEs either realized as ETH MEs (sublayer monitoring) or as server layer MEs (inherent monitoring).
- Upon detection of this fault, the black MEPs corresponding to the failed link generate ETH_AIS which is adapted by black "AIS Adaptation" associated with black MEPs ("AIS Adaptation" is represented by the trapezoid entity). ETH-AIS (represented by orange arrows) is forwarded by orange MIPs towards orange MEPs corresponding to orange ME.
- Upon receiving ETH_AIS, the orange MEPs generate higher level ETH_AIS which is adapted by orange "AIS Adaptation". ETH-AIS (represented by blue arrows) which are forwarded by blue MIPs towards blue MEPs corresponding to blue ME.
- ETH_AIS promoted to blue ME remains transparent to the purple ME, where the purple ME is at a lower level compared to blue ME. In above figure, purple ME is shown as the same level as the orange ME.
- Similarly, upon receiving ETH_AIS, the blue MEPs generate higher level ETH_AIS which is adapted by blue "AIS Adaptation". ETH-AIS (represented by green arrows), which are forwarded by green MIPs towards green MEPs corresponding to green ME.

Note: "AIS Adaptation" is responsible to replicate server layer ETH-AIS to per client layer ME multiplexed over server layer ME.

It may be noted that in Figure VIII-4.1, the green and blue MEs correspond to service level MEs while orange and black MEs correspond to network and/or facility level MEs. Therefore, it is conceivable that a network level failure could trigger ETH_AIS along the service level ME.

VIII.4.1 ETH-AIS Trigger Condition

ETH-AIS triggered by the following conditions:

- ETY/SRV defect or Signal Fail conditions

- Loss of CC (LOC) conditions

[EDITOR'S NOTE-Dec2004] For ETY/SRV layer specific defect enumerations please refer to G.8021.

VIII.4.2 ETH-AIS Insertion and Termination Scenario

The following figures illustrate the ETH-AIS insertion and termination in a p2p connection for different fault locations e.g. UNI, operator A domain, inter-operator NNI, operator B domain, access link and/or extension link in access provider device scenario.

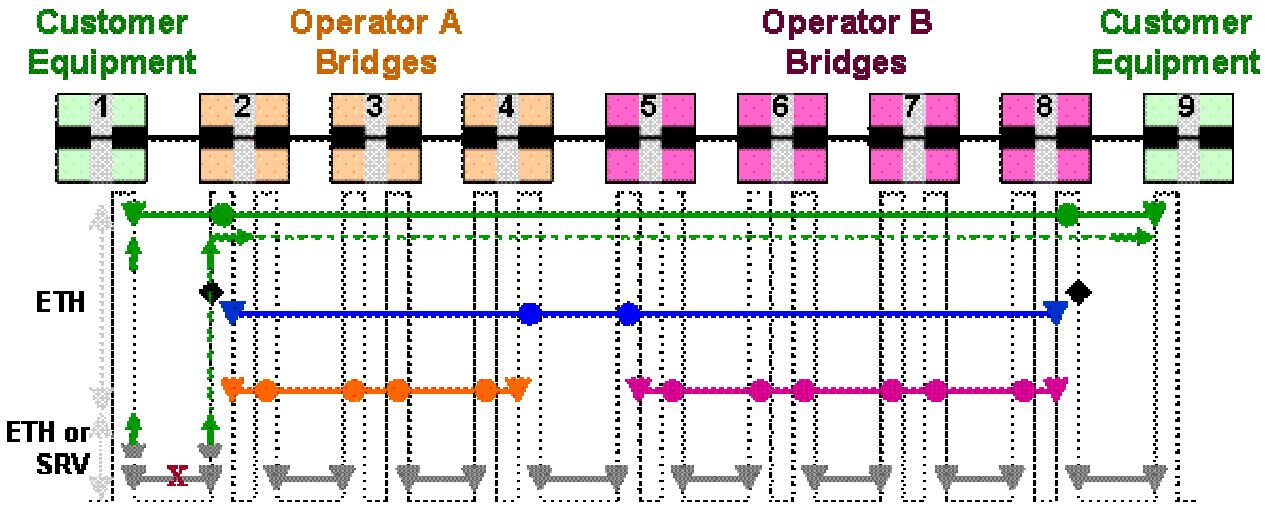

Figure VIII-4.2: ETH-AIS on p2p connection (failure on UNI)

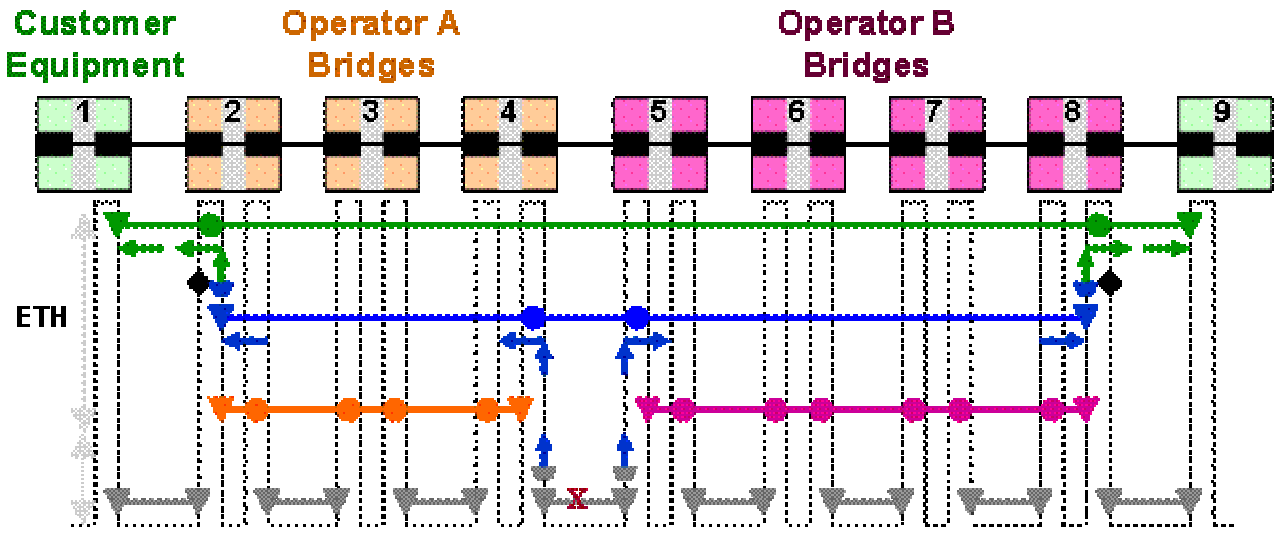

Figure VIII-4.3: ETH-AIS on p2p connection (failure on inter-provider NNI)

Interface ports with two or more MEP functions active will functionally terminate and re-generate ETH-AIS in each of the MEP Sink functions; as was mentioned in

Figure with the “AIS Adaptation”.

The termination and re-generation of ETH-AIS may increase the recovery time of the higher level MEs after the fault is repaired. Care should be taken with its processing definitions.

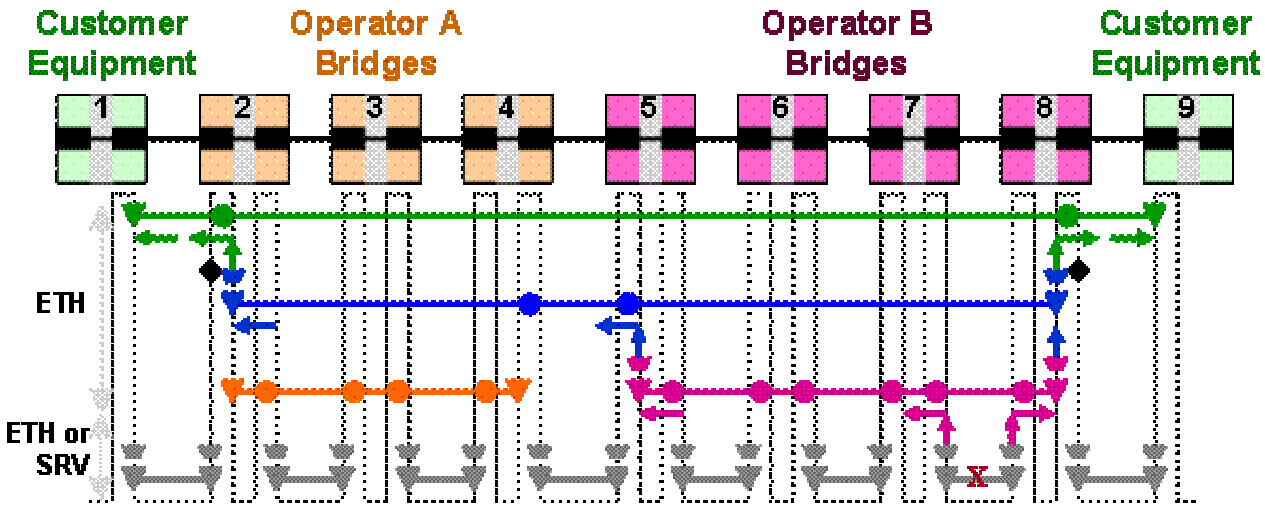

Figure VIII-4.3: ETH-AIS on p2p connection (failure in operator B domain)

When a FAILURE condition is detected in the final customer equipment's ingress port, the Server layer's MEP sink function will insert ETH-AIS, which will be terminated immediately in the next ETH layer MEP Sink function. If it is an ETHS_FT_Sk function then this function will also re-insert ETH-AIS to be forwarded through the customer domain towards the ETH flow termination. If it is an ETH_FT_Sk function (inside the LLC), then this MEP sink function will insert (if defined) client layer AIS.

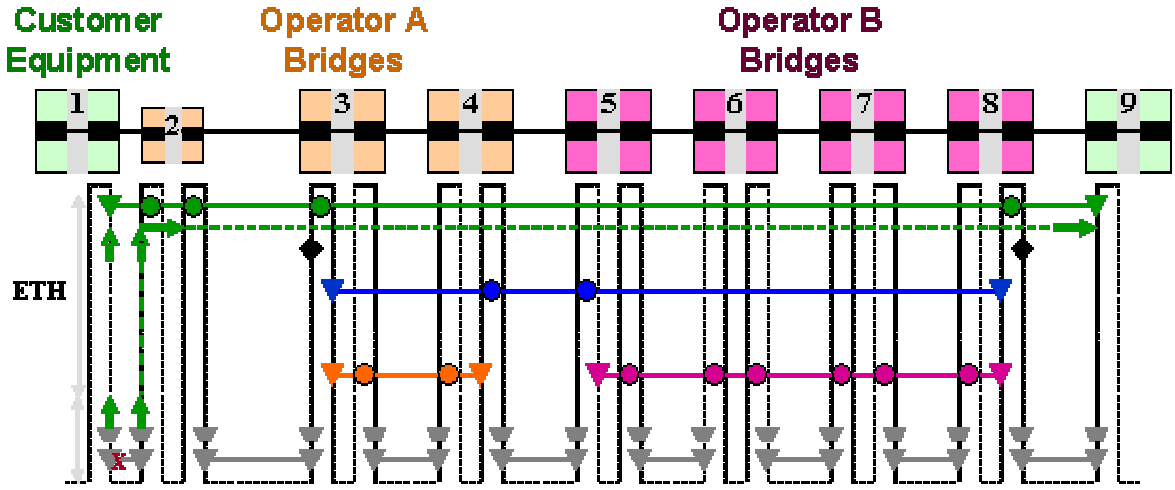

Figure VIII-4.4: ETH-AIS on p2p connections (failure on access link in access provider device)

EDITOR’S NOTE: TEXT NEEDS TO BE ADDED FOR ABOVE FIGURE. CONTRIBUTIONS ARE INVITED.

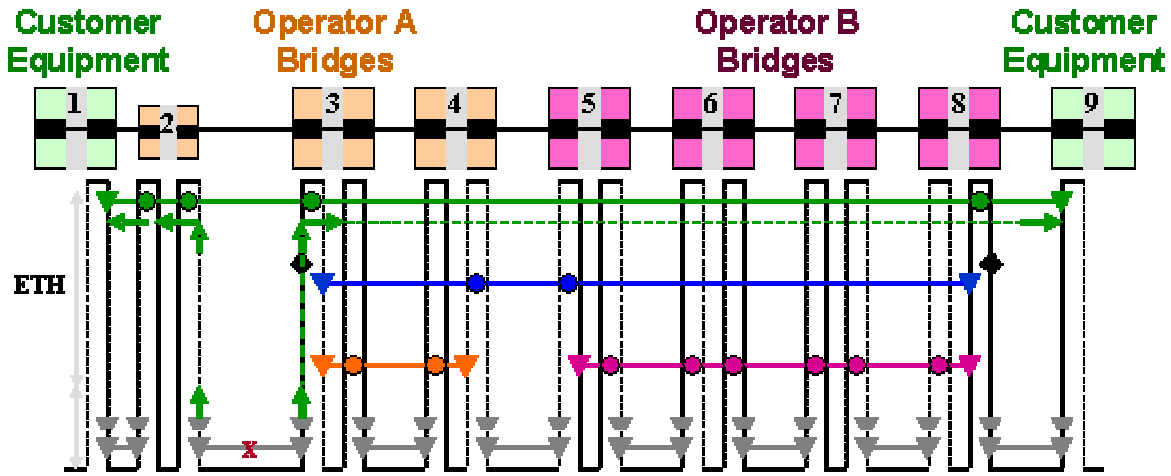

Figure VIII-4.5: ETH-AIS on p2p connections (failure on extension link in access provider device)

EDITOR'S NOTE: TEXT NEEDS TO BE ADDED FOR ABOVE FIGURE. CONTRIBUTIONS ARE INVITED.

VIII-5 ETH-RDI

APPENDIX IX: ME Level Assignment Considerations

[EDITOR'S NOTE-Dec2004] Specific scenarios to be added e.g. EPL case, Carrier in Carrier, Operator in Operator, with protection, Access Networks e.g. Internet Access with 2 VLANs and corresponding MEs. Some of these are expected to come from wd03.