International Telecommunication Union

ITU-T G.8032/Y.1344

(02/2012)

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

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

Packet over Transport aspects – Ethernet over Transport aspects

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

Internet protocol aspects – Transport

Ethernet ring protection switching

Recommendation ITU-T G.8032/Y.1344

ITU-T G-SERIES RECOMMENDATIONS

TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

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

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

Recommendation ITU-T G.8032/Y.1344

Ethernet ring protection switching

Summary

Recommendation ITU-T G.8032/Y.1344 defines the automatic protection switching (APS) protocol and protection switching mechanisms for ETH layer Ethernet ring topologies. Included are details pertaining to Ethernet ring protection characteristics, architectures and the ring APS (R-APS) protocol.

History

Edition	Recommendation	Approval	Study Group
1.0	ITU-T G.8032/Y.1344	2008-06-22	15
1.1	ITU-T G.8032/Y.1344 (2008) Amd. 1	2009-04-22	15
2.0	ITU-T G.8032/Y.1344	2010-03-09	15
2.1	ITU-T G.8032/Y.1344 (2010) Amd. 1	2010-06-11	15
2.2	ITU-T G.8032/Y.1344 (2010) Cor. 1	2010-10-07	15
2.3	ITU-T G.8032/Y.1344 (2010) Amd. 2	2011-02-25	15
3.0	ITU-T G.8032/Y.1344	2012-02-13	15

FOREWORD

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

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

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

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

NOTE

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

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

INTELLECTUAL PROPERTY RIGHTS

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

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

Table of Contents

Recommendation ITU-T G.8032/Y.1344

Ethernet ring protection switching

1 Scope

This Recommendation defines the automatic protection switching (APS) protocol and protection switching mechanisms for ETH layer Ethernet ring topologies. The protection protocol defined in this Recommendation enables protected point-to-point, point-to-multipoint and multipoint-to-multipoint connectivity within a ring or interconnected rings, called "multi-ring/ladder network" topology.

The ETH layer ring maps to the physical layer ring structure. Protection schemes for the other layers, including the ETY layer, are out of 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

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

[ITU-T G.806] Recommendation ITU-T G.806 (2012), *Characteristics of transport equipment – Description methodology and generic functionality*.

[ITU-T G.808.1] Recommendation ITU-T G.808.1 (2010), *Generic protection switching – Linear trail and subnetwork protection.*

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

[ITU-T G.870] Recommendation ITU-T G.870/Y.1352 (2012), *Terms and definitions for optical transport networks (OTN)*.

[ITU-T G.8001] Recommendation ITU-T G.8001/Y.1354 (2012), *Terms and definitions for Ethernet frames over transport.*

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

[ITU-T G.8021] Recommendation ITU-T G.8021/Y.1341 (2012), *Characteristics of Ethernet transport network equipment functional blocks*.

[ITU-T G.8013] Recommendation ITU-T G.8013/Y.1731 (2011), *OAM functions and mechanisms for Ethernet based networks*.

[IEEE 802.1Q] IEEE Std 802.1Q-2011, *IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks*.

3 Definitions

3.1 Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

3.1.1 Terms defined in [ITU-T G.805]:

a) adapted information
b) characteristic information
c) link
d) tandem connection
e) trail

3.1.2 Terms defined in [ITU-T G.806]:

a) defect
b) failure
c) server signal fail (SSF)
d) signal degrade (SD)
e) signal fail (SF)
f) trail signal fail (TSF)

3.1.3 Terms defined in [ITU-T G.808.1]:

a) transfer time (*Tt*):

3.1.4 Terms defined in [ITU-T G.809]:

a) adaptation
b) flow
c) layer network
d) network
e) port
f) transport
g) transport entity

3.1.5 Terms defined in [ITU-T G.870]:

a) APS protocol
b) hold-off time
c) non-revertive operation
d) protection
e) protected domain
f) revertive operation
g) signal
h) switch
i) switching time
j) transport entity:
    a) protection transport entity
    b) working transport entity
k) wait-to-restore time

3.1.6 Terms defined in [ITU-T G.8001]:

a) maintenance entity (ME)

b) maintenance entity group (MEG)

c) Ethernet ring

d) Ethernet ring node

e) ERP instance

f) interconnection node

g) major ring

h) R-APS virtual channel

i) ring MEL

j) ring protection link (RPL)

k) RPL neighbour node

l) RPL owner node

m) sub-ring

n) sub-ring Link

o) wait to block timer

3.1.7 Terms defined in [ITU-T G.8010]:

a) Ethernet characteristic information (ETH_CI)

b) Ethernet flow point (ETH_FP)

3.1.8 Terms defined and described in [ITU-T G.8010] and [ITU-T G.8013]:

a) maintenance entity group end point (MEP)

b) maintenance entity group level (MEL)

3.1.9 Terms described in [ITU-T G.8021]:

a) Ethernet connection function (ETH_C)

b) Ethernet MAC characteristic information server signal fail (ETH_CI_SSF)

c) Ethernet flow forwarding function (ETH_FF)

d) ETH to ETH multiplexing adaptation function (ETHx/ETH-m_A)

e) ETHDi/ETH adaptation function (ETHDi/ETH_A)

3.2 Terms defined in this Recommendation

None.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

AI Adapted Information

APS Automatic Protection Switching

BPR Blocked Port Reference

CCM Continuity Check Message

CI Characteristic Information

DNF Do Not Flush

ERP	Ethernet Ring Protection
ETH	Ethernet layer network
FDB	Filtering Database
FS	Forced Switch
ID	Identification
MEG	Maintenance Entity Group
MEL	Maintenance Entity group Level
MEP	Maintenance entity group End Point
MI	Management Information
MIP	Maintenance entity group Intermediate Point
MS	Manual Switch
NR	No Request
OAM	Operations, Administration and Maintenance
OUI	Organizationally Unique Identifier
PDU	Protocol Data Unit
R-APS	Ring APS
RB	RPL Blocked
RPL	Ring Protection Link
RSTP	Rapid Spanning Tree Protocol
SD	Signal Degrade
SF	Signal Fail
STP	Spanning Tree Protocol
TCM	Tandem Connection Monitoring
VID	VLAN Identifier
VLAN	Virtual LAN
VPLS	Virtual Private LAN Service
WTB	Wait To Block
WTR	Wait To Restore

5 Conventions

5.1 Representation of octets

Octets are represented as defined in [ITU-T G.8013].

When consecutive octets are used to represent a binary number, the lower octet number has the most significant value. The bits in an octet are numbered from 1 to 8, where 1 is the least significant bit and 8 is the most significant bit.

6 Introduction

This Recommendation specifies protection switching mechanisms and a protocol for Ethernet layer network (ETH) rings. Ethernet rings can provide wide-area multipoint connectivity more economically due to their reduced number of links. The mechanisms and protocol defined in this Recommendation achieve highly reliable and stable protection; and never form loops, which would fatally affect network operation and service availability.

Each Ethernet ring node is connected to adjacent Ethernet ring nodes participating in the same Ethernet ring, using two independent links. A ring link is bounded by two adjacent Ethernet ring nodes and a port for a ring link is called a ring port. The minimum number of Ethernet ring nodes in an Ethernet ring is two.

The fundamentals of this ring protection switching architecture are:

a) the principle of loop avoidance; and

b) the utilization of learning, forwarding and filtering database (FDB) mechanisms defined in the Ethernet flow forwarding function (ETH_FF).

Loop avoidance in an Ethernet ring is achieved by guaranteeing that, at any time, traffic may flow on all but one of the ring links. This particular link is called the ring protection link (RPL) and under normal conditions this ring link is blocked, i.e., not used for service traffic. One designated Ethernet ring node, the RPL owner node, is responsible to block traffic at one end of the RPL. Under an Ethernet ring failure condition, the RPL owner node is responsible for unblocking its end of the RPL, unless the RPL has failed, allowing the RPL to be used for traffic. The other Ethernet ring node adjacent to the RPL, the RPL neighbour node, may also participate in blocking or unblocking its end of the RPL.

The event of an Ethernet ring failure results in protection switching of the traffic. This is achieved under the control of the ETH_FF functions on all Ethernet ring nodes.

An APS protocol is used to coordinate the protection actions over the ring.

The Ethernet rings could support a multi-ring/ladder network that consists of conjoined Ethernet rings by one or more interconnection points. The protection switching mechanisms and protocol defined in this Recommendation shall be applicable for a multi-ring/ladder network, if the following principles are adhered to:

a) R-APS channels are not shared across Ethernet ring interconnections;

b) on each ring port, each traffic channel and each R-APS channel are controlled (e.g., for blocking or flushing) by the Ethernet ring protection control process (ERP control process) of only one Ethernet ring;

c) each major ring or sub-ring must have its own RPL.

7 Ring protection characteristics

7.1 Monitoring methods and conditions

Ring protection switching occurs based on the detection of defects on the transport entity of each ring link. The defects are defined within the equipment Recommendation [ITU-T G.8021]. For the purpose of the protection switching process, a transport entity, within the protected domain, has a condition of either failed (i.e., signal fail (SF)) or non-failed (OK).

Ethernet ring protection may adopt any of the following monitoring methods:

Inherent – The fault condition status of each ring link connection is derived from the status of the underlying server layer trail.

Sub-layer – Each ring link is monitored using tandem connection monitoring (TCM).

Test trail – Defects are detected using an extra test trail, i.e., an extra test trail is set up along each ring link.

The protection switching is agnostic to the monitoring method used, as long as it can be given (OK or SF) information regarding the transport entity of each ring link.

7.2 Ethernet traffic and bandwidth consideration

It is desirable that ring bandwidth accommodates all traffic that is protected, regardless of the ring protection switching state. Being different from linear protection, ERP does not separate working and protection transport entities, but reconfigures the transport entity during protection switching. Therefore care should be taken that ring link capacity can continue to support all ring APS (R-APS) and service traffic that is protected after protection switching.

7.3 Ethernet ring protection switching performance

In an Ethernet ring, without congestion, with all Ethernet ring nodes in the idle state (i.e., no detected failure, no active automatic or external command and receiving only "NR, RB" R-APS messages), with less than 1200 km of ring fibre circumference and fewer than 16 Ethernet ring nodes, the switch completion time (transfer time as defined in [ITU-T G.808.1]) for a failure on a ring link shall be less than 50 ms. On Ethernet rings under all other conditions, the switch completion time may exceed 50 ms (the specific interval is under study), to allow time to negotiate and accommodate coexisting APS requests. In case of interconnection of sub-rings with R-APS virtual channel to a major ring, the R-APS messages of the sub-ring that are inserted into the R-APS virtual channel take on performance characteristics (e.g., delay, jitter, packet drop probability, etc.) of the ring links and Ethernet ring nodes it crosses over the interconnected Ethernet ring. In this case, if the R-APS channel and R-APS virtual channel exceed the number of Ethernet ring nodes or fibre circumference defined above, the protection switching of the sub-ring may exceed 50 ms.

NOTE – The inclusion of the completion of FDB flush operation within the transfer time is for further study.

8 Ring protection conditions and commands

This Recommendation supports the following conditions of the Ethernet ring:

Signal fail (SF) – When an SF condition is detected on a ring link and it is determined to be a "stable" failure, Ethernet ring nodes adjacent to the failed ring link initiate the protection switching mechanism described in this Recommendation.

No request (NR) – The condition when no local protection switching requests are active.

The following administrative commands are supported (as possible values for ETH_C_MI_RAPS_ExtCMD):

Forced switch (FS) – This command forces a block on the ring port where the command is issued.

Manual switch (MS) – In the absence of a failure or FS, this command forces a block on the ring port where the command is issued.

Clear – The Clear command, at the Ethernet ring node, is used for the following operations:

a) Clearing an active local administrative command (e.g., Forced switch or Manual switch).

b) Triggering reversion before the WTR or WTB timer expires in case of revertive operation.

c) Triggering reversion in case of non-revertive operation.

The following commands are for further study:

Lockout of protection – This command disables the protection group.

Replace the RPL – This command moves the RPL by blocking a different ring link and unblocking the RPL permanently.

Exercise signal – Exercise of the R-APS protocol. The signal is chosen so as not to modify the position of the blocked ring port.

9 Ring protection architectures

In the ring protection architecture defined in this Recommendation, protection switching is performed at all Ethernet ring nodes.

The ring protection architecture relies on the existence of an APS protocol to coordinate ring protection actions around an Ethernet ring.

9.1 Revertive and non-revertive switching

In revertive operation, after the condition(s) causing a switch has cleared, the traffic channel is restored to the working transport entity, i.e., blocked on the RPL. In the case of clearing of a defect, the traffic channel reverts after the expiry of a WTR timer (see clause 10.1.4), which is used to avoid toggling protection states in case of intermittent defects.

In non-revertive operation, the traffic channel continues to use the RPL, if it has not failed, after a switch condition has cleared.

Since in Ethernet ring protection the working transport entity resources may be more optimized, in some cases it is desirable to revert to this working transport entity once all ring links are available. This is performed at the expense of an additional traffic interruption.

In some cases, there may be no advantage to revert to the working transport entities immediately. In this case, a second traffic interruption is avoided by not reverting protection switching.

9.2 Protection switching triggers

Protection switching shall be performed when:

a) SF is declared on one of the ring links and the detected SF condition has a higher priority than any other local request or far-end request; or

b) the received R-APS message requests to switch and it has a higher priority than any other local request; or

c) initiated by operator control (e.g., Forced switch, Manual switch) if it has a higher priority than any other local request or far-end request.

9.2.1 Signal fail declaration conditions

SF is declared when an ETH trail signal fail condition is detected. ETH trail signal fail is specified in [ITU-T G.8021].

9.3 Protection switching models on a single Ethernet ring

Figure 9-1 depicts an example of the Ethernet ring protection switching model defined in this Recommendation. Other network scenarios are permissible. In this example, four Ethernet ring nodes are depicted.

If the Ethernet ring is in its normal condition, one Ethernet ring node adjacent to the RPL is configured as the RPL owner node and, in this example, another Ethernet ring node adjacent to the RPL is configured as the RPL neighbour node. Both end nodes of the RPL are responsible for blocking the transmission and reception of traffic over the RPL when there is no request on the Ethernet ring.

In Figure 9-1 Ethernet ring node D is the RPL owner node and Ethernet ring node A is the RPL neighbour node. Both Ethernet ring nodes are responsible for blocking the traffic channel on the RPL. Figure 9-1 presents the case when no failure is present on any ring link. In this case, the ETH characteristic information (ETH_CI) traffic may be transferred over both ring links of any Ethernet ring node, except for the RPL on the Ethernet ring nodes where the RPL is blocked. In this figure, the traffic channel is illustrated as arrows being transmitted and received from the ring links. In subsequent figures only the ETH_FF function for a single VLAN is represented.

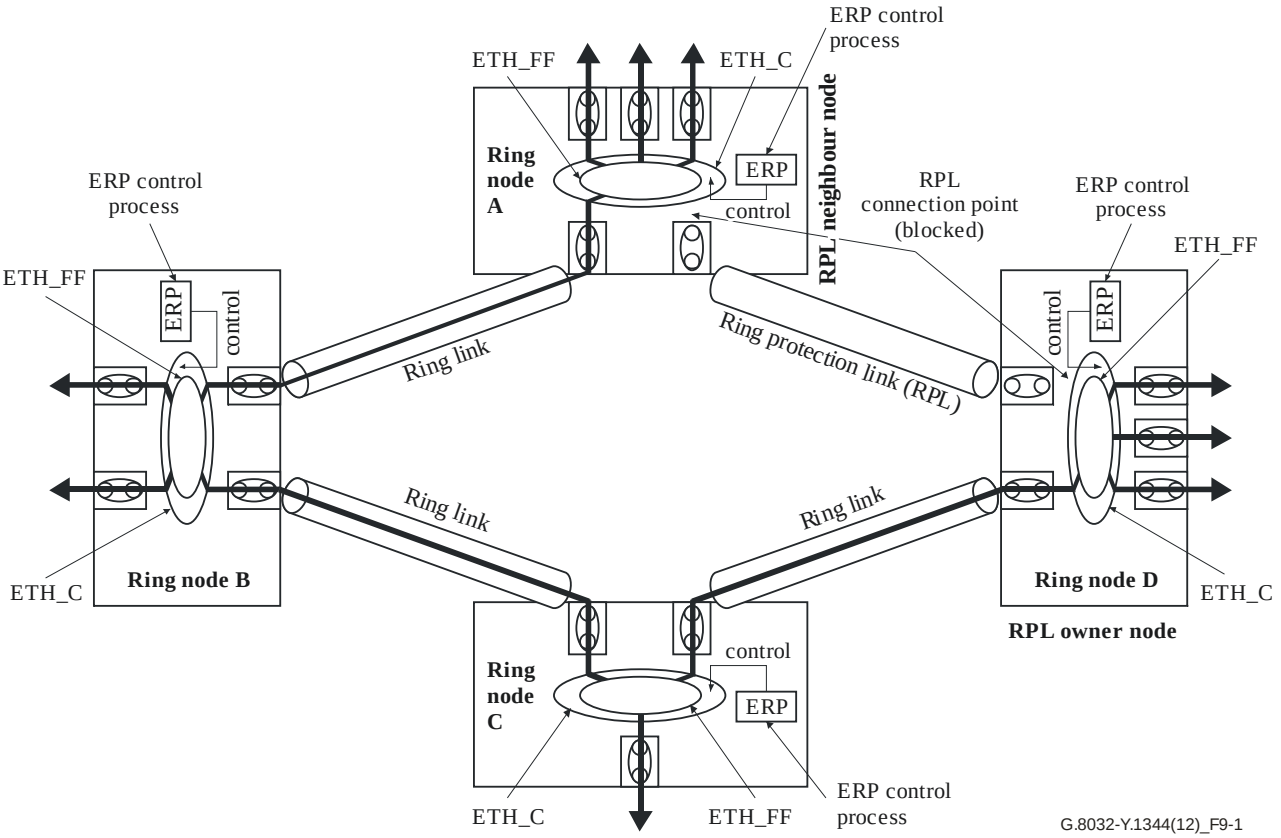

Figure 9-1 – Ethernet ring protection switching architecture – normal condition (single Ethernet ring)

Figure 9-2 illustrates a situation where a protection switch has occurred due to an SF condition on one ring link. In this case, the traffic channel is blocked bidirectionally on the ports where the failure is detected and bidirectionally unblocked at the RPL connection point.

In revertive operation, when the failure is recovered, the traffic channel resumes the use of the recovered ring link only after the traffic channel has been blocked on the RPL. On the other hand, in non-revertive operation, the traffic channel remains blocked on the recovered ring link and unblocked on the RPL even if the failure is recovered.

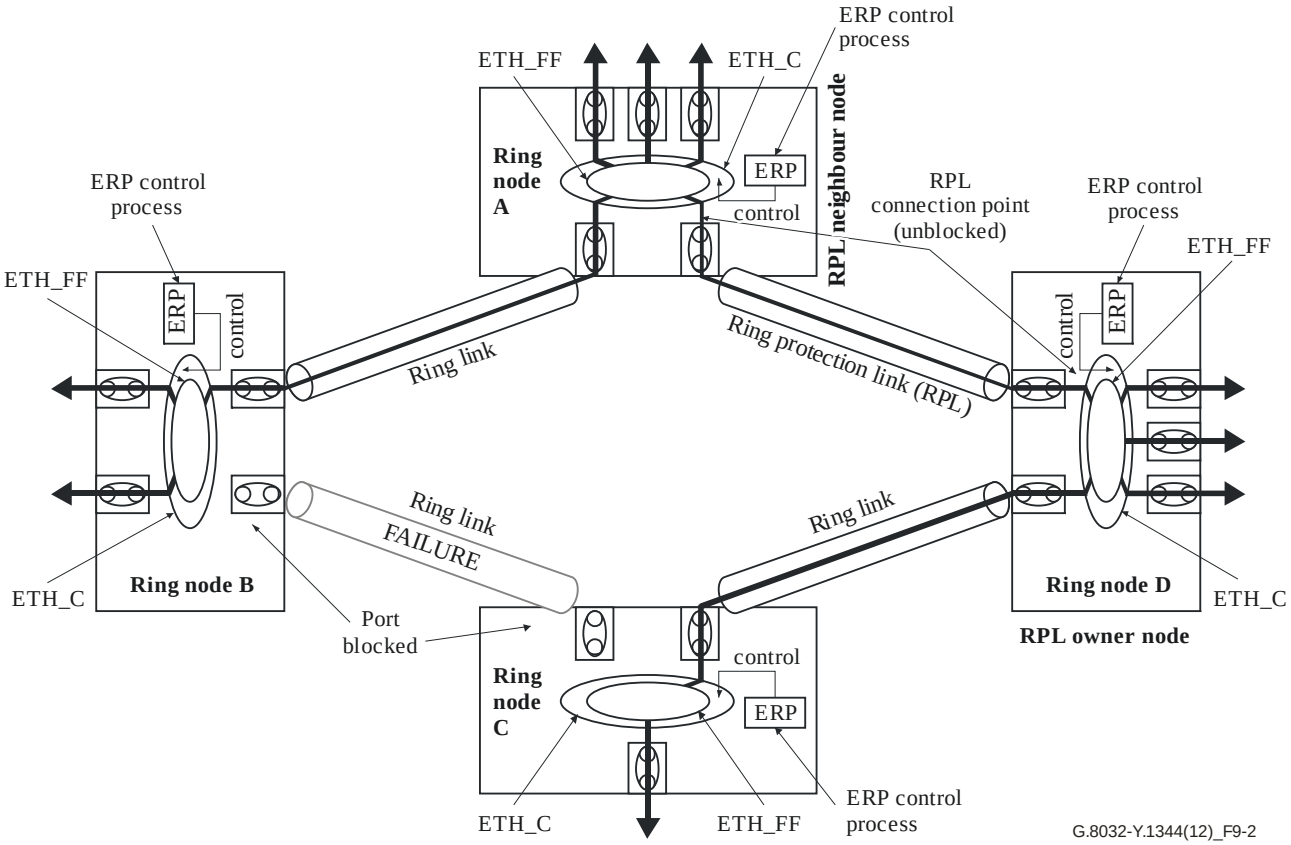

Figure 9-2 – Ethernet ring protection switching architecture – signal fail condition on one ring link (single Ethernet ring)

A model of the functionality of an Ethernet ring node is presented in Figures 9-3 and 9-4.

The ERP control process is instantiated to protect normal traffic over an Ethernet ring. Each instantiated ETH_FF function determines the specific output Ethernet flow point (ETH_FP) over which the ETH_CI is transferred. The ETH_CI may be forwarded over any ETH_FP corresponding to ring links or to non-ring links.

The ERP control process controls the ETH_FF function to perform actions such as disabling forwarding over any ETH_FP corresponding to blocked ring links and flushing the FDB.

As an example, the ring links of each Ethernet ring node may be monitored by individually exchanging continuity check messages (CCM) defined in [ITU-T G.8013] on the maintenance entity group end points (MEPs) illustrated in Figure 9-3.

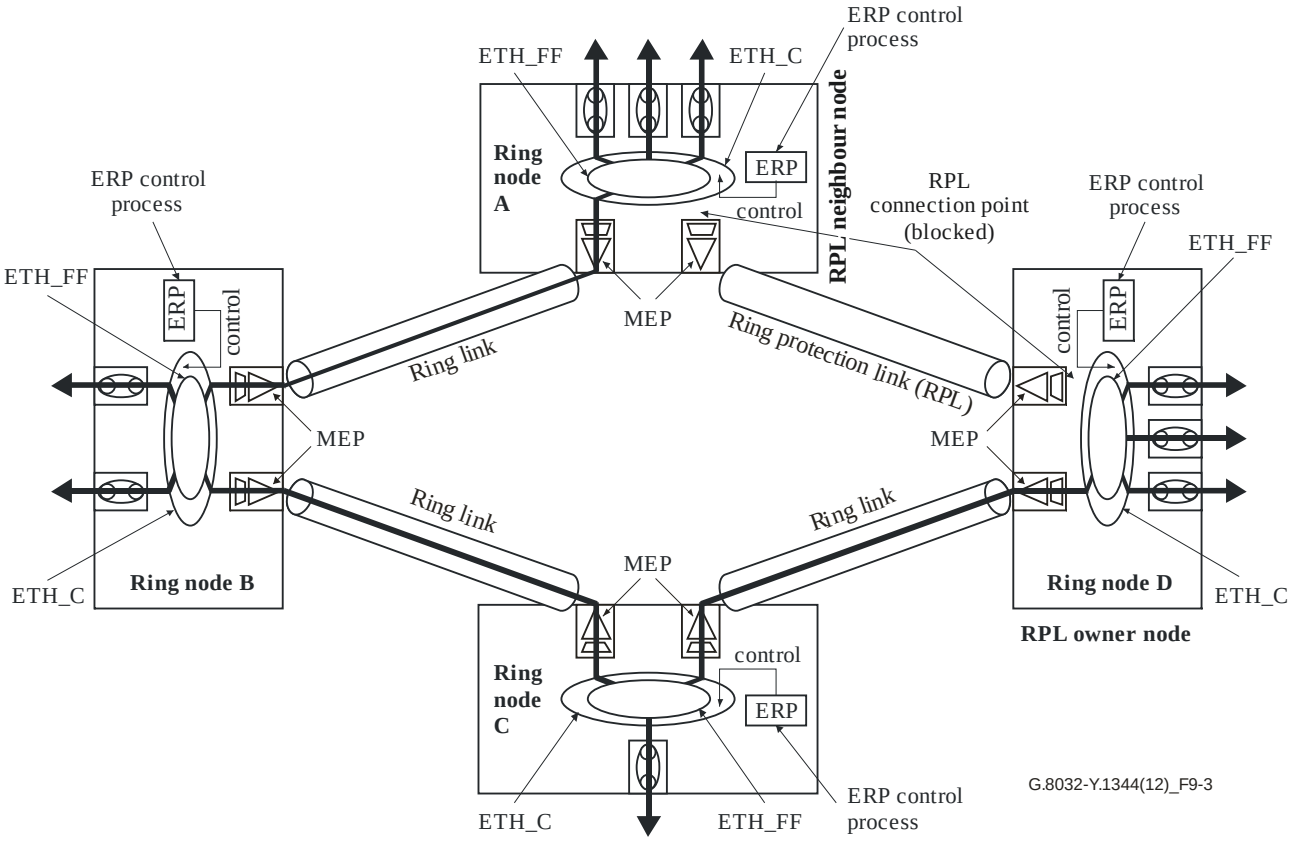

Figure 9-3 – MEPs in Ethernet ring protection switching architecture

Figure 9-4 represents the model of an Ethernet ring node. MEPs represented on each ring port are used for monitoring the ring link.

If an MEP detects a defect, which contributes to an SF defect condition, it informs the ERP control process that a failure condition has been detected. An ERP control function uses the ETH_CI_SSF information, forwarded from the ETHx/ETH-m_A_Sk, to assert the SF condition of the ring link.

The Ethernet ring protection switching mechanism requires the R-APS protocol to coordinate the switching behaviour among all Ethernet ring nodes. The R-APS protocol communication is performed using R-APS messages. R-APS messages are transmitted and received at an ERP control process. The ETHDi/ETH_A function in [ITU-T G.8021] extracts ETH_CI_RAPS information from a received R-APS message and sends the ETH_CI_RAPS information to the ERP control process. A received R-APS message is also forwarded to the ETH_FF. The ETHDi/ETH_A function also generates R-APS messages using the ETH_CI_RAPS information received from the ERP control process.

R-APS messages are forwarded using an ETH_FF function for R-APS traffic, represented in Figure 9-4 as R-APS_FF. Traffic, other than R-APS traffic, is forwarded by use of other ETH_FF functions, represented in Figure 9-4 as Service_FF. R-APS messages use a dedicated VLAN. Only one traffic VLAN is depicted in Figure 9-4. More traffic VLANs could be supported using multiple Service_FFs.

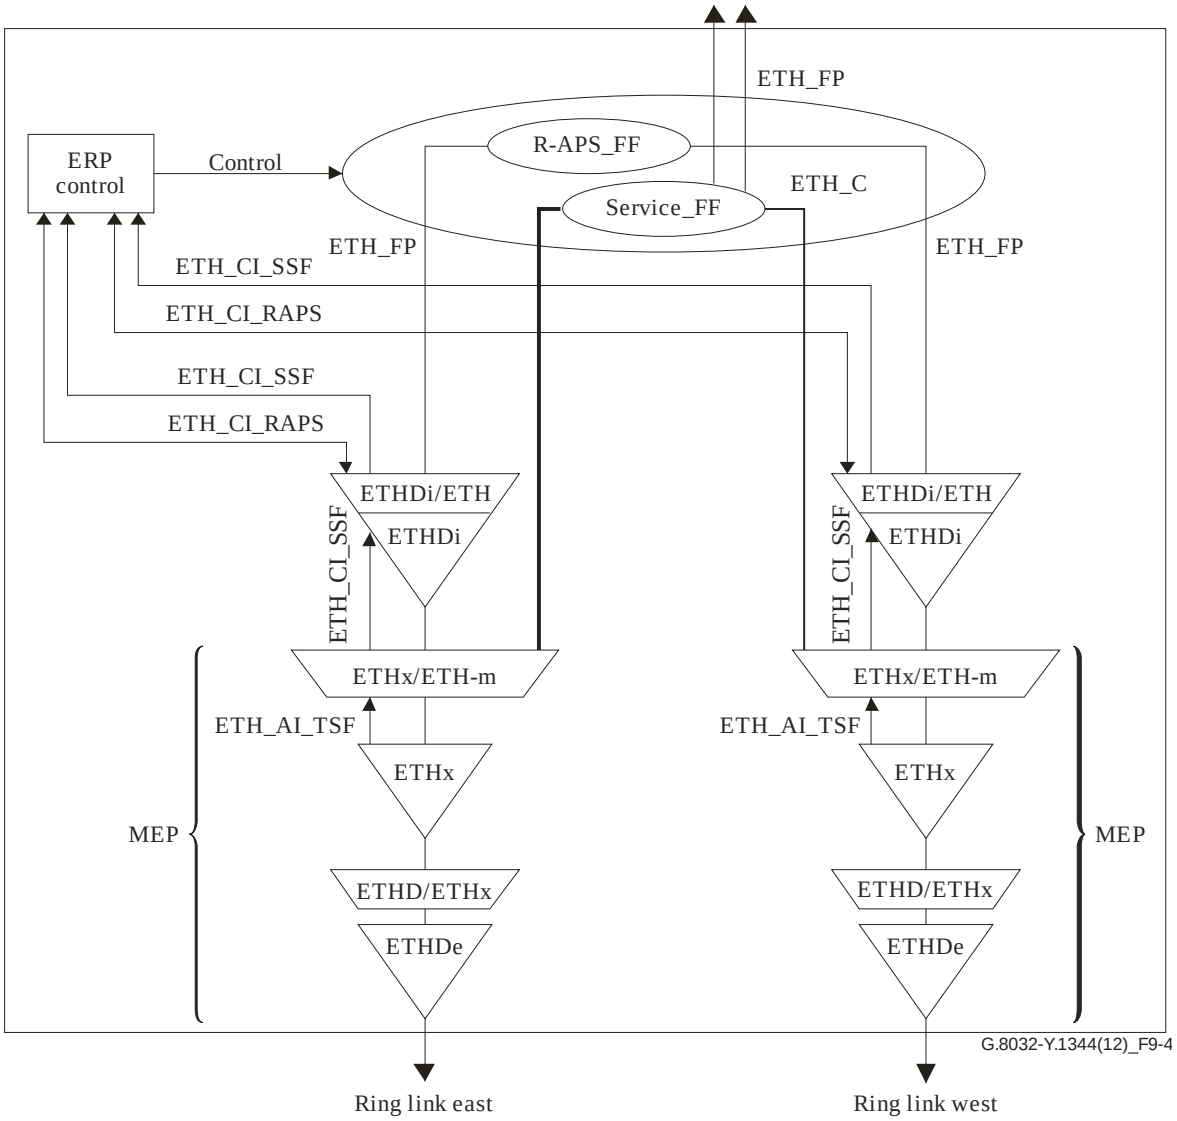

Figure 9-4 – MEPs and R-APS insertion function in Ethernet ring node (normal Ethernet ring node)

9.4 Traffic channel blocking

Blocking traffic is supported by excluding the connection point from the ETH_FF functions for the one or more VLAN IDs of the traffic channel controlled by the ERP instance. This is equivalent to VID filtering as defined in clause 8.13.10 of [IEEE 802.1Q]. This results in blocking the transmission and reception of traffic on one ring port. Each ERP instance shall only block or unblock the VLAN IDs of the traffic channels of the set of VLANs assigned for protection by that ERP instance.

9.5 R-APS channel blocking

R-APS channel VLAN traffic forwarding is always blocked at the same ring ports where the traffic channel is blocked, except on sub-rings without an R-APS virtual channel (see clause 9.7.2). It is supported by excluding the connection point from the ETH_FF function for the VLAN ID of the R-APS traffic and is equivalent to performing VID filtering as defined in clause 8.13.10 of [IEEE 802.1Q]. This:

a) only prevents R-APS messages received at one ring port from being forwarded to the other ring port;

b) does not prevent R-APS messages, locally generated at the ERP control process, from being transmitted over both ring ports;

c) allows R-APS messages received at each ring port to be delivered to the ERP control process. The ERP control process shall discard all received R-APS messages with a ring ID that does not match the configured ring ID of the current ERP instance.

Each ERP instance shall only block or unblock its R-APS channel. This is guaranteed by excluding the connection point from the ETH_FF for the VLAN ID of the R-APS traffic and is equivalent to performing group address filtering as defined in [IEEE 802.1Q]

On sub-rings without an R-APS virtual channel, the R-APS channel is never blocked on any of its sub-ring nodes. However, in this case, the R-APS channel is terminated at the interconnection nodes.

9.6 FDB flush

An FDB flush consists of removing MAC addresses learned on the ring ports of the protected Ethernet ring from the Ethernet ring node's filtering database.

Each ERP instance may flush only the FDB for the VLAN IDs of the traffic channels of the set of VLANs it is assigned to protect.

9.7 Ethernet ring protection switching models for interconnection

The Ethernet ring protection switching model for interconnection supports multi-ring/ladder topologies such as those illustrated in Appendix II.

Figure 9-5 depicts an example of the model on a multi-ring/ladder network defined in this Recommendation. If the multi-ring/ladder network is in its normal condition, the RPL owner node of each Ethernet ring blocks the transmission and reception of traffic over the RPL for that Ethernet ring. Figure 9-5 presents the configuration when no failure is present on any ring link.

In Figure 9-5 there are two interconnected Ethernet rings. Ethernet ring ERP1 is composed of Ethernet ring nodes A, B, C and D and the ring links between these Ethernet ring nodes. Ethernet ring ERP2 is composed of Ethernet ring nodes C, D, E and F and the ring links C-to-F, F-to-E, E-to-D. The ring link between D and C is used for traffic of Ethernet rings ERP1 and ERP2. On their own ERP2 ring links do not form a closed loop. A closed loop may be formed by the ring links of ERP2 and the ring link between interconnection nodes that is controlled by ERP1. ERP2 is a sub-ring. Ethernet ring node A is the RPL owner node for ERP1. Ethernet ring node E is the RPL owner node for ERP2. These Ethernet ring nodes (A and E) are responsible for blocking the traffic channel on the RPL for ERP1 and ERP2 respectively. There is no restriction on which ring link on an Ethernet ring may be set as RPL. For example the RPL of ERP1 could be set as the link between Ethernet ring nodes C and D.

Ethernet ring nodes C and D that are common to both ERP1 and ERP2, are called the interconnection nodes. The ring links between the interconnection nodes are controlled and protected by the Ethernet ring it belongs to. In the example of Figure 9-5, the ring link between Ethernet ring nodes C and D is part of ERP1 and as such, controlled and protected by ERP1. The ETH characteristic information (ETH_CI) traffic corresponding to the traffic channel may be transferred over a common ETH_C function for ERP1 and ERP2 through the interconnection nodes C and D. Interconnection nodes C and D have separate ERP control processes for each Ethernet ring.

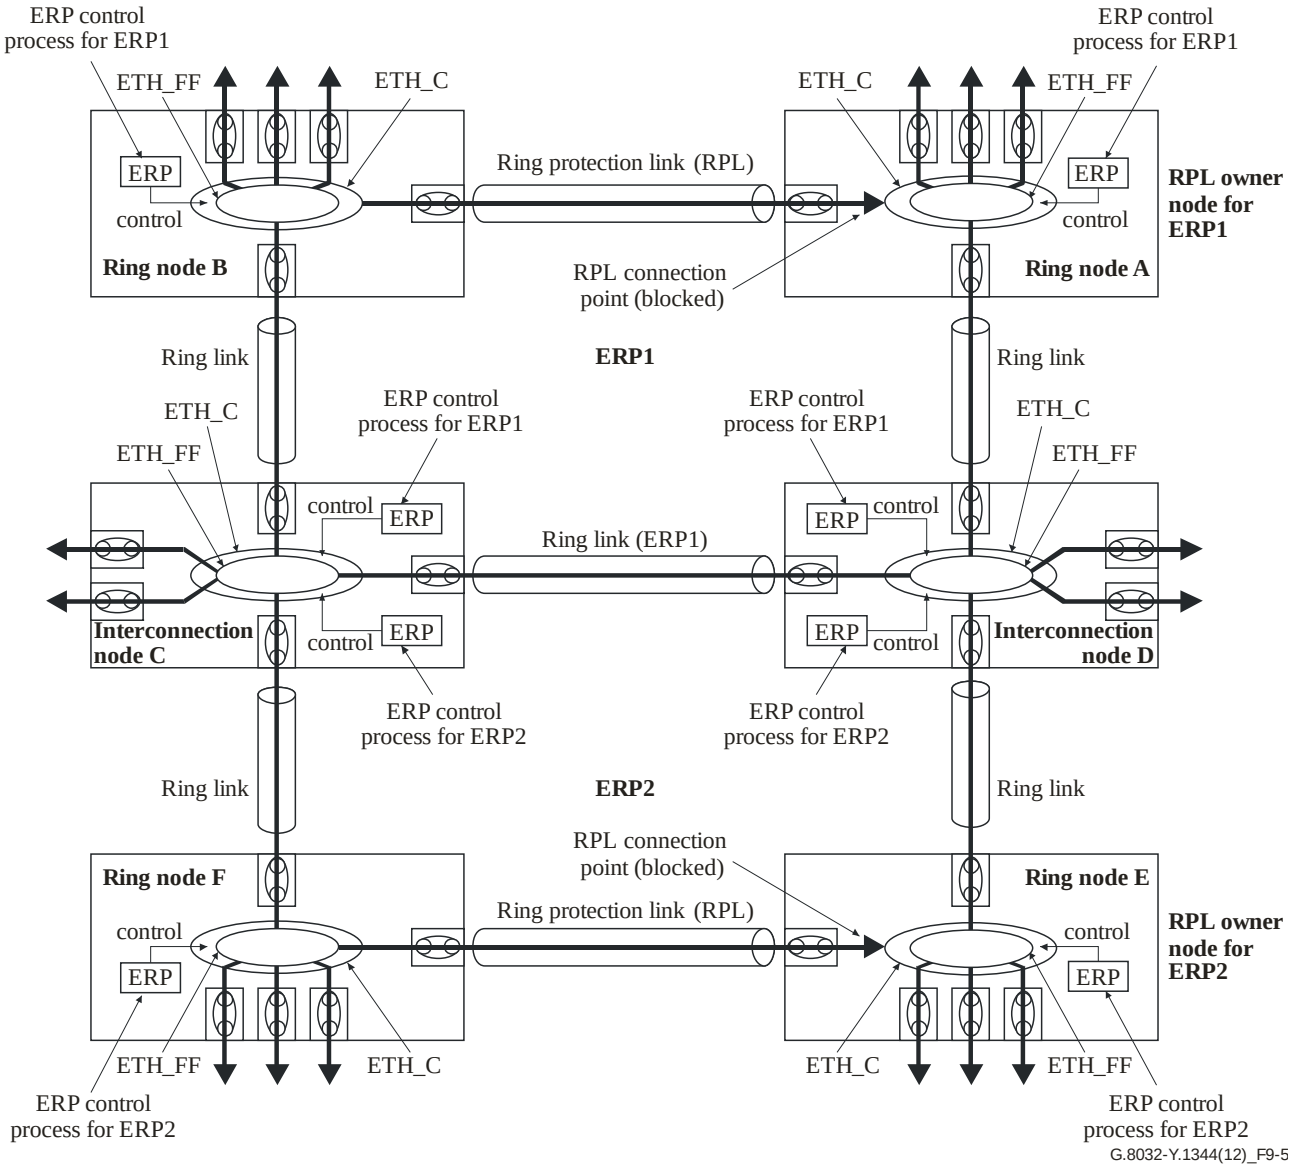

Figure 9-5 – Ethernet ring interconnection architecture – normal condition (multi-ring/ladder network)

Figure 9-6 illustrates a situation where protection switching has occurred due to an SF condition on the ring link between interconnection nodes C and D. The failure of this ring link triggers protection only on the Ethernet ring it belongs to, in this case ERP1. The traffic and R-APS channels are blocked bidirectionally on the ports where the failure is detected and bidirectionally unblocked at the RPL connection point on ERP1. The traffic channels remain bidirectionally blocked at the RPL connection point on ERP2. This prevents the formation of a loop.

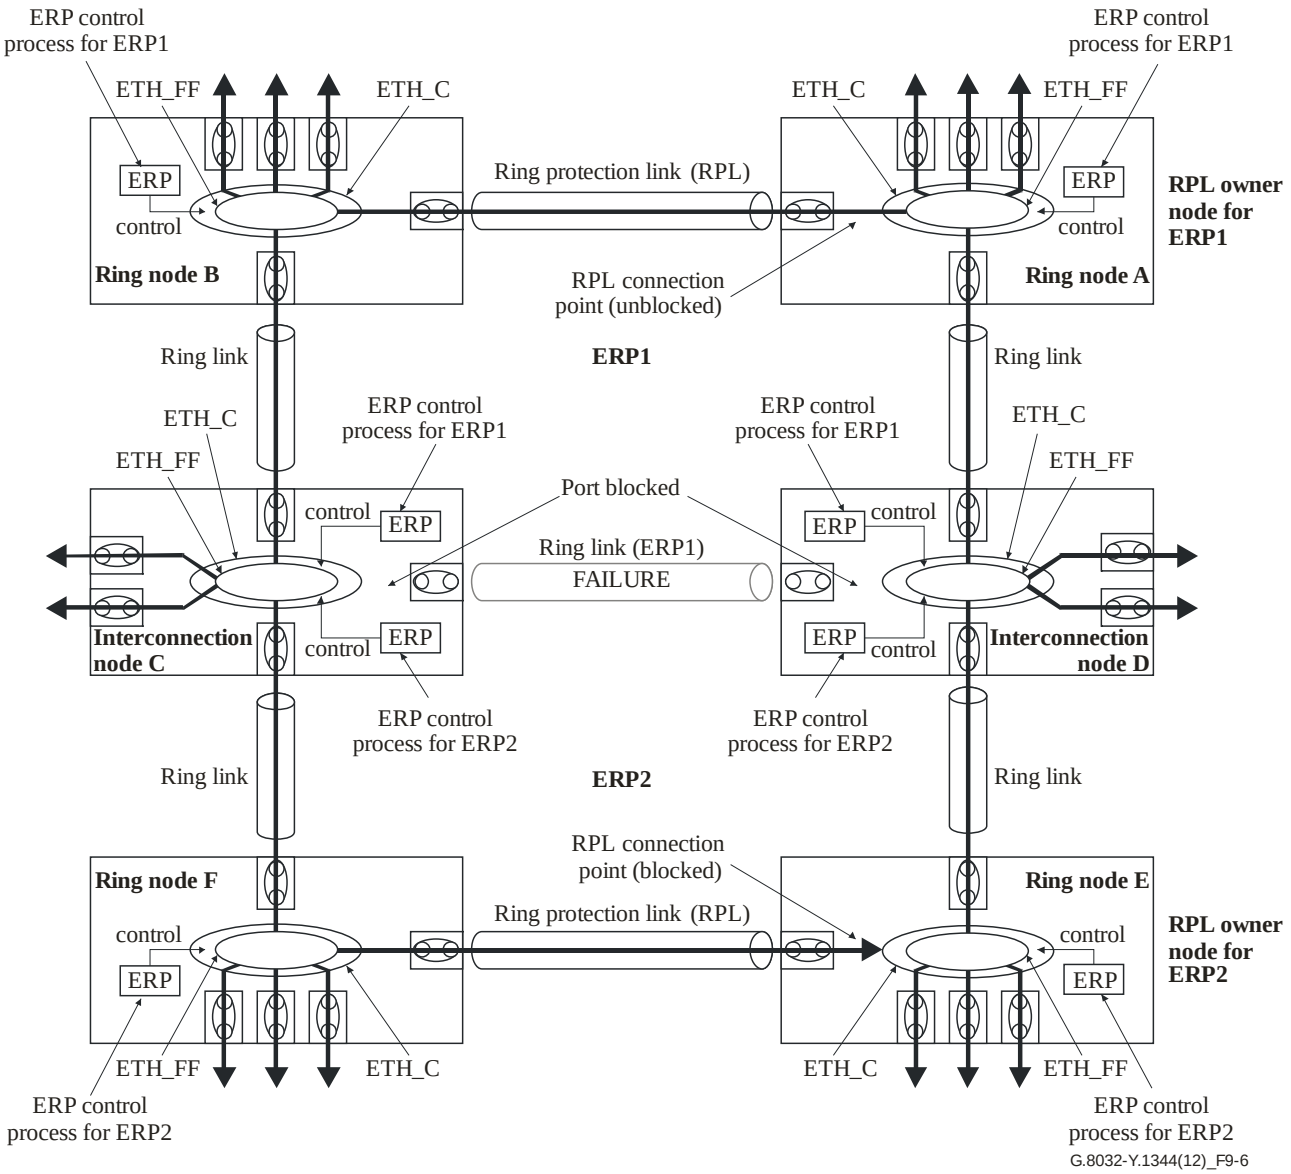

Figure 9-6 – Ethernet ring interconnection architecture – signal fail condition on a link between interconnection nodes (multi-ring/ladder network)

The interconnection nodes include functions to support the two Ethernet rings. Interconnection nodes C and D have a set of functions similar to Figure 9-4 to support Ethernet ring ERP1. Sub-ring ERP2 on these interconnection nodes only controls and protects one ring port, for this reason the model required to support sub-ring ERP2 on these interconnection nodes is as presented in the following clauses – clause 9.7.1 presents the model with an R-APS virtual channel and 9.7.2 presents the model without an R-APS virtual channel.

9.7.1 Ring interconnection model with an R-APS virtual channel

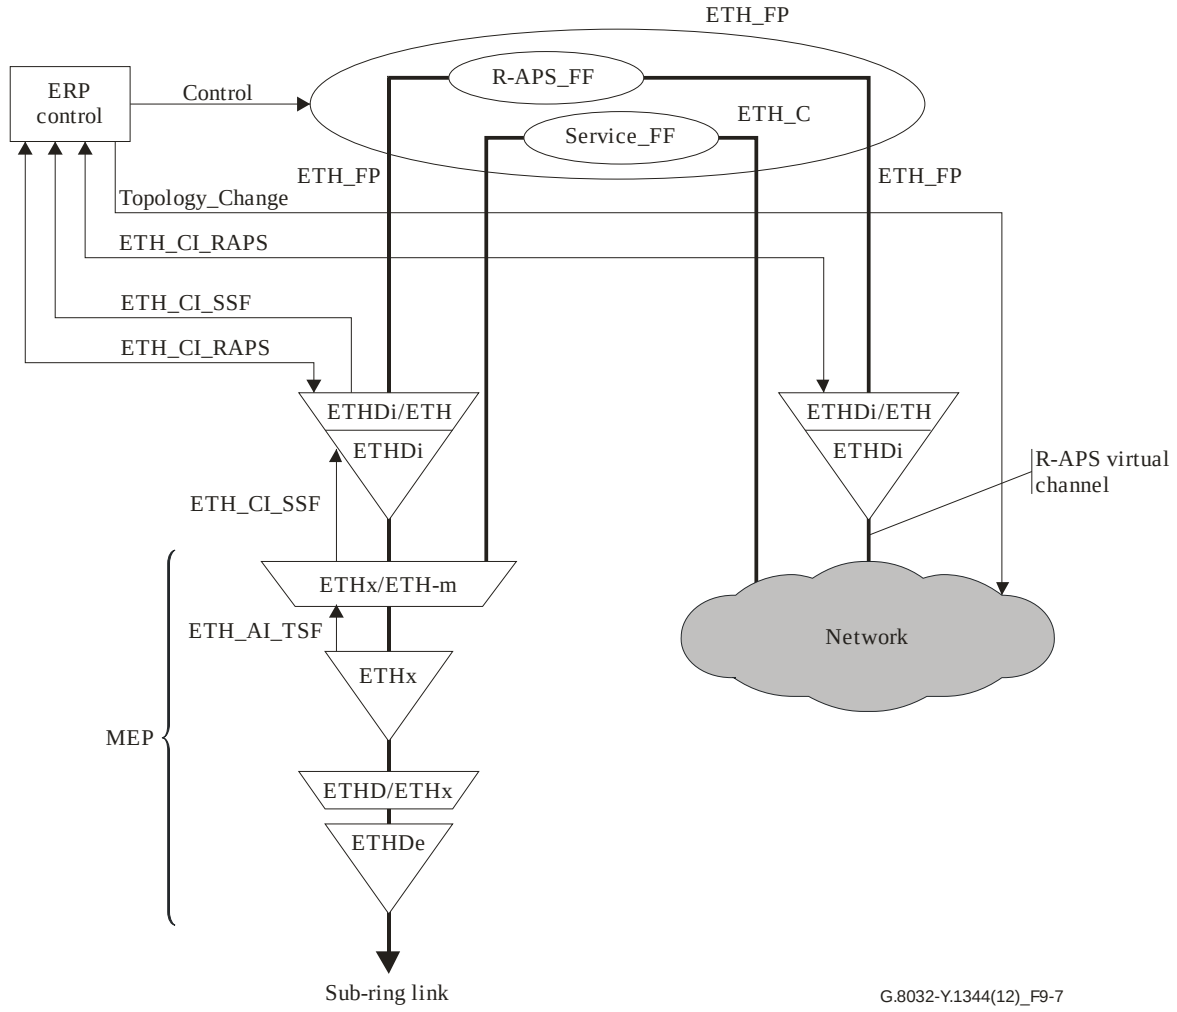

Figure 9-7 – MEPs and R-APS insertion function in an interconnection node (for a sub-ring connected to other network)

For the sub-ring, the connectivity at the interconnection node is provided between a sub-ring link and the domain of another network. In the example of Figure 9-5 this network corresponds to Ethernet ring ERP1. An R-APS virtual channel provides R-APS connectivity between this interconnection node and the other interconnection node of the same sub-ring, over the network.

An example of the functional model of an interconnection node for a sub-ring using the R-APS virtual channel is depicted in Figure 9-7.

The R-APS virtual channel may follow the same path as the traffic channel over the network. The ERP control process of the sub-ring is capable of receiving and inserting R-APS messages over the R-APS virtual channel.

R-APS messages of this sub-ring that are forwarded over its R-APS virtual channel are broadcast or multicast over the interconnected network. For this reason the broadcast/multicast domain of the R-APS virtual channel could be limited to the necessary links and nodes. For example, the R-APS virtual channel could span only the interconnecting Ethernet rings or sub-rings that are necessary for forwarding R-APS messages of this sub-ring. Care must be taken to ensure that the local R-APS messages of the sub-ring being transported over the R-APS virtual channel into the interconnected network can be uniquely disambiguated from those of other interconnected ring R-APS messages. This can be achieved by, for example, using separate VIDs for the R-APS virtual channels of different sub-rings.

Sub-ring topology changes may impact flow forwarding over the domain of the other (interconnected) network, as such topology change events are signalled to the domain of the other network using the Topology_Change signal. It is outside of the scope of this recommendation to define the use of Topology_Change signal by other technologies such as, STP or VPLS.

Figure 9-8 represents the model of an interconnection node combining the functions required to support the two Ethernet rings.

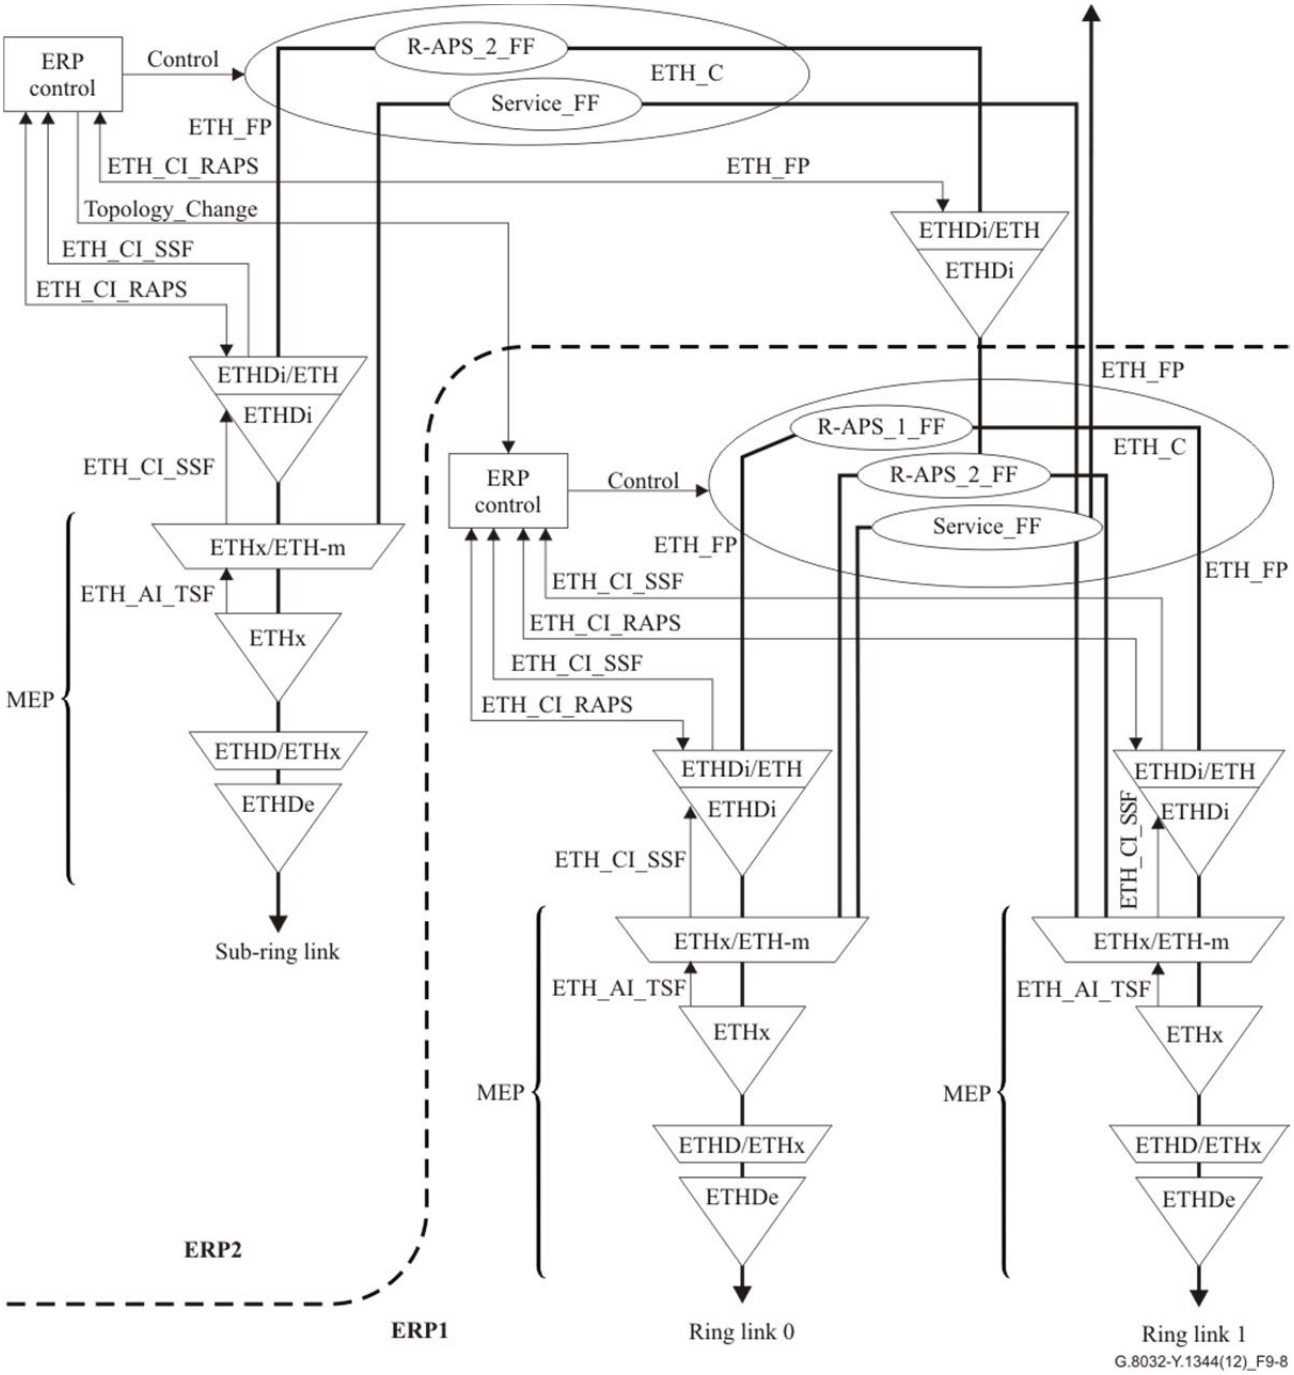

Figure 9-8 – MEPs and R-APS insertion function in an interconnection node with an R-APS virtual channel (different R-APS VIDs)

The MEPs on ring links 0 and 1 are used for monitoring the ring links of ERP1. The MEP on the sub-ring link monitors the ring link of the sub-ring, ERP2. In the model of this figure R-APS channels are separated in ERP1 using different R-APS VIDs. R-APS messages for ERP1 are received on ring links 0 or 1 and separated based on the VID used for the R-APS_1 flow at the ETHx/ETH-m_A function. The ETHDi/ETH_A functions extract ETH_CI_RAPS information from the received R-APS messages and send the ETH_CI_RAPS information to the ERP control process of ERP1. The R-APS messages of the sub-ring received on ring link 0 and on ring link 1 are separated based on the VID used for the R-APS_2 flow at the ETHx/ETH-m_A function and they are then forwarded by the R-APS_2_FF function to the ETHDi/ETH_A function where it extracts ETH_CI_RAPS information from the received R-APS messages and sends the ETH_CI_RAPS information to the ERP control process of ERP2. If not blocked at the ETH_C function of ERP2, these messages are then further transmitted to the sub-ring port.

The R-APS VID of ERP2 may be considered as protected traffic spanning all ring links of ERP1, being blocked on the ring links of ERP1 by the same function that blocks the traffic channel on the ring links of that Ethernet ring. Figure 9-8 is only one example, other options for the construction of the R-APS virtual channel may be used.

NOTE – Other solutions for the construction of the R-APS virtual channel are for further study.

Service traffic may be forwarded between any of the three ring ports, or even other ports. This forwarding is also subject to the blocking state of the Ethernet ring and sub-ring ports as defined by the respective ERP control processes.

Topology_Change signal is generated from ERP2 to ERP1 control process whenever sub-ring ERP2 performs a protection switching event that results in a topology change, this occurs when an FDB flush is generated for the ERP2 interconnection node. Depending on the configuration, this signal may be used by the ERP control process of ERP1 to initiate actions to also trigger a topology update over Ethernet ring nodes on Ethernet ring ERP1.

9.7.2 Ring interconnection model without an R-APS virtual channel

In certain network scenarios it may be desirable that the R-APS virtual channel of the sub-ring over the other network domain is not used.

An example of the functional model of an interconnection node for a sub-ring not using the R-APS virtual channel is depicted in Figure 9-9.

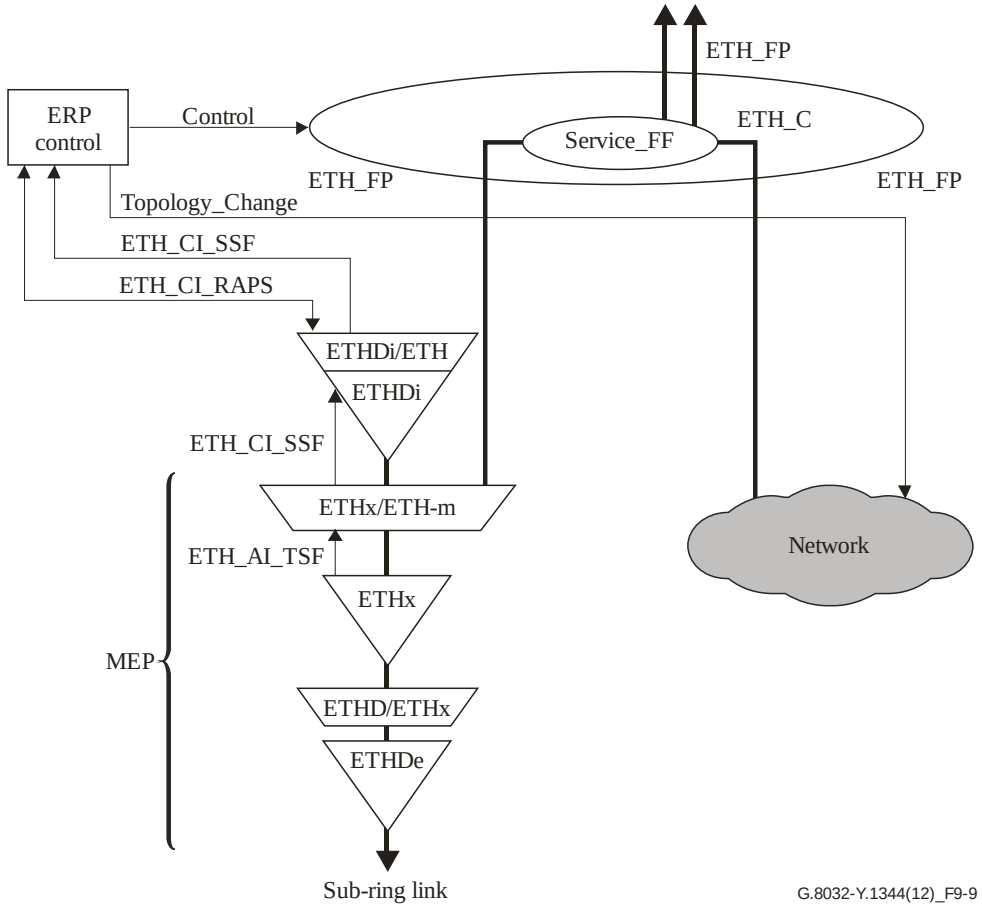

Figure 9-9 – MEPs and R-APS insertion function in a sub-ring interconnection node without an R-APS virtual channel (for a sub-ring connected to another network)

As depicted, the R-APS channel of the sub-ring is terminated at the interconnection nodes.

In order to prevent R-APS channel segmentation in the normal Ethernet ring condition, since there is neither an R-APS channel nor an R-APS virtual channel between the interconnection nodes of the sub-ring, the R-APS channel blocking (defined in clause 9.5) is not employed in these sub-ring configurations. In case of ring link failure of any ring link of the sub-ring, the R-APS channel of the sub-ring may be segmented, preventing R-APS message exchange between some of the sub-ring's Ethernet ring nodes.

Apart from R-APS channel specifics, the operation of the sub-ring without an R-APS virtual channel is identical to that of a sub-ring with an R-APS virtual channel. Interconnection nodes also perform the same functions to inform other networks of topology change and flush propagation.

In addition, in order to ensure correct operation of the FDB flush operation, there are changes to the operation of the flush logic (see clause 10.1.10).

Figure 9-10 represents the model of an interconnection node combining the functions required to support the two Ethernet rings.

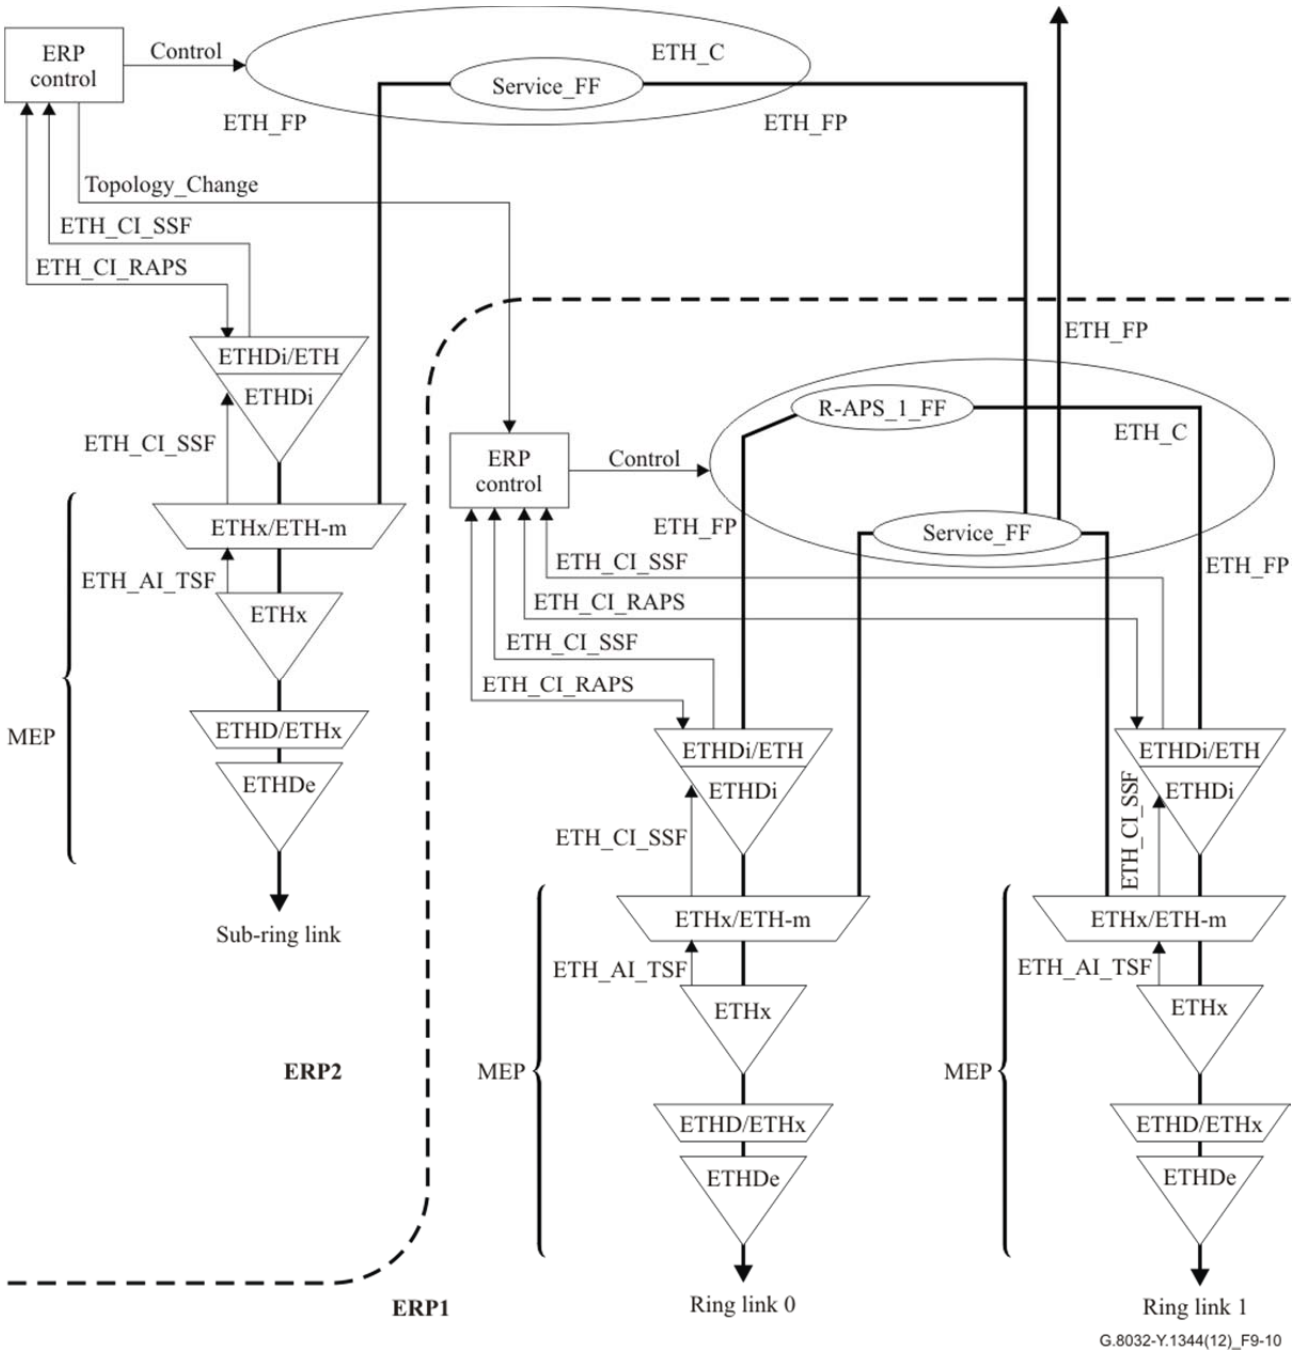

Figure 9-10 – MEPs and R-APS insertion function in a sub-ring interconnection node without an R-APS virtual channel (for a sub-ring connected to a major ring)

9.7.3 Guidelines for using the ring interconnection model with or without an R-APS virtual channel

This Recommendation defines two Ethernet ring interconnection options, as shown in Figure 9-11.

1) Sub-ring with an R-APS virtual channel: In this option, a virtual channel to tunnel R-APS messages from one interconnection node to the other interconnection node is established.

2) Sub-ring without an R-APS virtual channel: In this option, the R-APS channel is terminated at the interconnection nodes and its R-APS messages are not tunnelled between the interconnection nodes.

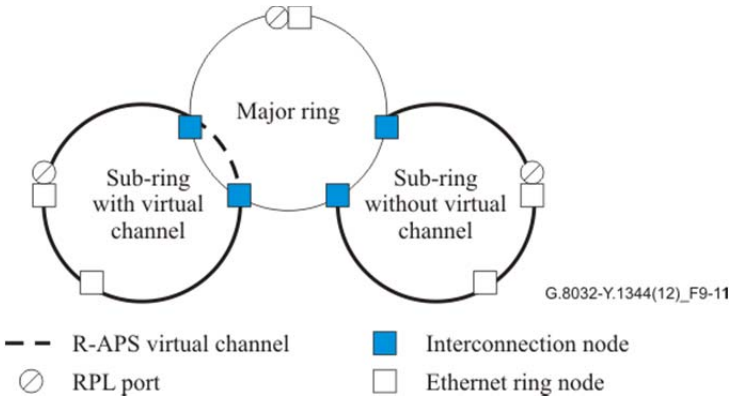

Figure 9-11 – Ring interconnection options

In option 1, the R-APS channel blocking mechanism as defined in clause 9.5 is the same for both single and multi-ring application. In addition, this option allows operators to interconnect multiple Ethernet rings (or non ITU-T G.8032 networks) without the need to reconfigure the major ring as a sub-ring (i.e., regarding the ERP control process and R-APS channel blocking mechanism). In the example of Figure 9-12, both major rings 1 and 2 can be interconnected via a newly configured sub-ring 3 with two R-APS virtual channels. However, it should be noted that the R-APS virtual channel requires a certain bandwidth to forward R-APS messages on the interconnected Ethernet ring(s) (or network) where a sub-ring is attached and it is necessary to allocate different VIDs to differentiate between each R-APS channel within a whole interconnected network. It should also be noted that the protection switching time of the sub-ring might be affected if R-APS messages traverse a long distance over an R-APS virtual channel. Major ring 1 might not be flushed due to protection switching in major ring 2 (and vice versa) and major rings 1 and 2 might be flushed due to protection switching in the sub-ring 3.

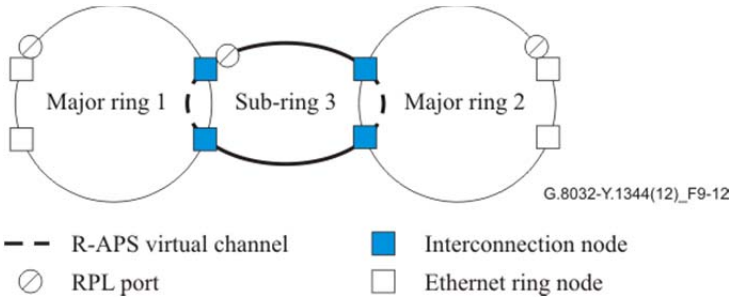

Figure 9-12 – Interconnection of two Ethernet rings with option 1

In option 2, no R-APS messages are inserted or extracted by another Ethernet ring(s) (or sub-ring(s)) at interconnection nodes where a sub-ring is attached. Hence there is no need for either additional bandwidth or different VIDs for the Ethernet ring interconnection. Furthermore, the protection switching time for a sub-ring is independent from the configuration of the interconnected Ethernet ring(s). In addition, this option always ensures that an interconnected network forms a tree topology regardless of its interconnection configuration. This means that it is not necessary to take precautions not to form a loop which is potentially composed of a whole interconnected network. However, the R-APS channel blocking mechanism is different from that of a single Ethernet ring as described in clause 10.1.14. In addition, if two Ethernet rings are interconnected using a sub-ring, the attributes of one of the Ethernet rings may need to be reconfigured to define it as a sub-ring. For example, major ring 2 of Figure 9-12 is reconfigured as a sub-ring (i.e., sub-ring 2 in Figure 9-13) for the interconnection. As a result, service interruption may occur during this reconfiguration and major ring 1 might perform FDB flushing due to protection switching in sub-rings 2 or 3.

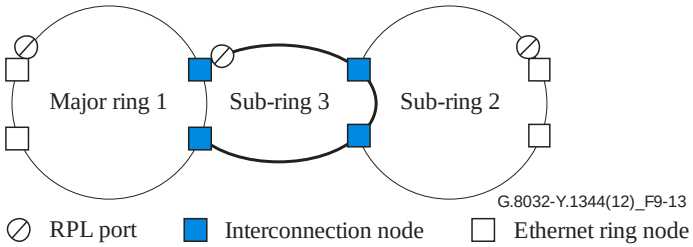

Figure 9-13 – Interconnection of two Ethernet rings with option 2

10 Protection control protocol

Ring protection is based on loop avoidance. This is achieved by guaranteeing that at any time traffic may flow on all but one of the ring links. From this principle the following rule is derived for the protocol:

> *Once a ring port has been blocked, it may be unblocked only if it is known that there remains at least one other blocked ring port in the Ethernet ring.*

This rule is used as the basis to control all actions of traffic channel unblocking in the Ethernet ring, as well as to define the information that is necessary to distribute between all Ethernet ring nodes.

10.1 Principles of operations

Figure 10-1 shows a decomposition of the ERP control process. This process is performed at all Ethernet ring nodes.

The protection algorithm is based on the transmission of local switch requests and local status to all Ethernet ring nodes via the R-APS specific information. Format and content of an R-APS message are described in clause 10.3.

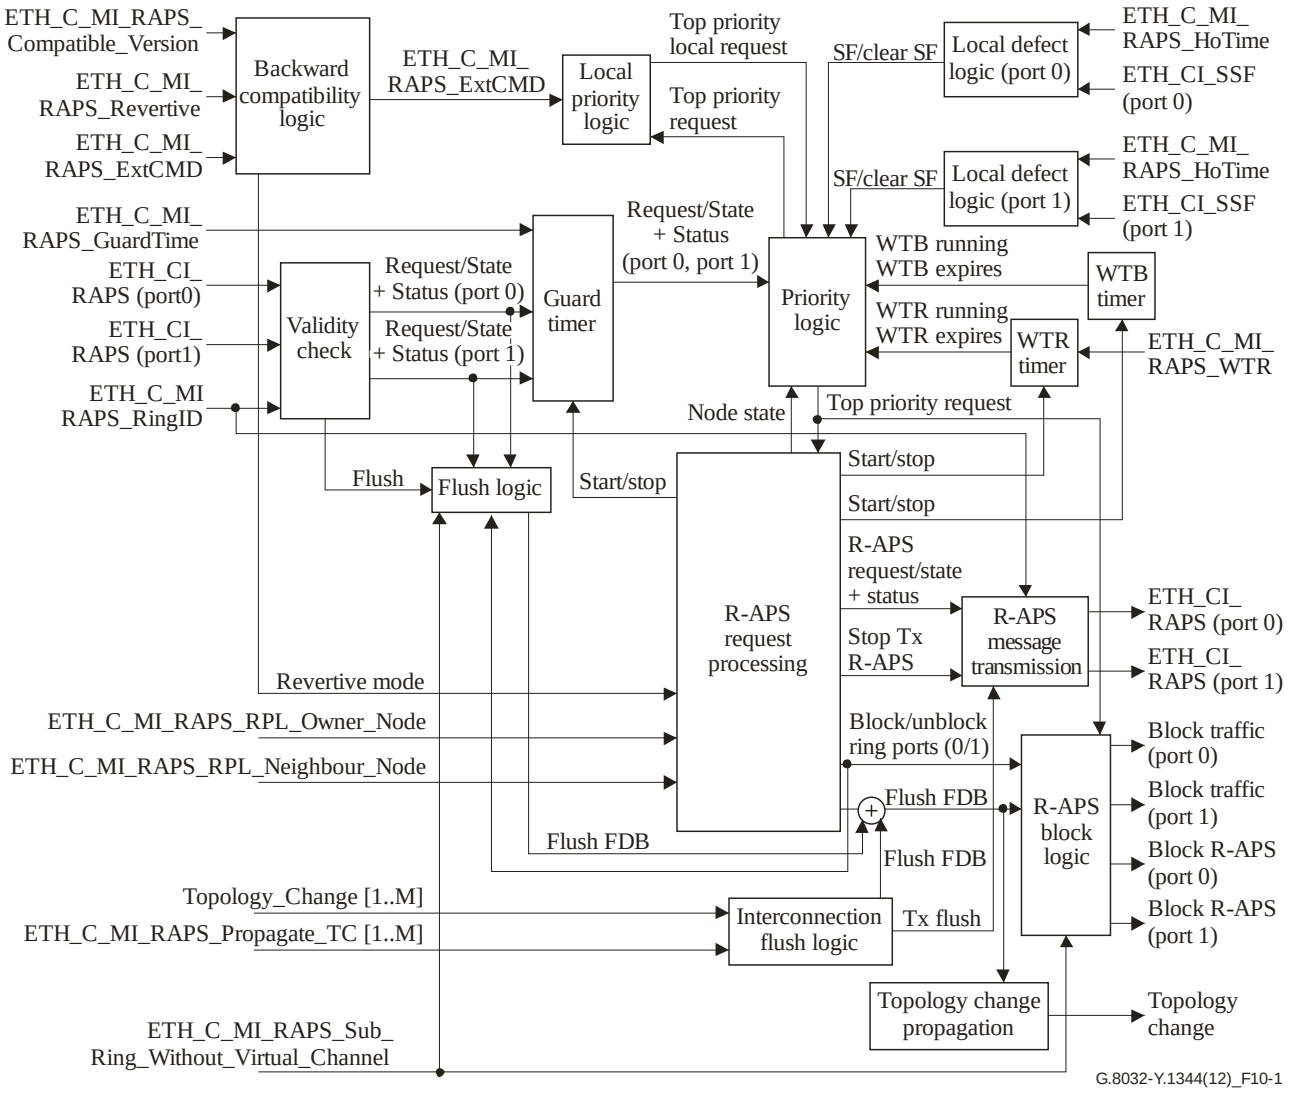

Figure 10-1 – Decomposition of ERP control process

The following is an overview of the ERP control process. The behaviour of each sub-process is described in detail in the following clauses.

At an Ethernet ring node, one or more local protection switching requests may be active. The local priority logic determines which of these requests is of top priority, using the priority order given in Table 10-1. This top priority local request information is passed to the priority logic.

The status of the local Ethernet ring node's ring ports is evaluated according to the methods defined in clause 9.2.1. This information is passed on to the local defect logic for each of the Ethernet ring node's ports. The local defect logic evaluates these signals, processes the hold-off timer and passes them to the priority logic. On the ERP control process for a sub-ring at an interconnection node only one local defect logic process exists, assigned to the sub-ring link of that Ethernet ring node. The local Ethernet ring node receives information from the other Ethernet ring nodes via R-APS messages. The validity check, as described in clause 10.1.6, verifies that the R-APS message is correctly constructed. The received request/state and status information (which indicates the top priority request and status of other Ethernet ring nodes) is then passed to the guard timer. At an interconnection node the R-APS messages may be received via an R-APS virtual channel.

The guard timer functionality is described in clause 10.1.5. While the guard timer is running, the received R-APS request/state and status information is not forwarded to the priority logic. If the

guard timer is not running, the R-APS request/state and status information is forwarded to the priority logic entity.

The functionality of the WTR timer is described in clause 10.1.4. While the WTR timer is running, the WTR Running signal is input to the priority logic. The expiration of the WTR timer is indicated by the WTR Expires signal and is passed to the priority logic entity.

The functionality of the WTB timer is described in clause 10.1.4. While the WTB timer is running, the WTB Running signal is input to the priority logic. The expiration of the WTB timer is indicated by the WTB Expires signal and is passed to the priority logic entity.

An R-APS message is defined as accepted if the message passes the validity check, is passed by the guard timer to the priority logic and is identified as the current top priority request signalled to the R-APS request processing logic.

The priority logic accepts as inputs (a) the R-APS request/state and status information (after screening by the validity check and the guard timer), (b) status and events from the WTR timer, (c) status and events from the WTB timer, (d) status of the local Ethernet ring node's ring ports, (e) top priority local request (from the local priority logic) and (f) the current node state from the R-APS request processing. It processes the priority according to Table 10-1 to determine the top priority signal.

ETH_C_MI_RAPS_RPL_Owner_Node represents management information that indicates if the local Ethernet ring node is an RPL owner node or not and in the case that this is an RPL owner node it specifies which ring port is attached to the RPL.

ETH_C_MI_RAPS_RPL_Neighbour_Node provides management information that indicates this Ethernet ring node to be adjacent to the RPL or not and in case it is an RPL neighbour node it also specifies which ring port is attached to the RPL. By default the ETH_C_MI_RAPS_RPL_Neighbour_Node indicates the Ethernet ring node as not being adjacent to the RPL.

Both ETH_C_MI_RAPS_RPL_Owner_Node and ETH_C_MI_RAPS_RPL_Neighbour_Node cannot be enabled at the same Ethernet ring node for a single ERP instance.

NOTE – In the case that ETH_C_MI_RAPS_RPL_Neighbour_Node is not configured for any Ethernet ring node on a ring, only one end of the RPL (i.e., only at the RPL owner node) is blocked.

The R-APS request processing receives the current top priority request and defines the necessary actions to take based on the local Ethernet ring node state. These actions may include transmission of R-APS messages, blocking or unblocking ring ports, flushing the FDB and starting or stopping the timers. The decision logic of the R-APS request processing is defined in clause 10.1.2 and represents the Ethernet ring protection behaviour described in the remaining subclauses of clause 10.

The Ethernet ring protection switching algorithm commences immediately after any of the input signals (see Figure 10-1) changes, i.e., when the status of any local request changes, or when a different R-APS message is received.

The flush logic is described in clause 10.1.10; it receives as inputs R-APS requests from the ring ports. Based on this information it infers whether the logical topology of the Ethernet ring has been changed and, in this case, triggers a flush of the local FDB.

The topology change propagation process is described in clause 10.1.12; it generates a signal to inform the entities of other network domains attached to a sub-ring of topology changes on the sub-ring. This process exists only on the ERP control processes of sub-ring interconnection nodes.

The interconnection flush logic is described in clause 10.1.11. It receives topology change notification information from other connected entities, such as a sub-ring's ERP control process and ETH_C_MI_RAPS_Propagate_TC management information. Based on this information, it may

initiate flushing of the FDB for the local ring ports and may trigger transmission of R-APS event requests to both ring ports. This logic is included on the ERP control processes of the interconnection nodes of Ethernet rings that sub-rings are connected to. This logic is not present on Ethernet ring nodes that are not interconnection nodes.

The backward compatibility logic is described in clause 10.1.13. It filters the configuration and requests of this version of this Recommendation when the Ethernet ring node is part of an Ethernet ring that is also composed of other Ethernet ring nodes which are implementing a previous version of this Recommendation.

The R-APS block logic is described in clause 10.1.14. It receives block/unblock ring ports (0/1) from the R-APS request processing, the top priority request from the priority logic and ETH_C_MI_RAPS_Sub_Ring_Without_Virtual_Channel signal. Based on these inputs, it decides to block or unblock the traffic channel and/or the R-APS channel on ring ports 0 and 1. This logic is present only in the ERP control process of sub-ring nodes.

10.1.1 Priority logic

This process receives requests from multiple sources. The request with the highest priority in Table 10-1, is declared as the top priority request. If an Ethernet ring node state is in Forced switch state, a local SF request is ignored.

The evaluation of the top priority request is repeated every time a local request changes or an R-APS message is received.

Ring protection requests, commands and R-APS signals have the priorities as specified in Table 10-1.

Table 10-1 – Request/state priority

Request/state and status	Type	Priority
Clear	local	highest
FS	local	|
R-APS (FS)	remote	|
local SF [a)]	local	|
local clear SF	local	|
R-APS (SF)	remote	|
R-APS (MS)	remote	|
MS	local	|
WTR Expires	local	|
WTR Running	local	|
WTB Expires	local	|
WTB Running	local	|
R-APS (NR, RB)	remote	|
R-APS (NR)	remote	lowest
[a)] If an Ethernet ring node is in the Forced switch state, local SF is ignored.		

As a result of this process, once an SF condition or operator command (e.g., FS, MS) is declared at one of the ring ports, the priority logic retains this condition request as the current top priority request, until either a new higher priority request or an appropriate Clear message (i.e., Clear for either FS or MS, local clear SF for SF) is signalled. The local clear SF condition is only signalled as the top priority request if it is the highest priority request present and there is not any higher priority request (such as local SF or local FS) still pending on the other ring port.

Received R-APS request/state and status are not stored in this process. As a result, after the change of a local request, R-APS request/state and status received previously are not taken into consideration for the definition of the new top priority request.

R-APS messages whose node ID field value corresponds to the local node ID are ignored by this process.

A ring ID in the range [1, .., 239] can be configured for each ERP instance. This ring ID is used in the "R-APS Message Transmission" function to determine the value of the last octet of the MAC destination address field of the R-APS PDUs generated by this ERP control process. It is also used by the "Validity Check" function to discard any R-APS PDU, received by this ERP control process with a non-matching ring ID.

With regard to the configuration of the ring ID, the following rules apply:

1. All ERP control processes instantiated in an ERP protected network composed of interconnected major rings and sub-rings must be identifiable by a unique (ring ID, R-APS VID) pair.
2. All ERP control processes instantiated on the same underlying physical major ring or sub-ring topologies must be assigned a different value of the R-APS VID. The same ring ID may be used for these ERP control processes.
3. ERP control processes instantiated on different physical major ring or sub-ring topologies may use different ring IDs and in that case their R-APS VIDs need not be different.

10.1.2 R-APS request processing

The R-APS request processing logic receives the current top priority request and defines the necessary actions to take, based on the local Ethernet ring node state. The R-APS request processing logic is defined in the format of a state machine. Table 10-2 has the following fields:

a) Node state – The current state of the Ethernet ring node.

b) Top priority request – The current top priority request as defined in clause 10.1.1. Each possible trigger is represented in a separate row.

c) Actions – A list of protection switching actions, in order of execution.

d) Next node state – The state to which the state machine transits.

Table 10-2 – State machine representation of the R-APS request processing logic

	Inputs		Outputs	
Node state	**Top priority request**	**Row**	**Actions**	**Next node state**

–	State machine initialization	1	Stop guard timer
			Stop WTR timer
			Stop WTB timer
			If RPL owner node:
			  Block RPL port
			  Unblock non-RPL port
			  Tx R-APS (NR)
			  If revertive:
			  Start WTR timer
			Else if RPL neighbour node:
			  Block RPL port
			Unblock non-RPL port
			  Tx R-APS (NR)
			Else:
			  Block one ring port
			  Unblock other ring port
			  Tx R-APS (NR)
				E

A (Idle)	Clear	2	No action	A

	FS	3	If requested ring port is already blocked:
			  Tx R-APS (FS,DNF)
			  Unblock non-requested ring port
			Else:
			  Block requested ring port
			  Tx R-APS (FS)
			  Unblock non-requested ring port
			  Flush FDB
				D

	R-APS (FS)	4	Unblock ring ports Stop Tx R-APS	D

	local SF	5	If failed ring port is already blocked:
			  Tx R-APS (SF,DNF)
			  Unblock non-failed ring port
			Else:
			  Block failed ring port
			  Tx R-APS (SF)
			  Unblock non-failed ring port
			  Flush FDB
				B

	local clear SF	6	No action	A
	R-APS (SF)	7	Unblock non-failed ring port Stop Tx R-APS	B
	R-APS (MS)	8	Unblock non-failed ring port Stop Tx R-APS	C

Table 10-2 – State machine representation of the R-APS request processing logic

	Inputs		Outputs	
Node state	**Top priority request**	**Row**	**Actions**	**Next node state**
	MS	9	If requested ring port is already blocked: Tx R-APS (MS,DNF) Unblock non-requested ring port Else: Block requested ring port Tx R-APS (MS) Unblock non-requested ring port Flush FDB	C
	WTR Expires	10	No action	A
	WTR Running	11	No action	A
	WTB Expires	12	No action	A
	WTB Running	13	No action	A
	R-APS (NR, RB)	14	Unblock non-RPL port If Not RPL owner node: Stop Tx R-APS	A
	R-APS (NR)	15	If neither RPL owner node nor RPL neighbour node, and remote node ID is higher than own node ID: Unblock non-failed ring port Stop Tx R-APS	A
B (Protection)	Clear	16	No action	B
	FS	17	If requested ring port is already blocked: Tx R-APS (FS,DNF) Unblock non-requested ring port Else: Block requested ring port Tx R-APS (FS) Unblock non-requested ring port Flush FDB	D
	R-APS (FS)	18	Unblock ring ports Stop Tx R-APS	D
	local SF	19	If failed ring port is already blocked: Tx R-APS (SF,DNF) Unblock non-failed ring port Else: Block failed ring port Tx R-APS (SF) Unblock non-failed ring port Flush FDB	B
	local clear SF	20	Start guard timer Tx R-APS (NR) If RPL owner node and revertive mode:	E

Table 10-2 – State machine representation of the R-APS request processing logic

	Inputs		Outputs	
Node state	**Top priority request**	**Row**	**Actions**	**Next node state**
			Start WTR	
	R-APS (SF)	21	No action	B
	R-APS (MS)	22	No action	B
	MS	23	No action	B
	WTR Expires	24	No action	B
	WTR Running	25	No action	B
	WTB Expires	26	No action	B
	WTB Running	27	No action	B
	R-APS (NR, RB)	28	No action	E
	R-APS (NR)	29	If RPL owner node and revertive mode: Start WTR	E
C (Manual switch)	Clear	30	If any ring port blocked: Start guard timer Tx R-APS (NR) If RPL owner node and revertive mode: Start WTB	E
	FS	31	If requested ring port is already blocked: Tx R-APS (FS,DNF) Unblock non-requested ring port Else: Block requested ring port Tx R-APS (FS) Unblock non-requested ring port Flush FDB	D
	R-APS (FS)	32	Unblock ring ports Stop Tx R-APS	D
	local SF	33	If failed ring port is already blocked: Tx R-APS (SF,DNF) Unblock non-failed ring port Else: Block failed ring port Tx R-APS (SF) Unblock non-failed ring port Flush FDB	B
	local clear SF	34	No action	C
	R-APS (SF)	35	Unblock non-failed ring port Stop Tx R-APS	B
	R-APS (MS)	36	If any ring port blocked: Start guard timer	E[a)]

Table 10-2 – State machine representation of the R-APS request processing logic

	Inputs		Outputs	
Node state	**Top priority request**	**Row**	**Actions**	**Next node state**
			Tx R-APS (NR) If RPL owner node and revertive mode: Start WTB	
	MS	37	No action	C
	WTR Expires	38	No action	C
	WTR Running	39	No action	C
	WTB Expires	40	No action	C
	WTB Running	41	No action	C
	R-APS (NR, RB)	42	No action	E
	R-APS (NR)	43	If RPL owner node and revertive mode: Start WTB	E
D (Forced switch)	Clear	44	If any ring port blocked: Start guard timer Tx R-APS (NR) If RPL owner node and revertive mode: Start WTB	E
	FS	45	Block requested ring port Tx R-APS (FS) Flush FDB	D
	R-APS (FS)	46	No action	D
	local SF	47	No action	D
	local clear SF	48	No action	D
	R-APS (SF)	49	No action	D
	R-APS (MS)	50	No action	D
	MS	51	No action	D
	WTR Expires	52	No action	D
	WTR Running	53	No action	D
	WTB Expires	54	No action	D
	WTB Running	55	No action	D
	R-APS (NR, RB)	56	No action	E
	R-APS (NR)	57	If RPL owner node and revertive mode: Start WTB	E
E (Pending)	Clear	58	If RPL owner node: Stop WTR Stop WTB If RPL port is blocked: Tx R-APS (NR, RB,DNF) Unblock non-RPL port	A

Table 10-2 – State machine representation of the R-APS request processing logic

	Inputs		Outputs	
Node state	**Top priority request**	**Row**	**Actions**	**Next node state**
			Else: Block RPL port Tx R-APS (NR, RB) Unblock non-RPL port Flush FDB	
	FS	59	If requested ring port is already blocked: Tx R-APS (FS,DNF) Unblock non-requested ring port Else: Block requested ring port Tx R-APS (FS) Unblock non-requested ring port Flush FDB If RPL owner node: Stop WTR Stop WTB	D
	R-APS (FS)	60	Unblock ring ports Stop Tx R-APS If RPL owner node: Stop WTR Stop WTB	D
	local SF	61	If failed ring port is already blocked: Tx R-APS (SF,DNF) Unblock non-failed ring port Else: Block failed ring port Tx R-APS (SF) Unblock non-failed ring port Flush FDB If RPL owner node: Stop WTR Stop WTB	B
	local clear SF	62	No action	E
	R-APS (SF)	63	Unblock non-failed ring port Stop Tx R-APS If RPL owner node: Stop WTR Stop WTB	B
	R-APS (MS)	64	Unblock non-failed ring port Stop Tx R-APS If RPL owner node: Stop WTR	C

Table 10-2 – State machine representation of the R-APS request processing logic

	Inputs		Outputs	
Node state	**Top priority request**	**Row**	**Actions**	**Next node state**
			Stop WTB	
	MS	65	If RPL owner node: Stop WTR Stop WTB If requested ring port is already blocked: Tx R-APS (MS,DNF) Unblock non-requested ring port Else: Block requested ring port Tx R-APS (MS) Unblock non-requested ring port Flush FDB	C
	WTR Expires	66	If RPL owner node: Stop WTB If RPL port is blocked: Tx R-APS (NR, RB,DNF) Unblock non-RPL port Else: Block RPL port Tx R-APS (NR, RB) Unblock non-RPL port Flush FDB	A
	WTR Running	67	No action	E
	WTB Expires	68	If RPL owner node: Stop WTR If RPL port is blocked: Tx R-APS (NR, RB,DNF) Unblock non-RPL port Else: Block RPL port Tx R-APS (NR, RB) Unblock non-RPL port Flush FDB	A
	WTB Running	69	No action	E
	R-APS (NR, RB)	70	If RPL owner node: Stop WTR Stop WTB If neither RPL owner node nor RPL neighbour node: Unblock ring ports Stop Tx R-APS If RPL neighbour node:	A

Table 10-2 – State machine representation of the R-APS request processing logic

	Inputs		Outputs	
Node state	**Top priority request**	**Row**	**Actions**	**Next node state**
			Block RPL port Unblock non-RPL port Stop Tx R-APS	
	R-APS (NR)	71	If remote node ID is higher than own node ID: Unblock non-failed ring port Stop Tx R-APS	E
[a)] If both ring ports are unblocked, next node state is C.				

NOTE 1 – Table 10-2 should not be interpreted independently of the other sub-processes of the ERP control process, including the priority logic.

NOTE 2 – In R-APS (msgtype, status_bits), "msgtype" indicates the request/state and "status_bits" indicates that the RB or DNF status bit is 1. If "status_bits" is 0, it is not included in R-APS (msgtype, status_bits). These fields and their possible values are defined in clause 10.3.

Row 1 represents the actions being triggered at the initialization of the state machine. Once those actions are performed the state machine shall transit to state E and eventually, when the network stabilizes to state A.

The possible actions triggered by this process and listed in the "Actions" column are:

a) Block requested ring port – blocks the traffic channel and R-APS channel (in accordance with the process described in clause 10.1.14) on the ring port for which an operator command was issued. If the ring port is already blocked, it remains blocked.

b) Unblock non-requested ring port – unblocks traffic channel and R-APS channel on the ring port for which no operator command is issued. If the ring port is already unblocked it remains unblocked.

c) Block failed ring port – blocks the traffic channel and R-APS channel (in accordance with the process described in clause 10.1.14) on the ring port which has an SF condition. If the ring port is already blocked it remains blocked.

d) Unblock non-failed ring port – unblocks the traffic channel and R-APS channel on either of the ring ports if it does not have an SF condition. If the ring port is already unblocked it remains unblocked. In case of an interconnection node of a sub-ring this action is only applied to the sub-ring port.

e) Block RPL port – blocks the traffic channel and R-APS channel (in accordance with the process described in clause 10.1.14) on the ring port which is connected to the RPL. If the ring port connected to the RPL is already blocked it remains blocked.

f) Unblock non-RPL port – unblocks the traffic channel and R-APS channel on the ring ports if it is not the RPL port. If the ring port is already unblocked it remains unblocked. In the case of an interconnection node of a sub-ring this action is only applied to the sub-ring port.

g) Block one ring port – blocks the traffic channel and R-APS channel (in accordance with the process described in clause 10.1.14) on one of the ring ports.

h) Unblock other ring port – unblocks traffic channel and R-APS channel on the second ring port where the port is not unblocked. In case of an interconnection node of a sub-ring this action is not applied.

i) Unblock ring ports – unblocks the traffic channel and R-APS channel on both ring ports. If a ring port is already unblocked it remains unblocked. In the case of an interconnection node of a sub-ring this action is only applied to the sub-ring port.

j) Start WTR – starts the WTR timer if it is stopped. If the WTR timer is already running, no action is taken.

k) Stop WTR – stops the WTR timer if it is running.

l) Start WTB – starts the WTB timer if it is stopped. If the WTB timer is already running, no action is taken.

m) Stop WTB – stops the WTB timer if it is running.

n) Start guard timer – starts the guard timer.

o) Stop guard timer – stops the guard timer if it is running.

p) Stop Tx R-APS – stops the transmission of any R-APS messages.

q) Tx R-APS (msgtype, status_bits) – starts the continuous transmission of R-APS messages on both ring ports as described in clause 10.1.3.

r) Flush FDB – Triggers an FDB flush as described in clause 9.6.

In the multi-ring/ladder network, a failure on the ring link connecting the interconnection nodes triggers the above actions only on the Ethernet ring that it is configured to be part of. In case of a link failure on one of the sub-ring links, this triggers the above actions only on that sub-ring.

10.1.3 R-APS message transmission

R-APS messages are transmitted with the request/state and status information defined by the R-APS request process and with the ring ID (configured via ETH_C_MI_RAPS_RingID) encoded in the MAC destination address.

The action Tx R-APS (msgtype, status_bits) starts the transmission of an R-APS message with the request/state field set to the value defined by msgtype and with the status bits enumerated in status_bits with value 1 and the remaining status bits with value 0. R-APS messages are transmitted over both ring ports. This also stops the continuous transmission of any other messages, with the exception of "event" messages described below.

The action Stop Tx R-APS, results in stopping transmission of any R-APS messages.

The R-APS messages are transported via an R-APS specific VLAN.

A new R-APS message should be transmitted immediately when required as an output action of Table 10-2.

If the R-APS information to be transmitted has been changed, a burst of three R-APS messages is transmitted as quickly as possible. This ensures that fast protection switching is possible even if one or two R-APS messages are lost or corrupted. For protection switching within 50 ms, the interval between the first three R-APS messages should be not more than 3.33 ms, which is the same interval as CCM messages for fast defect detection. For messages other than the "event" message, the R-APS message continues to be transmitted, after the first three messages are transmitted, with a frequency of one message every five seconds.

Unless otherwise stated, all R-APS messages are transmitted on both ring ports. In the case of interconnection nodes of a sub-ring with an R-APS virtual channel, the R-APS messages are always transmitted over the sub-ring link and the R-APS virtual channel. On interconnection nodes of a sub-ring without an R-APS virtual channel, the sub-ring R-APS messages are transmitted only to the sub-ring port. This is, in general, also applied in cases where transmission of messages is described to be performed on "both ring ports".

The transmission of R-APS "event" messages is performed only as a single burst of three R-APS messages, i.e., it is not continuously repeated beyond this burst. Contrary to other messages, the transmission of this R-APS message is done in parallel to other existing transmission. It does not stop the transmission of other messages and is not stopped by the transmission of other messages. Flush messages are R-APS "event" messages transmitted using a sub-code field (see clause 10.3) with value "0000" and with a status field (see clause 10.3) with value "00000000".

10.1.4 Delay timers

The RPL owner node uses a delay timer before initiating an RPL block in case of both revertive mode of operation or before reverting to idle state (state A) after clearing operator commands (FS, MS). In the revertive mode of operation, the 'wait to restore' (WTR) timer is used to prevent frequent operation of the protection switching due to intermittent signal failure defects. The 'wait to block' (WTB) timer is used when clearing Forced switch and Manual switch commands. As multiple Forced switch commands are allowed to co-exist in an Ethernet ring, the WTB timer ensures that clearing of a single Forced switch command does not trigger the re-blocking of the RPL. When clearing a Manual switch command, the WTB timer prevents the formation of a closed loop due to a possible timing anomaly where the RPL owner node receives an outdated remote MS request during the recovery process.

a) When recovering from a Signal fail, the delay timer must be long enough to allow the recovering network to become stable. This delay timer, called the WTR timer, may be configured by the operator (via ETH_C_MI_RAPS_WTR) in 1 minute steps between 1 and 12 minutes; the default value being 5 minutes.

b) When recovering from an operator command (i.e., FS or MS) the delay timer must be long enough to receive any latent remote FS, SF or MS. This delay timer called the WTB timer is defined to be 5 seconds longer than the guard timer (see clause 10.1.5). This is enough time to allow a reporting Ethernet ring node to transmit two R-APS messages and allow the Ethernet ring to identify the latent condition.

This delay timer is activated on the RPL owner node. When the relevant delay timer expires the RPL owner node initiates the reversion process by transmitting an R-APS (NR, RB) message. The delay timer, (i.e., WTR or WTB) is deactivated when any higher priority request pre-empts this delay timer.

The delay timers (i.e., WTR and WTB) may be started and stopped. A request to start running the delay timer does not restart the delay timer. A request to stop the delay timer stops the delay timer and resets its value. The Clear command can be used to stop the delay timer.

While a delay timer is running the WTR or the WTB Running signal is continuously generated, appropriately. After a delay timer expires, the WTR or WTB Running signal is stopped and the WTR or WTB Expires signal is generated, respectively. When a delay timer is stopped by the Clear command, neither the WTR nor WTB Expires signal is generated.

10.1.5 Guard timer

R-APS messages are transmitted as defined in clause 10.1.3. This forwarding method, in which R-APS messages are copied and forwarded at every Ethernet ring node, can result in a message corresponding to an old request, that is no longer relevant, being received by Ethernet ring nodes. Reception of an old R-APS message may result in erroneous ring state interpretation by some Ethernet ring nodes. The guard timer is used to prevent Ethernet ring nodes from acting upon outdated R-APS messages and prevents the possibility of forming a closed loop.

The guard timer is activated whenever an Ethernet ring node receives an indication that a local switching request has cleared (i.e., local clear SF, Clear). The period of the guard timer may be configured by the operator (via ETH_C_MI_RAPS_GuardTime) in 10 ms steps between 10 ms and 2 seconds, with a default value of 500 ms. This timer period should be greater than the maximum

expected forwarding delay in which an R-APS message traverses the entire ring. The longer the period of the guard timer, the longer an Ethernet ring node is unaware of new or existing relevant requests transmitted from other Ethernet ring nodes and therefore unable to react to them.

A guard timer is used in every Ethernet ring node. Once a guard timer is started, it expires by itself. While the guard timer is running, any received R-APS request/state and status information, except R-APS messages with request/state field = "1110" described in clause 10.1.6, is blocked and not forwarded to the priority logic. When the guard timer is not running, the R-APS request/state and status information is forwarded unchanged.

10.1.6 Validity check

The validity check verifies that the request/state field of the received R-APS message is one of the "Request/States" defined in Table 10-3. R-APS messages with request/state fields defined as "Reserved for future international standardization" are filtered. When an R-APS message is received with request/state field = "1110" and the sub-code field is "0000" and the status field has the value "00000000", the flush indication is signalled to the flush logic. The flush indication signal is disabled after a period of 10 ms. R-APS messages with request/state field = "1110" are not affected by the guard timer.

Additionally, the validity check verifies that the ring ID of the received R-APS message matches the ring ID of the ERP instance. R-APS messages with a non-matching ring ID are filtered.

10.1.7 Local defect logic

Local defect logic asserts the SF condition of one ring link based on the received ETH_CI_SSF information and the hold-off timer process. The reception of ETH_CI_SSF results in continuously signalling SF, after the hold-off timer process, until the ETH_CI_SSF is cleared.

Clearance of the ETH_CI_SSF results in producing the clear SF signal.

10.1.8 Hold-off timer

In order to coordinate the timing of protection switches at multiple layers, a hold-off timer may be required. Its purpose is to allow, for example, a server layer protection switch to have a chance to fix the problem before switching at a client layer.

Each ERP control process should have a configurable hold-off timer (configurable via ETH_C_MI_RAPS_HoTime). The suggested range of the hold-off timer is 0 to 10 seconds in steps of 100 ms with an accuracy of ±5 ms. The default value for the hold-off timer is 0 seconds.

When a new defect or more severe defect occurs (new SF), this event is not to be reported immediately to protection switching if the provisioned hold-off timer value is non-zero. Instead, the hold-off timer is started. When the hold-off timer expires, the trail that started the timer is checked as to whether a defect still exists. If one does exist, that defect is reported to protection switching. The reported defect need not be the same one that started the timer.

10.1.9 Local priority logic

Local priority logic evaluates the local operator commands (in ETH_C_MI_RAPS_ExtCMD) according to the current top priority request. The commands Clear, Manual switch and Forced switch from the operator, are forwarded to the priority logic.

The Clear command is only valid if:

a) a local Forced switch or Manual switch command is in effect (clear operation a) described in clause 8), or

b) a local Ethernet ring node is an RPL owner node and top priority request is neither R-APS (FS) nor R-APS (MS) (clear operations b or 0c described in clause 8).

If local command is overridden by a new top priority request of a higher priority, i.e., a local condition, local command or an R-APS request, that command is forgotten. For example, in case a higher priority request is received as the top priority request, any existing local Manual switch or Forced switch is removed and the previous command is no longer signalled as a top priority local request. In this case the command is automatically deleted without forwarding the specific Clear command to the priority logic.

10.1.10 Flush logic

The flush logic retains for each ring port the information of node ID and blocked port reference (BPR) of the last R-APS message received over that ring port. As part of the initialization of the ERP control process, this information pair should be reset at both ring ports to the following values:

- Node ID: 00:00:00:00:00:00
- BPR: 0

For each new R-APS message received over one ring port, it extracts the (node ID, BPR) pair and compares it with the previous (node ID, BPR) pair stored for that ring port. If it is different from the previous pair stored, then the previous pair is deleted and the newly received (node ID, BPR) pair is stored for that ring port; and if it is different from the (node ID, BPR) pair already stored at the other ring port, then a flush FDB action is triggered except when the new R-APS message has DNF or the receiving Ethernet ring node's node ID. An R-APS (NR) message received by this process does not cause a flush FDB, however, it causes the deletion of the current (node ID, BPR) pair on the receiving ring port. However, the received (node ID, BPR) pair is not stored. When the ring port is changed to be blocked – as indicated by the block/unblock ring ports signal – the flush logic deletes the current (node ID, BPR) pair on both ring ports.

For interconnected rings running the sub-ring without the virtual channel model the following procedure should be followed. For each new R-APS message received over one ring port, it extracts the (node ID, BPR) pair and compares it with the previous (node ID, BPR) pair stored for that ring port. If it is different from the previous pair stored, then the previous pair is deleted and the newly received (node ID, BPR) pair is stored for that ring port and a flush FDB action is triggered unless the new R-APS message has its DNF bit set. In addition, the (node ID, BPR) pair stored at the other ring port is deleted. An R-APS (NR) message received by this process does not cause a flush FDB, however it causes the deletion of the current (node ID, BPR) pair on the receiving ring port, while the received (node ID, BPR) pair is not stored. When a ring port's blocking status is changed to be blocked – as indicated by the block/unblock ring ports signal – the flush logic deletes the current (node ID, BPR) pair on both ring ports.

The flush logic triggers a flush FDB action when it receives a flush indication from the validity check.

10.1.11 Interconnection flush logic

The interconnection flush logic of an ERP control process that controls two ring ports (i.e., the target ERP instance) receives as inputs the topology change signal Topology_Change[1..M] from all ERP control processes for sub-rings, located at the same interconnection node. In addition, for each Topology_Change[1..M] signal there is a corresponding management information ETH_C_MI_RAPS_Propagate_TC[1..M] signal. When one of these Topology_Change signals toggles from disabled to enabled, a flush FDB action is triggered on the ring port of the target ERP instance. In addition to the Topology_Change signal, if the corresponding ETH_C_MI_RAPS_Propagate_TC management information is enabled, a transmission of a burst of three R-APS "event" messages is triggered over the R-APS channel of the target ERP instance.

ETH_C_MI_RAPS_Propagate_TC accepts the values enabled and disabled. The default value of the ETH_C_MI_RAPS_Propagate_TC shall be disabled.

10.1.12 Topology change propagation

The topology change propagation enables the Topology_Change signal when a flush FDB action is triggered by the ERP control process of a sub-ring's ERP instance. The Topology_Change signal is disabled after a period of 10 ms.

10.1.13 Backward compatibility logic

Backward compatibility logic accepts as inputs ETH_C_MI_RAPS_Compatible_Version, ETH_C_MI_RAPS_Revertive and ETH_C_MI_RAPS_ExtCMD, i.e., commands which are specific to this version of this Recommendation. If the ETH_C_MI_RAPS_Compatible_Version is set to the version number of this Recommendation, the inputs and commands are forwarded transparently. If the ETH_C_MI_RAPS_Compatible_Version is set to a previous version number than the version number of this Recommendation then some inputs and commands may not be forwarded. The default value of the ETH_C_MI_RAPS_Revertive shall be true. When the ETH_C_MI_RAPS_Revertive is set to false, the Ethernet ring is operated in non-revertive mode.

a) If the ETH_C_MI_RAPS_Compatible_Version is set to '1' then:
   1) Manual switch and Forced switch operator commands in ETH_C_MI_RAPS_ExtCMD are filtered and are not passed to the local priority logic.
   2) Revertive mode is set to the value true.

b) If the ETH_C_MI_RAPS_Compatible_Version is set to '2' then:
   1) Manual switch and Forced switch operator commands in ETH_C_MI_RAPS_ExtCMD are forwarded to the local priority logic.
   2) Revertive mode is set to the same value as the input ETH_C_MI_RAPS_Revertive.

c) ETH_C_MI_RAPS_Compatible_Version accepts the values '1' and '2'. The default value of the ETH_C_MI_RAPS_Compatible_Version shall be '2'. The ETH_C_MI_RAPS_Compatible_Version is set to '1' when an Ethernet ring node, supporting only functionalities of ITU-T G.8032 (2008) and ITU-T G.8032 (Amendment 1, 2009), exists on the same Ethernet ring.

10.1.14 R-APS block logic

The R-APS block logic receives the block/unblock ring ports (0/1) signal from the R-APS request processing, the top priority request from the priority logic and the ETH_C_MI_RAPS_Sub_Ring_Without_Virtual_Channel signal.

When the ETH_C_MI_RAPS_Sub_Ring_Without_Virtual_Channel is disabled, i.e., the sub-ring is configured to run with an R-APS virtual channel, both the traffic channel and the R-APS channel are blocked, when the block/unblock indicates the need to block a ring port.

When the ETH_C_MI_RAPS_Sub_Ring_Without_Virtual_Channel is enabled, i.e., the sub-ring is configured to run without an R-APS virtual channel and the top priority request is not a local SF or local FS request, then the traffic channel is blocked on the appropriate ring port (0/1) based on the block/unblock ring port (0/1) signal, however the R-APS channel is not blocked. If the top priority request is either a local SF or local FS then, depending on the value of the block/unblock ring port (0/1) signal, both the traffic channel and the R-APS channel are blocked for the appropriate ring port.

The default value of the ETH_C_MI_RAPS_Sub_Ring_Without_Virtual_Channel shall be disabled.

10.2 Protection switching behaviour

Protection switching behaviours on failure and recovery conditions are described in this clause.

NOTE – Scenarios illustrating the sequence of events in protection switching are included in Appendix III.

10.2.1 Protection switching – Link signal fail

An Ethernet ring with no SF request has a logical topology with the traffic channel blocked at the RPL and unblocked on all other ring links. In this situation, the detection of an SF condition on a ring link triggers protection switching as follows:

a) If no other higher priority request exists, an Ethernet ring node detecting an SF condition on one of its ring ports blocks the traffic channel and R-APS channel on the failed ring port.

b) If no other higher priority request exists, the Ethernet ring node detecting an SF condition transmits an R-APS message indicating SF on both ring ports. The R-APS (SF) message informs other Ethernet ring nodes of the SF condition and that the traffic and R-APS channels are blocked on one ring port. The R-APS (SF) message shall be continuously transmitted by the Ethernet ring node detecting the SF condition while this condition persists. For sub-ring interconnection nodes, the R-APS (SF) message is transmitted on the R-APS channel of the sub-ring port.

c) If no other higher priority request exists and assuming the Ethernet ring node was in an idle state before the SF condition occurred, upon detection of this SF condition the Ethernet ring node triggers a local FDB flush.

d) An Ethernet ring node accepting an R-APS (SF) message, without any local higher priority requests unblocks any blocked ring port that does not have an SF condition. This action unblocks the traffic channel on the RPL.

e) An Ethernet ring node accepting an R-APS (SF) message, without any local higher priority requests stops transmission of other R-APS messages.

f) An Ethernet ring node accepting an R-APS (SF) message without a DNF indication performs a flush FDB action by following the mechanism described in clause 10.1.10.

Protection switching is completed when the above actions are performed by each Ethernet ring node. At this point the conditions are created to allow the traffic flows to be steered around the Ethernet ring.

In the multi-ring/ladder network scenario, a failure on a ring link between interconnection nodes of a sub-ring triggers the above actions only on the Ethernet ring that the sub-ring is attached to. On the other hand, other ring link failures trigger the above actions within the Ethernet ring that the failed ring link belongs to.

Bidirectional link failures are detected by the two Ethernet ring nodes adjacent to the failed ring link. These two Ethernet ring nodes trigger protection switching and keep the traffic channel blocked at both ends of the failed ring link. Unidirectional link failures are detected by only one of the Ethernet ring nodes adjacent to the failed ring link. This Ethernet ring node is the only node triggering protection switching and keeps the traffic channel blocked at its end of the failed ring link. These ring port blocking behaviours are essential to prevent the Ethernet ring from forming loops when the link failure is recovered. A node failure situation is handled as the failure of both ring links of the Ethernet ring node. The two Ethernet ring nodes adjacent to the failed Ethernet ring node initiate protection switching by detecting the SF condition on ring links connected to the failed Ethernet ring node.

10.2.2 Protection switching – signal degrade on link

Protection switching behaviour in case of signal degrade condition is for further study.

10.2.3 Protection switching – recovery

An Ethernet ring node that has one or more ring ports in an SF condition, upon detection of clearance of the SF condition, keeps at least one of these ring ports blocked for the traffic channel and for the R-APS channel, until the RPL is blocked as a result of Ethernet ring protection

reversion, or until there is another higher priority request (e.g., an SF condition) in the Ethernet ring.

An Ethernet ring node that has one ring port in an SF condition and detects clearing of this SF condition continuously transmits the R-APS (NR) message with its own node ID as the priority information over both ring ports, informing that no request is present at the Ethernet ring node and initiates a guard timer as described in clause 10.1.5. Another recovered Ethernet ring node (or nodes) holding the link block receives the message and compares the node ID information with its own node ID. If the received R-APS (NR) message has the higher priority, the Ethernet ring node unblocks its ring ports. Otherwise, the block remains unchanged. There is only one link with one-end block.

The Ethernet ring nodes stop transmitting R-APS (NR) messages when they accept an R-APS (NR, RB), or when another higher priority request is received.

10.2.3.1 Revertive behaviour

When all ring links and Ethernet ring nodes have recovered and no external requests are active, reversion is the action to be taken. Reversion is handled in the following way:

a) The reception of an R-APS (NR) message causes the RPL owner node to start the WTR timer.

b) The WTR timer is cancelled if during the WTR period a higher priority request than NR is accepted by the RPL owner node or is declared locally at the RPL owner node.

c) When the WTR timer expires, without the presence of any other higher priority request, the RPL owner node initiates reversion by blocking its traffic channel over the RPL, transmitting an R-APS (NR, RB) message over both ring ports, informing the Ethernet ring that the RPL is blocked and performing a flush FDB action.

d) The acceptance of the R-APS (NR, RB) message causes all Ethernet ring nodes to unblock any blocked non-RPL link that does not have an SF condition. If it is an R-APS (NR, RB*)* message without a DNF indication, all Ethernet ring nodes perform a necessary flush FDB action by following the mechanism described in clause 10.1.10.

10.2.3.2 Non-revertive behaviour

In non-revertive operation, the Ethernet ring does not automatically revert when all ring links and Ethernet ring nodes have recovered and no external requests are active. Non-revertive operation is handled in the following way:

a) The RPL owner node does not generate a response on reception of an R-APS (NR) messages.

b) When other healthy Ethernet ring nodes receive the NR (node ID) message, no action is taken in response to the message.

c) When the operator issues a Clear command for non-revertive mode at the RPL owner node, the non-revertive operation is cleared, the RPL owner node blocks its RPL port and transmits an R-APS (NR, RB) message in both directions, repeatedly.

d) Upon receiving an R-APS (NR, RB) message, any blocking Ethernet ring node should unblock its non-failed ring port. If it is an R-APS (NR, RB*)* message without a DNF indication, all Ethernet ring nodes perform a necessary flush FDB action by following the mechanism described in clause 10.1.10.

10.2.4 Protection switching – Manual switch

An Ethernet ring with no request has a logical topology with the traffic channel blocked at the RPL and unblocked on all other ring links. In this situation, the operator-initiated Manual switch command triggers protection switching as follows:

a) If no other higher priority commands exist, the Ethernet ring node, where a Manual switch command was issued, blocks the traffic channel and R-APS channel (as described in clause 10.1.14) on the ring port to which the Manual switch command was issued. The Ethernet ring node shall unblock the other ring port.

b) If no other higher priority commands exist, the Ethernet ring node where the Manual switch command was issued transmits R-APS messages indicating MS over both ring ports. The R-APS (MS) message shall be continuously transmitted by this Ethernet ring node while the local MS command is the Ethernet ring node's highest priority command. The R-APS (MS) message informs other Ethernet ring nodes of the MS command and that the traffic channel is blocked on one ring port.

c) If no other higher priority commands exist and assuming the Ethernet ring node was in an idle state before the Manual switch command was issued, upon the Manual switch operator command the Ethernet ring node triggers a local FDB flush action.

d) An Ethernet ring node accepting an R-APS (MS) message, without any local higher priority requests unblocks any blocked ring port which does not have an SF condition. This action unblocks the traffic channel over the RPL.

e) The Ethernet ring node accepting an R-APS (MS) message, without any local higher priority requests stops transmission of R-APS messages.

f) The Ethernet ring node receiving an R-APS (MS) message performs a necessary flush FDB action by following the mechanism described in clause 10.1.10.

Protection switching on a Manual switch request is completed when the above actions are performed by each Ethernet ring node. At this point the conditions are created to allow the traffic flows to be steered around the Ethernet ring. From this point on, the following rules apply regarding processing of further Manual switch commands:

a) While an existing Manual switch request is present in the Ethernet ring, any new Manual switch request is rejected. The request is rejected at the Ethernet ring node where the new request is issued and a notification shall be generated to inform the operator that the new MS request was not accepted.

b) An Ethernet ring node with a local Manual switch command which receives an R-APS (MS) message with a different node ID shall clear its Manual switch request and start transmitting R-APS (NR) messages. The Ethernet ring node shall keep the ring port blocked due to the previous Manual switch command.

c) An Ethernet ring node with a local Manual switch command that receives an R-APS message or a local request of higher priority than R-APS (MS) shall clear its Manual switch request. The Ethernet ring node shall then process the new higher priority request.

10.2.4.1 Manual switch – clearing

A Manual switch command is removed by the operator by issuing a Clear command to the same Ethernet ring node where the Manual switch is presented. The Clear command removes existing local operator commands and triggers reversion in case the Ethernet ring is in revertive behaviour mode.

The Ethernet ring node where the Manual switch was cleared shall keep the ring port blocked for traffic channel and for the R-APS channel (as described in clause 10.1.14), due to the previous Manual switch command. This ring port is kept blocked until the RPL is blocked as a result of Ethernet ring protection reversion, or until there is another higher priority request (e.g., an SF condition) in the Ethernet ring.

The Ethernet ring node where the Manual switch was cleared continuously transmits the R-APS (NR) message on both ring ports, informing that no request is present at the Ethernet ring node. The

Ethernet ring nodes stop the transmission of R-APS (NR) messages when they accept an R-APS (NR, RB) message, or when another higher priority request is received.

If the Ethernet ring node where the Manual switch was cleared receives an R-APS (NR) message with a node ID higher than its own node ID, it unblocks any ring port which does not have an SF condition and stop the transmission of the R-APS (NR) message on both ring ports.

Revertive behaviour

Reversion is handled in the following way:

a) The RPL owner node, upon reception of an R-APS (NR) message and in the absence of any other higher priority request, starts the WTB timer and waits for expiration. While the WTB timer is running, any latent R-APS (MS) message is ignored due to the higher priority of the WTB Running signal.

b) When the WTB timer expires, it generates the WTB Expires signal. The RPL owner node, upon reception of the WTB Expires signal, initiates reversion by blocking the traffic channel on the RPL, transmitting an R-APS (NR, RB) message over both ring ports, informing the Ethernet ring that the RPL is blocked and performing a flush FDB action.

c) The acceptance of the R-APS (NR, RB*)* message causes all Ethernet ring nodes to unblock any blocked non-RPL that does not have an SF condition. If it is an R-APS (NR, RB*)* message without a DNF indication, all Ethernet ring nodes perform a necessary flush FDB action by following the mechanism described in clause 10.1.10. This action shall unblock the ring port which was blocked as a result of an operator command.

Non-revertive behaviour

Non-reversion is handled in the following way:

a) The RPL owner node, upon reception of an R-APS (NR) message and in the absence of any other higher priority request does not perform any action.

b) Then, after the operator issues a Clear command at the RPL owner node, this Ethernet ring node blocks the ring port attached to the RPL, transmits an R-APS (NR, RB) message over both ring ports, informing the Ethernet ring that the RPL is blocked and performs a flush FDB action.

c) The acceptance of the R-APS (NR, RB*)* message triggers all Ethernet ring nodes to unblock any blocked non-RPL which does not have an SF condition. If it is an R-APS (NR, RB*)* message without a DNF indication, all Ethernet ring nodes perform a necessary flush FDB action by following the mechanism described in clause 10.1.10. This action shall unblock the ring port which was blocked as result of an operator command.

10.2.5 Protection switching – Forced switch

An Ethernet ring with no request has a logical topology with the traffic channel blocked at the RPL and unblocked on all other ring links. In this situation, the operator-initiated Forced switch command triggers protection switching as follows:

a) The Ethernet ring node where a Forced switch command was issued blocks the traffic channel and R-APS channel (as described in clause 10.1.14) on the ring port to which the Forced switch command was issued. The Ethernet ring node shall unblock the other ring port.

b) The Ethernet ring node where the Forced switch command was issued transmits R-APS messages indicating FS over both ring ports. R-APS (FS) message shall be continuously transmitted by this Ethernet ring node while the local FS command is the Ethernet ring node's highest priority command. The R-APS (FS) message informs other Ethernet ring nodes of the FS command and that the traffic channel is blocked on one ring port.

c) An Ethernet ring node accepting an R-APS (FS) message, without any local higher priority requests unblocks any blocked ring port. This action unblocks the traffic channel over the RPL.

d) The Ethernet ring node accepting an R-APS (FS) message, without any local higher priority requests stops transmission of R-APS messages.

e) The Ethernet ring node receiving an R-APS (FS) message performs a necessary flush FDB action by following the mechanism described in clause 10.1.10.

Protection switching on a Forced switch request is completed when the above actions are performed by each Ethernet ring node. At this point, the conditions are created to allow the traffic flows to be steered around the Ethernet ring. From this point on the following rules apply regarding processing of further Forced switch commands:

a) While an existing Forced switch request is present in an Ethernet ring, any new Forced switch request is accepted, except for the Ethernet ring node having a prior local Forced switch request. The Ethernet ring nodes where further Forced switch commands are issued shall block the traffic channel and R-APS channel on the ring port to which the Forced switch was issued. The Ethernet ring node where the Forced switch command was issued transmits an R-APS message indicating FS over both ring ports. R-APS (FS) message shall be continuously transmitted by this Ethernet ring node while local FS command is the Ethernet ring node's highest priority command. As such, two or more Forced switches are allowed in the Ethernet ring. This may cause the segmentation of an Ethernet ring. It is the responsibility of the operator to prevent this effect if it is undesirable.

10.2.5.1 Forced switch – clearing

A Forced switch command is removed by the operator by issuing a Clear command to the same Ethernet ring node where the Forced switch is presented. The Clear command removes existing local operator commands and triggers reversion in case the Ethernet ring is in revertive behaviour mode.

The Ethernet ring node where the Forced switch was cleared shall keep the ring port blocked for traffic channel and for the R-APS channel (as described in clause 10.1.14), due to the previous Forced switch command. This ring port is kept blocked until the RPL is blocked as a result of Ethernet ring protection reversion, or until there is another higher priority request (e.g., an SF condition) in the Ethernet ring.

The Ethernet ring node where the Forced switch was cleared continuously transmits the R-APS (NR) message on both ring ports, informing that no request is present at the Ethernet ring node. The Ethernet ring nodes stop the transmission of R-APS (NR) messages when they accept an R-APS (NR, RB) message, or when another higher priority request is received.

If the Ethernet ring node where the Forced switch was cleared receives an R-APS (NR) message with a node ID higher than its own node ID, it unblocks any ring port which does not have an SF condition and stops the transmission of the R-APS (NR) message over both ring ports.

Revertive behaviour

Reversion is handled in the following way:

a) The reception of an R-APS (NR) message causes the RPL owner node to start the WTB timer.

b) The WTB timer is cancelled if during the WTB period a higher priority request than NR is accepted by the RPL owner node or is declared locally at the RPL owner node.

c) When the WTB timer expires, in the absence of any other higher priority request, the RPL owner node initiates reversion by blocking the traffic channel over the RPL, transmitting an

R-APS (NR, RB) message over both ring ports, informing the Ethernet ring that the RPL is blocked and performing a flush FDB action.

d) The acceptance of the R-APS (NR, RB*)* message causes all Ethernet ring nodes to unblock any blocked non-RPL that does not have an SF condition. If it is an R-APS (NR, RB) message without a DNF indication, all Ethernet ring nodes perform a necessary flush FDB action by following the mechanism described in clause 10.1.10. This action shall unblock the ring port which was blocked as a result of an operator command.

Non-revertive behaviour

Non-reversion is handled in the following way:

a) The RPL owner node, upon reception of an R-APS(NR) message and in the absence of any other higher priority request does not perform any action.

b) Then, after the operator issues a Clear command at the RPL owner node, this Ethernet ring node blocks the ring port attached to the RPL, transmits an R-APS (NR, RB) message on both ring ports, informing the Ethernet ring that the RPL is blocked and performs a flush FDB action.

c) The acceptance of the R-APS (NR, RB) message triggers all Ethernet ring nodes to unblock any blocked non-RPL which does not have an SF condition. If it is an R-APS (NR, RB) message without a DNF indication, all Ethernet ring nodes perform a necessary flush FDB action by following the mechanism described in clause 10.1.10. This action shall unblock the ring port which was blocked as result of an operator command.

10.3 R-APS format

R-APS information is carried in an R-APS PDU, which is one of a suite of Ethernet OAM messages. The OAM PDU format for each type of Ethernet OAM operation is defined in [ITU-T G.8013]. R-APS specific information is transmitted within specific fields in an R-APS PDU. An R-APS PDU is identified by the Ethernet OAM OpCode 40.

The R-APS messages will use the MAC address range allocated within ITU OUI for ITU-T G.8032 R-APS communication. The last octet of the MAC address is designated as ring ID (01-19-A7-00-00-[Ring ID]). The default ring ID is 01.

In this Recommendation, 32 octets in an R-APS message are used to carry R-APS specific information. This is illustrated in Figure 10-2 below. In addition, the TLV Offset field is required to be set to 32.

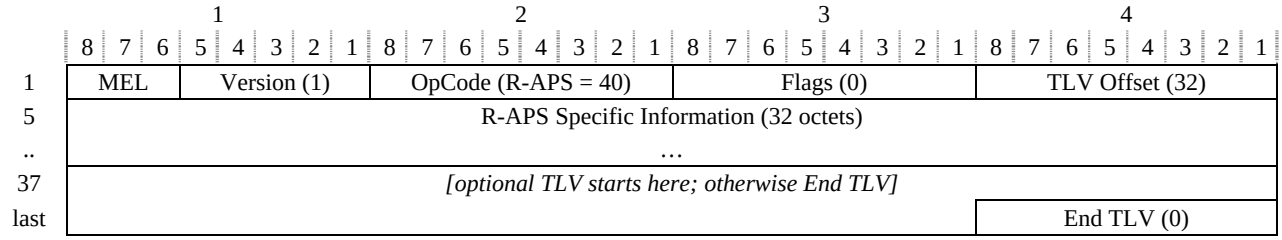

Figure 10-2 – R-APS PDU format

For other fields such as Version, OpCode, Flags and End TLV, the following values shall be used, as defined in [ITU-T G.8013].

a) **Version**: 0x01 shall be transmitted in the current version of this Recommendation.

b) **OpCode**: 40 shall be transmitted as defined in [ITU-T G.8013].

c) **Flags**: 0x00 shall be transmitted in the current version of this Recommendation. This field should be ignored upon reception.

d) **TLV Offset:** 0x20 (=32) shall be transmitted.

e) **End TLV**: 0x00 shall be transmitted.

This Recommendation does not define any R-APS specific TLVs.

In the MEL field, the MEG level at which the R-APS PDU is inserted.

The format of the R-APS specific information within each R-APS PDU is defined as per the following Figure 10-3:

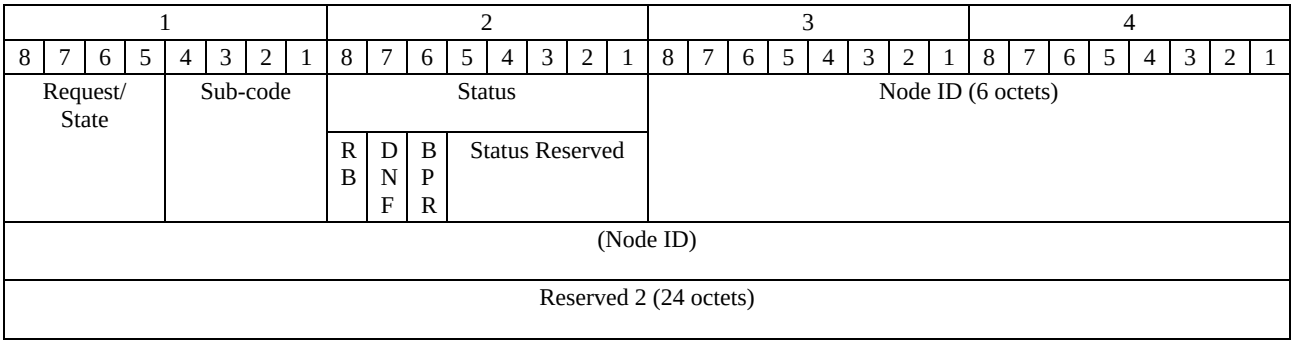

Figure 10-3 – R-APS specific information format

The fields of R-APS specific information:

a) Request/state (4 bits) – This field represents a request or state and is encoded as described in Table 10-3

Table 10-3 – Request/state values

Field	Value	Description
Request/state	1101	Forced switch
	1110	Event
	1011	Signal fail (SF)
	0111	Manual switch (MS)
	0000	No request (NR)
	Other	Reserved for future international standardization

b) Sub-code – encoding for sub-code for some of the request/states defined in the request/state field.

1) If Request/State Field = "1110" Event

I. Sub-code = "0000" – Flush Request.

II. Other values are reserved for future use.

2) For other request/state field coded the sub-code is transmitted as "0000" and ignored upon reception.

c) Status field – This includes the following status information:

1) RB – RPL Blocked

I. RB = 1 – represents that the RPL is blocked.

II. RB = 0 – represents that the RPL is unblocked.

This bit should be 0 when transmitted by non-RPL owner nodes.

2) DNF – Do Not Flush

   I. DNF = 1 – represents that an FDB flush should not be triggered by the reception of this message.

   II. DNF = 0 – represents that an FDB flush may be triggered by the reception of this message.

3) BPR – Blocked port reference

   I. BPR = 0 corresponds to ring link 0 blocked.

   II. BPR = 1 corresponds to ring link 1 blocked.

   This bit shall be set to 0 on messages transmitted from interconnection nodes on sub-ring's Ethernet ring nodes.

   If two ring links are blocked, the encoded value is can be either value.

4) Status reserved (5 bits) – For future specification. This field shall be transmitted encoded all zeroes. This field should be ignored upon reception.

d) Node ID (6 octets) – A MAC address unique to the Ethernet ring node.

e) Reserved 2 (24 octets) – This field is reserved for future extensions of the R-APS protocol. In the current version of this Recommendation, this field shall be transmitted encoded all zeroes. This field should be ignored upon reception.

10.4 Failure of protocol defect

Due to errors in provisioning, the ERP control process may detect a combination of conditions which should not occur during "normal" conditions. To warn the operator of such an event, a failure of protocol – provisioning mismatch (FOP-PM) is defined. The FOP-PM defect, detected if the RPL owner node receives one or more No request R-APS message(s) with the RPL blocked status flag set (NR, RB) and a node ID that differs from its own. The ERP control process must notify the equipment fault management process when it detects such a defect condition and continues its operation as well as possible. This is only an overview of the defect condition. The associated defect and its details are defined in [ITU-T G.8021].

The ERP control process must notify the equipment fault management process using the failure of protocol – time out (dFOP-TO) defect signal (as defined in [ITU-T G.8021]) if it fails to receive any R-APS messages on a ring port for a period exceeding K message cycles, as described in [ITU-T G.8021]. This defect signal should not be reported if the ring port is reporting a link level failure, or is either administratively disabled, or blocked from R-APS message reception. Some examples of these exceptions would be:

- SF reported on this ring port
- sub-ring ports when running sub-ring without virtual channel model
- when both ends of RPL blocked.

The ERP control process should continue its operation as well as possible.

This is only an overview of the defect condition. The mechanism for detection and clearance of dFOP-TO is defined in [ITU-T G.8021]. Such notification should allow the operator to take any proper corrective action. Such corrective action might, for example, include performing a Manual switch to allow resetting the Ethernet ring node to re-activate the ERP control process.

Appendix I

Ring protection network objectives

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

The following are the network objectives of the Ethernet ring protection.

I.1 The Ethernet ring protection mechanism shall prevent the creation of loops in an Ethernet ring topology under any circumstances (starting up the network, failure condition and switchover).

I.2 The ETH layer connectivity of ring links should be periodically monitored.

I.3 The ring link ETH layer monitoring should inform the Ethernet ring protection mechanism of SF or SD conditions (e.g., link bandwidth degradation and excessive error).

I.4 Server layer SF and SD conditions should be informed to the Ethernet ring protection mechanism.

Service restoration

I.5 Ethernet ring protection shall not contend with the protection mechanisms of the server layer.

General

I.6 The ring shall successfully recover multipoint connectivity in the event of a single ring link failure.

I.7 The ring shall successfully recover multipoint connectivity in the event of a single node failure, except for the traffic at that Ethernet ring node.

I.8 In the event of more than a single failure (e.g., of ring links or Ethernet ring nodes), the result should be ring segmentation with full connectivity within each segment.

I.9 Ethernet ring protection shall operate under all network load conditions.

I.10 Ethernet ring protection shall be independent of the capability of the server layer.

I.11 Ethernet ring protection shall support protection over multi-ring/ladder networks.

a) The protection mechanism shall enable the interconnection of rings using a single or dual Ethernet ring nodes. The mechanism shall protect services that are traversing interconnected rings. In the case of interconnected rings using dual Ethernet ring nodes, the mechanism shall ensure that a super loop is not formed in the event that there is ring link failure between interconnection nodes.

I.12 Ethernet ring protection control communication shall be performed using standard Ethernet messages (IEEE 802.3/802.1). The control messages of the Ethernet ring protection mechanism shall use the OAM message format defined in [ITU-T G.8013]. The OAM messages defined in [ITU-T G.8013] may be extended to support the protection control messages.

I.13 The protection process shall be deterministic. All Ethernet ring nodes in the Ethernet ring shall have the same view of the protection state.

I.14 The total communication bandwidth consumed by the protection mechanism shall be a very small fraction of the total available bandwidth and shall be independent of the total traffic supported by the network.

I.15 The protection mechanism shall not impose any limitation or requirements on the Ethernet relay and filtering function.

I.16 The mechanism should not impose any limitation on the number of Ethernet ring nodes that may form the Ethernet ring. From an operational perspective, the maximum number of Ethernet ring nodes supported should be in the range of 16 to 255 Ethernet ring nodes.

I.17 A switchover may be administratively triggered.

I.18 Revertive mode shall be supported.

I.19 Non-revertive mode should be supported.

I.20 In the event of a single Ethernet ring node or link failure, Ethernet ring protection shall support protection switching time (i.e., transfer time, Tt in clause 13 of [ITU-T G.808.1]) of no more than 50 ms.

I.21 Ethernet ring protection may support configurable hold-off times before triggering protection operation.

I.22 Ethernet ring protection may support configurable wait-to-restore times.

I.23 In the event of reversion, Ethernet ring protection shall support a revertive switching time (i.e., transfer time, Tt in clause 13 of [ITU-T G.808.1]) of no more than 50 ms.

I.24 In the event of administratively triggered switchover, Ethernet ring protection shall support a switching time (i.e., transfer time, Tt in clause 13 of [ITU-T G.808.1]) of no more than 50 ms.

I.25 The solution adopted for interconnected Ethernet rings, shall allow the operation of transforming one Ethernet ring into a sub-ring interconnected to another Ethernet ring without decommissioning the services already supported on the first Ethernet ring. It is acceptable that this operation may result in temporary traffic interruption due to protection switching events that result from reconfiguration of the Ethernet rings. It is also acceptable that during the operation, new link failures are not correctly protected.

Appendix II

Ethernet ring network objectives

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

The following are Ethernet ring network objectives:

II.1 An Ethernet ring shall be constructed from a set of Ethernet ring nodes, as defined in clause 3.2.1, which form a ring topology (i.e., a ring).

II.2 Traffic forwarding in an Ethernet ring and between a non-ring port and a ring port shall be based entirely on the forwarding rules defined by the IEEE 802.1 specifications.

II.3 Each Ethernet ring node shall have exactly two ring ports per logical ring.

II.4 The Ethernet ring nodes shall be connected in a closed loop.

II.5 The Ethernet ring shall provide direct or indirect communication between all Ethernet ring nodes in the Ethernet ring.

II.6 In Ethernet ring topology, each Ethernet ring node shall be connected to two other Ethernet ring nodes utilizing ring ports based on IEEE 802.3 MAC.

II.7 The Ethernet MAC may be transported over any server layer.

a) The Ethernet ring shall not preclude the use of any transport technology (e.g., SDH VCs using GFP mapping, Ethernet physical layer interfaces ETY, MPLS ETH pseudo-wires, Ethernet link aggregation [IEEE 802.3].

b) The capacity of each span in the ring (link) is dependent on the transport technology used. It shall not be a requirement that all ring links need to provide the same capacity.

II.8 The definition of an Ethernet ring shall be applicable to both physical ring topologies and logical ring topologies. Note these are not independent.

II.9 Shall support increased bandwidth utilization via concurrent transmissions, spatial reuse.

II.10 Shall utilize [ITU-T G.8013], [IEEE 802.1Q] and may use other Ethernet OAM specifications.

II.11 Each Ethernet ring node shall support MAC services and QoS according to the [IEEE 802.1Q] specification. The use of Ethernet ring resources at each ring link is controlled by the same rules.

II.12 Ethernet rings shall support E-Line, E-LAN and E-Tree services including EPL [b-ITU-T G.8011.1] and EVPL [b-ITU-T G.8011.2]

II.13 Ethernet ring topology shall support all types of communication: unicast, multicast and broadcast.

II.14 Normal Ethernet ring behaviour (i.e., without protection) shall prevent mis-ordering and/or duplication of transported client messages.

II.15 End-to-end services may traverse multiple interconnected rings.

II.16 Ethernet rings may be interconnected through an interconnection point (as depicted in Figure II.1), or through dual interconnection nodes with a ring link (as depicted in Figure II.2) or a multi-ring/ladder network that consists of conjoined Ethernet rings (as depicted in Figure II.3).

II.17 The logical rings shall be identifiable for management purposes.

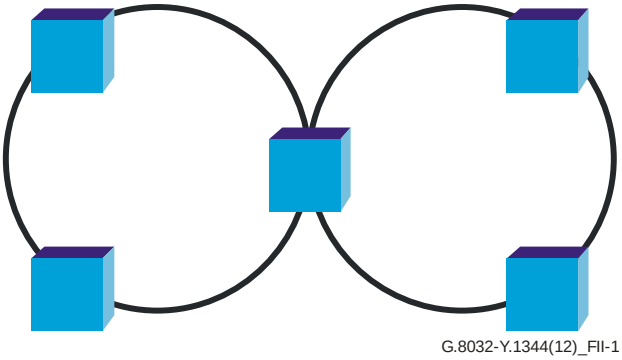

Figure II.1 – Interconnected Ethernet rings via an interconnection node

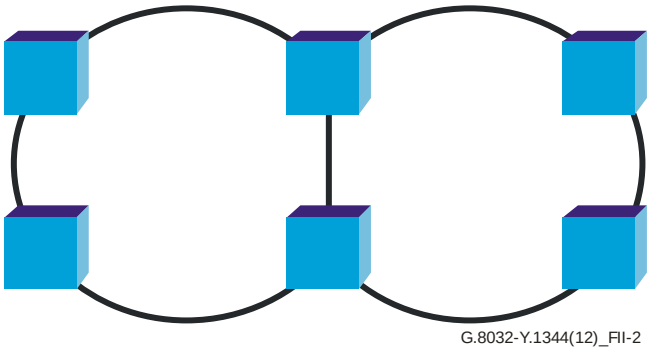

Figure II.2 – Interconnected Ethernet rings via dual Ethernet ring nodes with a ring link

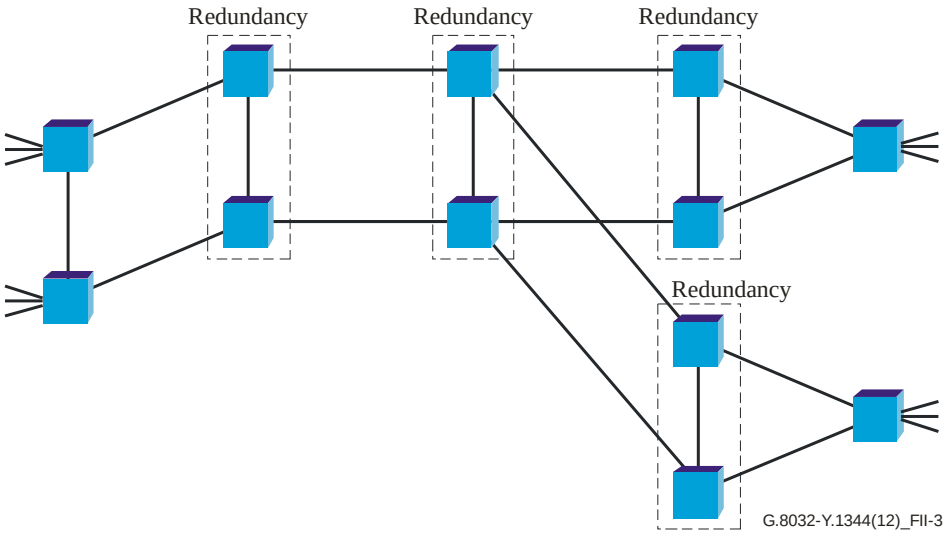

Figure II.3 – Example of multi-ring/ladder network

Appendix III

Ring protection scenarios

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

The following scenarios represent an Ethernet ring composed of seven Ethernet ring nodes. The RPL is the ring link between Ethernet ring nodes A and G. In these scenarios, both ends of the RPL are blocked. Ethernet ring node G is the RPL owner node and Ethernet ring node A is the RPL neighbour node.

NOTE – The scenarios described in Recommendation ITU-T G.8032 (2008) are fully supported by this version of this Recommendation. The following scenarios (that may extend the functionality described in previous versions) are also supported by this version of this Recommendation.

NOTE – In all of the following scenarios that show a <Node ID, BPR> pair, the node ID should be taken as a logical ID that is mapped to an actual node ID.

The following symbols are used:

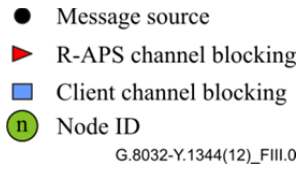

Scenario A – Single link failure

The following scenario represents protection switching in case of a single link failure.

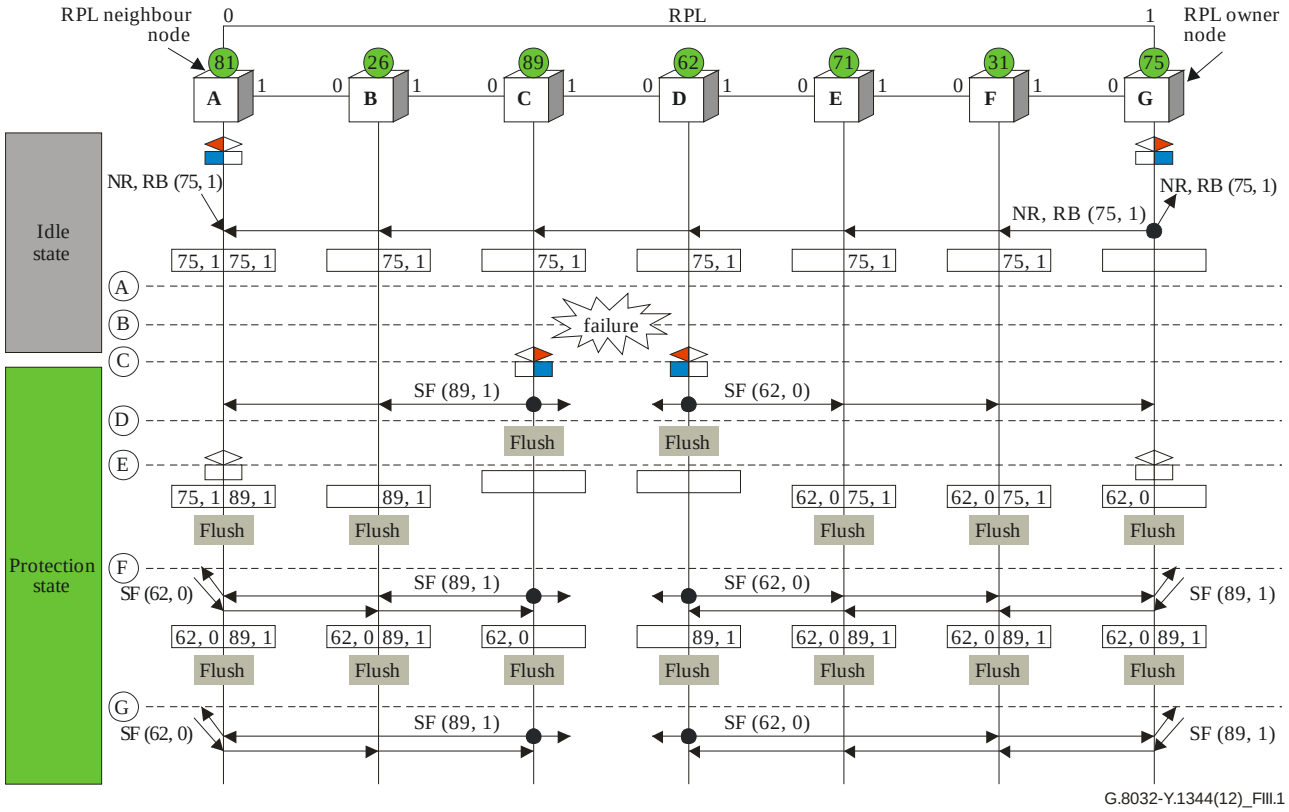

Figure III.1 – Single link failure

The following sequence describes the steps in the previous figure:

A. Normal condition.

B. Failure occurs.

C. Ethernet ring nodes C and D detect a local signal failure condition and after respecting the hold-off time, block the failed ring port and perform the FDB flush.

D. Ethernet ring nodes C and D start sending R-APS (SF) messages periodically with the (node ID, BPR) pair on both ring ports, while the SF condition persists.

E. All Ethernet ring nodes receiving an R-APS (SF) message perform an FDB flush. When the RPL owner node G and RPL neighbour node A receive an R-APS (SF) message, they each unblock their end of the RPL and perform the FDB flush.

F. All Ethernet ring nodes receiving a second R-APS (SF) message perform the FDB flush again due to the node ID and BPR-based mechanism.

G. Stable SF condition – R-APS (SF) messages on the Ethernet ring. Further R-APS (SF) messages trigger no further action.

The following scenario represents reversion in case of a single link failure.

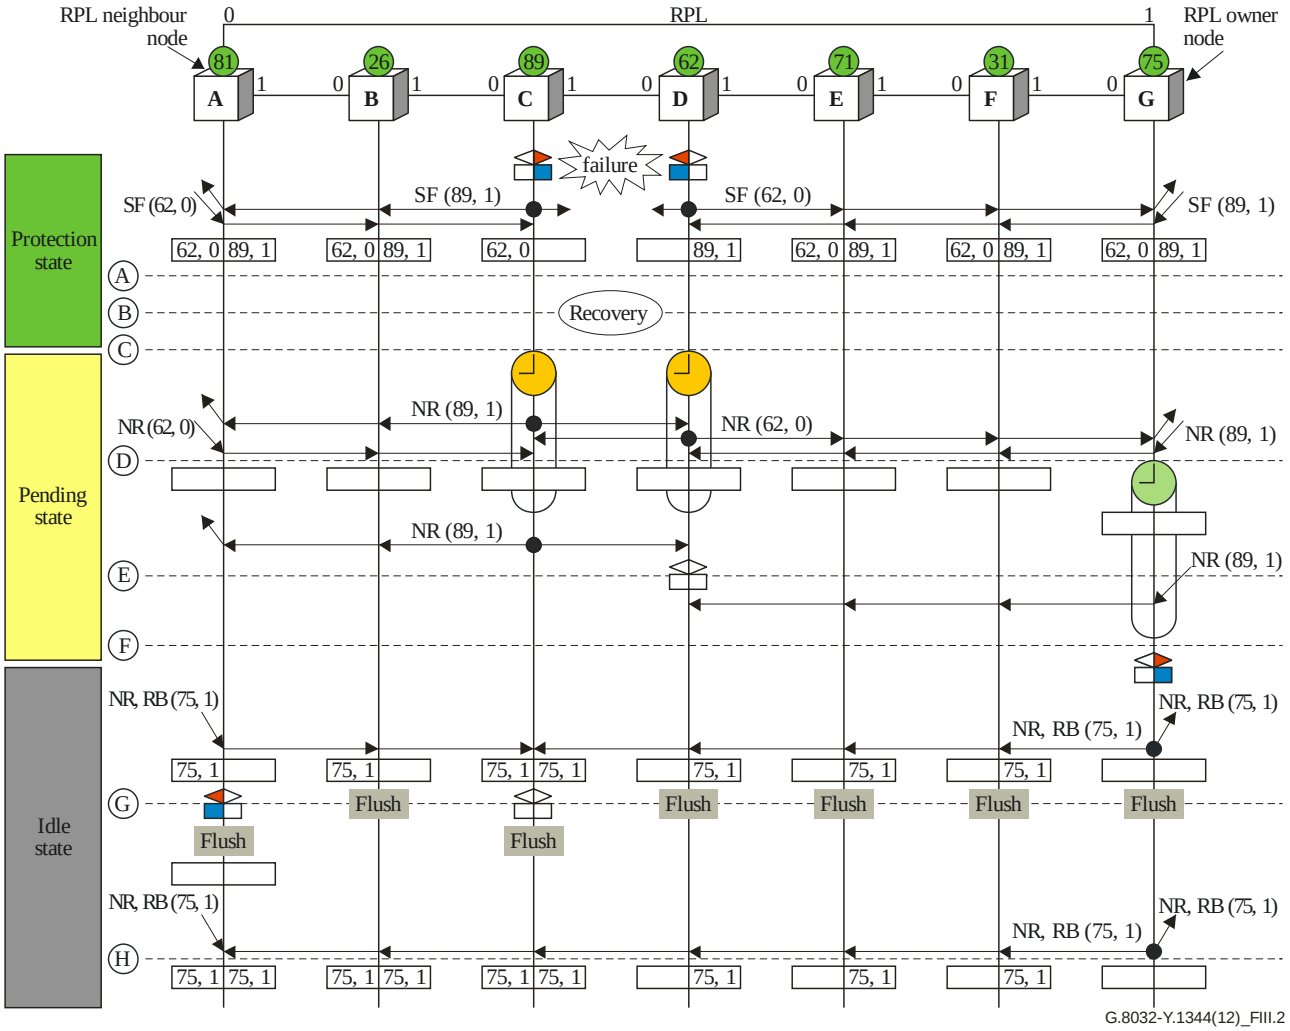

Figure III.2 – Single link failure recovery (revertive operation)

The following sequence describes the steps in Figure III.2:

A. Stable SF condition.

B. Recovery of link failure.

C. Ethernet ring nodes C and D detect clearing of SF condition, start the guard timer and initiate the periodical transmission of R-APS (NR) messages on both ring ports. (The guard timer prevents the reception of R-APS messages).

D. When the Ethernet ring nodes receive an R-APS (NR) message, the (node ID, BPR) pair of a receiving ring port is deleted and the RPL owner node starts the WTR timer.

E. When the guard timer expires on Ethernet ring nodes C and D, they may accept the new R-APS messages that they receive. Ethernet ring node D receives an R-APS (NR) message with a higher node ID from Ethernet ring node C and unblocks its non-failed ring port.

F. At expiration of the WTR timer, the RPL owner node blocks its end of the RPL, sends an R-APS (NR, RB) message with the (node ID, BPR) pair and performs the FDB flush.

G. When Ethernet ring node C receives an R-APS (NR, RB) message, it removes the block on its blocked ring ports and stops sending R-APS (NR) messages. On the other hand, when the RPL neighbour node A receives an R-APS (NR, RB) message, it blocks its end of the RPL. In addition to this, Ethernet ring nodes A to F perform the FDB flush when receiving an R-APS (NR, RB) message due to the node ID and BPR-based mechanism.

The following scenario represents the non-revertive operation in case of a single link failure.

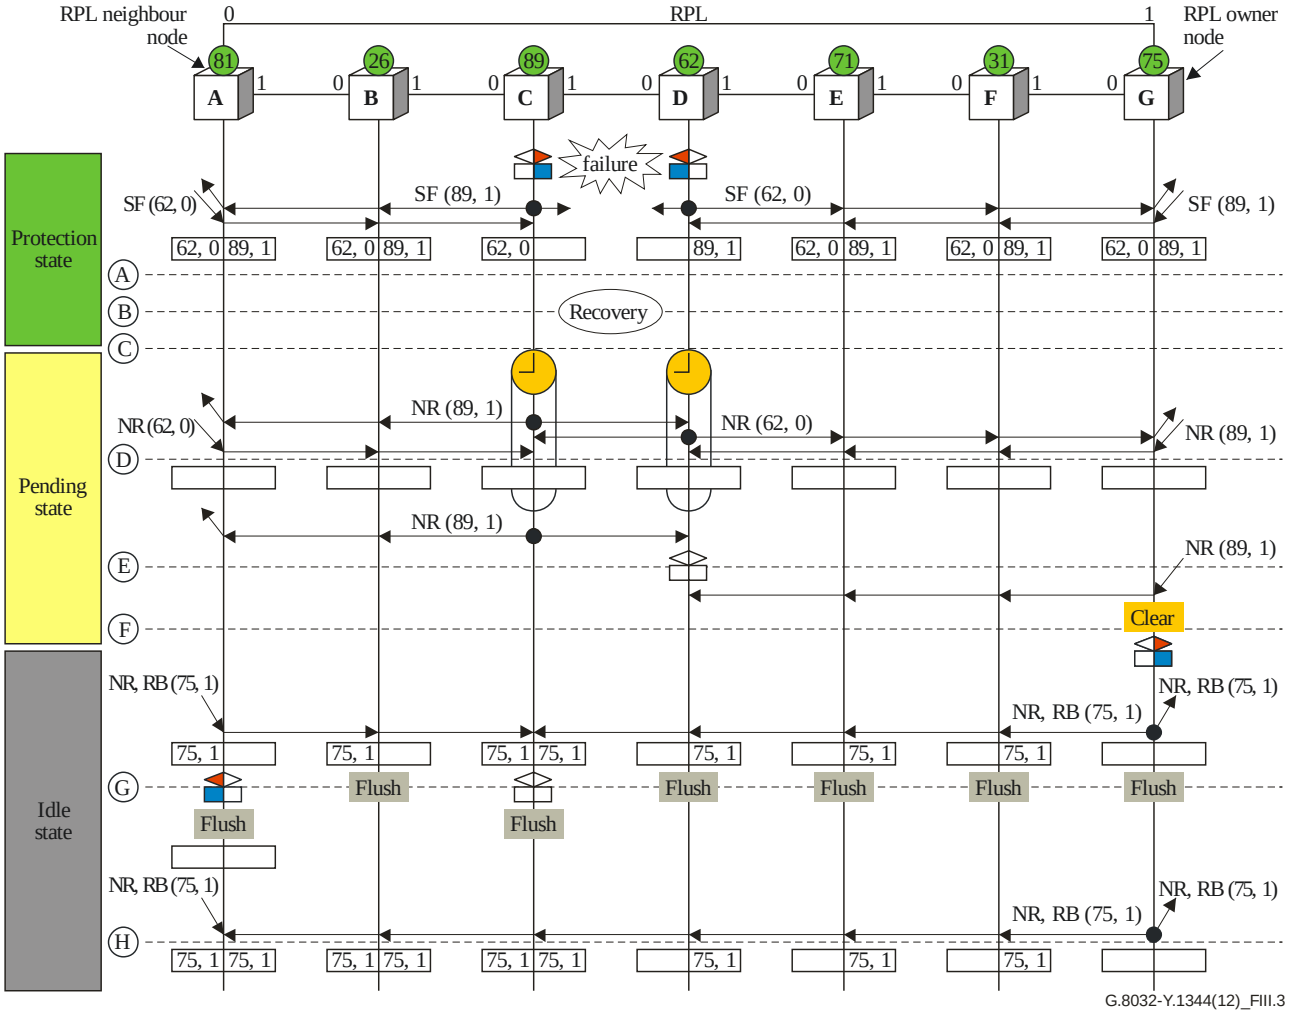

Figure III.3 – Single link failure recovery (non-revertive operation)

The following sequence describes the steps in Figure III.3:

A. Stable SF condition.

B. Recovery of link failure.

C. Ethernet ring nodes C and D detect clearing of SF condition, start the guard timer and initiate the periodical transmission of R-APS (NR) messages on both ring ports. (The guard timer prevents the reception of R-APS messages).

D. When the Ethernet ring nodes receive an R-APS (NR) message, the (node ID, BPR) pair of received ring port is deleted and the RPL owner node does not start the WTR timer.

E. When the guard timer expires on Ethernet ring nodes C and D, they may accept the new R-APS messages that they receive. Ethernet ring node D receives an R-APS (NR) message with a higher node ID from Ethernet ring node C and unblocks its non-failed ring port.

F. When the RPL owner node executes a Clear command, it blocks its end of the RPL, sends R-APS (NR, RB) message with the (node ID, BPR) pair and performs the FDB flush.

G. When Ethernet ring node C receives an R-APS (NR, RB) message, it removes the block on its blocked ring ports and stops sending R-APS (NR) messages. On the other hand, when the RPL neighbour node A receives an R-APS (NR, RB) message, it blocks its end of the RPL. In addition to this Ethernet ring nodes A to F perform the FDB flush when receiving an R-APS (NR, RB) message due to the node ID and BPR-based mechanism.

Scenario B – Single unidirectional link failure

This scenario is similar to scenario A with the difference that the link failure is unidirectional.

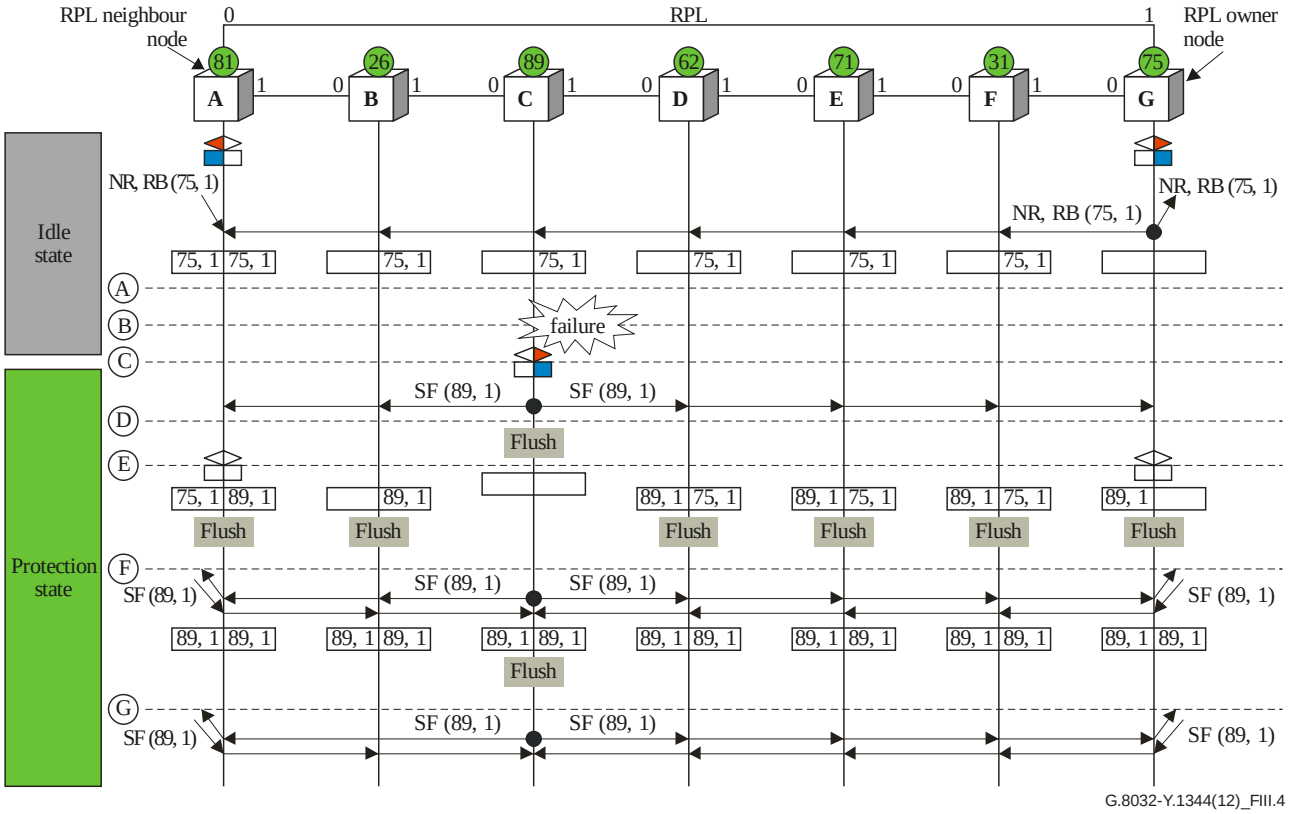

Figure III.4 – Single unidirectional link failure

The following sequence describes the steps in Figure III.4:

A. Normal condition.

B. Failure occurs in the direction of D to C, the direction of C to D is unaffected.

C. Ethernet ring nodes C detects a local signal failure condition and after respecting the hold-off period, blocks the failed ring port and performs the FDB flush (Ethernet ring node D performs no action).

D. Ethernet ring node C starts sending R-APS (SF) messages with the (node ID, BPR) pair on both ring ports, while the SF condition persists.

E. All Ethernet ring nodes receiving an R-APS (SF) message perform the FDB flush. When the RPL owner node G and RPL neighbour node A receive an R-APS (SF) message, they unblock their end of the RPL and perform the flush FDB.

F. Ethernet ring node C – on receiving a second R-APS (SF) message performs the FDB flush again due to the node ID and BPR-based mechanism.

G. Stable SF condition – R-APS (SF) messages on the Ethernet ring. Further R-APS (SF) messages trigger no further action.

The reversion for the unidirectional case is represented by Figure III.5.

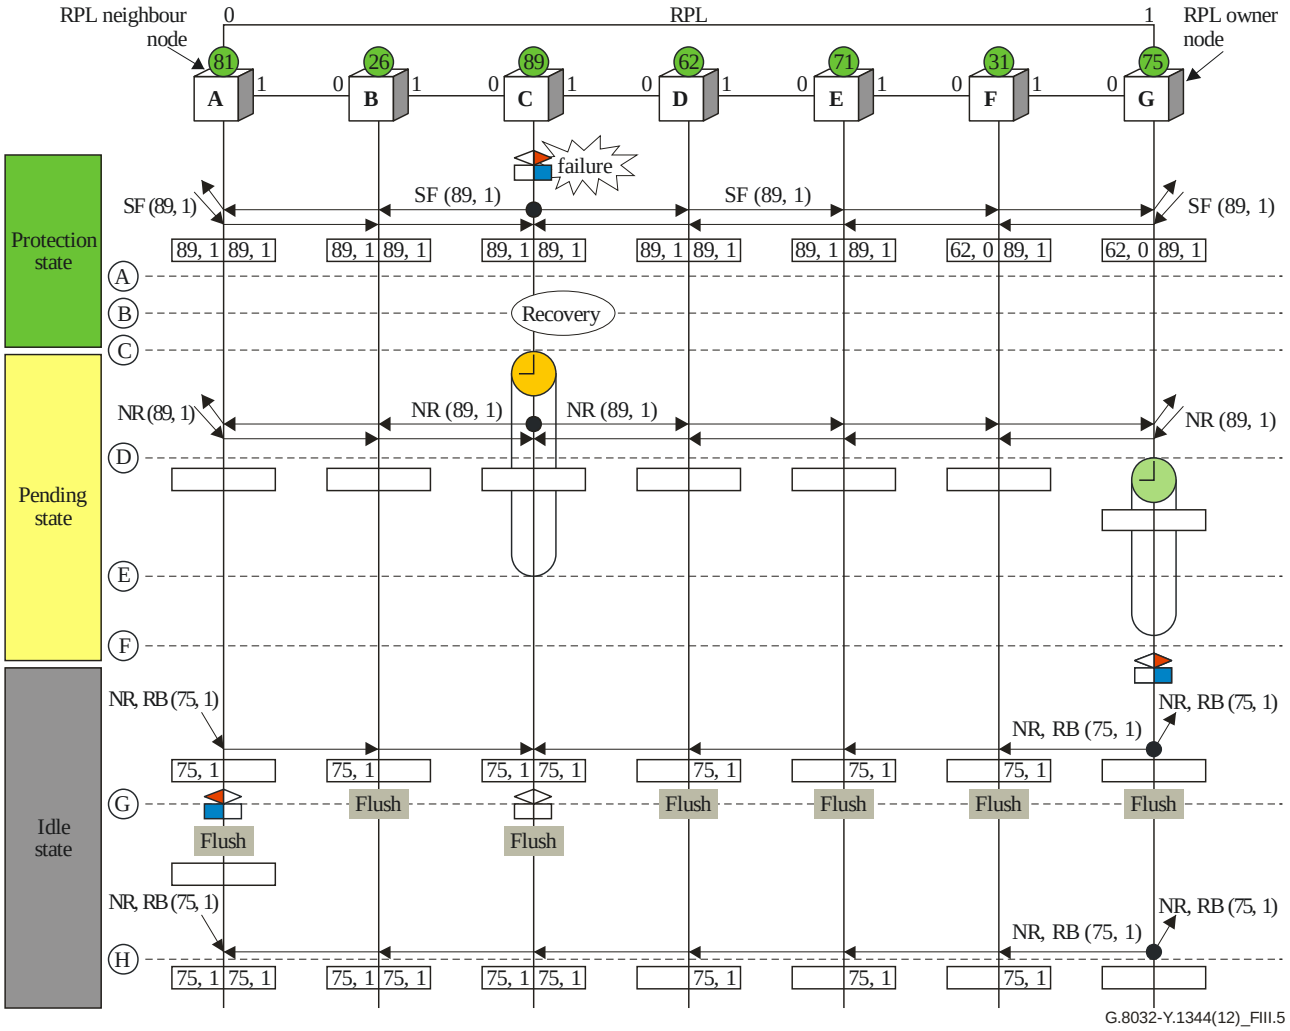

Figure III.5 – Single link failure unidirectional recovery

The following sequence describes the steps in Figure III.5:

A. Stable SF condition.

B. Recovery of link failure.

C. Ethernet ring node C detects the clearing of the SF condition, starts the guard timer and initiates the periodical transmission of R-APS (NR) messages on both ring ports. (The guard timer prevents the reception of R-APS messages.)

D. When the Ethernet ring nodes receive an R-APS (NR) message, the (node ID, BPR) pair of the receiving ring port is deleted and the RPL owner node starts the WTR timer.

E. When the guard timer expires on Ethernet ring node C, it may accept the new R-APS messages that it receives.

F. At expiration of the WTR timer, the RPL owner node blocks its end of the RPL, sends R-APS (NR, RB) messages with the (node ID, BPR) pair and performs the FDB flush.

G. When Ethernet ring node C receives an R-APS (NR, RB) message, it removes the block on its blocked ring port and stops sending R-APS (NR) messages. On the other hand, when the RPL neighbour node A receives an R-APS (NR, RB) message, it blocks its end of the RPL. In addition to this, Ethernet ring nodes A to F perform the FDB flush when receiving an R-APS (NR, RB) message due to the node ID and BPR-based mechanism.

Scenario C – RPL failure

The following figure represents the behaviour in case of RPL failure and shows an example of the possible use of the DNF status bit.

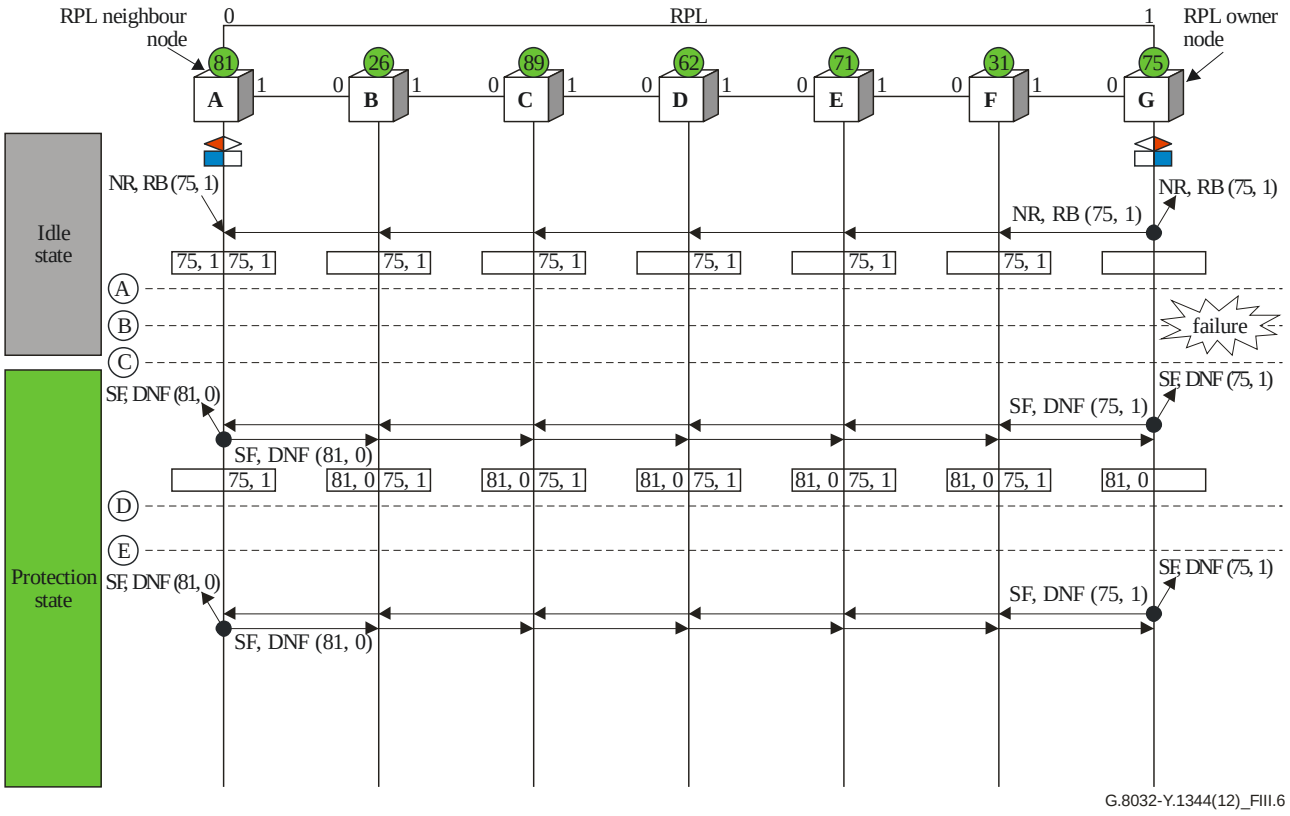

Figure III.6 – RPL failure

The following sequence describes the steps in Figure III.6:

A. Normal condition.

B. Failure occurs.

C. Ethernet ring nodes A and G detect a local SF condition and periodically start sending R-APS (SF) messages with the (node ID, BPR) pair on both ring ports, while the SF condition persists. The R-APS (SF) message includes a "do not flush" (DNF) indication and

this prevents all Ethernet ring nodes from performing the FDB flush, despite a transition from the idle to the protection state.

D. The RPL owner node receives an R-APS (SF) message, but it is ignored as there is a local higher priority request (local SF) [no transition]. All other Ethernet ring nodes receive the R-APS (SF) message with a DNF indication (flush is not performed), despite a transition from the idle to the protection state without flushing the FDB.

E. Stable SF condition – R-APS (SF) messages on the Ethernet ring with DNF indication. Further R-APS (SF) messages trigger no further action.

The actions after the repair of the RPL are represented in Figure III.7.

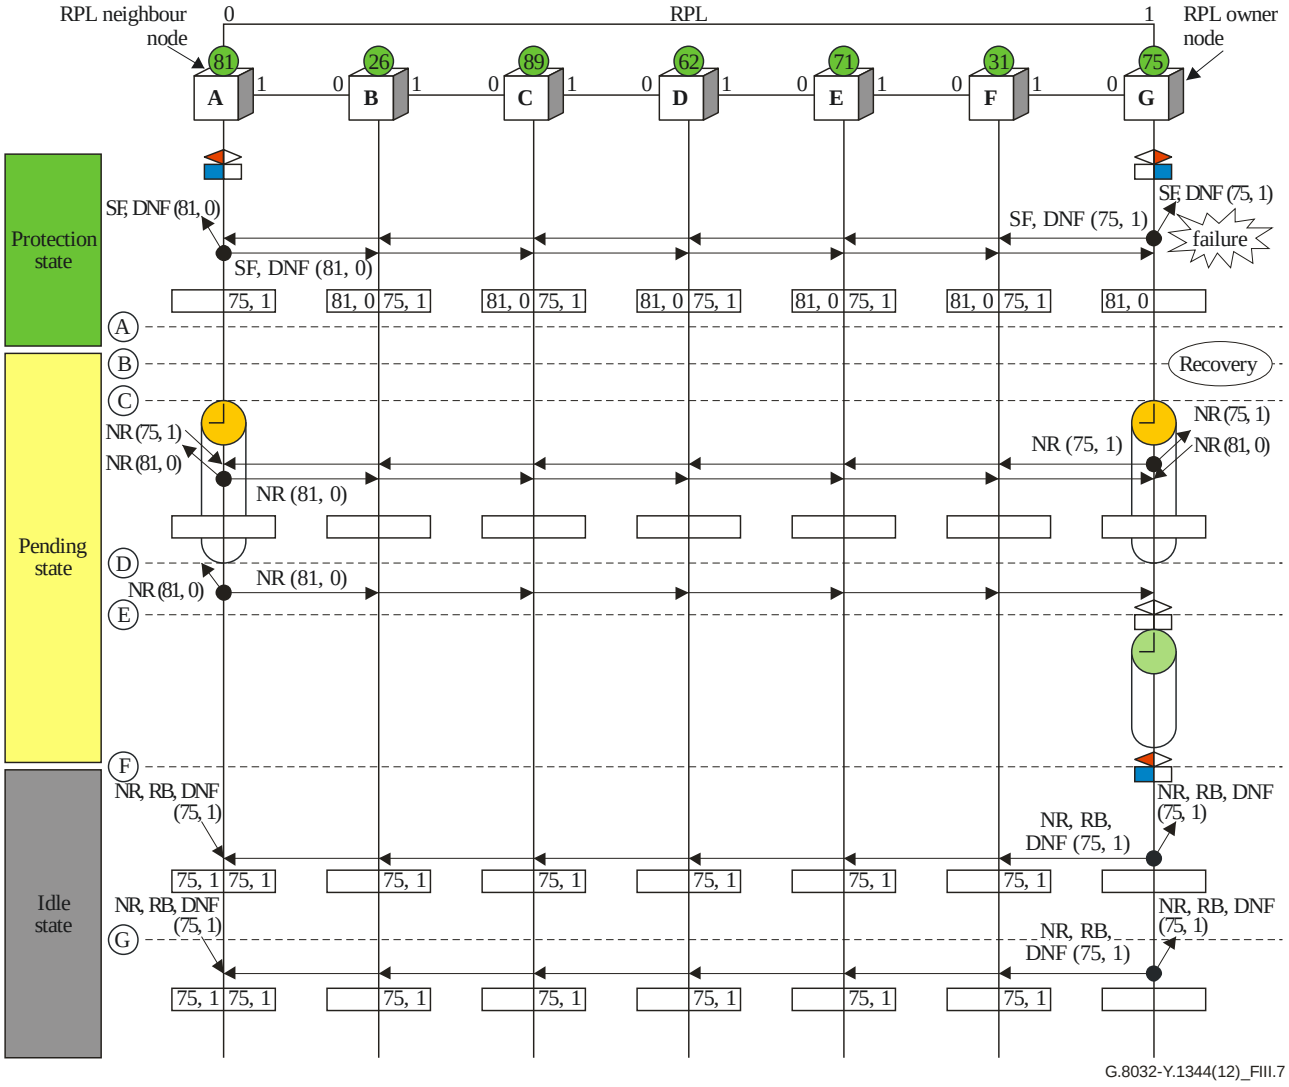

Figure III.7 – RPL failure recovery

The following sequence describes the steps in Figure III.7:

A. Stable SF condition.

B. Recovery of link failure.

C. Ethernet ring nodes A and G detect the clearing of the SF condition, start the guard timer and initiate the periodical transmission of R-APS (NR) messages on both ring ports. (The guard timer prevents the reception of R-APS messages).

D. When the guard timer expires on Ethernet ring nodes A and G, they may accept the new R-APS messages that they receive.

E. When the RPL owner node receives an NR message with a higher node ID, it unblocks the non-failed port and starts the WTR timer.

F. At expiration of the WTR timer, the RPL owner node blocks its end of the RPL (it was already blocked) and sends R-APS (NR, RB) messages. This message includes a "DNF" indication and this prevents all Ethernet ring nodes from performing an FDB flush, despite a transition from the pending to the idle state.

G. When Ethernet ring node A receives an R-APS (NR, RB) message, it keeps blocking its RPL port and stops sending R-APS (NR) messages. All Ethernet ring nodes receiving this message do not perform the FDB flush as the R-APS (NR, RB) messages include the "DNF" indication, despite a transition from the pending to the idle state without flushing the FDB.

Scenario D – Multiple failure case – recovery

The following scenario represents the case of sequential repair of multiple failures. In this case the failures between Ethernet ring nodes A and B and between Ethernet ring nodes E and F recover almost simultaneously. The SF condition remains on the ring link between Ethernet ring nodes C and D.

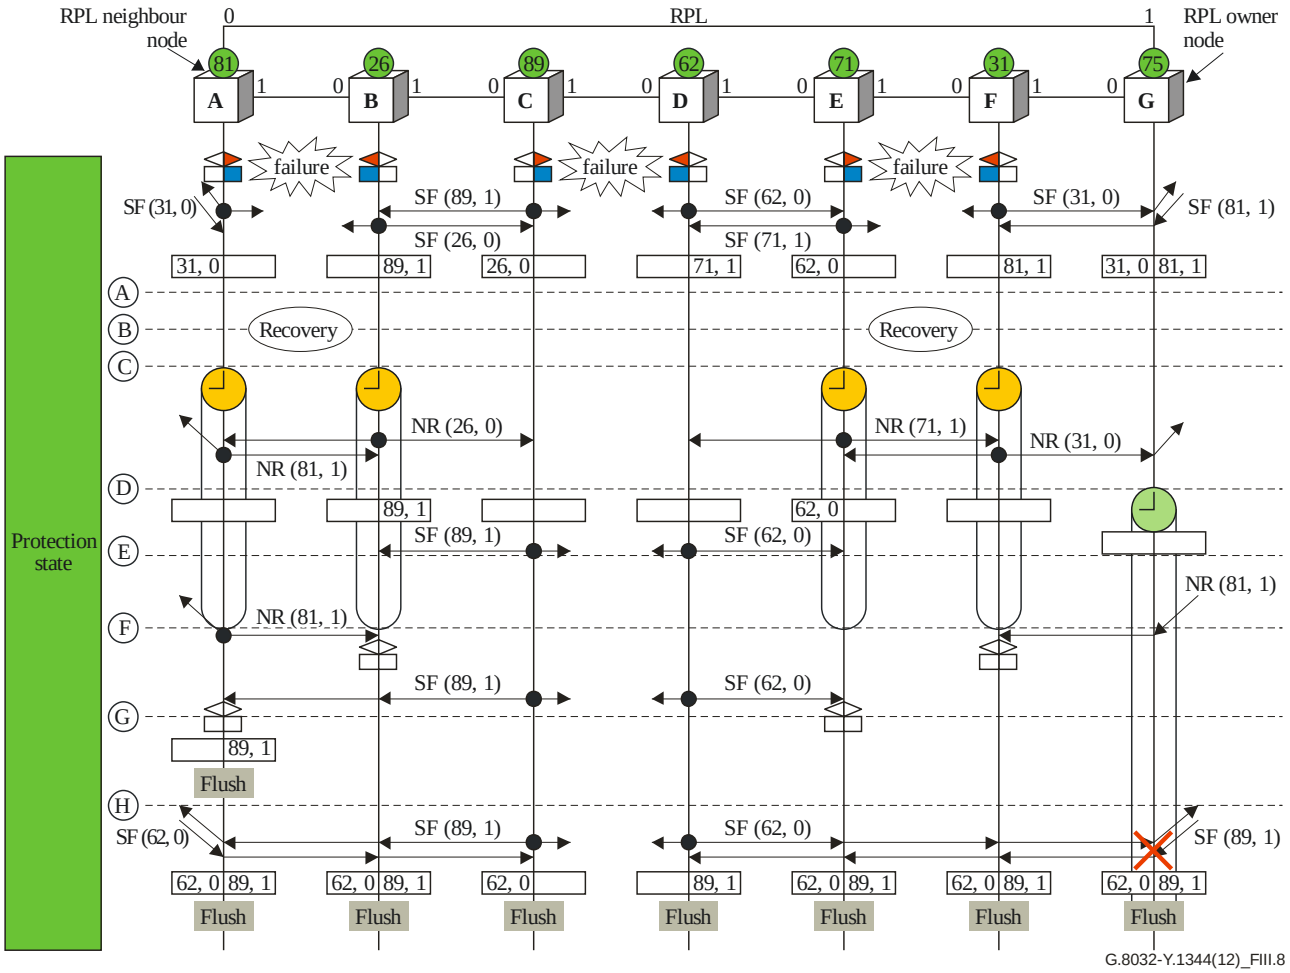

Figure III.8 – Multiple link failure

The following sequence describes the steps in Figure III.8:

A. Stable SF condition.

B. Recovery of link failures.

C. Ethernet ring nodes A, B, E and F detect the clearing of the SF condition, start the guard timer and initiate the periodical transmission of R-APS (NR) messages on both ring ports. The guard timer prevents the reception of R-APS messages, as is the case of an R-APS (SF) message transmitted by Ethernet ring nodes C and D, which are ignored by Ethernet ring nodes B and E.

D. When Ethernet ring nodes receive an R-APS (NR) message, the (node ID, BPR) pair on the receiving ring port is deleted and the RPL owner node starts the WTR timer.

E. Ethernet ring nodes B and E receiving an R-APS (SF) message do not perform the FDB flush due to the node ID and BPR-based mechanism.

F. When the guard timer expires on Ethernet ring nodes A, B, E and F, they may accept the new R-APS messages that they receive. The reception of an R-APS (NR) message with a higher node ID triggers unblocking of the blocked ring port and stops the transmission of R-APS (NR) messages at Ethernet ring nodes B and F.

G. The reception of an R-APS (SF) message triggers unblocking of the blocked ring port and stops the transmission of R-APS (NR) messages at Ethernet ring nodes A and E. Ethernet ring node A receiving an R-APS (SF) message performs FDB flush due to the node ID and BPR-based mechanism.

H. All Ethernet ring nodes receiving an R-APS (SF) message perform the FDB flush due to the node ID and BPR-based mechanism. The reception of an R-APS (SF) message informs the RPL owner node that an error is still present on the Ethernet ring. This results in the WTR timer being stopped.

NOTE – In rare cases where the link adjacent to the RPL owner is involved and recovers, the reversion process of this scenario may cause continued segmentation of the ring for the duration of the WTR/WTB timers running.

.

Appendix IV

Considerations for different timers

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

IV.1 State machine use of timers

There are four timers in this Recommendation – hold-off timer, guard timer, "wait to restore" (WTR) timer and wait to block (WTB) timer. These timers are described in clauses 10.1.8, 10.1.5, 10.1.4 and 10.1.4, respectively. According to Table 10-2, the different timers, except for the hold-off timer, are accessed (start or stop) in the following situations:

a) During initialization (row 1) – all timers are stopped to verify a clean situation.

b) During initialization (row 1) – the WTR timer is used by the RPL owner in revertive mode to verify that the node is stabilized before entering the idle state.

c) An Ethernet ring node that is recovering from an SF condition starts the guard timer (row 20).

d) An RPL owner node that is recovering from an SF condition starts the WTR timer (rows 20 and 29) – used to verify that the recovered SF is stabilized before reverting to the idle state.

e) An RPL owner node about to enter the pending state, after receiving an R-APS (NR) message, starts the WTB timer (rows 43 and 57) – used to cause the pending state to time out while the RPL owner node verifies that there are no additional live switching triggers in the Ethernet ring (e.g., two active FS conditions).

f) An Ethernet ring node that receives a Clear command (following an FS or MS) starts the guard timer (rows 30 and 44) – prior to entering the pending state to protect against stale R-APS messages.

g) An Ethernet ring node that has an MS command and receives an R-APS (MS) message from another Ethernet ring node in the Ethernet ring (row 36) starts the guard timer prior to entering the pending state.

h) An RPL owner node that has an MS command and receives either a Clear command or an R-APS (MS) message from another Ethernet ring node in the Ethernet ring (rows 30 and 36) starts the WTB timer prior to entering the pending state.

i) When the RPL owner node transits out of the pending state, it stops the WTR and WTB timers (rows 58, 59, 60, 61, 63, 64, 65, 66, 68 and 70).

IV.2 Guard timer use to block outdated R-APS messages

Two Ethernet ring nodes could transmit R-APS messages at the same time. In this case, the outdated R-APS message is transmitted by these Ethernet ring nodes until the Ethernet ring node receives the new R-APS message and it overwrites its state. For example, in Figure IV.1, Ethernet ring nodes A and B simultaneously detect local clear SF and start sending R-APS (NR) messages, and they transit to pending state [sequence B]. But soon after, they may receive an R-APS (SF) message from each other and unblock their recovered ring ports [sequence C]. Unblocking of non-failed ring ports at both Ethernet ring nodes may result in the formation of a loop. To avoid this, Ethernet ring nodes A and B need to discard the received R-APS message for a while. After this period, if they still receive the same R-APS (SF) message, they can properly identify the current SF condition. For this reason, a guard timer is mandatory to avoid forming a loop (rows 20, 30, 36, 44).

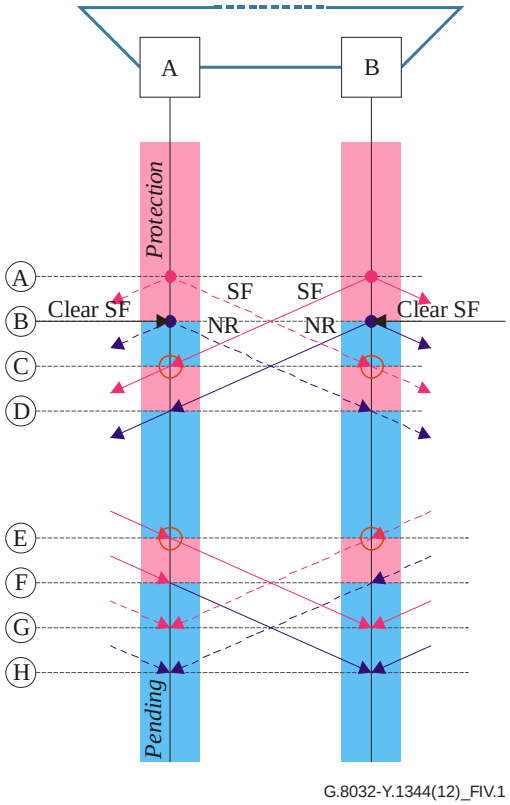

Figure IV.1 – Simultaneous requests from multiple Ethernet ring nodes

Appendix V

Interconnected rings example

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

V.1 Configuration for interconnected rings

The following Figure V.1 represents an example of a topology composed of two interconnected Ethernet rings. The lower Ethernet ring is a sub-ring. Figure V.2 represents an example of a topology composed of three interconnected Ethernet rings and the middle Ethernet ring is a sub-ring.

The R-APS channels of Ethernet rings A and B are consistent with the definition of this Recommendation.

When the sub-ring is operated with R-APS virtual channel, the R-APS channel of the sub-ring is complemented by the use of the R-APS virtual channel to enable R-APS channel connectivity between sub-ring ERP control processes of the two interconnection nodes. When the sub-ring is operated without an R-APS virtual channel, the R-APS channel of the sub-ring is terminated at the interconnection nodes as illustrated in Figure V.1 e).

The ring link between the two interconnection nodes is under the control of the ERP control processes of Ethernet rings A or B that are present at the interconnection nodes. These entities are responsible for triggering protection switching events upon the failure of this ring link and perform block and unblock operations for traffic on that ring link. The sub-ring is not aware of the existence.

The sub-ring is composed of at least one sub-ring link and one R-APS virtual channel in order to allocate the RPL on a sub-ring.

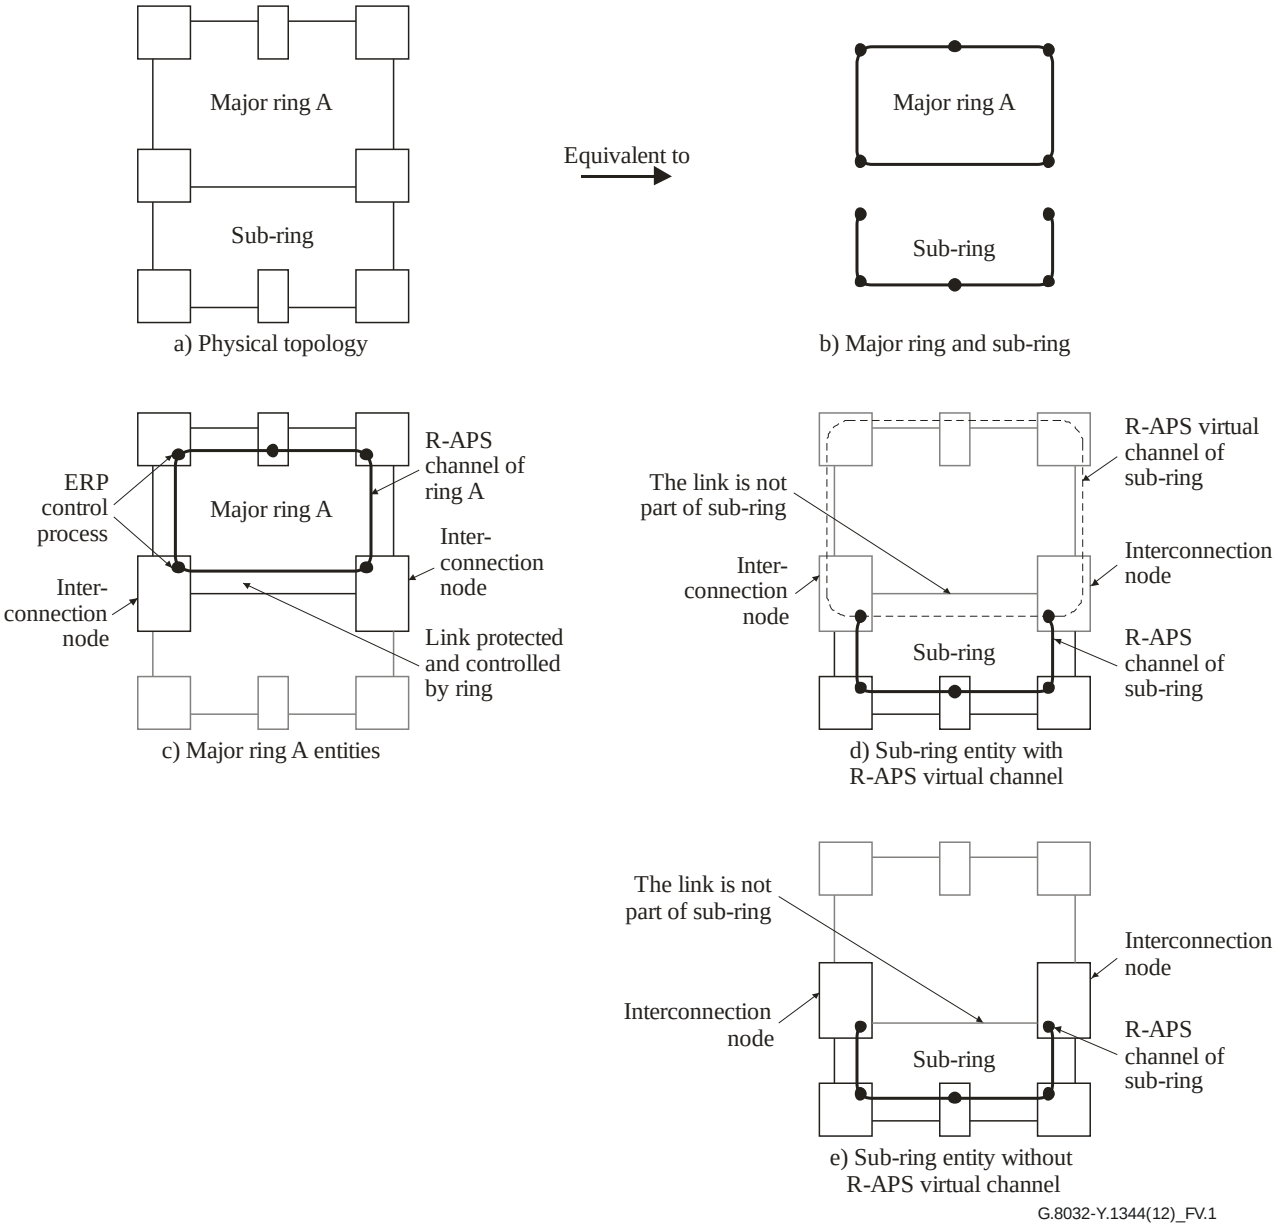

Figure V.1 – Configuration for interconnection between a major ring and a sub-ring

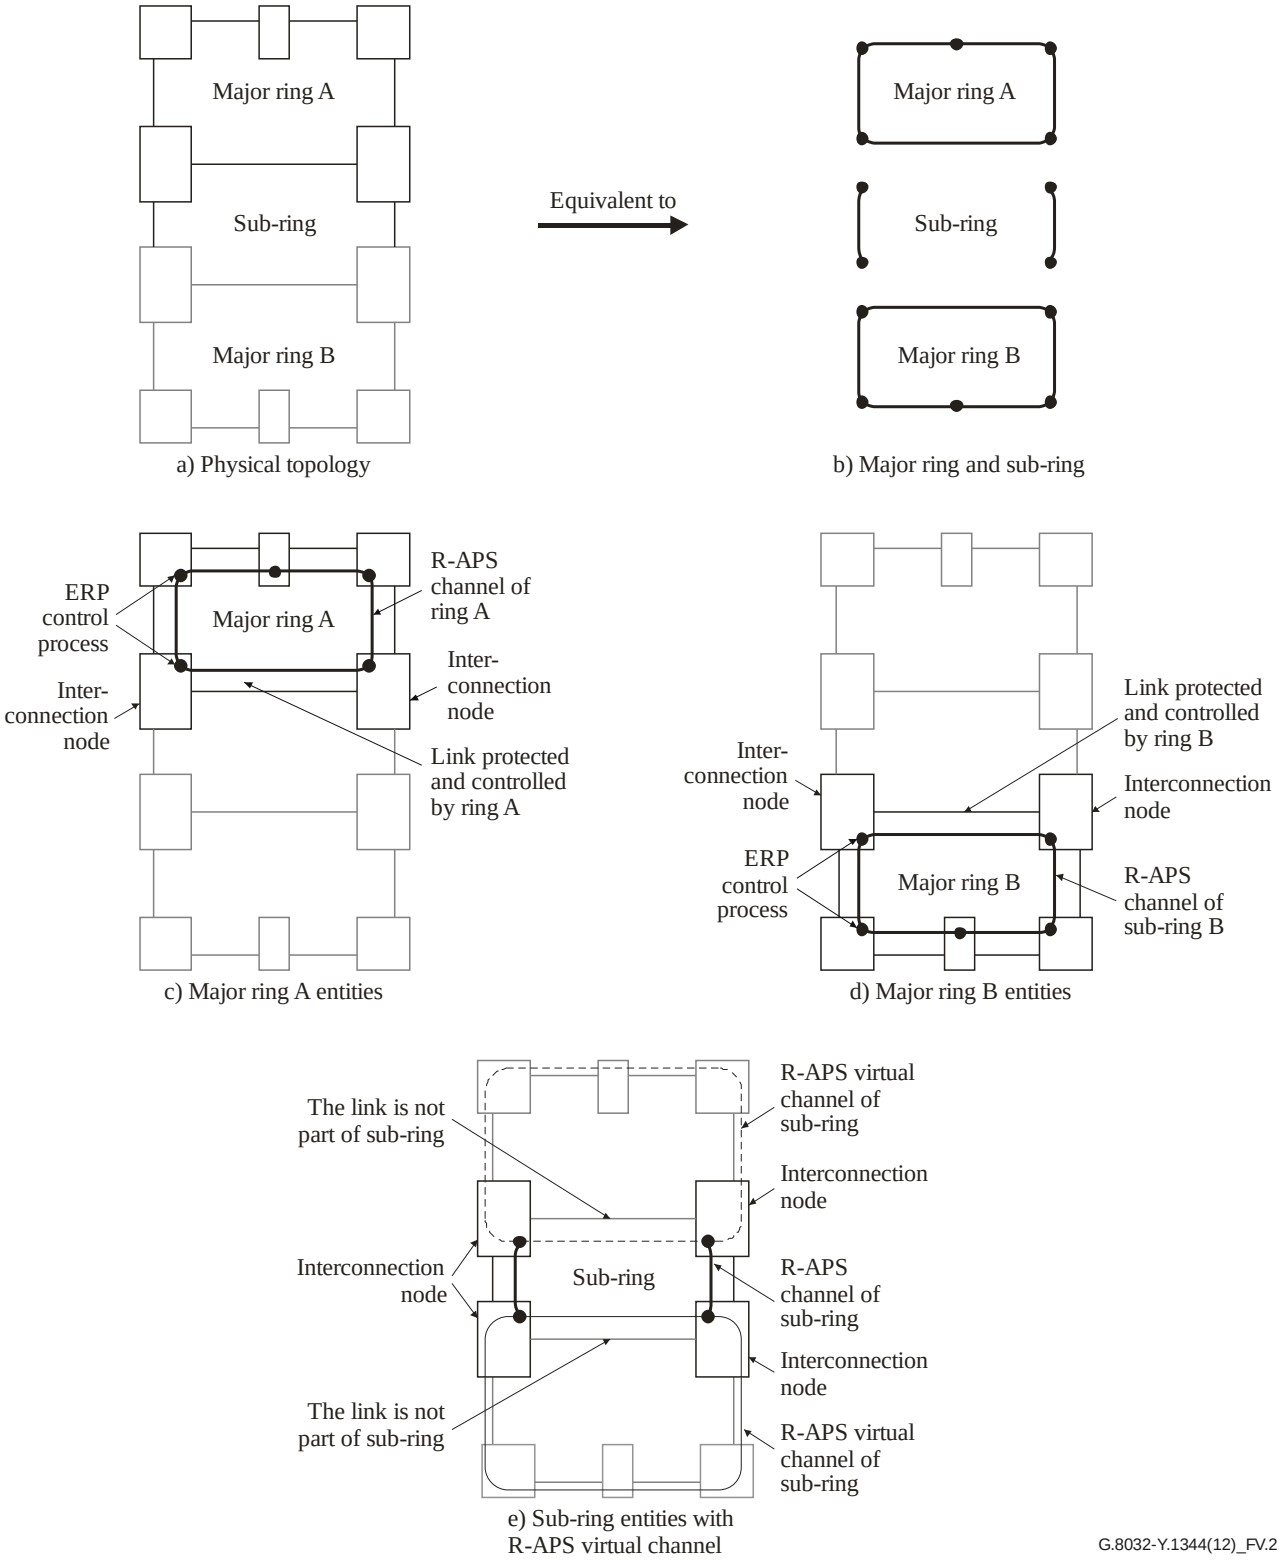

Figure V.2 – Configuration for interconnection between multiple major rings and a sub-ring

V.2 Topology examples for interconnected Ethernet rings

Figure V.3 represents examples of a topology composed of three or more interconnected Ethernet rings. The R-APS virtual channels are not depicted for simplification. When the sub-ring is operated with an R-APS virtual channel, it is deployed on an Ethernet ring that the sub-ring is connected to, as illustrated in Figures V.1 and V.2. There is no limit to the number of interconnected Ethernet rings.

a) Location of the RPL for a sub-ring

   The RPL can be placed on any ring link of a sub-ring. The RPL for a sub-ring cannot be placed on a major ring link between the interconnection nodes.

b) Intermediate Ethernet ring node(s) between interconnection nodes

   Ethernet ring node(s) that are part of a major ring can be placed between the interconnection nodes.

c) Multiple sub-rings connected to a major ring

   A major ring can accommodate multiple sub-rings. A pair of two interconnection nodes on a major ring can accommodate multiple sub-rings.

d) Sub-ring(s) interconnection

   A sub-ring can accommodate other sub-ring(s) on its ring link(s). The rules of b) and c) can be applied.

e) A sub-ring connected to multiple Ethernet rings

   A sub-ring can be accommodated in two or more different major rings or sub-rings. For example, sub-ring 2 is attached to a major ring and sub-ring 1, and sub-ring 5 is attached to both sub-ring 3 and sub-ring 4.

f) A sub-ring attached to multiple major rings

   A sub-ring can be attached to multiple major rings that are disjoint relative to each other. Multiple R-APS virtual channels are required (if using the sub-ring with an R-APS virtual channel model).

g) A sub-ring connected to a network that supports any technology network

   A sub-ring can be attached to a network that supports any other technology (e.g., xSTP, VPLS, etc.).

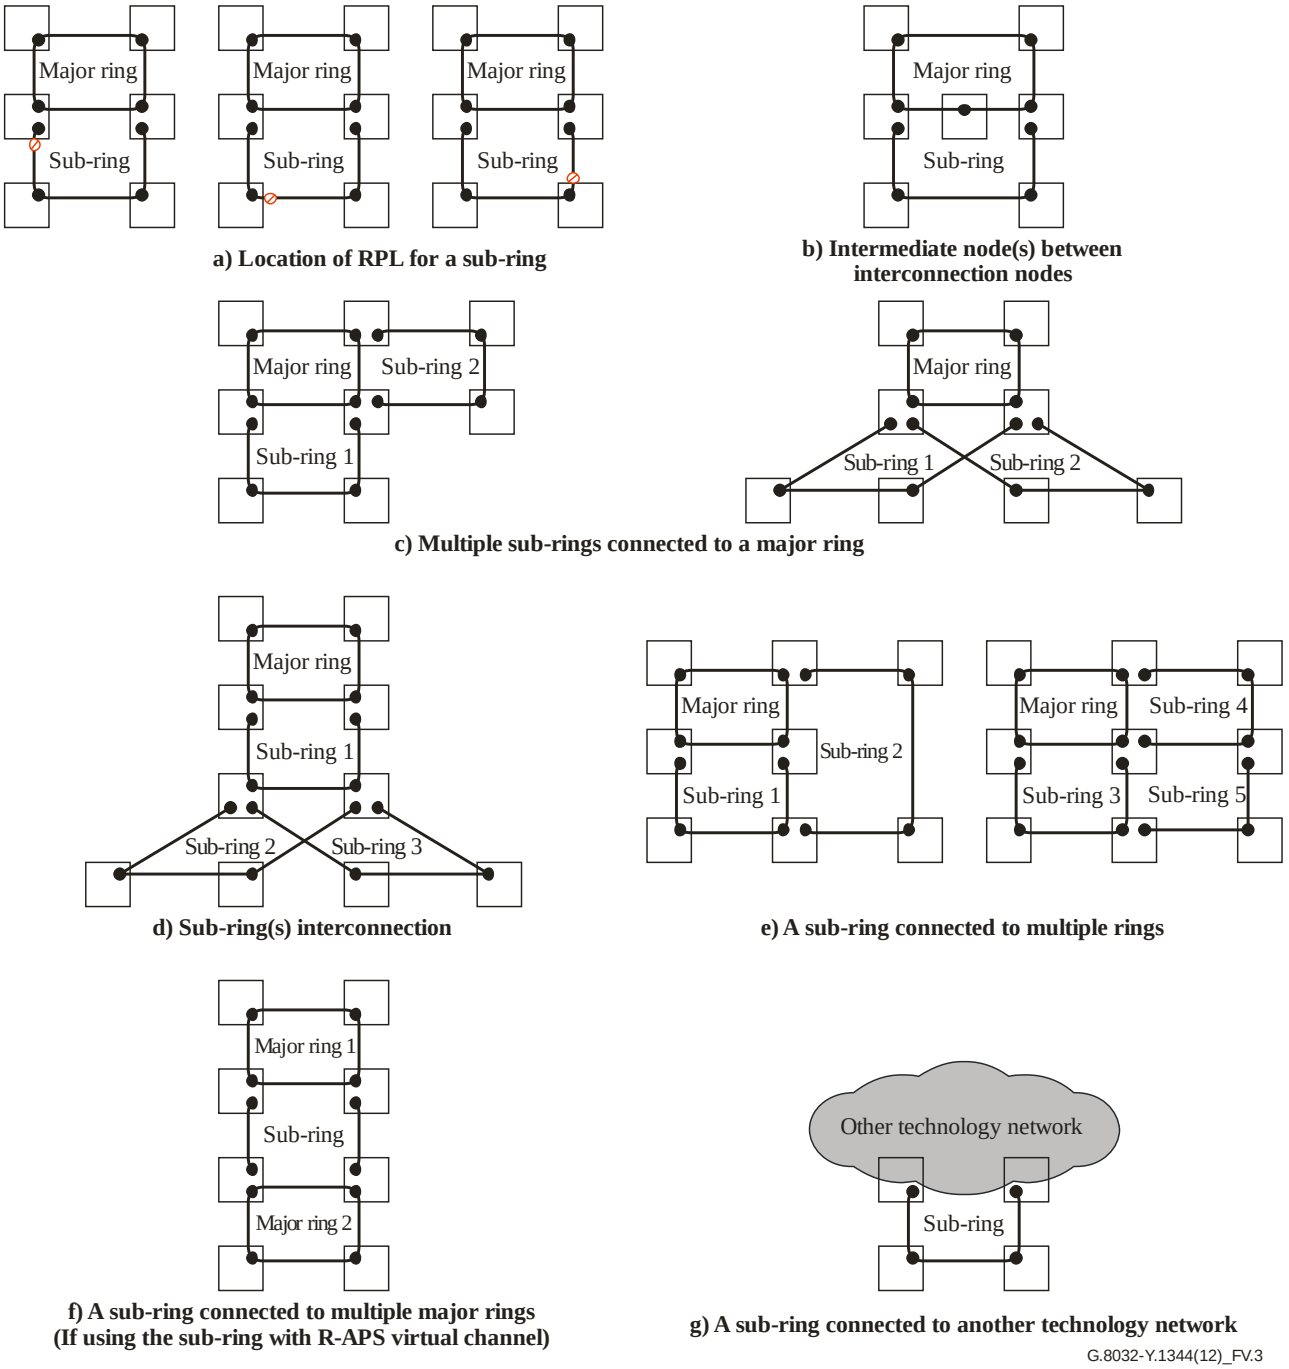

Figure V.3 – Topology examples for interconnected rings

Appendix VI

Protection switching for multiple ERP instances

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

VI.1 Multiple ERP instances

An Ethernet ring may support multiple traffic channels that may be grouped into different sets of VLANs. It is possible to define an ERPI instance as an entity that is responsible for the protection of a subset of the VLANs that transport traffic over the physical Ethernet ring. Each ERP instance is independent of other ERP instances that may be configured on the physical Ethernet ring. Each ERP instance independently applies the protection mechanism described in clause 10 of this Recommendation for the subset of the total traffic transmitted over the set of VLANs that the instance is configured for. For each ERP instance, an independent ERP control process exists.

Support of multiple ERP instances is optional for network elements supporting this Recommendation.

VI.2 Applying protection mechanisms to multiple ERP instances

When multiple ERP instances are configured for an Ethernet ring, each ERP instance should configure its own RPL, RPL owner node and RPL neighbour node. The ring link configured as the RPL may be (and generally is) different for each ERP instance supported.

VI.2.1 Addressing of multiple ERP instances

The protection mechanism defined in clause 10 is dependent upon the use of the R-APS protocol to notify the Ethernet ring nodes of the current condition of the Ethernet ring and control the protection switching operations. As stated in clause 10.3, the notification and control R-APS messages are transmitted using the MAC destination address 01-19-A7-00-00-01. When multiple ERP instances are activated, each ERP instance activates the protection switching procedures independently of each other. R-APS messages of different ERP instances are differentiated by the use of different R-APS VIDs.

VI.2.2 Protection switching – signal failure

If an SF condition is detected on an Ethernet ring supporting multiple ERP instances, then protection switching shall be invoked for each of the ERP instances configured. The R-APS messages should be transmitted on separate VIDs as specified in VI.2.1 over the ring links of the ERP instance.

Each ERP instance should perform protection switching under the control of the RPL owner node configured for that particular ERP instance. The functionality and state machine are consistent with those stated in clause 10 of this Recommendation.

VI.2.3 Protection switching – revertive and non-revertive

Support for revertive and non-revertive mode operation of the protection switching may be configured differently for each ERP instance configured in the Ethernet ring.

The recovery mechanism, when the SF condition is detected to be cleared, should be activated separately for each ERP instance in accordance with the revertive or non-revertive mode of the particular ERP instance.

VI.2.4 Protection switching – Manual switch and Forced switch

A Manual switch or Forced switch command is generated individually for each ERP instance. The ERP instance where the operator command (either FS or MS) is issued should transmit the R-APS message indicating the command over its R-APS channel. The operation of the protection switching and recovery should be compliant with the procedures described in clause 10 of this Recommendation.

VI.3 Protection switching model for multiple ERP instances

The protection mechanism for multiple ERP instances uses the same architecture as used for the single ERP instance case with the addition that this needs to be cloned for each ERP instance to transmit the R-APS messages for each ERP instance to the proper MAC address.

Figure VI.1 illustrates the model of an Ethernet ring node supporting two ERP instances. The MEP adaptation function is de-multiplexed, based on the VID to each ERP instance and informs the ERP control process for each ERP instance, that then asserts the proper condition for the ERP instance.

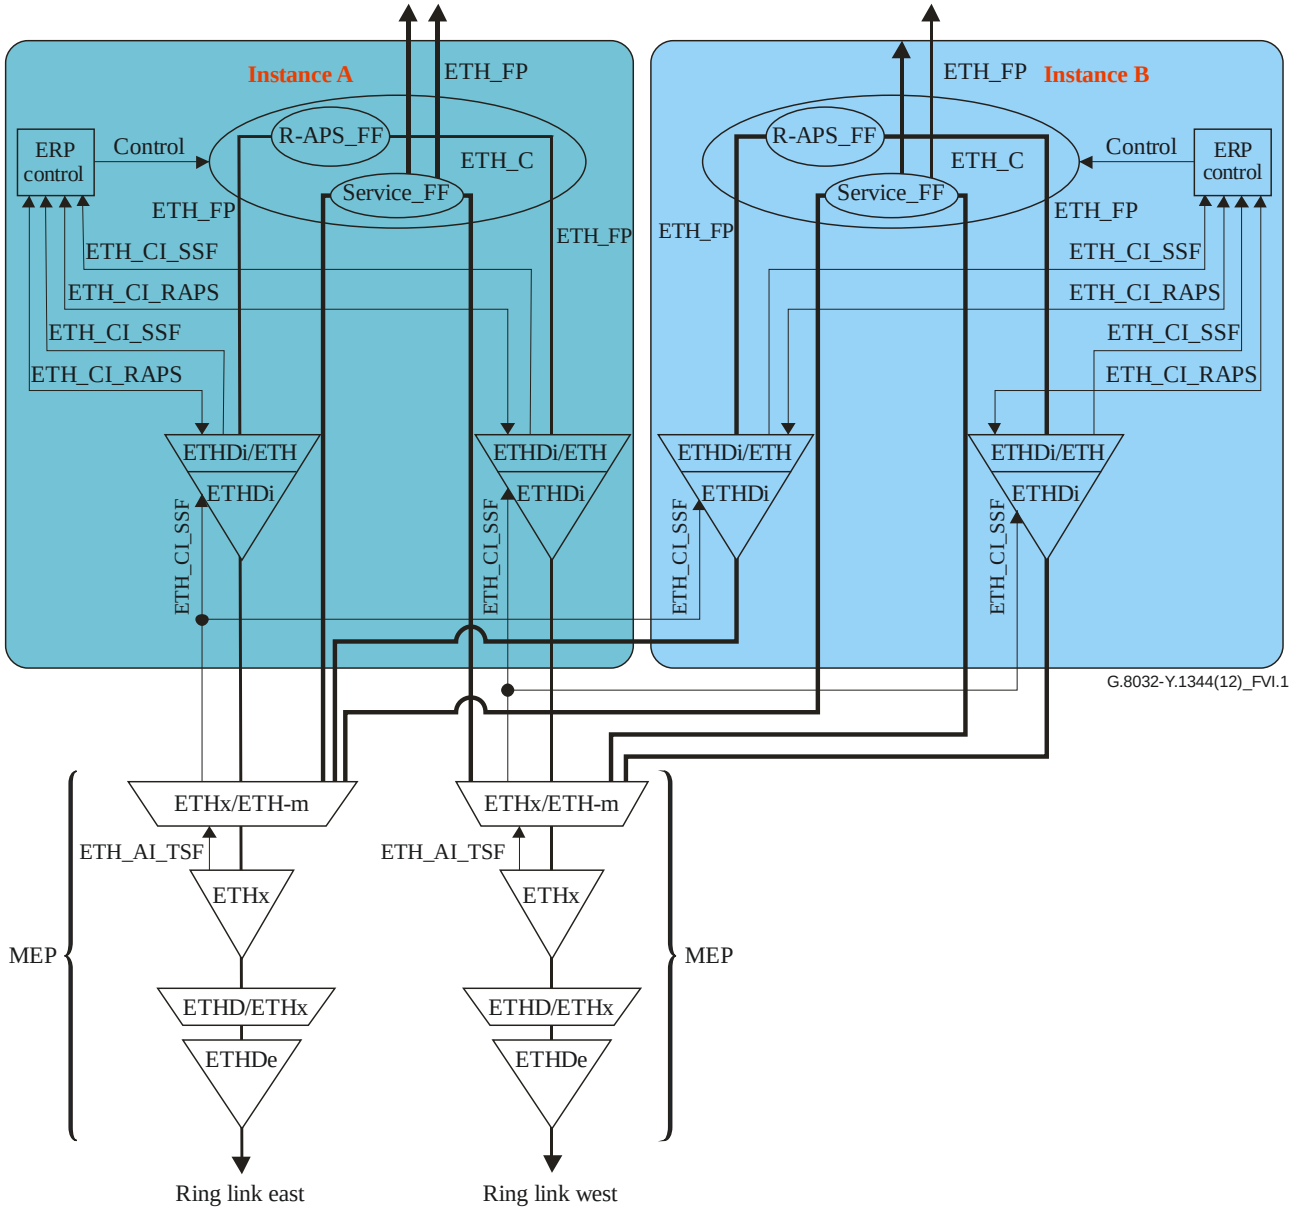

Figure VI.1 – MEPs and R-APS insertion function for Ethernet ring node supporting two ERP instances

VI.4 Multiple instances of interconnected rings

When the network includes interconnected rings, where neighbouring Ethernet rings are connected through a ring link between two interconnection nodes, it should be possible to configure multiple ERP instances in any of the following possible configurations:

a) Sets of VLAN may be limited to only one of the interconnected physical Ethernet rings. In this case an ERP instance is defined only on that physical Ethernet ring and is responsible for the protection of that set of VLANs.

b) A set of VLANS may span multiple interconnected Ethernet rings, in this case an ERP instance is defined on each physical Ethernet ring supporting that set of VLANs.

These possibilities are illustrated in Figure VI.2. In this figure, the network has two physical Ethernet rings that are connected. On the two Ethernet rings there are three groups of service traffic (red, blue, green) associated with four ERP instances (R1, B1, G1, G2):

a) G2 and B1 are ERP instances on the right-hand physical Ethernet ring.

b) G1 and R1 are ERP instances on the left-hand physical Ethernet ring.

c) G1 and G2 are ERP instances protecting the green group of service traffic that spans the interconnected Ethernet rings. These interconnected Ethernet rings shall not be two major rings since if the green group of service traffic is associated with both G1 and G2, the ring link between the interconnection nodes could be simultaneously blocked by the ERP control processes of both major rings resulting in a super loop on the group of service traffic.

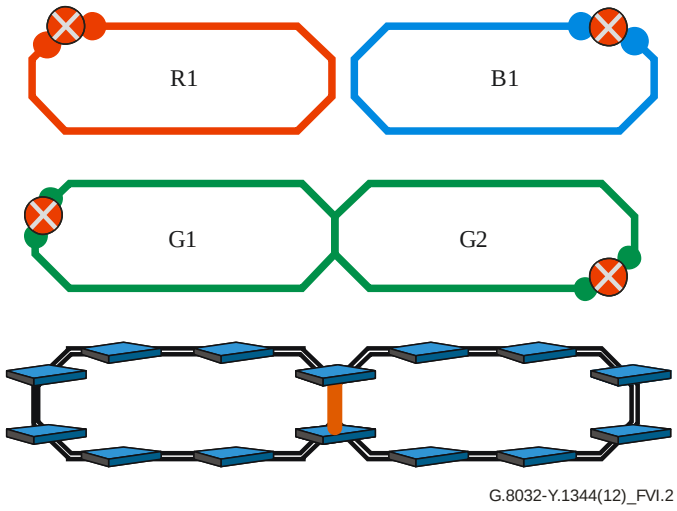

Figure VI.2 – Multiple ERP instances on interconnected physical Ethernet rings (normal condition)

When an SF condition is detected on the ring link between the interconnection nodes, the protection switching mechanism should be employed separately for each of the three groups of ERP instances:

a) For the G group protection switching is invoked on the Ethernet ring or sub-ring controlled by the R-APS channel of the G group that detected the ring link defect.

b) For the R group protection switching is invoked for R1.

c) For the B group protection switching is invoked for B1.

Figure VI.3 shows the ERP instances after the protection switching is invoked (on G2 for G group).

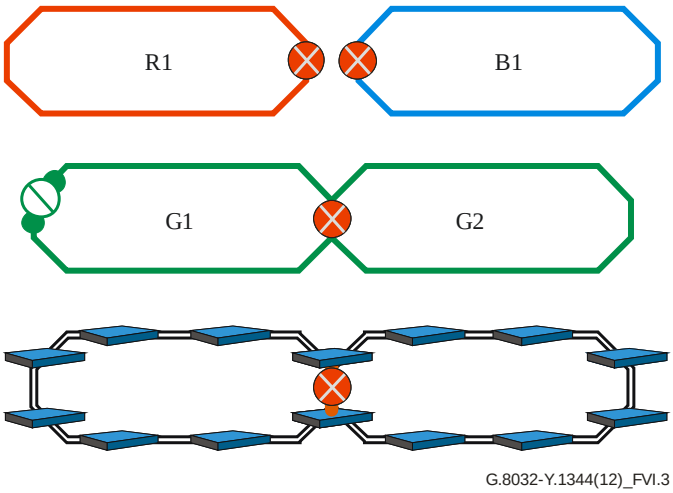

Figure VI.3 – Multiple ERP instances on interconnected physical Ethernet rings (the ring link between the interconnection nodes failure condition)

Appendix VII

Guidelines for the configuration of VIDs and ring IDs of R-APS channels

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

The following clauses contain guidelines on assigning VIDs and ring IDs for different sets of ERP instances configured as per this Recommendation. The different guidelines take into consideration the model of interconnected Ethernet ring support that the network operator employs.

VII.1 Sub-ring with an R-APS virtual channel

VII.1.1 Example 1: R-APS channel with different VIDs and the R-APS channel of a sub-ring and an R-APS virtual channel having different VIDs

Four ERP instances are deployed on interconnected physical Ethernet rings (rings 1 and 2) in the following figures. Each ERP instance has ring ID 1. ERP instance 1 is a major ring deployed on ring 1 and assigned ring ID "1". The R-APS channel of ERP instance 1 is identified by VID "1". ERP instance 2 is a sub-ring deployed on ring 2 and connected to ERP instance 1. The ring ID of ERP instance 2 is also "1". The R-APS channel of ERP instance 2 is identified by VID "2". The R-APS virtual channel is deployed on ring 1 and identified by VID "21" as data traffic associated with ERP instance 1. ERP instances 3 and 4 are similar but with the position of the major ring and sub-ring reversed, relative to ERP instances 1 and 2. ERP instances 3 and 4 have ring ID and VID as illustrated in Figure VII.1. Example 1 uses more VIDs in comparison to the other examples. However, a VID assigned to a sub-ring may be reused on Ethernet rings which are not immediately adjacent without translating VIDs. For example, VID 2 could be reused on an Ethernet ring connected to ring 1 on the opposite side, i.e., with no common interconnection node.

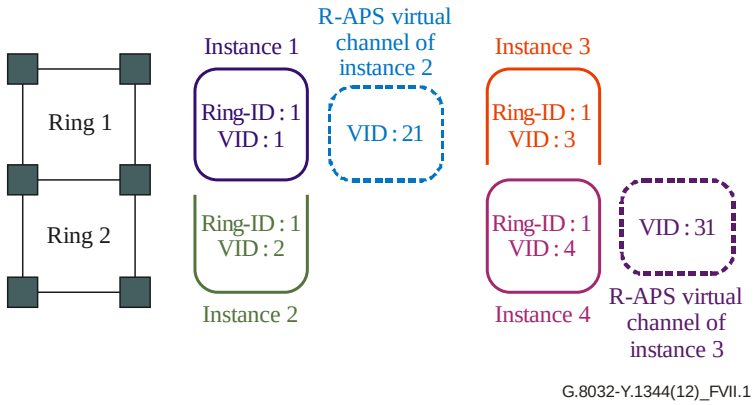

Figure VII.1 – Example 1: Different VIDs for different ERP instances; different VIDs for the ERP instances of a sub-ring and an R-APS virtual channel

The model of an interconnection node, which connects a sub-ring with an R-APS virtual channel and has different VIDs for each ERP instance, is represented in Figure 9-8. It is assumed that ERP instances 1 and 2 are depicted as ERP1 and 2 (in Figure 9-8), respectively. R-APS messages for ERP instance 1 are received on ring port 0 or ring port 1 and identified by its VID "1" at the ETHx/ETH-m_A function and forwarded to the ETHDi/ETH_A function. They are subsequently forwarded to the other ring port through the R-APS_1_FF function which is assigned VID "1". R-APS messages for ERP instance 2 are received on the sub-ring port and identified by its VID "2" at the ETHx/ETH-m_A function and forwarded to the ETHDi/ETH_A functions. They are subsequently forwarded through the R-APS_2_FF function which is assigned VID "2" to the

ETH_C function of ERP instance 1. On the ETH_C function, R-APS_2_FF assigned VID "21" is responsible for forwarding the R-APS messages from ERP instance 2 as service traffic.

VII.1.2 Example 2: R-APS channel with different VIDs and an R-APS channel of a sub-ring and an R-APS virtual channel having the same VID

The ring IDs and VIDs of ERP instance 1 and ERP instance 2 depicted in Figure VII.2 are the same as those in Example 1. The R-APS virtual channel of ERP instance 2 is deployed on ring 1 as data traffic associated with ERP instance 1 and identified by VID "2" which is the same as the VID of the R-APS channel of ERP instance 2 over ring 2. ERP instances 3 and 4 are similar but with the position of the major ring and sub-ring reversed, relative to ERP instances 1 and 2. ERP instances 3 and 4 have a ring ID and VID as illustrated in Figure VII.2. In Example 2, it seems to be easier to manage the VIDs than for the other examples. However, the same number of VIDs as in Example 1 may have to be used since a VID assigned to a sub-ring may not be reused on Ethernet rings which are not immediately adjacent without translating VIDs. For example, VID 2 could not be reused on a ring connected to ring 1 on the opposite side, i.e., with no common interconnection node and a different VID would need to be assigned.

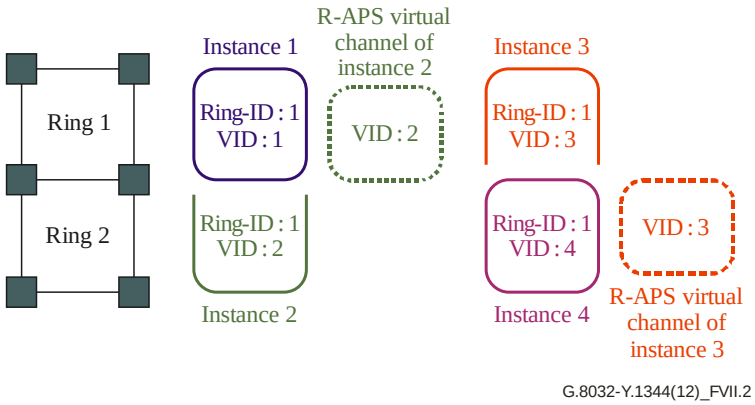

Figure VII.2 – Example 2: Different VIDs for different ERP instances; same VID for an ERP instance of a sub-ring and an R-APS virtual channel

It is assumed that ERP instances 1 and 2 are depicted as ERP1 and 2 respectively in Figure 9-8. R-APS messages for ERP instance 1 are received on ring port 0 or ring port 1 and identified by its VID "1" at the ETHx/ETH-m_A function and forwarded to the ETHDi/ETH_A function. They are subsequently forwarded to the other ring port through the R-APS_1_FF function which is assigned VID "1". R-APS messages for ERP instance 2 are received on sub-ring port and identified by its VID "2" at the ETHx/ETH-m_A function and forwarded to the ETHDi/ETH_A functions. They are subsequently forwarded through the R-APS_2_FF function which is assigned VID "2" to the ETH_C function of ERP instance 1. On the ETH_C function, R-APS_2_FF assigned VID "2" is responsible for forwarding the R-APS messages from ERP instance 2 as service traffic.

VII.1.3 Example 3: R-APS channel with the same VIDs and different multicast addresses with ring ID

In Figure VII.3, ERP instance 1 is a major ring deployed on the ring 1 and assigned ring ID "1". The R-APS channel of ERP instance 1 is identified by a set of VID "1" and the destination MAC address that includes ring ID "1". ERP instance 2 is a sub-ring deployed on ring 2 and connected to ERP instance 1. The ring ID assigned to ERP instance 2 is "2". The R-APS channel of ERP instance 2 is identified by a set of VID "1" and the destination MAC address that includes ring ID "2". The R-APS virtual channel is deployed on ring 1 and identified by a set of VID"1" and the destination MAC address that includes ring ID "2" for service traffic. ERP instances 3 and 4 are similar but

with the position of the major ring and sub-ring reversed in comparison with ERP instances 1 and 2. ERP instances 3 and 4 have a ring ID and VID as shown in Figure VII.3.

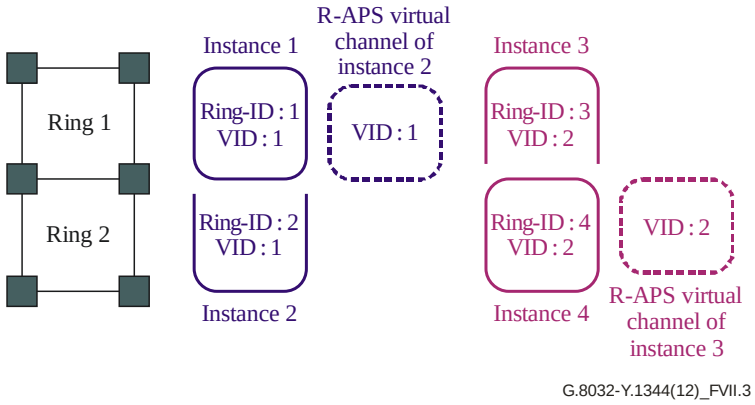

Figure VII.3 – Example 3: Same VIDs for connected ERP instances and an R-APS virtual channel

The model of an interconnection node, which connects a sub-ring with an R-APS virtual channel and has the same VIDs for each ERP instance, is represented in Figure 9-9. It is assumed that ERP instances 1 and 2 are depicted as rings 1 and 2 respectively. R-APS messages for ERP instance 1 are received on ring link 0 or ring link 1 and identified by its VID "1" at the ETHx/ETH-m_A function and forwarded to the ETHDi/ETH_A function. When the ETHDi/ETH_A function receives R-APS messagesRing ID, it forwards the R-APS messages to the ERP control process and the R-APS_1_FF function. When the ERP control process receives R-APS messages which have ring ID "1" in destination MAC address, it extracts the R-APS messages. When it receives R-APS messages with another ring ID, it just discards them. Ring IDOn the R-APS_FF function of ERP instance 1, the R-APS messages with VID "1" and ring ID "1" are forwarded to ring link 1 and 0, not to the ETHDi/ETH_A of ERP instance 2. R-APS messages for ERP instance 2 are received on the sub-ring link and identified by its VID "1" at the ETHx/ETH-m_A function and forwarded to the ETHDi/ETH_A functions. When the ETHDi/ETH_A function receives R-APS messagesRing ID, it forwards the R-APS messages to the ERP control process and the R-APS_2_FF function. On the R-APS_FF function of ERP instance 2, the R-APS messages with VID "1" and ring ID "2" are forwarded to the ETH_C function of ERP instance 1. They are subsequently forwarded through the R-APS_FF function of ERP instance 1 based on their VID "1" and ring ID "2" to ring link 0 or 1.

VII.2 Sub-ring without an R-APS virtual channel

VII.2.1 Example 4: Sub-ring without an R-APS virtual channel model, each R-APS channel with different VIDs

In Figure VII.4, ERP instance 1 is a major ring deployed on ring 1 and assigned ring ID "1". The R-APS channel of ERP instance 1 is identified by VID "1". ERP instance 2 is a sub-ring deployed on ring 2 and interconnected to ERP instance 1. The ring ID of ERP instance 2 is "2" and the R-APS channel of ERP instance 2 is identified by VID "2". ERP instances 3 and 4 are similar but with the position of the major ring and sub-ring reversed relative to ERP instances 1 and 2. ERP instances 3 and 4 have a ring ID and VID as illustrated in Figure VII.4. In Example 3 it seems to be easier to manage the VIDs than in the other examples. However, the same number of VIDs as in Example 2 is used and cannot be reassigned.

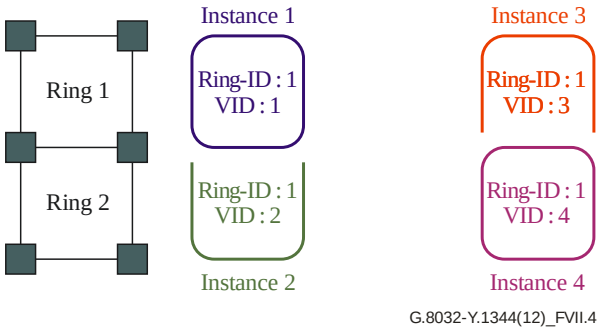

Figure VII.4 – Example 3: Different VIDs for each ERP instance

The model of an interconnection node, which connects a sub-ring without an R-APS virtual channel and has different VIDs for each ERP instance, is represented in Figure 9-11. It is assumed that ERP instances 1 and 2 are depicted as ERP1 and 2 respectively. R-APS messages for ERP instance 1 are received on ring port 0 or ring port 1 and identified by its VID "1" at the ETHx/ETH-m_A function and forwarded to the ETHDi/ETH_A function. They are subsequently forwarded to the other ring port through the R-APS_1_FF function which is assigned VID "1". R-APS messages for ERP instance 2 are received on the sub-ring port and identified by its VID "2" at the ETHx/ETH-m_A function and forwarded to the ETHDi/ETH_A functions. They are extracted at the ETHDi/ETH_A function and not forwarded to the ETH_C function.

VII.2.2 Example 5: R-APS channel with same VIDs and different multicast addresses with a ring ID

In Figure VII.5, ERP instance 1 is a major ring deployed on ring 1 and assigned ring ID "1". The R-APS channel of ERP instance 1 is identified by a set of VID "1" and the destination MAC address including ring ID "1". ERP instance 2 is a sub-ring deployed on ring 2 and connected to ERP instance 1. The ring ID of ERP instance 2 is "2". The R-APS channel of ERP instance 2 is identified by a set of VID "1" and the destination MAC address including ring ID "2". ERP instances 3 and 4 are similar but with the position of the major ring and sub-ring reversed in comparison with ERP instances 1 and 2. ERP instances 3 and 4 have a ring ID and VID as shown in Figure VII.5.

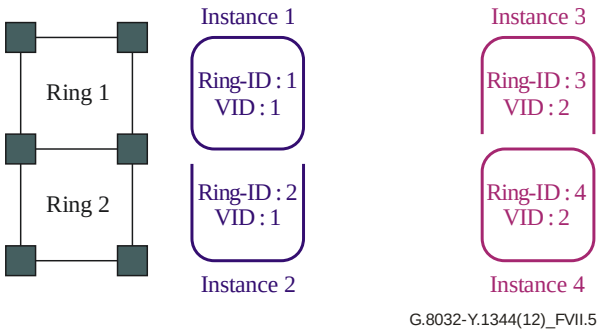

Figure VII.5 – Example 5: Same VIDs for connected ERP instances

The model of an interconnection node, which connects a sub-ring without an R-APS virtual channel and has the same VIDs for each ERP instance, is represented in Figure 9-10. It is assumed that ERP instances 1 and 2 are depicted as rings 1 and 2 respectively. R-APS messages for ERP instance 1 are received on ring link 0 or ring link 1 and identified by its VID "1" at the ETHx/ETH-m_A function and forwarded to the ETHDi/ETH_A function. When the ETHDi/ETH_A function receives R-APS messagesRing ID, it forwards the R-APS messages to the ERP control process and the R-APS_1_FF function. When the ERP control process receives R-APS messages which have

ring ID "1" in destination MAC address, it extracts the R-APS messages. When it receives R-APS messages with another ring ID, it just discards them. Ring IDOn the R-APS_FF function of ERP instance 1, the R-APS messages with VID "1" and ring ID "1" are forwarded to ring links 1 and 0. R-APS messages for ERP instance 2 are received on the sub-ring link and identified by its VID "1" at the ETHx/ETH-m_A function and forwarded to the ETHDi/ETH_A functions. When the ETHDi/ETH_A function receives the R-APS messagesRing ID, it forwards the R-APS messages and does not forward them to the ETH_C function of ERP instance 2.

VII.3 Backward compatibility

VII.3.1 Example 6: co-existence on an Ethernet ring of Ethernet ring nodes which support the 2010 version of this Recommendation (v2) and the 2008 version (v1) of this Recommendation

When Ethernet ring nodes running ITU-T G.8032v1 and ITU-T G.8032v2 co-exist on an Ethernet ring, note that the ring ID of each Ethernet ring node is configured as "1". Figure VII.6 is an example of the case of a sub-ring with an R-APS virtual channel. An Ethernet ring node running ITU-T G.8032v1 always transmits R-APS messages with the destination MAC address "01-19-A7-00-00-01". Interconnection nodes running ITU-T G.8032v2 should recognize the interconnected rings as ring ID "1" in order to extract or transmit R-APS messages from the Ethernet ring nodes running ITU-T G.8032v1. The R-APS channels and the R-APS virtual channels are indicated by a VID. In this figure, a single ERP instance can be deployed on each Ethernet ring (rings 1 and 2) because ITU-T G.8032v1 does not support the multiple ERP instance capability.

When a sub-ring with an R-APS virtual channel is used as illustrated in Figure VII.6, the behaviour of the blocked ring port (e.g., whether it forwards R-APS messages or not), defined in ITU-T G.8032v2, is the same as that specified in ITU-T G.8032v1. On the other hand, when a sub-ring without an R-APS virtual channel is used, the behaviour of the blocked ring port is different between ITU-T G.8032v1 and ITU-T G.8032v2 specifications as specified in clause 10.1.14. Therefore, when an Ethernet ring node running ITU-T G.8032v1 and ITU-T G.8032v2 co-exist on an Ethernet ring, the sub-ring should be deployed with the R-APS virtual channel.

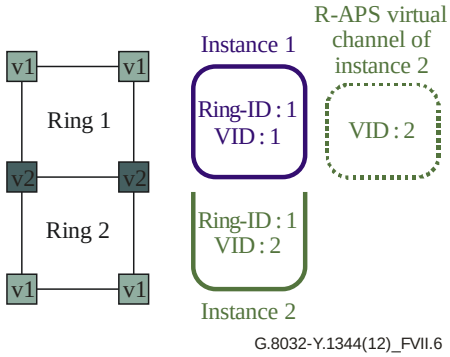

Figure VII.6 – Example 6: Co-existence of Ethernet ring nodes running G.8032v1 and G.8032v2 (with an R-APS virtual channel)

The model of the interconnection node depicted in Figure VII.6 is represented in Figure 9-8 and the behaviour of functions is the same as described in VII.1.2.

Appendix VIII

Flush optimization

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

VIII.1 Flushing FDB consideration

The ERP mechanism requires flushing the FDB with the goal of re-learning the correct filtering entries when protection switching has executed. However, in cases where the logical topology of a client channel has not changed as a result of failure, recovery or administrative operation, it is not necessary to flush FDB entries. A flush operation causes traffic flooding on the Ethernet ring and a consequent transient broadcast storm may occur. It is possible to reduce the occurrence of these broadcast storms by avoiding unnecessary FDB flushing.

VIII.2 Scenarios of unnecessary FDB flushing

The following are scenarios of protection switching that do not require FDB flushing. In these scenarios, all blocked ring ports continue to be blocked and the logical topology of a client channel is not changed.

a) Do not flush when RPL fails or recovers.

b) Do not flush when the RPL owner node or the RPL neighbour node fails or recovers.

c) Do not flush when the currently blocked ring port fails or recovers in non-revertive mode.

d) Do not flush when a request that results in blocking an already blocked ring link is issued (e.g., MS on RPL owner node).

The latter two scenarios are extensions beyond the scenarios described in the main text. These point to cases where FDB flushing may be omitted.

VIII.3 Example of FDB flush optimization

The following are rules for FDB flush optimization. Ethernet ring nodes connected to the RPL owner node or RPL neighbour node, need to be configured as the RPL next-neighbour node. The ring ports connected to the RPL owner node or the RPL neighbour node are called RPL next-neighbour ports.

Rule 1: If detecting an RPL link failure in [idle state], transmit R-APS (SF, DNF).

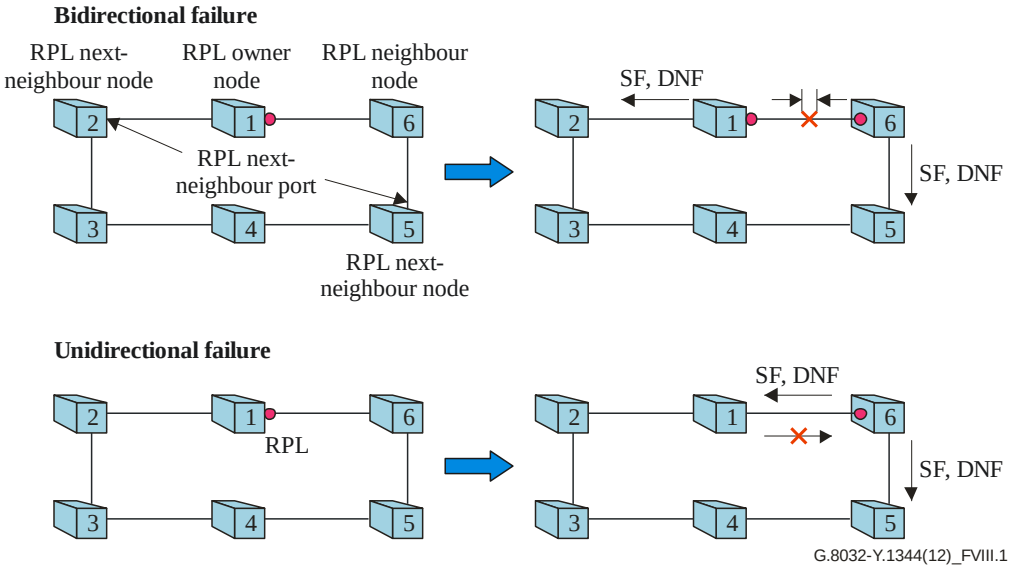

Figure VIII.1 – RPL failure case

Rule 2: When detecting a failure from an RPL next-neighbour port, in idle state, transmit an R-APS (SF) message only on the RPL next-neighbour port and do not transmit R-APS messages on the other ring port.

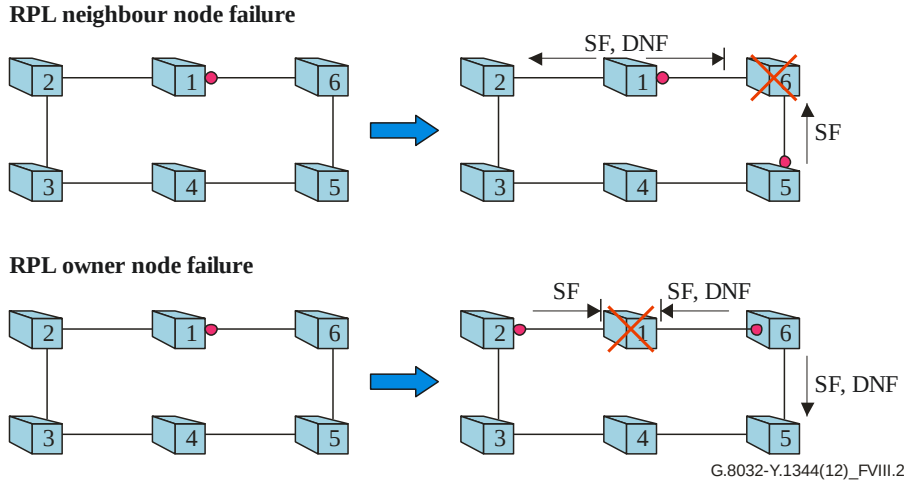

Figure VIII.2 – RPL owner node or RPL neighbour node failure case

Rule 3: If the RPL recovers, transmit the R_APS (NR, RB, DNF) message from the RPL owner node after the WTR timer expires.

Rule 4: If the RPL owner node detects ring recovery in the R-APS (SF, DNF) condition, transmit R_APS (NR, RB, DNF) after the WTR timer expires.

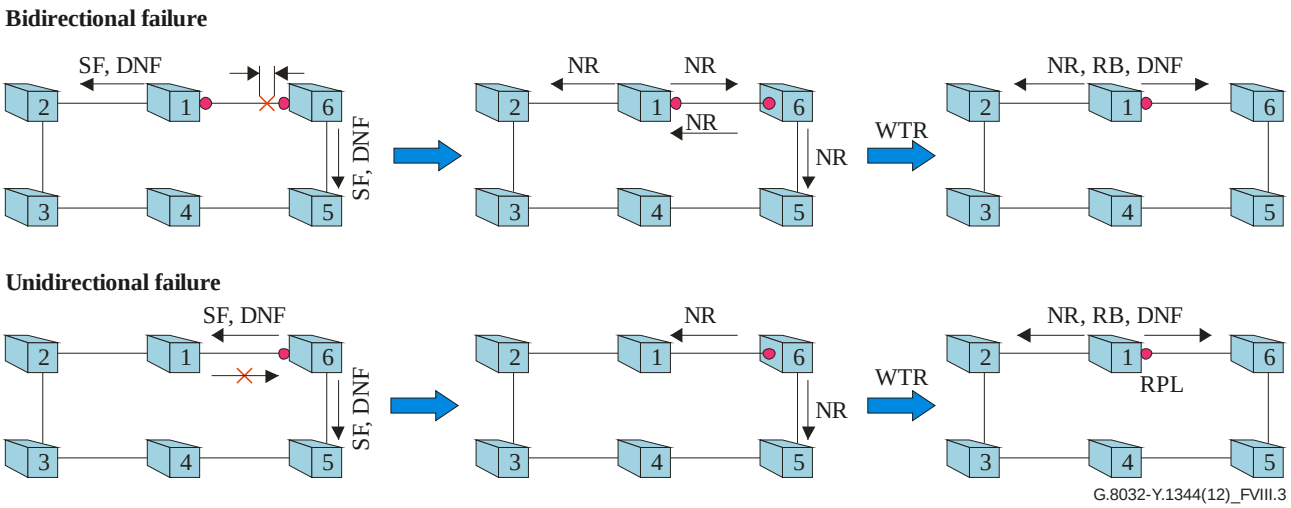

Figure VIII.3 – RPL recovery case

VIII.4 Additional definition of the ERP control process model and state machine

Rules 2 and 4, mentioned in the previous clause, require additional functionality in the ERP control process model and modification to the state machine. It should be noted that rules 1 and 3 are addressed in the basic functionality described in this Recommendation. In particular, rule 4 requires a "history" of DNF to be maintained and a "store/clear DNF status" process to be included in the ERP control process model as illustrated in Figure VIII.4.

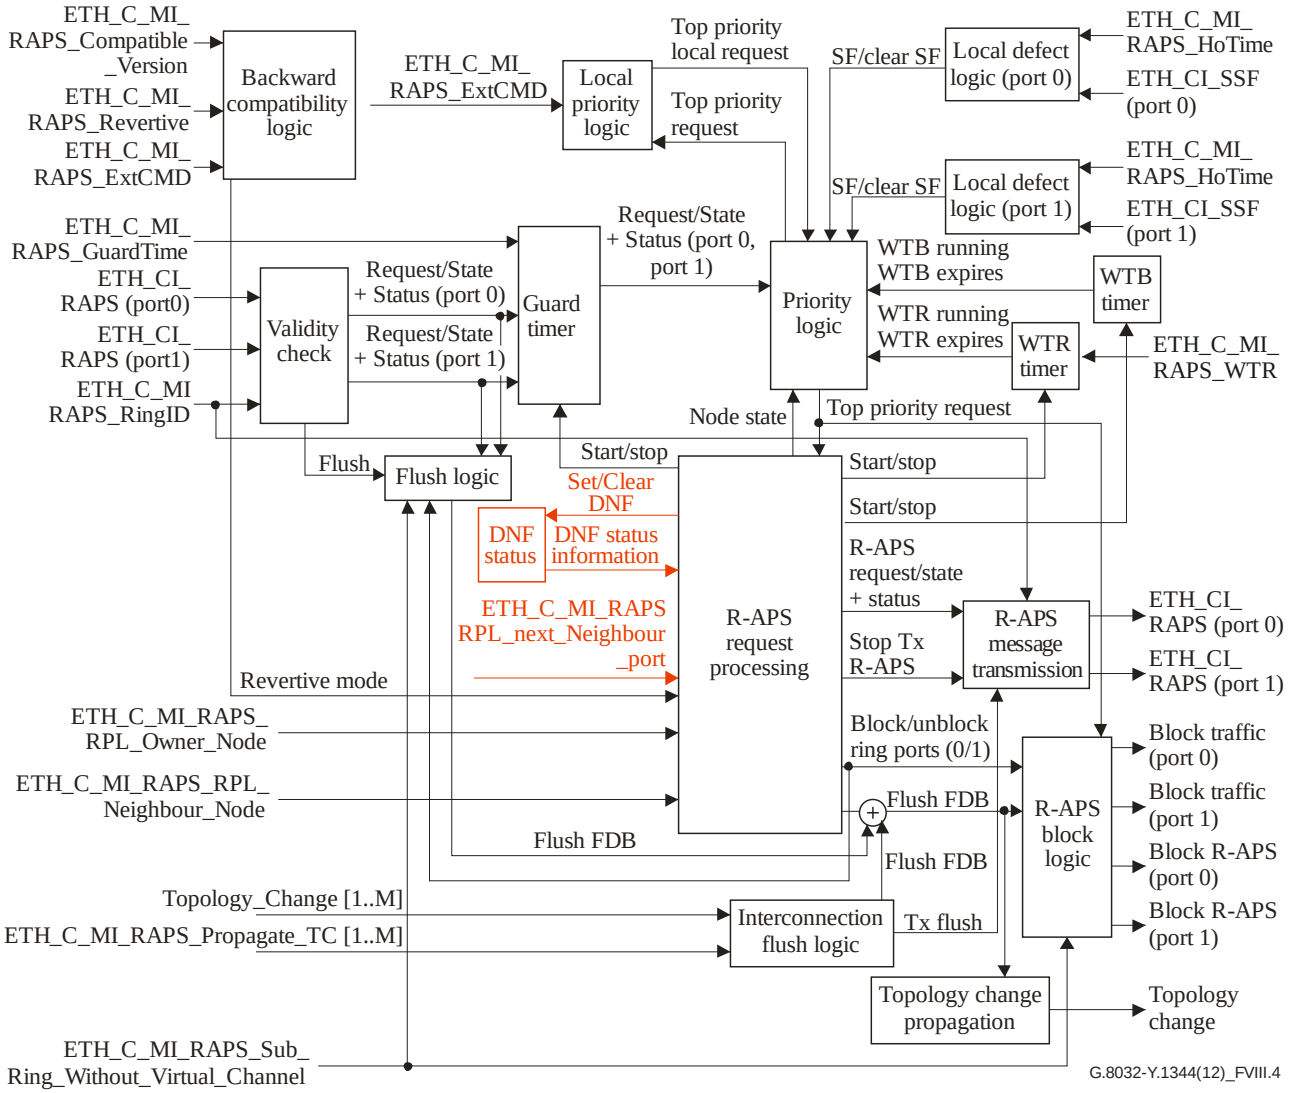

Figure VIII.4 – ERP control process model for flush optimization

In addition to the elements already defined in clause 10, the following are introduced for the specific support of flushing optimization.

The DNF status functionality is described in clause VIII.5 and it represents a memory element which retains the information of whether the protection switching was performed with flush optimization.

ETH_C_MI_RAPS_RPL_next_Neighbour_port represents the management information describing which ring port is connected to an RPL neighbour node or RPL owner node. By omission neither ring ports are considered RPL next-neighbour ports. If one ring port is an RPL next-neighbour port, ETH_C_MI_RAPS_RPL_next_Neighbour_port holds the information of which ring port is the RPL next-neighbour port.

The following Table VIII.1 presents the modification to the state machine (Table 10.2) in compliance with the above rules.

Table VIII.1 – State machine modification

Node state	Top priority request	Actions	Next node state

–	State machine initialization	Stop guard timer
		Stop WTR timer
		Stop WTB timer
		Clear DNF
		If RPL owner node:
		  Block RPL port
		  Unblock non-RPL port
		  Tx R-APS (NR)
		  If revertive:
		  Start WTB timer
		Else if RPL neighbour node:
		  Block RPL Port
		  unblock non-RPL port
		  Tx R-APS (NR)
		Else:
		  Block one ring port
		  unblock other ring port
		  Tx R-APS (NR)
			E
			(Pending)

A
(idle)
	local SF	If failed ring port is RPL port:
		  Block failed ring port
		  Tx R-APS (SF, DNF)
		  Unblock non-failed ring port
		  Set DNF status
		Else if failed ring port is RPL next-neighbour port:
		  Block failed ring port
		  Tx R-APS (SF) from failed ring port
		  Unblock non-failed ring port
		Else:
		  Block failed ring port
		  Tx R-APS (SF)
		  Unblock non-failed ring port
		  Flush FDB
			B
			(Protection)

Table VIII.1 – State machine modification

Node state	Top priority request	Actions	Next node state
A (idle)	R-APS (SF)	Unblock non-failed ring port Stop Tx R-APS If not DNF flush FDB If RPL next-neighbour node Tx three R-APS (SF) message If RPL owner node Tx three R-APS (SF) message clear DNF status Else: If RPL owner node set DNF status	B (Protection)
E (Pending)	WTR Expires	If RPL owner node: Stop WTB If RPL port is blocked: Tx R-APS (NR, RB,DNF) Unblock non-RPL port Else: Block RPL port If DNF status Tx R-APS (NR, RB, DNF) Else: Tx R-APS (NR, RB) Flush FDB Unblock non-RPL port clear DNF status	A (idle)
E (Pending)	R-APS (SF)	Unblock non-failed ring port Stop Tx R-APS If RPL owner node and not DNF clear DNF status If RPL owner node: Stop WTR Stop WTB	B (Protection)

NOTE – The highlighted actions in Table VIII.1 represent the changes relative to Table 10.2.

The following actions triggered by this process are introduced to support flush optimization:

a) Clear DNF status – triggers the action "clear DNF" of the DNF status.

b) Set DNF status – triggers the action "set DNF" of the DNF status.

c) Transmit three R-APS (msgtype, status bits) messages – Triggers the transmission of the initial burst of three R-APS messages over the two ring ports as described in clause 10.1.3.

d) Transmit R-APS (msgtype, status bits) from failed ring ports – Triggers the continuous transmission of R-APS messages over the failed ring port as described in clause 10.1.3.

VIII.5 DNF status

The DNF status retains the information on the "do not flush" condition so as to support flush optimization, for example, during protection reversion operations. The DNF status information takes the logical values "true" or "false".

The DNF status may be set or cleared. If set, the DNF status information input to the R-APS request processing takes the logical value "true"; if cleared, the DNF status information input to the R-APS request processing takes the logical value "false".

Appendix IX

Guidelines for management procedures

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

IX.1 An example procedure for removing an Ethernet ring node

When an operator wishes to remove an Ethernet ring node, it is recommended to issue FS commands at the Ethernet ring nodes adjacent to the Ethernet ring node that is being removed as illustrated in Figure IX.1. FS commands are issued at the ring ports of Ethernet ring nodes B and D, adjacent to the target Ethernet ring node, C [Step (b)]. Clear commands are also later issued at Ethernet ring nodes B and D in order to revert from the FS condition to the idle condition [Step (d)]. If an FS command is issued at the target Ethernet ring node directly, the additional procedure introduced in clause IX.2 is required.

NOTE – An MS command may also be used for removing an Ethernet ring node by issuing an MS command at the target Ethernet ring node or at one of the adjacent Ethernet ring nodes.

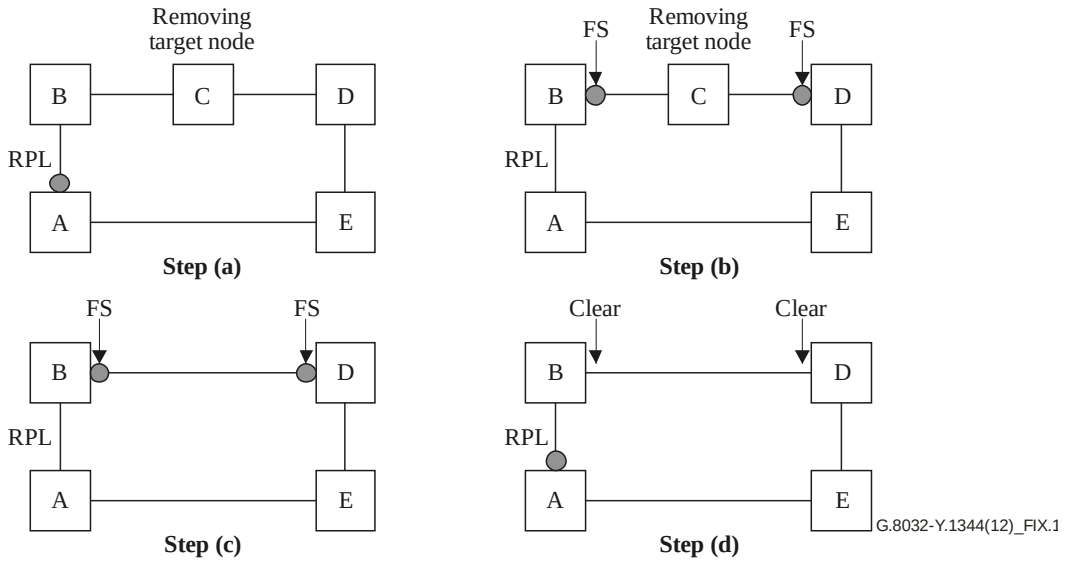

Figure IX.1 – Example procedure for removing an Ethernet ring node

IX.2 Management procedures to exit the FS state in case of failure of an Ethernet ring node under an FS condition

When an Ethernet ring is under an FS condition and the Ethernet ring node on which an FS command was issued is removed or fails, the Ethernet ring remains in the FS state because the FS command is to be cleared only at the Ethernet ring node where the FS command was issued. In Figure IX.2, even if Ethernet ring node C where an FS command was issued fails, Ethernet ring nodes B and D that are adjacent to the failed Ethernet ring node C, do not react to the local SF because a remote FS has a higher priority than a local SF. Additionally, there is no Ethernet ring node in the FS state where a Clear command could be issued to revert from the FS condition. This results in an inextricable FS condition.

When an operator has to perform a maintenance procedure (e.g., replacing, upgrading, etc.) on an Ethernet ring node (or a ring link), it is recommended that FS commands be issued at the two adjacent Ethernet ring nodes (e.g., B and D) instead of directly issuing an FS command at the Ethernet ring node (e.g., C) under maintenance in order to avoid falling into the aforementioned problematic situation.

Even if the FS command is issued at the Ethernet ring node under maintenance, it is possible to circumvent the problematic situation by following the procedure starting at step 4 in Figure IX.2.

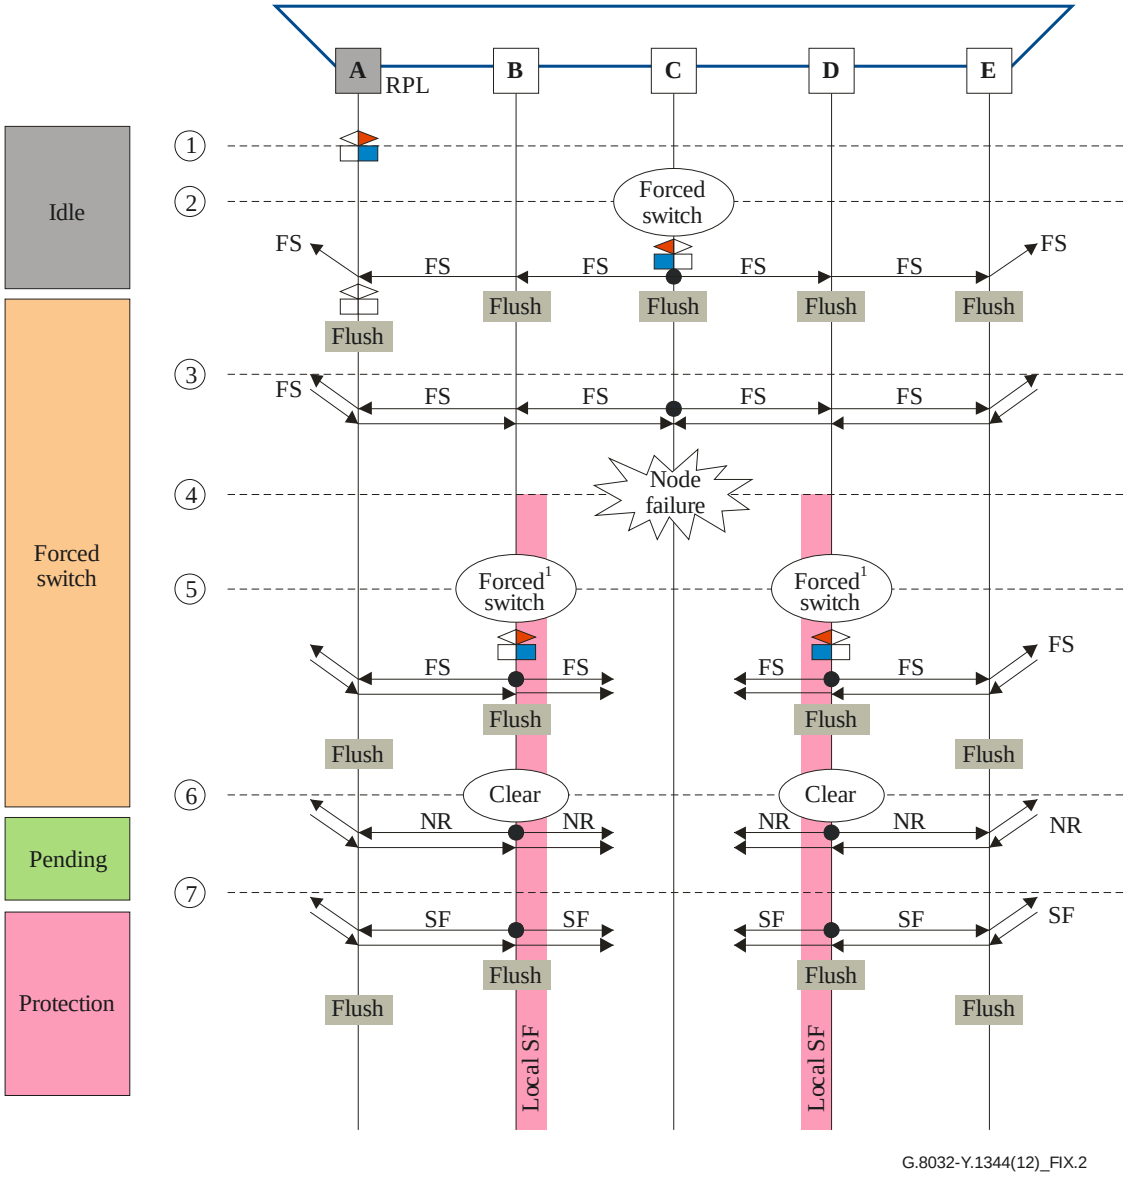

Figure IX.2 – Ethernet ring node failure scenario with management procedure

When the failure of an Ethernet ring node where a local FS command was issued (e.g., Ethernet ring node C in Figure IX.2, step 4) is detected, new FS commands are manually issued at Ethernet ring nodes (e.g., B and D in Figure IX.2, step 5) adjacent to the failed Ethernet ring node. At this point, these adjacent Ethernet ring nodes retain the local FS command. If Clear commands are then issued at these adjacent Ethernet ring nodes (e.g., B and D in Figure IX.2, step 6), these Ethernet ring nodes transit to the pending state and begin transmitting R-APS (NR) messages. This allows detection of the local SF condition. As a result of these actions, the FS state is successfully cleared within the Ethernet ring.

IX.3 Replacing an ITU-T G.8032 (2008) v1 Ethernet ring node with an ITU-T G.8032 (2010) v2 Ethernet ring node

When an Ethernet ring, already deployed using Ethernet ring nodes supporting only the functionalities of ITU-T G.8032 (2008) and Amendment.1 (2009) (Ethernet ring nodes running ITU-T G.8032v1), is upgraded with Ethernet ring nodes supporting the functionalities of this ITU-T G.8032 (2010) (v2 Ethernet ring nodes), an RPL owner node should be upgraded to become an Ethernet ring node running ITU-T G.8032v2 ahead of other Ethernet ring nodes deployed on the same Ethernet ring. Otherwise, differences between ITU-T G.8032v1 and ITU-T G.8032v2 flush behaviour might be exposed in the case of unidirectional failure on the non-RPL ring link attached to the RPL owner node.

Appendix X

Minimizing segmentation in interconnected rings

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

X.1 Characterization of the segmentation issue

When considering traffic that is being transmitted over a network of interconnected rings, there are situations that have been identified that can cause segmentation of the network as a result of dual failures in one of the rings. While it is not possible to address all cases of segmentation that are caused by multiple failures and such situations should be addressed by the operator before they become problematic, there is a need to characterize the situations in which the resulting segmentation may be avoided for some portion of the traffic being transported.

X.1.1 Problem statement

Figure X.1 shows a simplified network of two interconnected rings – the major ring includes nodes A, E, H, G and F, while the sub-ring includes nodes A, B, C, D and E. Nodes A and E are the interconnection nodes. The RPL for the major ring is the link G<->F and the RPL owner is node G. The RPL for the sub-ring is the link C<->B and the RPL owner is node C.

Figure X.1(a) shows that there is a service that is defined between Host X and Host Y that enters the major ring at node G and must crossover into the sub-ring and exit at node D. In the idle case (when there are no SF conditions in either ring) the traffic traverses through nodes G-H-E in the major ring and then crosses over to the sub-ring and traverses E-D and connects to Host Y.

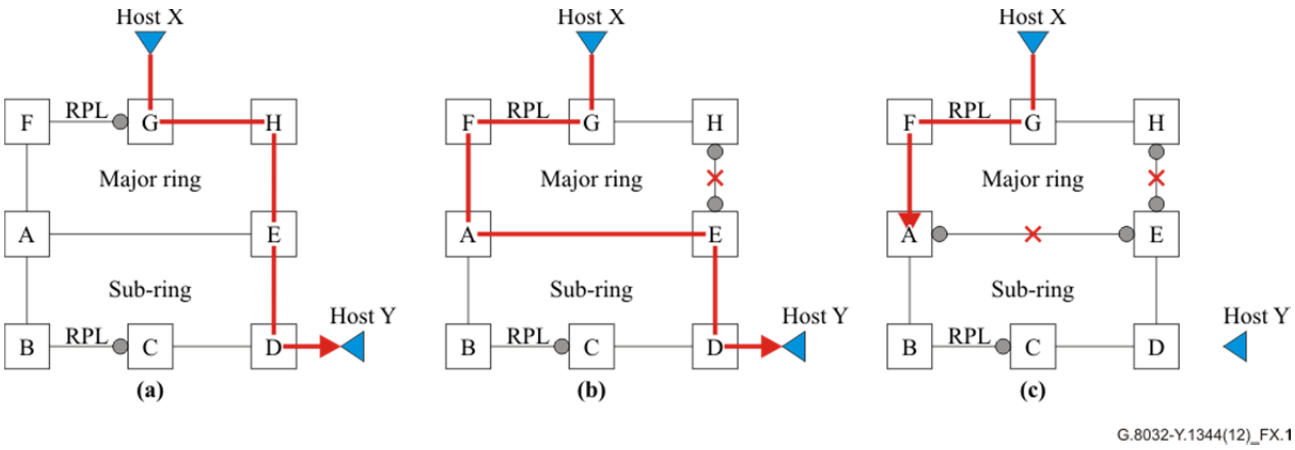

Figure X.1 – Scenario showing loss of connectivity between major ring and sub-ring

Figure X.1(b) shows the protection switching effect when an SF is identified on link H-E. Ring protection switching is invoked in the major ring, i.e., the RPL (G-F) is unblocked and the service traffic is now rerouted over the path G-F-A-E in the major ring and then E-D in the sub-ring.

However, if an additional SF is identified on link A-E in the major ring, then the service traffic will not be able to be transported. However, in theory it would be possible to reach Host-Y by following the route G-F-A-B-C-D, except that the sub-ring does not apply protection switching and the link B-C remains blocked to service traffic. As a result of this discontinuity, there are sub-ring nodes that are unreachable from the major ring, i.e., nodes C and D in Figure X.1(c). This situation is shown in Figure X.1(c) above.

The problem that exists in the basic application of ring protection is a result of the simplicity of the general procedures. Since each ring in the network addresses the switching triggers locally without propagating the trigger to any neighbouring ring, there is no way to cause the sub-ring to unblock the RPL in the given situation.

The following clauses outline a possible way of overcoming the segmentation that is created in a restricted class of scenarios. The next clause characterizes the class of scenarios that are addressed and the following clause presents a method for propagating the switching trigger to the sub-ring for that particular class of scenarios.

X.1.2 Relationship to interconnection models

It should be clarified that these scenarios are relevant to both interconnection models presented in clause 9.7 of the Recommendation. For the "ring interconnection model with R-APS virtual channel", it should be noted that even though the R-APS virtual channel is used to transmit the R-APS control information, a loss of connectivity between the interconnection nodes within the major ring does not trigger protection switching within the sub-ring, and the segmentation described above, may occur. Protection for this loss of connectivity is assumed to be controlled by the ERP control process of the major ring exclusively. For the "ring interconnection model without R-APS virtual channel", there is no correlation between the data path over the major ring and the ERP control process of the sub-ring and therefore segmentation may occur.

X.2 Class of double faults addressed

When considering the scenarios described in Figure X.1, we can characterize the segmentation as occurring when there is a double fault in the major ring that causes a break in connectivity between the two interconnection nodes. If the interconnection nodes are still able to transport traffic between them, then there will always be a path to reach all of the sub-ring nodes from any major ring node that is still connected to one of the interconnection nodes. However, if there is no connected path between the two interconnection nodes, then there is a problem for traffic that arrives at one interconnection node to reach the nodes that are beyond the sub-ring RPL (from the perspective of the interconnection node).

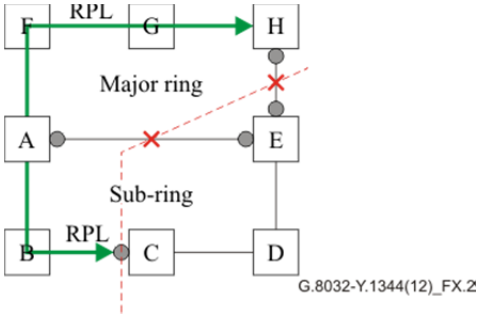

Figure X.2 – No connectivity between interconnection nodes

For example, in Figure X.2, if there is no connectivity between nodes A and E, then traffic that arrives at node A has no path to reach node D and traffic that arrives at node E has no path to reach node B – because the RPL (link B-C) is blocked.

However, if the sub-ring were to perform protection switching and unblock the RPL and block traffic at one of the interconnection nodes, e.g., node E, then all traffic that arrives at the non-blocked interconnection node, e.g., node A, could reach all of the sub-ring nodes.

In this and the following clause we will present a methodology for identifying the cases where traffic segmentation can be avoided in spite of multiple failures.

X.2.1 Detection of interconnection segmentation

The first step in minimizing the segmentation effect is to identify that the major ring is currently disconnected from the sub-ring. To facilitate this identification it is recommended to use the available tools to determine the connectivity of the two paths in the major ring between the two interconnection nodes. For this, a unicast UP MEP (as defined in [802.1ag]) can be used from the sub-ring port of the interconnection node to the sub-ring port of the other interconnection node. This can use the VID used by the major ring R-APS channel with a higher MEL or use any one VID that is controlled by the major ring protection mechanism. It should be noted that in the idle state, only one of these paths should be connected, the second path should be blocked by the RPL. In Figure X.3, the two tandem connections [A-I-E] and [A-F-G-H-E] are tested for connectivity. The latter path will be blocked by the RPL (on link F-G).

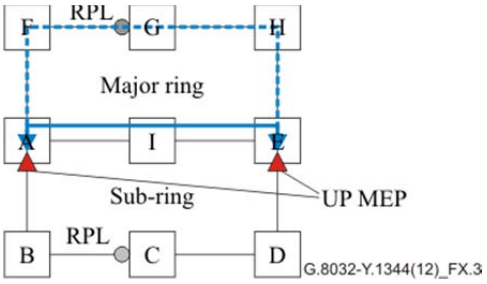

Figure X.3 – Interconnected rings connectivity verification

X.3 Procedure for minimization of segmentation

X.3.1 Management configuration

To apply the procedure outlined in the following clause, the operator should supply additional management information that will be used to determine the actions taken by the procedure.

The following management information items should be configured for each interconnection node of the sub-ring ERP control process:

- ETH_C_MI_RAPS_Interconnection_Node
  - Values: "primary", "secondary", or "none"
  - Default value: "none"
- ETH_C_MI_RAPS_Multiple_Failure
  - Values: "primary", "secondary", or "disabled"
  - Default value: "disabled"

In addition, the management system should configure a tandem connection between the two interconnection nodes. If the ETH_C_MI_RAPS_Interconnection_Node is set to either "primary" or "secondary", then an UP-MEP (as defined in [802.1ag]) should be configured at the sub-ring port of the interconnection node.

X.3.2 Block indication logic procedure

The following procedure will be used to minimize the segmentation of the traffic from the major ring to the sub-ring.

1. Employ connectivity verification for the tandem connection between the two UP MEPs. If there is connectivity, continue.
2. If there is a loss of connectivity between the two interconnection nodes, the UP MEP sends an indication to the block indication logic through the ETH_CI_SSF (see Figures X.4 and X.5).

3. The block indication logic (see Figure X.6) of the interconnection node sub-ring port accepts the two management information items as input. Compare the values of the two items:

   a) If the two values are identical (either both "primary" or both "secondary") – then perform the MS command to the sub-ring port.

   b) If the two values are different – then ignore.

4. When the connectivity of the tandem connection is restored, the UP MEP sends an indication through the ETH_CI_SSF to the block indication logic and the block indication logic should again compare the values of the two management information items.

   a) If the two values are identical (either both "primary" or both "secondary") – then clear MS to the sub-ring port (either port 0 or port 1 as the case may be).

   b) If the two values are different – then ignore.

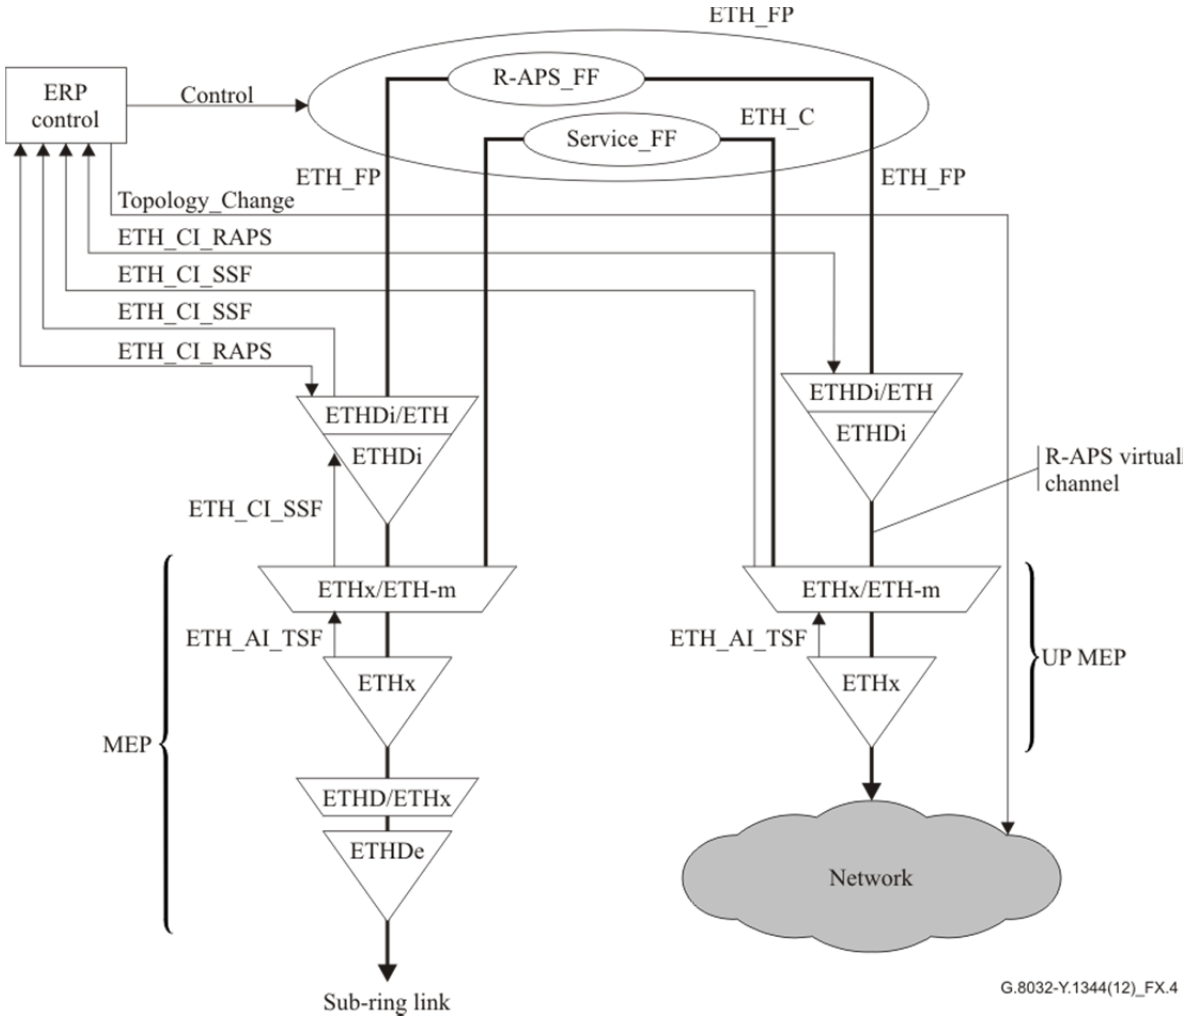

Figure X.4 – MEPs and R-APS insertion function in an interconnection node for the minimization of ring segmentation in interconnected rings

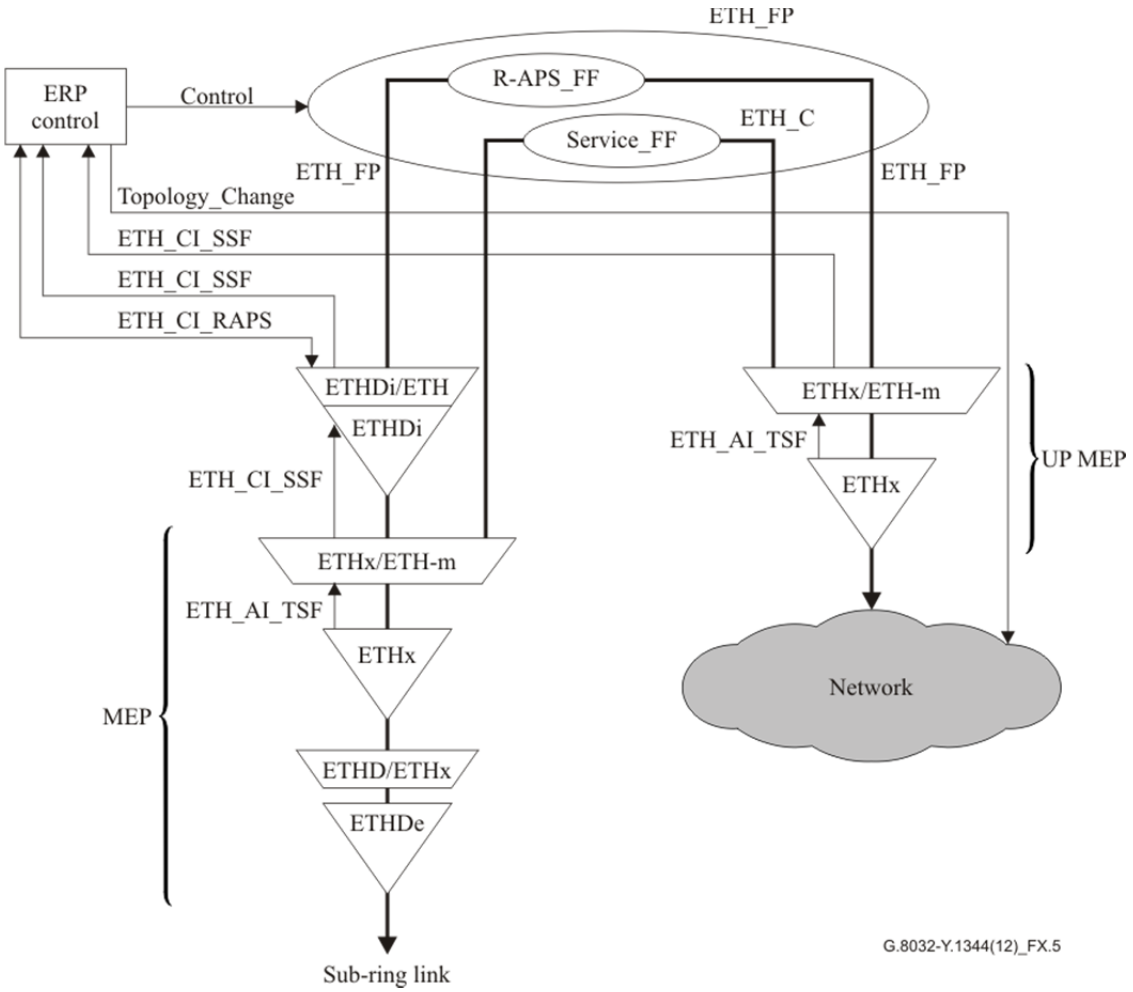

Figure X.5 – MEPs and R-APS insertion function in an interconnection node without an R-APS virtual channel for the minimization of ring segmentation in interconnected rings

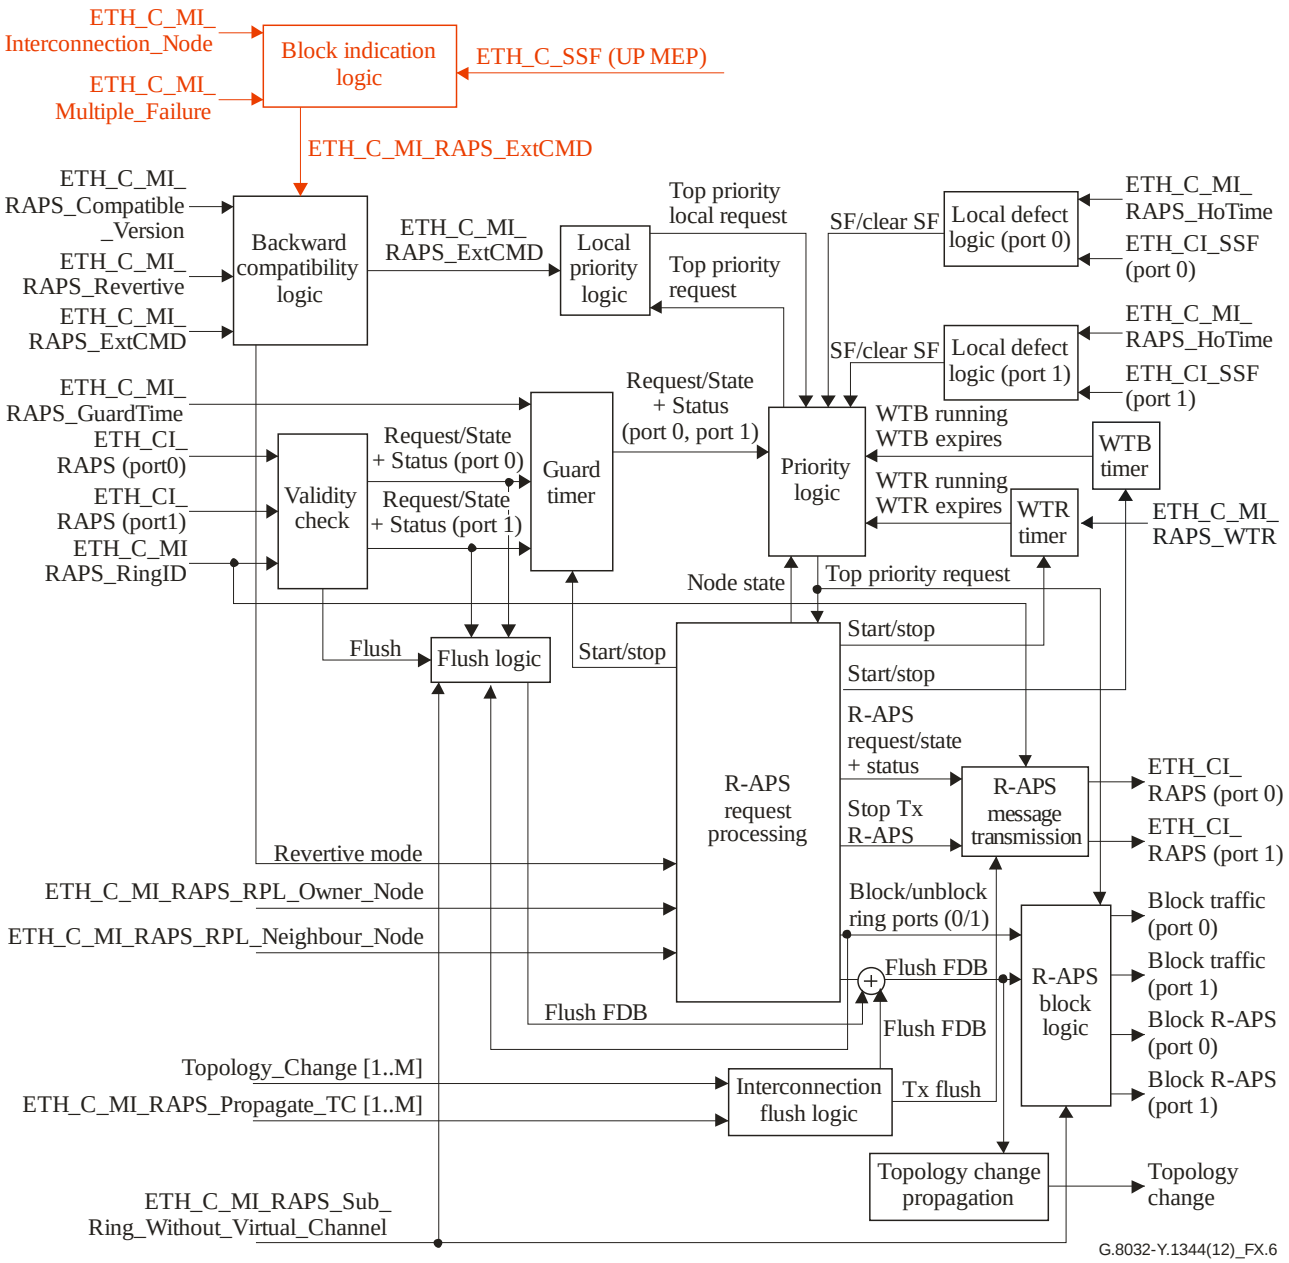

Figure X.6 – ERP control process for the minimization of ring segmentation in interconnected rings

Appendix XI

End-to-end service resilience

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

XI.1 Generic end-to-end service resilience

End-to-end service resilience may require protection based on the protection provided, as described in this Recommendation. However, additional protection may be required for the access links. This can be achieved by duplicating the access links as shown in Figure XI.1.

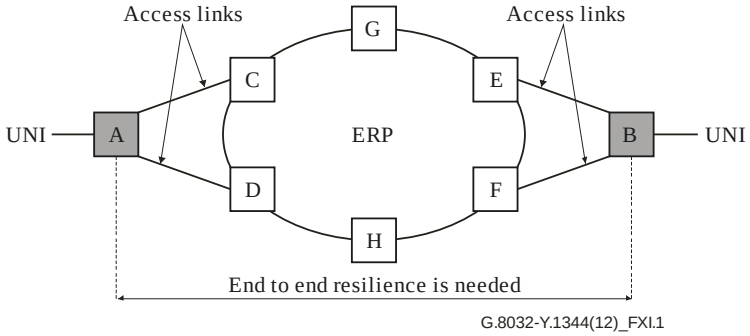

Figure XI.1 – A network model example for end-to-end service resilience

The protection mechanism used on the access links to provide end-to-end resilience could use the protection mechanisms described in [b-ITU-T G.8031], [b-IEEE 802.1D] or some other similar protection mechanism.

XI.2 Layering ITU-T G.8031 protection over ITU-T G.8032

For the purposes of this clause we pre-suppose that the protection mechanism employed for the end-to-end service is [b-ITU-T G.8031].

Referring to the service shown in Figure XI.1 above, we can imagine that the end-to-end protection would configure a working path that traverses the nodes [A-C-E-B] and a protection path that traverses the nodes [A-D-F-B].

XI.2.1 Basic guidelines for the layering of ITU-T G.8031 over ITU-T G.8032

When the protection of the end-to-end service, for example when the service runs between nodes A and B in Figure XI.1, is based on ITU-T G.8031 Ethernet linear protection, where part of the working and/or the protection path crosses a logical ring that is protected by an ITU-T G.8032 ERP, then the following guidelines are recommended:

- The working/protection path that crosses the ERP protected ring should only include two Ethernet ring nodes, at the points where the ring is entered and exited.
- The "link" between these two nodes can be considered a logical link, in the sense that the exact path that connects these two nodes is determined by the ERP mechanism, i.e., the ERP protection mechanism may determine that the connection may traverse the ring on either the shorter path or in the opposite direction along the longer path.
- The hold-off timer of the Ethernet linear protection mechanism should be configured with a value large enough to allow the ERP mechanism to complete its procedures prior to triggering linear protection as a result of a failure condition of this logical link.

- The working and protection paths (whichever cross the ring) should use different VIDs that are protected by ERP instances of the ring. Both of these VIDs may be protected by the same ERP instance or by separate ERP instances, at the operator's discretion.

  NOTE – When there are multiple services that are protected by ITU-T G.8031, it may be possible to reuse these same VIDs for the additional services based on the method of service identification.

This scenario may also be applied to the layering of RSTP [b-IEEE 802.1D] over ITU-T G.8032 by connecting two Ethernet private line (EPL) services (as defined in [b-ITU-T G.8011.1], where EPLs are separated by VIDs) between nodes A and B (in Figure XI.1).

XI.2.2 End-to-end service that traverses interconnected rings

If the end-to-end service crosses a network of interconnected rings, as shown in Figure XI.2, below, then the entire network of interconnected rings may be considered the underlying layer in the sense of the previous clause. Similar guidelines as stated above would apply, with the following generalization:

- The working/protection path that crosses the network of ERP-protected rings should only include two Ethernet ring nodes, at the points where the chain of Ethernet rings is entered and exited (for example, nodes C and E for the working path and nodes D and F for the protection path in Figure XI.2, i.e., nodes G, H, J, K, L, M would be transparent to the ITU-T G.8031 protection mechanism).

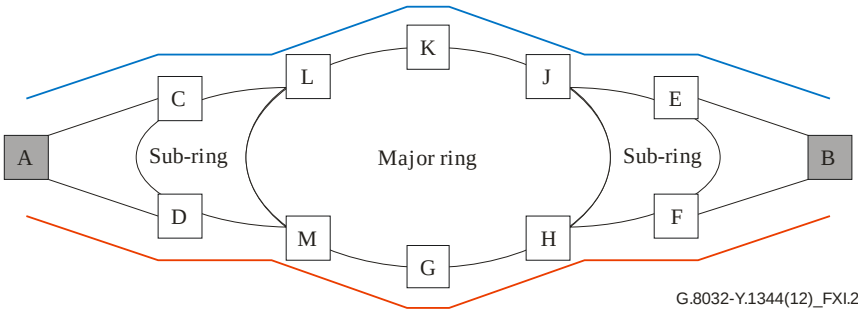

Figure XI.2 – End-to-end service resilience over interconnected rings

XI.3 Access sub-ring connected to major ring

Referring to the service shown in Figure XI.2 above, we can also imagine that the end-to-end protection is realized by using sub-ring C-A-D connected to major ring C-D-H-F-E-G. However, in this sub-ring, node A does not support ERP functionality and, as a result, is excluded from R-APS communication. Therefore, this sub-ring is a modified version of the sub-ring as presented in the main body of this Recommendation and is referred to as an access sub-ring.

XI.3.1 Basic configuration

Figure XI.3 shows the access network with access sub-ring B-A-C connected to major ring B-C-D-E-F-B. Their ERP instances are shown in Figure XI.4.

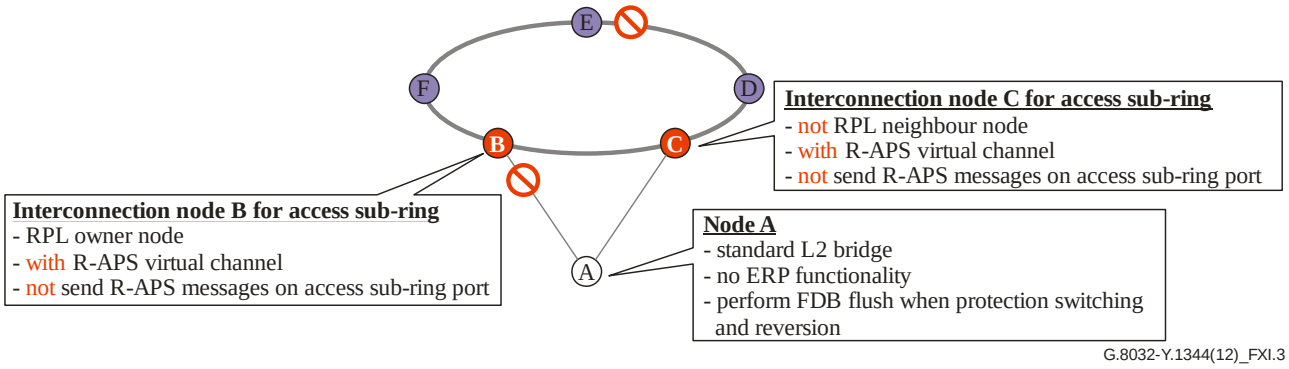

Figure XI.3 – Access sub-ring connected to major ring

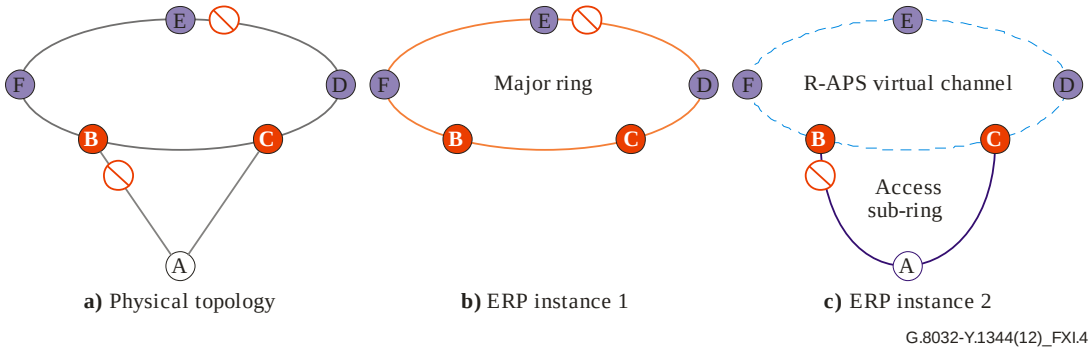

Figure XI.4 – ERP instances

In this network, Ethernet ring nodes B and C are the interconnection nodes connecting the major ring and access sub-ring; node A is a user node that does not support an ERP control process. RPLs of the major ring and access sub-ring are located on E-D link and B-A link, respectively. The RPL owner node of access sub-ring is node B.

The characteristics of the access sub-ring are as follows:

I. Node A is excluded from R-APS communication, i.e., does not generate nor transfer R-APS messages.

II. Interconnection nodes B and C do not send any R-APS messages on their access sub-ring port.

III. Interconnection nodes B and C must be able to notify node A when protection switching and reversion is invoked on the access sub-ring. (The notification is needed to trigger the FDB flush when a failure occurs or when it is recovered. This generic flush request should comply with the standard requests that are supported by node A.)

IV. Node A should be able to perform an FDB flush when protection switching and reversion are invoked.

V. Access sub-ring must configure an R-APS virtual channel on the major ring.

(This is because node A cannot receive and transfer any R-APS messages, so the R-APS message cannot be received from the access sub-ring ports of interconnection nodes B and C.)

VI. Access sub-ring should not configure an RPL neighbour node.

(If interconnection node C is configured as RPL neighbour node, both access links are blocked.)

XI.4 Non-ERP node connected in a major ring

Ensuring end to end resilience when connecting into a different technology domain can be made simpler by allowing for a non-ERP node to be located within a major ring. This clause provides guidelines on how to support such a configuration.

XI.4.1 Basic configuration

Figure XI.5 shows a network with Ethernet ring A-B-C-D-E-F-G where node D does not support ERP functionality, but is a VLAN bridge:

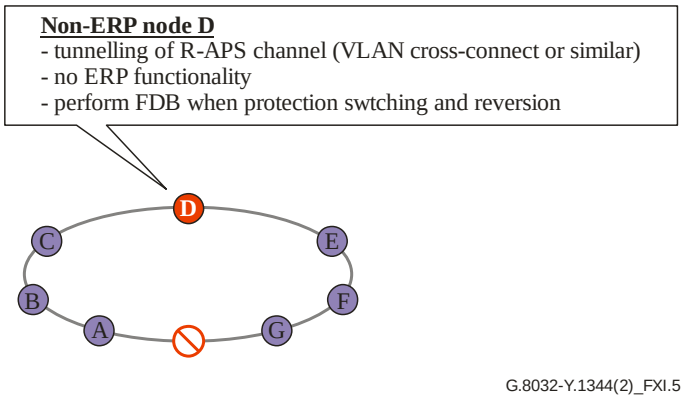

Figure XI.5 – Non-ERP node as part of a major ring

The characteristics of a non-ERP node in an Ethernet ring are as follows:

I. Non-ERP node D tunnels R-APS communication between Ethernet ring nodes C and E.

II. Ethernet ring nodes C and E must be able to notify node D when protection switching and reversion is invoked on the Ethernet ring. (The notification is needed to trigger the FDB flush when a failure occurs or when it is recovered. This generic flush request should comply with the standard requests that are supported by node D.)

III. Non-ERP node D performs an FDB flush when protection switching and reversion is invoked.

XI.4.2 Principles of operation

The principles of clause 10 are used. It should be noted that when a failure occurs on a ring link between an ERP and non-ERP node, all three burst messages triggered by the state machine on a "Tx R-APS()" action are required to ensure successful operation of block/unblock and flush operations. One example of this is the requirement of the RPL owner node to unblock the RPL upon receipt of R-APS(SF) – there is no guarantee that the RPL port would be unblocked for R-APS frames fast enough that the second and third R-APS(SF) frames would pass through.

In this case, implementations may optionally include a management configuration option ETH_C_MI_RAPS_Pass_Thru that is applied to the R-APS block logic and R-APS request processing of Figure 10-1. Note that this implies the presence of the R-APS block logic in the ERP control process of ring nodes on the Ethernet ring. The additions to Figure 10-1 to support this functionality are shown in red (and as dashed lines) in Figure XI.6 below:

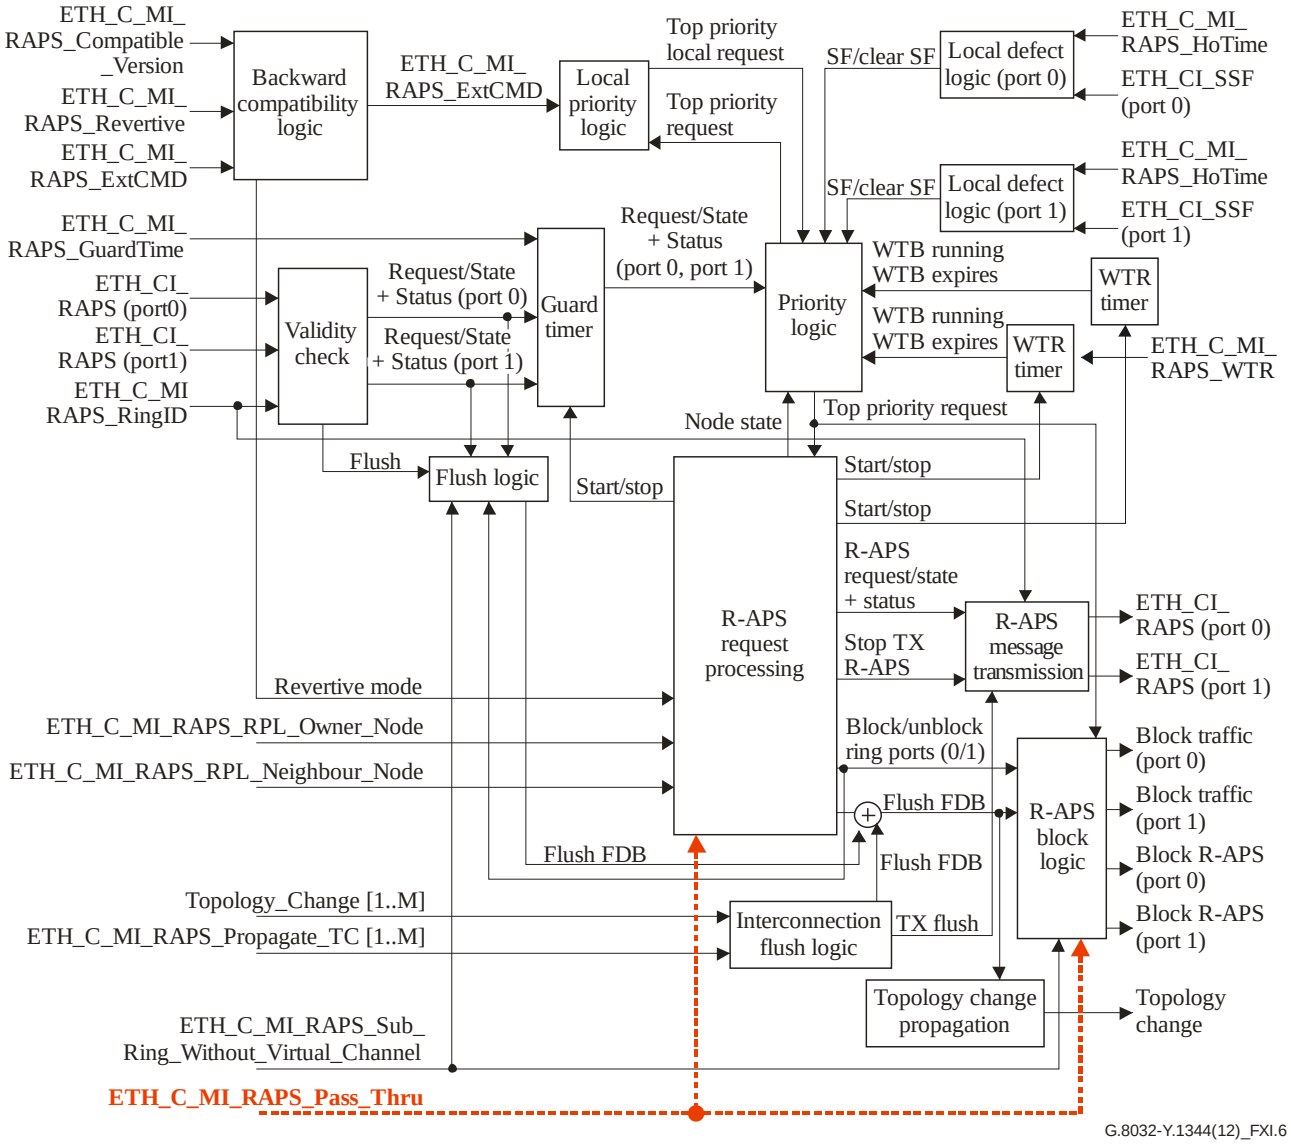

Figure XI.6 – Decomposition of ERP control process

The management information supplied in ETH_C_MI_RAPS_Pass_Thru may affect the blocking of the R-APS channel at the RPL. This MI may be either disabled or enabled. When disabled, the RPL is blocked and R-APS messages are not transmitted. This is the default value and behaviour. When the MI is enabled, the RPL owner node and the RPL neighbour node block the RPL traffic channel and will terminate any R-APS messages intended for the RPL, however the R-APS channel will not be blocked. If the action indicated in the state machine, Table 10-2, indicates that the RPL owner node or the RPL neighbour node should unblock the RPL link then, if the MI is enabled, they will additionally transmit a copy of the last received R-APS message over the RPL immediately.

Bibliography

[b-ITU-T G.8011] Recommendation ITU-T G.8011/Y.1307 (2009), *Ethernet service characteristics*.

[b-ITU-T G.8011.1] Recommendation ITU-T G.8011.1/Y.1307.1 (2009), *Ethernet private line service*.

[b-ITU-T G.8011.2] Recommendation ITU-T G.8011.2/Y.1307.2 (2009), *Ethernet virtual private line service*.

[b-ITU-T G.8031] Recommendation ITU-T G.8031/Y.1342 (2009), *Ethernet linear protection switching.*

[b-IEEE 802.1D] IEEE Std 802.1D™-2004, *IEEE Standard for Local and Metropolitan Area Networks – Media Access Control (MAC) Bridges.*

[b-IEEE 802.3] IEEE Std 802.3-2008, *IEEE Standard for Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications.*

ITU-T Y-SERIES RECOMMENDATIONS

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

GLOBAL INFORMATION INFRASTRUCTURE	
General	Y.100–Y.199
Services, applications and middleware	Y.200–Y.299
Network aspects	Y.300–Y.399
Interfaces and protocols	Y.400–Y.499
Numbering, addressing and naming	Y.500–Y.599
Operation, administration and maintenance	Y.600–Y.699
Security	Y.700–Y.799
Performances	Y.800–Y.899
INTERNET PROTOCOL ASPECTS	
General	Y.1000–Y.1099
Services and applications	Y.1100–Y.1199
Architecture, access, network capabilities and resource management	Y.1200–Y.1299
Transport	**Y.1300–Y.1399**
Interworking	Y.1400–Y.1499
Quality of service and network performance	Y.1500–Y.1599
Signalling	Y.1600–Y.1699
Operation, administration and maintenance	Y.1700–Y.1799
Charging	Y.1800–Y.1899
IPTV over NGN	Y.1900–Y.1999
NEXT GENERATION NETWORKS	
Frameworks and functional architecture models	Y.2000–Y.2099
Quality of Service and performance	Y.2100–Y.2199
Service aspects: Service capabilities and service architecture	Y.2200–Y.2249
Service aspects: Interoperability of services and networks in NGN	Y.2250–Y.2299
Numbering, naming and addressing	Y.2300–Y.2399
Network management	Y.2400–Y.2499
Network control architectures and protocols	Y.2500–Y.2599
Smart ubiquitous networks	Y.2600–Y.2699
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899
Carrier grade open environment	Y.2900–Y.2999
Future networks	Y.3000–Y.3099

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

SERIES OF ITU-T RECOMMENDATIONS

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

Printed in Switzerland
Geneva, 2012