Ethernet Ring Protection for Carrier Ethernet Networks

Jeong-dong Ryoo, ETRI
Hao Long and Yang Yang, Huawei Technologies
Marc Holness, Nortel Networks
Zahir Ahmad and J. Kevin Rhee, Information and Communications University

ABSTRACT

Ethernet technologies are rapidly gaining importance as a prevailing solution for carrier networks. Ethernet ring protection switching, defined in ITU-T Recommendation G.8032, provides a means of achieving carrier network requirements for ring network topologies. This article outlines the novel concepts in this Ethernet ring protection switching mechanism and discusses the fundamental operating principles by which the automatic protection switching (APS) protocol works. In addition, several feature enhancements in protection behavior and solutions being considered for the next phase of standardization are discussed.

INTRODUCTION

Ethernet as a carrier-class technology continues to make considerable progress within carrier networks. In fact, the majority of service providers consider Ethernet technology mature enough and cost effective for carrier-class network deployments. According to a recent market study report [1], more than 75 percent of respondents from service providers have a strategy of using Ethernet instead of synchronous digital hierarchy (SDH) or synchronous optical network (SONET) for accessing and collecting customer traffic. Currently, Ethernet technology development has been challenged by service providers (as well as enterprises) that need rapid restoration capability to guarantee carrier-grade availability attainable in ring network architectures. Although SDH/SONET rings undeniably provide fast protection switching, synchronous equipment is more expensive and less flexible than Ethernet-based equipment.

International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) Question 9 of Study Group 15 has developed a technically feasible, economically viable, and scalable solution to provide rapid service restoration that delivers SDH/SONET-grade resilience at the cost of Ethernet. Such an effort was successfully introduced by the recent Recommendation G.8032 for Ethernet ring protection switching [2], as well as G.8031 for Ethernet linear protection switching [3]. The advent of G.8032 seems to trigger wide acceptance of Ethernet ring networks by service providers, since 93 percent of surveyed service providers use Ethernet collector rings to connect to customer sites [1], and approximately three quarters have deployed Ethernet overlay networks to deliver Ethernet services.

Resilient packet ring (RPR), defined in IEEE 802.17 [4], is a competing technology, which is a metropolitan area network (MAN) technology supporting data transfer among stations interconnected in a dual-ring configuration. The RPR network provides connectivity across many sites using a shared packet aware infrastructure, which enables a large reduction in fiber requirements compared with mesh networks. The values of RPR include a 50-ms protection switching time and better management of excess information rate (EIR) traffic under traffic congestion and protection scenarios.

A new medium access control (MAC) of IEEE 802.17 was developed to achieve its protection switching objectives. It introduces a new MAC header, which is not compatible with ubiquitous Ethernet, and a new set of complex protocols and algorithms (topology discovery, fairness, etc.). These features contribute to its complexity and development/deployment cost, and thus lack of economic viability.

Ethernet ring protection (ERP) defined by G.8032 has been developed on a principle of utilizing generic mechanisms inherited from the traditional Ethernet MAC and bridge functions. The objective of fast protection switching is achieved by integrating mature Ethernet operations, administration, and maintenance (OAM) functions and a simple automatic protection switching (APS) protocol for Ethernet ring networks. In addition, since ERP is based on standard Ethernet, it can take advantage of the rapidly increasing Ethernet bandwidth-cost merits of 1–100 Gb Ethernet (GbE), due to their wide commoditization. In addition, since Ethernet, and thus ERP, is virtually agnostic to all

0163-6804/08/$25.00 © 2008 IEEE

physical/server layer technologies, it can be supported by any carrier's network infrastructure. As a result, it is a more deployable and economically viable solution than the RPR technology.

The G.8032 protocol is designed for ring topologies and developed as a standardized alternative to replace the spanning tree protocol (STP) to change the port status without requiring complex computation, provisioning overhead, and excessive information exchange, so as to achieve much faster (i.e., sub-50-ms) protection switching. As STP is a general (mesh) protocol applicable to any kind of network, it does not have any optimization for ring topologies. STP requires much more time (i.e., on the order of seconds) to rebuild topology because it needs extensive information exchange for tree computing. Particularly for a ring, tree computing simply results in the selection of one port to be blocked. In order to single out just one port to be blocked, such time-consuming tree computation should not be necessary. The G.8032 protocol is focused on producing an optimized process to handle ring protection and should perform better than all variants of STPs, including Rapid STP (RSTP) and Multiple STP (MSTP), in ring topologies.

In summary, G.8032 ERP is developed to meet the following objectives:

- To provide efficient network connectivity
- To provide rapid service restoration (sub-50 ms)
- To support multiple Ethernet services (e.g., E-Line, E-Tree, and E-LAN [5])
- To be client- and server-layer agnostic; that is, G.8032 ERP can be supported over (virtually) any physical and server layer and can transport (virtually) any Ethernet client
- To utilize existing IEEE 802.1 bridging and IEEE 802.3 MAC hardware, thus being a simple software increment on existing Ethernet switching equipment
- To support flexible deployment models in access, metro, and core network applications
- To leverage Ethernet's broad PHY bandwidth (e.g., 1/10/40/100GbE), low cost, and time-to-market competency in support of cost-effective and large bandwidth rings
- To introduce lower operational expenses (OPEX) and capital expenses (CAPEX) for service providers
- To utilize standardized Ethernet OAM functionalities as defined in Y.1731 [6] and IEEE 802.1ag [7]

Ethernet Ring Protection Mechanism

Ring Topology

An Ethernet ring is a collection of ring nodes forming a closed loop whereby each node is connected to two adjacent nodes via duplex communications links. The topology of an ERP network can be a single ring or a collection of interconnected rings. Mitigation of frame duplication is required to maintain the quality of service provided by the MAC service. Consequently, the prevention of loops being formed by the logical Ethernet ring over the physical closed loop is required. Since a time-to-live (TTL) field is not defined within a native Ethernet frame, traffic looping transported over the ring is prevented by blocking traffic at one of the ring ports. Therefore, a physical ring will maintain a logical (nonlooping) linear MAC topology with dynamic assignment of end nodes by the ring APS (R-APS) protocol. Recommendation G.8032 does not limit the number of nodes on a ring, but recommends the maximum number of nodes to be in the range between 16 and 255 nodes from an operational perspective. Spanning tree protocols are not used on the Ethernet ring network; they are replaced by R-APS.

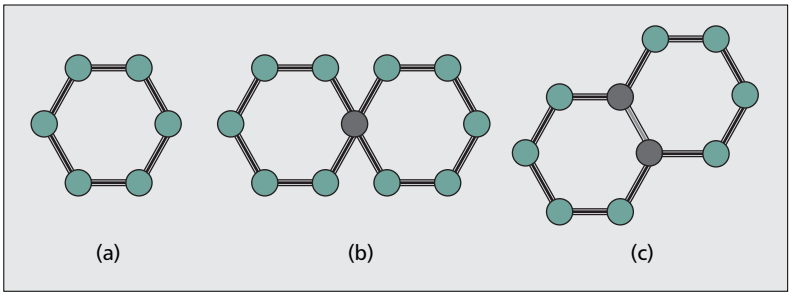

Figure 1. *Ring topology variants: a) single ring; b) ring interconnect via shared node; c) ring interconnect via shared link.*

Figure 1 shows the possible variants of Ethernet ring topology. The current G.8032 Version 1 supports a single ring, as depicted in Fig. 1a. As for ring interconnect scenarios depicted in Figs. 1b and 1c, ring interconnection can be achieved via either a shared node or a shared link. Subsequent development of G.8032 (e.g., G.8032 Version 2) will address the aforementioned ring interconnect scenarios.

Loop Avoidance

In a normal state, one of the ring links is designated as the ring protection link (RPL), which blocks Ethernet traffic to avoid traffic looping. An RPL block is provided by port blocking at either end of the RPL. The node that sets the block is referred to as an RPL owner. An RPL owner plays a very important role in G.8032 ERP, as it is responsible for preventing loops on the ring in a normal operating state. When a link failure occurs, each node that detects the failure blocks the port for the failed link and sends R-APS messages with signal fail indication (R-APS(SF)). The messages are disseminated over the ring. When the RPL owner receives the R-APS(SF) message, it removes the block from the port to RPL, resulting in changes to the ring topology to achieve maximal connectivity of ring nodes.

Filtering Database Flush

Once ring port blocks are relocated due to failure or recovery, the filtering database (FDB) at every ring node is flushed, meaning that all MAC addresses and their port associations for traffic forwarding are cleared from the FDB. After an FDB flush, new FDB entries are added as a result of source address learning from the traffic using the new ring topology. Unlike traditional linear protection and SDH multiplex section protection ring, the protection entity is not

The Node ID field consists of 6 bytes which contain the source node MAC address. After the Node ID field, the Reserved 2 field is defined, which is 24 octets long and set to zero for this version. Any TLVs are not defined in current version of G.8032.

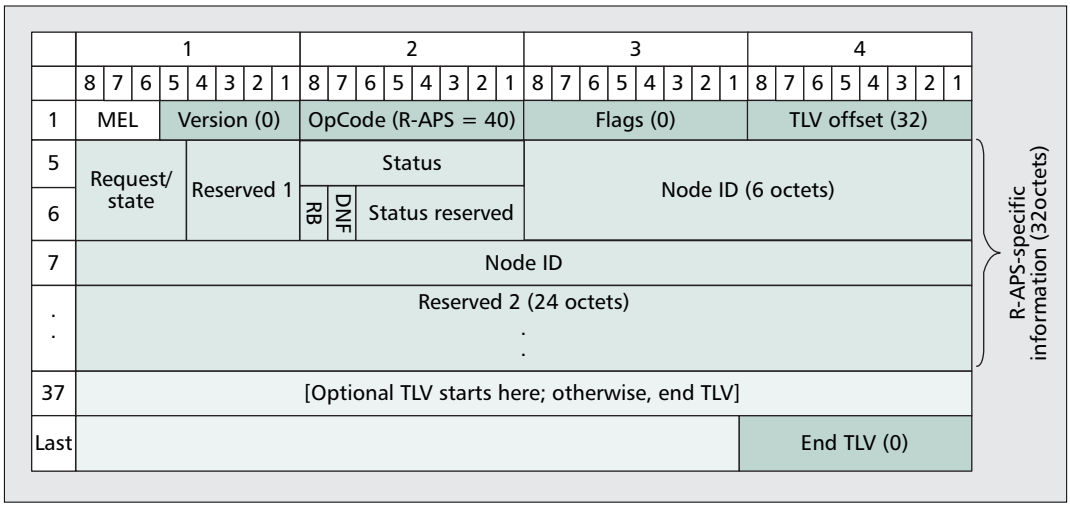

Figure 2. *R-APS PDU format. MEL: maintenance entity group level, TLV: type, length and value, RB: RPL blocked, DNF: do-not-flush.*

Request	Type	Priority
Local SF	Local	Highest
Local clear SF	Local	|
R-APS (SF)	Remote	|
WTR expires	Local	|
WTR running	Local	|
R-APS (NR, RB)	Remote	|
R-APS (NR)	Remote	Lowest

Table 1. *Protection switching requests ordered by priority.*

preconfigured on the Ethernet ring. G.8032 makes use of traffic flooding and the address learning mechanism of generic Ethernet bridging. Whenever the position of the ring port block changes, an FDB flush operation should occur at each ring node.

Frame Format of R-APS

Figure 2 describes the R-APS protocol data unit (PDU) format of G.8032, which is framed in the Ethernet OAM-PDU format used in ITU-T Recommendation Y.1731 [6]. In the MEL field, the maintenance entity group (MEG) level (MEL) of the R-APS PDU is specified. The Version, Flags, and END TLV fields are set to 0×00 in the current version, and these fields should be ignored upon reception. The OpCode field of an R-APS PDU is set to 40. The TLV Offset field contains the value of the offset to the first TLV, and its value is determined to be 32.

For R-APS-specific information, 32 octets are allocated. The first 4 bits are for request/state information; the value '1' '0' '1' '1' represents the signal fail (SF), and '0' '0' '0' '0' represents no request (NR). The rest of he assignments are reserved for future standardization. The Reserved 1 field consists of 4 bits, which are currently set to '0' '0' '0' '0' and reserved for future extension of requests or indication of the protection type. The Status field includes the status information. Currently, two status bits are defined and the remaining six bits reserved. The RPL blocked (RB) status bit, which is set by the RPL owner, is enabled when the RPL is blocked. The do-not-flush (DNF) status bit is enabled when FDB flush is not necessary.

The Node ID field consists of 6 bytes that contain the source node MAC address. After the Node ID field, the Reserved 2 field is defined, which is 24 octets long and set to zero for this version. No TLVs are defined in the current version of G.8032.

Requests of Protection Switching

Protection switching is triggered by R-APS requests, which are specified in R-APS message PDUs, and the requests generated by local events. The four defined local events — local signal failure (local SF), local clear signal failure (local clear SF), wait-to-restore expire (WTR Expire), and wait-to-restore running (WTR Running), — are general events for most other protection switching technologies. The first version of G.8032 defines three types of R-APS messages:

- R-APS(SF) is sent from the node detecting link failure (i.e., from the node encountering a local SF request).
- R-APS(NR) is transmitted by the node that detects link recovery (i.e., the node that gets a local clear SF request).
- R-APS(NR, RB) is sent by the RPL owner, and indicates that the ring is in the normal state and the RPL is blocked.

All these messages are periodically transmitted on the ring until any new higher-priority request is present.

Table 1 lists all protection switching requests ordered by priority. When the ERP control process of a ring node receives multiple outstanding requests, it only responds to the one with the highest priority.

Ring Node and R-APS Channel Models

Messaging R-APS protocol requires a virtually isolated and designated transmission channel defined as an R-APS channel. As shown in Fig. 3, which is an illustration of a schematic model of ring node architecture, an R-APS channel is separated from a service traffic channel. An R-APS channel is configured by use of a separate virtual LAN (VLAN) assignment so that R-APS messages can be processed separately from service traffic.

Both R-APS and service traffic channel blocks are placed at the same location. In Fig. 3, if a ring node installs a port block on a ring port, the corresponding R-APS and service channel VLAN identifier (VID) filtering functions are set to drop R-APS PDUs and service data units (SDUs), respectively. The bridge function for VID filtering is defined by IEEE 802.1Q-2005 [8]. In this arrangement there are two important implications. First, even though the R-APS PDUs are blocked from being forwarded, R-APS message signals (ETH_CI_RAPS) are extracted to and inserted from the ERP control process by an ETHDi/ETH adaptation function on the R-APS channel of the blocked port, regardless of the existence of a port block. This feature is important and useful when ports are blocked due to a reason other than link failure because it enables this node to exchange R-APS messages with other ring nodes. Note that the ETHDi function operating at the MEG for R-APS channel is an MEG intermediate point (MIP) function, so the R-APS message is forwarded toward the next nodes. Second, the service traffic can be forwarded between the client network link and ring link with no block even when the node has a block, so the service for client networks can be provided unceasingly.

The ERP control process controls the node behavior based on R-APS and OAM signaling from the server-level MEG endpoint (MEP) function and ETHDi/ETH adaptation functions (Fig. 3). When the server-level MEP detects a signal failure, it reports it to the ERP control process as an ETH_CI_SSF signal. The ETH_CI_RAPS signals from and to the ETHDi/ETH adaptation function provide protection switching information exchange with other nodes.

Failure Detection

Protection switchover occurs on detection of failure on a link of a ring. The link failure can be monitored by an Ethernet continuity check (ETH-CC) function [6]. Physical layer (or server layer) failure conditions can also be reported to the ERP control process. When ETH-CC is utilized, two end ports of a link form an MEG, and an MEP function is installed on each ring port. Periodic exchange of continuity check messages (CCMs) every 3.3 ms is activated between pairing MEPs to monitor link health. For link monitoring, the MEL, which defines the server-client relation in the Ethernet OAM model, is assigned to be lower than the MEL of the R-APS channel when the MEG levels are shared, so the link monitoring MEG becomes the server for an R-APS process. For further discussions and applications of Ethernet OAM, refer to [9]. When an MEP detects such a failure, it signals the ERP control process to initiate protection switching. A node failure is regarded as failure of the two links attached to the node. The two nodes adjacent to the failed node detect failures on the two links connected to the failed node and trigger protection switching.

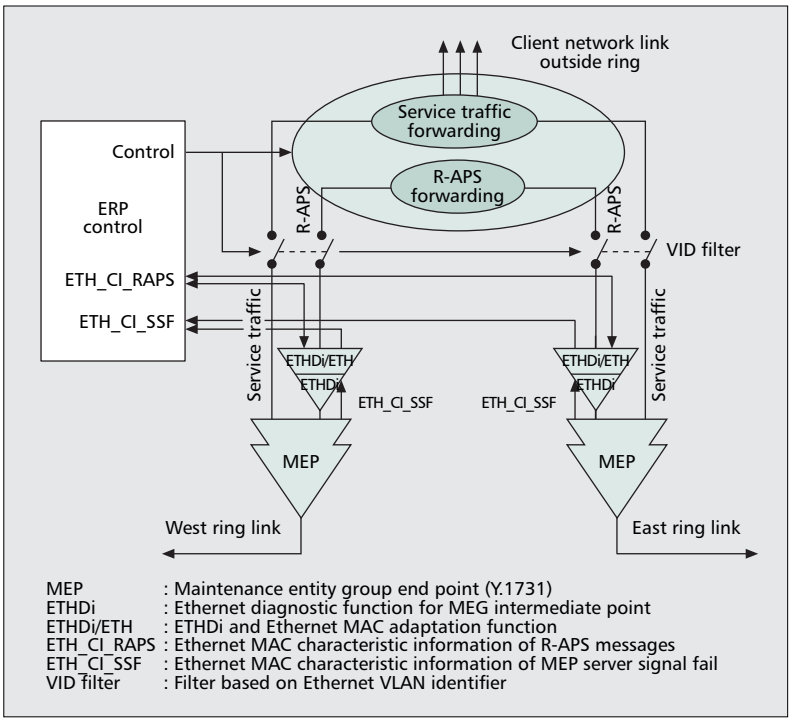

Figure 3. *A schematic model of a ring node and R-APS channel.*

Reversion on Recovery

After recovery from all failures, the port blocking returns to the RPL owner in revertive-mode operation. Since the position of the RPL on the ring is selected to optimize the use of network resources, the revertive operation is desirable. However, it costs additional disruption to traffic services. In order to avoid an erroneous switching operation that may be caused by intermittent failures, a wait-to-restore (WTR) timer is adopted. When the RPL owner recognizes a failure recovery by reception of an R-APS(NR) message from a node at one end of a recovered link, it starts its WTR timer. If any local or remote failure is detected before expiration of the WTR timer, the WTR timer and reversion process are aborted. When the WTR timer expires, the RPL owner blocks its port for the RPL and instructs the node connected to the recovered link to remove its block by sending an R-APS(NR,RB) message.

A Sample Scenario with a Single Failure and Recovery

Figure 4 illustrates a scenario with a single failure and recovery. In the normal state the RPL owner (node A) block is in place at its port connected to the RPL (link between nodes A and B). When a ring link failure occurs between nodes D and E, nodes D and E detect the "local SF" condition; both nodes block the failed port and transmit R-APS(SF) messages on both ring

Manual switch is used to move the port block from the RPL to a port on a different ring link while there is no failure on the ring. Forced switch moves the port block from the RPL to a different ring link no matter whether or not there exists any type of failure on the ring.

ports. In turn, both nodes flush their FDBs. R-APS(SF) messages are transmitted periodically while the SF condition persists. Other nodes flush FDBs on receiving an R-APS(SF) message. When the RPL owner receives an R-APS(SF) message, it flushes its FDB and unblocks its port on the RPL. After the failed link between nodes D and E recovers, nodes D and E send R-APS(NR) messages periodically and start guard timers, which are used to prevent ring nodes from receiving outdated R-APS messages generated before starting guard timers. When the RPL owner receives an R-APS(NR) message, it starts the WTR timer. When the WTR timer expires, the RPL owner blocks its port on the RPL, sends R-APS(NR, RB) messages, and flushes the FDB. Each node flushes its FDB on reception of the first R-APS(NR, RB) message after recovery. Upon receiving an R-APS(NR,RB) message, nodes D and E remove blocks from their recovered ports, stop transmitting R-APS(NR) messages, and flush the FDBs. Finally, the ring returns to the normal state.

Future Enhancements for Ethernet Ring Protection

To enhance features and performance of the first version of G.8032 Ethernet ring protection switching, there are several areas identified for future work. The following future study areas are listed in order of the work priority agreed on by the working party for G.8032 Version 2 at the time of publication of this article.

Non-Revertive Mode Operation on Recovery

In the non-revertive mode, the blocked ports are not returned to the RPL owner even though all failure links have recovered. However, the operator can trigger reversion by an external command. In a single link failure, the blocked ports on the failed and recovered link can remain blocked, which can be achieved without any complication. However, when multiple links recover concurrently from a failure, all but one of the recovered links should be unblocked. In order to select the only node that does not unblock its port, a Node ID-based non-revertive mechanism is proposed and being considered for future study of G.8032. In this scheme nodes adjacent to a recovery send R-APS(NR) messages with the Node ID of the node. Another node receiving this R-APS(NR) compares its Node ID with the Node ID of the source node of the message. If it is lower, the node can unblock the port. Accordingly, the Node ID can behave as the priority assignment for node reversion on concurrent multiple failure recovery.

FDB Flush Optimization

An FDB flush operation always causes traffic flooding on the ring. Since traffic flooding requires more link capacity, flush operations, in principle, should be avoided as much as possible. The current state machine model of G.8032 describes how to trigger FDB flush operations, mainly to secure protection switching itself. In fact, it has been recognized that flush operations are not required for some scenarios, including:

- Failure or recovery of the RPL
- Failure or recovery of the nodes adjacent to the RPL

In both scenarios, the protection switching will not change the active logical topology of the ring as described in Fig. 5; thus, flush is not necessary.

As a result, for the cases in Fig. 5, the DNF status bit is set in the R-APS(SF) or R-APS(NR, RB) message PDU to suppress FDB flush operation. The detailed usage of this indication will be defined in the future. As an example, when a node detects an RPL link or RPL owner failure, it will send R-APS messages with DNF indications that just trigger state transition without FDB flush. In a multiple failure more delicate optimization can be sought.

Manual Switch and Forced Switch

Besides protection switching due to failures, an operator control can initiate protection switching. Such administratively triggered switchovers are manual switch and forced switch. Manual switch is used to move the port block from the RPL to a port on a different ring link while there is no failure on the ring. Forced switch moves the port block from the RPL to a different ring linkwhether or not there is any type of failure on the ring. Manual switch and forced switch are removed by a Clear command issued by an operator control. The ring will then switch back to a normal state. Manual and forced switch are considered temporary commands and do not change the assignment of the RPL permanently.

Interconnected Rings

There are two kinds of interconnected rings, as shown in Fig. 1. Rings interconnected by sharing a node can operate isolated ERPs independent of each other. In this case the inter-ring traffic service is not protected from shared node failure. When multiple rings are interconnected by sharing a link or links with more than one shared node, inter-ring traffic can be protected from a shared node failure. However, the complexity of ERP control increases in order to handle shared node and link failures. An example of such complexity is a super loop problem that happens when two interconnected rings locate the port blocks on the shared link simultaneously.

Multiple Incidences per Single Ring

A physical Ethernet ring can provide multiple ERP service rings separated by VLANs, so a ring can serve multiple clients that share the resources of the ring. In this case possible configurations are not limited to having each ERP incidence protect a single client ring; a multiple VLAN client ring can be protected by an ERP incidence. In addition, there can be multiple ERP incidences that can protect multiple groups of VLAN clients. One of the benefits of having multiple incidences is that ring traffic can be evenly distributed in a normal state when the RPLs of different ERP incidences are distributed, so a ring can serve more client traffic. How-

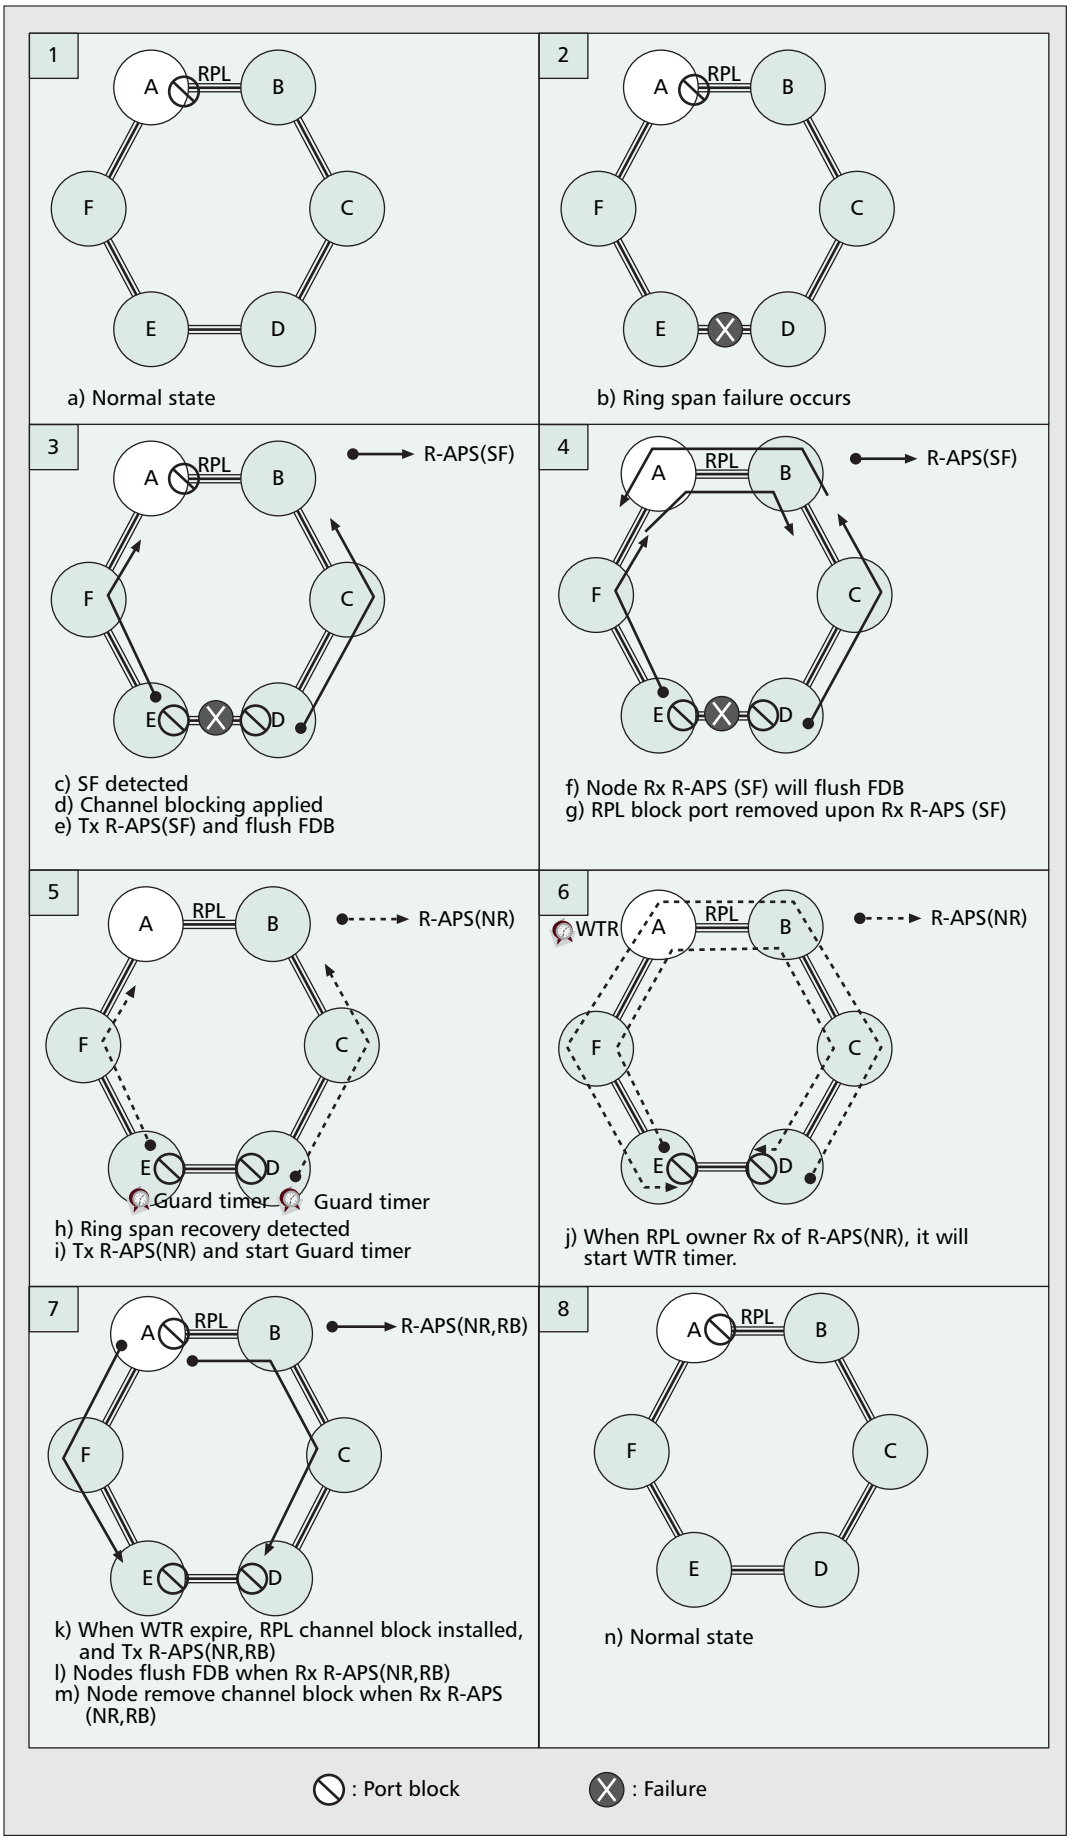

■ **Figure 4.** *An example ERP sequence for a single link failure and recovery.*

The complexity of the ERP control increases in order to handle shared node and link failures. An example of such complexity is a super loop problem that happens when two interconnected rings locate the port blocks on the shared link simultaneously.

In order to avoid flooding, the RPL is blocked at both ends in a normal state. When a failure occurs, both blocked ports will be opened for traffic protection. Similarly, when unidirectional failure occurs, both of the ends also should be blocked.

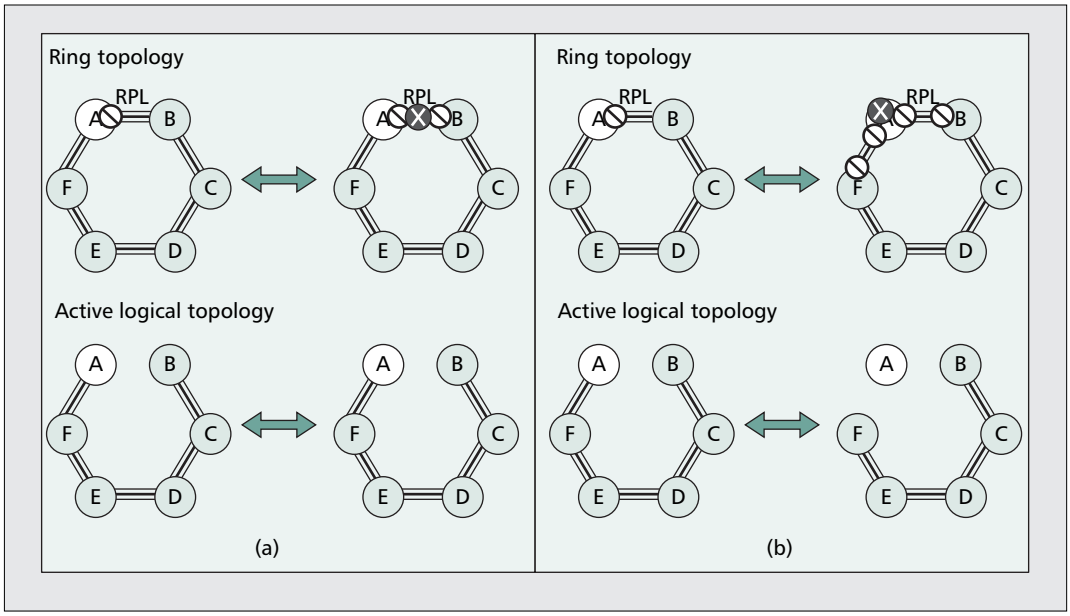

Figure 5. *Example scenarios for flush optimization: a) failure and recovery of the RPL; b) failure and recovery of the RPL owner.*

ever, this benefit disappears in a protection state as the block positions of all ERP incidences coincide at the failed link.

Dual End Blocking

In G.8032 Version 1, one end of the RPL is blocked for breaking the loop in a normal state. In this case the traffic would be flooded on the link from the unblocked end. Thus, the RPL capacity is always wasted by flooded traffic. It is not a problem when the ring is occupied by only one ERP incidence. But in many cases, there may be multiple ERP incidences or other kinds of services sharing the link. The traffic flooded on the RPL link would compete for bandwidth resources with other services.

In order to avoid flooding, the RPL is blocked at both ends in a normal state. When a failure occurs, both blocked ports will be opened for traffic protection. Similarly, when unidirectional failure occurs, both of the ends should also be blocked. In this case Ethernet remote defect indication (ETH-RDI) [6] or other functions may be used to notify the other end of the unidirectional defect. To synchronize the configuration of the two-port block at both ends of the RPL, it is suggested that a block at one end of the RPL be provisioned in a normal state and the other block at the other end automatically determined by the network management system (NMS) or other initialization mechanism. The new mechanism should be compatible with that of a single-end blocking architecture.

RPL Replacement

The RPL can be changed permanently to any link other than the original RPL. The RPL replacement command moves the position of the RPL by blocking a newly designated link and unblocking the original RPL permanently. Also, the functionality of the RPL owner is transferred to the corresponding node that is connected to the new RPL.

Protection for Link Aggregation Group Failure

When a logical link is composed by a link aggregation group (LAG) and one or more members of the LAG fail, the LAG continues to deliver traffic with reduced capacity. However, this may not be the desired behavior for a protected network. In such an LAG partial failure, the traffic would be better served by switching to the protection. For the detection of an LAG failure, ETH-CC may verify only one of the physical links, so it may not detect reduction of link capacity by failure of some other physical links. There are two proposed solutions under study:

- The ring link monitoring runs per physical port.
- When the Ethernet link aggregation function detects a failure from one or more links in the group, the whole group should be shut down; this approach is similar to that used in SDH virtual concatenation without the Link Capacity Adjustment Scheme (LCAS).

Filtering Database Flip

FDB flip is an alternative technique that can replace FDB flush. On FDB flush, after placing a block in a new position, either due to failure or for reversion, all client traffic is flooded as there is no address forwarding information in the FDB. This traffic flooding creates traffic volume several times greater than the steady-state traffic that can be reached only after the FDB completes address learning [10]. When such flooded traffic volume is far greater than the link capacity, a majority of frames can be lost or delayed due to queuing in a buffer. In this situation a twofold network impairment manifests itself as extended delay and increased loss of client traffic. In addition, the burst of traffic flooding extends the address learning period. The combination of all of these impairments can

make R-APS switching and settling time longer than 50 ms. The phenomenon can be critical when a ring provides services to a large number of hosts. In the proposed FDB flip method an R-APS flip message provides information on how the FDBs at other nodes should be modified so that the FDBs provide optimized forwarding immediately, as described in [10].

CONCLUSIONS

Benefiting from recent progress in Ethernet technology, standards, applications, and initial deployment, carrier Ethernet is well poised as the next packet transport network infrastructure for metro service providers. Given the wide deployment of rings with wavelength-division multiplexing (WDM) and SONET/SDH systems, for these networks an Ethernet ring protection network is an attractive upgrade strategy for efficient packet transport networking.

Recommendation G.8032 for ERP technology can realize an immediate and economic solution to provide carrier-class protection for Ethernet ring networks without requiring any new Ethernet forwarding and filtering functions on the data path. Ethernet ring protection can be implemented simply by an incremental software-level change that allows the service provider to leverage economics of utilizing installed Ethernet switches. As ERP is designed to be independent of the capability of the server-layer transmission media, this new ring protection for Ethernet can run over any server-layer network that any network operator might have. Since the G.8032 ERP can support heterogeneous rings, which means not all ring spans need to be of the same bandwidth or physical layer, an upgrade strategy with ERP becomes an attractive solution. In addition, ERP can also achieve efficient bandwidth utilization of ring traffic by means of spatial reuse.

Taking all the aforementioned advantages into account, it is expected that the G.8032 ERP will be an efficient, deployable, and economically viable solution for carrier Ethernet ring networks.

ACKNOWLEDGMENT

This work at ICU was partly supported by the IT R&D program of MKE/IITA [2008-F017-01] and the ERC program of MOST-KOSEF [R11-2000-074-02006-0].

REFERENCES

[1] Infonetics Research, "Service Provider Plans for Metro Optical and Ethernet: North America, Europe, and Asia Pacific 2007," Sept. 2007.
[2] ITU-T Rec. G.8032/Y.1344, "Ethernet Ring Protection Switching," 2008.
[3] ITU-T Rec. G.8031/Y.1342, "Ethernet Linear Protection Switching," 2006.
[4] IEEE Std. 802.17, "Part17: Resilient Packet Ring (RPR) Access Method and Physical Specifications," 2004.
[5] MEF 6.1, "Ethernet Services Definitions — Phase 2," 2008.
[6] ITU-T Rec. Y.1731, "OAM Functions and Mechanisms for Ethernet Based Networks," 2006.
[7] IEEE Std. 802.1ag, "Local and Metropolitan Area Networks, Virtual Bridged Local Area Networks, Amendment 5: Connectivity Fault Management."
[8] IEEE Std. 802.1Q, "Virtual Bridged Local Area Networks," 2005.
[9] J.-D. Ryoo *et al.*, "OAM and Its Performance Monitoring Mechanisms for Carrier Ethernet Transport Networks," *IEEE Commun. Mag.*, Mar. 2008, pp. 97–103.
[10] J. Im, J.-D. Ryoo, and J.-K. K. Rhee, "Managed FDB Algorithm and Protection in Ethernet Ring Topology," *Proc. COIN-ACOFT '07*, Paper WeC1-1, June 2007.

BIOGRAPHIES

JEONG-DONG RYOO (ryoo@etri.re.kr) is a principal member of research staff with the Electronics and Telecommunications Research Institute (ETRI), South Korea. He holds Master's and Ph.D. degrees in electrical engineering from Polytechnic University, Brooklyn, New York, and a Bachelor's degree in electronic engineering from Kyungpook National University, South Korea. After completing his Ph.D. study in the area of telecommunication networks and optimization, he started working for Bell Labs, Lucent Technologies, New Jersey, in 1999. While he was with Bell Labs, he was mainly involved with performance analysis/evaluation/enhancement study for various wireless and wired network systems. Since he left Bell Labs and joined ETRI in 2004, his work has been focused on next-generation network and carrier class Ethernet technology research, especially participating in OAM and protection standardization activities in ITU-T. He co-authored *TCP/IP Essentials: A Lab-Based Approach* (Cambridge University Press, 2004). He is a member of Eta Kappa Nu.

HAO LONG (lonho@huawei.com) is a member of network research staff at Huawei. In the past couple of years he has been focused on research and development of packet transport networks. He is an active contributor in several ITU-T SGs in the transport network area including SG15 Q9 and Q11. He received his M.S. degree in computer software from Huazhong University of Science and Technology, Wuhan, Hubei, China.

YANG YANG (healthinghearts@huawei.com) is a member of network research staff at Huawei. He has eight years of experience in the area of transport networks. He actively participates in several standards groups, including ITU-T and IETF. He is currently focusing on development of packet transport networks and next-generation optical transport networks. He received his B.S. degree in telecommunication engineering from Xidian University, Xi'an, Shanxi, China.

MARC HOLNESS (holness@nortel.com) is a member of scientific staff at Nortel. He has 20 years of telecommunications industry experience spanning voice and data networking, system architecture, design engineering, project management, and development of real-time memory constraint systems. His current focus area is the application of Ethernet technologies within the carrier's network. He received his B.A.Sc. degree in computer science engineering from the University of Toronto.

ZAHIR UDDIN AHMAD (zahir@icu.ac.kr) is pursuing his M.S. degree at the Information and Communications University (ICU), Korea. He received his B.S. degree in computer engineering from American International University-Bangladesh in 2004. He actively participates in the standardization efforts on Ethernet ring protection in ITU-T. His research interests include high-speed network architecture, protocol development, and related areas.

JUNE-KOO KEVIN RHEE (rhee.jk@ieee.org) is an associate professor at ICU, Korea, and a graduate of Seoul National University with B.E. (1988) and M.Sc. degrees (1990), and the University of Michigan, Ann Arbor, with a Ph.D. degree (1995), all in electrical engineering. Prior to his current position, he was with Princeton University (1995–1996), NEC Research Institute (1996–1998), Corning Incorporated (1998–2002), and Samsung Advanced Institute of Technology (2003–2005). Early in his career he made substantial contributions in the area of optical communications including the first demonstration of the R-OADM concept, the first introduction of DPSK optical data transmission in WDM systems, and the first demonstration of a WDM optical protection switching network. His publications have been cited more than 500 times by other technical articles. In recent years he has been an active contributor to the ITU-T SG13 and SG15 working parties, particularly on G.8032 ERP technology. His current research interests include carrier-class Ethernet, 100 GbE, wireless mesh networking, and optical packet switching.

Since the G.8032 ERP can support heterogeneous rings, which means not all ring spans need to be of the same bandwidth or physical layer, an upgrade strategy with ERP becomes an attractive solution. In addition, ERP can also achieve efficient bandwidth utilization of ring traffic by means of spatial reuse.