Protection Proposal

To IEEE 802.17

Edited by

Jason Fan, Luminous

John Lemon, Lantern

Vittorio Mascolo, Alcatel

Harry Peng, Nortel

Frederic Thepot, Dynarc

1 Overview

This section describes the services provided by the RPR Protection Protocol which implements a reliable, accurate, efficient, and quick means for all RPR stations on a ring to discover a broken segment in a ringlet of an RPR and to steer RPR frames away from the break until the break is healed. The RPR Protection Protocol uses local physical layer triggers, manual triggers, and loss of Keep_Alive control messages to start a protection action. Changes in the link status information are broadcast to all stations on the ring via a Link_Status message. The topology image generated by the RPR Topology Discovery Protocol is used to locate a segment failure within the topology image of the ring.

2 Scope

This RPR Protection Protocol implements a reliable, accurate, efficient, and quick means for all RPR stations on a ring to discover a broken segment in a ringlet of an RPR and to re-route RPR frames away from the break, until the break is healed. The protocol is intended to scale from 1 to 100's of stations, to cause minimal overhead for ring traffic, and to cause minimal impact on software and ASICs. The services and features provided are:

- Quick dissemination of loss of connectivity information on the ring
- Tolerance of message loss
- Operation without any master station on the ring
- Operation independent of and in the absence of any management systems
- Operation with dynamic addition and removal of stations to/from the ring
- Minimal overhead

3 Algorithm Overview

The end result of the RPR Protection Protocol is that each station on the ring knows of a ring segment failure and steers ring traffic away from the failure within 50 ms of the failure.

Broadcast of a Link_Status message is initiated by the station directly downstream from the failed or healed link (detected via physical layer trigger or keep alive control). The broadcast messages sent as part of the RPR Protection Protocol are indicated in the RPR frame header as control packets with high priority class of service.

3.1 Link_Status message

The Link_Status message reports changes in link status. They are sent as MAC Control messages, as broadcast frames, and with TTL of Max_Ring_Size. They are removed by the source station.

Byte 0	Link_Status opcode
Byte 1	ringlet_id: unsigned 8-bit integer
Byte 2	upstream_link_status: {DISCONNECTED_FS = 1, DISCONNECTED_SF = 2, DISCONNECTED_SD = 3, DISCONNECTED_MS = 4, DISCONNECTED_WTR= 5, CONNECTED = 6, UNKNOWN = 0}

Link_Status opcode: The MAC Control opcode value for a Link_Status message.

ringlet_id: The ringlet_id parameter carries the ID of the ringlet on which the message is sent.

upstream_link_status: The upstream_link_status parameters carry the current status of the in coming links from the stations clockwise to the sending station.

3.2 Keep_Alive message

The Keep_Alive message is used to detect MAC layer failure. They are exchanged periodically (TBD msec.) between neighbors to check continuity at the data link level. They are sent as MAC Control messages, as broadcast frames, and with TTL of 1. This guarantees that they will be received by any neighbor and removed from the ring immediately.

Byte 0	Keep_Alive opcode

Keep_Alive opcode: The MAC Control opcode value for a Keep_Alive message.

3.3 Initialization

At station initialization, the local topology image is learned through the RPR Topology Discovery Protocol. No protection state is enabled until the local topology image is generated. Once the topology is learned, any link failure detected locally or received from a remote station initiates the protection state.

4 Link_Status Triggers

4.1 Physical Layer Triggers

Link failure and restoration are initially detected by the physical layer, according to its capability, and then indicated to the MAC layer.

4.1.1 SONET/SDH Triggers

Link failure on a SONET/SDH link is triggered by Signal Failure (SF) or Signal Degradation (SD) alarms provided by the SONET/SDH physical layer. Clearing of these alarms clears the link failure indication.

4.1.2 Ethernet Triggers

Link failure on an Ethernet link is triggered by Signal Failure (SF) or Signal Degradation (SD) alarms provided by the Ethernet physical layer, or inferred from statistics from the Ethernet physical layer. For Ethernet physical layers that do not provide SF or SD alarms, these shall be calculated from threshold values of the coding violations. Clearing of these alarms clears the link failure indication.

4.2 Manual Triggers

For the purpose of maintenance, an RPR station or RPR link can be taken off-line. The operator shall be able to signal the off-line state via a forced Status_Change message indicating a down station or a down link. The operator command will be translated into a protection message that simulates a link or station failure.

4.3 Keep_Alive Triggers

The Keep_Alive message provides a means of detecting a link failure or station failure. Any time a configurable number of Keep_Alive messages in a row are missed (as indicated by receiving 0 Keep_Alive messages within the configured number of expirations of the Keep_Alive_Timer), the link upon which they were to have been received is indicated to be down, and a Link_Status message is broadcast. The indicated failure is DISCONNECTED_SF. The default value of the number of messages that must be missed is 3. 2 Keep_Alives must be received in succession (within 3 expirations of the Keep_Alive_Timer) to clear the link failure indication.

4.4 Protection Hierarchy

The protection switching events hierarchy is used to handle multiple, concurrent events. The hierarchy is shown below from the highest priority to the lowest one:

FS, Force Switch - operator originated (highest priority)

SF, Signal Fail (e.g LOS, LOF, EXBER, LOK (Loss Of Keep_Alive)) - automatically originated

SD, Signal Degrade - automatically originated

MS, Manual Switch - operator originated

WTR, Wait Time To Restore - automatically originated

NR, No Request present (lowest priority)

4.5 Hold-Off Time

For RPR deployments over existing a physical medium providing its own protection (such as SONET/SDH Automatic Protection Switching or optical protection), the underlying protection for the RPR bandwidth cannot always be turned off. In such cases, the underlying protection already in place will be used instead of the RPR protection.

In order not to use RPR protection when an underlying protection should be used as first resort, a configurable Hold-Off Time (HOT) is provided for the RPR protection. By setting the HOT to a value greater than zero, the underlying protection switching is given a chance to work. RPR protection will take place only if the link is not declared up before the HOT expires.

When the RPR protection is the “lowest layer” protection scheme, the HOT is simply set to zero.

4.6 Wait To Restore

To avoid frequent changes in reported link status over unstable links, a configurable Wait To Restore (WTR) timer is provided. By setting the WTR to a value greater than zero, the failed link is required to show no failure for the given period of time before being declared working. RPR protection will be ceased only if the link is not declared down before the WTR expires.

4.7 Multiple Link Failure

Multiple link failures between two neighbor stations are treated in the same manner as a single link failure. Each link failure triggers a separate protection for the ringlet that experienced the failure. For dual-ringlet rings, each of the two ringlets protects the other ringlet. For N-ringlet rings, each ringlet is protected by 1, or optionally more, ringlets with a counter rotation. Each flow, conversation, aggregate, or whatever means is used to segregate and route traffic (henceforth called the **routable unit**) is protected by a single ringlet. Protection does not need to be one to one if 100% protection is not needed.

5 Algorithm Details

5.1 Normal Operation

During normal operation, each station listens for Link_Status messages and examines them for any change in link status. Any link status change from CONNECTED to any of the DISCONNECTED states triggers protection for the ring segment on which it is reported.

5.2 Unicast Protection

During protection, for each routable unit, if a sourcing station can reach the intended destination through the normal route (either through the assigned ringlet or locally), then it uses the normal route. Otherwise, if the routable unit is protected and it can reach the intended destination through an alternate route (through the protected ringlet), then it uses the protection route.

5.3 Multicast Protection

If a sourcing station can reach the all intended destinations through the normal route (either through the assigned ringlet or locally), then it uses the normal route. Otherwise, if the routable unit is protected and it can reach all intended destinations through the alternate route (through the protected ringlet), then it uses the protection route. Otherwise, it uses both routes (both directions). In all of these cases, the TTL would be set to the distance to the failed link on each ringlet. A station can make a local optimization to set the TTL to the shorter of distance of the farthest station for the routable unit or the distance to the failed link, for each ringlet.

5.4 Recovery From Protection

During protection operation, each station listens for Link_Status messages and examines them for any change in link status. Any change of link status from a DISCONNECTED state to the CONNECTED state triggers healing for the ring segment on which it is reported. At this point the system returns to its state prior to the protection event.