1. Protection

Editors' Notes: To be removed prior to final publication.

TOM - There is a numbering problem for Fig. 11.7. I have manually forced it to be correctly numbered.
There is a numbering problem for Table 11.6. I have manually forced it to be correctly numbered.
I have two cross-references needed in 11.4 to 7.2.3.

References:
To be added.

Definitions:
To be added.

Abbreviations:
To be added.

Revision History:
Draft 0.1, February 2002Initial draft document for RPR WG review.
Draft 0.2, April 2002Addressed most comments from draft 0.1. Those not yet addressed are specifically mentioned in editor's notes.
Added Scope and moved normative portions of introduction to later in clause.
Merged initial descriptive clauses into Overview.
Added Steering and Wrapping clauses with detailed state diagrams for each. Copied requirements relevant to each from Draft 0.1 into each of the clauses.
Draft 0.3, June 2002Addressed most comments from draft 0.2. Those not yet addressed are specifically mentioned in editor's notes.
Added general protection overview.
Merged common steering and wrapping requirements into a common protection requirements clause.
Added extra definition to the protection hierarchy clause.
Added sequence number to the protection packet format.
Included new detailed state diagrams in tabular form for steering and wrapping clauses.
Draft 1.0, August 2002Addressed most comments from draft 0.3. Those not yet addressed are specifically mentioned in editor's notes. The protection ad-hoc (PAH) is meeting to address the list of issues described in the editorial note immediately following this one.
Major modifications made in D1.0 include:
Clean-up of figures
Clean-up of protection message section
Addition of new wrapping state diagram

Editors' Notes: To be removed prior to final publication.

The protection ad-hoc (PAH) is meeting to address the following list of issues in a re-written protection clause to be made available to the working group prior to the September meeting.
1. Determine the functional partitioning between topology and protection.
Status: The PAH is continuing to work on this. An annex is in progress to define representative scenarios illustrating the interworking of topology and protection.

2. Comment #565: Determine if and how non-revertive operation will be supported, and what the configuration options are for reversion (per interface and/or per station per class of service). If non-revertive operation is supported for wrapping on an interface of a station, it will apply to locally initiated wraps only. If a station on one end of the span wraps/unwraps, the station on the other end of the span will passively wrap/unwrap independent of whether it is configured as revertive or non-revertive on that interface.
Status: The PAH has not addressed this issue yet.

3. Comment #567: Include description of what is covered by the scenario descriptions to be added later in this clause to resolve #581. Description will also be included of how protection switching interacts with topology discovery and ringlet selection.
Status: The PAH is continuing to work on this, via the annex containing scenarios.

4. Comment #576: Specify the behavior of a ring configured as a wrapping ring that has a steering-only station inserted into it.
Status: The PAH is continuing to work on this. The annex containing scenarios describes suggested behavior for this case.

5. Comments #578 and #579: Determine the requirements of retransmission of protection messages. The options are exponential back-off and constant periodic transmission.
Status: Completed. The PAH determined that a bi-level periodic timer will be used.

6. Comment #581: Reorganize subclauses 1.3, 1.6.1, and 1.7.1 to illustrate key scenarios that illustrate the relevance of the underlying rules. These scenarios will include (for example) protection commanding, protection triggering and hold-off time, protection clearing and wait to restore time, protection hierarchy scenarios, and revertive/non-revertive behavior. The rules and the need for them will be described within the scenario descriptions. In addition, the scenario descriptions will indicate the relationship between protection switching, topology discovery, and ringlet selection.
Status: The PAH is continuing to work on this, via the annex containing scenarios.

7. Comment #582: Examine the issue of whether the protection protocol can by itself determine whether a newly added station is the same as the station that was previously connected on the interface, or whether this is fundamentally a function that should be provided by topology discovery.
Status: Completed. This is addressed by the former MAC address entries in the topology database.

8. Comment #585: Determine how to tie configured absolute time for loss of keepalives to variability in the interval for duration of fairness messages.
Status: The PAH has not addressed this issue yet.

9. Comment #589: Determine if the wrapping status code bit should be kept for informational and diagnostic purposes, and if there is direct need for this bit in the protection protocol itself.
Status: Completed. The wrapping status code bit will be kept for informational and diagnostic purposes.

10. Comment #594: Determine the exact behavior of revertive and non-revertive modes for steering protection.
Status: The PAH has not addressed this issue yet.

11. Comment #595: Determine whether a station in a steering ring reports its local link protection state to the rest of the ring irrespective of protection requests present on the rest of the ring.
Status: The PAH has not addressed this issue yet.

12. Comment #596: Determine whether rule 5 from 1.7.1.2 should apply to steering or not.
Status: The PAH has not addressed this issue yet.

Editors' Notes: To be removed prior to final publication.

13. Comment #601: Determine the desired behavior and provide a scenario description for multiple SD conditions on a ring.
Status: Completed. The desired behavior is that stations detecting a SD condition will wrap immediately and will report SD immediately to the rest of the ring (if there is not already a wrap due to SD or a higher priority defect elsewhere on the ring). In the event that there is a simultaneous SD detection elsewhere on the ring, then there will be wraps on both degraded spans for a short interval, followed by unwrapping on both spans. The scenario description is in progress and will be included in the annex containing scenarios.

1.1 Scope

This clause describes the RPR protection protocol, which implements resiliency, an important RPR objective. The RPR protection protocol provides reliable mechanisms for sub-50 ms protection switching for all protected traffic on an RPR ring. It consists of a mandatory protection mechanism called steering, and an optional protection mechanism called wrapping. The protection protocol is completely defined for both mechanisms, and either by itself is sufficient to meet RPR requirements. The protocol resides in the MAC control sublayer, as shown in the shaded region of Figure 1.1.

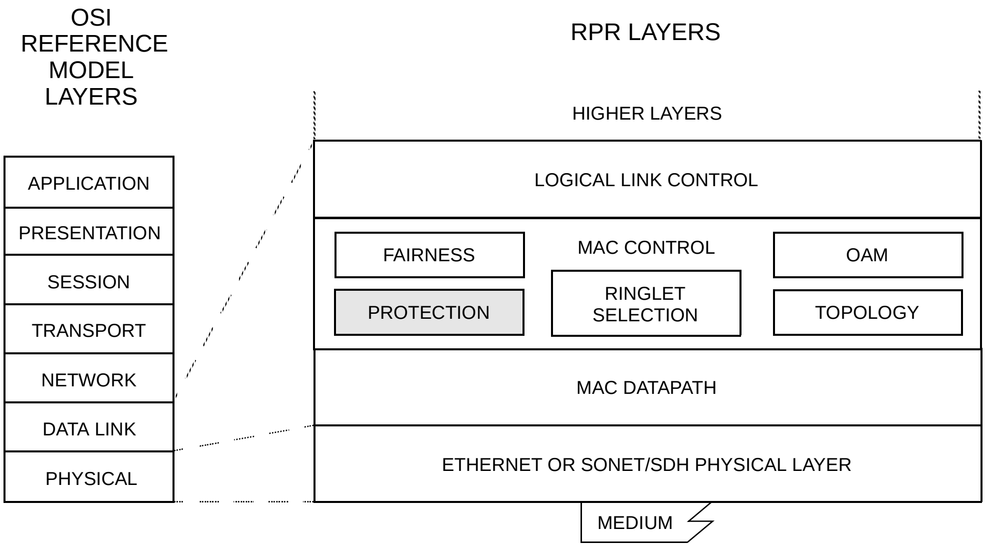

Figure 1.1—RPR layer diagram

The services and features provided are:

a) Sub-50 ms protection switching for unicast and multicast traffic based on media "hard" or "soft" failures or based on operator invoked command
b) Quick dissemination on the ring of information indicating media failures or operator invoked commands impacting use of the media
c) Support of a standard protection hierarchy

d) Tolerant of message loss but with minimal overhead
e) Greater bandwidth efficiency for unprotected services than path switching or bidirectional line switching
f) Operate independently of and in the absence of any management systems
g) Support dynamic addition and removal of stations to/from the ring
h) Scalable from one to hundreds of stations
i) Support of revertive and non-revertive protection switching operational modes
j) Operate without any master station on the ring

1.2 Overview

Editors' Notes: To be removed prior to final publication.

As per comment #567 from March, the PAH will include description of what is covered by the scenario descriptions to be added later in this clause to resolve #581. Description will also be included of how protection switching interacts with topology discovery and ringlet selection.

1.2.1 General protection overview

A protection switching protocol that provides the services and features described in must have the following building blocks explicitly defined:

a) A hierarchy of protection requests that requires local action by a station and reporting of the protection state to the rest of the ring. 1.4 includes description of the protection hierarchy, and of the definition of triggers resulting in specific protection requests.
b) The local actions that can be taken by stations on the ring to prevent traffic loss. As described in the rest of 1.2, these consist of steering or wrapping.
c) A mechanism and control packet format for reporting local protection state to the rest of the ring is given in 1.5.
d) State diagrams and related requirements are given in 1.6 and 1.7. These state diagrams define in detail what protection state information is reported to the rest of the ring, exactly when it is reported, and exactly when action is taken at a station to protect traffic.

The protection protocol does not directly switch traffic on the data plane. Rather, it supplies essential information to entities within the MAC that do, such as ringlet selection, or directly controls when wrapping occurs. It ensures that parameters such as hold-off time and wait to restore time specified by the network operator translate into correct station behavior. It ensures the robust interaction of operator-initiated protection requests and protection requests triggered by physical layer and MAC layer monitoring.

Editors' Notes: To be removed prior to final publication.

There are aspects of the protection protocol that exhibit similarities with other protocols within this draft, such as topology discovery or fairness. The similarity is in the need for control messaging that can go quickly around the ring, and that is robust in the event of message loss. There is a great deal of ongoing discussion about whether separate protection messaging is required, or whether it can be combined with messaging used for topology discovery or for fairness.
The PAH will determine the functional partitioning between topology and protection and provide re-written clauses for the September meeting.

1.2.2 Steering protection

For steering protection, a station shall not wrap a failed segment when a failure is detected. Instead, a protection request message is sent to every station to indicate a link failure, as in the wrap protection scheme. When stations receive a protection request message indicating a failure, their steering database will be updated accordingly. It is the responsibility of each source station to direct its traffic onto ringlet 0 or ringlet 1, whichever avoids the failed link.

Packets destined to a station beyond the point of failure that have been transmitted onto the ring before the steering database is updated at the source station will be dropped at the failure point, since there is no delivery mechanism available.

.

Editors' Notes: To be removed prior to final publication.

The description of setting of TTL for multicast packets in the event of protection is covered in 6.2.4.8 (ringlet selection). As a result, the below description formerly contained in 11.2.4 will be deleted from this editorial note in the next version of the draft:
For steering rings, when a link fails, it is simplest for multicast packets to be sent in both directions, with the TTL set to the number of stations on the ring between the source station and the defective span on each direction. Duplicate packets must not be allowed to arrive at any station on the ring as a result of a protection condition ceasing to exist. For wrapping rings, multicast packets continue to be sent in a single direction, with the TTL set to the number of stations on the ring**.**

Editors' Notes: To be removed prior to final publication.

The steering protection description above will be expanded to correspond to the wrapping description as stipulated in comments 89 and 613 from May.

1.2.3 Wrap protection

An RPR wrap capable ring is composed of two counter-rotating ringlets. If an equipment or facility failure is detected, traffic going towards and from the failure is wrapped (looped) back to go in the opposite direction on the other ring (subject to the protection hierarchy). Wrapping takes place on the stations adjacent to the failure, under control of the protection switch protocol. The wrap re-routes the traffic away from the failure.

.

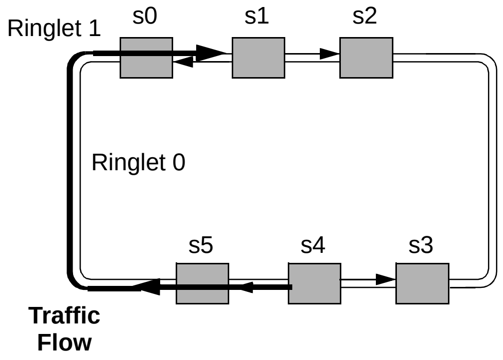

Figure 1.2—Data flow before fiber cut.

An example of the data paths taken before and after a wrap is shown in Figure 1.2 and Figure 1.3, respectively. Before the fiber cut, Station 4 sends to Station 1 via the path S4->S5->S0->S1.

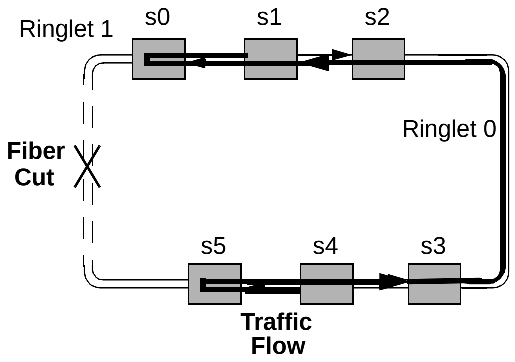

Figure 1.3—Data path after wrap

If there is a fiber cut between Station 5 and Station 0, Station 5 will wrap traffic on ringlet 1 to ringlet 0, and Station 0 will wrap traffic on ringlet 0 to ringlet 1. After the wraps have been set up, traffic from Station 4 to Station 1 initially goes through the non-optimal path S4->S5->S4->S3->S2->S1->S0->S1. S1 does not strip packets from the protection ringlet (ringlet 0 in this case) to avoid packet mis-ordering during the protection event.

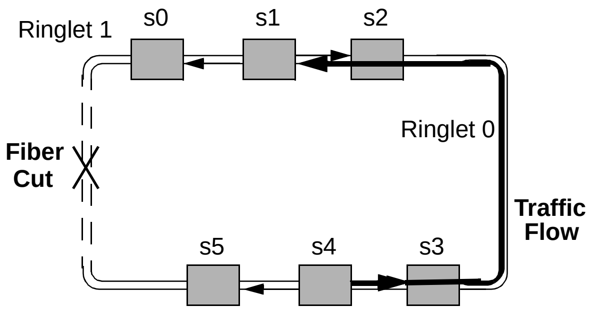

Figure 1.4—Data path after new topology discovery

Subsequently a new ring topology is discovered and a new optimal path S4->S3->S2->S1 is used, as shown in Figure 1.4. Note that the topology discovery and the subsequent optimal path selection are not part of the protection switch protocol.

Editors' Notes: To be removed prior to final publication.

As per comment 357 from March, references will be fixed later based on where references appear (in Clause 2 or Appendix A).

The ring wrap is controlled through SONET BLSR [3][4] style protection switch signaling. It is an objective to perform the wrapping at least as fast as in SONET equipment.

1.3 Common protection rules

Editors' Notes: To be removed prior to final publication.

As per comment #581from March, the PAH will reorganize subclauses 1.3, 1.6.1, and 1.7.1 to illustrate key scenarios that illustrate the relevance of the underlying rules. These scenarios will include (for example) protection commanding, protection triggering and hold-off time, protection clearing and wait to restore time, protection hierarchy scenarios, and revertive/non-revertive behavior. The rules and the need for them will be described within the scenario descriptions. In addition, the scenario descriptions will indicate the relationship between protection switching, topology discovery, and ringlet selection.
The PAH is targeting to include this in the modified protection clause to be distributed to the working group prior to the September meeting.

Editors' Notes: To be removed prior to final publication.

The rules below need to be reformatted to conform with the IEEE editorial template and style.

Fundamental protection rules in RPR include:

1. An RPR ring shall provide protection within 50ms of detection of a link or station failure.

2. All RPR stations shall provide support for steering, with support for wrapping optional.

3. All stations within the same RPR ring shall choose the same protection mechanism. Via the topology discovery protocol, every RPR station shall indicate if it supports wrapping protection or not. If all stations on an RPR ring are able to support wrapping protection, the choice of wrapping or steering for the protection mechanism shall be based on configuration by the system operator; accordingly no wrapping. protection can be initiated until the full topology is learned. Otherwise, steering shall be selected as the default protection scheme on the RPR ring.

Editors' Notes: To be removed prior to final publication.

As per comment #576 from July, the PAH will specify the behavior of a ring configured as a wrapping ring that has a steering-only station inserted into it.

4. Revertive operation shall be the default mode of operation in RPR rings.

Editors' Notes: To be removed prior to final publication.

As per comment #565 from July, the PAH will determine if and how non-revertive operation will be supported, and what the configuration options are for reversion (per interface and/or per station per class of service). If non-revertive operation is supported for wrapping on an interface of a station, it will apply to locally initiated wraps only. If a station on one end of the span wraps/unwraps, the station on the other end of the span will passively wrap/unwrap independent of whether it is configured as revertive or non-revertive on that interface.

1.3.1 RPR protection packet transfer mechanism

1. Protection packets are transferred in a broadcast packet format between stations on the ring. A received packet is passed to the receiving station's MAC control sublayer.

2. All protection messages are triggered by self-detect or user request. Protection messages are sent when the local protection state changes, and are repeated based on a bi-level periodic timer. The fast rate timer is used for the first 8 messages sent after triggering. The period of the fast timer ranges between 1 ms and 20 ms with 1 ms resolution and with a default value of 5 ms. The slow timer is used for all messages following the first 8 messages after triggering. The period of the slow timer ranges between 50 ms and 1000 ms with 50 ms resolution and with default value of 100 ms.

3. Protection messages shall continue to be delivered on links that are in non-idle protection states.

1.3.2 RPR protection signaling mechanism

1. Protection switch signaling is performed using protection control packets as defined in Figure 1.5.

2. A station executing a local request signals the protection request on both short (opposite ringlet from the link for which the state is being reported) and long (same ringlet as the link for which the state is being reported) paths.

1.3.3 RPR protection protocol rules:

1. The protection request hierarchy values are listed in Table 1.3 (listed from highest priority to lowest priority). In general, a higher priority request preempts a lower priority request within the ring, with exceptions noted as rules. The 4 bit values shown in the table correspond to the protection message request type (PMRT) field in the protection packet.

Editors' Notes: To be removed prior to final publication.

Requirements must be added to the protection clause to cover its interface with the topology and status database defined in Clause 10.

2. When a station which initially detected a failure discovers the disappearance of the failure, it enters WTR (for a user-configured WTR time-period). The configurable range for WTR is defined in 1.4.

3. When a station is in WTR mode, and detects (via matching the MAC address of the new neighbor to the former neighbor MAC address on that interface stored in the topology database) that a new neighbor is not the same as the old neighbor (stored while the old neighbor was still recognized as part of the topology), the station drops the WTR.

The information from the topology database is needed so that protection knows when to apply WTR without requiring protection to separately store topology information.

Editors' Notes: To be removed prior to final publication.

Rule 4 in this subclause no longer seems to be necessary based on the rewording of rule 3. As per comment #582 from July, it will be deleted in the next version of the draft unless justification is provided to keep it prior to then. The PAH will examine the issue of whether the protection protocol can by itself determine whether a newly added station is the same as the station that was previously connected on the interface, or whether this is fundamentally a function that should be provided by topology discovery.

4. When a station is in WTR mode, and the source of a long path request is not equal to its neighbor on the opposite side (as stored at the time of protection switch initiation), the station drops the WTR. This is the case when a new neighbor is added to the ring.

5. When a station receives a local protection request of type SD or SF and it cannot be executed (according to protocol rules), it keeps the request pending. (The request can be kept pending outside of the protection protocol implementation.)

6. If a local non-failure request (WTR, MS, FS) clears, and if there are no other requests pending, the station enters idle state.

7. If there are two link failures and two resulting WTR conditions on a single segment, the second WTR to time out brings both the links up. (After a WTR time expires, a station does not unprotect automatically, but waits until it receives idle messages from its neighbor on the previously failed segment.)

8. A configurable hold-off time must be supported with a range of the hold-off time is zero to at least 200 ms with 10 ms resolution. The default value is zero.

Editors' Notes: To be removed prior to final publication.

A subclause is needed here to put this clause in context with respect to the MAC reference model. This subclause will cover description of how protection switching relates to other components in the MAC reference model, the service interface to the client (to address comments 93 and 351 from March), and interactions between components of the MAC control sublayer. This will address comment 91 from May.
The description of how the components of the MAC control sublayer interact and make information available to each other may tie to a common topology/protection database. The determination of whether such a database is sufficient for this purpose depends on the definition of this database, which is tied to topology discovery as per comment 611 from March.
The MAC data path module may need notification from the protection module when a wrapping-based ring has returned to a fully normal operating state, as per comment 707 from May. This notification is required to enable the capability to immediately strip packets from the protection ringlet upon restoration of the complete bi-directional ring without requiring TTL to decrement fully.
Protection may provide ringlet selection a preferred ringlet for each destination station, or may provide ringlet selection the actual protection state of each link in the ring via the topology and status database. A default mechanism for selecting a ringlet during a protection event should consider the highest protection state applying to links on each direction between each source and destination and select the direction with the lower protection state in the hierarchy.

1.4 Protection hierarchy and triggers

The protection switch protocol processes the following request types (in the order of priority, from highest to lowest). The triggers for these requests are defined below. All requests are signaled using protection control messages defined in 1.5.

1) Forced Switch (FS): Operator originated. Results in a protection switch away from a requested link (steering or wrapping all traffic away from the link). An example use of this request type is to add another station to the ring in a controlled fashion.
2) Signal Fail (SF): Automatic. Caused by a media Signal Failure or RPR keep-alive failure. A media SF shall be detected based on the PHY_LINK_STATUS.indicate defined in 7.2.3. Results in a protection switch away from the impacted link (steering or wrapping all traffic away from the link). SONET examples of SF triggers are: Loss of Signal (LOS), Loss of Frame (LOF), Line Bit Error Rate (BER) above a preselected SF threshold, and Line Alarm Indication Signal (AIS). Explicit definition of the SF triggers and SF clearing criteria for SONET are provided in the Telcordia GR-253-CORE and ANSI T1.105-01 standards, among others. An RPR keep-alive failure is defined later in this clause. The configurable hold-off time, if non-zero, starts after the physical SF conidtion is detected and shall elapse before SF triggers. As described in Clause 10, incorrect connection of interfaces of neighboring stations through miscabling shall also result in a MAC layer SF condition.
3) Signal Degrade (SD): Automatic. Caused by a media Signal Degrade (e.g. excessive Bit Error Rate). A media SD shall be detected based on the PHY_LINK_STATUS.indicate defined in 7.2.3. Results in a protection switch away from the impacted link (steering or wrapping all traffic away from the link). SONET examples of SD triggers are: Line BER or Path BER above a preselected SD threshold. Explicit definition of the SD triggers and SD clearing criteria for SONET are provided in the Telcordia GR-253-CORE and ANSI T1.105-01 standards, among others. The configurable hold-off time, if non-zero, starts after the physical SD condition is detected and shall elapse before SD triggers.
4) Manual Switch (MS): Operator originated. Like Forced Switched but of a lower priority. An example use of this request type is to force down a marginally operating link, but allow it to come back into use in the case of a more serious failure elsewhere on the ring.
5) Wait to Restore (WTR): Automatic. Entered after a link meets the restoration criteria following exit from a SF or SD condition. The protection switch protocol waits for the WTR time-out

before restoring traffic. An example use of this request type is to prevent protection switch oscillations. The configurable range for WTR is defined later in this clause.

The protection module finds out about operator originated protection requests and clearing from the Layer Management Entity defined in Clause 13. It finds out about automatic protection requests via link status primitives defined in Annex C and Annex D for the Ethernet and SONET/SDH reconciliation sublayers, respectively. The protection module may also utilize automatic protection triggers mapped into SF or SD not explicitly defined in Annex C and Annex D.

The purpose of RPR keep-alives is to provide an indication that a station has an acceptable degree of operational capability. The transmission of RPR keep-alives ideally should occur only if all implemented checks of normal MAC operation are fulfilled.

An RPR keep-alive failure on a receive interface is defined by the failure to receive a configurable number of consecutive type A fairness messages (as defined in Clause 9) . The recommended value should correspond to the number of type A fairness messages transmitted in 3 milliseconds. For example, a fairness message interval of 100 microseconds translates to 30. The configurable range is from 16 to 512 with resolution of 1.

Editors’ Notes: To be removed prior to final publication.

The PAH will ensure that the protection clause definition works if there is no signal degrade indication available, e.g. for PHYs that do not support signal degrade.

Editors’ Notes: To be removed prior to final publication.

As per comment #585 from July, the PAH will determine how to tie configured absolute time for loss of keepalives to variability in the interval for duration of fairness messages. Therefore, please do not comment on the above paragraph..

A SF condition due to loss of keep-alives is not cleared solely by the re-start of reception of keep-alives from a newly connected or re-connected neighbor station. Upon receipt of a single keep-alive, the link transitions to WTR state. For a new neighbor, if a topology status message (defined in Clause 10) is received from the neighbor before the WTR expires, the link is brought up upon receipt of the topology status message. If the WTR expires prior to receipt of a topology status message, then the link is brought up based on WTR regardless of whether the neighbor is a new neighbor or unchanged from before the SF condition.

WTR shall be configured with values in the range of 0-1440 sec. with resolution of 1 second and with a default value of 10 seconds.

> **Editors' Notes:** To be removed prior to final publication.
>
> The below description of the use of protection hierarchy for steering rings may be modified based on the outcome of ongoing discussions on ringlet selection. Ringlet selection may be provided a preferred ringlet per destination station from the protection module. In this case the protection module needs to utilize the protection hierarchy to determine this. Ringlet selection may also use the complete map of the topology with protection state per link, as provided in the topology and status database. In this case ringlet selection may utilize the protection hierarchy as one part of its algorithm.
>
> Since all protection switches are performed bidirectionally, both links on a span are considered to have a protection request status equal to the highest protection request of either link on the span. For example, if one link of a span has protection request SF and the other has no request (Idle), then both links are considered to be SF from the perspective of steering protection.

For wrapping rings, as protection requests travel around the ring, the protection hierarchy is applied. If a requested protection switch is of a higher priority (e.g. a Signal Fail request is of higher priority than a Signal Degrade request), then the requested protection switch takes place and any lower priority protection switches elsewhere in the ring are superseded and ignored. When a lower priority request is presented, it is not allowed if a higher priority request is present in the ring. The only exception is multiple SF and FS switches, which can coexist in the ring.

For steering rings, protection requests from any station travel around the ring and are stored in the topology and status database at all stations on the ring. To determine which ringlet is preferred for the transmission of packets from a given source to a given destination, the highest protection request on each path connecting the source and destination must be compared. For example, if there are links with protection requests SD and MS on the portion of ringlet 0 connecting station A to station B, and there is a link with protection request SF on the portion of ringlet 1 connecting station A to station B, then ringlet 0 would be selected for steering traffic from station A to station B.

All protection switches are performed bidirectionally (protect at both ends of a segment for both transmit and receive directions, even if a failure is only unidirectional).

1.5 Protection message packet format

The protection switch message is used for signaling various link failures and degradations. The protection switch message packet format is outlined in Figure 1.5.

	RPR CONTROL HEADER
1 BYTE	PROTECTION MESSAGE BYTE
1 BYTE	SEQUENCE NUMBER
4 BYTES	FCS

Figure 1.5—Protection switch packet format

The protection switch specific fields are detailed below.

1.5.1 Destination MAC Address within RPR header

The Destination MAC Address is the broadcast address defined in Clause 6. Protection switch messages are broadcast in order to minimize the transmission delay.

1.5.2 Source MAC Address within RPR header

The Source MAC Address is the MAC address of the originator of the protection message.

1.5.3 Protection message byte

The protection message byte is shown in Figure 1.6. The subfields of the protection message byte are described in the subclauses below.

	MSB				LSB
	7	6	5	4	3:0
Byte 1	RSVD	WSC	RSVD	PI	PMRT

Figure 1.6—Protection message byte format

1.5.3.1 Reserved (RSVD)

Bits 7 and 5 are reserved for future use.

1.5.3.2 Wrapping status code (WSC)

The Status Code is used on rings utilizing wrapping protection to indicate whether a wrap is present on an any interface to a station. More precisely, it indicates whether a station is in Wrapping state in 1.7.

The values of WSC are defined in Table 1.1.

Table 1.1—Wrapping status code values

Value	Description
0	Idle
1	Protection switch completed - traffic wrapped (W)

1.5.3.3 Path indicator (PI)

The Path Indicator indicates on which path a protection message is sent, long and short. Short path messages are sent towards the failed link through the opposite ringlet. They indicate a failure on the other ringlet on the link immediately preceding the station whose address is given in the source address of the protection request message. For example, a protection message with a short path indicator received on ringlet 0 from a station indicates that the protection state is for the ingress link to that station on ringlet 1. Long path messages are sent away from the failed link on the same ringlet. They indicate a failure on the same ringlet on the link immediately preceding the station whose address is given in the source address of the protection request messages. For example, a protection message with a long path indicator received on ringlet 0 from a station indicates that the protection state is for the ingress link to that station on ringlet 0.

The values of WSC are defined in Table 1.2.

Table 1.2—Path indicator values

Value	Description
0	Short (S)
1	Long (L)

1.5.3.4 Protection message request type (PMRT)

A definition of the protection message request types is given in 1.4. The values of PMRT are defined in Table 1.3

Table 1.3—Protection message request type values

Value (bin)	Description
1011	Forced switch (FS)
1101	Signal fail (SF)
0001	Signal degrade (SD)
0110	Manual switch (MS)
1010	Wait to restore (WTR)
0000	No request (IDLE)

The protection control messages are shown in this document as:

{requestType, sourceAddress, wrapStatus, pathIndicator}

1.5.4 Sequence number

The sequence number is an 8-bit field that is incremented upon the bidirectional broadcast of a protection control message, e.g. the same sequence number is used for the two copies of a given protection control message, one sent on ringlet 0 and the other sent on ringlet 1. This is needed to prevent processing of out of order protection control messages. This can occur due to messages being received on the short path and on the long path in a transient scenario.

Editors’ Notes: To be removed prior to final publication.

The text of the below subclauses as written are generally descriptive rather than normative. The appropriate portions of the text need to be made normative.

1.6 Steering

1.6.1 Listing of rules

Editors’ Notes: To be removed prior to final publication.

The rules below need to be reformatted to conform with the IEEE editorial template and style.

Editors’ Notes: To be removed prior to final publication.

Comments are solicited on whether the below requirements are worded correctly for steering, and on whether additional requirements should be included. It is important for the wording of these requirements to be consistent with the state diagram.

1.6.1.1 RPR protection protocol rules

Rules in this subclause are applicable to steering protection only.

1. A protection control packet shall be accepted only if the sequence number contained within the control header as defined in 1.5 meets the conditions defined in the state diagram in 1.6.3.

This rule is needed to prevent processing of out of order protection control messages. This can occur due to messages being received on the short path and on the long path in a transient scenario.

1.6.2 Parameters

a) Revertive (revertive switch operation mode)
Indicates that a station is in revertive mode for a given local ingress link. Revertive mode means that a station enters WTR state upon clearing of an SF or SD condition for a given local ingress link to a station.

b) Non-revertive (non-revertive switch operation mode)
Indicates that a station is in non-revertive mode for a given local ingress link. Non-revertive mode means that a station remains in SF or SD state even upon clearing of an SF or SD condition at the underlying physical layer for a given local ingress link to a station.

c) FS
Forced switch is in effect for a given local ingress link to a station.

d) SF
Signal fail is in effect for a given local ingress link to a station.

e) SD
Signal degrade is in effect for a given local ingress link to a station.

f) MS
Manual switch is in effect for a given local ingress link to a station.

g) IDLE
Idle (no protection request).

h) SELF or Station
Source MAC address for station sending protection message.

i) I
Not wrapped (unused indication for steering)

j) S
Short path indication, e.g. protection message sent on other ringlet from that on which the link with the protection request lies.

k) L
Long path indication, e.g. protection message sent on same ringlet on which the link with the protection request lies.

l) Neighbor
MAC address of station on the other end of the affected span.

m) NewNbr
MAC address of station on the other end of the affected span after connectivity is regained on the span.

n) BEGIN
A Boolean variable that is set to TRUE when the System is initialized or reinitialized, and is set to FALSE when (re-)initialization has completed.
Value: Boolean

o) storedProtReqState
The stored protection request state (if the active request is higher in the protection hierarchy) of an ingress link of given station.

1.6.3 State diagram

Editors' Notes: To be removed prior to final publication.

In this version of the draft a complete tabular description of the state diagram is provided, with textual description. It has been difficult to formulate a simple flow chart that does not convey a lot of inaccuracies when compared to the complete state diagram (as requested in comment 312 from May). Contributions are invited.

Editors' Notes: To be removed prior to final publication.

As per comment #594 from July, PAH will determine the exact behavior of revertive and non-revertive modes for steering protection.
As per comment #595 from July, PAH will determine whether a station in a steering ring reports its local link protection state to the rest of the ring irrespective of protection requests present on the rest of the ring.
As per comment #596 from July, PAH will determine whether rule 5 from 1.7.1.2 should apply to steering or not.
Therefore, please do not comment on the above listed issues.

The states in the steering state diagrams shown in Table 1.4 and Table 1.5 are described in this subclause.

The state diagram in Table 1.4 shows the state transitions for determination of the protection state to be reported to the rest of the ring (and used internally to the station to update the topology and status database defined in Clause 10) for each incoming or ingress link to a given station. The state transitions reflect the state that is reported for the link to the rest of the ring by the station. As there are two ingress links to the station, two instances of this state diagram are being executed independently and in parallel, one for each ingress link. The state diagram takes into account wait to restore time and revertive/non-revertive operation for a given link. The state diagram also takes into account a protection event register per link for existing protection requests underlying the protection request reported to the rest of the ring (that is highest in the protection hierarchy).

The state diagram in Table 1.5 shows the state transitions for the actual sending of protection control messages on the ring and receipt of protection control messages from other stations on the ring. It also shows the state transitions for updates to the topology and status database defined in Clause 10 and when new preferred directions are computed for steering. The transmit and receive processing of protection control messages, database updates, and computation of preferred directions for steering are running in parallel with the determination of protection state for each link described above.

A simple example can help illustrate the operation of the state diagrams. Imagine the following series of events at a station:

1) SF is detected on ingress link on ringlet 0.
2) A FS request is received on the span on which the SF was detected.
3) Later in the day, the FS is cleared and a MS is put in place to prevent traffic from using the span until the operator allows it.
4) The SF clears.
5) A SD appears on another link somewhere else in the ring later in the day.

In terms of Table 1.4, the transitions are as follows:

1) When the SF is detected and the built-in hold-off time has elapsed, the link protection state transitions from Idle to SF.
2) When the FS request is received, the link protection state transitions immediately to FS (so now FS is reported to the rest of the ring). SF is now stored in the protection event register.
3) When the FS clears, SF state is immediately re-entered. When MS is set on the link, MS is stored in the protection event register.
4) When the SF clears, MS state is entered (as MS is higher in the hierarchy than WTR).
5) When the SD is reported from elsewhere on the ring, this does not impact the link protection state (still in MS state).

In terms of Table 1.5, the transitions are as follows:

1) When the SF is detected, the new protection state is broadcast on the ring and the database is updated as part of the transition from Wait to Steering. In steering state, a new preferred direction is computed for all destination stations that does not utilize the span with SFon it.
2) When the FS request is received, FS state is entered. The same process as for when the FS request is received is executed (but this time no traffic is rerouted).
3) When the FS clears, SF state is entered. The same process as for when the FS request is received is executed (but this time no traffic is rerouted). When the MS is set, the request is stored in the protection event register and no action takes place in this state diagram..
4) When the SF clears, MS state is entered. Again the same process as above occurs (but no traffic is rerouted)..
5) When the SD is reported, the sequence Wait -> Check_Seq -> Check_DB -> Steering -> Wait occurs. Check_Seq ensures that the received protection control message is in sequential order from previous received control messages from that station. Check_DB determines whether the database has actually changed (which it has). Steering results in the rerouting of traffic since SD is higher in the protection hierarchy than MS. Then the state diagram returns to Wait state.
6) All through the day, periodic transitions are taking place from Wait state back to Wait state for rebroadcasts of local protection state to the rest of the ring.

Table 1.4—Steering: Local link protection state

Current state		Row	Next state	
state	condition		action	state
Idle	Local FS request	1	—	FS
	Local SF request	2	—	SF
	Local SD request	3	—	SD
	Local MS request	4	—	MS
FS	Local FS cleared * (storedProtReqState = NULL)	5	—	Idle
	Local SF request	6	storedProtReqState = SF	FS
	Local SD request	7	storedProtReqState = SD	FS
	Local MS request	8	storedProtReqState = MS	FS
	Local FS cleared * (storedProtReqState = SF)	9	storedProtReqState = NULL	SF
	Local FS cleared * (storedProtReqState = SD)	10	storedProtReqState = NULL	SD
	Local FS cleared * (storedProtReqState = MS)	11	storedProtReqState = NULL	MS
	Stored request cleared	12	storedProtReqState = NULL	FS

Table 1.4—Steering: Local link protection state

Current state		Row	Next state	
state	**condition**		**action**	**state**
SF	Local FS request	13	Set existing request (SF) into protection event register	FS
	Local MS request	14	Set MS request into protection event register	SF
	Local SF cleared * (storedProtReqState = NULL) * revertive on span	15	—	WTR
	Local SF cleared * (storedProtReqState = NULL) * non-revertive on span	16	—	SF
	Local SF cleared * (storedProtReqState = SD)	17	storedProtReqState = NULL	SD
	Local SF cleared * (storedProtReqState = MS)	18	storedProtReqState = NULL	MS
	Stored request cleared	19	storedProtReqState = NULL	SF
SD	Local FS request	20	storedProtReqState = SD	FS
	Local MS request	21	storedProtReqState = MS	SD
	Local SD cleared * (storedProtReqState = NULL) * revertive on span	22	—	WTR
	Local SD cleared * (storedProtReqState = NULL) * non-revertive on span	23	—	SD
	Local SD cleared * MS request present	24	storedProtReqState = NULL	MS
	Local SD cleared * replaced by SF	25	—	SF
	Stored request cleared	26	storedProtReqState = NULL	SD

Table 1.4—Steering: Local link protection state

Current state		Row	Next state	
state	**condition**		**action**	**state**
MS	Local MS cleared	27	—	Idle
	Local FS request	28	storedProtReqState = MS	FS
	Local SF request	29	storedProtReqState = MS	SF
	Local SD request	30	storedProtReqState = MS	SD
	Stored request cleared	31	storedProtReqState = NULL	MS
WTR	WTR timer expired	32	—	Idle
	Local FS request	33	—	FS
	Local SF request	34	—	SF
	Local SD request	35	—	SD
	Local MS request	36	—	MS
	~(NewNbr == Neighbor)	37	—	Idle

Row 1.4-1: Local FS request received. Transition into FS state.

Row 1.4-2: Local SF request received. Transition into SF state.

Row 1.4-3: Local SD request received. Transition into SD state.

Row 1.4-4: Local MS request received. Transition into MS state.

Row 1.4-5: Local FS cleared and no subsequent requests stored in protection event register. Transition into Idle state.

Row 1.4-6: Local SF request received. Set SF request into protection event register. Remain in FS state.

Row 1.4-7: Local SD request received. Set SD request into protection event register. Remain in FS state.

Row 1.4-8: Local MS request received. Set MS request into protection event register. Remain in FS state.

Row 1.4-9: Local FS cleared and SF request stored in protection event register. Transition into SF state.

Row 1.4-10: Local FS cleared and SD request stored in protection event register. Transition into SD state.

Row 1.4-11: Local FS cleared and MS request stored in protection event register. Transition into MS state.

Row 1.4-12: A request stored in the protection event register is cleared. Remove the request from the protection event register and remain in FS state.

Row 1.4-13: Local FS request received. Set existing request (SF) into protection event register. Transition into FS state.

Row 1.4-14: Local MS request received. Set MS request into protection event register. Remain in SF state.

Row 1.4-15: Local SF cleared, no subsequent requests stored in protection event register, and revertive on the span. Transition into WTR state.

Row 1.4-16: Local SF cleared, no subsequent requests stored in protection event register, and non-revertive on the span. Remain in SF state.

Row 1.4-17: Local SF cleared and SD request stored in protection event register. Transition into SD state.

Row 1.4-18: Local SF cleared and MS request stored in protection event register. Transition into MS state.

Row 1.4-19: A request stored in the protection event register is cleared. Remove the request from the protection event register and remain in SF state.

Row 1.4-20: Local FS request received. Set existing request (SD) into protection event register. Transition into FS state.

Row 1.4-21: Local MS request received. Set MS request into protection event register. Remain in SD state.

Row 1.4-22: Local SD cleared, no subsequent requests stored in protection event register, and revertive on the span. Transition into WTR state.

Row 1.4-23: Local SD cleared, no subsequent requests stored in protection event register, and non-revertive on the span. Remain in SD state.

Row 1.4-24: Local SD cleared and MS request stored in protection event register. Transition into MS state.

Row 1.4-25: Local SD cleared and replaced by SF request. Transition into SF state.

Row 1.4-26: A request stored in the protection event register is cleared. Remove the request from the protection event register and remain in SD state.

Row 1.4-27: Local MS request cleared. Transition into Idle state.

Row 1.4-28: Local FS request received. Set existing request (MS) into protection event register. Transition into FS state.

Row 1.4-29: Local SF request received. Set existing request (MS) into protection event register. Transition into SF state.

Row 1.4-30: Local SD request received. Set existing request (MS) into protection event register. Transition into SD state.

Row 1.4-31: A request stored in the protection event register is cleared. Remove the request from the protection event register and remain in MS state.

Row 1.4-32: WTR timer expired. Transition into Idle state.

Row 1.4-33: Local FS request received. Transition into FS state.

Row 1.4-34: Local SF request received. Transition into SF state.

Row 1.4-35: Local SD request received. Transition into SD state.

Row 1.4-36: Local MS request received. Transition into MS state.

Row 1.4-37: The MAC address of the station on the other end of the affected span after connectivity is regained (NewNbr) is not the same as the previously stored MAC address of the station at the other end of the span (Neighbor). The MAC address is determined from the topology discovery message defined in Clause 10. Transition directly to Idle state without waiting for the WTR timer to expire.

Table 1.5—Steering: Messaging and database updates per local link

Current state		Row	Next state	
state	condition		action	state
Wait	Transmit timer expired	1	Tx{local state,SELF,I,S}*Tx{local state,SELF,I,L}	Wait
	Local protection state changed	2	Tx{local state,SELF,I,S}*Tx{local state,SELF,I,L} Update link status entry in topology and status database	Steering
	Protection control packet received	3	—	Check_Seq
Check_Seq	{((seq_num_stored >= 128) * (seq_num_stored <= 255)) * ((seq_num_rcvd > seq_num_stored) + (seq_num_rcvd <= seq_num_stored - 128))}	4	Accept protection control message seq_num_stored = seq_num_rcvd	Check_DB
	{~((seq_num_stored >= 128) * (seq_num_stored <= 255)) * ((seq_num_rcvd > seq_num_stored) * (seq_num_rcvd <= seq_num_stored +128))}	5	Accept protection control message seq_num_stored = seq_num_rcvd	Check_DB
	—	6	Delete protection control message	Wait
Check_DB	Link status in protection control message different from link status stored in topology and status database	7	Update link status entry in topology and status database	Steering
	—	8	—	Wait
Steering	—	9	Compute preferred direction per station for unicast and for multicast, where span status is the highest priority request across both links of the span Send information to ringlet selection	Wait

Row 1.5-1: Exponential back-off timer described in R S.2 expires for an individual local ingress link to the station. Then broadcast local protection state for that ingress link on both ringlets.

Row 1.5-2: Local protection state changes based on state diagram in Table 1.4. Then broadcast local protection state for that ingress link on both ringlets, and update the corresponding link status entry in the topology and status database.

Row 1.5-3: Protection control packet received from another station on the ring.

Row 1.5-4: Accept the received protection control packet if the sequence number in the packet is “greater” than stored sequence number in a circular sense. Since the field in the protection control packet is 8 bits, this means that the received packet is valid if its sequence number is within the 128 sequence numbers following the stored sequence number including rollovers.

Row 1.5-5: Same as Row 4.

Row 1.5-6: Else discard the received protection control packet.

Row 1.5-7: Link status in received protection control packet is different from the link status for that link stored in the topology and status database. If so, update the link status entry in the topology and status database.

Row 1.5-8: Else return to the Wait state

Row 1.5-9: Compute preferred direction per station for unicast and for multicast, where span status is the highest priority request across both links of the span. Then send the resulting information to ringlet selection.

1.7 Wrapping

1.7.1 Listing of rules

Editors’ Notes: To be removed prior to final publication.

The rules below need to be reformatted to conform with the IEEE editorial template and style.

Editors’ Notes: To be removed prior to final publication.

Comments are solicited on whether the below requirements are worded correctly for wrapping, and on whether additional requirements should be included. It is important for the wording of these requirements to be consistent with the state diagram.

1.7.1.1 RPR protection signaling and wrapping mechanism

1. Protection control packets are never wrapped.

2. If the protection protocol calls for sending both short and long path requests on the same segment (for example if a station has all fibers disconnected), only the short path request should be sent.

This rule is needed for wrapping in the event of failures on both ingress links to a station. The station may still be able to transmit to other stations via functioning egress links.

This rule is not needed for steering. If there is a failure of a single ingress link to a station on a steering ring, the station will send a protection message on each ringlet. If there is a failure of both ingress links to a station on a steering ring, the station will send a protection message on each ringlet for each link failure. The messages will contain fields as described in 1.4.

3. A station on a wrapping RPR ring wraps only as a result of a local request or a short path request. A station never wraps as a result of a long path request. On wrapping rings, long path requests are used only to maintain the protection hierarchy. (Since the long path requests do not trigger wrapping protection, there is no need for destination addresses and no need for topology maps). A station may unwrap as a result of a long

path request. For example, if the local request at a station in wrapping state has lower priority than the long path request, the station should unwrap.

1.7.1.2 RPR protection protocol rules

1. Requests >= SF can coexist. All requests above SF need to be cleared before the state is transferred into idle state.

2. Requests < SF can not coexist with other requests. A higher priority request will preempt a lower priority request.

For wrapping, two different spans with SF conditions will result in simultaneous wraps on both spans. The desired behavior for two different spans with SD conditions is that the two spans are equally usable, and therefore that no wraps should result. If two SD conditions occur nearly simultaneously, then there may be a transient condition where there are wraps on both SD spans for a brief interval, but stations on both of these spans will unwrap when they find out that there is more than one SD condition on the ring. There is no corresponding condition for steering, as ringlet selection for traffic entering a ring with two different spans with SD conditions can be handled independently at each source station.

3. A station always honors the highest of {short path request, self detected request} if there is no higher long path message passing through the station.

4. When there is more than one request of priority < SF, the first request to complete long path signaling will take priority. However, a higher priority request can preempt the request as long as its long path signal is completed.

5. If a short path FS request is present on a given segment, and a SF/SD condition takes place on the same segment, a station shall accept and process the SF/SD condition ignoring the FS. (Without this rule, a single ended wrap condition could take place, wrapping on only one end of a segment.)

1.7.2 Parameters

a) FS
Forced switch is in effect for a given local ingress link to a station.

b) SF
Signal fail is in effect for a given local ingress link to a station.

c) SD
Signal degrade is in effect for a given local ingress link to a station.

d) MS
Manual switch is in effect for a given local ingress link to a station

e) WTR
Wait to restore

f) IDLE
Idle (no protection request).

g) REQ
protection request except IDLE.

h) SELF
Source MAC address for station sending protection message.

i) I
Not wrapped on the interface.

j) W
Wrapped on the interface.

k) S
Short path indication, e.g. protection message sent on other ringlet from that on which the link with the protection request lies.

l) L
Long path indication, e.g. protection message sent on same ringlet on which the link with the protection request lies.

m) Neighbor
MAC address of station on the other end of the affected span.

n) NewNbr
MAC address of station on the other end of the affected span after connectivity is regained on the span.

o) BEGIN
A Boolean variable that is set to TRUE when the System is initialized or reinitialized, and is set to FALSE when (re-)initialization has completed.
Value: Boolean

p) sideProtectionState
The protection state of an ingress link of given station.

q) otherSideProtectionState
The protection state of the other side’s ingress link of given station.

r) neighborProtectionState
The protection state for the neighbor of the ingress link of given station.

s) otherStationProtectionState
The highest protection state for other stations on the ring.

t) Other
Other station on the ring.

u) localFailureClear
Self detect failure like SF, SD clears.

v) localRequestClear
Local protection request llike FS, MS clears.

w) localProtectionClear
No local failure or local protection event.

x) sidePendingReq
The pending request for the ingress link.

1.7.3 Functions

a) oneSideWrapAvoid(REQ)
TRUE: (neighborProtectionState==FS)*((REQ==SF)+(REQ==SD))

b) isConxistedProtection(REQ)
TRUE: (REQ >=SF)*(sideProtectionState<=SF)

c) isNonConxistedProtection(REQ)
TRUE: (REQ < SF)*(REQ > sideProtectionState)*(REQ > otherSideProtectionState)*(REQ > neighborProtectionState)*(REQ > otherStationProtectionState)

d) isProtectionGrant(REQ)
TRUE : oneSideWrapAvoid(REQ)+
isConxistedProtection(REQ)+
isNonConxistedProtection(REQ)

e) isProtectionPreemptive(REQ)
TRUE: ((sideProtectionState < SF) * (sideProtectionState <= REQ))

1.7.4 State diagram

The wrapping state diagram in Table shows the state transitions for determination of wrapping protection state for each incoming or ingress link to a given station. As there are two ingress links to the station, two instances of this state diagram are being executed independently and in parallel, one for each ingress link.

Table 1.6—Wrapping state diagram for an ingress link

Current state		**Row**	**Next state**	
state	**condition**		**action**	**state**
IDLE	Local REQ request * isProtectionGrant(REQ)	1	Tx{REQ,SELF,W,S}*Tx{REQ,SELF,W,L} * sideProtectionState = REQ	Wrapping
	Local REQ request * !isProtectionGrant(REQ)	2	sidePendingReq = REQ	IDLE
	Rx{REQ,Neighbor,W,S}	3	Tx{IDLE,SELF,W,S}*Tx{IDLE,SELF,W,L} * neighborProtectionState = REQ	Wrapping

Table 1.6—Wrapping state diagram for an ingress link

Current state		Row	Next state	
state	**condition**		**action**	**state**
Wrapping	((sideProtectionState == SF)+(sideProtectionState == SD))*localFailureClear	4	Tx{WTR,SELF,W,S}*Tx{WTR,SELF,W,L} * Start WTR timer* sideProtectionState = WTR	WTR
	((sideProtectionState == FS)+(sideProtectionState == MS))*localRequestClear	5	Tx{IDLE,SELF,I,S}*Tx{IDLE,SELF,I,L}* sideProtectionState = IDLE	IDLE
	localProtectionClear * (sidePendingReq >= otherSideProtectionState)	6	Tx{sidePendingReq,SELF,W,S}*Tx(sidePendingReq,SELF,W,L}* sideProtectionState = sidePendingReq	Wrapping
	localProtectionClear * (sidePendingReq < otherSideProtectionState)	7		IDLE
	Rx(WTR,Neighbor,W,S)	8	Tx{IDLE,SELF,W,S}*Tx{IDLE,SELF,W,L} * sideProtectionState = WTR	WTR
	Rx(IDLE,Neighbor,I,S)	9	Tx{IDLE,SELF,I,S}*Tx{IDLE,SELF,I,L} * sideProtectionState = IDLE	IDLE
	(Rx(WTR,Neighbor,W,S)+Rx(IDLE,Neighbor,I,S))*(sideProtectionState >= otherSideProtectionState)	10	Tx{sideProtectionState,SELF,W,S}*Tx(sideProtectionState,SELF,W,L}	Wrapping
	(Rx(WTR,Neighbor,W,S)+Rx(IDLE,Neighbor,I,S))*(sideProtectionState >= otherSideProtectionState)	11		IDLE
	Rx(REQ, Other,W,L) * *isProtectionPreemptive(REQ)	12	Tx{IDLE,SELF,I,S}*Tx{IDLE,SELF,I,L} * sideProtectionState = IDLE	IDLE
	REQ from other side *isProtectionPreemptive(REQ)	13		IDLE

Table 1.6—Wrapping state diagram for an ingress link

Current state		Row	Next state	
state	condition		action	state
WTR	Rx{IDLE,Neighbor,W,S}* WTR expired	14	Tx{IDLE,SELF,I,S}*Tx{IDLE,SELF,I,L}* sideProtectionState = IDLE	IDLE
	Rx{WTR,Neighbor,W,S}* WTR expired	15	Tx{IDLE,SELF,W,S}*Tx{IDLE,SELF,W,L}	WTR
	Rx{IDLE, NewNbr,I,S}	16	Tx{IDLE,SELF,I,S}*Tx{IDLE,SELF,I,L}* sideProtectionState = IDLE	IDLE
	Local REQ request * isProtectionGrant(REQ)	17	Tx{REQ,SELF,W,S}*Tx{REQ,SELF,W,L}* sideProtectionState = REQ	Wrapping
	Rx{REQ,Neighbor,W,S}	18	Tx{IDLE,SELF,W,S}*Tx{IDLE,SELF,W,L} * neighborProtectionState = REQ	Wrapping
	REQ from other side *isProtectionGrant(REQ)	19	sideProtectionState = IDLE	IDLE

Row 1.6-1: Local protection request or local failure detect.

Row 1.6-2: If local request is not granted then keep this request in sidePendingReq.

Row 1.6-3: Received a short path protection request message from the neighboring station.

Row 1.6-4: Station is wrapped due to local failure, then local failure clears.

Row 1.6-5: Station is wrapped due to local protection event, then local protection request clears.

Row 1.6-6: Station is wrapped due to local protection event. After local protection event clears, side pending request takes precedence.

Row 1.6-7: Station is wrapped due to local protection event. After local protection event clears, the other side's pending request takes precedence.

Row 1.6-8: Station is wrapped due to neighbor station's local failure detect, then the failure clears.

Row 1.6-9: Station is wrapped due to neighbor station's local protection request, then the protection request clears.

Row 1.6-10: Station is wrapped due to neighbor station's local failure detect. After neighbor's failure clears, side pending request takes precedence.

Row 1.6-11: Station is wrapped due to neighbor station's local protection request. After neighbor's protection request clears, the other side's pending request takes precedence.

Row 1.6-12: A higher priority protection event occurs at another station on the ring.

Row 1.6-13: A higher priority protection event occurs on the other side of the station.

Row 1.6-14: WTR timer expires and IDLE status is received from neighbor station.

Row 1.6-15: WTR timer expires but neighbor's WTR timer has not yet expired.

Row 1.6-16: A new neighbor is detected and the WTR timer is dropped on that interface

Row 1.6-17: A local failure is detected or a local protection request appears.

Row 1.6-18: Received a failure notification from neighbor.

Row 1.6-19: Other side is granted a protection request.

.

Editors' Notes: To be removed prior to final publication.

Note B1
The examples below between this editor's note (B1) and the next editor's note (B2) will be modified as per work done by the PAH to resolve the request for scenario descriptions in comment #581 from July.

1.7.5 Example

In Figure 1.7, Station A detects SF (via local request or self-detected request) on the span between Station A and Station B and starts sourcing {SF, A, W, S} on ringlet 1 and {SF, A, W, L} on ringlet 0. Station B receives the protection request from Station A (short path request) and starts sourcing {IDLE, B, W, S} on ringlet 0 and {IDLE, B, W, L} on ringlet 1 periodically.

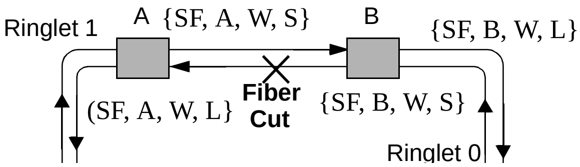

Figure 1.7—RPR protection switch signaling

Editors' Notes: To be removed prior to final publication.

Note B2
See Editor's Note B1 prior to 1.7.5.

Editors' Notes: To be removed prior to final publication.

Note A1
The examples below between this editor's note (A1) and the next editor's note (A2) will be modified as per work done by the PAH to resolve the request for scenario descriptions in comment #581 from July.

1.8 Failure examples

1.8.1 Signal failure - single fiber cut scenario

This sample scenario is a ring of six stations: A, B, C, D, E, and F, with a unidirectional failure on a link from A to B, detected on B. The ring is in the Idle state (all stations are Idle) prior to the failure.

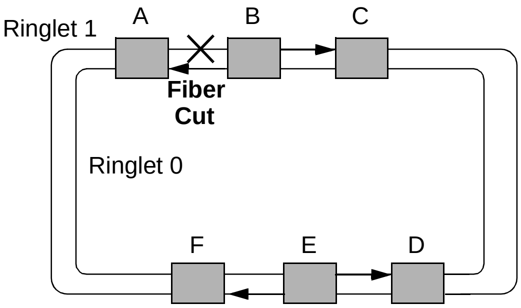

Figure 1.8—An RPR ring with a fiber cut in ringlet 1

1.8.1.1 Signal fail scenario

1) Station B detects SF on ringlet 1, transitions to Wrapped state (performs a wrap), transmits {SF, B, W, S} towards A on ringlet 0, and transmits {SF, B, W, L}on ringlet 1.
2) Station A receives a protection request on the short path, and transitions to Wrapped state.
3) Steady state is reached.

1.8.1.2 Signal fail clears

1) SF on Station B clears, Station B does not unwrap if it is in wrap state, sets WTR timer, transmits {WTR, B, W, S} on ringlet 0, and transmits {WTR, B, W,L} on ringlet 1.
2) Station A receives a WTR request on the short path, and does not unwrap.
3) Steady state is reached.
4) WTR times out on Station B. Station B transitions to idle state (unwraps), transmits {IDLE, B, I, S} on ringlet 0, and transmits{IDLE,B,I.L} on ringlet 1.
5) Station A receives {IDLE, B, I, L} and transitions to Idle.
6) Steady state is reached.

1.8.2 Signal failure - bidirectional fiber cut scenario

This sample scenario is a ring of six stations: A, B, C, D, E, and F, with a bidirectional failure between A and B. The ring is in the Idle state (all stations are Idle) prior to the failure.

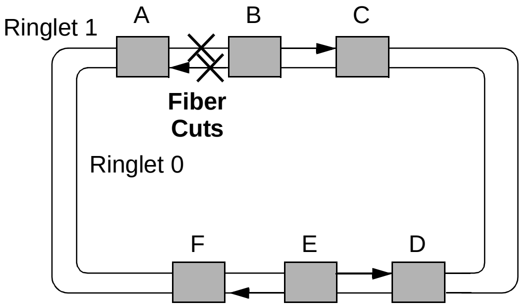

Figure 1.9—An RPR ring with bidirectional fiber cut

1.8.2.1 Signal fail scenario

1) Station A detects SF on ringlet 0, transitions to Wrapped state (performs a wrap), transmits {SF, A, W, S} towards B on ringlet 1, and transmits {SF, A, W, L} on ringlet 0.
2) Station B detects SF on ringlet 1, transitions to Wrapped state (performs a wrap), transmits {SF, B, W, S} towards A on ringlet 0, and transmits {SF, B, W, L} on ringlet 1.
3) Steady state is reached.

1.8.2.2 Signal fail clears

1) SF on A clears, A does not unwrap, A sets the WTR timer, transmits {WTR, A, W, S} through ringlet 1towards B, and transmits {WTR, A, W, L} on the long path through ringlet 0.
2) SF on B clears, B does not unwrap. Since it now has a short path WTR request from A, it acts upon this request. It keeps the wrap, transmits {IDLE, B, W, S} towards A, and transmits {WTR, B, W, L} on the long path.
3) Steady state is reached.
4) WTR times out on A. A enters the idle state (drops wrap) and starts transmitting idle on both ringlets.
5) B sees an idle request on the short path and enters idle state.
6) Steady state is reached.

1.8.3 Failed station scenario

This sample scenario is a ring of six stations: A, C, B, D, E, and F, where station C fails. The ring is in the Idle state (all stations are Idle) prior to the failure.

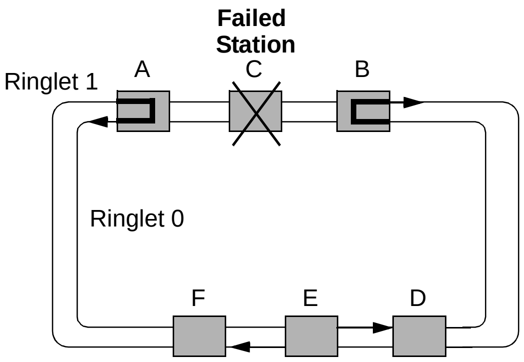

Figure 1.10—An RPR ring with a failed station

1.8.3.1 Station failure (or fiber cuts on both sides of the station)

1) B detects SF on ringlet 1, transitions to Wrapped state (performs a wrap), transmits {SF, B, W, S} towards C on ringlet 0, and transmits {SF, B, W, L} on ringlet 1.
2) A detects SF on ringlet 0, transitions to Wrapped state (performs a wrap), transmits {SF, A, W, S} towards C on ringlet 1, and transmits {SF, A, W, L} on ringlet 0.
3) Steady state is reached.

1.8.3.2 Failed station and one span return to service

Note: A station will usually return to service with one segment coming up after the other (with the time delta potentially close to 0). In this scenario, a station is powered up with the links connected and fault free.

1) Station C and a segment between A and C return to service (SF between A and C disappears).
2) Station C, not seeing any faults, starts to source idle messages {IDLE, C, I, S} in both directions.
3) The fault disappears on A, and A enters WTR state (briefly).
4) Station A receives an idle message from station C. Because the long path protection request {SF, B, W, L} received over the long span is not originating from the short path neighbor (C), station A drops the WTR.
5) Steady state is reached.

1.8.3.3 Second span returns to service

The scenario is like the bidirectional fiber cut fault clearing scenario described in 1.8.2.

This sample scenario is a ring where stations A and B are initially connected. Subsequently, the links between stations A and B are disconnected, and a new station, C, is inserted.

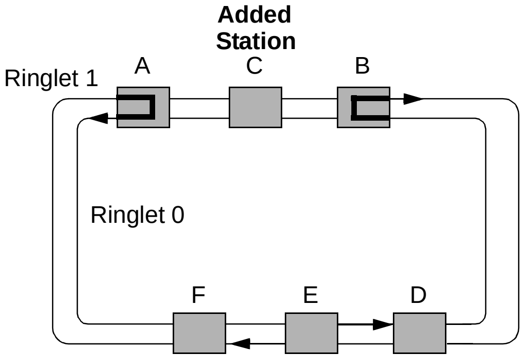

Figure 1.11—An RPR ring with a failed station

1.8.3.4 Bidirectional fiber cut

1) Links are removed between stations A and B.
2) B detects SF on ringlet 1, transitions to Wrapped state (performs a wrap), transmits {SF, B, W, S} towards A on ringlet 0, and transmits {SF, B, W, L} on ringlet 1.
3) A detects SF on ringlet 0, transitions to Wrapped state (performs a wrap), transmits {SF, A, W, S} towards B on ringlet 0, and transmits {SF, A, W, L} on ringlet 1.
4) Steady state is reached.

1.8.3.5 Station C is powered up and fibers between stations A and C are reconnected

This scenario is identical to the returning a failed station to service scenario described in 1.8.3.2.

1.8.3.6 Second span put into service

Stations C and B are connected. The scenario is identical to the bidirectional fiber cut fault clearing scenario described in 1.8.2.2.

Editors' Notes: To be removed prior to final publication.

Note A2
See Editor's Note A1 prior to 1.8.