

1. Topology discovery

Editors' Notes: To be removed prior to final publication.	1
	2
	3
	4
TOM: Can't figure out why Table 1.2 refuses to autonumber correctly without override <n=2>.	5
Cross-reference needed in 10.2.6 to Table 11.3.	6
Cross-reference needed in 10.4.5 to Clause 6.3.	7
	8
References:	9
To be added.	10
	11
Definitions:	12
	13
Abbreviations:	14
To be added.	15
	16
Revision History:	17
Draft 0.3, June 2002 Initial draft document for RPR WG review. Editorial notes added for clarification.	18
Draft 1.0, August 2002I Addressed 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.	19
Major modifications made in D1.0 include:	20
Addition of layer diagram	21
Clarifications in topology database and isolated station subclauses	22
Reorganization and additions to topology discovery message subclause	23
	24
	25
	26
	27
	28
	29
	30
	31
	32
	33
	34
	35
	36
	37
	38
	39
	40
	41
	42
	43
	44
	45
	46
	47
	48
	49
	50
	51
	52
	53
	54

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

This is the topology clause submitted by the protection ad-hoc (PAH) to the working group for approval in the September meeting.

1. Comment #528: Provide topology state machine.

Status: Completed. Kshitij Kumar provided the state machine.

2. Comment #529: Provide scenarios illustrating whether there are topology discovery related requirements for the initiation of wrapping and/or steering protection.

Status: Completed. The initiation of wrapping and/or steering protection does not require that topology discovery has completed, in the event, for example, that a protection event overlaps a period during which new stations have just been added to the ring. The determination of a preferred direction on the ring for steering protection at each station is done only for stations contained in the topology database. The determination of whether to wrap at a station is not dependent on the contents of the topology database.

Comment #535: Determine whether the topology module needs to convey a topology change indication along with the corresponding old and new neighbor information so that protection switching can determine whether a link coming up should be subject to wait to restore (WTR), or whether this can be handled by the protection protocol.

Status: Completed. The topology module will store former neighbor MAC information so that protection switching can determine whether a link coming up should be subject to WTR.

4. Comment #538: Determine the criteria for topology consistency checking, the algorithm used for the checking, and the intent for the checking (to inform the operator via events).

Status: Completed. Jim Kao provided the criteria for topology consistency checking. Algorithms for performing this checking may be added for informative purposes to either the clause or to an annex.

5. Comments #541, #557, and #558: Define the usage of the downstream link reserved subclassA0 bandwidth, how it is computed, and its units. The PAH, in conjunction with the RAH, also needs to define how each station gets the necessary information to compute this bandwidth.

Status: Completed. The units for the bandwidth and how it is reported (via the topology extended status message) have been defined.

6. Comment #560: Clarify the interworking of topology and protection. It will become clear through this exercise what the actual list of triggers is for generation of topology discovery messages.

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.

7. Comment #562: Modify 10.2.2 to illustrate the key scenarios for topology discovery. These include (for example) station initialization, addition of a station, removal of a station, isolation of a station, handling of bandwidth allocation, and handling of misconfigurations. Description will be included as to how topology interacts with protection and fairness. The topology database description in 10.2.6 will be reorganized to apply to the different scenarios. The topology discovery message description in 10.2.4 will also be reorganized to apply to the different scenarios.

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.

1.1 Scope

This clause describes the RPR topology discovery protocol, which implements a reliable and accurate means for all RPR stations on a ring to discover the initial topology of the stations on the ring and any changes to that topology. The protocol is intended to be very scalable, to cause insignificant overhead for ring traffic, and to cause insignificant overhead on software and ASICs. The protocol resides in the MAC control sub-layer, as shown in the shaded region of Figure 1.1.




Figure 1.1—RPR layer diagram

The services and features provided are:

- a) Determine/validate connectivity and ordering of stations on the ring.
- b) Ensure all stations on the ring will converge to a uniform and current image of the topology normally within 1 ringlet circulation time.
- c) Tolerant of message loss.
- d) Operate without any master station on the ring.
- e) Operate independently of and in the absence of any management systems.
- f) Usable with all supported topologies: ring, bus (broken ring), and isolated station.
- g) Support dynamic addition and removal of stations to/from the ring.
- h) Detect mis-cabling between stations.
- i) Provide means of sharing additional information between stations.
- j) Cause insignificant overhead.

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

Determination of connectivity and ordering of stations on the ring includes the detection of physical connectivity on "half-alive" spans. It is an objective of topology discovery, for example, that station F in the partial string A=B=C->D=E=F be able to provide a user performing diagnostics at F with the complete view of the physical connectivity of the string.

Detecting mis-cabling (above) assumes static local assignment of ring IDs (as opposed to dynamic discovery). How the assignments are made is a local, non-standardized decision. This might need to be discussed in the OAM clause.

The RPR topology discovery protocol is used to discover the physical link configuration between stations. It is not within the scope of the RPR topology discovery protocol to determine the dynamic link status information, i.e. which ringlet links are up or down, ring segment failures, etc. The discovered topology is used by other protocols such as the RPR protection protocol and the RPR congestion avoidance protocol.

1.2 Algorithm overview

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

As per comment #562 from July, the PAH will modify this subclause as needed (given the annex containing topology and protection scenarios) to illustrate the key scenarios for topology discovery. These include (for example) station initialization, addition of a station, removal of a station, isolation of a station, handling of bandwidth allocation, and handling of misconfigurations. Description will be included as to how topology interacts with protection and fairness.

The RPR topology discovery protocol provides each station on the ring with knowledge of the number and arrangement of other stations on the ring. This collection of information is referred to as the topology image. Each station maintains its own local copy of the topology image for the entire ring. Initially, the station's topology image contains information only about itself.

Ring topology discovery is initiated as needed and periodically. No station acts as a master for the topology image or for the protocol. All ringlet segments that can be discovered are included. A fully connected ring is not needed for the protocol.

In addition to station identifiers and physical connectivity relationships, the topology discovery protocol is also used to propagate additional station information, for use in other parts of this standard.

The messages sent as part of the RPR topology discovery protocol are indicated in the RPR frame header as control frames.

1.2.1 At initialization

At station initialization, the local topology image is initialized to contain only the local station and no links, and the addresses of the neighboring stations are initialized to all 0's. The station starts the topology algorithm by broadcasting a topology discovery message on all ringlets. Then it continually listens for topology discovery messages broadcast on the ring, and broadcasts topology discovery messages periodically and whenever there is a local topology change, including a detected change in neighbor identity.

1.2.2 At addition of a station

When a station is inserted into an existing ring, it initializes itself as described above. Its neighbors will detect the station as a new neighbor by receiving its topology discovery message with a TTL set to MAX_STATIONS (indicating the message has traveled exactly one hop) and with the SA set to a value other than that previously stored (if any). The neighbors will respond by sending an immediate topology discovery message providing topology information to the new station. All other stations will detect the new station by receiving its topology discovery message, and each will send an immediate topology discovery message, providing topology information to the new station.

1.2.3 At span failure

When a span fails, the stations detecting the failure store the knowledge of their neighbor station address on the failed interface and clear their current knowledge of their neighbor station address on that interface. The stations send protection messages that report exactly which links failed to the other stations on the ring. The stations also are triggered to send topology messages due to the change in their neighbor identity on the failed interface.

1.2.4 At removal of a station

When a station is removed from a ring, the stations detecting the removal of the station take the same actions as described in 1.2.3 in the event of detection of signal fail on incoming interfaces.

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

Passthrough mode, if required, can be handled by topology discovery in an a similar fashion to what is described in 1.2.4, except that the trigger is now the detection of a new neighbor on an interface rather than a signal fail indication. In this case, the current neighbor address is replaced, and there is no need to store the identity of the previous neighbor.

1.2.5 Topology discovery

Upon triggers defined in 1.5.1.1, a station broadcasts a topology discovery message to all stations on both ringlets. The topology discovery message contains information about the local station relevant to connectivity, including its links to its neighbors. When a station receives a topology discovery message, it updates its local topology image.

The topology discovery message contains the ringlet ID on which it is sent. When a station receives a topology discovery message with a TTL set to MAX_STATIONS (indicating the message has traveled exactly one hop), it verifies that the ringlets of both stations are connected with the same ringlet ID value (as indicated in the received message and by local configuration).

The topology discovery message, like all RPR frames, contains the source MAC address of the station from which it is sent. When a station receives a topology discovery message with a TTL set to MAX_STATIONS (indicating the message has traveled exactly one hop), it verifies that it knows who its neighbor is.

1.2.6 Topology database and hop count determination

Under certain conditions such as a failed span in a bidirectional ring, a station will only be able to receive topology discovery messages from the other stations on the ring on one ringlet. It is a requirement to be able to construct a complete topology database reflecting connectivity information on both ringlets. Hence, a single topology database showing information for both ringlets can meet this requirement as well as facilitate maintainance of a consistent database.

The hop count to another station can be calculated through examination of the topology image. For each ringlet, the topology image must be completed up to the station whose hop count is desired.

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

As per comment #537 from July, the formal definition of the information in the topology database should be present in the OAM and/or MIB clauses. This will be determined in upcoming versions of the draft.

The values of each entry are: the downstream hop count on ringlet 0, the downstream hop count on ringlet 1, the local MAC address, the MAC address of the neighboring station connected to the west interface of the station, the MAC address of the neighboring station connected to the east interface of the station, the MAC address of the neighboring station previously connected to the west interface of the station, the MAC address of the neighboring station previously connected to the east interface of the station, the protection status of the receive link from the west station, the protection status of the receive link from the east station, the reserved subclassA0 bandwidth on the transmit link to the west station, the reserved subclassA0 bandwidth on the transmit link to the east station, data reachability downstream on ringlet 0, and data reachability

downstream on ringlet 1. The receive link's availability can be any of the protection message request values given in Table 11.3, and are filled in by the protection protocol. The reserved subclassA0 bandwidth is described in 1.5.2.

The topology database also contains the MAC address information of the neighbor previously connected to each interface of the local station. This information is stored in the topology database for the use of the protection protocol. The former neighbor entries will be rewritten upon a change in the value of the current neighbor entry for each respective interface.

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

Former neighbor MAC information may be put in a separate table since this information applies only to the local station.

An example table for one ringlet is shown in Table 1.1.

Table 1.1—Topology and status database example

hop count for ringlet 0	hop count for ringlet 1	local MAC	west neighbor's MAC	east neighbor's MAC
0	0	00-10-A4-97-A8-DE	00-10-A4-97-A8-EF	00-10-A4-97-A8-BD
1	3	00-10-A4-97-A8-EF	00-10-A4-97-A8-AC	00-10-A4-97-A8-DE
2	2	00-10-A4-97-A8-AC	00-10-A4-97-A8-BD	00-10-A4-97-A8-EF
3	1	00-10-A4-97-A8-BD	00-10-A4-97-A8-DE	00-10-A4-97-A8-AC

Table 1.2—Topology database example continued

former west neighbor's MAC	former east neighbor's MAC	west receive link availability	east receive link availability
00-10-A4-97-A8-GH	00-10-A4-97-A8-BD	IDLE	IDLE
N/A	N/A	IDLE	IDLE
N/A	N/A	IDLE	IDLE
N/A	N/A	IDLE	IDLE

Table 1.3—Topology database example continued

west transmit link reserved subclassA0 bandwidth	east transmit link reserved subclassA0 bandwidth	reachability, ringlet 0	reachability, ringlet 1
50	40	N/A	N/A
100	80	Data	Data

Table 1.3—Topology database example continued

150	120	Data	Data
200	160	Data	Data

1.2.6.1 Modification of the topology database

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

A flow chart or other more detailed description of the rules for this section is desirable.

The topology database is modified upon receipt of topology or protection messages that result in a change in information contained in the database. Protection messages report changes in link status. Topology discovery messages may report a change in the address of a station a given number of hops away on a given ringlet, a change in the address of a neighboring station of a station a given number of hops away, a change in reserved subclassA0 bandwidth reported by that station, or a change in station capabilities (such as whether it is wrap capable). Changes in reserved subclassA0 bandwidth or in station capabilities by themselves result only in modification of single entries within the database.

Changes in the address of a station a given number of hops away or in the address of a neighboring station of a station a given number of hops away force, for the sake of consistency, the deletion of all entries in the database corresponding to stations on the ringlet beyond the point of the change. For example, if a station 2 hops away reports a signal fail on receive of ringlet 0, this implies that stations further upstream from that station on ringlet 0 are also no longer connected to the ringlet. Therefore, the hop count entries for ringlet 1 downstream of values 3 or greater are cleared. If multiple failures occur on a ring such that both the hop count values for ringlet 0 and ringlet 1 are cleared for a given station, then that station is removed from the topology database.

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

Passthrough mode, if required, can be handled by topology discovery in an a similar fashion to what is described above. If a station 2 hops away reports a new neighbor, this implies that stations further away than that station on the given ringlet currently stored in the topology database must be removed. Stations upstream must then reconfirm their connection to the ringlet so that they can be added back to the topology database.

Upon detection of signal fail on a given interface (west for example), the west neighbor MAC address is copied into the former west neighbor MAC address, then the west neighbor MAC address field is cleared. Upon detection of a new neighbor on the west interface, the new west neighbor MAC address is written into the west neighbor MAC address field.

1.2.7 Data/control reachability

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

The PAH needs to discuss whether topology messages should be forwarded through spans containing one failed link. If topology messages are not forwarded, then the topology database at a station will contain information about data reachable stations only, and about the immediate neighbors of the station. Another mechanism will then be required to enable determination of the complete physical connectivity of the topology.

A station S2 is considered downstream data reachable from a station S1 on a given ringlet when data can flow from S1 to S2 on that ringlet. If S2 is data reachable from S1 on ringlet 0, for example, then it is ensured that S1 is data reachable from S2 on ringlet 1 in a correctly connected ring.

A station S1 is considered control reachable from another station S2 when topology packets can be received at station S1 from station S2. It is possible that a station can be control reachable from another station, but not vice versa. If a station S2 appears in the topology database of station S1, then S1 is control reachable from S2.

The purpose of explicitly showing downstream data reachability in the topology database is to provide a clear indication of stations that are downstream data reachable from a given station, while at the same time enabling the complete physical connectivity of the ring topology to be shown. It is beneficial for the complete physical connectivity of the ring to be locally available at each station for field diagnostic purposes.

When a station is not data reachable or control reachable from any other station, it determines itself to be isolated. This will occur if the receive links on both the east and west interfaces of a station have a protection state of signal fail (SF). An isolated station shall remove all entries from its local topology image, but shall transfer its own east neighbor and west neighbor information to the former east neighbor and former west neighbor fields in the database before clearing the east neighbor and west neighbor fields. The information in the former east neighbor and former west neighbor fields is needed to determine whether a link should be brought up immediately, as described in 11.3.3.

1.2.8 Topology consistency check

It can be easily determined when an image is complete and consistent by examining the image contents. When the contents of the local topology image show station information for each station described in the link information of another station, then the image is complete. For each ringlet, when the contents of the local topology image show that all stations on that ringlet are connected to each other in a logical ring, bus (broken ring), or isolated station, then the topology image is consistent.

A canonical form for the topology image allows all the stations to eventually arrive at the same image for the topology.

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

Algorithms for performing this checking may be added for informative purposes to either the clause or to an annex.

1.2.8.1 Determination and validation of ringlet ID

Each station determines which interface is associated with which ringlet and assigns the corresponding ringlet_id either through fixed mapping between hardware locations or through configuration. Each topology control message is sent separately on each ringlet, identifying the ringlet on which it is being sent. Any topology message with a TTL set to MAX_STATIONS received on a ringlet different from the ringlet on which it is identified as being sent shall cause the link to be declared non-operational and trigger a mis-configuration alarm. This will result in alarm-triggering events being generated only from stations at which the physical misconfiguration has occurred. That shall trigger a signal fail for the link in protection. When a matched ringlet ID topology message from its neighbor is received, that shall clear the signal fail in protection.

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

The marking of a link as non-operational due to misconfiguration will use a status of signal fail (SF).

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

A subclause needs to be added to put this clause in context with the remainder of the standard in terms of the MAC reference model.

1.2.8.2 Neighbor inconsistency

The following will trigger a neighbor inconsistency alarm or indication:

- a) East station's west neighbor MAC address is not equal to local station MAC address, or
- b) West station's east neighbor MAC address is not equal to local station MAC address.

The zero MAC address is an exceptional case for this inconsistency check. The inconsistency check shall be performed when the neighbor information is updated due to a new received topology discovery message. An alarm will be triggered to operator if this neighbor inconsistency is detected.

1.2.8.3 Duplicate MAC address

A MAC address is the unique identification for a station on an RPR ring. If a station receives a topology discovery message from another station with the same MAC address as its MAC address, an alarm will be triggered to the operator for this duplicate MAC address.

1.2.8.4 The number of stations exceed MAX_STATIONS

If the total number of entries in the topology database exceeds MAX_STATIONS, an alarm or indication shall be triggered to the operator.

1.3 Topology discovery process

1.3.1 Topology discovery process description

An exact definition of the topology discovery process is specified in Table 1.4, where rows are evaluated in top-to-bottom order.

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

Additions to this state diagram are needed to show:

1. Actions taken for generation and upon receipt of topology extended status message
2. When different types of validation are performed

Updates to this state diagram are needed to show:

3. Neighbor MAC addresses are not contained in the topology discovery message, but rather in the topology extended status message
4. Now that the weight field has been moved to the extended status message, are there any possible invalid station capabilities?
5. Additional states may be needed to show how the mis-configuration notifications mentioned are cleared.
6. Topology discovery messages are triggered only on change in the content of the messages (line 10)
7. Is NeighbourChanged notification to operator required? (line 39)
8. The state diagram may need to include the triggering of a topology discovery message from each station on the ring upon receipt of a topology discovery message from a non-data-reachable station. If both receive links of a station are in SF state, then that station clearly has no knowledge of its neighbors. However, in other conditions (such as both receive links of a station are in FS state), the station will have knowledge of its immediate neighbors only.

After some of these modifications are completed, then row descriptions will be added.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Table 1.4—Topology discovery state table

Current State		Row	Next State	
State	Conditions		Action	State
Dormant (Waiting for initiation)	System receives a wakeup signal (or any other event)	1	Start Topology Timer. Clear Topology Database	Init
Init	Incoming TD message on Ringlet 0 (TTL == MAX_TTL)	2	Insert West Neighbour information into Topology Database: Source MAC, station capabilities, and with Ringlet 0 hop count = 1. Send TD Message on both Ringlets, East Neighbour MAC = 0, West Neighbour MAC as received. Initiate Fast Topology Timer	WNHF
	Incoming TD message on Ringlet 1 (TTL == MAX_TTL)	3	Insert East Neighbour information into Topology Database: Source MAC, station capabilities, and with Ringlet 1 hop count = 1. Send TD Message on both Ringlets, West Neighbour MAC = 0, East Neighbour MAC as received. Initiate Fast Topology Timer	ENHF
	Incoming TD message on Ringlet 0 (TTL != MAX_TTL) and source node MAC IS NOT in DB	4	Insert entry for this node in DB, with hop count = (MAX_TTL minus TTL) for Ringlet 1. Send TD Message on both Ringlets, Neighbour MACs = 0, Initiate Fast Topology Timer	Init
	Incoming TD message on Ringlet 1 (TTL != MAX_TTL) and source node MAC IS NOT in DB	5	Insert entry for this node in DB, with hop count = (MAX_TTL minus TTL) for Ringlet 0. Send TD Message on both Ringlets, Neighbour MACs = 0, Initiate Fast Topology Timer	Init
	Incoming TD message on Ringlet 0 (TTL != MAX_TTL) and source node MAC EXISTS in DB	6	If Station Capabilities have changed, Update DB with entry for this node hop count = (MAX_TTL minus TTL) on Ringlet 1, Initiate Fast Topology Timer? Else - No Action.	Init
	Incoming TD message on Ringlet 1 (TTL != MAX_TTL) and source node MAC EXISTS in DB	7	If Station Capabilities have changed, Update DB with entry for this node, hop count = (MAX_TTL minus TTL) on Ringlet 0, Initiate Fast Topology Timer? Else – No Action.	Init

Table 1.4—Topology discovery state table

	Fast Topology Timer Expired	8	Send TD Message on both Ringlets, Neighbour MACs = 0, If Timer has expired 8 times, start Slow Topology Timer, Else restart Fast Topology Timer	Init
	Slow Topology Timer Expired	9	Send TD Message on both Ringlets, Neighbour MACs = 0, Start Slow Topology Timer	Init
	Locally detected Signal Fail Indication on either Ringlet	10	Send TD Message on both Ringlets, Neighbour MACs = 0, Initiate Fast Topology Timer	Init
	Incoming TD message with Invalid Station Capabilities or Ringlet ID parameter	11	Send Mis-configuration Notification	Init
WNHF (West Neighbour heard from)	Incoming TD message on Ringlet 0 (TTL = MAX_TTL), and MAC SA or Station Capabilities have not changed	12	No Action	WNHF
	Incoming TD message on Ringlet 0 (TTL = MAX_TTL), where MAC SA or Station Capabilities have changed	13	Send Neighbour Changed Notification to Operator, Update Topology DB with new neighbour SA or capabilities. Send TD message on both Ringlets with new neighbour information. Initiate Fast Topology Timer.	WNHF
	Incoming TD message on Ringlet 1 (TTL = MAX_TTL)	14	Update Database with East Neighbour Send TD message on both Ringlets with new East Neighbour address as well as West Neighbour address. Initiate Fast Topology Timer.	BNHF
	Incoming TD message on Ringlet 0 (TTL != MAX_TTL) where source node MAC EXISTS in DB	15	If Station Capabilities have changed, Update DB with entry for this node hop count = (MAX_TTL minus TTL) on Ringlet 1. Send TD Message on both Ringlets, Initiate Fast Topology Timer? Else - No Action.	WNHF
	Incoming TD message on Ringlet 1 (TTL != MAX_TTL) where source node MAC EXISTS in DB	16	If Station Capabilities have changed, Update DB with entry for this node hop count = (MAX_TTL minus TTL) on Ringlet 0. Send TD Message on both Ringlets, Initiate Fast Topology Timer? Else - No Action.	WNHF

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Table 1.4—Topology discovery state table

	Incoming TD message on Ringlet 0 (TTL != MAX_TTL) where source node MAC does NOT exist in DB	17	Insert entry for this node in DB, with hop count = (MAX_TTL minus TTL) for Ringlet 1. Send TD Message on both Ringlets, Initiate Fast Topology Timer	WNHF
	Incoming TD message on Ringlet 1 (TTL != MAX_TTL) where source node MAC does NOT exist in DB	18	Insert entry for this node in DB, with hop count = (MAX_TTL minus TTL) for Ringlet 0. Send TD Message on both Ringlets, Initiate Fast Topology Timer	WNHF
	Fast Topology Timer Expired	19	Send TD Message on both Ringlets, If Timer has expired 8 times, start Slow Topology Timer, Else restart Fast Topology Timer	WNHF
	Slow Topology Timer Expired	20	Send TD Message on both Ringlets. Restart Slow Topology Timer	WNHF
	Locally detected Signal Fail Indication, Ringlet 0 unavailable	21	Send TD Message on both Ringlets, with West Neighbour's MAC = 0, East Neighbour's MAC = 0. Initiate Fast Topology Timer	Init
	Locally detected Signal Fail Indication, Ringlet 1 unavailable	22	No action	WNHF
	Incoming TD message with Invalid Station Capabilities or Ringlet ID on Ringlet 0	23	Send Mis-configuration Notification	Init
	Incoming TD message with Invalid Station Capabilities or Ringlet ID on Ringlet 1	24	Send Mis-configuration Notification	WNHF
ENHF (East Neighbour heard from)	Incoming TD message on Ringlet 0 (TTL = MAX_TTL), and MAC SA or Station Capabilities have not changed	25	No Action	ENHF
	Incoming TD message on Ringlet 1 (TTL = MAX_TTL), where MAC SA or Station Capabilities have changed	26	Send Neighbour Changed Notification to Operator, Update Topology DB with new neighbour SA or capabilities. Send TD message on both Ringlets with new neighbour information. Initiate Fast Topology Timer.	ENHF

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Table 1.4—Topology discovery state table

Incoming TD message on Ringlet 0 (TTL = MAX_TTL)	27	Update Database with West Neighbour Send TD message on both Ringlets with new East Neighbour address as well as West Neighbour address. Initiate Fast Topology Timer.	BNHF
Incoming TD message on Ringlet 0 (TTL != MAX_TTL) where source node MAC EXISTS in DB	28	If Station Capabilities have changed, Update DB with entry for this node hop count = (MAX_TTL minus TTL) on Ringlet 1. Send TD Message on both Ringlets, Initiate Fast Topology Timer? Else - No Action.	ENHF
Incoming TD message on Ringlet 1 (TTL != MAX_TTL) where source node MAC EXISTS in DB	29	If Station Capabilities have changed, Update DB with entry for this node hop count = (MAX_TTL minus TTL) on Ringlet 0. Send TD Message on both Ringlets, Initiate Fast Topology Timer? Else - No Action.	ENHF
Incoming TD message on Ringlet 1 (TTL != MAX_TTL) where source node MAC does NOT exist in DB	30	Insert entry for this node in DB, with hop count = (MAX_TTL minus TTL) for Ringlet 0. Send TD Message on both Ringlets, Initiate Fast Topology Timer	ENHF
Fast Topology Timer Expired	31	Send TD Message on both Ringlets, If Timer has expired 8 times, start Slow Topology Timer, Else restart Fast Topology Timer	ENHF
Slow Topology Timer Expired	32	Send TD Message on both Ringlets. Restart Slow Topology Timer	ENHF
Locally detected Signal Fail Indication, Ringlet 1 unavailable	33	Send TD Message on both Ringlets, with West Neighbour's MAC = 0, East Neighbour's MAC = 0. Initiate Fast Topology Timer	Init
Locally detected Signal Fail Indication, Ringlet 0 unavailable	34	No action	ENHF
Incoming TD message with Invalid Station Capabilities or Ringlet ID on Ringlet 1	35	Send Mis-configuration Notification	Init
Incoming TD message with Invalid Station Capabilities or Ringlet ID on Ringlet 0	36	Send Mis-configuration Notification	ENHF

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Table 1.4—Topology discovery state table

BNHF (Both Neighbours heard from)	Incoming TD message on Ringlet 0 (TTL = MAX_TTL), where MAC SA or station capabilities have not changed	37	No action	BNHF
	Incoming TD message on Ringlet 1 (TTL = MAX_TTL), where MAC SA or station capabilities have not changed	38	No action	BNHF
	Incoming TD message on Ringlet 1 or 0 (TTL = MAX_TTL), where MAC SA or Station Properties have changed	39	Send Neighbour Changed Notification to Operator Update Topology DB with new neighbour. Send TD message on both Ringlets with new neighbour information. Initiate Short Topology Timer.	BNHF

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Table 1.4—Topology discovery state table

Incoming TD message on Ringlet 0 (TTL != MAX_TTL) where source node MAC EXISTS in DB	40	If Station Capabilities have changed, Update DB with entry for this node hop count = (MAX_TTL minus TTL) on Ringlet 1. Send TD Message on both Ringlets, Initiate Fast Topology Timer? Else - No Action.	BNHF
Incoming TD message on Ringlet 1 (TTL != MAX_TTL) where source node MAC EXISTS in DB	41	If Station Capabilities have changed, Update DB with entry for this node hop count = (MAX_TTL minus TTL) on Ringlet 0. Send TD Message on both Ringlets, Initiate Fast Topology Timer? Else - No Action.	BNHF
Incoming TD message on Ringlet 1 (TTL != MAX_TTL) where source node MAC does NOT exist in DB	42	Insert entry for this node in DB, with hop count = (MAX_TTL minus TTL) for Ringlet 0. Send TD Message on both Ringlets, Initiate Fast Topology Timer	BNHF
Incoming TD message on Ringlet 0 (TTL != MAX_TTL) where source node MAC does NOT exist in DB	43	Insert entry for this node in DB, with hop count = (MAX_TTL minus TTL) for Ringlet 1. Send TD Message on both Ringlets, Initiate Fast Topology Timer	BNHF
Fast Topology Timer Expired	44	Send TD Message on both Ringlets, If Timer has expired 8 times, start Slow Topology Timer, Else restart Fast Topology Timer	BNHF
Slow Topology Timer Expired	45	Send TD Message on both Ringlets. Restart Slow Topology Timer	BNHF
Locally detected Signal Fail Indication, Ringlet 1 unavailable	46	Send TD Message on both Ringlets, with West Neighbour's MAC = 0, East Neighbour's MAC = 0. Initiate Fast Topology Timer	WNHF
Locally detected Signal Fail Indication, Ringlet 0 unavailable	47	Send TD Message on both Ringlets, with West Neighbour's MAC = 0, East Neighbour's MAC = 0. Initiate Fast Topology Timer	ENHF
Incoming TD message with Invalid Station Capabilities or Ringlet ID on Ringlet 0	48	Send Mis-configuration Notification	ENHF
Incoming TD message with Invalid Station Capabilities or Ringlet ID on Ringlet 1	49	Send Mis-configuration Notification	WNHF

1.4 Message formats

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

Need MAC Control message type for topology extended status message.

There are two messages within topology discovery. The first, referred to as the topology discovery message, is used for reporting changes in neighbor identity. It is a fixed length message and is sent as a MAC Control message with a control type value of 1. The second is referred to as the topology extended status message. It is a variable length message and is sent as a MAC Control message with a control type value of TBD. The topology extended status message is not used for discovery of the physical topology of the ring, but rather is used to convey additional information that needs to be reported by a station to the rest of the ring on a slower timeframe, such as reserved bandwidth configuration information.

1.4.1 Topology discovery message format

The topology discovery message is used for reporting changes in neighbor identity. It is sent as a MAC Control message with a control type value of 1, for topology discovery. It is sent as a broadcast frame on all ring-lets, with a TTL of MAX_STATIONS, removed by the source station, and with the source MAC set to the actual MAC of the sending station.

The topology discovery message packet format is outlined in Figure 1.2.




Figure 1.2—Topology discovery message packet format

The topology specific fields are detailed below.

1.4.1.1 Station capabilities

The station capabilities bytes contain the capabilities of the sending station, as shown in Figure 1.3. The sub-fields of the station capabilities bytes are described in the subclauses below.

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

It is not necessary to add values for ingress link rates on the east and west interfaces to the MAC unless it is specified elsewhere in the standard that different available rates can coexist on the same ring.




Figure 1.3—Station capabilities field format

1.4.1.1.1 Reserved (RSVD)

Bits 7 through 0 of byte 1 and bits 7 through 2 of byte 2 are reserved for future use.

1.4.1.1.2 Jumbo frame receive capable (JC)

The JC bit is used to indicate if a station is jumbo frame (see 8.1) receive capable.

The values of JC are defined in Table 1.5.

Table 1.5—Jumbo frame receive capable (JC) values

Value	Description
0	Not capable of receiving jumbo frames
1	Capable of receiving jumbo frames

1.4.1.1.3 Wrap protection capable (WC)

The WC bit is used to indicate if a station is wrap protection capable.

The values of WC are defined in Table 1.6.

Table 1.6—Wrap protection capable (WC) values

Value	Description
0	Not capable of wrap protection
1	Capable of wrap protection

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

1.4.2 Topology extended status message format

The topology extended status message is used for reporting changes in station status that are not as time critical as the those reported in the topology discovery message. The topology extended status message is sent as a MAC Control message with a control type value of TBD. It is sent as a broadcast frame on all ringlets, with a TTL of MAX_STATIONS, removed by the source station, and with the source MAC set to the actual MAC of the sending station.

The topology extended status message frame format is outlined in Figure 1.4.




Figure 1.4—Topology extended status message frame format

The topology specific fields are detailed below.

1.4.2.1 Station capabilities

The station capabilities bytes contain the capabilities of the sending station, as shown in Figure 1.3 and as described in 1.4.1.1.

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

It needs to be determined if the station capabilities field needs to be repeated in both the topology discovery and topology extended status messages.

1.4.2.2 TLV entries

A Type-Length-Value (TLV) encoding scheme is used to encode additional station information in topology extended status messages. The TLVs may appear in the topology frame in any order.

A TLV is encoded as a 2 byte Type field that uses 14 bits to specify a Type and 1 bit to specify the behavior when a station doesn't recognize the Type, followed by a 2 byte Length field, followed by a variable length Value field. The general format of a TLV is shown in Figure 1.5.




Figure 1.5—TLV format

1.4.2.2.1 Unknown (U) bit

Upon receipt of an unknown TLV, if U is clear ($==0$), a notification must be returned to the management and the entire message must be ignored; if U is set ($==1$), the unknown TLV is silently ignored and the rest of the message is processed as if the unknown TLV did not exist. The sections following that define TLVs specify a value for the U-bit.

1.4.2.2.2 Reserved (R) bit

The R bit is reserved for future use. It shall be set to 0 and shall be ignored by the receiving stations.

1.4.2.2.3 Type

The Type field encodes how the Value field is to be interpreted.

1.4.2.2.4 Length

The Length field specifies the length of the Value field in bytes.

1.4.2.2.5 Value


The Value field is a byte string, of Length bytes, as specified in the Length field, that encodes information to be interpreted as specified by the Type field.

1.4.2.3 Defined TLV encodings

The following subclauses define the TLV encodings supported by this standard. Any additional TLV encodings used are beyond the scope of this standard.

1.4.2.3.1 Weight TLV

The Weight TLV encodes the station weight in each ringlet. The TLV format for the Weight TLV is shown in Figure 1.6.


**Figure 1.6—Weight TLV format**

The Ringlet 0 Weight field carries the weight of the station in Ringlet 0.

The Ringlet 1 Weight field carries the weight of the station in Ringlet 1.

1.4.2.3.2 Total reserved bandwidth TLV

The total reserved bandwidth TLV encodes the total station reserved subclassA0 bandwidth in each ringlet. The TLV format for the total reserved bandwidth TLV is shown in Figure 1.7.


**Figure 1.7—Total reserved bandwidth TLV format**

The Ringlet 0 reserved bandwidth field carries the total reserved subclassA0 bandwidth of the station in Ringlet 0.

The Ringlet 1 reserved bandwidth field carries the total reserved subclassA0 bandwidth of the station in Ringlet 1.

1.4.2.3.3 Multicast groups TLV

The Multicast group TLV encodes the multicast group addresses of the station. The TLV format for the Multicast groups TLV is shown in Figure 1.7.


**Figure 1.8—Multicast groups TLV format**

The Multicast address #1 field carries the first Multicast address assigned to the station.

The Multicast address #N field carries the Nth Multicast address assigned to the station.

1.4.2.3.4 Neighbor address TLV

The neighbor address TLV encodes the MAC addresses of the neighbors of the sending station. The TLV format for the neighbor address TLV is shown in Figure 1.9.

**Figure 1.9—Neighbor address TLV format**

The Ringlet 0 neighbor MAC address field carries the MAC address of the neighboring station downstream on Ringlet 0. If the station's MAC address is unknown, it shall be all 0's.

The Ringlet 1 neighbor MAC address field carries the MAC address of the neighboring station downstream on Ringlet 1. If the station's MAC address is unknown, it shall be all 0's.

1.4.2.3.5 Individual reserved bandwidth TLV

The individual reserved bandwidth TLV encodes the individual station reserved subclassA0 bandwidth for each link on each ringlet. The TLV format for the individual reserved bandwidth TLV is shown in Figure 1.10.




Figure 1.10—Individual reserved bandwidth TLV format

The Ringlet 0 reserved bandwidth fields carry the amount of reserved subclassA0 bandwidth reserved by the sending station for each link on Ringlet 0. The fields are ordered by the number of hops away, such that field 1 is for the link 1 hop away, and field N is for the link N hops away.

The Ringlet 1 reserved bandwidth fields carry the amount of reserved subclassA0 bandwidth reserved by the sending station for each link on Ringlet 1. The fields are ordered by the number of hops away, such that field 1 is for the link 1 hop away, and field N is for the link N hops away.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

1.4.2.3.6 Vendor specific TLV

The Vendor specific TLV encodes the vendor specific information of the station. The TLV format for the vendor specific TLV is shown in Figure 1.11.




Figure 1.11—Vendor specific TLV format

The vendor organizationally unique identifier (OUI) field contains the IEEE company identifier for the vendor.

The vendor specific data field includes vendor defined data. The exact definition of the vendor specific data field is outside of the scope of this standard.

1.5 Message handling

1.5.1 Topology discovery message handling

1.5.1.1 When generated

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

The PAH is discussing adding a topology discovery message trigger upon topology validation failure at a station.

The topology discovery message is broadcast on the initial start of RPR topology discovery, at any point that a station detects a change in local status or in the identity of a neighbor, at any point that a station detects a new station on the ring, and periodically. The topology discovery messages are sent on a bi-level periodic timer, starting with the fast period. The fast topology discovery message period is configurable from 1 ms to 1000 ms with 1 ms resolution and a default value of 15 ms. 8 copies of the message are sent with the fast period. The slow topology discovery message period is configurable from 50 ms to 10 seconds with 50 ms resolution and a default value of 100 ms. Each time a topology discovery message is triggered by a change in the information contained in the message or upon initialization, the fast topology discovery message period is used, followed by the slow topology discovery message period until a new message is triggered from the station.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

1.5.1.2 Effect of receipt

The receipt of this message on the same ringlet as which it was sent (as indicated in the ringletID field) from any station causes the MAC Control sublayer to update its current local topology image.

The receipt of this message on the same ringlet as which it was sent (as indicated in the ringletID field) from a neighbor station (as determined by $TTL == MAX_STATIONS$) causes the MAC Control sublayer to validate and (if needed) update the identity of its neighbor.

The receipt of this message on a ringlet other than that on which it was sent (as indicated in the ringletID field) from a neighbor station (as determined by $TTL == MAX_STATIONS$) causes the MAC Control sublayer to discard the message, place the link upon which the message was received into a non-operational state, and generate a miscabling error.

1.5.1.3 Handling of topology discovery messages during protection

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

The implications of the rule in this clause that a wrapping station shall strip topology discovery messages upon receipt must be considered, given that no comparable rule is in effect for steering stations (see the topology and protection scenario annex). This may result in incomplete physical topologies being reported by some stations in wrapping rings in some double failure cases. However, it is guaranteed that all stations will see the same data reachable topology, and that these data reachable topology will be consistent with that seen in steering rings.

A station in wrapped protection state shall not wrap a topology discovery message, and shall strip it after receiving it. The wrap eligible (WE) bit in the RPR header shall be set to zero for topology discovery messages to ensure that they are not wrapped. Topology discovery messages shall continue to be delivered and received on links that are in non-idle protection states.

1.5.2 Topology extended status message handling

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

The editors will add this in pending receipt of original contribution from Leon Bruckman. It is suggested that the topology extended status message follow the same triggering rules as topology discovery messages, except that there be no fast message period, only a slow message period. The range for the period may also be slower than that of the slow message period for topology discovery messages.

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

The PAH is discussing the potential inclusion of a paragraph describing receive buffer size. This paragraph would read as follows:

To ensure rapid and bounded completion of topology discovery, each MAC control queue shall be capable of storing one topology discovery message from each of the possible sources on the ring; the minimal storage for one source shall not be compromised by the offered load from other stations.

The implication of including a statement like this in this clause is that similar statements would need to be included in other clauses defining MAC control protocols.

The above statement would not apply to topology extended status messages.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54