

Improving Bandwidth Efficiency When Bridging on RPR

November 2001

Marc Holness, Nortel Networks


Anoop Ghanwani, Lantern Communications

Raj Sharma, Luminous

Robin Olsson, Vitesse

CP Fu, NEC

Components of a Complete RPR Proposal



Objectives

- Improve efficiency of BW utilization on an 802.17 Ring when Bridging interaction
- No impact to 802.17 MAC
- Minimal impact to existing 802.17 Frame Structure

Outline

- Problem Definition
- Solution Characterization
- Solution Architecture
- Solution Comparison
- Conclusion
- Proposal Details

802.17 Spatial Re-Use


- Spatial re-use is achieved on the RPR due to destination station stripping of packets
- 802.17 MAC reception rules currently strip a packet off the Ring using the following rules:
 - 802.17 MAC address (source and/or destination)
 - TTL
 - RPR frame header error checking
 - RPR control message identification
 - Combinations of the above

802.17 MAC Reception Rules

Implications

- Numerous RPR network scenarios result in the MAC address of the packet traversing the Ring, not belonging to RPR address space
 - Examples include frames supplied by Transparent Bridges
 - In these cases, the 802.17 MAC is essentially acting as a Proxy for the off Ring MACs
- The 802.17 MAC currently *flood* frames around the Ring, when the MAC DA is unknown unicast
- Such scenarios, result in continuous flooding, which leads to poor BW utilization of the RPR

Example #1: WITHOUT BW Efficiency MAC Extensions



Assume R1, R2, R3, R4 are Routers on the RPR. Sx and Sy are 802.17 Relay Stations.


- a) R1 sends ARP Request to R3
- b) 802.17 MAC (@ R1) broadcast packet around Ring.
- c) R2 sees packet, but does not process.
- d) Station x (Sx) sees packet and relays it out 802.3 port.
- e) 802.1D Transparent Bridge forwards packet.
- f) 802.3 ingress port on Sy receives packet.
- g) 802.17 MAC (@ Sy) broadcasts packet around Ring.
- h) R4 sees packet and discards.
- i) R3 sees packet and processes.
- j) R3 send ARP Reply back to R1. R1's MAC address is known.
- k) **802.17 MAC (@R3) floods packet around Ring, since unknown unicast.**
- l) R4 sees packet, but does not process.
- m) Sy sees packet and relays out 802.3 port.

- n) 802.1D Bridge forwards packet.
- o) 802.3 ingress port on Sx receives packet.
- p) 802.17 MAC (@ Sx) forwards unicast packet to R1.
- q) R1 receives packet and processes.

All subsequent communication from R1 to R3, results in the packet being flooded over Ring A. **Sustained poor BW utilization over Ring A is realized.**

All subsequent communication from R3 to R1, results in the packet being flooded over Ring B. **Sustained poor BW utilization over Ring B is realized.**

Example #2: WITHOUT BW Efficiency MAC Extensions



All Stations on both Rings are 802.1D/Q compliant 802.17 Bridges.
Assume all Bridge FDBs have learnt all appropriate Host MAC addresses.

- Host A sends packet destined to Host B.
- Station S1 receives packet and Bridges on Ring.
- 802.17 MAC (@ S1) floods packet over Ring A, since unknown unicast DA.**
- S2 sees packet and does not forward out.
- Sx see packet and Bridges onto Ring B.
- 802.17 MAC (@ Sx) floods packet over Ring B, since unknown unicast DA.**
- S4 sees packet and does not forward out.
- S3 sees packet and Bridges packet out 802.3 egress port.
- Host B receives packet and processes it.

All communication between Host A and Host B results in the packets being flooded over Ring A and Ring B. **Sustained poor BW utilization over the RPR Rings is realized.**

Solution Characterization

A generalized mechanism is required to identify the RPR station acting as the Proxy (specifically when the DA is not in the domain of the RPR Station MAC address space)

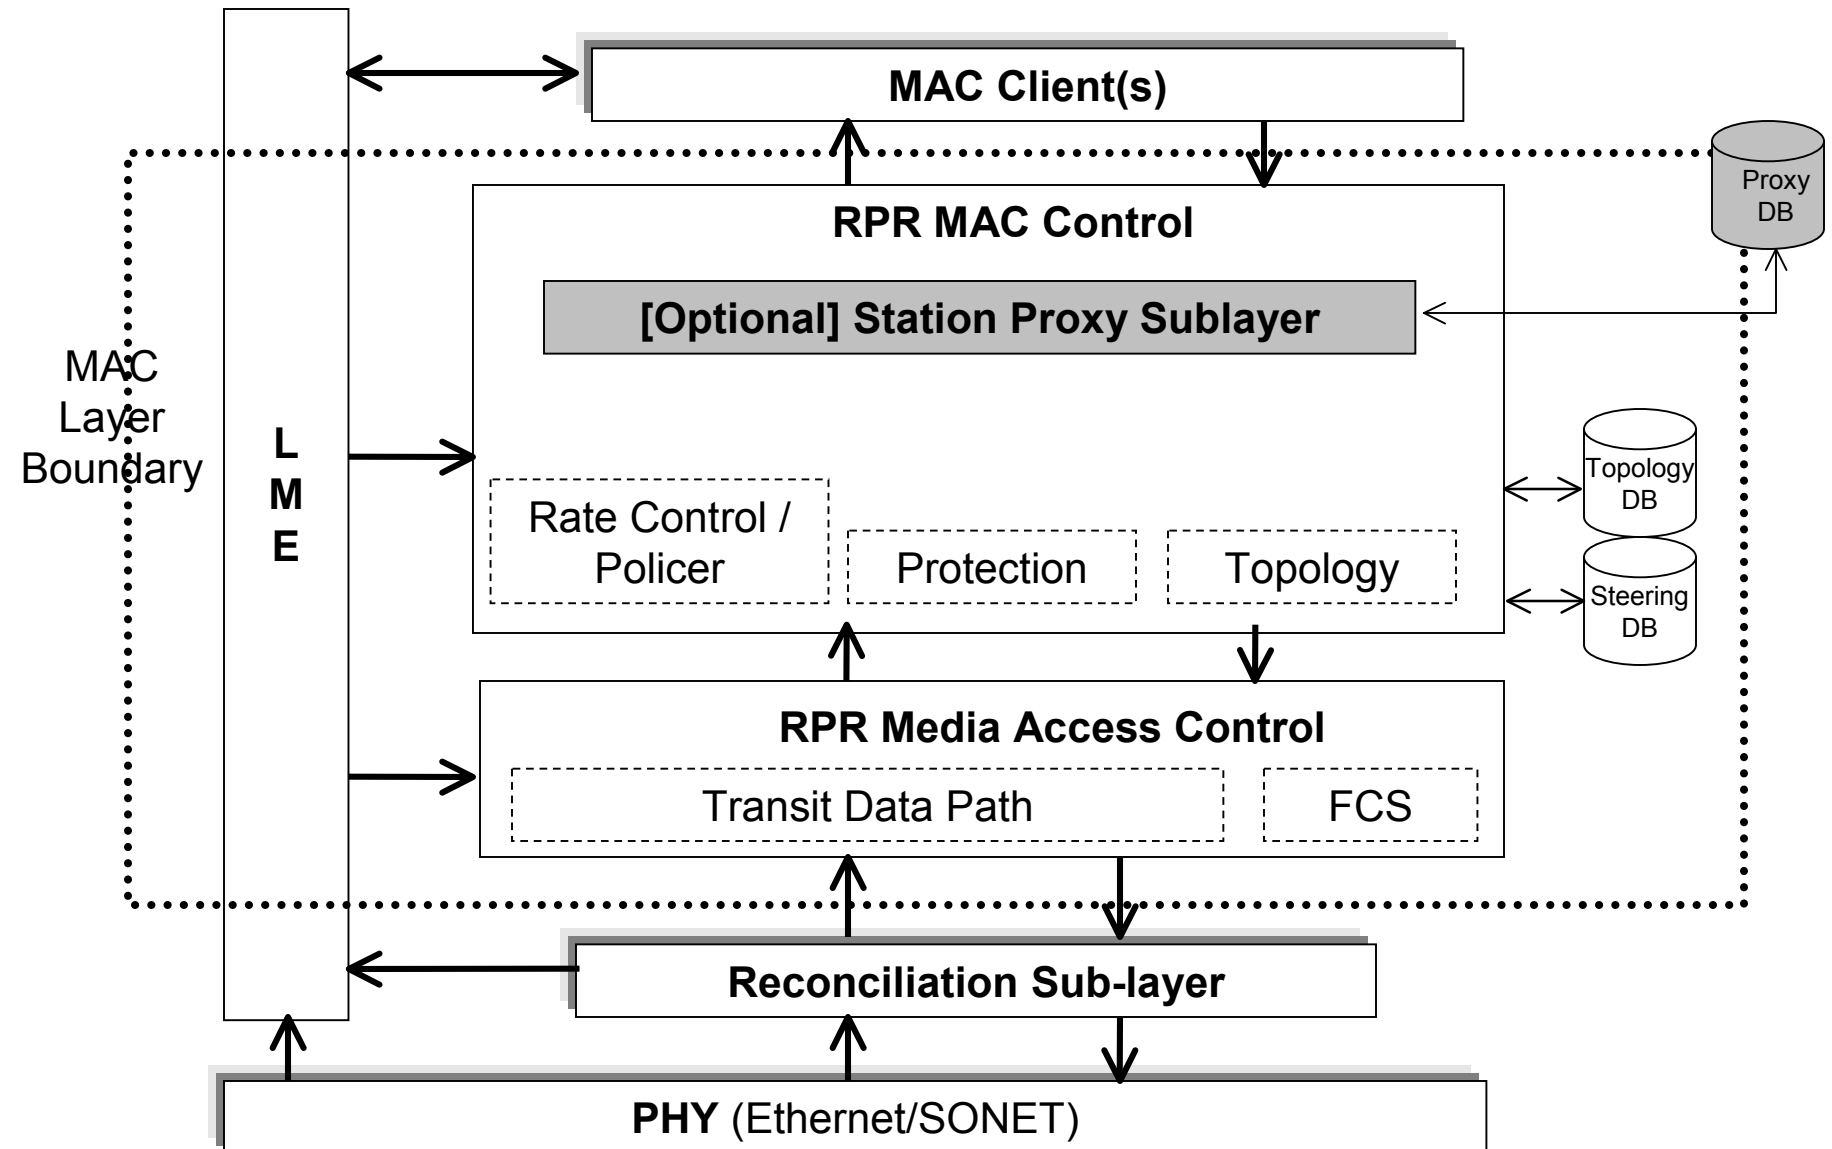
The solutions to address this problem can be characterized as follows:

1. **Explicit Value**: These mechanisms explicitly identify the station on the RPR within the RPR frame structure
 - a) Station MAC Address identification. Station MAC address encapsulation is used for frames on the RPR
 - b) Station Identifiers. Station identifiers are embedded in the RPR frame traversing the RPR
2. **Explicit Reference**: These mechanism provide a *pointer* to other stations on the RPR. The RPR frame structure passes a pointer to the receiving Station
3. **Implicit Station Proxy**: Each RPR Station snoops the DA address on the Ring, and with the use of a Proxy table, determines if it is acting as a Proxy for that DA


Architecture of Proposed Solution

- Introduce an **optional** 802.17 MAC Control sublayer, Station Proxy Sublayer (SPS), to provide MAC extensions to improve BW efficiency
 - SPS will not impact the operation of the RPR MAC
 - SPS can be enabled or disabled via LME
 - SPS will associate (off Ring) MAC addresses with the (source) Station in a Proxy DB. This gets done whenever a RPR Frame is sent from 802.17 MAC to a MAC Client
 - SPS will encapsulate/embed Station identifier (e.g., MAC address or Station pointer) on all frames sent by MAC client to 802.17 MAC
- Capitalize on existing 802.17 MAC packet removal rules to achieve BW efficiency. No impact to MAC transit path

802.17 MAC Components



Example #1: WITH Station Proxy Sublayer




- a) R1 sends ARP Request to R3
- b) 802.17 MAC (@ R1) broadcast packet around Ring.
- c) R2 sees packet, but does not process.
- d) Station x (Sx) sees packet and relays it out 802.3 port.
- e) 802.1D Transparent Bridge forwards packet.
- f) 802.3 ingress port on Sy receives packet.
- g) 802.17 MAC (@ Sy) broadcasts packet around Ring.
- h) R4 sees packet and discards.
- i) R3 sees packet and processes.
- j) R3 send ARP Reply back to R1. R1's MAC address is known.
- k) 802.17 MAC (@R3) forwards packet to Sy.**
- l) R4 sees packet, but does not process.
- m) Sy sees packet and relays out 802.3 port.
- n) 802.1D Bridge forwards packet.
- o) 802.3 ingress port on Sx receives packet.
- p) 802.17 MAC (@ Sx) forwards unicast packet to R1.
- q) R1 receives packet and processes.

All subsequent communication from R1 to R3 results in the packets being directly forward to Sx on Ring A. **BW efficiency is maintained over Ring A.**

All communication from R3 to R1 results in the packets being directly forwarded to Sy on Ring B. **BW efficiency is maintained over Ring B.**

Example #2: WITH Station Proxy Sublayer



All Stations on both Rings are 802.1D/Q compliant 802.17 Bridges. Assume all Bridge FDBs and all 802.17 StationID DBs have learnt all appropriate Host MAC addresses.

- Host A sends packet destined to Host B.
- Station S1 receives packet and Bridges on Ring.
- 802.17 MAC (@ S1) directly forwards packet to Sx.**
- Sx see packet and Bridges onto Ring B.
- 802.17 MAC (@ Sx) directly forwards packet to S3.**
- S3 sees packet and Bridges packet out 802.3 egress port.
- Host B receives packet and processes it.

All communication between Host A and Host B results in the packets being forwarded between S1 and Sx on Ring A, and between S3 and Sx, on Ring B. **BW efficiency is maintained over both RPR Rings.**

Failure Scenario Handling

- If a Ring topology change (either due to Ring span failure, or manual intervention) is detected by the Station Proxy Sublayer, the Station Proxy DB will be cleared
- RPR topology changes are determined by the Station Proxy Sublayer when a CONTROL Indication primitive is received containing a *Protection Event* opcode
- RPR topology changes may occur as a result of Ring span failures (e.g., Ring segment failure, loss of signal, signal degradation, etc.), or manual intervention (e.g., force switch or manual switch operations)

Solution Comparisons

Explicit Value		Explicit Reference	Implicit
Station MAC Encapsulation	Station Ids	Station Pointers (e.g., Hop Count)	Station Proxy DB Look-up on Transit Path
<ul style="list-style-type: none"> ✗ At least, 12 octet frame tax on RPR frames around Ring. ± Depending on design, may not be able to claim 802.1D/Q compliance. ± Allows source Station packet stripping (during exception handling). 	<ul style="list-style-type: none"> ✗ 2 octet frame tax. ✗ Distribution mechanism (protocol) of Station IDs across RPR. ✗ Ensuring uniqueness of Station Ids, within Station ID space, adds complexity to the distribution algorithm. ✗ Some sort of RPR Audit may be required to detect duplication of Station Ids. ✗ Changes to the (existing) 802.17 MAC reception rules. Specifically, stripping packets based on Station ID is now required on Transit path. ✗ Potential changes to the Steering DB. May need to index into this DB now using Station ID to determine Ringlet selection, TTL value setting, etc. ✗ Potential changes to the RPR Topology Discovery protocol to incorporate Station ID values. ± Allows source Station packet stripping (during exception handling). ☑ Does not impact conformance to 802.1D/Q during Bridging. 	<ul style="list-style-type: none"> ✗ 1 octet frame tax. (Perhaps can be overlayed in existing 802.17 Frame field.) ± Imposes the use of TTL <i>scoping</i> when transmitting 802.17 packets. ☑ Does not impact conformance to 802.1D/Q during Bridging. 	<ul style="list-style-type: none"> ✗ RPR MAC Transit path needs to update and check Proxy DB for all frames on transit. ✗ RPR MAC reception rules would need to access Proxy DB to determine whether packet should be stripped or not. ✗ Solution does not scale, since need to access (large) DB at line rate.

Conclusions

- RPR network scenarios resulting in the MAC address of the packet traversing the Ring, not belonging to the local RPR MAC address space, can still benefit from efficient BW utilization over RPR
- Bridging interactions with RPR, topologies supporting RPR Router clients, etc can maintain improved BW efficiency over the Ring when the 802.17 MAC Station Proxy Sublayer is enabled

Back Up Charts

Proposal #1: Summary

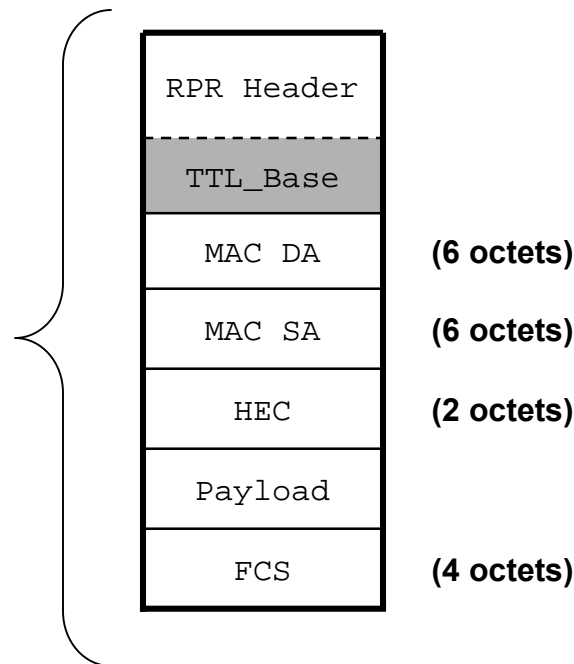
This proposal explores an Explicit Reference model to address the problem

- The Station Proxy Sublayer determines the station source of the received packet via a hop count (TTL based) pointer
- Capitalize on existing 802.17 MAC packet TTL removal rules to achieve BW efficiency

Proposal #1: RPR Frame Structure

- Utilize 1 octet in the RPR Frame Format. The field is the TTL_Base
- The Station Proxy Sublayer will set the TTL_Base field in the RPR frames. The TTL_Base field never gets decremented or altered

Example 802.17
Frame Structure



Proposal #1: RPR MAC Station Proxy Database

- The source Station is identified by keeping track of the hop count (I.e., number of Stations away and direction). The hop count is derived from $TTL_Base - TTL$
- The Station Proxy sublayer updates the Station Proxy DB each time a frame is passed to a MAC Client

802.17 Station
Proxy DB

MAC	Proxy Station	
	Hop Count	Ringlet
MAC _x	3	E
MAC _y	2	W
MAC _z	1	E

Proposal #1: RPR MAC Frame Transmission Extensions

802.17 MAC Frame Transmission procedures get extended to include the following general sequence:

1. The UNITDATA Request primitive (and corresponding parameters) supplied by the MAC Client is passed to the 802.17 MAC Station Proxy sub-layer
2. If the DA is determined to be an *"unknown Ring unicast address"*, the Station Proxy DB is indexed using the DA (supplied by the Request primitive)
 - If the DA is found in the Station Proxy DB, the TTL and TTL_Base gets set based upon the corresponding entry hop count, and the appropriate direction/ringlet is selected
3. If the DA is a known unicast DA, multicast DA, or broadcast DA, the TTL_Base value is set to the same value as the TTL, based on standard 802.17 MAC procedures

Proposal #1: RPR MAC Frame Reception Extensions

802.17 MAC Frame Reception procedures get extended to include the following general sequence:

1. The UNITDATA Indication primitive (and corresponding parameters) gets passed to the MAC Client, from the 802.17 MAC Entity
2. Before the frame gets sent to the MAC Client, the Station Proxy Sublayer will update the Station Proxy DB with the (*“off Ring”*) source MAC, computed hop count ($TTL_Base - TTL$), and Ringlet

Proposal #1: Failure Scenario Handling

- The Station Proxy Sublayer will discard (I.e., not send to MAC Client), any packet where the pointer ($TTL_Base - TTL$) exceeds the number of Stations on the Ring

Proposal #1: Cost-Benefit Analysis

- ☑ RPR network scenarios can utilize the BW efficiency capabilities of the RPR. Example scenarios include Bridging interactions with RPR, etc.
- ☑ Does not break 802.17 Bridging conformance to 802.1D/Q
- ☑ The Station Proxy DB is not accessed for Tandem traffic around the Ring. Consequently, there is no throughput impact to tandemed traffic
- ☑ The BW efficiency feature provided by the 802.17 Station Proxy Sublayer is optional, and can be enabled or disabled without impact 802.17 MAC operation
- ☑ Can be used to simplify the identification of Source Stations during monitoring of traffic around the RPR
- ✗ Each entry in the Station Proxy DB is a MAC address, Station Reference (hop count, ringlet) combination. This is ~7 octets (6 +1). Worse case, the Station Proxy DB needs to support the total number of MAC address transmitted or received by the Station. The number of entries supported can be Vendor supplied. Assume the default number of entries in the Station Proxy DB is 2K.

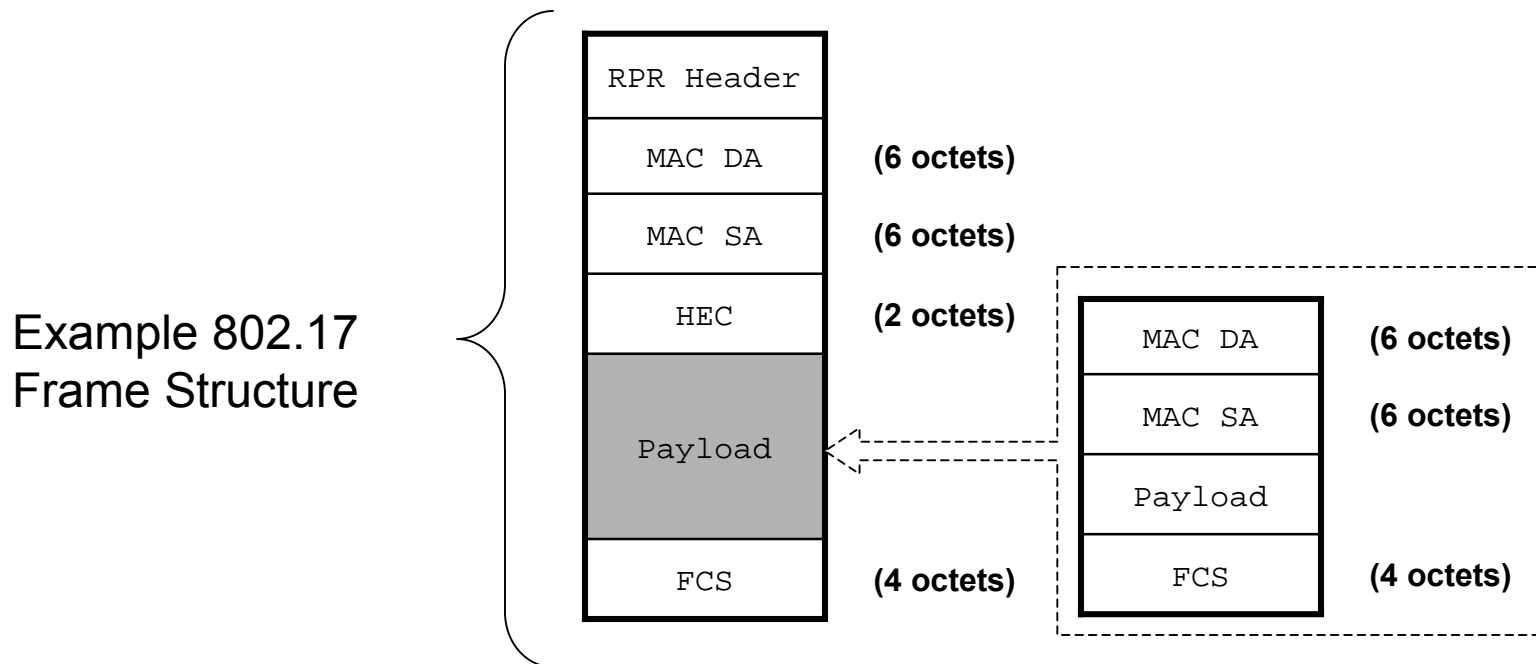
Proposal #2: Summary

This proposal explores an Explicit Value: Station MAC Encapsulation model to address the problem

- Encapsulate off Ring (originated or destined) frames with Station MAC addresses
- Capitalize on existing 802.17 MAC packet DA removal rules to achieve BW efficiency

Proposal #2: RPR Frame Structure

- Frame provided by MAC Client will be embedded in the 802.17 Frame payload
- The Station Proxy Sublayer will encapsulate the frame with RPR Station addresses



Proposal #2: RPR MAC Station Proxy Database

- The Station Proxy Sublayer updates the Station Proxy DB each time a frame is passed to a MAC Client
- The source Station is identified by the MAC SA of the received frame

802.17 Station
Proxy DB

Off Ring MAC	Station MAC
MAC _x	RPR MAC1
MAC _y	RPR MAC2
MAC _z	RPR MAC3

Proposal #2: RPR MAC Frame Transmission Extensions

802.17 MAC Frame Transmission procedures get extended to include the following general sequence:

1. The UNITDATA Request primitive (and corresponding parameters) supplied by the MAC Client is passed to the 802.17 MAC Station Proxy sublayer
2. If the DA/SA is determined to be an *"unknown Ring unicast address"*, the Station Proxy DB is indexed using the DA/SA (supplied by the Request primitive)
 - A new frame structure is created
 - The entire frame requested by the MAC Client is placed in the payload
 - A bit in the RPR Header (or PT) is set to indicate RPR_encap
 - If DA is off Ring MAC, index DA into the Station Proxy DB. The returned Station MAC is placed in the DA of the new frame. Otherwise, #FF is placed in DA
 - If DA is on Ring MAC, then use DA in encapsulated DA
 - The MAC address of this Station is placed in the SA

Proposal #2: RPR MAC Frame Reception Extensions

802.17 MAC Frame Reception procedures get extended to include the following general sequence:

1. The UNITDATA Indication primitive (and corresponding parameters) gets passed to the MAC Client, from the 802.17 MAC Entity
2. Before the frame gets sent to the MAC Client, the Station Proxy Sublayer will update the Station Proxy DB with the SA (RPR Station address space) of the frame, and the SA (e.g., off Ring address space) found in the payload, if the RPR Header bit (or PT) indicates `RPR_encap`

Proposal #2: Cost-Benefit Analysis

- ☑ RPR network scenarios can utilize the BW efficiency capabilities of the RPR. Example scenarios include Bridging interactions with RPR, etc
- ☑ The Station Proxy DB is not accessed for Tandem traffic around the Ring. Consequently, there is no throughput impact to tandemed traffic
- ☑ The BW efficiency feature provided by the 802.17 Station Proxy Sublayer is optional, and can be enabled or disabled without impact 802.17 MAC operation
- ☑ Can be used to simplify the identification of Source Stations during monitoring of traffic around the RPR

- ✗ Breaks 802.17 MAC conformance to 802.1D/Q
- ✗ Each entry in the Station Proxy DB is a MAC address, Station Reference (RPR MAC address) combination. This is 12 octets. Worse case, the Station Proxy DB needs to support the total number of MAC address transmitted or received by the Station. The number of entries supported can be Vendor supplied. Assume the default number of entries in the Station Proxy DB is 2K