802.17 MAC and Link Layer Encapsulation

March 2002

Marc Holness, Nortel Networks

Robert Castellano

Outline

- Issue Definition
- Solution Overview
- Benefits of Utilizing StationID Encapsulation
- Costs of Utilizing StationID Encapsulation
- Network Scenario Analysis
- Observations
- Conclusions

THE ISSUE

- How does the 802.17 MAC handle the transmission of frames where the destination address (and/or source address) is **not** in the domain of the RPR local station address space?

Example: RPR HostA send packet destined to RPR HostB. RPR Station S2 acting as a Relay.

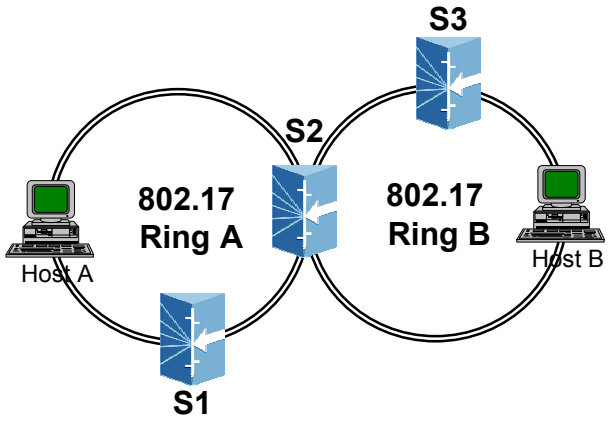

This situation occurs during Bridging interactions. For example, Host A dispatching packets destined to Host B, will result in packets being consistently *flooded* over Ring A.

Link Layer Encapsulation Solution Overview

- MAC implements DB that associates (off Ring) MAC, VLAN, and potentially other information (e.g., port) with Station identifier (StationID)
  - Solution needs to support 802.17 Host (e.g., Router, Server) and Bridge Clients
- DB gets updated when MAC Client dispatches packet to the Ring, and when packet gets stripped from Ring
- MAC Client sends *request* primitive to MAC, results in
  - MAC indexes into DB (I.e., StationDB) using SA and DA found in *request* primitive
  - Retrieves StationIDs associated with SA and DA
  - Encapsulate StationIDs associated with SA and DA in the packet to be dispatched
- MAC strips packet based upon StationID

Station Identification Encapsulation Benefits

- Facilitates increased BW utilization when transporting RPR packets where the MAC address space is outside the domain of the Ring local
  - Packets no longer need to be persistently flooded over the Ring during Bridging interactions
  - Once StationID association to MAC is made, Bridged packets can be sent unicast to (proxy) RPR Station
  - Station ID can be MAC address or label

Station Identification Considerations

- Changes to MAC reception rules to include stripping of packets based on station ID
- Protocol (or Protocol extension) potentially required to uniquely distribute and identify the StationID
- StationDB entries need to be updated upon Ring Topology changes
- DB (StationDB) required to associate station ID with the (off Ring) MAC addresses and VLAN (and potentially port)
  - Size of DB is an implementation concern
  - Accessing the DB can not impact achieving line rate by the MAC

Station Identification Considerations

- 802.17 configurations may exist, that results in persist “black holing” of packets being sent through the Network
- Synchronization of the StationDB and the Bridging (802.1D/Q) FDB needs to be done
  - Especially true for Spanning Tree interactions when Topology changes

Network Scenario #1

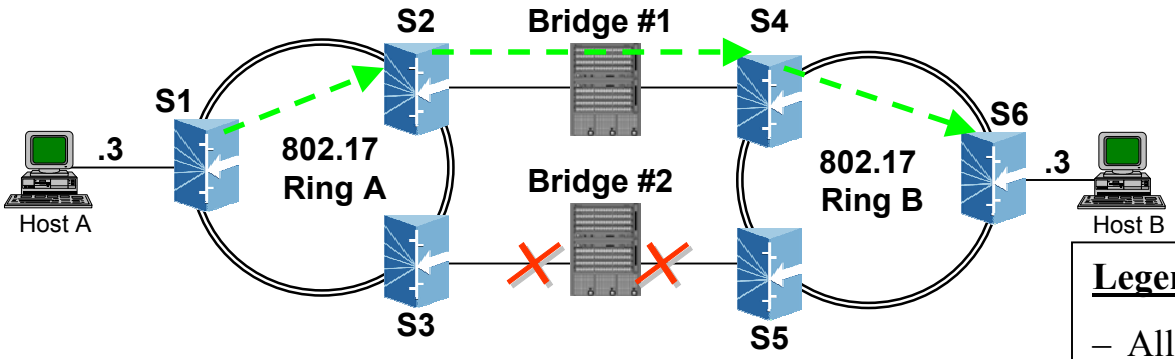

Legend

– All RPR Stations (S1-S6) are 802.1D/Q compliant devices

1. Bridge #1 and Bridge #2 running Spanning Tree. Network topology change
2. STP invoked. Links between S2-S4 move from "*forwarding*" to "*blocking*" state
3. Link between S3-S5 moved from "*blocking*" to "*forwarding*" state
4. Traffic between HostA to HostB should not be adversely impacted. No persistent ***black hole*** should occur

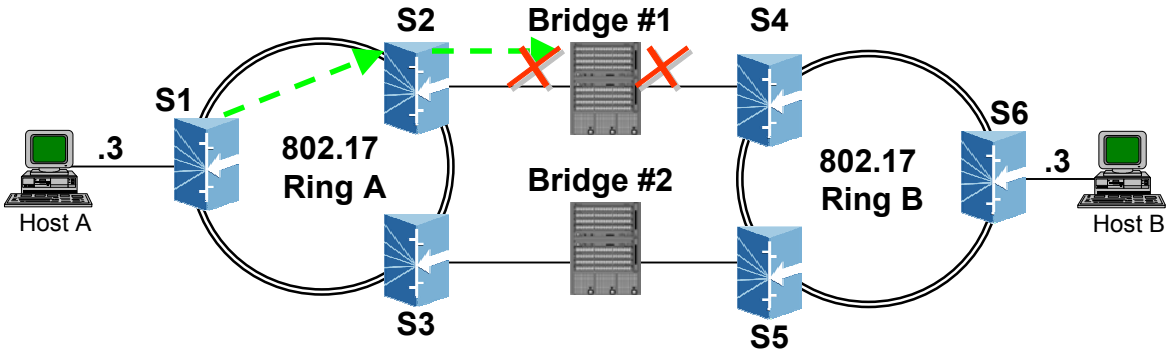

Network Scenario #1: High Level Walk-Thru

- Steady state: Host A sending traffic to Host B
  - StationDB at S1 associates MAC_HostB with station S2
  - At Ring A, S1 sends directly to S2
- Steady state: Host B sending traffic to Host A
  - StationDB at S6 associates MAC_HostA with station S4
  - At Ring B, S6 sends directly to S4
- Topology change occurs and STP converges
- Host A sending traffic to Host B
  - StationDB at S1 still associating MAC_HostB with station S2
  - At Ring A, S1 sends directly to S2
  - Bridged link between S2 to S4 is now in *blocking* state
  - This traffic gets dropped by Bridge #1
  - Persistent ***black hole*** present!
- Similar persistent ***black hole*** exist for traffic from Host B to Host A at Ring B exist

Network Scenario #1: Analysis

1. To avoid/address the persistent black hole in this scenario, 802.17 MACs must essentially synchronize the appropriate StationDB entries with the Bridging FDBs (that change due to STP)
   - Options considered:
     a) Remove entry in StationDB at S1 that associates MAC HostB with S2. The trigger for this action is based on reception (and snooping of) BPDUs sent during STP
     b) Timeout entry in StationDB at S1 that associates MAC HostB with S2. The trigger for this action is based on reception (and snooping of) BPDUs sent during STP

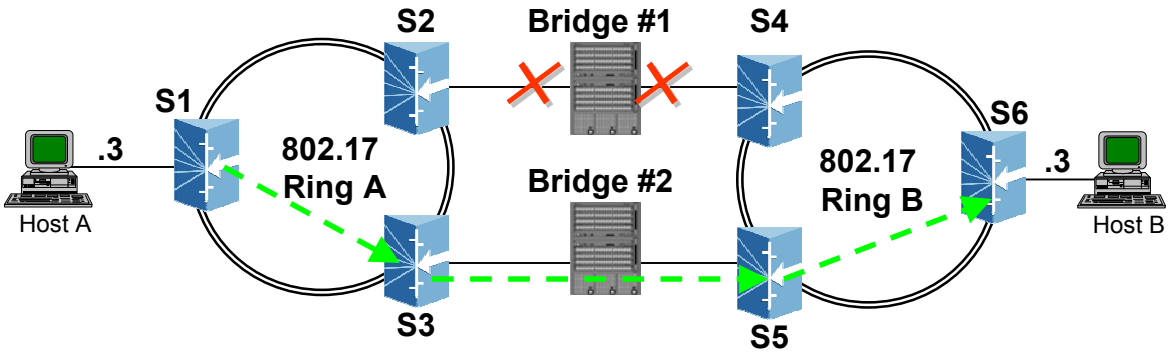

Network Scenario #1: Analysis

2. To avoid/address the persistent black hole in this scenario, 802.17 MACs must periodically timeout entries in the StationDB. The timeout of the entries is independent of any Network Topology changes.

   - Options considered:

     a) Let entries in StationDB at RPR Stations that associates MAC Host address with RPR local stations timeout. This would revert the operations of the S1 station to flooding non local destination addresses around the Ring. The flooding would cease once the association between HostB MAC and S3 is re-learnt.

        – The engineering of the Network is complicated due to StationDB entry timers expiries, resulting in Flooding heart beat.

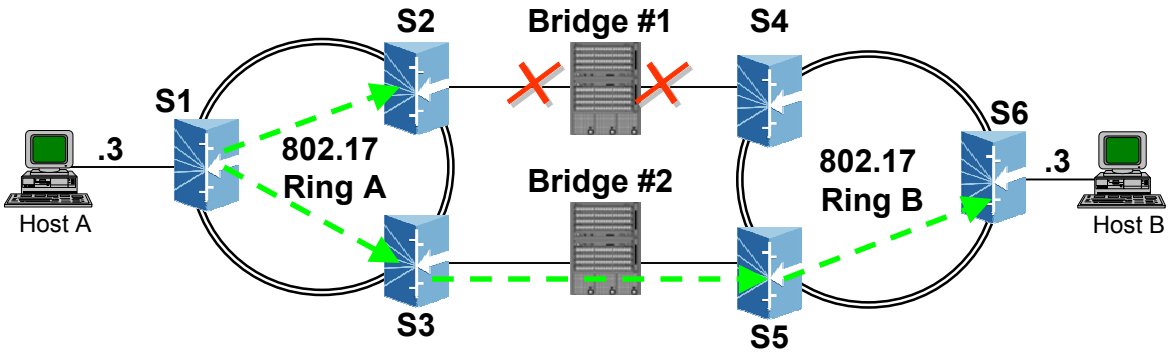

Network Scenario #1: Solution Impact

- Overall StationDB management essentially involves embedding Bridging like functionality within the 802.17 MAC
- StationDB needs to be managed by the 802.17 MAC
  - Each entry requires an aging timer
  - Each entry needs to have a MAC address by VLAN (or Customer Identification) key
  - Size of DB needs to be engineered (on MAC)
- If Network Topology changes trigger the synchronization of StationDB entries, then 802.17 MAC needs to snoop BPDUs and interpret contents
  - Applicable to RPR Router clients, RPR Host clients, as well as RPR Bridging clients.

Network Scenario #2

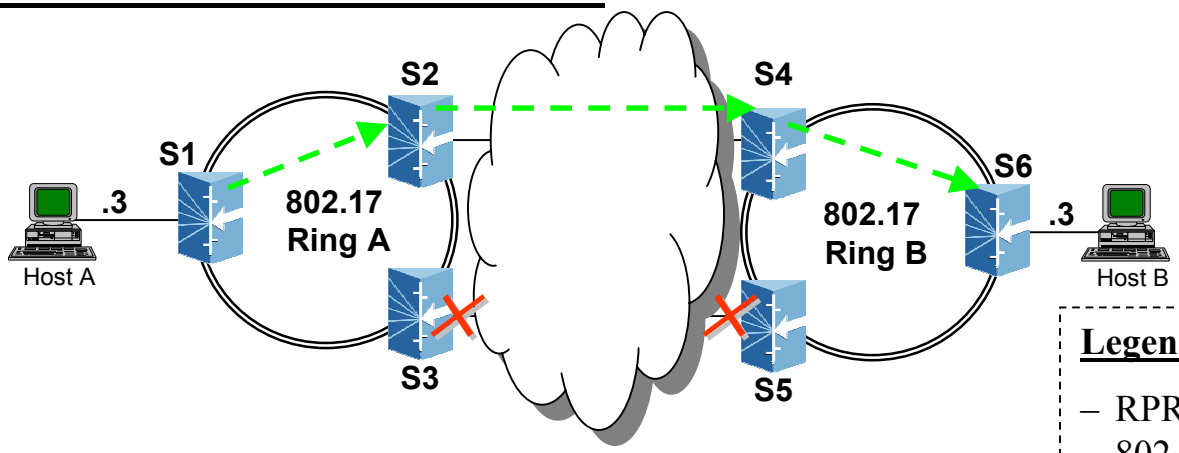

Legend

- RPR Stations (S1 and S6) are 802.1D/Q compliant devices
- RPR Stations (S2 to S5) may run MAC Client protocol to address loops/redundancy within the network

1. Network topology (in cloud) change
2. Link between S2-S4 move from "*active*" to "*stand-by*"
3. Link between S3-S5 moved from "*stand-by*" to "*active*"
4. Traffic between HostA to HostB should not be adversely impacted. No persistent ***black hole*** should occur

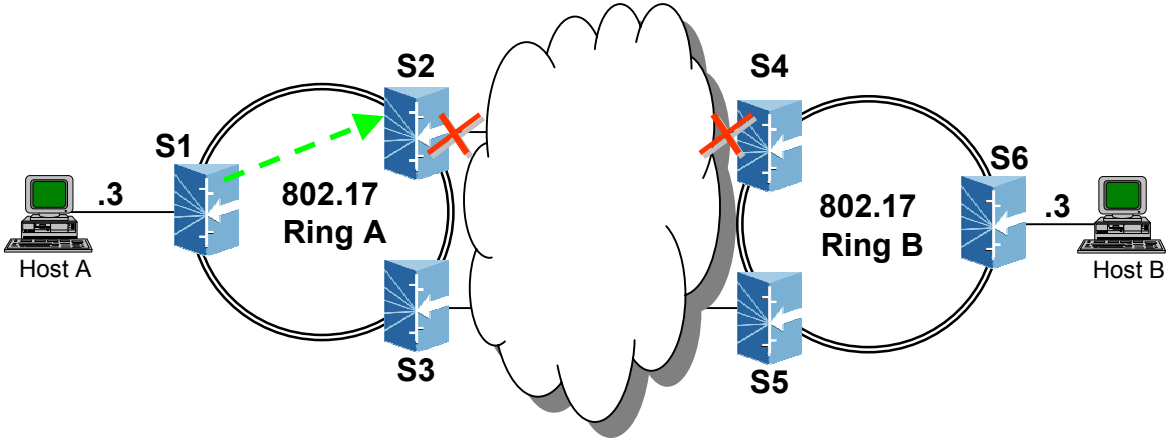

Network Scenario #2: Impacts

- Similar Analysis and Solution Impact as described in Network Scenario #1
- 802.17 MAC needs to snoop Client specific generated frame indicating topology change to synchronize the StationDB entries
  - Applicable to RPR Router clients, RPR Host clients, as well as RPR Bridging clients
  - Not something we want the MAC to do. This snooping is Client specific

Observations

- Synchronization of the StationID DB with Network Topology changes is required (to avoid persistent "black holing" of frames)
  - MAC Client specific logic (e.g., Spanning Tree Topology Change BPDU recognition) should not be embedded in the 802.17 MAC
  - The 802.17 MAC should support other MAC Client protocols that manage Network Topology changes. Consequently, 802.17 MAC would need to be aware of other MAC Client protocol Topology Change notifications. This should not be embedded in the 802.17 MAC
- Supporting the management of the StationDB to achieve BW efficiency over the Ring, results in a complex 802.17 MAC
  - The 802.17 MAC essentially needs to implement a type of Encapsulation Bridge

Observations

- Clause 5.5 (Conformance) of IEEE 802.1D Standard states:

  ".. MAC-specific bridging methods may exist. Use of a MAC-specific bridging method and the method specified in this standard on the same LAN shall

  a) Not prevent communication between stations in a Bridged LAN.

  b) Preserve the MAC Service.

  c) Preserve the characteristics of each bridging method within its own domain.

  d) Provide for the ability of both bridging techniques to coexist simultaneously on a LAN without adverse interaction*… "*

  – Adverse interactions would occur when 802.17 MAC specific StationID (encapsulation) Bridging interacts with 802.1D specified Bridging.

Conclusions

- Link Layer Encapsulation should not be supported in the 802.17 MAC
- If support considered:
  - Place functionality in an 802.17 MAC Client
  - The functionality should be made optional
  - Need to interwork with other local RPR stations that do not support this functionality

Next Steps

- Ensure 802.17 frame structure that will not preclude MAC Clients that want to support the BW Efficiency feature (via StationID)
- Specify the flooding mechanisms that the 802.17 MAC can use