Resilient Rings

Addressing markets
that are not good fits for RPR

John Lemon

Contents

- What Limits RPR's Success
- What Is Needed For Mass Success
- Technical Approach
- Technical Issues
- Technical Recommendation
- Forum Issues
- Forum Recommendation
- Marketing Issues
- Marketing Recommendation
- Open Issues

What Limits RPR's Success

- RPR is not Ethernet
  - RPR doesn't have the Ethernet brand
  - RPR's advantages are not often compelling enough improvements where Ethernet is mostly 'good enough'
- Cost
  - RPR has large development costs
  - RPR has large material costs
- Complexity
  - RPR has more complexity than is often needed
  - RPR has more complexity than is often used
  - RPR has more complexity than often can be understood
  - Closely tied to cost

RPR Is Not Ethernet

- RPR doesn't have the Ethernet brand
  - The Ethernet brand is universally recognized and valued
  - Every new layer 2 technology wants to be known as Ethernet
    - 802.3ah has little more than the frame format in common with previous versions of Ethernet
    - WiFi is often referred to as "wireless Ethernet"
    - WiMax is often referred to as "wireless metro Ethernet"
- RPR's advantages are not often compelling enough
  - Wireless (e.g., 802.11 or 802.16) is successful because of the uniqueness of the medium, not any special features
- Non-Ethernet solutions add costs
  - Additional development costs for new tools and training
  - Additional capex for new provisioning, management, and diagnostics tools
  - Additional opex for different management paradigm

Cost

- RPR has large development costs
  - RPR complexity reduces interest in implementing
  - RPR complexity reduces ability to implement (at least correctly)
  - Existing Ethernet chips, Ethernet cores, and embedded Ethernet MACs can't be used
- RPR has large material costs
  - RPR silicon solutions are scarce and expensive
    - Choice of one (relatively) expensive ASIC, turning Xilinx "solution" into something that works, or roll your own
    - Barriers to entry limit providers, which limits volume, which limits cost amortization in ASICs
  - RPR requires a custom solution, without use of off the shelf chips, cores, or embedded Ethernet MACs
    - Either new silicon or additional silicon is required

Complexity

- RPR has more complexity than is needed
  - Conservative fairness is not necessary
  - Multichoke fairness is unneeded
  - Spatially-aware fairness is not often needed
- RPR has more complexity than is used
  - Most of fairness is not fully utilized
  - Classes of services are rarely used as true classes
- RPR has more complexity than can be understood
  - Fairness and service classes are almost universally not (fully) understood by equipment providers, service providers, or customers
  - Most providers do not fully understand RPR's capabilities, features, or areas of best applicability
  - Much of what providers do understand, they are not able (or willing) to explain to their customers

What Is Needed For Mass Success

- Ethernet
  - The (most important) advantages, features, and capabilities must be made available via Ethernet
- Low cost
  - The economics of Ethernet must be leveraged
    - Additions beyond base Ethernet must be accomplished with very little or no additional hardware
  - There must be (reasonably) wide adoption
- Low complexity
  - Only what is needed (and explainable) should be added
  - Everything should be doable in software or at least with very minimal additional hardware

Technical Approach — Frame

- Shim added to Ethernet (or other 802) frame
  - Contains the essential portion of the RPR header
    - Contents, if any, of ring tag yet to be determined
  - Similar to VLAN shim
  - Includes / preceded by 802.1ah header/shim

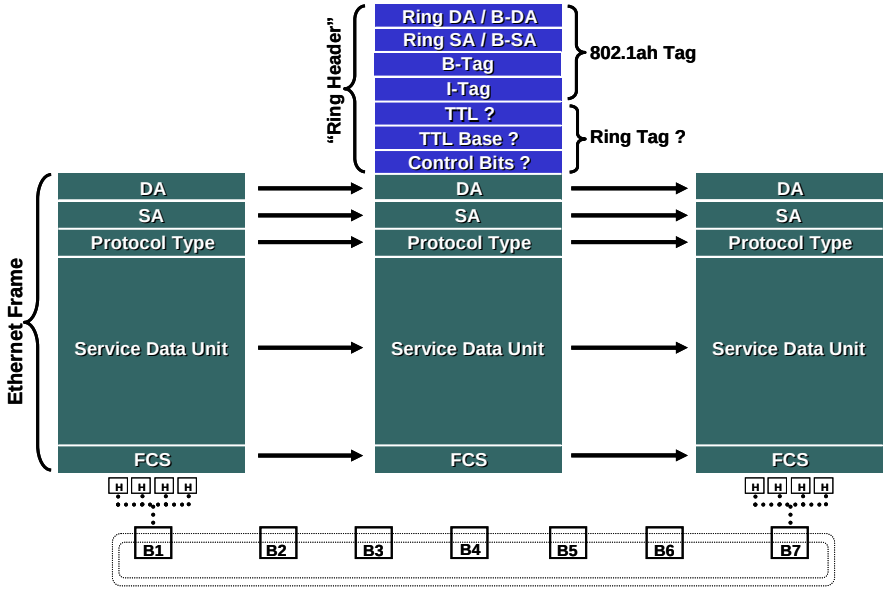

Technical Approach — Bridge Shim

- Following applies only if a ring tag is used
- Ring shim processed at each station as part of REISS (in other words, new bridge shim layer for this)
  - Implies each ring station has two MACs, instead of RPR's one MAC per station
    - Would make topology discovery more complicated, or
    - Would necessitate including peer MAC address in TP frame
  - Transit path becomes path between bridge ports
    - Allows for processing of frame between ring ingress and ring egress
  - Includes / preceded by 802.1ah header/shim

Technical Approach — Fairness

- Drop multichoke
- Drop conservative mode
- Run (or allow for running) at software speeds
- Fairness applies to any frame with DEI bit set

Technical Approach — QoS

- Replace service classes with 802.1Q priorities

Technical Issues

- Is a ring tag/header/shim needed
- Where to stick the header/shim
  - If the ring shim is added before other VLAN tags, then new/special hardware would be needed to parse this header in order to see VLAN tags
  - If the ring shim is added after other VLAN tags, then new/special hardware would be needed to parse through the VLAN tags to get to the ring shim, and this probably violates layering
- TTL
  - A TTL provides benefits, especially for rings
  - A TTL provides differentiation from service layer only approaches
  - A TTL in the header requires checking and decrementing the header field at every ring station
    - FCS can be modified instead of being recomputed
    - Simple hardware between two Ethernet MACs, if possible to stick something between them

Technical Recommendations

- Tag location
  - Use 802.1ah I-tag followed by possible ring tag, followed by 802.1Q VLAN tag
  - Assume that if this starts soon enough, and if we decided a ring tag is needed, that vendors adding support for 802.1ah will take this extra tag into account
- TTL
  - Use TTL at least on management frames
    - To be processed by software off the bridge relay path
  - No opinion yet on whether to use on data frames

Forum Issues

- Possibility 1: RPR Lite
  - A new RPR MAC standard would be created
  - It would probably be named 802.17.1
  - The work would take place entirely with the 802.17 WG
- Possibility 2: Ring bridge
  - 2a: It could be a new type of bridge, a la 802.1AP
  - 2b: Or it could be an addition to 802.1Q, a la 802.1ah
  - It would have an 802.1<something> name
  - The work would probably take place partially within the 802.17 WG and partially within the 802.1 WG
    - The 802.17 WG would (hopefully) take the lead role

RPR Lite

- Advantages:
  - The new MAC would be designed and controlled entirely by the 802.17 WG
  - The new MAC would clearly be something unique
- Disadvantages:
  - Involvement by only the 802.17 WG greatly reduces the industry exposure and involvement needed for wide adoption
  - Being unique is good only when clearly better
  - There is perhaps some marketing baggage for 802.17 because of its lack of wide success
  - Anything bad about 802.17 would be assumed to exist for 802.17.1

Ring Bridge

- Advantages:
  - Standardization by the 802.1 WG greatly increases the industry exposure, hopefully leading to wide adoption
  - 802.1 is more widely recognized and perhaps has better valuation than 802.17
  - This would be a chance to have a fresh start, without any real or perceived baggage from 802.17
- Disadvantages:
  - The new standard would not be entirely designed and controlled by the 802.17 WG
    - Those who don't understand the benefits could mess up the result
    - Those who have competing interests could mess up the result

Forum Recommendations

- Ring bridge standard from 802.1
- Amendment to 802.1Q
  - Most likely no modification to bridge relay, only to higher level control and management
- 802.17 WG to do work for 802.1
  - Pitch as being similar to how 802.1AX is being created by 802.3
    - May be viewed as not qualifying for such independence
  - Fall back to 802.1 TG which automatically includes all 802.17 WG members
    - Would run in parallel to other 802.1 TGs

Marketing Issues

- RPR Lite
  - Probably be the worst name possible
  - Carries all baggage of RPR
  - Sounds like a weak subset of RPR
- 802.1az
  - Does not roll of the tongue or convey any meaning
- Resilient Ring Relay (RRR)
  - Doesn’t take advantage of Ethernet brand
- Resilient Ethernet Ring (RER)
  - Rides Ethernet brand popularity
  - Same approach as ITU’s “Protected Ethernet Ring”

Marketing Recommendation

- Refer to as
  Resilient Ethernet Ring (RER) - 802.1Qaz
  - Emphasis is on key attributes:
    - Resilient/protected
    - Ethernet
    - Ring-specific

Open Issues (1)

- How to roll this out in a manner least likely to garner resistance
- PHYs
  - Should this support SONET/SDH PHYs?
    - Not sure how to do this under auspices of 802.1
    - Not doing so would leave out substantial portion of the market
  - Any reason why this wouldn't automatically work with other 802 PHYs than 802.3?
- TTL
  - Can existing bridge chips recompute FCS when passing from one bridge port to another?
  - Can existing bridge chips do simple substitutions (to be used as a replacement for decrement of TTL)?

Open Issues (2)

- Fairness
  - How well will fairness work if it runs at microprocessor speeds (e.g., 10 ms)?
  - Does stripping fairness down to aggressive, single choke make it simple enough, or should it be further simplified?
  - Should each priority contain its own fairness domain?
- Topology and attribute discovery
  - Any reason to use anything other than existing topology discovery and attribute discovery messages?

Open Issues (3)

- Protection
  - If topology discovery messages are being used, any reason to use any other standard messages such as CC for break detection?
  - Should both steering and wrapping be supported?
- OAM
  - Keep Echo frame, or use LBM/LBR frames?
    - RPR Echo allows choice of paths. Not sure how this could be done with LB frames.
  - Keep Flush frame (but rename it), or what?
  - Keep SAS Notify frame?
  - Keep org-specific frame?