Response to 802.3 (ad-hoc) on RPRSG PAR / 5 Criteria

November 2000
Hyatt Regency Tampa

Key Features	RPR	Bridged Ethernet
Bounded Delay / Jitter on Ring	Y	Y
Low Latency	Y	Y
Support for BW Multiplication	Y	Y
Support for Shortest Path between any two nodes on the ring	Y	spanning tree limits
Can work with 1 fiber cut	Y	protocol requirement for duplex link
less than 50 ms restoration	Y	N
Dynamic BW Mgt (fairness / unfairness)	Y	N
Mechanism for a loss-less tx path for some data	Y	N
Avoids Downstream / Upstream unfairness per queue	Y	N
"Add Drop" MAC Layer Function	Y	N
Service Specific Protection	Y	N
End Station connected directly onto the Ring	Y	N
Different Header than Ethernet eg. TTL	Y	

The physical network topology of our target customers is primarily rings. Therefore it is important that our solution be optimized for that topology and for our particular customer requirements.

A key requirement is to maximize resiliency and network capacity by using the traffic handling capabilities in both directions. In order to minimize delay it is imperative that the shorted path can be chosen.

A ring of ethernet bridges does not close the ring due to 802.1D cutting the ring, thus violating the requirement of shortest path selection. The convergence of 802.1D violates the goal of 50 ms restoration which is in our PAR based on customer requirements.

802.1s is not scalable as a method of providing multiple spanning cuts (a method to allow each node to send in either direction by breaking the ring at the furthest point).

-Scalability. The number of spanning trees required for a dual ring is 4 per node, hence the total number of ST for a ring is 4N.

-Management of the VLANs is excessively complex. It lacks dynamic configuration which violates our goal of plug and play. A protocol that could configure these VLANs automatically would not be a simple extension to existing protocols.

-If VLANs are being used for traffic management they cannot be used for user segregation.

An Ethernet bridge operating as a ring node can be viewed as a 2+P port bridge at each hop. 2 ports are ring attachment ports, and the remaining ports are for ingress/egress. A key difference of optimizing for the ring case is that there is no bridge in the transit path on an RPR MAC.

Moreover, each Ethernet bridge in the Ethernet bridged ring implements a local congestion management mechanism, while RPR requires a distributed congestion management to provide end-to-end QoS.

Customer requirements have shown that the RPR ring is not just a collection of switching elements. It must also include content servers in order to eliminate layers of equipment; improve resiliency and reduce latency.

A host operating on a ring requires a MAC layer that understands how to pass through traffic. In the case where a ring of Ethernet bridges is used all, hosts must become bridges if they are to reside on the ring.