

If You Want a Solid Network... Put a Ring on It

The Redundant Path Technology of ERPS

By: David Zaveski, Field Applications Engineer

February 26, 2016

White Paper

If You Want a Solid Network...Put a Ring on It

ERPS (Ethernet Ring Protection Switching), sometimes referred to as G.8032 and defined in ITU-T G.8032, provides redundant path technology at the Ethernet layer with very low failover times. As stated in the name, the technology uses Ethernet switches to create a highly reliable and stable ring topology from which a robust network can be built. The technology permits the incorporation of QoS (Quality of Service) which allows network designers to comply with SLAs (Service Level Agreements) indicative of the service provider industry.

This technology provides a solution to common link failures found in networks everywhere, but it is not the only, first, or even best solution in some applications. RPR (Resilient Packet Ring), SONET (Synchronous Optical Network) and SDH (Synchronous Digital Hierarchy) are ring technologies that have been around for years and are loved for their robust abilities. These technologies, however, do not use Ethernet and with that comes a cost. The drive to use Ethernet, as seen dramatically in the carrier networks, is one of cost. Ethernet technology is used everywhere and that has driven down the cost of both hardware installation and the maintenance. Almost every network engineer understands Ethernet but almost none have ever worked with SONET. This means that engineers with the skill set to understand Ethernet are available everywhere which is not the case for many other technologies. Combine that with economical hardware and the bottom line is Ethernet's solutions to problems are preferable over other technologies.

To better understand Ethernet Ring Technologies we must first understand what a loop is. A loop is a configuration where an Ethernet frame travels around a network and ends up back where it started. The image to the right in figure 1 is an example of a simple loop. If we assume that a loop preventing protocol has not been employed on these Ethernet switches we could see where a frame could be sent from one switch to the next. It would keep going around and around eating up bandwidth. Even more troubling, the frames passing around the loop will wreak havoc with the MAC address tables. The result is a network that will not allow packets to pass through it.



Figure 1: Ethernet Loop

Now that we understand loops, we also understand that loops are bad but we would like to have the path redundancy that a loop provides. ERPS is one way of accomplishing this. An Ethernet ring consists of switches that form a closed physical loop. Each ring switch is connected to two adjacent Ethernet ring switches. This is exactly what we said to avoid earlier, but since each ring switch is running ERPS on the ports which interconnect them, it works.

White Paper

That was too simple of an answer of course but it does describe it in about as concise of a way as possible. Let's look at it a little deeper. Diagram 1.0 below shows four switches configured in a ring. Switches A through D each have ports P1 and P2 configured to participate in a ring, specifically ERPS. One of the switches needs to have one of its ports configured as the owner of the ring. In the case of the diagram below Switch A, port P1 is the Owner. Switch B, port P1 is directly attached to the Owner port and needs to be designated as the Neighbor. This is a way of letting switch B, port P1 know that it is directly connected to the owner of the ring. This is important because to prevent a loop, one of the links in this ring must not pass data. All other ports involved in the ring must be defined as part of the ring, but not the owner or neighbor. While the rest of the loop is intact, the Owner blocks traffic from passing thus preventing the loop. Switch B will also start blocking traffic going out of the P1 port because it has determined that the rest of the ring is intact. This is the ring in a converged state under normal conditions. The word converged is used to refer to a ring that has settled into a particular configuration. It is no longer in a transitional state.




Diagram 1.0: ERPS Ring

Now let's look at a ring with a fault. In diagram 1.1 below, a fault has occurred between switches B and C. Switch A and B have been notified of the fault and start allowing traffic across the link they share. ERPS can make this change in under 50ms. Switch B and C will start blocking traffic passing between them but continue to check to see if the link fault is there.




Diagram 1.1: ERPS Ring Setup at Fault

Once the broken connection is fixed, the ring will transition back to the state where the Owner port and Neighbor port are blocking traffic if the ring is configured to be revertive. If the ring is configured to be non-revertive the ring will remain blocking the ports where the link was broken waiting for an administrator to reset the ring back to its normal state. This can help with links that may be flapping or where the fault may be intermittent.

Let's look at how to configure ERPS on an Antaira switch. When configuring a member of a ring you will need to know which two ports will be participating in the ring, what roles they will be playing, the Ring ID, the APS Channel and finally if you

White Paper

want the ring to be revertive. Below in figure 2 is a screen shot from an Antaira switch showing how this information might be entered.

The port selection and role should be pretty obvious at this point but the Ring ID and APS channel are a little more interesting. The APS channel is a VLAN tag that is given to the frames that contain the ring control commands called R-APS (Ring Automatic Protection Switching) messages. All members of the ring must have the same APS Channel. By encapsulating these frames in a VLAN the protocol can contain the R-APS messages to the ring and not pollute the rest of the network with useless frames.

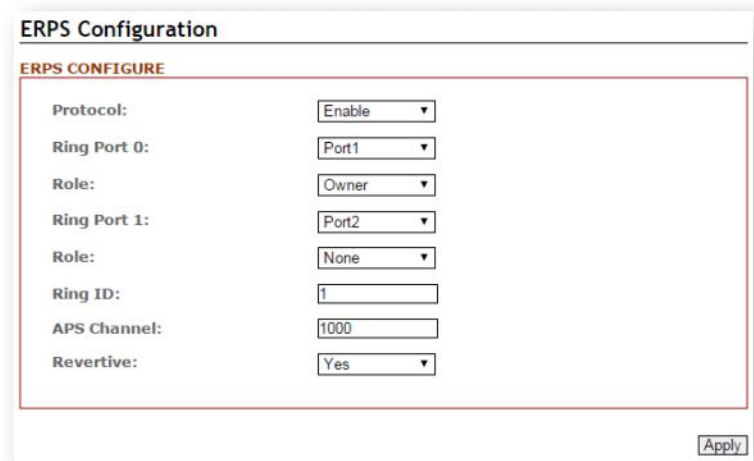


Figure 2: Antaira Web Interface – ERPS Configuration

Additionally since the control frames only need to run around the ring, changes to the ring’s status can quickly be disseminated across the switches participating in the ring. This leads to very quick changes to fix faults in the ring. The Ring ID is appended to each of the frames of the R-APS control messages preventing messages from one ring being confused with another ring when using the same APS channel.

A revertive ERPS configuration will let the ring converge back to its normal state after a fault has been repaired. In most cases this is a preferred behavior but sometimes when a link is having intermittent issues and can’t be fixed right away a non-revertive configuration is needed. This will leave the ring in the failover state until an administrator can reset the ring back to its normal state

But how do the switches participating in the ring determine link status? It is possible that the link at the PHY (physical) layer can remain connected while data is not able to pass so simply relying on physical link layer status is not ideal. CFM (Connectivity Fault Management) and line status messages are used to detect ring link and switch failure. This provides a truer test of connection and permits faster detection of faults. The CFM uses CCM (continuity check messages) from the 802.1ag standard which is set at an interval of 3.3ms as specified in the g.8032 standard providing a very fast detection of fault and convergence of topology.

There are a few R-APS control messages worth noting here. A failure along the ring will initiate an R-APS Signal Failure (R-APS SF) message while also blocking the failed port. The message will be sent out by both of the switches participating in the ring that have identified the fault. On obtaining this message, the ring owner and ring neighbor will unblock their ports.

When the failure has been repaired, the switches connected to the restored link send R-APS NR (R-APS No Request messages), assuming the ring is configured to be revertive. The ring protection link (RPL)

White Paper

Owner and Neighbor blocks the RPL port and sends R-APS NR and R-APS RPL (R-APS NR, RB) messages. These messages cause all other switches, other than the RPL Owner and Neighbor in the ring, to unblock all blocked ports. In a non-revertive configuration the repair of the fault would not cause the R-APS NR message to be sent and would wait for an administrator to initiate the convergence back to its normal state.

The interesting thing is that Ethernet Ring Protection (ERP) protocol works for both unidirectional failure and multiple link failure scenarios in a ring topology.

A single ring can withstand a single failure before connectivity is lost between switches. The larger the ring the greater the potential of having more than one failure at a time. To address this, it is possible with G.8032 (version 2) to have multiple rings on a single switch. In diagram 2.0 below, switches B and C have two different rings associated with them. Connectivity between A and F could sustain failure in two places if each failure occurred on different rings.

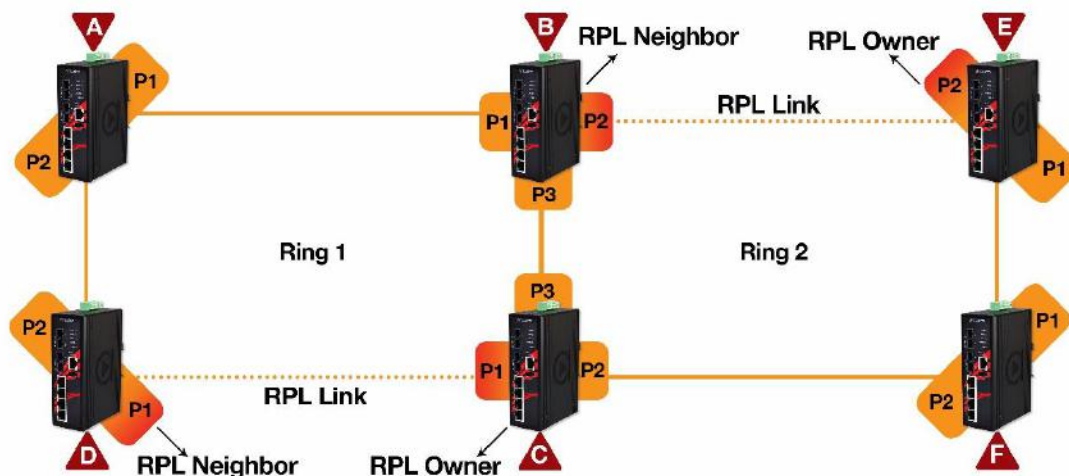


Diagram 2.0: Dual ERPS Ring

By adding more rings the network becomes more resilient to failure and starts to take on the look of a ladder as in diagram 2.1. This is referred to as a multi-ring/ladder network.

ERPS is not alone in providing link redundancy in the Ethernet world. There are other types of rings and technologies based on Spanning tree. They are worth noting here because ERPS is not always the best solution for providing a resilient network.

Proprietary rings have been around for a while, but as the name denotes, they are not viable in a heterogeneous networks or networks with multiple hardware manufacturers. Heterogeneous networks are becoming more prevalent as more electronics manufacturers enter the networking market. There are ring technologies that were originally proprietary but have been opened up for other manufacturers such as MRP (Media Redundant Protocol). MRP operates at the MAC layer of the Ethernet switches and

White Paper

is a direct evolution of the HiPER-Ring protocol used by Hirschman switches since 2003. This protocol is used extensively in markets where Hirschman switches dominated for years.

Spanning tree based redundancy is typically slower to converge than rings but can require less configuration depending on how it is implemented. When used with VLANs (MSTP-Multiple Spanning Tree Protocol) this technology will use all paths for maximum bandwidth while all links remain available. This is when the configuration can be more complex than ERPS. Generally speaking spanning tree technologies are best used in smaller networks or in small groups of devices as it can become slow when too many devices are participating in the same spanning tree.

No single technology is ideal for every application. ERPS rings have found a niche where they not only excel but currently provide the best option. Are you still interested in knowing more? Please contact a Sales Engineer at Antaira to find out more or to arrange a demo.




Diagram 2.1: Multiple ERPS Ring

About Antaira Technologies, LLC

Antaira Technologies is a leading developer and supplier of industrial networking and communication product solutions. Antaira has offered a full spectrum of product lines that feature reliable Ethernet infrastructures, extended temperature tolerance, and rugged enclosure designs. Our product lines range from industrial Ethernet switches, industrial wireless devices, Ethernet media converters, and serial communication devices. Our vast professional experience allows us to deploy a wide array of products worldwide in mission-critical applications across various markets, such as, automation, transportation, security, oil & gas, power/utility, and medical.