



GarrettCom, Inc.
213 Hammond Ave.
Fremont, CA 94539
USA

Phone: (510) 438-9071
Fax: (510) 438-9072
Email: mktg@garrettcom.com
Web site: www.GarrettCom.com

Redundancy with Standards in Industrial Ethernet LANs

A White Paper for Network Engineers in Factories, Transportation
Systems, Utilities, and Other Heavy Duty Networking Applications

by

Frank Madren, President, GarrettCom, Inc.

May 2003

© Copyright 2003 GarrettCom, Inc.

Redundancy with Standards in Industrial Ethernet LANs

Industrial Ethernet networking has inherent advantages for a multitude of industrial applications. By utilizing a standards-based solution that supports multi-vendor implementations, Industrial Ethernet users enjoy highly reliable systems, reduced costs of deployment, and a guaranteed upgrade strategy as needs evolve. New product offerings for Industrial Ethernet and transportation networks offer hardened skins, expanded temperature ranges, DC power, fiber-built-in media, and other embellishments over the traditional office LAN products that have used standards-based deployment for over a decade. Now, Ethernet is tough enough for the heavy-duty applications in environments that were previously judged to be too hostile.

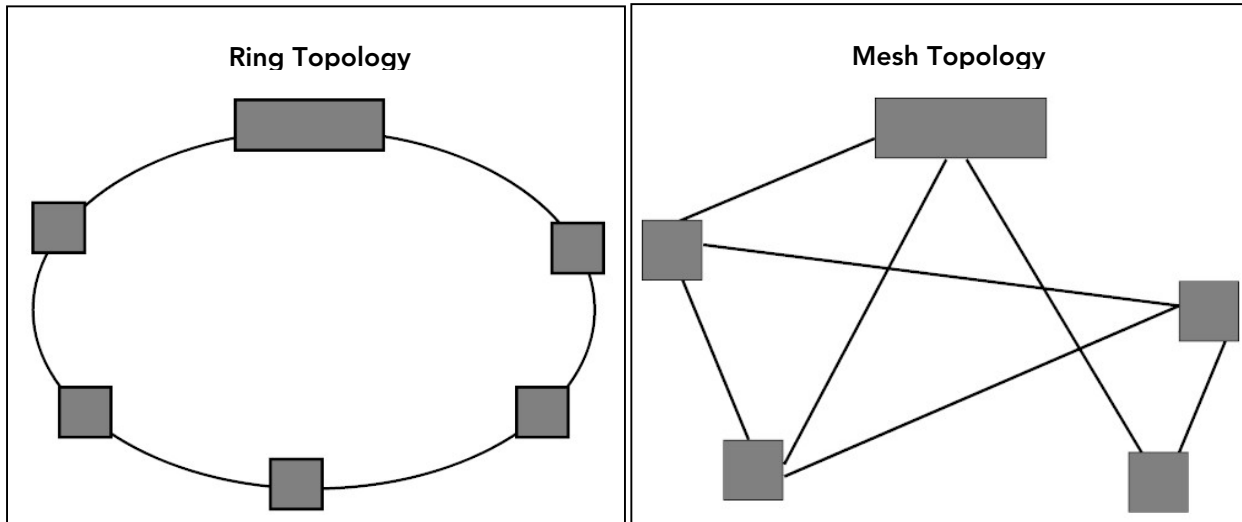
High availability, achieved through redundancy and fault tolerance, is a critical component of many industrial network deployments. Where loss of an enterprise network for a few minutes is inconvenient, loss of an industrial network can have disastrous consequences. Many transportation systems and production processes are highly integrated, so that a fault at one location travels rapidly upstream and downstream. Interruptions to factory operations can cost thousands of dollars per hour, easily justifying the extra expense for hardened and highly reliable control and information systems.

In deploying redundant Industrial Ethernet applications, companies have depended mainly upon custom systems developed to suit individual applications. There is relatively little practical experience since working installations were first being deployed in the year 2000. The industry, which has used costly serial connection protocols and proprietary solutions for years, is attempting to adapt standard Ethernet to fault-tolerant industrial networks. The standards benefits of flexibility and interoperability are obvious. However, there is a challenge in adopting software standards that were designed for a more forgiving office environment to the harsh realities of the tough and skeptical industrial market.

What Standard Software Is Available?

The IEEE 802.1d standard Spanning Tree Protocol (STP) has been available for use with managed switches and bridges for several years. This software provides a mechanism for resolving redundant physical connections in order to maintain operation of standard Ethernet LANs that does not allow more than one path for a packet to be in use at a given time. The Spanning Tree Protocol is included with the managed switch software provided by all major Ethernet managed switch product suppliers, and is widely available in the marketplace. Further, STP has proven in general use over many years to be interoperable, and commercial systems utilizing products from multiple vendors are routinely implemented. Standard STP supports redundant configurations of any type: meshes or rings or combinations.

More recently, Rapid Spanning Tree Protocol (RSTP) has appeared and is beginning to be used in some redundant LAN systems. It was designed to be faster than STP in fault recovery time for small- or medium-sized meshes and some small rings. Experience with RSTP interoperability is not generally available because it is new, and it appears that some vendors who support it have made proprietary modifications.



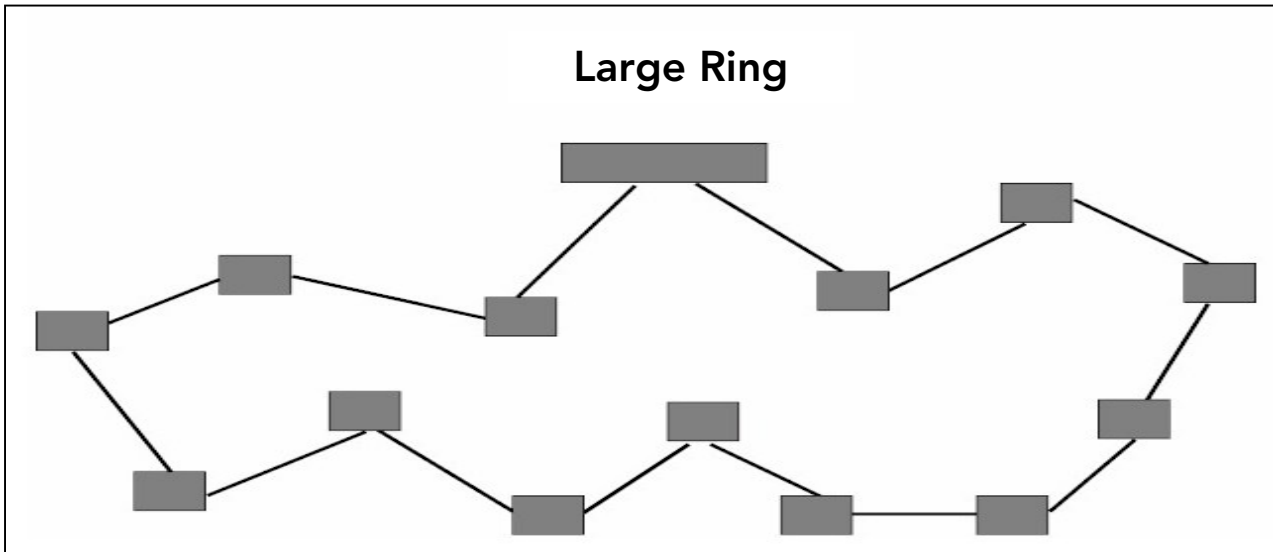
Standard software for redundancy, especially STP, is sometimes denigrated by vendors who offer proprietary alternatives and want to promote the advantages they claim for their products. Because STP does not have significant vendor sponsorship, data on its use is not widely available and training materials and applications experience is somewhat limited. As STP experiences increasing use in Industrial LANs, more general information on applications is emerging. Redundant LAN configurations can be constructed in a variety of ways. While mesh configurations are a more general topological case, ring configurations for redundancy are especially useful and cost-effective in industrial LAN systems, and will be treated in this paper in more detail.

Ring Structures in Industry

Simpler cabling makes rings ideal where industrial facilities cover extended areas, such as in energy and transportation systems. Railroads, pipelines, windmill farms, oil and gas producing fields, waterways and canals, tunnels, highways and city traffic control systems are all good examples of redundant ring applications covering large areas and long distances. Other industrial facilities that benefit from rings include water treatment plants, mines and quarries, forest product mills, agricultural buildings, and warehouses.

A ring topology is the most practical and most widely used redundant LAN configuration for industrial applications because of the distance between nodes that is typical of applications such as those listed above. A mesh structure would be impractical and too expensive because of the high costs of constructing the interconnect cabling.

Ethernet is the preferred protocol because of the plentiful supply of industrial-grade switches and hubs running at 10 and 100 Mb/sec speed that provide more than adequate bandwidth for industrial applications. Use of a "daisy-chain" or sequential point-to-point topology is optimal for minimizing the cabling expenses that dominate overall installation cost. In most cases, routing the end of the cable string back to the switch that manages the daisy-chained units is fairly easy. This enables the creation of a ring structure with redundant capabilities.



The trick is finding a standards-based LAN software solution to support the redundant ring. It is in this area that major strides are being made.

There are three options for managing a redundant ring:

1. Make do with current standards, even though they were not optimally designed for ring applications.
2. Use a vendor-proprietary solution for the ring, which may provide faster fault recovery but at the cost of a more expensive and less flexible implementation. Proprietary solutions restrict the choices of devices that can be used in the ring and lock out other vendors' offerings.
3. Choose a standards-based implementation with ring-specific extensions that speeds up fault resolution, and also allows for any standard Ethernet product to be utilized within the topology.

Option 1, Use STP, the Established Redundancy Standard

STP has been used in redundant LAN applications for more than 10 years. Initially designed to support redundancy using Bridge technology in 10Mb/sec. bus topologies, STP has evolved to be used in multi-port switches with star topologies in high-speed LANs. The STP functional standards, established in an earlier networking era, have served well and are still performing in many high-availability networks today.

Rings are a simple subset of the mesh topology where STP excels (See diagrams above). As previously mentioned, there are delays inherent in STP's collection and analysis of data prior to executing a recovery that are necessary to support complex mesh structures. However, in some edge applications, the standard STP delay is too long to be acceptable.

The obvious advantages of STP are its maturity, its proven reliability, and the inherent interoperability achieved by using an accepted industry standard. STP is available on most, if not all Ethernet managed switches, which can then be mixed and matched in a deployment. STP was designed to support a variety of LAN topologies and works well with both hubs and switches.

The seeds of STP's weaknesses reside in its strengths: it is not inherently ring-oriented, and the complexity that allows it to support a variety of topologies limits its performance in a relatively simple redundant ring. When there is a fault in a ring, the obvious solution is to treat the interrupted ring as two separate strings until the fault is repaired. Since there is only one fault recovery solution, the time it takes standard STP to collect data and make an analysis of the situation that conforms to the standard is both unnecessary and a weakness when fast fault recovery is an objective.

STP's other weakness is that it does not easily scale up to handle large rings. STP passes messages among switch members that resolve redundancy conflicts, which works well when all members are within a couple of hops of the "root" switch decision maker. It is not good in rings where the switch members are strung out and may number 10 or 25 or even a hundred in a ring with each member passing messages along down the line. A large ring becomes too unwieldy for STP to handle well. A simple ring structure is best handled by one decision-maker switch handling the two "top" ends of the ring, and with ring members following the standard Ethernet packet-processing protocol. The complex structures that STP was designed to handle makes it overkill for simple rings.

Option 2, Go With a Proprietary Solution

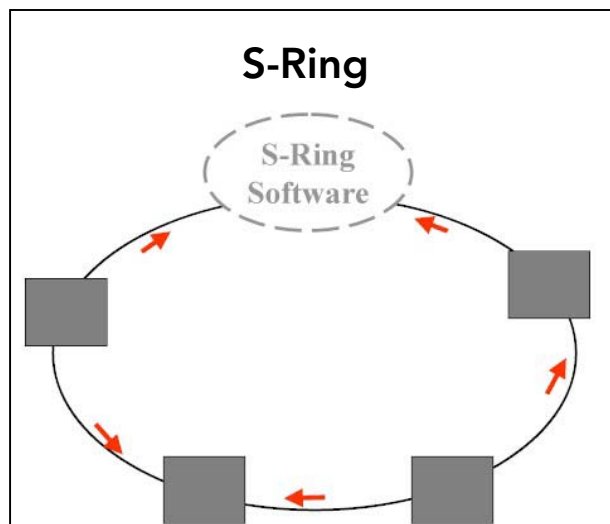
In fast-moving industries, the need for an immediate, practical solution often initially outweighs the perceived benefits of waiting for a standard to develop. Companies with an immediate problem, or companies that like to be on the leading edge of new technologies are willing to take the risk of working with vendors that offer a proprietary resolution to their problem. For the last several years, such has been the case with rings in redundant LANs where fault recovery faster than standard STP can provide is needed.

The benefits of going proprietary, so long as the solution works, are: a leg up on resolving the problem and more rapid fault resolution than waiting for a standards-based solution. The downside is the risk and cost associated with a proprietary solution and becoming locked in to a single source. For proprietary ring solutions, there is limited – or no – interoperability with other products on the market, and the solution is more costly – both in initial purchase price and in the lifecycle costs as standards-based solutions evolve. Most users rightly feel that proprietary solutions are a last resort, only to be chosen if nothing else is available that will do the job.

Option 3, Choose Standard STP with Ring Enhancements

The third option combines the strengths of the other two solutions options while managing or eliminating their downsides. By developing a faster ring-based fault recovery process that takes advantage of the features and protocols of the STP standard, a vendor can provide customers with a safe solution that works fast today. By ensuring interoperability and the ability to take advantage of evolution in industry standards, this option will continue to protect the application investment because it is open to innovation that may make it work even better tomorrow. Because such a solution retains the benefits of the standard, multi-vendor implementations remain a viable option, providing the competition and product selection that will keep the cost and the vendor risk factors of the deployment low.

GarrettCom™, Inc. has developed a ring fault recovery protocol that provides for fast, reliable fault recovery while preserving interoperability with standard STP products. The company's S-Ring™ product, available on Magnum 6K Managed Switches, uses the standard STP status-checking multi-cast packets (called Bridge Protocol Data Units or BPDUs) to determine the occurrence of a fault, but takes the initiative to override the STP analysis step, immediately forcing the reconfiguration of the ring to recover from the fault. Without interfering with standard STP operation, the S-Ring software can be selected to operate on a port-pair that supports a ring to reduce the fault recovery time from minutes (one-half minute to 5 minutes for standard STP) to seconds (less than 2 seconds, ring switch buffers permitting). The speed of an S-Ring solution is competitive with proprietary ring recovery methods but implemented on a standard STP base that allows multi-vendor, industry-standard Ethernet hubs and switches to be used.



Without interfering with standard STP operation, the S-Ring software can be selected to operate on a port-pair that supports a ring to reduce the fault recovery time from minutes (one-half minute to 5 minutes for standard STP) to seconds (less than 2 seconds, ring switch buffers permitting). The speed of an S-Ring solution is competitive with proprietary ring recovery methods but implemented on a standard STP base that allows multi-vendor, industry-standard Ethernet hubs and switches to be used.

Special Ring Timing Issue – Switch Address Buffer Behavior

Ethernet switches learn MAC addresses in order to switch packets to their destination port, and save the addresses in their memory as long as they are active. If a MAC address ceases to be active, it is aged out of the switch memory after a few minutes. This switch-address-aging delay presents a problem when a LAN needs to be reconfigured quickly. While repeaters (hubs) have no address buffer and therefore do not create a recovery bottleneck, the switch's stored addresses prevent packets from going via a new recovery route until the addresses have aged out of switch memory.

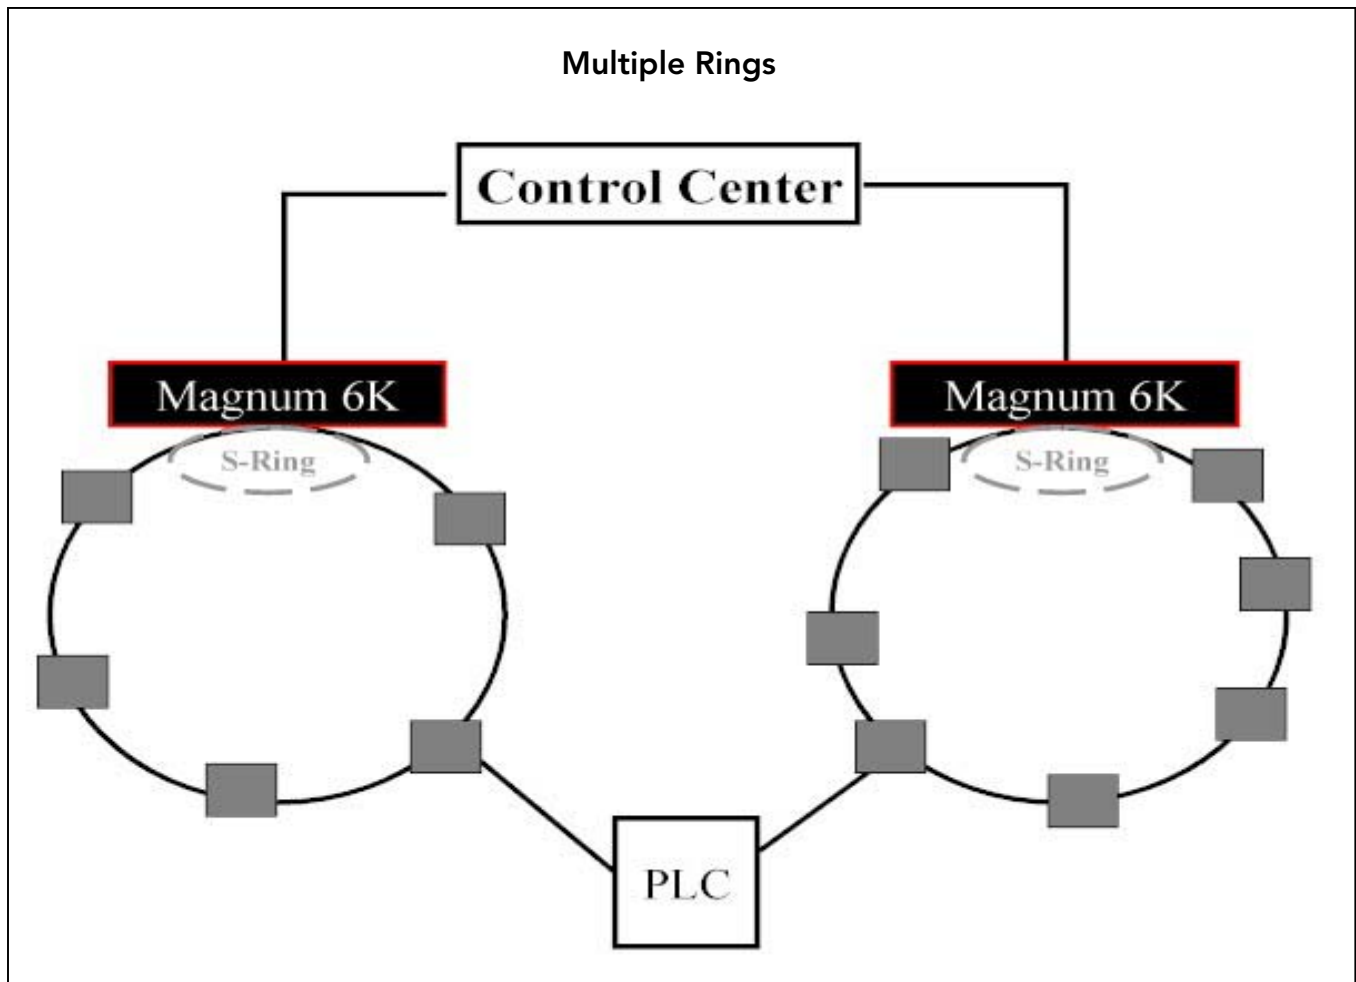
There is no benefit to having a fast ring recovery technique if the switch members of the ring prevent Ethernet traffic from moving to the recovery traffic path. Different switch vendors implement different address buffer aging times. In a multi-vendor implementation, the slowest aging time in the recovery path will govern the ring recovery time. It is critical to build in switch address buffer aging times when calculating the time-to-recovery in a redundant ring.

In the Magnum mP62 edge switches, a feature called Link-Loss-Learn™ (LLL) can be activated to immediately flush its address buffer and relearn the MAC addresses that route packets around the fault. This procedure, which is similar to switch initialization, occurs within milliseconds, resulting in fast ring recovery. An S-Ring implementation watches for Link-loss as well as for STP BPDU packet failures and responds to whichever occurs first. In most instances, the Link-loss will be detected faster than the two-second interval at which the BPDU packets are successfully passed around the ring. Typical ring recovery times using S-Ring software and mP62 edge switches with the LLL feature enabled on the ring ports is less than 250 milliseconds, even with 50 or more mP62 switches in a ring structure. Without LLL activation, the mP62 Switch's address buffer aging time (3 minutes default) could be the gating factor in ring recovery time.

Multiple Rings for Dual Redundancy

In industrial networks, dual-ported PLCs maximize uptime by requiring two independent paths through the LAN without a common point of failure. While very expensive to install and operate, such a redundant system can be justified where the cost of downtime is extremely high.

A ring structure has a common point of failure, namely the top switch in the ring. For dual systems needing maximum redundancy and uptime, no single point of failure, and fast fault recovery time, two rings operated from two Magnum 6K Switches running S-Ring software may provide both performance and security.



A full economic analysis of the cost of redundant media and active products along with technical trade-offs would be necessary to determine the best dual redundant solution. It is useful to know, however, that dual redundant ring structures can be supported based on industry-standard interoperable platforms such as Ethernet with STP, and they can incorporate multi-vendor hubs and switches if desired. Dual redundant ring structures supported by S-Ring technology can achieve sub-second fault recovery times, making them a candidate for use in any fault tolerant LAN design.

Conclusions

Enhancements to standard Spanning Tree Protocol can add robustness and fault recovery speed to a LAN ring structure without sacrificing the benefits of a standards-based redundancy solution. It is particularly important for industrial applications that such enhancements can support the large rings and long distances typical of industrial applications because the fault detection and recovery delays are considerably shortened. Redundancy with enhanced standards addresses the need of industrial customers for a reliable, mature solution that avoids proprietary vendor lock-in but provides competitive or superior fault recovery time.

© Copyright 2003 GarrettCom, Inc. Magnum, Link-Loss-Learn, S-Ring, and GarrettCom are trademarks of GarrettCom, Inc. Ethernet is a trademark of Xerox Corp.

Bibliography:

IEEE 802.1d Standard

GarrettCom Link-Loss-Learn Technical Brief, 2002

GarrettCom S-Ring Technical Brief, 2003

Networking as a 2nd Language; Understanding Spanning Tree Protocol -- the Fundamental Bridging Algorithm, Michael Norton, O'Reilly Network, 03/30/01

Achieving Fault-Tolerance with PC-Based Control, David W. Cawfield, ISA Automation & Control Subsystems Committee