

Standards-based Approaches to Redundancy and Fault Tolerance using Industrial Ethernet LANs

A White Paper for Network Engineers in Factories, Transportation Systems, Utilities, and other Industrial Networking Applications

Redundancy and Fault Tolerance Using Industrial Ethernet LANs

Industrial Ethernet networking has inherent advantages for many industrial applications. By utilising 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, power utility substations, and transportation networks offer hardened packaging, optional sealed chassis and anti-corrosive coatings, expanded temperature ranges, DC power, fibre-built-in media, and other adaptations that make them distinctly different from the traditional office LAN products that have been using standards-based deployment for nearly 20 years.

Today, 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. In substations, transportation systems, video surveillance, access control and production environments, processes are highly integrated; a fault at one location can travel rapidly upstream and downstream. Interruptions to factory operations can be very costly, they can easily justifying the extra expense for hardened and highly reliable control and information systems.

Redundant Industrial Ethernet applications date back to the turn of the century, and are becoming increasingly commonplace. As IP networking becomes the solution of choice for networking applications, no matter how hostile the environment, creative and cost-effective solutions that provide redundancy for virtually 100 percent availability are frequently used. Ceding to the benefits of standards-based IP solutions, serial connection protocols and WAN interfaces are being integrated into fault-tolerant Ethernet industrial networks. The standard benefits of flexibility and interoperability are obvious.

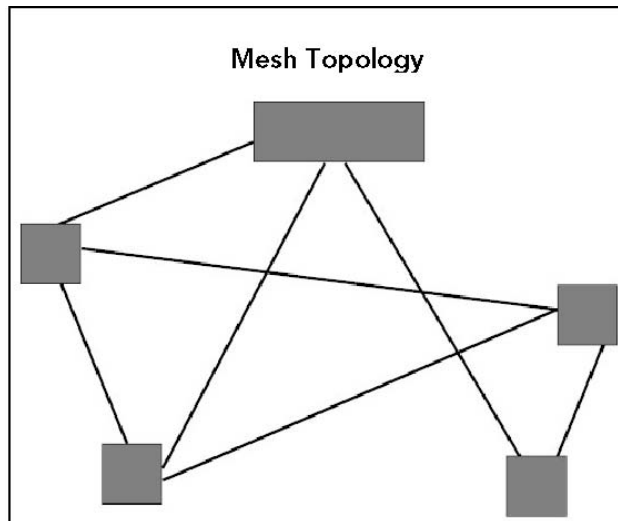
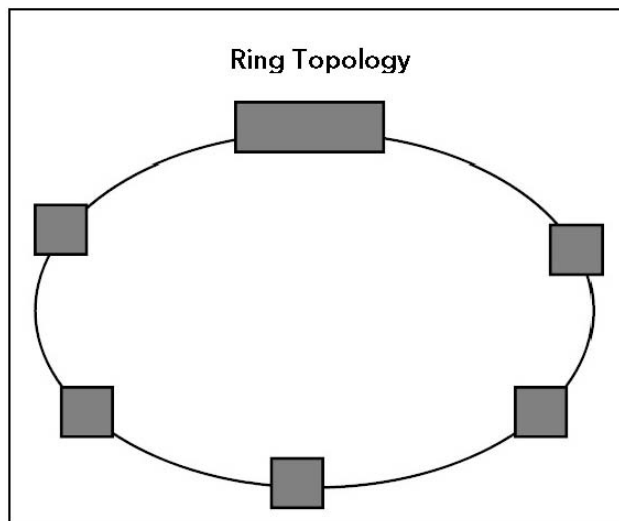
What Standard Software Is Available?

The IEEE 802.1d standard Spanning Tree Protocol (STP), adopted in 1990, was the original standard for Ethernet fault recovery. It provides a mechanism for resolving redundant physical connections to maintain the operation of standard Ethernet LANs but does not allow more than one path for a packet to be in use at a given time. STP has been deemed too slow for most modern industrial applications. Some vendors offer proprietary alternatives, however Rapid Spanning Tree Protocol (RSTP), adopted in 1998 and described in IEEE 802.1w, was up to 50 times faster than STP and is widely accepted today.

IEEE 802.1D-2004 further revised the standard and now offers a higher-speed implementation of RSTP that can also support larger networks. Redundant LAN configurations can be constructed in a variety of ways. While mesh configurations are a more general topology, ring configurations

White paper

for redundancy are especially useful and cost-effective in industrial LAN systems, and will be treated in this paper in more detail. In addition to RSTP, which requires a managed switch or router at each node, this paper also addresses the GarrettCom standards-based S-Ring[†] protocol that can utilise both managed and unmanaged switches.




Ring Structures in Industry

Ethernet is the preferred protocol for redundant industrial applications because of the plentiful supply of industrial-grade switches and hubs running at 10/100/1000 Mb/sec and higher speeds that provide more than adequate bandwidth. Use of a "daisy-chain" or sequential point-to-point topology is optimal for minimising 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.

Ring topologies are ideal where industrial facilities cover extended areas, such as transportation systems and power utilities, as well as other industrial applications. 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 long distances. Other industrial facilities that benefit from rings include power substations, water treatment plants, mines and quarries, forest product mills, agricultural buildings, seaports and airports, and warehouses. A mesh structure will usually be impractical and too expensive because of the high costs of constructing the interconnect cabling.

RSTP goes a long way toward providing a universal solution for LAN redundancy. However, it still contains complex reconciliation testing required for mesh networks that can slow down the less complex redundant rings.



As with all standards, the evolution from first attempts to a highly-tuned implementation takes time. Companies attempting solutions and the vendors that supply those solutions must make hard decisions over the evolution of a standard.

The obvious benefits of standards approaches are interoperability, wide availability, lower cost, and the highest possible assurance of forward and backward compatibility. However, from time to time, there is the temptation, particularly in the earlier stages of adoption of standards, to adopt a proprietary solution because it can be streamlined to meet a particular objective. It may, therefore, provide a temporary advantage – but at a cost as the standard progresses. In the case of Spanning Tree Protocols for redundancy support, it is useful to evaluate the three options that have been available to companies in industrial environments.

1. Adopt a standards-only policy,
2. Use a vendor-proprietary solution that may provide a temporary advantage but may restrict options as the standard evolves.
3. Choose a standards-based implementation with offers extensions to the standard, and yet remains compliant as the standard evolves.

Option 1: Use the Established Redundancy Standard

Since the initial adoption of the STP standard in 1990, there has been a dramatic increase in performance. Early STP implementations took minutes to resolve faults. When RSTP was adopted in 1998, it reduced this time to seconds, and RSTP-2004 resolves faults in milliseconds. Over the years, required fault resolution times for industrial applications often far outstripped the capabilities of the standard, and there was pressure to find alternatives. RSTP-2004's fault resolution speeds are similar to or better than those achieved through proprietary alternatives.

Today, RSTP is the standard of choice for redundant LAN applications, although STP is still in use to support legacy LAN hubs and switches. For completeness, we will describe RSTP standards in this section, with the understanding that STP is simply an earlier and slower implementation.

The obvious advantages of RSTP are maturity, proven reliability, and the inherent interoperability

achieved by using an accepted industry standard. RSTP is widely available on Ethernet managed switches, which can then be mixed and matched in a deployment. It works well with both hubs and switches, and supports a variety of LAN topologies including complex star technologies using multi-port switches in high-speed LANs. Rings are a simple subset of the mesh topology where RSTP and STP excel (See diagrams above).

Initially, the decision process necessary to resolve faults in mesh topologies added unacceptably high overhead for fault resolution in the simpler ring applications. 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. A simple ring structure is best handled by a single decision-maker switch handling the two “top” ends of the ring, and with ring members following the standard Ethernet packet-processing protocol. Until RSTP-2004, the complex structures that RSTP and STP were designed to handle made the standard solutions overkill for ring topologies. GarrettCom has measured the recovery time for RSTP networks utilising the 2004 standard as 2 milliseconds per hop in even large rings,

Another weakness with STP and RSTP, which has been largely addressed in RSTP-2004, was that they could not easily scale up to handle large rings. Spanning Tree protocols pass messages among switch members that resolve redundancy conflicts. This works well when all members are within a couple of hops of the “root” switch decision maker, but in the past this has limited performance in large rings of 10 to 100 nodes where the switch members are deployed at a distance from each other and each member must handle messages passed down the line.

With the adoption of RSTP-2004, performance utilising a standard solution is a highly competitive choice for most applications.

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, or applies when proposed standards do not address specific requirements of a user. Industrial companies that operate on the leading edge of new technologies may be willing to take the risk of working with vendors that offer a proprietary resolution to their problem. For several years, such had been the case with rings in redundant LANs where fast fault recovery was needed.

The downside is the risk and cost associated with a proprietary solution, including becoming locked in to a single source and not being able to take advantage of standards as they evolve. For most 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. Companies that chose a proprietary solution for earlier implementation of rapid fault recovery are now faced with a dilemma on how to take advantage of the RSTP-2004 standard. With the adoption of RSTP-2004, proprietary redundancy solutions are effectively obsolete.

Option 3: Choose Standards-based Solutions with Extensions

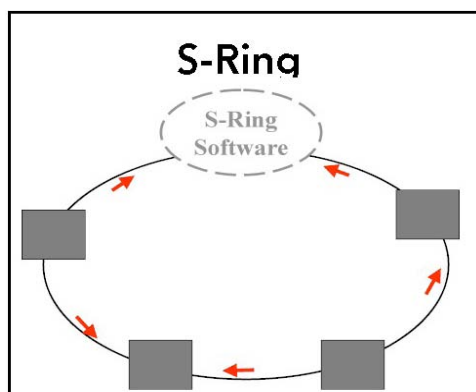
A third option exists when addressing standards that are still evolving. That is to look for standards-based solutions that provide extensions to handle certain requirements, while remaining compatible with the underlying standards. It has taken almost 20 years to develop a standard that meets the performance requirements of industrial applications.

By developing a faster ring-based fault recovery process that takes advantage of the features and protocols of the initial STP standard (and later, RSTP-1998), vendors, such as GarrettCom, have been able to provide customers with fast, safe solutions that worked in situations where performance requirements exceeded the rated performance of the standards. GarrettCom's S-Ring™ solution provided performance levels that exceeded those of both STP and RSTP-1998. At the same time, S-Ring ensured interoperability and the ability to take advantage of evolution in industry standards. This kind of solution retains the benefits of the standard and keeps multi-vendor implementations as a viable option; it preserves the competitive environment that keeps costs and vendor risk factors of the deployment low.

Without interfering with standard STP (or RSTP) 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 earlier spanning tree protocols) to seconds (less than 2 seconds, ring switch buffers permitting). Even with the performance improvement in the RSTP-2004 standard, the S-Ring solution is still viable and cost-effective because it supports both managed and unmanaged switches within a ring. RSTP requires a managed switch at each node.

The 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 (necessary with mesh topologies), immediately forcing the reconfiguration of the ring to recover from the fault.

Utilising a feature found in the Magnum 6K and mP62 managed switches called Link-Loss-Learn™ (LLL), S-Ring software can immediately force the flushing of switch address buffers so that they can relearn the MAC addresses that route packets around the fault. This procedure, which is similar to switch initialisation, occurs within milliseconds. 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 address buffer aging time (up to several minutes) could be the gating factor in ring recovery time.



White paper

This table provides a convenient comparison for the fault recovery protocols discussed above.

	RSTP-2004	S-Ring with LLL	RSTP	STP
License	Included in MNS-6K	A license key is needed. One key per ring manager switch	Included in MNS-6K	Included in MNS-6K
Spanning Tree	--	Works with RSTP or STP devices	--	--
Topology	Mesh topology – can have multiple paths	Single ring, multiple rings, no overlapping rings or ring of rings	Mesh topology – can have multiple paths	Mesh topology – can have multiple paths
Recovery time	Fastest, 2 ms per hop	Fast	Medium – sub second to a few seconds	Slow – in tens of seconds
Recovery decision	Typically done using BPDU. Rapid resolution.	Centralised to “Ring Manager”. LLL provides triggers to recomputed topology for ring members. Also works with RSTP or STP.	Typically done using BPDU. Can take time.	Typically done using BPDU. Can take time.
Resiliency	Multiple points of failure – each connected node can be in stand-by	Fast recovery from a single point of failure. Ring Master is responsible for decision making	Multiple points of failure – each connected node can be in stand-by	Multiple points of failure – each connected node can be in stand-by
Inter-operability	Multi-vendor; managed switches only	Managed or certain non managed Magnum switches; some hubs. Requires at least one Magnum 6K switch as ring manager	Multi-vendor; managed switches only	Multi-vendor; managed switches only
Software Cost	Included in MNS-6K	Licensed per ring	Included in MNS-6K	Included in MNS-6K
Hardware cost	Multiple switch types and vendors	One Managed 6K per ring. Multiple choices for members of the ring	Multiple switch types and vendors	Multiple switch types and vendors
Software Alarm	No	Yes	No	No
Ring Size	100+ nodes	50+ nodes	NA	NA
Dual Homing	Supports dual homed device to devices in the network	Supports dual homing to members in the ring	Supports dual homed device to devices in the network	Supports dual homed device to devices in the network


Multiple Rings for Dual Redundancy

The redundancy solutions that have been described above focus on faults within a single ring topology. Additional redundancy at the edge of the network, where the actual measurement and management is being done, is called for in some applications, and this paper will briefly review redundancy through dual-porting (or dual-homing) either a PLC or an edge switch. Dual homing provides two independent paths through a LAN without a common point of failure. Typically, one access point is the operating connection, and the other is a standby or back-up connection that is activated in the event of a failure of the operating connection.

The first strategy is to use dual-ported PLCs. While expensive to install and operate, such a redundant system can be justified where the cost of downtime is extremely high. The illustration demonstrates such an application using two rings operated from two Magnum 6K Switches running S-Ring software for 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 or RSTP, or by S-Ring technology.

Another dual-homing strategy is for use at the edge of the network. Because high availability is a key component in many industrial environments, shut-down manufacturing lines, power outages, and other system failures are becoming much too expensive and visible to tolerate. Practical ways to provide for recovery from faults for edge devices and nodes can be difficult. As discussed above, the software required to manage computers and other devices that have dual connectivity for redundant connections into the network is complex and costly. GarrettCom offers dual-homing technology (patent pending) in small industrial Ethernet switches that greatly simplifies the process. Simple, unmanaged Magnum ESD42 Switches offer convenient plug-and-play dual connectivity in a physically small package (about the size of a fist), and they are hardened and rugged for use in any industrial environment. With a MTBF of more than 30 years, they provide high reliability to enable redundancy for nodes at the edge of the network at a low cost.

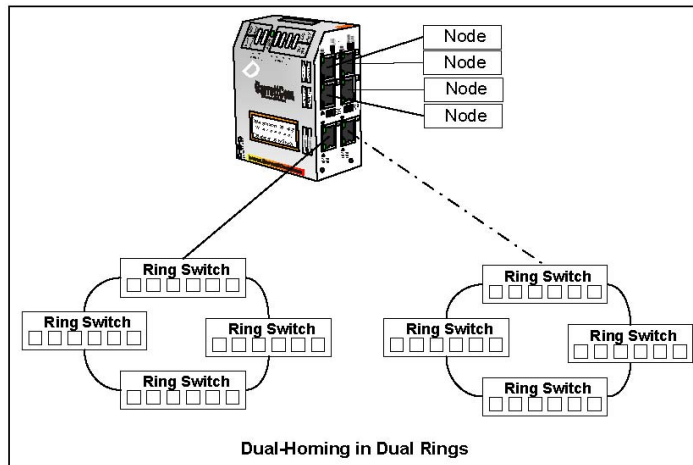


A dual-homing switch, with two attachments into the network, offers two independent media paths and two upstream switch connections. Loss of the Link signal on the operating port connected upstream indicates a fault in that path, and traffic is quickly moved to the standby connection to accomplish a fault recovery.

Conclusions

A standards-based redundancy strategy addresses the need of industrial customers for a reliable, mature solution that avoids proprietary vendor lock-in, but, from time to time, with the aid of standards-based enhancements may provide competitive or superior fault recovery time. When the fastest recovery speed for leading edge applications is the prime criteria, vendor enhancements to a standard may be the best choice.

As this paper has described, GarrettCom provides a variety of redundancy strategy options. By allowing customers to mix and match redundancy technologies, GarrettCom offers:



- Software-based standard RSTP solutions for managed network structures, including both routers and switches;
- Software-based standard STP solutions for use with legacy LAN hubs and switches;
- Standards-based proprietary solutions, such as S-Ring that allows for fast recovery in rings comprising both managed and unmanaged switches for a lower cost redundancy solution;
- Link-Loss-Learn recovery feature for managed edge switches; and
- Hardware-based Dual-Homing Switches for high-availability edge devices

Using a GarrettCom solution will make it possible to implement high-availability strategies that meet the unique requirements of any industrial application.