


Introduction

GarrettCom's S-Ring™ product enables Magnum™ 6K Managed Switches to simplify and speed up recovery from faults in Ethernet LAN configurations that use a ring structure. It is built upon networking software standards such as IEEE 802.1d Spanning Tree Protocol (STP). The user configures and controls S-Ring (patent pending) as part of the 6K Switch's management software. The S-Ring product can be used in multi-vendor LANs running standard STP on the Magnum 6Ks along with other switches and hubs in any redundant LAN topology, including rings and meshes. S-Ring makes Ethernet ring-topology more reliable and faster without sacrificing the benefit of standards-based interoperability.

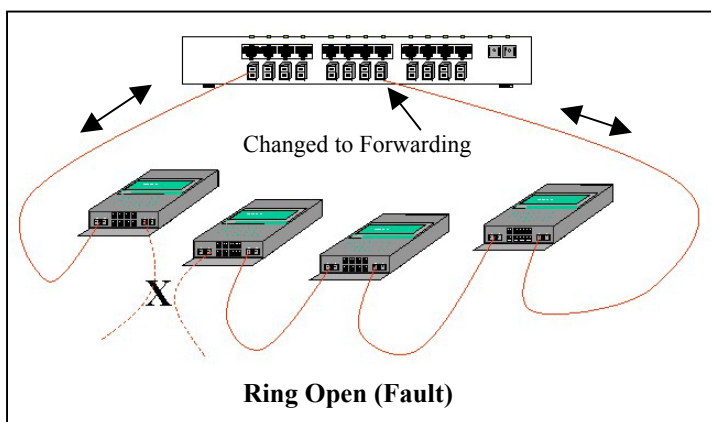
The S-Ring product operates from specifically defined port pairs that have ring-topology Ethernet devices attached. It builds upon the foundation of STP, but offers an additional option related specifically to ring topologies. Each of the two ends of a ring must be connected to two ports in a Magnum 6K Switch that is running the S-Ring software. To understand the S-Ring product, first examine the fault recovery of an Ethernet LAN ring configuration with only standard STP (no S-Ring) in operation.

STP Operation without S-Ring

The two top-of-the-ring ports form an otherwise-illegal redundant path, and standard STP causes one of these two ports to block incoming packets in order to enable normal Ethernet traffic flow. All ring traffic goes through the non-blocking port for normal LAN operation. Meanwhile, there is a regular flow of status-checking multi-cast packets (called BPDUs or Bridge Protocol Data Units) sent out by STP that move around the ring to show that things are functioning normally. This normal status is designated as RING_CLOSED. Operations will continue this way indefinitely until a fault occurs.



A fault anywhere in the ring will interrupt the flow of standard STP status-checking BPDU packets, and will signal to STP that a fault has occurred. According to the standard STP-defined sequence, protocol packets are then sent out, gathered up and analyzed to enable STP to calculate how to re-configure the LAN to recover from the fault. After the standard STP reconfiguration time period (typically 20 to 30 seconds), the STP analysis concludes that recovery is achieved by changing the blocking port of the ring port-pair to the forwarding state.




When this change is made by STP and both of the ring manager switch's ring ports are forwarding, the fault is effectively bypassed and there is a path for all LAN traffic to be handled properly. This abnormal status is designated RING_OPEN, and may continue indefinitely, until the ring fault is repaired. At that time, STP will change one of the ring control ports to be a blocking port again.

STP Operation with S-Ring

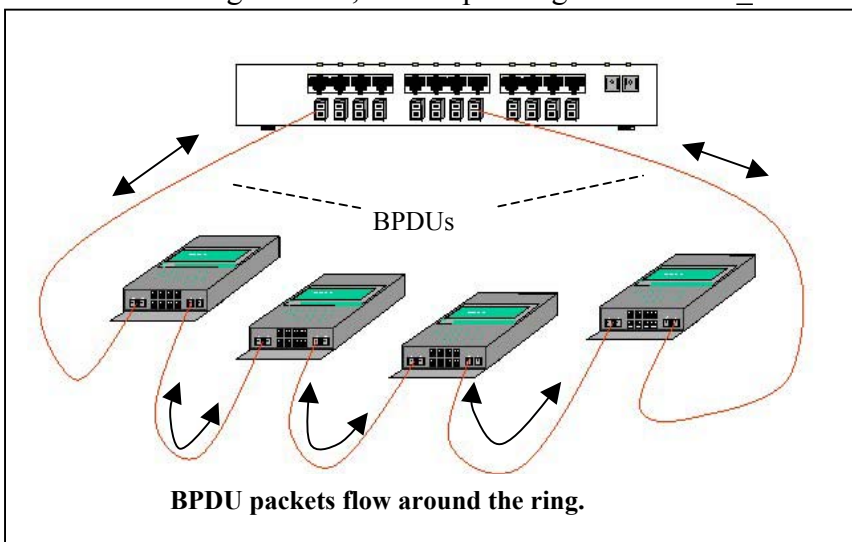
When the Magnum S-Ring product is enabled, the result of a ring-fault is the same but the recovery is faster. The S-Ring product overrides the normal STP analysis for the ring-pair ports of the ring manager (or ring-control) switch, providing quick recovery of the ring fault without conflicting with standard STP.

Magnum 6K Managed Fiber Switches with the MNS-6K software offer users the choice of selecting S-Ring when STP is configured and in use. For the S-Ring product option, the user must select two ports of one 6K switch to operate as a pair in support of each Ethernet ring, and attach to the two “ends” of each ring as it comes together at the ring control switch. More than one S-Ring port-pair may be selected per ring control switch. Each port-pair will have its own separate attached ring, and each port-pair operates on faults independently. The port-pairs may be of any media type, and the media type does not have to be the same for the pair. For Magnum 6K Switches, a port operating at any speed (10Mb, 100Mb, Gb) may be designated as part of an S-Ring port-pair . . . given proper Ethernet configuration of the ring elements, of course.



The user, having selected a port-pair for a ring, enables S-Ring on the selected port-pairs via S-Ring software commands. One command (enable / disable) turns S-Ring on and off. Another command adds / deletes port-pairs. Other commands provide for status reporting on the S-Ring product. The Magnum MNS-6K software package provides for remote operation, access security, Event Logs, and other industry-standard managed network capabilities suitable for industrial applications requiring redundancy.

When S-Ring is enabled for a port-pair, fault detection and recovery are armed for the associated ring. The standard STP functions are performed by the 6K Switch for other ports in the same manner as they would be without S-Ring enabled, when operating in the RING_CLOSED state. During this state, S-Ring is also



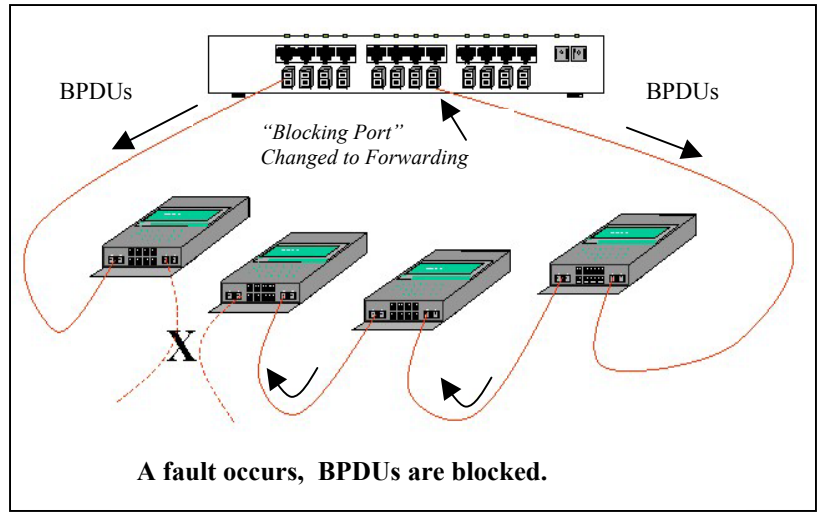
watching the flow of the BPDUs packets that move around the ring between the designated port-pair.

The extra capability of S-Ring comes into play when a fault occurs. When the flow of BPDUs packets around the ring is interrupted (or when Link-Loss is sensed on one of the ports of the ring port-pair), S-Ring quickly acts to change the blocking port's state to forwarding. No waiting for STP analysis. No checking for other possible events. No other ports to look at. No 30-second delay before

taking action. S-Ring takes immediate corrective action for quick recovery from the fault in the ring. The ring becomes two strings topologically, and there is a path through the two strings for all normal LAN traffic to move as needed to maintain LAN operations.

The BPDU packets continue to be sent out on the port-pair by STP and S-Ring while there is a fault, but none get through from one top-switch ring port to the other because the ring is in the RING_OPEN state. Meanwhile, S-Ring and STP look for the resumption of the BPDU packet flow that will signal a recovery from the fault.

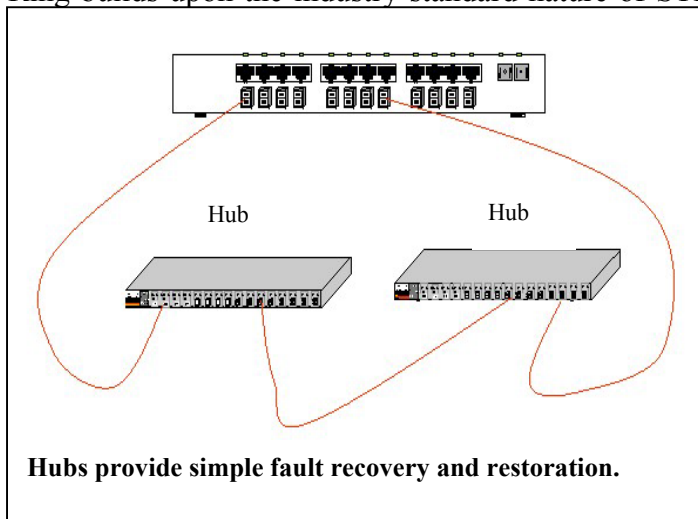
When the fault is cured, the re-emergence of the ring structure enables the BPDU packets to flow again between the ring's port-pair. This is recognized by S-Ring (and STP), and one of the ports in the ring's port pair is changed to the blocking state. S-Ring takes the recovery action immediately, not waiting for the 30-second STP analysis. Rings are simple structures. Either one port of a pair is forwarding or both are. Not complicated; not much to go wrong.



A Link-loss on one of the 6K Switch's ring ports (or a Link-loss on an mP62 Managed Hardened Switch somewhere in the ring if that is what the ring is composed of) is an alternative trigger for S-Ring to initiate fault recovery. The Link-loss trigger almost always comes quicker after a fault (a few milliseconds) than the loss of a BPDU packet which is gated by the standard STP 2-second "hello time" interval. So, the Link-loss trigger will almost always provides faster fault detection and faster recovery accordingly.

LAN Products in an S-Ring

The Ethernet products employed in an S-Ring, referred to as "edge" devices because they operate at the outer portion of the LAN where nodes and NICs attach, can be any standard Ethernet hubs or switches. S-Ring builds upon the industry standard nature of STP, and can work with Magnum products in the ring, and with mixed-vendor ring members.



If hubs (repeaters) are used in an S-Ring structure, then the hop-count rules pertaining to Ethernet hubs must be observed. Hubs (repeaters) have no memory and always broadcast each packet. They are easy to re-configure in a ring accordingly. They do not have to change the packet-flow direction when it is time to re-configure for fault recovery and ring restoration.

If switches are used in an S-Ring structure, then the address memory decay time must be taken into account for fault recovery and ring restoration.

Switches remember where to forward packets, and cannot adapt to accommodate a fault re-configuration until their address memory is changed or deleted.

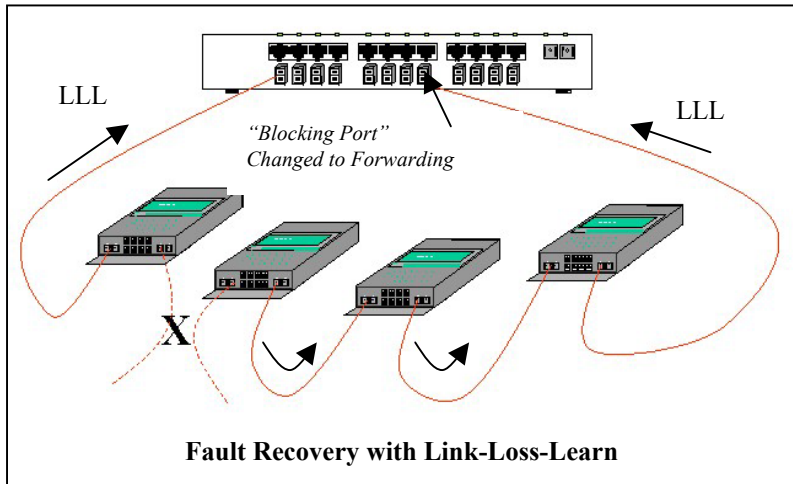
The standard STP status-checking BPDU packets flow in an S-Ring structure through any standard Ethernet hub or switch products. The detection of a fault, based on interrupting the BPDU flow, is the same no matter what equipment is used in the ring, but the fault recovery time and ring size will vary

considerably. In general, hubs can recover quickly but switches typically are delayed by the addresses stored in their memory. More importantly, hubs may only be used in small (two to five) unit rings while rings made of switches can number a hundred switches or more spread over great distances.

Switches with Link-Loss-Learn in an S-Ring

GarrettCom's Link-Loss-Learn™ feature, available on mP62 Managed Hardened Switches, can significantly reduce switch address memory decay time, resulting in more rapid reconfiguration. With

Link-Loss-Learn, mP62 Switches in a ring can flush their address memory buffer and quickly re-learn where to send packets, enabling them to participate in a very quick recovery or restoration.



Note that a Link-loss on an mP62 Managed Hardened Switch port somewhere in the ring is an alternative trigger for S-Ring to act for either fault recovery or ring restoration. The interruption (or the restoration) of the flow of BPDU packets is one trigger, Link-loss is another, and action is taken by S-Ring

based on whichever occurs first. See Technical Brief, the Magnum mP62 Link-Loss-Learn Feature.

Standard STP Interoperability with S-Ring

The same industry-standard BPDU packets used by STP are used by S-Ring for operations logic. This enables S-Rings to be constructed from industry standard Ethernet switches and hubs, and to co-exist with standard STP in a multi-vendor Ethernet LAN. In addition, standard STP can be in operation on the non-S-Ring ports of a Magnum 6K Managed Switch, participating in a multi-level standard STP redundant LAN hierarchy.

Spurious Operations Avoided

When configured correctly, S-Ring does not create any violation of standard protocols, nor does it cause degradation of network performance. The only spurious action possible would be caused by one of the edge switches in the ring generating a false Link-loss signal when the ring is not actually broken. This will result in an illegal loop for a theoretical maximum duration of one "hello time" period of the STP protocol, Packet loss / recovery will be handled as applicable by the Ethernet protocol's acknowledge / retry procedures.

Ring Learn Feature

One of the S-Ring software commands, S-RING LEARN, causes the scanning of all ports in the 6K Switch for the presence of rings. This command can be a handy tool in setting up the S-Ring product for correct initial operation. During a ring-learn scan, if any port receives a BPDU packet that was also originated by the same switch, the source and destination ports are designated as a ring port-pair and they are automatically added to the S-Ring port-pair list for that 6K Switch. The user can enable or disable ports pairs that are on the S-Ring list by CLI commands in order to exercise final control if needed.

Summary

S-Ring provides fault-tolerant rings in Ethernet LANs with very fast recovery and restoration, while using industry-standard STP software as a foundation. It is simple, fast, dependable, user-controlled, standards-based, and more economical than proprietary alternatives. It can handle small rings as well as very large rings with the same speed and reliability. It is what you can expect from GarrettCom – the place you need to look to for safe, broad-based, cost-effective Industrial Ethernet products.

© 2003 GarrettCom, Inc. Printed in United States of America 05/03. GarrettCom, Inc. reserves the right to change specifications, performance characteristics and/or model offerings without notice. GarrettCom, Magnum, Link-Loss-Learn (patent pending), S-Ring (patent pending), and Personal Switch are trademarks and Personal Hub is a registered trademark of GarrettCom, Inc. NEBS is a registered trademark of Telcordia Technologies. UL is a registered trademark of Underwriters Labs.



GarrettCom, Inc.

47823 Westinghouse Drive • Fremont, CA 94539 • PH: (510) 438-9071 • FAX: (510) 438-9072
Email: mktg@garrettcom.com • Web: www.GarrettCom.com

Appendix A

S-Ring on Magnum 6K Managed Ethernet Switches . . . a Technical Brief

S-Ring Timing Data, Summary

Definitions: 1) The STP “hello time” interval between BPDUs is user controlled, and can be selected within the limits of the STP IEEE 802.1d standard of 2 seconds to 10 seconds. For rings, the minimum of 2 seconds is usually selected, meaning that path (or ring) recovery time is 2 seconds maximum, 1 second typical or average

2) Path Recovery (or “ring recovery”) is defined as the operating state such that a new node can come on and find a working path enabling use of the ring elements to communicate with another new node.

3) Fault Recovery is defined as the operating state such that all existing nodes that previously communicated using the ring elements can communicate again.

4) S-Ring Restoration times are the same as Recovery times. S-Ring Recovery and S-Ring Restoration (i.e., ring break and ring make) are mirror events, requiring the same action by S-Ring and taking the same amount of time.

I. Standard Hubs in the ring

Path (or Ring) Recovery time, MNS-6K with STP and S-Ring: One STP “hello time”, 1 second typical

Fault Recovery time, same as Path Recovery time.

II. Standard Switches in the ring, no LLL.

Path (or Ring) Recovery time, MNS-6K with STP and S-Ring: One STP “hello time”, 1 second typical.

Fault Recovery time, one STP “hello time”, or the ring member switch’s address buffer memory aging time-out, whichever is greater.

III. All mP62 Switches in the ring with Link-Loss-Learn (LLL) enabled on the ring ports

Path (or Ring) Recovery time, MNS-6K with STP and S-Ring and LLL: 150 to 200 milliseconds reset time, plus 2 milliseconds per mP62 Switch in the ring for LLL propagation time.

Fault Recovery time, same as Path Recovery time.

For a single round number that represents a large number of timed experiences but is on the “safe side” of the average experience, 250 milliseconds for S-Ring Fault Recovery time is typically used

IV. For additional timing details, see the S-Ring product datasheet, or email sales@garrettcom.com

S-Ring, Ring Size and Distance Data

I. Standard Hubs in the ring

Ring size may be limited to the number of hubs in a collision domain per IEEE 802.1 configuration rules:

A standard Ethernet collision domain may be up to 4.5 Km at 10 Mb or 450 m at 100 Mb speed. Maximum number of hubs in a daisy-chain is 5. (Use of a switch between hub daisy chains may increase the number of hub units in the ring and provide greater ring distances, but will introduce switch complexities into the fault recovery process).

II. Standard Switches in the ring

Ring size may be unlimited in theory, i.e., any number of switches may be connected into a ring.

See your Switch manufacturer for test data or specifications on switches operating in a string or ring, including buffer packet delay times, address buffer memory aging time parameters, and link-fault-indication propagation times (if any).