Link-Loss-Learn on Magnum Managed and Un-Managed Ethernet Switches . . . a Technical Brief

Introduction

GarrettCom's Link-Loss-Learn™ feature, designed to simplify and speed up recovery for Ethernet switches used in redundant LAN topologies, is implemented in Magnum™ managed and un-managed Ethernet switches. It is a unique way of getting the benefits of fault tolerance in the network while retaining the economy and simplicity and convenience of un-managed switches such as the Magnum ES42-Series for connection of network nodes.

Link-Loss-Learn addresses the time delay associated with changing the network addresses that are normally stored in a switch's memory. With patented Link-Loss-Learn, Magnum switches are able to handle most fault recovery situations, and they may improve network reliability and provide faster fault recovery accordingly. Without a redundant network topology application, the Link-Loss-Learn feature has no significant benefit and should be turned off via the applicable software command in 6K managed switches.

All of the managed Magnum 6K Ethernet Switches support Link-Loss-Learn and S-Ring. The ES42 Edge Switch Series are the first un-managed Magnum products to offer the Link-Loss-Learn feature. In typical Link-Loss-Learn fault-tolerant LAN applications, combinations of managed and unmanaged switches are used.

The Link-Loss-Learn feature is placed into operation in a managed switch under the users control. It is not "automatic," but instead it is turned on with software commands by the user at initial set-up of the management software. If Link-Loss-Learn is not turned on, the managed switch operates and functions the same as an ordinary Ethernet Switch, learning the MAC addresses of attached nodes and retaining those addresses in memory until they eventually age out or power is turned off. The Link-Loss-Learn feature will improve fault recovery in ring topologies, but it can never hurt by going into operation unexpectedly.

The user enables Link-Loss-Learn on a managed switch-by-switch basis using commands in the management software. For un-managed ES42's, the selection of ports 1 and 2 enables Link-Loss-Learn since these ports are the ones used to connect the ES42 into a redundant ring structure network. The Link-Loss-Learn feature may be enabled on any combination of copper and fiber ports in ES42's. The media choices depend upon the desired switch media connections to best serve the application.

How does Link-Loss-Learn assist with Fault Recovery?

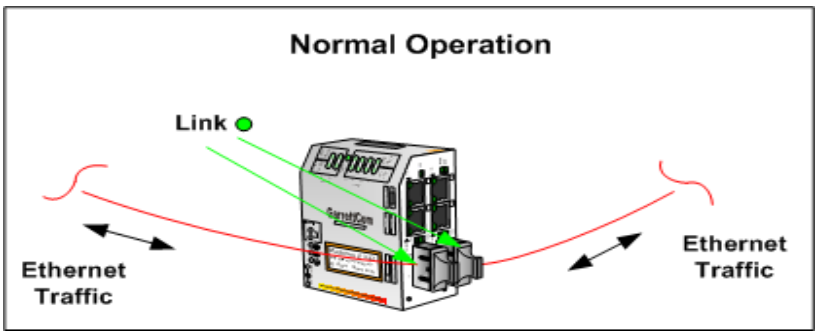

When a LAN is functioning normally, the LINK indicator is present for each port in use. A fault (a cable cut, or a unit losing power, or a unit failing while in operation) will usually cause LINK to be lost on one or more ports. Therefore, the loss of LINK during normal operation is interpreted as a signal that something has gone wrong, and in a redundant LAN, recovery operations must be brought into action to restore Ethernet traffic to its expected performance level.

© 2009, GarrettCom, Inc.

When LINK fails on a port in a redundant LAN, another back-up port is expected to take over and keep the network packets flowing. The back-up port is connected and ready to provide service. However, a normal Ethernet Switch engine will continue to use its old address table, and will continue to try to forward packets to the failed port. This will go on until the address table aging time expires for the addresses whose connection was lost (which can be as much as 3 to 5 minutes) or until nodes re-announce their presence for service purposes.

When standard 802.1d Spanning Tree Protocol is implemented, once a topology change is detected, the STP aging timer is set to 15 seconds until the topology is re-computed and / or reconfigured. The process of re-computing as well as reconfiguring the LAN can take equally as long, even in a simple set-up. (Note - complex set-ups such as multi-level meshes take much longer). For some industrial networks, this time of less than a minute for fault recovery is an acceptable delay, and standard Spanning Tree is an acceptable solution. For faster fault recovery and restoration times, Link-Loss-Learn and S-Ring can help. (RSTP-2004 provides the fastest fault recovery, typically sub-second, but requires a managed switch at each location in the network and economical un-managed switches cannot participate).

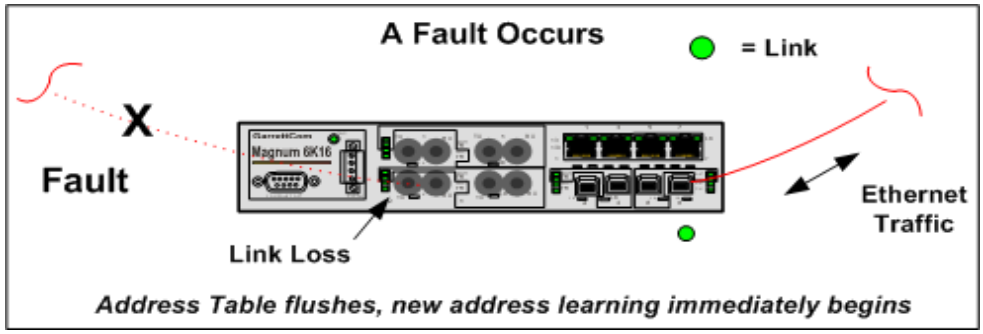

The Link-Loss-Learn feature improves the recovery time by forcing the switch's address table to be flushed when LINK is lost on any designated port. The effect on the operation of the Switch is the same as upon power up. The first packet is broadcasted and its address is learned. This continues rapidly until all addresses are learned and operation is normal . . . but with new information now in the address table on how to switch packets. Some bandwidth is used unnecessarily during the re-learning, but the recovery process is not delayed. Thus, the immediate re-learning of the addresses of attached devices results in fast re-routing of the network traffic passing through the Link-Loss-Learn switch.

Because the Link-Loss-Learn feature is very fast (it takes just a few milliseconds), the Link-Loss-Learn switch will not be the gating item for fault recovery in a redundant LAN. Whether the redundant paths upstream are controlled by 802.1d Standard Spanning Tree Protocol (STP), or by 802.1s Tagged VLAN Spanning Tree Protocol, or manually such as in a bench-test situation, the Link-Loss-Learn switch can reset its address table and participate in the LAN configuration change and network recovery faster than the other Ethernet elements.

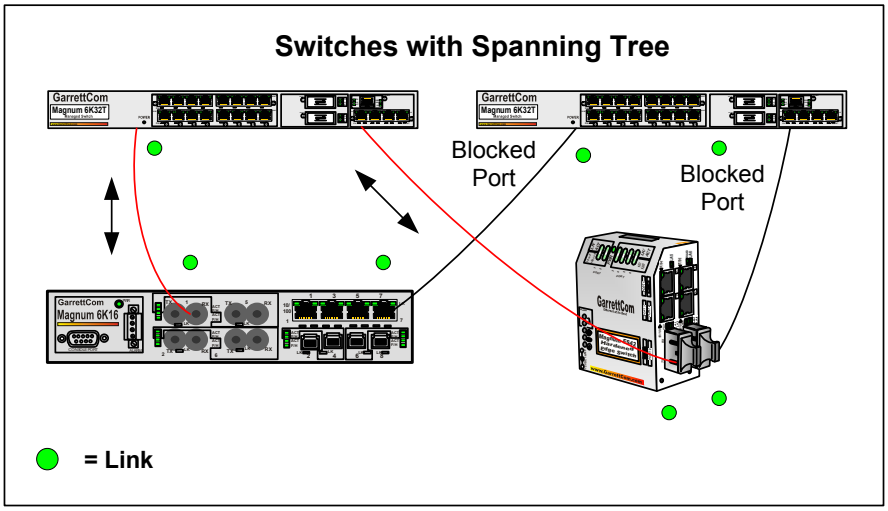

For redundant systems, simplicity is a virtue.

© 2009, GarrettCom, Inc.

Redundant LANs and fail-over scenarios are necessarily complicated, but complexity can also add to the risk that not all will go well when the critical time arrives. It is better to keep things as simple as possible and minimize complexity. The Link-Loss-Learn feature fits in with this philosophy. Typically, the Link-Loss-Learn switch is used as an edge switch in a redundant LAN configuration. While a managed 6K switch can run RSTP-2004 or Spanning Tree, and can participate in failure recovery schemes accordingly, it can also perform its role in a redundant LAN recovery scenario via the Link-Loss-Learn feature simply and independently of other things that are going on. One less thing to go wrong. Most applications are more economically served with the combination of managed and unmanaged switches running the simple Link-Loss-Learn feature than with the switches running the more complex RSTP and Spanning Tree Protocols. The Link-Loss-Learn capability allows the user to choose.

Rings and Strings, and Link-Loss-Learn with Propagation

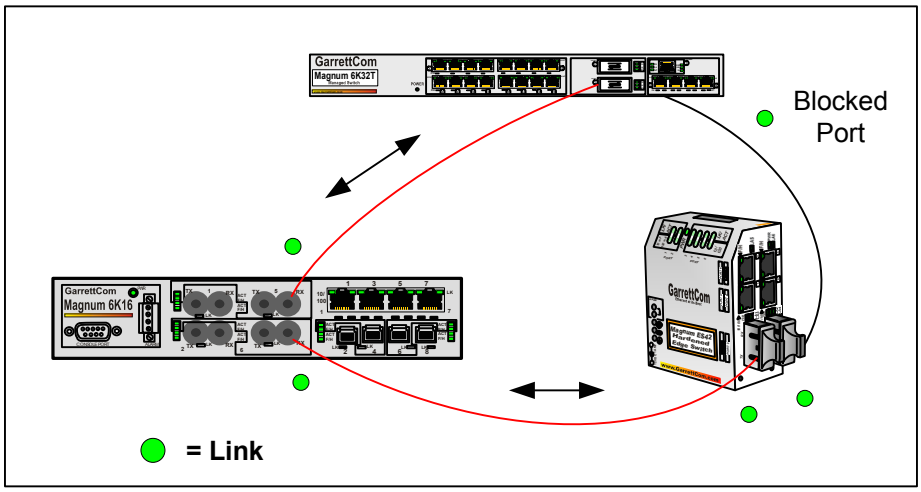

Frequently, a redundant system design using Magnum Switches as edge devices has the switches deployed over a distance and interconnected using the fiber ports connected in a string, i.e., daisy chained from one switch to the next with fiber media. It is common to continue from the last unit in a string, connecting the 2nd fiber port of the last switch in a given string back into a redundant LAN connection, thus forming a "ring" of Link-Loss-Learn units. Such a ring must have a port somewhere in the series operating blocked (i.e., not passing packets) so that a correct Ethernet topology exists. The Spanning Tree or similar logic (such as S-Ring™) controls which port is blocked in order to manage operating the network and to facilitate recovery from faults.

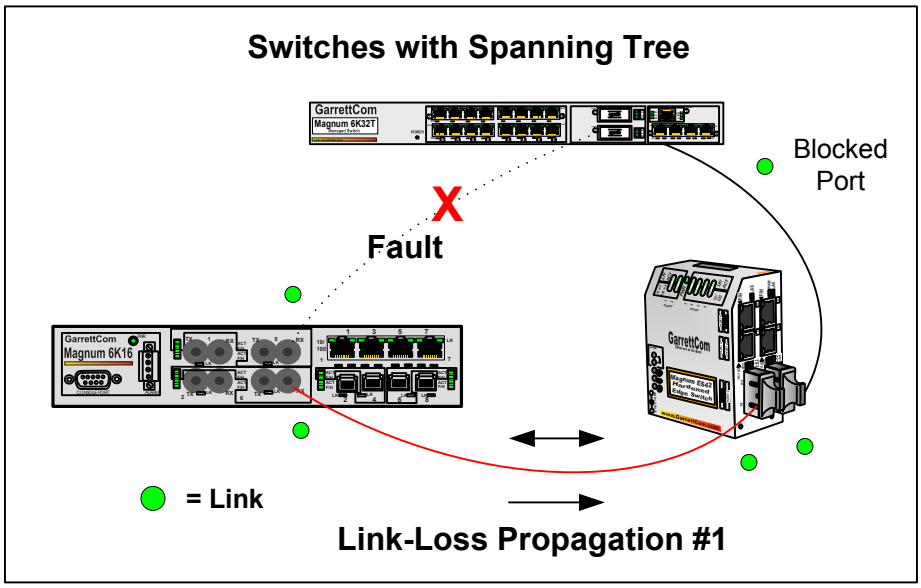

A Link-Loss-Learn switch in a string or ring could experience a link loss on a fiber port, indicating a fault has occurred and the string has been broken. Recovery action needs to take place, and a switch with the Link-Loss-Learn feature enabled would immediately dump its address table and be ready to re-learn addresses and operate in a changed network configuration. But, what about the other Link-Loss-Learn units in the string or ring . . . how will they know to do the same thing?

© 2009, GarrettCom, Inc.

For this reason, the Link-Loss-Learn feature includes a "propagation" function that, upon Link loss on one enabled port, temporarily drops Link on any other Link-Loss-Learn enabled ports to propagate the action in the units in the string or ring.

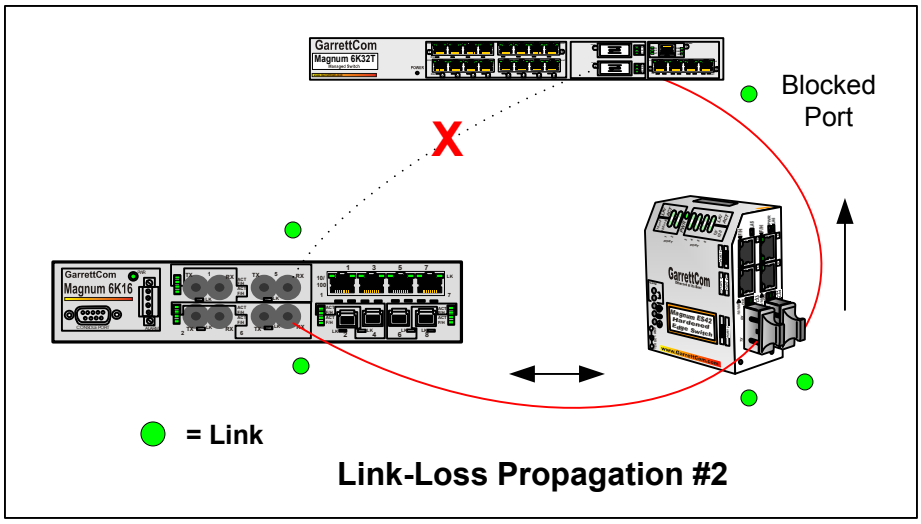

The Propagation function associated with the Link-Loss-Learn feature is always present, and its operation only affects those ports that have the Link-Loss-Learn feature enabled. Users have control over the propagation accordingly, and can control their redundant LAN set-up with Link-Loss-Learn switches to best suit their application.

Summary

Using the Link-Loss-Learn feature with Propagation, the Magnum 6K-Series with management and ES42-Series un-managed Ethernet Switches can simplify and speed up recovery from faults in redundant Ethernet LAN configurations. The Link-Loss-Learn feature is applicable to both mesh and ring or string topologies. Typically, using S-Ring and the simple Link-Loss-Learn feature will be more economical than running Spanning Tree or Rapid Spanning Tree on every switch in a redundant LAN ring structure, increasing reliability economically by reducing complexity.

©2009 GarrettCom, Inc. Printed in United States of America 10/09. GarrettCom, Inc. reserves the right to change specifications, performance characteristics and/or model offerings without notice. GarrettCom is a registered trademark, Magnum, Link-Loss-Learn (US patent # 7,593,319), S-Ring, Dymec, DynaStar, and Personal Switch are trademarks of GarrettCom, Inc. NEBS is a registered trademark of Telcordia Technologies. UL is a registered trademark of Underwriters Labs.