the Availability Digest

www.availabilitydigest.com

Shore Micros' 100-Microsecond Link Failover

March 2011

How do you fail over to a redundant component in 100 microseconds? With hardware, not software or firmware. But isn't this an expensive approach? Not with field-programmable gate arrays. This is the secret behind Shore Microsystems' Network Protection System (NPS) Link Protectors.

Ethernet Link Failures

There are many current technologies today that contribute to meeting availability requirements, all depending upon redundancy – redundant servers, redundant storage subsystems, redundant power sources, redundant networks.

Redundant networks pose a special problem. If redundant servers, storage units, or power are down, no data is being generated. However, if redundant networks are down, valid data is not being distributed and may be lost.

A redundant network may be down for two reasons – all network links are down, or an active link has failed; and the network is in the process of failing over to a backup link. Considering the latter case, link failover times are typically very fast, much faster than server or storage failover times. Link failover is typically measured in milliseconds, not seconds or minutes as is typical with highly available redundant server/storage architectures. When a ten-megabit link was once considered fast, these failover times were more than adequate. Unfortunately, in today's world of gigabit and faster links, millisecond failover times are not fast enough; and they are becoming troublesome.

Consider a one-gigabit link. A single bit is sent every nanosecond. During a forty-millisecond link failover, forty million bits, or five million bytes, will be lost. At 1,000 bytes per packet, this is a loss of 5,000 packets, or probably hundreds if not thousands of messages. Though TCP/IP messages may be recoverable, UDP messages are certainly lost. And the problem is ten times worse with the newer ten-gigabit link technologies.

Shore Microsystems' Link Protectors greatly reduce this problem by providing link failover times of 100 microseconds rather than milliseconds.

Link Protection Devices

Many network devices support dual connections to a backbone Ethernet network, such as dual NIC cards on the attached device, load balancers, switches, and hubs. Routers can detect failed links and reroute to an alternate link according to their routing tables. These devices monitor traffic on a primary link and switch over to a backup link if the primary link fails.

There are two ways to determine that a link has failed:

- There is no link signal (essentially a carrier) being received from the network.
- There is no packet traffic even though a link signal is being received (the fault may be further down the network).

Link protectors are redundant transceivers that provide a full-duplex link connection for a network-attached device such as a server, a storage subsystem (SAN or NAS), or a client terminal. They route traffic over either of two links that provide connection to the backbone network. One link is the primary link that is used during normal operation. The other link is a backup link that carries the protected device's traffic should the primary link fail.

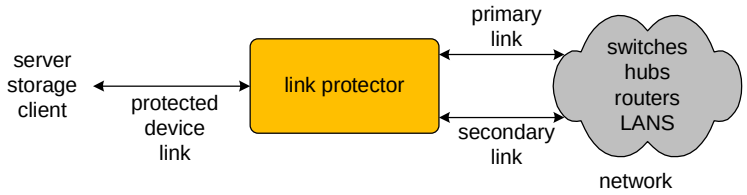

Shore Microsystems' Link Protector

Among its other products, Shore Microsystems' Ethernet Link Protectors have extremely fast failover times – 100 microseconds - that provide redundant links for mission-critical Ethernet connections. Link Protectors are a key component for designers that must respond to high-availability requirements. The company's Link Protectors support 1000BaseT, 1000BaseF, 100BaseTX, 100BaseFX, and 10BaseT Ethernet connections.

In addition, Link Protectors can broadcast simultaneous streams of data from the protected device over both primary and backup links.

Using the company's *Straight Through Wire*® technology, the Link Protectors will continue to pass traffic even in the event of a total power loss.

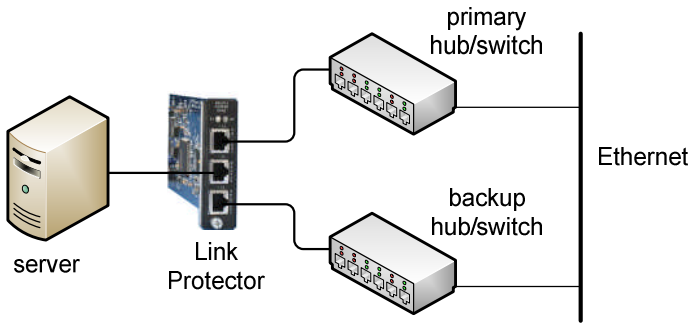

Turbo Switching

Shore Microsystems achieves such fast failover times in its Link Protectors by moving from firmware-controlled devices to Field-Programmable Gate Arrays (FPGA) to implement the Link Protector logic. An FPGA is an integrated circuit that can be configured and reconfigured in the field. Its configuration is generally specified via a high-level hardware description language. Rather than storing firmware instructions that must be executed by an operating system, FPGAs route signals through reconfigurable hardware. Thus, logic execution times are measured in nanoseconds rather than in microseconds. Failover time is measured in microseconds rather than in milliseconds.

A Link Protector monitors the link signal on its primary and backup links and will alarm if it senses a loss of signal on any link. It checks signal quality every five microseconds. If it determines that the primary port has lost its link signal, the Link Protector switches traffic to the backup link. Total time to detect a signal loss, to decide to fail over, and to switch over is accomplished in under 100 microseconds.

The Link Protector will also failover if it detects no packet traffic. A downstream network failure will not cause a link signal loss, but will prohibit any packets from being passed.

When the primary link is restored, the Link Protector will reconnect to the primary link. Alternatively, using Link Protector's Latchmode feature, the connection can be locked to the backup port following a failover. This eliminates the possibility of the connection switching back and forth if the primary link is erratic.

Power Safe

Most commercially available redundant transceivers regenerate signals electronically. Thus, these network components are a single point of failure in the event of a power interruption. Using the company's *Straight Through Wire* technology, Shore Microsystems' Link Protectors continue to pass traffic even in the event of a total power failure, provided that the primary link is still operational. Networks will typically continue to operate in the event of a power failure since the network devices include battery backups.

For copper links, the Link Protectors use high-frequency relays for power protection. Should the Link-Protector power fail, the deenergized relay directly connects the protected link to the primary link.

For fiber links, an optical switch supports the same function.

Link Protector Packaging Options

Link Protector Card

Each Link Protector is a separate card. The card provides three ports – one for the protected link and two for the primary and backup links. Multiple cards are housed in a chassis. There are two chassis options – the NPS 2 and the NPS 12. A chassis comprises a Network Protection System (NPS).

NPS 2

The NPS 2 provides space for two removable Link Protector cards in a 1U rack-mountable configuration. The chassis includes dual fixed power supplies. Link connections can be a mix of copper and fiber. The chassis contains onboard management of the Link Protectors with a port for the Network Management Console.

NPS 12

The NPS 12 provides space for twelve removable Link Protector cards in a 2U rack-mountable configuration. The chassis includes hot-swappable power supplies. Link connections can be a mix of copper and fiber. The chassis contains onboard management of the Link Protectors with a port for the Network Management Console.

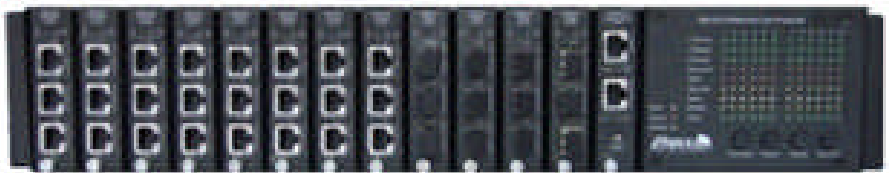

Configuration and Control

Each chassis includes a system controller (software-implemented) to configure and control the Network Protection System. All configuration can be accomplished via a Network Management Console that attaches to the management port on the chassis.

Via the Network Management Console, ports can be configured, disabled, and enabled; and the Link Protector can be forced to its backup link. Forced failover can also be commanded via SNMP, Telnet, or RS232 links.

The management facility polls each Link Protector card every few milliseconds. Alarm conditions result in visual and audible alarms in the chassis and can be sent to network management facilities via SNMP (Simple Network Management Protocol).

Headless operation

The system controller in the chassis is not needed for the Link Protectors to provide their link monitoring and failover functions. The Link Protectors will continue to provide link backup protection and power failure protection even in the absence of system management.

Restoration or rebooting of the system controller has no impact on the links. The links do not have to be reinitialized; thus, user data is unaffected.

Upgrades

System Controller

Because of the headless capability of an NPS, the system controller can be upgraded without affecting network communication. It can be taken out of service, upgraded, and rebooted with no effect on the active links.

Link Protector Cards

If the FPGA on a Link-Protector card needs to be updated, the card can be put into the power-fail bypass mode so that communication is not lost. The FPGA on the card can then up upgraded and the card returned to service without affecting communication.

Appliance Bypass Switch

Inline appliances are in use in today's networks for many reasons – security, intrusion detection, sniffing, data deduplication. These appliances generally represent a single point of failure in the network.

Via its bypass technology, Shore Microsystems protects networks against appliance failure. Using its Ethernet Bypass Switch, a card similar to its Link Protectors, a faulty appliance can be bypassed. In contrast to the Link Protector card, a Bypass Switch card has four ports. One connects to the upstream network and the other to the downstream network. The other two ports connect to the appliance.

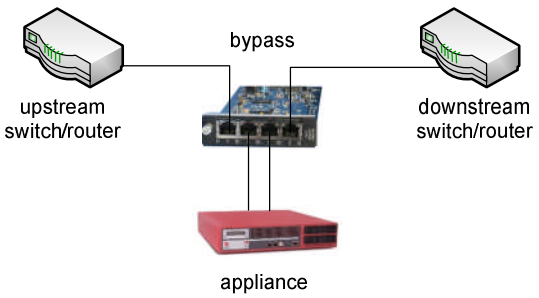

Normally, traffic flows from the upstream network through the appliance to the downstream network. However, if the appliance should fail, as determined by no outgoing traffic in the presence of incoming traffic, the Bypass Switch will disconnect the appliance and will route traffic directly to the downstream network.

Shore Microsystems' bypass technology has been used by many appliance manufacturers to protect their customers from inline appliance faults.

Shore Microsystems

Located in Long Branch, New Jersey, Shore Microsystems (www.shoremicro.com) was formed in 1984 by two Bell Laboratories engineers for the purpose of designing and manufacturing networking equipment. Since 1994, the company has focused on network devices to deliver high availability to mission-critical Ethernet networks.

The company has numerous clients in the commercial, industrial, and military communities.

Summary

Shore Microsystems specializes in network devices that add significant fault tolerance to mission-critical Ethernet networks. Its Link Protector products provide 100-microsecond failover for redundant links to ensure data flow during network failures. Shore's Bypass Switches route around failed inline network appliances, such as those used for security, to prevent network downtime due to appliance faults.

The company's products currently support link speeds up to one gigabit. It is looking to extend its product line to ten gigabit speeds.