Fibre Channel and the Autonomous SAN

By Howard Johnson, FCIA Member, Technology Architect, Broadcom

Automation, machine learning, DevOps and digital transformation are all the buzz of the modern IT infrastructure. Heralding advancements in capabilities and appearing seemingly everywhere with the promise of simplifying network operations; but what are these advancements really all about? To get an idea of "how stuff works," let's take a look at a new technology that helps these emerging technologies integrate into Fibre Channel.

Storage networks are the mainstay of major industries such as finance, healthcare, manufacturing, insurance, government, and many more. As such, Fibre Channel Storage Area Networks (SANs) have stringent reliability and availability requirements that distinguish Fibre Channel solutions from alternative technologies. The impact of an outage in mission critical industries is calculated to be about $100 every second. Imagine going "out-to-lunch" for an hour and returning to a parking meter that read over $300,000 – I'd say it's time to skip lunch!

With SANs at the core of every major business segment, the INCITS T11 standards body that creates the Fibre Channel industry standards, is continually pursuing improvements in reliability and availability. Supporting this focus, Fibre Channel architecture teams have embarked on a method of enhancing the ability of the SAN infrastructure to heal itself, such as identifying when a simple oversubscription condition exists (for example, a 32GFC storage device overdriving a 16GFC server). The first phase of this capability is embodied in a new element of the architecture known as Fabric Notifications.

What's the Buzz?

Fabric Notifications effectively creates a mechanism to surface events occurring in the fabric to the end devices that need the information to make resiliency-based decisions. For example, today when a multi-path solution determines that an IO has been impacted, it cannot tell if the impact was due to a logic error or a physical error. Logic errors are often recovered through retry or logical reset operations that are relatively non-disruptive to the system; however, physical errors often require the administrator to intervene to complete the repair. When the physical error is persistently intermittent, the task of the administrator becomes especially difficult to perform.

With Fabric Notifications, the fabric, or end device, detects the occurrence of the intermittent physical issue. It monitors the issue to see if persists, and if it does, it generates a message to be sent to the devices that are affected by the event. Instantly, the multi-path solution knows the location and nature of the physical error and can "route around" it by utilizing good, alternate paths. A clear benefit of this feature is that the administrator is not involved in the identification, isolation, and recovery of the error.

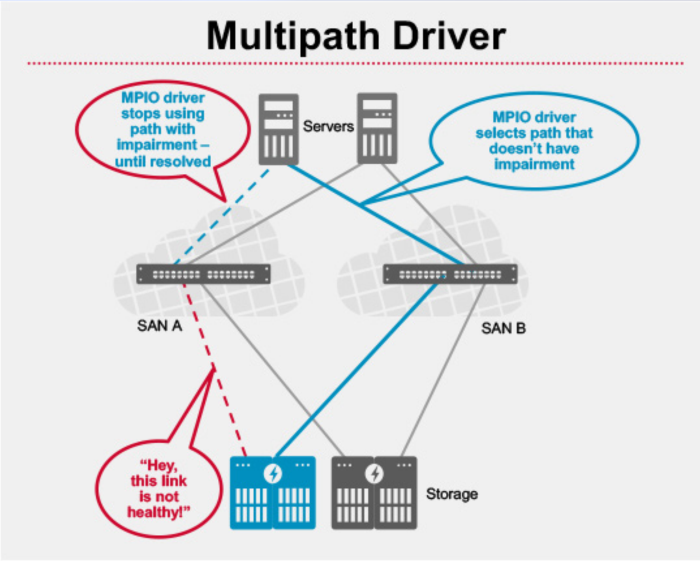
Multipath Driver
MPIO driver stops using path with impairment – until resolved
Servers
MPIO driver selects path that doesn't have impairment
SAN A
SAN B
"Hey, this link is not healthy!"
Storage

Tell Me More, Tell Me More

Let's break down the elements of Fabric Notifications to see how it actually works. There are two main functions provided with this architecture – signals and notifications. The signals are physical layer operations that allow the two ends of a link to indicate when transmission resources are being consumed to the point of causing problems in the fabric or internal to the device. The notifications are logical layer operations that allow the fabric or devices to notify each other of events they have detected and how those events might affect their logically connected peers.

These mechanisms are fully controlled by the end devices through an exchange of capabilities and the registration of operations. This approach ensures that only the devices that are capable of managing and responding to these new messages ever receive them. Hence, a device that wants to enable the Congestion Signal mechanism, exchanges that capability with the fabric port and the fabric port does the same. Upon successful completion of the exchange of capabilities, the two sides enable their logic to generate the signals based on their specific implementations. Likewise, a device that wants to enable the Notification mechanism registers to receive the notification Extended Link Service (ELS) and includes a list of the specific sub-functions they want to receive. In this way, the end device completely manages its participation level as well as controls the degree of participation.

When the exchange and registration process is completed, the devices are ready to receive signals and notifications from the fabric that will help them effectively evaluate and isolate errors. The devices can now make intelligent decisions for automated corrective actions (that is, they can engage in machine learning techniques that improve reliability, availability, and resiliency). Essentially, Fabric Notifications adds more information into the system and distributes that information to the participating devices. The devices can then make more informed decisions and respond intelligently to the events in the system. Furthermore, this mechanism allows solutions to manage their level of participation and only react to the events they care about – and only to the degree to which it helps, without impacting their operations.

What's a Device to Do?

The beauty of the Fibre Channel Fabric Notifications architecture is its simplicity and control. Since the participating devices control the level and degree of participation, solutions can evolve over time and adjust participation based on experience in the enterprise.

Basic solutions may choose to simply register to receive all notifications and just log them. For a SAN administrator, this simple action reduces problem determination and isolation considerably because the fabric and device logs describe exactly where the event occurred in the fabric and which devices are affected.

A slightly more advanced solution may choose to handle congestion conditions in the fabric and register for Peer Congestion Notifications. The device may then take the simple action of speed matching for flows to a device that exhibits congested behavior (for example, our 32GFC storage to 16GFC server connection automatically adjusts the flow to be 16GFC to 16GFC).

In our multi-path example, a solution may register for Link Integrity Notifications to allow the device to avoid paths that are automatically identified by the fabric to be intermittently impaired without requiring the SAN administrator to intervene.

Fibre Channel Autonomous SANs

The position of the Fibre Channel switch in the data path allows it to have visibility into the surrounding components and gather intelligence not only on the infrastructure, but the entire storage network, including the attached devices. This intelligence is exchanged amongst the switches of a fabric to create the vision of an autonomous SAN that is comprised of self-learning, self-optimizing, and self-healing activities. Fabric Notifications enable these characteristics and provide the foundation for delivering autonomous infrastructures. These capabilities eliminate countless wasted cycles of manual analysis, which becomes very challenging with the scale of modern data centers. Furthermore, Fabric Notifications reduce the period of time that applications perform at sub-par levels resulting in lost revenue or degraded service.

The Fibre Channel industry's Fabric Notifications architecture is the technology that supports Autonomous SAN technology with the capability to automatically identify and resolve issues. The fabric identifies data traffic congestion and facilitates automatic failover or traffic adjustment by notifying the devices. The devices in turn, take automatic corrective actions to mitigate the impact of the congestion or persistent, intermittent failures.

Imagine an environment where the network mitigates issues automatically and flags the troubled component for resolution later. It's technologies like Fibre Channel Fabric Notifications that enable infrastructures to manage themselves and brings us closer to truly autonomous SANs.

[i] *The capabilities exchange occurs when the N_Port sends the Exchange Diagnostic Capabilities (EDC) ELS command with a Congestion Signal Capabilities (CSD) descriptor in the payload to enable Congestion Signal processing with the F_Port to which it is attached. The registration process occurs when the N_Port sends the Register Diagnostic Functions (RDF) ELS with the Fabric Performance Impact Notification (FPIN) descriptor that contains a list of descriptor tags corresponding to the specific FPIN sub-functions it wants to receive. The sub-functions include descriptors that enable Link Integrity Notification (FPIN-LI), Delivery Notification (FPIN-DN), Peer Congestion Notification (FPIN-PN), and Congestion Notification (FPIN-CN) sub-functions.*