Technical Review

Ensuring Application Performance and Reliability with Brocade Gen 7 SAN

Date: May 2021 **Author:** Tony Palmer, Senior Validation Analyst

Abstract

This ESG Technical Review evaluates Brocade Gen 7 next generation SAN, which uses autonomous technology to ensure application performance and reliability across an organization's SAN. Hands-on testing was used to validate how Brocade Gen 7 Autonomous SAN technology automatically detects and takes action against issues to improve application availability and reliability through self-learning, self-optimizing, and self-healing capabilities. ESG used three common use cases to demonstrate how proactive traffic optimization and end-device fabric performance impact notifications to address common congestion and physical layer issues that can arise in a SAN.

The Challenges

According to ESG research, three-quarters (75%) of organizations surveyed view their IT environments as more complex than they were two years ago (see Figure 1). Although numerous factors contribute to such complexity, ESG research uncovered that more than a third (38%) of respondents cited higher data volumes as the reason for their organization's increased complexity.[1] While organizations may not welcome the resulting complexity in their environments, they do realize that storing and managing the data is important to the business as 71% consider data storage technology to be strategic to their operations.[2]

Figure 1. Most Primary Storage Remains On-premises

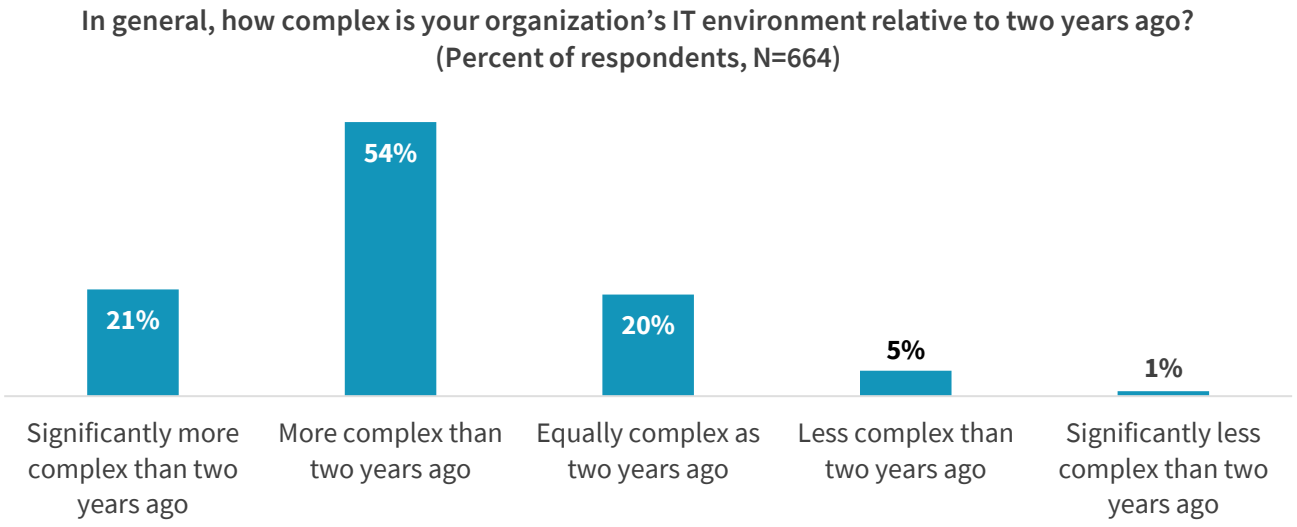

Source: Enterprise Strategy Group

Fibre Channel has been the go-to architecture for high-performance SANs in mission-critical environments for decades. ESG research confirms that more than half (53%) of organizations surveyed said that they expected spending on Fibre Channel

[1] Source: ESG Research Report, *2021 Technology Spending Intentions Survey*, January 2021.
[2] Source: ESG Master Survey Results, *2019 Data Storage Trends*, November 2019.

storage area networks (SANs) located in company data centers to increase. More than one in three organizations of all types and sizes rely on Fibre Channel (FC) SANs.[3]

The environments where they are deployed and the requirements for availability, stability, and performance that FC networks are expected to satisfy make them extremely sensitive to various network congestion and component or endpoint failure or degradation issues. The storage network must be able to maximize the data performance and efficiency of the entire IT storage ecosystem, including multiple generations of storage technology, in order to meet expectations.

Congestion occurs when the rate of frames entering the fabric exceeds the rate of frames exiting the fabric. FC SANs can experience congestion for a variety of reasons, including applications with high bursts of storage activity and accidentally oversubscribed SAN ports. New storage technologies and application workloads, including faster host bus adaptors (HBA), SAN switches, and all-flash storage arrays, mixed with older servers, can move storage bottlenecks into SANs and servers. In addition, pervasive NVMe adoption to improve performance of existing applications and support new, more demanding workloads can further shift bottlenecks.

FC SAN congestion that impacts application-level performance is a hard problem to identify and fix. Multipath I/O (MPIO) addresses hard errors and effectively fails traffic over to redundant paths. MPIO systems struggle to address intermittent errors, which significantly slow down application I/O due to the sequential nature of many application operations.

Troubleshooting and resolving SAN congestion has historically been manual and labor-intensive, such as manually adjusting a queue depth limit in an HBA or setting a hard performance limit with a quality of service (QoS) algorithm. The new Brocade Gen 7 Autonomous SAN is designed to address this market demand by providing FC SAN automation and intelligence. Recently, ESG evaluated a Brocade Gen 7 Fibre Channel environment and its ability to automatically mitigate three common SAN issues that can significantly impact application performance:

- Speed mismatch issues that propagate through the storage network.
- Congestion caused by oversubscription of host or storage endpoints.
- Physical link layer impairment that impacts application performance.

The Solution: Brocade Gen 7 Autonomous SAN

Brocade Gen 7 SAN solutions provide a number of enhanced network monitoring and analytics features. As shown in Figure 2, the Gen 7 product set includes Brocade FC Switches, Emulex HBAs, and supporting orchestration/management software. These solutions are engineered with the performance needed to support end-to-end NVMe technology and the ability to produce and utilize comprehensive telemetry data across the fabric to analyze and take actions based on that data to optimize the storage network automatically.

[3] Ibid.

Figure 2. Brocade Gen 7 Autonomous Self-healing SAN

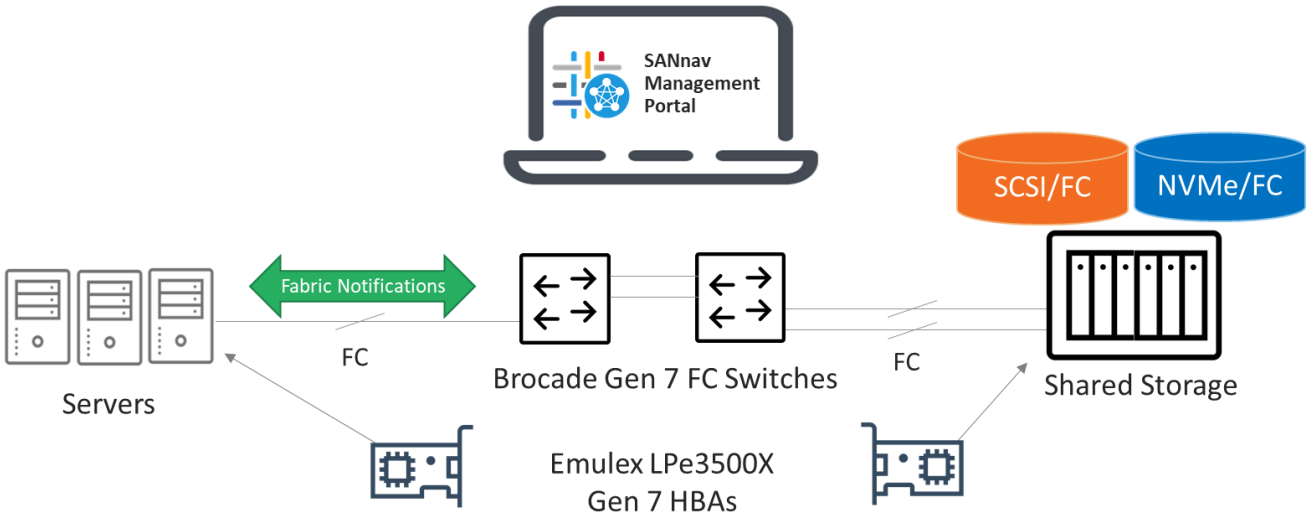

Source: Enterprise Strategy Group

Implemented in Brocade Fabric OS 9.x software, Brocade Gen 7 hardware, equipped with the latest Brocade ASICs and in Emulex Gen 7 HBAs autonomous SAN technology, enables a Brocade switch that automatically detects SAN congestion issues and propagates Fabric Performance Impact Notifications (FPINs) to tell the HBAs which paths are congested and need to be remediated.

The Fibre Channel SAN use cases evaluated in this report all leverage the intelligence built into Brocade Gen 7 Autonomous SANs to solve common storage network congestion issues. When congestion occurs, it can be pervasive, sometimes affecting hundreds of unrelated flows. SANs can be challenging to troubleshoot because of the magnitude and density of flows commonly seen from hypervisors, which abstract the applications (VMs), and end devices don't have visibility into the SAN. Within the SAN, Brocade has introduced FPIN, a hardware, software, and management solution for achieving real-time congestion identification, notification, and mitigation.

In collaboration with end devices, FPIN is designed to solve fabric performance issues. Brocade switches that detect SAN congestion using Fabric Notifications send FPINs to tell the HBAs which paths are congested and need to be remediated. Broadcom played a key role in the development of the FPIN specification. The motivation for this effort was simple: Brocade, a Broadcom Company was commonly told by customers that finding and fixing SAN congestion problems was a top SAN management challenge because the impacted applications are often not the root cause of the congestion. This often leads to investigations taking longer because they start in "the wrong place." FPINs are a key technology involved in two of the three Fibre Channel SAN use cases analyzed and evaluated by ESG.

ESG Tested

Broadcom provided access to a Brocade Gen 7 SAN environment, allowing ESG analysts to exercise, analyze, and evaluate the self-learning, self-optimizing, and self-healing capabilities of the Brocade Gen 7 Autonomous SAN. To exercise the Gen 7 functionality in ways that offer better evaluation and deeper insights, ESG analysts focused on three common issues that often lead to congestion and can impact application performance within Fibre Channel SANs.

Brocade Gen 7 Traffic Optimizer

First, we looked at speed mismatch congestion issues propagating through the SAN. Speed mismatch is a real issue in data centers today, where multiple generations of HBAs supporting speeds from 8G to 64GFC all need to operate side by side without impact. The testing was performed in two environments: one with Brocade Gen 6 FC switches and servers with

Emulex Gen 5 and Gen 7 HBAs and the other with Brocade Gen 7 FC switches and servers with Gen 5 and Gen 7 HBAs. The industry-standard FIO benchmark utility was used to simulate the large block (1MB) sequential read traffic of the "culprit" (high-throughput demand) and "victim" (lower-throughput demand) workloads.

Figure 3. Speed Mismatch Congestion

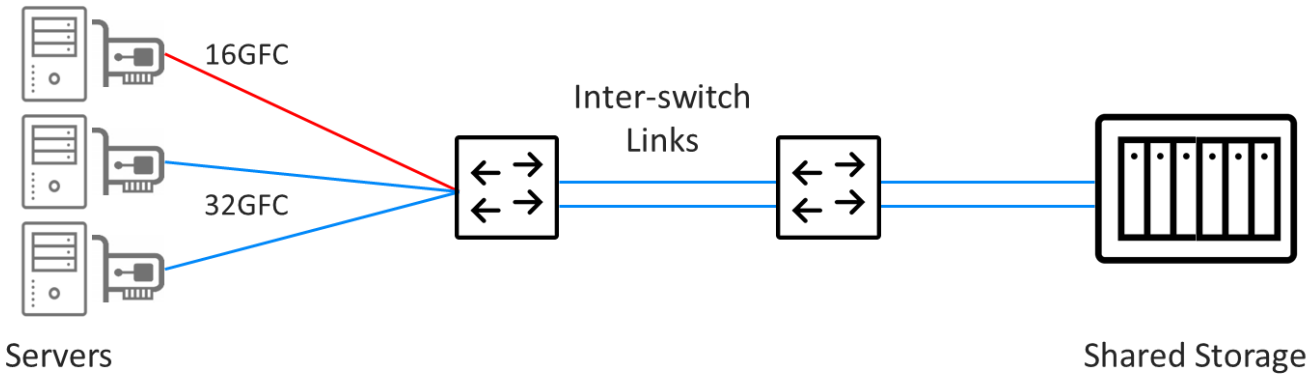

Source: Enterprise Strategy Group

In Brocade SANs, virtual channels are used to divide traffic and minimize the impact when congestion happens to a single virtual channel. In Gen 6 (and earlier versions), virtual channels were statically assigned based on the physical location where devices were connected, so SAN admins needed to design their SANs carefully to ensure storage ports were on different virtual channels. In Gen 7, there is a dynamic assignment of virtual channels based on traffic speeds and traffic characteristics to proactively avoid mismatches, and only the same type of traffic will share the same virtual channels.

ESG used Brocade SANnav Management Portal to monitor performance and SAN congestion during each phase of this test. Congestion was detected when performance for the higher-throughput workloads dropped substantially.

SAN congestion was flagged as a problem through the Brocade SANnav Management Portal Health Summary as an event that needed investigation. Through the network topology view, Brocade SANnav enabled ESG to drill down and explore the cause of the alerts and identify which HBA ports and switch paths were fighting for bandwidth.

Figure 4. Speed Mismatch Congestion—Automatic Remediation

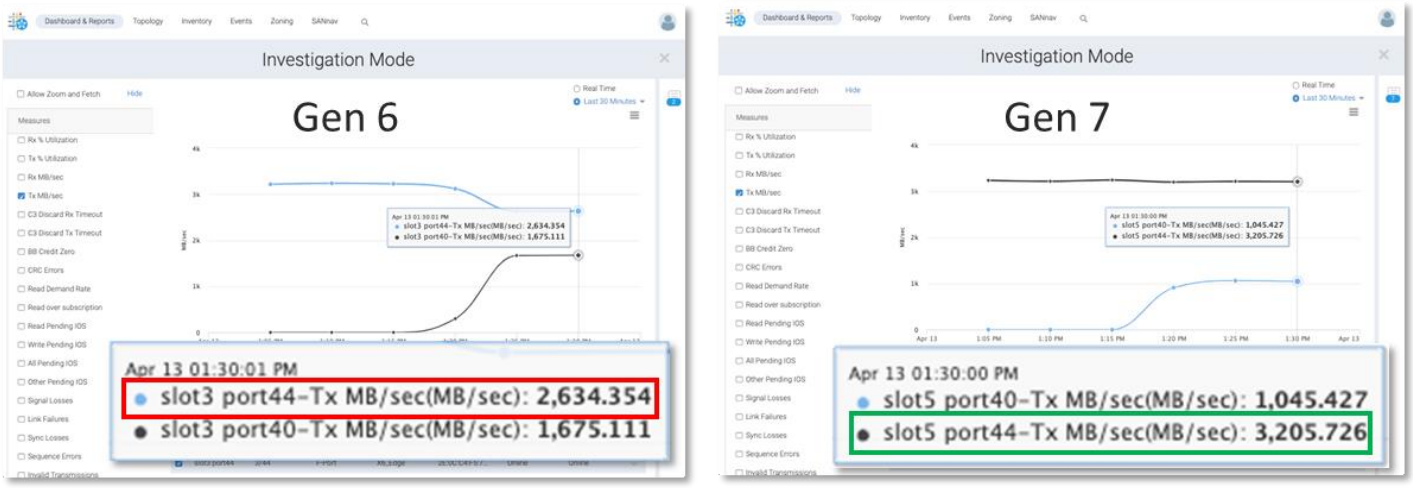

Source: Enterprise Strategy Group

Investigation revealed that in the Gen 6 environment, as soon as the workload started on the 16G host, the 32G host's throughput began to drop—as shown by the black line—and settled at 2.6GB/sec, as seen in the image on the left in Figure 4, because the two servers were sharing the same virtual channel.

In contrast, in the Gen 7 environment, Traffic Optimizer proactively places the 16G and 32G traffic on separate virtual channels. As seen in the image on the right, this happened completely autonomously, with no intervention on ESG's part required. As a result, the 32G host was not impacted by the slower and congested 16G flow.

Why This Matters

SAN performance issues that result from speed mismatches between servers and storage are a common challenge in modern Fibre Channel storage networks supporting multiple generations of equipment. Previous generation SAN environments require administrators to carefully plan the physical topology and manually throttle hosts in order to optimize traffic to high priority applications—a labor-intensive effort with decidedly less than real-time responsiveness.

ESG confirmed that Traffic Optimizer functionality in the Brocade Gen 7 Autonomous SAN proactively separates traffic with different profiles to avoid congestion on Interswitch Links (ISLs). The functionality worked flawlessly in real time.

Traffic Optimizer proactively optimizes applications' storage performance, a sophisticated approach to avoid SAN congestion compared to manually configuring quality of service (QoS) settings on the SAN or in a storage array to set a hard performance limit for a lower-priority application. The Brocade self-optimizing approach is not only more sophisticated; it also doesn't require any manual intervention and works regardless of server and storage vendor.

FPIN Congestion Notifications

Next, ESG examined how Brocade Gen 7 SAN handles congestion caused by oversubscription. Oversubscription occurs when the incoming data rate is faster than the outgoing data rate, causing the receiving device to return buffer credits at a slower pace than necessary to run at full throughput. An insufficient buffer credit release pace from a server (or storage) slows data transmission, causing upstream backpressure. For example, the aggregate data from multiple sending ports exceeds some port's ability along the path to accommodate it. When oversubscription impedes multiple flows, a fabric-wide problem could occur.

FPIN congestion notifications are valuable information for any devices that support FPIN and can act accordingly. A common congestion scenario for oversubscription is when a server receives more data than it can process at the pace the data is received. This is instantly identified in the fabric and FPIN congestion notifications are sent to the server and storage to inform them that their data transmissions are creating congestion and have an adverse impact on network performance. In general, congestion notifications happen when applications perform sluggishly—or in the worst case are non-responsive, indicating why exchange completion times may be long.

ESG carefully examined what happens when a Brocade Gen 7 SAN encounters congestion. First, we confirmed that all hosts' HBAs had full performance, then we started up a workload on a machine, which did not create congestion but soon would become the victim of congestion. It was processing data at 1.1 GB/sec. Next, we started up a workload on machine 2, which created congestion (the culprit) and was driving 3.4 GB/sec of throughput (nearly full line rate). We then see Brocade FPI (Fabric Performance Impact monitor) detect oversubscription via the trigger of the port bandwidth oversubscribed condition and its visualization on the Network Port Traffic Conditions (NPTC) dashboard (see Figure 5).

Figure 5. The Network Port Traffic Conditions (NPTC) Dashboard

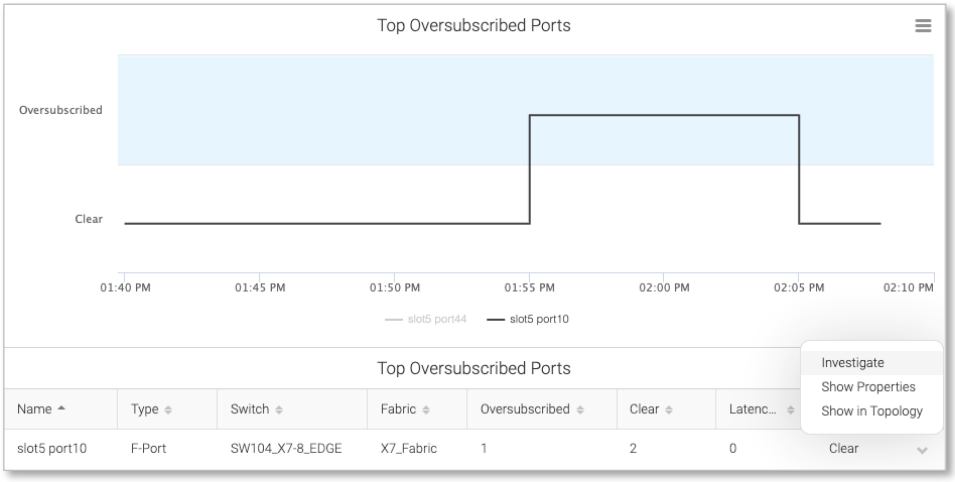

Source: Enterprise Strategy Group

When we checked the victim machine, we found that its performance was heavily impacted. Its throughput had been cut by nearly half to 620 MB/sec, while response time increased. Brocade was then able to send FPIN Congestion Notifications from the fabric to the HBA to take appropriate action and mitigate the impact to neighboring devices. The HBA throttled the rate at which the culprit device was requesting data enough to completely relieve the impact to the victim.

Figure 6. Insights from Brocade SANnav on Why Oversubscription Occurred

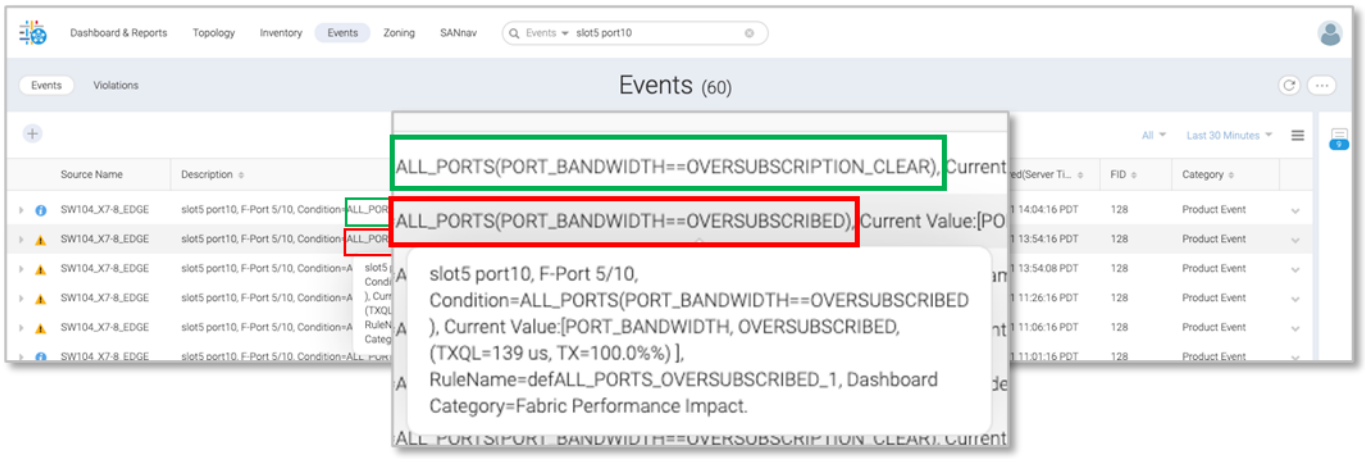

Source: Enterprise Strategy Group

Brocade SANnav Management Portal was able to visualize the oversubscription state via the NPTC Dashboard, presenting options to further investigate the impact as seen in Figure 5, and detect congestion caused by oversubscription when it occurred, issue relevant alerts (Figure 6).

With congestion management enabled on the HBAs, everything happened completely automatically, and no further intervention on ESG's part was required. In a production environment, this feature can be enabled by default.

FPIN Link Integrity

The impact of questionable components along a SAN path can lead to application degradation, crashes, outages, and lost revenue. The integrity of the link between switch ports and an end device is vital for proper operation. Finding and correcting a failed component is not overly tricky, but finding and fixing components with intermittent issues is quite

challenging. Link integrity symptoms frequently manifest fabric-wide rather than at the actual problem location. Moreover, unidirectional link integrity from a switch port to a device port cannot be detected by the fabric itself; errors are detected by the receiver and must then be reported to the fabric.

FPIN link integrity notifications address link integrity issues by identifying degraded paths throughout the SAN. ESG tested and evaluated the ability of the Brocade Gen 7 SAN to detect link impairments and take actions to mitigate the application performance impacts that result from them. Figure 7 shows a simple diagram of the simulated network with a link impairment between the machine and the first switch on one of its paths across the SAN.

Figure 7. Link Integrity Issues

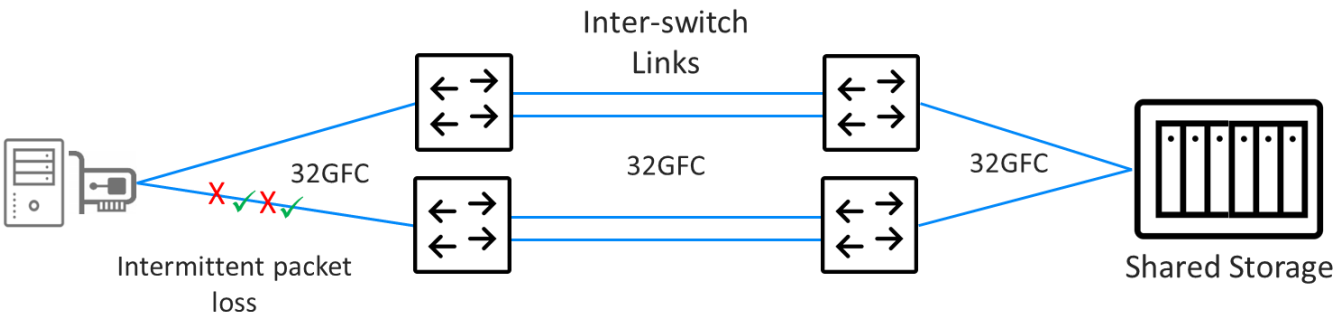

Source: Enterprise Strategy Group

To highlight the effects of a network link issue, ESG testing used the Medusa Labs Test Tools Suite, an application workload test tool generating storage traffic between server and storage array, then set up a link impairment simulation device at the network location shown in Figure 7 to simulate the issue. The device caused the network to drop a single data packet every ten seconds. An initial test was run for three minutes without FPIN enabled.

Figure 8. SAN and Application Performance without FPIN

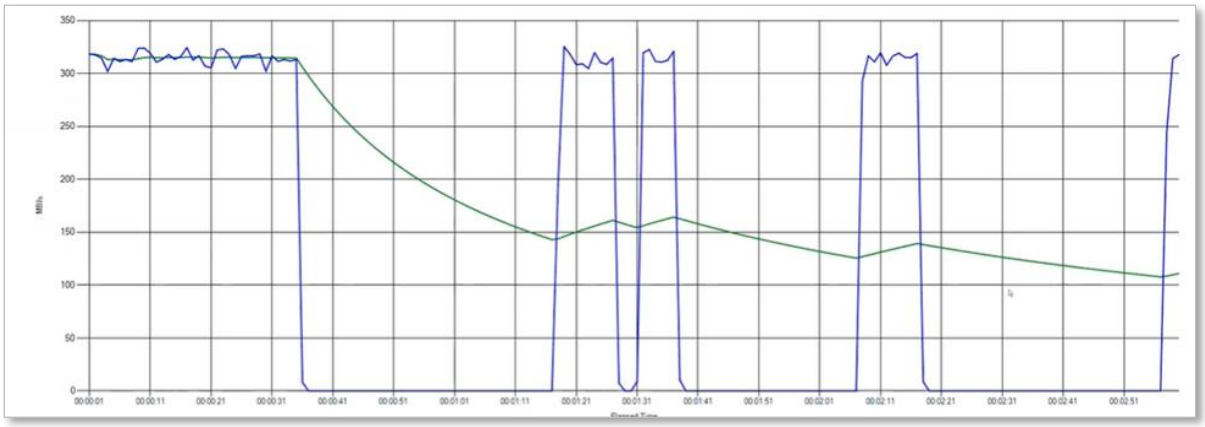

Source: Enterprise Strategy Group

Figure 8 shows the network and application performance results of this initial test. The very first instance of network impairment caused a SCSI timeout; the network traffic literally dropped off the graph. It eventually resumed, but each time a frame was dropped, the consequences for the network were severe and ongoing because the issues at the link level were not being detected by the MPIO driver, which would otherwise reroute network traffic. The green trace—average performance—clearly demonstrates application performance degradation.

The situation illustrated in the graph of Figure 8 would be difficult to diagnose in a traditional SAN environment and cause real impacts to application users. In short, link impairments cost organizations business, money, and reputation. When

applications don't respond fast or reliably enough, it will drive customers to a competitor or an alternate solution. ESG then ran the same test profile again, this time with FPIN enabled. The network and performance results are shown in **Error! Reference source not found.**.

Figure 9. SAN and Application Performance with FPIN Enabled

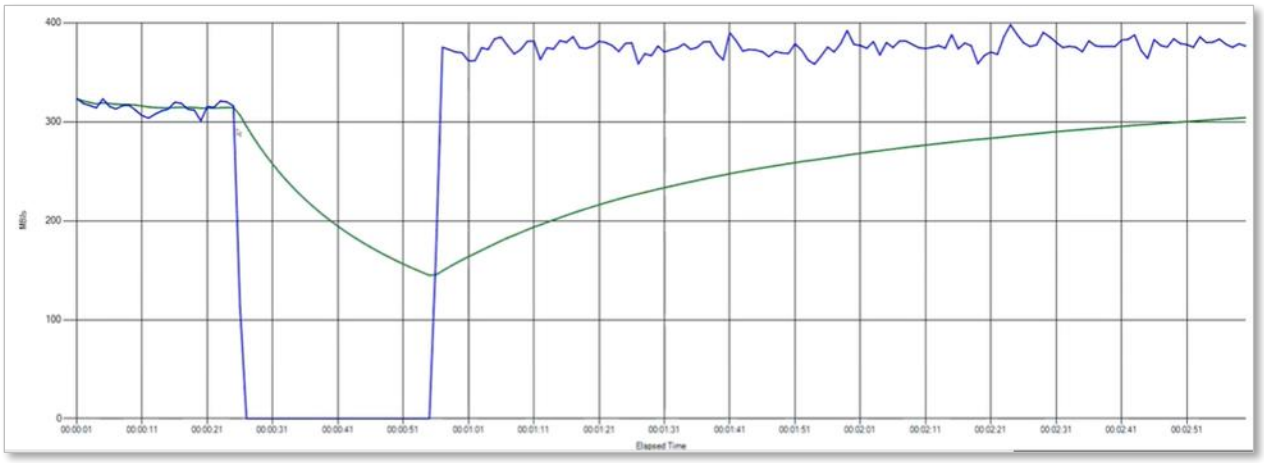

Source: Enterprise Strategy Group

As before, the first network link impairment caused a single SCSI timeout, but FPINs were sent by network devices to the MPIO layer, which responded by rerouting network traffic to a healthy path. Network performance quickly recovered, then slightly improved because it was no longer impacted by the slight cost of load balancing between network paths. Application performance also dipped after the first frame drop but then steadily recovered and stabilized, returning to pre-impairment levels. Application users might never notice any performance degradation. Meanwhile, SAN administrators wouldn't have to respond to a crisis; they receive the FPIN alert logs, then can troubleshoot and fix the problem when appropriate.

Why This Matters

The impact of questionable components along a SAN path can lead to application degradation, crashes, outages, and lost revenue. The integrity of the link between switch ports and an end device is vital for proper operation. Finding and correcting a failed component is not overly tricky, but finding and fixing components with intermittent issues is quite challenging and time consuming.

ESG Link Integrity testing clearly demonstrates that the self-learning, self-optimizing, and self-healing capabilities of the Brocade Gen 7 SAN can make a real difference in business operations. Before the autonomous SAN, a link impairment such as the one simulated in the ESG test environment (one dropped frame every ten seconds) would have been difficult to detect. Long before SAN administrators could successfully troubleshoot and fix the problem, business operations could have been substantially affected, depending on the criticality of the application involved.

Brocade's Gen 7 enhancements enable the SAN to detect this issue automatically, alert network management, and automatically mitigate the issue until it can be permanently resolved by administrators. No applications are impacted, and no operations or revenue streams are disrupted.

The Bigger Truth

In high performance storage networks—like all networks— congestion happens and can become a barrier to business efficiency, reliable application performance, and in extreme cases, application availability. Administrators needed to monitor the SAN to see if, when, and where SAN congestion was happening. More often than not, the application performance issue doesn't look like it has anything to do with the SAN, which makes it more difficult to pinpoint. After a SAN congestion issue has been detected and isolated, most of the traditional methods for remediation are manual and can't be adjusted quickly with changing traffic conditions.

ESG has validated that the latest Brocade Gen 7 switch hardware, Brocade Fabric OS version 9.0 software, and Emulex Gen 7 HBAs have leveraged the FPIN specification to automatically find and fix SAN congestion problems. Detecting and troubleshooting SAN congestion with Brocade SANnav Management Portal was intuitive and easy. Less than a minute after enabling self-healing, workloads had been throttled back precisely enough to restore the performance of a simulated business-critical data warehouse application.

ESG looks forward to seeing how customers respond to this new self-healing SAN capability. Any type of IT infrastructure change that introduces automation of an existing manual process is usually deployed with a touch of human observation and a set of approval processes. Brocade's new SAN congestion management feature can be tested in monitor mode before enabling autonomous self-healing. Based on the results of our validation testing, ESG believes that autonomous self-healing mode should be quickly embraced by IT, SAN, and storage administrators.

In the version we tested, the Brocade Gen 7 solution addressed speed mismatches quickly and automatically. Grouping traffic based on protocol is the next logical step, since SCSI and NVMe are such different protocols, and NVMe devices have much lower response time than SCSI-based devices.

If your organization relies on a Fibre Channel SAN to keep your business-critical applications running at peak performance levels, ESG believes that you should consider the IT productivity and bottom-line business benefits of eliminating SAN congestion problems with Brocade Gen 7 SANs and Emulex Gen 7 HBAs.