

Technical Review

Broadcom Autonomous Self-healing SANs

Date: September, 2020 **Author:** Brian Garrett, VP, Validation Services

Abstract

This ESG Technical Review demonstrates how Broadcom has leveraged a new Fibre Channel (FC) standard and collaboration between Brocade switches and Emulex HBAs (Brocade and Emulex are divisions of Broadcom) to deliver autonomous, self-healing SAN capabilities that minimize the application performance impact of SAN congestion.

Background

Data storage is playing a more strategic role in enterprise IT. Storage networks are being pushed to the limit as the vast majority of organizations (84%) leverage flash storage to meet the demanding performance requirements of legacy and emerging applications (e.g., analytics, artificial intelligence, and machine learning).¹ Better storage automation is the capability that the largest percentage of organizations indicated would increase the likelihood that they would repatriate their public cloud-resident storage workloads back on-premises—due to a variety of public cloud data storage concerns including security, compliance, performance, and cost.²

Figure 1. Data Storage Trends



71%

Organizations that indicated that data storage plays a **strategic** role in their IT and business operations



84%

Organizations that indicated that they leverage **flash storage** for their on-premises applications/workloads



#1

The most desired storage capability for moving cloud-resident storage on-premises: **better automation**

Source: Enterprise Strategy Group

SAN congestion that impacts application-level performance is a hard problem to identify and fix. The lossless FC storage area network that sits between servers and storage systems is a shared resource that can get congested for a variety of reasons, including applications with high bursts of storage activity (aka “noisy neighbors”) and accidentally oversubscribed SAN ports. The SAN congestion problem is getting worse for many organizations due to a perfect storm of new storage technologies and application workloads, including faster HBAs and SAN switches, all-flash storage arrays that are moving storage bottlenecks into SANs and servers, server virtualization, AI, and ML.

SAN congestion has historically been addressed with a heavy-handed approach, such as manually adjusting a queue depth limit in a host bus adapter driver or setting a hard performance limit with a quality of service (QoS) algorithm.

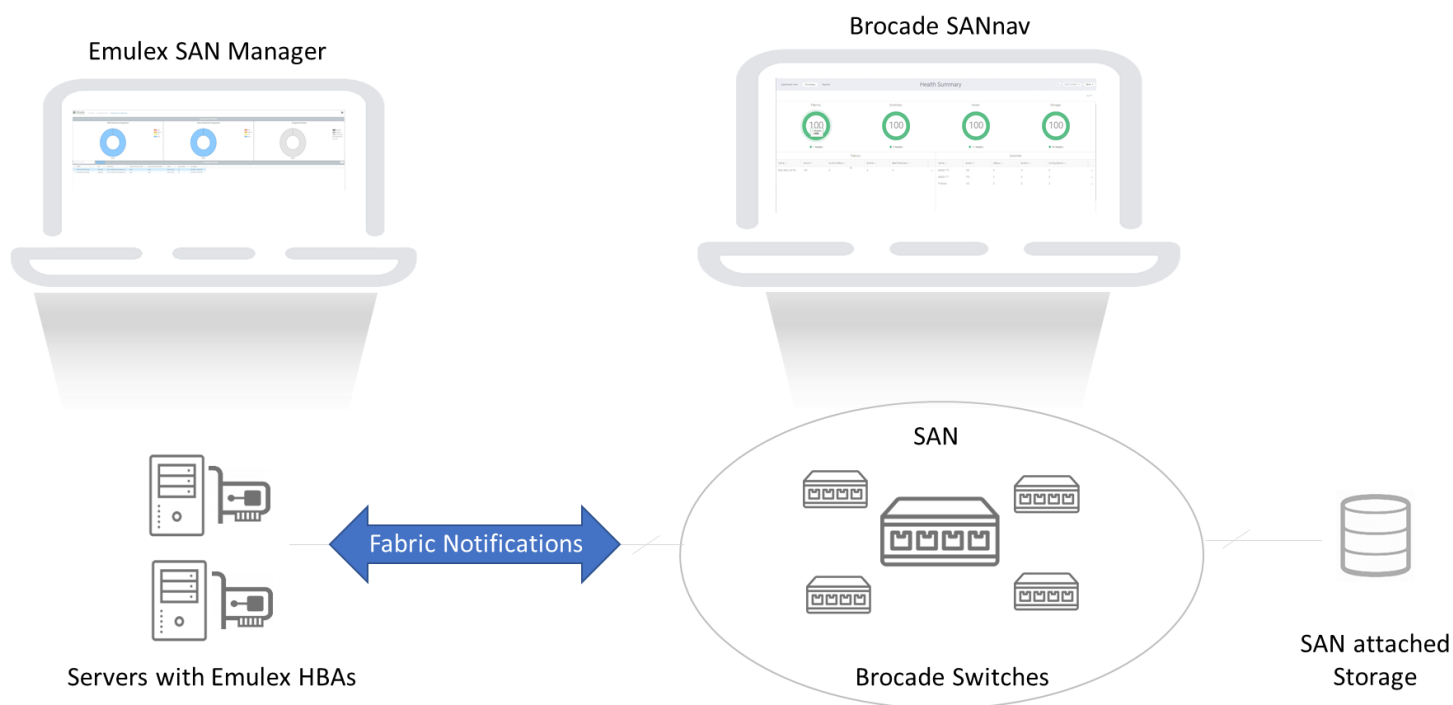
¹ Source: ESG Master Survey Results, [2019 Data Storage Trends](#), November 2019.

² *ibid*

Autonomous Self-healing SAN Management

Broadcom’s autonomous SAN self-healing capability detects and automatically fixes SAN congestion problems. Autonomous SAN technology is implemented in Brocade Fabric OS 9.x software and Brocade GEN 7 hardware that are equipped with the latest Brocade ASICs. Autonomous SAN technology is implemented in Emulex Gen 7 HBAs. A Brocade switch that has detected a SAN congestion issue uses the recently approved INCITS/T11 specification updates to include Fabric Performance Impact Notification (FPIN)³ to tell Emulex HBAs which paths are congested and need to be remediated. As shown in Figure 2, Brocade switches and Emulex HBAs work together to detect, diagnose, and fix SAN congestion problems with FPINs.

Figure 2. Broadcom Autonomous Self-healing SAN Management



Source: Enterprise Strategy Group

Broadcom played a key role in the development of the FPIN specification. The motivation for this effort was simple: Broadcom customers consistently reported that finding and fixing SAN congestion problems was a top SAN management challenge.

ESG Testing

ESG tested Broadcom’s ability to detect and autonomously recover from SAN congestion by simulating the performance impact of an I/O-intensive workload running on the same SAN as a business-critical data warehouse application. The testing was performed on Brocade Gen 7 FC switches, storage arrays, and Emulex Gen 7 Host Bus Adapters. A deliberate speed mismatch between the flash array and server HBAs was used to simulate a congestion slowdown for a higher priority data warehouse application (aka “the victim”) after a throughput-intensive workload had started (aka “the bully”). The industry-standard FIO benchmark utility was used to simulate the large block (1MB) sequential read traffic of the bully and victim workloads.

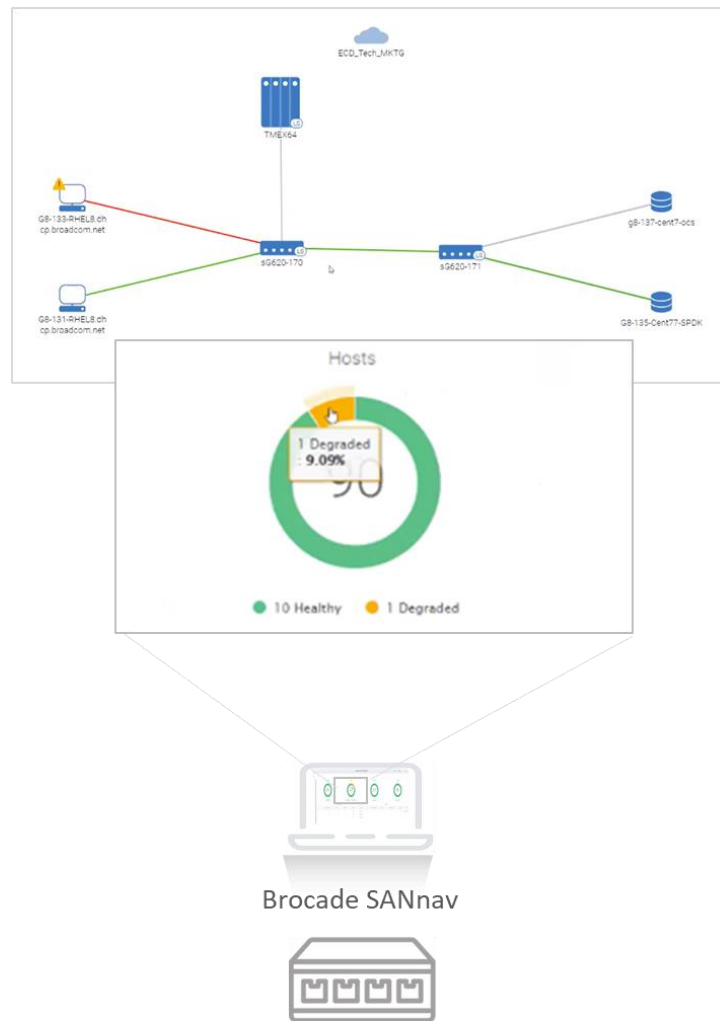
Brocade SANnav Management Portal and Emulex SAN Manager were used to investigate SAN congestion, enable autonomous self-healing to fix the congestion, and monitor application-level performance during testing.

³ Fabric Performance Impact Notifications (FPINs) are defined by INCITS/T11.

Detecting SAN Congestion

SAN congestion was flagged as a problem through the Brocade SANnav Management Portal Health Summary score ring as an event that needed investigation. Brocade SANnav made it easy to visually investigate the error with a single click to drill down to explore the cause of the alerts. As shown in Figure 3, Brocade SANnav made it easy to detect SAN congestion and investigate which HBA ports and switch paths were fighting for bandwidth through the topology view. Once Brocade’s switches detected and identified the congestion impact, the FPIN was sent to the end device, in this case the Emulex HBA.

Figure 3. Detecting SAN Congestion (Brocade SANnav Management Portal)

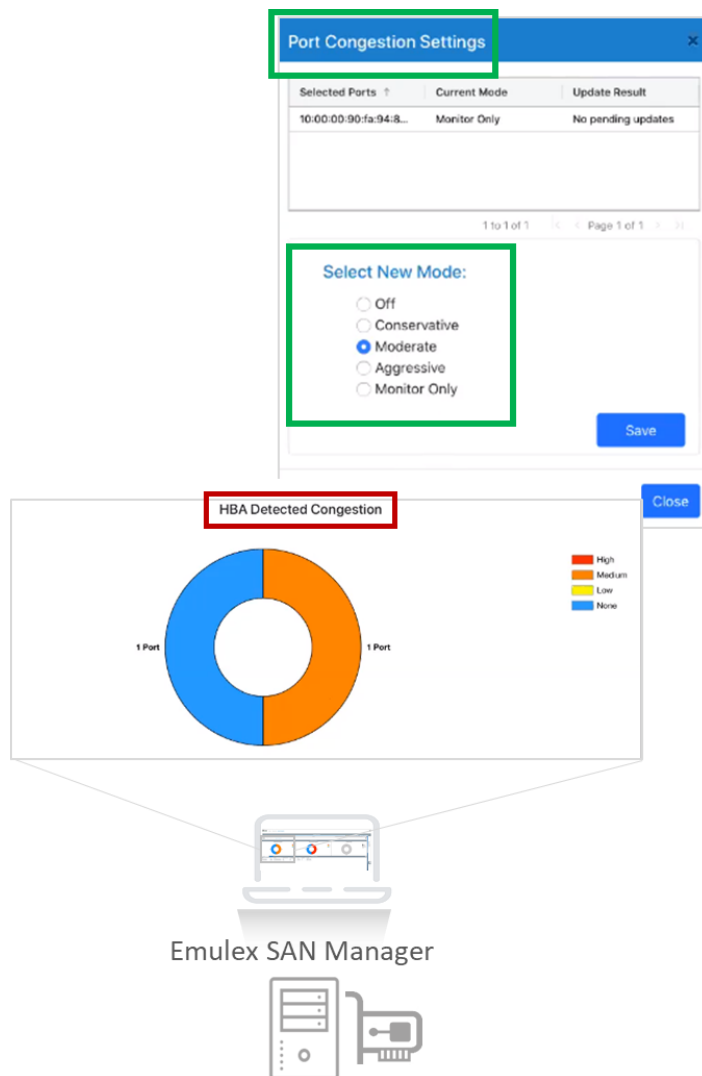


Source: Enterprise Strategy Group

Enabling Autonomous Self-healing

Next, ESG used the Emulex SAN Manager interface to visualize the FPIN notifications sent by Brocade switches, as shown toward the left in Figure 4. ESG then turned on a moderate level of port congestion management for the lower priority application. This setting, which was previously set to “monitor only,” tells the Emulex HBAs to slow down the lower priority “bully” traffic with a goal of eliminating the port congestion issue.

Figure 4. Enabling Autonomous Self-healing (Emulex SAN Manager)



Source: Enterprise Strategy Group

ESG noted that the autonomous SAN technology detected the SAN congestion issue in Brocade switches and used the FPIN protocol to notify Emulex HBAs about the congestion and to start using an adaptive congestion management algorithm to fix the problem. Slowing down a lower priority application with an adaptive congestion algorithm is a sophisticated approach for SAN congestion management compared to the legacy approaches of manually changing host adapter queue depth settings or using a QoS setting to set a hard performance limit for a lower priority application. The Broadcom self-healing approach doesn’t require agents on the hosts and it works with operating systems available today. The Broadcom self-healing approach is not only more sophisticated, but also operates without intervention in real time, and constantly adjusts performance levels to maximize bandwidth usage if the congestion problem is transient.

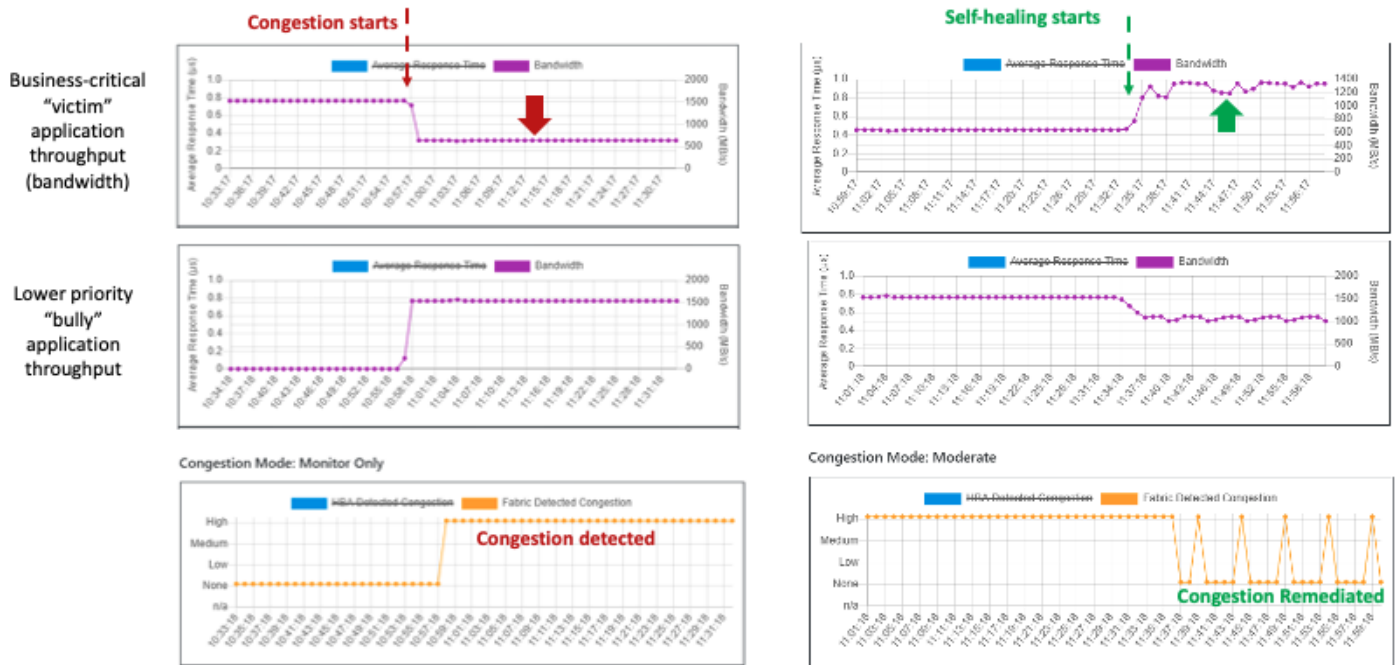
Autonomous Self-healing in Action

ESG used both Brocade SANnav Management Portal and Emulex SAN Manager GUI to monitor performance and SAN congestion during each phase of the self-healing SAN test. As shown on the left in Figure 5, the business-critical data warehouse workload was running at approximately 1,500 MB/sec of sustained throughput before the lower priority backup job kicked in. Congestion was detected as performance for the business-critical data warehouse workload dropped by more than 50%.

Autonomous self-healing is shown on the right in Figure 5. Note how business-critical performance picked up as the congestion algorithms were used to slow down the lower priority application. Also note how congestion remediation is happening in real time. Self-healing periodically kicks in to make real-time adjustments with a goal of providing a fair share of SAN bandwidth for lower priority applications if and when the congestion goes away.

Figure 5. Autonomous Self-healing in Action (Emulex SAN Manager)

Source: Enterprise Strategy Group



Validating Autonomous Self-healing

Brocade SANnav Management Portal then indicated that the congestion impact had disappeared and performance latency had gone back to normal. The original congestion violation message and the message that indicates that the congestion has been cleared are highlighted in Figure 6.

Figure 6. Validating Autonomous Self-healing (Brocade SANnav Management Portal)

Violations Congestion							
Object Name	Prod...	Rule Name	Rule ...	Measure	Measure Value	Port Type	Last Occurred
port48	10.231.1...	defALL_PORTS_IO...	ALL_POR...	Device Impact latency (DEV_LATENCY_IMPACT)	IO_LATENCY_CLEAR	E-Port	Sep 01, 2020 1
port48	10.231.1...	defALL_PORTS_IO...	ALL_POR...	Device Impact latency (DEV_LATENCY_IMPACT)	IO_PERF_IMPACT	E-Port	Sep 01, 2020 1

Close

Source: Enterprise Strategy Group



Why This Matters

SAN congestion that causes application performance problems is a notoriously hard problem to diagnose and fix. Slow application performance that is caused by (or mistakenly suspected to have been caused by) SAN congestion can lead to a loss of employee productivity, customer satisfaction, and revenue.

Finding, fixing, and automatically eliminating SAN congestion problems saves time and money for IT professionals and improves customer satisfaction and productivity for application users.

The Bigger Truth

SAN congestion problems that block the flow of the data have been a challenge for IT organizations since the T11 technical committee first defined the Fibre Channel standard in 1988. Fibre Channel HBAs, switches, and drivers don't typically log errors that can be easily correlated to SAN congestion. The SAN is simply working overtime and needs to be monitored to see if, when, and where SAN congestion is happening. And more often than not, the application performance problem doesn't have anything to do with the SAN. After a SAN congestion problem has been detected and isolated, the traditional methods for fixing SAN congestion are manual and can't react quickly to changing traffic conditions.

ESG has validated that latest Brocade SAN Gen 7 switch hardware, Brocade Fabric OS version 9.0 software, and Emulex Gen7 HBAs have leveraged the FPIN specification to automatically find and fix SAN congestion problems. Detecting and troubleshooting SAN congestion with Brocade SANnav Management Portal and displaying FPIN notifications with Emulex SAN Manager was intuitive and easy. Five minutes after enabling self-healing with a single mouse click, a "noisy neighbor" workload had been muted and the performance of a business-critical data warehouse application had recovered.

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. Broadcom's new SAN congestion management feature can be turned off, left on in monitor mode only, or configured for autonomous self-healing. Based on the results of our validation testing, ESG believes that autonomous self-healing mode will be quickly embraced by IT, SAN, and storage administrators.

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 Broadcom autonomous self-healing SAN management.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

The goal of ESG Validation reports is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Validation reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.