Fibre Channel Self-Healing Networks

By Brandon Hoff, FCIA member, Principal Architect, Broadcom

The good old days: when we talked about congestion but didn't really see it in the wild. When we did, it was a big deal. We would sit down with the customer to work out a solution. By we, I mean the storage, server, HBA and switch vendors to name a few. This was a long and costly process. Teams of engineers going though endless data to pinpoint the root cause so it could be addressed. Then, addressing the issue took another team a lot of time and effort to fix. It is important to note that congestion isn't a Fibre Channel problem, it is a networking problem.

Congested Storage Networks

Today it is worse. All Flash Arrays and NVMe/FC storage arrays can stream 128GFC and soon 256GFC of data from a single HBA. Congestion happens when the server(s) cannot download the data fast enough from the storage network. In a lossy network, the congestion ends up in packet drops that either incur a protocol-level retry or an application-level retry, which add significant latency to the flow and cause application performance problems.

For lossless networks, the data fills up the network and consumes hardware resources across the fabric. The network of the past could easily identify when device behavior had an adverse impact on the performance of the specific device performance and other performance, but it was only able to do some of the mitigation by itself. Now with the ability to communicate behavior with adverse impact to the devices and the ability to react accordingly, the end-to-end system becomes self-healing.

The problem occurs when a workload's "eyes are bigger than its stomach" and the workload asks for too much data and becomes the Bully workload. This workload consumes hardware resources across the network. For lossless networks, buffer credits in the fabric switches are consumed by the Bully workload. Victim workloads, other workloads connected to the SAN, are starved and their performance is significantly impacted. It is becoming common for network administrators to plan for Bully Workloads to pop up in their storage network.

Causes of Congestion

When a new installation is architected, it is right-sized. That means it has enough compute, memory, storage, and bandwidth to handle the workloads deployed on the hardware platform.

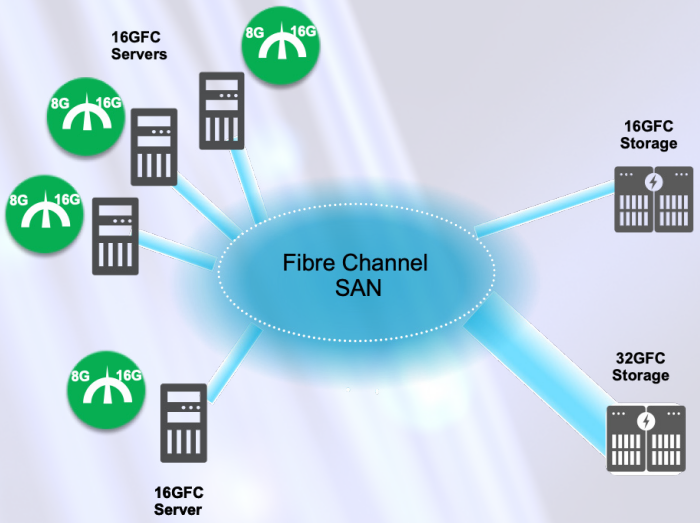

Figure 1: Well-architected data center

Over time, things change and there are two key causes of congestion:

1. A workload gets too big for its hardware footprint and results in an over-utilized server. It consumes either 100% of the CPU, memory, PCIe bandwidth, or Fibre Channel bandwidth. Examples of this are an administrator moving too many VMs onto a server or from the normal growth of data used in applications.

2. A 32GFC All Flash Array is installed on the network and now can feed the demands of over-utilized servers.

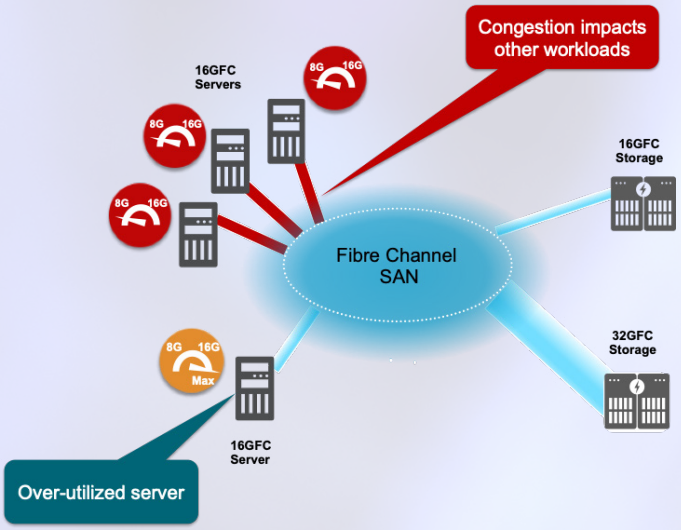

Figure 2: Congestion impacting multiple workloads

The over-utilized workload, or the Bully workload, looks like it is running well, and the lossless network is doing its job, so it isn't obvious that it is causing problems.

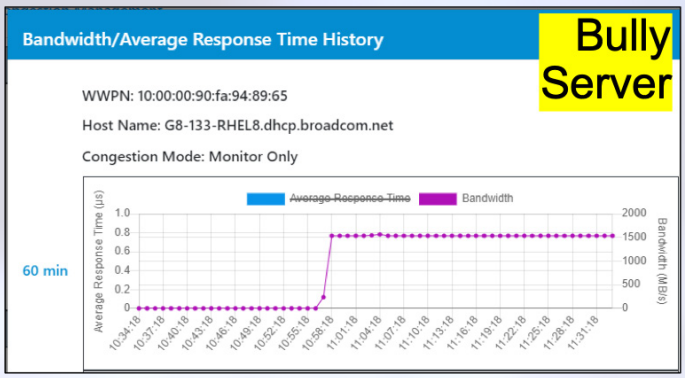

Figure 3: The Bully Workload comes online

Here is the challenge: the hardware resources of the SAN are holding data and waiting for the Bully workload to download them. If you connect more workloads to the network, Victim workloads are created. Victim workloads take a significant performance hit. In Figure 4, the performance degradation of the Victim workload is cut in half.

At this point, many other workloads on the SAN are being negatively impacted. Customer wait times go up as application administrators file trouble tickets. Everybody points to the SAN as being the bad guy.

This scenario applies to any lossless network: IP, InfiniBand, and Fibre Channel. Identifying where the problem is and how to solve it is very difficult, until now.

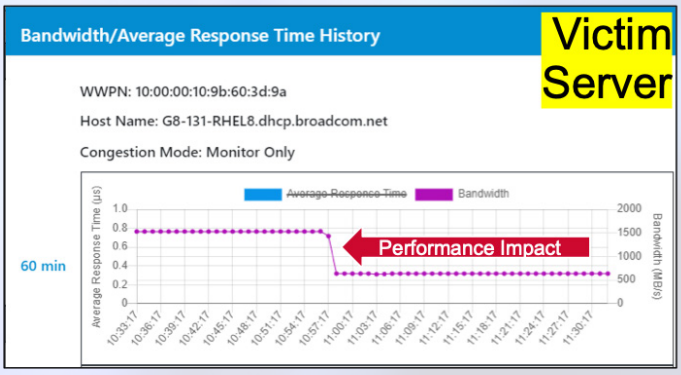

Figure 4: Performance impact on Victim Workloads

Based on the requests from enterprise customers, the Fibre Channel industry developed Fabric Notifications. Fabric Notifications is a new set of standards from INCITS/T11 that enables Fibre Channel Fabric to collaborate with Fibre Channel end-points (HBAs) and Fibre Channel end-points to collaborate with the Fibre Channel Fabric. This article discusses how Fabric Notifications can be used to detect congestion and remediate it. The rule is that the Fabric is the best place to detect congestion, and the end-point is the best place to fix congestion.

Detecting Congestion in the Fabric

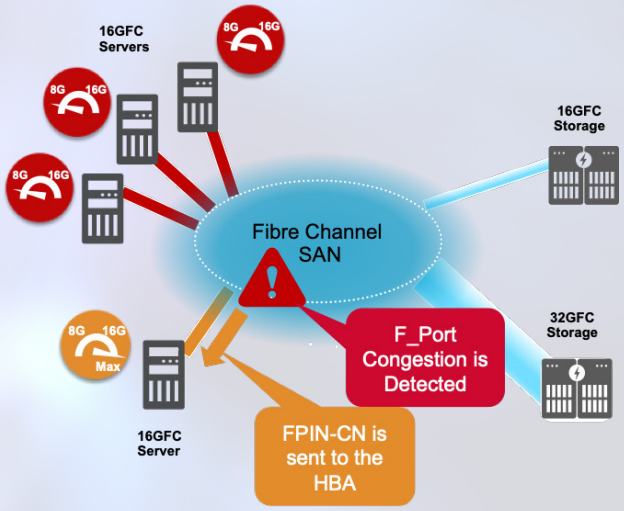

Part of the Fabrics Notifications are two new signals. One is called FPIN-CN that stands for Fabric Performance Impact Notification, specifically, Congestion Notification. FPIN-CN is sent to the HBA as an ELS Frame and requires a buffer credit. The second is congestion signaling that is part of the 64GFC Gen 7 Fibre Channel standard which provides the same information, but at lower level in the networking stack and doesn't require a buffer credit. Both signals tell the HBA that its server is causing congestion.

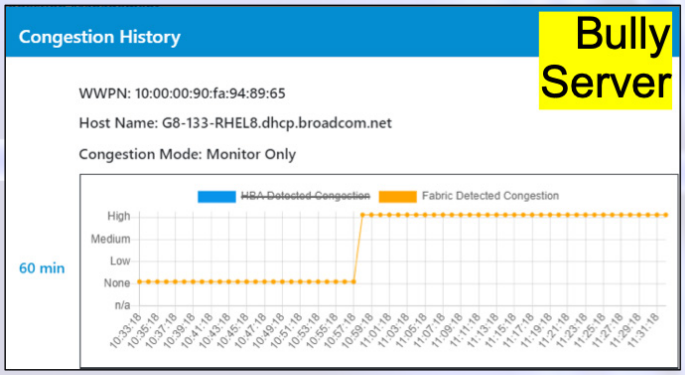

Figure 5: FPIN-CNs received by the HBA

Once the switch detects congestion it sends FPIN-CNs to the HBA. The Fabric now collaborates with the endpoint, a new networking technique first deployed in Fibre Channel networks.

Fabric Notifications is a new technology, and as with any technology, the administrator needs to review data from the Fabric and the HBA. They can set the policy to monitor only and review performance over time. They can turn on congestion management for low priority workloads, like noisy neighbors, and monitor high priority workloads.

Monitoring and Remediating Congestion at the HBA

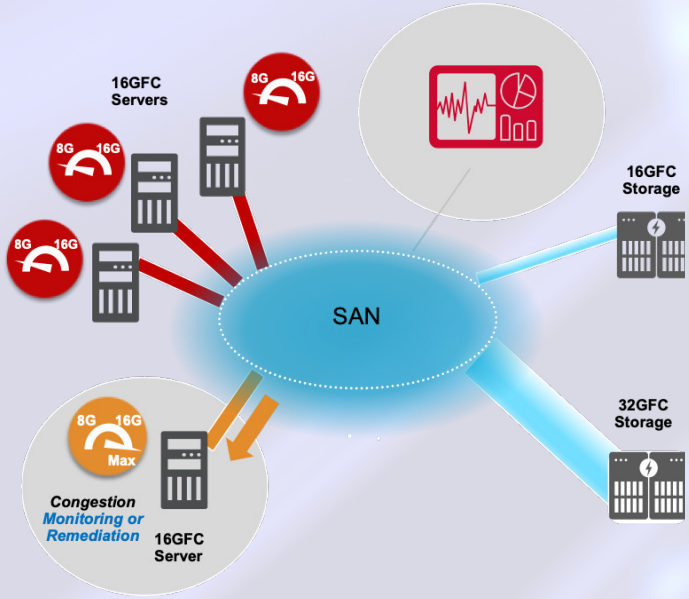

Figure 6: Fabric Notifications at work

Monitoring Congestion

When deploying a new technology, it is essential to be able to see the before and after of a solution like congestion management. So, the first step is to monitor congestion, identify impacted workloads, and decide how to proceed. For a tier 1 workload, you may decide not to touch it and just monitor it over time. Or you may decide to turn on congestion management and measure the effects of this new feature. For low priority workloads, you may turn on congestion management by default and monitor those workloads for performance issues. Fabric Notifications give administrators a lot of flexibility around detecting and remediating congestion.

Remediating Congestion

If you decide to turn on congestion management for the Bully workloads, congestion management can restore performance to the Victim workloads. This is done by slowing down the Bully workload at the Fibre Channel HBA.

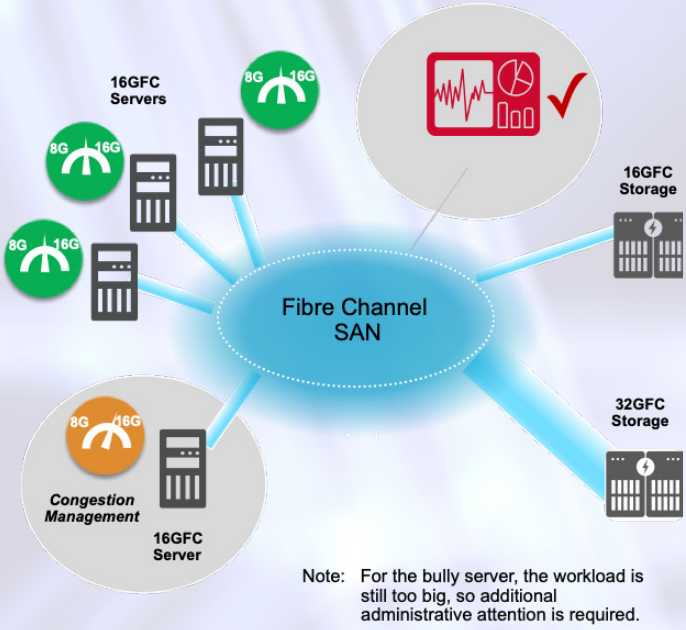

Figure 7: Performance is restored

With congestion management, the demands on the storage network are balanced to maximize resource utilization on the SAN. Victim workloads are no longer impacted by the Bully workload. The Bully workload has slowed down, which may seem like a bad thing, but, if you dig into the data, you'll see that the latency for the Bully workload has dropped significantly, thus, even the Bully server runs better.

The result is balanced resource utilization maximizing performance on the hardware footprint. Note that the oversized workload still needs to be addressed by either moving the workload to a faster machine or by moving VMs off the server.

The Proof is in the Data

Figure 8 and Figure 9 show the complete cycle from congestion, enabling congestion management, and remediating congestion.

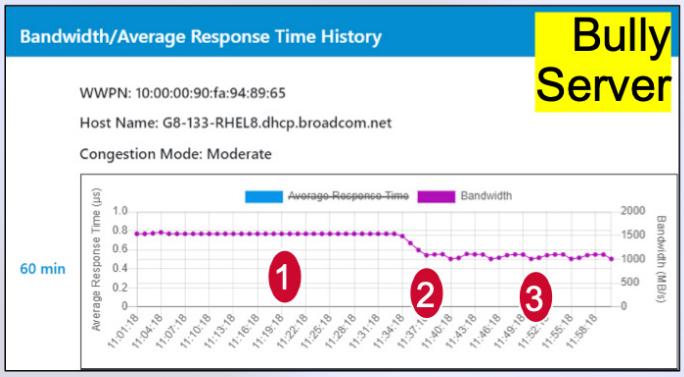

Figure 8: Bully workload before and after congestion management is enabled

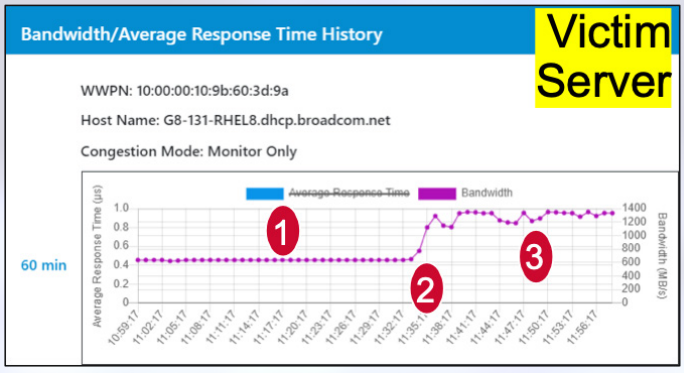

Figure 9: Victim workload before and after congestion management is enabled

The different points on the graphs are:

1. The Bully workload is causing congestion and the Victim's performance is cut in half.
2. Fabric Notifications is turned on in the Fibre Channel Fabric Congestion management is enabled on the Fibre Channel HBA.
3. Congestion management reduces the throughput on the Bully workload by about 30%. Performance is restored for the Victim workloads on the SAN. The throughput on the Victim flow is doubled and near line-rate.

Note that the Bully workload, at point 3 on the graph, bumps up and down a little bit. This is because this solution is adaptive and can address transient and steady state congestion events.

Of course, this process doesn't need to be manual. It can be integrated with enterprise management tools and automated making it easy to deploy and manage.

Summary

Today, cogestion is becoming more common. Driven by All Flash Arrays, the velocity of data has increased, and congestion creates performance problems around the data center.

In a lossless network, that data fills up the network and consumes all of the network's resources, impacting other workloads and causing performance problems across the data center.

Fabric Notifications is a powerful set of new tools to solve data center networking issues for Fibre Channel storage networks. Congestion management directly address a key customer point of pain in their storage networks – the performance degradation due to congestion.

Availability:

- This solution works transparently with any Fibre Channel storage array.
- This solution is supported in VMware, RedHat, SUSE, and Windows today.
- This solution is available on servers from many server vendors including HPE, Lenovo, Dell, and others.

For more information on congestion:

gblogs.cisco.com/in/slow-drain-in-fibre-channel-better-solutions-in-sight

fibrechannel.org/wp-content/uploads/2018/02/FCIA_Fibre_Channel_Performance_Congestion_Slowdrain_and_Over_Utilization_Final.pdf