

Marvell® QLogic® Fibre Channel USCM - Universal SAN Congestion Mitigation

Effective and efficient identification and mitigation of SAN congestion by leveraging intelligent fabric – end point collaboration

Background

Fibre Channel Storage Area Networks (SANs) are still considered the most reliable and secure storage architectures for today's enterprise environments. With a proven track record of over 20 years, Fibre Channel is the only protocol in the data center that is purpose-built for storage networking.

We have seen some great leaps in Fibre Channel technology over the years including increased IOPS and bandwidth, lower latency, simplified management, and advanced diagnostics. This is even more evident today with the introduction of FC-NVMe or Non-volatile Memory Express over Fibre Channel for Solid State Disks (SSD). Just as Fibre Channel was purpose built for storage, NVMe was designed from the ground up for SSDs. NVMe is a much more efficient interface using fewer commands, providing greater scalability and lower latency than the legacy SCSI interfaces.

As next generation All Flash Arrays (AFA) are developed using NVMe drives, it is the natural progression to use Fibre Channel as the transport for NVMe storage commands across the storage area network. Hence the FC-NVMe standard was developed, and we are now seeing native NVMe storage arrays connected to servers across the SAN.

With all this in mind, imagine the FC SANs of today incorporating the latest speeds along with the lowest latency available with FC-NVMe.

The Challenge

To deliver on the high-performance expectations associated with native FC-NVMe storage arrays, the SAN infrastructure must be optimized to not only transfer data faster, but to also deal with potential congestion, link integrity and delivery failures. The ability to identify and deal with these issues is crucial to delivering the performance benefits of NVMe in the SAN.

Why is congestion an issue? In most enterprise environments, there are many legacy servers and storage devices connected to the SAN at 8/16/32GFC. One of the benefits to Fibre Channel SANs is that they are always backward compatible two generations, enabling customers to keep legacy connections while adding new storage to their environment. FC SAN environments today can have 8GFC, 16GFC, 32GFC all running in the same fabric. The problem is these legacy connections are not as efficient or as quick to respond as the new NVMe native all flash arrays (AFA) and may create congestion in the SAN that can impact the performance of the entire SAN. When native 32G FC-NVMe AFA and 32GFC devices as well as 16 and 8GFC devices are all connected to the same fabric, there can be some congestion on the fabric caused by "slow drain" related to

buffer credit starvation of the slower devices. A slow device can't respond fast enough and as a result, buffer credit back-pressure is created in the fabric, and even in overprovisioned inter-switch links (ISLs). This can impact performance of all the devices on the SAN over time (Figure 1).

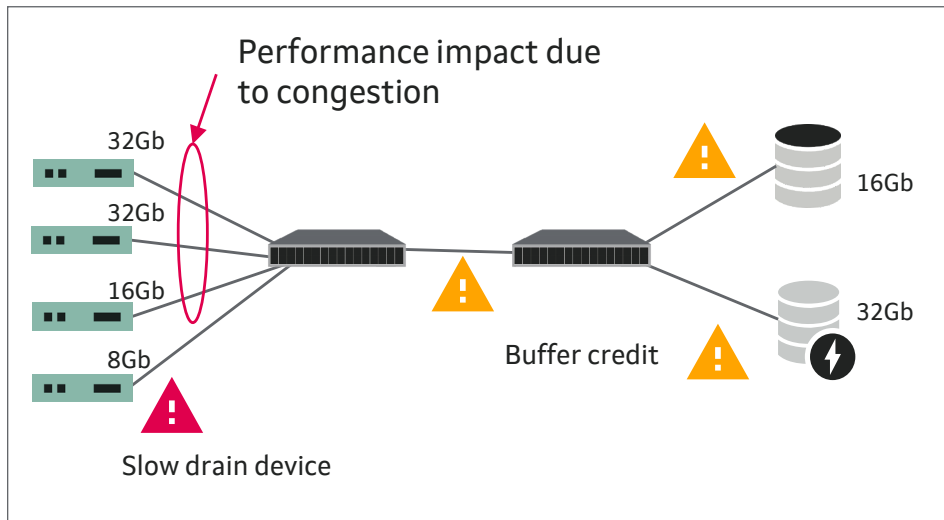


Figure 1

Even in a well architected data center an oversubscribed or over-utilized server can cause SAN congestion starving other servers in the fabric of resources. This is where Marvell StorFusion® USCM (Universal SAN Congestion Mitigation) technology comes into play.

The Solution

USCM is based on Fabric Performance Impact Notification (FPIN) which was developed by the collaborative efforts of the Fibre Channel community and known as the T11 standard FC-SW-8 (Fabric Notifications Overview & Scope). There is a different standard for each of the notifications from congestion signals and frame discard timeout values to link integrity issues. Table 1 breaks it down in an easy-to-understand format.

FPIN Type	What it does
Congestion Notification (FPIN-CN)	Fabric notifies the port that is causing congestion
Peer Notification (FPIN_PN)	When a FPIN-CN is sent to the end point causing congestion, a Peer Notification (PN) is sent to all the other end points in the zone.
Link Integrity (FPIN-LI)	Fabric notifies the end point (server or target) that the link it is connected to is 'sick but not dead.' Typically sent when physical layer errors exceed a defined threshold
Delivery Notification (FPIN_DL)	Delivery Notification - Best example is if a switch drops a frame due to egress hold time expiring (220ms) and unable to deliver the frame.

Table 1

These notifications allow the HBAs and switches in the SAN to identify and notify administrators of an issue within the fabric being monitored. Not only does it provide strong awareness to the user that there is an issue, these notifications will be used by the HBA driver to dynamically correct the issue or recommend a course of action requiring user intervention. For example, if there is an over-utilized server causing congestion on the fabric, USCM can have the HBA in that server implement IO leveling (throttling) and reduce the impact it is making on the overall SAN data flow.

QLogic Fibre Channel HBAs from Marvell work in conjunction with either Brocade® or Cisco® Fibre Channel SAN switches to provide all the FPIN notifications listed above. Hence the QLogic implementation is referred to as **Universal** SAN Congestion Mitigation, or USCM. Additionally, the standard firmware/driver for QLogic HBAs support all four FPIN notification types with no need for additional software licenses that may be required in competitor offerings.

Decisive Actions for USCM

Over time, USCM will move beyond just supporting notifications of SAN congestion. HBA Driver enhancements will be forthcoming to add I/O throttling, automatic transitions to alternate paths in the SAN, and the ability to segregate Fibre Channel flows into virtual lanes that drive traffic prioritization.

The result is, using QLogic HBAs with USCM enables customers running mission critical workloads to be able to add the latest FC-NVMe AFAs to their existing storage environments and get the most out of their storage investments. The new FC-NVMe storage can co-exist with legacy servers and storage and still deliver optimal performance with potential congestion issues eliminated by the SAN infrastructure with the capabilities from QLogic Universal SAN Congestion Mitigation and future enhancements. This form of “Self-Driving” networks will be the new paradigm in storage area networks (Figure 2).




Figure 2

Requirements

USCM requires a server with a 16GFC Enhanced or 32GFC Marvell QLogic Fibre Channel HBA with the latest driver and firmware, Brocade 32G and 64G switches running FOS 9 or later or Cisco 16G/32G switches running NXOS 8.5.1 or later.

Summary

With over 20 years of service, the QLogic FC HBA from Marvell is one of the top-rated adapters in the industry with a reputation for reliability... “It just works!” Our commitment and contributions to the Fibre Channel community for technical innovation and forward thinking has made Marvell a respected leader in the industry. Our unique port-level isolation architecture for reliability, performance stability and security provide the end user with peace of mind.

Marvell architecture and innovation deliver the ultimate reliability to meet the needs of mission-critical enterprise applications with less fabric congestion all while maintaining peak performance. For more information on Universal SAN Congestion Mitigation (USCM) please visit <https://www.marvell.com/QLogic>.



To deliver the data infrastructure technology that connects the world, we're building solutions on the most powerful foundation: our partnerships with our customers. Trusted by the world's leading technology companies for 25 years, we move, store, process and secure the world's data with semiconductor solutions designed for our customers' current needs and future ambitions. Through a process of deep collaboration and transparency, we're ultimately changing the way tomorrow's enterprise, cloud, automotive, and carrier architectures transform—for the better.

Copyright © 2022 Marvell. All rights reserved. Marvell and the Marvell logo are trademarks of Marvell or its affiliates. Please visit www.marvell.com for a complete list of Marvell trademarks. Other names and brands may be claimed as the property of others.