

HOW NVME™ OVER FIBRE CHANNEL IMPACTS SAN PRINCIPLES

Curt Beckmann, Broadcom Corporation

The availability of the NVMe standard has radically changed the landscape for solid-state storage, driving commoditization of the media along with aggressive competition for density and performance. This media revolution is causing a secondary pair of disruptions in the storage array space. The first disruption is the move toward NVMe-based SSDs (rather than SAS- or SATA-based SSDs) as the media of choice on the All-Flash array back end. The second disruption is a move toward NVMe over Fabrics, and particularly NVMe over Fibre Channel, as the emerging high-performance protocol for accessing enterprise storage.

Both disruptions promise dramatically higher performance, in terms of lower latency, higher IOPS and higher bandwidth. Vendors can offer support for NVMe SSDs as back-end media usually without need for changes to other elements of your SAN. Moving to NVMe over Fibre Channel as the transport for communicating to your arrays (as an alternative to SCSI over Fibre Channel, often called FCP) will also be smooth and straightforward, since Gen 6 SANs and host bus adapters (HBAs) support NVMe over Fibre Channel concurrently with SCSI FCP. NVMe over Fibre Channel is also supported by 16 Gbps SANs, which offers SAN administrators flexibility in evaluating NVMe over Fibre Channel, although Gen 6 is a better choice for delivering the performance benefits of NVMe in production environments.

As storage teams consider the benefits of NVMe-capable storage arrays, many have questions about whether this new transition will impact the “best practices” with which they are familiar. The good news is that NVMe technology, whether as back-end array media or fabric transport protocol on the front-end connecting hosts, does not

significantly alter SAN principles, with one small but important caveat. The caveat is that before NVMe, SAN teams may have felt they had some wiggle room in applying recommended best practices. That is to say, the need for compliance with best practices is even more important as we begin adopting NVMe technologies in our storage environments.

Let’s do a quick survey of some familiar SAN principles in the context of NVMe.

End-to-end node speed alignment: A long-time best practice for Fibre Channel SANs has been to limit port speeds differences to no more than a factor of four. For example, a SAN with 32-Gbps devices can include 16-Gbps and 8-Gbps devices, but not 4-Gbps devices. Larger speed differences, such as the 10x difference between 10 GbE and 100 GbE, can create head-of-line blocking when a fast target responds to large read requests by sending more frames into the fabric than the fabric can quickly forward to a much slower initiator. When this happens, the data frames pile up in the fabric, hogging frame buffers and impeding the flow of other traffic. With higher bandwidth (up to 30 Gbps) and enhanced queuing of NVMe SSDs, it is more important than ever to adhere to this principle.

Storage traffic isolation: In situations where there is a strong need to support mismatched end nodes (either host or target), usually driven by a need to support a legacy application, the best practice is to isolate the traffic between the end nodes involved. Your SAN infrastructure should be capable of isolating traffic for this use case. Like the others listed, this best practice is not new, but it is likely to become more applicable to older SCSI platforms over time as storage volumes migrate to NVMe.

Use the latest fabric services: Almost since the beginning of Fibre Channel SANs, fabric-based services running on switches have been a cornerstone of interoperability. Indeed, the choice to provide built-in fabric services (versus external add-ons like DNS) has become a key differentiator relative to other technologies. This will continue to be the case as NVMe over Fibre Channel adoption grows. The use of fabric-based services ensures early multi-vendor interoperation of essential functions like discovery, zoning and feature detection.

Full matrix, end-to-end storage qualification: Another enduring best practice among top-tier Fibre Channel vendors (and buyers issuing requests for information) has been to insist on full qualification of products across a large matrix of SAN infrastructure and equipment vendors. This testing proactively eliminates most issues before deployment and therefore delivering the reliability that FC promises, as well as preparing all parties to promptly resolve issues that occasionally arise after deployment. Customers have come to expect and demand this level of testing, ensuring a culture of “it’s gotta just work” across the SAN industry. This expectation is so ingrained that Fibre Channel Plugfests are consistently fast-paced multi-vendor break-fix events hosted by the Fibre Channel Industry Association (FCIA), instead of choreographed one or two vendor beauty contests that often occur in some adjacent markets. The rise of NVMe over Fibre Channel is a good time for all members of the Fibre Channel community to remind ourselves that it doesn’t matter how good it looks, it’s gotta just work! And so, we all must demand and support full matrix testing.

Granular monitoring: In our daily lives we are all frequent consumers of best effort networking, where things get flaky and the solution is to “reboot and try again.” Perhaps this is just reality for a consumer-grade, multi-function network, but imagine what customers of your always-on storage network would say to that! Fortunately, because your Fibre Channel SAN is focused on storage, the economics have encouraged a best practice of embedding highly granular storage-aware monitoring tools for detecting, understanding and resolving minor anomalies before they grow into anything disruptive. Multifunction commodity networks, by contrast, are simply not in a position to offer comparable diagnostic features across their user space.

Automation: As NVMe over Fibre Channel brings new applications to the SAN, there is a “new kid” on the SAN principles block; that is, the rise of DevOps-driven agility and cloud computing in the enterprise datacenter have led to an increasing need for automation and orchestration. Thus, the newest “best practice” is to ensure that your SAN infrastructure supports the development of process automation through APIs, especially REST APIs based on industry standard modeling mechanisms such as NETCONF and Yang.

As mentioned earlier in this article, most of these SAN principles are familiar. Indeed, these concepts can become so ingrained that they are like the water we swim in. But as we anticipate a technology transition, such as the adoption of NVMe over Fibre Channel, it behooves us to take a step back and remember what we’re about, and make sure that we consciously consider all of the practices that have served us well over the decades.