



Flash Memory Summit

Low latency and high throughput storage access

Journey from SCSI to NVMe in FC fabrics

Raj Lalsangi



Protocol neutral architecture principles

- Reduce / avoid interrupts
- Avoid context switches
- Avoid locks
- Lean code path length



NVMe performance features

- NVMe Namespace state vs. SCSI LUN state in data path (device server, task set, ordered commands, ACA and so on)
- NVMe multi-queue model
- NVMe lends itself for lower code path length



Flash Memory Summit

FC-NVMe benefits

- Built for low latency and high throughput from ground up
- Better latency and throughput over existing 32G fabric and host HBA ecosystem with software upgrade
- Seamless fabric management – name services, zoning, and auto discovery



Flash Memory Summit

The Needs of NVMe over Fabrics (NVMe-oF)

Dennis Martin
Founder, Demartek



Demartek NVMe-oF Rules of Thumb

To achieve maximum throughput in a storage target (without oversubscription):

- At least one 25Gb or faster network port for each NVMe drive (PCIe 3.0 x4) for large-block sequential I/O
- Dual-port 25GbE or 32GFC adapters need PCIe 3.0 x8
- For every two NVMe drives and network ports 16 lanes of PCIe 3.0 are needed (FC has more headroom)
- Prospects are better with PCIe 4.0



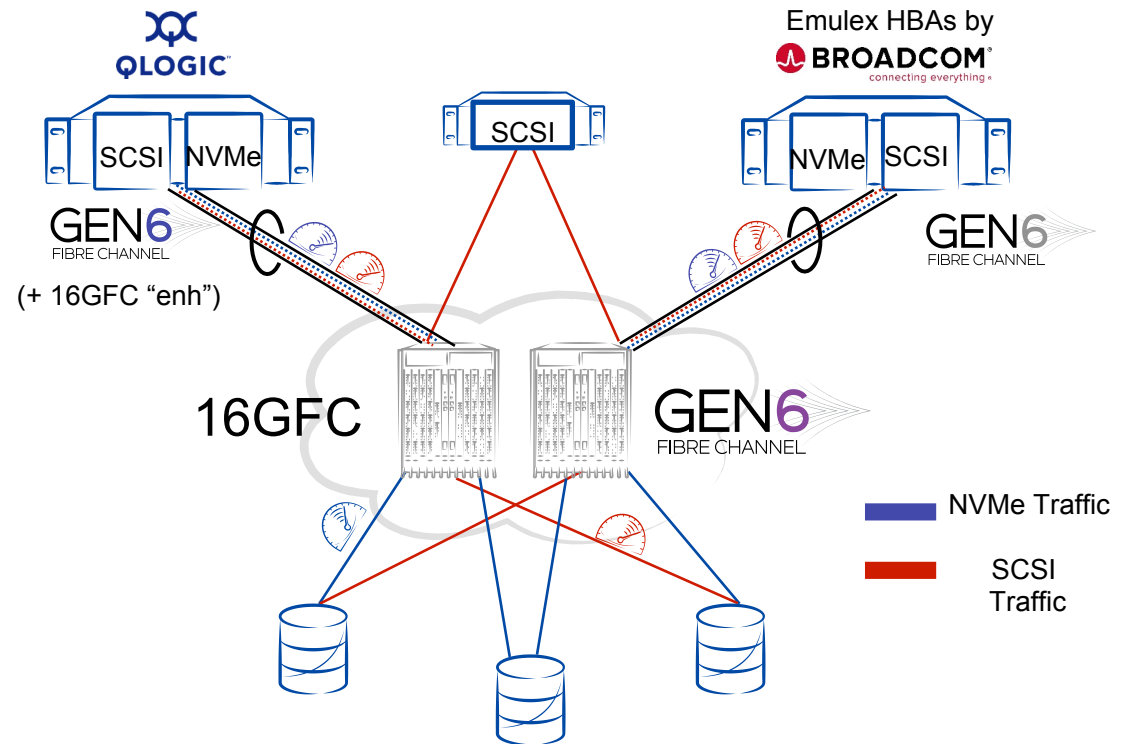
Benefits of FCP + FC-NVMe

Curt Beckmann
Principal Architect
Brocade



The Benefits of Dual Protocol FCP / FC-NVMe SANs

- Enhances performance of existing SAN; no need for costly, disruptive infrastructure duplication / replacement
- Migrate application volumes 1 by 1 with clean, easy rollback options
- Makes interesting dual-protocol use cases (DB mining) available
- Full fabric awareness, visibility and manageability with familiar tools





Flash Memory Summit

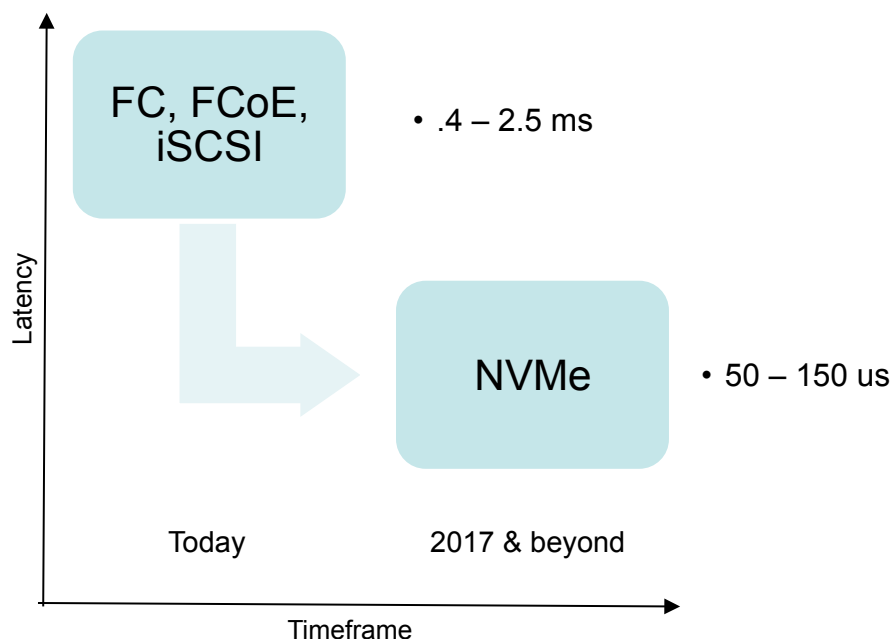
NVMe over Fabrics Comparisons

Rupin Mohan, Director R&D, HPE
August 2016



Shared Storage with NVMe

RDMA-Based NVMe-oF	NVMe over Fibre Channel
Completely new fabric protocol being developed and standardized	Uses Fibre Channel as a base, existing fabric protocol, shipping, standardized by T11
Standards group dealing with same type of challenges, shipping I/O commands/status, data over distance.	Fibre Channel solved these problems when FCP protocol was developed to ship SCSI commands/status over distance over a FC network
RDMA is available as per protocol	RDMA is not available, uses FCP
Zero-copy capability Transport options are iWARP, RoCE (v1 or v2) and even considering TCP/IP	Zero-copy capability Transport is FC. No changes to switching infrastructure / ASICs required to support NVMe over FC
Complex integrated fabric configuration Could be lower cost if onboard NIC's on servers and cheap non-DCB switches are used	FC fabrics are well understood Higher cost, especially the newer generations of FC are expensive
New I/O protocol, New transport	New I/O protocol, Existing reliable transport
Lower latency if RDMA is enabled	Latency improvements with hardware assists on Adapters. No RDMA option.





Flash Memory Summit

FC “Goodness of Fit”

Praveen Midha
Cavium

Is FC a Good Fit/Transport for NVMe?



FC-NVMe is the Perfect Harmony between Low Latency NVMe Flash and a High Performance FC Transport