Fundamentals of Fibre Channel

Rupin Mohan
Earl Apellanes
June 15, 2017
11:00 am PT

FIBRE CHANNEL INDUSTRY ASSOCIATION

Today's Presenters

Rupin Mohan
www.linkedin.com/in/rupin
HPE

Earl Apellanes
@FabricPrincipal
www.linkedin.com/in/earl-apellanes-a851482/
Brocade

Agenda

- Market and business context
- Customer SAN requirements
- What is Fibre Channel?
  - FC protocol features refresher
- NVMe over FC
  - Fabric advantages
  - Latency
- Summary

What are we going to discuss today?

- We will discuss
  - What is Fibre Channel?
  - What are customer requirements from a SAN?
  - What makes FC so reliable?
  - What makes FC different?
  - What are the best practices for designing FC solutions?

- We will NOT discuss
  - N-port, E-port, etc
  - Classes of service
  - FC Frame, Exchanges, Sequences, etc
  - FC detailed addressing
  - 8B/10B, 64B/66B encoding
  - Blah blah blah

  (Geeks, so sorry to disappoint you)

Connecting Businesses to Important Data

Fibre Channel fabrics are the common thread

Reliability

Keep running no matter what

Performance

Unleash application performance

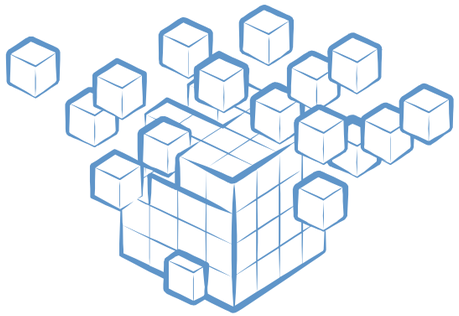

Scalability

Adapt to your business

Security

Secure your data and mitigate risk

Critical applications and data require purpose-built networking for storage

What are SAN Requirements?

- Reliable network
  - Minimal to zero packet loss
- Deterministic network behavior
  - Robust error recovery mechanisms
  - NSPF / redundancy
- Security and Authentication
- Predictive performance, ability to handle –
  - Bursty storage traffic
  - Unbalanced network flows
  - Bad/Degraded or Greedy N nodes and slow legacy devices
- Manageability
  - Rich in error detection and reporting
- Seamless scalability

The future of storage protocols

Fibre Channel is data center storage protocol of choice for the next decade

Figure 1. Comparison of Storage Protocols

Protocol	Performance		Cost		Choice	Reliability, Availability and Serviceability	Ease of Administration	Future Upgrade Path
	Throughput and Latency	Bandwidth	Purchase	Operational				
Fibre Channel	●	◔	◔	◔	◕	●	◔	●
iSCSI	◔	◔	●	●	●	◔	●	●
NFS	◐	◐	●	●	●	◔	●	●
FCoE	◐	◔	◔	◐	○	◐	◐	◐
SMB	◐	◐	●	●	◕	◔	●	●
PCIe	●	●	◔	◔	○	◔	◕	◔
InfiniBand	●	●	○	◐	○	◔	◕	◔
Object Storage	◕	◐	●	◔	◐	◔	●	◐

Lowest rating ○ ◕ ◐ ◔ ● Highest rating

Source: Gartner (June 2016)

Fibre Channel Refresher

- Foundational Fabric Services
  - FC-CT protocol
- Fabric Controller
  - Principal switch, FSPF (Fabric Shortest Path First) Routing, ISL initialization
- Name Server
  - Fabric, Distributed, Phone Book Addressing
- Zoning
  - Security and Access controls, Distributed
- High Performance
  - Flow Control with Buffer Credits, 2K packets, In order Delivery, Congestion Control Algorithms that actually work
- Low Latency
  - Highly optimized (ASIC does I/O transfer, very little software overhead)
- Multi generational Interoperability
  - Current and two back (3 speeds)
- High Availability
  - Extensibility over long distances to different data centers or NSPF with redundant fabrics with MPIO
- Scalability
  - Easy scalability with up to 1000s of nodes per fabric

Foundational Fabric Services

- FC Generic services protocol
  - For practical purposes, all fabric services exist in the switch (embedded)
  - Provides the control plane for information transfer between devices
  - Technical name is FC-CT protocol
  - Distributed services
- Principal switch and Domain ID assignment
- FSPF routing protocol

Name and Event Server

- FC Name server is like the telephone exchange or post office in every town
  - Distributed database in every switch/fabric
  - Devices in a fabric login and logout of the fabric
  - Provides the foundation of a reliable network
- All members, devices register with the event server
  - Updates name server, change registrations and notifications
- Can invoke fabric readjustments if required to maintain proper access between servers and storage
  - Fabric reconfigurations
  - Path failovers and failbacks
  - Fast convergence
  - Etc.

Zoning

- Zone: Consists of one or more members allowed to access each other
  - Storage is unique as we're dealing with file systems, data access, etc
  - This feature was required as data stores (LUNs) can get corrupted
- Has a unique alphanumeric name
- Devices can be part of single or multiple zones
- Two types of zoning – Soft and Hard
- Soft Zoning
  - Lightly enforced
  - Access map to each device is tailored to its zone membership assignment
  - Can be implemented by the name server
- Hard Zoning
  - Physical access is restricted by the hardware
  - More like access key card, more secure
- Redundant and distributed in a fabric (all switch members)
- Latest development is TDPZ - Auto zoning

High Performance

- Flow control with BB Credits
  - Difference between highway and air traffic systems
- ASIC driven, high performance DMA engines
- 2K packets, no jumbo frames required
- Virtual Channels for differentiating and prioritizing traffic
- FC fabrics can guaranty In-order delivery even during path failovers, etc
- MPIO, adds redundancy and high performance load balancing algorithms to take advantage of multiple paths to storage arrays
- Easier to design long distance replication solutions

NVMe: The Future of Storage Networking

- Improves performance and value of solid state storage by an order of magnitude
- Increases scalability, while reducing latency and CPU overhead
- Standards-based technology, connected over a fabric

NVMe and the Data Center

Next Generation Flash Networks

Latency

- No translations between protocols
- Native for NAND
- Lowest in industry

Lower TCO

- Native protocol designed for flash
- Increased density
- Lower system power

Scale

- Massive parallel I/O capabilities
- Broad industry support including, HW, OS and NAND vendors

Fabrics

- Can be transported on different interconnects for best performance
- FC is ideal choice (80% networked flash deployed on FC)

Why Fibre Channel for NVMe?

Industry Leading Low Latency

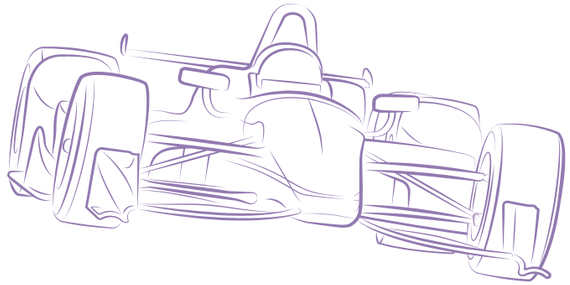

NVMe over Fibre Channel yields 55% latency reduction

Scalability Beyond the Rack

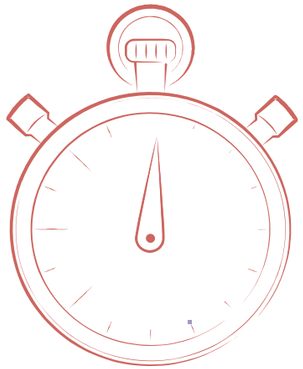

Optimized for heavy storage workloads

Low Risk Deployment

Leverages existing FC infrastructure and Concurrently run NVMe

Faster than 25Gb Ethernet

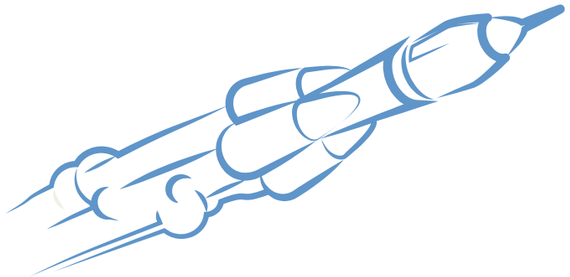

Speeds matched to the PCIe bus for maximum performance

Gen 5 and Gen 6 Fibre Channel is NVMe Ready Today

Why NVMe over Fibre Channel Fabrics?

- NVMe is native to flash
  - Low latency
  - Low overhead
  - High performance
- Natural extension to run it over fabrics
  - Fibre Channel is the predominant fabric for storage
    - Will run on existing Gen 5/Gen 6 fabrics
    - Networked storage is a must for large customers
    - Makes clustering, VM mobility etc a reality i.e. requires networked storage
    - Direct Attached SSD (PCI-E based) doesn't scael
- Fibre Channel is key, as the vast majority of "all flash storage" is deployed on FC
    - Ethernet will also be supported but still some unknowns (new move away from RDMA etc..)

Gartner “Future of Storage Protocols”

Fibre Channel is data center storage protocol of choice for the next decade

- **Orders of magnitude performance improvement**, low latency requires higher-throughput protocols
- **Bottlenecks exist**: 10GbE, 8 Gbps Fibre Channel
- **16 Gbps Fibre Channel will be too slow** for the next generation of storage arrays
- **Plan for higher throughput**, e.g. 32 Gbps Fibre Channel

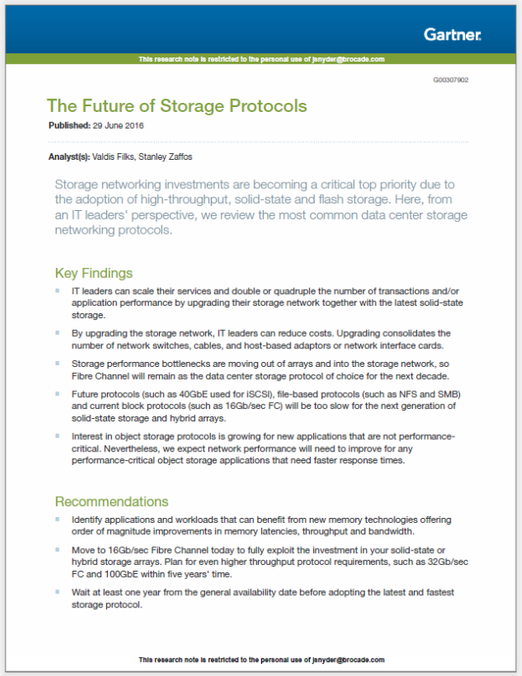

Gartner

G00307902

The Future of Storage Protocols

Published: 29 June 2016

Analyst(s): Valdis Filks, Stanley Zaffos

Storage networking investments are becoming a critical top priority due to the adoption of high-throughput, solid-state and flash storage. Here, from an IT leaders' perspective, we review the most common data center storage networking protocols.

Key Findings

- IT leaders can scale their services and double or quadruple the number of transactions and/or application performance by upgrading their storage network together with the latest solid-state storage.
- By upgrading the storage network, IT leaders can reduce costs. Upgrading consolidates the number of network switches, cables, and host-based adaptors or network interface cards.
- Storage performance bottlenecks are moving out of arrays and into the storage network, so Fibre Channel will remain as the data center storage protocol of choice for the next decade.
- Future protocols (such as 40GbE used for iSCSI), file-based protocols (such as NFS and SMB) and current block protocols (such as 16Gb/sec FC) will be too slow for the next generation of solid-state storage and hybrid arrays.
- Interest in object storage protocols is growing for new applications that are not performance-critical. Nevertheless, we expect network performance will need to improve for any performance-critical object storage applications that need faster response times.

Recommendations

- Identify applications and workloads that can benefit from new memory technologies offering order of magnitude improvements in memory latencies, throughput and bandwidth.
- Move to 16Gb/sec Fibre Channel today to fully exploit the investment in your solid-state or hybrid storage arrays. Plan for even higher throughput protocol requirements, such as 32Gb/sec FC and 100GbE within five years' time.
- Wait at least one year from the general availability date before adopting the latest and fastest storage protocol.

Gartner®

Gartner Research: **The Future of Storage Protocols**, G00307902 June 2016

Summary

- The world's most mission critical applications run on Fibre Channel
  - Healthcare, Finance, Retail, Airlines
- Fibre Channel is purpose built for storage traffic
  - Deterministic reliable network optimized for performance and utilization
  - Network efficiency with multipath load balancing
- Future proof investment protection
  - GEN 6 Fibre Channel supports 4/8/16/32/128G
  - Most All Flash Arrays are attached to Fibre Channel
  - NVMe ready while supporting existing FCP SCSI investments and legacy FICON

Q&A

After this Webcast

- Please rate this event – we value your feedback
- We will post a Q&A blog at http://fibrechannel.org/ with answers to all the great questions we received today
- Follow us on Twitter @FCIAnews
- Join us for our next live FCIA webcast:

Deep Dive into NVMe over Fibre Channel

August 29, 2017

10:00 am PT

Register at https://www.brighttalk.com/webcast/14967/265459

Thank you!