

# Networking Flash Storage? Fibre Channel Was Always The Answer!

Mark Jones Fibre Channel Industry Association Broadcom Limited, ECD



- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

### NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.



## Networking Flash Storage? Fibre Channel Was Always The Answer!

 Fibre Channel has been the ubiquitous connection of choice for connecting storage within the datacenter for nearly twenty years. The start of the sixth generation is being celebrated this year by introducing a staggering leap in performance and new features. We will discuss why Fibre Channel holds the enduring popularity it has as well as an in-depth look at the new Gen 6 features and what the future holds. We will discuss how Fibre Channel fits in with key datacenter initiatives such as virtualization, the pervasive adoption of SSD's/Flash and the movement towards cloud architectures.

# Fibre Channel: Timeline



# Fibre Channel SAN Attach Capacity Growth



Enterprise Storage Systems Capacity 2015-2019 60000 Fibre Channel 50000 NAS 40000 iSCSI 30000 Direct (DAS) FCoE 20000 InfiniBand 10000 Mainframe (ESCON/FICO 0

Source: IDC Worldwide 2015 External Storage Forecast

2019

2018

2017

Networking Flash Storage? Fibre Channel Was Always the Answer Approved SNIA Tutorial © 2016 Storage Networking Industry Association. All Rights Reserved.

2016

2015

FC #1 High Capacity Protocol

**2.5X** Fibre Channel Capacity Growth

N SAN)

# Why Fibre Channel?



#### **Non-Stop**



- Fibre Channel is the only purposebuilt, data center proven network infrastructure for storage that keeps running, no matter what
- Enables resilient IT infrastructure that optimizes availability and minimizes application disruptions
- Industry leading network reliability minimizes management resources and costs

#### **High Performance**



- Fibre Chanel delivers 128GFC
  Fibre performance ideal for high density virtualization, cloud infrastructure, and SSD storage
- Lossless and deterministic networking ensures predictable performance under high utilization
- FC dedicated networks are inherently low latency and secure

#### Scalable and Simple



- Fibre Channel fabrics are simple, and elastic networks that easily scale up and down as needed
- Backward compatibility enables scalability with new technology while leveraging legacy infrastructure
- Extensible architecture that can adapt and carry new protocols such as NVMe alongside FCP and FICON

Networking Flash Storage? Fibre Channel Was Always the Answer

Approved SNIA Tutorial © 2016 Storage Networking Industry Association. All Rights Reserved.

# How the FC Industry Innovates





# Fibre Channel Roadmap







Networking Flash Storage? Fibre Channel Was Always the Answer Approved SNIA Tutorial © 2016 Storage Networking Industry Association. All Rights Reserved.

8

# Fibre Channel speed mix





## Speed vs Cable Distance





Approved SNIA Tutorial © 2016 Storage Networking Industry Association. All Rights Reserved.

10

### Gen 6 -32GFC



WHATS NEW?

- 32GFC Specification completed
- Doubles data bandwidth over 16GFC to 6400MB/s1
- Backwards compatible two generations

- 1x single-lane
- 28.05GBaud with 64b/66b encoding
- 100 Meter on OM4
- Forward Error Correction
  - Performance / Reliability



- ANSI T11 specification complete
- Major HBA and Switch vendors are shipping today!

### Gen 6 -128GFC



WHATS NEW?

- 128GFC based on 4 lanes of 32GFC
- 25,600MB/s 4x the bandwidth of 32GFC
- Backwards compatible to single lane 32GFC or 16GFC

HOW	

- 4x 28.05GBaud lanes with 64b/66b encoding
- QSFP cable connectors
- Forward Error Correction across striped lanes
  - Performance / Reliability



- ANSI T11 completed specification in 2015
- Major Switch vendors are shipping 128GFC today!

1. Full Duplex data transmission

# **Enterprise Flash Storage Growth**



## Why Flash?

- Drives more IOPS with Lower Latency
- Guarantees Quality of Service (QoS)
- Lower cost per operation on a \$ per IO basis



Source: IDC September 2015 WW Quarterly Disk Storage Systems Forecast

# **Flash-based Arrays Growing Fastest**

## **Elements of Application Performance** Flash Storage is the Super Catalyst for Gen6





- Performance is balancing system, network, and storage capabilities
- System and network capabilities are increasing by a *linear* rate
- Storage (Flash specifically) capabilities are increasing by an *exponential* rate
- As Flash Array growth increases, it has a multiplier effect on Gen6 adoption

# Host – Switch Link Performance Profile



Networking Flash Storage? Fibre Channel Was Always the Answer

Approved SNIA Tutorial © 2016 Storage Networking Industry Association. All Rights Reserved.

**Global Education** 

# Examples of Gen6 Workload Acceleration





### FC-NVMe



WHATS NEW?

HQW

# WHEN

- New T11 Project to define NVMe over existing Fibre Channel
- Using existing Fibre Channel networks is a natural fit
  - Trusted, Proven, Ubiquitous....
- Connecting SSD storage is as easy as SAN!
- Use native FC protocol not over SCSI
- Leverage new "NVMe over Fabrics" specification
- Leverage Fibre Channel Strengths:
  - Lossless, High Perf, Zoning, Name Server, Secure
- INCITS T11 kicks off FC-NVMe working group 8/2014
- FCIA specification to be completed ~2016
- Plugfest and Joint Industry Demos Summer 2016

# What is NVMe and NVMe over Fabrics?

#### NVMe – "Non-Volatile Memory Express"

- PCI Express storage device interface SSD's connect directly to PCIe
- Simple connection interface
  - High Performance Low Latency
- Becoming the interface of choice for high performance flash drives



**Global Education** 

#### NVMe over Fabrics

- **Extend beyond PCI-Express** •
- Allow scale out over fabrics ightarrow
- Eliminate unnecessary protocol lacksquaretranslation
- NVMe Host-Controller Architecture  $\bullet$ Model for all fabric types
- Standard completion expected 6/2016



Approved SNIA Tutorial © 2016 Storage Networking Industry Assoc

## **FC-NVM**e





- FC-NVMe Standard defining the mapping of the NVMe command set over Fibre Channel
  - Device Management
  - Process Login/Logout management
  - FC Link management
- Compliments FC proven core attributes
  - NVMe protocol fits into FC4
  - Multi-Host, Multi-Port
  - Low Latency
  - Seasoned Fabric Management
  - Proven large-fabric congestion control for lossless traffic
  - Dual-Mode, NVMe can coexist with traditional SCSI traffic
  - Inherent security

Networking Flash Storage? Fibre Channel Was Always the Answer

Approved SNIA Tutorial © 2016 Storage Networking Industry Association. All Rights Reserved.

# FC-NVMe in the SAN





Approved SNIA Tutorial © 2016 Storage Networking Industry Association. All Rights Reserved.

### Switch



# WHATS NEW?

- Distributed Switch Architecture
- Makes use of Controlling Switches and FCDFs
- Looks like a single switch to the attached Fabric
- New FC Ports (A\_Port)
- New Protocols and data payloads within Distributed Switch
- Redundancy Protocol for Controlling Switches



- Complete by T11 Committee
- Completed public review
- Will be Published in 2H2016

### Fabric Services



WHATS NEW?

- Complete by T11, Letter Ballot in June
- Supports the Topology Mapping for Distributed Switches
- Companion to FC-SW-6

- Adds Controlling Switches and FCDFs in Topology
- Updated Port Attributes and Characteristics



- In public review
- Will be published in 2H2016



- FC-PI-7: 64/256GFC Underway
- FC-MSQS-3: New technical report. Test Specification for 64GFC
- Virtualization Support (VM\_IDs in Frame)
  - Proposals integrated into existing FC draft standards
  - New Frame Header and Modifications to CS\_CTL Modifications

## New Projects

- FC-SW-7 Switch Enhancements
- FC-GS-8 Fabric Services, Virtualization and Cloud
- FC-FS-7 Framing and Signaling
- FC-LS-4 Link Services



### www.t11.org

- Recently Published
  - > FC-SB-5, (INCITS BSR 485)
  - > FC-BB-6, (INCITS BSR 509)
  - > FC-MSQS-2 (INCITS TR-50)
  - > FC-PI-6 (INCITS BSR 512)
  - > FC-SP-2 AM1 (INCITS 496:2012/AM1-2015))
- Work In Progress
  - > FC-SB-6 FICON In T11 Letter Ballot
  - > FC-SW-6 Fabric Enhancements (Public Review completed, Soon to be Published)
  - > FC-GS-7 Management Enhancements (At public review)
  - > FC-LS-3 Link Service (Nearing T11 Letter Ballot)
  - > FC-FS-4 Framing and Signaling Protocol (Public Review completed)
  - > FC-NVMe Fibre Channel specific NVMe over fabrics, (1<sup>st</sup> working draft complete)
  - > FC-PI-6p 128GFC (Public Review completed, Soon to be Published)
  - > FC-PI-5 AM1 (Public Review completed, Soon to be Published)
  - > FC-PI-7 64GFC/256GFC Well underway
  - > FC Workstudy group Complete
  - > New Projects, FC-FS-5, FC-SW-7, FC-GS-8, FC-LS-4



- Fibre Channel is still the dominant protocol for Enterprise Storage
- Speed Roadmap continues to lead the industry
- Gen6: 32GFC/128GFC
- Flash Storage is the catalyst to upgrade to Gen6
- NVMe over Fabrics Fibre Channel is the smart way to bring NVMe into the storage fabric
- Industry innovations continue to drive fibre channel into the future



# Thank You www.Fibrechannel.org www.T11.org #TrustFC #FCIAnews

## Mark.Jones@Broadcom.com



# **BACKUP / ARCHIVE**



- 128GFC will require two 4-lane ports
- The ports can be either QSFP28, CFP2, CFP4 or some future 4 lane interface

12 fiber ribbon with MPO connectors 4 Tx, 4 Rx and 4 dark fibers







AOC – Active Optical Cable Up to 50 meters



DAC – Direct Attach Cable Active up to 5 meters



128GFC port can be broken out to individual 16GFC or 32GFC lanes with a breakout cable





MPO connector – 12 fi**b**ers 4 Tx, 4 Rx and 4 dark



16GFC SFP+ -2 fibers

32GFC SFP+ -2 fibers

32GFC SFP+ -2 fibers

16GFC SFP+ -2 fibers



DAC – Direct Attach Cable Active up to 5 meters



# The SNIA Education Committee thanks the following Individuals for their contributions to this Tutorial.

#### **Authorship History**

Name/Date of Original Author here: Mark Jones May 13 2016

**Updates:** 

Name/Date Name/Date Name/Date

### **Additional Contributors**

FCIA Marketing Committee

Please send any questions or comments regarding this SNIA Tutorial to <u>tracktutorials@snia.org</u>