

SNIA DEVELOPER CONFERENCE



*BY Developers FOR Developers*

**September 16-18, 2024**

**Santa Clara, CA**

# Fibre Channel Gen8 Update

128GFC, Fabric Notifications, and Managing NVMe NQNs

Howard L. Johnson, chair INCITS/Fibre Channel TC

# The SNIA Community



**200**  
Corporations,  
universities, startups,  
and individuals



**2,500**  
Active  
contributing  
members



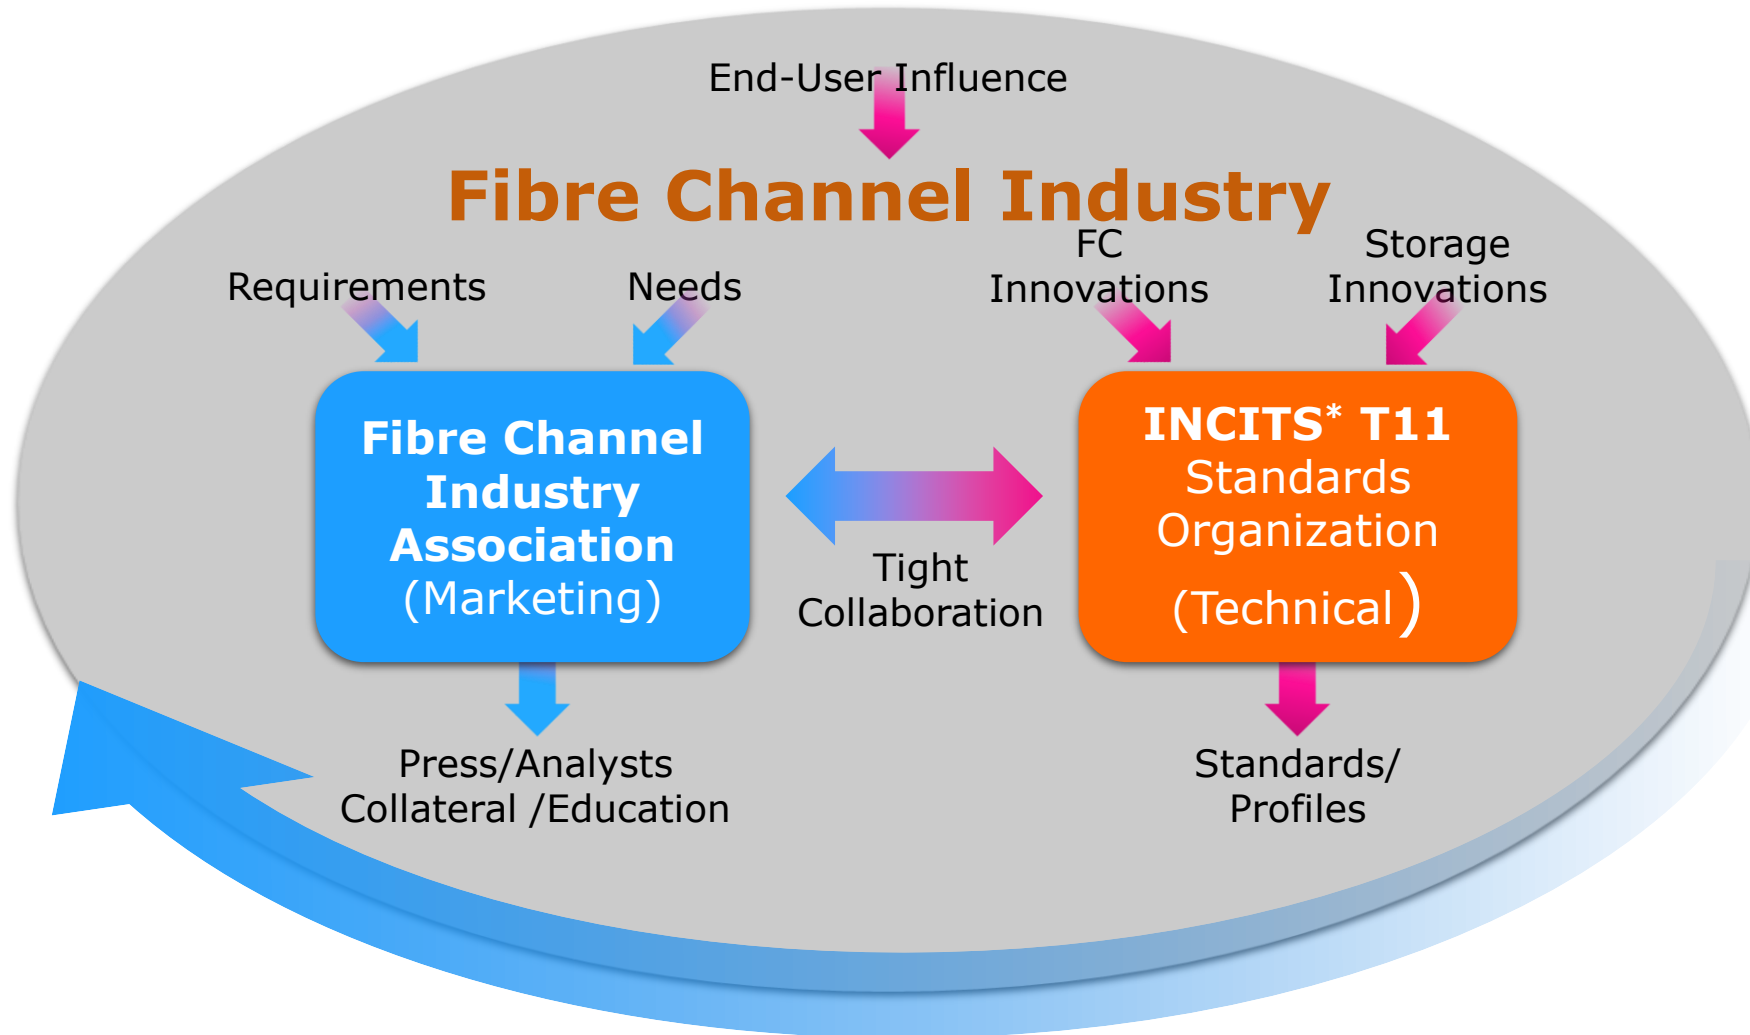
**50,000**  
Worldwide  
IT end users and  
professionals

# Agenda

- Thank you for participating
  - SNIA Storage Developer Conference 2024!
- Fibre Channel Gen 8
  - The path to 128GFC
- Fabric Notifications
  - An ecosystem-wide resiliency feature
- Managing NQNs
  - NVMe Qualified Names for use in SANs
- Summary
  - All this and more!
  - A peak into the future 😊



# How the Fibre Channel industry innovates



# Gen 8 Fibre Channel

128GFC

# Fibre Channel: Timeline



# Fibre Channel Speeds

Product Naming	Throughput (Mbytes/s)*	Line Rate (Gbaud)	T11 Specification Technically Complete (Year) †	Market Availability (Year) †
8GFC	1,600	8.5 NRZ	2006	2008
16GFC	3,200	14.025 NRZ	2009	2011
32GFC	6,400	28.05 NRZ	2013	2016
64GFC	12,800	28.9 PAM-4	2017	2020
128GFC	24,850	56.1 PAM-4	2022	2024
256GFC	49,700	112.2 PAM-4	2025	Market Demand
512GFC	TBD	TBD	2029	Market Demand
1TFC	TBD	TBD	2033	Market Demand

“FC” used throughout all applications for Fibre Channel infrastructure and devices, including edge and ISL interconnects. Each speed maintains backward compatibility at least two previous generations (I.e., 32GFC backward compatible to 16GFC and 8GFC)

\*These numbers are representative throughput values for the line rate and are payload dependent

† Dates: Future dates estimated

# Fibre Channel 128GFC Requirements

- **Throughput**
  - 128GFC doubles the throughput of 64GFC
- **Error correction**
  - Corrected bit-error-rate (BER) target of **1e-15**
- **Compatibility**
  - 128GFC had to be backward compatible to 64GFC and 32GFC
  - Backward compatibility and “plug and play” to utilize existing infrastructure with new speeds is always a “**must have**” for FC development
- **Distance**
  - 100 meters
    - Multi-mode short reach optical variant using OM4/OM5 cable plants
  - 10 kilometers
    - Single mode optical variant
- **Reuse existing cable assemblies**
  - Must plug into 128GFC capable products
  - LC (connector) and SFP+ (form factor)





# 128GFC Modulation

## ■ Modulation


- Refers to the signal levels that are on the “wire” (physical interface) whether optical or electrical
- For 64GFC studies by the Fibre Channel committee and other committees determined that moving to PAM4 modulation would be “easier” from a component and IP perspective than staying with the NRZ/PAM2 modulation and doubling the “wire” rate to 57.8Gb

## ■ 64GFC and 128GFC use PAM4

- 128GFC uses PAM4 modulation and is 112.2Gbps (56.1Gb)
- 64GFC also uses PAM4 modulation and is 57.8Gbps (28.9Gb)

## ■ All lower FC speeds use PAM2 (NRZ)

- 32GFC has a NRZ/PAM2 line rate of 28.05Gb



As serial data rates surpass 32Gb/s per channel, signal impairments caused by increasing bandwidth necessitate the high-speed serial data technology to shift from simple NRZ (non-return to zero PAM2) signal modulation to the bandwidth efficient PAM4 (4-level pulse amplitude modulation).

# 128GFC Forward Error Correction

- Forward Error Correction (FEC) is mandatory for all types of 128GFC links
- How it works
  - The transmitter encodes the data stream in a redundant way using an error correcting code
- 128GFC uses a Reed Solomon block code
  - The code used for 128GFC is RS(544,514)
    - Allows correction of single bit errors or burst errors for 15 ten bit symbols out of 5140 bits sent
- FEC has been used in previous FC variants
  - 64GFC FEC uses RS(544,514)
  - 32GFC FEC uses RS(528,514)
  - 16GFC FEC uses (2112,2080) and is optional

# 128GFC Speed Negotiation

- Fibre Channel link bring up phases
  - Link Speed Negotiation -> Optical Module bring-up -> Transmitter Training -> Mission Mode
  - Fibre Channel devices typically support three speeds (e.g. 32/64/128)
- Link Speed Negotiation (LSN)
  - Baud rate 32G (28.05Gb) NRZ
  - Advertise supported speeds in Extended Marker (32G, 64G, 128G)
  - See IEEE Clause 72 Training Signal
    - Exchanged between link partners uses Differential Manchester Encoding (DME), which runs at a lower signaling rate
      - High probability of error free exchange
- Training Signal negotiates capabilities between transmitter and receiver
  - 128G FEC type
    - FEC transmission mode (interleaved vs single, symbol forward vs bit interleave)
  - Speed Negotiation (SN) field
    - 0 = LSN finished
    - 1 = LSN in progress

# 128GFC Milestones

## ■ FC-PI-8

- 128 GFC single lane specification
- INCITS Fibre Channel/T11 complete October 2022
- ANSI accepted October 2023

## ■ FC-FS-6

- 128 GFC single lane framing and signaling specification
- INCITS Fibre Channel/T11 complete October 2023
- ANSI accepted June 2024

## ■ FC-PI-9 project

- 256 GFC single lane specification
- INCITS Fibre Channel/T11 TC work started December 2022

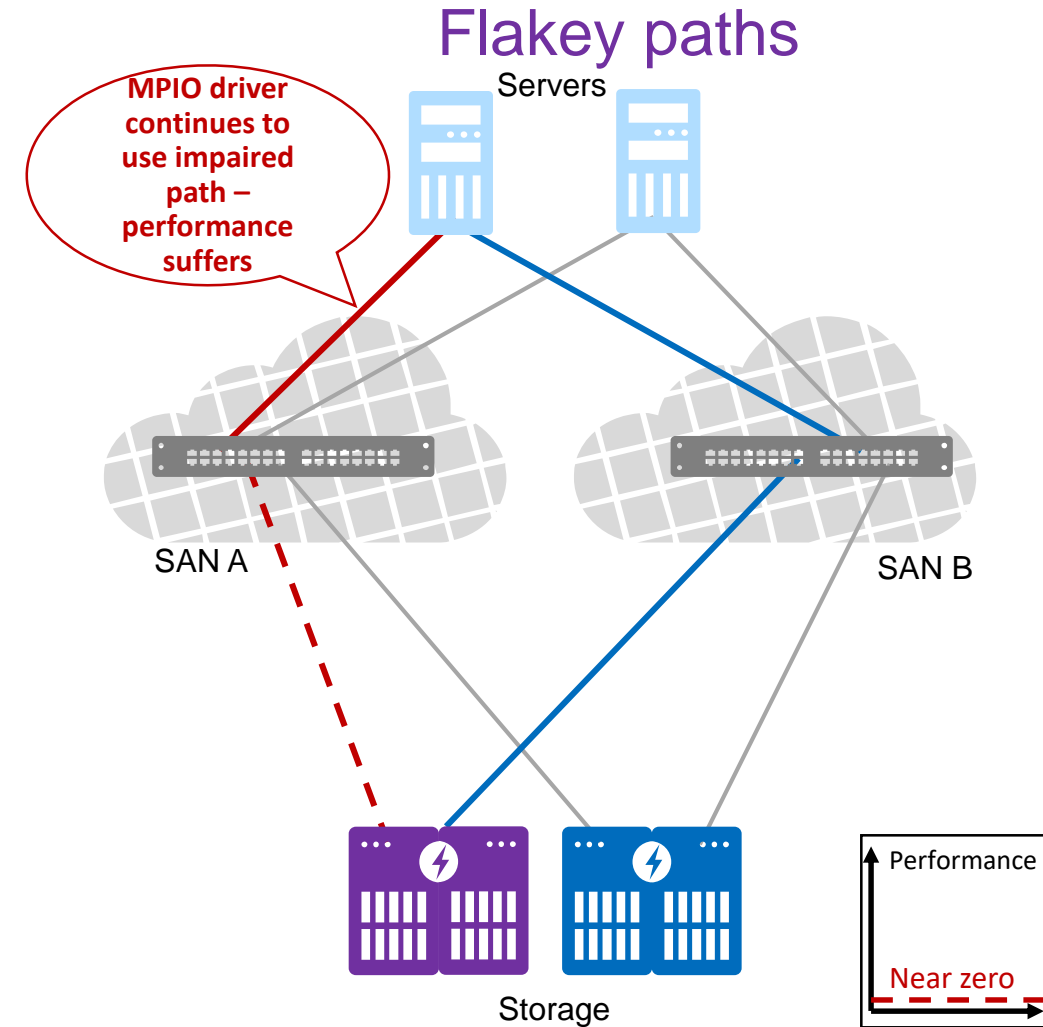


# Gen 8 Fibre Channel

Fabric Notifications

# The Problem

- Persistent, intermittent errors
  - Significant role in customer escalations
  - Difficult for traditional solutions to resolve
  - Required manual intervention increases mitigation costs
  - MPIO solutions struggle with resolution, which impacts the dual fabric paradigm
- Causes
  - Marginal cables, SFPs, connections, etc
  - Congestion due to lost credit, credit stall, or oversubscription
- Why now?
  - Fibre Channel solutions are mature and diversified
  - Identification and mitigation tools have evolved
  - Customers are demanding more automation



# The Solution

- Fabric Notifications

- Notifications and signals
  - Generated by the fabric
  - Inform devices of impairments

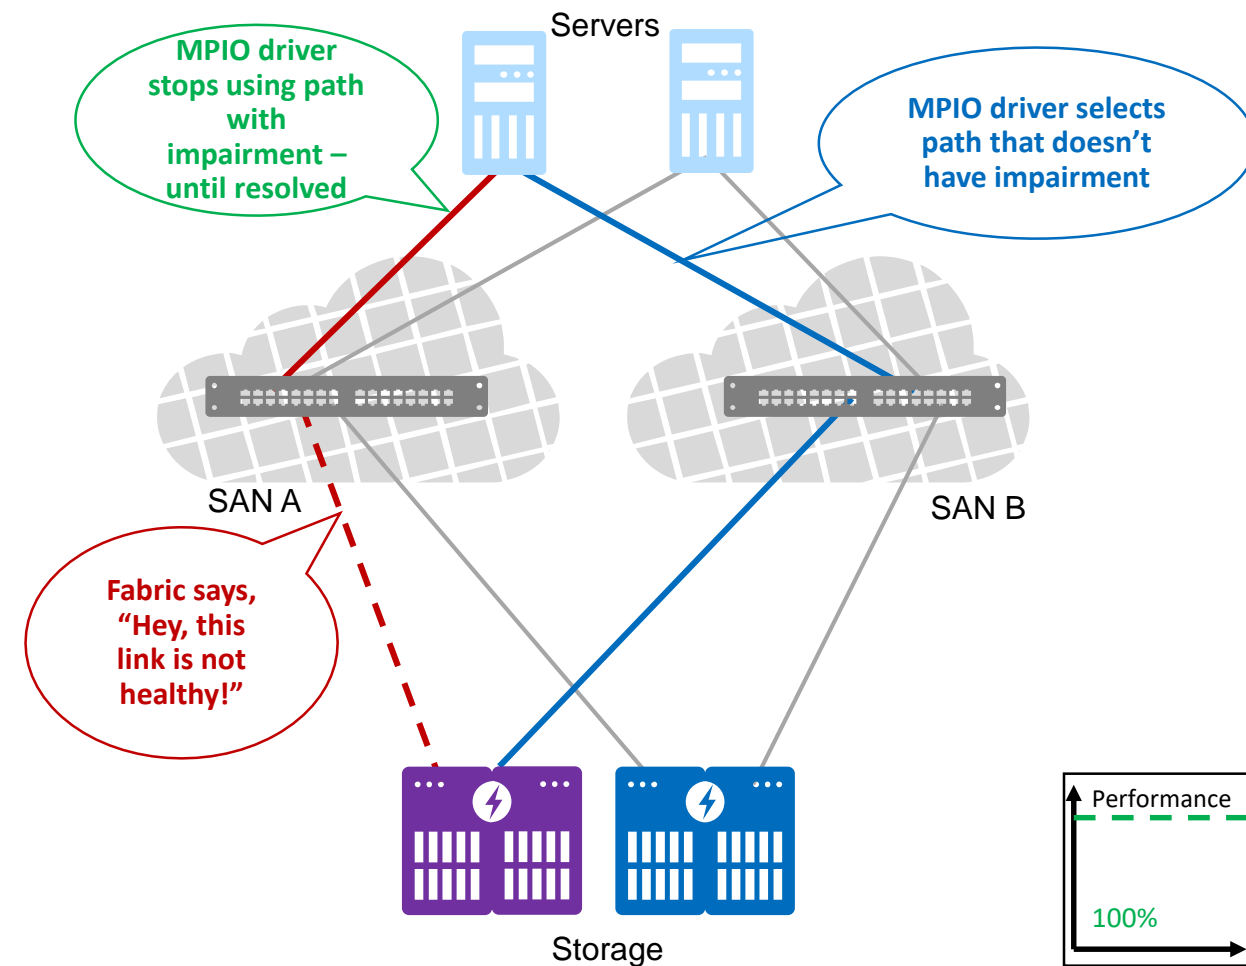
- Notifications

- Reporting: Events sent to registered devices
- Diagnostics: Helps efficiently evaluate errors
- Operation: Extended Link Services (ELS)

- Signals


- Signaling: Report resource depletion to registered devices
- Diagnostics: Transmitter indicates resource usage
- Operation: Link level Primitive Signal

## Fabric Notifications



# Fibre Channel Standards


- Standards History
  - Began in December 2018
  - Fully specified in April 2022
  - Standards complete in June 2024
- Initial introduction into draft standards
  - FC-FS-6: Congestion Signals (r0.3)
    - [ANSI Standard](#)
  - FC-LS-5: Notifications (r5.01)
    - [INCITS final draft](#)
  - FC-SW-8: Fabric detection and generation (r1.01)
    - [INCITS final draft](#)





# Fabric Notifications

- Software-based FPIN
  - Extended Link Services commands
  - Fabric Performance Impact Notification (FPIN)
- Hardware-based Congestion Signal primitives
  - Defined as Primitive Signal characters
  - Warning and Alarm Signals



# Fabric Notifications

- **Link Integrity Notifications**

- Link Integrity notifications are received by MPIO drivers, which update the path selection to avoid the impaired path
- The Link Integrity notifications allow the MPIO driver to take the appropriate action for errors (e.g., CRC, ITW)

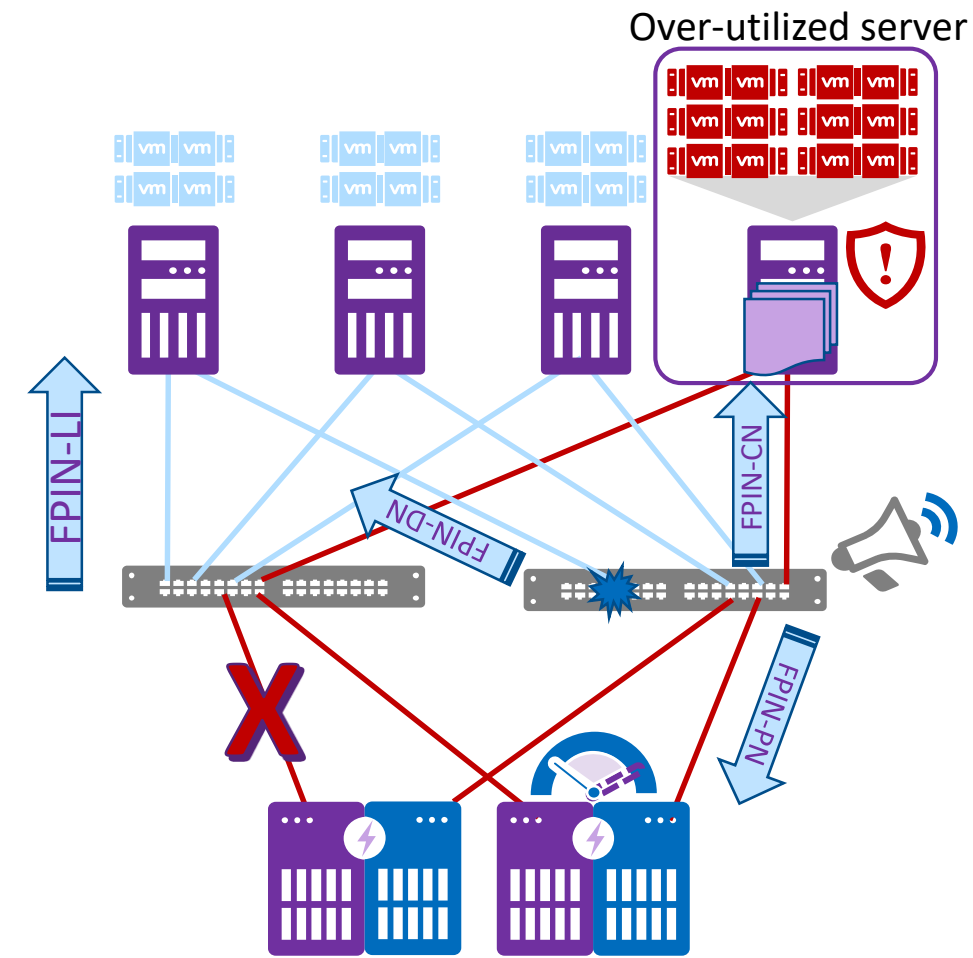
- **Congestion and Peer Congestion Notifications**

- Congestion notifications are the software equivalent of the Congestion Signal and are sent to congesting end devices
- Peer congestion notifications are sent to registered and “in-zone” peers of end devices that are experiencing congestion

- **SCSI Command Delivery Notifications**

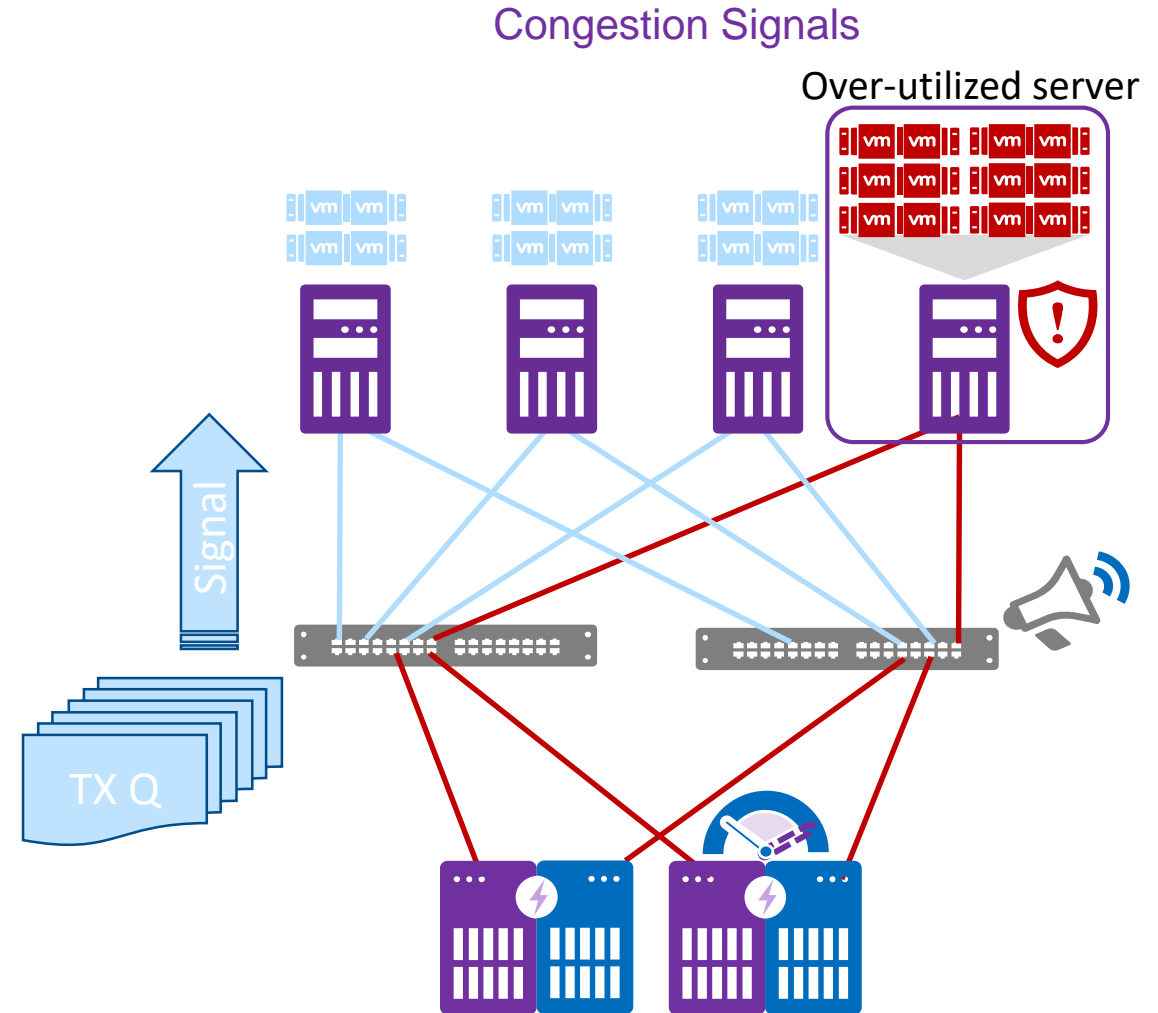
- Delivery notifications are sent when a fabric discards a SCSI command or status frame to notify the initiator of the failure

## Fabric Performance Impact Notifications (FPIN)



# Fabric Notifications

- Congestion Signals
  - Immediate feedback mechanism
  - Indicates transmission resources are consumed
- Link level communication
  - Transmitter to receiver



# Gen 8 Fibre Channel

Managing NQNs (NVMe Qualified Names)

# The Problem

## Summary

- Fibre Channel devices reference World Wide Port Names (WWPN)
- NVMe devices reference NVMe Qualified Names (NQN)
- Both must be configured correctly for FC-NVMe solutions to function

## Example

- NVMe Storage define access control using the host NQNs
- NVMe Hosts “connect” to specific storage Sub-NQNs to access Namespaces
- FC-NVMe requires both WWPNs (zoning) and NQNs to be configured correctly

## Challenge

- SAN administrators access/view WWPNs from Fabric management tools
- However, the NQNs at the NVMe layer are not visible to the Fibre Channel Fabric

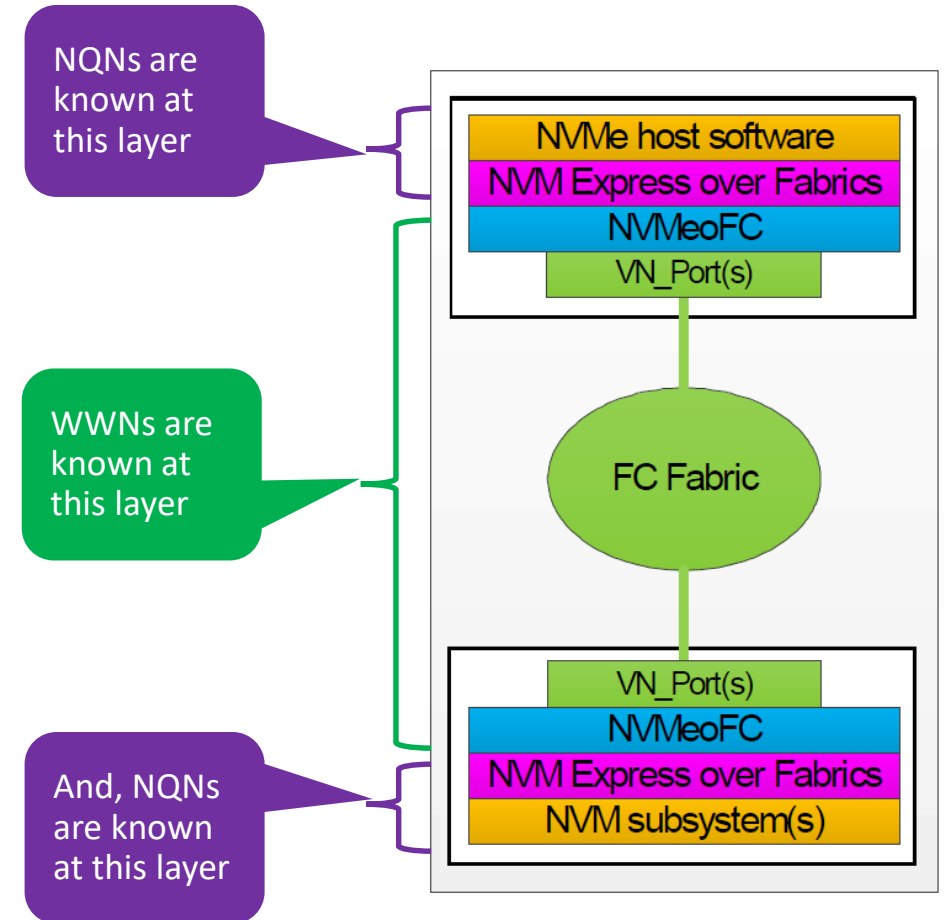
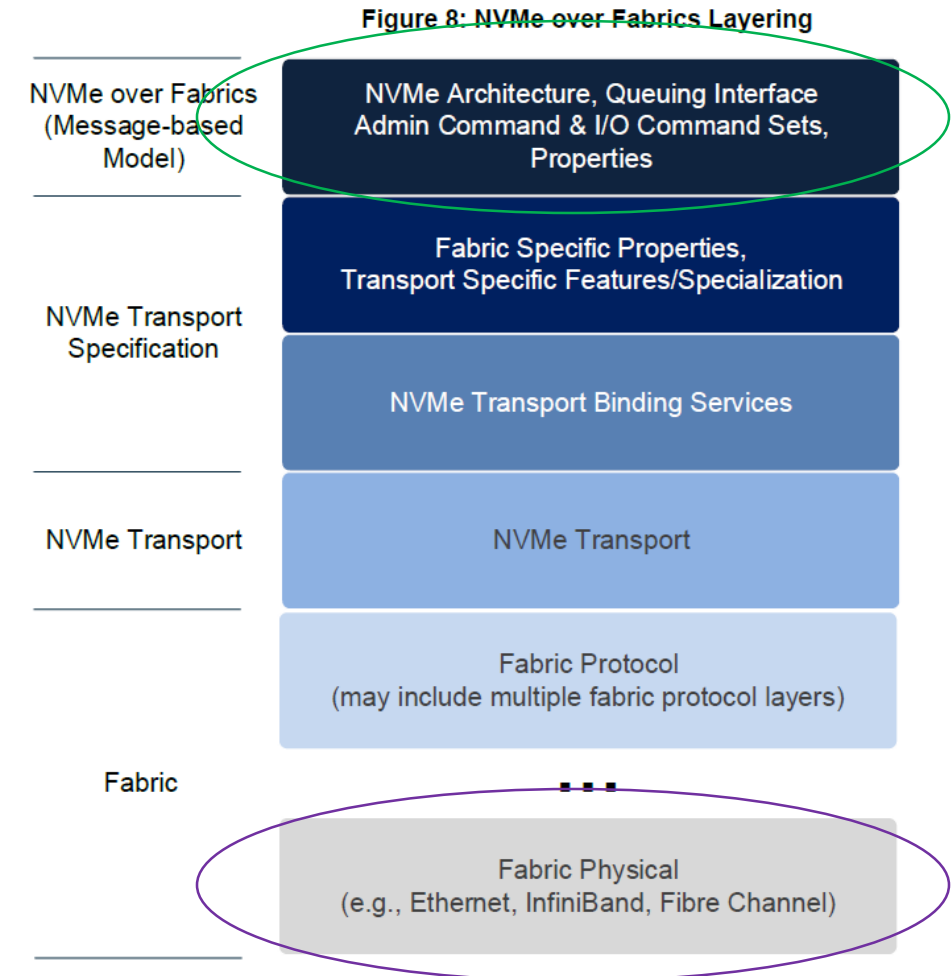


Figure 1 – NVMeoFC protocol layers

# The conundrum – NQNs and FC-NVMe

- Connection requirements
  - NVMe connection requirements (parameters)
    - CONNECT(HOSTID, CNTLID, SUBNQN, HOSTNQN)
  - Fabric Physical layer (Fibre Channel)
    - ZONE(WWN0, ..., WWNn)
- NQNs are opaque data to FC-NVMe
  - Provided by the ULP (NVMeoF)
  - Used in Create Association NVMe\_LS (FC-NVMe)
    - Validated for consistency, not for authenticity
    - See FC-NVMe clause 4.4.1
- NVMe connections require “both”
  - NVMeoF and Fabric Physical layer connectivity parameters are independent of each other
  - WWNs for zoning and NQNs for connection requests
    - If either is incorrect, IO is not possible



# The Solution

- Management server
  - Maintains NQNs as ULP Name objects
    - Creates associations between WWPN and accessible NQNs
  - Fibre Channel devices register and query for ULP Names
  
- FC-GS-9
  - Upper Layer Protocol (ULP) Name object
  - Upper Layer Object Server
    - Registration/deregistration
    - Query
  
- FC-NVMe-3
  - NVMe\_Port registration of the NQNs associated with the NVMe\_Port

Table 524 – ULP Name object

Item	Size (bytes)
FC-4 TYPE Code	1
ULP Name Type	1
Reserved	1
ULP Name Length	1
ULP Name	n

Table 10 – ULP Name Type field for FC\_NVMe

Value	ULP Name Type
0	Not specified
1	NVM Subsystem NQN
2	Host NQN
3	Discovery Service NQN
4 to 255	Reserved



# Gen 8 Fibre Channel

Miscellaneous Fun Stuff



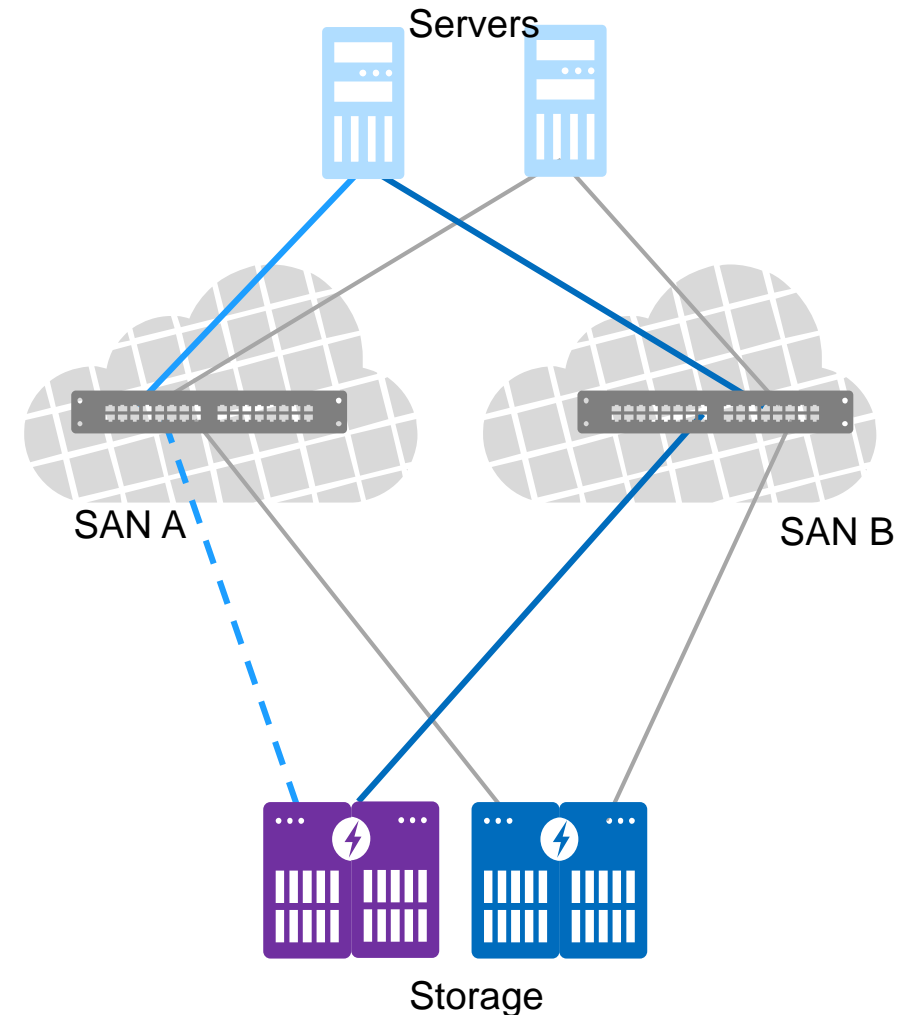
# Revving up for Fibre Channel 256GFC

- **Throughput**
  - 256GFC doubles the throughput of 128GFC
- **Error correction**
  - Corrected bit-error-rate (BER) target of **1e-15**
- **Compatibility**
  - 256GFC backward compatibility with 128GFC and 64GFC
- **Distance**
  - **100 meters (watch this space)**
    - Multi-mode short reach optical variant using OM4/OM5 cable plants
  - 10 kilometers
    - Single mode optical variant
- **Reuse existing cable assemblies**
  - Must plug into 256GFC capable products
  - LC (connector) and SFP+ (form factor)
- **FC-PI-9 project is underway!**



# Fibre Channel Automation

- Platform Name\_Identifier
  - WWN for the enclosure
  - Establishes Port->Node->Platform hierarchy
- Port Notifications
  - Extension of Fabric Notifications
  - Supports Port Decommission/Recommission
- Virtual Machine Identification
  - VM attributes for VE and ULP servers



# About the Fibre Channel Industry Association ([FCIA](#))



**25+ Years\***  
Promoting Fibre  
Channel Technology



**Industry Leading**  
Member Companies



**142M+ FC Ports**  
Shipped Since 2001

Working in cooperation with the Storage Networking Industry Association ([SNIA](#)) to promote [storage solutions](#)!

# Summary

- Thank you for attending SNIA Storage Developer Conference 2024!
- Fibre Channel speed
  - The path to 128GFC
- Fibre Channel reliability
  - An ecosystem-wide resiliency feature
- Fibre Channel flexibility
  - NVMe Qualified Names for use in SANs
- Fibre Channel future
  - Gen 9 and 256GFC
  - Security for the CNSA 2.0 world
  - Integration and automation

## Call to Action

Demand stability, reliability, speed, and flexibility from  
your Storage Area Networking environment – demand  
Fibre Channel! 😊

[howard.johnson@broadcom.com](mailto:howard.johnson@broadcom.com)

# Thank You!

Fibre Channel Gen 8 Update

128GFC, Fabric Notifications, and Managing NQNs



Please take a moment to rate this session.

Your feedback is important to us.

# References

Fibre Channel Gen 8 Update

128GFC, Fabric Notifications, and Managing NQNs

# References

- FCIA BrightTalk [Channel](#)
- Highlighted
  - “128GFC: A Preview of the New Fibre Channel Speed” ([BrightTalk](#))
  - “Fibre Channel Data Center Interconnects (DCI): 64G FC and More” ([BrightTalk](#))
  - “NVMe over FC: Deep Dive in Protocol, Architecture and Use Cases” ([BrightTalk](#))
  - “Introducing Fabric Notifications, From Awareness to Action” ([BrightTalk](#))
- Of Interest
  - “Inside a Modern Fibre Channel Architecture” (BrightTalk [Part 1](#) & [Part 2](#))
  - “Advance Your Career with Fibre Channel Knowledge” ([BrightTalk](#))



# Fibre Channel Gen 8 Standards References

- Fibre Channel 128GFC
  - FC-FS-6 ([ANSI](#)) ([INCITS 562-2024](#))
  - FC-PI-8 ([ANSI](#)) ([INCITS 560-2023](#))
- Fibre Channel Fabric Notifications
  - FC-LS-5 ([INCITS](#)) ([INCITS 569-2024](#))
  - FC-SW-8 ([INCITS](#)) ([INCITS 568-2024](#))
- Fibre Channel NQNs
  - FC-GS-9 ([INCITS](#)) ([INCITS 570-2024](#))
  - FC-LS-5 ([INCITS](#)) ([INCITS 569-2024](#))

# Fabric Notifications Solutions

## Fabrics and Storage

- Fabrics
  - Brocade
    - FOS 9.0.0, FOS 9.2.1
  - Cisco
    - NX-OS 9.2(1), NX-OS 9.4(2a)
  - Emulex
    - LPe3100x, LPe3200x, LPe3500x-M2
  - Marvell
    - QLE269x, QLE274x, QLE277x, QLE28xx
- Storage
  - Dell
    - PowerMax InfoScale 10.1
  - NetApp
    - OnTap 9.10
  - PureStorage
    - Oxygen

## Multipath solutions

- Operating systems
  - IBM AIX
    - 7.2 TL5, 7.3 TL2
  - Redhat
    - RHEL 8.3 / EPEL 8, RHEL 9.0 / RHEL 8.7, RHEL 9.2 / RHEL 8.8, RHEL 9.3 / RHEL 8.9, RHEL9.4 / RHEL 8.10
  - SuSE
    - SLES15 SP4, SLES 15 SP5, SLES 15 SP6
  - Vmware
    - ESXi 8.0, ESXi 8.0U1, ESXi 8.0U2
- Multipath software
  - Dell
    - PowerPath 7.4
  - Veritas
    - InfoScale 8.0.2 DMP

# Fabric Notifications References

## ■ Webinars

- "Introducing Fabric Notifications, From Awareness to Action" ([FCIA BrightTalk presentation](#))
  - [SNIA SDC 2021 EMEA](#) virtual session ([Part One](#) and [Part Two](#))
  - [SNIA SDC 2021](#) virtual session ([Presentation](#))
- "Fabric Notifications – An Update from Awareness to Action"
  - [SNIA SDC 2022](#) live session ([Presentation](#))
- "The Evolution of Congestion Management in Fibre Channel"
  - [SNIA BrightTalk](#) webinar ([Presentation](#))

## ■ Industry

- IBM Power Community – [AIX Support for Fabric Congestion Notification](#)
- PureStorage [blog](#)
- Marvell SAN congestion mitigation [Video](#)

## ■ Articles

- The Autonomous SAN ([FCIA Solutions guide](#))
- Fabric Notifications Technical Brief ([Brocade Whitepaper](#))
- MPIO Load Balancing Recommendations ([Brocade Whitepaper](#))
- Cisco Fabric Notifications [Blog](#)
- Dell Fabric Notifications [Technical Brief](#)
- Emulex Fabric Notifications [Product Brief](#)

## ■ Videos

- Fabric Notifications Primer ([Brocade video](#))
- Fabric Notifications using RHEL 8.3 ([Brocade video](#))
- Fabric Notifications using IBM AIX 7.2 TL5 ([Brocade video](#))

# Fibre Channel Signaling – Rate Abbreviations

Abbreviation	Signaling rate	Number of Lanes	Data rate
1GFC	1.0625 MBd	1 (NRZ)	100 MB/s
2GFC	2.125 MBd	1 (NRZ)	200 MB/s
4GFC	4.250 MBd	1 (NRZ)	400 MB/s
8GFC	8.500 MBd	1 (NRZ)	800 MB/s
16GFC	14.025 MBd	1 (NRZ)	1600 MB/s
32GFC	28.050 MBd	1 (NRZ)	3200 MB/s
64GFC	28.900 MBd	1 (PAM4)	6400 MB/s
128GFC	56.1 MBd	1 (PAM4)	12425 MB/s
256GFC	112.200 MBd	1 (PAM4)	24850 MB/s



# EOF

That's All Folks!