SNIA ESF | ETHERNET STORAGE

FCoE vs. iSCSI vs. iSER
A Great Storage Debate

Live Webcast
June 21, 2018
10:00 am PT

Today's Presenters

Tim Lustig
Mellanox Technologies

J Metz
SNIA Board of Directors
Cisco

Rob Davis
Mellanox Technologies

Saqib Jang
Chelsio Communications

© 2018 Storage Networking Industry Association. All Rights Reserved.

SNIA-At-A-Glance

SNIA-At-A-Glance

170
industry leading organizations

2,500
active contributing members

50,000
IT end users & storage pros worldwide

Learn more: **snia.org/technical** **@SNIA**

© 2018 Storage Networking Industry Association. All Rights Reserved.

SNIA Legal Notice

SNIA ESF | ETHERNET STORAGE

- The material contained in this presentation 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.
- 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.

© 2018 Storage Networking Industry Association. All Rights Reserved.

Agenda

- A Brief Background
  - FCoE – J Metz
  - iSCSI – Saqib Jang
  - iSER – Rob Davis
- Compare and Contrast
- How do you decide?
  - Scalability, in-house expertise, use case

© 2018 Storage Networking Industry Association. All Rights Reserved.

J Metz

Fibre Channel over Ethernet - FCoE

© 2018 Storage Networking Industry Association. All Rights Reserved.

In The Beginning…

- There were two philosophies
  - Deterministic Networks
  - Non-Deterministic Networks
- Similar, but not compatible

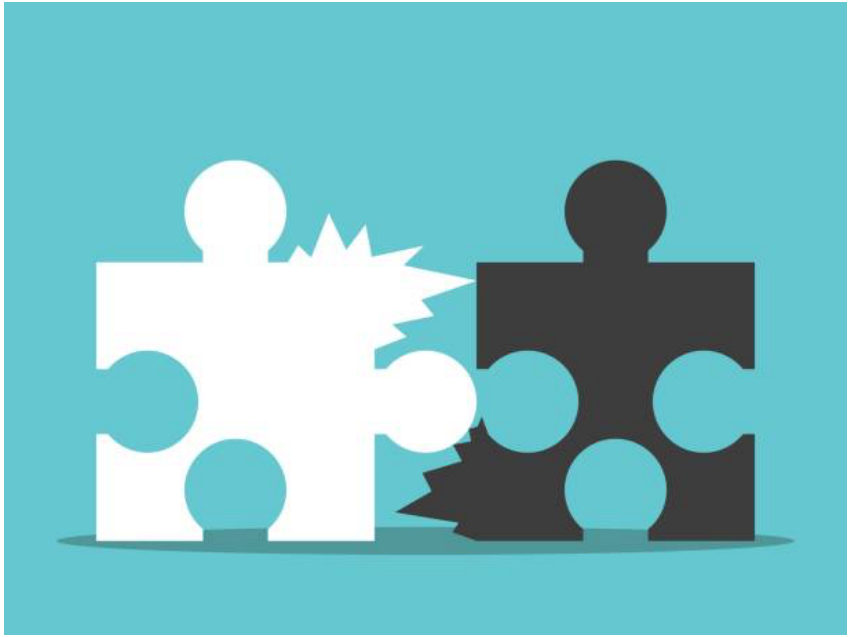

© 2018 Storage Networking Industry Association. All Rights Reserved.

What's the Problem?

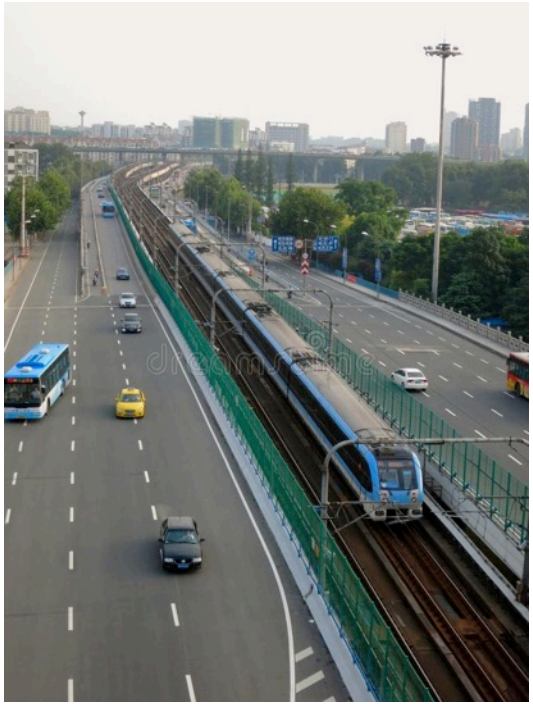

- Ethernet is non-deterministic
  - Flow control is destination-based
  - Relies on TCP drop-retransmission / sliding window
- Fibre-Channel is deterministic
  - Flow control is source-based (B2B credits)
  - Services are fabric integrated (no loop concept)

© 2018 Storage Networking Industry Association. All Rights Reserved.

Standards for Unified I/O with FCoE

- FCoE is fully defined in FC-BB-5 standard
- FCoE works alongside additional technologies to make I/O Consolidation a reality

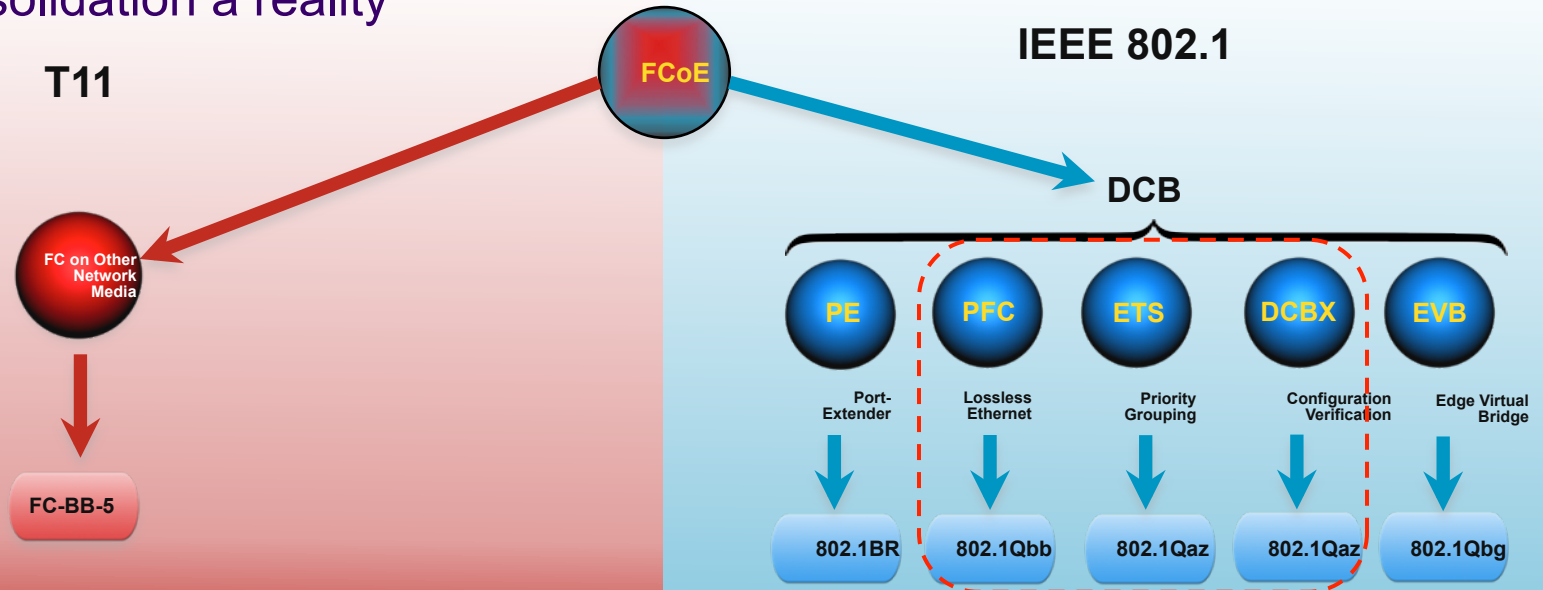

Best of Both Worlds

- From a Fibre Channel standpoint its still Fibre Channel
  - FC connectivity over a Ethernet cable
- From an Ethernet standpoint its
  - Yet another ULP (Upper Layer Protocol) to be transported

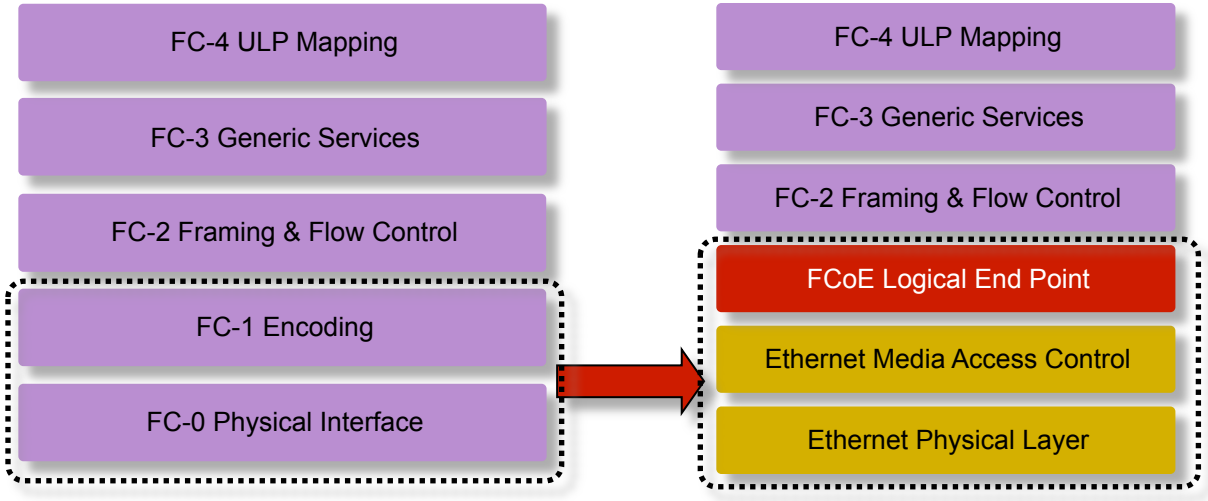

© 2018 Storage Networking Industry Association. All Rights Reserved.

Fibre Channel Encapsulation

- From a Fibre Channel standpoint its still Fibre Channel
  - FC connectivity over a Ethernet cable
- From an Ethernet standpoint its
  - Yet another ULP (Upper Layer Protocol) to be transported

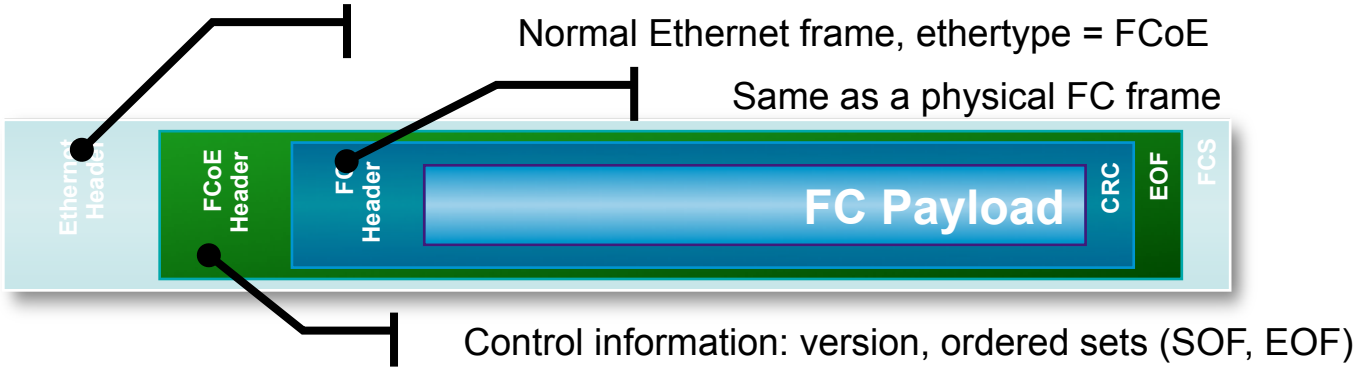

© 2018 Storage Networking Industry Association. All Rights Reserved.

Solving the Problem – Priority Flow Control

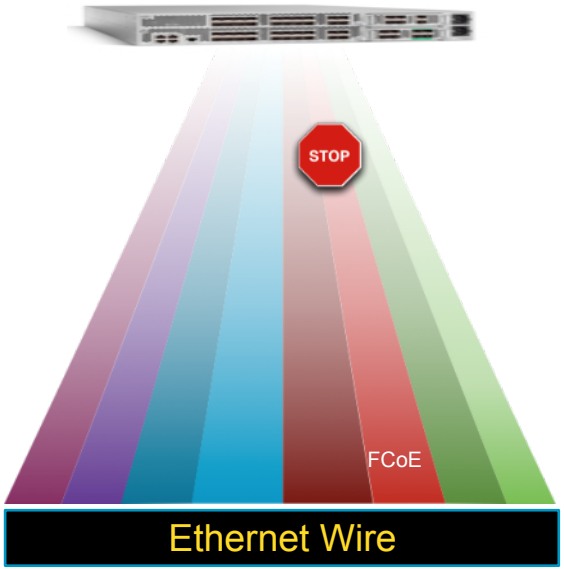

- Lossless Ethernet
  - Not only for FCoE traffic
- PFC enables Flow Control on a Per-Priority basis
  - Ability to have lossless *and* lossy priorities at the same time on the same wire
  - Allows FCoE to operate over a lossless priority independent of other priorities
- Other traffic assigned to other CoS will continue to transmit and rely on upper layer protocols for retransmission

© 2018 Storage Networking Industry Association. All Rights Reserved.

History of Unified Fabrics

Using FCoE to Provide Consolidated I/O

SNIA ESF | ETHERNET STORAGE

- Traditional Data Centers had separation at the host
  - Separate Ethernet-based networks and Fibre Channel-based networks
  - Multiple cards per server
    - 2 HBAs
    - Average of 6 (or more!) NICs per server
    - High underutilization drives up unnecessary power, cooling, and asset costs

© 2018 Storage Networking Industry Association. All Rights Reserved.

History of Unified Fabrics

Using FCoE to Provide Consolidated I/O

SNIA ESF | ETHERNET STORAGE

- Access-Layer Convergence
  - Consolidate I/O on 10G links
  - Drastically reduced CapEx and OpEx
  - Multiprotocol connectivity eased purchasing decisions for server refreshes
  - Prepared Data Centers for VM mobility requirements
    - Any VM could connect to FC storage if necessary, not just the ones with HBAs pre-installed

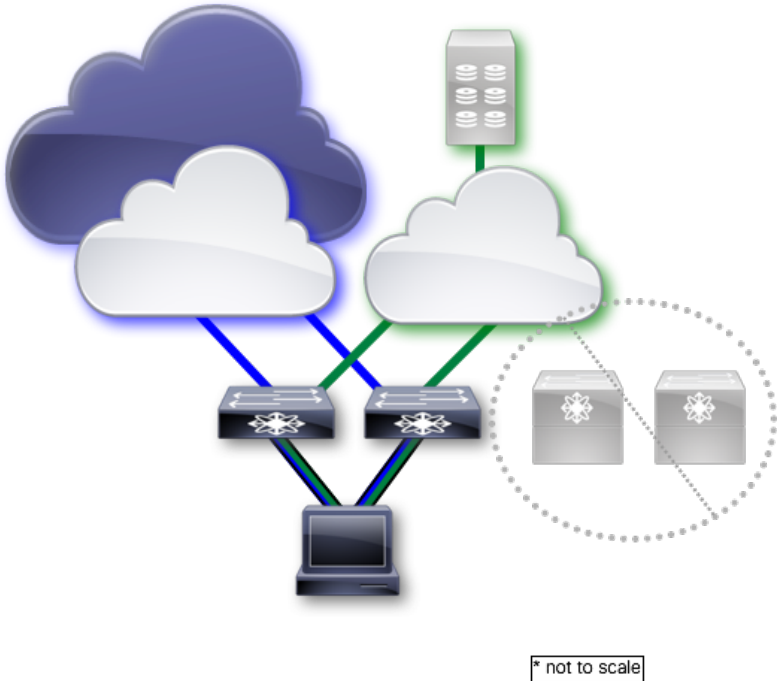

© 2018 Storage Networking Industry Association. All Rights Reserved.

History of Unified Fabrics

Using FCoE to Provide Consolidated I/O

SNIA ESF | ETHERNET STORAGE

- Multihop Convergence
  - Standardize on Ethernet assets
    - One physical infrastructure
  - Keeps best practices for both Ethernet and Fibre Channel
- Reduction of Additional Equipment
- Protected investment and future-proofed deployments

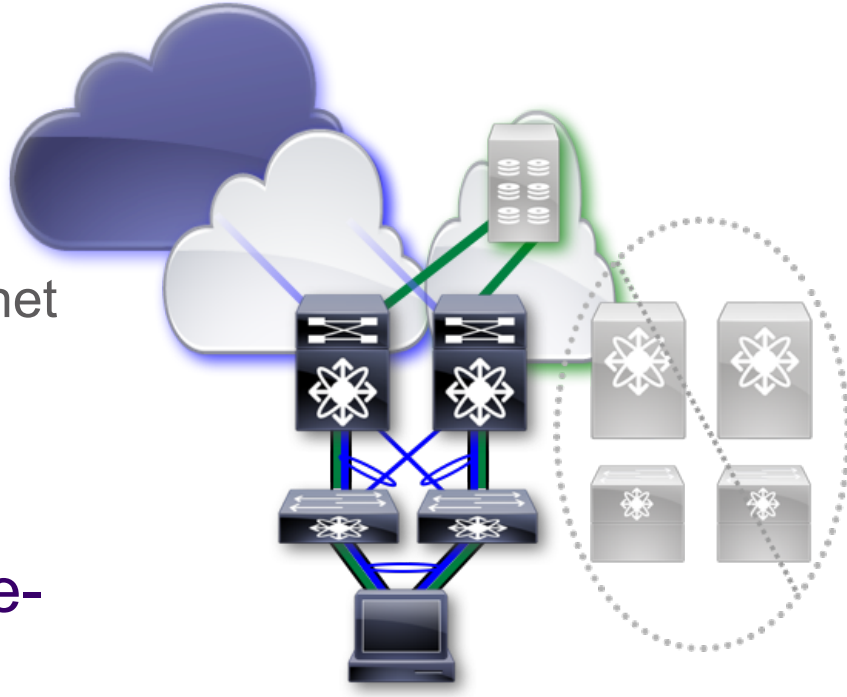

* not to scale

© 2018 Storage Networking Industry Association. All Rights Reserved.

Advanced Design

- Dynamic FCoE – Clos networks
  - Use Ethernet Equal-Cost Multipathing (ECMP) to provide load-balanced traffic across entire topology
- Greater resiliency and robustness across core (spine)
- Dynamic configuration of Inter-Switch Links (ISLs)

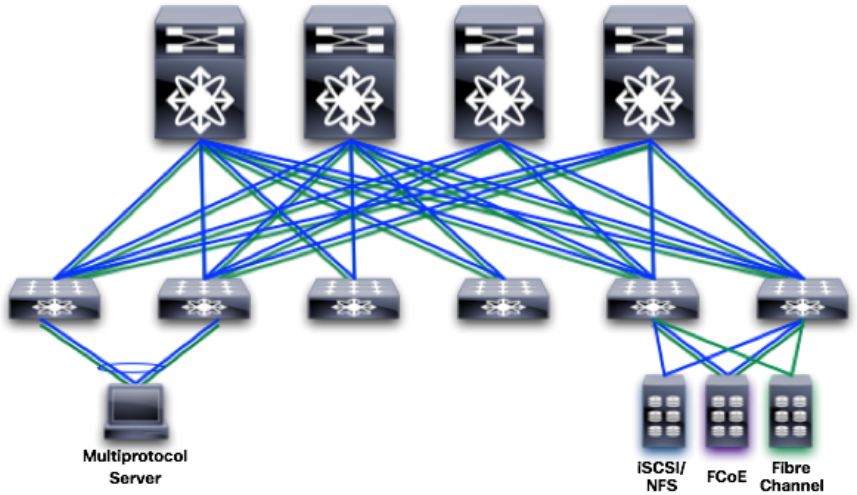

© 2018 Storage Networking Industry Association. All Rights Reserved.

Flexibility

- Can be used anywhere FC is used
  - Server-to-TOR
  - Server-to-Storage
  - Massive bandwidth ISLs
- Runs both SCSI and NVMe ULPs
- Zero-Copy transfers (same as RDMA)
- Flexible topology considerations (edge/core, edge/core/edge, Clos)
- End-to-End Qualification

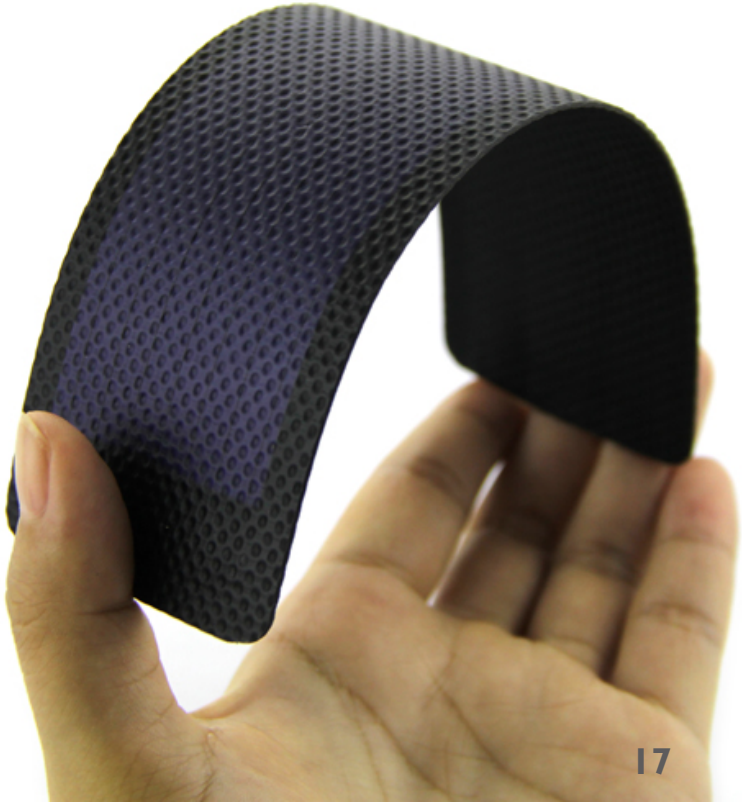

© 2018 Storage Networking Industry Association. All Rights Reserved.

Saqib Jang

iSCSI

© 2018 Storage Networking Industry Association. All Rights Reserved.

What is iSCSI (Internet SCSI)?

- Mature and widely supported Ethernet block storage network protocol
  - Standardized by IETF: RFCs 3721, 3722, 4018, 4056, 7143, etc.
- Built-in support in mainstream server operating systems
  - Windows Server, Linux, and BSD (Initiator and Target)
  - Major Hypervisors: Hyper-V, Xen, and ESX
- iSCSI offload initiator/target adapters for performance-sensitive applications
  - Complementing servers using newer multi-core CPUs and target performance scalability

© 2018 Storage Networking Industry Association. All Rights Reserved.

iSCSI Benefits for Enterprise Deployment

- Use of TCP/IP simplifies LAN/WAN deployment and operational requirements
- iSCSI software initiator is widely supported in-box capability
- Ethernet SAN protocol with built-in offload adapter "second source"
- Decoupling of server and storage SAN hardware upgrade cycles

© 2018 Storage Networking Industry Association. All Rights Reserved.

iSCSI I/O Path

- iSCSI Software
  - Software-based Protocol Processing
- iSCSI Offload
  - Protocol Bypass
  - RDMA

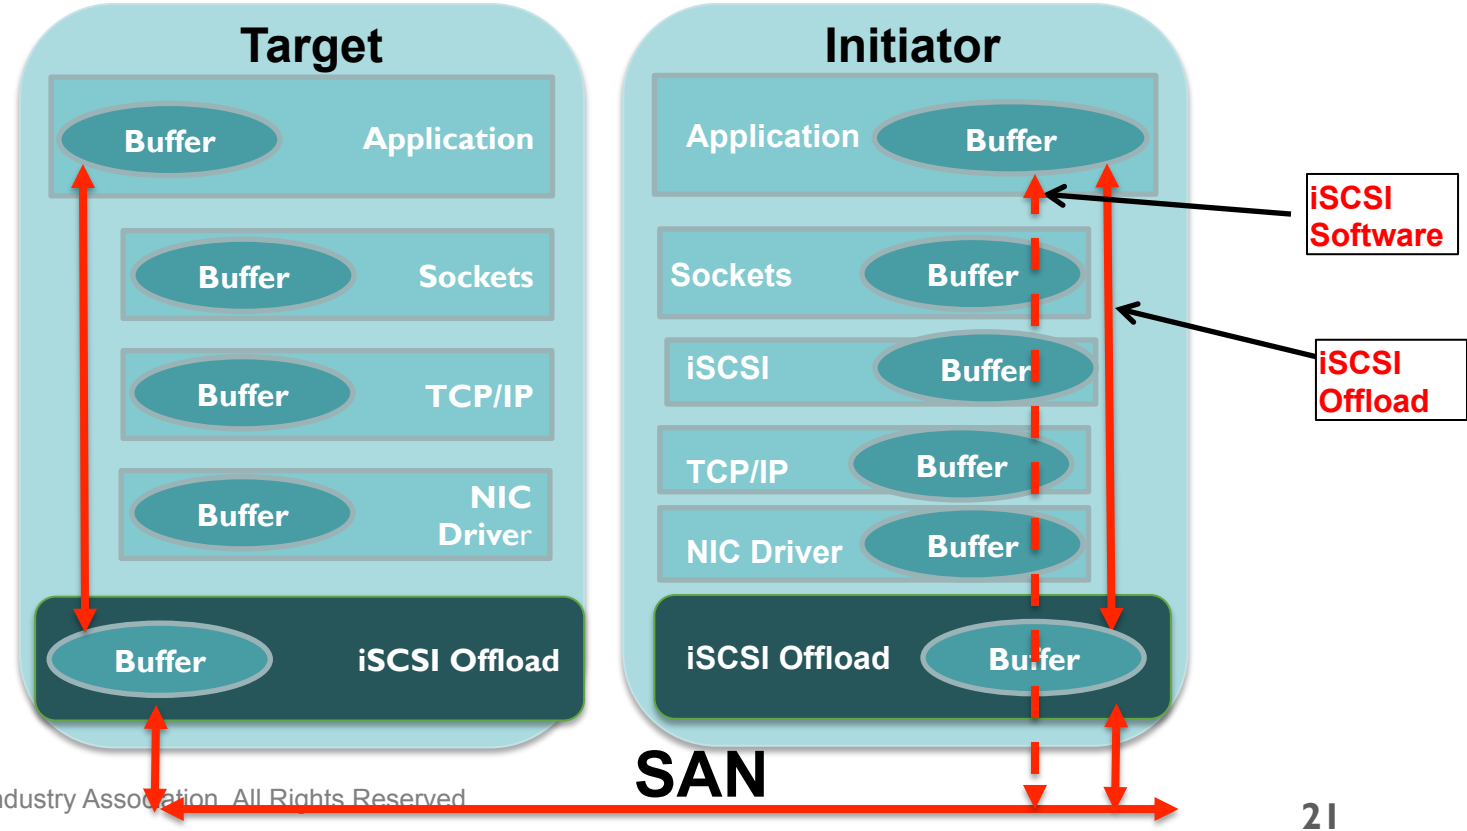

© 2018 Storage Networking Industry Association. All Rights Reserved.

100GbE iSCSI Performance

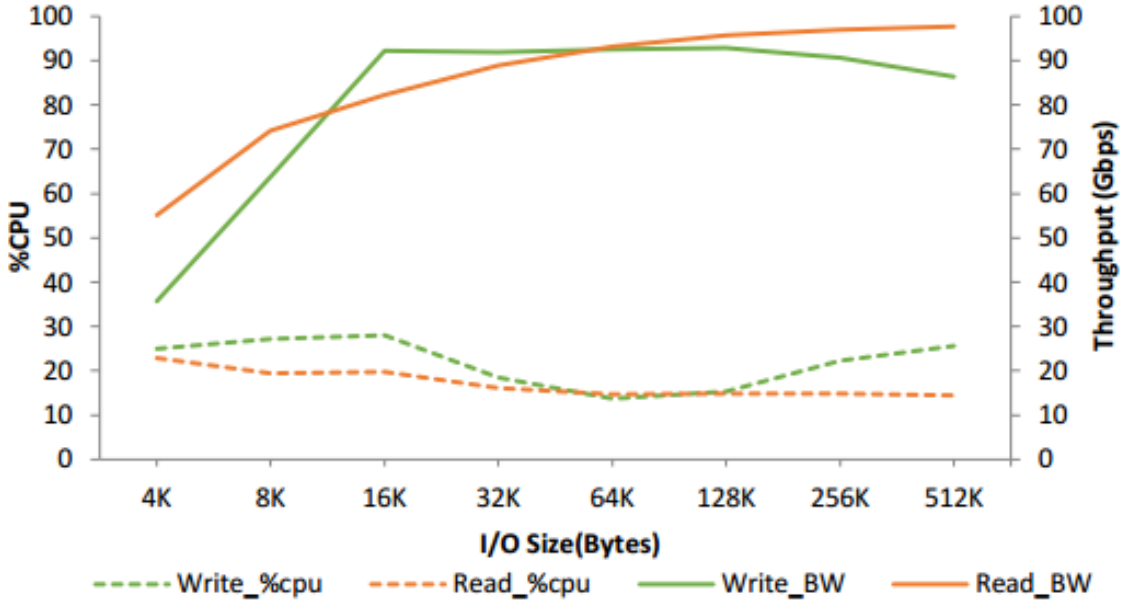

iSCSI Target - 2 Intel Xeon CPU E5-2687W v4 12-core processors @ 3.00GHz (HT enabled), 128GB RAM and RHEL 7.2 operating system using 100GbE T6 iSCSI Offload

iSCSI Initiators - 1 Intel Xeon CPU E5-1620 v4 4-core processor @ 3.50GHz (HT enabled) and 16GB of RAM using 40GbE T5 iSCSI Offload

© 2018 Storage Networking Industry Association. All Rights Reserved.

iSCSI Services

- Management
  - Mostly distributed (in clients and targets)
  - Ethernet, TCP/IP-based monitoring/troubleshooting tools
- Servers/targets/network reusable
  - Can concurrently other storage protocols: NFS, SMB, NVMe-oF
  - Object storage or scale-out filesystems
  - Compute traffic or hyper-converged infrastructure
- Reliability
  - iSCSI digest (CRC)
  - Ethernet CRC, TCP/IP checksums

© 2018 Storage Networking Industry Association. All Rights Reserved.

iSCSI Services

- Redundancy/Availability
  - Protocol: Link aggregation (LACP) or iSCSI multi-pathing
  - Physical: Duplicate Ethernet networks (optional)
- Zoning or isolation options
  - Physically or logically separate networks
  - ACLs (access control lists), VLANs (virtual LAN), VPN (virtual private network)
- Security
  - IPSEC for encryption, CHAP or RADIUS for authentication

© 2018 Storage Networking Industry Association. All Rights Reserved.

iSCSI Speeds and Feeds

- Today
  - 1G / 2.5G / 5G / 10G / 25G / 40G / 50G / 100G
- Futures
  - Coming in 2018: 200GbE (4x50) and 400GbE (8x50)
  - In the plan: 800G, 1.6T, 3.2T (dates TBD)
  - NVMe-oF: Ethernet network also supports NVMe

© 2018 Storage Networking Industry Association. All Rights Reserved.

Rob Davis

iSER

© 2018 Storage Networking Industry Association. All Rights Reserved.

iSER – iSCSI Extensions for RDMA

- Officially a standard in 2007 when IETF issued RFC-7145

L7	**Applications**		
L6	SCSI		
L5	iSCSI	iSER iWARP	iSER RoCE
L4	TCP		UDP
L3	IP (Network)		
L2	Ethernet (Link)		
L1	Ethernet (Physical)		

© 2018 Storage Networking Industry Association. All Rights Reserved.

iSER – iSCSI Extensions for RDMA

- Officially a standard in 2007 when IETF issued RFC-7145
- Features and characteristics are almost the same as iSCSI
  - Leverages management and tools (security, HA, discovery...)

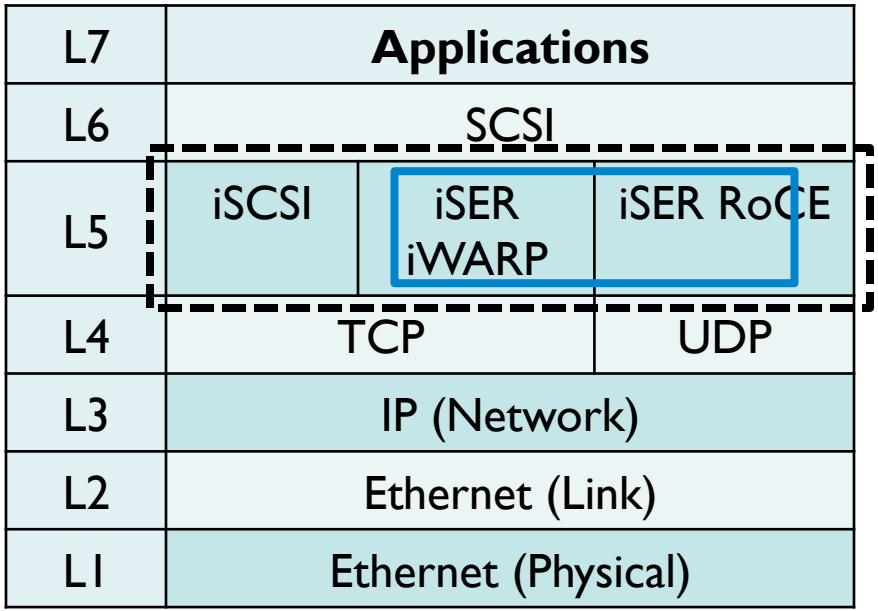

© 2018 Storage Networking Industry Association. All Rights Reserved.

iSER – iSCSI Extensions for RDMA

- Officially a standard in 2007 when IETF issued RFC-7145
- Features and characteristics are almost the same as iSCSI
  - Management and tools (security, HA, discovery...)
- Major difference is performance

L7	**Applications**		
L6	SCSI		
L5	iSCSI	iSER iWARP	iSER RoCE
L4	T		
L3			
L2			
L1	E		

© 2018 Storage Networking Industry Association. All Rights Reserved.

Performance Difference – iSCSI vs. iSER

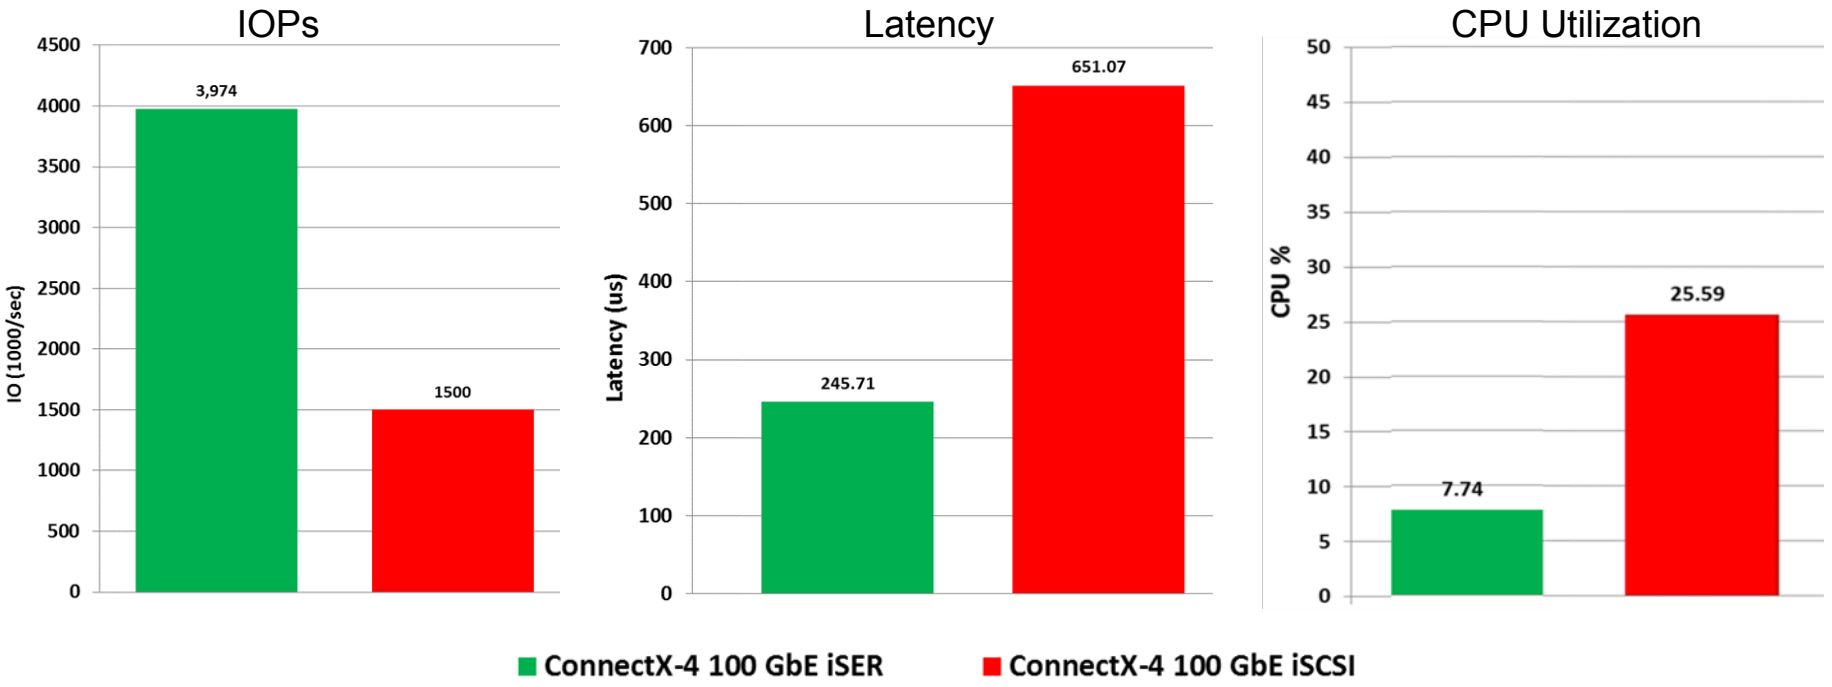

© 2018 Storage Networking Industry Association. All Rights Reserved.

Performance Difference – iSCSI vs. iSER

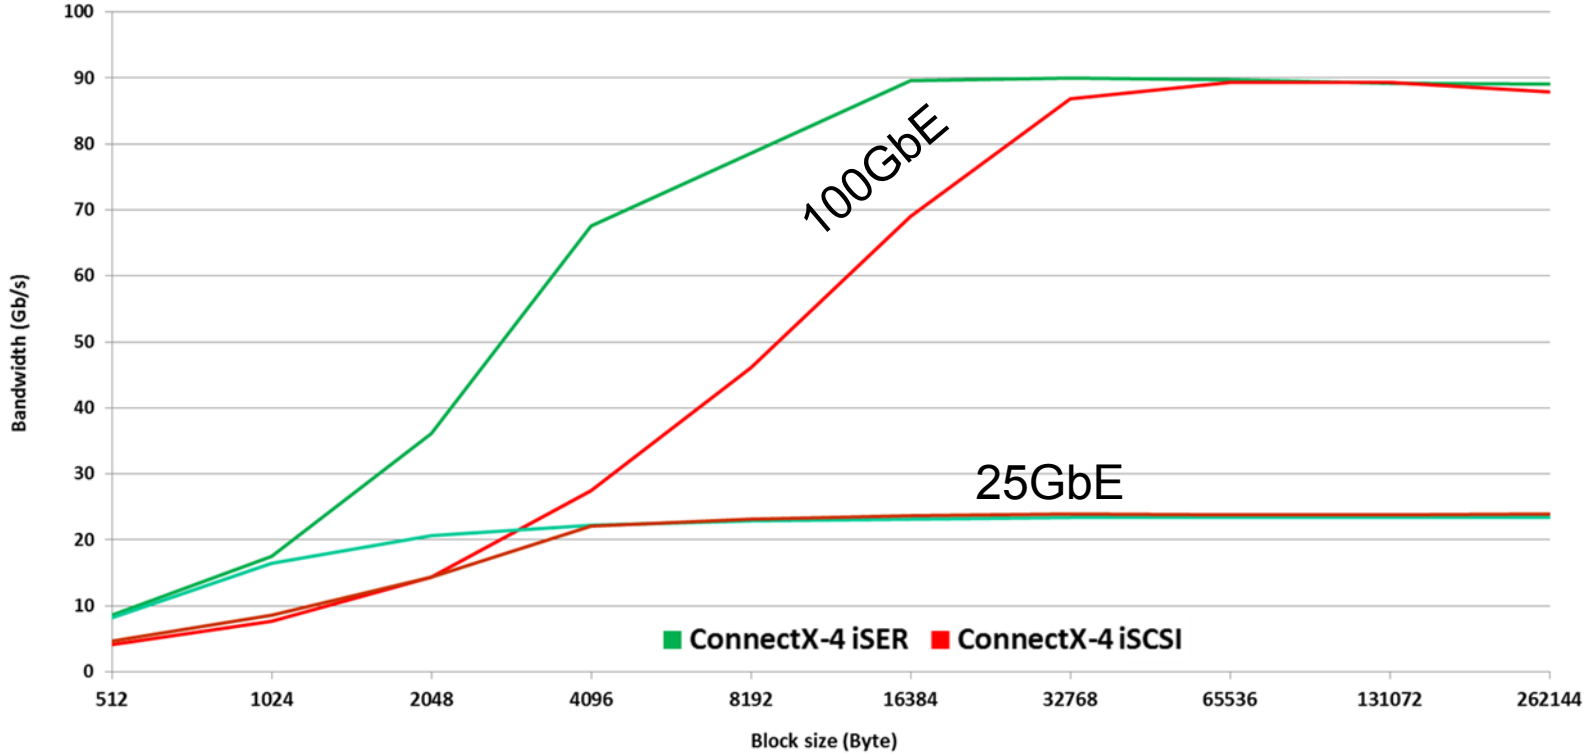

© 2018 Storage Networking Industry Association. All Rights Reserved.

Why Should We Care About Performance?

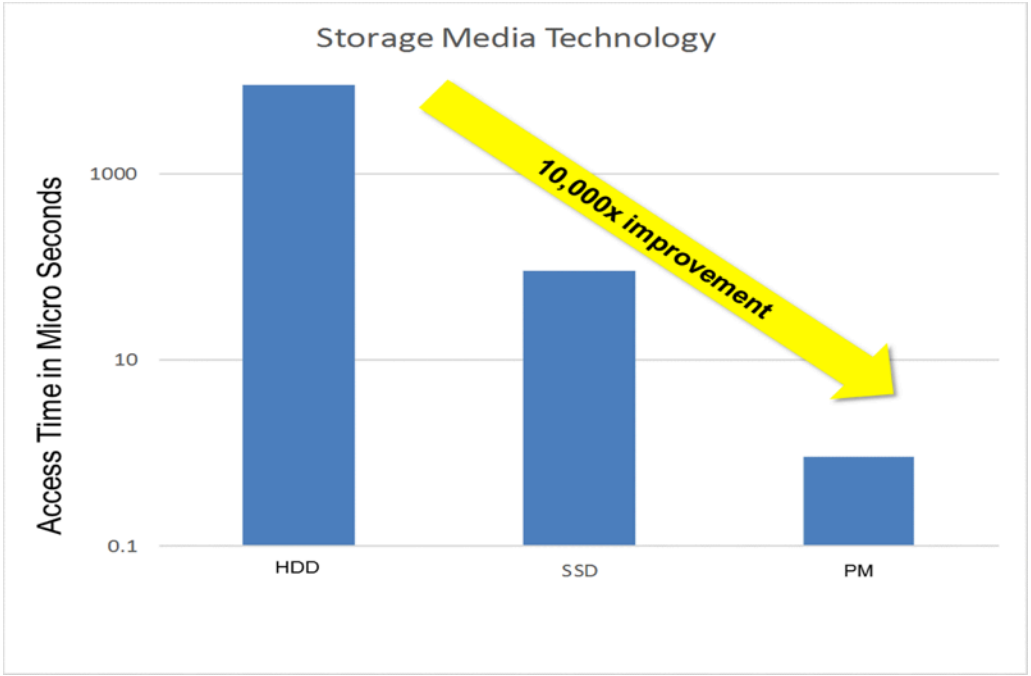

Because Faster Storage Needs a Faster Network!

© 2018 Storage Networking Industry Association. All Rights Reserved.

Faster Storage Needs a Faster Network

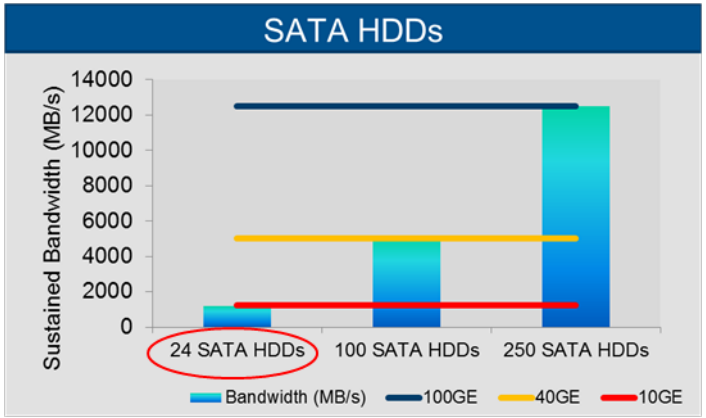

© 2018 Storage Networking Industry Association. All Rights Reserved.

Faster Storage Needs a Faster Network

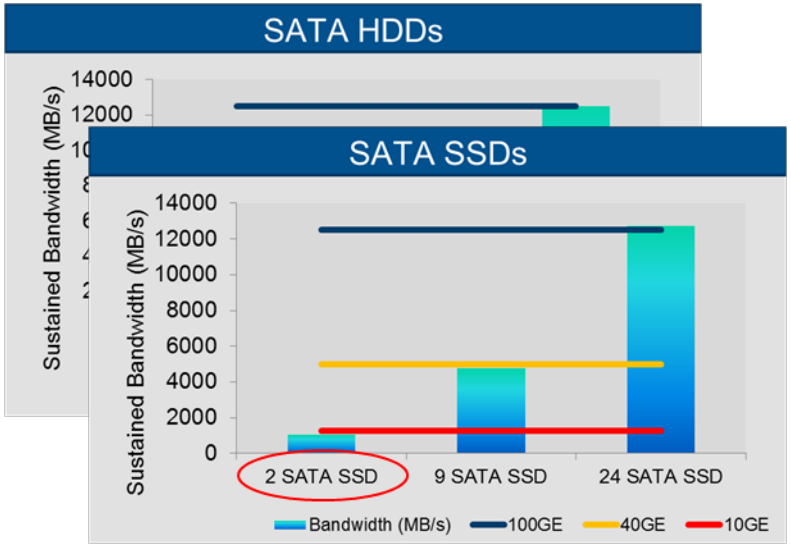

© 2018 Storage Networking Industry Association. All Rights Reserved.

Faster Storage Needs a Faster Network

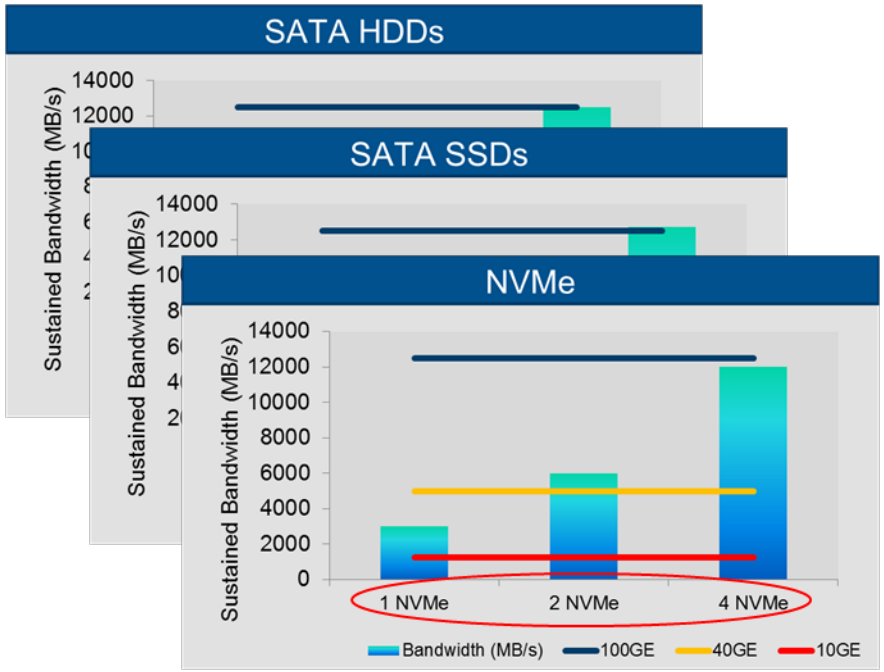

© 2018 Storage Networking Industry Association. All Rights Reserved.

Faster Storage Needs a Faster Network…and a Faster Protocol

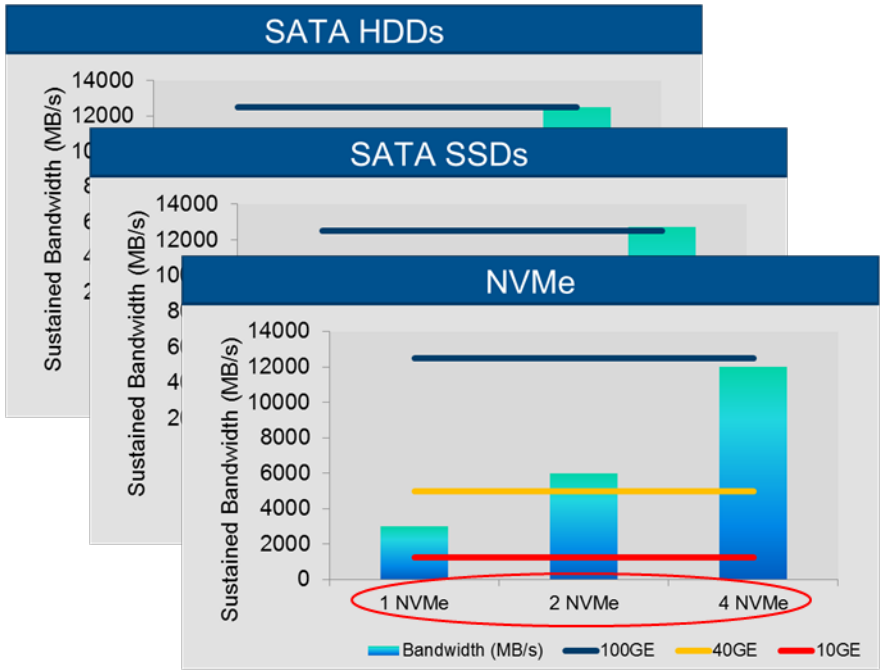

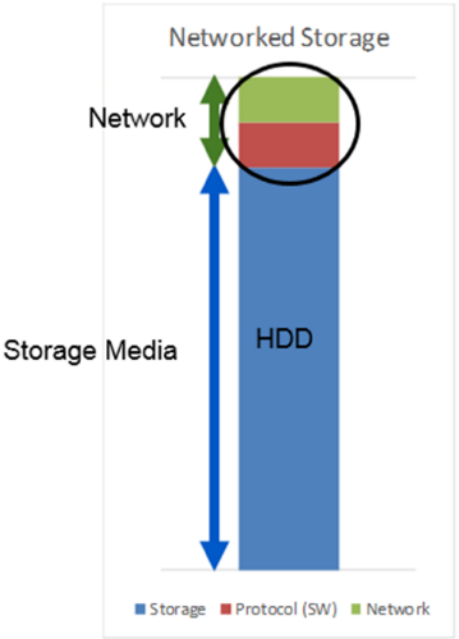

© 2018 Storage Networking Industry Association. All Rights Reserved.

Faster Storage Needs a Faster Network…and a Faster Protocol

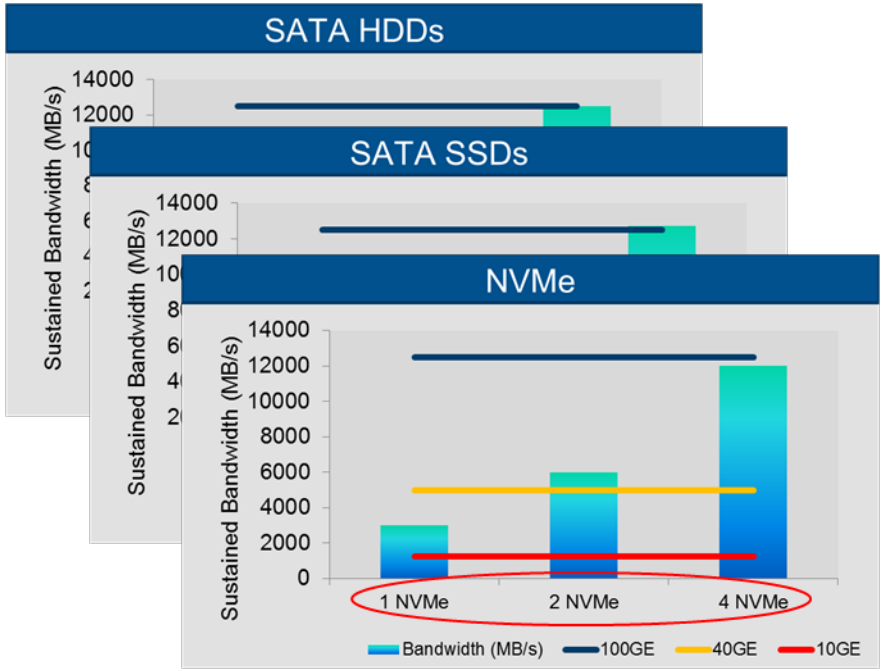

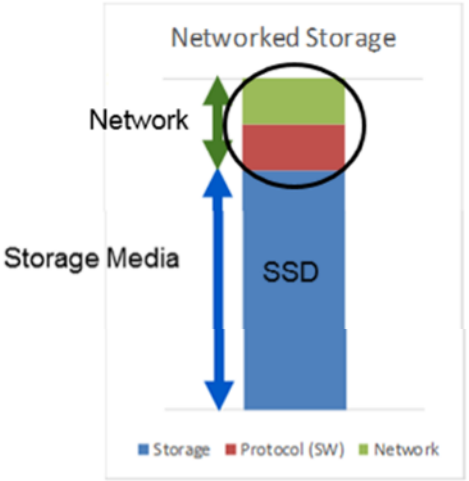

© 2018 Storage Networking Industry Association. All Rights Reserved.

Faster Storage Needs a Faster Network…and a Faster Protocol

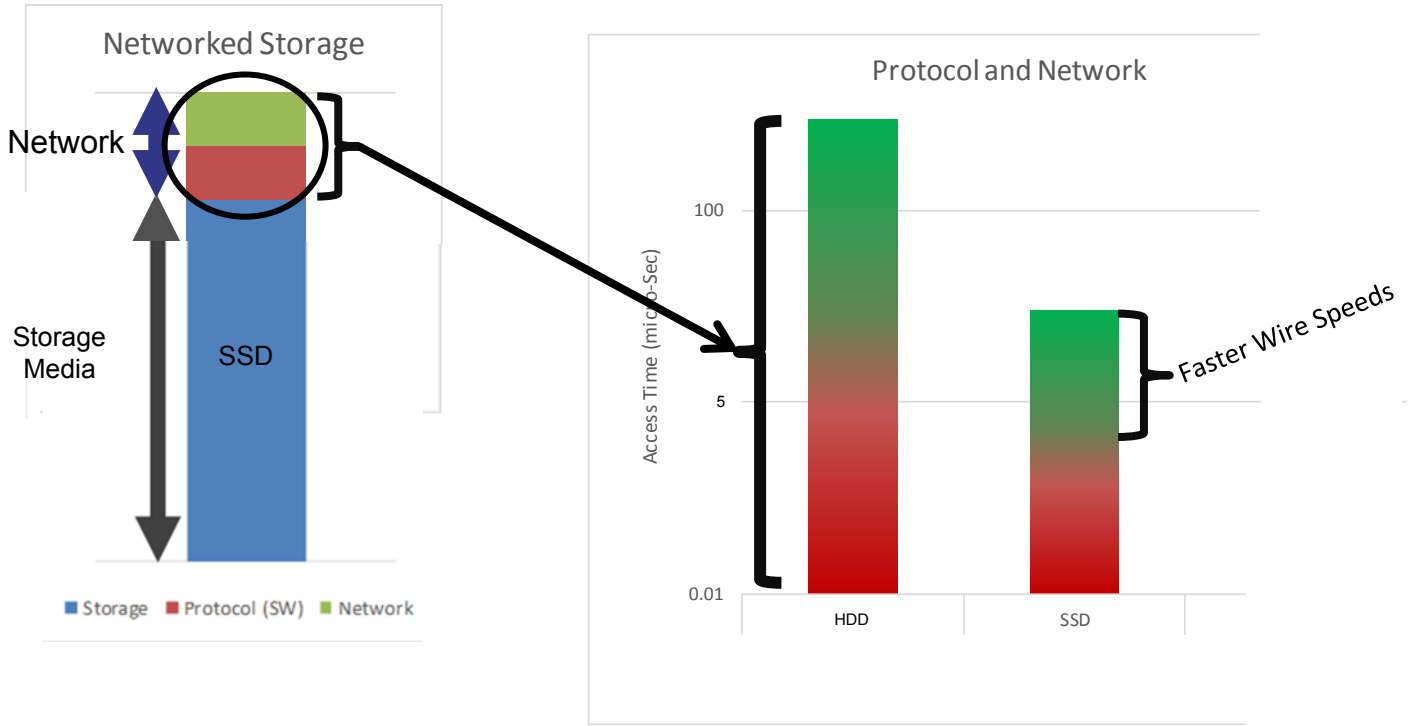

© 2018 Storage Networking Industry Association. All Rights Reserved.

Faster Network and Protocol

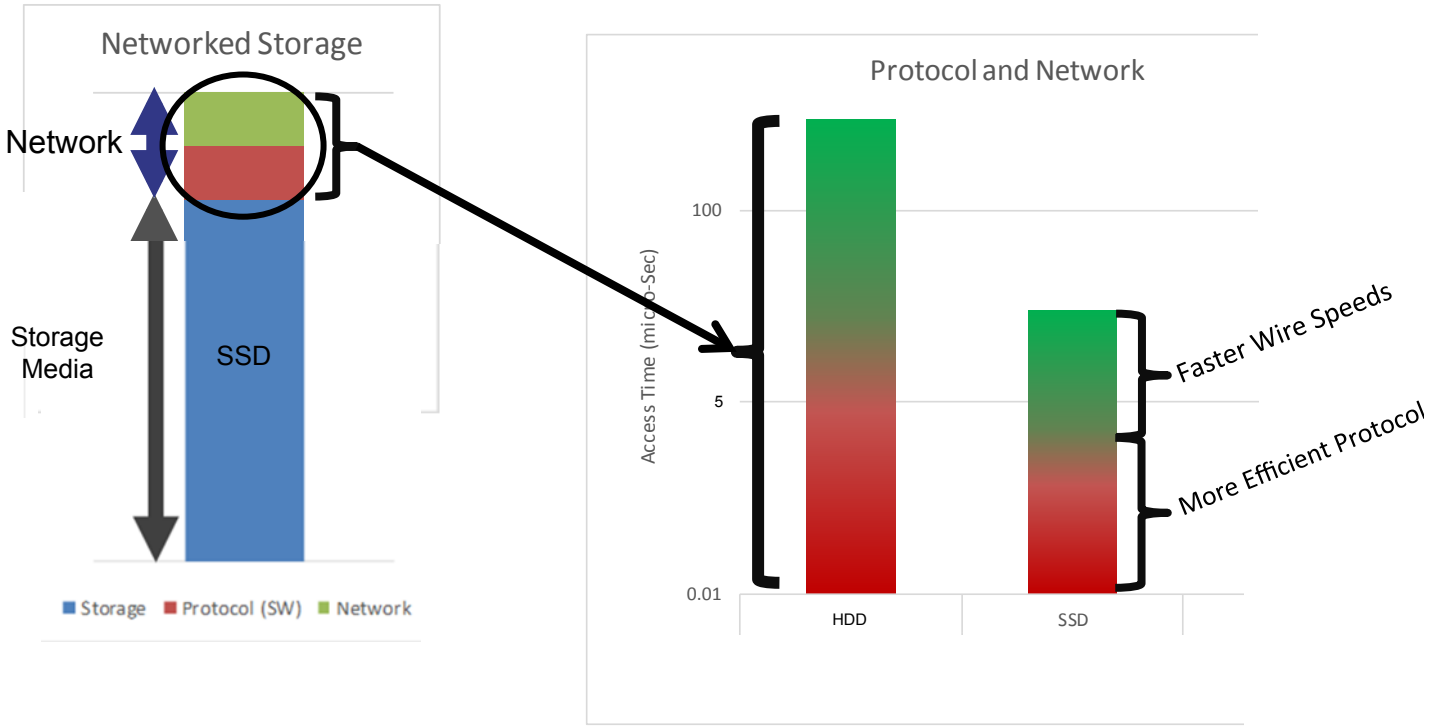

© 2018 Storage Networking Industry Association. All Rights Reserved.

iSER and RDMA Protocols

- Any applications that uses SCSI and iSCSI can use iSER
- iSER uses RDMA to avoid unnecessary data copying on the target and initiator
- For Ethernet the RDMA can be RoCE or iWARP

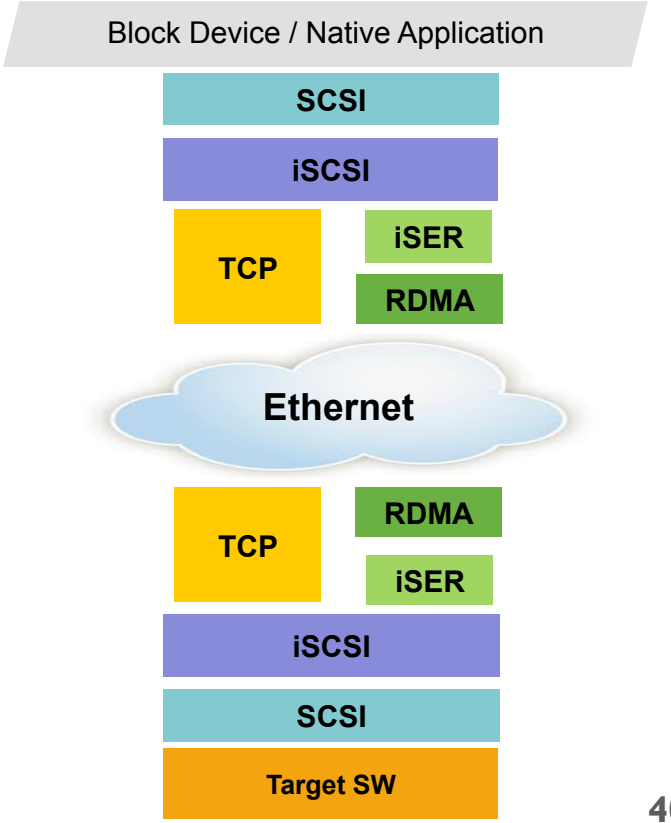

© 2018 Storage Networking Industry Association. All Rights Reserved.

What is RDMA

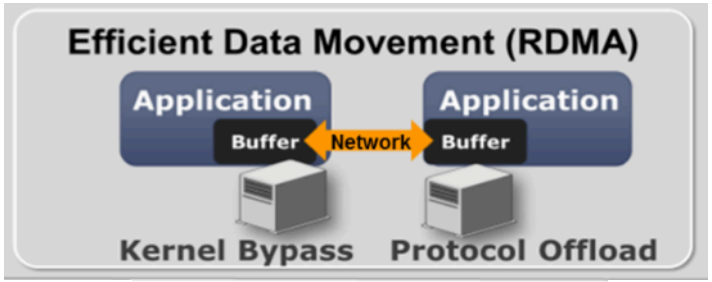

© 2018 Storage Networking Industry Association. All Rights Reserved.

What is RDMA

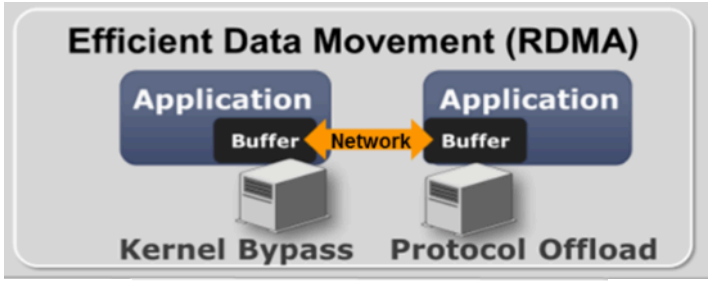

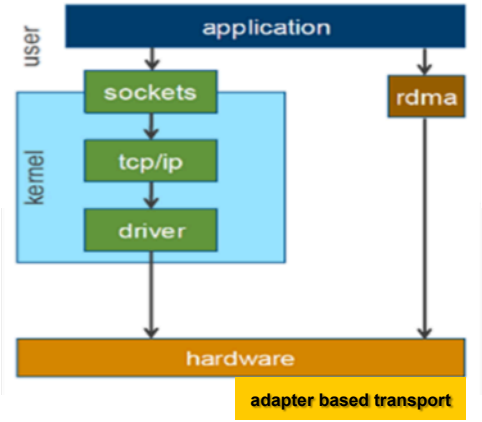

© 2018 Storage Networking Industry Association. All Rights Reserved.

What is RDMA

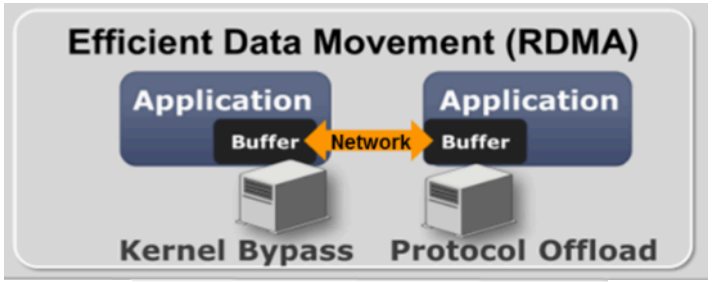

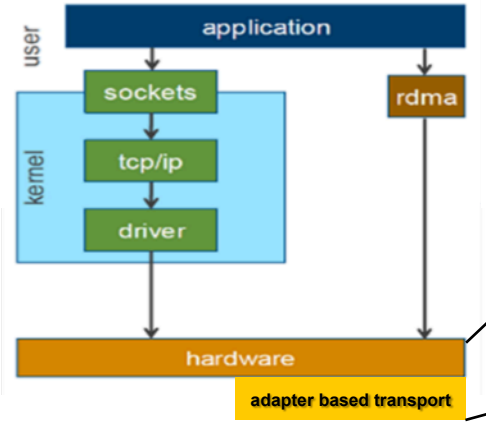

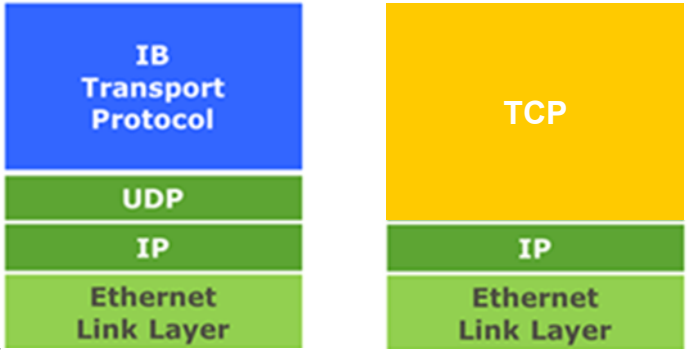

© 2018 Storage Networking Industry Association. All Rights Reserved.

Compare and Contrast…

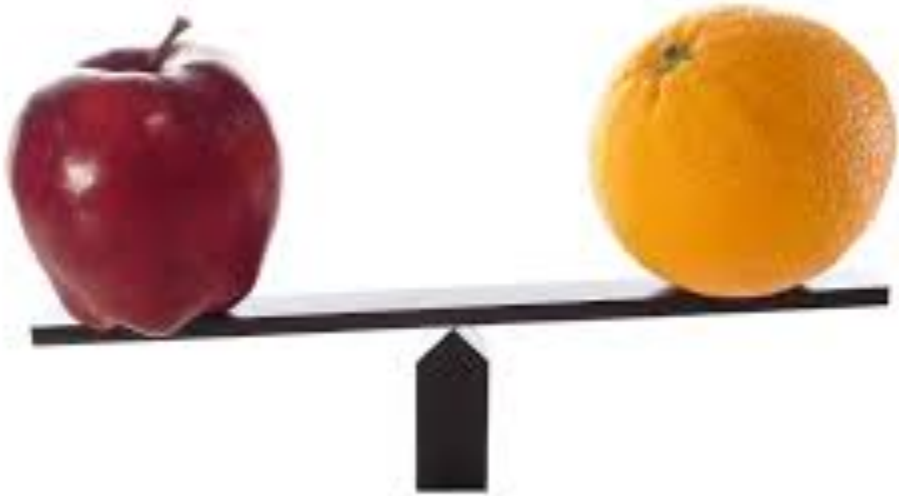

© 2018 Storage Networking Industry Association. All Rights Reserved.

Compare and Contrast

- FCoE

+ Maps FC frames over Ethernet
+ Enables FC on lossless Ethernet
+ Less infrastructure expense (cables, adapters and switches)
- FC expertise required

© 2018 Storage Networking Industry Association. All Rights Reserved.

Compare and Contrast

iSER

+ Provides mapping of the iSCSI protocol on to iWARP/RoCE protocol suites
+ Provides lower latencies and higher bandwidth with lower CPU utilization
+ Leverages iSCSI management infrastructure
- Requires newer adapters and switches that support RDMA

© 2018 Storage Networking Industry Association. All Rights Reserved.

Compare and Contrast

iSCSI

+ Ubiquitous, native support across major operating systems and hypervisors
+ Builds low-cost SAN across standard Ethernet & TCP/IP
+ Offload/TOE availability
- Overhead (packetization, out-of-order delivery, etc.)

© 2018 Storage Networking Industry Association. All Rights Reserved.

How do you decide?

- How do you decide?
  - Do you want an isolated dedicated storage network?
  - How big / complex is your environment?
  - What is your inhouse expertise?
  - Future scale?
  - What are the applications/use cases?

© 2018 Storage Networking Industry Association. All Rights Reserved.

More Webcasts

- Other Great Storage Debates
  - Fibre Channel vs. iSCSI: https://www.brighttalk.com/webcast/663/297837
  - File vs. Block vs. Object Storage: https://www.brighttalk.com/webcast/663/308609
- On-Demand “Everything You Wanted To Know About Storage But Were Too Proud To Ask” Series
  - https://www.snia.org/forums/esf/knowledge/webcasts-topics
- SNIA resources on iSCSI
  - Evolution of iSCSI: https://www.brighttalk.com/webcast/663/197361
  - Comparing iSCSI and NVMe-oF blog: http://sniaesfblog.org/?p=647

© 2018 Storage Networking Industry Association. All Rights Reserved.

After This Webcast

SNIA ESF | ETHERNET STORAGE

- Please rate this webcast and provide us with feedback
- This webcast and a PDF of the slides will be posted to the SNIA Ethernet Storage Forum (ESF) website and available on-demand at www.snia.org/forums/esf/knowledge/webcasts
- A full Q&A from this webcast, including answers to questions we couldn't get to today, will be posted to the SNIA-ESF blog: sniaesfblog.org
- Follow us on Twitter @SNIAESF

© 2018 Storage Networking Industry Association. All Rights Reserved.

Thank You

© 2018 Storage Networking Industry Association. All Rights Reserved.