

Fibre Channel vs. iSCSI

January 31, 2018


Today's Presenters



Fred Knight
NetApp



John Kim
SNIA ESF Chair
Mellanox



Alex McDonald
SNIA ESF Vice Chair
NetApp

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



- ◆ 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.

- A Brief History
 - ◆ Fibre Channel
 - ◆ iSCSI
- Similarities and Differences
- Is one better than the other? Challenge Topics
 - ◆ Performance, manageability, security, cost, etc.
 - ◆ The future

What is Fibre Channel?

➤ Block storage network technology

- ◆ Upper level protocols (SCSI, FICON, NVMe, etc.)
- ◆ Protocol mapping (FCP, FC-SB, FC-NVMe, etc.)
- ◆ Lower/signaling protocols and Physical network

➤ Standardized by ANSI/INCITS T11

- ◆ First standardized 1994
- ◆ First available 1997

➤ “Fibre Channel” usually means SCSI on FCP over FC



Fibre Channel Network Model

- 5-Layer Model
- 2x wire speed every ~4 years (4/8/16/32 Gb/s, etc.)


ULPs	SCSI	FICON	NVMe
Mapping FC-4	FCP-4	FC-SB	FC-NVMe
Services FC-3		FC	
Signaling FC-2		FC	
Transmission FC-1		FC	
Physical FC-0		FC	

➤ FC: Purpose-Built for Storage

- ◆ Optimized out of the box
 - › Fire it up and Forget it

➤ Uses Different Adapters & Switches

- ◆ Dedicated Network with one job to do – storage
- ◆ Familiar configurations and topologies
 - › Point to Point or Switched Fabric



➤ FC: Purpose-built for Storage

- ◆ Centrally Managed (in the network/switch)
- ◆ Lossless in-order delivery
- ◆ Reusable protocol
 - SCSI (traditional LUNs)
 - FICON (Mainframe block storage)
 - NVMe (Namespaces)
 - Even Aircraft (FC-AE/FC-AV)




➤ Reliability

- ◆ Strong CRC
- ◆ Fewer retransmissions

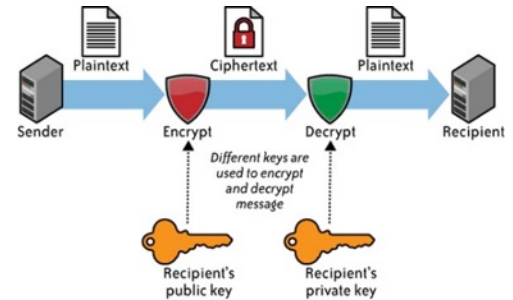
➤ Redundancy

- ◆ Built using duplicate SANs
- ◆ Simplifies network protocol
- ◆ Complexity in Host S/W and Storage device
 - Not in the network



➤ FC Services

- ◆ Zoning
 - Manual zoning
 - Peer zoning (automatic)
- ◆ Name Server
 - Registration / Discovery
 - Notifications of network state changes
- ◆ Security
 - FC-SP-2 authentication
 - FC-FS-4 encryption



➤ Speeds and Feeds

- ◆ 1G / 2G / 4G / 8G / 16G / 32G / 64GFC / 128GFC (x4)

➤ Futures

- ◆ In Development: 256GFC (x4) / 128GFC (x1)
- ◆ Roadmap: 256GFC (TBD), 512GFC (TBD) / 1TFC (TBD)
- ◆ FC-NVMe runs NVMe protocol over existing FC networks

- A Brief History
 - ◆ Fibre Channel
 - ◆ iSCSI
- Similarities and Differences
- Is one better than the other? Challenge Topics
 - ◆ Performance, manageability, security, cost, etc.
 - ◆ The future

What is iSCSI?

- **Block storage network technology**
 - ◆ SCSI protocol running (usually) on TCP/IP
 - ◆ Can run on RDMA: InfiniBand, iWARP, RoCE
 - ◆ Almost always over Ethernet
- **Standardized by IETF**
 - ◆ RFCs 3721, 3722, 4018, 4056, 7143, etc.
 - ◆ First available in 2003/2005/2006 (Windows/Linux/VMware)
- **“iSCSI” usually means SCSI on TCP/IP over Ethernet**

iSCSI Network Model

- OSI 7-layer model
- Was 10x wire speed every ~5 years
 - ◆ 10Mb, 100Mb, 1/10/100GbE
 - ◆ Now 2x every 3-4 years (200/400GbE)
- Ethernet supports NVMe-oF
 - ◆ But not over iSCSI

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

- **iSCSI: Protocol is purpose-built for storage**
 - ◆ Underlying Ethernet network is all-purpose
 - ◆ iSCSI just works out of the box
 - › But discovery requires configuration
 - › Optimization or tuning required for best performance
- **Can have dedicated or shared network**
 - ◆ Shared network for lower cost, maximum flexibility
 - ◆ Dedicated storage network for highest performance



- Management
 - ◆ Mostly distributed (in clients and targets)
- Transport can be configured as lossless or lossy
- Underlying network is reusable
 - ◆ Other storage protocols: NFS, SMB, NVMe-oF
 - ◆ Object storage or scale-out filesystems
 - ◆ Compute traffic or hyperconverged infrastructure

➤ Reliability

- ◆ iSCSI digest (CRC)
- ◆ Ethernet CRC, RoCE/InfiniBand CRC, TCP/IP checksums

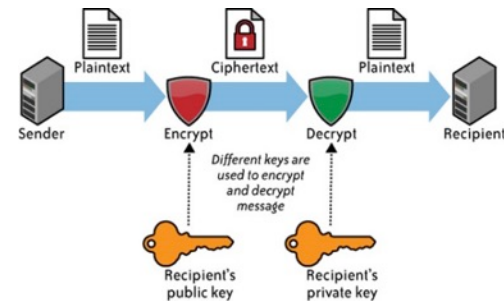
➤ Redundancy

- ◆ Protocol: Link aggregation (LACP) or iSCSI multipathing
- ◆ Physical: Duplicate Ethernet networks (optional)



➤ iSCSI Services

- ◆ Zoning or isolation options
 - Physically separate network
 - ACLs (access control lists), VLANs (virtual LAN), VPN (virtual private network)
- ◆ Internet Storage Name Service (iSNS)
- ◆ Ethernet QoS and monitoring tools
- ◆ Security
 - IPSEC for encryption
 - CHAP or RADIUS for authentication



◆ 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

➤ FCoE

- ◆ SCSI on FCP over lossless Ethernet
- ◆ Often used for blade servers

➤ SRP – SCSI RDMA Protocol

- ◆ Only runs on InfiniBand—usually for HPC customers

➤ NVMe over Fabrics

- ◆ New kid on the (protocol) block
- ◆ NVMe running over RDMA (Ethernet) or Fibre Channel



Important to Remember

- Applications can use either FC or iSCSI
 - ◆ Both deliver SCSI commands to LUNS
 - ◆ Most OSes and Hypervisors support both


- Most arrays support both FC and iSCSI
 - ◆ Both protocols have initiator and target
 - ◆ The optical cabling is often the same

Key Differences

- Underlying Network
- Adapter Offload Options
 - ◆ FC almost completely offloaded
 - ◆ iSCSI can be offloaded or S/W based
- Connection Management Model
- Security & Cost
 - ◆ We'll explore shortly



Fibre Channel vs. iSCSI Challenge Questions



Challenge Question #1 — Performance

➤ Which offers better performance?

- ◆ Bandwidth?
- ◆ Latency? Deterministic latency?



➤ Who is winning the speed race?

- ◆ Does it matter? Who needs that kind of speed anyway?
- ◆ What about distance?
- ◆ How can a server support 100/128/200 Gb/s?
- ◆ In reality, both use the same groups for their “speed” work


Challenge Question #2 — Reliability & Manageability

- Which is More Reliable?
- Which is Easier to Manage?
 - ◆ Setup and Expansion
 - ◆ Monitoring
 - ◆ Troubleshooting



Challenge Question #3 — Cost & Vendors


- What about price?
 - ◆ Acquisition cost
 - ◆ Operational cost
- Is it better to have 2 vendors or 6?
 - ◆ Competition
 - ◆ Innovation
 - ◆ Interoperability



Challenge Question #4 — Security

➤ Which is More Secure?

- ◆ Zoning vs. ACLs and VLANs
- ◆ How does iSCSI do LUN masking?
- ◆ How about encryption on the wire?



➤ Should storage always use a separate network?

- ◆ Does a network “air gap” guarantee storage security?

Challenge Question #5 — Use Cases


➤ What are the Best Applications for Each?

- ◆ Where is Fibre Channel used most?
- ◆ Where is iSCSI used most?



➤ Industry/Vertical Preferences

- ◆ Who prefers Fibre Channel?
- ◆ Who prefers iSCSI?



Challenge Question #6 — The Future

➤ What Will Happen with NVMe and NVMe-oF?

- ◆ How will NVMe SSDs affect block storage networking?
- ◆ Will NVMe-oF replace one more than the other?



➤ What about Hyperconverged Infrastructure?

- ◆ Does HCI make both FC and iSCSI SANs obsolete?



Fibre Channel Advantages

- **Fast, Dedicated Hardware**
 - ◆ Always lossless with credit-based flow control
 - ◆ HBA offloads; High-bandwidth switches
- **Mature Ecosystem**
 - ◆ Proven interoperability, Simple setup
 - ◆ Integration with storage and management software
- **Storage Services**
 - ◆ Name services, zoning, monitoring, diagnostics

➤ Ubiquity of Ethernet

- ◆ Lower price alternatives, many vendors, popularity in the Cloud
- ◆ Same network can support other storage/compute traffic

➤ Flexible Feature-vs.-cost Tradeoffs

- ◆ Hardware acceleration: iSCSI offload, TOE, RDMA
- ◆ Lossless, flow control, congestion management

➤ Higher Bandwidth, Lower Latency

- ◆ 100GbE now, 200GbE soon

Decision time...

- How do you decide?
- Compare apples to apples
 - ◆ Don't compare 1GE iSCSI to 32GFC
- Do you want an isolated dedicated storage network?
- Do you want a converged shared network?
- How big / complex is your environment?
- What is your inhouse expertise?

- Next Live Webcast: Storage Performance Benchmarking: Workloads
 - ◆ February 14th 10:00 am PT
 - ◆ Register at: <https://www.brighttalk.com/webcast/663/297859>
- 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>
- FCIA webcast: Fibre Channel Performance: Congestion, Slow Drain, and Over-Utilization, Oh My!
 - ◆ February 6th 10:00 am PT
 - ◆ Register at: <https://www.brighttalk.com/webcast/14967/295141>
- 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>

After This Webcast

- 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



Thank You