SNIA ESF | ETHERNET STORAGE

Re-Introduction to Ethernet Networked Storage

August 4, 2016

SNIA Presenters

Rob Davis
VP of Technology
Mellanox

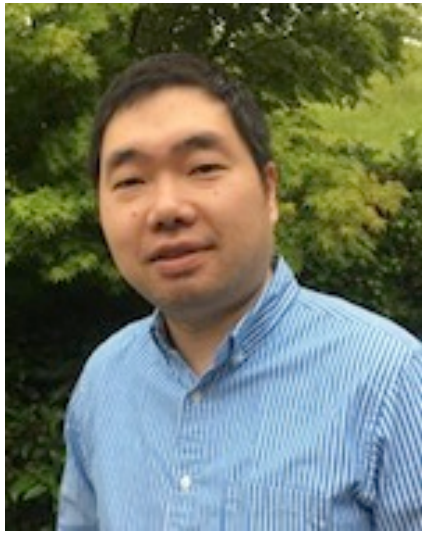

Fred Zhang
Product Marketing Manager
Intel

About SNIA

SNIA ESF | ETHERNET STORAGE

SNIA-At-A-Glance

160
unique member companies

3,500
active contributing members

50,000
IT end users & storage pros worldwide

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

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.

The Evolution of Storage Devices

The Evolution of Storage Devices

IBM Almaden Research Facility

First Disk Drive - RAMAC 350 - 1956

Almaden Lobby

First Disk Drive - RAMAC 350 - 1956

Almaden Lobby

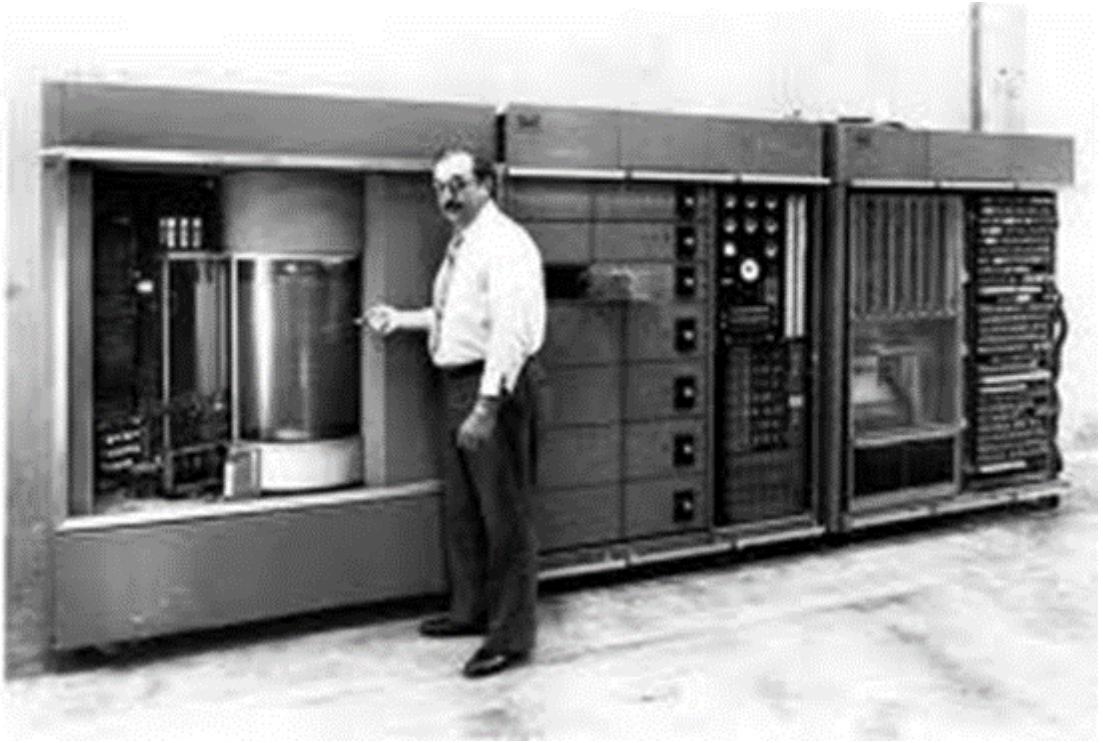

First Disk Drive - RAMAC 350 - 1956

Almaden Lobby

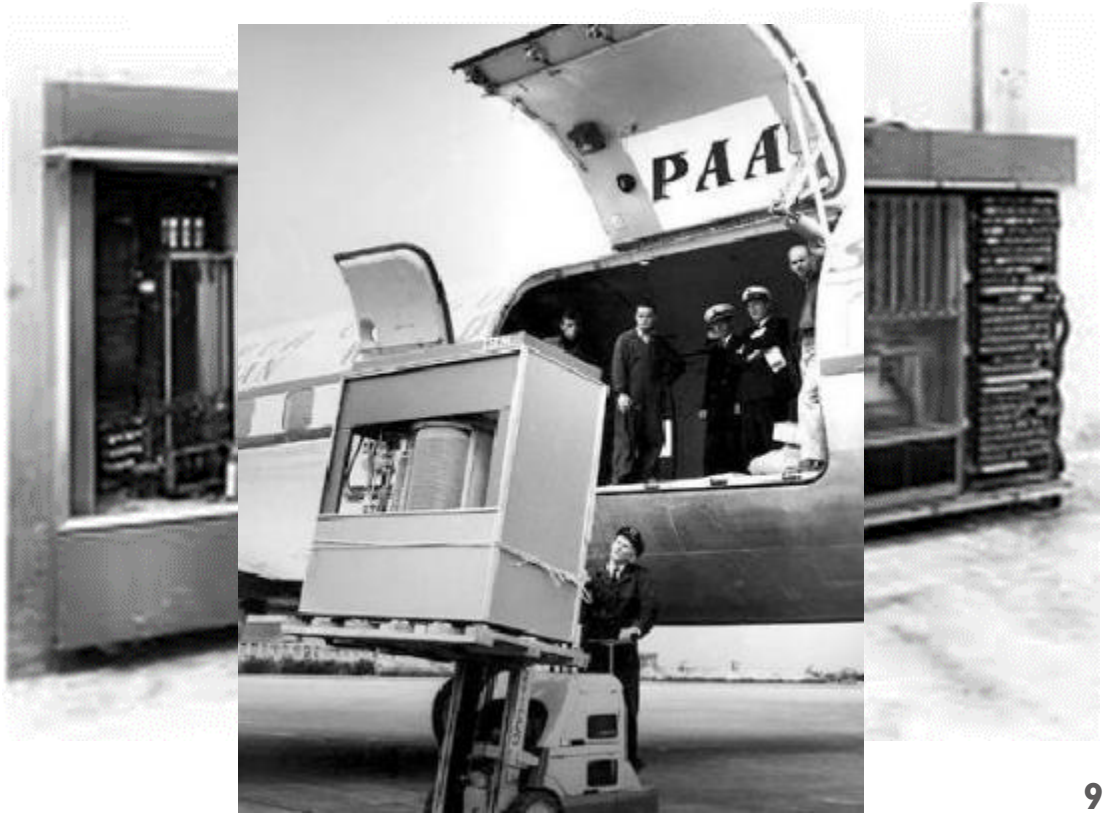

Evolution of Disk Drive

~$81,000/MB

5 MB

Latency ~25MS

60 years

~$60/TB

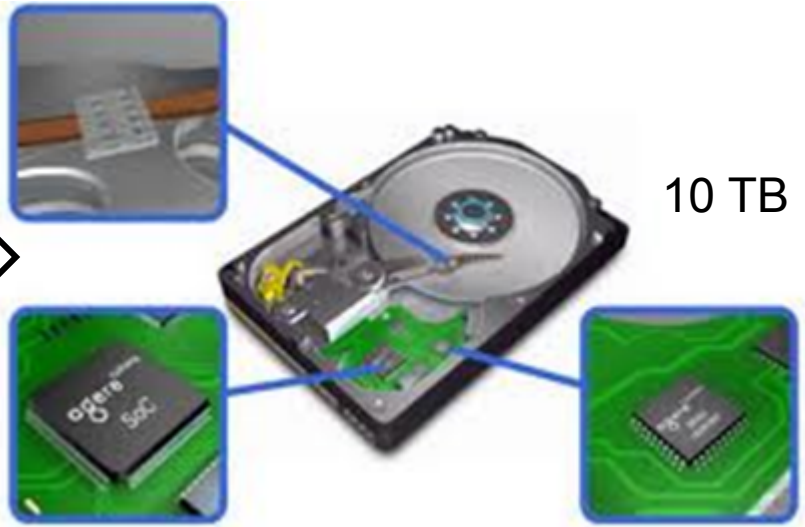

10 TB

Latency ~4MS

First Solid-State Drive(SSD) - Bulk Core - 1976

- 8 memory boards, 256KB of RAM, total 2MB
- Access time 750usec
- Cost $9700 in 1977, equivalent to $40,000 today

Evolution of SSDs

~$20,000/MB today

2MB

40 years

~$300/TB

4TB

RAM

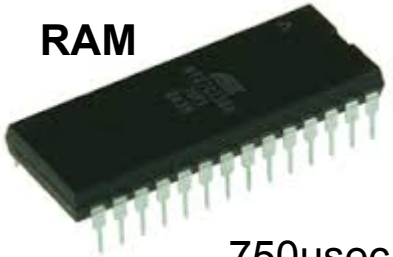

750usec

NAND

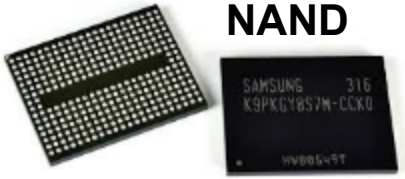

~25usec read, ~250 write

At First SSDs Fill the Performance Gap

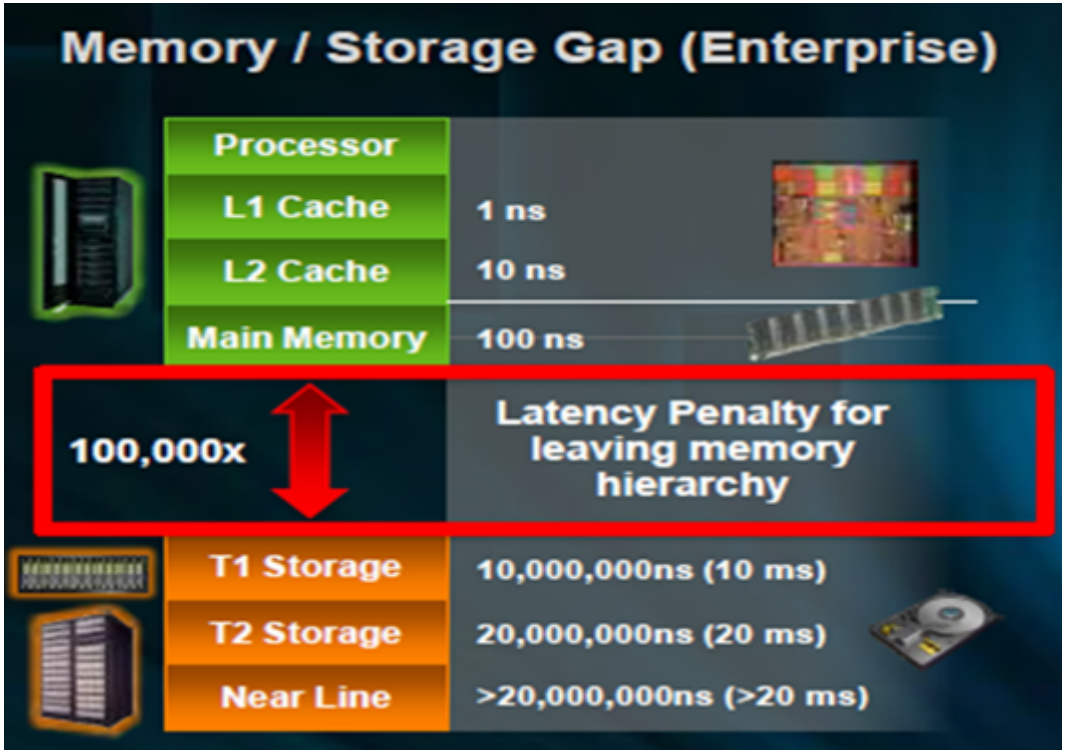

Next Step - SANs

- Sharing storage across multiple servers
  - Better utilization: capacity, rack space, power
  - Scalability, management, fault isolation

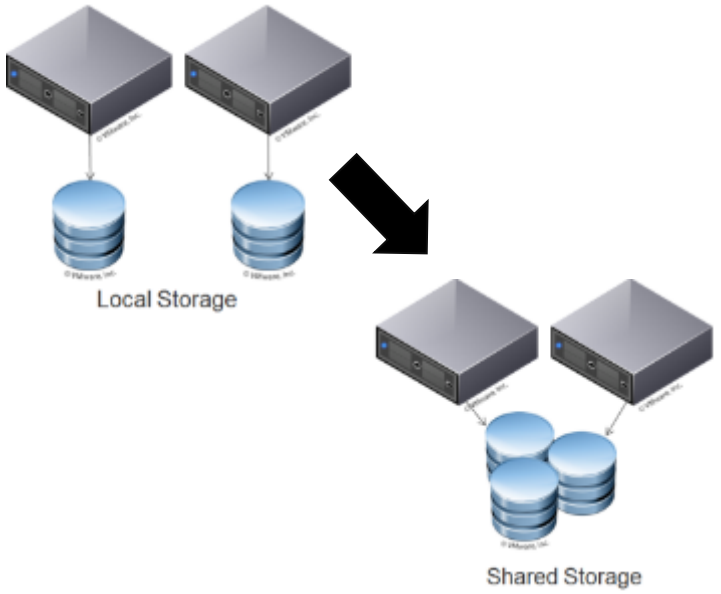

Next Step - SANs

- Sharing storage across multiple servers
  - Better utilization: capacity, rack space, power
  - Scalability, management, fault isolation
- Fibre Channel standard started in 1988, with ANSI standard 1.0 approval in 1994
- First Fibre-channel connected storage product SPARCstorageArray, 1Gb/s

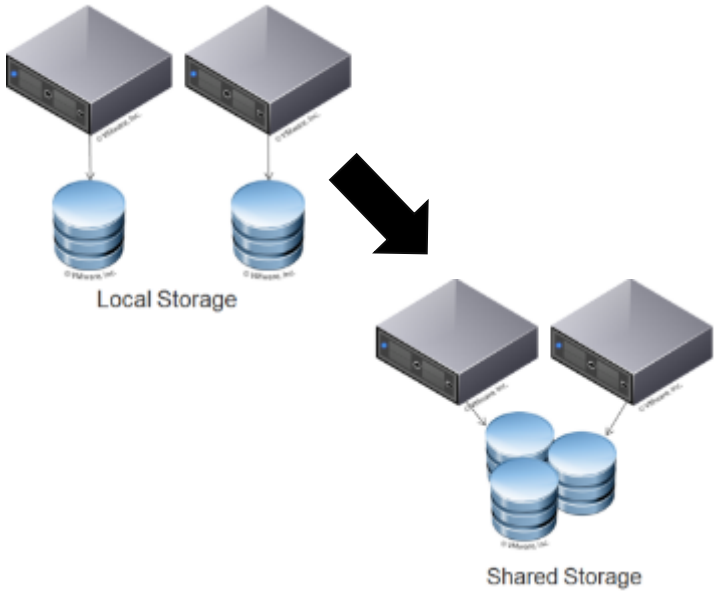

Fibre Channel SANs

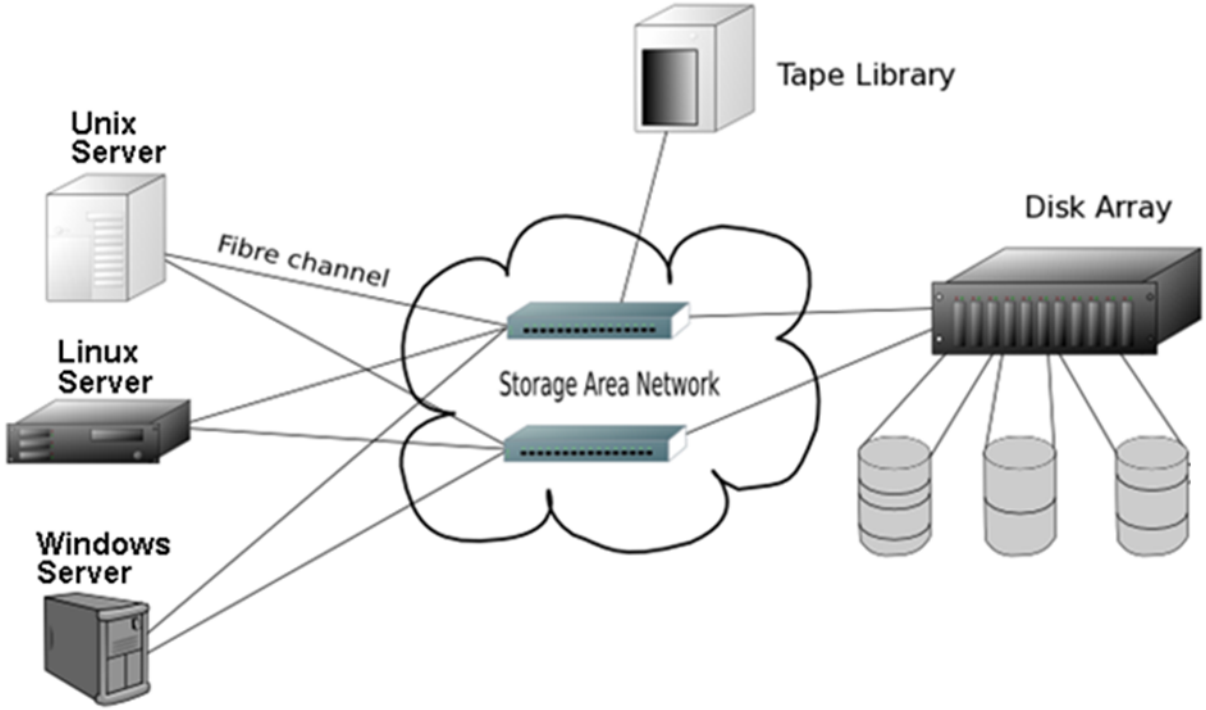

SANs and LANs

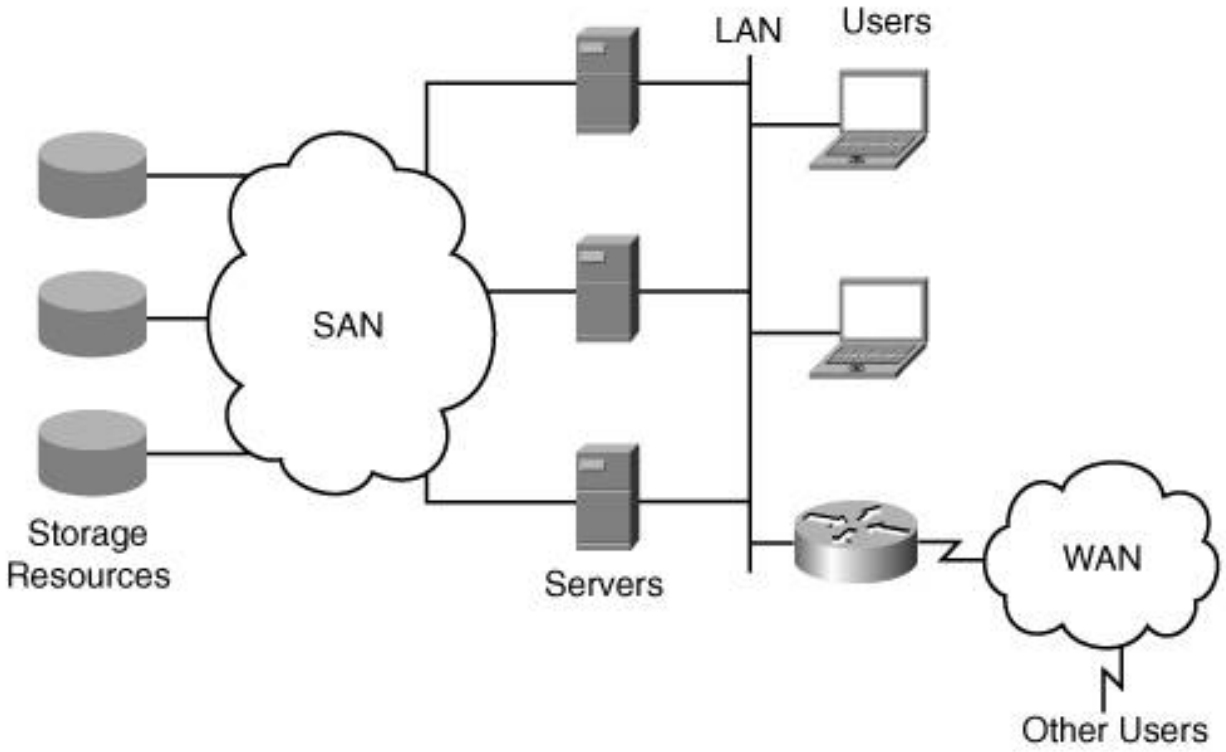

iSCSI SANs

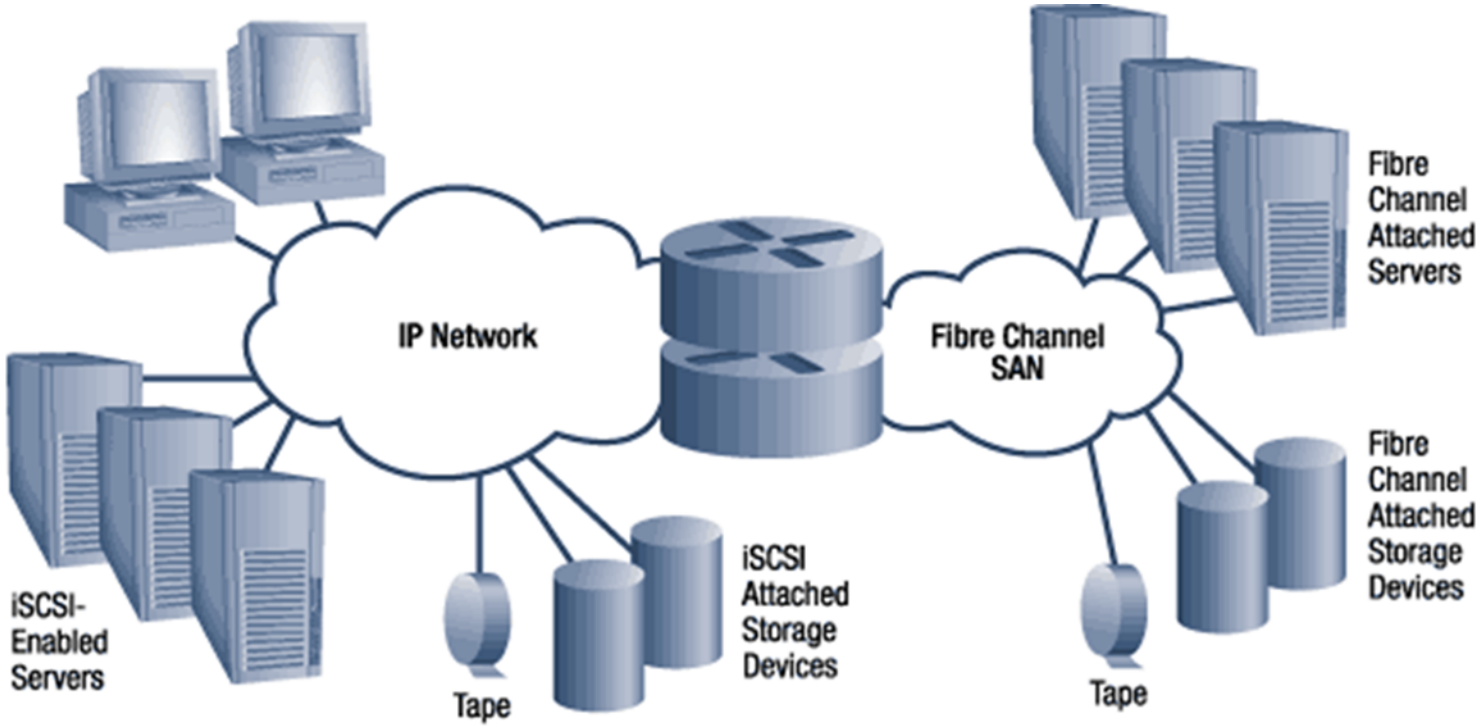

IETF ratified the iSCSI standard 2003

iSCSI SANs

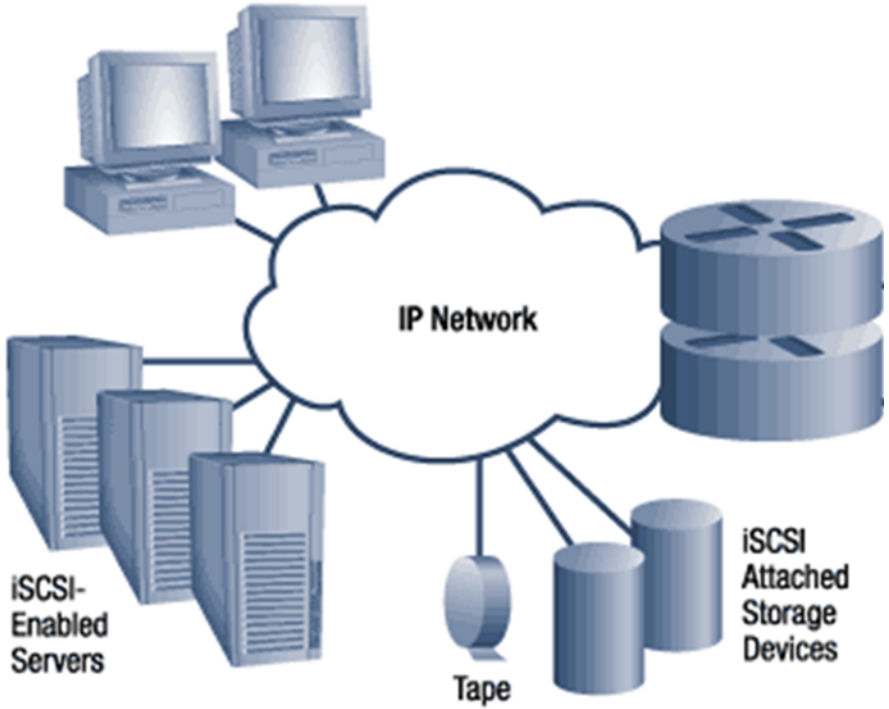
IP Network
iSCSI-
Enabled
Servers
Tape
iSCSI
Attached
Storage
Devices

iSCSI SANs

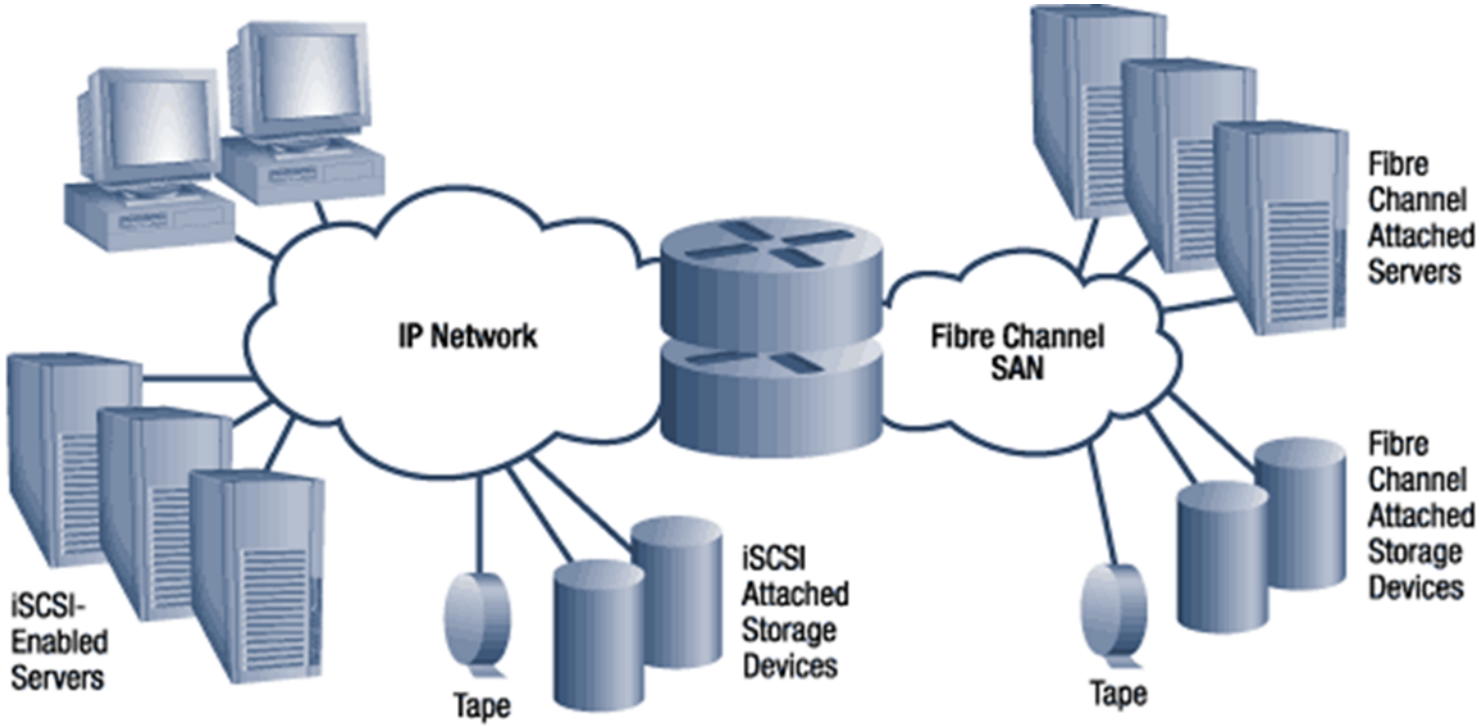
IP Network
Fibre Channel SAN
Fibre Channel Attached Servers
iSCSI-Enabled Servers
iSCSI Attached Storage Devices
Tape
Fibre Channel Attached Storage Devices
Tape

FCoE SANs

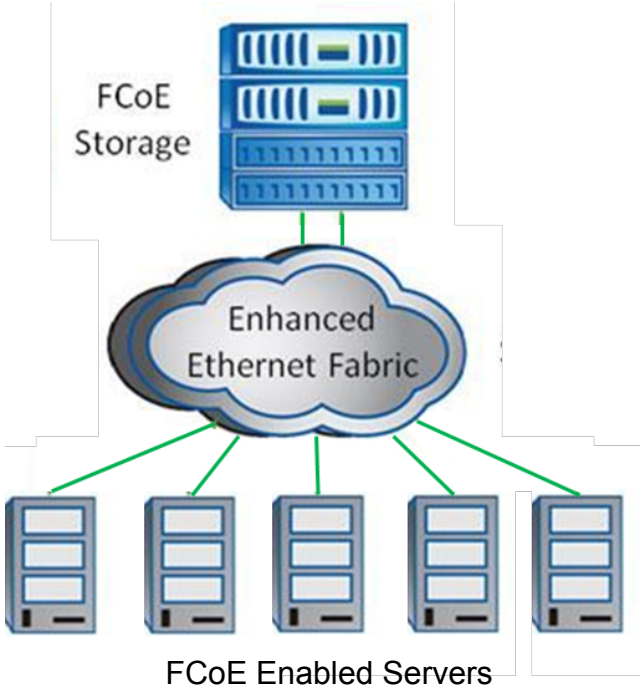

FCoE part of T11 FC-BB-5 was forwarded to ANSI for publication in 2009

FCoE SANs

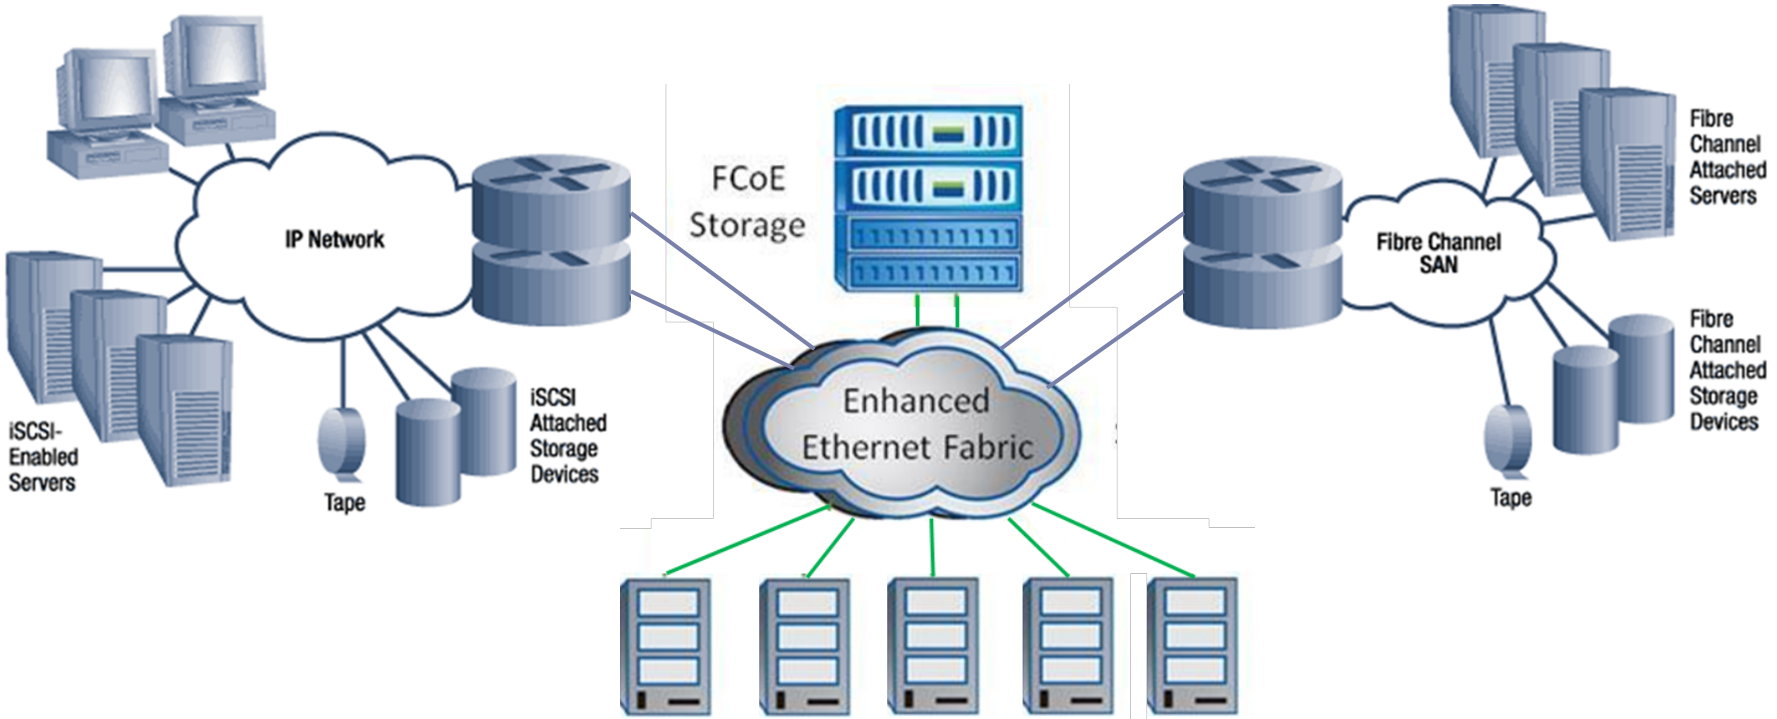
IP Network
iSCSI-Enabled Servers
Tape
iSCSI Attached Storage Devices
FCoE Storage
Enhanced Ethernet Fabric
Fibre Channel SAN
Fibre Channel Attached Servers
Fibre Channel Attached Storage Devices
Tape
FCoE Enabled Servers

FCoE SANs

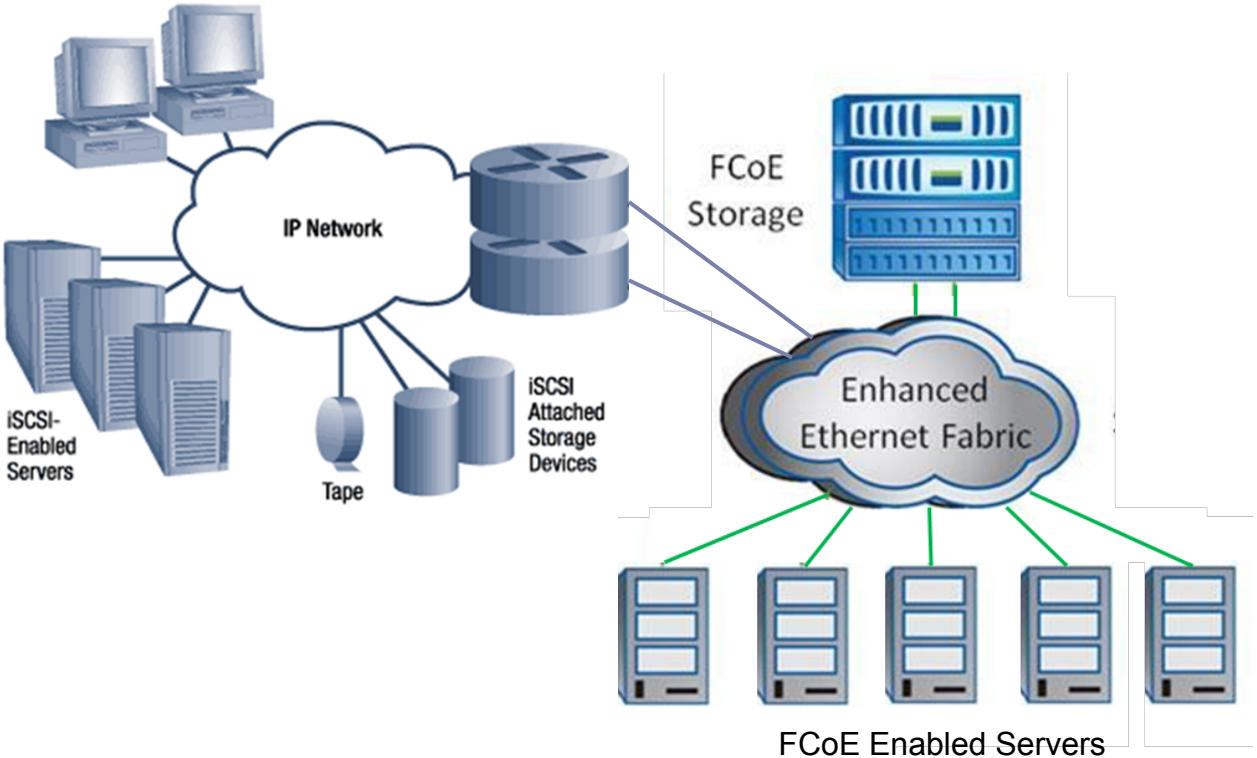
IP Network
iSCSI-
Enabled
Servers
Tape
iSCSI
Attached
Storage
Devices
FCoE
Storage
Enhanced
Ethernet Fabric
FCoE Enabled Servers

FCoE SANs

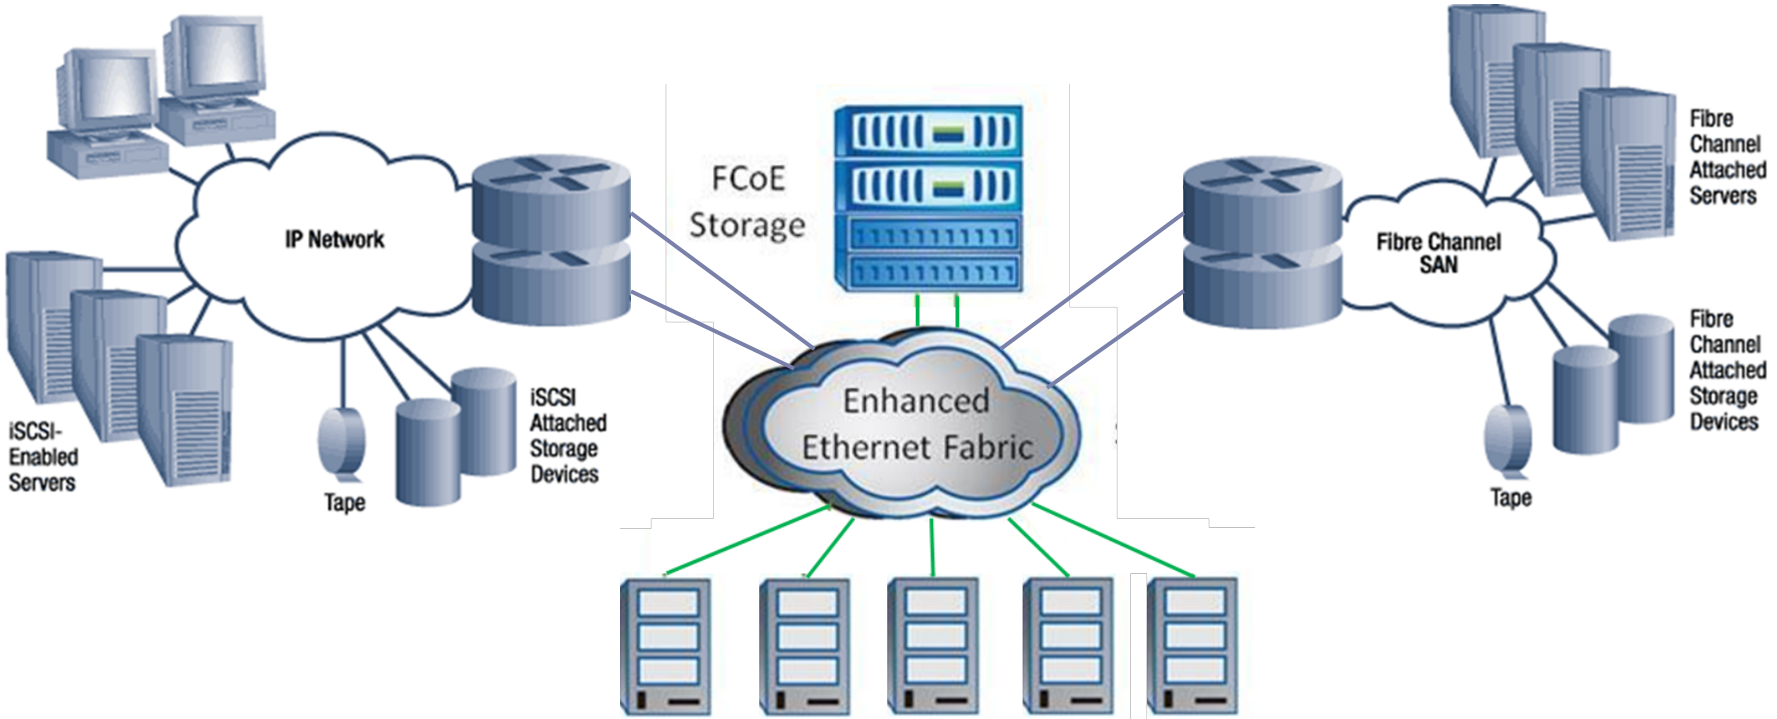
IP Network
iSCSI-Enabled Servers
Tape
iSCSI Attached Storage Devices
FCoE Storage
Enhanced Ethernet Fabric
Fibre Channel SAN
Fibre Channel Attached Servers
Fibre Channel Attached Storage Devices
Tape
FCoE Enabled Servers

Why Ethernet Storage Revival Now?

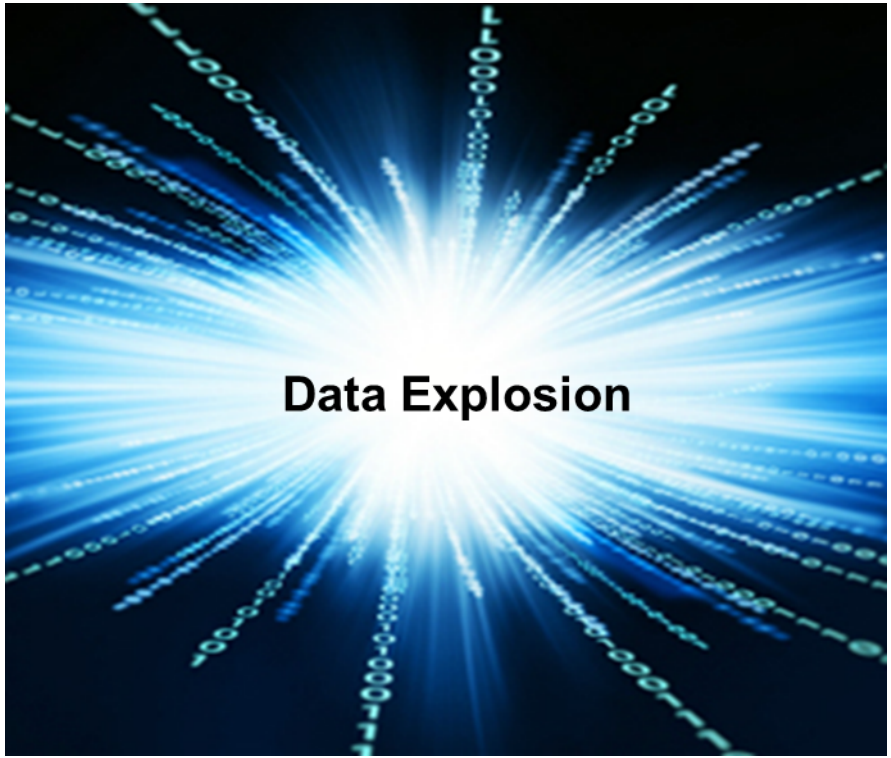

Why Ethernet Storage Revival Now?

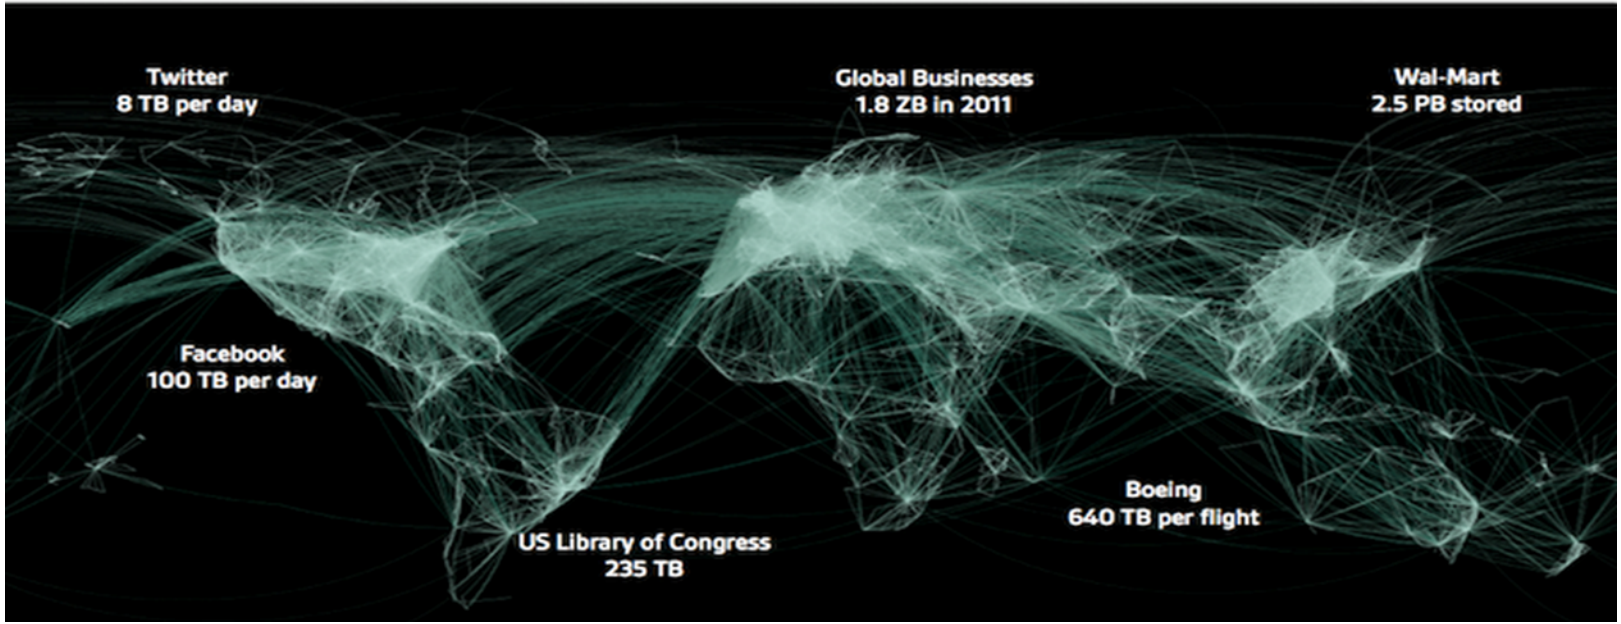

Capacity Needs Driving Lower Cost Networked Storage

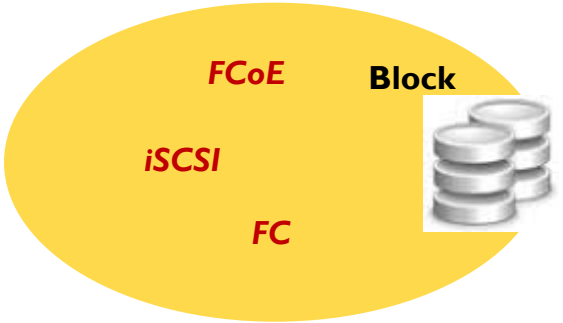

Capacity Needs Driving Lower Cost Networked Storage

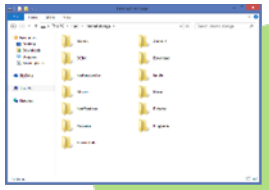

SMB (CIFS)

File

NFS

FCoE

Block

iSCSI

FC

Ceph

Object

Capacity Needs Driving Lower Cost Networked Storage

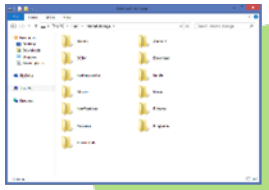

SMB (CIFS)

File

NFS

FCoE

Block

iSCSI

FC

Ceph

Object

Capacity Needs Driving Lower Cost Networked Storage

SNIA ESF | ETHERNET STORAGE

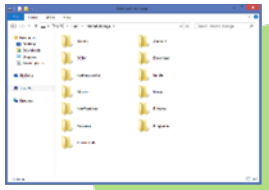

SMB (CIFS)

File

NFS

FCoE

Block

iSCSI

FC

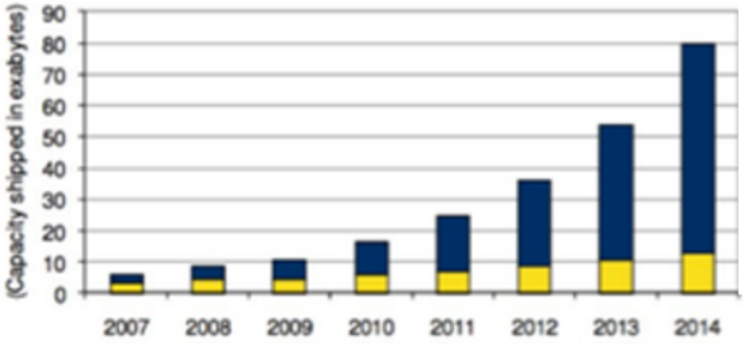

Ceph

Object

Capacity Needs Driving Lower Cost Networked Storage

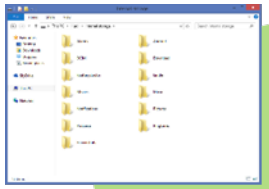

SMB (CIFS)

File

NFS

FCoE

Block

iSCSI

FC

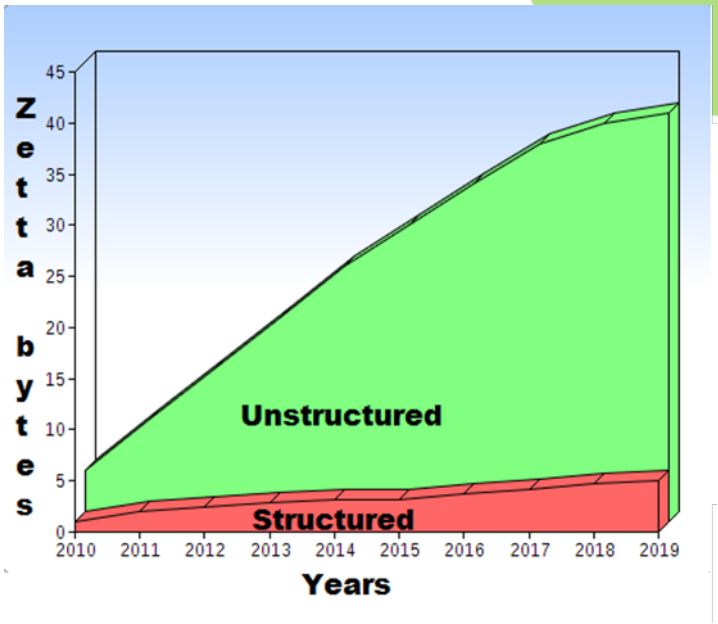

Ceph

Object

Capacity Needs Driving Lower Cost Networked Storage

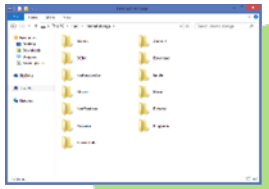

SMB (CIFS)

File

NFS

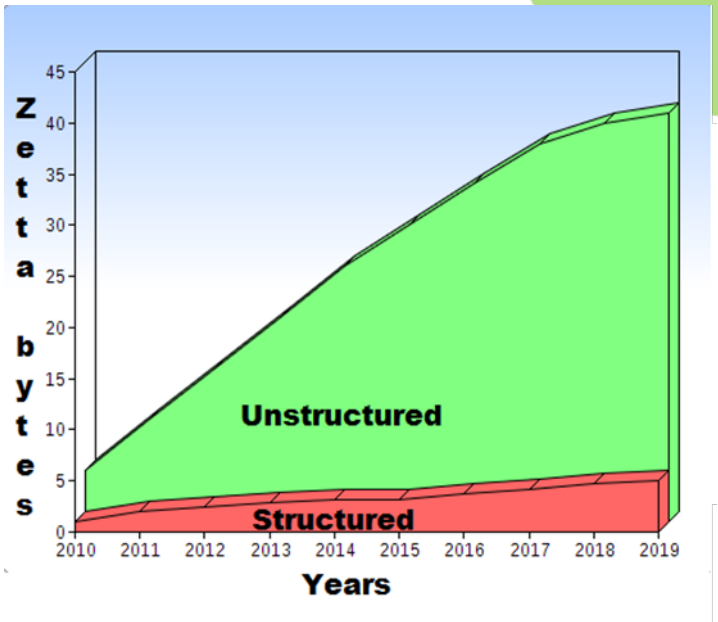

FCoE

Block

iSCSI

FC

Ethernet

Ceph

Object

Why Ethernet Storage Revival Now?

Hard Disk Drive

Solid State Drive(SSD)

Storage Arrays

Scale-Out Storage

- Three major storage technology transitions
  - SSDs replacing hard disk drives
  - Scale-out, software-defined, or “Server-SAN” storage
  - Hyperconverged

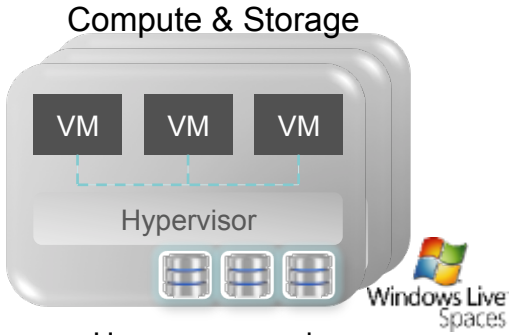

Hyperconverged

Cloud Driving Scale Out Storage into Data Center

Cloud Driving Scale Out Storage into Data Center

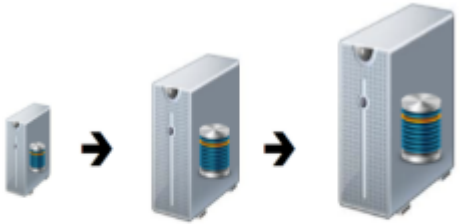

Scale Up

Scale Out & Software-Defined

- Clouds abandoned traditional SAN
  - There is no Fibre Channel in the cloud
  - Different networking and storage models
- Enterprises are migrating to a cloud model
- New storage applications use scale-out

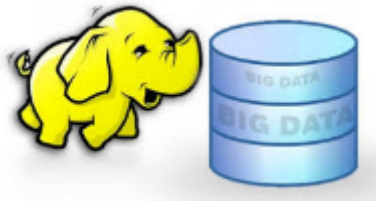

Scale Out Storage Growth

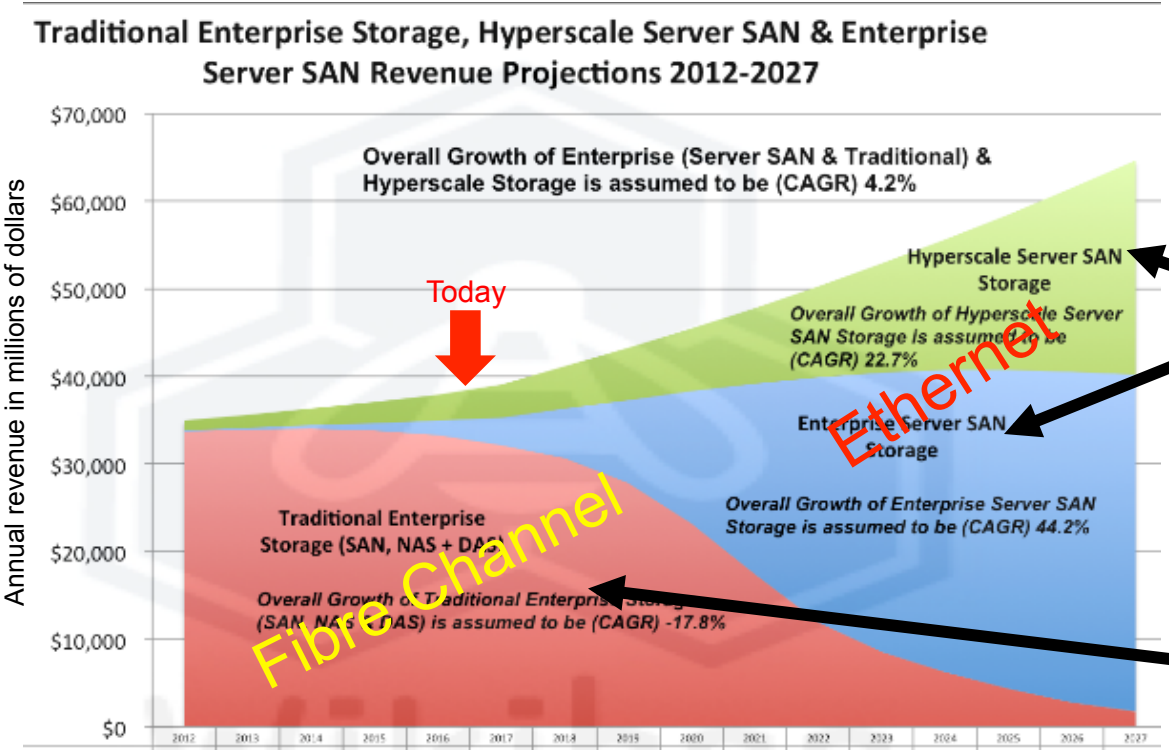

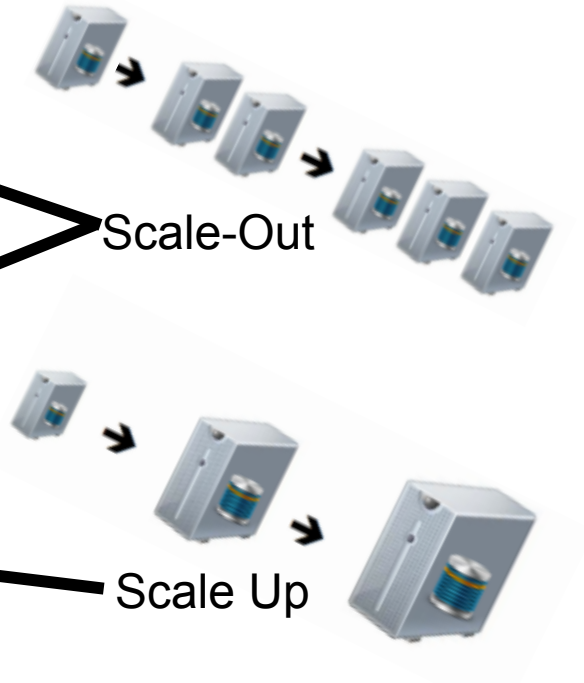

http://wikibon.org/wiki/v/The_Rise_of_Server_SAN#The_Business_.26_Technology_Drivers_behind_Server_SAN

Hyperconverged

- Adaptation of scale-out storage
- Collapse separate compute & storage into one layer
  - Integrated compute-storage nodes
- IO requirements
  - Low latency, high bandwidth, **Ethernet**

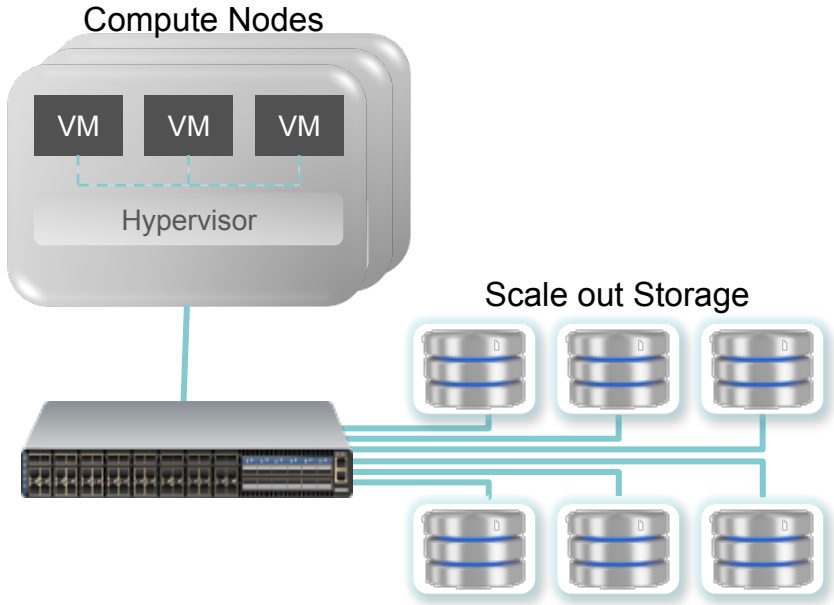

Hyperconverged

- Adaptation of scale-out storage
- Collapse separate compute & storage into one layer
  - Integrated compute-storage nodes
- IO requirements
  - Low latency, high bandwidth, **Ethernet**

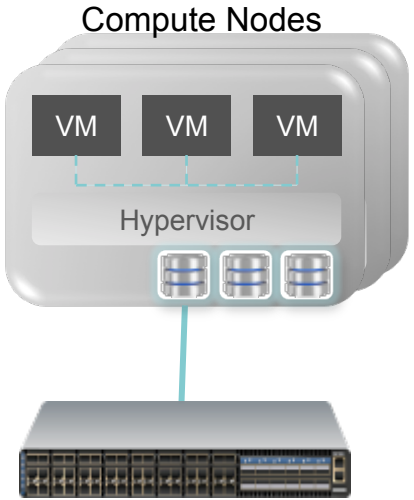

Hyperconverged Market

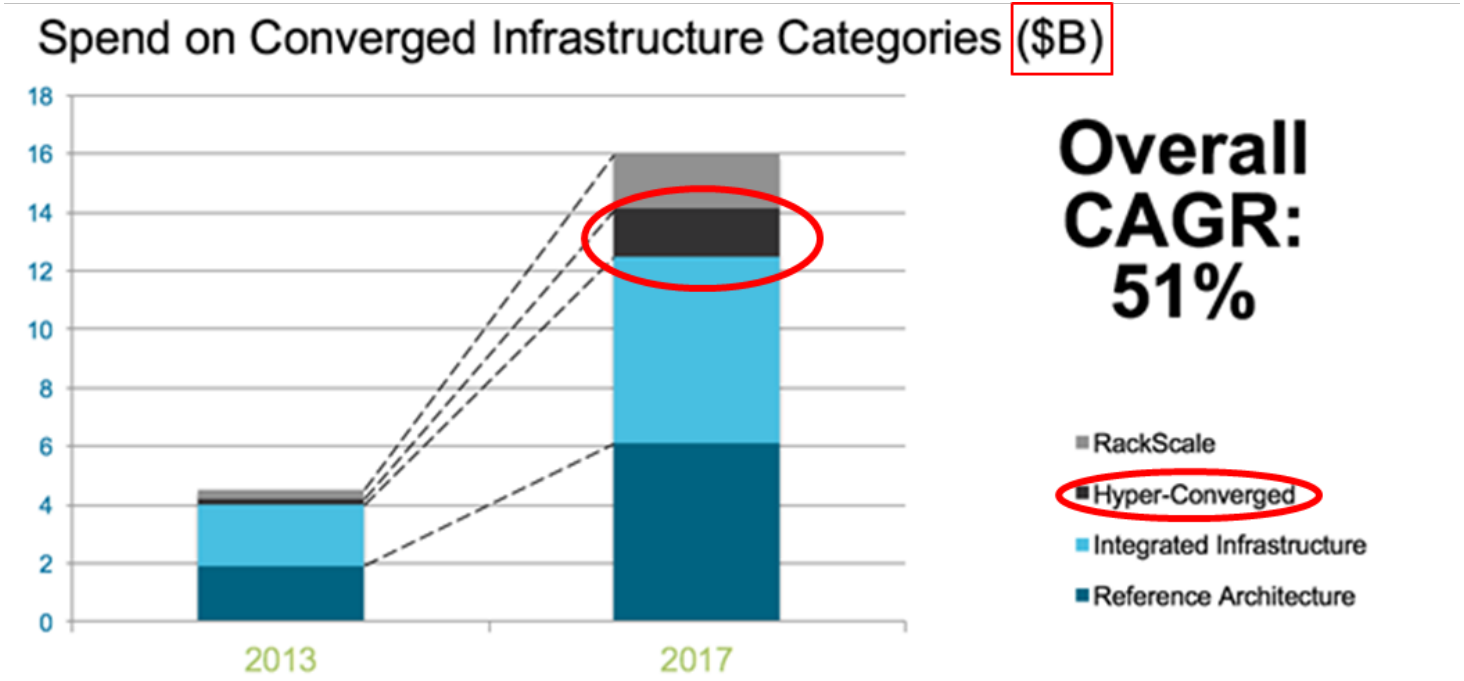

Source: IDC #244924
Source: EMC Market Intelligence. TAM excludes Internal Enterprise Storage, DIY Hardware and Services in excess of EMC share

Why Ethernet Storage Revival Now?

SNIA ESF | ETHERNET STORAGE

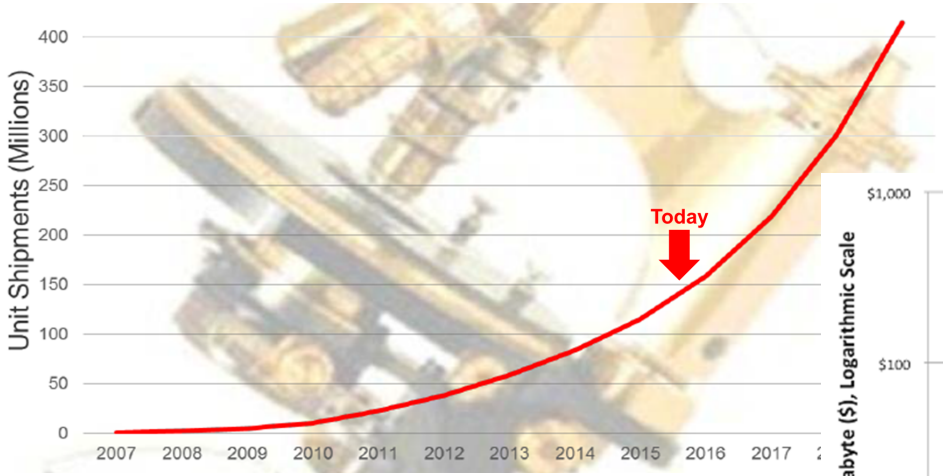

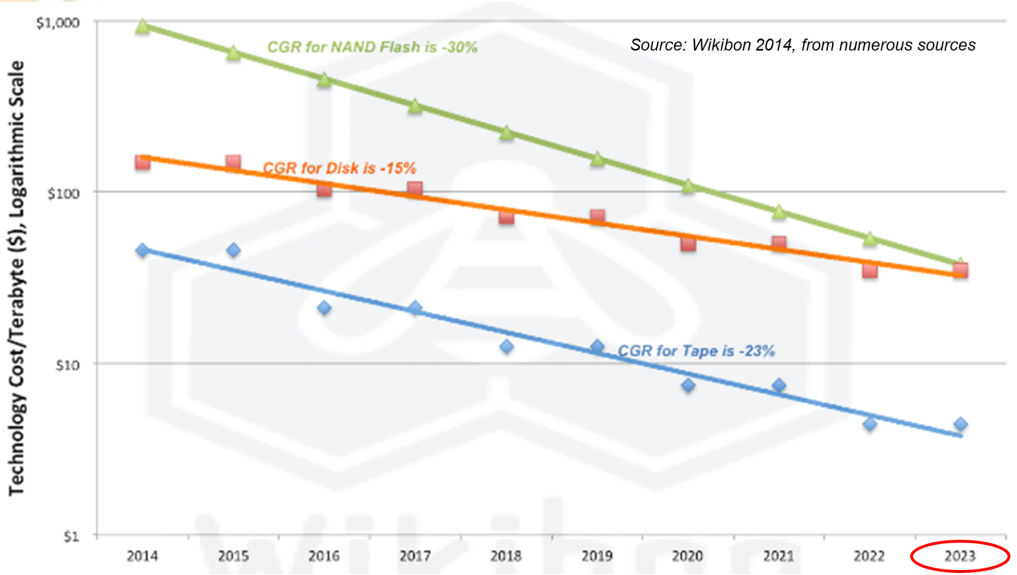

Why Ethernet Storage Revival Now?

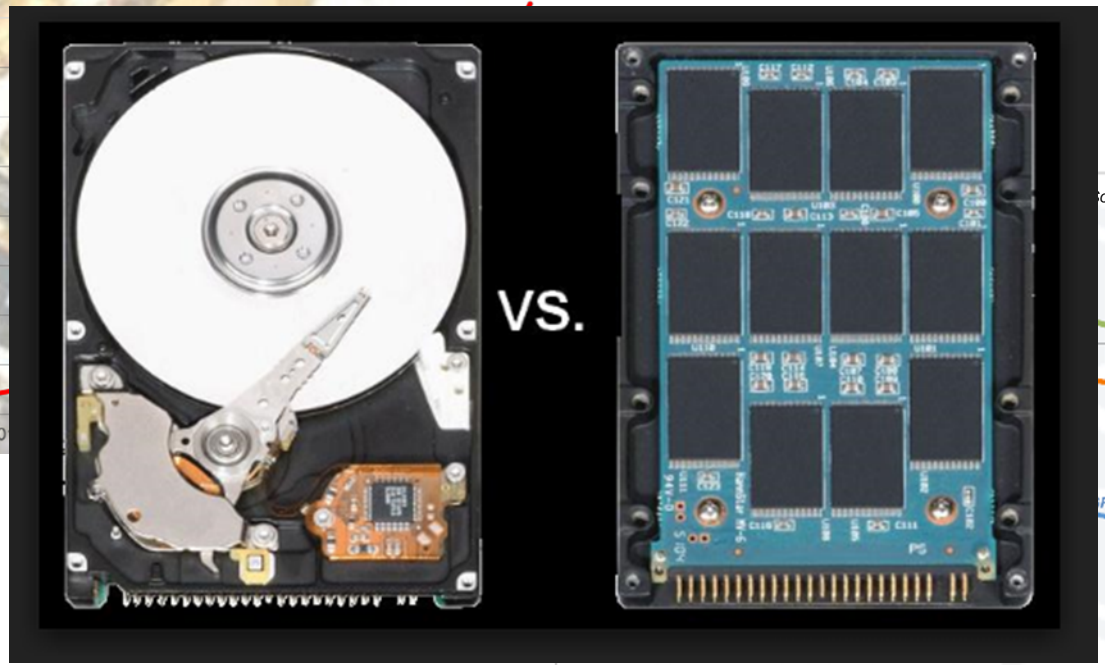

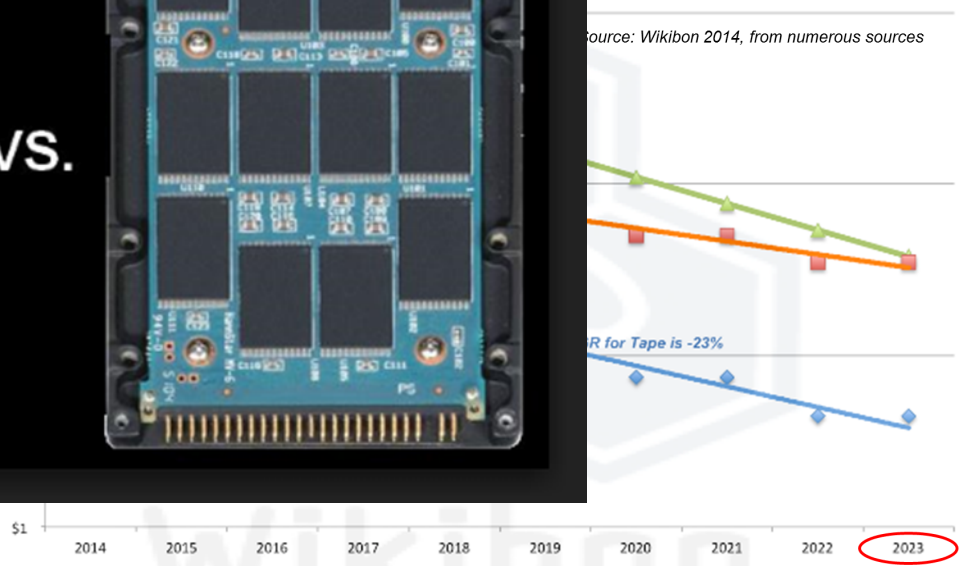

SSD are Getting Faster – Much Faster

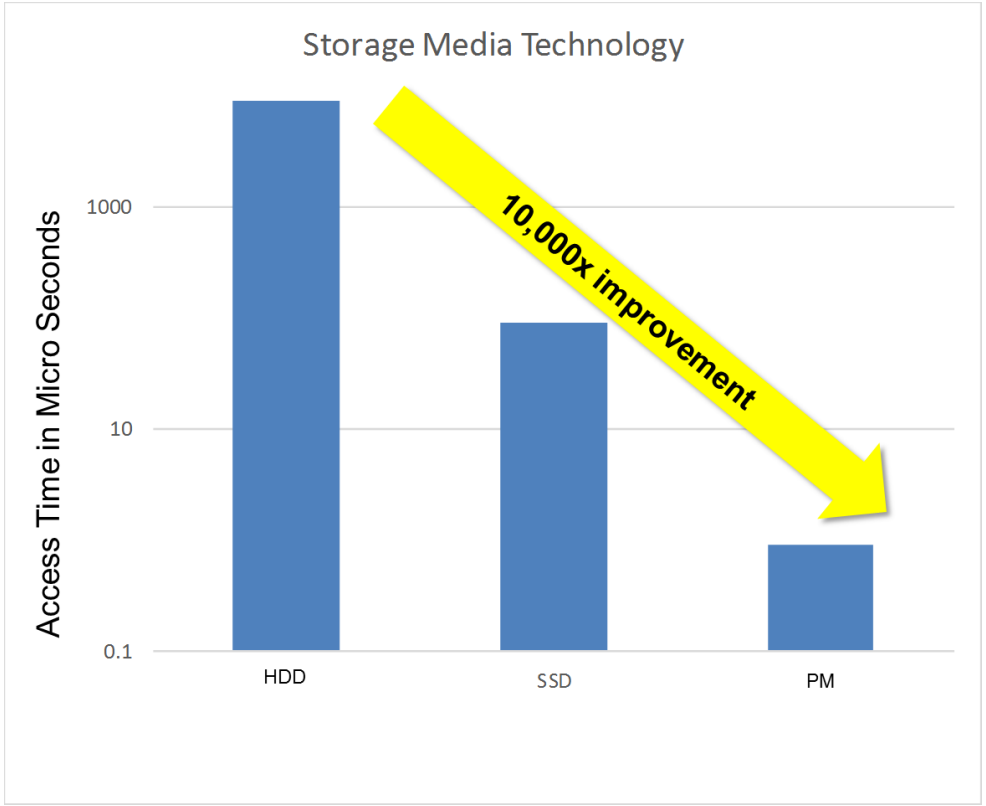

And so is Ethernet – Much Faster

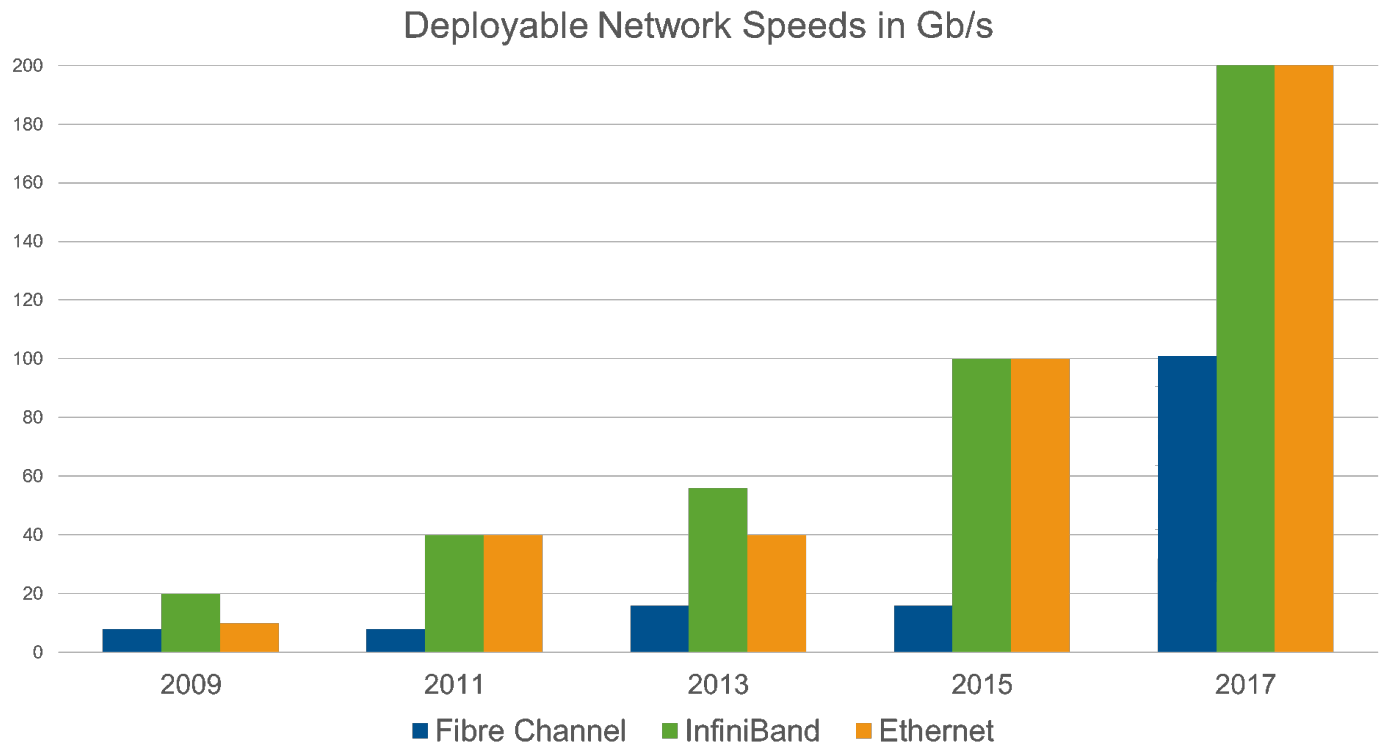

And Ethernet Storage is Also Innovating

- Pure Bandwidth – up to 100Gb/s
  - Flash over Block, File and Object

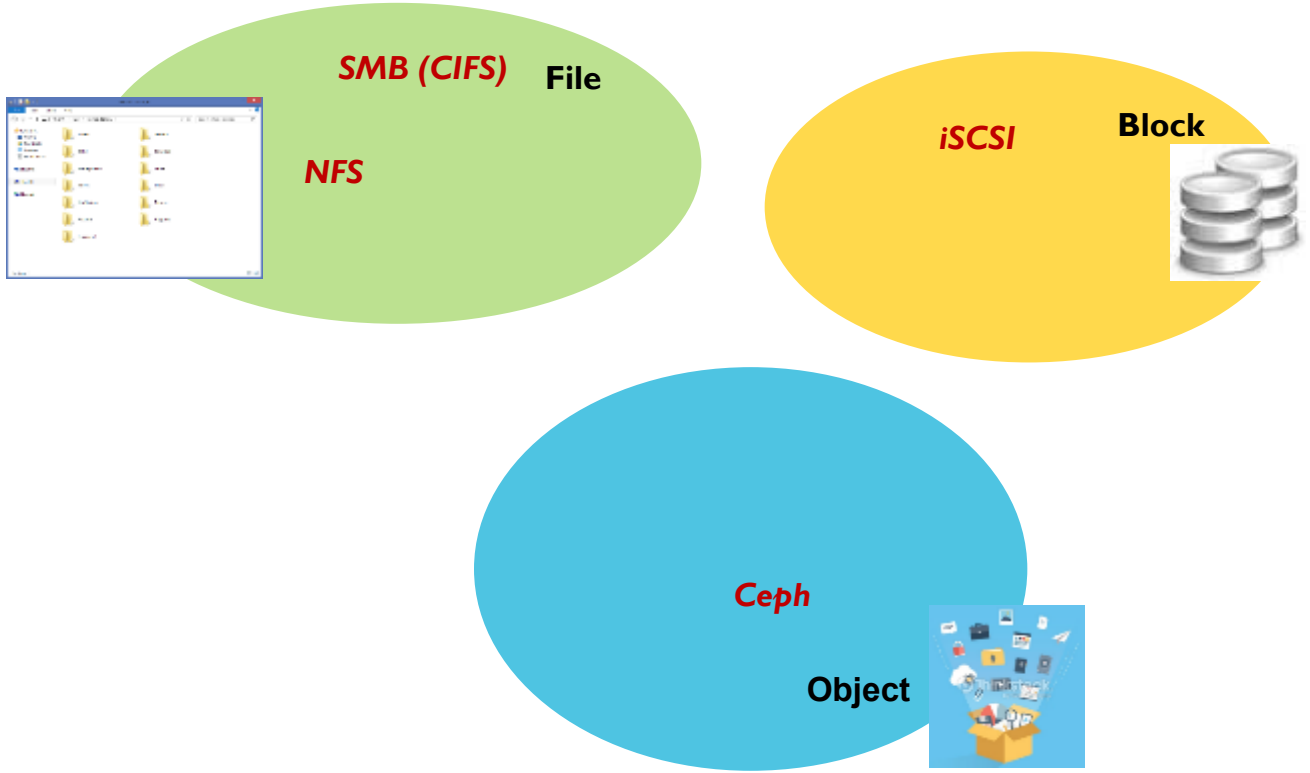

And Ethernet Storage is Also Innovating

- Pure Bandwidth – up to 100Gb/s
  - Flash over Block, File and Object
- RDMA
  - RoCE, iWARP, InfiniBand
  - iSER
  - SMB Direct, NFSoRDMA
  - Ceph over RDMA

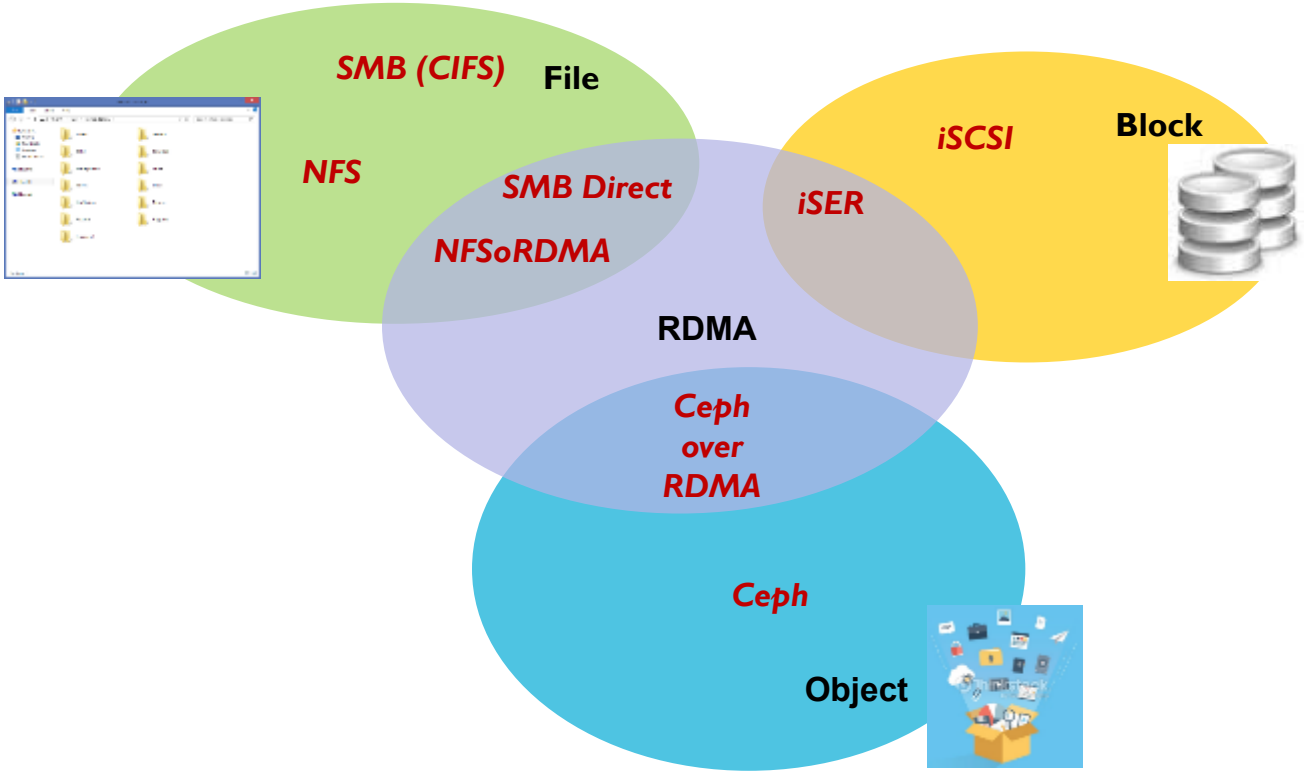

And Ethernet Storage is Also Innovating

- Pure Bandwidth – up to 100Gb/s
  - Flash over Block, File and Object
- RDMA
  - RoCE, iWARP, InfiniBand
  - iSER
  - SMB Direct, NFSoRDMA
  - Ceph over RDMA
- Persistent Memory(PM)
  - NVMe over Fabrics (NVMf)
  - PMf (NVDIMM, 3D-Xpoint, etc.)

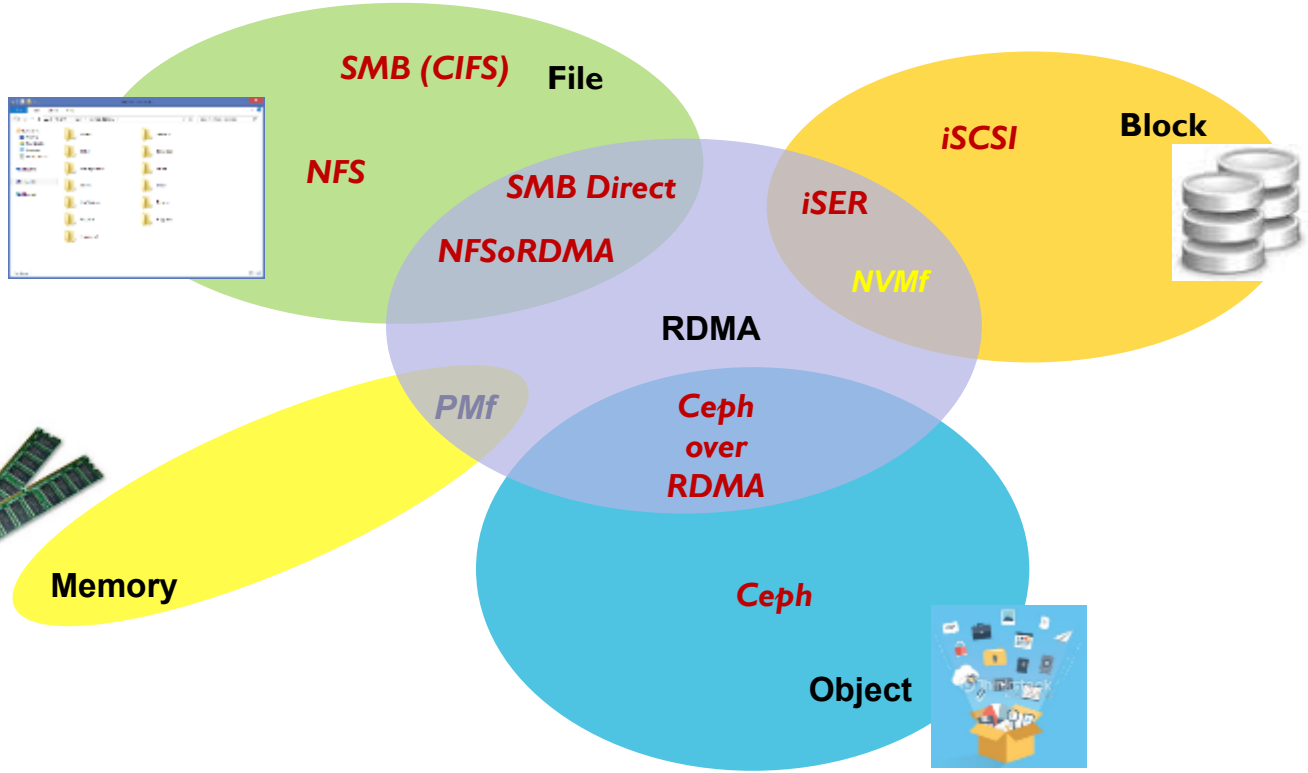

Faster Wire Speeds is only ½ the solution

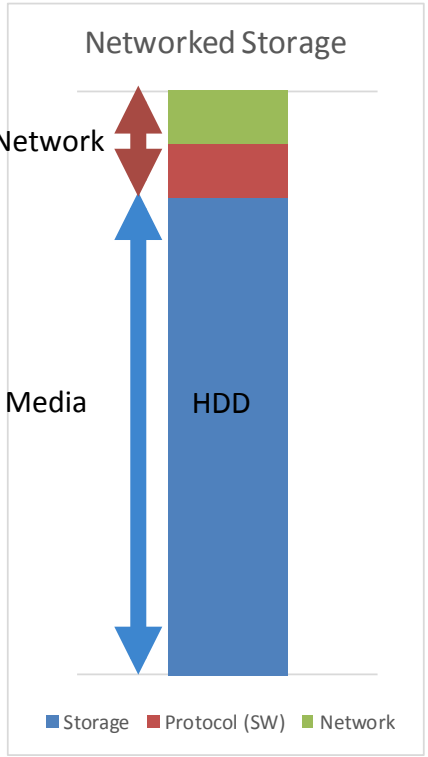

Faster Wire Speeds is only ½ the solution

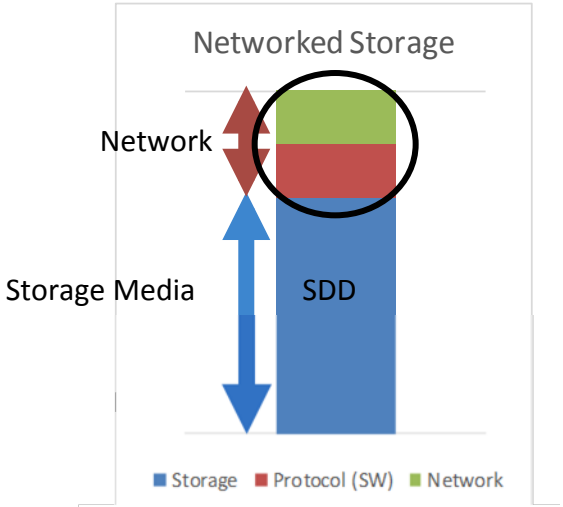

Faster Wire Speeds is only ½ the solution

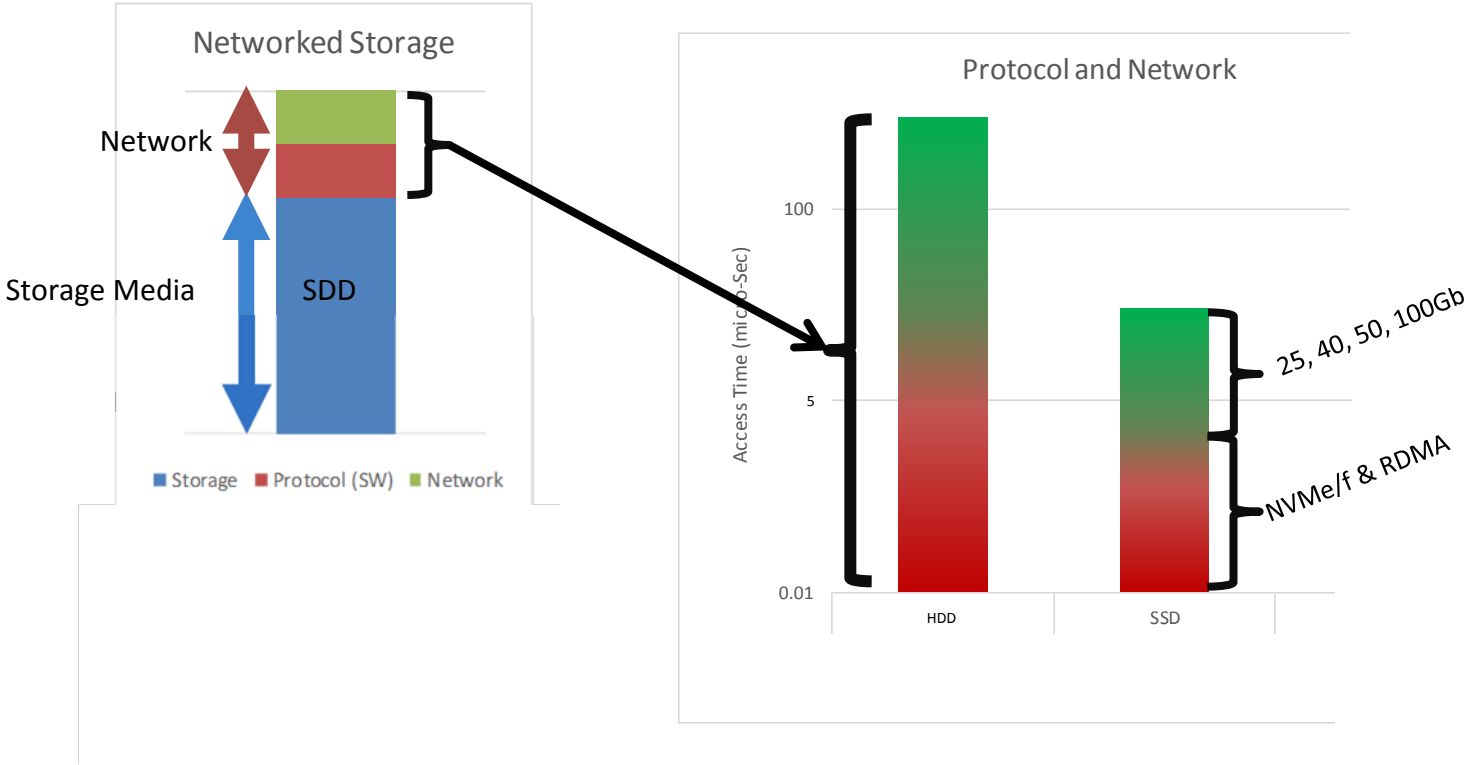

Faster Protocol: NVMe

- NVMe: Optimized for flash and next-gen NV-memory
  - Traditional SCSI interfaces designed for spinning disk
  - NVMe bypasses unneeded layers

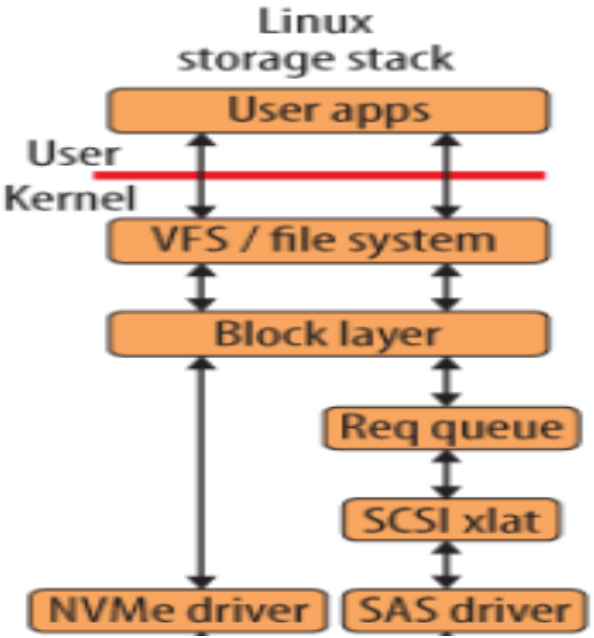

Faster Protocol: NVMe

- NVMe: Optimized for flash and next-gen NV-memory
  - Traditional SCSI interfaces designed for spinning disk
  - NVMe bypasses unneeded layers
- NVMe Flash Outperforms SAS/SATA Flash
  - 2x-2.5x more bandwidth, 40-50% lower latency, Up to 3x more IOPS

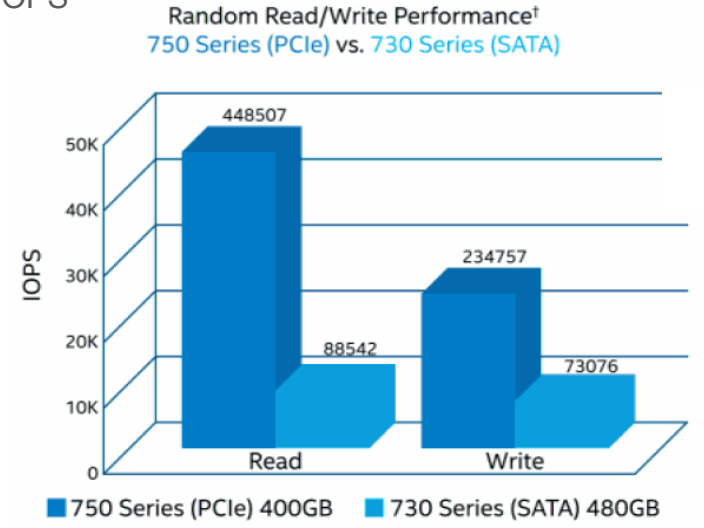

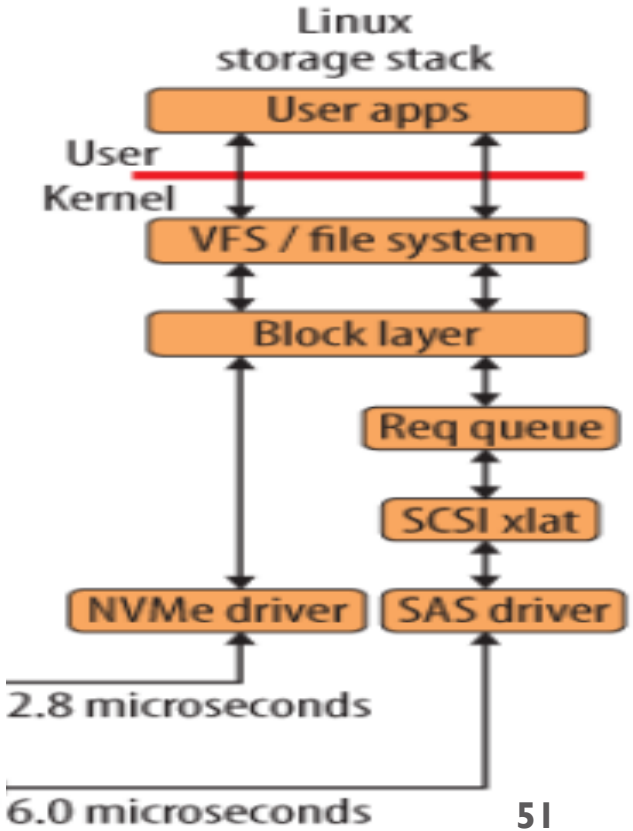

Faster Protocol: NVMf

- The idea is to extend the efficiency of the local NVMe interface over a fabric
  - Ethernet or IB
  - NVMe commands and data structures are transferred end to end
- Capitalizes on RDMA for performance

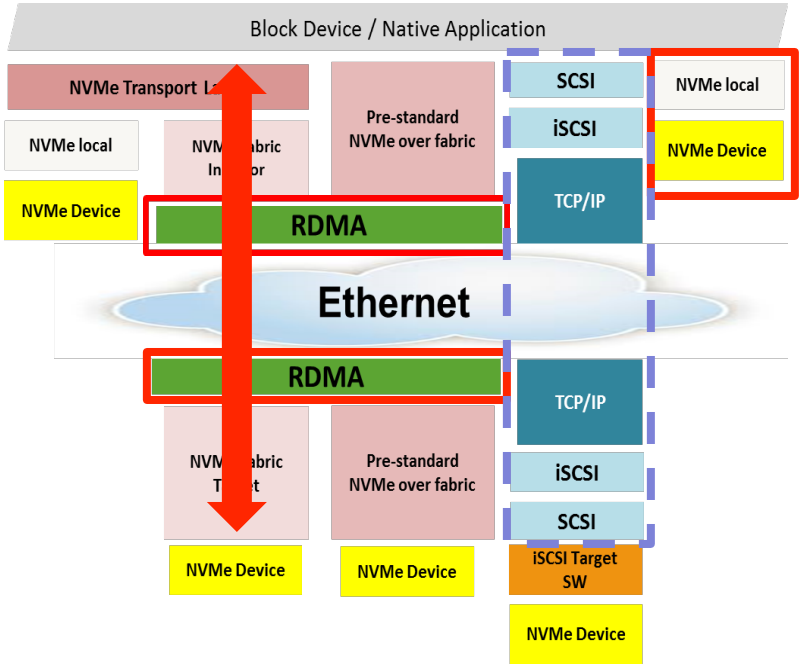

https://www.brighttalk.com/webcast/663/185909

Faster Protocol: RDMA

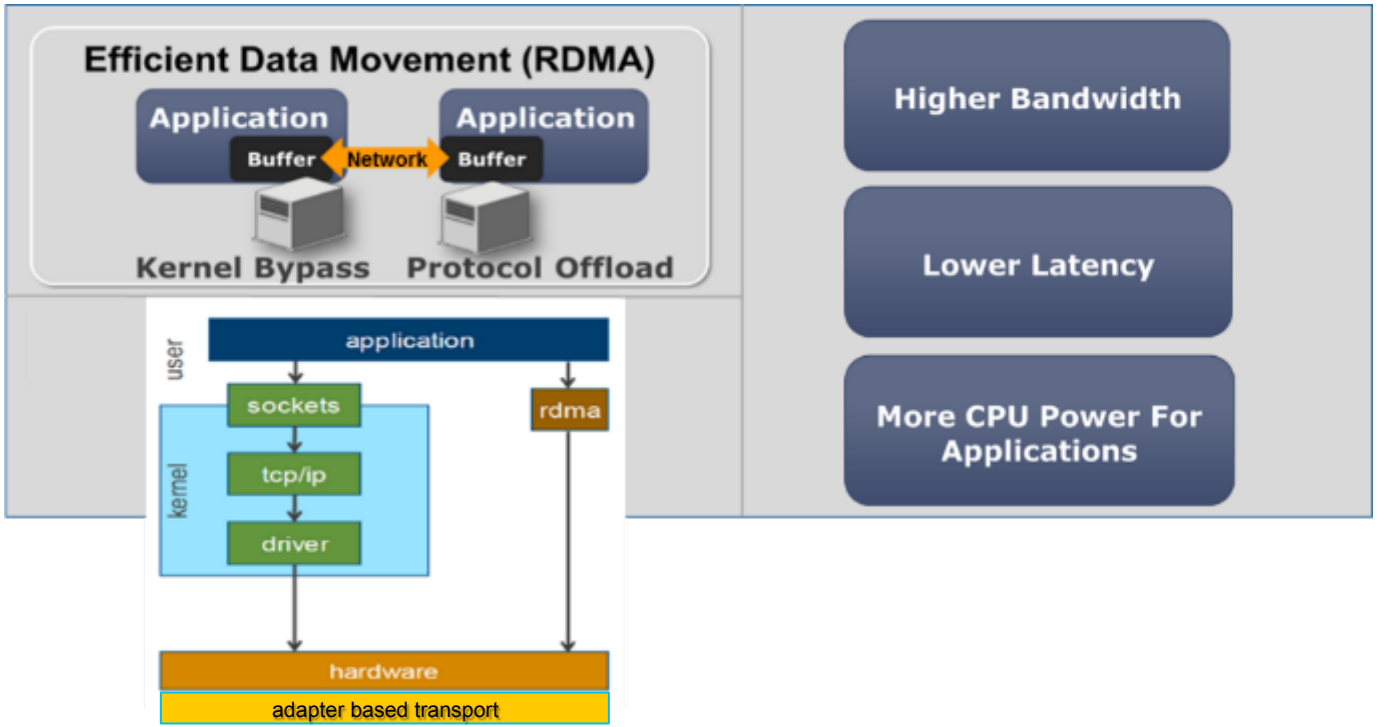

Faster Protocol: RDMA

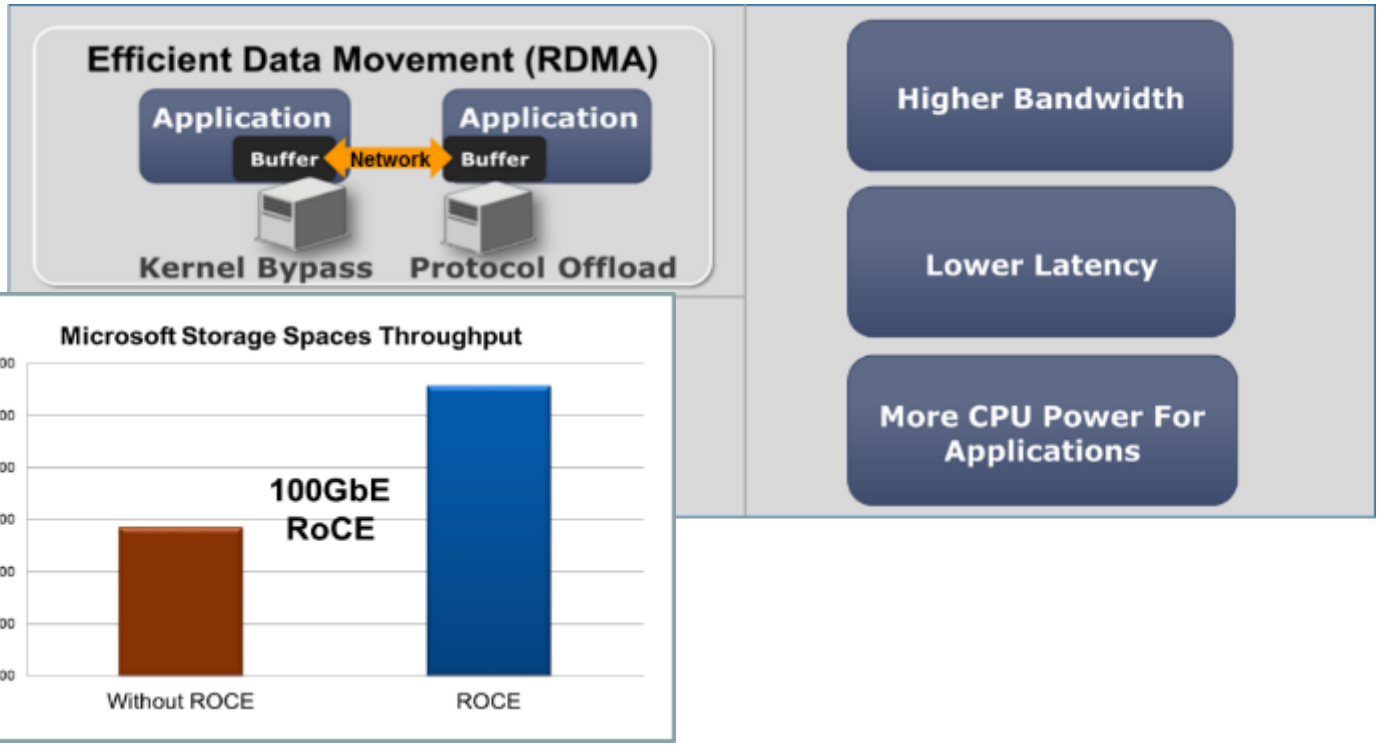

https://www.youtube.com/watch?v=u8ZYhUjSUoI

Summary

- Many factors are changing the direction of storage networking in Ethernet's favor
  - Accelerating capacity requirements driving lower cost solutions
  - SSD performance driving higher performance solutions
  - Scale-Out storage and Hyperconverged driving single network solutions
- Ethernet is VERY well positioned to take advantage of the new direction
  - Already dominates low cost File and Object solutions
  - Performance roadmap lines up well with SSD performance
  - Innovative new protocols that tightly couple to new storage directions

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
- 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

SNIA ESF | ETHERNET STORAGE

Thanks!