SNIA NSF | NETWORKING STORAGE

Storage Technologies & Practices Ripe for Refresh: Part 2

Live Webcast

June 22, 2021

10:00 am PT / 1:00 pm ET

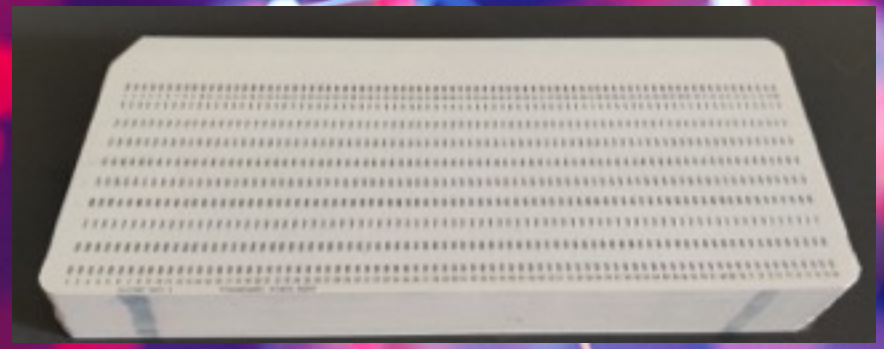

Today's Presenters

Alex McDonald
Moderator
Independent Consultant
Vice Chair SNIA NSF

Tom Friend
Illuminosi

John Kim
Chair, SNIA NSF
NVIDIA

SNIA-at-a-Glance

180
industry leading organizations

2,500
active contributing members

50,000
IT end users & storage pros worldwide

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

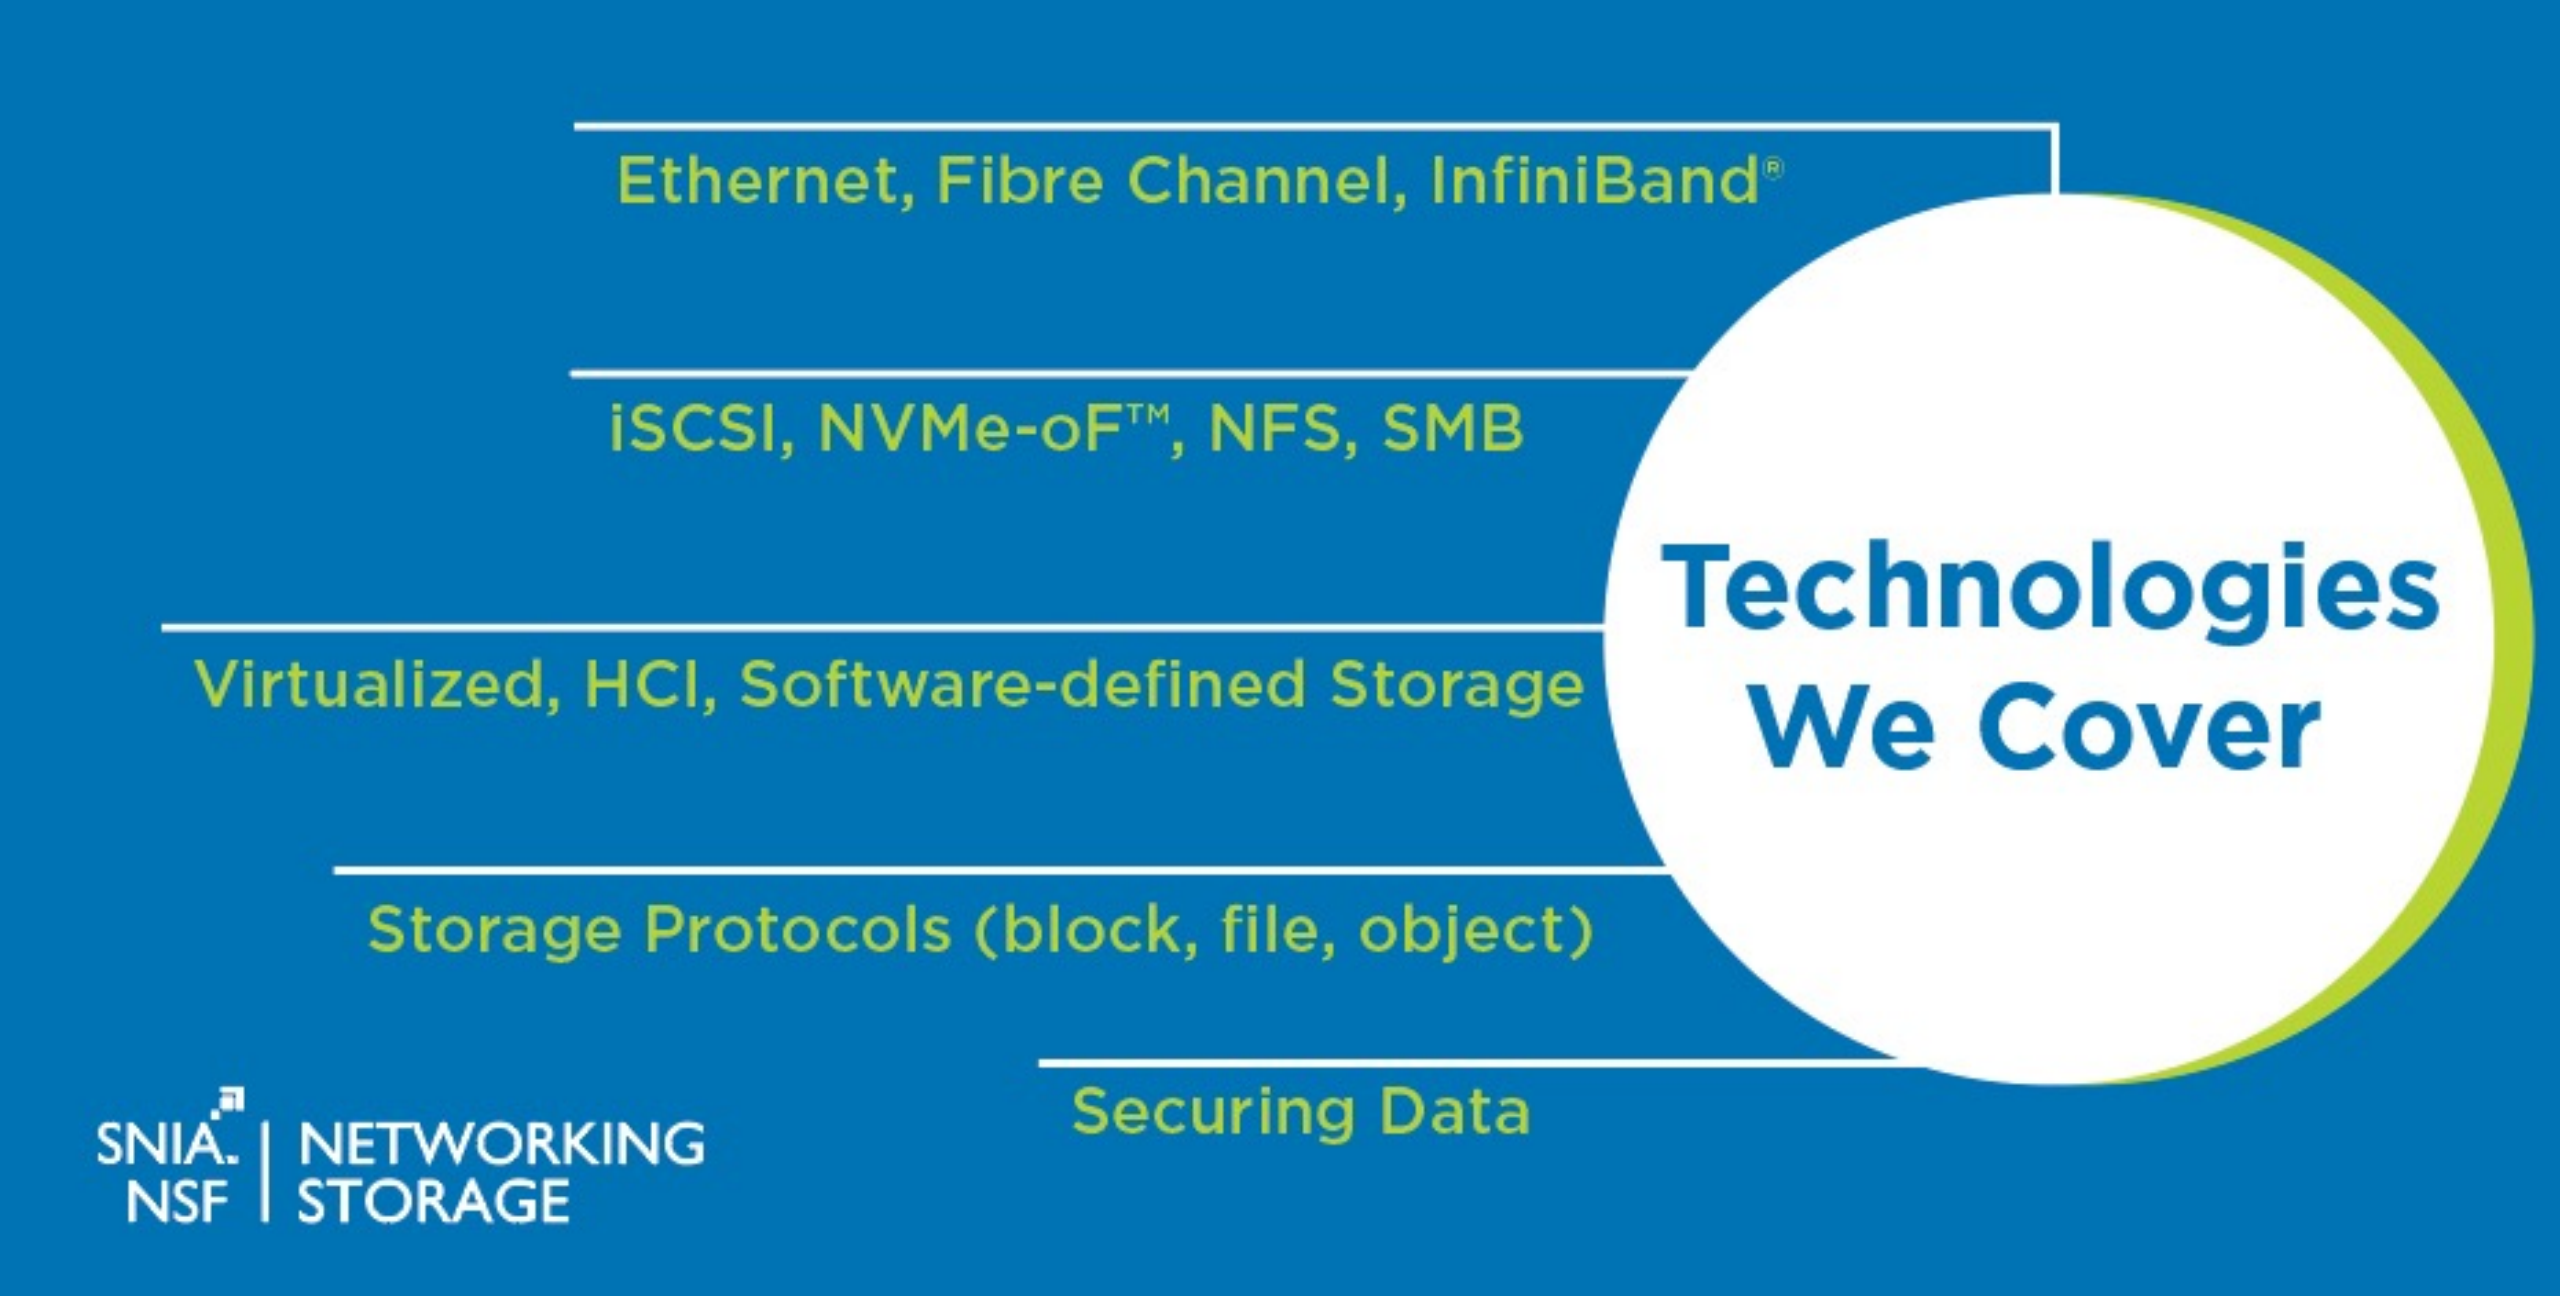

SNIA Legal Notice

- 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.

Agenda

- Obsolete HDD interfaces
- Problem file systems to review
- Not quite retired network technologies

Obsolete HDD Interfaces

Can you still get replacements/spares?

Tom Friend

HDD Interfaces to Replace

- Several types to consider:
  - SCSI Parallel Interface (SPI)
  - IDE / PATA
- Why?
  - SPI drives go bad:
    - You can't find replacement drives (of the same kind) any more
    - There are a myriad of subtypes here-
      - SCSI-1, Fast SCSI, Fast-Wide SCSI, Ultra SCSI, Ultra Wide SCSI, Ultra2 SCSI, Ultra3 SCSI, Ultra-320 SCSI, Ultra-640 SCSI, 50/68/80 pins
  - Parallel SCSI host controllers go bad:
    - You may not find a host controller that works with your old motherboard (ISA, MCA, EISA or PCI)
    - Same problems exist with SCSI sub-types
  - IDE/PATA interfaces are usually on the motherboard- can you replace that when it fails?

Old file systems

When you upgrade the hardware, ditch these

Windows File Systems to Replace

- Windows environment:
  - File Allocation Table variants: FAT12, FAT16, FAT32, exFAT
  - Replace with NTFS
- Why change?
  - FAT is designed for single user systems
    - No way to limit user access
    - No advanced features like redundancy, journaling or encryption
  - NTFS features include:
    - Access Control Lists (ACLs) define file & directory security
    - Encrypting files
    - Journaling (for efficient backup file revisions)
    - Redundant Master File Table for recovering damaged MFT
- What to change: disk must be taken off line for copy-style conversion

Linux File Systems to Replace

- Linux environment:
  - Early ext variants: ext, ext2, ext3
  - Replace with ext4
- Why change?
  - Ext features: Created in 1992, Single timestamp, 255 character file name, 2GB file size
  - Ext2 features: Created in 1993, 2TiB files, 16TiB file system, triple timestamps
  - Ext3 features: Adopted in 2001, 32k subdirectories, journaling for reliability, somewhat resistant to fragmentation
  - Ext4 features
    - Available in 2006
    - 16TiB files, 1EiB file system theoretical, 50TiB practical (Red Hat limit), unlimited subdirectories
    - Higher performance, lower fragmentation, journal checksums, improved timestamps, online defrag
    - In place (live) upgrade from ext3 to ext4
- What to change: disk must be taken off line for copy-style conversion between ext, ext2 & ext3

Linux Advanced File stems

- Linux file systems- beyond ext4
  - XFS
    - Default for Red Hat 7 and beyond
    - Highly performant in large concurrent use cases
    - Address limits beyond 50TiB
  - OpenZFS
    - Zettabyte volume size, 16 EiB file size
    - Volume management
    - Block-level checksum, async incremental replication, redundancy, snapshots, clones, defrag, self-healing
    - CDDL, not GPL licensing
  - btrfs
    - Many of the same features as ZFS, but uses GPL licensing
    - 16 EiB file & volume size

Not Quite Retired Networking

More Networking You Should Consider Upgrading

John Kim

Aging Network Technologies

- Twisted Pair Ethernet Wiring
- Switches with a Locked NOS
- iSCSI

Twisted Pair Ethernet

- Ethernet mainstay -- unshielded copper in twisted pairs
  - Established 1986/1990 (StarLAN/10BASE-T)
  - Speeds from 10 Mb/s to 10 Gb/s, low-cost cabling
- Hits speed limitations beyond 10Gb/s
  - Range limitations, high power consumption
  - Very limited adoption of 25GBASE-T (IEEE 802.3bq in 2016)
- Copper cabling in general hits distance limits
  - At 25Gb/s per lane, effective limit = 5m; at 50Gb/s, limit = 3m
  - At 100Gb/s per lane, copper limit might be 1-2m
  - Move to optical technology but only for higher speeds

Switches with Locked Network Operating System

- Traditional switch: Bundled hardware and NOS
  - Simple to buy and support but beholden to switch vendor
  - Hard to change, inflexible updates or fixes
- New option: Decouple hardware and software
  - Choose a NOS, commercial or open source
  - Install NOS yourself or bundled by switch vendor
  - May require higher level of expertise and self-support
- Open switches to dominate very large networks

Traditional switch

Open Ethernet switch

iSCSI—Is its Future Limited?

- iSCSI—The midrange block storage protocol
  - Widely supported, economical, usually in software
- Squeezed by NVMe-oF, improved file/SAN options
  - NVMe-oF is more efficient but not as widely supported
  - NAS protocols gain RDMA boost, support more apps
  - Fibre Channel speed boost: NVMe-oF option, 64G
- Prognosis: iSCSI will slowly decline
  - Will continue in limited use for years

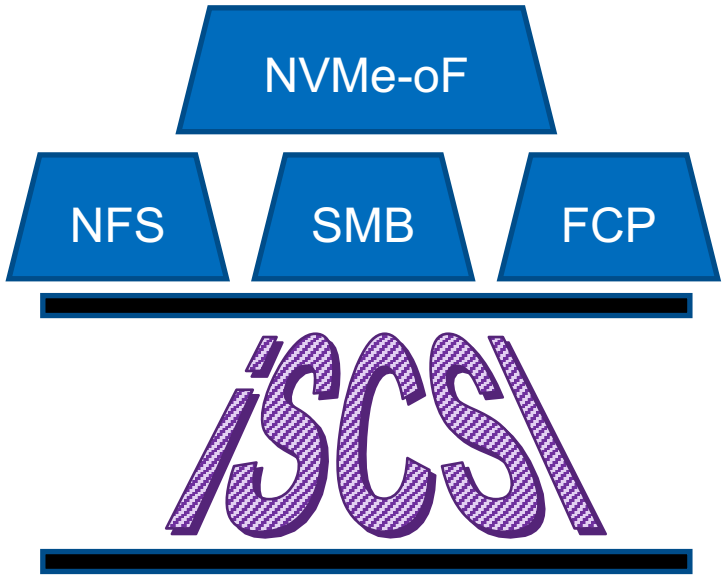

“Ripe for Refresh” Summary

- Review/replace parallel HDD interface devices:
  - SPI, IDE…
- Replace ext, ext2, ext3 with ext4 or better
- Replace FAT12, FAT16 & FAT32 with NTFS
- Consider btrfs or ZNS raid solutions
- Change cabling for higher-speed networking
- Open Ethernet switches in very large networks
- iSCSI gradually replaced by alternatives

“Part 1” of this Series On-Demand

- Storage Technologies & Practices Ripe for Refresh covers other technologies whose time is running out:
  - Security Problems
  - Not quite retired network protocols
  - NAS protocols
- Watch on-demand at: http://bit.ly/RipeforRefresh

After this Webcast

- Please rate this webcast and provide us with your feedback
- This webcast and a copy of the slides will be available at the SNIA Educational Library https://www.snia.org/educational-library
- A Q&A from this webcast, including answers to questions we couldn't get to today, will be posted on our blog at https://sniansfblog.org/
- Follow us on Twitter @SNIANSF

Thank You

