

SNIA

STORAGE NETWORKING INDUSTRY ASSOCIATION

EDUCATION

Storage Networking Standards: Future Directions

David L. Black, EMC Corporation,
Member, SNIA Technical Council

Steven Wilson, Brocade Communications,
Member, SNIA Technical Council

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individuals may use this material in presentations and literature under the following conditions:
 - Any slide or slides used must be reproduced without modification
 - The SNIA must be acknowledged as source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee and the SNIA Technical Council

Storage Networking Standards: Future Directions

Interoperability standards play a vital role in customer adoption and advancement of storage networking technologies and systems. Storage networking is based on a broad spectrum of standards (developed by multiple standards organizations) in areas such as Fibre Channel (INCITS T11), SCSI (INCITS T10), iSCSI (IETF), and storage management (SNIA, DMTF, IETF). The current state and future direction of standards development can provide useful insights into technology developments. This tutorial covers storage networking standards and the role that the resulting standardized interfaces/functionality play in networked storage infrastructure. The tutorial presenters are members of the SNIA Technical Council who are actively involved in development of many storage networking standards, including those listed in this abstract.

About the Authors

David L. Black

Senior Technologist, EMC
Corporation

Member, SNIA Technical Council

Chair, T11 FC-SP-2 (Fibre
Channel Security-2) Working
Group

Chair, IETF IP Storage (IPS)
Working Group


Chair, SNIA Fixed Content Aware
Storage (FCAS) Technical
Working Group

Steven Wilson

Director, Technology and
Standards, Brocade
Communications

Member, SNIA Technical Council
Chair, T11 FC-IFR (Fibre
Channel Inter-Fabric Routing)
Working Group

Editor, T11 FC-SW-5 (Fibre
Channel Switch Fabric-5)



SNIA

STORAGE NETWORKING INDUSTRY ASSOCIATION

EDUCATION

Storage Networking Standards: Future Directions

David L. Black, EMC Corporation,
Member, SNIA Technical Council

Steven Wilson, Brocade Communications,
Member, SNIA Technical Council

Overview

- Interoperability standards for storage
 - Connect products from different vendors
- Standards can provide technology insight
 - Emerging technology
 - Evolution of existing technology
- This talk: Standards developments and directions
 - Implications for technology vendors and users

Standards: End User Benefits

- Protect technology investment
- Ensure a base level of interoperability
- Provide choice among products
- Ensure continuing innovation
- Commonality leads to less training, simpler deployment

Storage Networking Standards

Data Protocols

- SCSI (T10)
- Fibre Channel (T11)
- ATA and SATA (T13)
- NAS (IETF, Microsoft)
- IP Storage (IETF)


Management

- SMI-S (SNIA)
 - Uses CIM (DMTF)
- SNMP (IETF)
- Web (IETF, W3C)

Stored Data

- RAID Layout (SNIA)
- Encryption (IEEE)
- Fixed Content (SNIA)

Storage Protocol Classes



Network Communication

- Fibre Channel Fabric
- IP Storage (iSCSI, FCIP, iFCP)
- Network Attached Storage (NAS)

Drive Interface

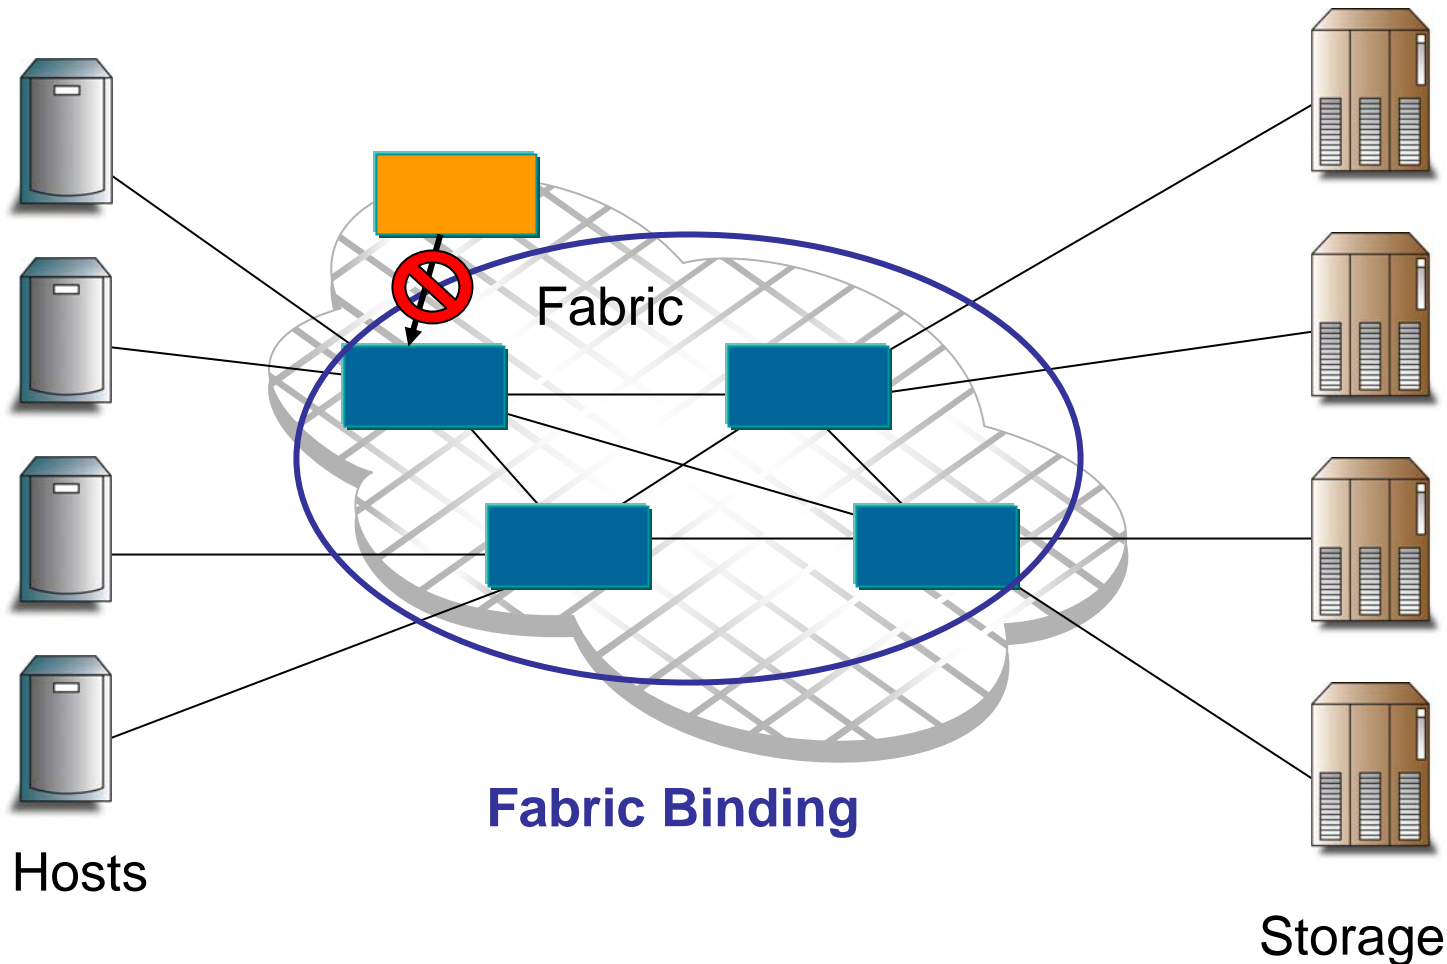
- FC-AL (Arbitrated Loop)
- Parallel SCSI and Serial Attach SCSI (SAS)
- ATA and Serial ATA (SATA)

Fibre Channel Developments 1: Security (T11 FC-SP)

- Fibre Channel (FC) Fabric access and config controls
 - Control fabric structure and what can join the fabric
- In-band Authentication
 - Secret (CHAP, DH-CHAP)
 - CHAP = Challenge Handshake Authentication Protocol
 - DH-CHAP = Diffie-Hellmann CHAP (adds a DH exchange)
 - Public Key (FCAP) and Password (FCPAP)
 - FCAP = Fibre Channel Authentication Protocol
 - FCPAP = Fibre Channel Password Authentication Protocol
- Secure communication channels
 - Adaptation of IPsec subset to Fibre Channel

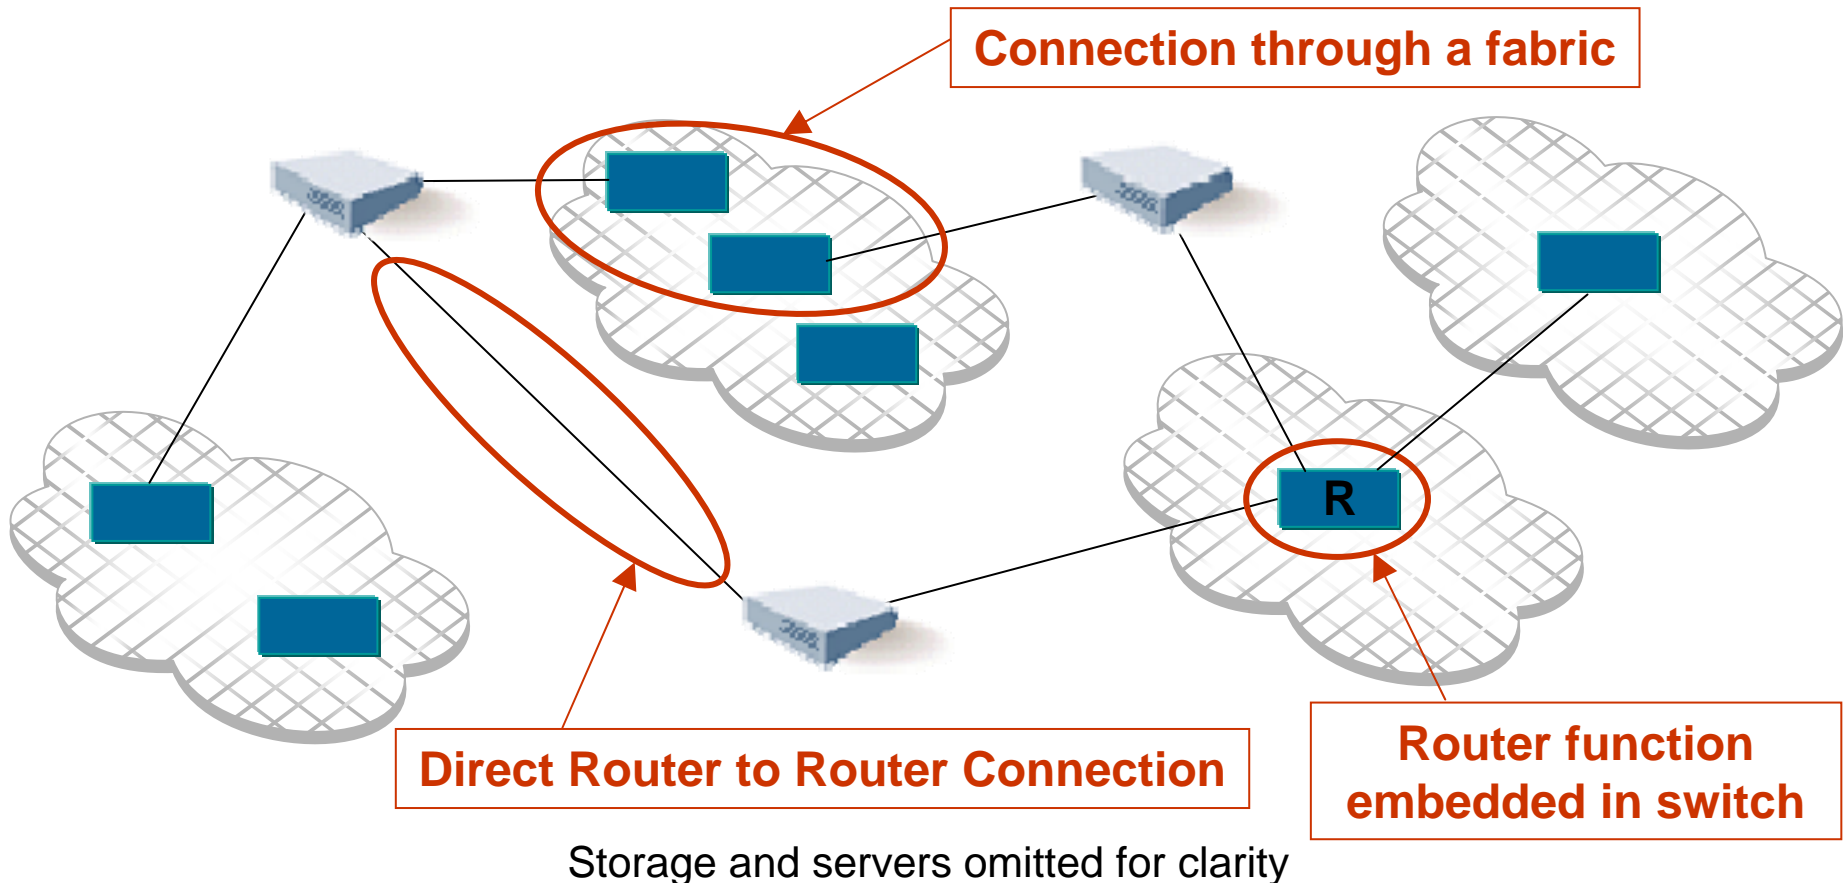


Fabric Access Control Example: Block Unauthorized Switch



Fibre Channel Developments 2: Inter-Fabric Routing

- FC Routers inter-connect FC Storage Area Networks (SANs)
 - Particularly useful for isolated SANs (SAN islands)



Inter-Fabric Routing Properties

- Routing interconnects physical **and** virtual fabrics
 - Virtual fabrics can be in the same or different physical SAN
 - Routes can pass through existing fabrics and switches
- The interconnected fabrics **do not merge**
 - Prevents some disruption propagation (e.g., RSCN)
 - RSCN = Registered State Change Notification
 - Address translation used as necessary
- Routing is transparent to servers and storage
 - Zoning, name service, etc. continue to work
- Routing function can be packaged in a separate router **or** combined with a fabric switch

Fibre Channel Developments 3: Communication Media

- 4 Gbit/sec Fibre Channel speed
 - Compatible upgrade, easy fabric transition
 - Similar to 1 Gbit/sec to 2 Gbit/sec transition
 - Next speed upgrade will be to 8 Gbit/sec
 - Limited deployment of 10 Gbit/sec Fibre Channel
- FC Base-T: Fibre Channel over twisted pair
 - Twisted Pair: Category 5 and 6 (Cat5 and Cat6) cable
 - Same cable as Gigabit and 10 Gigabit Ethernet
 - 1, 2, 4 Gbit/sec speed – distance varies
 - Cat 6 & 6a cable for longer runs, not Cat 5 & 5e

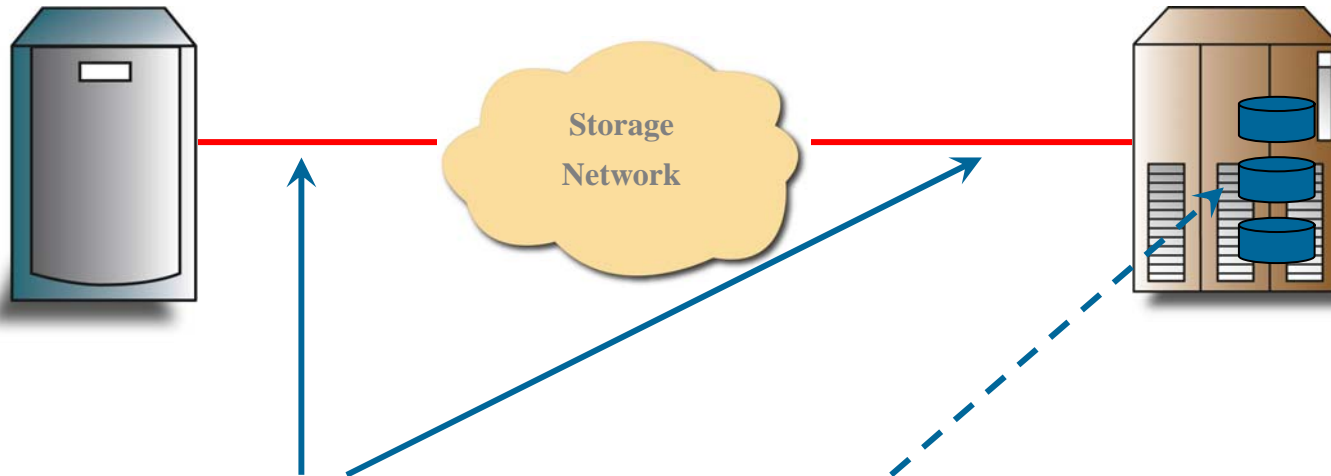


Fibre Channel Developments 4: Protocol Topics

- FAIS API to program intelligent FC switches
 - FAIS = Fabric Application Interface Specification
 - Enables storage applications within the fabric
 - First version of FAIS nearing completion
- N_Port (usually server) identity virtualization (NPIV)
 - Logical N_Port per virtual server supports server virtualization
 - Replaces questionable use of multiple FC-AL identifiers
- Fibre Channel adaptation of IP networking concepts
 - Fibre Channel MIBs (for SNMP)
 - MIB = Management Interface Base
 - SNMP = Simple Network Management Protocol
 - Fibre Channel *ping* and *traceroute*



Storage Protocol Classes



Network Communication

- Fibre Channel Fabric
- IP Storage (iSCSI, FCIP, iFCP)
- Network Attached Storage (NAS)

Drive Interface


- FC-AL (Arbitrated Loop)
- Parallel SCSI and Serial Attach SCSI (SAS)
- ATA and Serial ATA (SATA)

IP Storage Developments

- RFC specifications are being published
 - iSCSI (RFC 3720), FCIP (RFC 3821), iFCP (RFC 4172)
 - MIBs have been approved; appearing during 2006
- RDMA for iSCSI (iSER)
 - RDMA = Remote DMA over a TCP/IP network
 - iSER = iSCSI Extensions for RDMA
 - InfiniBand: iSER emerging as an alternative to SRP (SRP = SCSI RDMA Protocol)
- Fibre Channel Pseudo-Wire over MPLS
 - MPLS: Multi-Protocol Label Switching
 - Primarily for carrier MPLS infrastructure
 - Will be jointly developed by IETF and T11



Storage Protocol Classes



Network Communication

- Fibre Channel Fabric
- IP Storage (iSCSI, FCIP, iFCP)
- Network Attached Storage (NAS)

Drive Interface


- FC-AL (Arbitrated Loop)
- Parallel SCSI and Serial Attach SCSI (SAS)
- ATA and Serial ATA (SATA)

Network Attached Storage (File Serving Protocols)

- NFS, primarily for Unix clients (IETF)
 - Transition to NFSv4 is underway
 - Parallel NFS (pNFS): SAN filesystem standardization
 - Parallel (e.g. striped) access across NFS servers
 - RDMA for NFS
 - Draft of NFS 4.1 specification is available
 - draft-ietf-nfsv4-minorversion1-06.txt
 - Sessions, directory delegations and pNFS
- CIFS, primarily for Windows clients (Microsoft)
 - Please ask Microsoft



Parallel NFS - pNFS




- NFS file naming, management, and administration
- Parallel high bandwidth file access (via Storage Network)

IP Infrastructure: Selected Developments

- IPv6 deployment increasing
- Many proposed TCP changes for high bandwidth
 - IRTF research group will make recommendations
 - IRTF = Internet Research Task Force
- 10 Gbit Ethernet: 10GBase-T
 - For Category 6 (Cat6) copper twisted pair cable
 - Will increase interest in TCP/IP offload engine chips (TOEs)
 - And other offload techniques (e.g., dedicated CPU core)
- IPsec: Better Than Nothing Security (btns)
 - Remove IPsec (IKE) identities to simplify management
 - Bind IPsec security to higher layer identities (e.g., iSCSI, NFS)
 - In progress at IETF

Storage Protocol Classes



Network Communication

- Fibre Channel Fabric
- IP Storage (iSCSI, FCIP, iFCP)
- Network Attached Storage (NAS)

Drive Interface

- FC-AL (Arbitrated Loop)
- Parallel SCSI and Serial Attach SCSI (SAS)
- ATA and Serial ATA (SATA)

Drive Interface Developments

- ATA drives: Serial ATA (SATA)
 - Replacing Parallel ATA
- SCSI drives: Serial Attached SCSI (SAS)
 - Replacing parallel SCSI, gradually
 - Lots of parallel SCSI products still available
 - Connection-based protocol (not packet switched)
 - Can carry SATA traffic, attach to SATA drives
 - Includes inband management protocol
 - SAS Zoning under development
 - Will allow shared SAS infrastructure and storage
- Fibre Channel drives: FC-AL
 - 4 Gbit/sec drives coming, but after switched fabric
 - FATA/FC-LC: ATA-class disk with FC-AL interface
 - FC-SATA standard will enable FC to attach to SATA drives



SCSI Command Set Developments

- SCSI Command Sets are used for:
 - Storage Networks (e.g., FC Fabric, iSCSI)
 - Disk and Tape drives (e.g., FC-AL, SAS)
- 1. Data Integrity Field (DIF)
 - 16 bit CRC + data offset within LUN (volume)
 - Detects corruption and mistakes/errors (typically software)
- 2. Security Protocol commands
 - Important target: Encrypting tape drives (using SSC commands)
 - Also support for Trusted Computing Group protocol (ask TCG)
- 3. Management: SCSI MIB for SNMP (recently completed)



Storage Networking Standards

Data Protocols

- SCSI (T10)
- Fibre Channel (T11)
- ATA and SATA (T13)
- NAS (IETF, Microsoft)
- IP Storage (IETF)

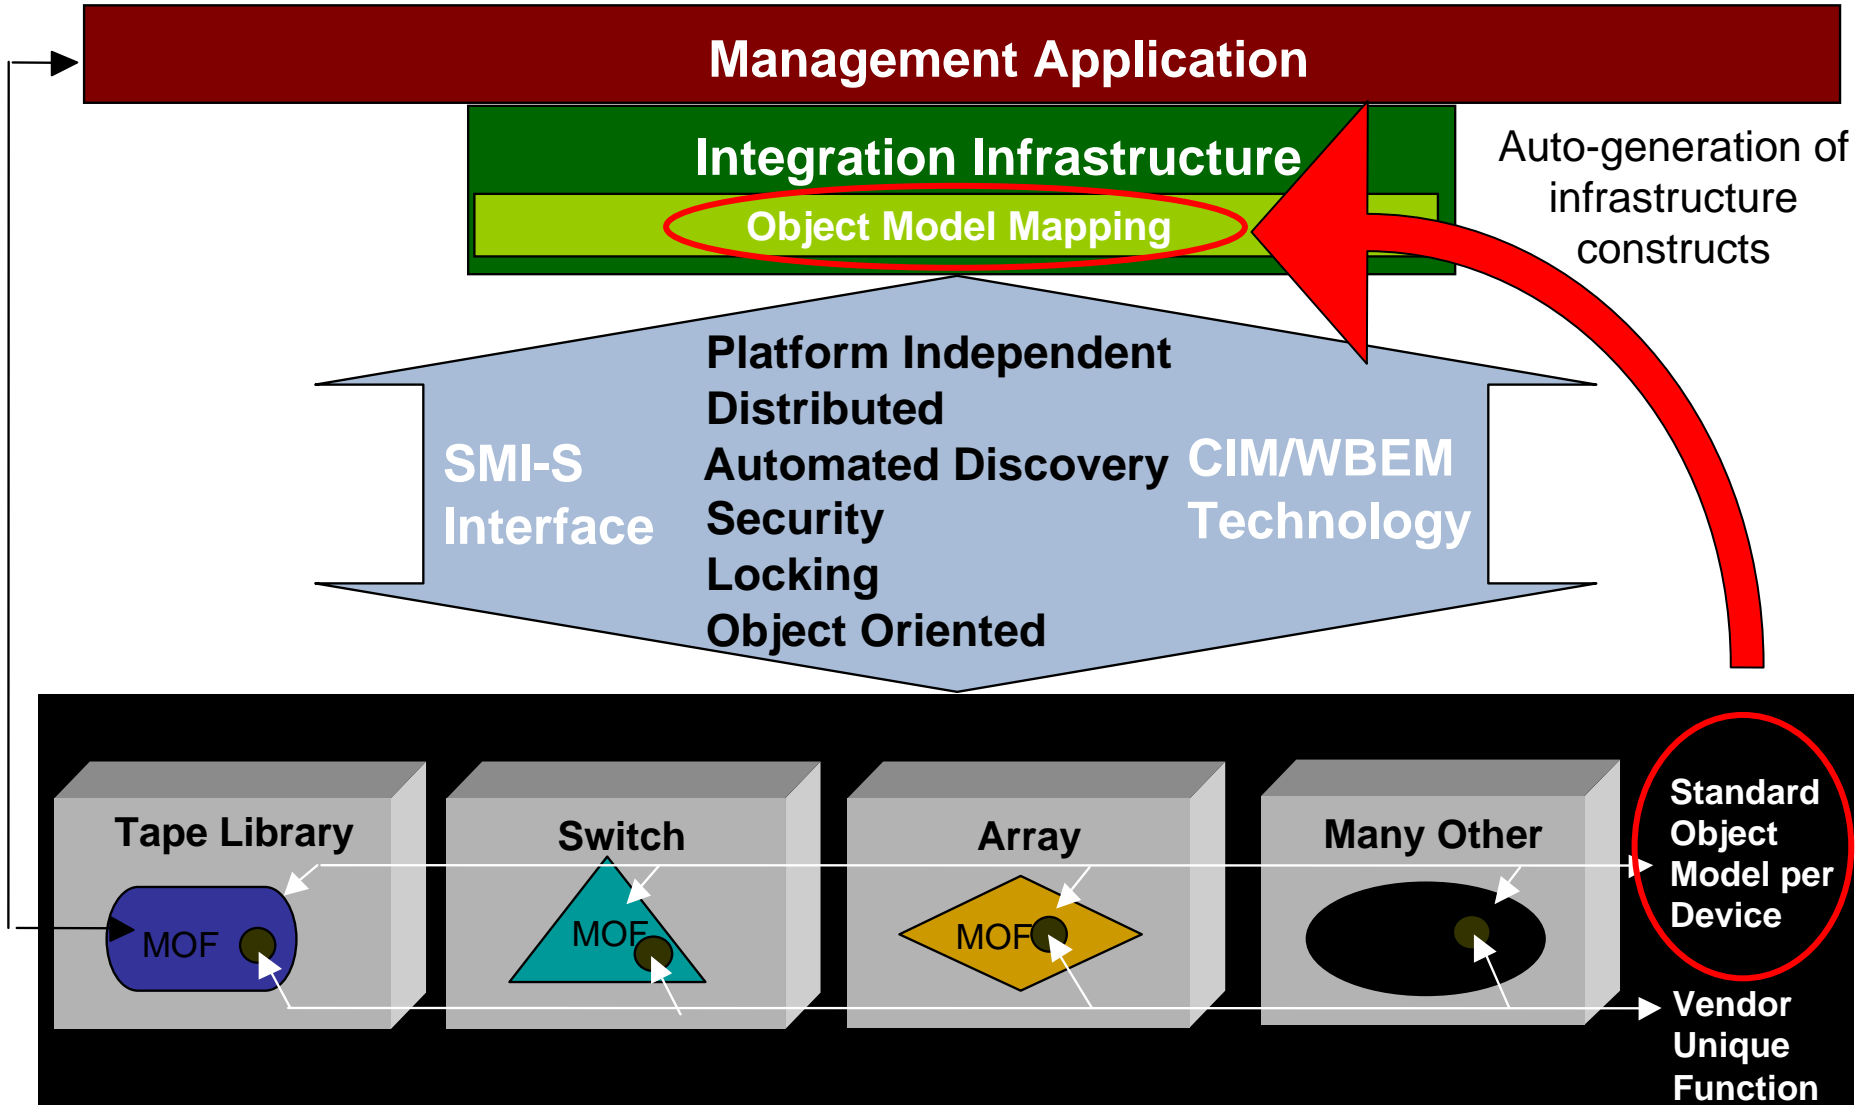
Management

- SMI-S (SNIA)
 - Uses CIM (DMTF)
- SNMP (IETF)
- Web (IETF, W3C)

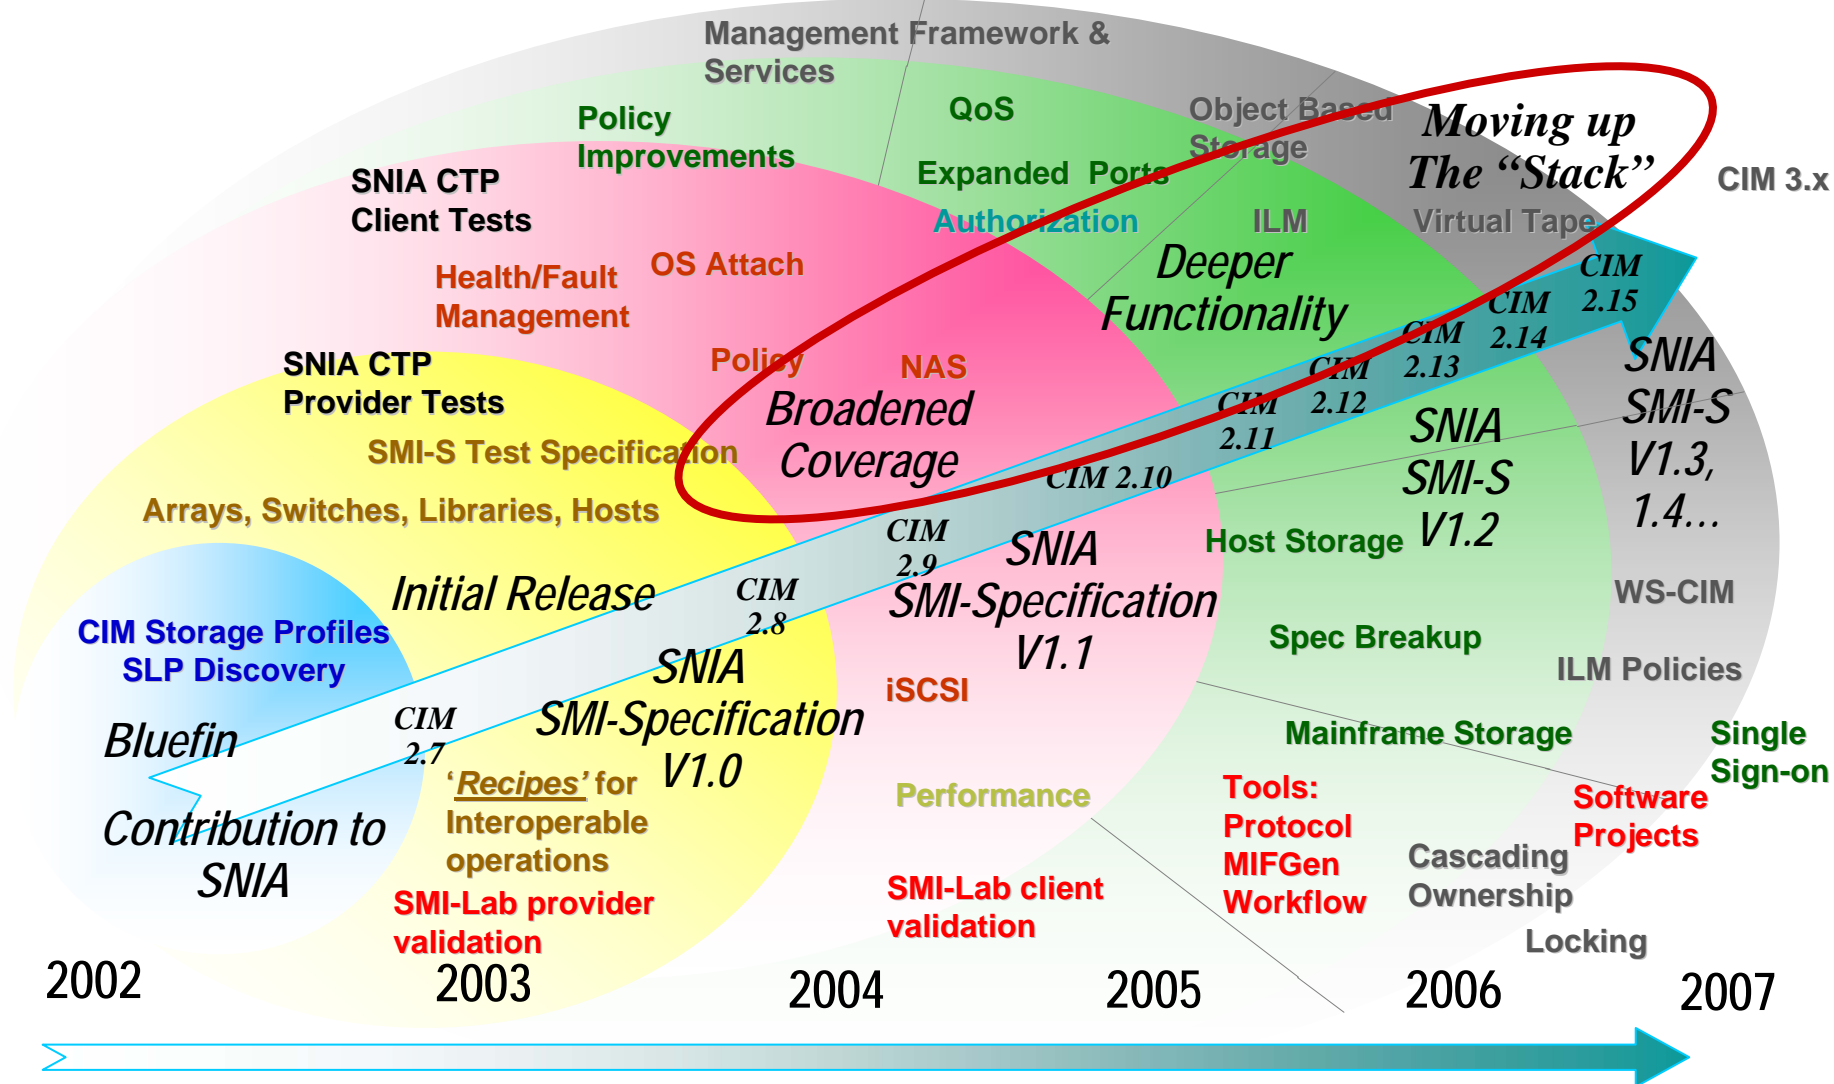
Stored Data

- RAID Layout (SNIA)
- Encryption (IEEE)
- Fixed Content (SNIA)

SMI-S Architecture: Model Based Storage Management



SMI-S Technology Road Map



SMI-S Progress

Version	SMI-S v1.0	SMI-S v1.1	SMI-S v1.2
Release Date (Dependent on Conformance Test Availability)	March 2004	April 2005	October 2006
Functionality:	Basic Discovery and Configuration of Fibre Channel SANs	System Management of Heterogeneous Storage Environments	Basic Data Management and Advanced System Management of Heterogeneous Storage Environments
Products:	<ul style="list-style-type: none"> • SAN Storage • FC Fabric & switches • FC HBAs 	<ul style="list-style-type: none"> • iSCSI Storage • Tape Libraries • Storage-related host software (e.g., Vol. Mgr) 	<ul style="list-style-type: none"> • NAS Storage

SNIA Management APIs

- Two recent SNIA API standards
 - iSCSI HBA Management API (IMA)
 - Multipath Management API (MMA)
- These are server APIs
 - Analogous to the Fibre Channel HBA API
 - Can be used by SMI-S providers

IP Management Standards

- SNMP (IETF): SNMPv3
 - SNMP versions before SNMPv3 are not secure!
 - SSH protocol will be applied to SNMPv3 security
 - SSH = Secure Shell (e.g., remote secure CLI)
- Web (IETF, W3C, OASIS)
 - Storage management can use web standards
 - Web services generating increasing interest
 - Two management stacks (WSDL, WS-MAN)
 - Convergence whitepaper has been published

Storage Networking Standards

Data Protocols

- SCSI (T10)
- Fibre Channel (T11)
- ATA and SATA (T13)
- NAS (IETF, Microsoft)
- IP Storage (IETF)

Management

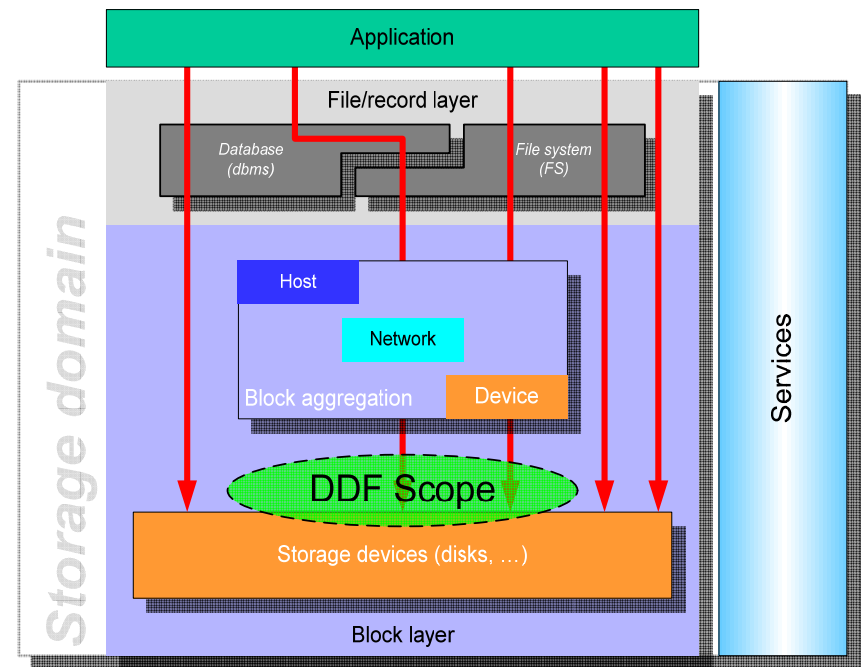
- SMI-S (SNIA)
 - Uses CIM (DMTF)
- SNMP (IETF)
- Web (IETF, W3C)

Stored Data

- RAID Layout (SNIA)
- Encryption (IEEE)
- Fixed Content (SNIA)

SNIA Disk Data Format (DDF)

- DDF - Data structures describing how data is distributed across the drives in a RAID implementation.
- Primary intended scope: RAID controllers for internal and direct attach storage.
- Does not standardize operating system/RAID controller interface or create a single driver



SNIA Shared Storage Model

IEEE P1619 Encrypted Media: Encrypt Stored Data at Rest

- **Storage Attack: Move encrypted blocks**
 - Attacker may know where data is stored
 - Disks: no room to store additional integrity checks
- **Encryption mode prevents block swapping**
 - Move encrypted blocks: Decryption produces random gibberish
 - New tweaked encryption mode: LRW-AES
- **Disk encryption and key backup format: close to done**
 - LRW-AES definition being revised (tweak function)
- **Tape encryption: work underway (P1619.1)**
 - Tapes have room to store additional integrity checks
 - Combined modes (encryption + cryptographic integrity):
 - AES-GCM (Galois Counter Mode) – hardware friendly
 - AES-CCM (Counter with CBC-MAC) – simpler alternative

XAM API Context and Goals

- ‘Fixed Content’ has a long shelf life
 - E.g. regulations routinely require 7-20 years
- Storage technology is transient
 - Content needs to be able to transcend physical, geographical, technological, organizational boundaries
- Content needs to be preserved in a self-describing way
 - Recognize “record” (data + metadata) as a storage type
- Enable “records” to “migrate” across storage devices
 - from different vendors, across different technologies
- Enable “records” to be shared among applications
 - E.g., content management and backup/restore
- Manage billions if not trillions of “records”

SNIA: XAM API status

- Initially proposed to SNIA in 2005
 - by EMC, IBM, HDS, HP and Sun
- Now under development in FCAS Technical WG
 - Significant design changes and extensions
- First version of standard target: end of 2006
 - Limit addition of functionality to maintain schedule
- Additional possible program elements
 - Reference implementation
 - Conformance test

Storage Networking Standards

Data Protocols

- SCSI (T10)
- Fibre Channel (T11)
- ATA and SATA (T13)
- NAS (IETF, Microsoft)
- IP Storage (IETF)

Management

- SMI-S (SNIA)
 - Uses CIM (DMTF)
- SNMP (IETF)
- Web (IETF, W3C)

Stored Data

- RAID Layout (SNIA)
- Encryption (IEEE)
- Fixed Content (SNIA)

Standards: End User Benefits

- Protect technology investment
- Ensure a base level of interoperability
- Provide choice among products
- Ensure continuing innovation
- Commonality leads to less training, simpler deployment

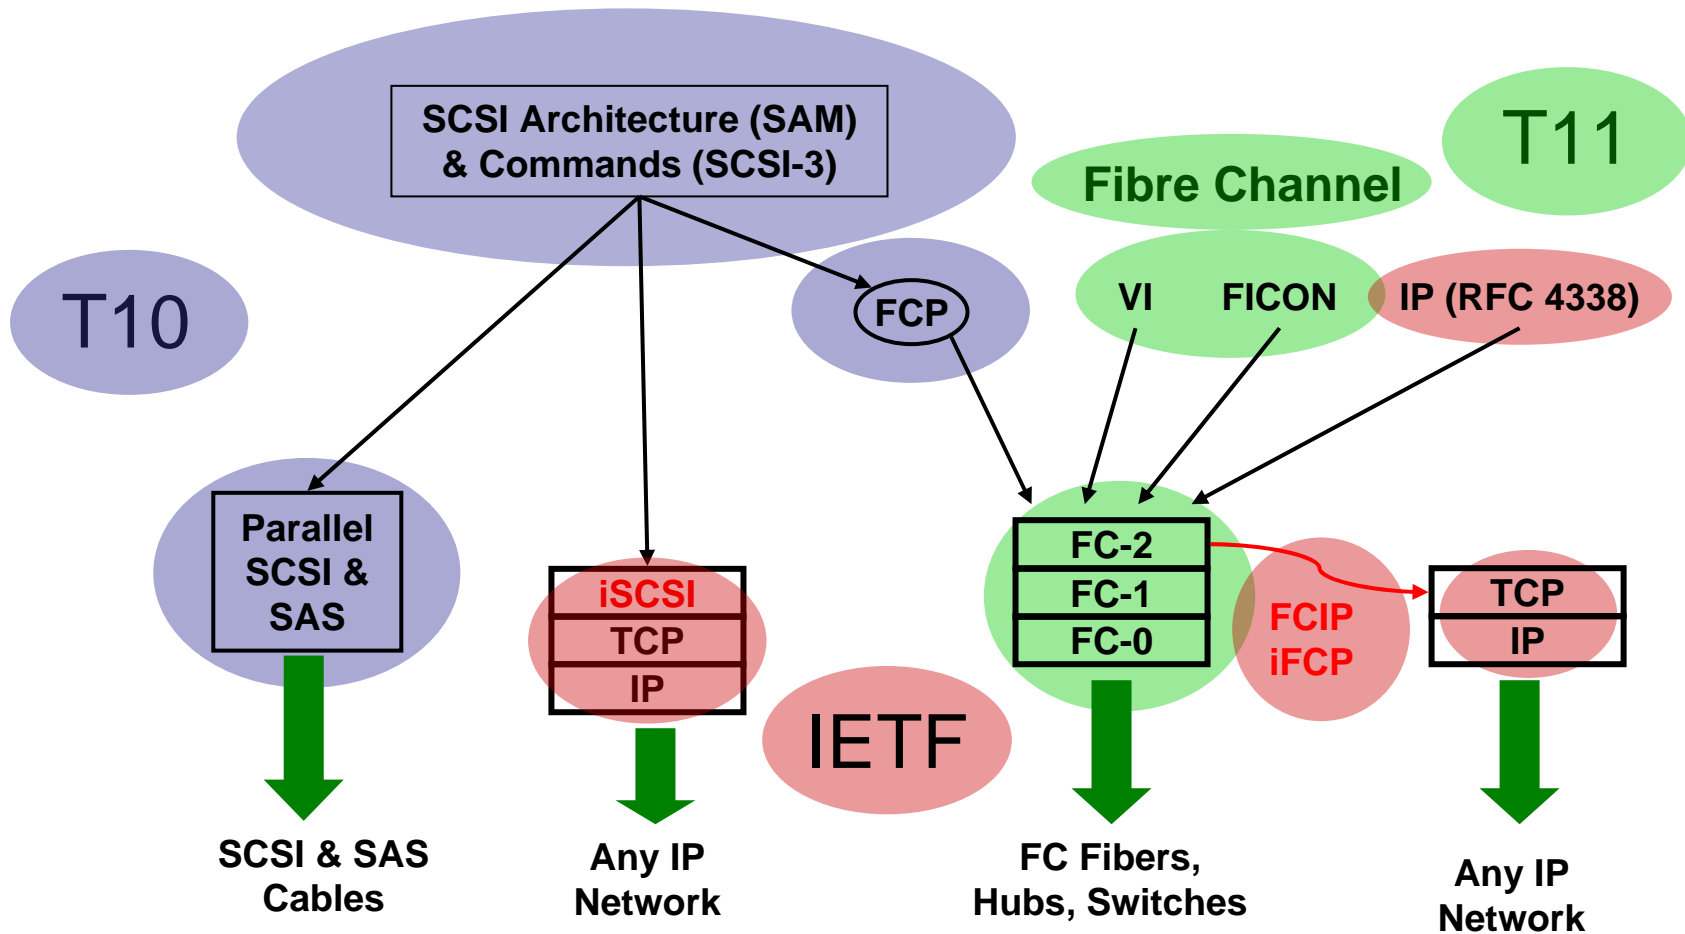
National and International Standards

- Standards Progression Path:
 1. Development in standards body (e.g., T11 for Fibre Channel),
 2. Becomes a national standard (e.g., ANSI)
 3. Becomes an international standard (ISO)
- INCITS: Umbrella Standards Organization
 - Umbrella for T10 (SCSI), T11(FC), T13 (ATA) and SNIA
 - Shepherds completed standards to ANSI and ISO
 - Usual path: Completed standard to INCITS to ANSI to ISO
- Not all standards follow this path
 - IETF RFCs are internationally recognized without ISO approval
 - Industry consortia standards (e.g., PCI)
 - Vendor de-facto standards (e.g., CIFS)

Getting Involved in Standards

- End users are always welcome
 - Remind participants why the standard matters
 - Help make the end result usable and useful
- Lots of opportunities to participate
 - Voting or contributing member
 - Observer: Still very important
 - No Substitute for hallway conversations

Block Storage Communication Protocols



Storage-Related Standards Organizations

- Storage Networking Industry Association (www.snia.org)
 - Storage Management (SMI-S) and other topics
- Distributed Management Task Force (www.dmtf.org)
 - Systems Management
- INCITS Technical Committees
 - SCSI: T10 (www.t10.org)
 - Fibre Channel: T11 (www.t11.org)
 - ATA: T13 (www.t13.org)
- IETF: Internet Engineering Task Force (www.ietf.org)
 - IP and Internet-related protocols, including IP Storage and NFS
- IEEE: Institute of Electrical and Electronics Engineers
 - Encrypted Media: P1619 (siswg.ieee.org)
- World Wide Web Consortium (www.w3c.org)
 - Web content representation (e.g., HTML, XML)

Q&A / Feedback

- Please send any questions or comments on this presentation to SNIA: trackvirtualization@snia.org and the SNIA Technical Council: snia-tc@snia.org

**Many thanks to the following individuals
for their contributions to this tutorial.**

SNIA Education Committee

**David L. Black
Mark Carlson
Zoran Cakeljic
Eric Hibbard**

**Steven Wilson
Bill Dawkins
SW Worth**