Everything You Wanted To Know About Storage (But Were Too Proud To Ask) The Basics

Today's Presenters

Bob Plumridge
HDS
Chief Technology Officer - EMEA

Alex McDonald
NetApp
CTO Office

© 2017 Storage Networking Industry Association. All Rights Reserved.

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.

© 2017 Storage Networking Industry Association. All Rights Reserved.

SNIA in Action

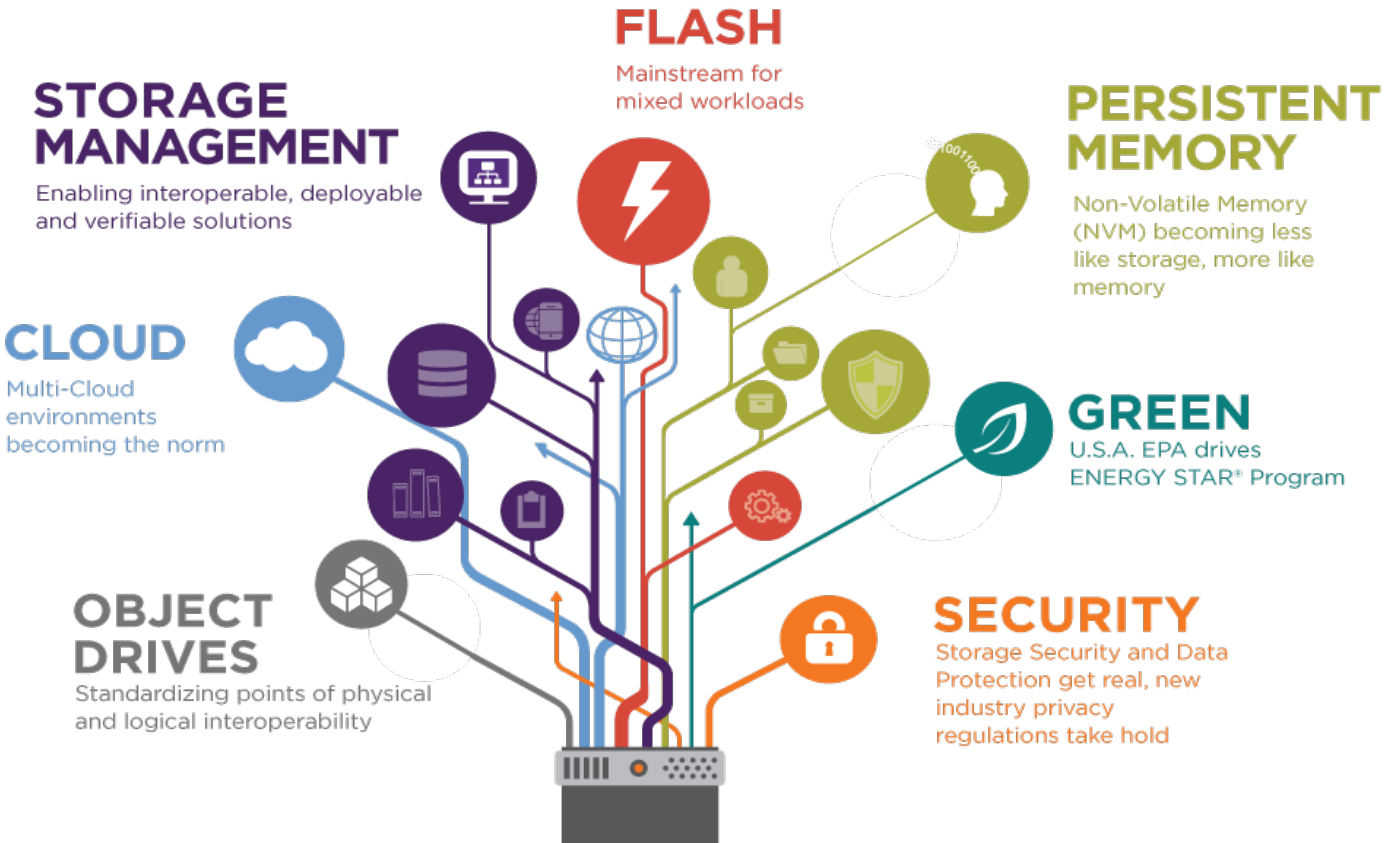

SNIA Global Presence

SNIA Europe™

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

© 2017 Storage Networking Industry Association. All Rights Reserved.

Agenda

- Naming of the Parts
  - **Initiators and Targets**
  - RAID and RAID Controllers
  - Storage Controllers
  - Volume Managers
  - Storage Stack
- Summary and Conclusion

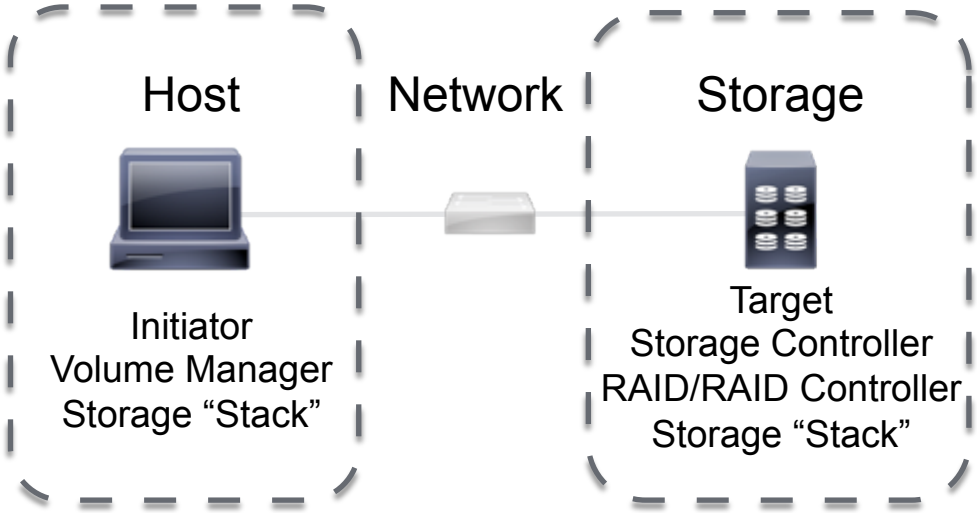

© 2017 Storage Networking Industry Association. All Rights Reserved.

Initiators & Targets

Definition

- Initiators- Starts the SCSI session, sends SCSI commands. Usually a client or server
- Target- Waits for commands sent from the initiator, then provides the I/O transfers. Usually a storage sub system

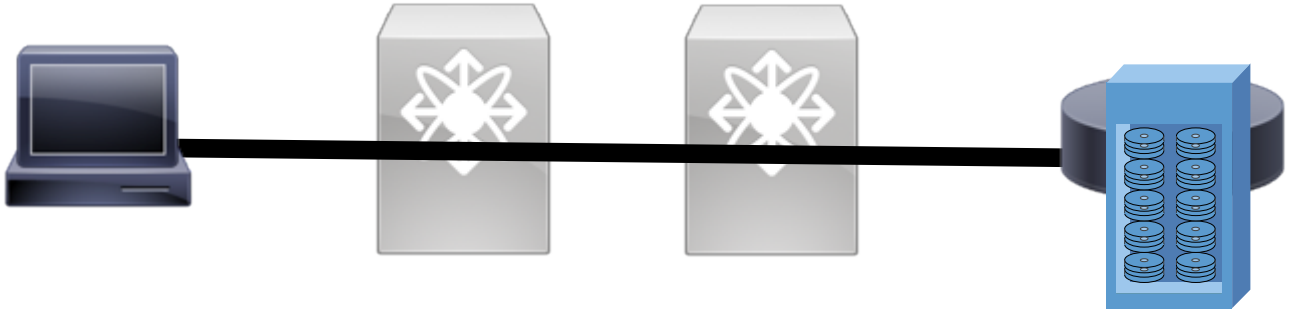

© 2017 Storage Networking Industry Association. All Rights Reserved.

SCSI Commands

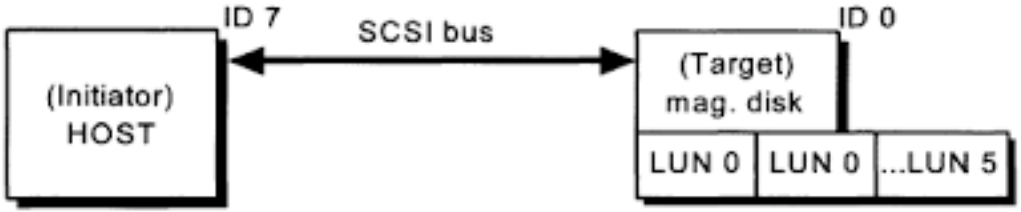

SCSI system: one Initiator, and one Target with six LUNs

- Initiator selects SCSI target and Sends Command Descriptor Blocks (CDB)-SCSI Commands
- If CDB contains READ command, the target retrieves data from disk and responds.
- If CDB contains WRITE command, the target prepares buffers and returns indication of ready. Initiator receives ready command, data is written.
  - When complete, SCSI responds as successful or unsuccessful.

© 2017 Storage Networking Industry Association. All Rights Reserved.

Agenda

- Naming of the Parts
  - Initiators and Targets
  - **RAID and RAID Controllers**
  - Storage Controllers
  - Volume Managers
  - Storage Stack
- Summary and Conclusion

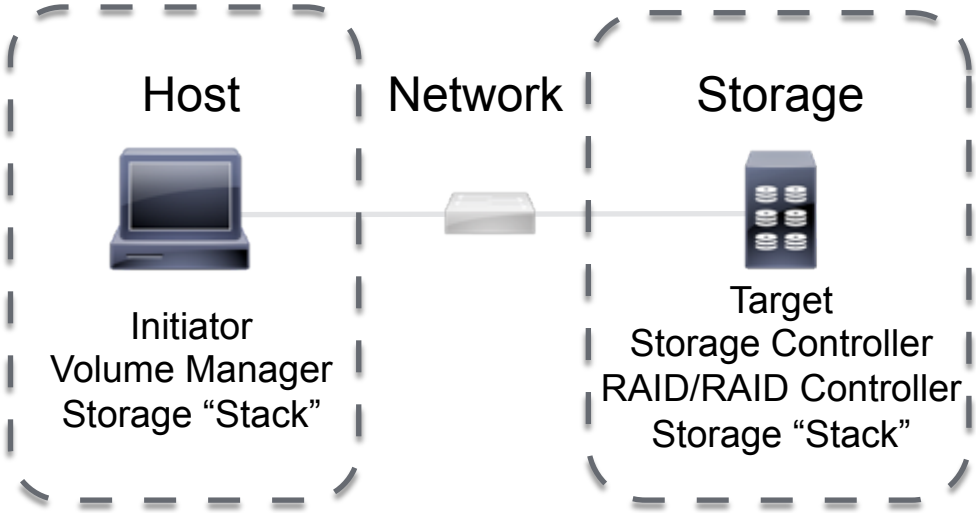

© 2017 Storage Networking Industry Association. All Rights Reserved.

RAID

Definition

- ”Redundant Array of Inexpensive Disks”*
- Ability to read and write to multiple disks as a single entity, increasing performance and reliability over a single, large expensive disk

*Patterson, Gibson & Katz 1988 http://www.cs.cmu.edu/~garth/RAIDpaper/Patterson88.pdf

© 2017 Storage Networking Industry Association. All Rights Reserved.

RAID History

- **Performance:** increase the # of targets for write IO, decreasing queuing and latency; normally does nothing for reads, since data only written to a single disk, but see RAID-0
- **Reliability:** Add in redundancy to provide superior error checking
- **Cost:** do so with cheap disks

© 2017 Storage Networking Industry Association. All Rights Reserved.

Types of RAID

- RAID still in use today, even (especially!) in Flash era
  - Can be done in hardware or software
- Various types of RAID
  - Each has an associated number; RAID-n
  - Most common are RAID-5, and RAID-6 (or RAID-DP)
  - Then there's **Erasure Coding** (EC)…
    - RAID levels standardized by SNIA: http://www.snia.org/tech_activities/standards/curr_standards/ddf

© 2017 Storage Networking Industry Association. All Rights Reserved.

Some RAID Levels

- RAID is k data and p parity (k,p)
- RAID-0 (k,0)
  - Alternating writes to two or more devices
  - No data protection
- For others, parity is an important concept
- RAID-5 (k,1)
  - Block-level striping with distributed parity
  - Parity is XOR across drives, tolerates 1 drive failure

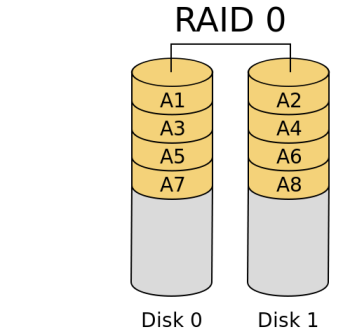

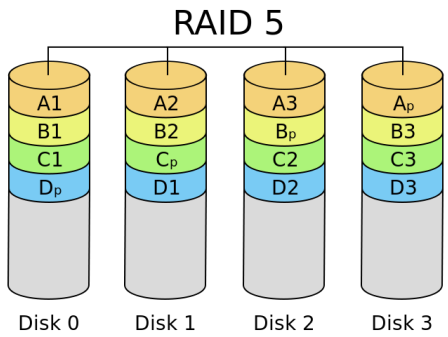

© 2017 Storage Networking Industry Association. All Rights Reserved.

Diagrams by en:User:Cburnett - Own work, this vector image was created with Inkscape., CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=1509075

Some RAID Levels

- RAID-1 (1,1)
  - Mirroring
  - Like RAID-5 with 1 data 1 parity; except parity is an exact copy
- RAID-6 (k,2)
  - Two parities calculated, tolerates 2 drive failures
- Combinations possible
  - Eg RAID-10 is RAID-1 and RAID-0
- Lots of wrinkles; performance, capacity and data protection are all compromises

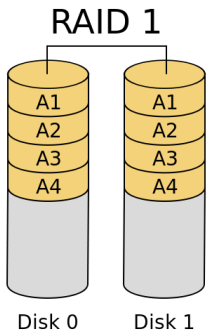

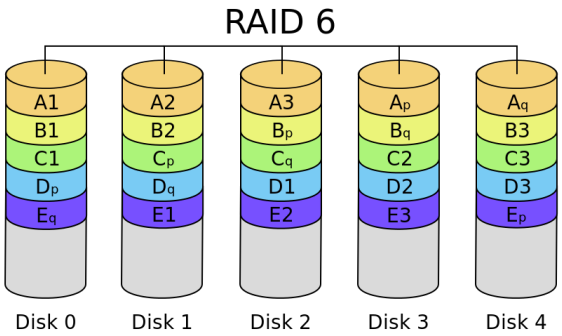

© 2017 Storage Networking Industry Association. All Rights Reserved.

Diagrams by en:User:Cburnett - Own work, this vector image was created with Inkscape., CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=1509075

RAID and Erasure Coding (EC)

- RAID is specialized case of Erasure Coding (EC)
  - You may have heard of Galois fields, Reed-Solomon encodings, Tornado or turbo codes; various related techniques for doing EC
  - Math is based on polynomials
  - Can be expensive to calculate, but powerful for distribution of data & protection; strips & stripes data
  - EC is k “data” and p “parity” (k,p)
- Any k out of the k+p disks are sufficient to recover
  - Or, put another way, we can lose p disks
  - RAID-5 is (k,1), RAID-6 is (k,2)
  - EC (13,5) allows 5 disks to be lost out of 18
- EC generally used in object stores, and for geographically or node based dispersed storage
  - RAIN – Redundant Array of Inexpensive Nodes
- RAID & EC is now about data protection
  - Performance a given (& networks now the bottleneck)
  - Still required for Flash (if not more so)

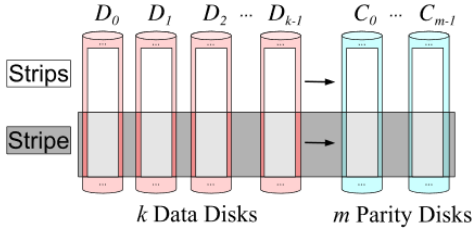

Figure 1. A typical storage system with erasure coding.

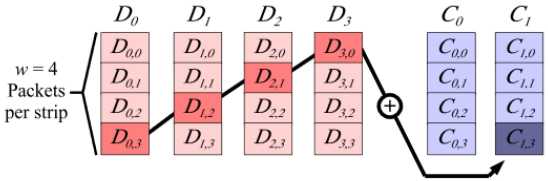

Figure 2. An example of one stripe where $k = 4$, $m = 2$ and $w = 4$.

Credit: http://nisl.wayne.edu/Papers/Tech/dsn-2009.pdf

© 2017 Storage Networking Industry Association. All Rights Reserved.

Agenda

- Naming of the Parts
  - Initiators and Targets
  - RAID and RAID Controllers
  - **Storage Controllers**
  - Volume Managers
  - Storage Stack
- Summary and Conclusion

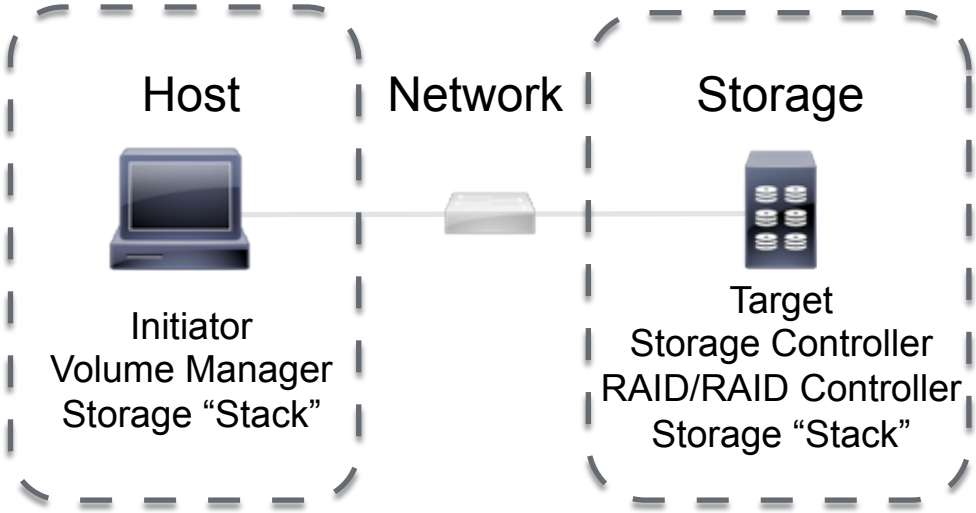

© 2017 Storage Networking Industry Association. All Rights Reserved.

What is a “Controller”?

- Definition (from https://www.snia.org/education/dictionary)
  - SNIA (1): The control logic in a disk or tape that performs command decoding and execution, host data transfer, serialization and deserialization of data, error detection and correction, and overall management of device operations
  - SNIA (2): A disk array whose control software executes in a disk subsystem controller
  - SNIA (3): An integrated collection of (a.) storage controllers and/or host bus adapters, (b.) storage devices, CD-ROM drives, tape drives, and libraries, and (c.) any required control software, that provides storage services to one or more computers

© 2017 Storage Networking Industry Association. All Rights Reserved.

Types of Controllers

- Controller is a confusing historical term. Often used synonymous with adapter. Used as a “last” name - as in - video controller, network controller, keyboard controller, USB controller, storage controller, RAID controller.

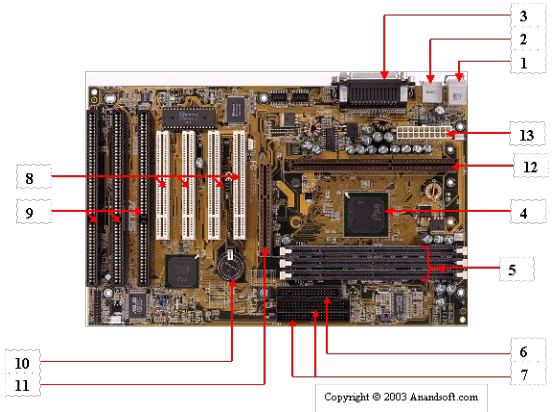

© 2017 Storage Networking Industry Association. All Rights Reserved.

Okay, so what's a *storage* controller?

- But we're not talking about those kinds of controllers.
- Storage Controllers or Array Controllers are devices that connect between a Host and a large number of physical devices.
- They **virtualize** the physical devices, while providing advanced storage features

© 2017 Storage Networking Industry Association. All Rights Reserved.

What does it look like?

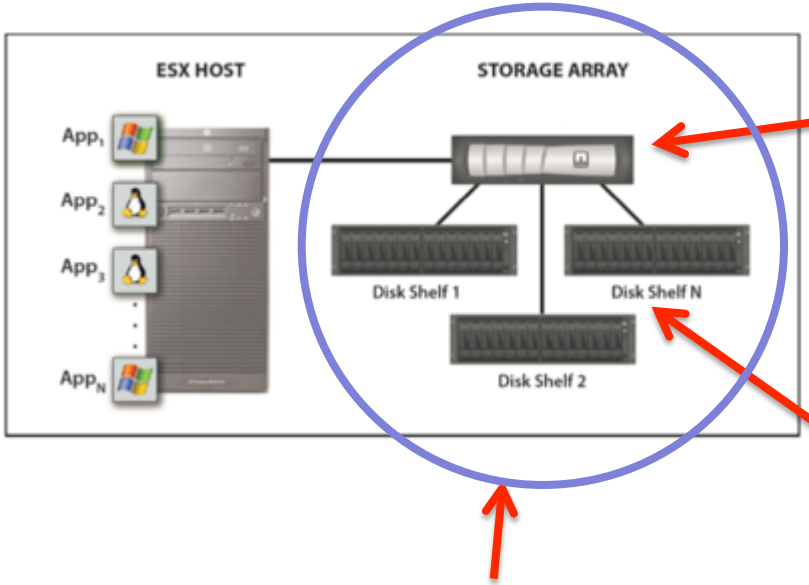

Controllers

Physical Storage Devices

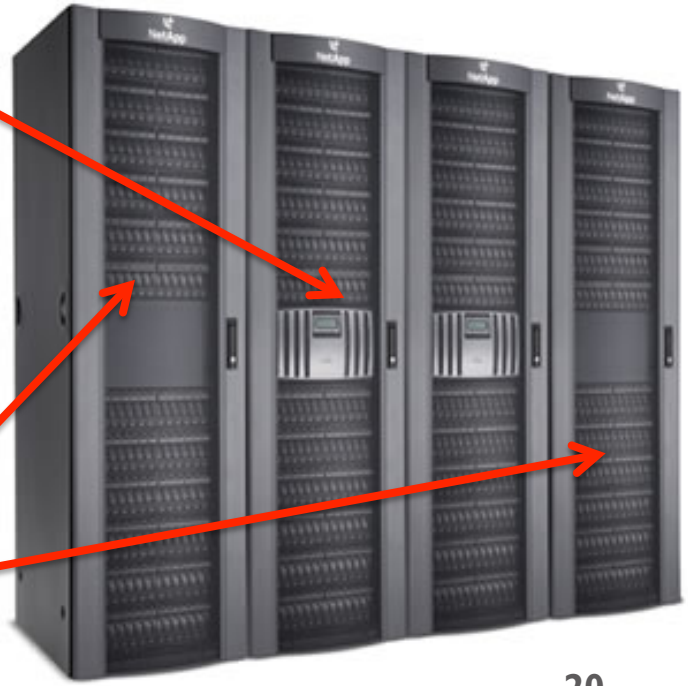

© 2017 Storage Networking Industry Association. All Rights Reserved.

Storage Controller Features

- May include things like:
  - Management
  - RAID, High Availability/Fault Tolerance,
    - Multipathing (aka SCSI: ALUA)
  - Backups (snapshots/clones),
  - Thin Provisioning, Compression, de-duplication,
  - Long distance replication (disaster recovery)
- Controllers manage trade-offs between capabilities, performance, and scale

© 2017 Storage Networking Industry Association. All Rights Reserved.

Where Do Controllers Fit?

- They contain front end ports that listen to hosts
  - These are the target ports we just heard about
- They contain back end ports that talk to devices
  - These are the initiator ports we just heard about
- Controlling negotiations
  - Have processors (CPUs) and an OS (of some kind, usually proprietary) that manages the requests from the frontend and turns them into requests to the backend

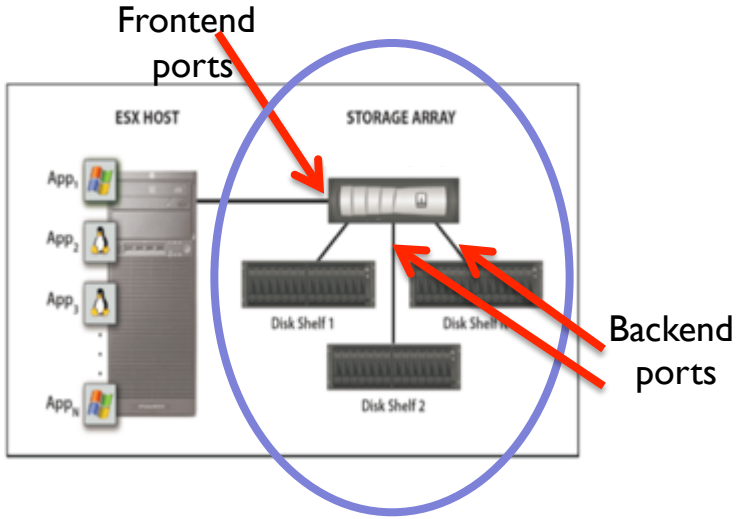

© 2017 Storage Networking Industry Association. All Rights Reserved.

All About the Context

- Context is important – is it the whole array (controller + storage) or is it just the controller or is it the adapter card stuck into a PCI slot?
- Storage Array (usually the whole thing)
- Array Controller (usually the controller in the array)
- Storage Controller (usually the controller in the array – or a simple PCI adapter card)

© 2017 Storage Networking Industry Association. All Rights Reserved.

Agenda

- Naming of the Parts
  - Initiators and Targets
  - RAID and RAID Controllers
  - Storage Controllers
  - **Volume Managers**
  - Storage Stack
- Summary and Conclusion

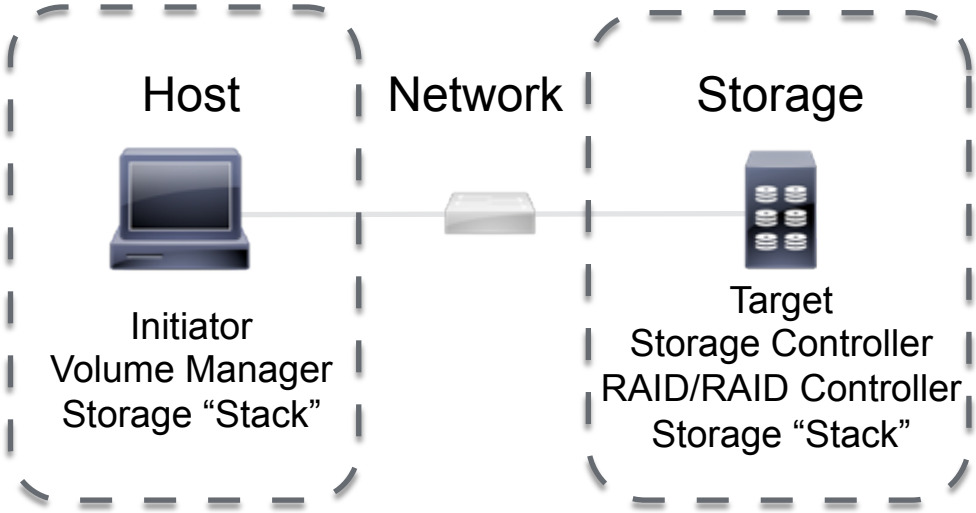

© 2017 Storage Networking Industry Association. All Rights Reserved.

Logical Volume Managers

SNIA Europe™

- **Definition (SNIA)**
  - A body of software that provides common control and management for one or more disk arrays or tape arrays.
  - Presents the arrays of disks or tapes it controls to its operating environment as one or more virtual disks or tapes. It may execute in a disk controller or intelligent host bus adapter, or in a host computer. When it executes in a disk controller or adapter, it is often referred to as firmware.
- **Definition (Whatis.com)**
  - A volume manager is software within an operating system (OS) that controls capacity allocation for storage arrays. Most operating systems provide an option for managing volumes.

© 2017 Storage Networking Industry Association. All Rights Reserved.

Brief History

- Devices had small capacity (apps wanted large)
- Applications didn’t want to break up their data
- Volume managers made small disks into large disks
- Applications wanted extra resiliency
- Volume managers added RAID to existing devices
- But now, they are everywhere

© 2017 Storage Networking Industry Association. All Rights Reserved.

So what does it do?

- In its most basic form, a Volume Manager virtualizes storage to the application.

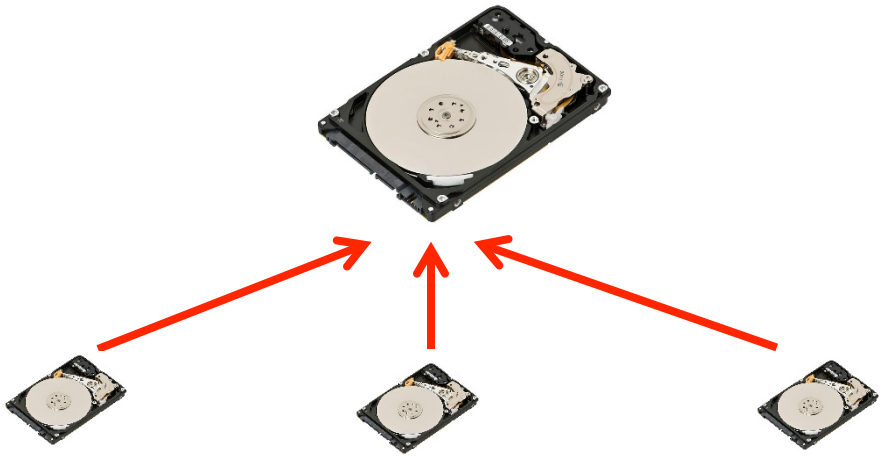

Three small disks are made to look like one large one

© 2017 Storage Networking Industry Association. All Rights Reserved.

Image Evan-Amos (Own work) [CC BY-SA 3.0 (http://creativecommons.org/licenses/by-sa/3.0) or GFDL (http://www.gnu.org/copyleft/fdl.html)], via Wikimedia Commons

So what does it do?

- In its most basic form, a Volume Manager virtualizes storage to the application.

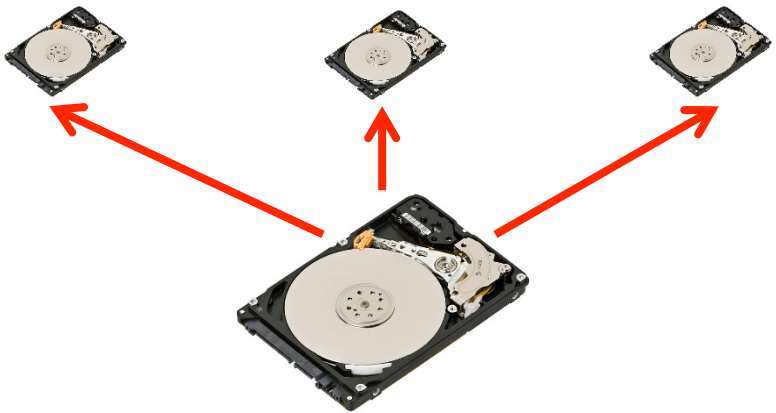

One large disk is made to look like three small ones – such as 3 VMs sharing the same disk

© 2017 Storage Networking Industry Association. All Rights Reserved.

So what does it do?

- In its most basic form, a Volume Manager virtualizes storage to the application.

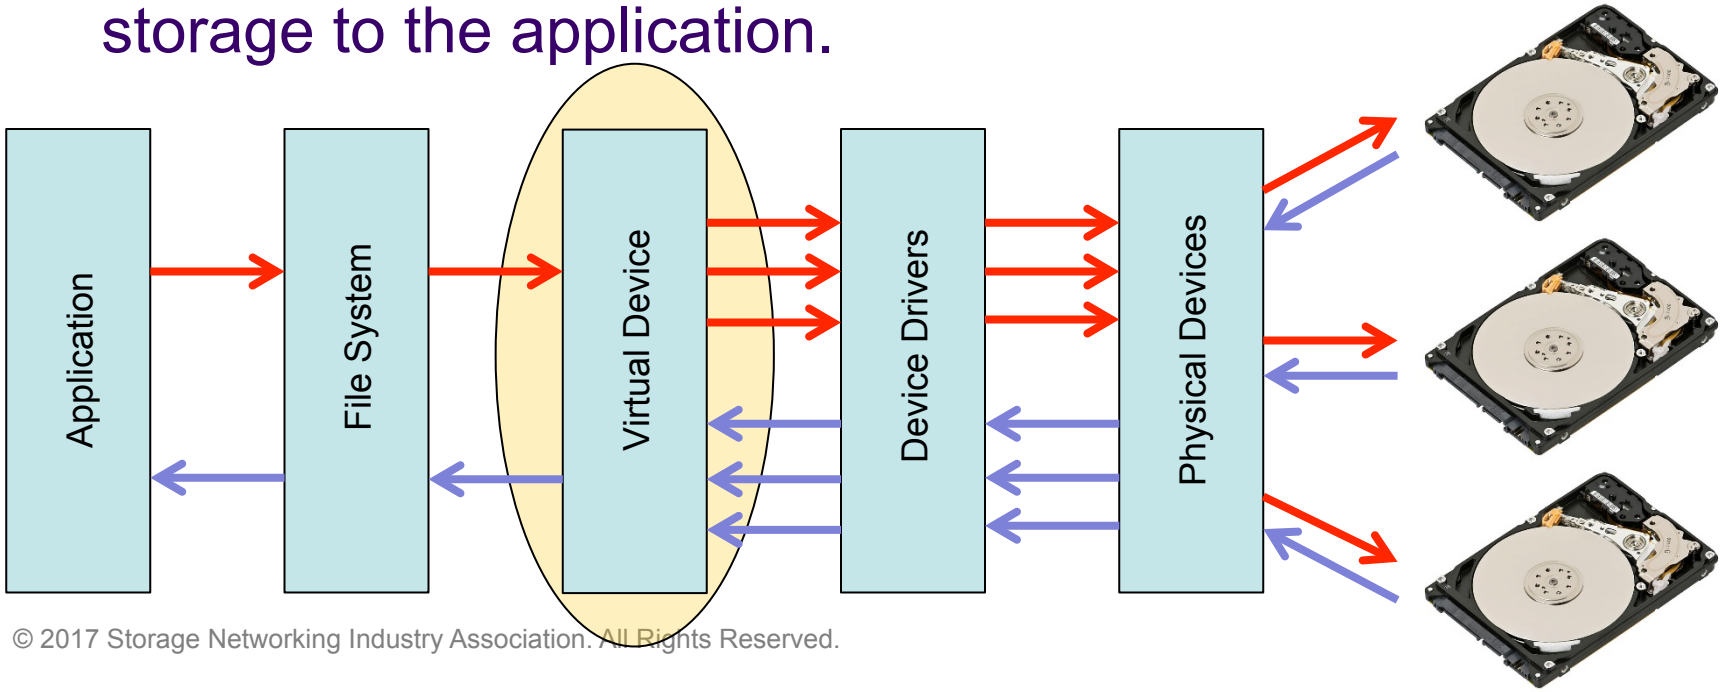

© 2017 Storage Networking Industry Association. All Rights Reserved.

So what does it do?

- In its most basic form, a Volume Manager virtualizes storage to the application.

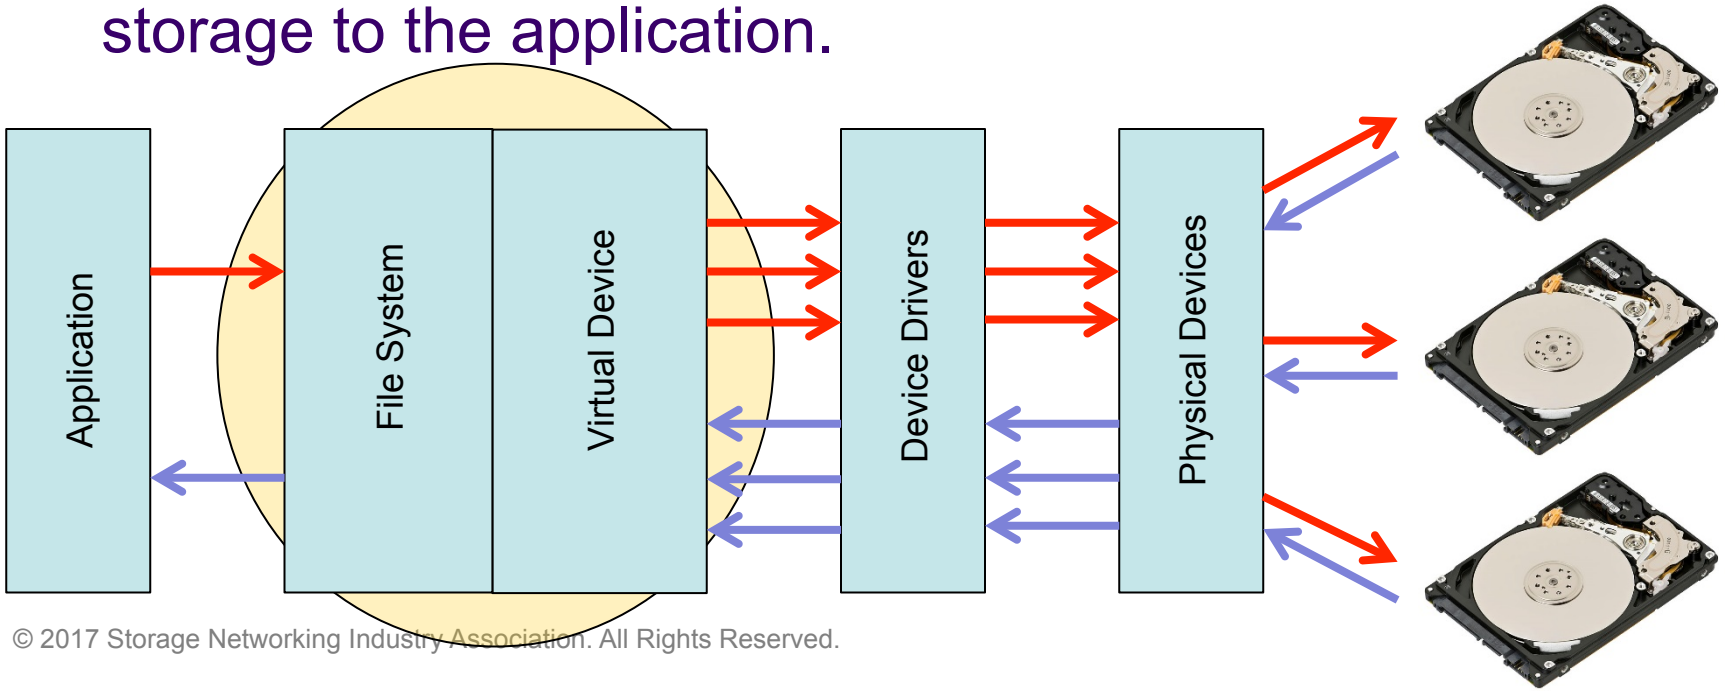

© 2017 Storage Networking Industry Association. All Rights Reserved.

How does it add resilience?

- In its most basic form, a Volume Manager virtualizes storage to the application.

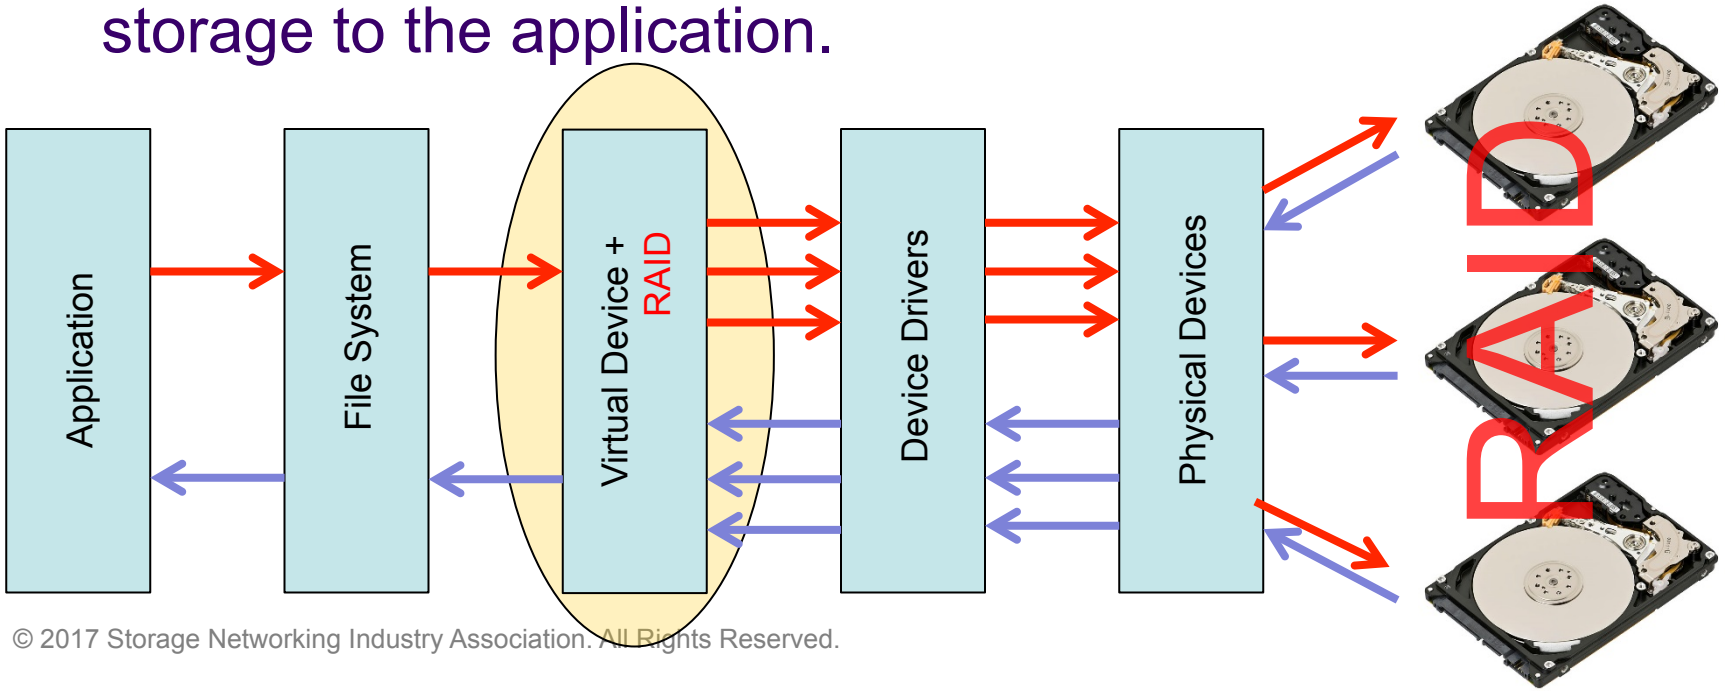

© 2017 Storage Networking Industry Association. All Rights Reserved.

How it works – one example

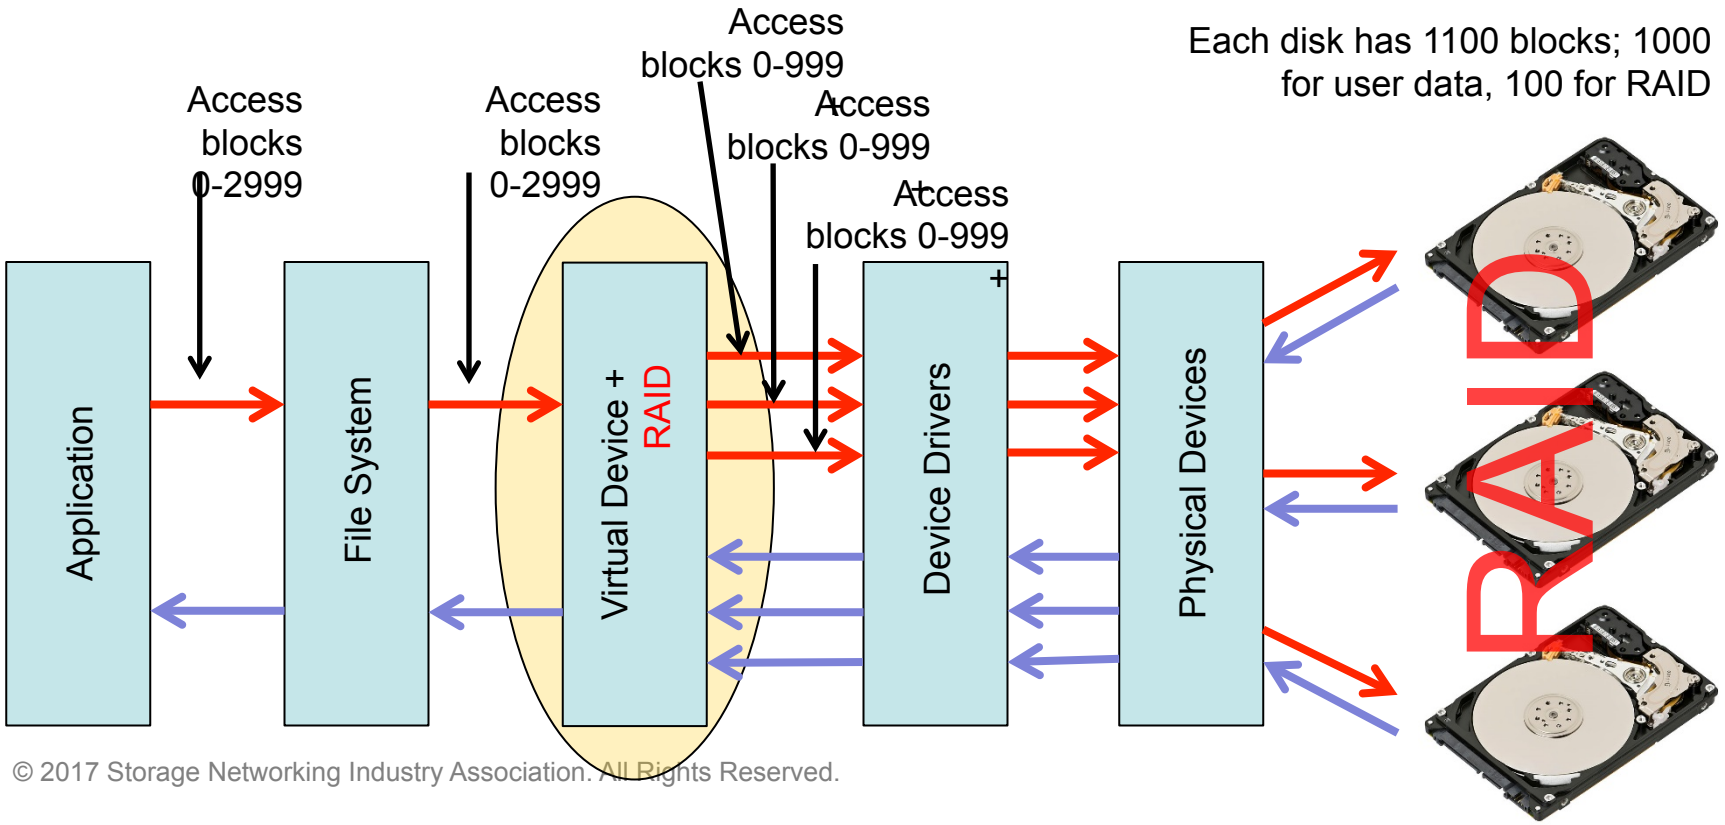

© 2017 Storage Networking Industry Association. All Rights Reserved.

Where are volume managers

- In host O/S software
- In Storage Arrays
- In the Cloud
- Virtualizing physical storage into a form that more directly meets the requirements of a particular application

© 2017 Storage Networking Industry Association. All Rights Reserved.

Agenda

- Naming of the Parts
  - Initiators and Targets
  - RAID and RAID Controllers
  - Storage Controllers
  - Volume Managers
  - **Storage Stack**
- Summary and Conclusion

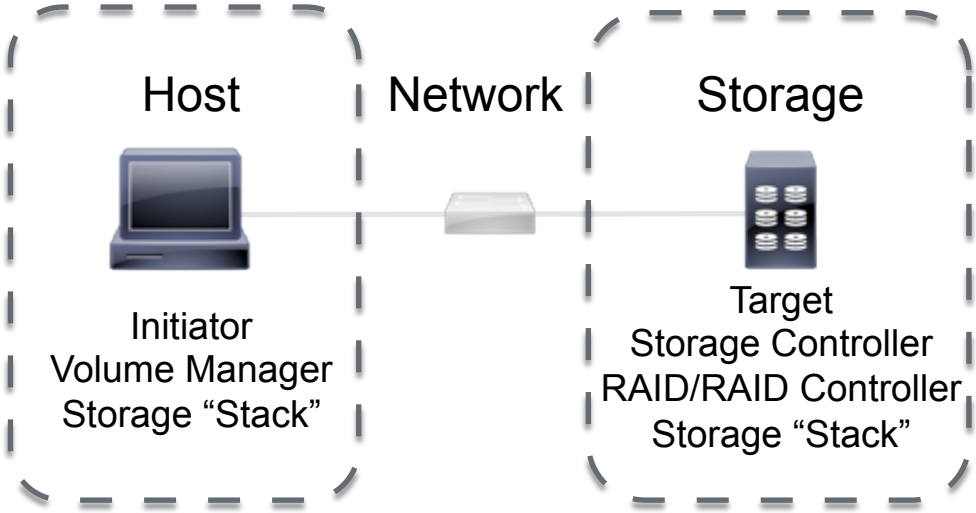

© 2017 Storage Networking Industry Association. All Rights Reserved.

Storage Stacks

SNIA Europe™

Two Definitions

- **Solution**: “Layers of software and hardware which interoperate to deliver storage functionality to a user or application.”
- **Software**: “Set of layered features, protocols, and services in the storage controller and/or initiator.”

© 2017 Storage Networking Industry Association. All Rights Reserved.

Solution vs. Software Stack

Solution Stack

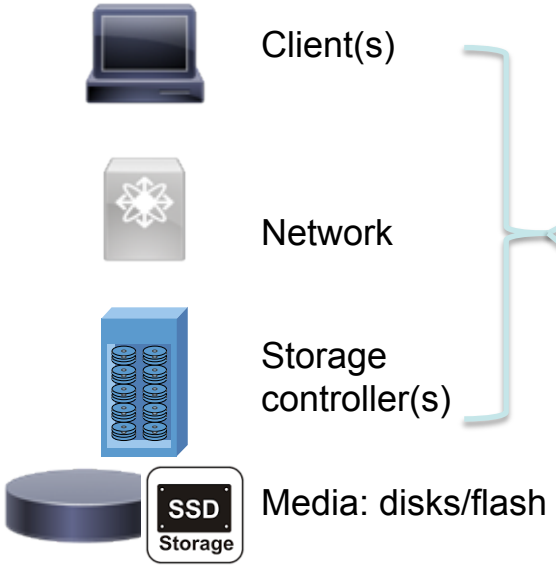

Software Stack

- Client Software or Agent (optional)
- Storage Protocols
- Other Storage Features & App/hypervisor Integration
- RAID & Volume Management
- Hardware and Media Management

© 2016 Storage Networking Industry Association. All Rights Reserved.

Role of the Storage Stack

- Translates between user and media
  - From user/app requests to disk commands
  - From disk output to user/app-requested data

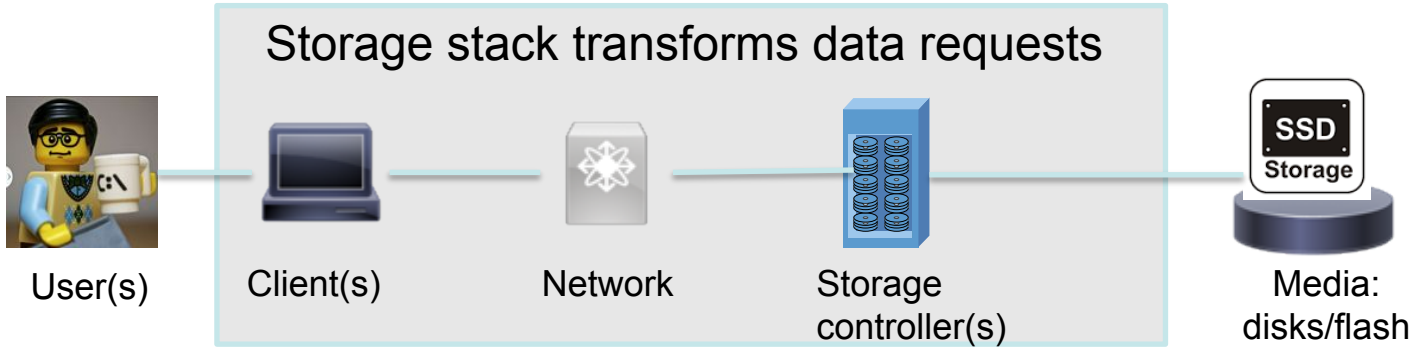

© 2016 Storage Networking Industry Association. All Rights Reserved.

Software Stack Hierarchy

- Management
- Integration
- Protocols
- Security
- Data Protection
- Efficiency
- Virtualization
- Availability
- Physical

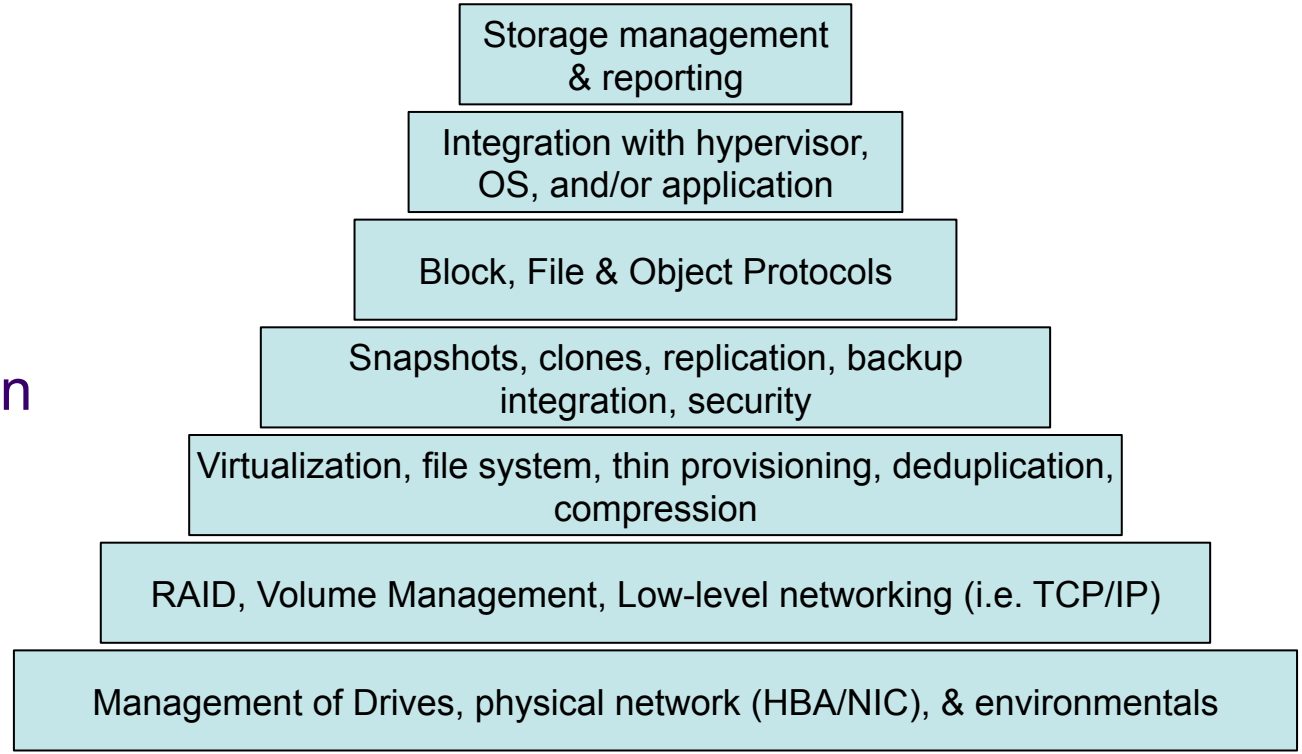

© 2017 Storage Networking Industry Association. All Rights Reserved.

Agenda

- Naming of the Parts
  - Initiators and Targets
  - RAID and RAID Controllers
  - Storage Controllers
  - Volume Managers
  - Storage Stack
- **Summary and Conclusion**

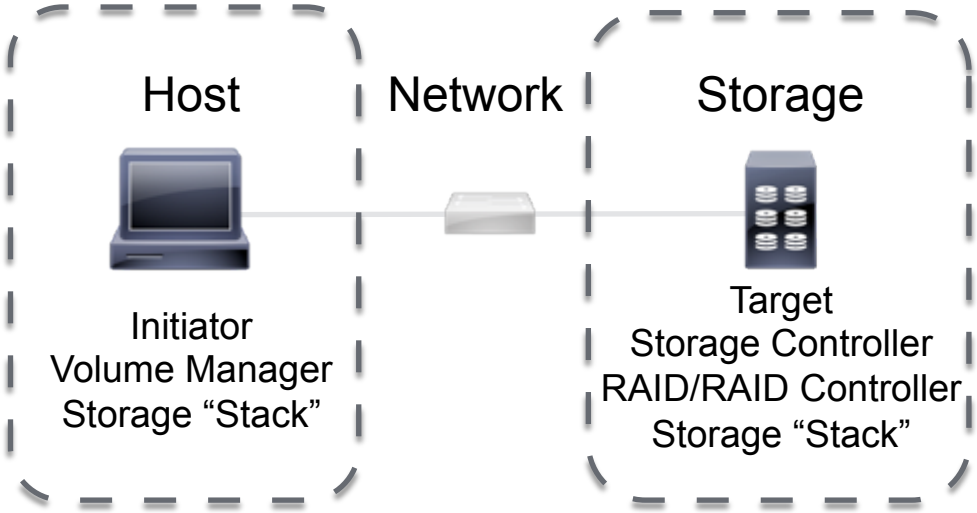

© 2017 Storage Networking Industry Association. All Rights Reserved.

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 Europe website (https://www.snia.org/Europe) and will be available in an on-demand format
- A full Q&A from this webcast, including answers to questions we couldn't get to today, will be posted to the same website
- Follow us on Twitter @SNIA_Europe

© 2017 Storage Networking Industry Association. All Rights Reserved.