SNIA ESF | ETHERNET STORAGE

EVERYTHING YOU WANTED TO KNOW ABOUT STORAGE BUT WERE TOO PROUD TO ASK

Part Aqua

Storage Controllers

May 15, 2018
10:00 am PT

Today's Presenters

SNIA ESF | ETHERNET STORAGE

J Metz
Cisco

Craig Carlson
Cavium

John Kim
Mellanox

Peter Onufryk
Microsemi

Chad Hintz
Cisco

SNIA-At-A-Glance

SNIA ESF | ETHERNET STORAGE

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** **@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.

Controllers!

- So many things to control, so little time!

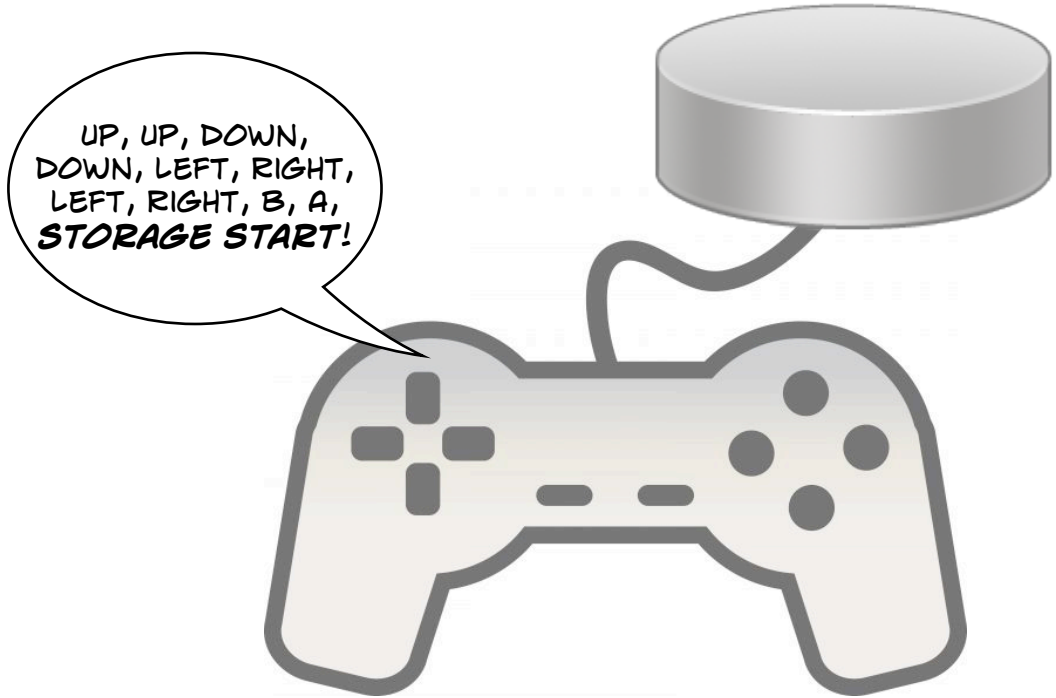

Why Controllers?

- When we talk about “Controllers,” what do we mean?
  - Do DevOps people talk about *controllers* the same way networking people do? Storage people? Is a *domain controller* the same thing as a storage “fault domain”?
- What, exactly, do “Controllers” *control?*
  - Hardware? Software? Media? Logic? Direct-Attached or Network-Attached?
- What happens when you have multiple controllers? Are they required?
  - Are I/O controllers the same thing as disk controllers? If not, do you need both?

Goal of This Webinar

- Introduce the concept of storage controllers
- Illustrate how controllers work in various contexts
- Provide examples of different types of storage controllers
  - Hardware, software, logical, media, protocol
- Clarify some of the aspects of how controllers work in their respective contexts

Agenda

- Introduction
- Storage Controllers 101
- SCSI Controllers
- Fibre Channel Controllers
- NVMe Controllers
- Networking SDN Controllers and Storage SDS Controllers
- Summary

SNIA ESF | ETHERNET STORAGE

Storage Controllers 101

Craig Carlson

What is a Storage Controller

- Controller - A logical or physical entity that manages or directs the flow of data between two entities
- A Storage is controller is the device that converts between OS storage commands on the system bus to the storage device
- A controller also resides on the storage device to manage storage media

Storage Controller Functions

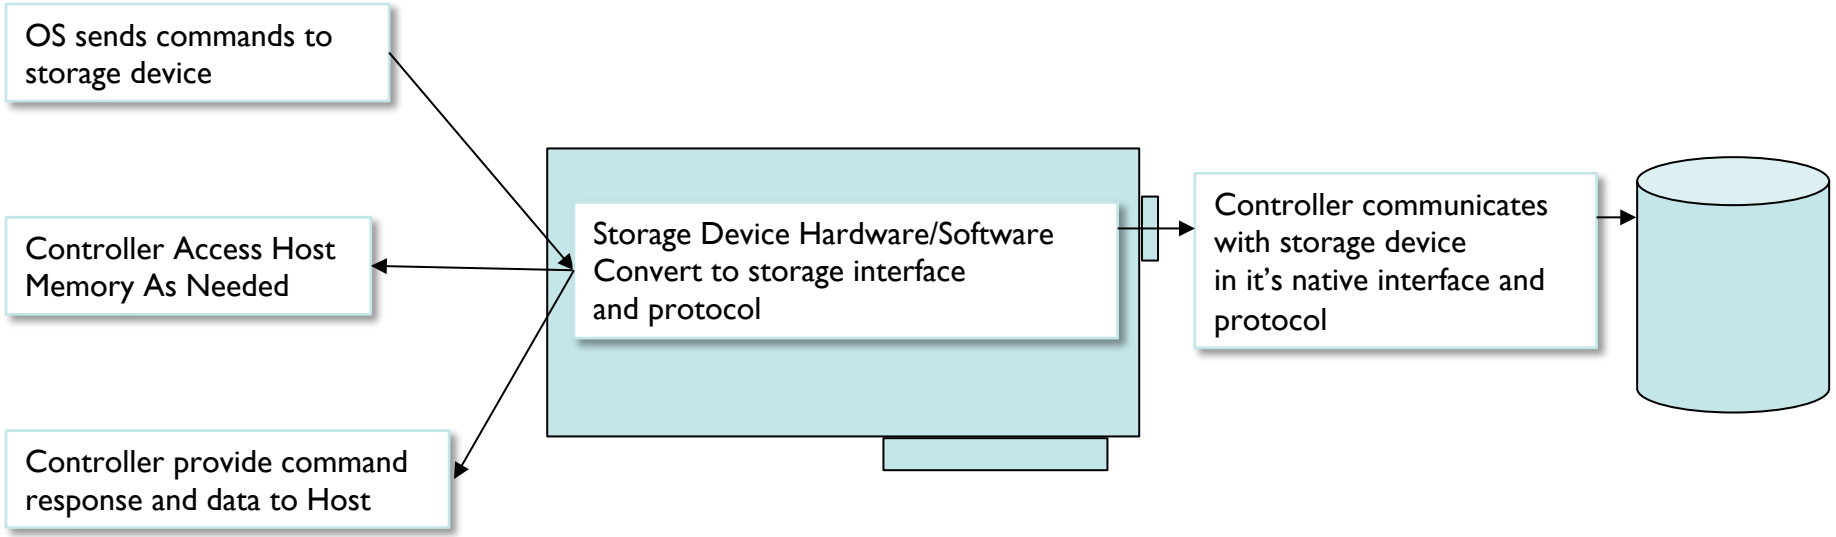

Front-end vs. Back-end controller

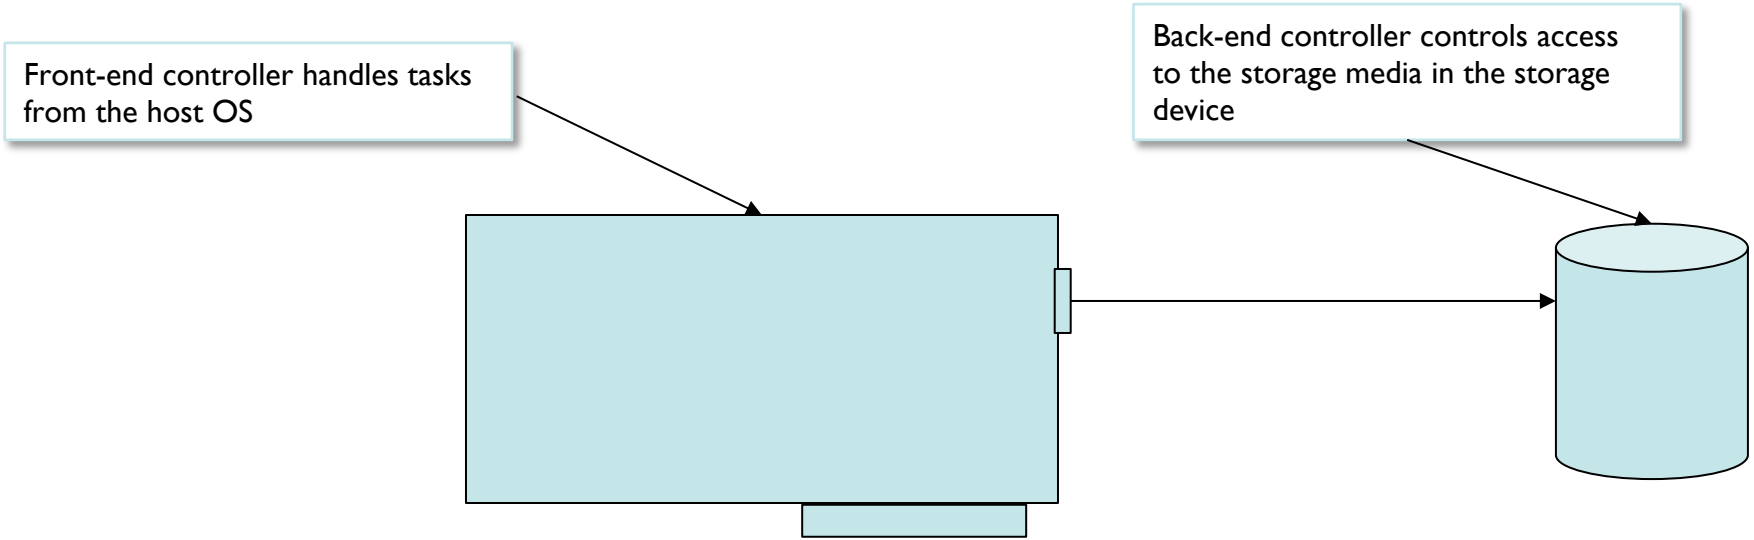

Memory Access

- Controllers typically access host memory through a Scatter Gather List (SGL)
  - For data writes, the SGL tells the controller the memory to read from
  - For data reads, the SGL tells the control the memory to write the received data to
  - Performed by Direct Memory Access (DMA) to/from card

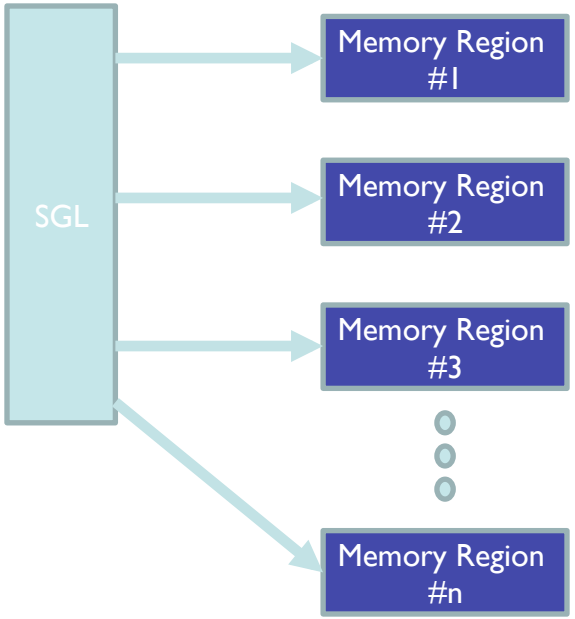

Hardware Offload

- Many times, controllers use Hardware Offload to process commands
  - Provides a hardware accelerated path for routine command operations
    - Non-routine events such as error processing still done in software or firmware
  - Advantages include:
    - Lower latency for command processing
    - Potential lower power usage
    - May be able to offload processing from the host CPU

Native Protocol

SNIA ESF | ETHERNET STORAGE

- Controller processes OS commands to the Native Storage Protocol, these could be one of:
  - SCSI
  - NVMe
  - FICON
  - SATA
  - NFS

Native Interface

Controller sends commands on the native interface used for a particular storage system, these could be

- Fibre Channel
- SAS
- Ethernet
- SATA
- And many others

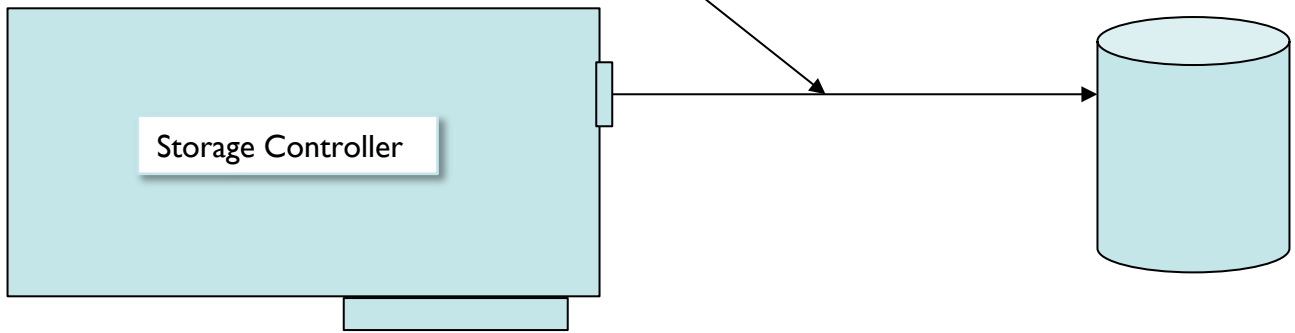

Key Takeaways

- “Controller” can mean different things depending upon context
- Storage controllers can be an interface between devices, can manage physical devices and media, and/or handle I/O
- Controllers may support one or more storage protocols

SNIA ESF | ETHERNET STORAGE

SCSI Controllers

John F. Kim

SCSI Controller

- Adapter that manages connections to SCSI devices
  - Connect to hard/flash/optical drives, scanners, etc.
  - E.g., SAS host bus adapter (HBA)
  - Chip on motherboard or add-in PCIe card
- Or… controller within the SCSI device
  - Responds to SCSI commands
  - Talks with SCSI controller adapter
- Network adapters can carry SCSI commands
  - But are not themselves SCSI controllers

SCSI Controllers: Adapter and Device

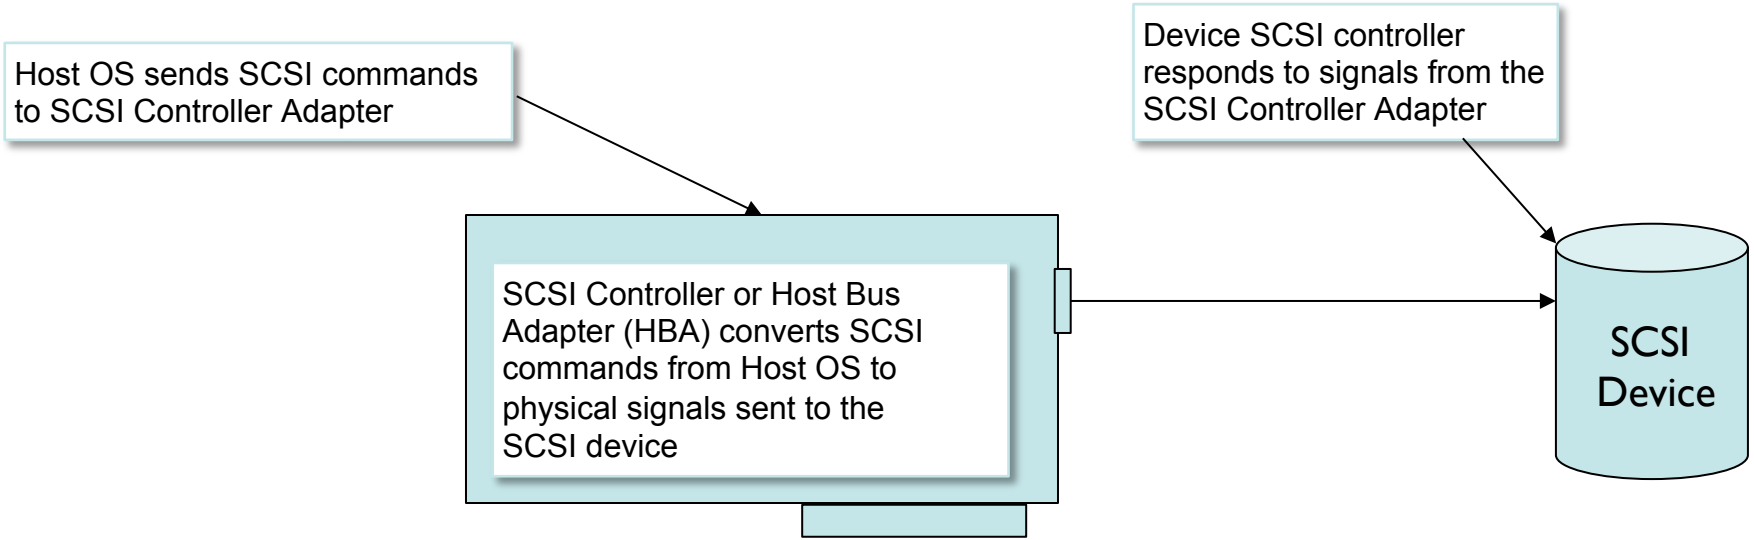

Defining SCSI

- SCSI = Small Computer Systems Interface
  - Standardized command set
  - Different physical connections
- Physical Connection Options
  - Parallel SCSI (1981)
  - Fibre Channel (1994), iSCSI (2002)
  - Serial Attached SCSI – SAS (2002)
  - Other: SRP, iSER, SCSI on USB

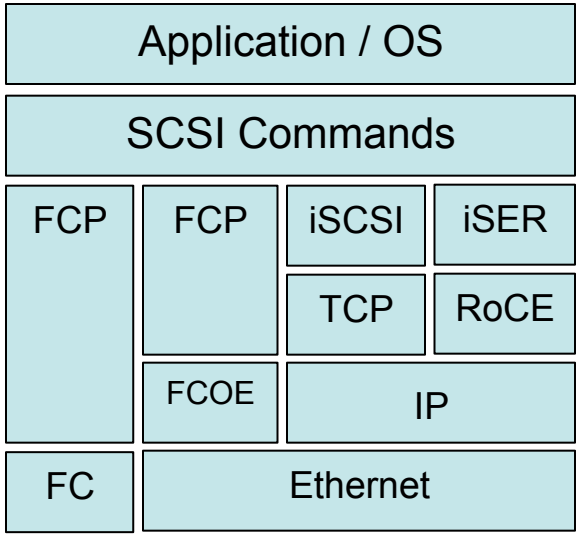

Common Networked SCSI Options

SCSI Controller Use Case

SNIA ESF | ETHERNET STORAGE

- Local SCSI
- Networked SCSI

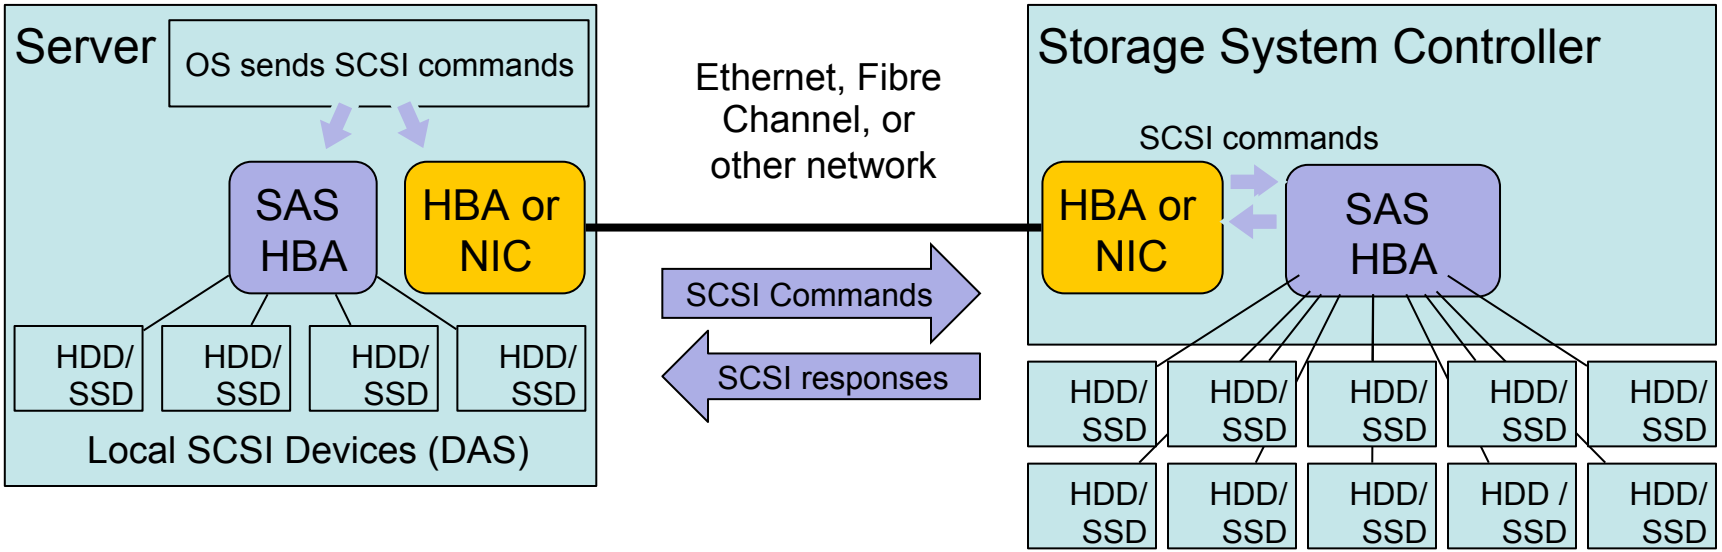

SCSI Controller Adapter Examples

- PCIe card
  - Or chip on motherboard
- SAS or SATA connectivity
  - Some models do RAID
  - Some also support NVMe
  - SAS expander allows more than 4 devices per connector

Adaptec (Microsemi) SCSI controller

Marvel SAS/SATA controller chip

Broadcom SAS/ SATA HBA

Adaptec (Microsemi) SAS Expander

Key Takeaways

- "SCSI Controller" means the SCSI adapter (HBA)
- Can send SCSI commands over different connections
  - Fibre Channel, Ethernet, SAS, InfiniBand, USB, etc.
- Local SCSI (DAS) uses SCSI devices
- Networked SCSI storage might not use SCSI devices
  - Storage controller can translate SCSI commands

SNIA ESF | ETHERNET STORAGE

FC Storage Controllers

Craig Carlson

Tasks of an FC Controller

- Provide Discovery of devices attached to FC Fabric
  - Accessed through the FC Name Server
- Provide access to the FC Management Server
  - Allows access to zoning configuration and other management tasks
- Process commands and provide HW offload
  - Many FC controllers provide HW offload
- Provide interface to the FC Physical transport
  - 64GFC or other speeds

FC Protocol stack

- Multiple protocols supported

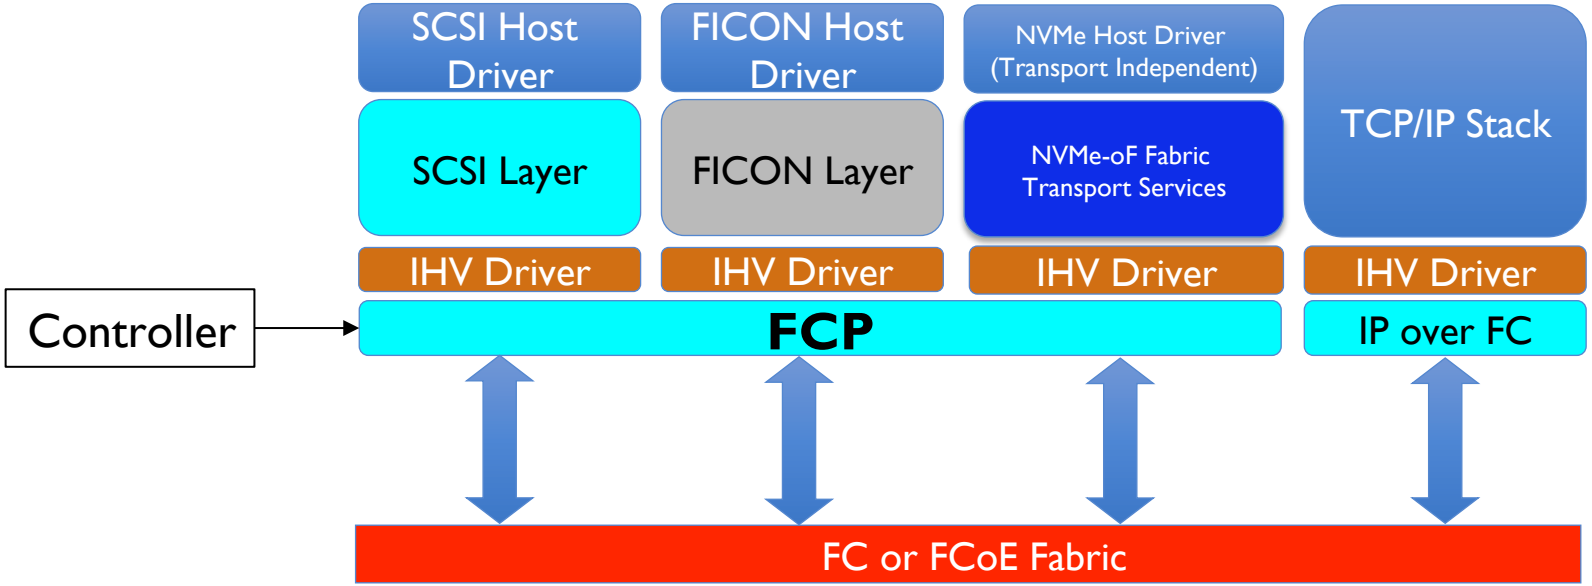

IHV = Independent Hardware Vendor

NVMe over Fabrics

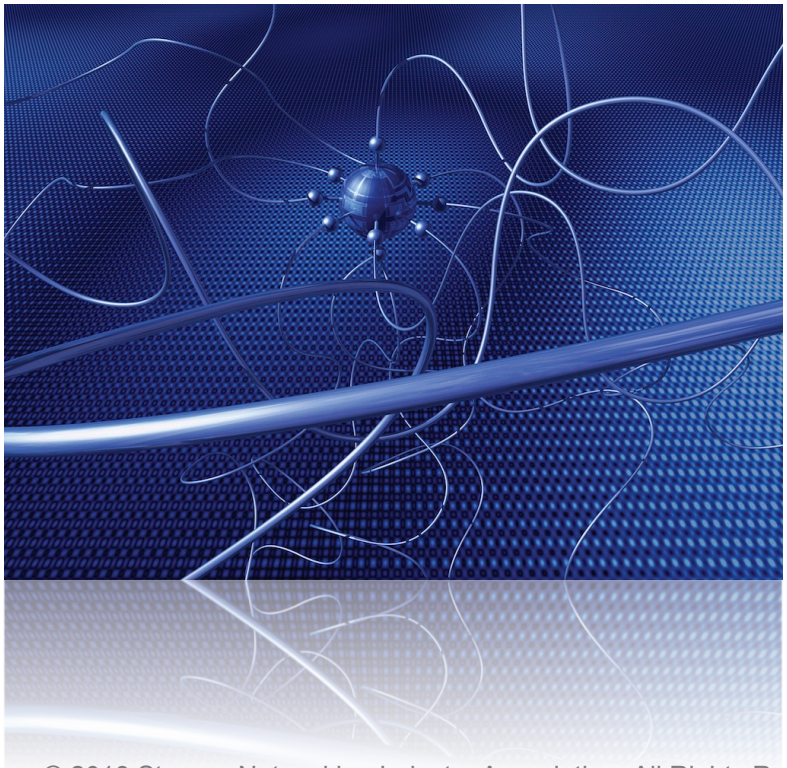

- Most recently supported protocol – NVMe over Fabrics
  - Defined by the FC-NVMe project
    - First version completed end of 2016
  - Now working on second version – FC-NVMe-2
    - Main item is to define Enhanced Error Recovery

Traditional Protocols

- Long time storage protocol support for
  - SCSI
    - Supported since 1994
  - FICON
    - Supported since 1996

FC Controller Tasks

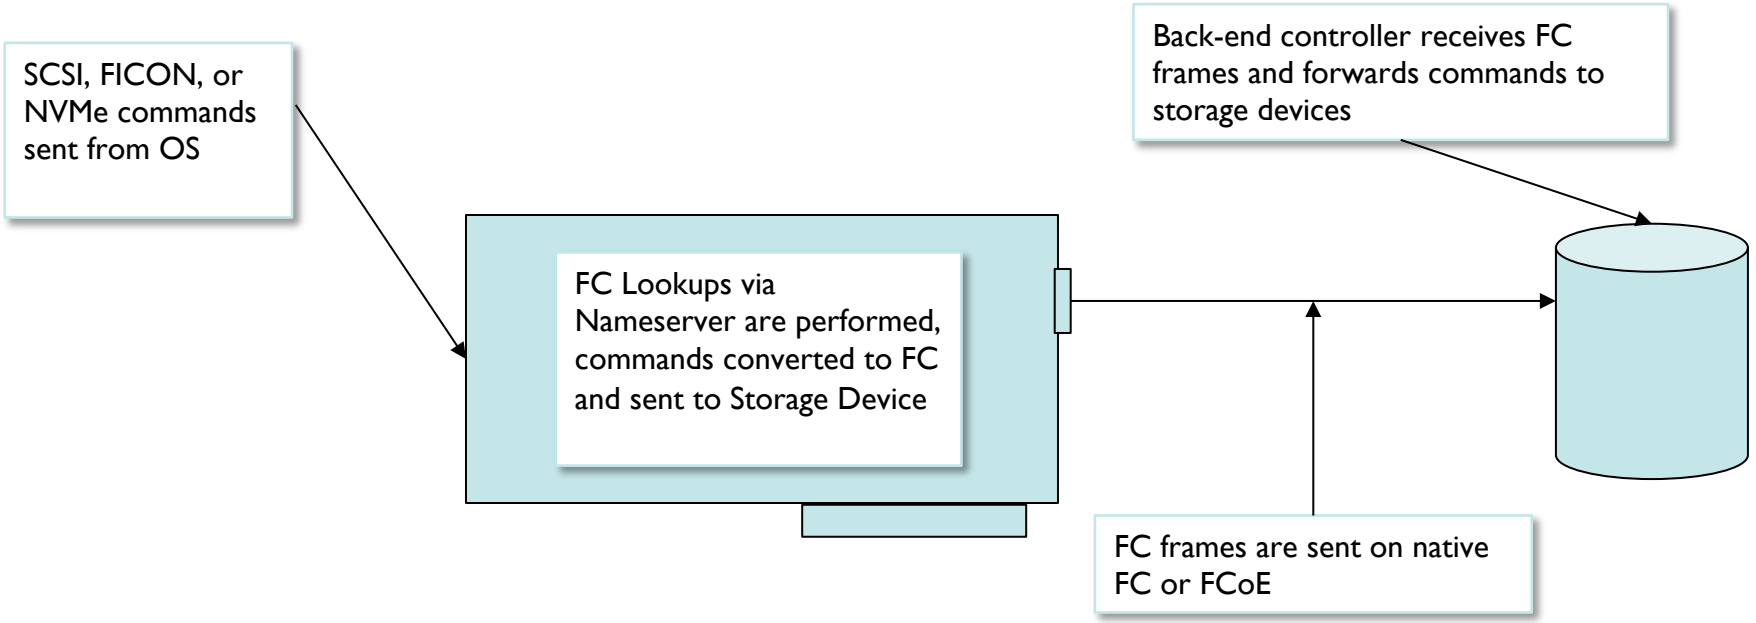

Key Takeaways

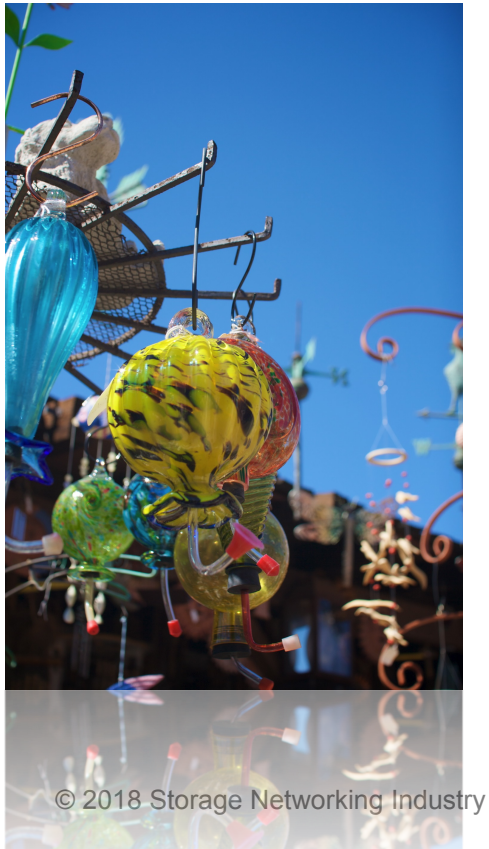

- Fibre Channel has supported SCSI and FICON since mid nineties
- NVMe over FC is the newest addition
- Multiple storage protocols can be supported at the same time on a single port

SNIA ESF | ETHERNET STORAGE

NVMe Controllers

Peter Onufryk

NVMe Controllers

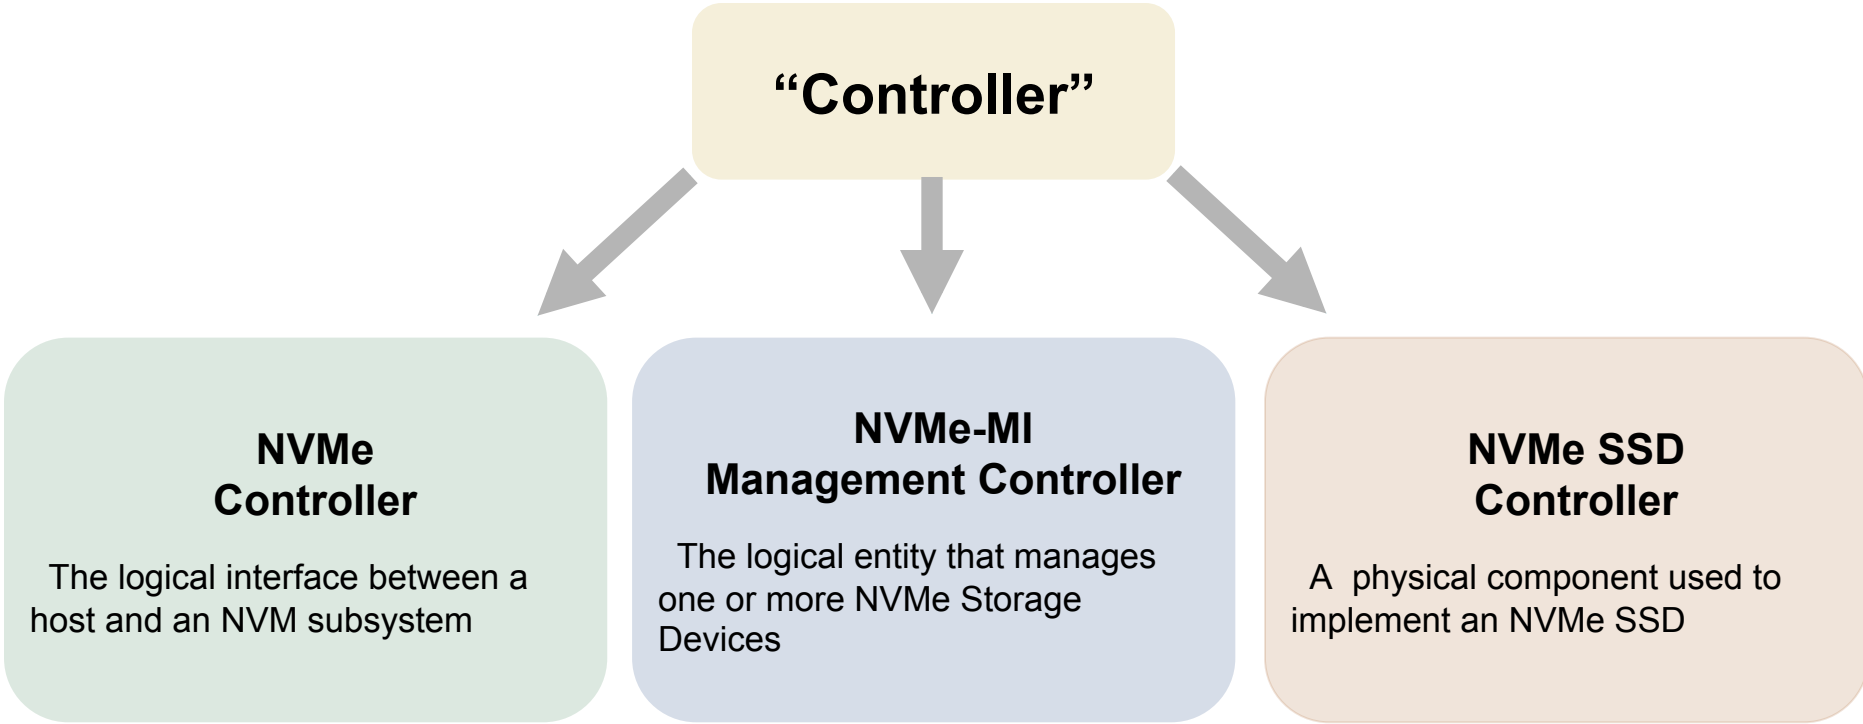

NVMe Controller

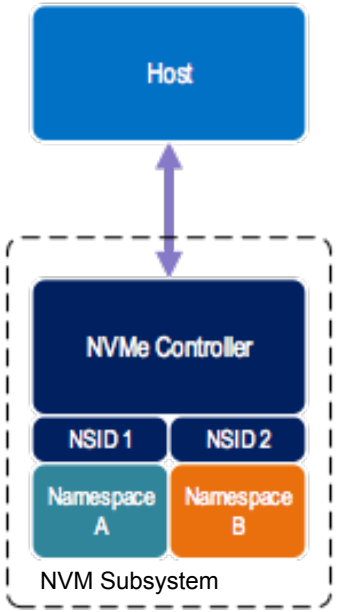

- An NVMe controller is the logical interface between a host and an NVM subsystem
  - Implements one Admin Submission Queue and Completion Queue
  - Implements one or more I/O Submission Queues and Completion Queues
  - Processes commands submitted on a Submission Queue and posts a Completion on a Completion Queue
  - When PCI Express is used as the transport, then a controller is a PCI function
    - PCI Function, SR-IOV Physical Function, or SR-IOV Virtual Function
  - May expose non-volatile memory storage medium to a host through one or more namespaces

Dual Ported NVMe Storage Device

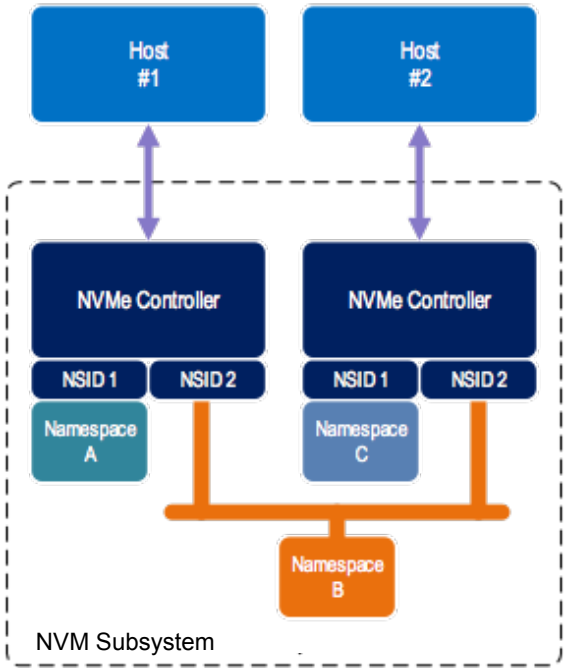

- A dual ported NVMe Storage Device contains an NVM subsystem with two controllers
- Each host is associated with its own independent controller
  - Independent Admin Submission and Completion queue
  - Independent I/O queues
  - Independent command processing

NVM Subsystem with SR-IOV

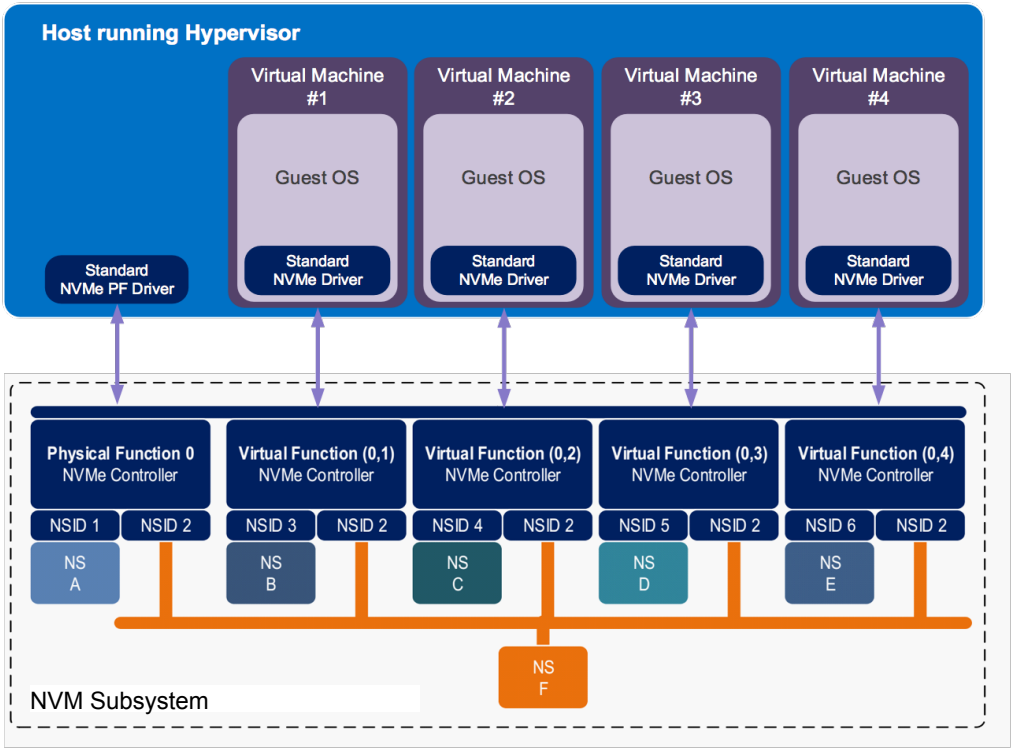

- An NVMe Storage Device that supports PCIe SR-IOV has one controller for the Physical Function (PF) and one controller for each Virtual Function (VF)
- Each virtual machine and the hypervisor is associated with its own independent controller
  - Independent Admin Submission and Completion queue
  - Independent I/O queues
  - Independent command processing

NVMe Specifications

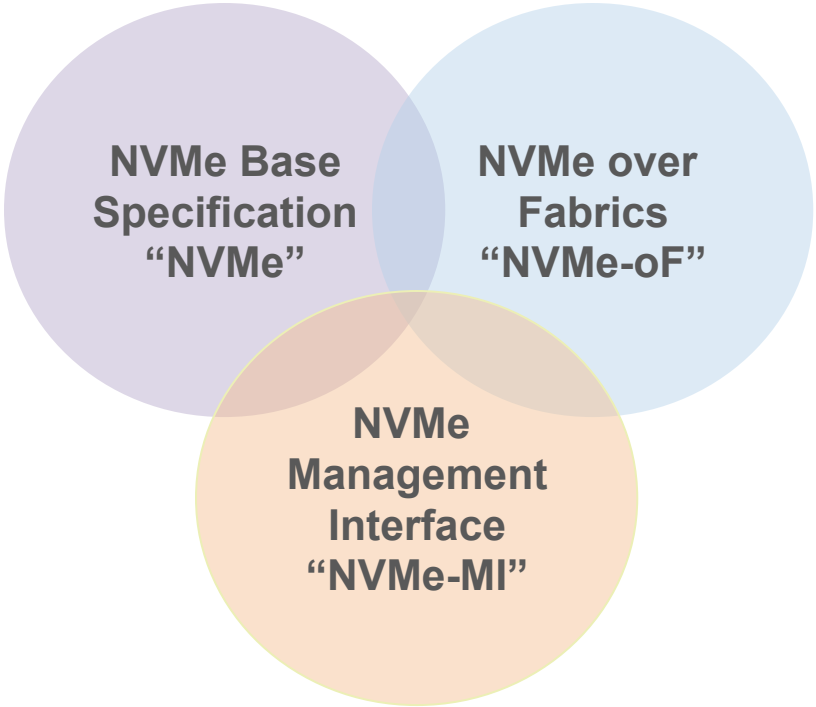

- NVMe specification
  - NVMe architecture and command set
  - NVMe over PCIe transport
- NVMe-oF specification
  - Extends NVMe architecture and command set to general interconnects
  - NVMe over RDMA transport
- NVMe-MI specification
  - Architecture and command set for out-of-band management of NVMe Storage Devices

NVMe-MI Out-of-Band Management

- Out-of-Band Management – Management that operates with hardware resources and components that are independent of operating system control

- Management Controller
  - Sometime referred to as a Baseboard Management Controller (BMC)
  - Management Responsibilities
    - Inventory
    - Configuration
    - Monitoring
    - Change Management

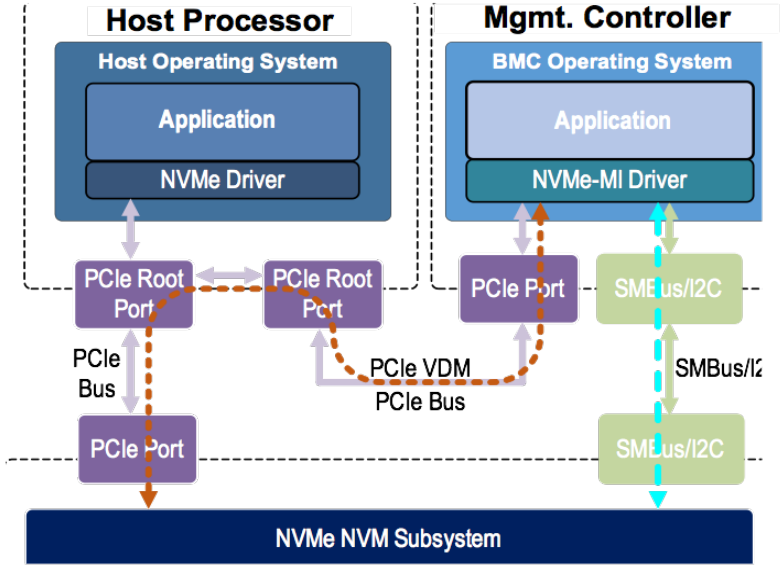

NVMe-MI Protocol Layering

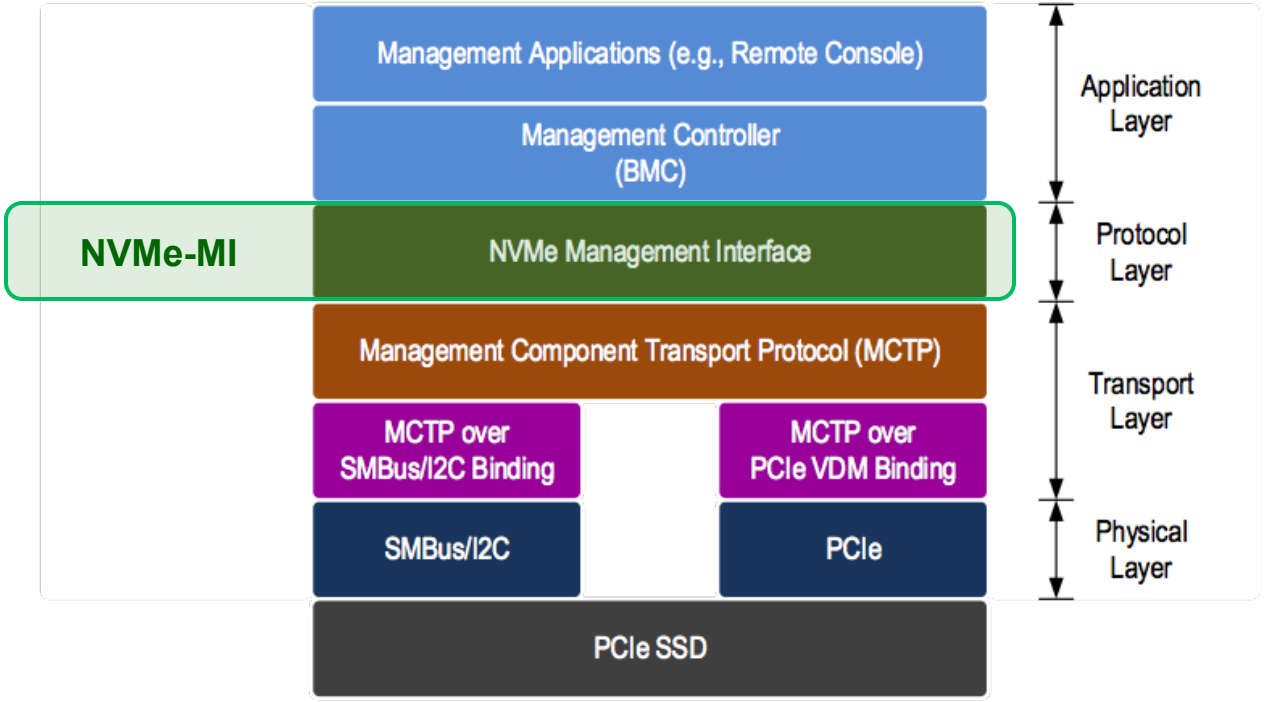

NVMe-MI Protocol Layering

NVMe SSD Controller

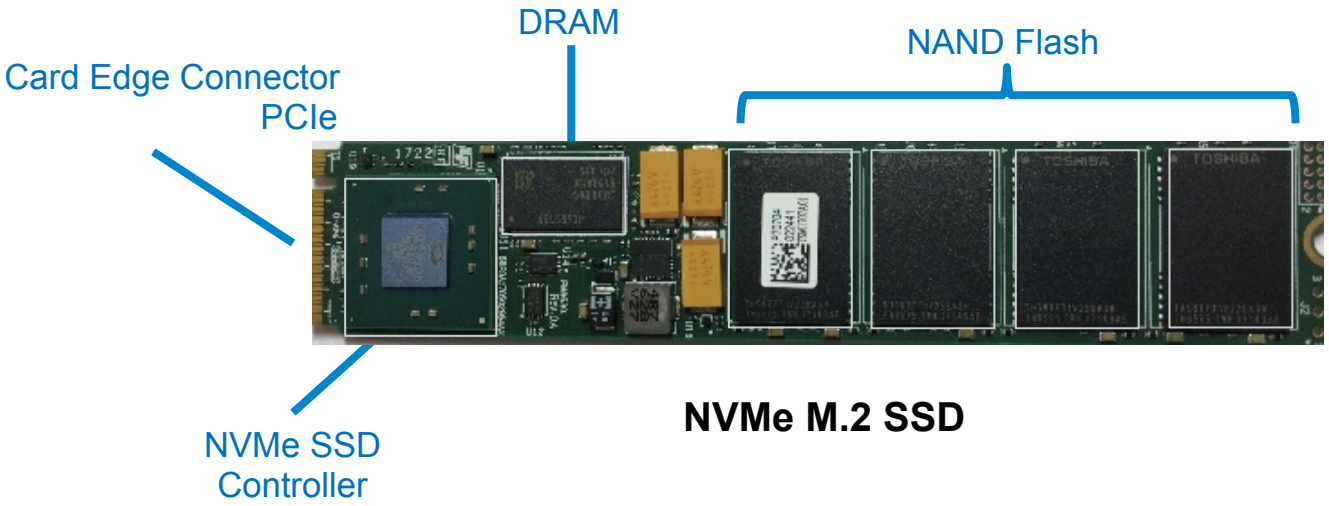

NVMe SSD Controller - A physical component used to implement an NVMe SSD

NVMe SSD Controller Internals

Internals of an NVMe SSD Controller

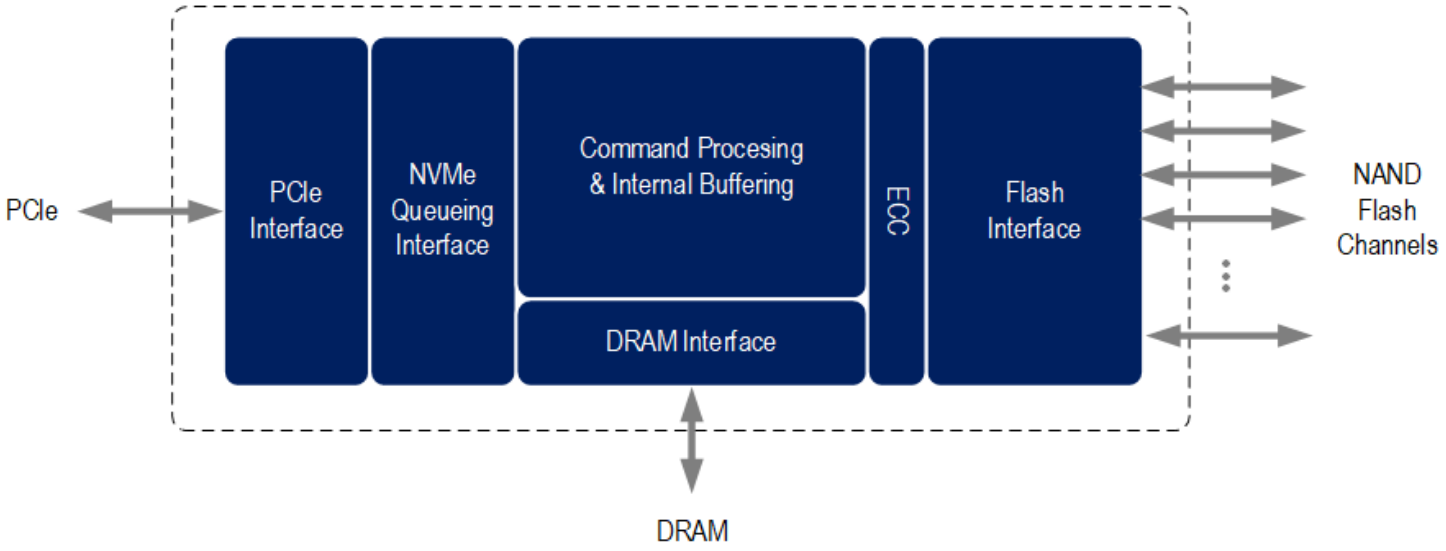

Key Takeaways

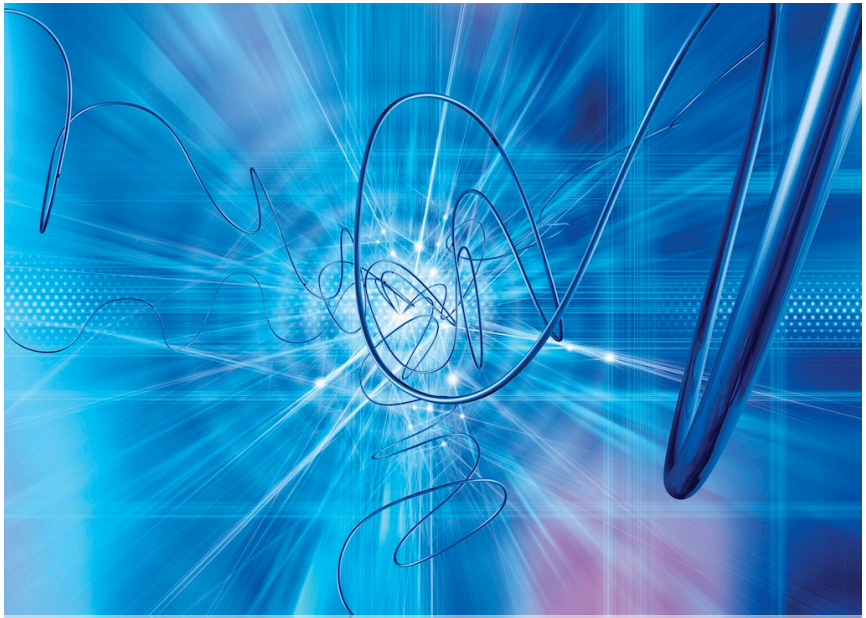

- This section shows examples of how different controllers can work in concert to provide holistic functionality in a system

SNIA ESF | ETHERNET STORAGE

SDN Controllers

Chad Hintz

What is SDN you ask!

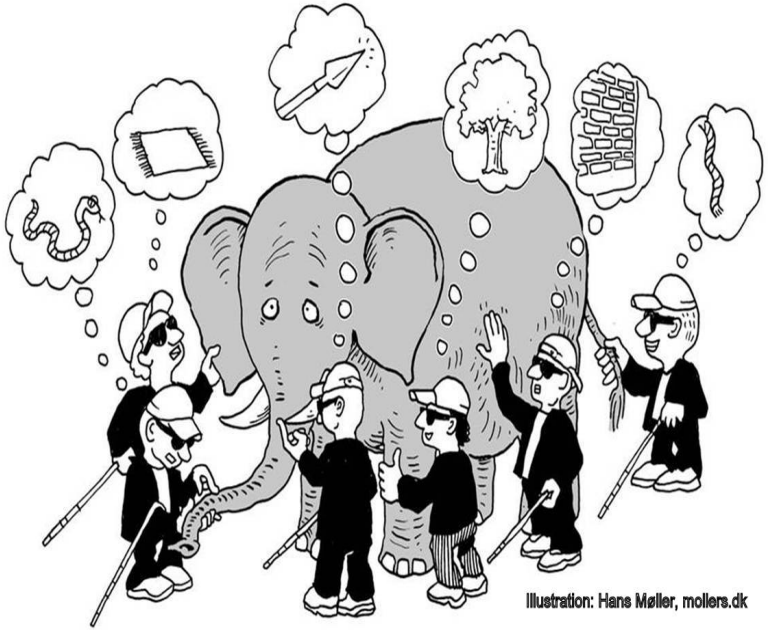

Basic Definitions

- An SDN controller is an application in software-defined networking (SDN) that manages flow control to enable intelligent networking. SDN controllers are based on protocols, such as OpenFlow, that allow servers to tell switches where to send packets.
  - Source: https://searchsdn.techtarget.com/definition/SDN-controller-software-defined-networking-controller
- In the SDN architecture, the **control and data planes are decoupled,** network intelligence and state are logically centralized, and the underlying network infrastructure is abstracted from the applications.
  - Source: www.opennetworking.org

The Traditional Network…

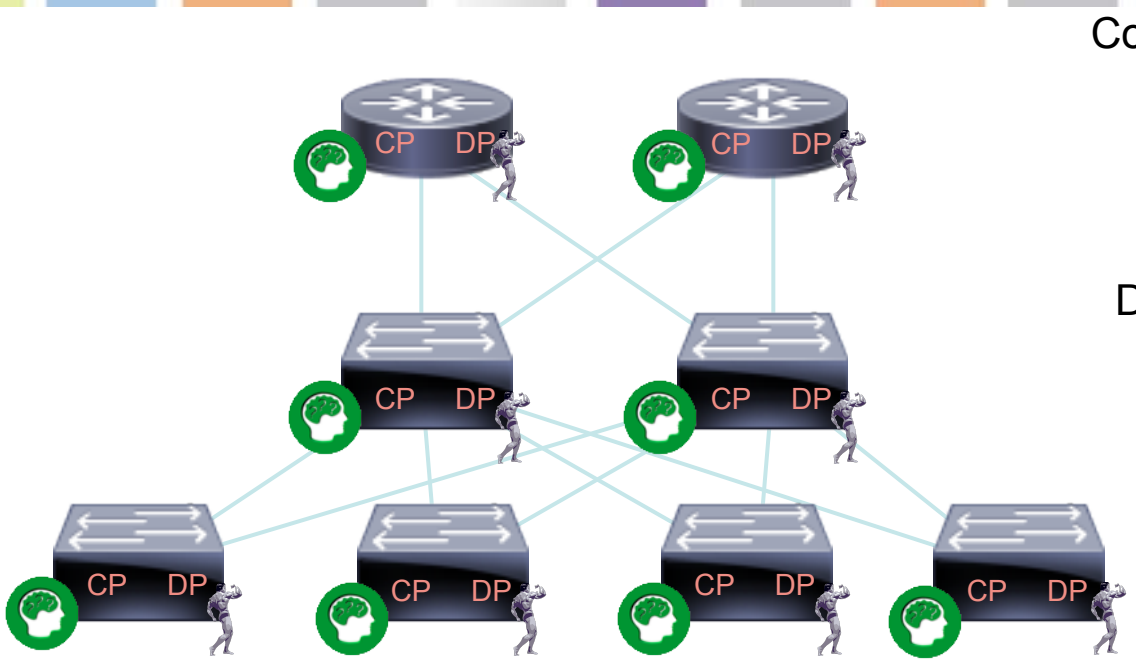

Control Plane (CP)

Data Plane (DP)

Control and Data Plane resides within Physical Device

Control plane learns/computes forwarding decisions

Data plane acts on the forwarding decisions

The Network As It Could Be…to an SDN 'Purist'

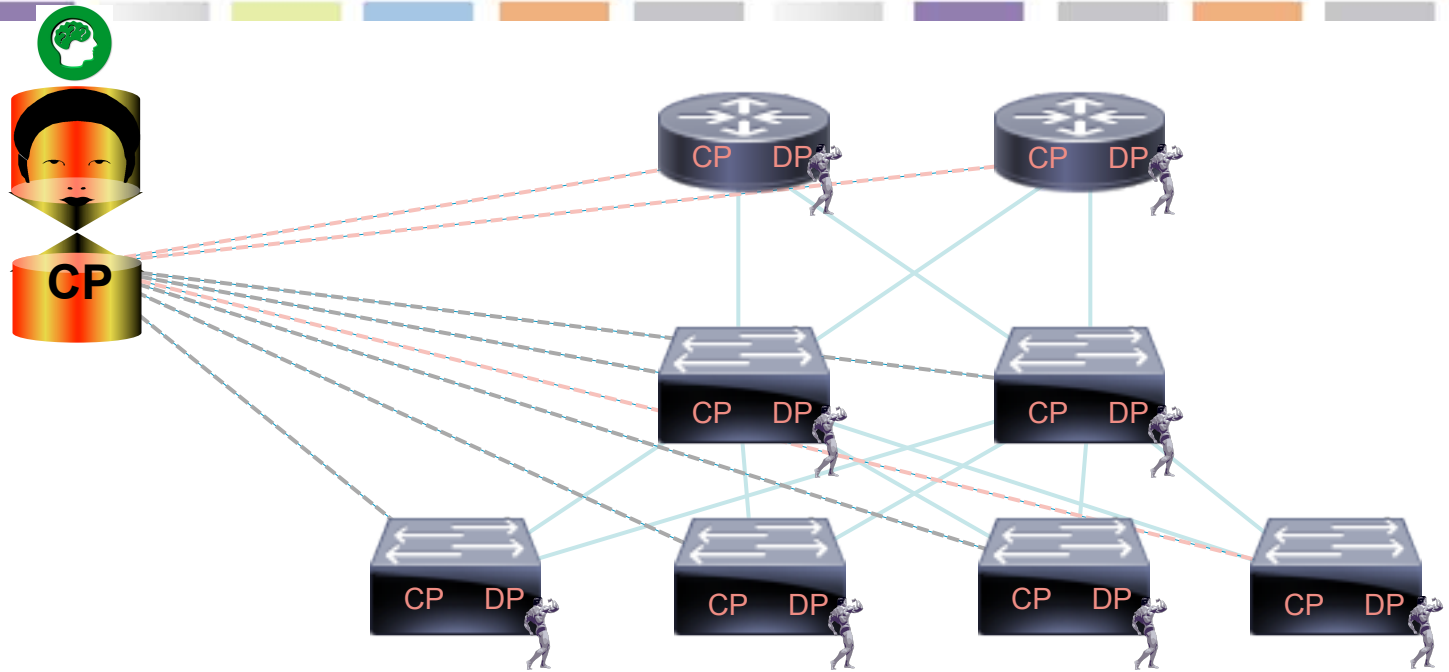

Control plane becomes centralized
Physical device retains Data plane functions only

Evolution to SDN

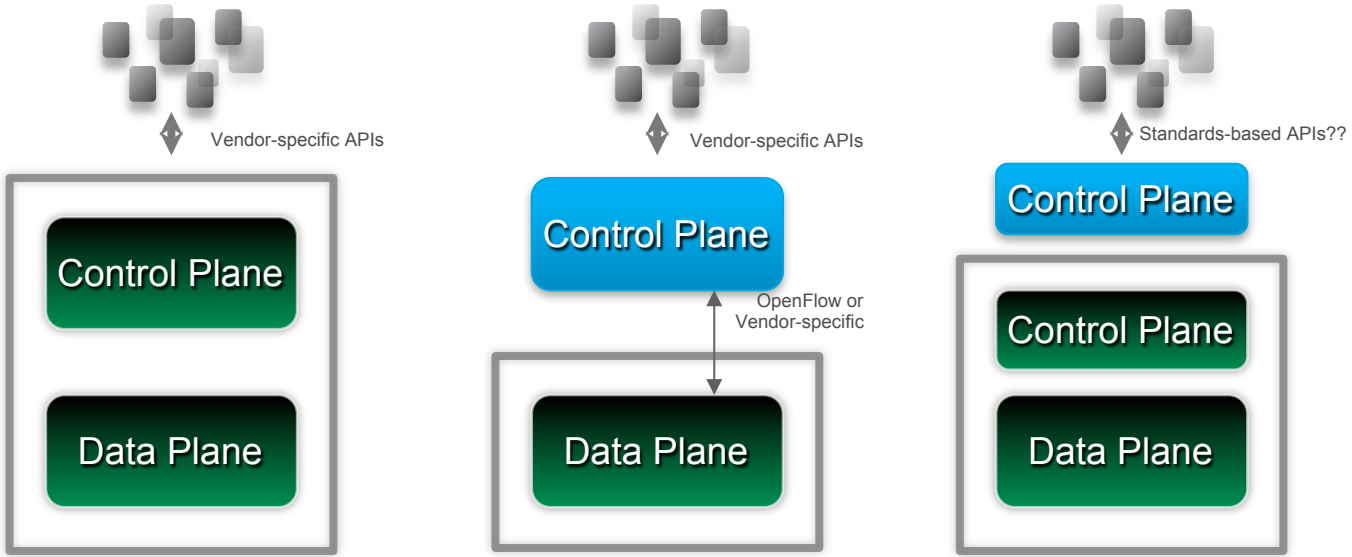

The Network As It Could Be…In a ‘Hybrid SDN’

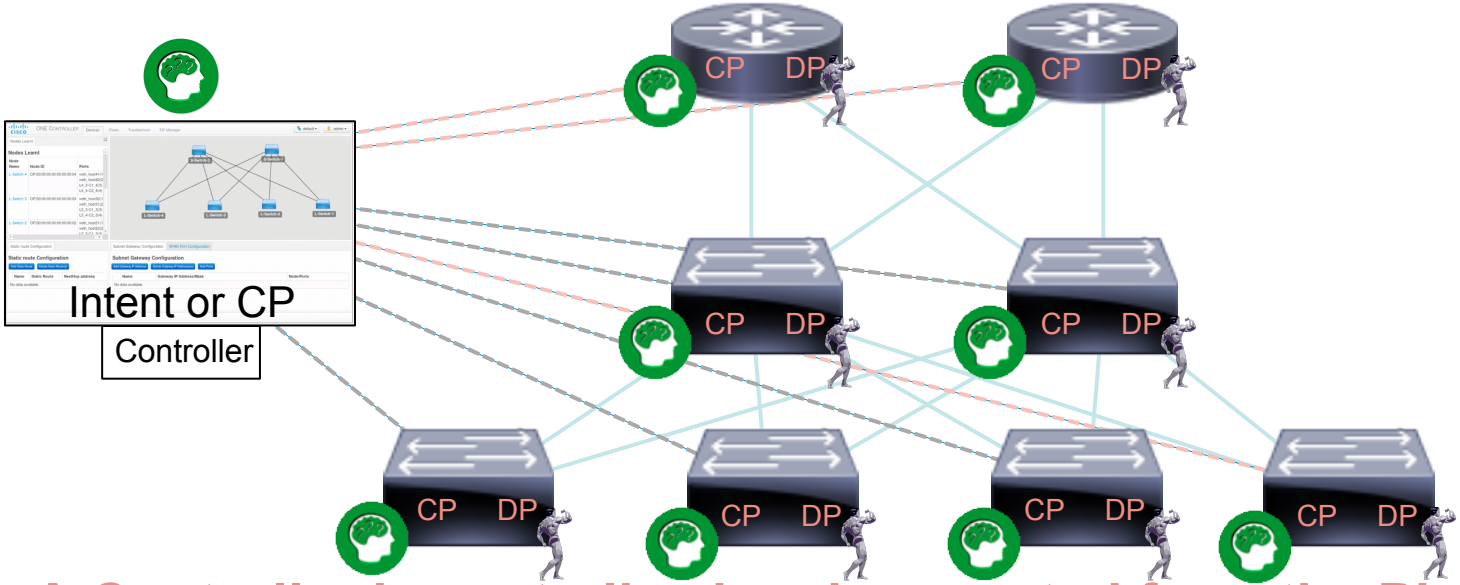

A Controller is centralized and separated from the Physical Device, but devices still retain a localized Control plane intelligence

Software Defined Networking and Software Defined Storage do they work together?

SDN & SDS

- Two ways of abstracting underlying hardware from a management/control plane perspective for their respective supported technologies
- SDN and SDS do not have anything to do with each other
  - SDN controllers do allow storage application to interact with one control plane (the controller) versus every networking device they are running on.

What is an SDS Controller

SNIA ESF | ETHERNET STORAGE

Data Services
- Application that runs in data plane to optimize storage
- Ex: Predictive Analytics
- Ex: De-Duplication
- Ex: Tiering

SDS Controller
- Visibility and Control of ALL storage resources
- Communication between Apps, Orchestrator and Storage Systems
- Allocates storage resources to meet SLA's

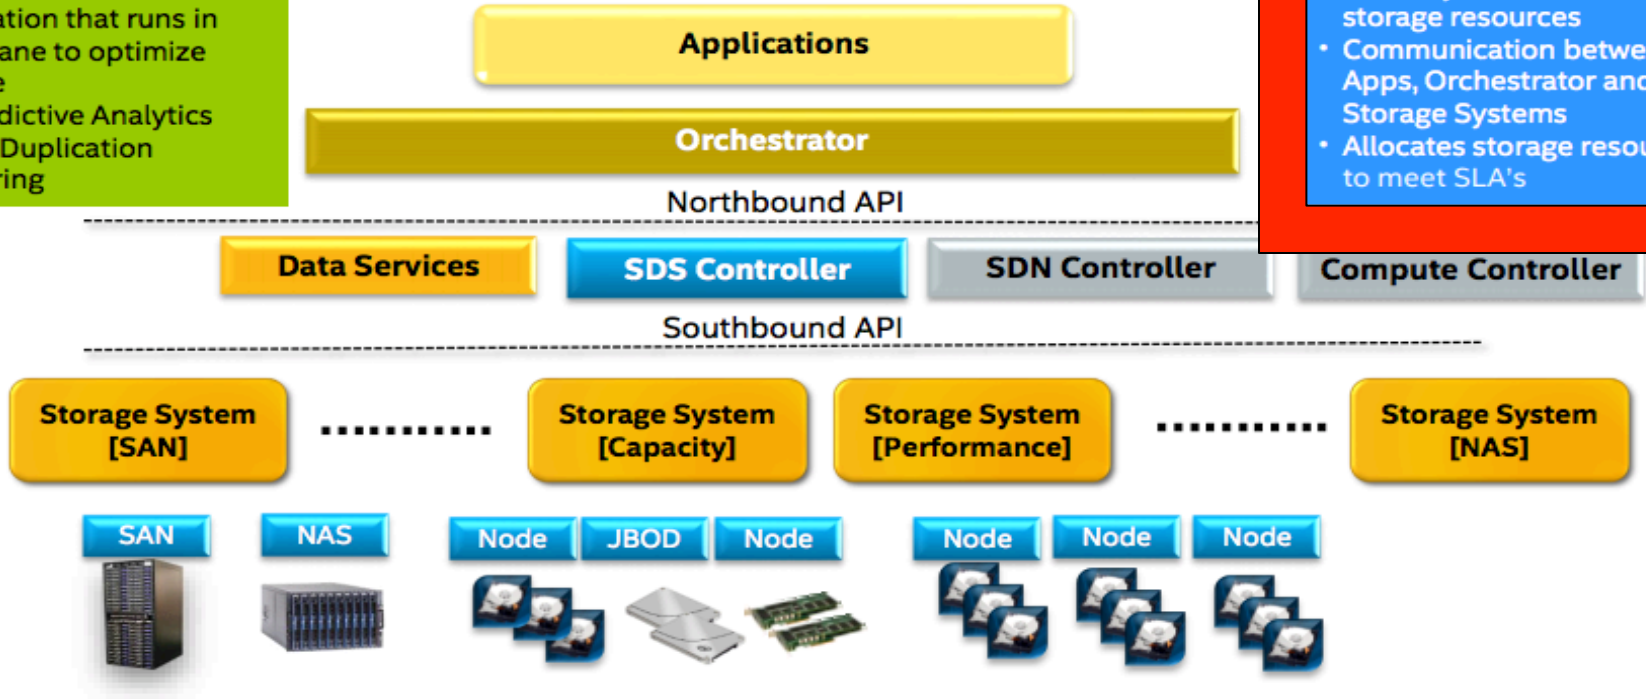

Where do SDS and SDN interact?

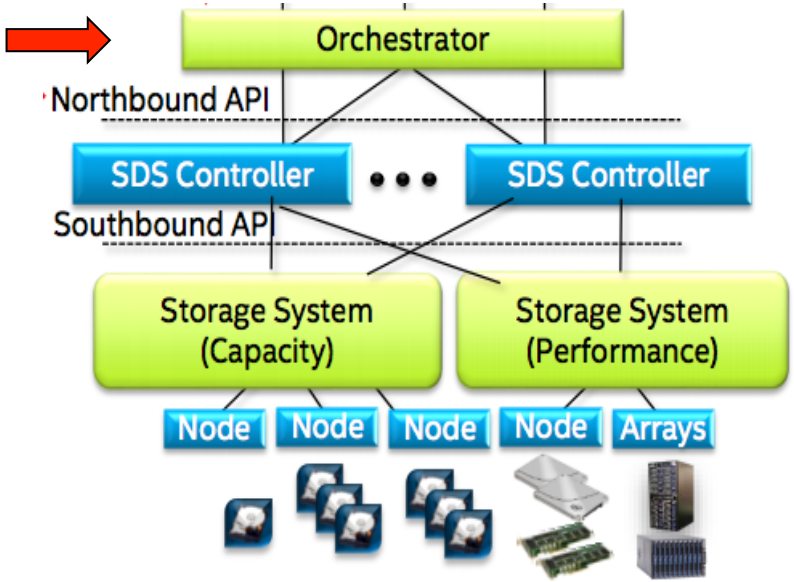

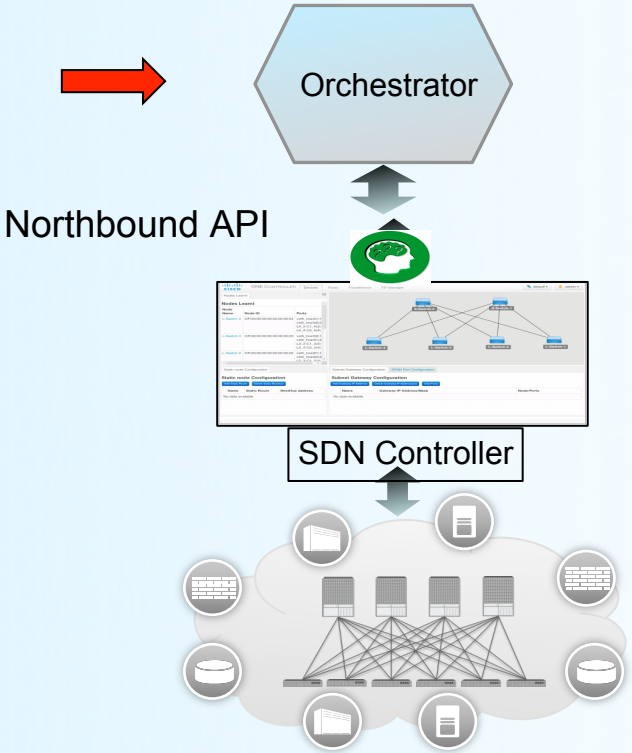

Key Takeaways

- SDN and SDS controllers are not the same thing
  - SDN controllers primarily control the management and flow creation of networking devices
  - SDS controllers primarily control the visibility and provisioning of all storage resources
- Both use northbound API to tie into applications and orchestration.

Summary

- Controller is a confusing historical term.
- 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)
  - Logic Controller (as used in Software-Defined Storage systems)

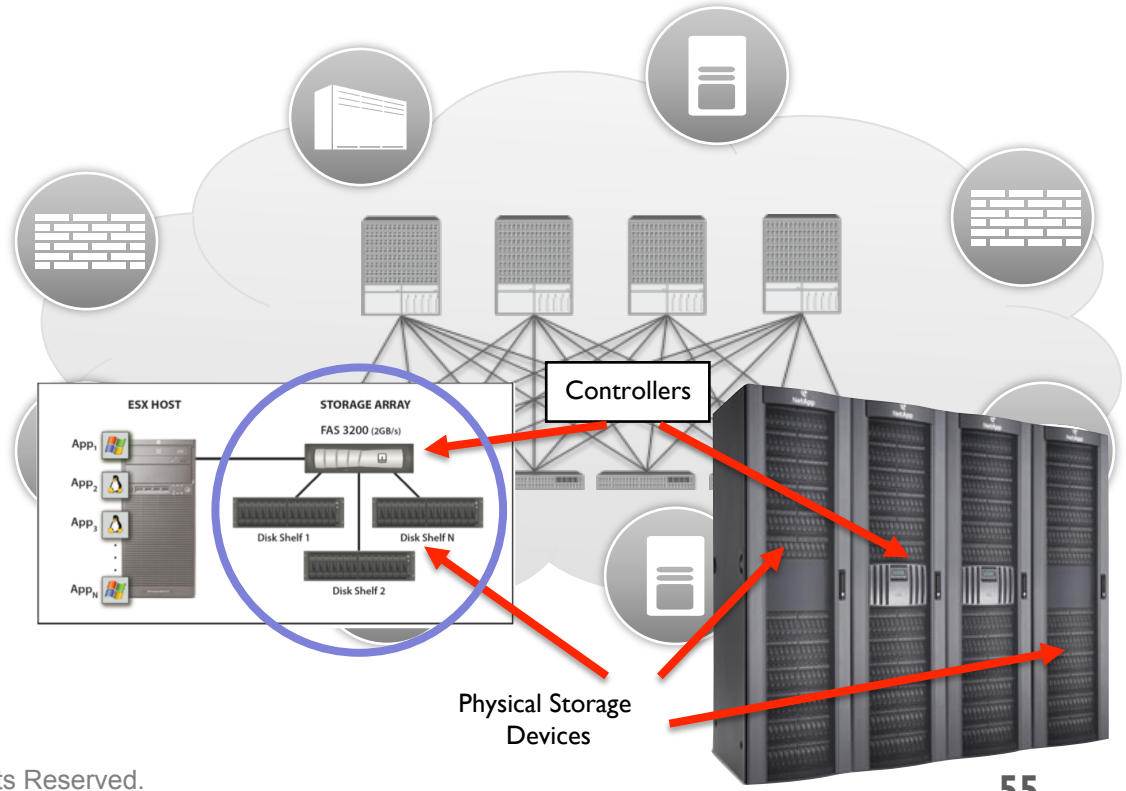

More Webcasts

SNIA ESF | ETHERNET STORAGE

- This is our 9th "Too Proud To Ask" webcast in this series
- 8 More "Everything You Wanted To Know About Storage But Were Too Proud To Ask" on-demand at:
  - https://www.snia.org/forums/esf/knowledge/webcasts-topics
- Next Live Webcast:
- The Great Storage Debate: FCoE vs. iSCSI vs. ISER
  - June 21, 2018, 10:00 am PT
  - Register at: https://www.brighttalk.com/webcast/663/318003

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 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