



**SDC** 18

September 24-27, 2018  
Santa Clara, CA

[www.storagedeveloper.org](http://www.storagedeveloper.org)

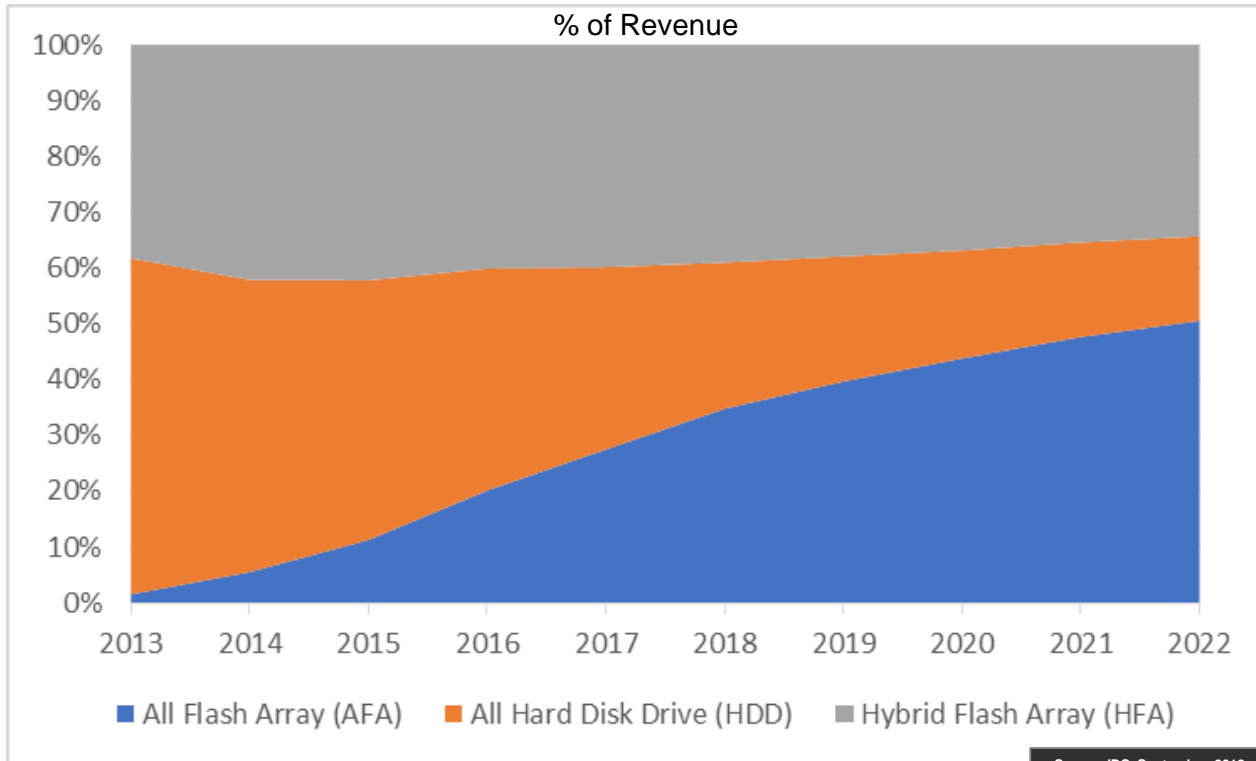
# Building an All-Flash Array with SAS, NVMe or SATA

**Cameron T Brett**

**Board of Directors, SCSI Trade Association**

**Director, SSD and Storage Solutions, Toshiba Memory America**

# All-Flash Array Type Forecast (2013-2022)

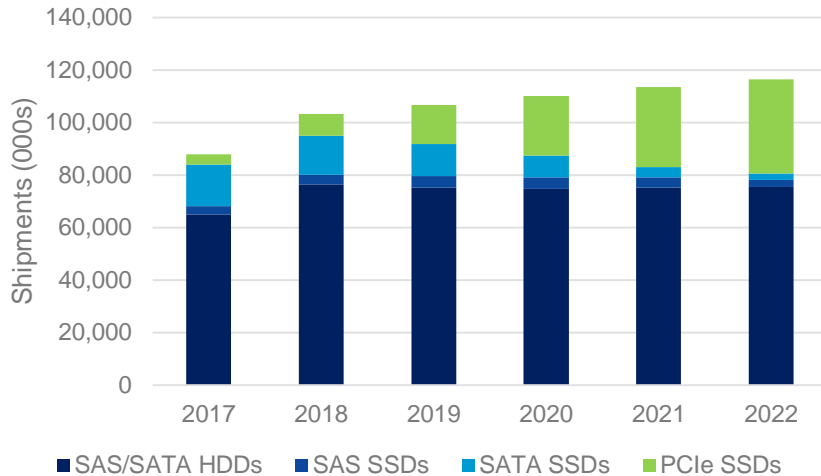


Source: IDC, September 2018


- ❑ 2017 AFA \$6.2B
- ❑ IDC sees AFAs growing to 51% of all arrays by 2022
- ❑ Steep decline of HDD-only arrays
- ❑ Slow decline of hybrid arrays

# SAS Is the Main Enterprise Storage Interface

WW Enterprise Drive Unit Forecast (2017-2022)



WW Enterprise Drive Capacity Forecast (2017-2022)



Source: IDC, May 2018

**SAS Infrastructure Enables >70% of Enterprise Storage Drives and >85% of Enterprise Storage Capacity thru 2022**

# All-Flash Arrays Built on SAS


# SAS Overview

- ❑ Developed in early 2000's. Based on SCSI protocol (1986)
- ❑ Data rates<sup>1</sup>: Up to 12Gb/s, with 24G SAS coming soon
- ❑ Cable lengths: Up to 100m optical; up to 10m copper
- ❑ Typical deployment: server; enterprise storage; high-performance applications; large scale requiring up to 1000s of devices
- ❑ Supports both SAS and SATA HDDs and SSDs
- ❑ SCSI protocol is supported with other storage technologies, such as FC, USB, 1394/FireWire, iSCSI
- ❑ Top SAS SSD suppliers – Toshiba, HGST/WD, Seagate, Samsung
- ❑ 33.07M SAS drives sold in 2017<sup>2</sup> (3.0M SSDs and 30.07M HDDs)


<sup>1</sup> Single-lane performance    <sup>2</sup> Source: TRENDFOCUS, 2018

# SAS AFA Enterprise Data Center Use

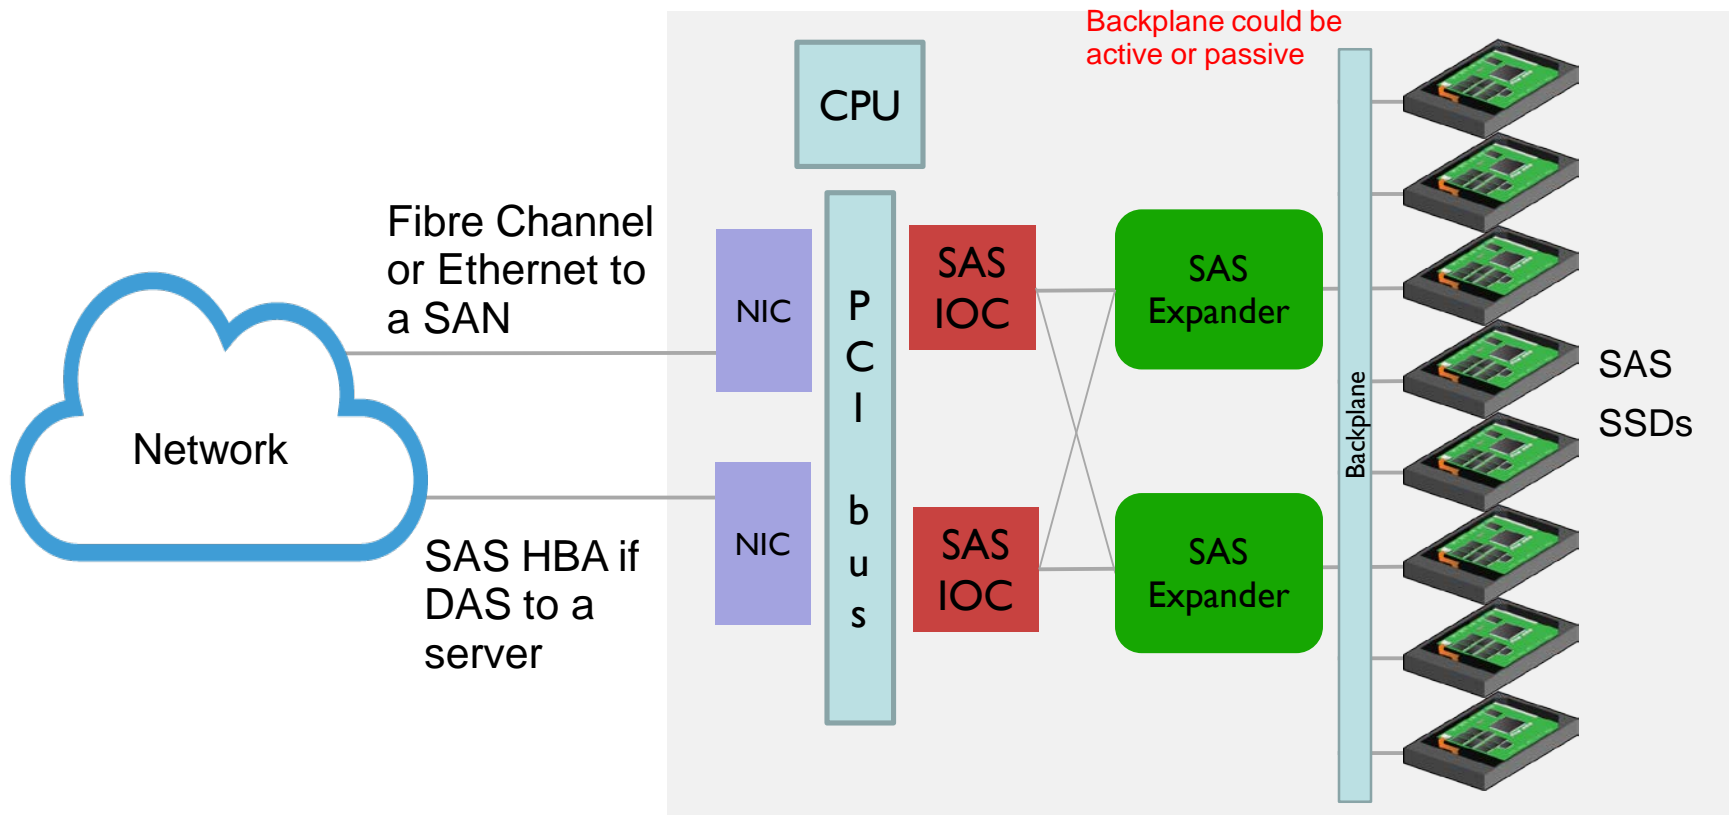
## Storage Area Network



## Direct Attached Storage



# SAS AFA – Closer Look Inside



# The Value Pillars of SAS

**Reliable**

**Secure**

**Trusted**

**Data Center  
Ready**

**SAS and  
SATA**

**SSDs and  
HDDs**

**Backward  
Compatibility**

**1000s of End  
Devices**

**Capacity  
Storage**

**Enables  
“Fabric”  
Solutions**

**\$ / Performance**

**\$ / GB**

**Highest SSD  
capacities**

**Huge  
Established  
Ecosystem**

**Development**

**Dependability**

**Flexibility**

**Scalability**

**Economics**



# What's Next? 24G SAS Highlights

## Physical Layer Enhancements

Double the effective single-lane bandwidth of 12Gb/s SAS

Higher throughput and IOPS performance

Enhanced 20-bit Forward Error Correction (FEC)

More robust data reliability and connectivity

SAS-4 enhanced transmitter training algorithm

Continuous optimal signal tuning

## Protocol & Block Level Enhancements

Fairness enhancements

Performance consistency across large and mixed protocol topologies

Storage intelligence and persistent connections

Improves SSD efficiency, latency and QoS

SMP priorities

Prioritizes management-class communications for complex, deep topologies

# All-Flash Arrays Built on NVMe

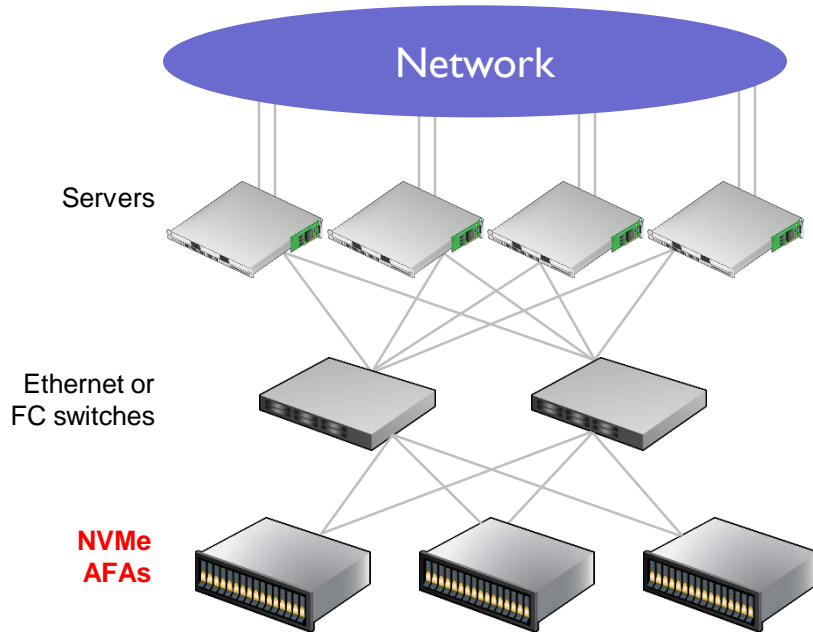
# NVMe Overview

- ❑ Introduced in 2011 with large industry backing; designed for SSDs
- ❑ Works over PCIe. Ethernet, Fibre Channel, IB for NVMe-oF.
- ❑ Robust roadmap plans for protocol, NVMe-oF and management
- ❑ Data rates<sup>1</sup>: Up to 16Gb/s
- ❑ Connection lengths: 1m internal, external not defined for PCIe
- ❑ Typical deployment: server, HFA/AFA, cache tier, high-performance, lowest-latency applications
- ❑ Many, many companies developing NVMe storage products
- ❑ 3.93M NVMe drives sold in 2017<sup>2</sup>


<sup>1</sup> Single-lane performance    <sup>2</sup> Source: TRENDFOCUS, 2018

# NVMe AFA Enterprise Data Center Use

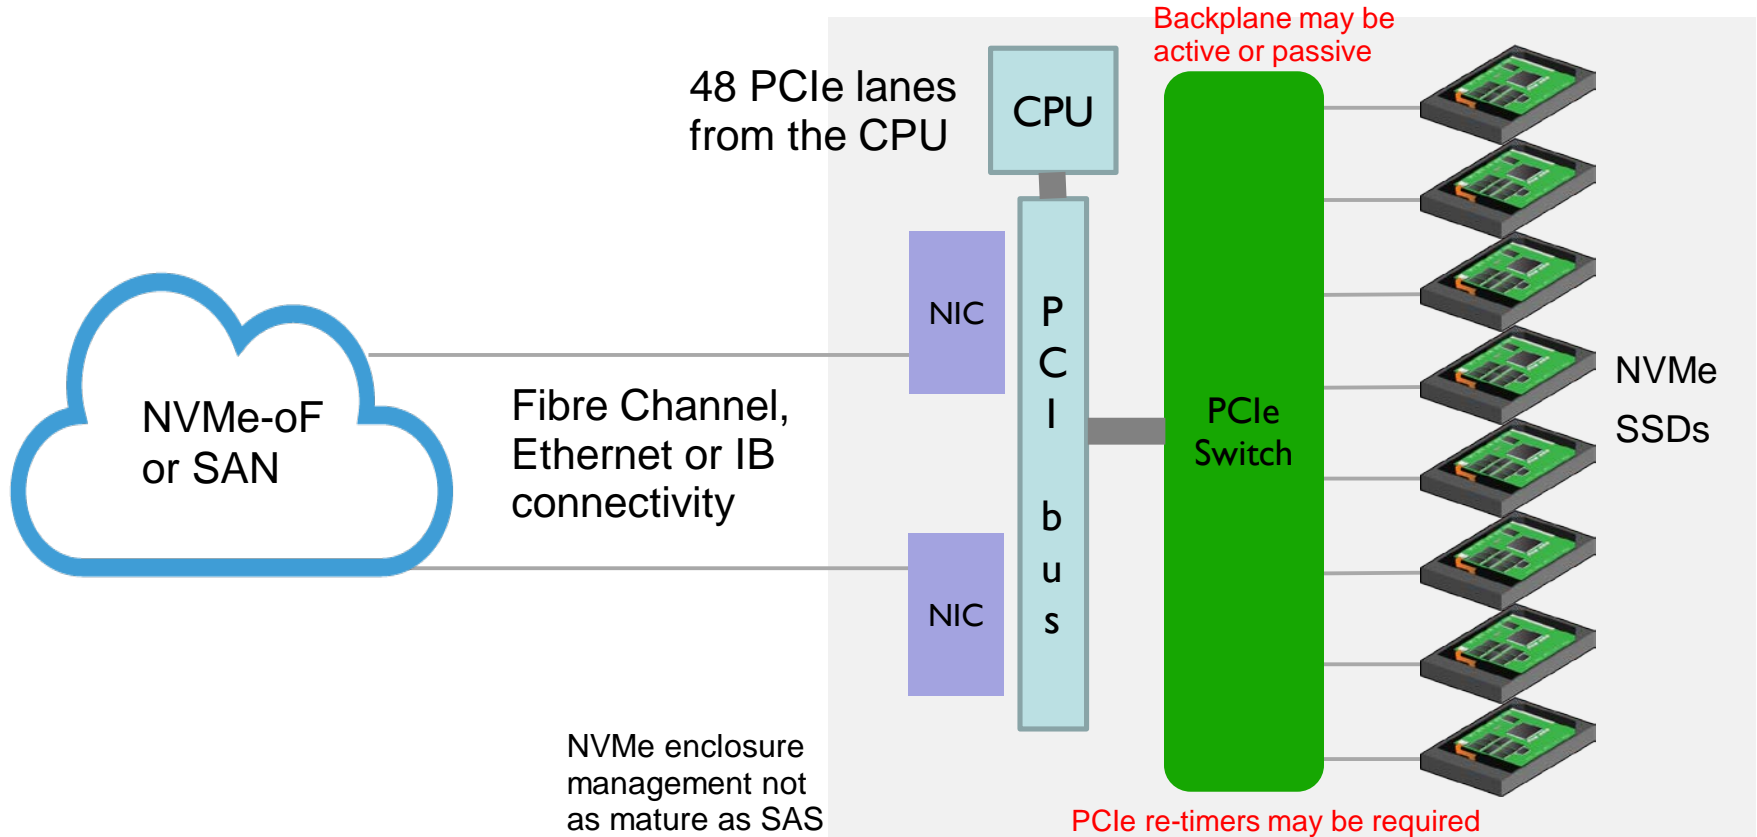
## Storage Area Network



## NVMe-oF Network (Disaggregated)



# NVMe AFA – Closer Look Inside



# NVMe Value Proposition

- ❑ **Large industry backing; multiple sources**
- ❑ **Optimized for flash storage and next generation NVM**
- ❑ **Lowest latency and highest performance for typical deployments**
- ❑ **Lower AFA BOM cost for smaller topologies (no IOC required)**
  - ❑ Large and HA topologies require additional switches
- ❑ **Best for high performance storage tiers (i.e. Tier-0)**
- ❑ **Form factor and performance variations for client, data center and enterprise applications**

# All-Flash Arrays Built on SATA

# SATA Overview

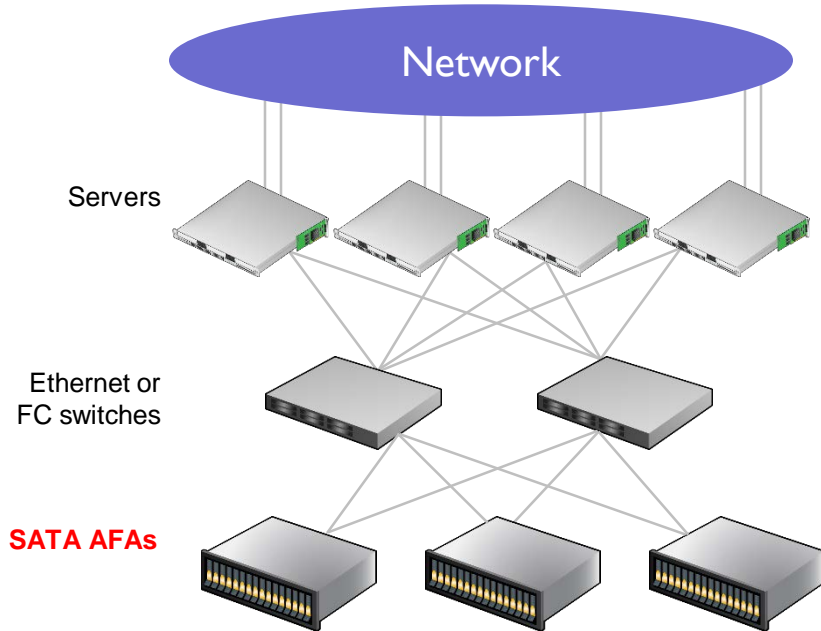
- ❑ Developed in early 2000s
- ❑ Roadmap ends with 6Gb/s; no further physical layer development
- ❑ Data rates<sup>1</sup>: 6Gb/s, 3Gb/s, 1.5Gb/s
- ❑ Cable lengths: Up to 1m
- ❑ Typical deployment: desktop, server and storage; near-line and cold storage; low-cost, high-capacity applications,
- ❑ In an AFA or enterprise server, SATA devices commonly deployed behind SAS infrastructure
- ❑ 50.61M SATA drives shipped in 2017 (15.78M SSDs and 34.83M HDDs)<sup>2</sup>

<sup>1</sup> Single-lane performance    <sup>2</sup> Source: TRENDFOCUS, 2018




# SATA AFA Enterprise Data Center Use

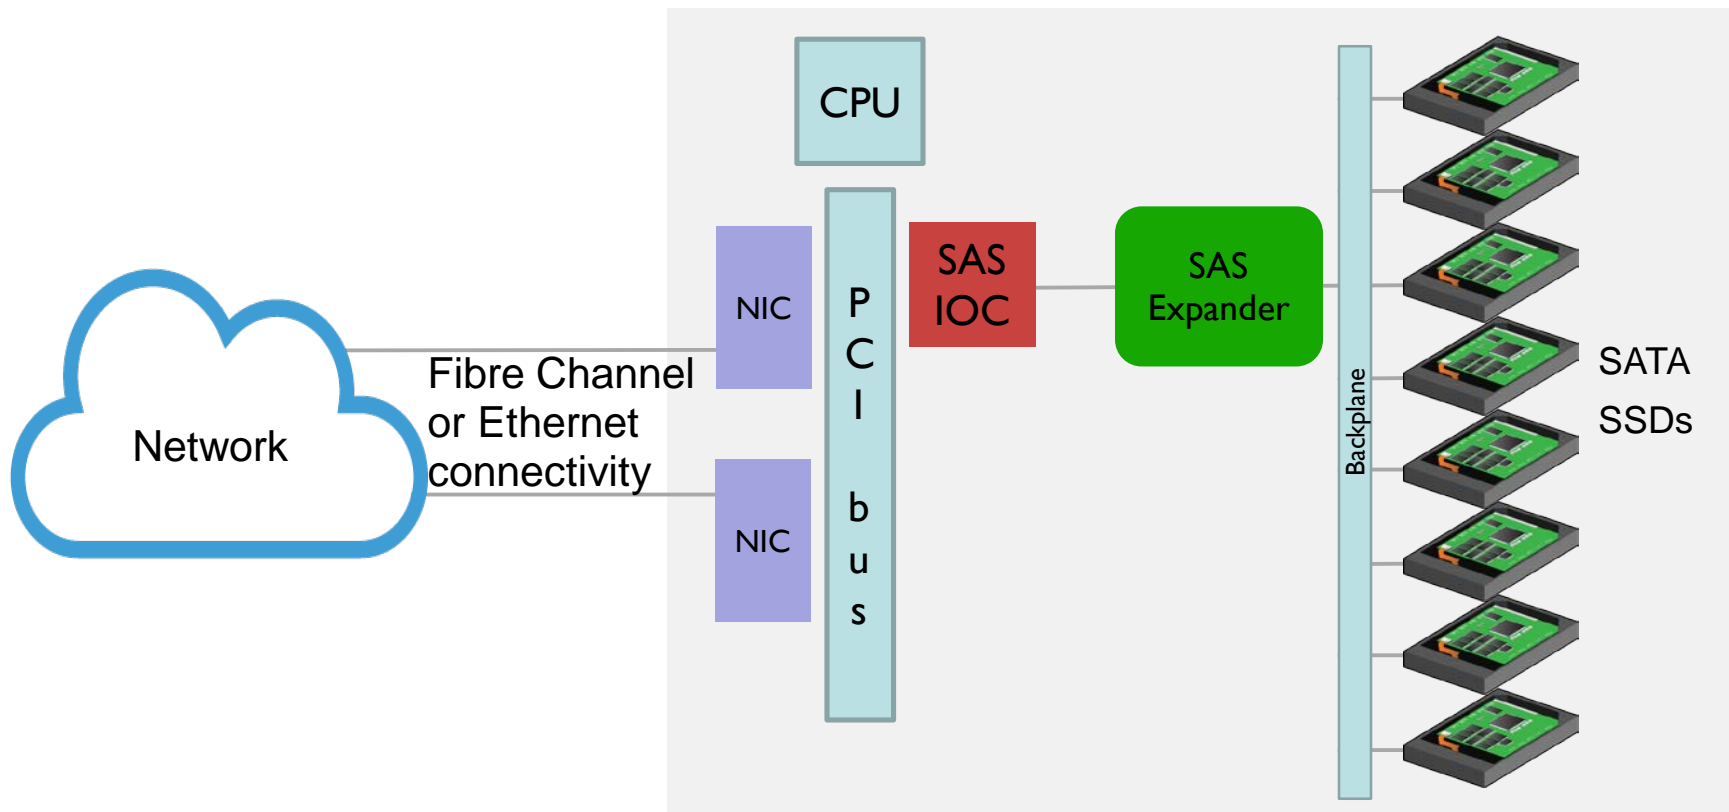
## Storage Area Network



## Direct Attached Storage



# SATA AFA – Closer Look Inside



# SATA Value Proposition

- ❑ **It's cheap - leverages economies of scale to lower cost**
- ❑ **Lowest power consumption**
- ❑ **Leverages SAS infrastructure**
- ❑ **Both SSDs and HDDs supported**
- ❑ **Will have a lengthy tail and will be around for a long time.....**

# Summary Comparison of Typical Deployment

	SAS	NVMe	SATA
Performance (IOPS, GB/s)	Better (x1 lane)	Best (x4 lanes)	Good
Performance (Read Latency*)	Good	Best	Good
Scalability	Best	Good	Better (SAS infrastructure)
Power (per drive)	Better (9-12W)	Good (up to 25W)	Best (~6W)
Flexibility	Best (SAS, SATA, HDDs, SSDs)	Good (SSDs)	Better (SSDs, HDDs)
Manageability	Best (most mature)	Good (recent spec)	Better (SAS infra)
Reliability	Best	Better	Good
System cost	Higher	Highest (performance premium)	Lowest
Roadmap future	Long-term	Long-term	Limited

\*Latency includes OS, driver, HBA (if required) and flight time, media access times not included

# Summary

- ❑ SAS is a modern, and mature interconnect
  - ❑ Actively being developed with long roadmap plans
  - ❑ New features to enhance SSD functionality and performance
  - ❑ Highly scalable and flexible
  - ❑ Value SAS may cannibalize SATA
  - ❑ 3Gb/s, 6Gb/s, 12Gb/s, 24G (sampling today)
- ❑ NVMe is still new, but will mature over time
  - ❑ Highest performance with lowest latency
  - ❑ Very active development and robust, long-term roadmap
  - ❑ Development of specifications for networked NVMe storage and management
- ❑ SATA has no future roadmap
  - ❑ Unit shipments of SATA devices still strong; will have a long tail