SNIA NSF | NETWORKING STORAGE

Security of Data on NVMe over Fabrics, The Armored Truck Way

Live Webcast

May 12, 2021

10:00 am PT / 1:00 pm ET

Today’s Presenters

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. NSF | NETWORKING STORAGE

Technologies We Cover

Ethernet, Fibre Channel, InfiniBand®

iSCSI, NVMe-oF™, NFS, SMB

Virtualized, HCI, Software-defined Storage

Storage Protocols (block, file, object)

Securing Data

SNIA. NSF | NETWORKING STORAGE

SNIA. NSF | NETWORKING STORAGE

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.

Storage Security Overview

Nishant Lodha

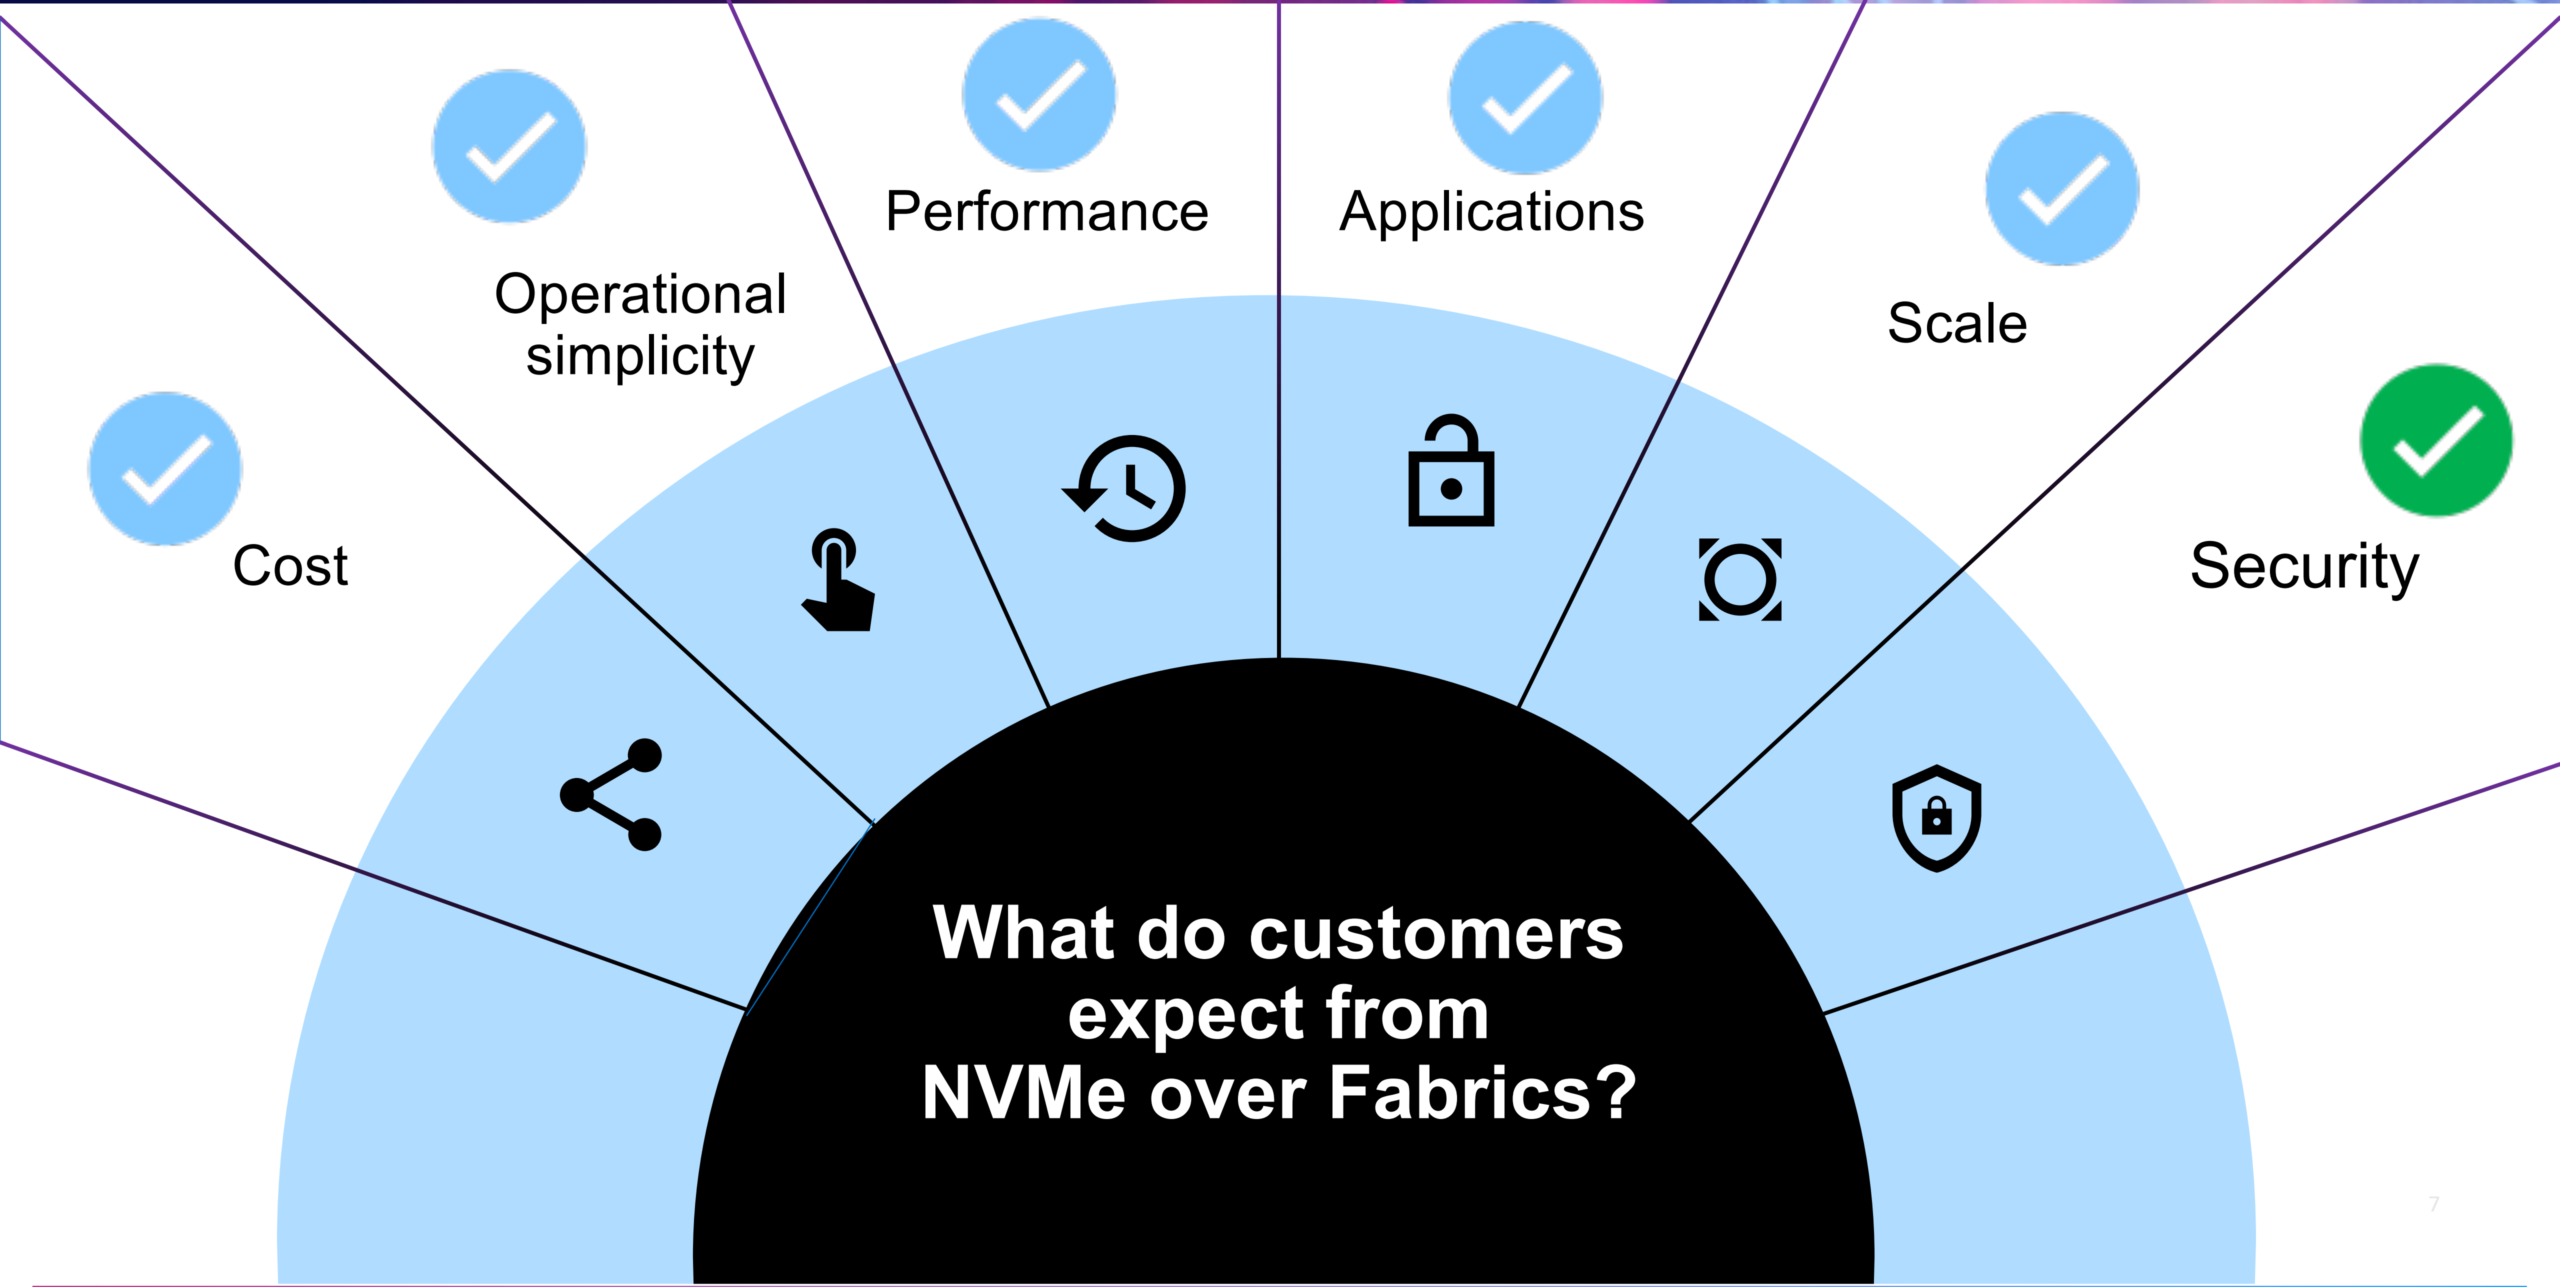
Cost
Operational simplicity
Performance
Applications
Scale
Security
What do customers expect from NVMe over Fabrics?

Data Protection vs. Security

There is often confusion between storage/data security and protection.

Data security refers specifically to measures taken to protect the integrity of the data itself.

Data security primarily involves keeping private information out of the hands of anyone not authorized to see or modify it.

Unauthorized access and access control, auditing

Intentional or accidental loss/corruption of sensitive data

Data Security measures include encryption of data both at rest and in motion, physical and logical access controls that prevent unauthorized access etc.

Data Protection, refers to the mechanism of making copies of your data to restore in the event of a loss or corruption.

SNIA. NSF | NETWORKING STORAGE

Datacenter Security and Standards

Incidents and Security Standards

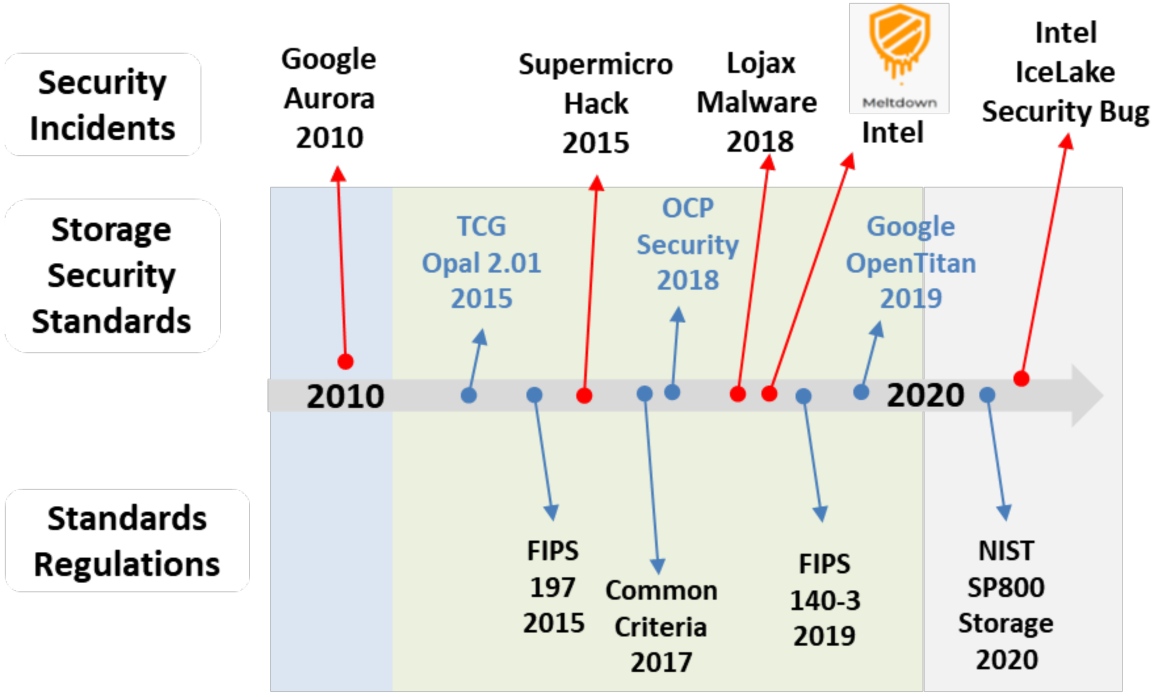

- Standards, Security threats growing in past 10 yrs.
- New Security Standards organizations emerged

Datacenter Security Considerations

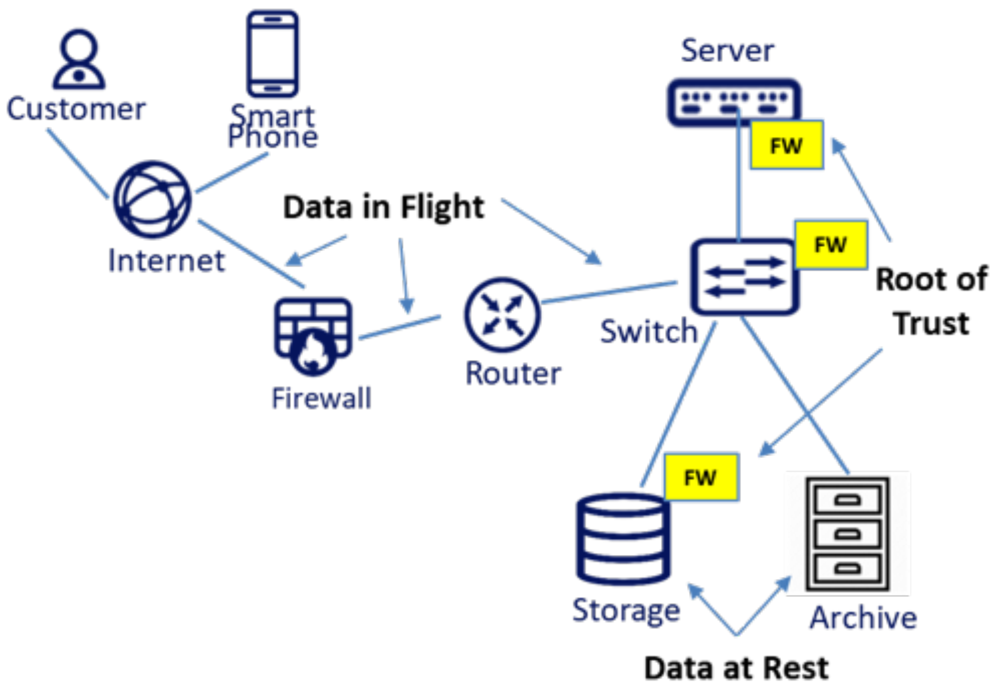

- **Data in Flight**: Network security (especially applicable for shared infrastructure)
- **Data at Rest**: Against theft of data or keys, and ransomware (esp. SSD media and key encryption with SED)
- **HW Root of Trust** : Dedicated security engine to ensure Secure Boot, Secure FW, and Key Management across all peripherals

SNIA. NSF | NETWORKING STORAGE

Storage Fabrics: Where to Secure – 10K View

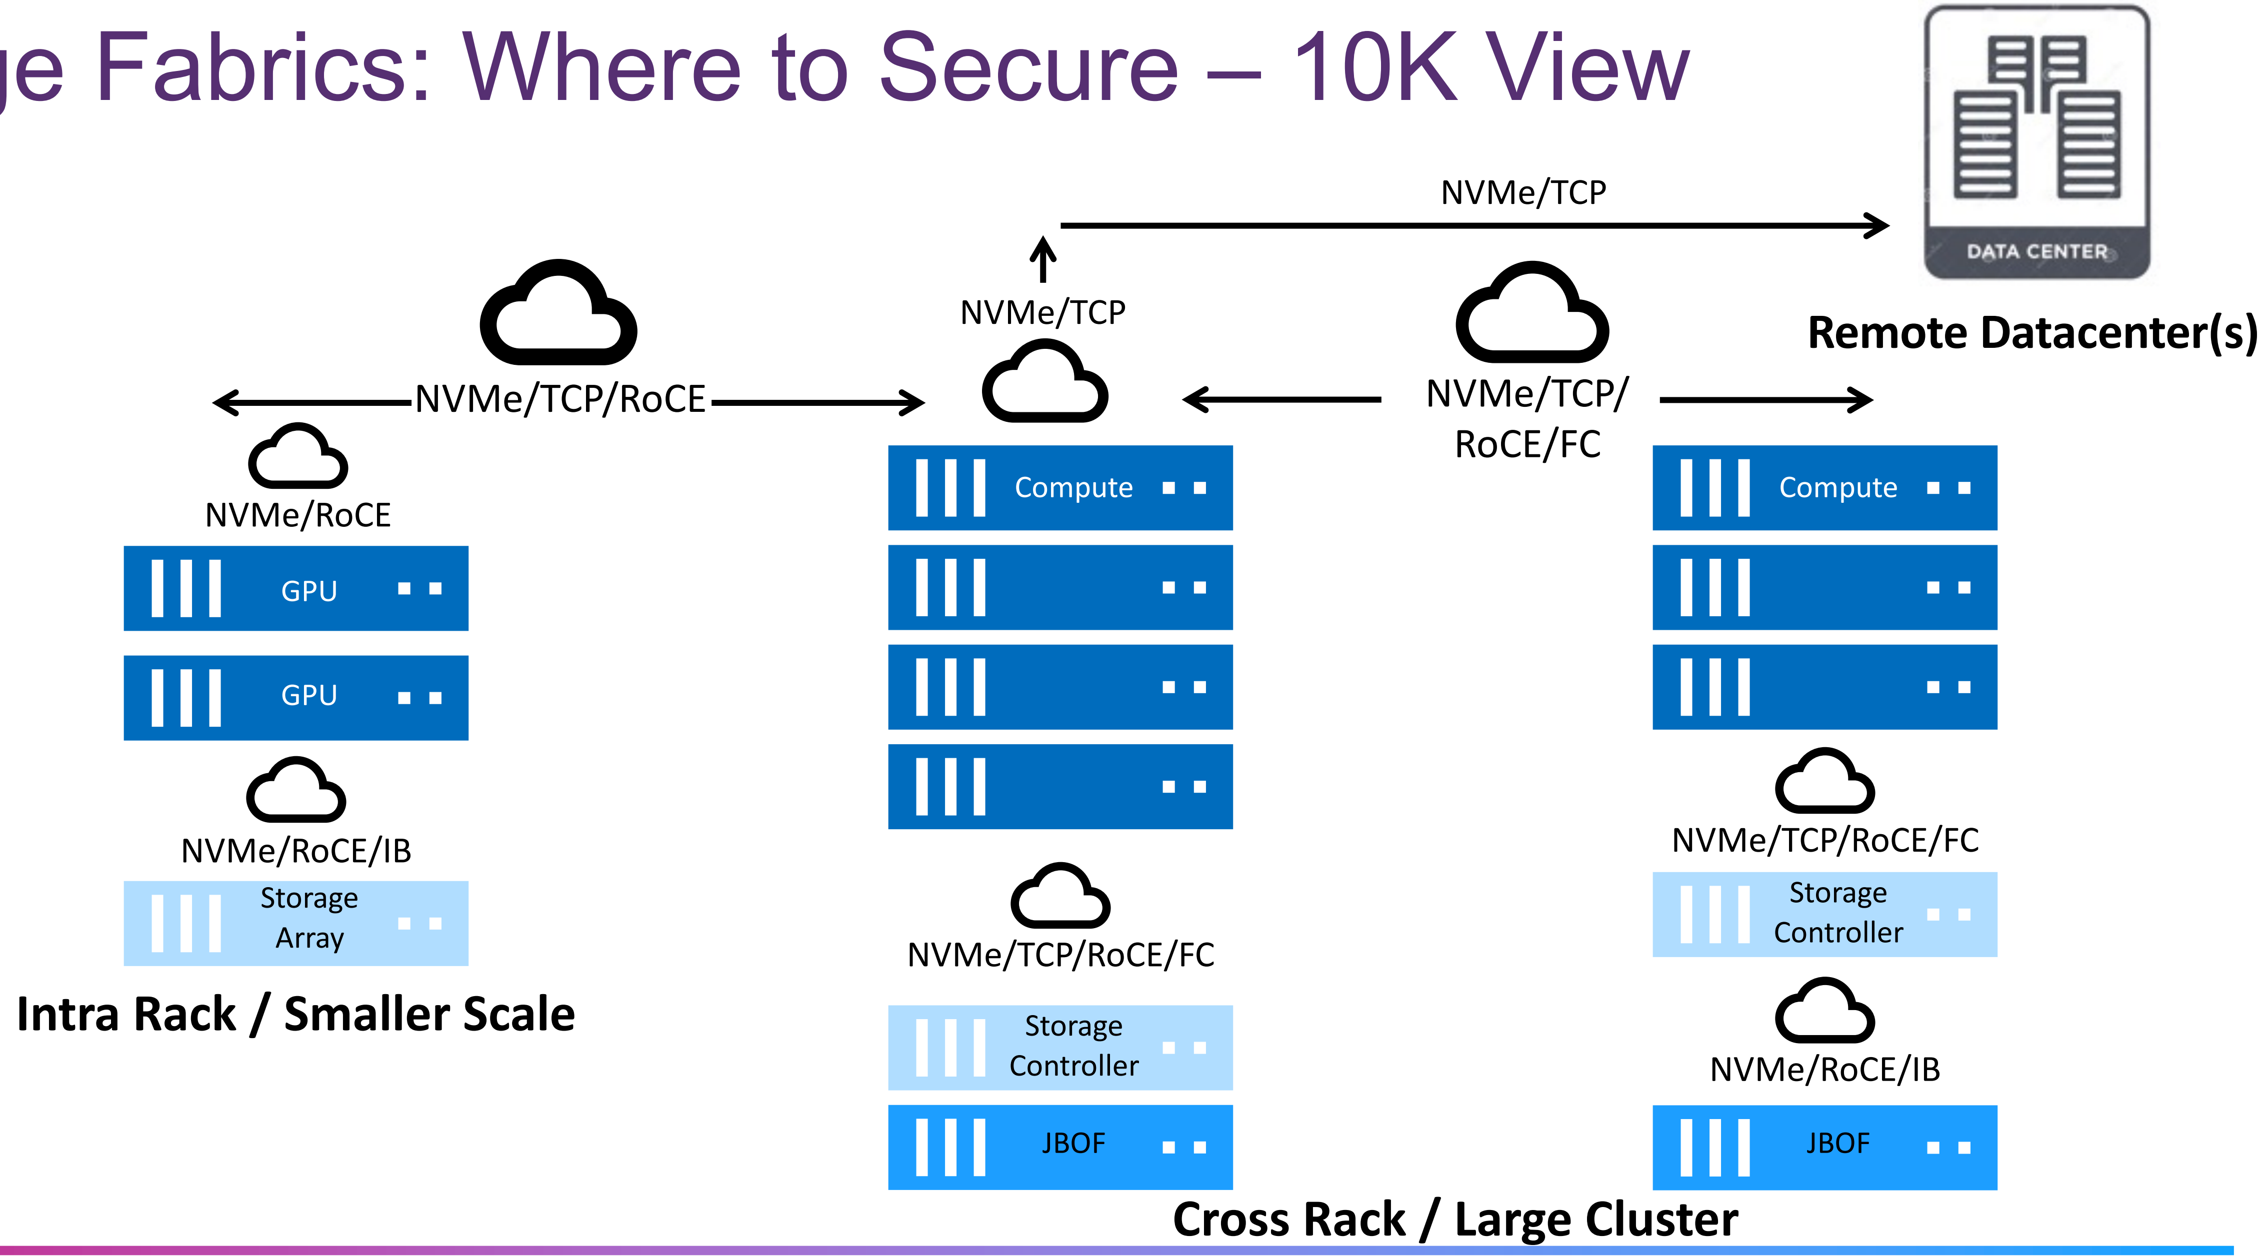

Drivers for Storage Security

Implementing a “ZERO Trust” framework requires advanced technologies

Sensitive Verticals	New Deployment Use Cases	Regulatory
Healthcare Finance Defense Government	Multi-tenancy DR/Cloud Storage Malicious Insiders	HIPAA GDPR ISO270001

Potential DC NVMe-oF Security Threats

Sniffing Storage Traffic

Storage Masquerading

Ransomware

Session Hijacking

Must secure NVMe payloads in flight and rest

Securing Storage Area Networks Focus on NVMe-oF

Claudio DeSanti

SAN Protocols - Security Mechanism Comparison

	iSCSI	Fibre Channel	NVMe over Fabrics/IP
Storage Endpoint Authentication	**CHAP** (strong secret) **SRP** (weak secret, e.g., password) [not used in practice]	**DH-CHAP** (strong secret) **FCPAP** (weak secret, e.g., password) **FCAP** (certificates) **FC-EAP** (strong secret)	**DH-HMAC-CHAP** (strong secret)
Secure Channel (authenticated encryption & cryptographic integrity)	**IPsec** (e.g., in security gateway)	**FC ESP_Header** (IPsec-like)	**TLS** (pre-shared key) – for TCP only (not usable with RoCEv2 or iWARP) **IPsec** is also an option

SNIA. NSF | NETWORKING STORAGE

NVMe-oF Authentication Example

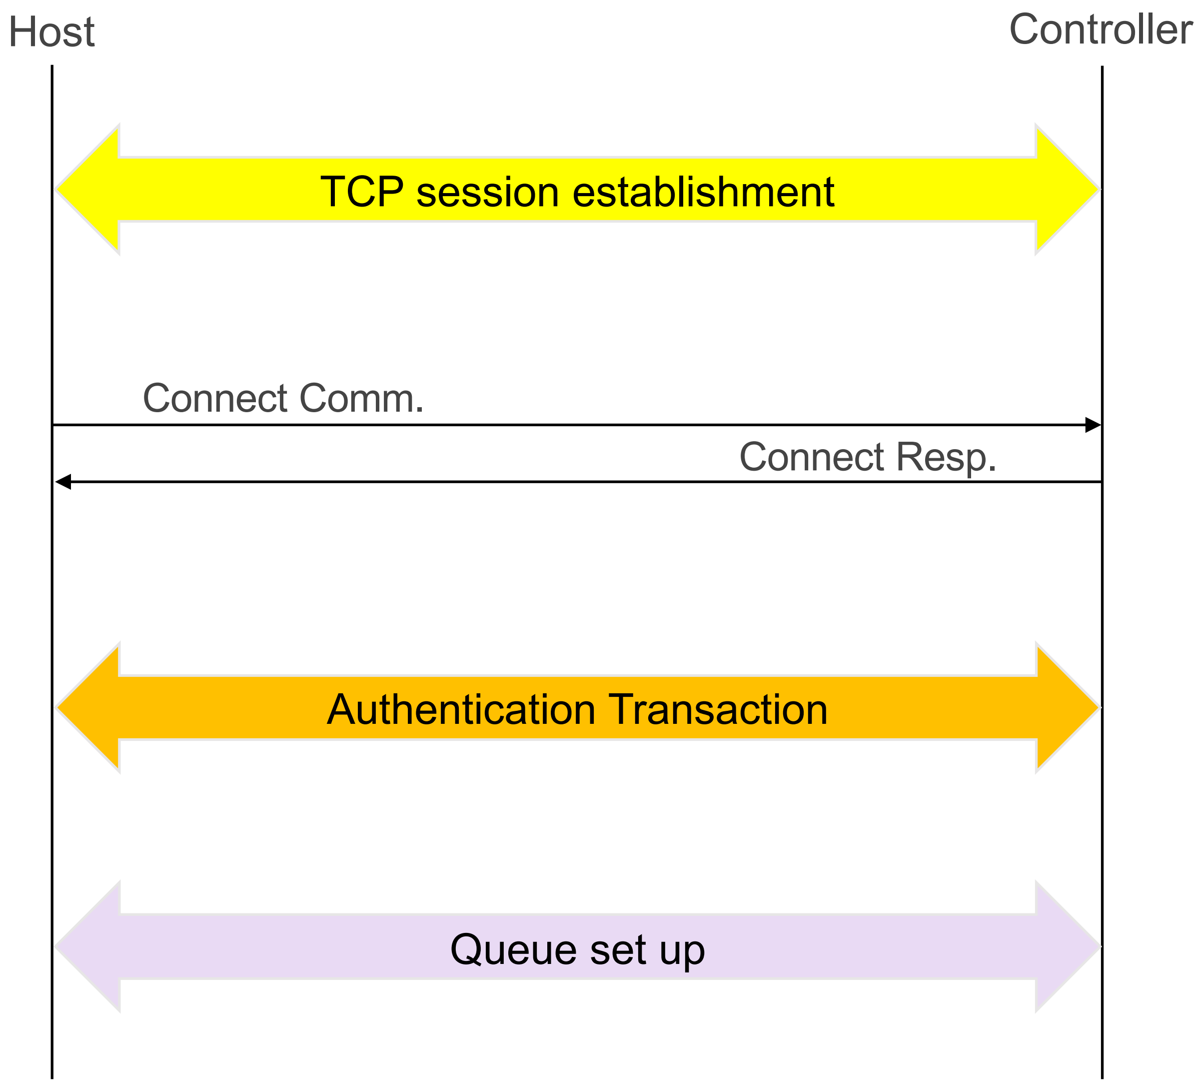

1. A TCP session is established
2. The Connect exchange is performed to set up NVMe queue and associate host to controller
3. The host performs an authentication transaction with the controller to authenticate the end-points
4. Queue is ready for subsequent operations

NVMe-oF Authentication: DH-HMAC-CHAP Protocol

- Defined in TP 8006
- Based on keys that need to be different for each NQN
- Challenge/response protocol: CHAP
  - the authenticator sends a challenge C;
  - the responder computes a response R = Hash(C || $key_{responder}$ || other things);
  - the authenticator verifies the response (or delegates verification)
- DH-HMAC-CHAP: Strengthened version of CHAP
  - DH: Diffie-Hellman, adds (optional) key exchange to frustrate eavesdroppers
  - HMAC: Hashed MAC, uses secure hash twice to improve security
- Bidirectional authentication

SNIA. NSF | NETWORKING STORAGE

DH-HMAC-CHAP Protocol: Bidirectional Authentication

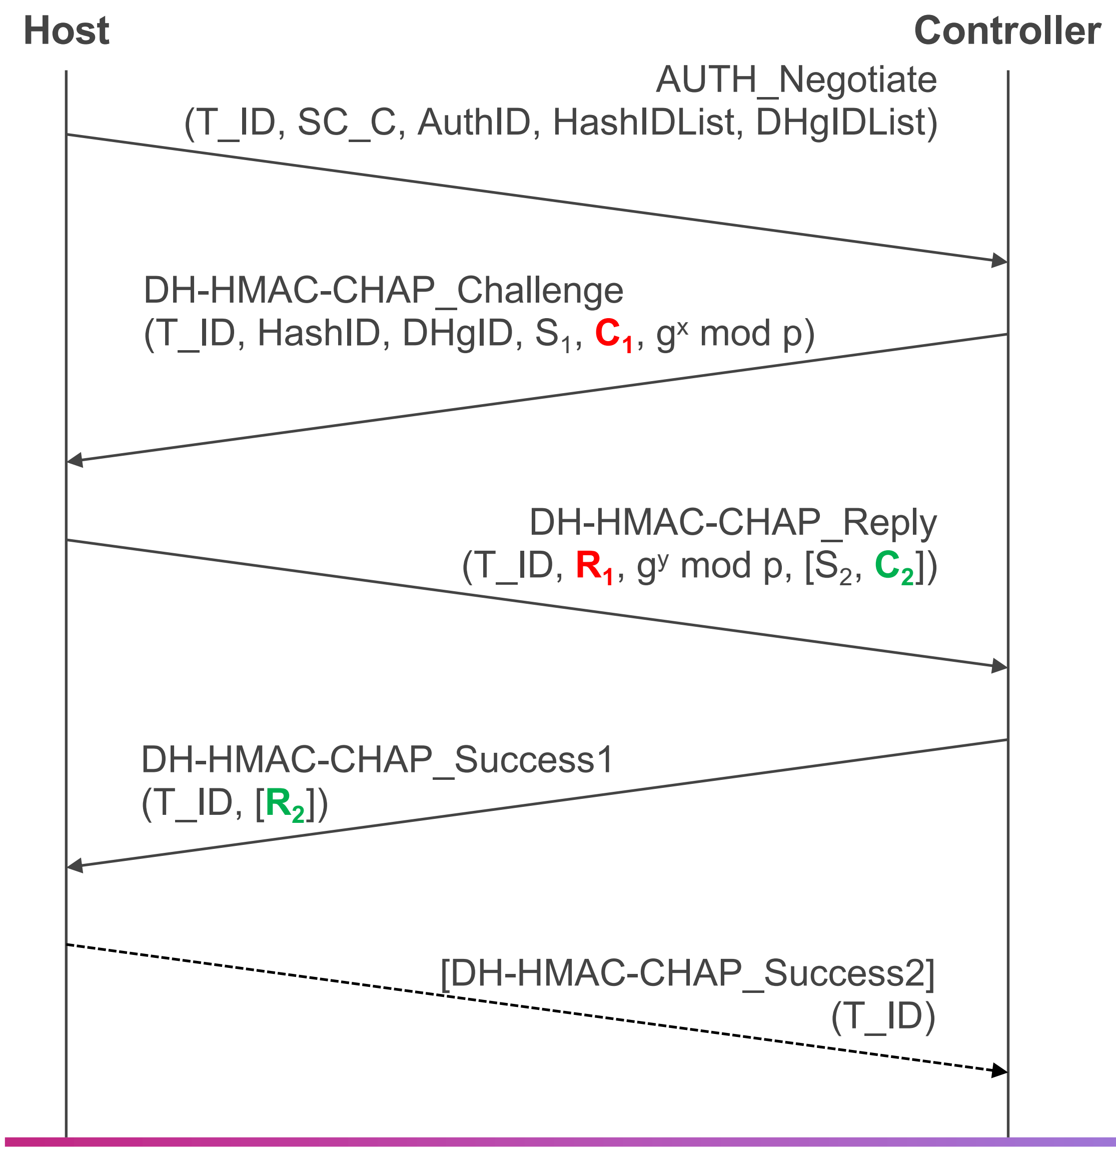

Unidirectional challenge/response protocol

- Controller C sends a challenge $\mathbf{C_1}$
- Host H computes a response $\mathbf{R_1}$ = HMAC($\mathbf{K_h}$, $\mathbf{C_1}$ || other things)
- Controller C verifies the response
- Unidirectional authentication (controller authenticates host)

Getting bidirectional authentication

- H sends a challenge $\mathbf{C_2}$
- C computes a response $\mathbf{R_2}$ = HMAC($\mathbf{K_c}$, $\mathbf{C_2}$ || other things)
- H verifies the response
- Unidirectional authentication (host authenticates controller)

Verification:

- Controller computes $\mathbf{R_1}$**'** and check if it matches the received $\mathbf{R_1}$
- Host computes $\mathbf{R_2}$**'** and check if it matches the received $\mathbf{R_2}$

Authentication Responses

K_S = H((g^x mod p)^y mod p)
C_{a1} = (DHgID == 0) ? $\mathbf{C_1}$: HMAC(K_S, $\mathbf{C_1}$)
$\mathbf{R_1}$ = HMAC($\mathbf{K_h}$, C_{a1} || S_1 || T_ID || SC_C || "HostHost" || NQN_h || 00h || NQN_c)

K_S = H((g^y mod p)^x mod p)
C_{a2} = (DHgID == 0) ? $\mathbf{C_2}$: HMAC(K_S, $\mathbf{C_2}$)
$\mathbf{R_2}$ = HMAC($\mathbf{K_c}$, C_{a2} || S_2 || T_ID || SC_C || "Controller" || NQN_c || 00h || NQN_h)

SNIA. NSF | NETWORKING STORAGE

SAN Protocols - Security Mechanism Comparison

	iSCSI	Fibre Channel	NVMe over Fabrics/IP
Storage Endpoint Authentication	**CHAP** (strong secret) **SRP** (weak secret, e.g., password) [not used in practice]	**DH-CHAP** (strong secret) **FCPAP** (weak secret, e.g., password) **FCAP** (certificates) **FC-EAP** (strong secret)	**DH-HMAC-CHAP** (strong secret)
Secure Channel (authenticated encryption & cryptographic integrity)	**IPsec** (e.g., in security gateway)	**FC ESP_Header** (IPsec-like)	**TLS** (pre-shared key) – for TCP only (not usable with RoCEv2 or iWARP) **IPsec** is also an option

SNIA. NSF | NETWORKING STORAGE

Secure Channel: TLS

- TLS (Transport Layer Security): Widely used secure channel protocol
  - Secure channel = authentication, confidentiality, cryptographic integrity (primary properties)
  - Typical (web) usage: Server uses certificate with TLS, client authenticates after TLS setup (e.g., TLS-protected HTTP)
- TLS versions:
  - TLS 1.0 and 1.1: Obsolete - should not be used
  - TLS 1.2: Baseline TLS version, widely implemented and used, getting replaced by TLS 1.3
  - TLS 1.3: New version, complete protocol redesign (TLS 2.0 in practice), usage rolling out
    - Support available in some security libraries (e.g. OpenSSL & LibreSSL), expanding to more
- TLS 1.3 specification for NVMe-oF/TCP: completed in TP 8011
  - TLS not specified for other NVMe-oF IP-based protocol (i.e., RDMA, e.g., RoCEv2)
  - Based on pre-shared keys (PSKs)
- NVMe-oF/TCP TLS 1.2: discouraged by TP 8011
  - Usage was specified by NVMe-oF 1.1 standard

Why TLS 1.3?

TLS 1.2

- IANA TLS registry has 300+ cipher suite code points
  - Uncertain security properties, difficult interoperability
- Encryption starts late in the handshake
  - Client cert and target site are sent in the clear
  - Poor privacy
- Many features with known security flaws

TLS 1.3

- 5 cipher suites, all with PFS and modern algorithms
  - Consistent security properties
- Encryption starts as early as possible, hiding content length
  - Minimal set of cleartext protocol bits on the wire
  - Less user information visible to the network
- All those features are omitted from 1.3

TLS 1.3

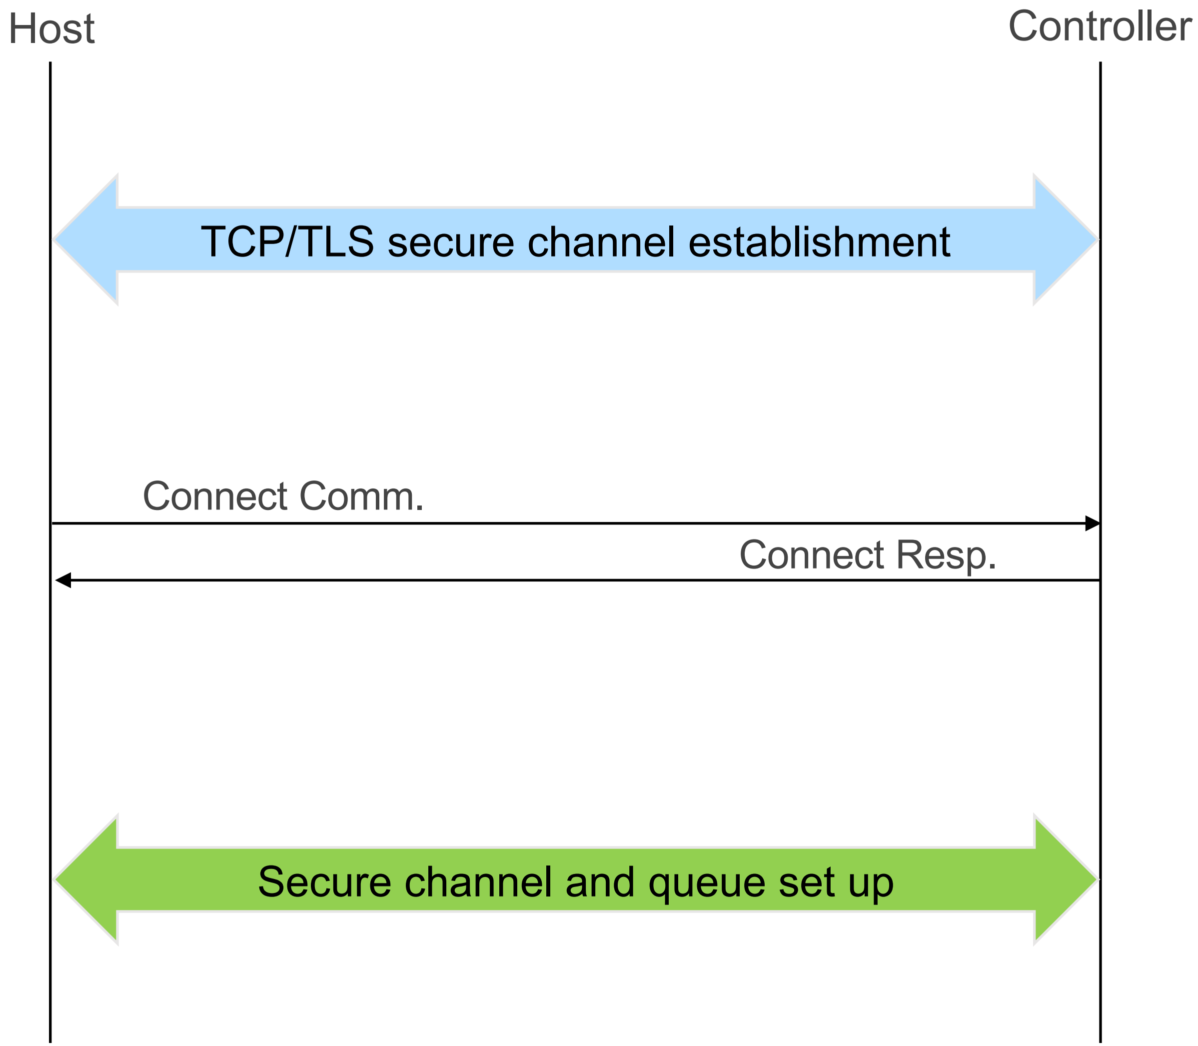

1. A TCP/TLS session negotiation is performed and a secure channel is established
2. The Connect exchange is performed to set up NVMe queue and associate host to controller
3. Secure channel and queue are set up, ready for subsequent operations

TLS Credentials

- TLS secure channel for NVMe-oF/TCP is based on pre-shared keys (PSKs)
  - In order to authenticate and establish a secure channel between themselves, two NVMe entities need to be configured with the same PSK
  - This can lead to a deployment option called 'group PSK': all NVMe entities share the same PSK
  - Big security concern (compromising a single node may allow an attacker to access all secure channels)
  - The proper way would be to have a PSK per each pair of entity that can communicate (n^2 problem)
- Authentication protocols to the rescue
  - Upon successful completion of an authentication exchange, the two involved NVMe entities generate an ephemeral shared session key (e.g., a 'PSK' computed on the fly)
  - The TLS negotiation can then be performed using a PSK derived from that shared key
    - No more need for 'group PSK'
  - Implementation result: the TCP connection begins unsecured and then transitions to secured
    - Opportunistic TLS
  - Linear problem (not anymore n^2): need just one secret per entity

SNIA. NSF | NETWORKING STORAGE

Authentication Followed by TLS

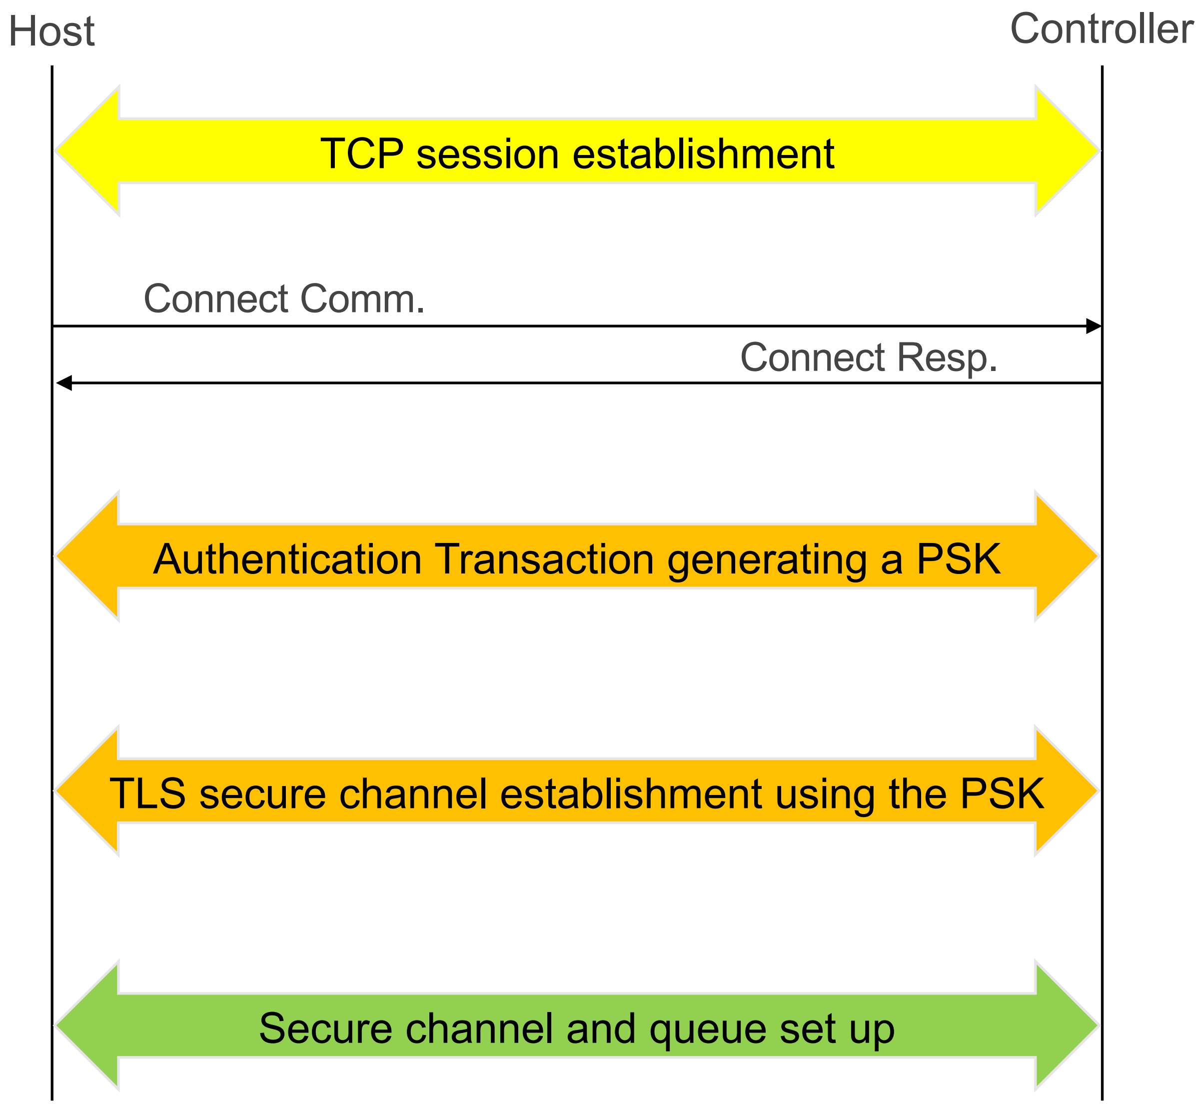

1. A TCP session is established
2. The Connect exchange is performed to set up NVMe queue and associate host to controller
3. The host performs an authentication transaction with the controller, transaction that generates a pre-shared key PSK between host and controller
4. The pre-shared key PSK is used to perform a TLS negotiation and to establish a secure channel
5. Secure channel and queue are set up, ready for subsequent operations

SNIA. NSF | NETWORKING STORAGE

NVMe-oF Security in Action
A Use Case in E-SSDs

Hrishikesh Sathawane and Eric Hibbard

NVMe Over Fabrics Architectures for Disaggregated Storage

JBOF with x86

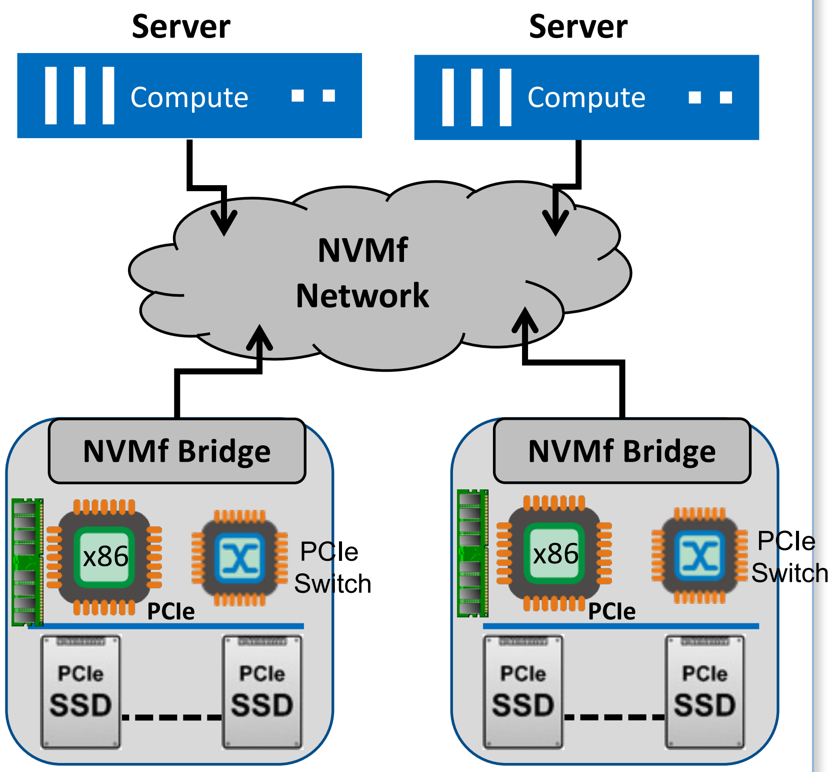

Pros
Current production
Established Ecosystem
PCIe SSDs in production

Cons
BW bottleneck
Added PCIe latency
High power

JBOF with SmartNIC

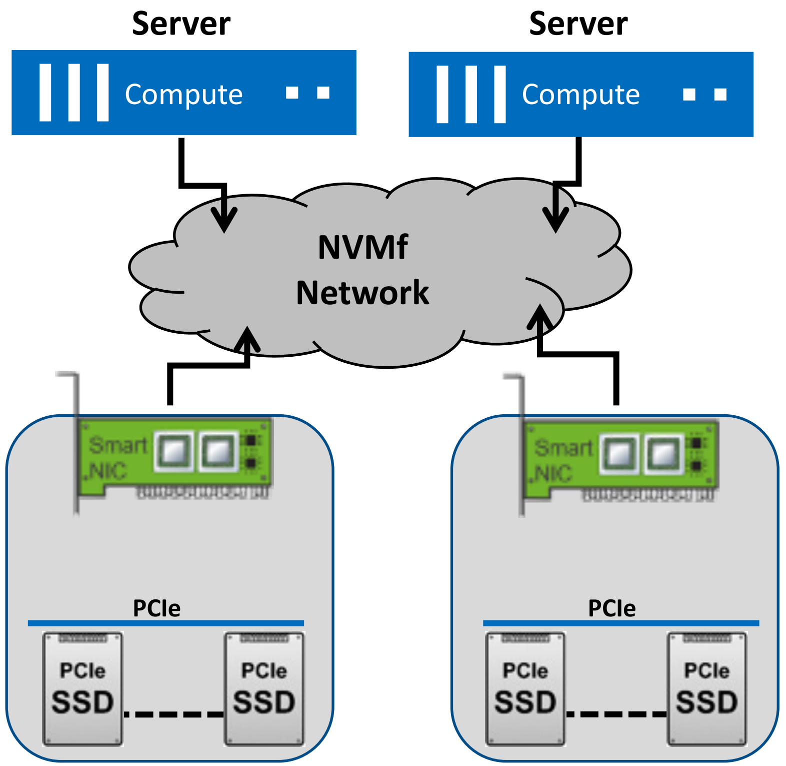

Pros
Finding some use cases
Emerging Ecosystem
PCIe SSDs in production
Lower Power

Cons
Added PCIe latency
Needs SmartNIC on both sides

E-BOF

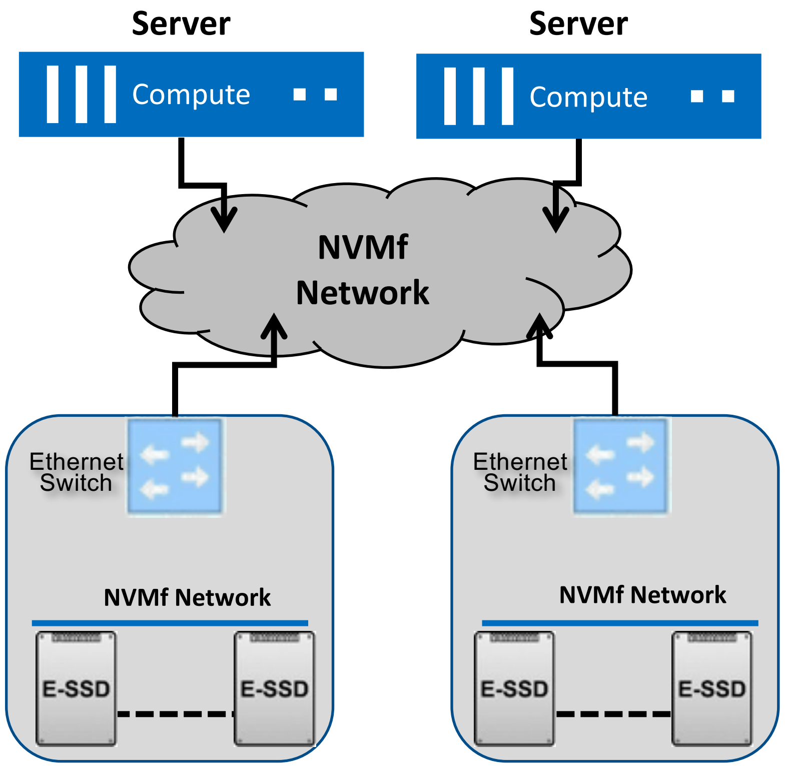

Pros
BOM cost savings
Lower latency/power
True disaggregation possible

Cons
No Ecosystem yet
E-SSDs are in PoC

Production | Engage-Short Term | Engage-Long Term

SNIA. NSF | NETWORKING STORAGE

E-SSD and the Hype Cycle

E-SSD is in the early phase of the hype cycle.

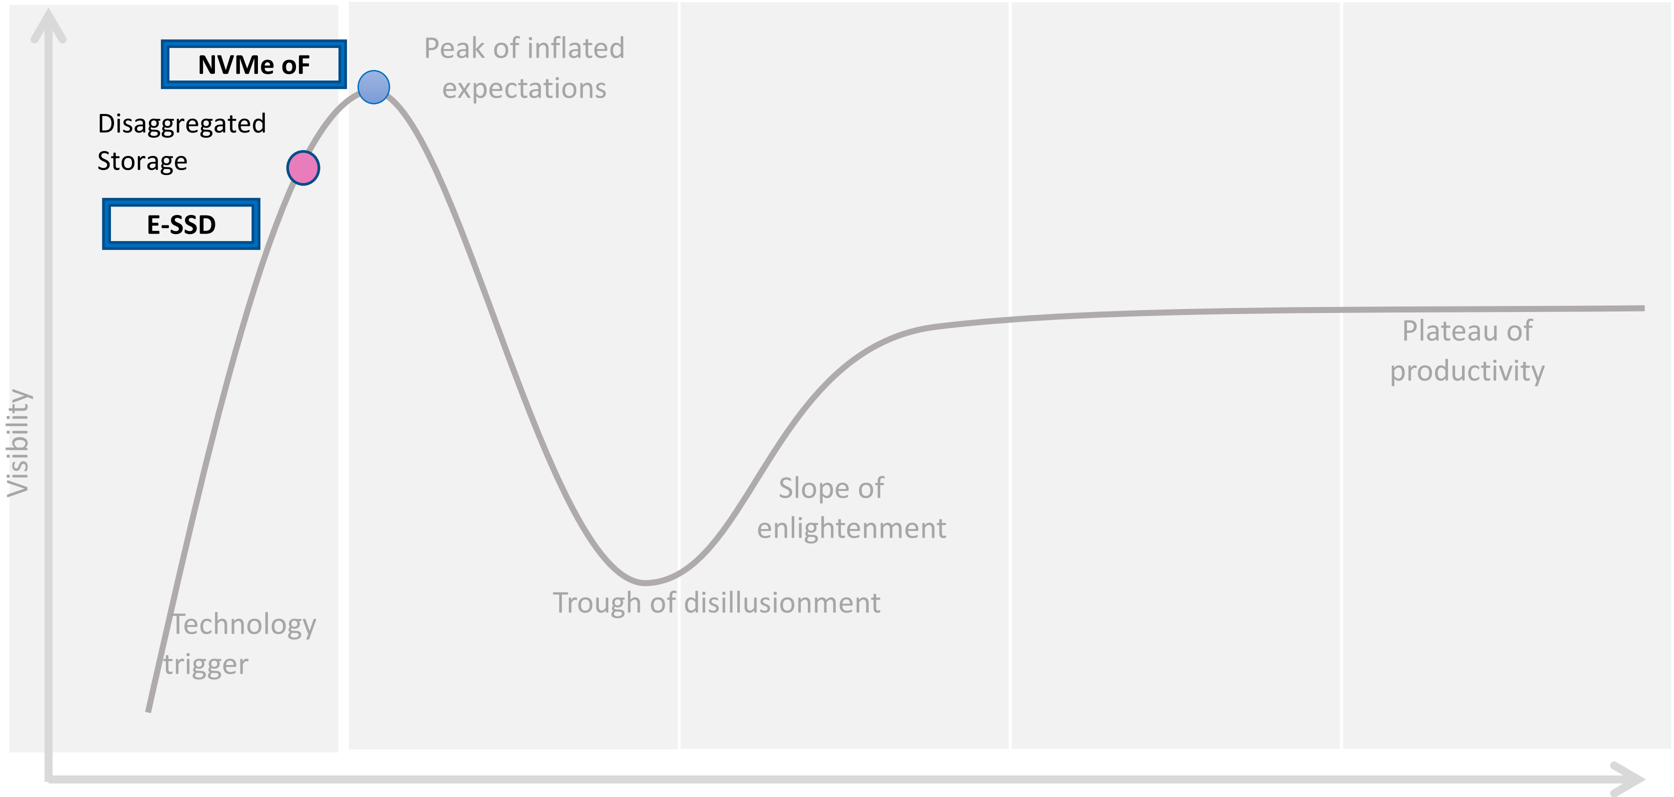

Ethernet SSD Introduction and Value Proposition

- Ethernet SSD can be one of the solution for scalable performance of JBOF
  - Also targeting for TCO saving through replacement from high CPU & BOM to Ethernet switch
  - Samsung is also continuing to study the architecture, ecosystem and benefits of Ethernet SSD

NVMe-oF Capable JBOF

- Limited performance
- High CPU power / cost

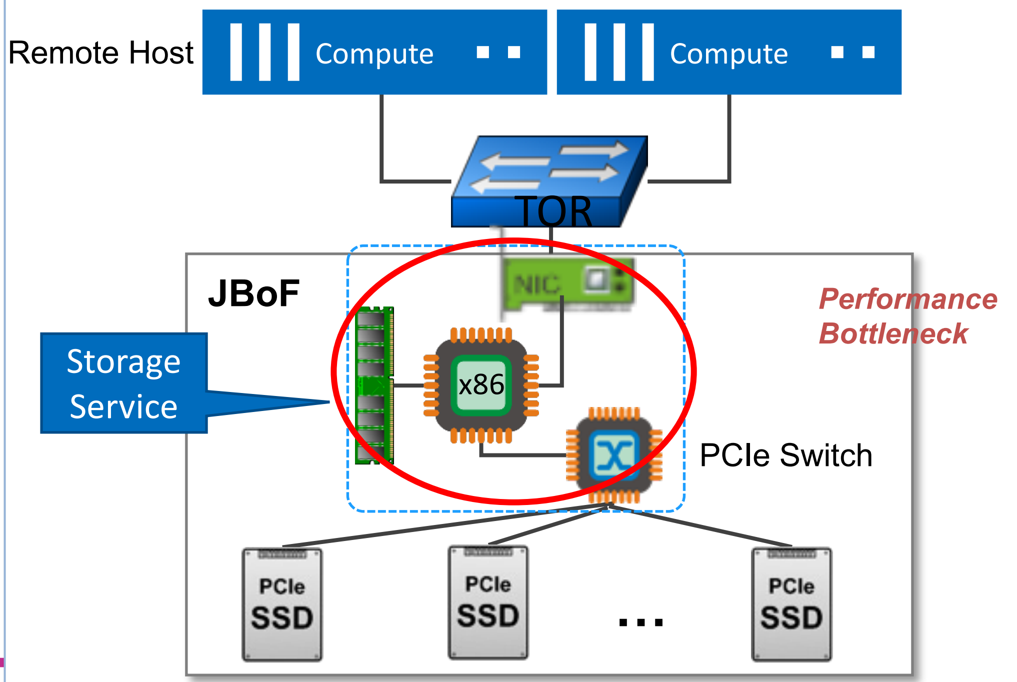

Ethernet BOF (EBOF)

- Scalable performance
- Storage service should be moved to upward

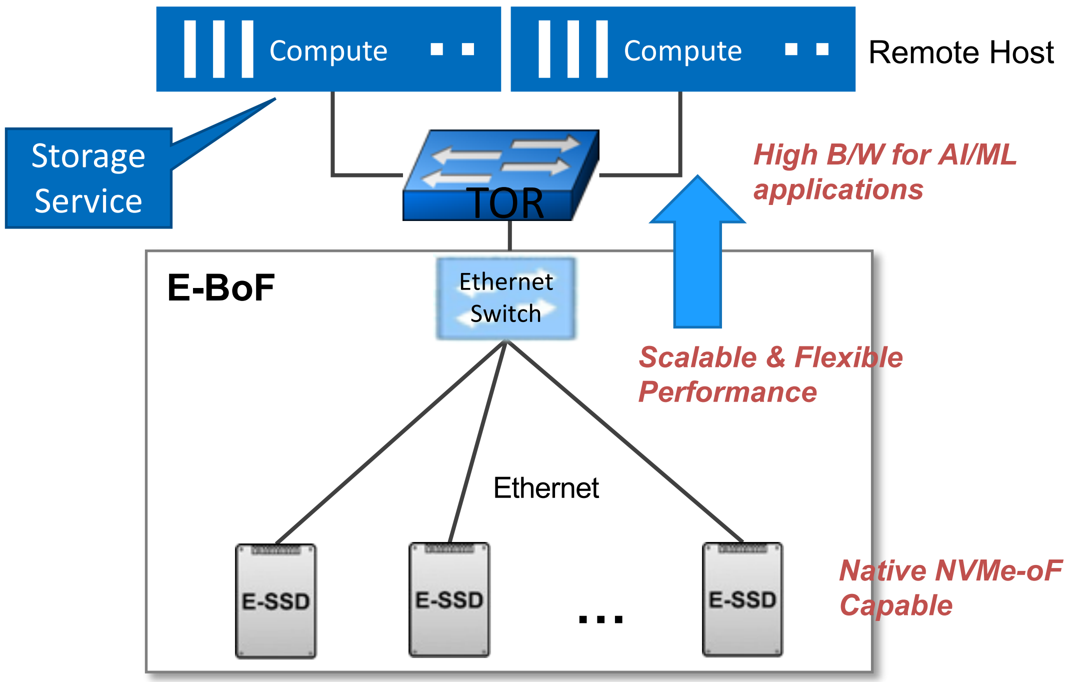

SNIA. | NETWORKING NSF | STORAGE

Advantage of NVMe-oF SSD

- NVMeoF JBOF can solve performance, scalability, and flexibility
  - Scalable and flexible data center solution through Ethernet-only infrastructure

Conventional SSD System

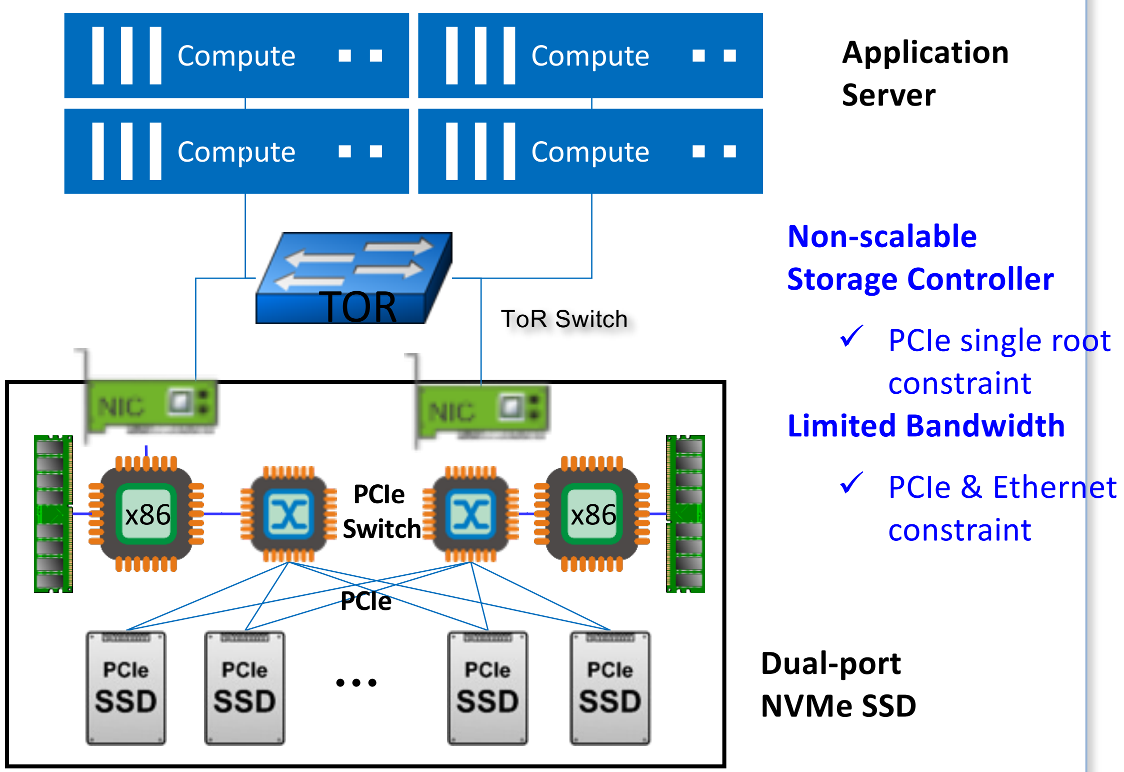

NVMeoF SSD System

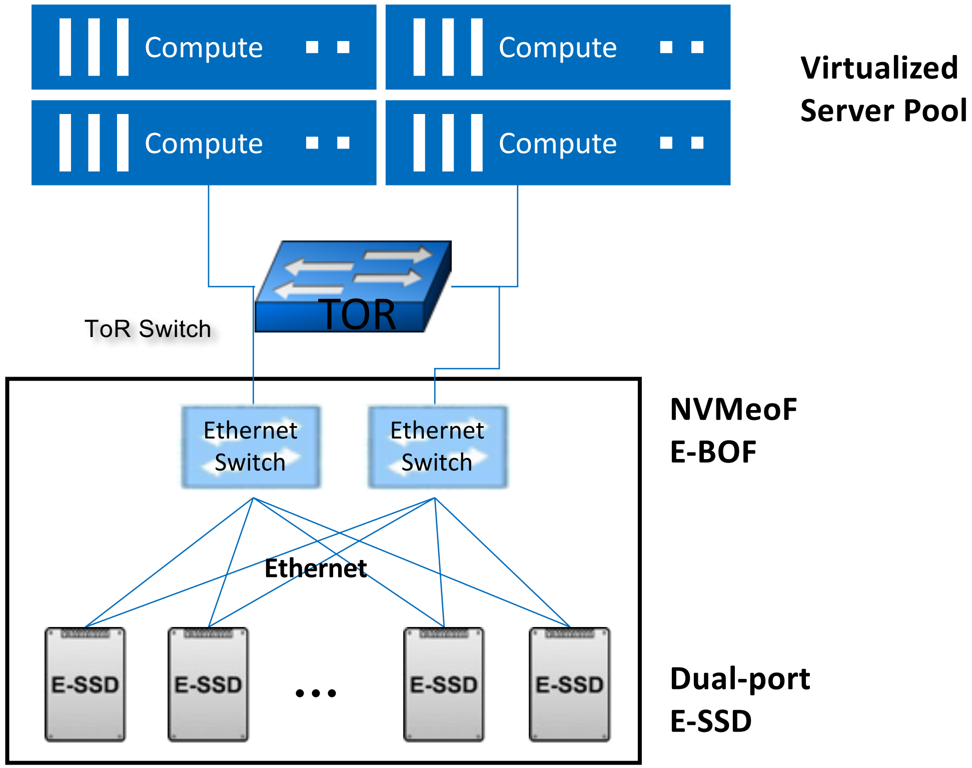

SNIA. NSF | NETWORKING STORAGE

Disaggregated Architecture

- **NIC card's essential features are offloaded to E-SSD.**
- **Storage controller can communicate with multiple E-BOFs as it is done with JBOFs.**
  - Connection and discovery are added to maintain connection in NVMe-oF specification.
  - NVMe I/O command processing is equal as PCIe based NVMe SSD for E-SSD.
  - Additional target configuration may not be needed, since configured information are saved in E-SSD.
- **SW modification**
  - Schema of JBOF for Redfish can be changed.

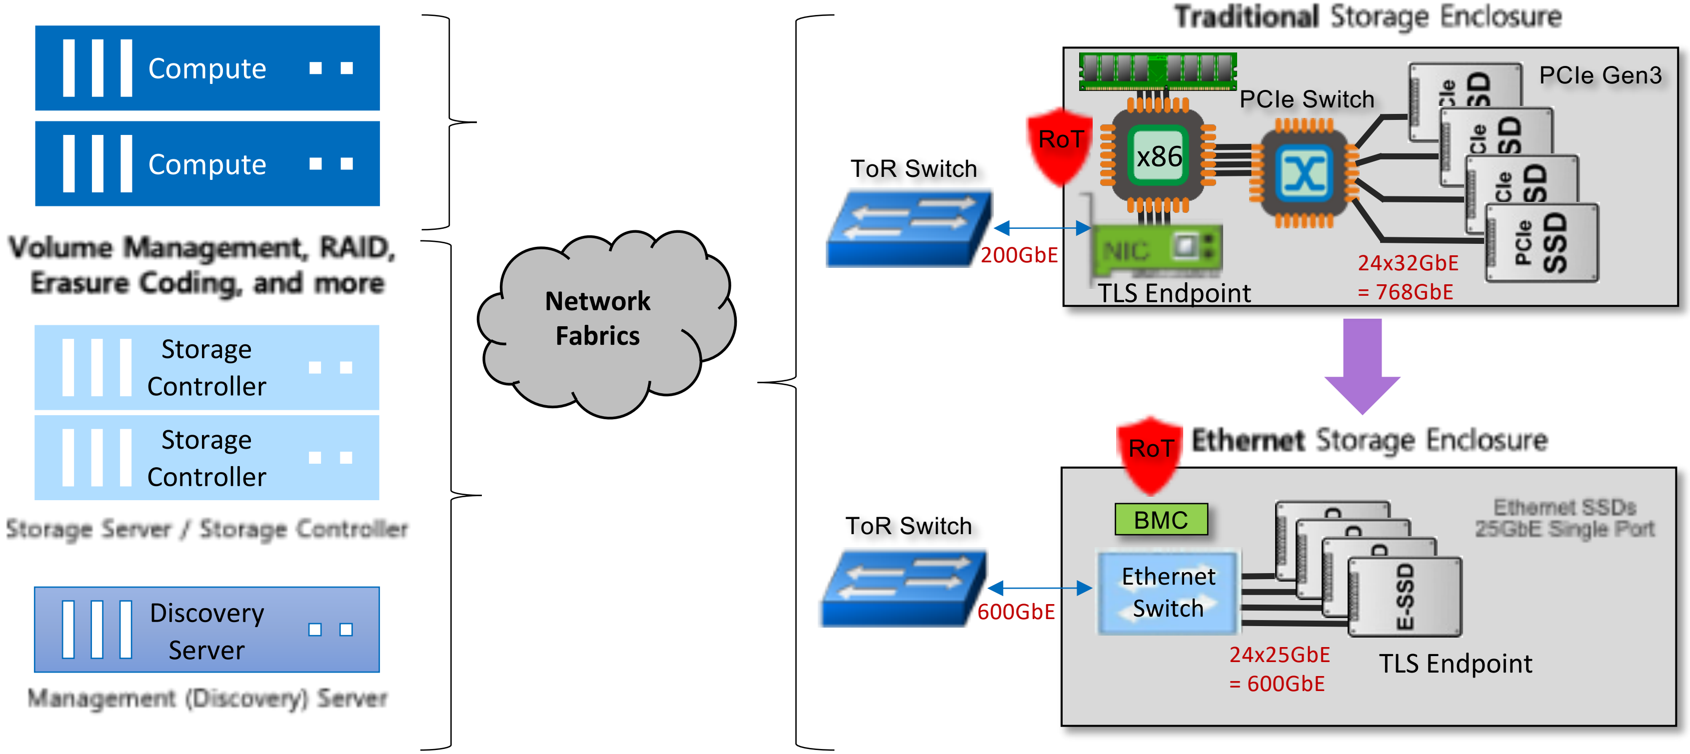

NIC + CPU
- IP address configuration
- Manage Connections
- Network Management

PCIe SSD
- NVMe I/O Processing

BMC (Expected)
- E-SSD IP address configuration

E-SSD (offloaded)
- IP address configuration (manual)
- Manage Connections
- Network Management
- NVMe I/O Processing

SNIA. NSF | NETWORKING STORAGE

Disaggregated Architecture

- **NIC card's essential features are offloaded to E-SSD.**
- **Storage controller can communicate with multiple E-BOFs as it is done with JBOFs.**
  - Connection and discovery are added to maintain connection in NVMe-oF specification.
  - NVMe I/O command processing is equal as PCIe based NVMe SSD for E-SSD.
  - Additional target configuration may not be needed, since configured information are saved in E-SSD.
- **SW modification**
  - Schema of JBOF for Redfish can be changed.

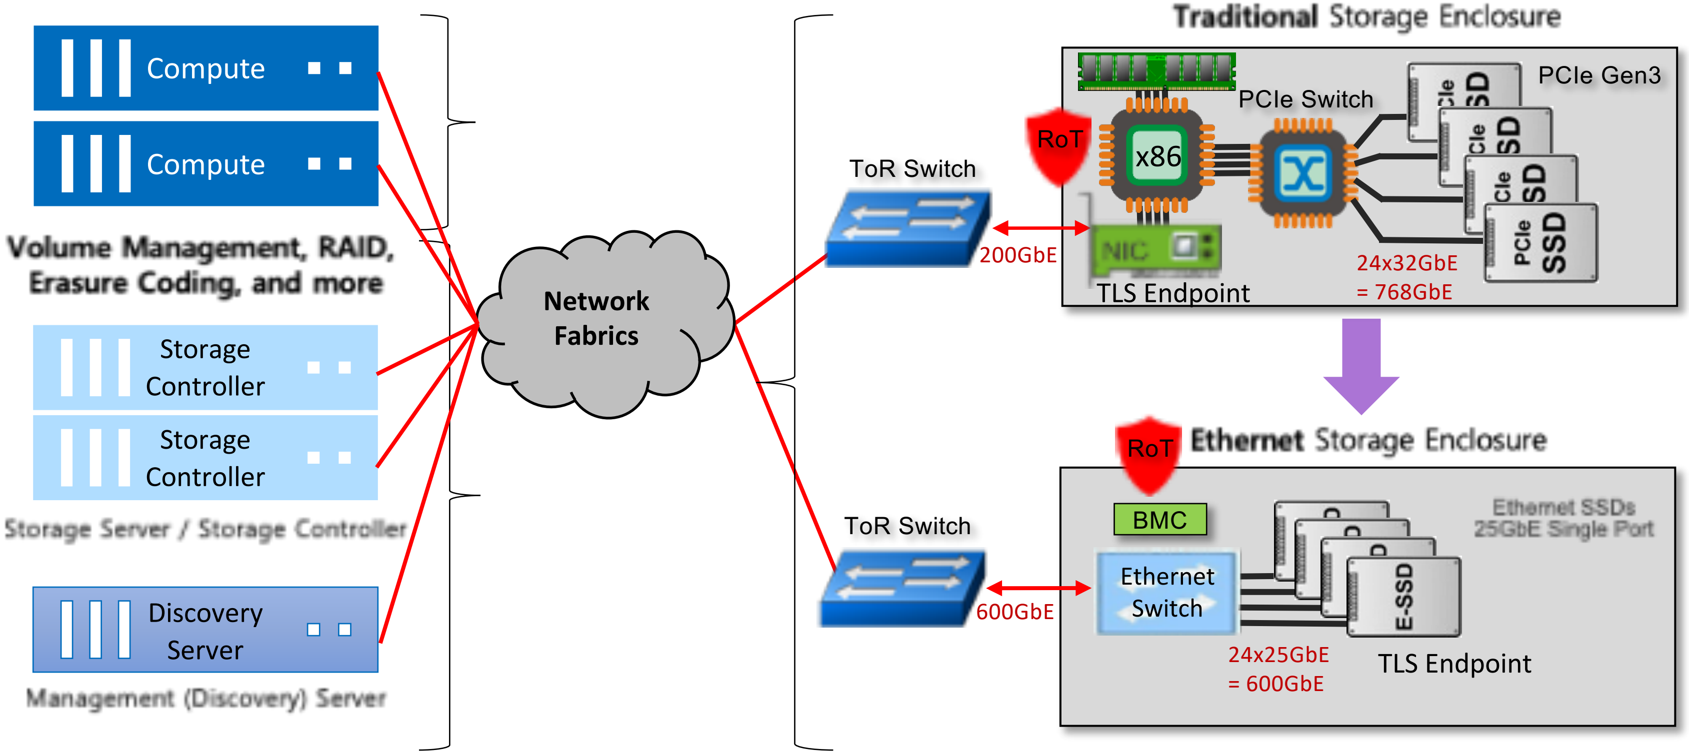

NIC + CPU
- IP address configuration
- Manage Connections
- Network Management

PCIe SSD
- NVMe I/O Processing

BMC **(Expected)**
- E-SSD IP address configuration

E-SSD **(offloaded)**
- IP address configuration (manual)
- Manage Connections
- Network Management
- NVMe I/O Processing

SNIA. NSF | NETWORKING STORAGE

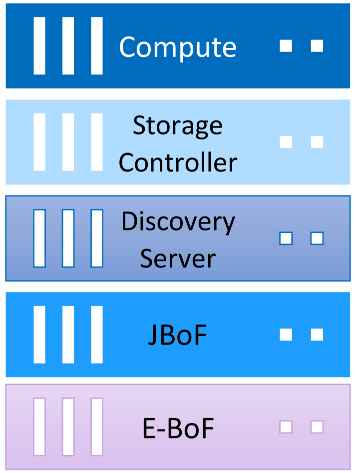
Compute
Storage Controller
Discovery Server
JBoF
E-BoF

Questions

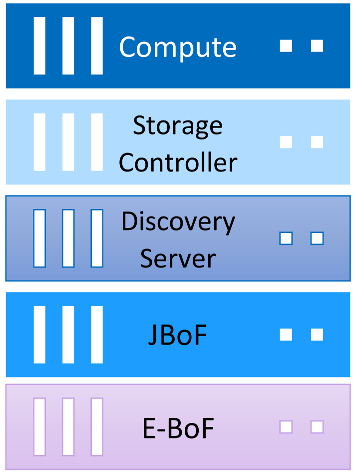
Compute
Storage Controller
Discovery Server
JBoF
E-BoF

Other Resources

- Webcast: NVMe over Fabrics: Looking Beyond Performance Hero Numbers
  - http://bit.ly/NVMeoFHero
- Multiple resources: SNIA Geek Out on NVMe-oF
  - https://bit.ly/GeekNVMeoF
- Blog: NVMe over Fabrics for Absolute Beginners
  - https://bit.ly/3aKf3JS

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

SNIA. NSF | NETWORKING STORAGE