



Open. Together.



OCP
SUMMIT

Storage

Ethernet Is The New Fabric-of-Choice for Storage Expansion

Shahar Noy
Sr Director Marketing
Marvell



OPEN
COMMUNITY®

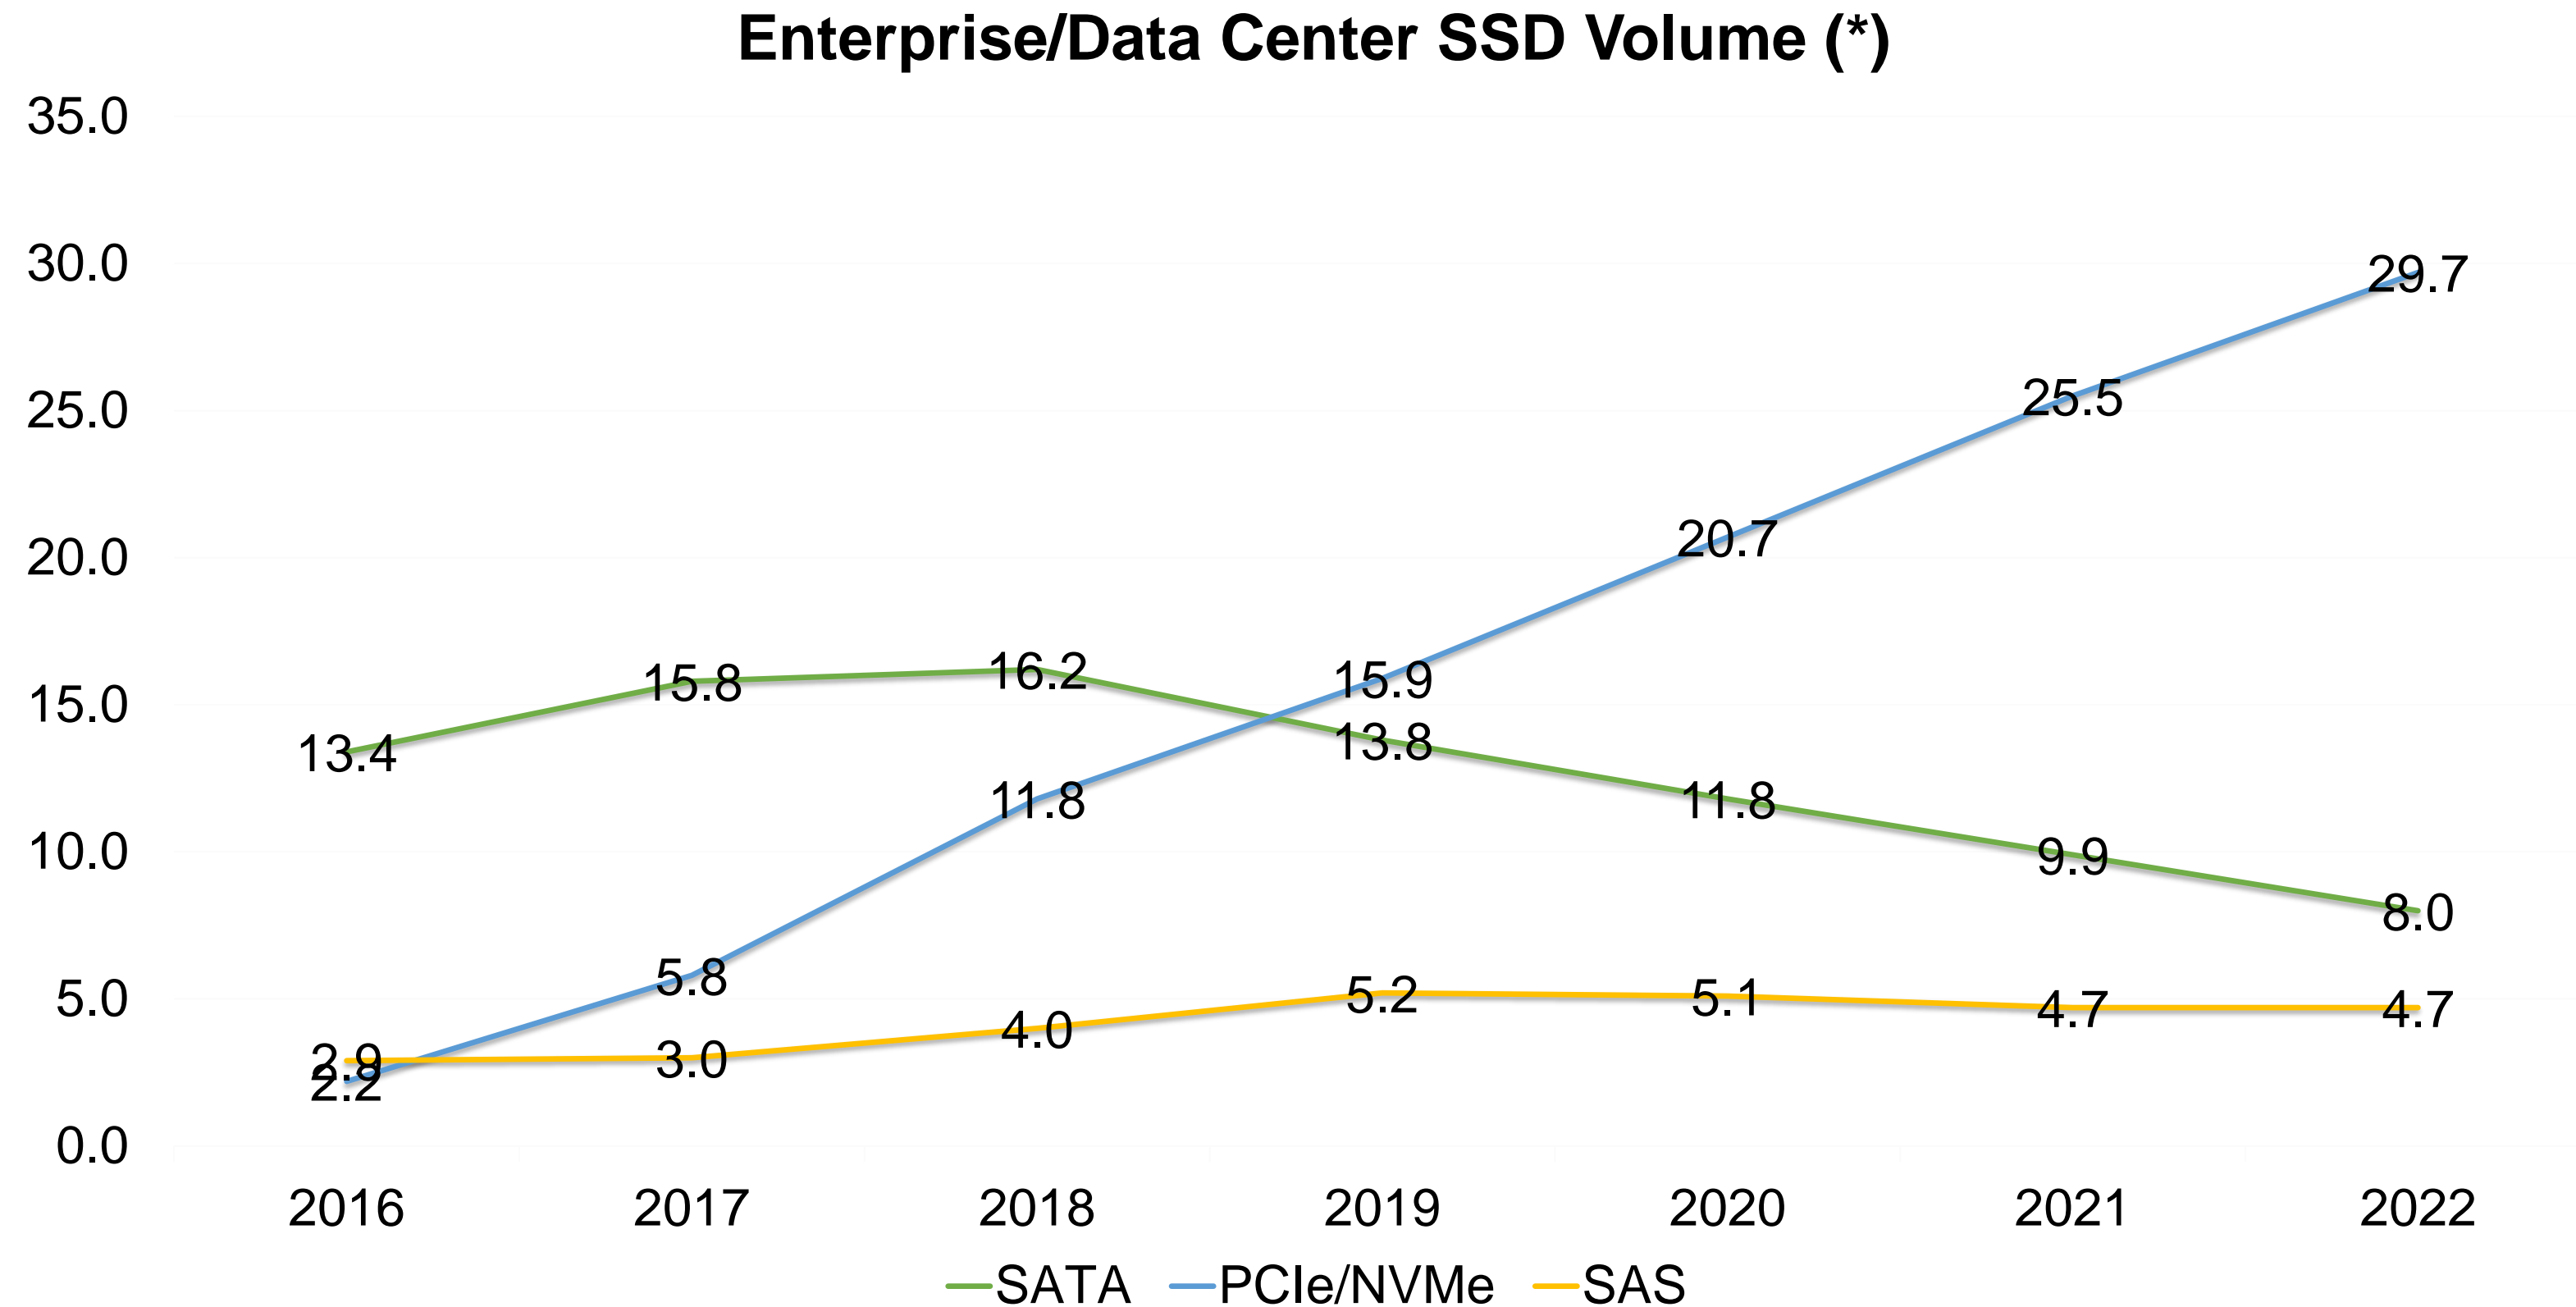


Open. Together.

Who is the king of SSDs?



STORAGE

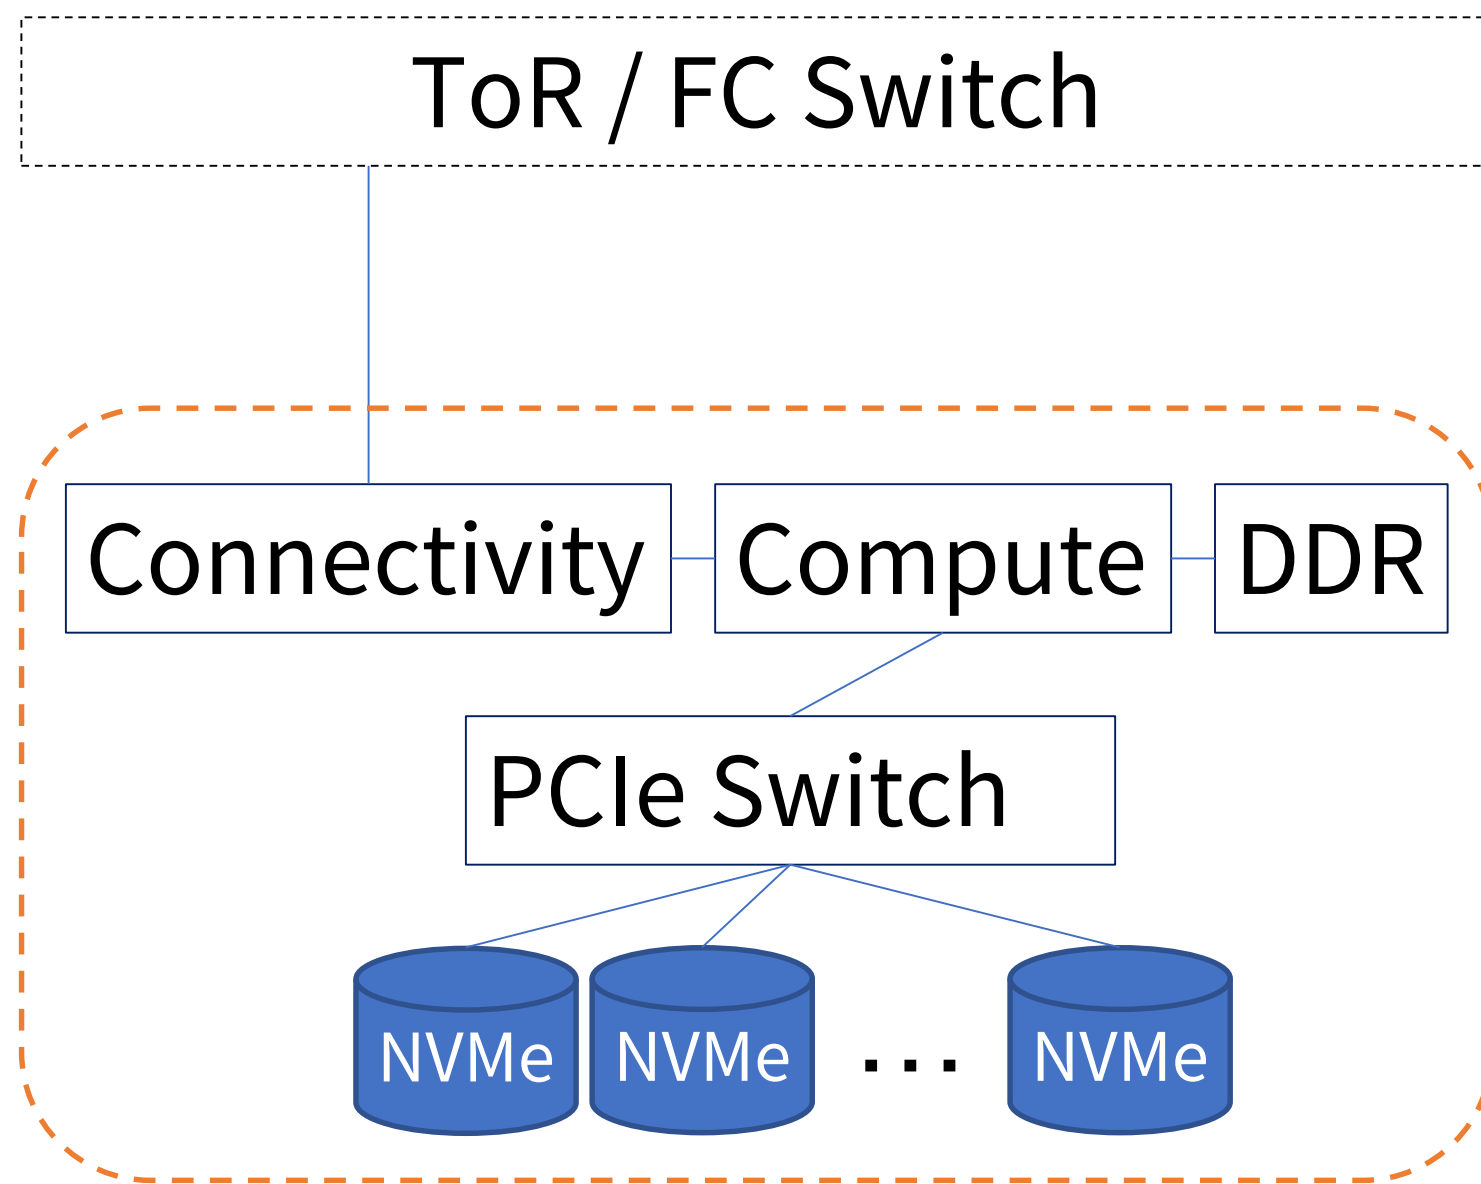


If NVMe is the future, what fabric would replace the traditional SAS/SATA backplane?

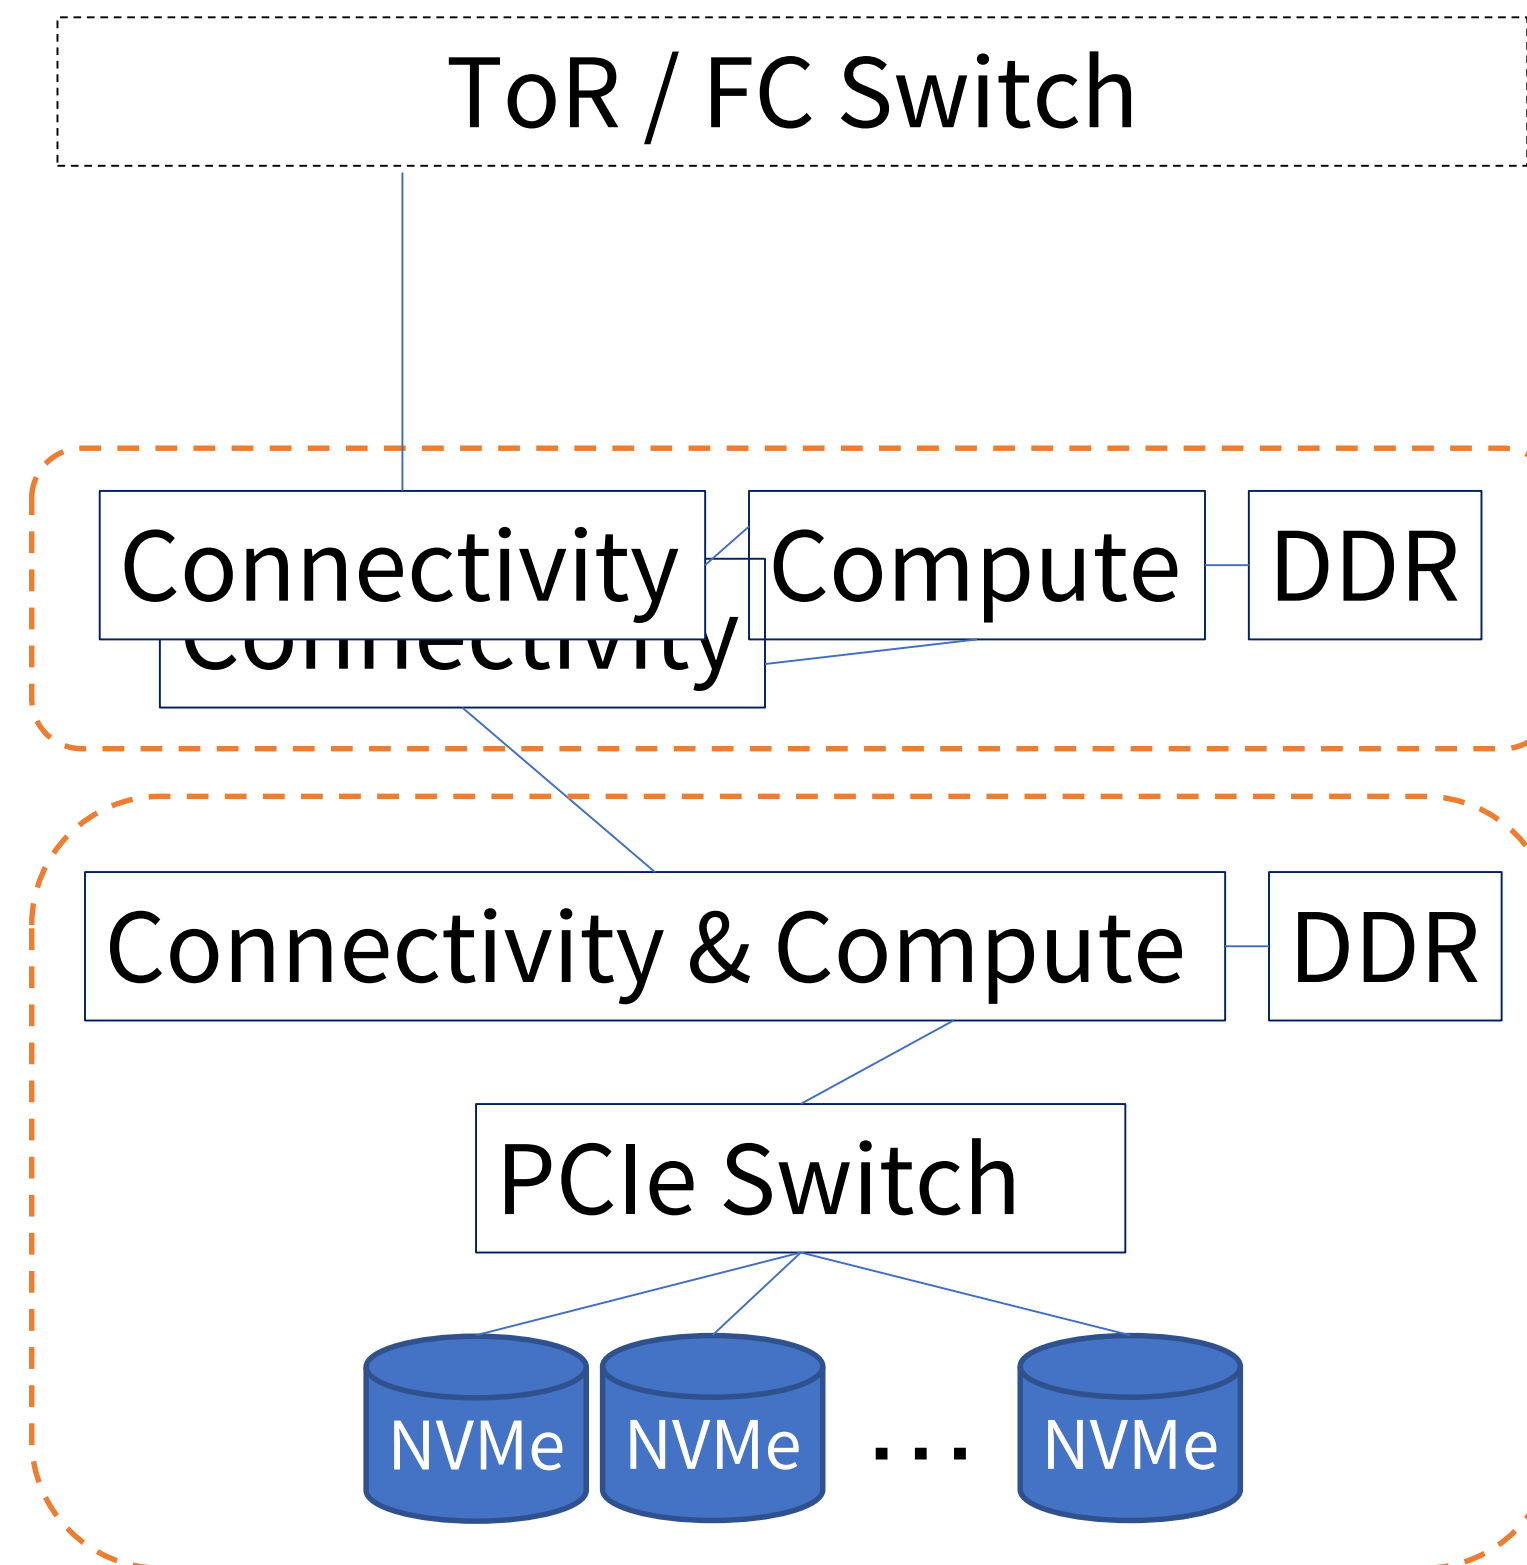
* Marvell market view

Options for NVMe Expansion

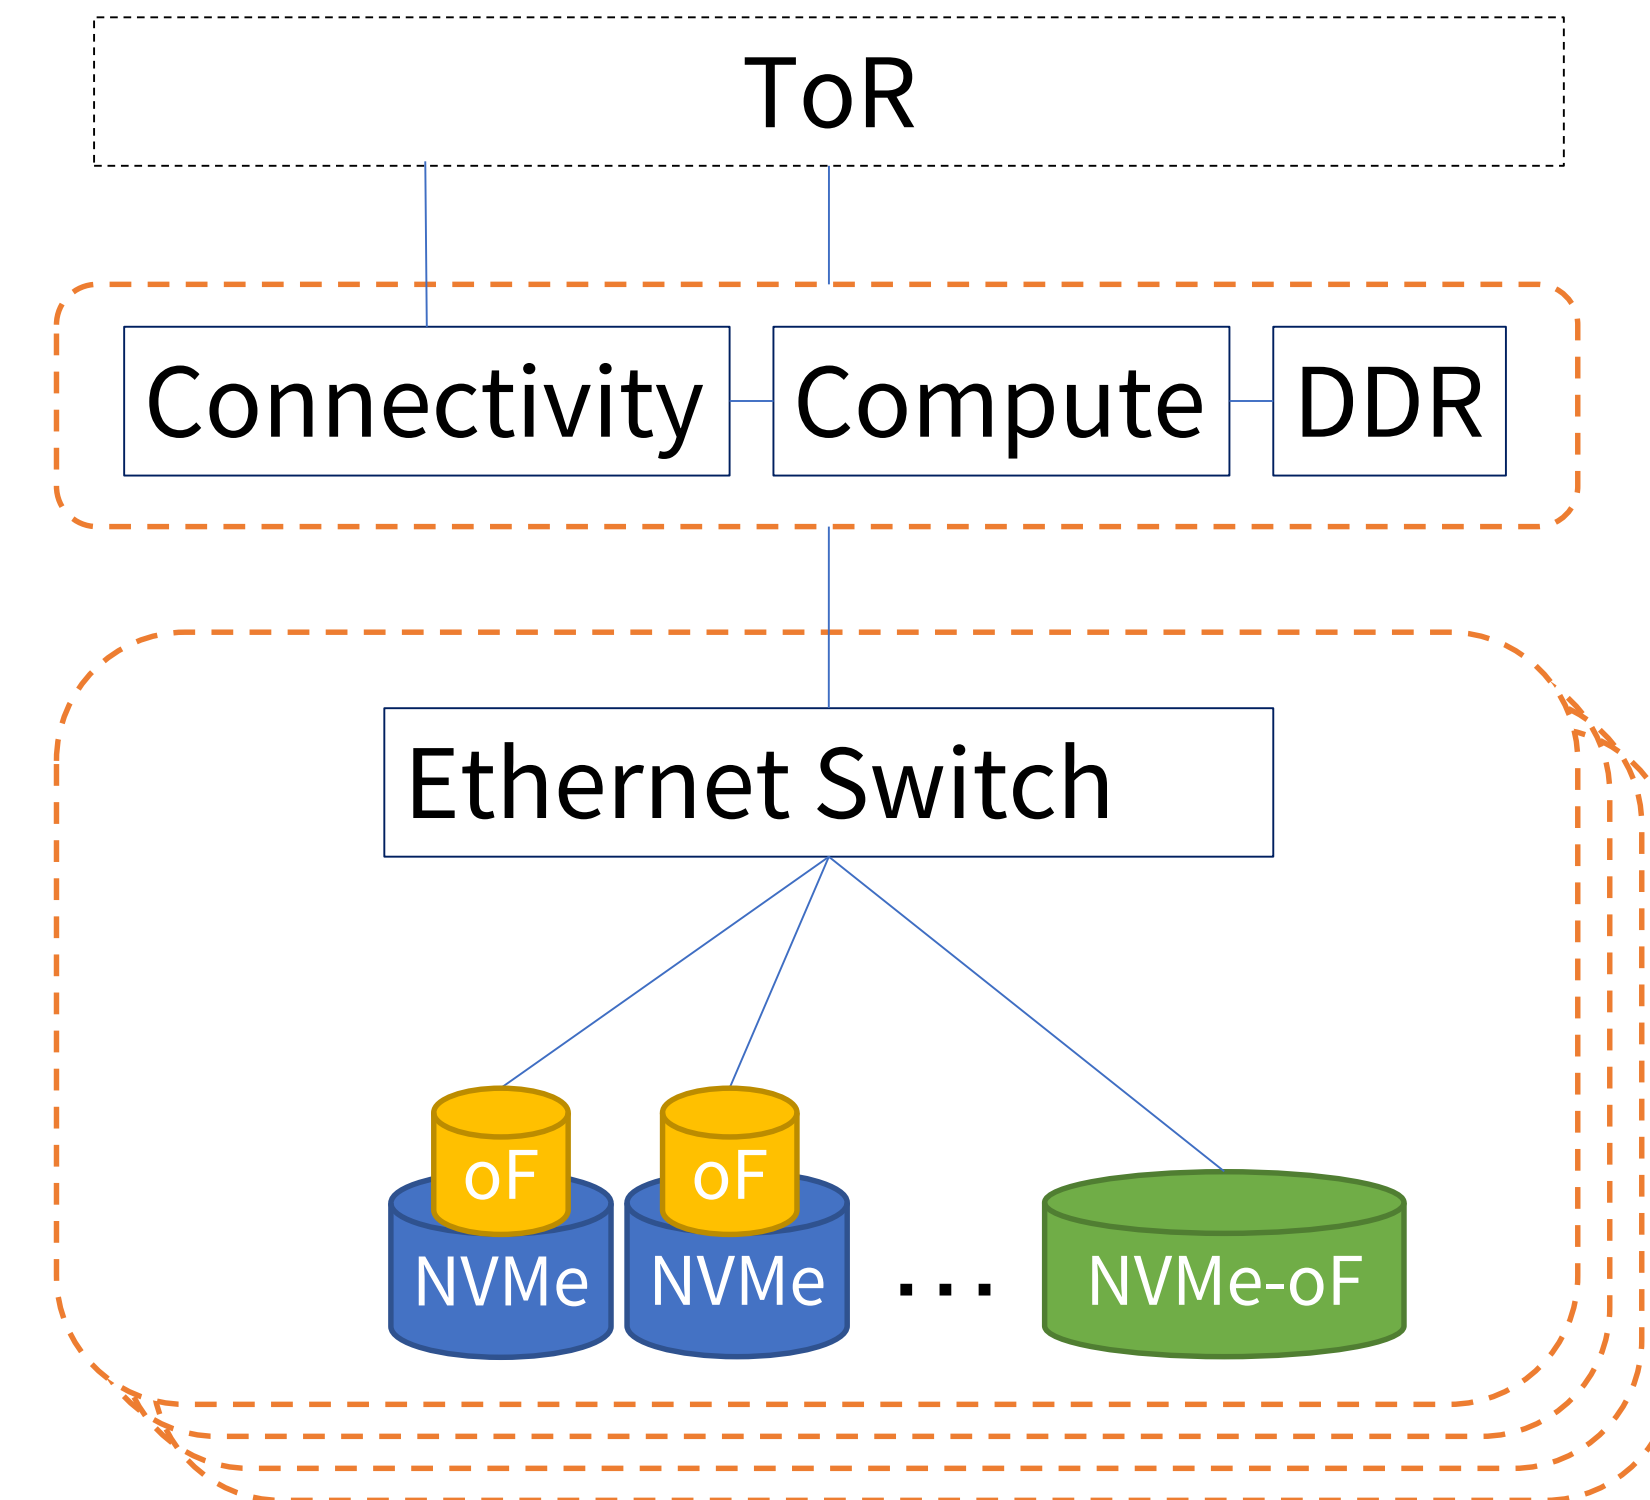
NVMe All-Flash-Array



Composable NVMe (1.0) FBOF



Composable NVMe (2.0) **EBOF**



NVMe expansion: Fabric Comparison

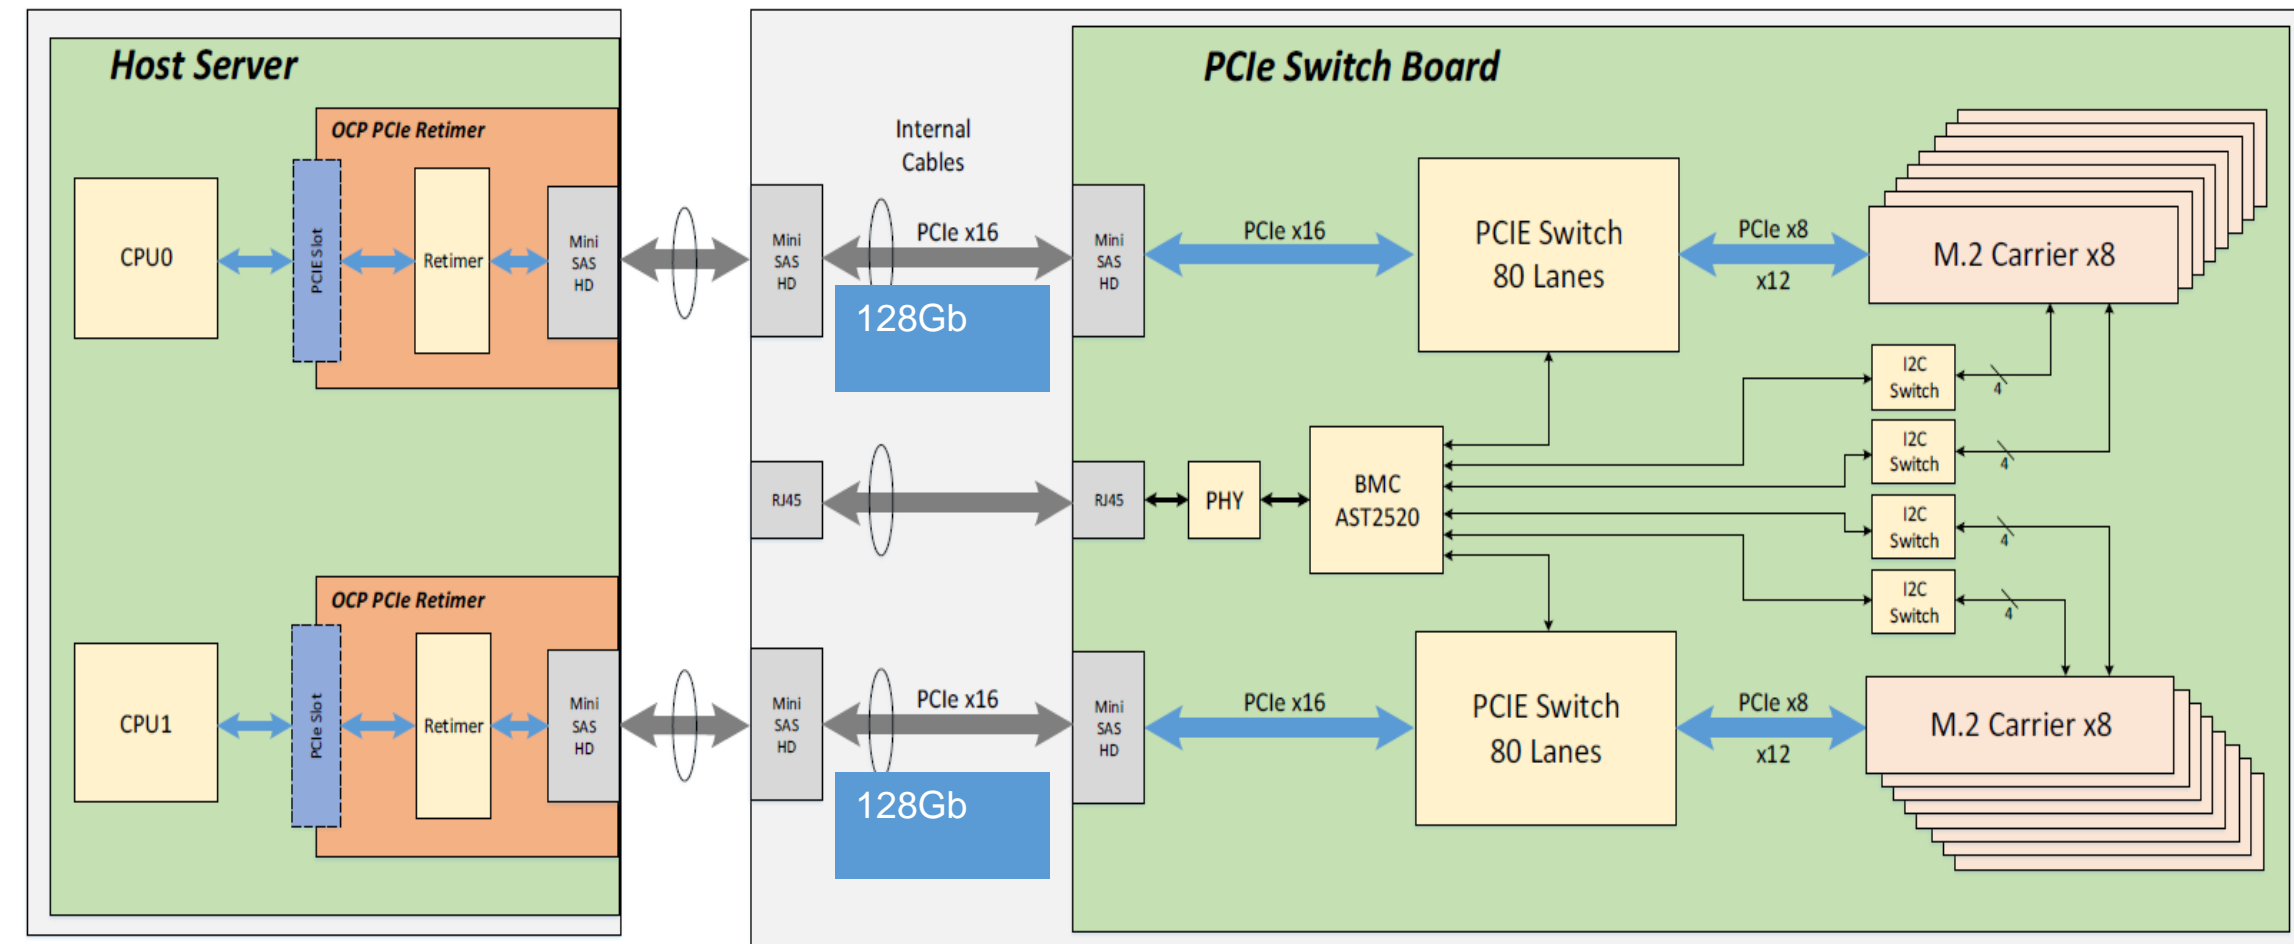
	Ethernet (EBOF)	PCIe (FBOF)
Connectivity	NVMe-oF SSD Ethernet Switch	Compute, DDR, PCIe Switch
Switch Throughput	6.4Tb/s (25Gb/s ports) 12.8Tb/s (50Gb/s ports)	3TB/s (Gen3 x96) 6TB/s (Gen4 x96)
Latency	SSD + Transport (80~400us)	SSD only (60~200us)
IOPS	SSDs (24SSDs = 16MIOPS)	Bound by Compute/DDR (24SSDs = 10MIOPS)
Oversubscription	No	Yes
Daisy chain shelves	Yes	No

Proposed Architectural Changes

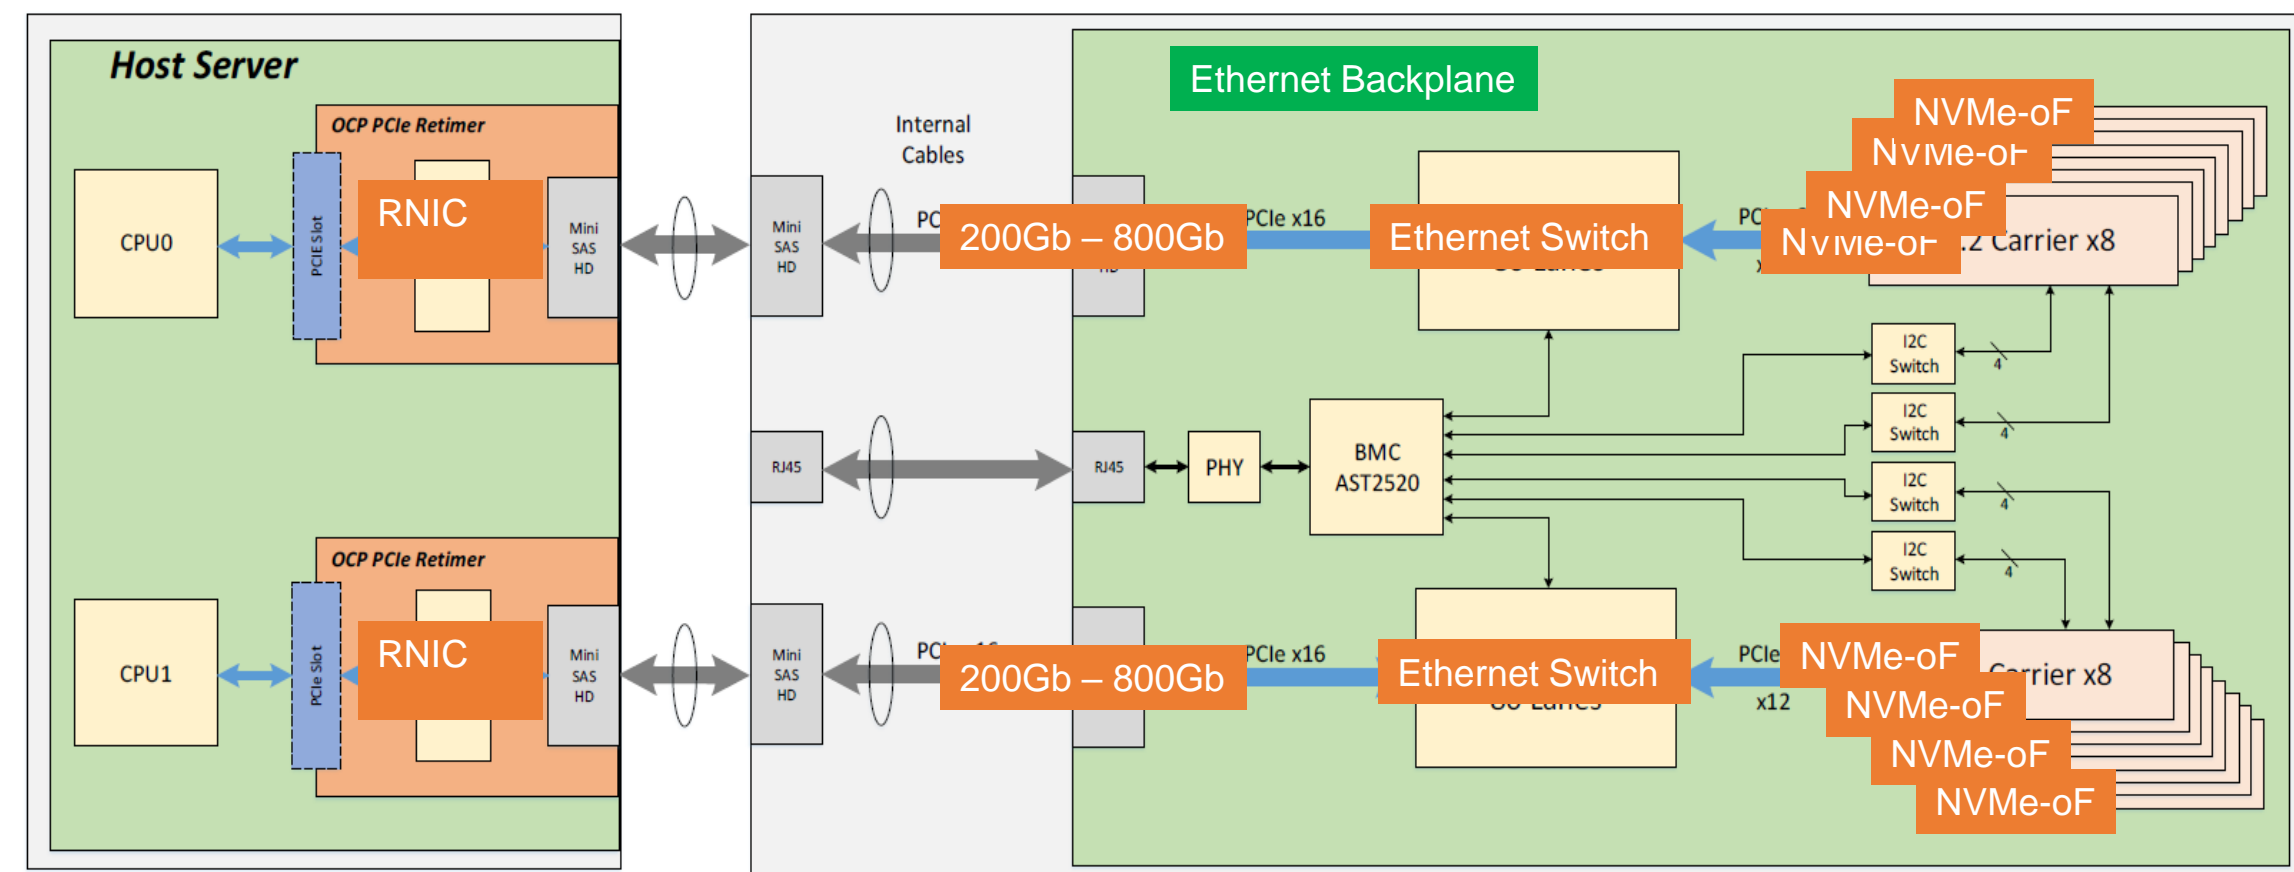


Reference Architecture

Existing FX-16
(PCIe fabric):



Future **Ethernet eFX-16**
(RDMA/TCP fabric):

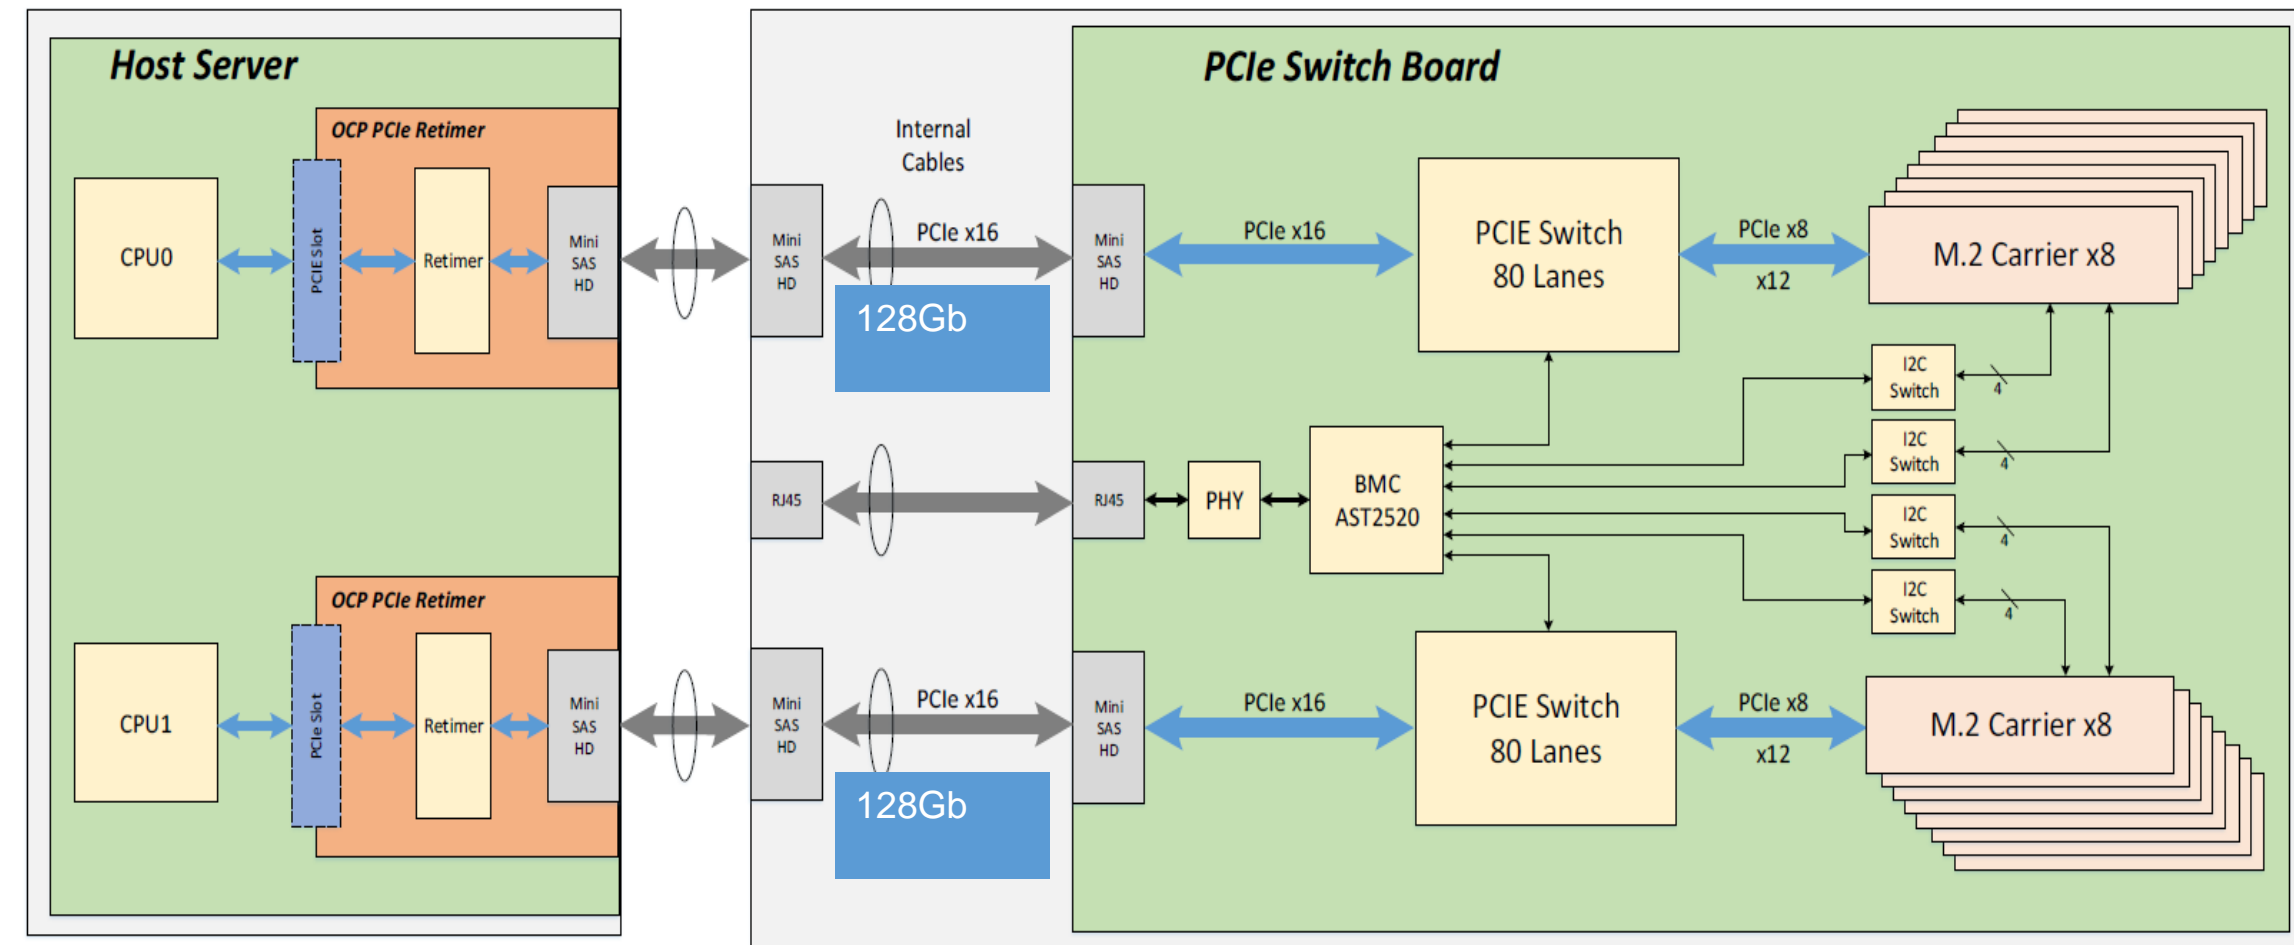


Proposed Architectural Changes



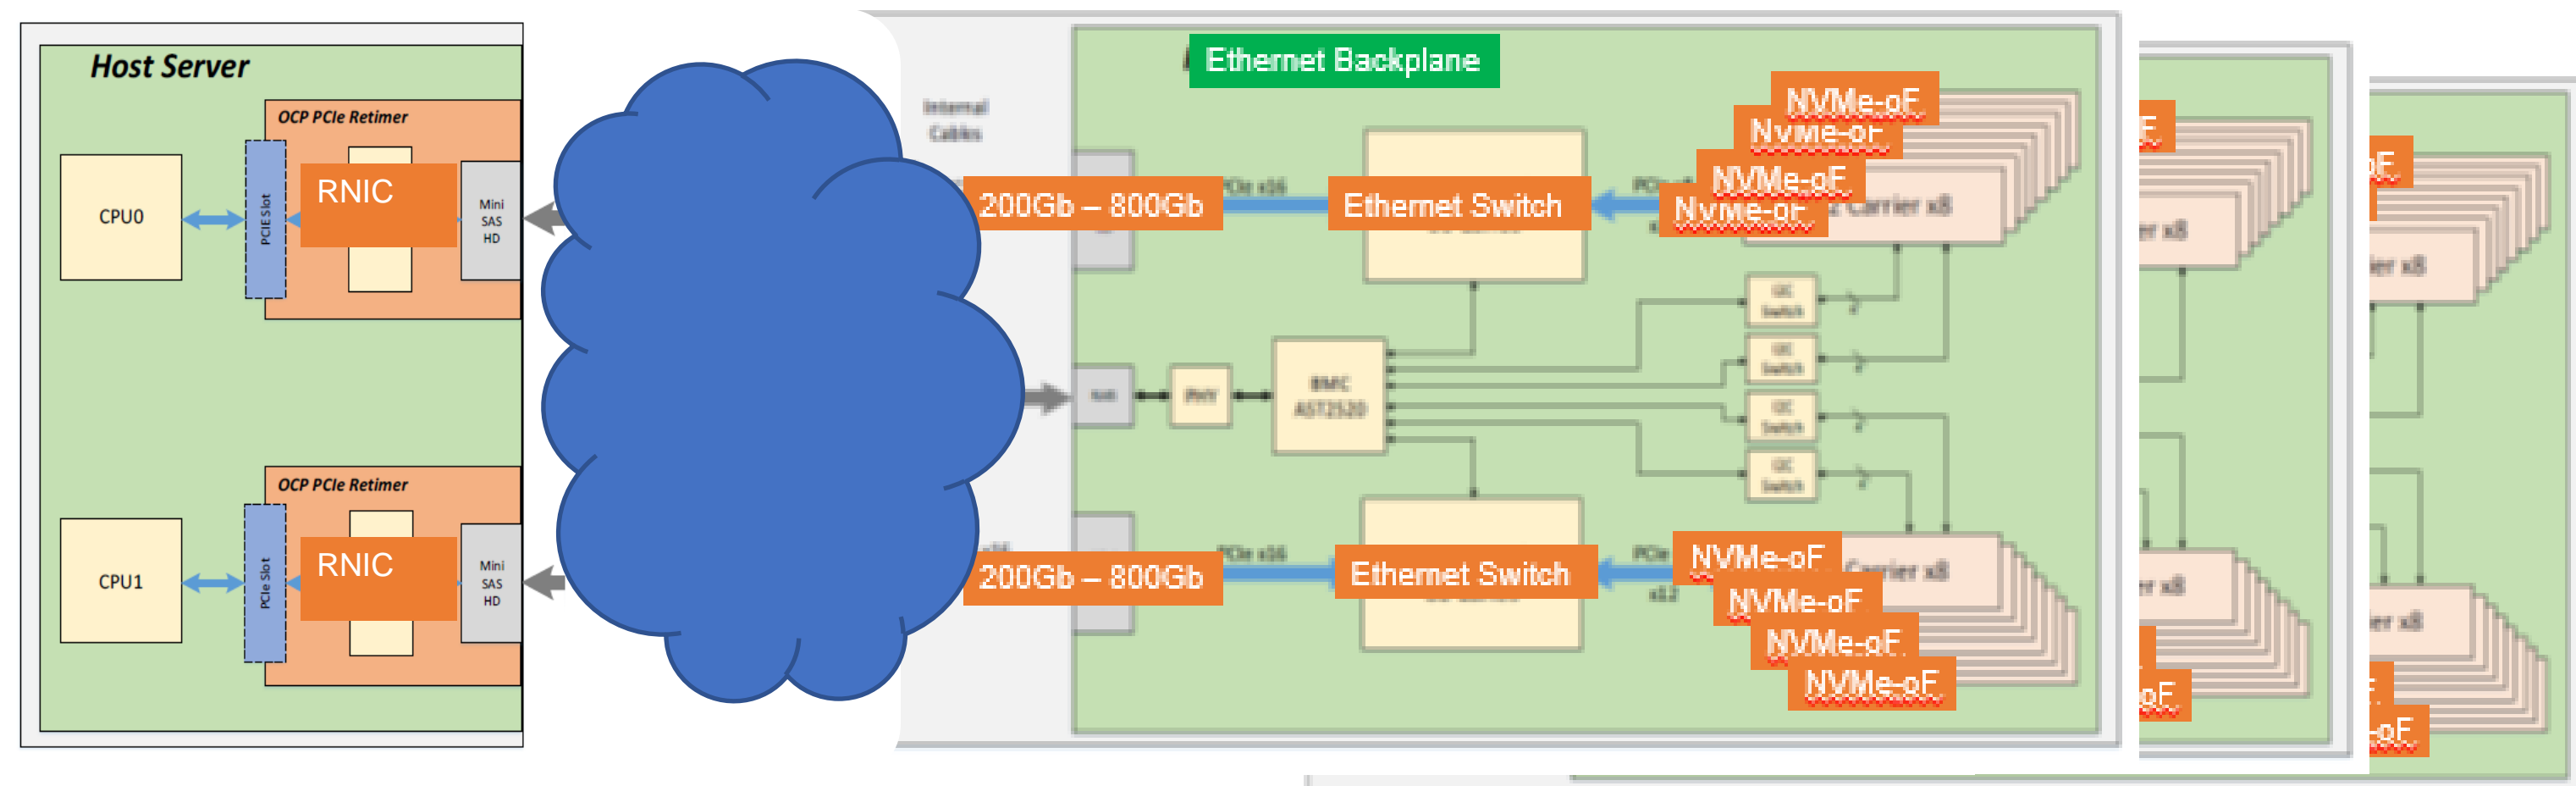
Reference
Architecture

Existing FX-16
(PCIe fabric):



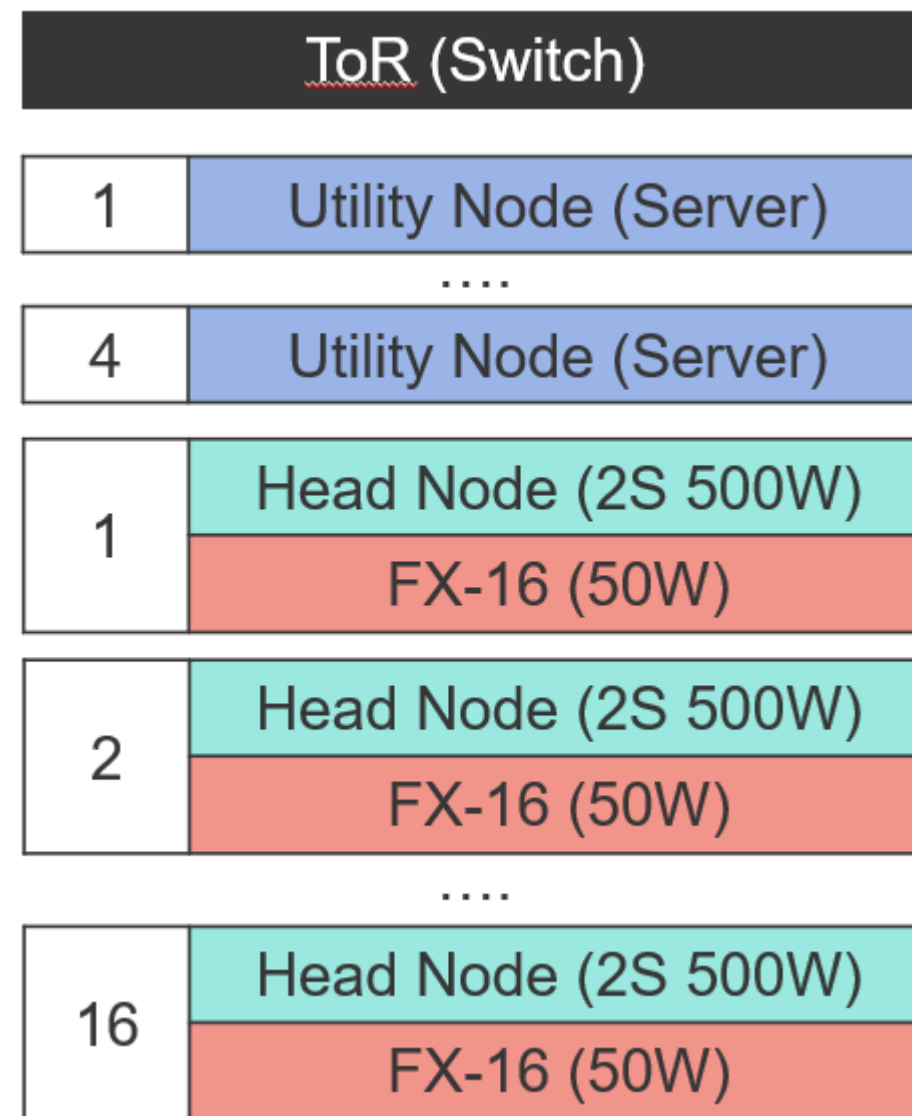
Future **Ethernet eFX-16**
(RDMA/TCP fabric):

Better Scalability!

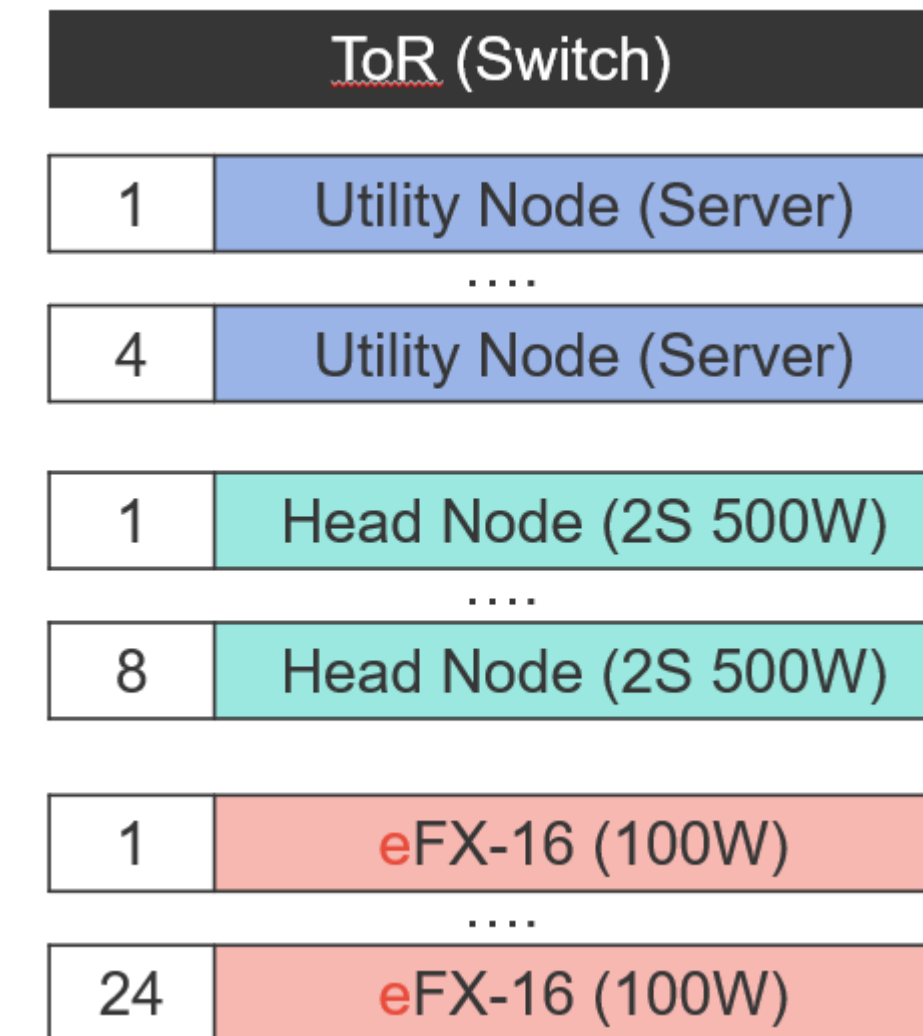


Better Rack TCO

Existing FX-16 (PCIe fabric)



Future eFX-16 (Ethernet fabric)



50% more capacity



30% less power



20% lower rack cost

Summary

- Ethernet has many advantages for NVMe Storage expansion
- Call for action:
 - Future EBOF FX-16 proposal (eFX-16?)
 - Possible EBOD DX-88 proposal (eDX-88?)
 - Standardization of SSD connectors with ethernet
- Technology demonstration at Marvell booth



Open. Together.

OCP Global Summit | March 14–15, 2019

