



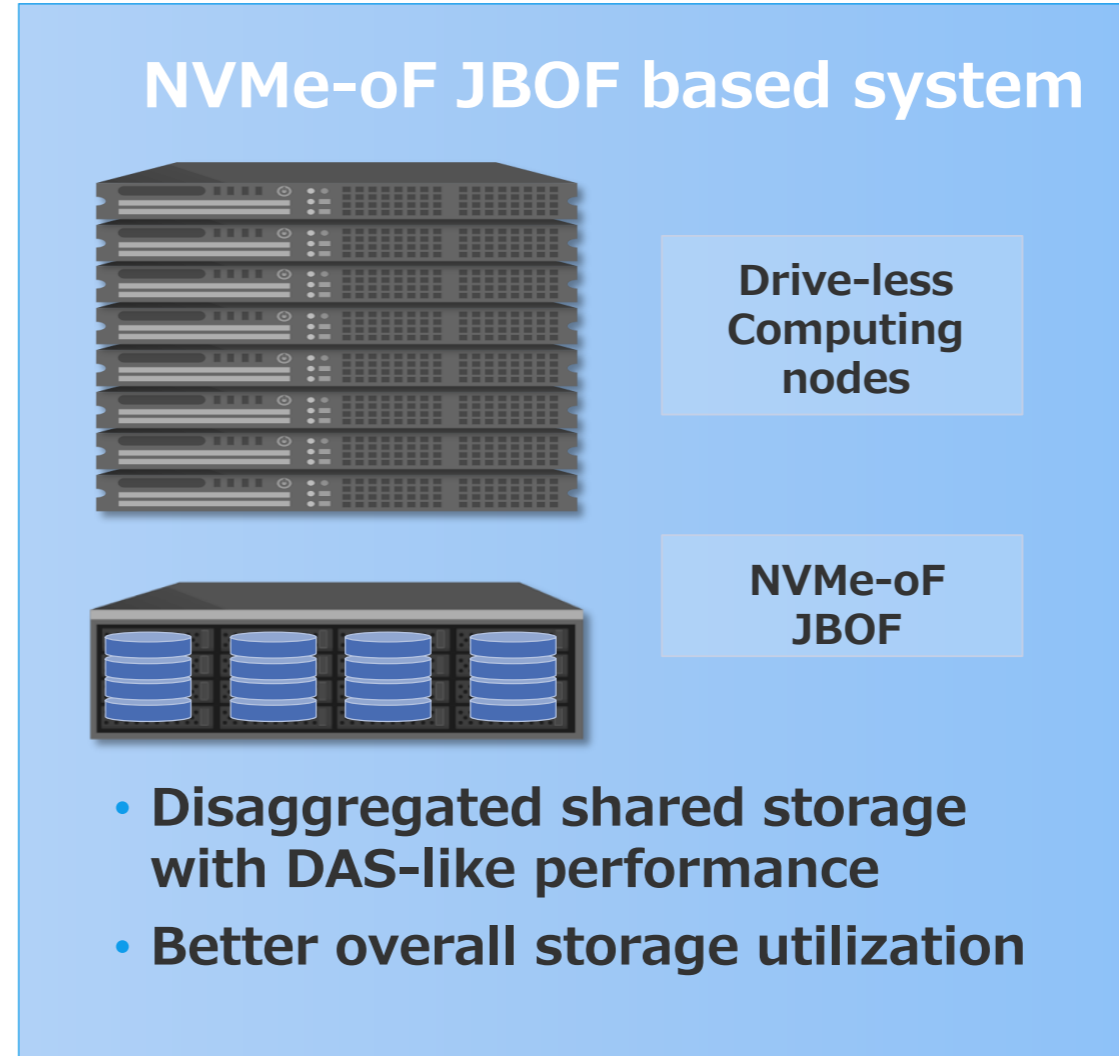

Flash Memory Summit

Enabling Scalable Ethernet JBOD with a Native NVMe-oF™ SSD

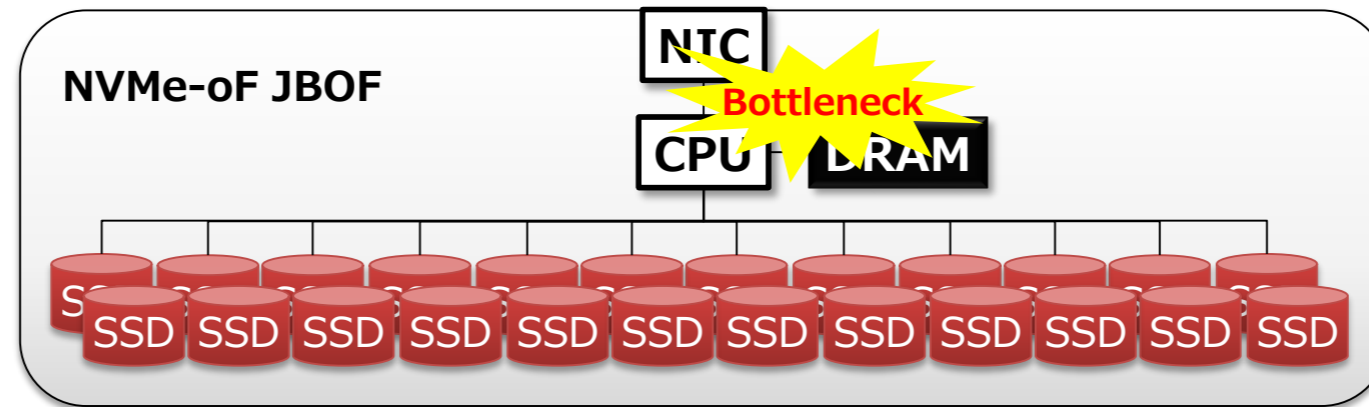

Shingo Tanaka

Toshiba Memory Corporation

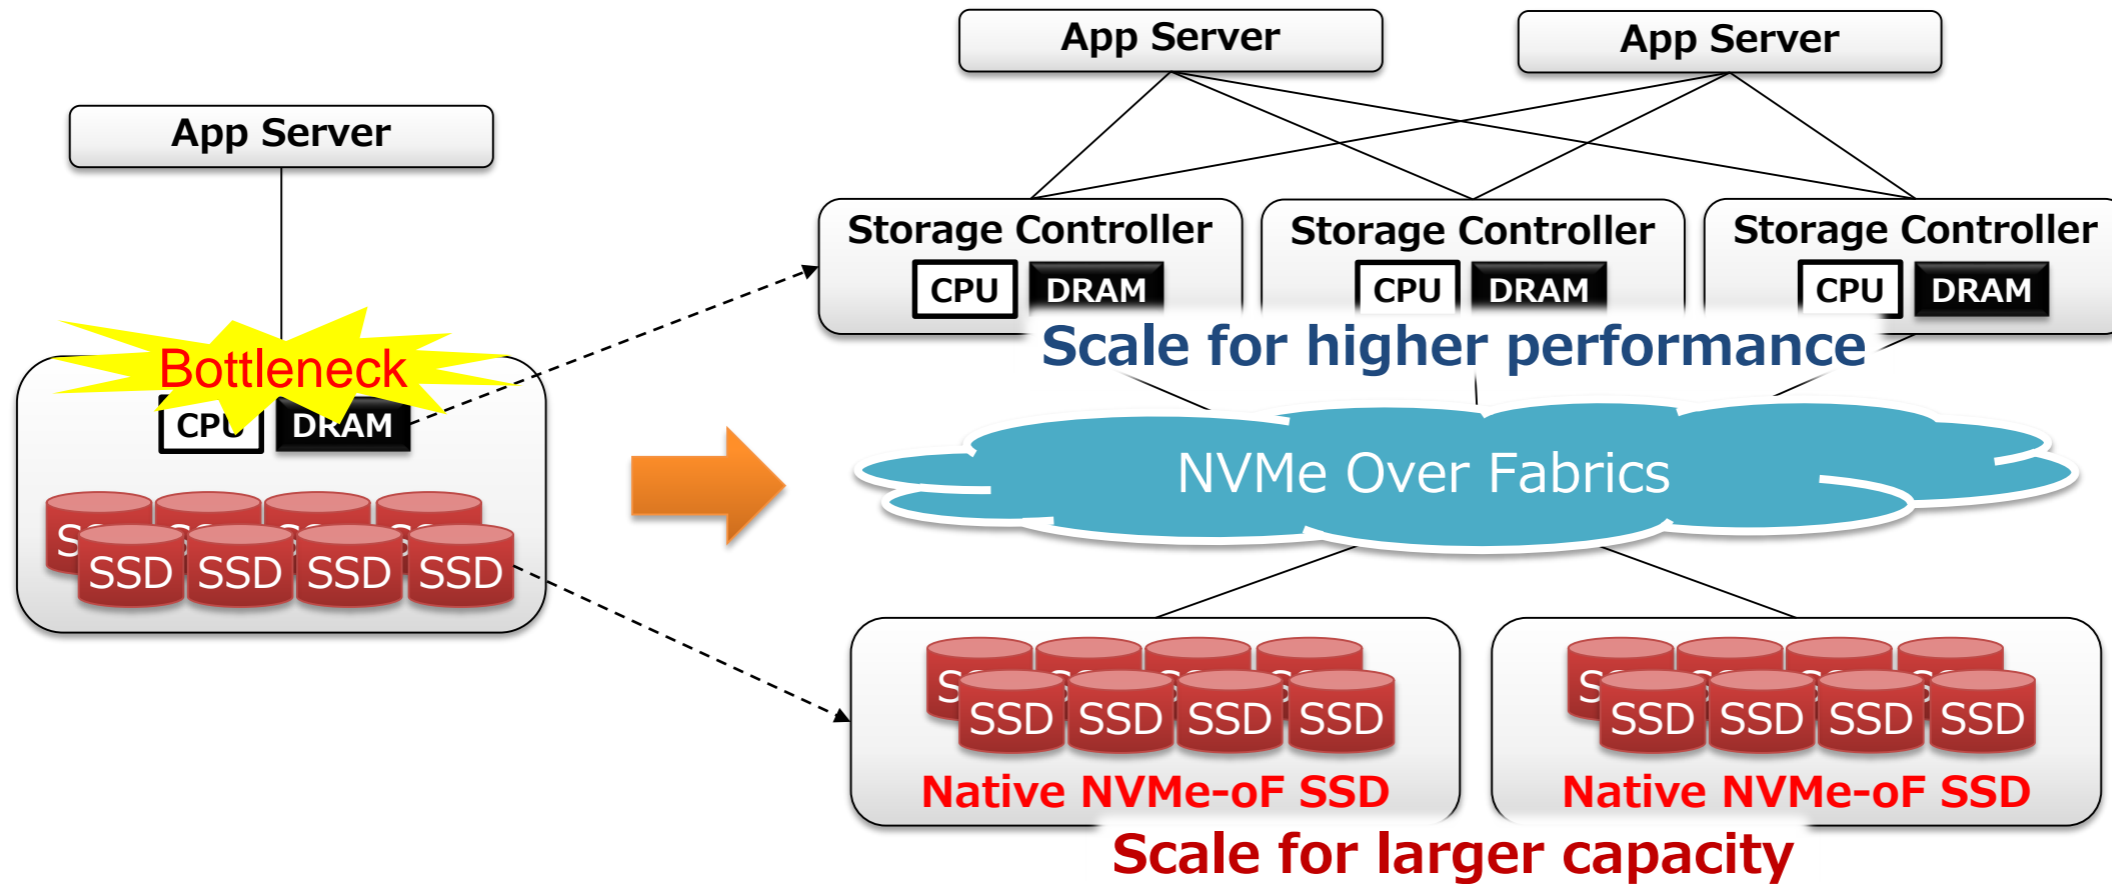
NVMe™ over Fabrics



Challenge: Break the Performance Bottleneck



Storage “Drive” Disaggregation

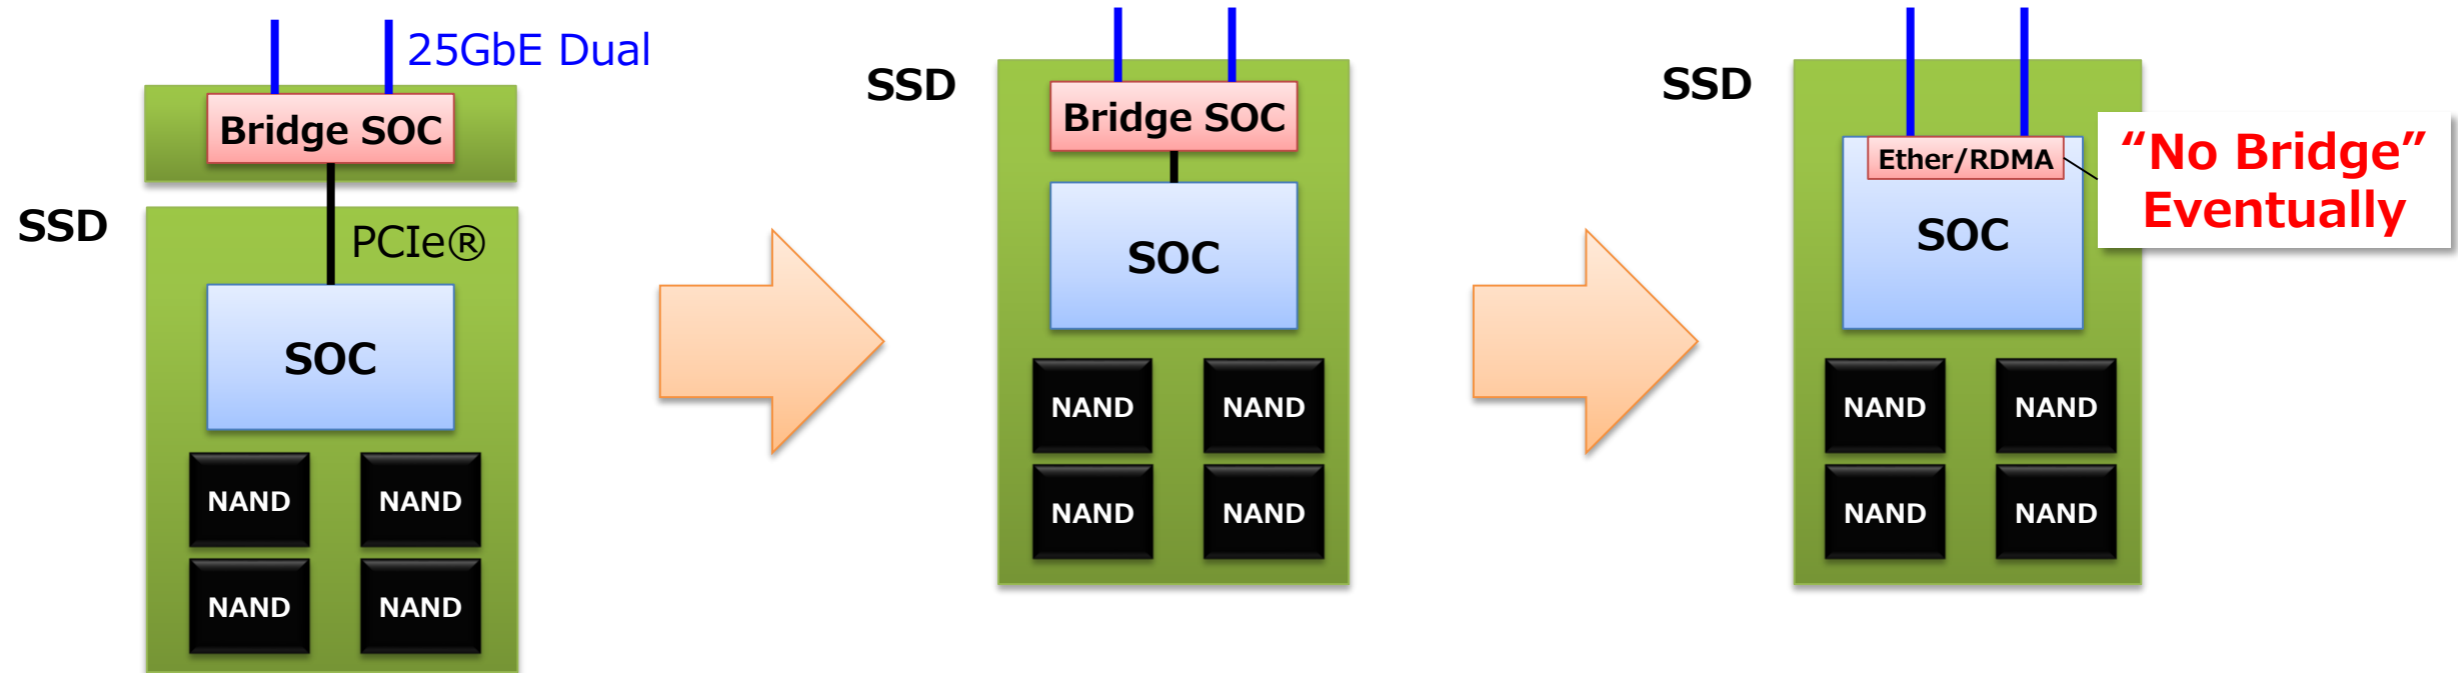


Solution: Native NVMeoF SSD and Ethernet JBOF




Native NVMe-oF SSD

- NVMe-oF operation is guaranteed at SSD level
 - Overall cost reduction for end customers
- HW optimization with eliminating NVMe-oF bridge in the middle



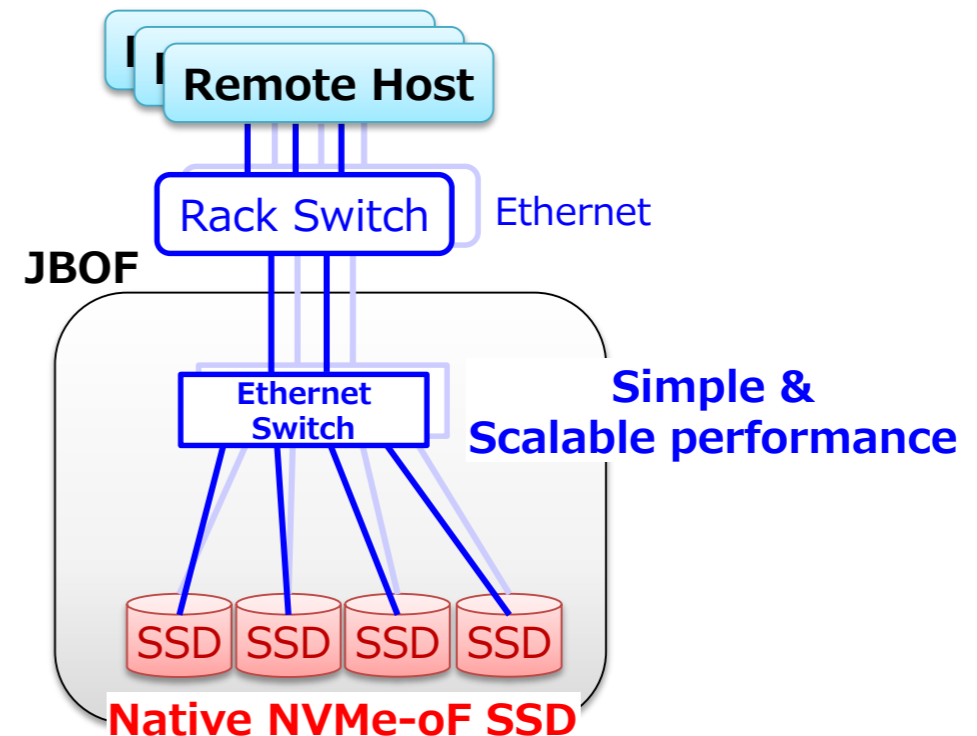
Ethernet JBOF

Existing JBOF



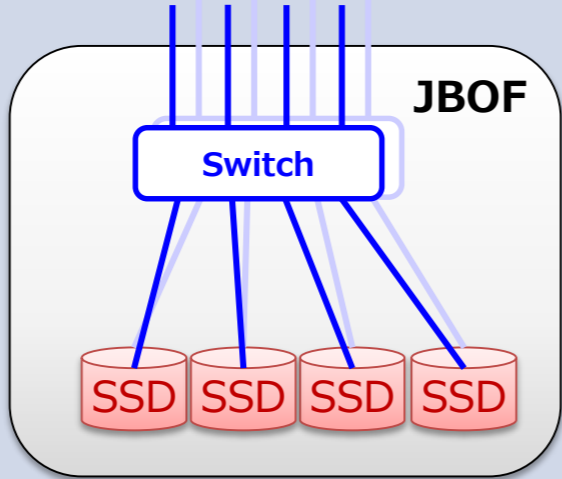
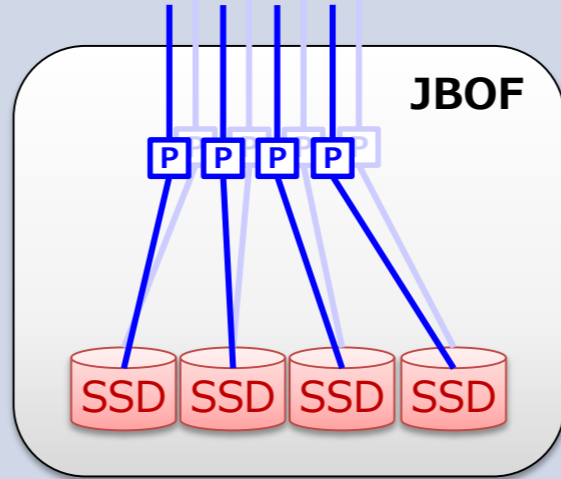
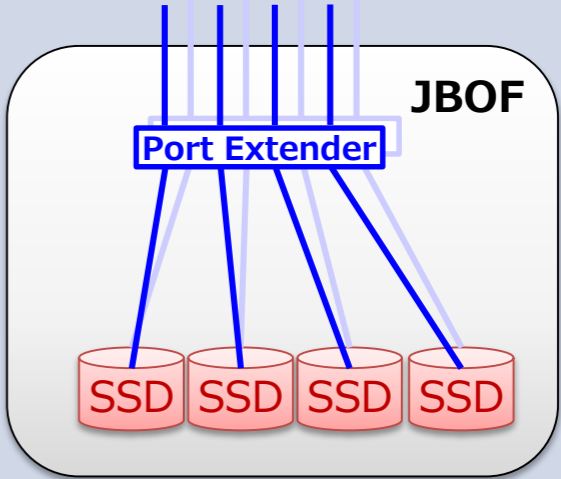
Limited performance
High CPU & NIC power consumption

Ethernet JBOF



Native performance
Extremely lower power consumption

Ethernet JBOF – Frontend Options –

	Ethernet Switch	Ethernet PHY	Ethernet Port Extender
			
Cost (BOM & Power)	High	Low	Low
Port Flexibility	High	Low (Need to connect all ports to access all SSDs)	High
Remarks	Could omit rack switch in small deployments	Port/cable can be physically aggregated (c.f. 4x25G in a single QSFP28)	Best option for large deployments



Proof-of-Concept prototype

- Best-in-class performance w/ 24x SSDs
- Low cost & power for NVMe-oF bridging
- Scalable system performance

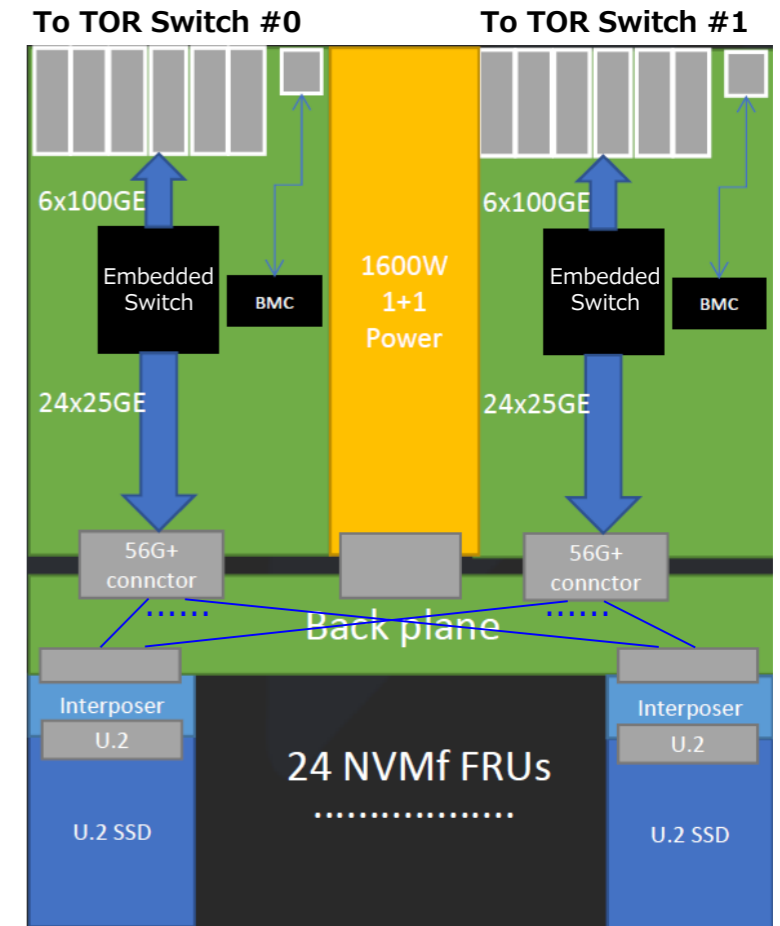


2.5" U.2 Native NVMe-oF SSD
w/ NVMe-oF interposer
Dual 25GbE I/F

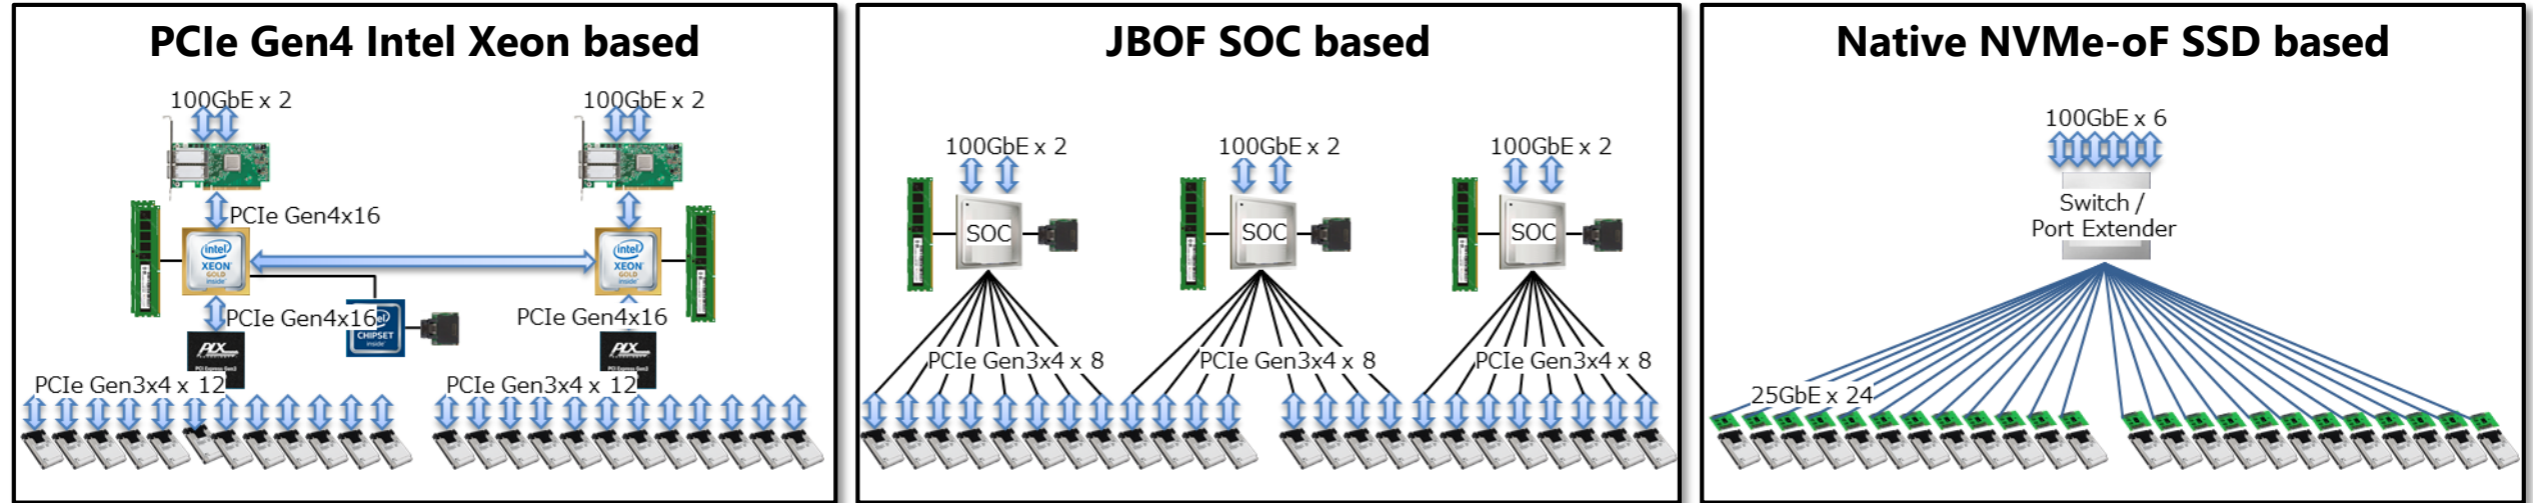


2RU 24xSSD HA Ethernet JBOF
Dual 6x 100GbE I/F


Demo at Toshiba Booth #307, Hall A



Preliminary TCO analysis



TCO comparison of NVMe-oF bridging & aggregation overhead (excluding SSDs)



$TCO = CAPEX + OPEX$

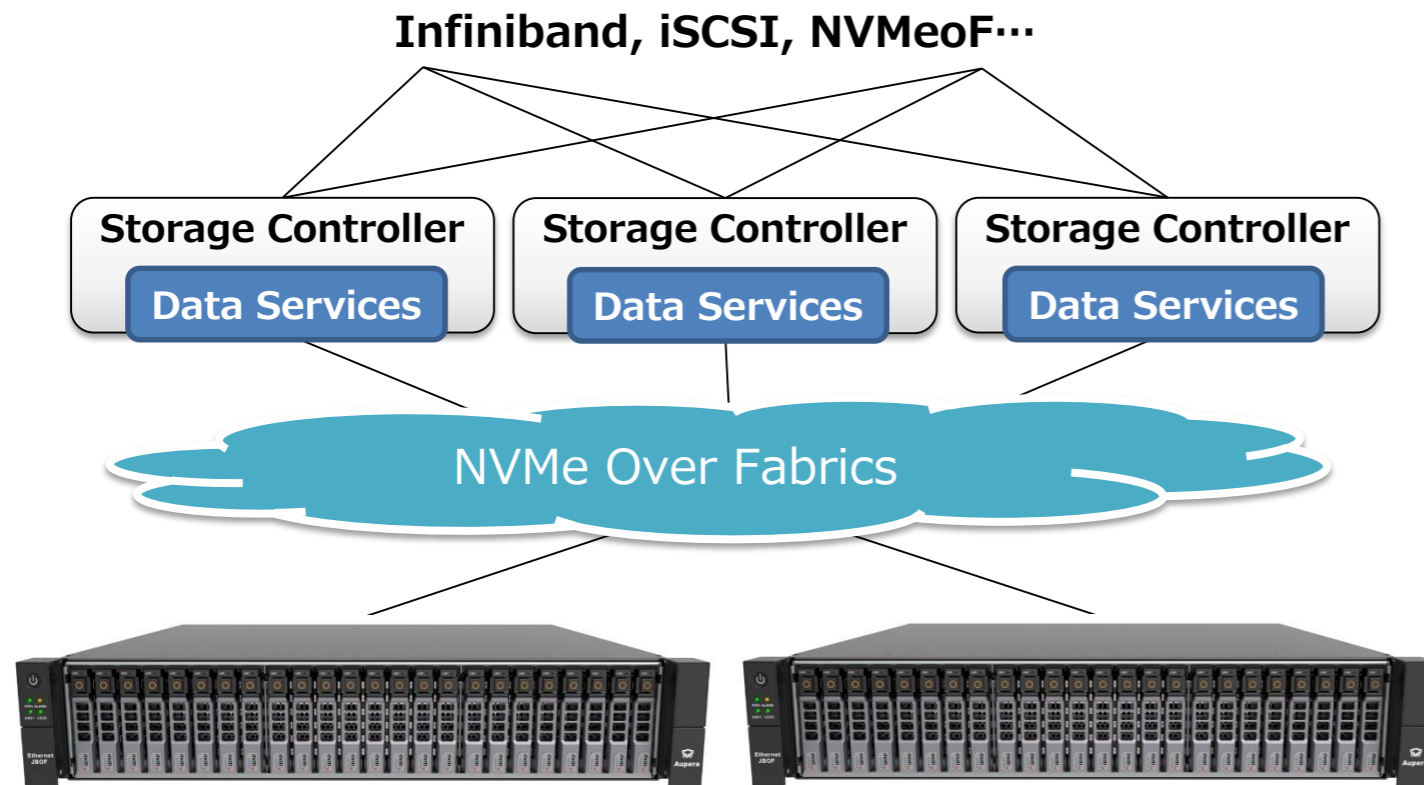
CAPEX: based on market prices in small quantity

OPEX: based on running cost estimate with power cost 0.1 \$/KWh, system power efficiency 0.8, Lifetime 5 yrs



System level POC concept

- High performance storage system based on Ethernet JBOF with data services



NVMe over Fabrics Solution
 Best-in-class performance
w/ Data Services



- Data Services
- Multi-tenancy
 - High-availability
 - Inline Deduplication
 - Inline Compression
 - Snapshot
 - Thin-provisioning





Flash Memory Summit

Thank You

Contact: shingo3.tanaka@toshiba.co.jp