# Datacenter Designs Using the EDSFF Form Factors

Mark A. Shaw

*Microsoft Azure*

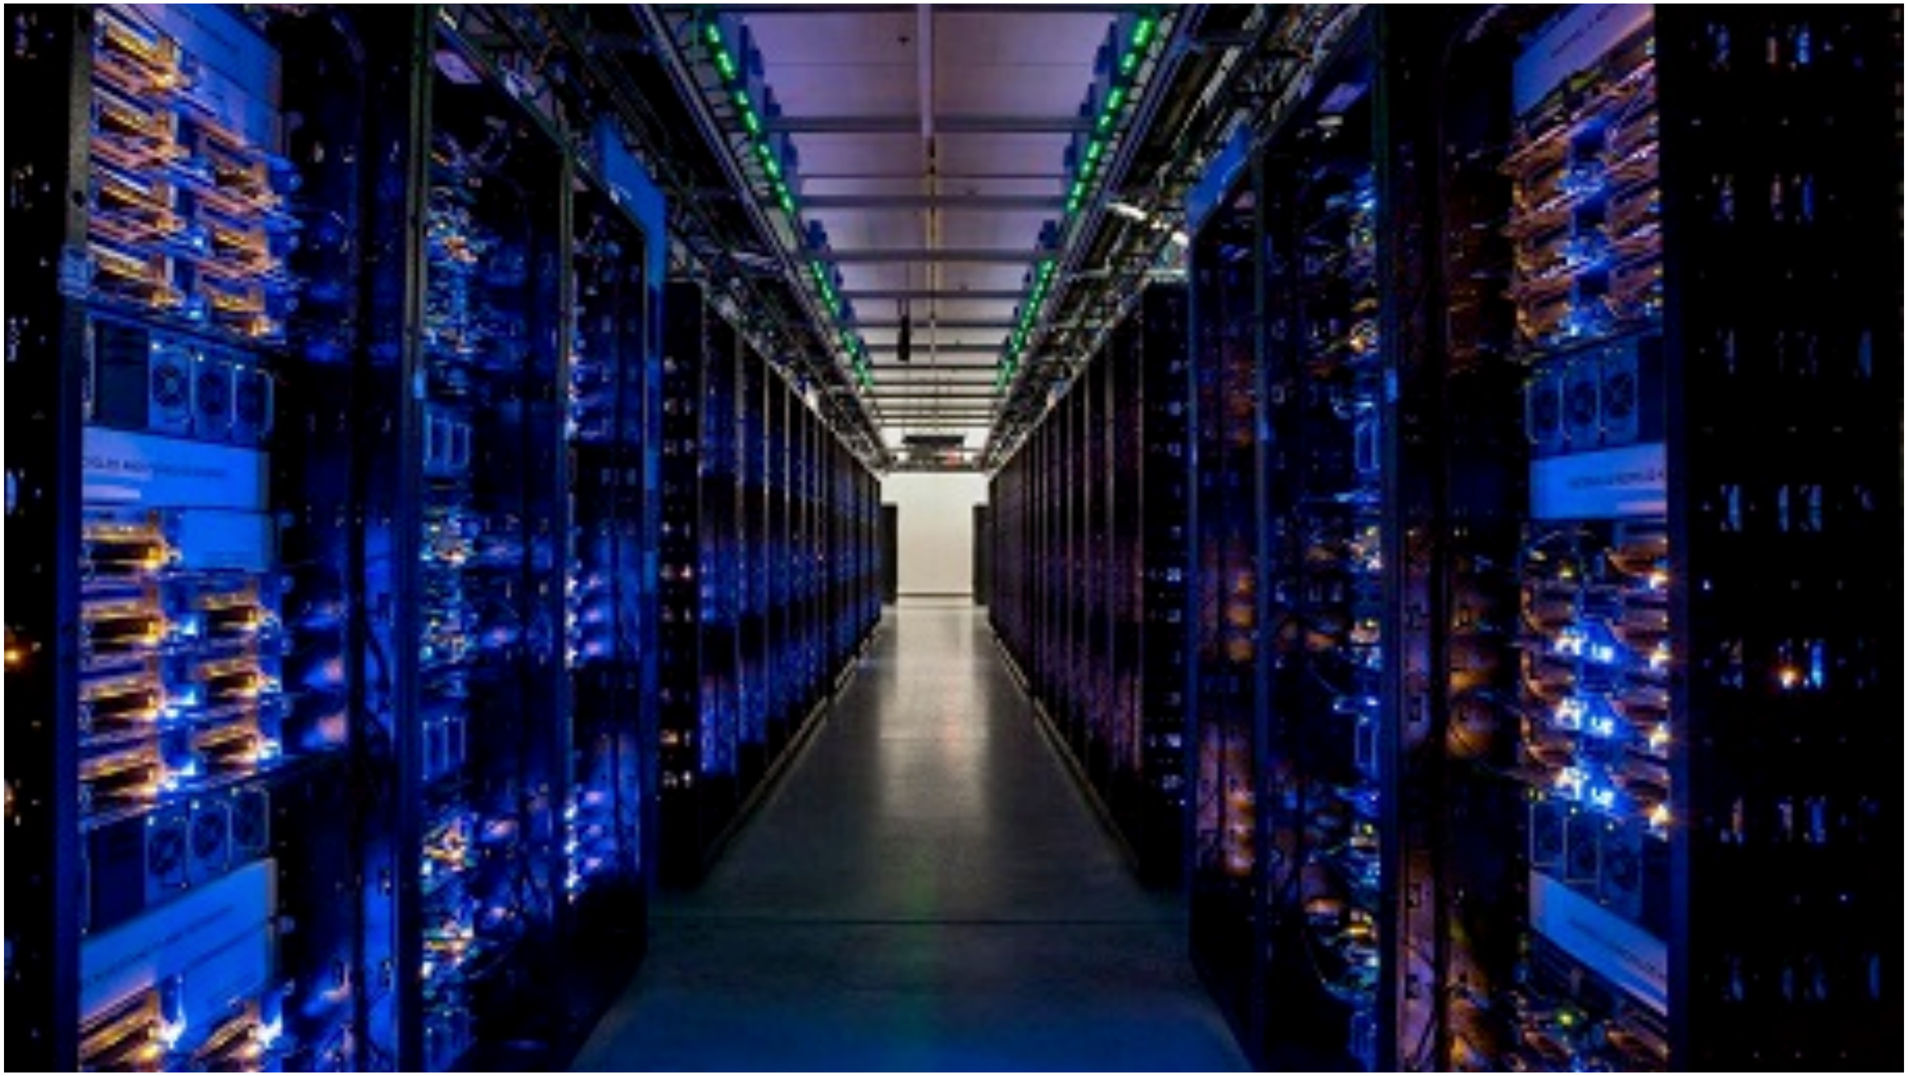

# EDSFF – Decoupling for workloads

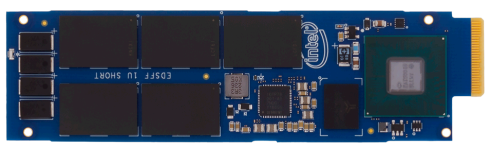

**SFF-TA-1006 (E1.S)**

Performance Density Optimized Services
High Power SCM (i.e. PCM/ReRAM/FRAM)
Cache Coherent Devices (i.e. Gen-Z)
Targets Compute nodes

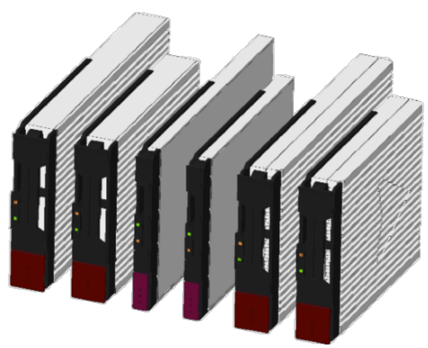

**SFF-TA-1008 (E3)**

Ultra High-Performance Applications
FPGA or Computational Accelerations
AIC / PCIe HHHL replacement

**SFF-TA-1007 (E1.L)**

Capacity Density Optimized Services
Cheapest Consumer memories (i.e. QLC)
Denali-based pools of flash
Targets Storage

# Azure Design Considerations

## Performance Storage

High Density
Low Latency
Low $/GB
On-Line Service

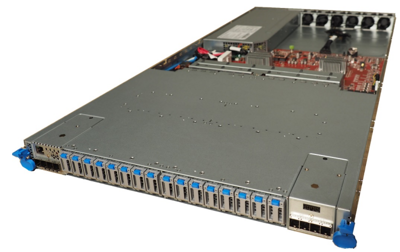

## Compute

Increasing # of VMs per node
Challenges keeping up with IOP Density
Varying IOP and Density requirements
Power, Space and Thermal constrained

## Cold Storage

High Density
Lowest $/GB
Low $/Watt
On-Line Service

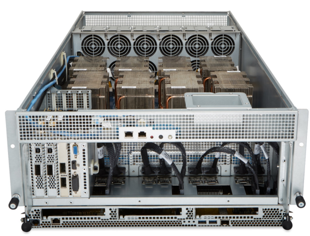

## GPU

Ultra High-Performance Applications
FPGA or Computational Accelerations
PCIe HHHL replacement

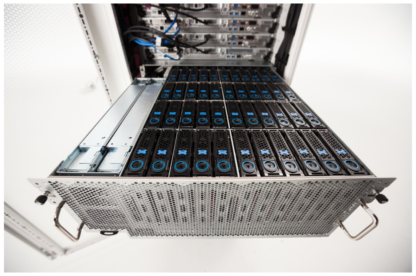

# Performance Storage – FX16

- Lowest $/TB
- High Density (256 TB or higher)
- Low latency
- Moderate BW/SSD
- Expandable capacity
- On-line Service
- Modular to support all head nodes
- Leverage existing components and architectures

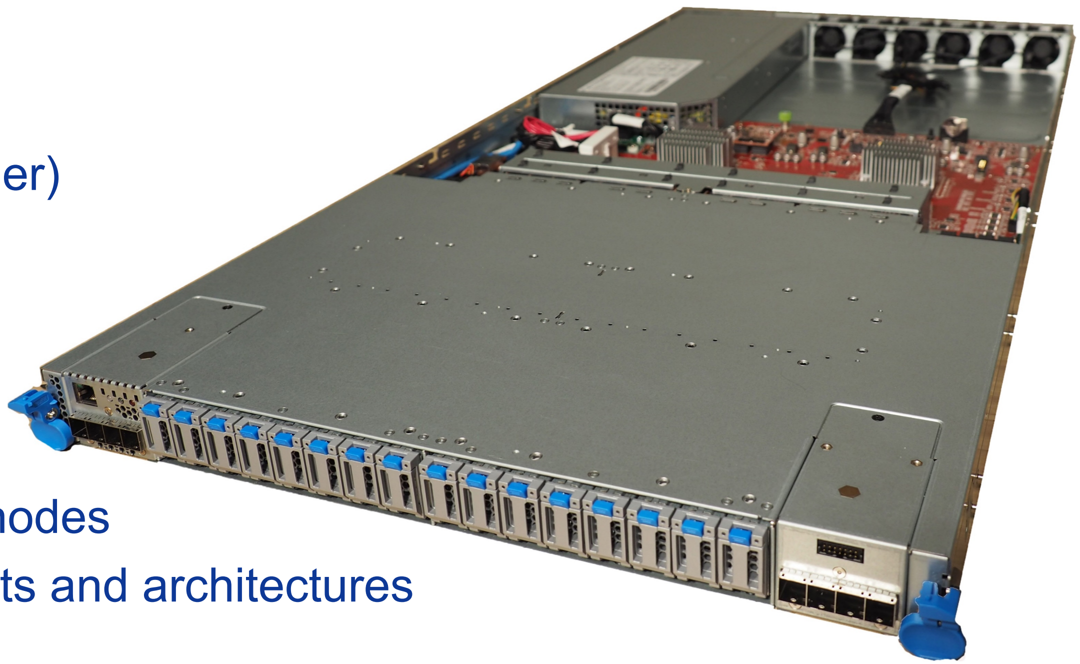

# M.2 Carrier in FX16


## Advantages

- TTM with existing tech
- Meets current performance requirements

## Disadvantages

- Mechanical complexity
- Thermal performance
- 3.3V power scaling
- Supply Chain complexity
- Service complexity
- Increased “spares” support
- Lack of QLC support ($$)

# E1.L in FX16

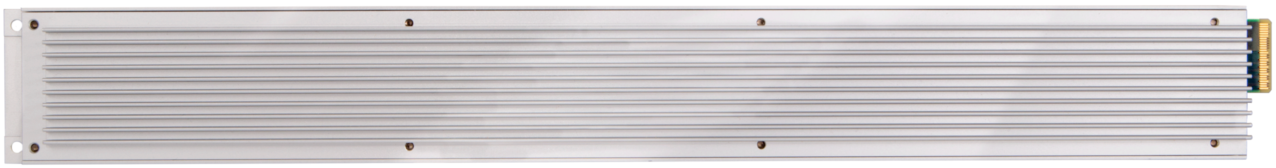

## Advantages

- Meets performance requirements
- Reduced mechanical complexity
- Thermal Performance
- Scaleability
- Reduce Supply Chain complexity
- Reduced service complexity
- QLC support $$

## Disadvantages

- TTM with new technology
- Risk with new technology

# Cloud SSD E1.L

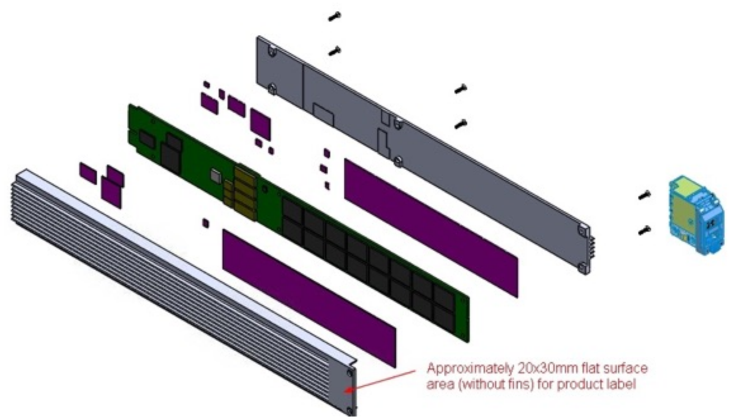


MULTI-SOURCED!!

- E1.L Form Factor
- Configurable platform for multiple work loads
- Single or Dual controller configurable
- Supports LL-NAND, eTLC, QLC
- Supports LBA and Denali Media SSD

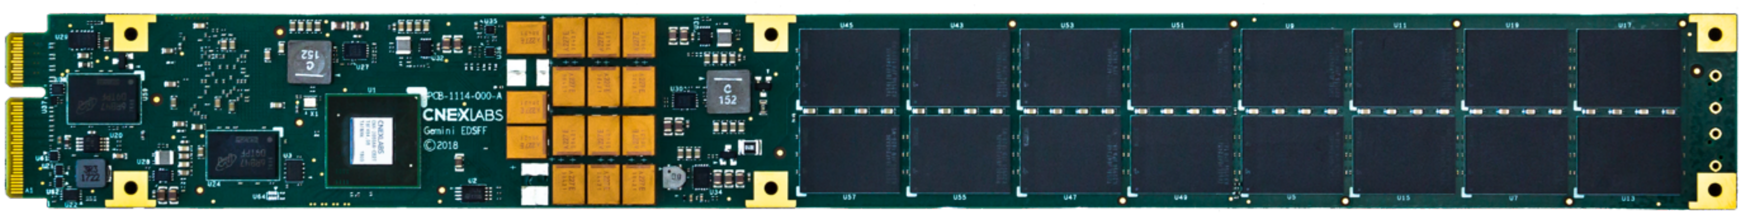

# Going Forward

| | PCIE Gen3 | PCIE Gen4 | PCIE Gen5 |
|---|---|---|---|
| **M.2** | 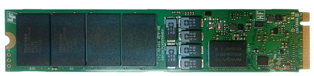 | ✔ Degraded connector performance<br>Carrier topology risks<br>Bent but not broken | ✖ High risk connector performance<br>Higher risk carrier topologies<br>Likely broken |
| **PCIe CEM** |  | ✔ Degraded connector performance<br>Reduced trace lengths<br>Riser topology risks<br>Increased board material cost<br>Bent but not broken | ✖ Degraded connector performance<br>Reduced trace lengths<br>Riser topology risks<br>Increased board material cost<br>Likely broken |
| **EDSFF** | 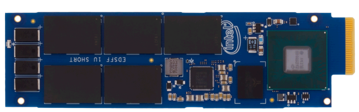 | ✔ Connector rated for Gen5<br>Demonstrated 2 connector topology | ✔ Connector rated for Gen5<br>Demonstrated 2 connector topology |

# Thank You