



SDC 18

September 24-27, 2018
Santa Clara, CA

www.storagedeveloper.org

Various All Flash Array Architectures

Mike Scriber
Supermicro

Agenda



- ❑ Architectural Considerations
- ❑ Direct Connect Solutions
- ❑ Switched Solutions
- ❑ JBOF Solutions
- ❑ High Availability Solutions

Architectural Considerations

- ❑ Limited PCIe lanes for NVMe solutions
 - ❑ 96 lanes for dual socket Intel
 - ❑ 128 lanes for AMD
- ❑ Physical Limitations
- ❑ Cable Limitations
- ❑ Thermal Limitations
- ❑ Balanced Designs



PCIe Tradeoffs




- ❑ NVMe SSDs require PCIe lanes
- ❑ Network requires PCIe lanes
- ❑ How should you divide up the PCIe lanes?



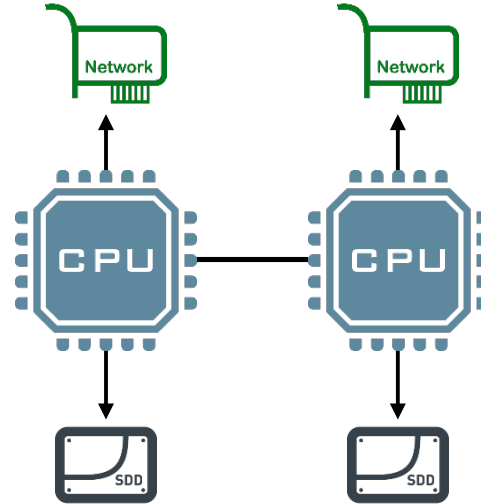
Balanced Designs



Asymmetric



Symmetric



Direct Connect Solutions




- ❑ 1U 10 SSD
- ❑ 1U/2U 20 SSD
- ❑ Multi-node Systems




1U 10 SSD

- ❑ 40 PCIe lanes for SSD
- ❑ 40 PCIe lanes for slots
- ❑ 16 PCIe lanes LOM, etc
- ❑ Physical limitation out the back



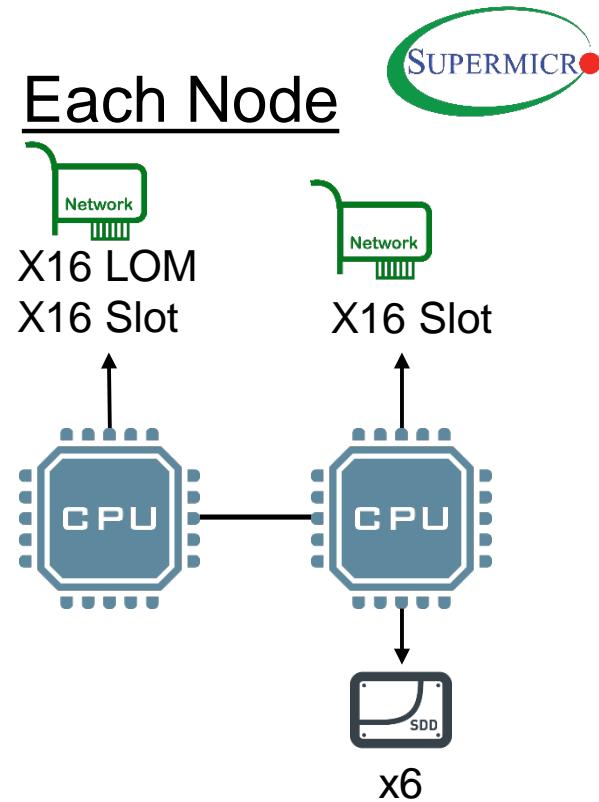
1U/2U 20 SSD

- ❑ 80 PCIe lanes for SSD
- ❑ 8 PCIe lanes for slots
- ❑ 8 PCIe lanes LOM
- ❑ Limited I/O out of the box
- ❑ 5:1 ratio SSD to network



Multi-node Systems

- ❑ 4 nodes in a 2U system with 24 SSDs
- ❑ 24 PCIe lanes for SSD
- ❑ 32 PCIe lanes for slots
- ❑ 16 PCIe lanes LOM
- ❑ Physical limitation out the back
- ❑ Low blast radius



Switched Solutions

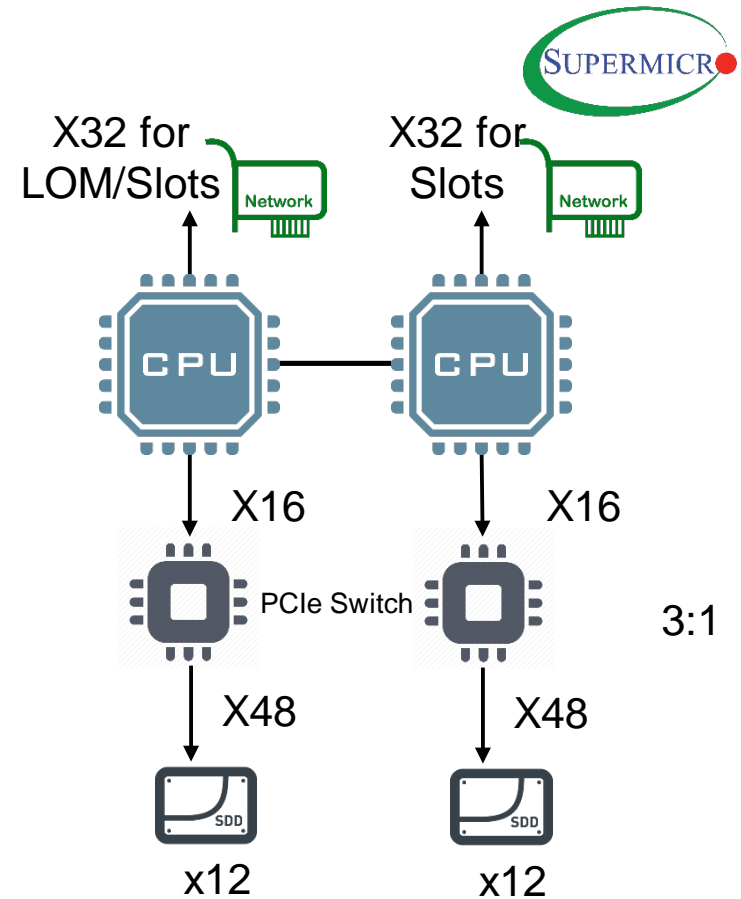


- ❑ 2U 24 SSD
- ❑ 1U 32 SSD
- ❑ 1U 36 SSD



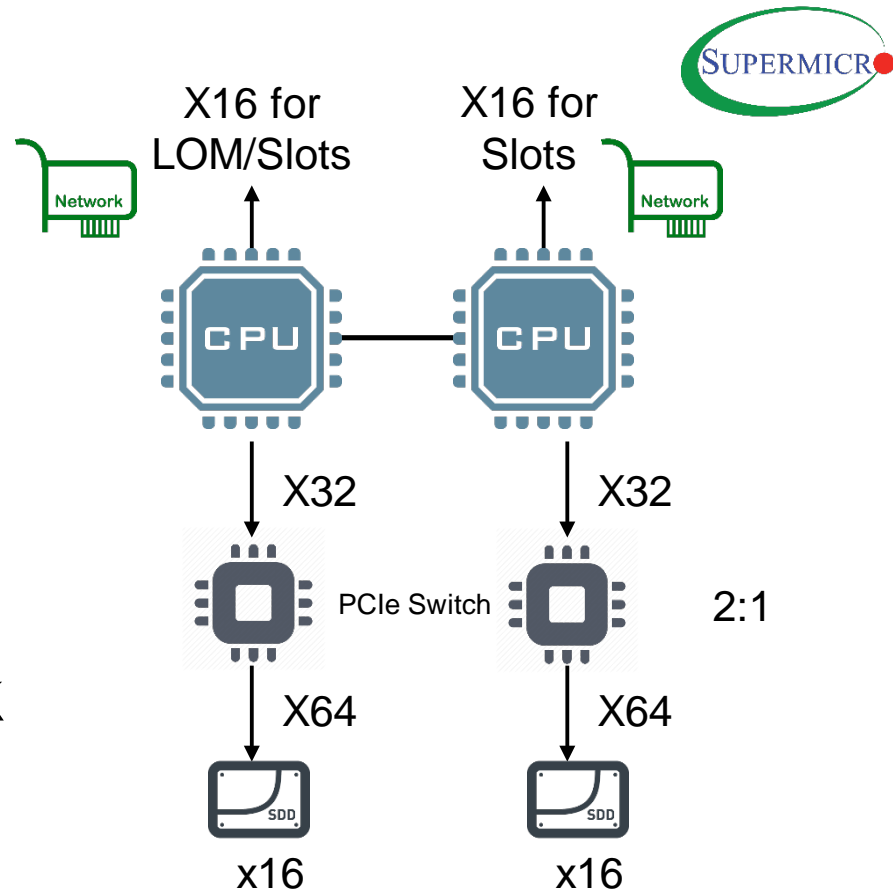
2U 24 SSD

- ❑ 32 PCIe lanes for SSD
- ❑ 3:1 ratio at the switch
- ❑ 64 PCIe lanes for LOM and slots
- ❑ 1:2 ratio SSD to network



1U 32 SSD


- ❑ 64 PCIe lanes for SSD
- ❑ 2:1 ratio at the switch
- ❑ 32 PCIe lanes for LOM and slots
- ❑ 2:1 ratio SSD to network



1U 36 SSD

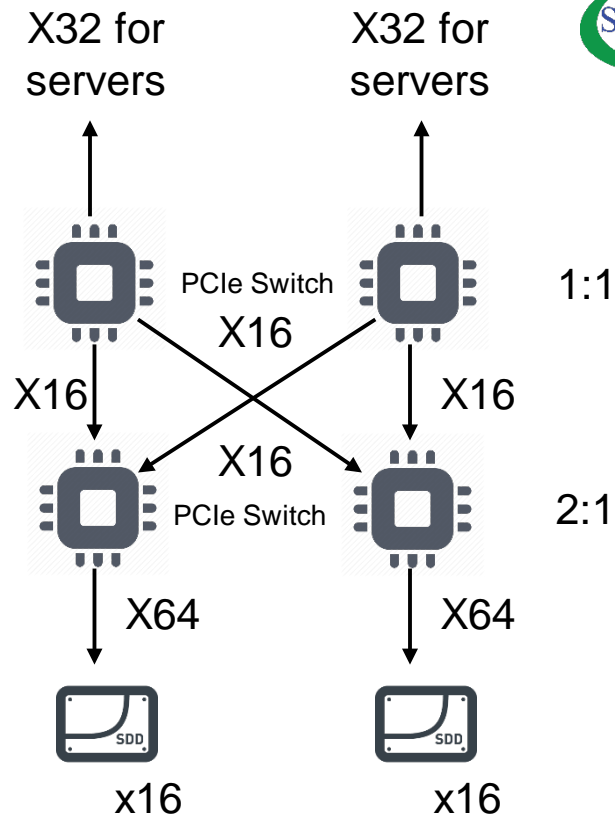


- ❑ 48 PCIe lanes for SSD
- ❑ 3:1 ratio at the switch
- ❑ 48 PCIe lanes for LOM and slots
- ❑ 1:1 ratio SSD to network



1U 32 JBOF Solution


- ❑ 64 PCIe lanes for SSD
- ❑ 2:1 ratio at the switch
- ❑ High bandwidth to multiple servers



Multi-node & JBOF System



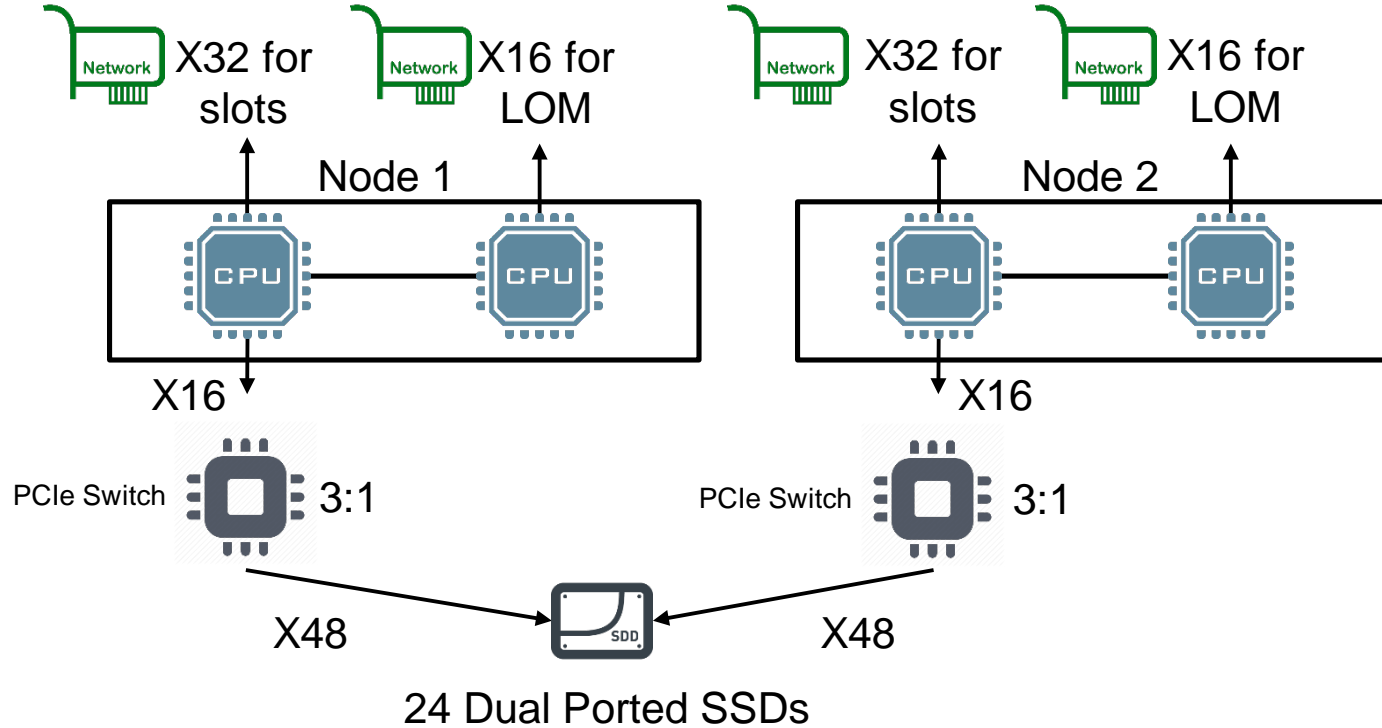
- ❑ 4 nodes in a 2U system with 56 SSDs
- ❑ 40 PCIe lanes for SSD
- ❑ 16 PCIe lanes for slots
- ❑ 16 PCIe lanes LOM
- ❑ 800GbE total



2U 24 HA Solution



- ❑ 16 PCIe for SSD
- ❑ 3:1 ratio at switch
- ❑ 48 PCIe lanes for LOM and slots
- ❑ 1:3 ratio SSD to network



Which Architecture is Right?

- ❑ Depends on your configuration
- ❑ Depends on your application
- ❑ Depends on your usage model
- ❑ Depends on what is important to you.