Pre-Conference Seminar E: Flash Storage Networking

Rob Davis, Chris DePuy, Tameesh Suri, Saurabh Sureka, Gunna Marripudi, and Asgeir Eiriksson

Agenda

- Networked Flash Storage Overview
  - Mellanox – Rob Davis
- The Effects on Networks Caused by Including Flash Storage
  - Avago – Saurabh Sureka
- The Effects on Flash Storage Systems Caused by Networking
  - Samsung – Tameesh Suri and Gunna Marripudi
- Competing Technologies and Architectures for Networked Flash Storage
  - Chelsio – Asgeir Eiriksson
- The Market for Networked Flash Storage
  - DellOro – Chris DePuy

Networked Flash Storage Overview

- Why networked flash storage?
- What's involved in networking flash storage?
- What do typical implementations look like?
- What are the tradeoffs?

Why Network Flash Storage?

- There are advantages to shared storage
  - Better utilization of capacity
  - Scalability
  - Easier to manage
  - Cluster applications
  - Server Virtualization
  - Fault Isolation
- Shared Storage requires a Network

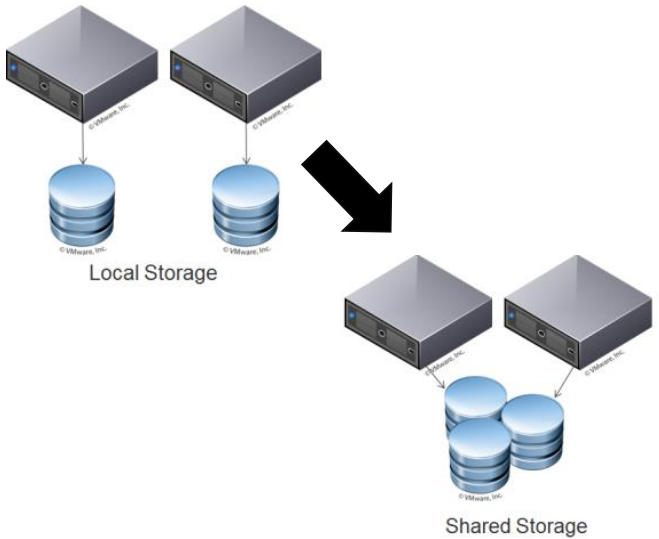

Better Utilization, Scalability, and management

- Driven by Hyper Scale market need compute efficiency
  - Best possible utilization of capacity, rack space, power, cost
- Scalability
- Easier to manage

Dense Server Rack

Enabled by Shared Storage

- Cluster applications
  - Oracle Rack
  - DB2
- Server Virtualization
- Fault Isolation

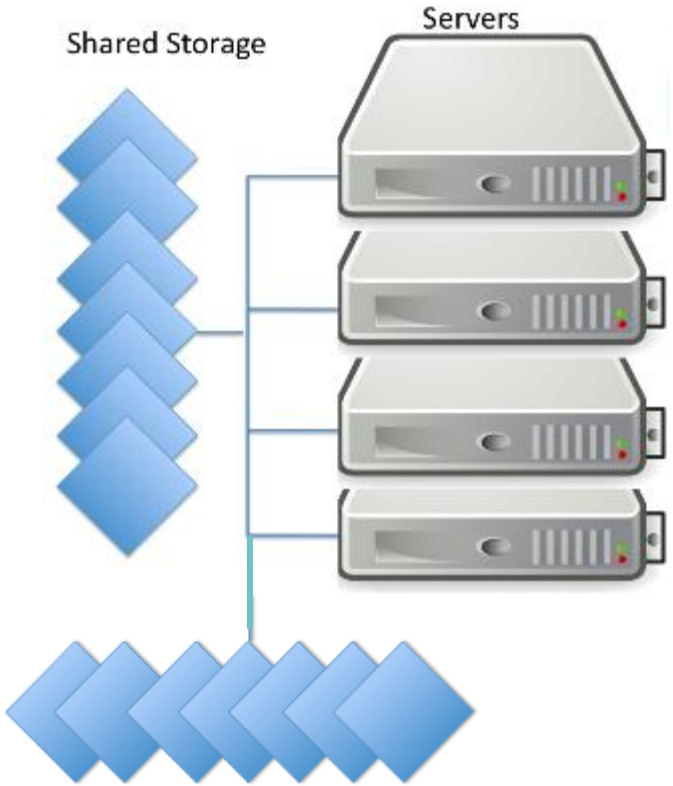

What's involved in networking flash storage?

- Equipment/Hardware
- Software
- Vendors/Suppliers

Equipment/Hardware needed to networking flash storage

- Pick a Network Technology
  - Ethernet, Fibre Channel, InfiniBand, SAS, PCIe
- Make sure your storage device supports this
  - N/A if doing Hyper Converged or Scale Out
- Network Adaptors for your servers
  - NIC or LOM – Ethernet
    - iSCSI
    - iSER, NVMeOF
      - RDMA support
  - HBA – Fibre Channel or SAS
  - HCA – InfiniBand
  - Bus Extender – PCIe
- Switches

Software needed to networking flash storage

- Driver software for you Adapters that matches the OS on your servers
- Storage Management software
  - Often part of the OS
  - Open Source – SDS
- Switch and/or Fabric Management software
  - Usually from switch supplier
  - Open Source – SDN

Who are the Storage Networking Vendors

- Adapters
  - Mellanox, Avago, QLogic, Intel, PMC, others
- Switches
  - Cisco, Arista, Mellanox, Brocade, Dell, HP, Lenovo, others
- Arrays
  - EMC, NetApp, HP, Dell, IBM, HGST, Violin, many others
- Software
  - VMWare, Microsoft, Red Hat, Symantec, open source, others

What do typical implementations look like?

- Components, Boxes, Topology, etc.

Some assembly required

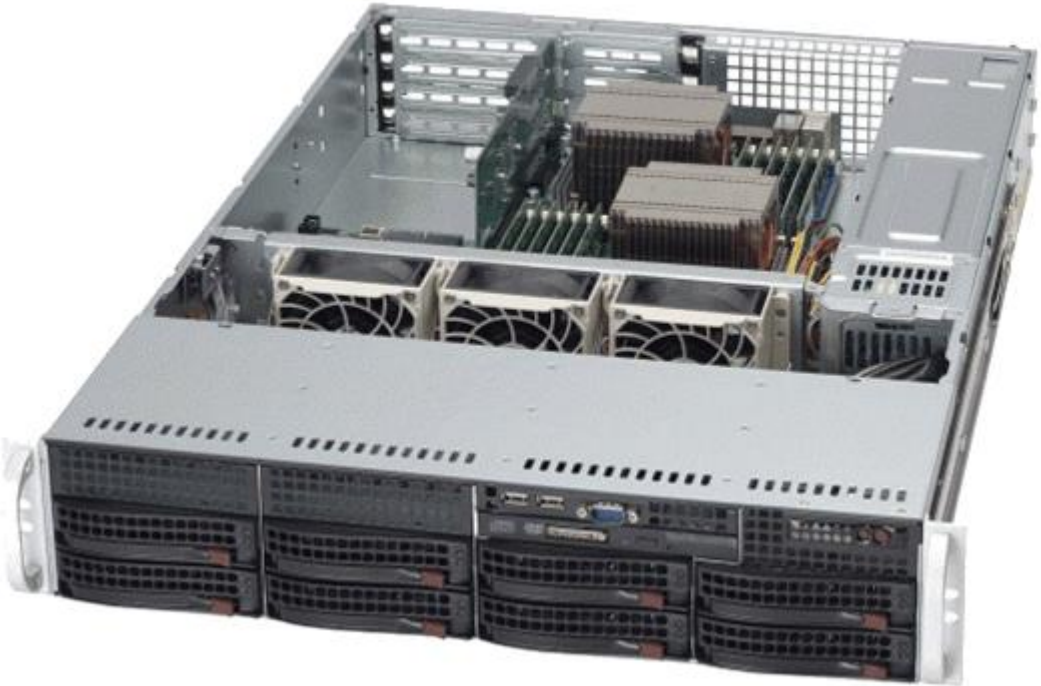

Where best to plug in?

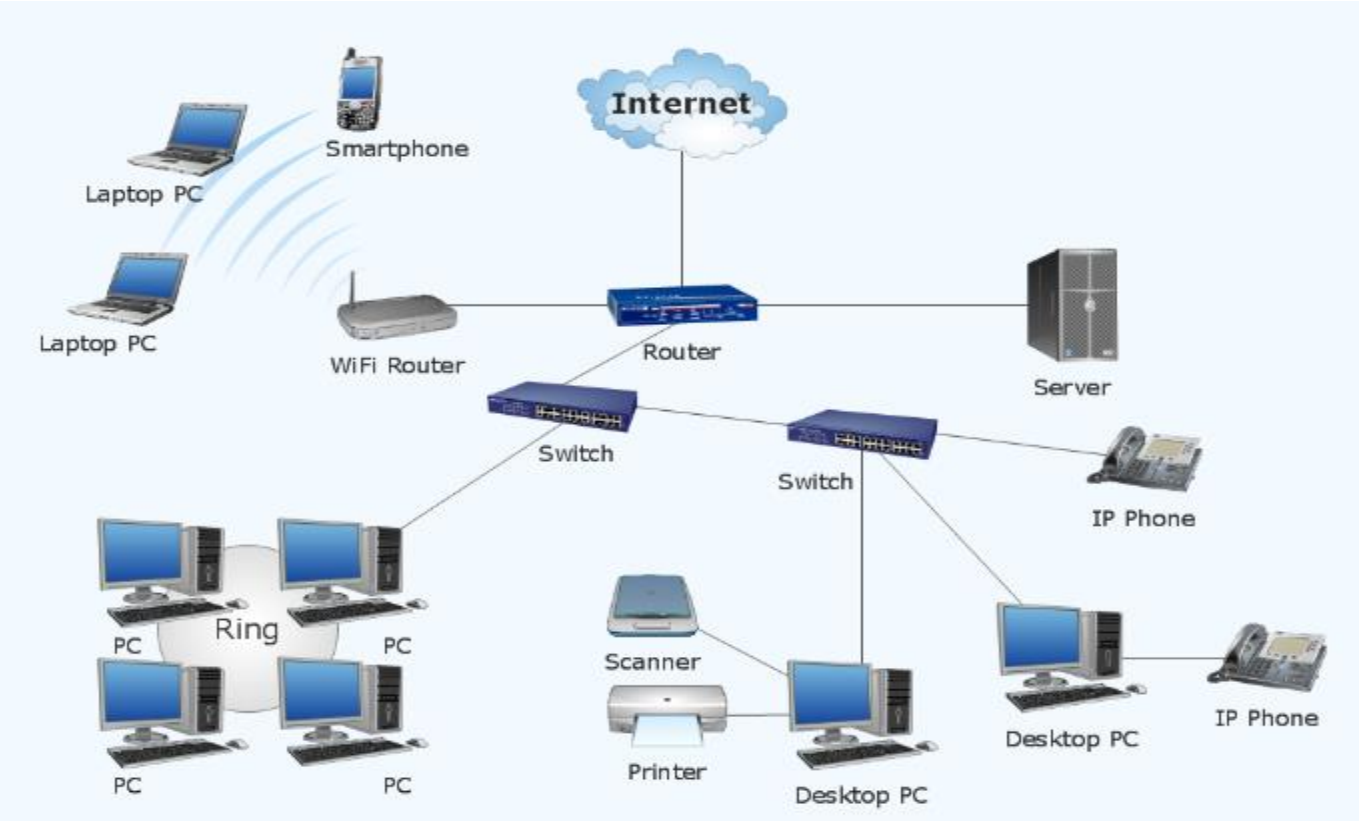

Some assembly required

Classic Network Architecture

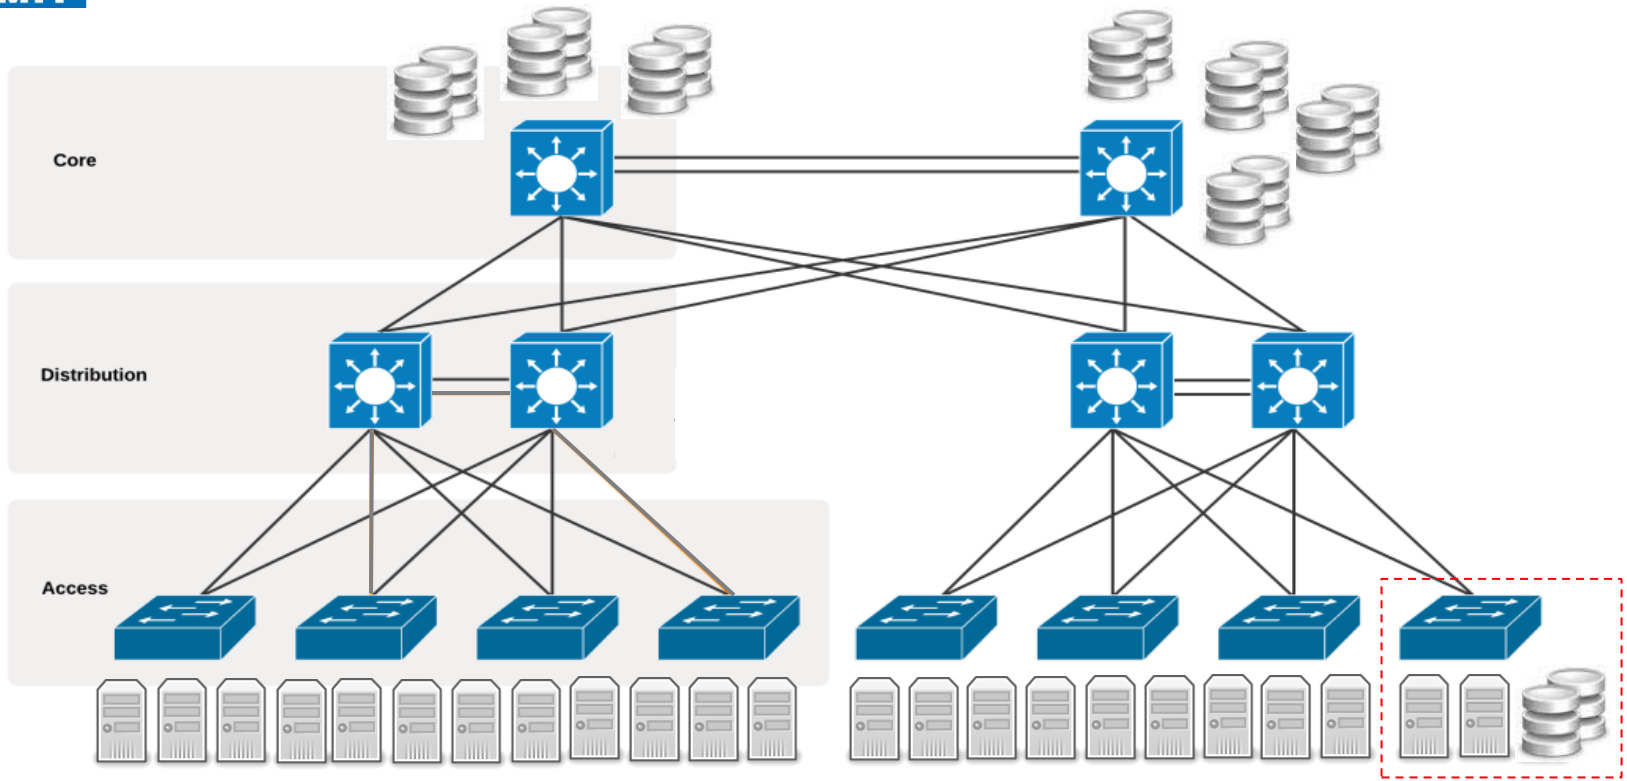

New Leaf-Spine Architecture

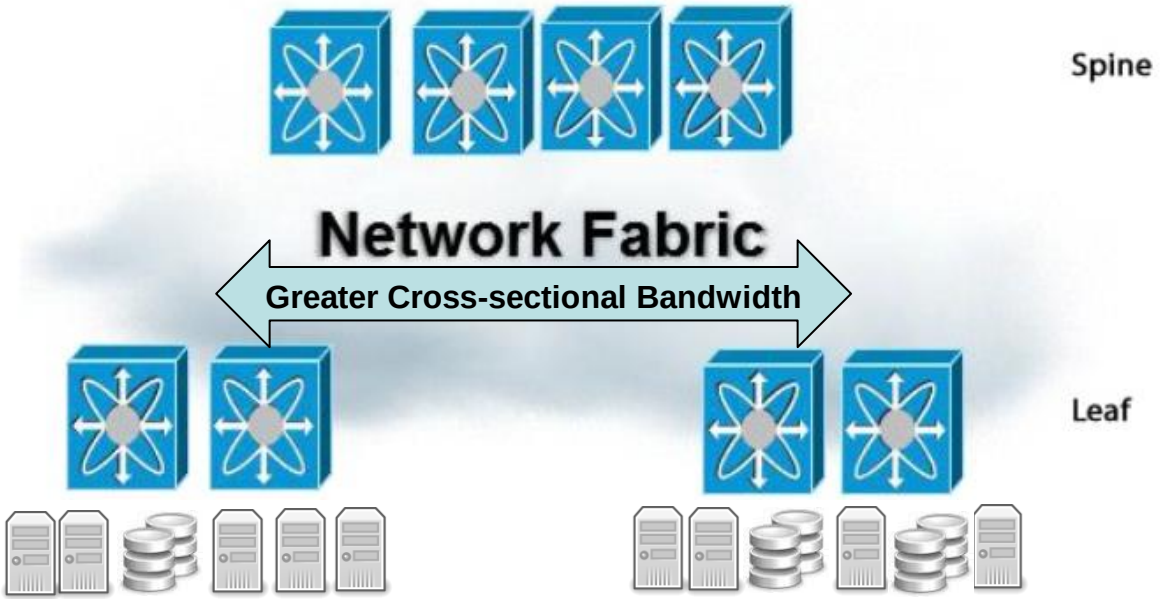

SAS/PCI Architecture

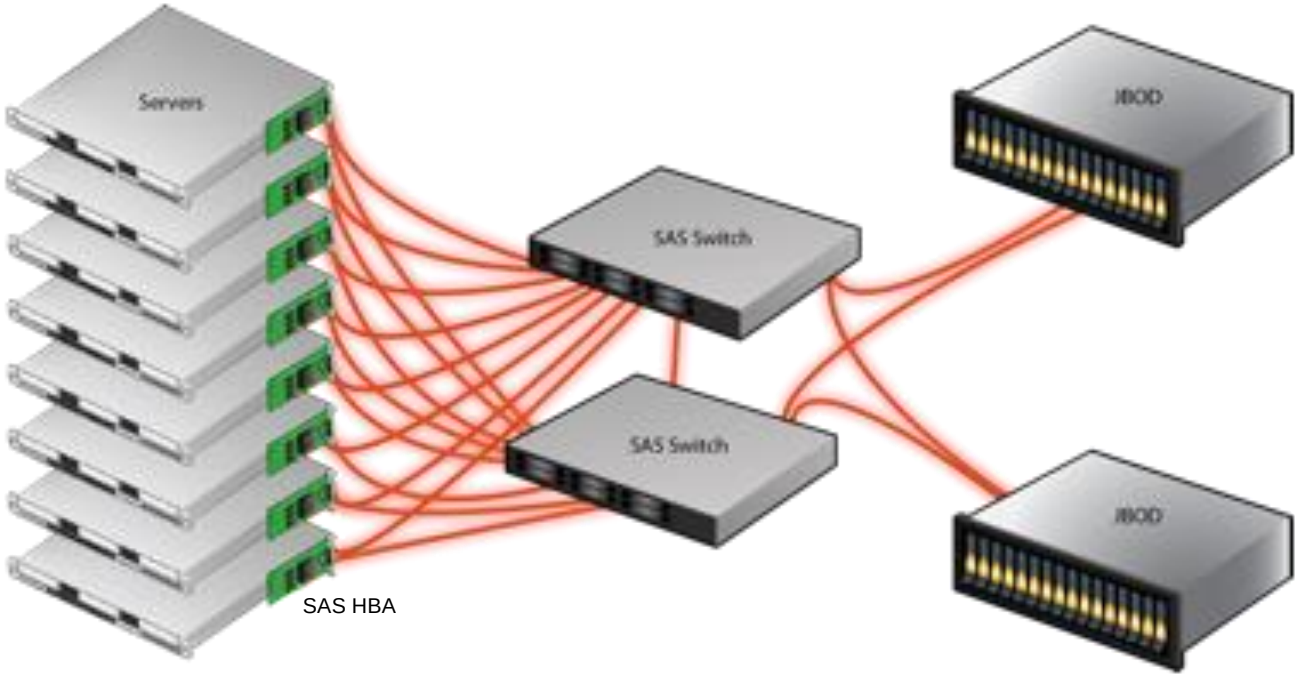

Scale-Out Architecture

- Scale-out grows capacity and performance in parallel
- White Box Servers with a clustered storage application
  - Ceph, ScaleIO(now EMC)
- Need high performance network
  - High bandwidth and low latency
- Flash storage is added to augment or replace disks for performance

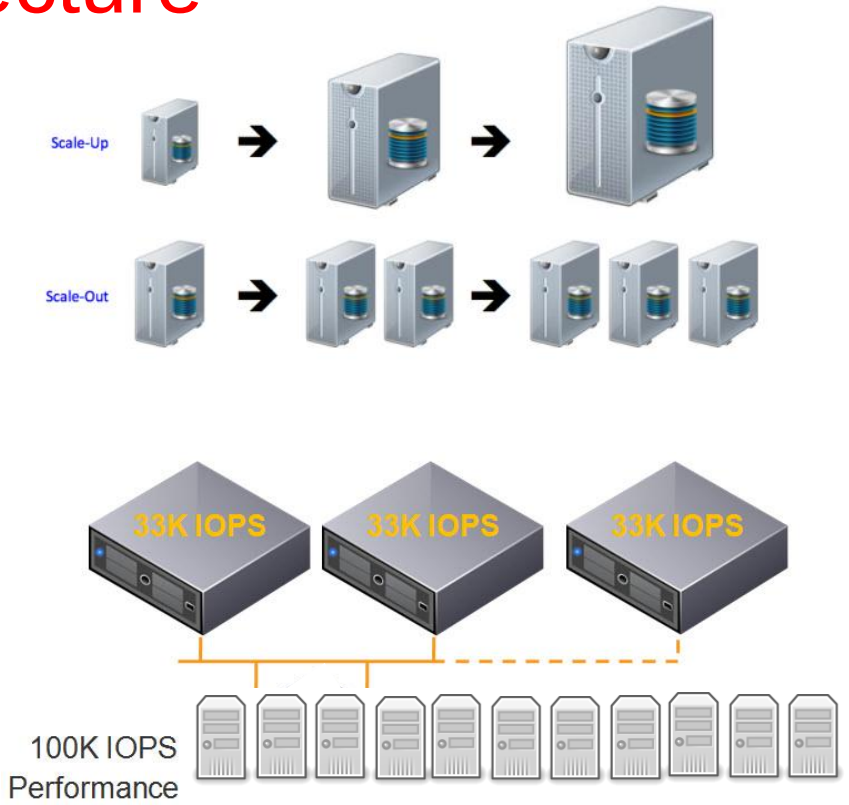

Hyper-Converged

- Collapse Separate Compute & Storage
  - Integrated Compute-Storage nodes
- Integrated workload
  - Hadoop, MongoDB, Nutanix
  - SMB-D, V-SAN

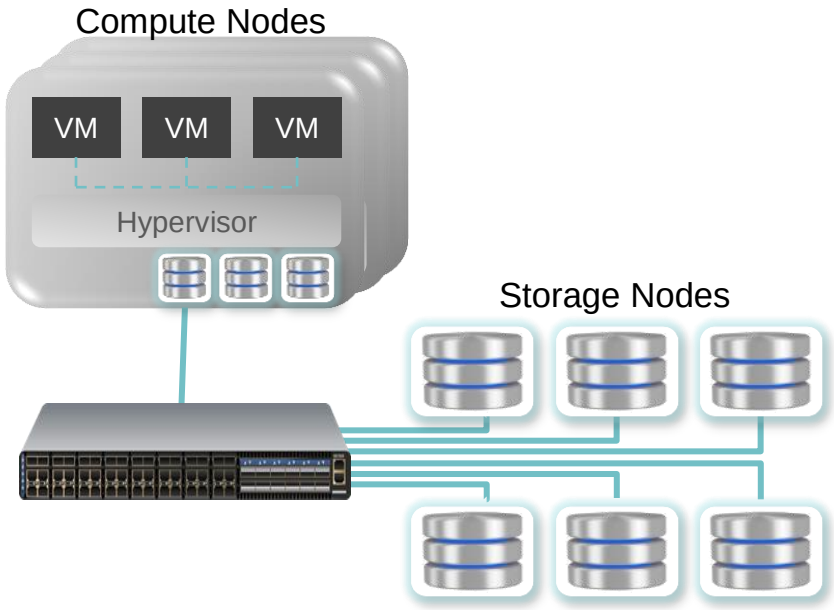

What are the tradeoffs: Local vs. Shared

- Performance: Depends on remote and local controller, network performance and remote controller load
- Trade-offs:
  - Better utilization of capacity
  - Scalability
  - Management ease
  - Applications
    - Server Virtualization
    - Cluster
  - Fault Isolation

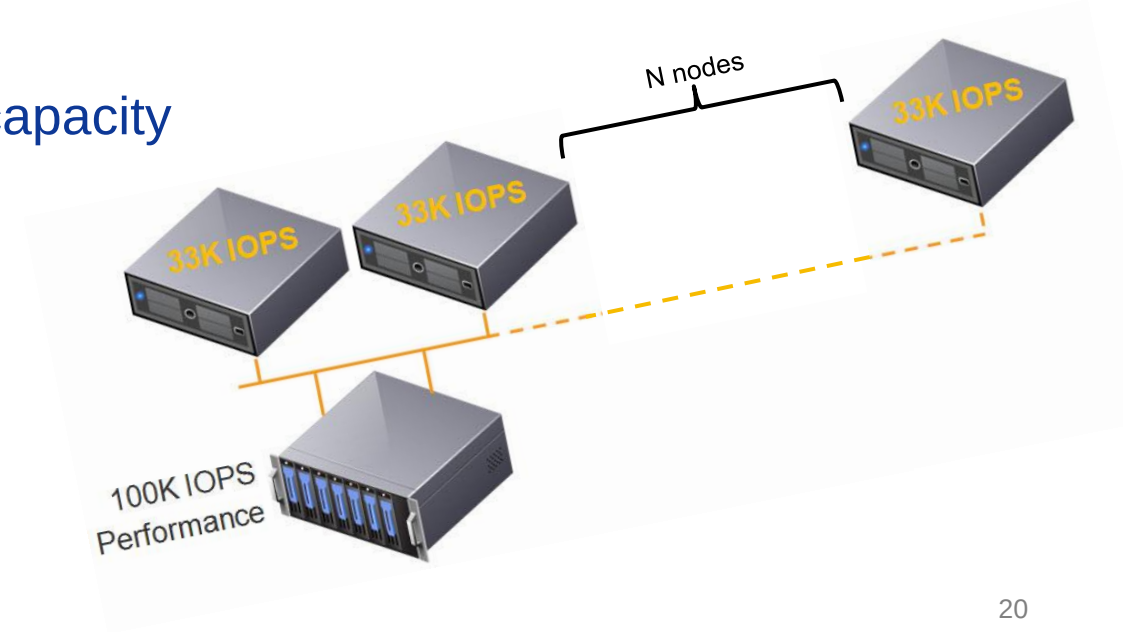

Questions?

Rob Davis

robd@mellanox.com