Benefits of Double Data Rate (DDR)

OPENFABRICS
ALLIANCE

Steve Lyness
Vice President, Worldwide Support
SilverStorm Technologies

Session Agenda

- InfiniBand SDR Interconnect … Well Known Benefits
- 20Gb/s InfiniBand DDR … Taking it to the Next Level
- DDR Customer Deployments …Realizing the Benefits

Benefits of InfiniBand SDR

- High Performance Interconnect
  - 2.5 to 30Gbit/sec
  - 3 microsecond or less end-to-end latency
  - RDMA based
- Highly Reliability
  - Loss less, self managing end-to-end fabric
  - Multi-path redundancy
  - No single point of failure
  - Standards based
- Extremely Scalable
  - Up to 48K local nodes, up to 2^{128} total
  - Copper and Fiber interface support
- Supports Virtual Networking
  - Multiple protocol networks on a single wire
  - Built-in QOS
  - Gateway support to 10 GigE and Fibre Channel

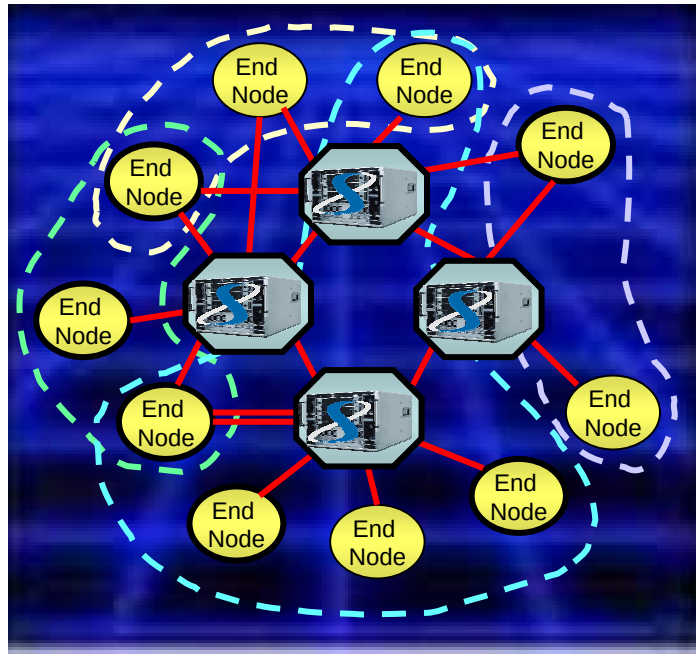

Benefits of InfiniBand SDR

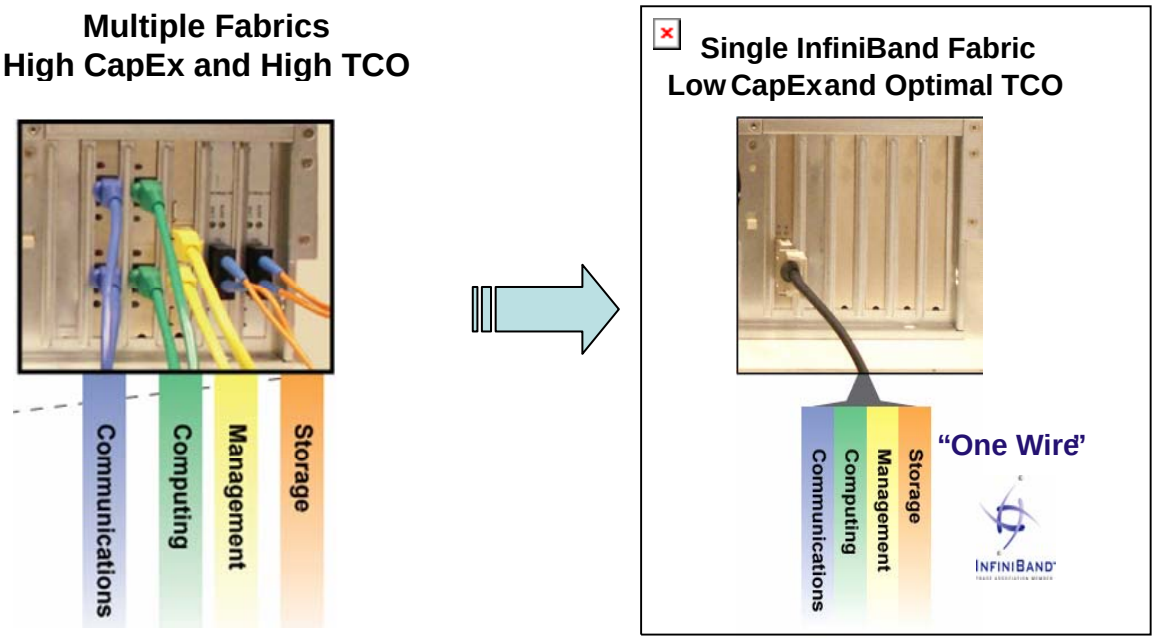

- "One Wire" Virtual I/O Interconnect
- Simplified Management
- Ultimate Scalability
- Fabric I/O Instead of Device I/O
- Optimal Total Cost of Ownership

Today's Cluster Computing Environment

- Number of CPU Cores Increasing in Density
- PCI Express Adoption Expanding Rapidly
- More Options Available for Highly Parallel Application Software
- Cost Pressures are Driving Network Consolidation and On-demand HPC networks
- Demand for Very Large Clusters Increasing

InfiniBand DDR - Taking it to the Next Level

- Industry's Highest Performance Interconnect Solution
  - 20 Gb/s node-to-node, 60 Gb/s switch to switch
  - Only Interconnect that can saturate an 8x PCI Express link
- Requires Less Infrastructure, Power and Space
  - Highest throughput interconnect silicon available
  - Supports more bandwidth and switching capacity per unit of rack space
  - Money saved on interconnect can be spent on more compute
- Minimizes Cluster Network Congestion
  - Big fabric pipes mean less path contention
- Supports Multi-Thousand Node FBB Scale-out
- Industry's Best Interconnect Price Performance

DDR Reduces Infrastructure Costs

1024 Server Node 10Gb/s Compute Cluster

Typical Fat Tree Topology

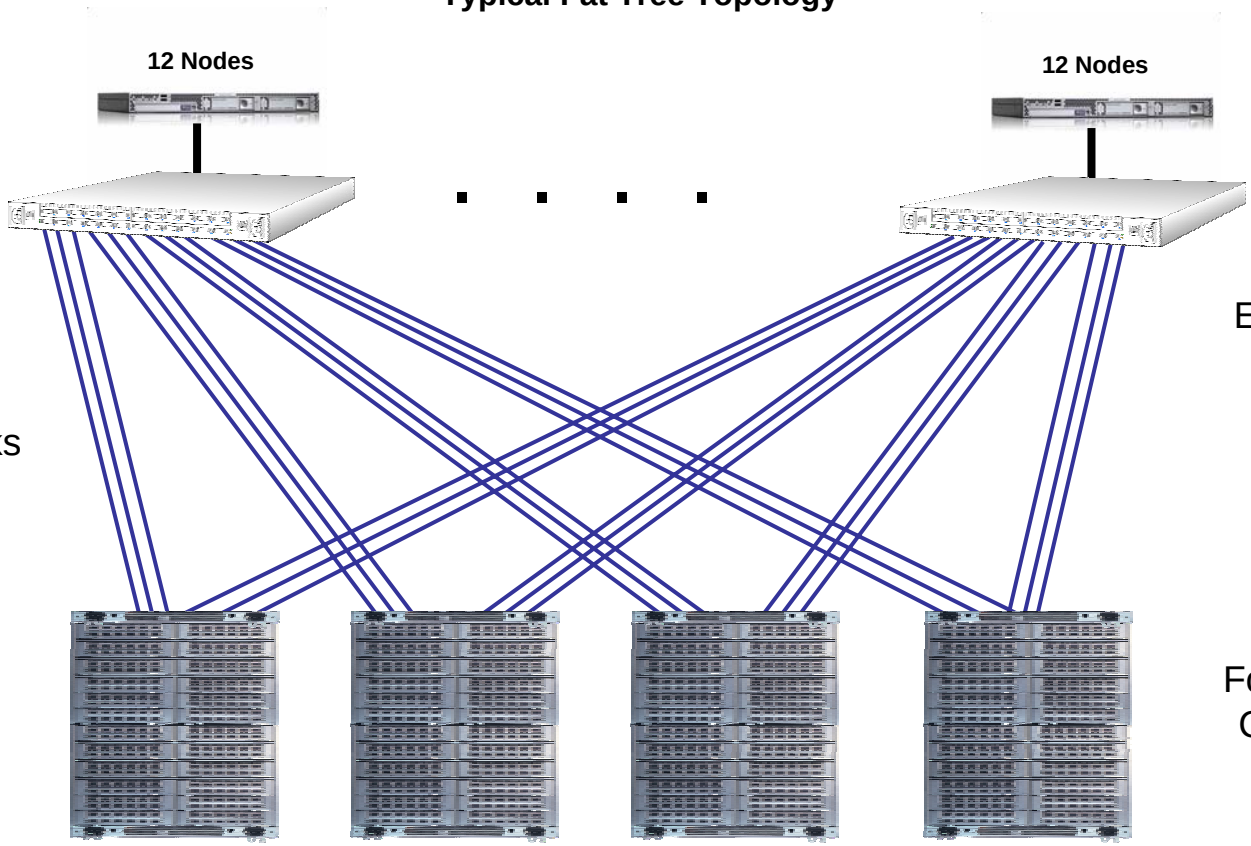

SDR
4x Links

86 x 24-Port
Edge Switches

Four x 288-Port
Core Network
Switches

DDR Reduces Infrastructure Costs

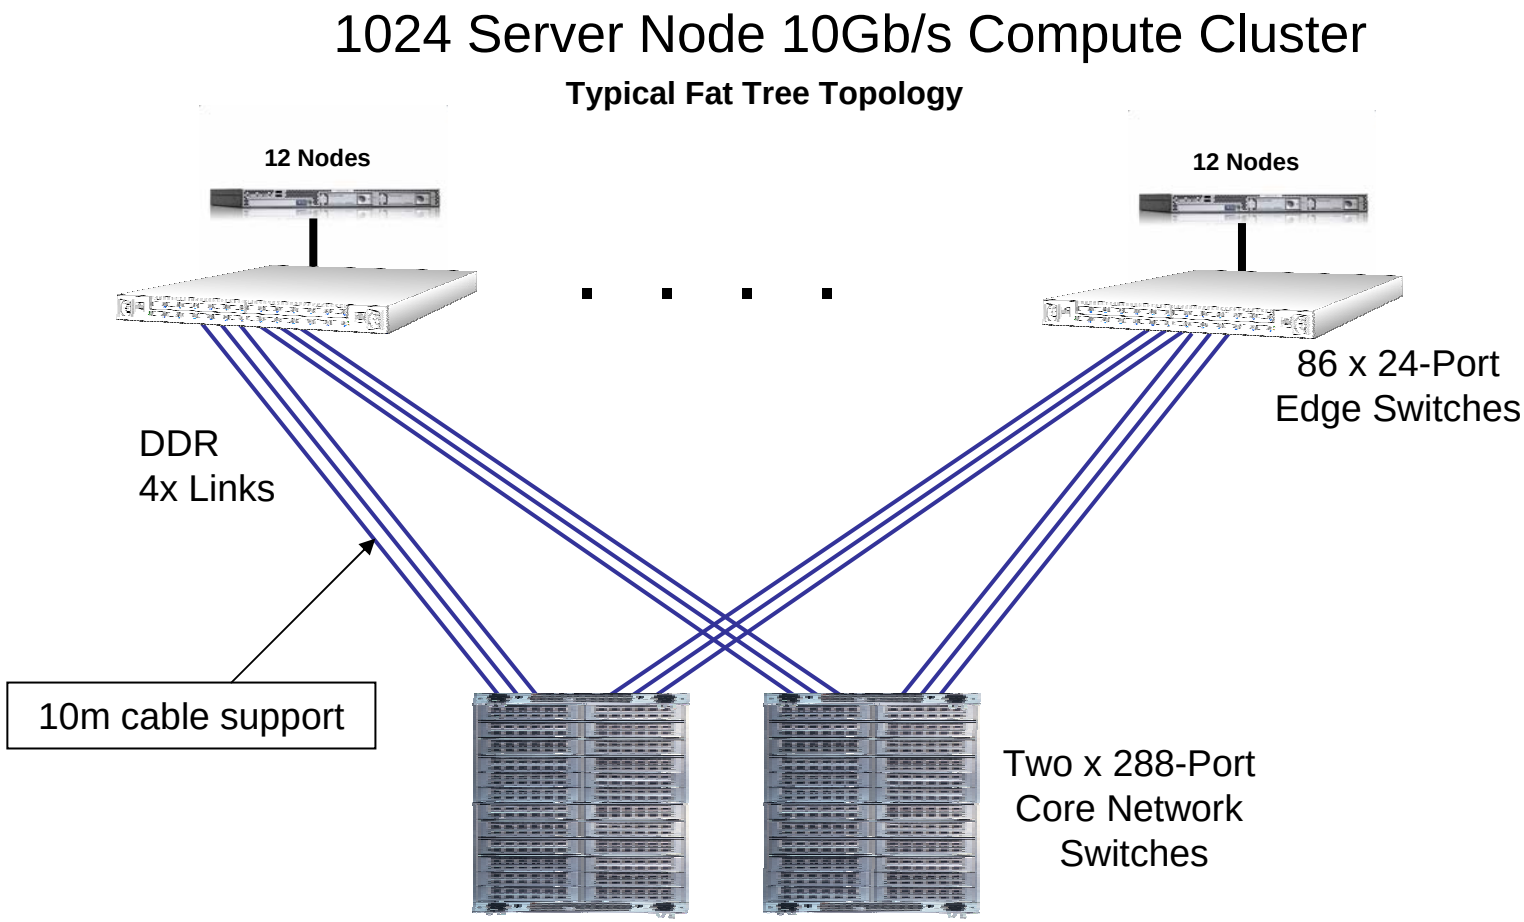

30% Less ports, cables, power & space at 20% Lower Cost

DDR Minimizes Network Congestion

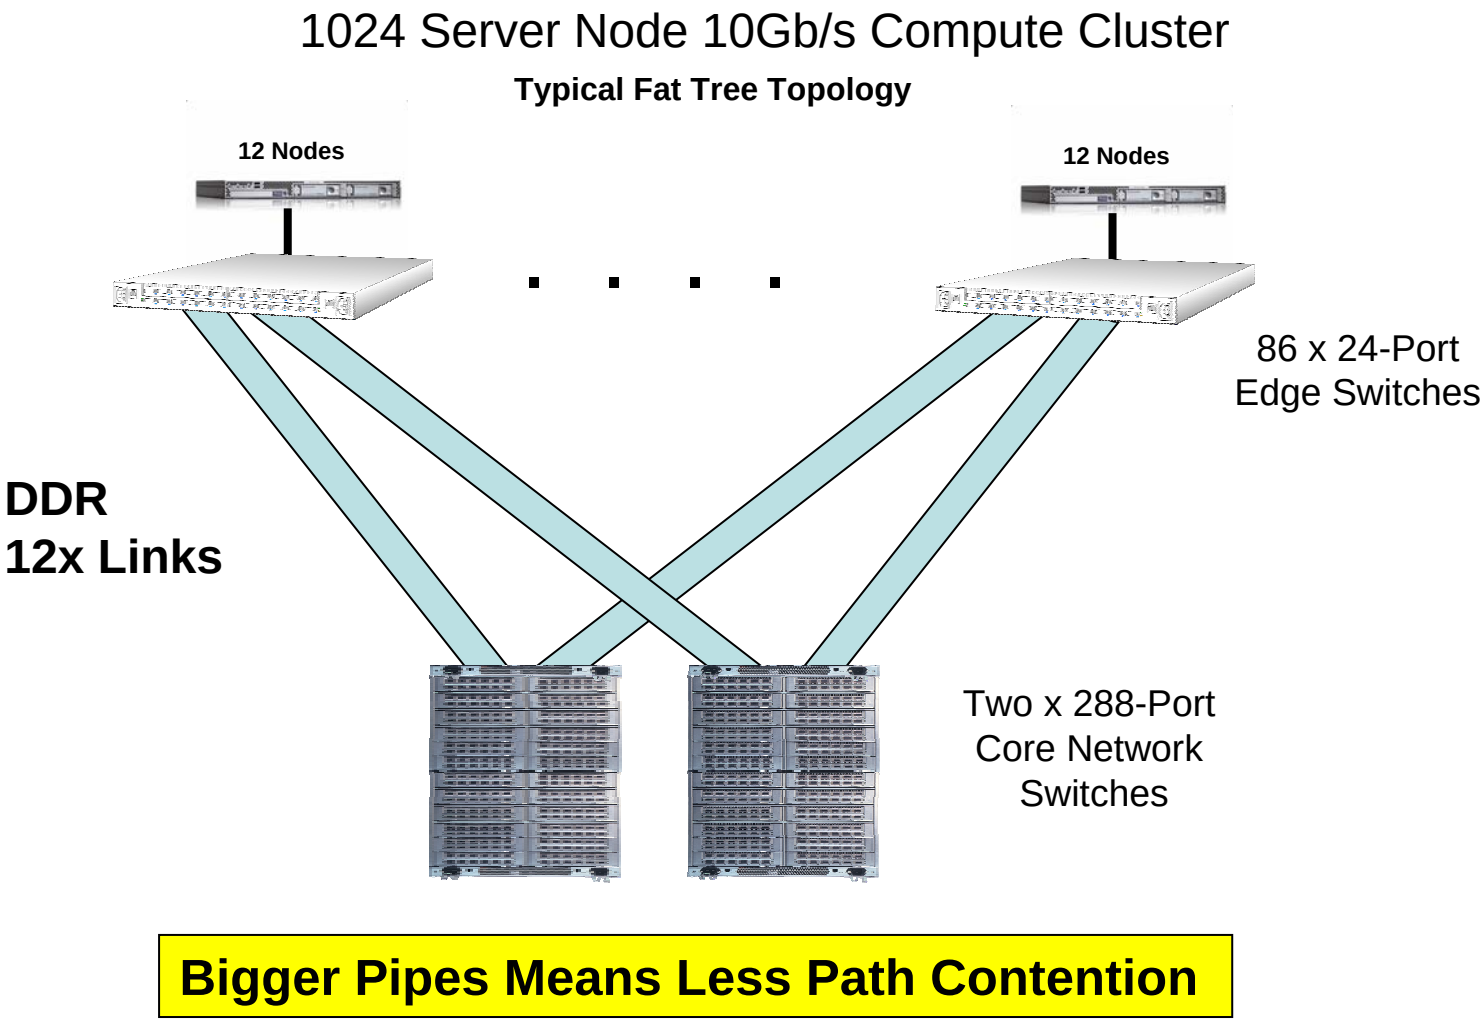

DDR Reduces Infrastructure Costs

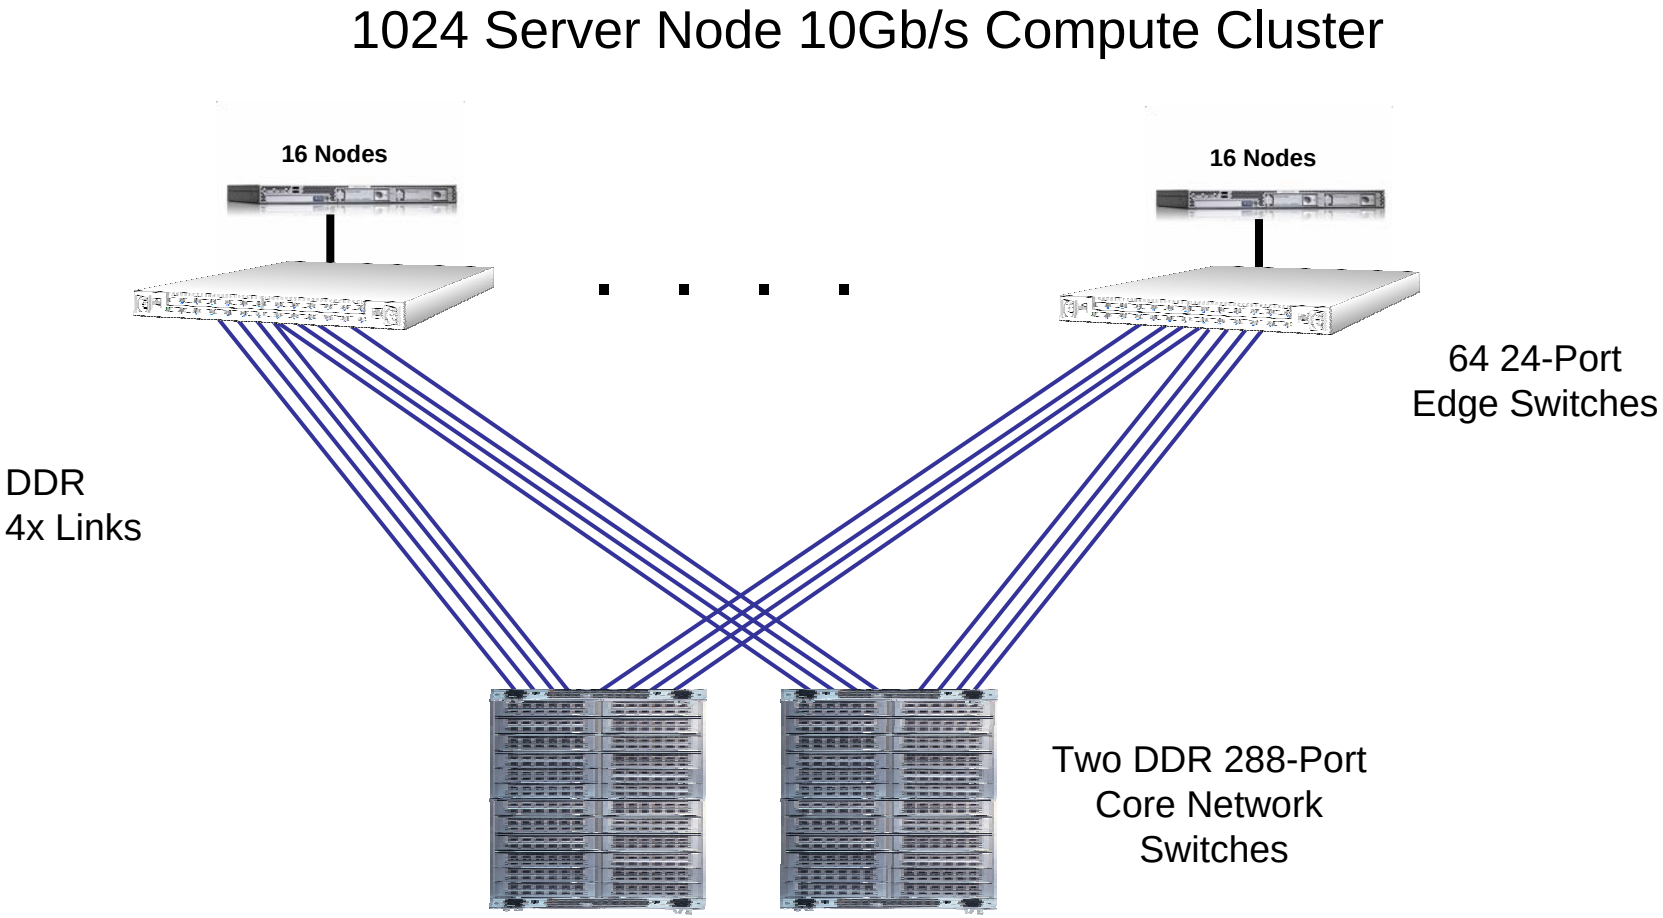

30% Less ports, cables, power & space at 20% Lower Cost

Lends well to Scalable Unit Concept

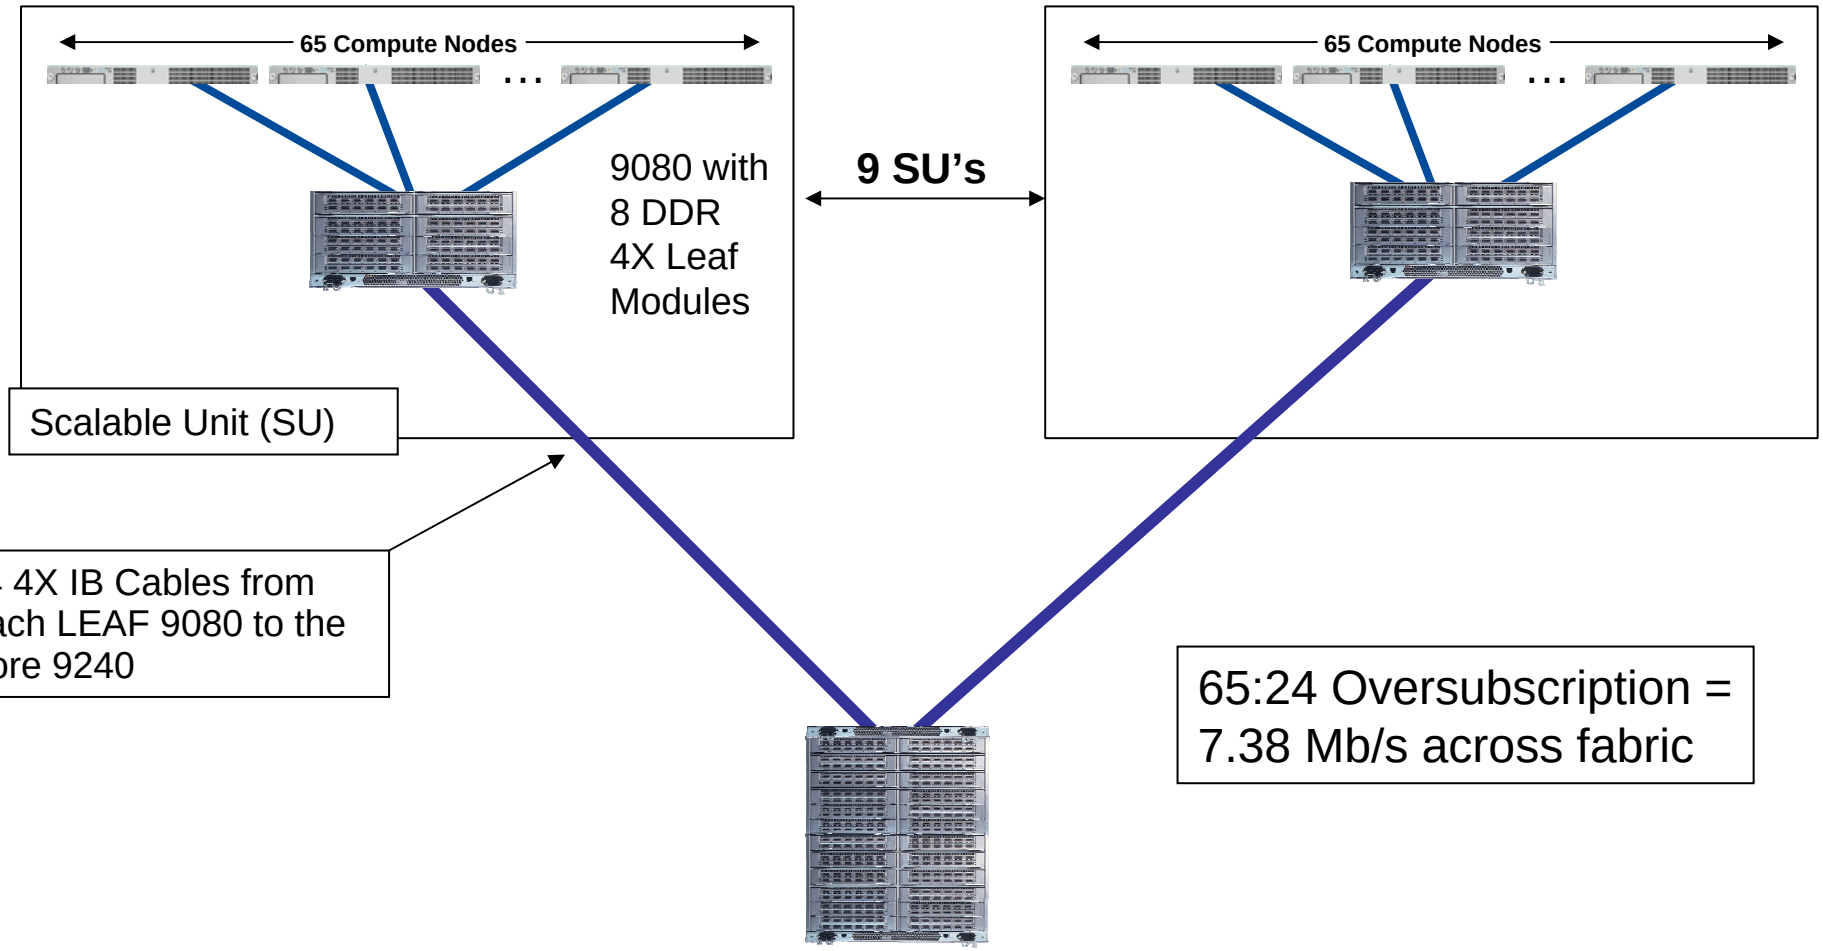

One 9240 Core Switch with Eighteen (18) 4x DDR Leaf Modules

Lends well to Scalable Unit Concept

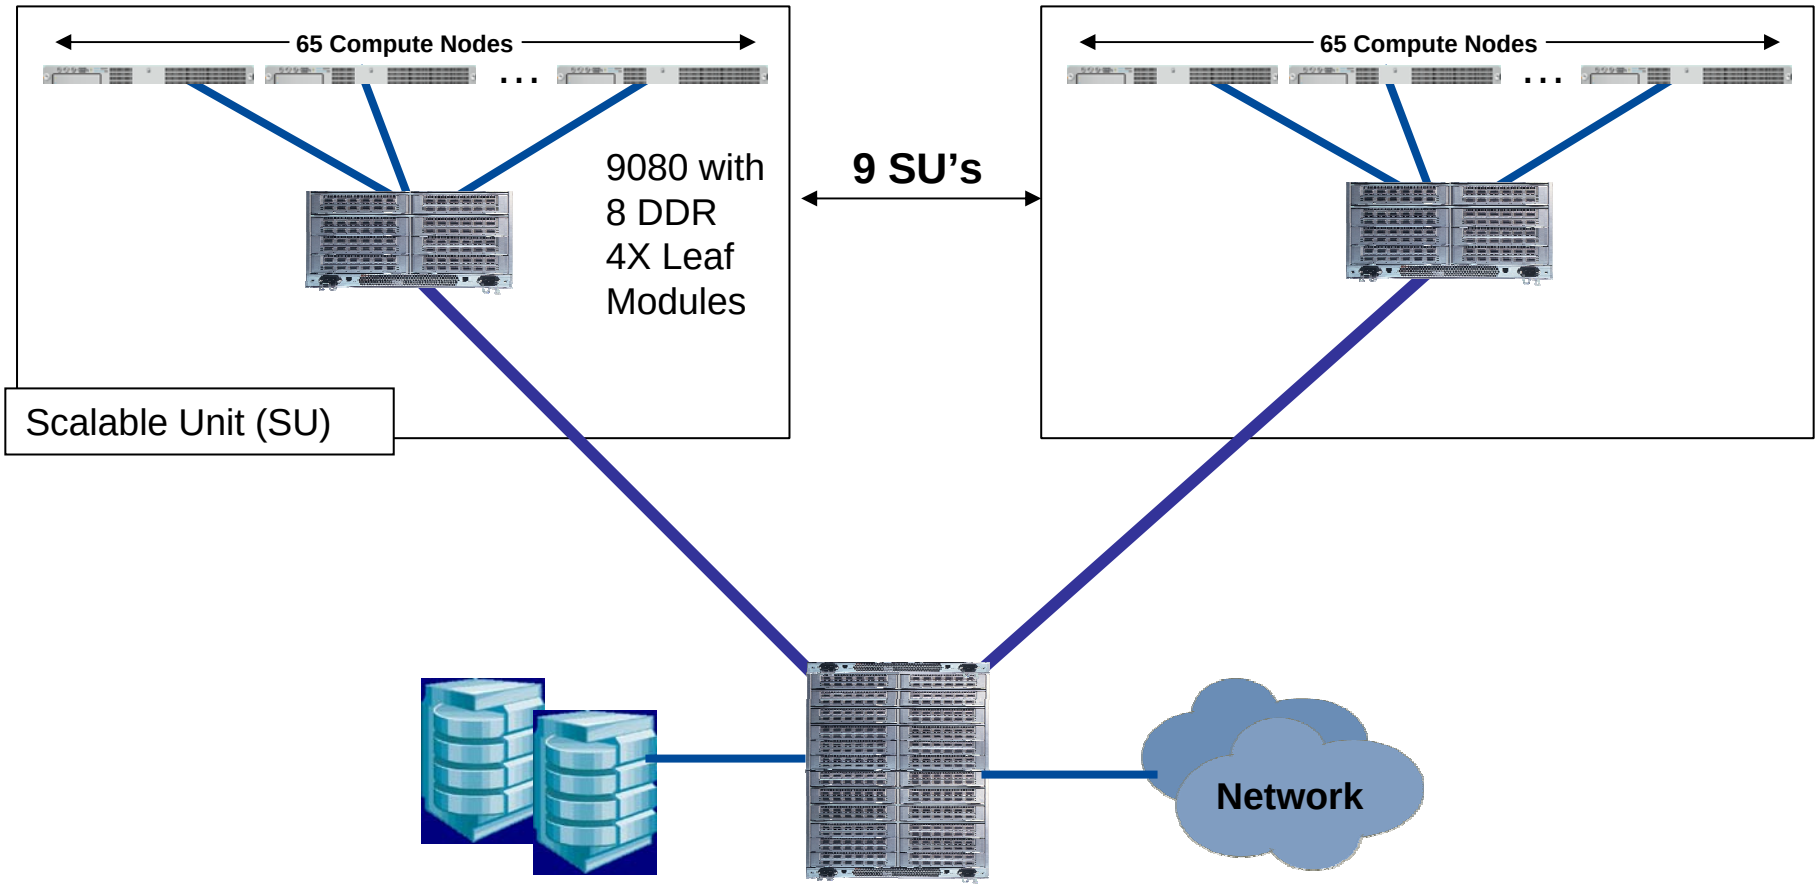

One 9240 Core Switch with Eighteen (18) 4x DDR Leaf Modules

PLUS 10 GigE Gateway

PLUS Additional ports for storage AND/OR Fibre Channel Gateway

Perfect for High Density VIO

- Using only two of the DDR IB ports allows you to add 2x 10GigE or 8x4Gb FC gateway ports on a leaf leaving eight (8) 4X DDR IB ports as well!

- 288 port IB switch can then support
  - 256 x 4X IB ports PLUS 32 x 10 GigE ports
  - 256 x 4X IB ports PLUS 128 x 4Gb FC ports

Large Scale DDR Deployments

- **US Department of Defense**
  - 5 different Supercomputing Clusters
  - Weapons and Classified Research
  - Supporting 80 TFLOPS of Compute Power
  - Two Rank in Top Twenty of Worlds Largest

- **NASA**
  - Goddard Space Center Supercomputer
  - Weather, Climate and Astrophysical Simulations and Research
  - Designed to Scale to 40 TFLOPs

THANK YOU!