



# HPC Update: Engagement Model

**MIKE VILDIBILL**

Director, Strategic Engagements  
Sun Microsystems

[mikev@sun.com](mailto:mikev@sun.com)

# Our Strategy

## Building a Comprehensive HPC Portfolio that Delivers Differentiated Customer Value

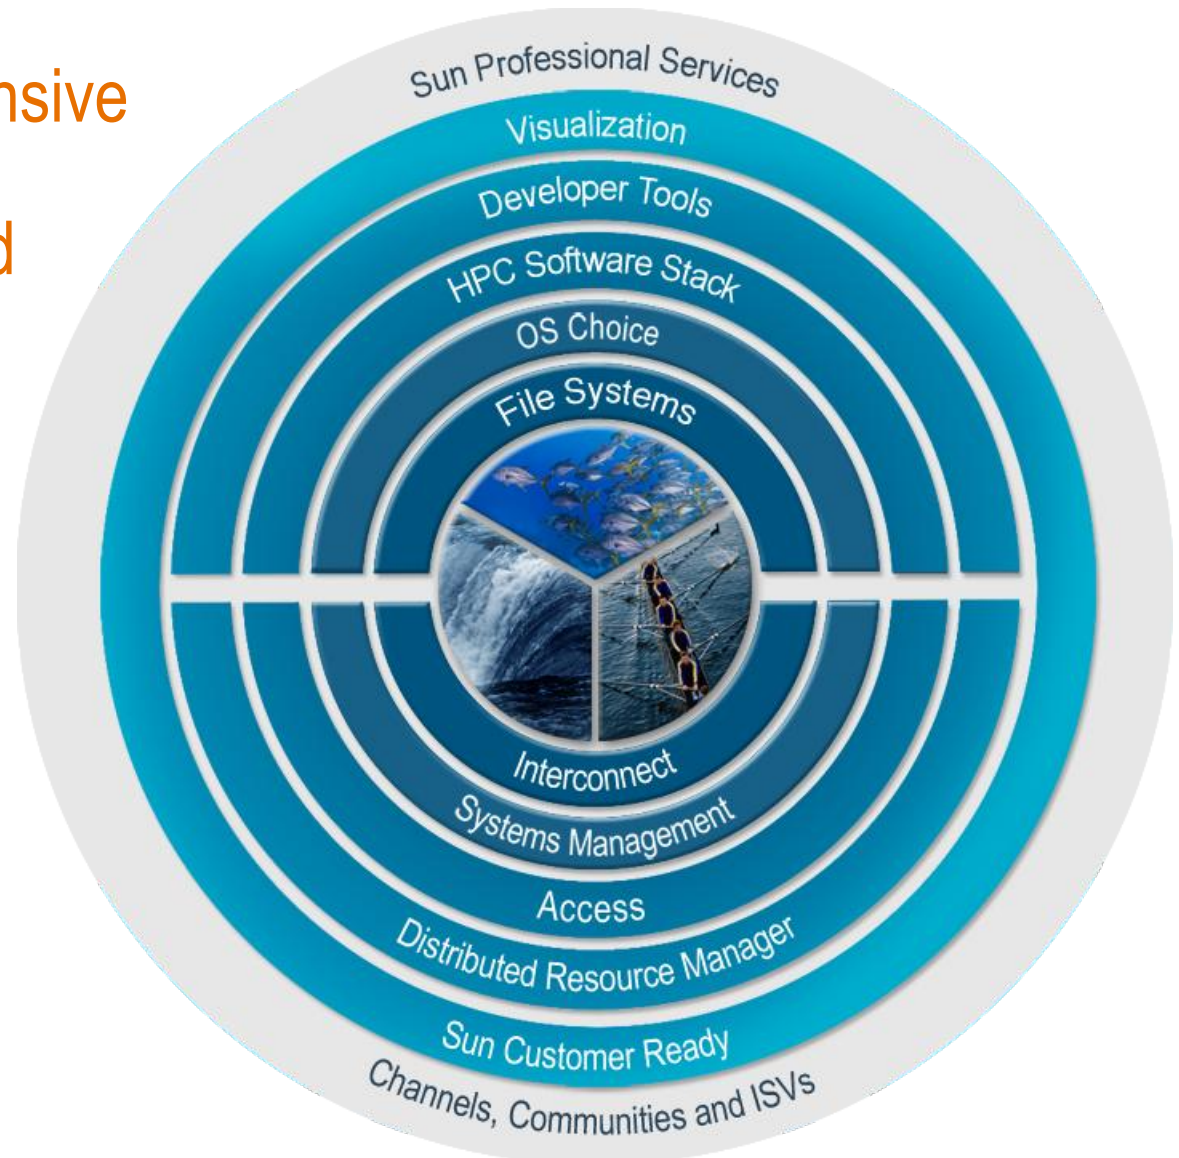
Providing complete and integrated solutions for targeted markets by combining unique Sun IP with best-in-class products

Innovation built on Open Standards at the component, product & solution level

Making HPC simpler and easier

Ensuring design and delivery excellence using strategic engagement teams

Extending value through partnerships



# “Strategic Engagements” Team

- 1.5 years of effort focusing on definition, design, modeling, prototyping and build of a portfolio of highly integrated standards-based products that address high-end HPC and Enterprise needs
- Emphasis on technology adoption; engagement with HPC customers on bids; 20+ large systems pre-sold
- Strong alignments into engineering, manufacturing, marketing, services and delivery organizations to ensure “integrated offerings” as opposed to just integrated systems

# Sun Constellation System Success

**TACC:** World's largest  
Open Supercomputer



**KISTI:** Largest supercomputing  
center in Korea



**Jülich:** Germany's largest HPC  
center



**Sandia:** Worlds first & largest  
QDR Torus configuration



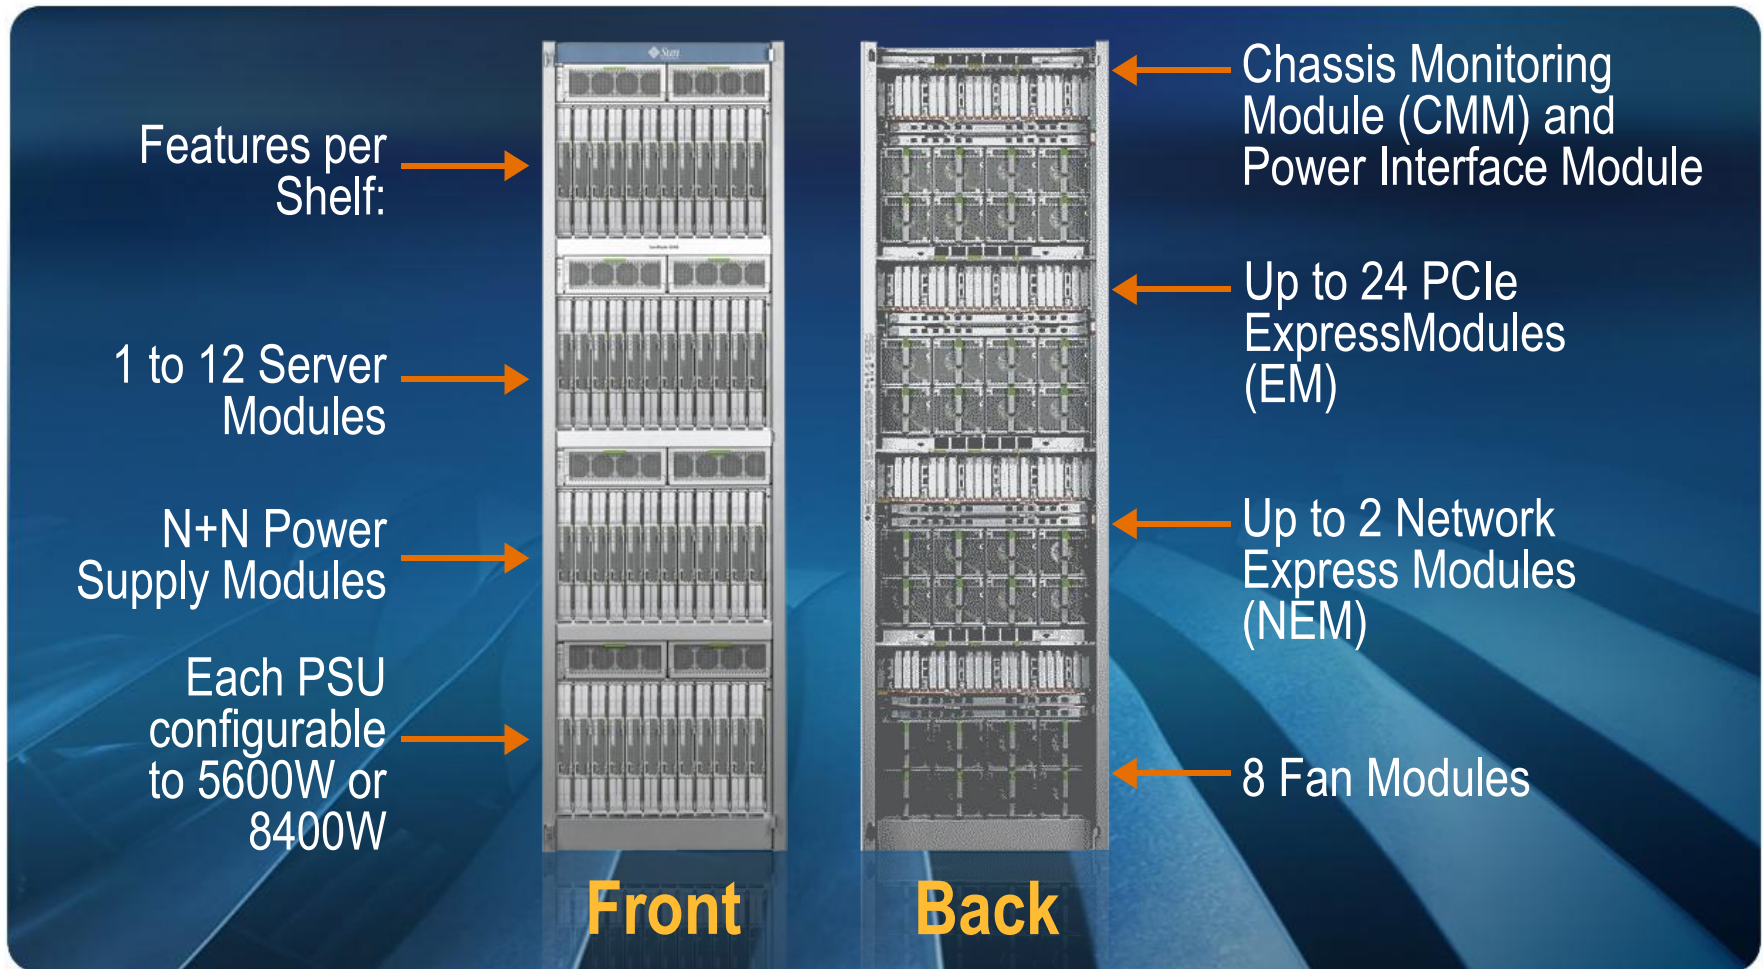
# Launched; Deliveries are underway

- Blade racks that support industry standard PCIgen2 EMs and 48 “no compromise” blades; integrated cooling doors (water or refrigerant) that achieve room neutral cooling at 32-35kW per rack
- 2x2 HPC blades with two dual-socket 95w Nehalem nodes & two on-blade QDR (10 GHz) InfiniBand host adapter chips
- QDR 48-port leaf switch (NEM) integrated into the rack
- 648-port QDR IB switch for non-blocking Clos; optical 12x cables
- JBOD w/ SW RAID running Lustre
- Configuration options include cost optimized mesh/Torus, dual rail IB, non-IB, variety of Intel, AMD and SPARC blades, 3<sup>rd</sup> party products

A popular rack config: 192 Nehalem CPUs (9 peak TFLOPs), 6.1 TBytes/sec peak memory bandwidth, and 0.77 TBytes/sec of peak InfiniBand bandwidth (only 32 cables), with simple, non-blocking scaling to 486+ TFLOPs and custom configs scaling to 2+ PFLOPs

# The Chassis is the Rack

Up to 48 Server Modules, Up to Eleven TFLOPS per Rack



# Looking Forward


- There exists tremendous opportunity for innovative, open, industry standard HPC solutions
- Dramatic cost, performance, and operational efficiencies are achievable through intelligent system design
- The synergies between broad open community adoption and high-end HPC offerings remains strong



[www.sun.com/  
sunconstellationsystem](http://www.sun.com/sunconstellationsystem)

# Sun Constellation System

The World's First Open Petascale Environment



Scales to over  
two Petaflops

20% floor space reduction

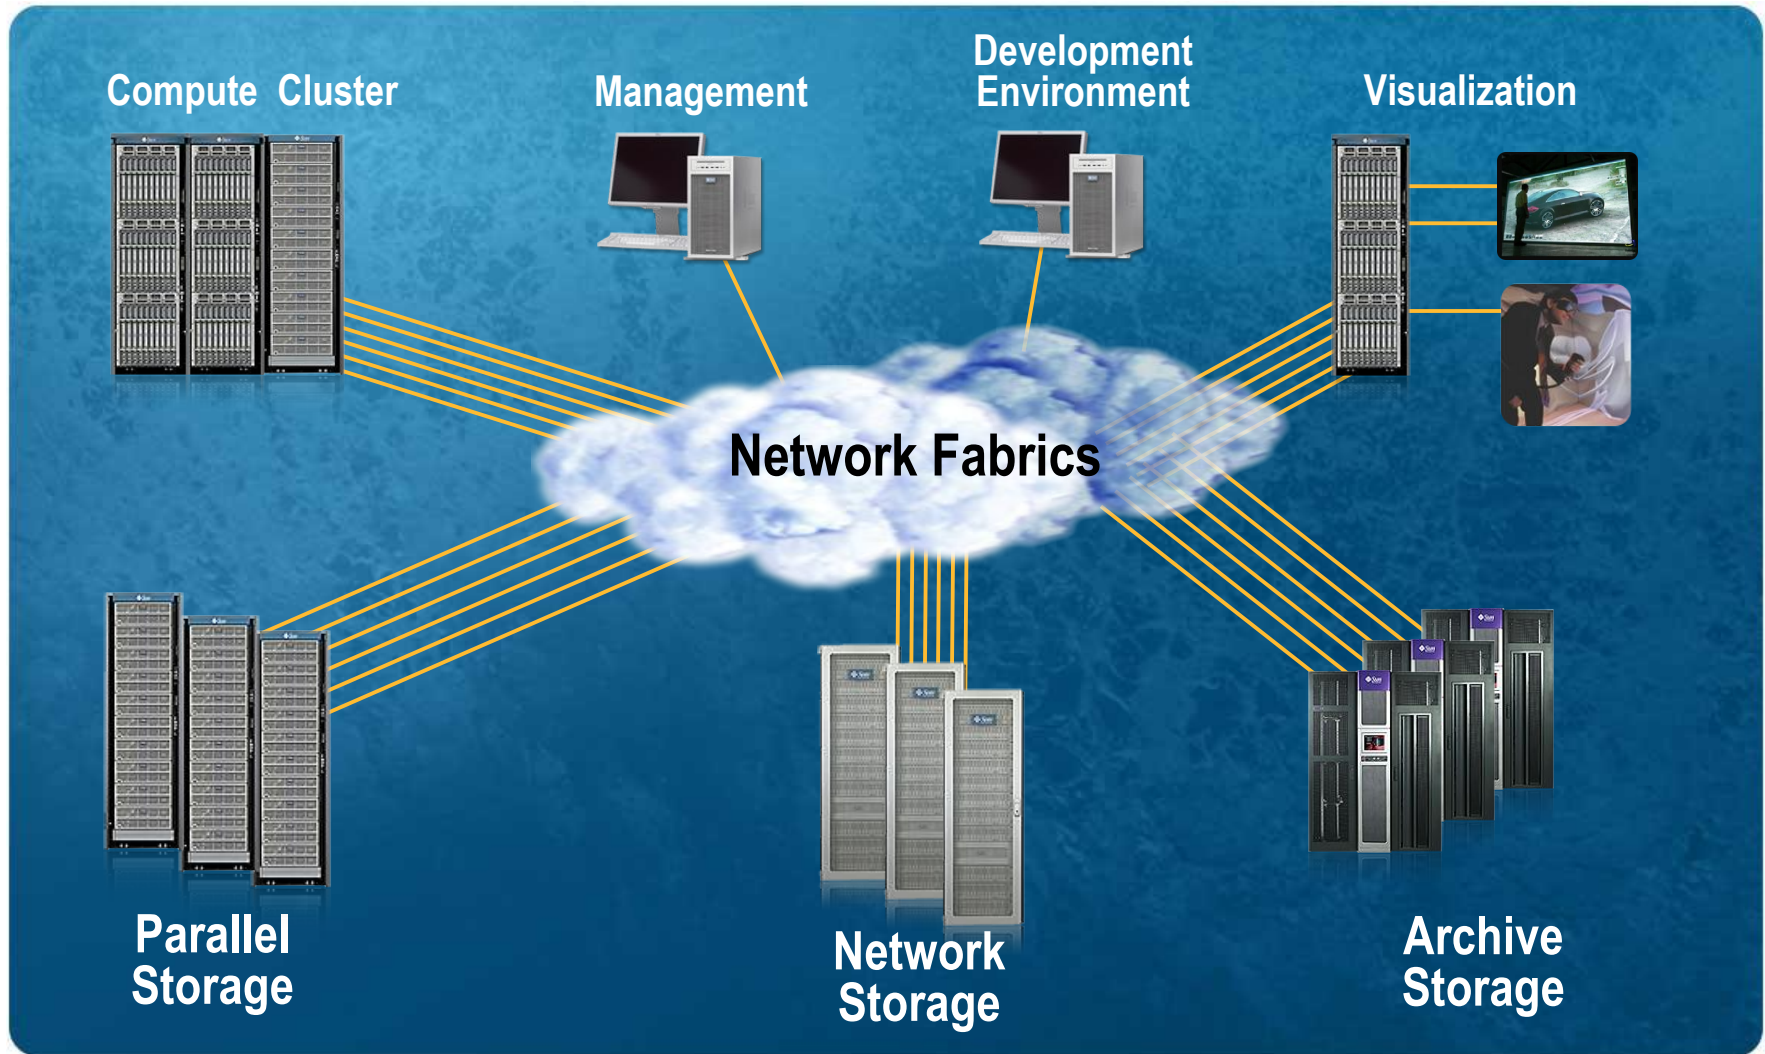
6x fewer cabling required

50% more compute density\*

[sun.com/sunconstellationsystem](http://sun.com/sunconstellationsystem)

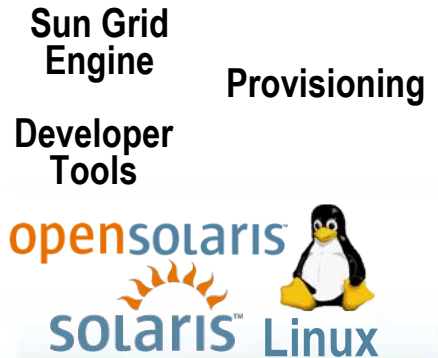
\*Using quad socket blades in a 42RU Rack

# The Complete Picture



# Sun Constellation System Open Petascale Architecture

## Building Blocks



Compute	Networking	Storage	Software
<p><b>Ultra-Dense Blade Platform</b></p> <p>Fastest processors: SPARC, AMD Opteron, Intel Xeon</p> <p>Highest compute density</p> <p>Fastest host channel adaptor</p>	<p><b>Ultra-Dense and Ultra-Slim Switch Solutions</b></p> <p>72 and 3456 port InfiniBand switches</p> <p>Unrivalled cable simplification</p> <p>Most economical InfiniBand cost/port</p>	<p><b>Ultra-Dense Storage Solution</b></p> <p>Most economical and scalable parallel file system building block</p> <p>Up to 48 TB in 4RU</p> <p><b>Up to 2TB of SSD</b></p> <p>Storage controllers in the rack</p>	<p><b>Total HPC Cluster Management</b></p> <p>Integrated developer tools</p> <p>Scalable Distributed Resource Management</p> <p>Provisioning, monitoring, patching</p> <p>Simplified inventory management</p>

# Sun Cooling Door Systems

Up 6x More Efficient Than Traditional Cooling Systems

New

## Passive Rear Door Heat Exchanger Design

No Additional Fans means greater efficiency  
Up to 35KW Capacity

## Pumped Refrigerant Door (5600)

Datacenter safe R134A refrigerant  
Compatible with Liebert XD systems  
Highest Energy Efficiency and smallest footprint

## Chilled Water Door (5200)

Low investment for those already with water in the datacenter  
Economical for smaller installations  
Connects to bottom (raised floor) or top (ceiling) water supply source

## Fits in the Rear of

Sun Blade 6048 chassis



Sun Cooling  
Door 5600

Sun Cooling  
Door 5200

# Versatile Sun Blade Portfolio: Intel Xeon



	X6250	X6270	X6275	X6450
<b>CPU</b>	Up to 2 Xeon 5200/5400 Series Quad-Core	Up to 2 Xeon 5500 Series Quad-Core	2X2 Xeon 5500 Series Quad-Core	Up to 4 Xeon 7400 Series 4 or 6-Core
<b>Memory</b>	Up to 64GB (16 FBDIMM)	Up to 144GB (18 DDR3 DIMM)	Up to 192GB (2X12 DDR3 DIMM)	Up to 192GB (24 FBDIMM)
<b>On-blade GbE</b>	2	2	2	2
<b>On-blade SAS</b>	4X SAS 1.0	4X SAS 2.0	N/A	4X SAS 1.0
<b>On-blade InfiniBand</b>	N/A	N/A	2X QDR IB	N/A
<b>PCIe I/O</b>	2X PCIe 1.0 X8, 2X PCIe X4	4X PCIe 2.0 X8	2X PCIe 2.0 X8	2X PCIe 1.0 X8, 2X PCIe X4




# Versatile Sun Blade Portfolio: Intel Xeon



## Sun Blade X6270

- Two Nehalem-EP 4-core CPUs
- 18 DDR3 DIMM Slots, up to 144GB (8GB DIMMs)
- 4x PCI Express Gen2 x8 interfaces
- 2x GbE and 4x SAS2 via NEM
- Four 2.5-inch disks (SATA , SAS , SSD)
- Compact Flash Media
- Hardware RAID
- Full onboard ILOM Service Processor



## Sun Blade X6275

- New, high density, dual node blade server
- Each node contains
- Two Nehalem-EP 4-core CPUs
- 12 DDR3 DIMM Slots, up to 96GB (8GB DIMMs)
- 1x PCI Express Gen2 x8 interface
- 1x Gigabit Ethernet port via NEM
- 1x Sun Flash Module (24GB SATA) option
- 1x Dual Port QDR IB HCA (optional\*); interfaces with IB QDR Switch NEM (6048 chassis only)
- Full onboard ILOM Service Processor



# Versatile Sun Blade Portfolio: AMD Opteron



## Sun Blade X6440 Server Module

Four Quad-core AMD Opteron processors  
2.2, 2.5, 2.7 GHz with 6M L3 Cache  
32 DDR2 DIMM, Up to 256GB  
Hardware RAID: optional 0, 1, 5, 6





## Sun Blade X6240 Server Module

Two Quad-core AMD Opteron processors  
1.9, 2.1, 2.3, 2.5GHz  
16 DDR2 DIMM, up to 64GB  
Four 2.5-inch SATA or SAS disks  
Hardware RAID 0, 1, optional 5, 6



# Massive Scaling, Performance and Reduced Complexity



# Sun Blade 6048 InfiniBand QDR Switched Network Express Module



**New**

## Industry's Only Chassis QDR Integrated Leaf Switch

- Dual Port DDR FEM or onboard QDR IB HCA per server module
- 2 Passthru GigE ports per server module
- Leverages x8 PCIe 2.0 midplane
- Compact design maximizes chassis utilization

## Designed for Highly Resilient Fabrics

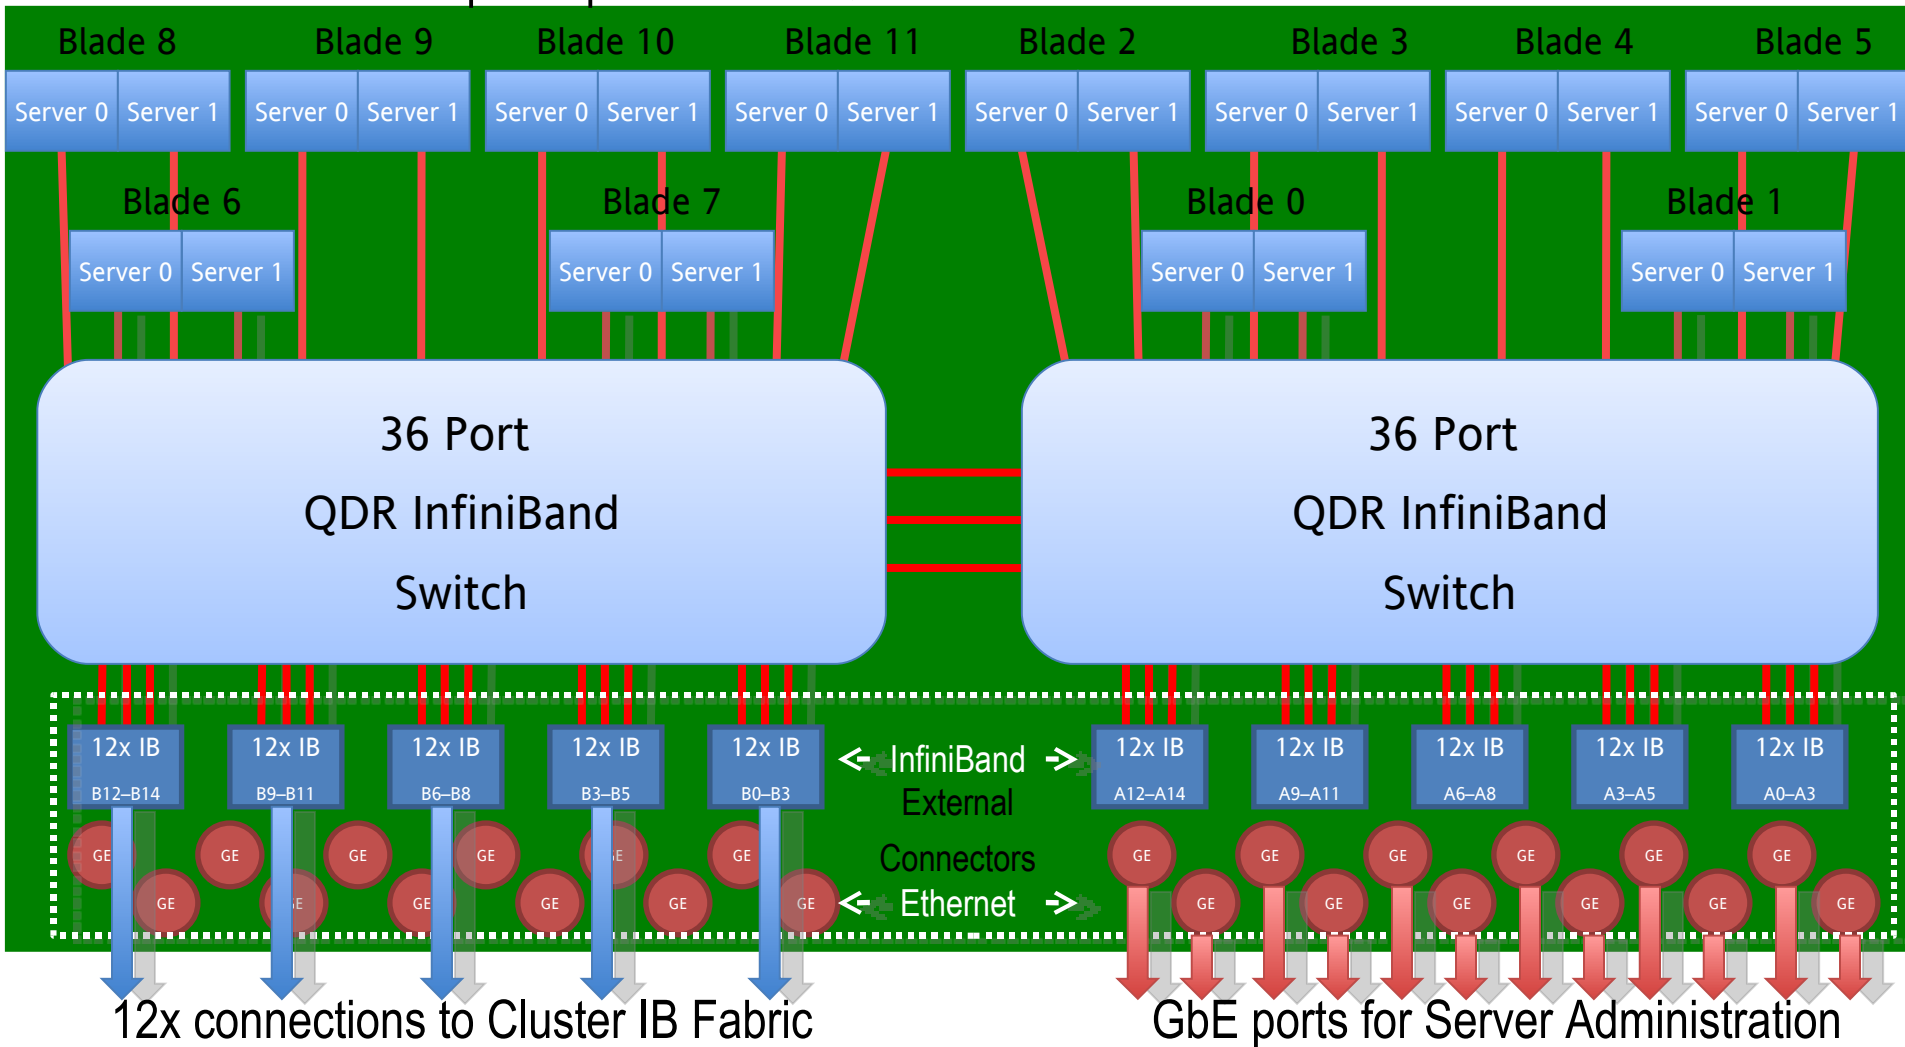
- Two onboard 36-Port QDR IB switches
- Increased resiliency allowing connectivity; up to 8 Ultra-dense switches

## 3:1 Reduction in Cabling Simplifies Cable Management

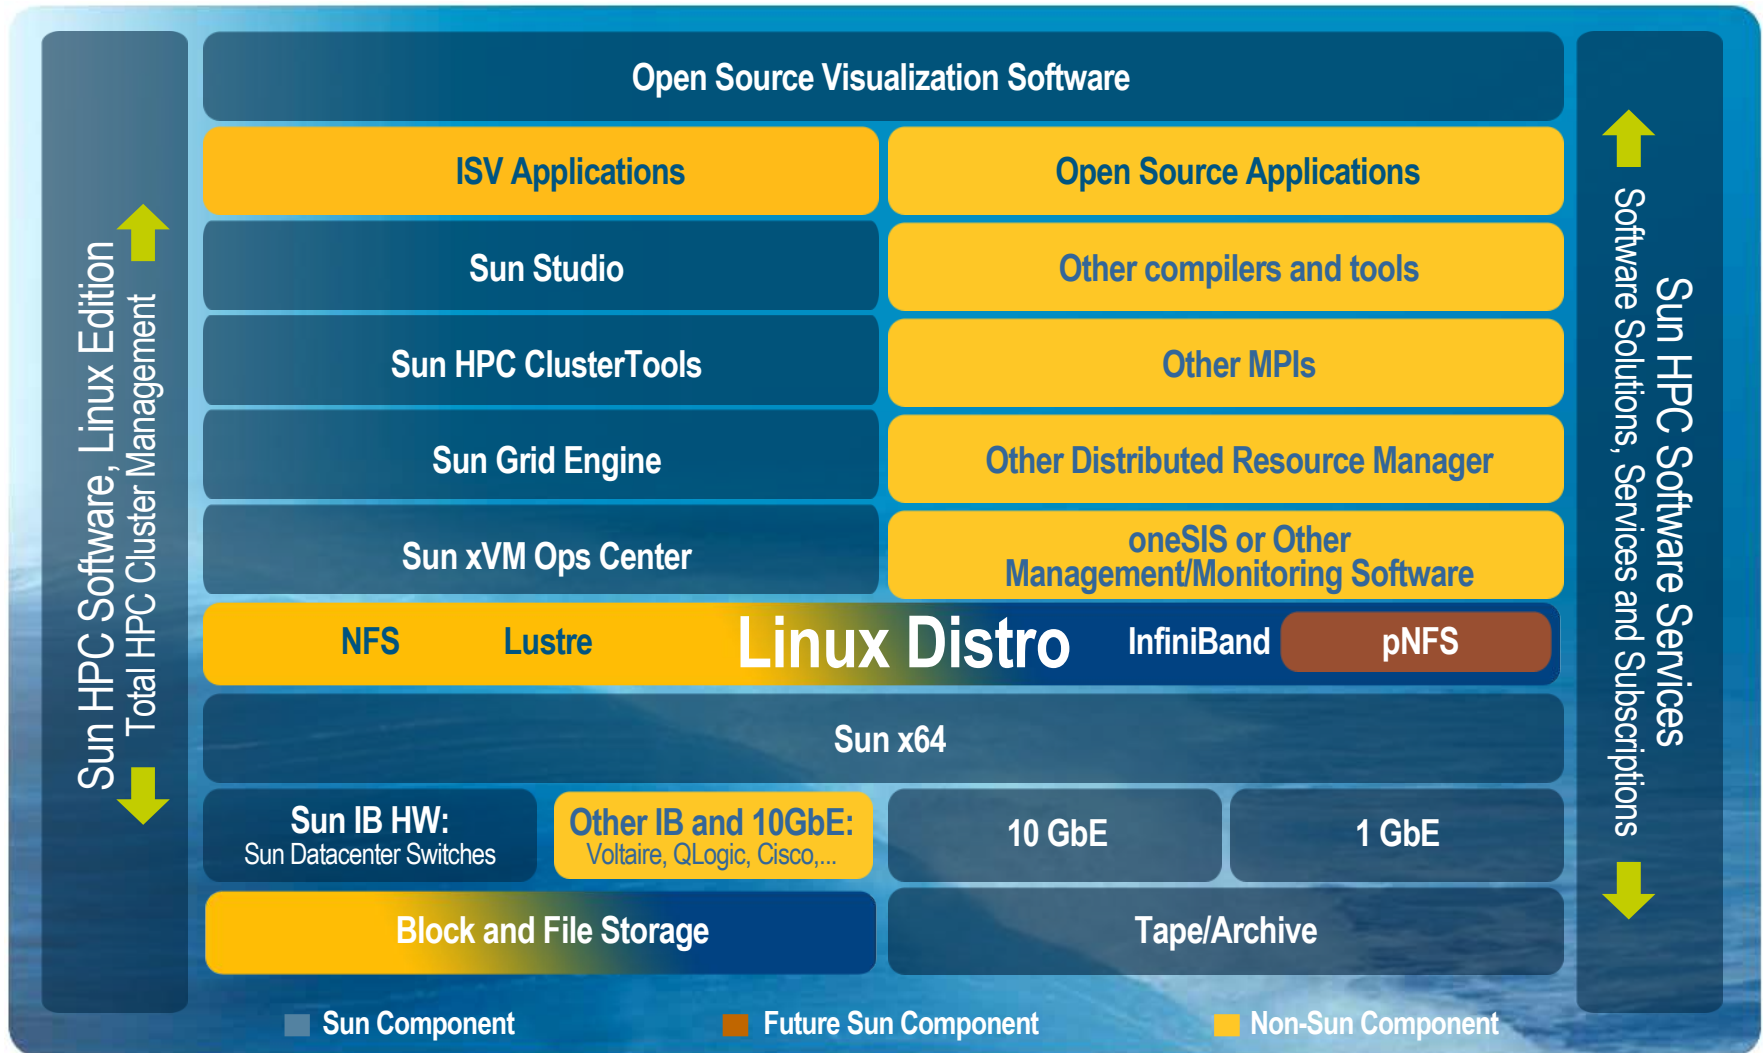
- 30 ports of 4x IB QDR connectivity realized with only 10 physical 12x connectors

# Sun Blade 6048 IB QDR Switched NEM

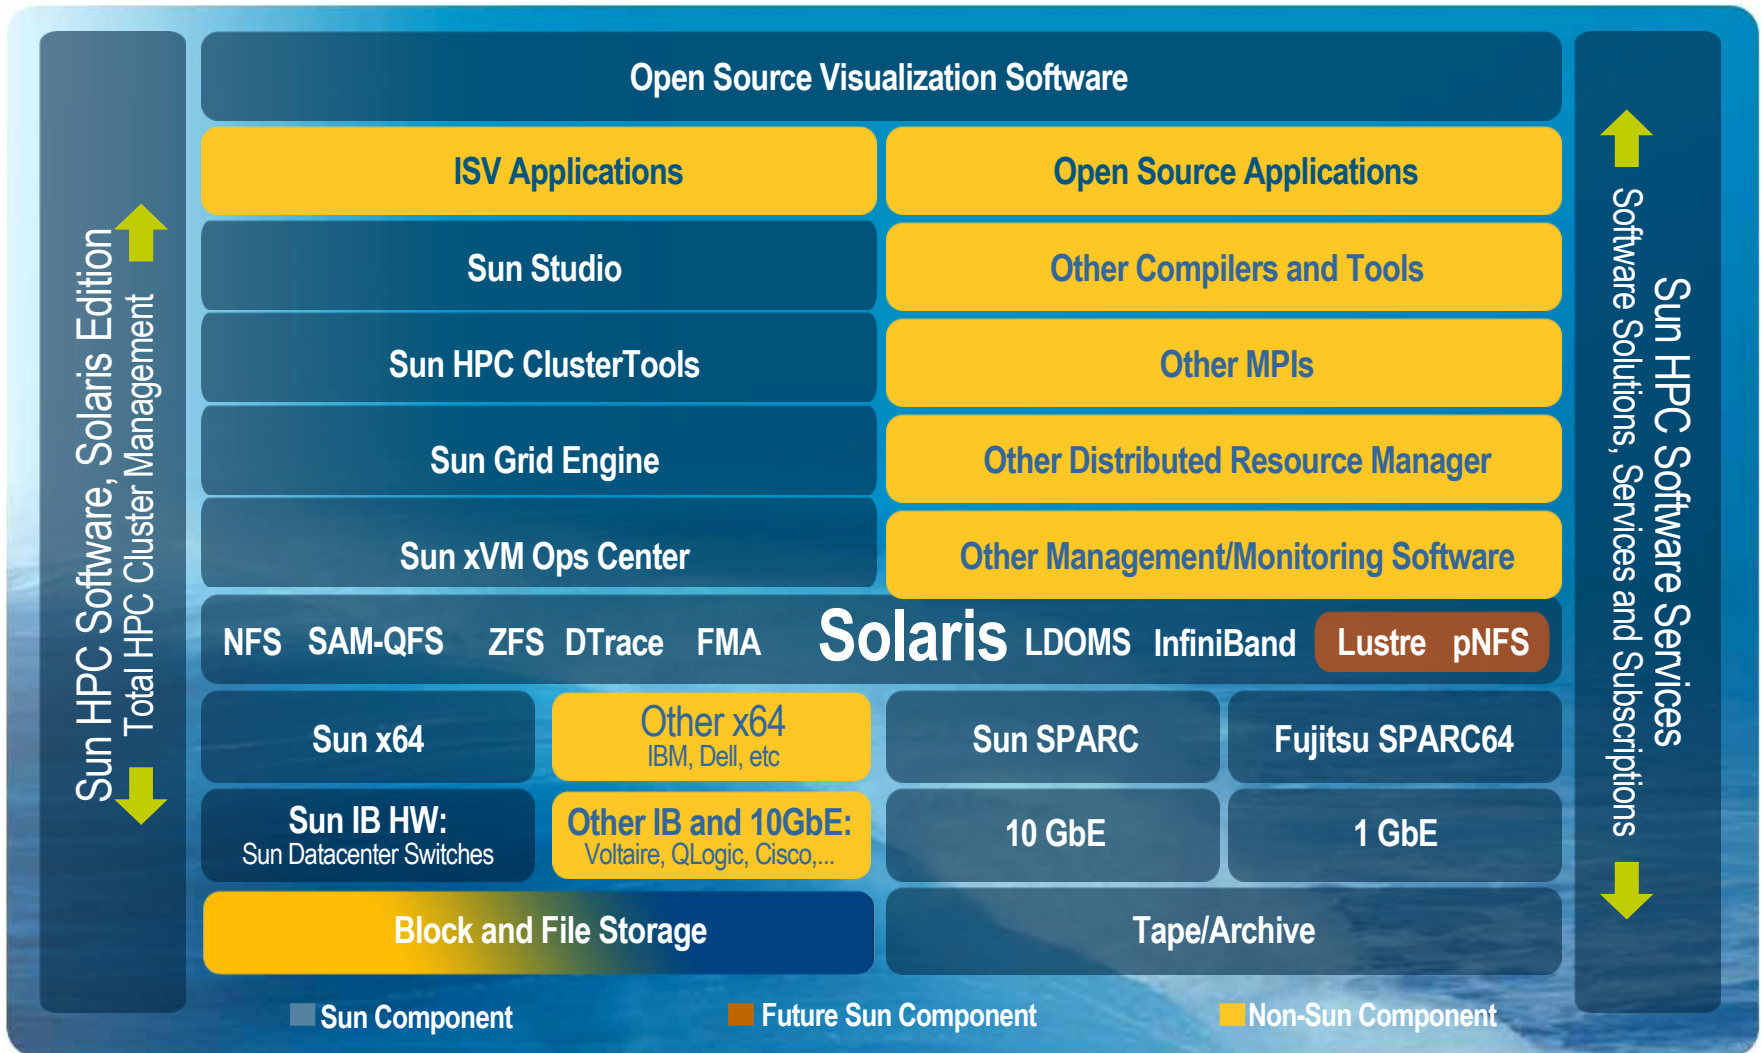
Mid-plane provides 4x QDR InfiniBand connection to 12 Server Modules



# Sun HPC Software, Linux Edition

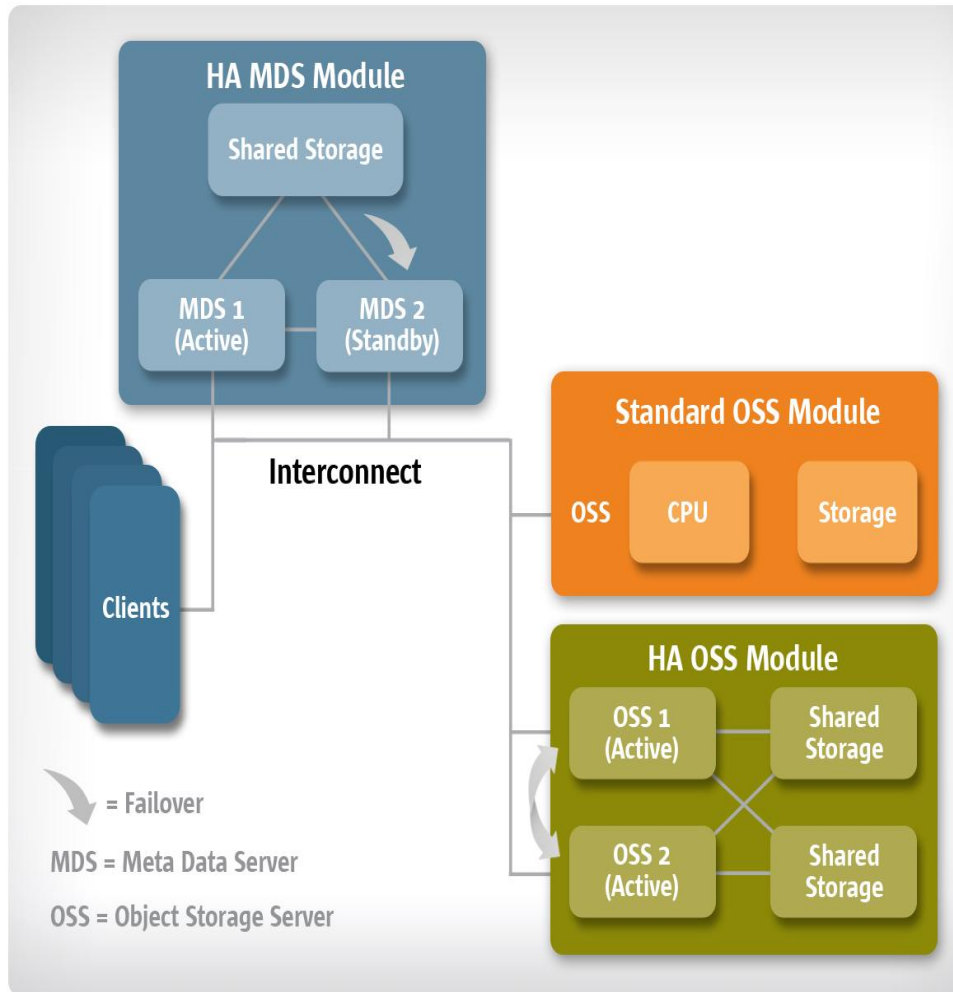


# Sun HPC Software – Solaris 10 and OpenSolaris



# Parallel Storage

## Sun Storage Cluster



Lustre and Open Storage<sup>®</sup> providing a wide range of Scaling and Performance for a wide range of cluster sizes

Breakthrough economics


Up to 50% cost savings vs. competitors

Simplified deployment and management of Lustre storage



## TACC Ranger

1.73 PB storage, 35GB/s I/O throughput,  
3,936 quad-core clients



## Sandia Red Storm

340 TB Storage, 50GB/s I/O throughput,  
25,000 clients



## Framestore

“The Tale of Despereaux”

200TB Lustre file system – averaged 1.2GB/s in sustained reads, peaks of 3GB/s, 5TB data generated per night from cluster of 4,000 cores interfacing with Lustre file system



## German Climate Research Data Centre (DKRZ)

272 TB, Lustre and Storage Archive  
Manager, 256 nodes with 1024 cores



# THANK YOU

**Mike Vildibill**  
mikev@sun.com