



HIGH PERFORMANCE COMPUTING FROM SUN

Paulo Vilela

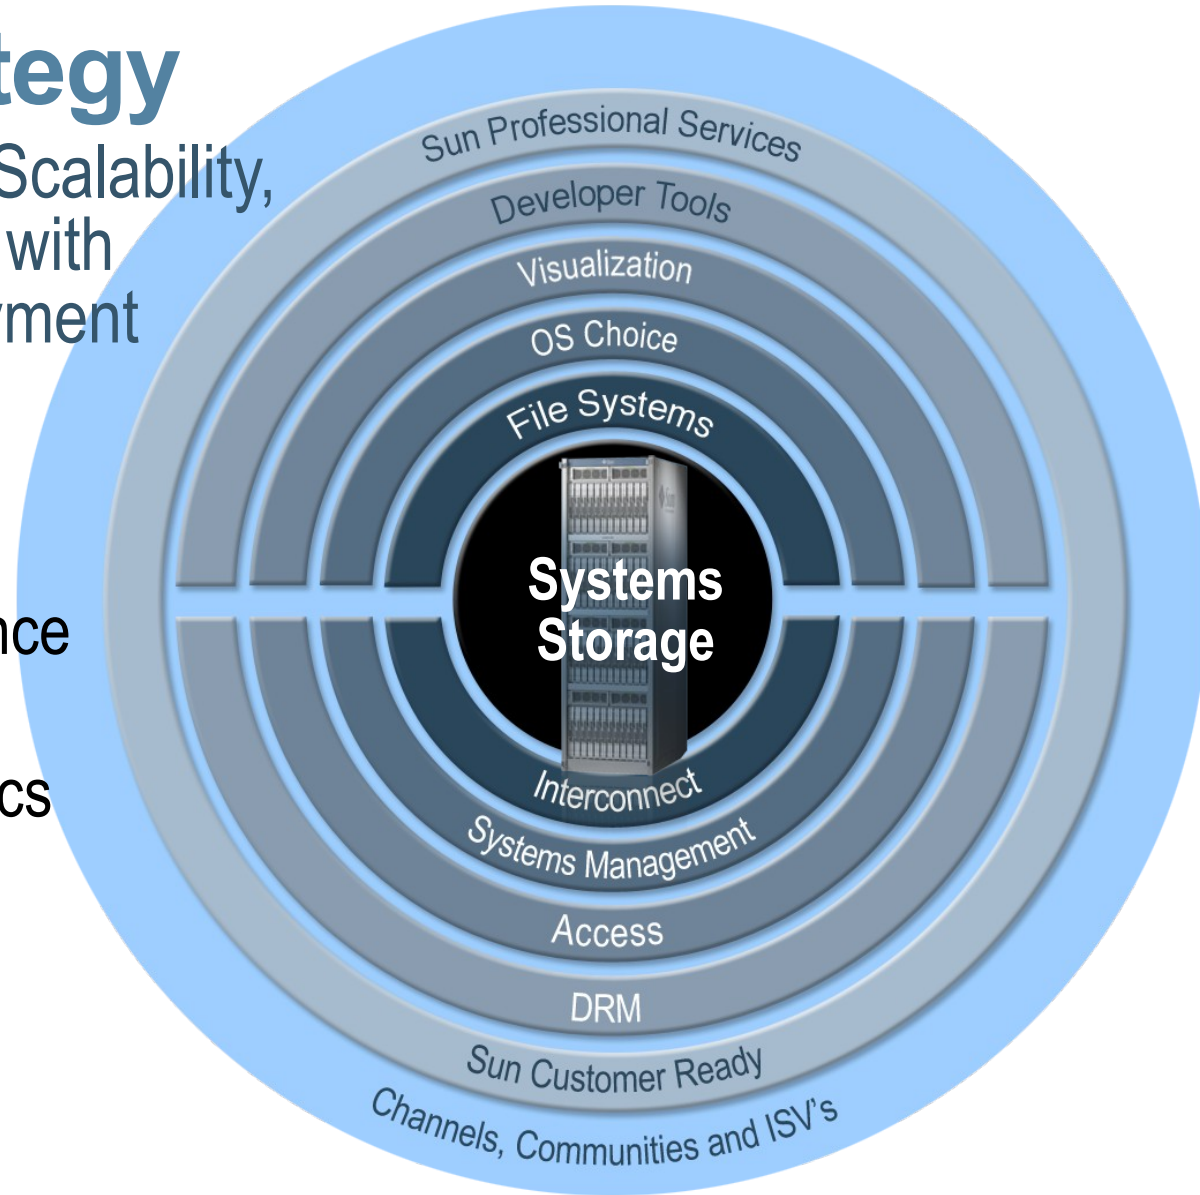
Sun Microsystems Portugal



Our HPC Strategy

Offering: Performance, Scalability, Capacity and Efficiency with Rapid, Low Risk Deployment

- Easy to deploy
- Un-matched performance and scalability
- New levels of economics
- Offers unprecedented choice and flexibility



TACC




Sun Fire X4500

- 72 Systems
- 1.7 petabytes
- 64.8 GB/sec total bandwidth



Sun Fire X4600

- 25 systems
- 800 cores



SunBlade 6048

- 3,936 blades
- 15,744K CPUs
- 62,976 cores
- 125 TB/RAM



Sun Data Center Switch 3456


- Dual redundant
- 110 TB/sec bisectional bandwidth

- The world's 4th largest compute cluster. Based on Sun Constellation System
- 326 Tflops peak performance
- Sun is the sole hardware supplier

Sun is Leading the Way with the Sun Constellation System

The World's Largest & Most Scalable General Purpose Computer

- New levels of performance and scalability
 - > Up to 2 PetaFLOPs
 - > Up to 1.8 PetaBytes RAM
 - > Up to 0.7 ExaBytes disk
- Open industry standards
 - > Linux, Solaris, OpenMPI, open InfiniBand interfaces and management
 - > X64 computing architecture
 - > InfiniBand DDR interconnect
- New levels of efficiency
 - > Provides a 6:1 reduction in physical ports and cables
 - > 20% smaller footprint than competition
 - > Eliminates 100s of discrete switching elements
- New levels of reliability
 - > Dramatically reduced complexity: 6x fewer cables, up to 50% fewer racks



Massive Scaling, Performance and Reduced Complexity

AMD, Intel or SPARC CPU's



Single
Blade

Sun
Blade 6000

Single Rack
768 Cores
7 TFLOPS
Peak

4 Core
Switches
13,824 Servers
2 PFLOPS
Peak

10 x Scaling

4.8 x Scaling

288 x Scaling

The Sun Cluster Portfolio

Open, Seamless and Comprehensive

Access	Developer	Management	OS	Inter-connect	Storage/Archive	Systems
Visualization, Workstation, Thin Clients, Remote Access	Compilers, Debuggers, Optimization Tools, Libraries	Workload, Systems and Cluster Management	Linux, Solaris	InfiniBand or Ethernet	Cluster Storage, Backup, Archive, File Systems, HSM	Racks or Blades, Variety of CPU Architectures



Sun Services

Storage



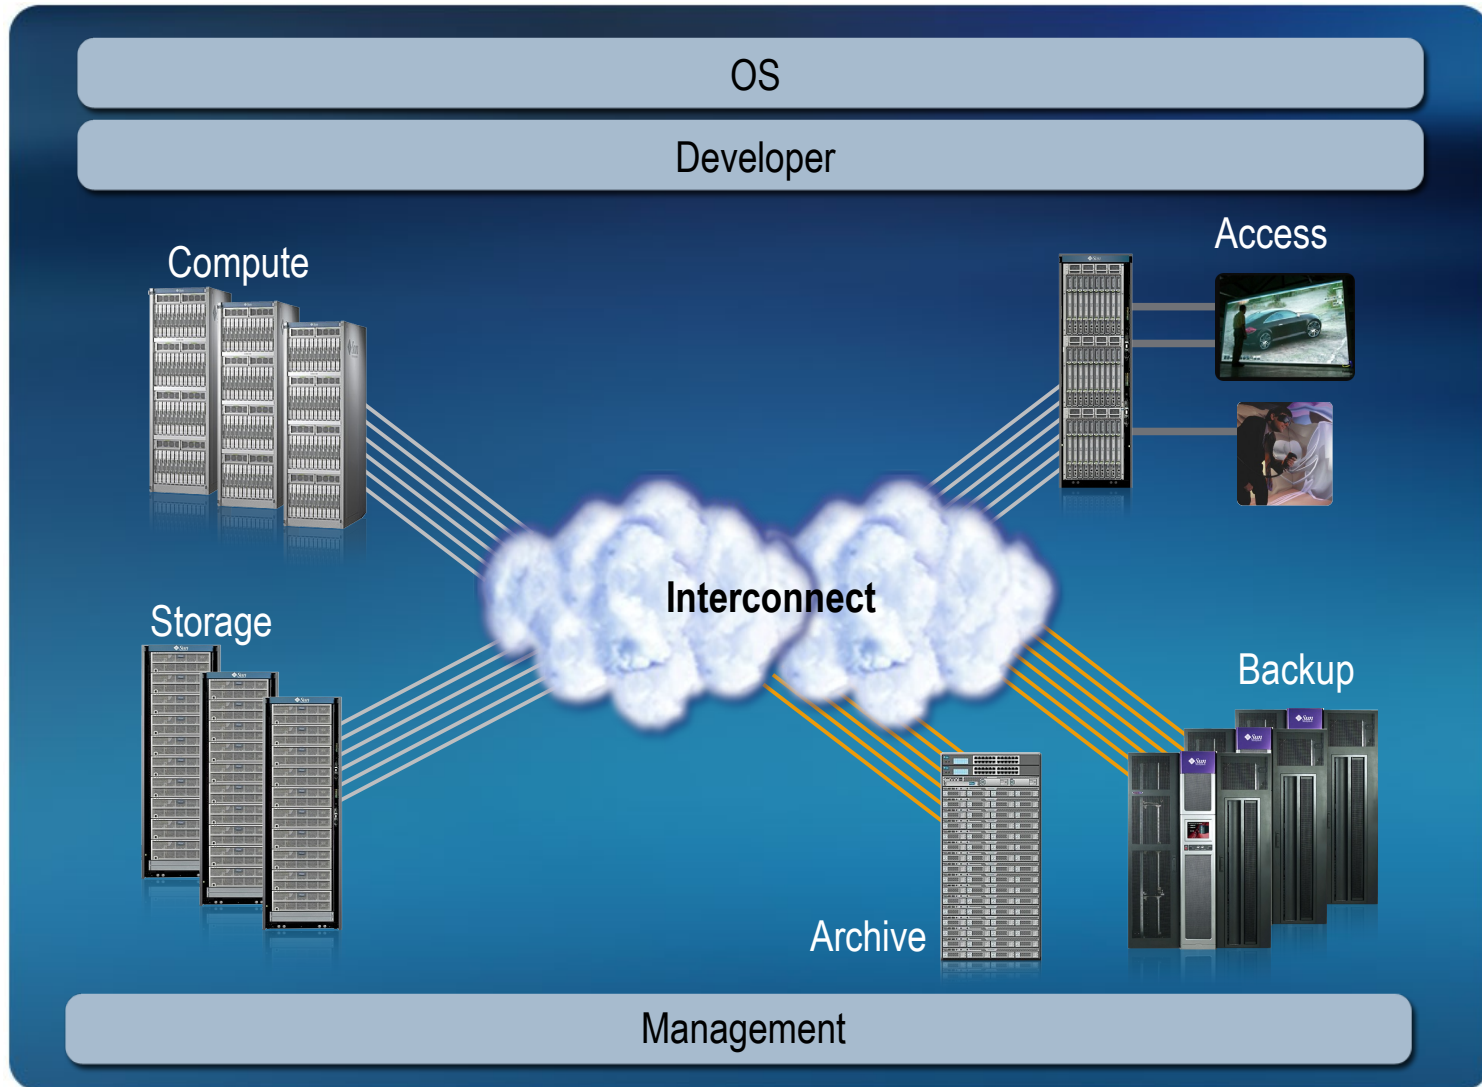
Networking

Compute

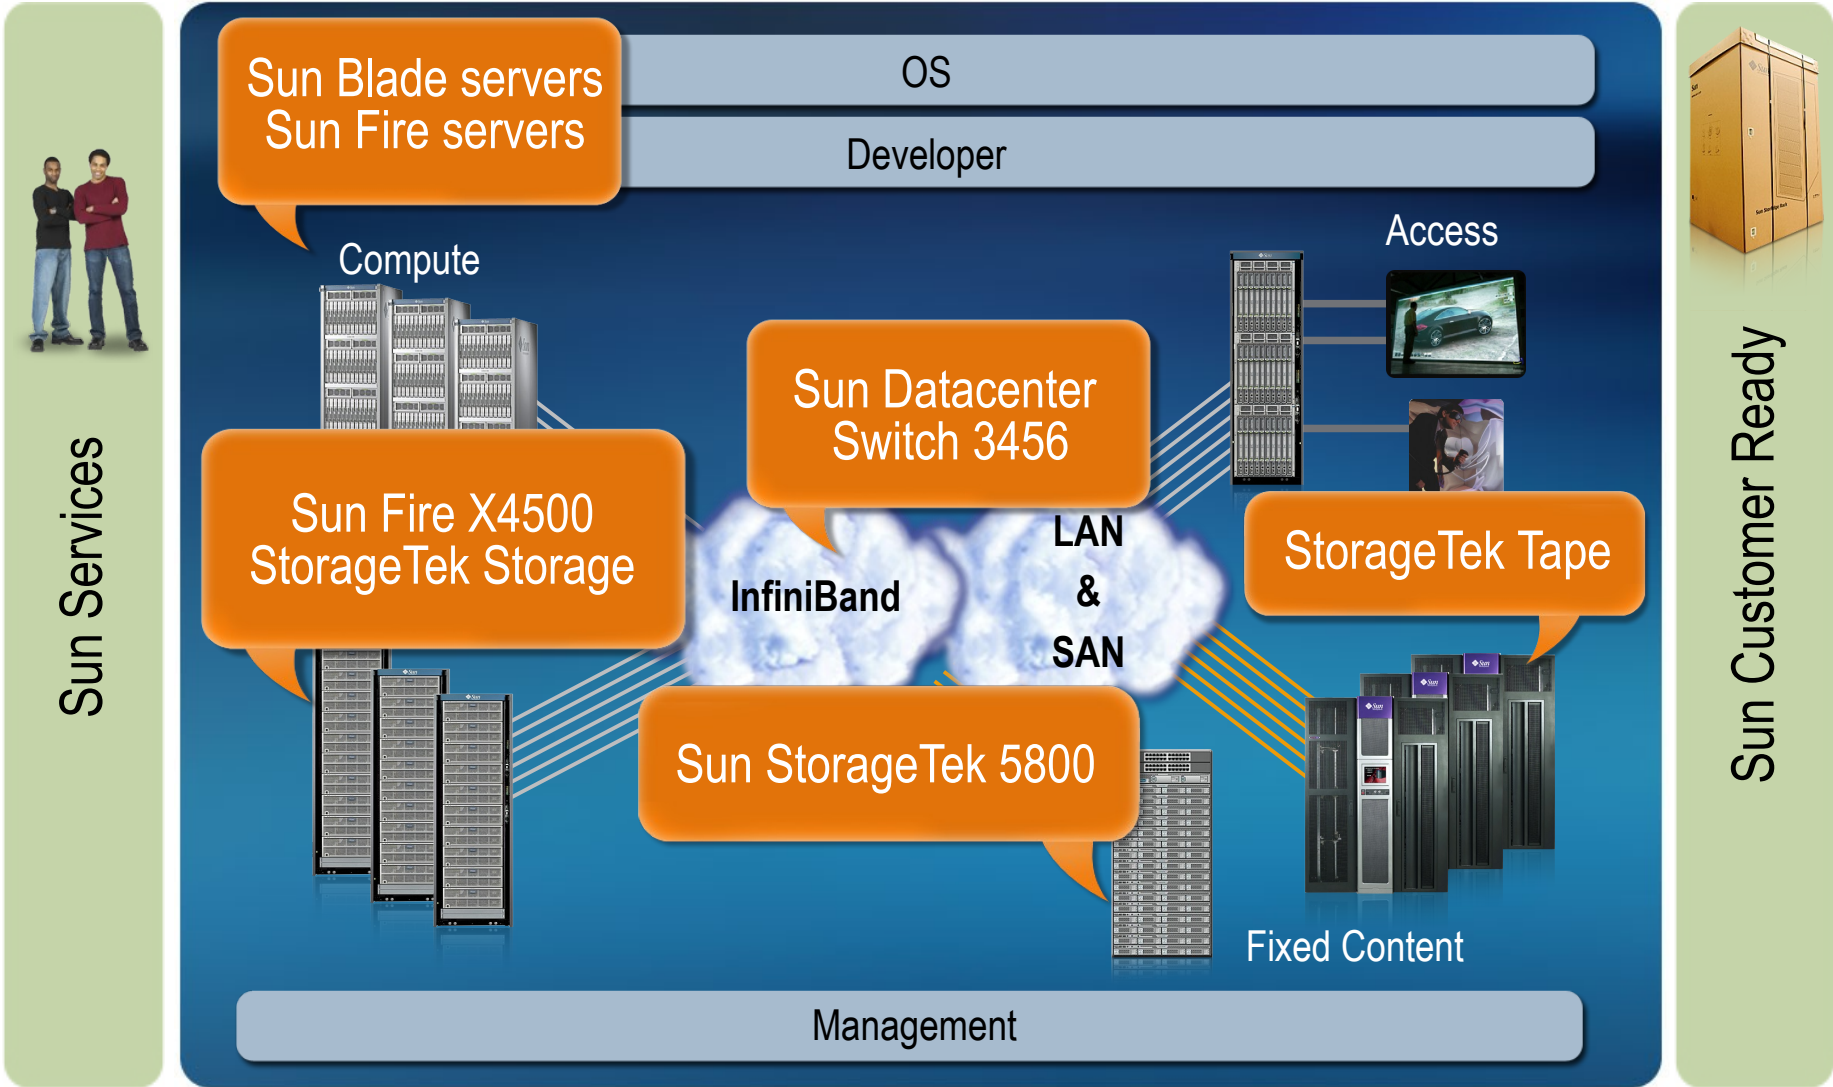
Sun Customer Ready



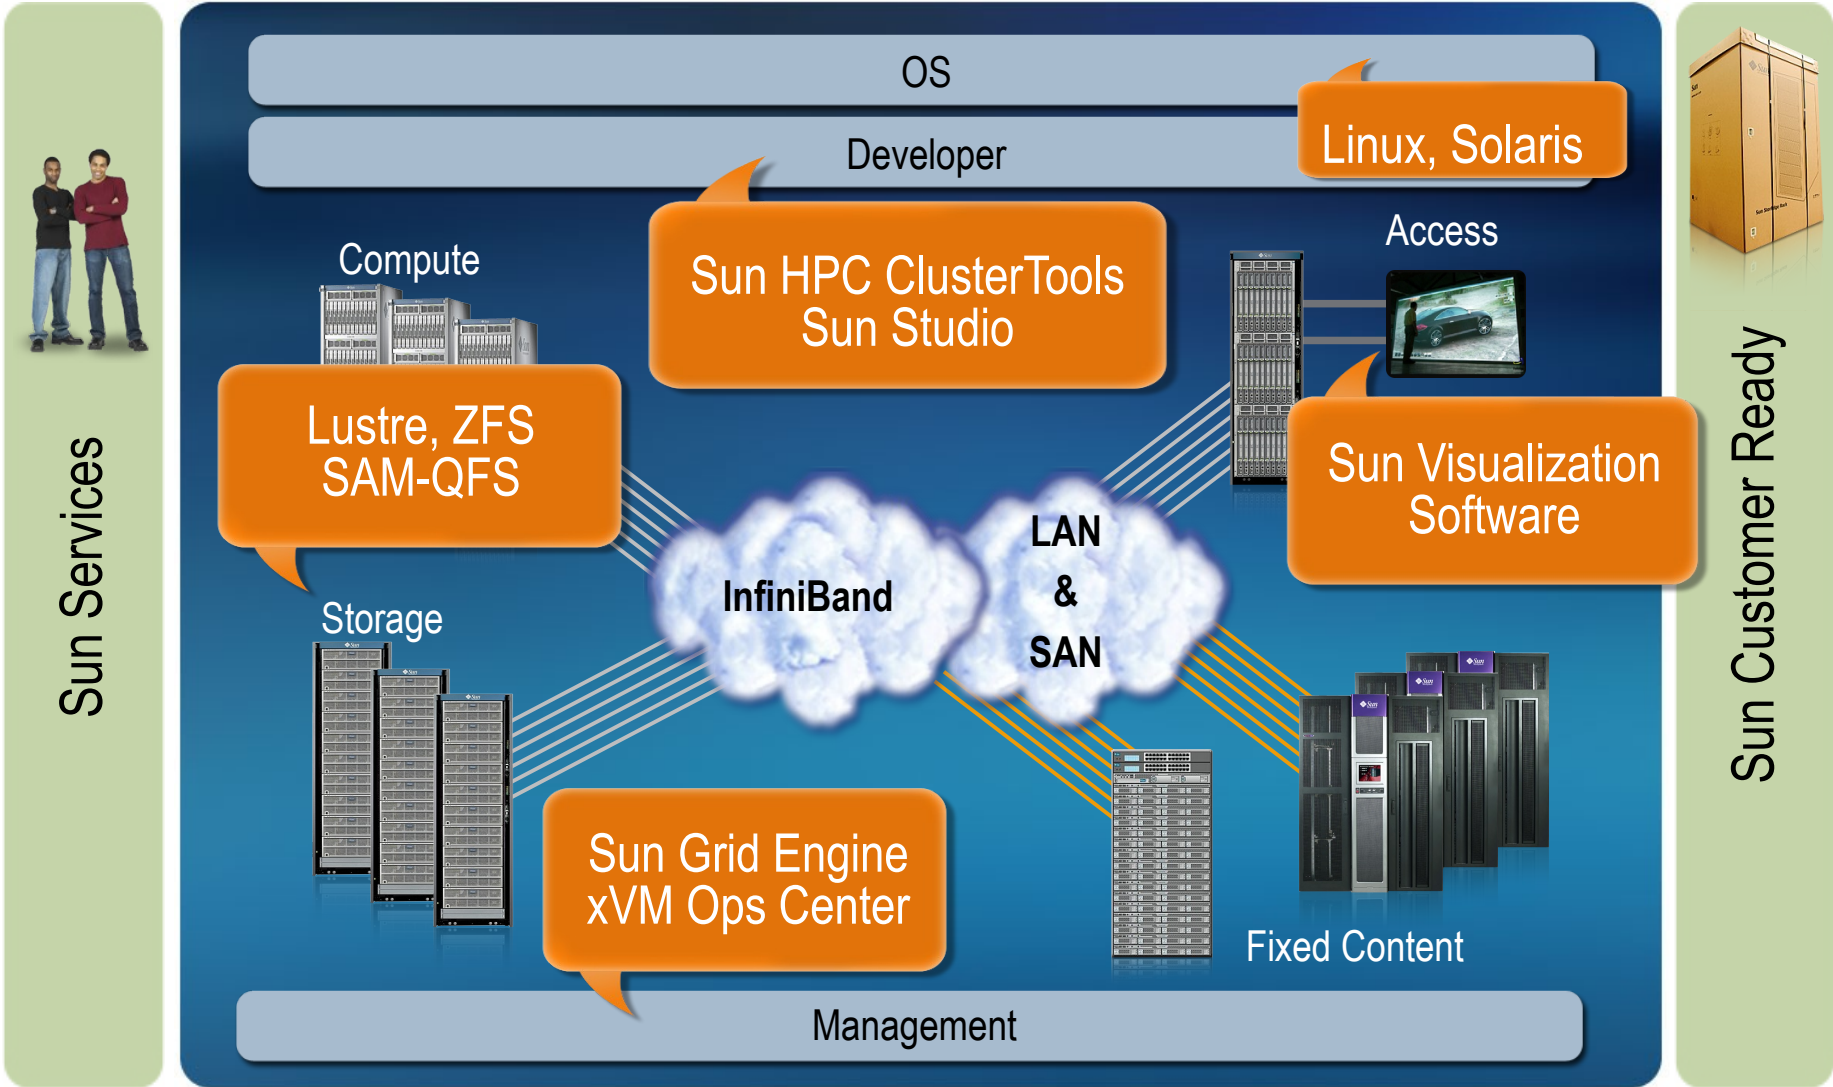
The Sun HPC Cluster Platform



The Sun HPC Cluster Platform




The Sun HPC Cluster Platform



Sun Constellation System Open Petascale Architecture


Radical Simplicity, Faster time to Deployment



Alternative Open Standards Fabric

- 300 switching elements
- 6912 cables
- 92 racks

Sun Constellation System Cluster



Reduced Cabling

Constellation System Open Super Computer

- 1 switching element **300:1 reduction**
- 1152 cables **6:1 reduction**
- 74 racks **20% smaller footprint**

HPC Clusters Arrive Ready to Run

Your Choice of
Sun and 3rd Party
Hardware and Software



Hardware Racked
and Cabled



Software Installed
and Configured



Sun Customer Ready Program Deliver “Ready-to-Deploy” HPC Systems

Agile Development


- Up to 90% faster
- Tools to tailor to specific needs
- You concentrate on your core business

Higher Quality, Lower Risk

- Up to 80% less installation issues
- Integration and testing in the factory
- Uses Sun’s ISO9000 certified manufacturing processes

Sun Blade 6000 Modular System

10 Server Module Capacity



Compute

- 10 Server Modules

I/O

- 1.42Tb/s I/O per Chassis
- 20 Ho1t-Plug PCIe ExpressModules^(TM) for granular blade I/O configuration
- 2 PCIe Network Express Modules^(TM) for multi blade I/O configuration

Availability

- Hot Swap fans
- Hot Swap power supplies
- Redundant power grid connection capability

Management

- SNMP, SSH, CLI

Form Factor


- 10 Rack Units
- 40 Opteron Cores
- 80 Intel Cores
- 80 SPARC cores



Chassis – Rear View




- 20 x PCI Express ExpressModules
 - > Two per Server Module
- 2 x PCI Express Network Express Modules (NEM)
- 1 x Chassis Monitoring Module
- Cooling Fans
 - > 6 x Rear Fan Modules
 - > Hot-Swap
 - > Redundant (N+1)
- 4 x power inlets with cable holders
 - > Two inlets per power supply
 - > Cable holders prevent accidental loss of power



Chassis – Air Flow

- Air flow in the chassis is front to rear
- Two separate air flows
 - > One is powered by the Front Fan Modules within the Power Supplies
 - > Second is powered by the Rear Fan Modules which cools the Server Modules



Sun Blade 6048 Modular System



- The first blade platform designed for extreme density and performance
 - > 7 TFLOPS, 768 cores per chassis/42U
 - 50% more compute power than HP C-Class
 - 71% more compute power than IBM BladeCenterH
 - > 4 InfiniBand Leaf Switch Network Express Modules
 - Lowest cost per port with ultra-dense switch solution
- Pay as you grow platform ideal for fast growing businesses
 - > Choose among SPARC, AMD Opteron and Intel Xeon CPU technologies
- Runs general purpose software
 - > Custom compiles and tuning are not required
- Realize economies of scale savings in power and cooling

Massive Horizontal Scale

Sun Blade 6048 Chassis



Compute

- 4 x Shelves per chassis
- 48 x Sun Blade 6000 series server modules
 - Sun UltraSPARC
 - Intel Xeon
 - AMD Opteron

I/O

- 96x Hot-Plug PCIe ExpressModules
- 8x Hot-Plug PCIe Network Express Modules
 - 12 1xGbit Ethernet connections per NEM

Availability

- Hot swap redundant fans
 - Eight Rear Fan modules per shelf
 - Two Front Fan Modules per shelf
- Hot swap power supply modules
 - Two 9000 W power supplies per shelf
 - Total of 8 power supplies per chassis






Management

- 4 x Chassis Monitoring Modules
- Solaris, Linux, or Windows OS, VMware

Density

- **48 Servers**
- **192 Sockets / 768 cores**

Versatile Blade Portfolio



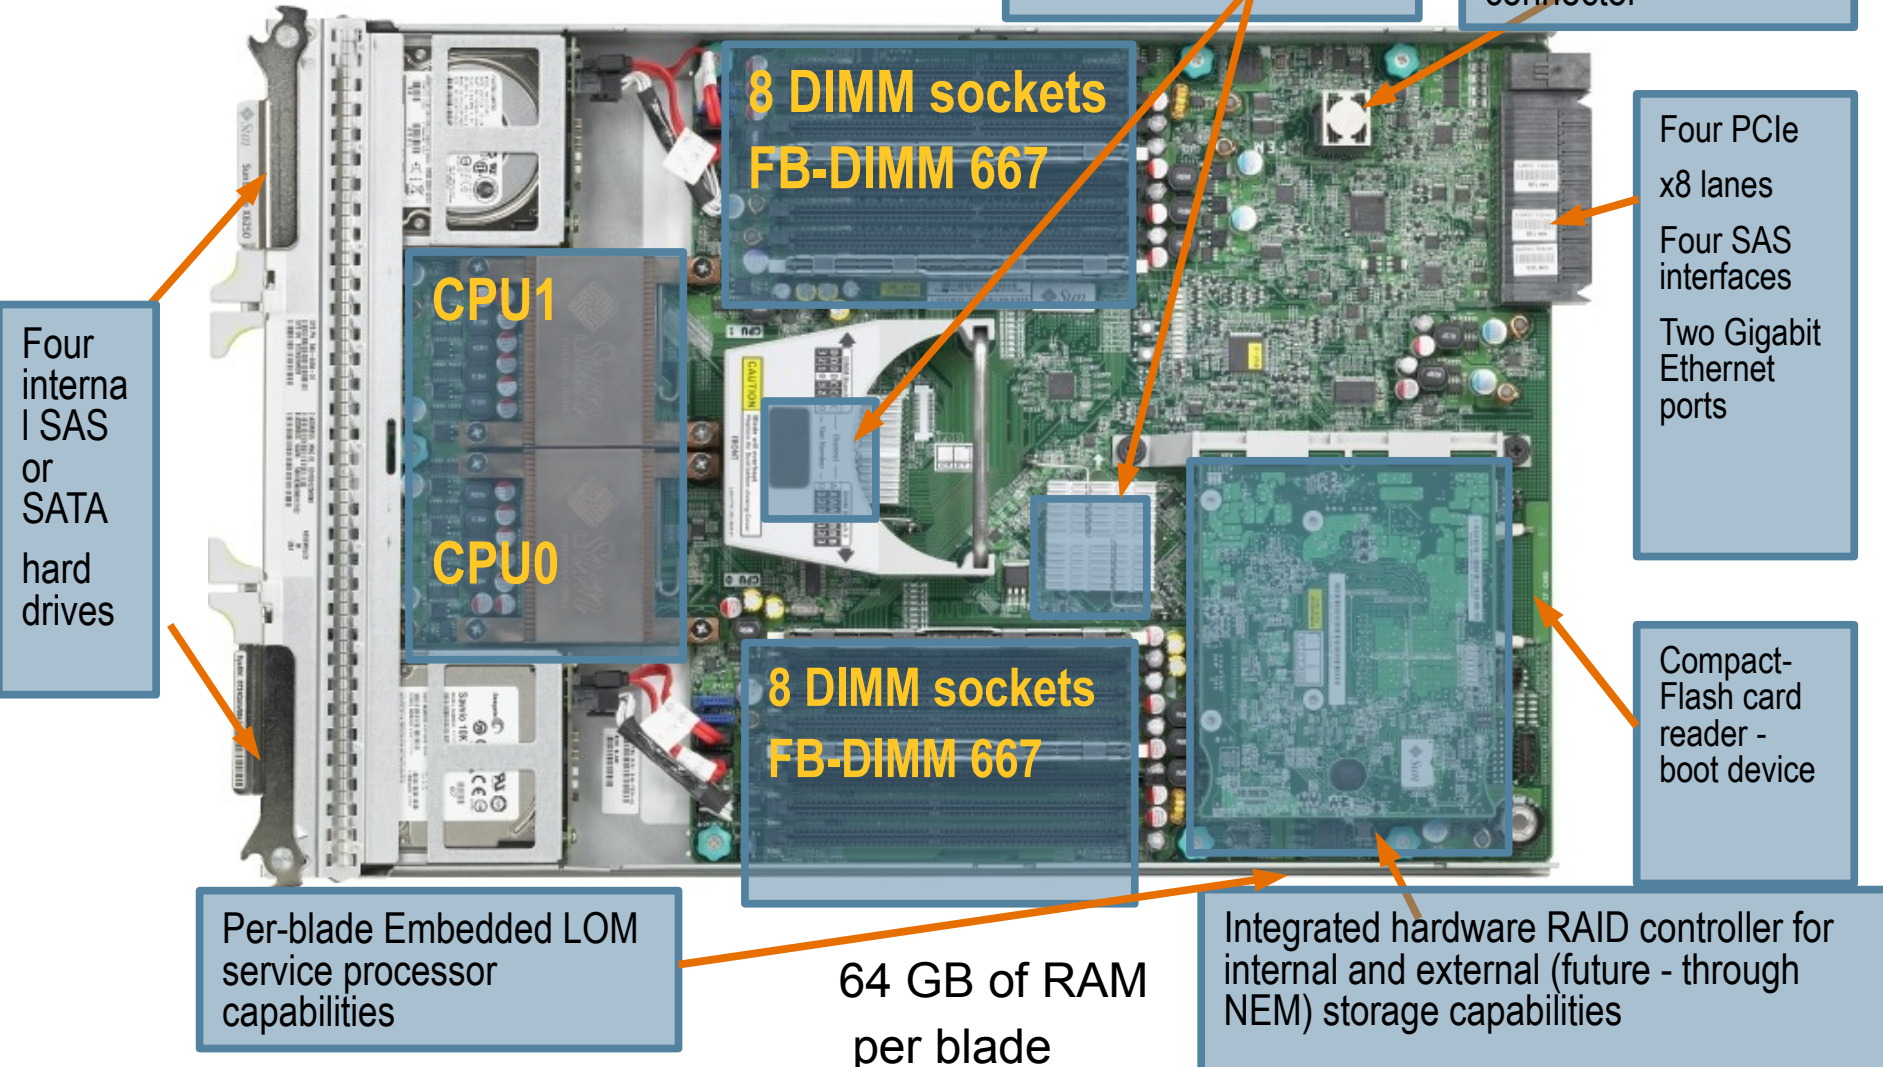
Sun Blade T6300

Sun Blade T6320

Sun Blade X6220

Sun Blade X6250
Sun Blade X6450

Sun Blade X6250 Server Module Device Map



Sun Blade X6450 Server Module

24 FB-DIMM slots. Up to 96 GB of RAM

FEM connector

Four PCIe interfaces, four SAS links and two 1GbEth ports

Four Intel 7000 processors. Dual or quad-core.


Intel 7000 MCH (Clarksboro) Northbridge

Fault remind button

Compact Flash Storage


REM connector

ESB2 I/O bridge Battery location for the REM



Sun Blade 6048 InfiniBand switched NEM

Double width NEM – takes up both slots of a shelf



- Twelve dual port Host Channel Adapters (HCA)
 - > Two 4x InfiniBand uplinks per server module
- Two 24 port InfiniBand switches
 - > 12 x4 ports used for the HCA connection
 - > 4 x12 ports used for the uplink

Sun Customer Ready HPC Cluster


Clusters Made Simpler

- Sun™ Customer Ready HPC Cluster
 - > Designed specifically for science, engineering and compute intensive applications
 - > Choice of: Sun servers/blades based on AMD or Intel processors, Gigabit Ethernet or InfiniBand, the Solaris 10 OS, or Linux. Optional Sun xVM and Sun Grid Engine software
 - > Web-based configurator tool available

- Sun™ Customer Ready HPC Cluster for Crash Analysis
 - > Ideal for computer-aided engineering crash analysis applications
 - > Sun Fire rackmount servers, the Solaris 10 OS
 - > Reference configurations as starting points



Sun Datacenter Switch 3456



- Unique in the industry
- 12x Capacity of largest IB switch available today
- 6x Reduction in Cables, Space and Weight
- 300:1 Switch Reduction

The World's First Petascale Single Chassis Switch

Sun Visualization System

Scalable, Sharable and Secure



- Virtualizing graphics resources
- Massive scale in compute density and Display resolution
- Share across multiple users with different devices
- Data is secured on central server

Scalable Clustered Storage

Sun Fire X4500 Server



- Industry's first data server
- Best server data throughput and storage density
- Runs Lustre parallel file system
- Standard platform and common systems management capabilities



Sun Modular Data Center

The Virtualized Datacenter

High Capacity

- 820 CPU's, 3,280 Cores
 - > Sun Blade 6000 systems and Sun Fire X2200 servers
- 2,240 cores and 17,920 compute threads!
 - > UltraSPARC T2 servers
- 3 petabytes of disk
 - > Sun Fire X4500 servers



High Performance

- 31 Teraflops Peak
 - > 210 Sun Blade 6000 nodes
- ~170,000 web ops/sec
 - > 78 x Sun Fire T2000 servers

Current Deployments

Stanford Linear Accelerator
Center (SLAC) High
Performance Computing Node



Sun HPC Grid - Menlo
Park, CA



Sun HPC Resources

On-Demand Computing

Visit: network.com
 access compute
 cycles and
 applications over
 the network



Developers

Visit:
[http://opensolaris.org/os/
 community/mmunity/hpcdev](http://opensolaris.org/os/community/mmunity/hpcdev)
 Join the opensolaris HPC
 community. Download the
 latest opensolaris HPC distro



Radio HPC

Subscribe via iTunes
 and get regular
 updates on HPC
 technology from Sun
 and our partners



Website

Visit: sun.com/hpc
 Read the latest news,
 view the latest offers,
 download the latest
 white papers and more



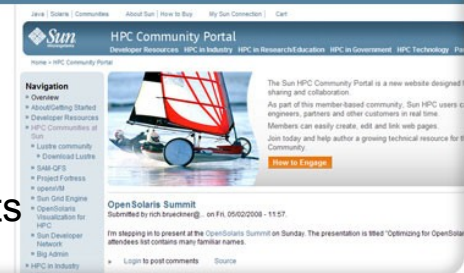
Watercooler

Visit the HPC
 Watercooler at:
blogs.sun.com/hpc and
 get the latest HPC news
 from around the globe



HPC Community

Join the online
 HPC community at:
hpc.sun.com and
 collaborate with Sun
 engineers and experts





HIGH PERFORMANCE COMPUTING FROM SUN.

Paulo Vilela

