HIGH PERFORMANCE COMPUTING FROM SUN

Paolo Sestini
HPC Senior Architect - Solaris Ambassador
System Practice
Sun Microsystems Italia S.p.A.

The New Era of HPC

Innovations in HPC Meet the Commercial World

Today HPC Customers Want More:

- **Radical Efficiency** – Performance, power, cooling, space, cost
- **Super Scalability** – Paving the way to Petaflops
- **Open Systems** – Open interfaces, industry standard components and community
- **Industrial Robustness** – High availability and reliability
- **Production-Ready** – Time-to-results or time-to-production

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

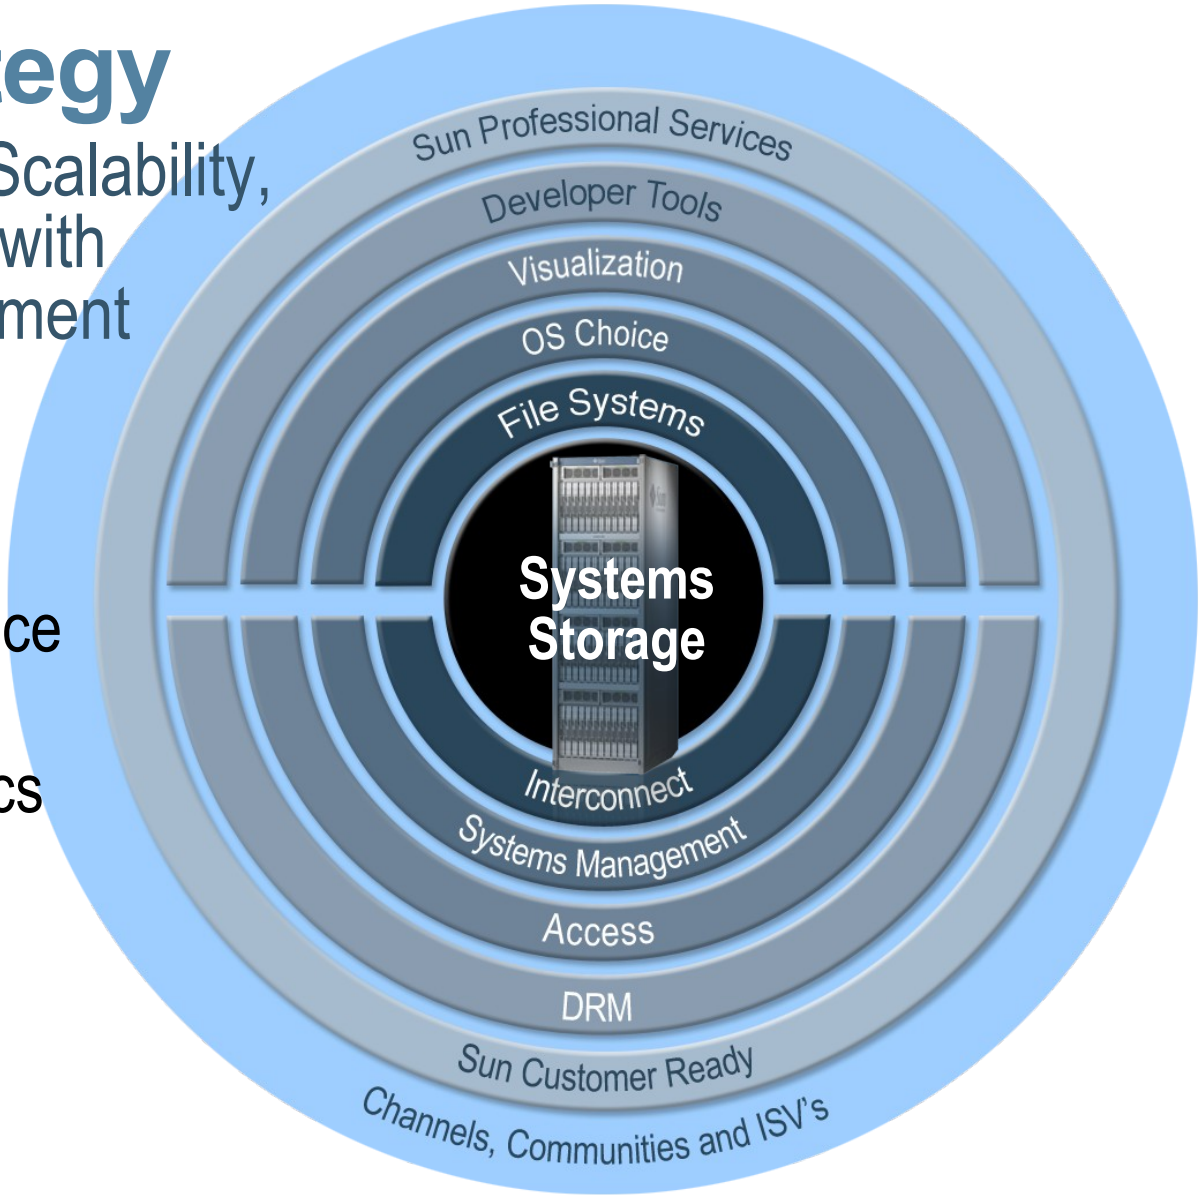

Sun Constellation System Open Petascale Architecture

Eco-Efficient Building Blocks

Compute

Ultra-Dense Blade Platform

- Fastest processors: SPARC, AMD Opteron, Intel Xeon
- Highest compute density
- Fastest host channel adaptor

Networking

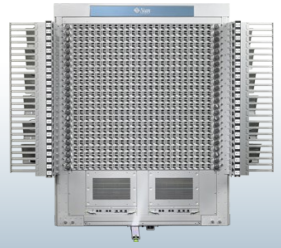

Ultra-Dense Switch Solution

- 3456 port InfiniBand switch
- Unrivaled cable simplification
- Most economical InfiniBand cost/port

Storage

Ultra-Dense Storage Solution

- Most economical and scalable parallel file system building block
- Up to 48 TB in 4RU
- Direct cabling to IB switch

Software

Developer Tools | Grid Engine | Provisioning

solaris™ | Linux

Comprehensive Software Stack

- Integrated developer tools
- Integrated Grid Engine infrastructure
- Provisioning, monitoring, patching
- Simplified inventory management

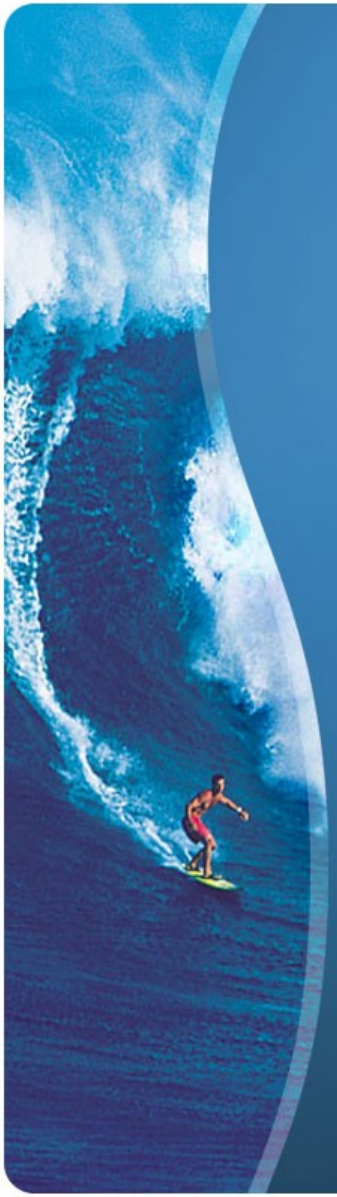

Key Innovations

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 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

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

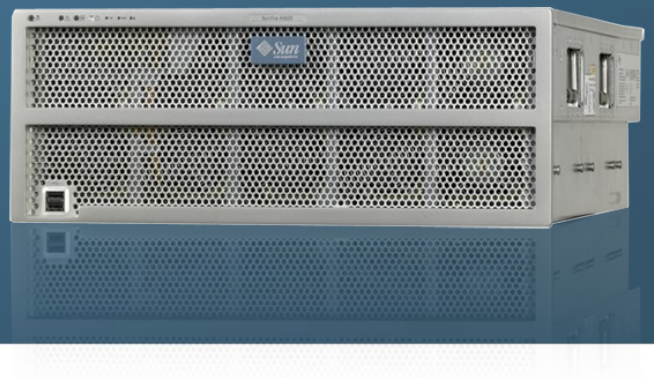

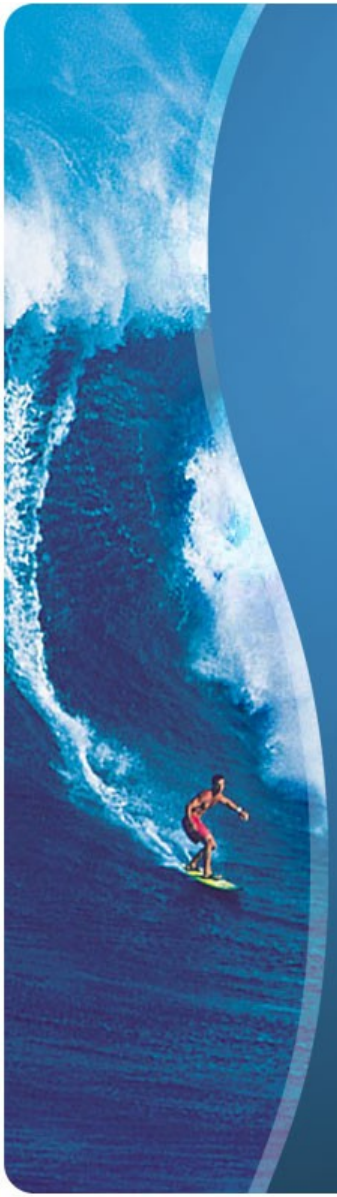

Software

Choice and Flexibility

Run the Environment You Need

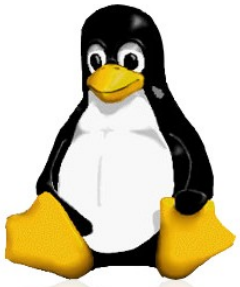

Red Hat and SuSE Enterprise Linux

Available for All Sun HPC Clusters

solaris™

Leading Enterprise Operating System

Available for All Sun HPC Clusters

Microsoft® Windows®

Windows OEM Certification and Fully Supported on All Sun x86 Systems

Lustre™ File System

World's Largest Network-Neutral Data Storage and Retrieval System

- The worlds most scalable parallel filesystem
- 10,000's of clients
- 50% of Top 30 run Lustre
- 15% of Top 500 run Lustre
- Open Source, multi-platform

Sun xVM Ops Center

Ops Center Provides an Amazingly Easy-to-Use Interface to Manage Rapid Growth in Your IT Environment

DISCOVER	Scan and identify servers across your network, even when powered off
PROVISION	Hands-off installation of Linux and Solaris onto both bare metal servers and virtual environments
UPDATE	Stay secure and up-to-date with patch management tools for RedHat, SUSE and Solaris
MANAGE	Securely manage users and assets with the best knowledge about your Sun gear
REPORT	Assure compliance with the industry's first compliance auditing solution

Developer Software

- Scalable Parallel Debuggers
  - Totalview, Totalview Technologies
  - DDT, Allinea Software
- Cache-use Performance Analysis
  - Virtual Performance Expert, Acumem
- **Sun Studio 12** application development
  - **C, C++, Fortran development:** Compilers support industry and defacto standards to enable portable, maintainable, extensible code
  - **High performance:** Optimized for target systems: UltraSPARC, X86 and x64
  - **Multi-core development:** Powerful debugger, auto-parallelizer, advanced perf analysis tools
  - **Full featured graphical development environment**
  - **SunPerf lib:** Optimized mathematical libraries
  - Supports a variety of **MPI choices**
- **Sun HPC ClusterTools 8** for parallel program development based on OpenMPI

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

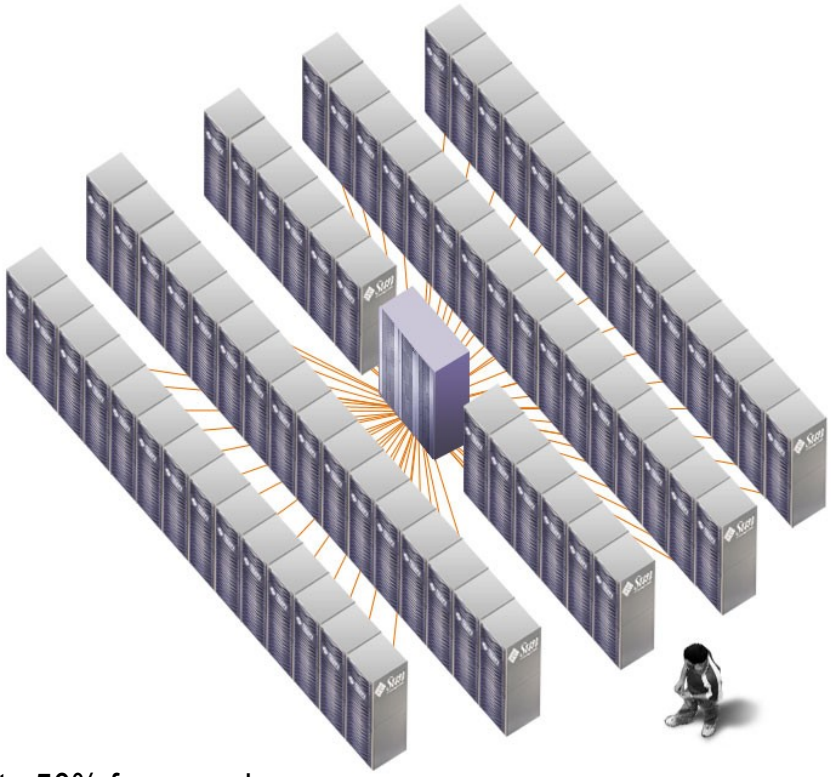

Flexible Configurations Scale to Meet Computing Needs

1 Core Switch

Servers: 3,456
PFLOPS: 0.5

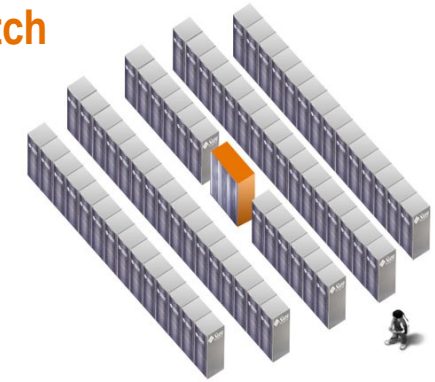

2 Core Switches

Servers: 6,912
PFLOPS: 1.0

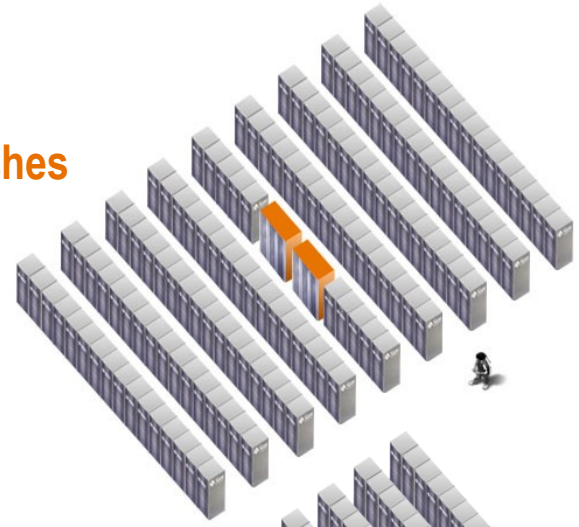

3 Core Switches

Servers: 10,368
PFLOPS: 1.5

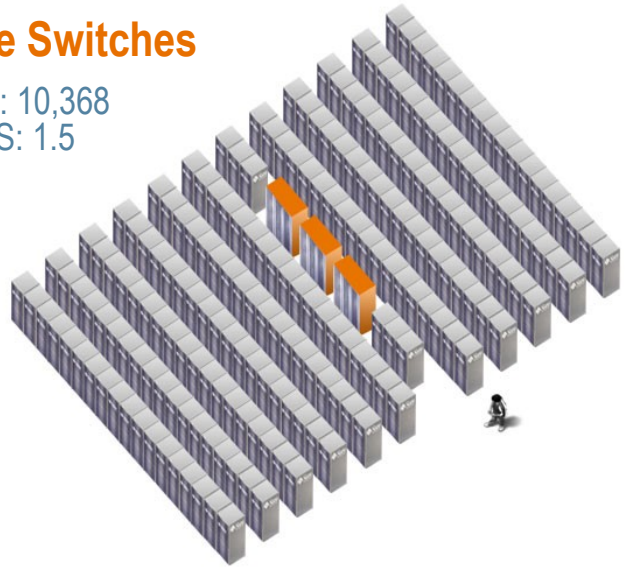

4 Core Switches

Servers: 13,824
PFLOPS: 2.0

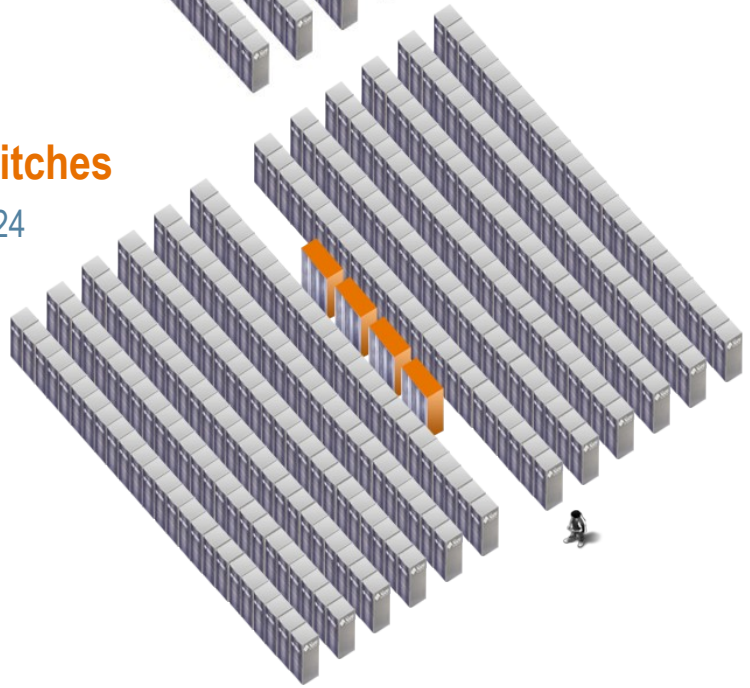

TACC: World's Top Supercomputer

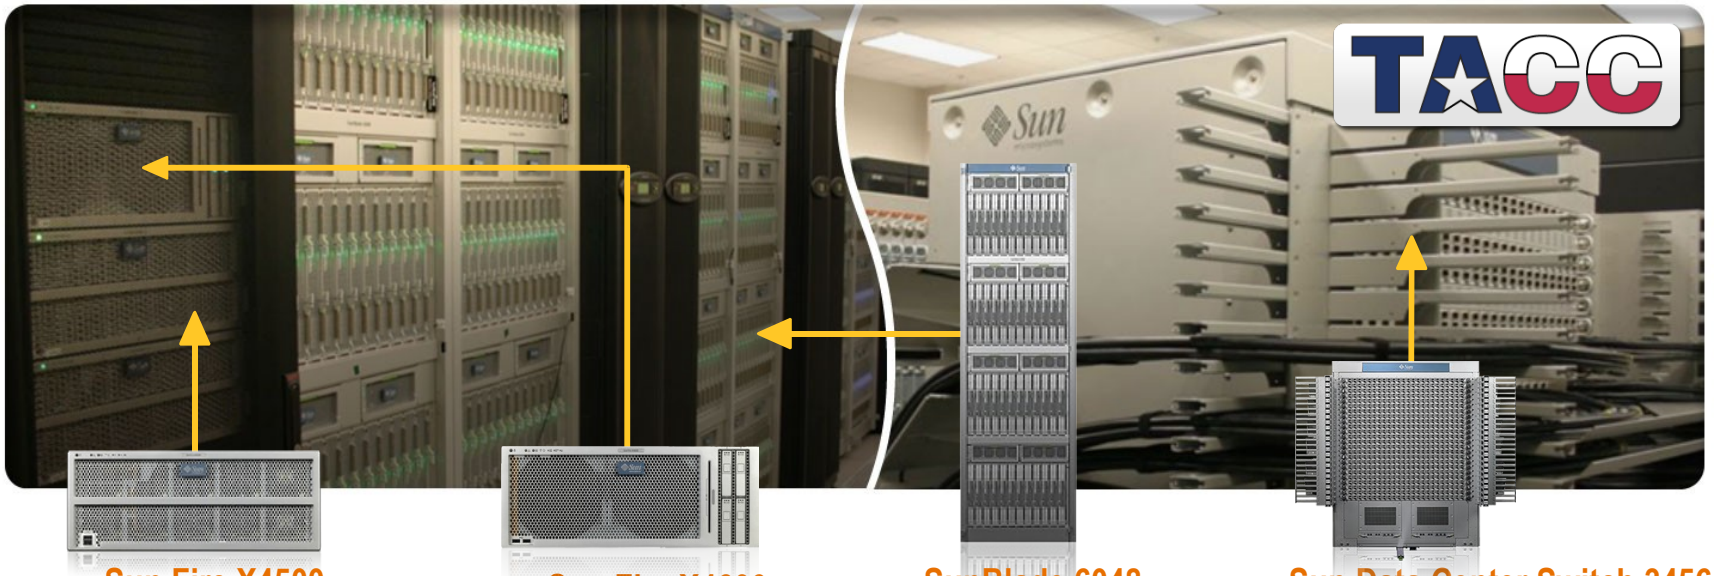

Sun Fire X4500
- 72 Systems
- 1.7 petabytes
- 64.8 GB/sec total bandwidth

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 largest general purpose compute cluster. Based on Sun Constellation System
- 504 Tflops peak performance
- Sun is the sole hardware supplier
- In production since 4th February, 2008 on the Terragrid enabling new levels of science

TACC

TEXAS ADVANCED COMPUTING CENTER

THE UNIVERSITY OF TEXAS AT AUSTIN

HPC Systems

☆Sun Linux Cluster ☆Dell Linux Cluster ☆IBM Power5 System ☆TACC Stampede Cluster ☆ Operations

Sun Constellation Linux Cluster

System Name:	Ranger
Host Name:	ranger.tacc.utexas.edu
IP Address:	129.114.50.163
Operating System:	Linux
Number of Nodes:	3,936
Number of Processing Cores:	62,976
Total Memory:	123TB
Peak Performance:	504TFlops
Total Disk:	1.73PB (shared) 31.4TB (local)

Description:

"Ranger" is the largest computing system in the world for open science research. As the first of the new NSF Track2 HPC acquisitions, this system provides unprecedented computational capabilities to the national research community and ushers in the petascale science era. Ranger will enable breakthrough science that has never before been possible, and will provide groundbreaking opportunities in computational science & technology research – from parallel algorithms to fault tolerance, from scalable visualization to nextgeneration programming languages.

Ranger went into production on February 4, 2008 using Linux (based on a CentOS distribution). The system components are connected via a full-CLOS InfiniBand interconnect. Eighty-two compute racks house the quad-socket compute infrastructure, with additional racks housing login, I/O, and general management hardware. Compute nodes are provisioned using local storage. Global, high-speed file systems will be provided, using the Lustre file system, running across 72 I/O servers. Users will interact

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
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

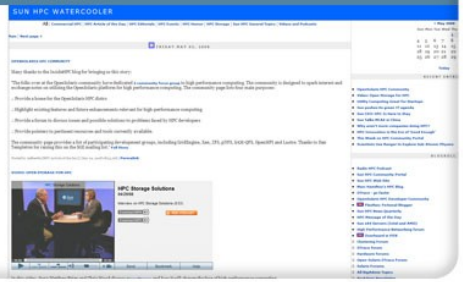

HPC Community

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

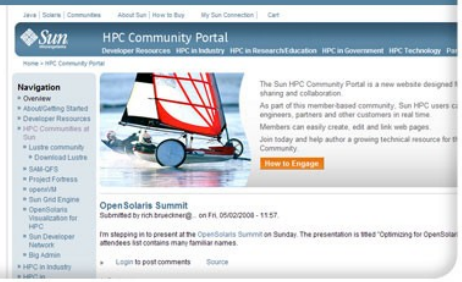

HIGH PERFORMANCE COMPUTING FROM SUN.

Paolo Sestini
paolo.sestini@sun.com