



Open Source Efforts in Grid Computing and HPC

Bob Porras

Vice President,

Solaris Data, Availability, Scalability & HPC

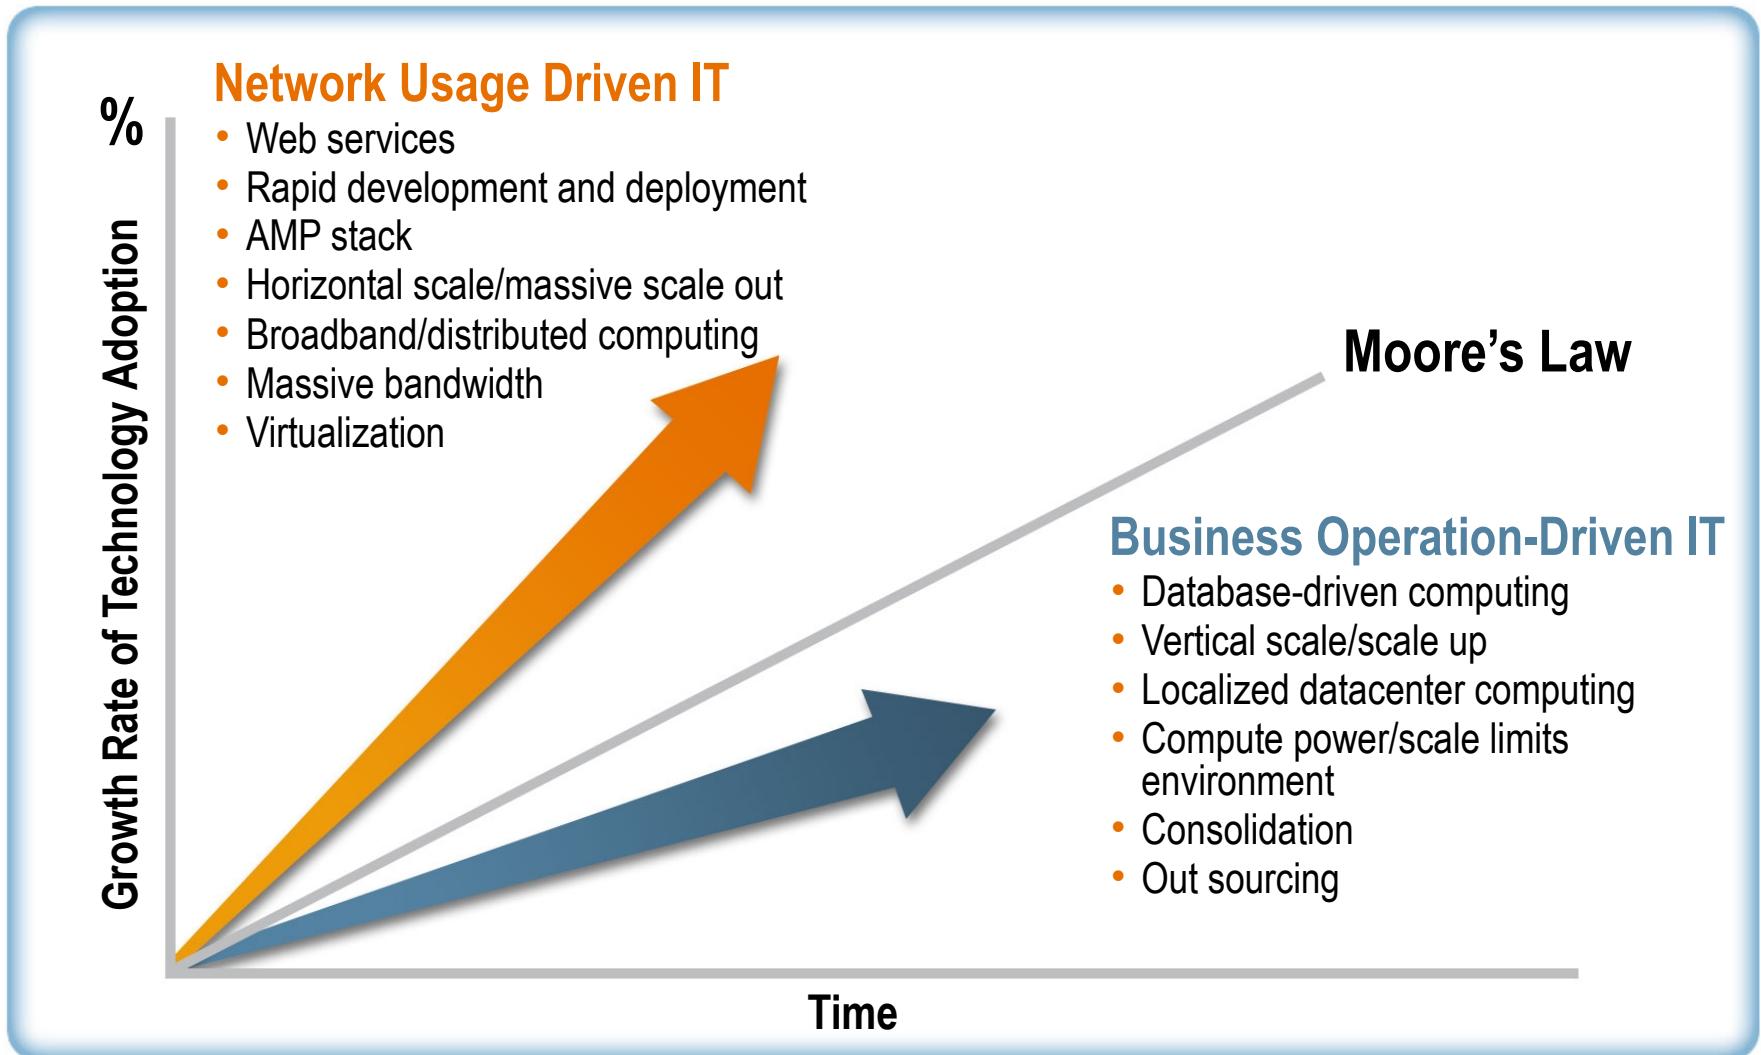
Sun Microsystems, Inc.

blogs.sun.com/bobp

Keynote Topics

- The **RedShift** Opportunity
- Open Source @ Sun
 - > OpenSolaris
 - > Open HPC and Grid Software Stacks
 - > Open Development Platforms and Tools
 - > OpenStorage
 - > Open Virtualization
 - > Open Java Platform
- Call for Participation

Grid & HPC – The *RedShift* Evolution



Sun's Vision

The Journey: Open Innovations to Create Breakthrough Economics

Open Servers

- Leveraging innovative product design and packaging
- Common components
- Open-source software
- Wide interoperability to deliver breakthrough economics

Open Storage

- Open software
- An open architecture
- Common components
- Open interoperability to create innovative storage products
- Delivers breakthrough economics

Open Networks

- Unified datacenter network that utilizes common components
- Open-source software
- Seamless integration with existing environments
- Delivers breakthrough economics

Sun's Open Stack

Flexible and Heterogeneous with Zero Barrier to Exit

Developer Environment		
Database/ Storage Platform		
Application Infrastructure		
Virtualization	<p>xVM, Ops Center, Grid Engine</p>	
Operating System		
Systems Servers Storage Networking		
Microprocessor		

Sun's Contribution to Linux

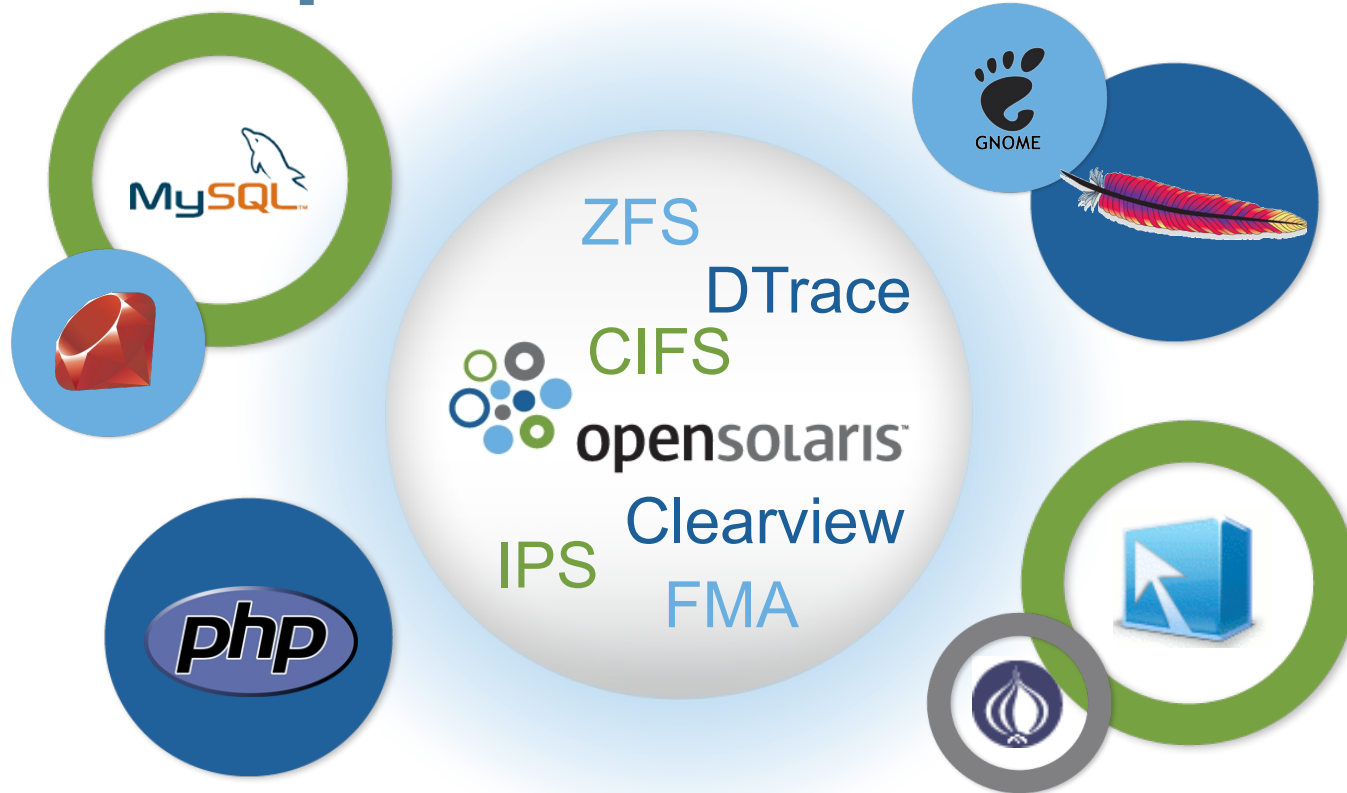
Rank	Company	Estimated \$ value
1	Sun Microsystems Inc	404 m
2	IBM Corp	116 m
3	Red Hat Corp	76 m
4	Silicon Graphics Corp	61 m
5	SAP AG	60 m
6	MySQL AB	45 m
7	Netscape Communications Corp	41 m
8	Ximian Inc	39 m
9	RealNetworks Inc	35 m
10	AT&T	34 m

Estimated Substitution Cost of Sun's contribution to Debian GNU/Linux

includes code in
GNOME
Linux kernel
Mozilla
OpenOffice.org
X.org
and other projects

Source: UNU-MERIT report for the European Commission
 "Economic impact of FLOSS on innovation and competitiveness of the EU ICT sector"
 January 2007

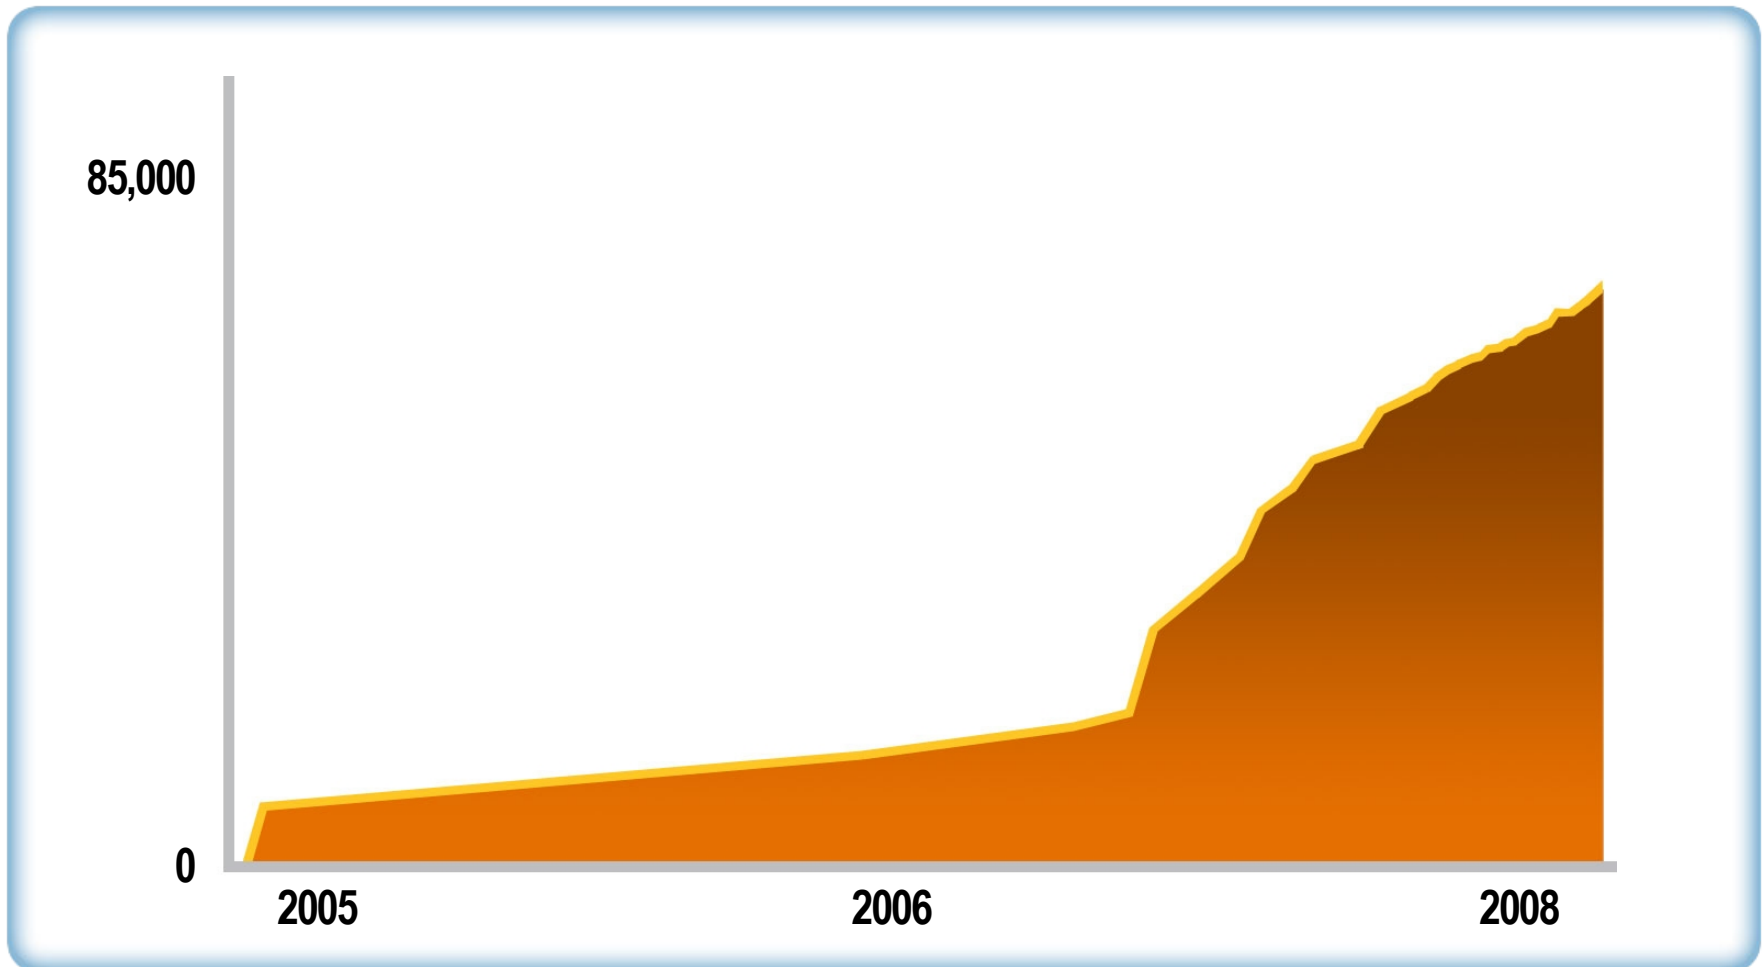
What is OpenSolaris?



**Community Participation + Solaris =
Innovation**

OpenSolaris Community Growth


Registered Community Members



Cumulative Growth

OpenSolaris 2008.05

- First release of the OpenSolaris OS
 - > First OS with ZFS as default file system
 - > Enhanced DTrace with D-Light
 - > Fast in kernel CIFS server
- New OpenSolaris package repositories
 - > Integrated & optimized AMP/MARS stack
- Fully supported for production deployments



OpenSolaris Repositories

Community

Sun

Subscription

Latest Releases

3rd Party Open Source Packages

Community Open Source Packages

Current Release Updates

Free Sun and 3rd Party Binaries

Developer Showcase

Partner Marketplace

Bug Fixes

Security Updates

Back Ports

Early Access



World Class Support

OpenSolaris Essentials

OpenSolaris Production

Telephone, Online Tech Support

Email only 48-hour Response

24/7 P1 Live Transfer

OpenSolaris full release and major updates	✓	✓
Knowledge base access	✓	✓
Sun alerts and notifications	✓	✓
Installation assistance	✓	✓
Sun Spectrum eLearning Library	✓	✓
Administration support*	✓	✓
Solaris fix and errata patch updates		✓

OpenSolaris Features

OpenSolaris 6 month releases	✓	✓
Free Sun and OpenSolaris Community IPS server	✓	✓
Sun OpenSolaris Premium IPS Server	✓	✓
OpenSolaris Bug Logging	✓	✓
OpenSolaris Bug Escalation		✓
Open Stack and Postgres Support		*promo inclusion


OpenSolaris. Start It Up!

Get It



- Available everywhere
- Smaller faster download

Experience It



- Live CD
- See the best in OpenSolaris, risk free

Install It



- Easy
- Graphical
- Supported on ~1000 systems
- Runs in popular virtualization environments

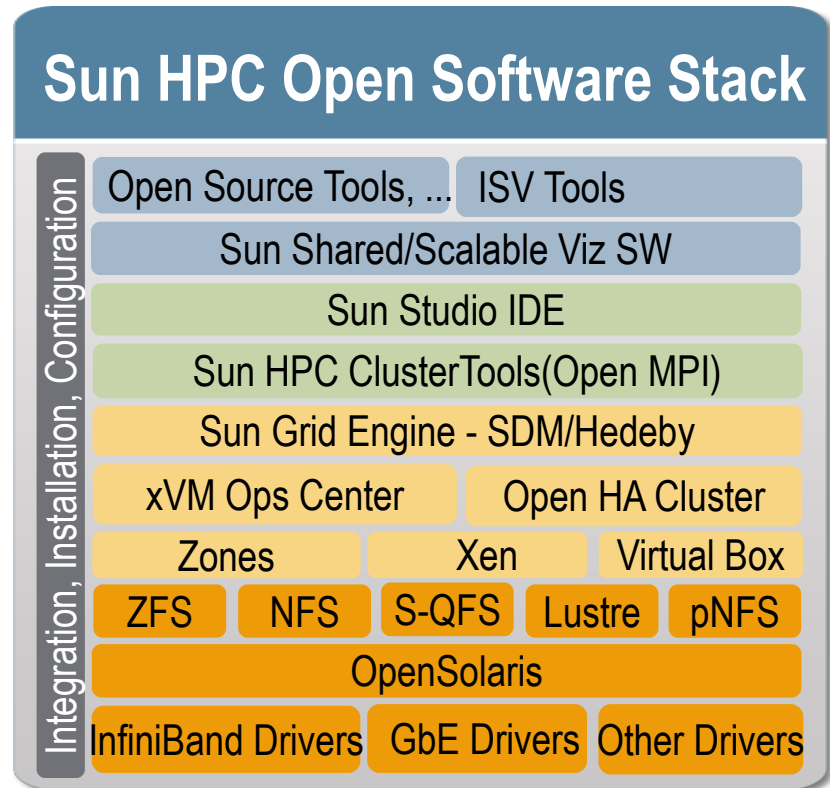
OpenSolaris.com

- Learn
 - > Explore the learning opportunities
 - > Get the guides and videos
- Collaborate
 - > Join the communities
 - > Make your mark; help shape the next generation of the Solaris OS
 - > Connect with Sun distinguished engineers and other leading lights of the OpenSolaris community
- Get it, use it, learn it
 - > opensolaris.com/get



OpenSolaris/Solaris HPC Distro

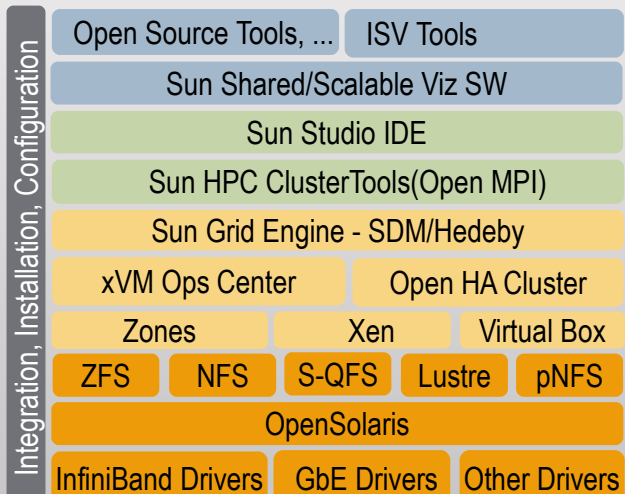
- HPC Developer Community & HPC Stack Project
- Objective:
 - > Build a tightly integrated, low-barrier-to-entry stack
- Two life lines – two platforms Solaris 10 & OpenSolaris*
 - > HPC Developer Preview
 - Centered around IDE
 - Seamless developer experience
 - > HPC vertical distros
 - Deployment stack targeted at specific vertical market segment needs
 - Integrated with OpenSolaris and Solaris 10
 - Also available unbundled



* A Sun-supported Linux HPC stack is also underway

HPC Stack- Download, Join, Succeed


Sun HPC Open Software Stack



- We need your input!
 - > What tools do you need?
 - > What features do you need?
 - > What applications do you have to get integrated?
 - > What platforms do you need?
 - > What support model do you need?
- Participate with us.....
 - > Port applications
 - > Test
 - > Get the word out
- opensolaris.org/os/community/hpcdev/
- opensolaris.org/os/project/hpc-stack/

Award-winning Sun Grid Engine

Thousands of successful Grids



2004

FROST & SULLIVAN

Excellence in Technology Award

ClusterWorld™

REDEFINING HIGH PERFORMANCE COMPUTING

Excellence in Cluster Technology

Grid Engine & Sun Grid Engine

The Ultimate Try & Buy



Grid Engine Open Source

- Full source code access; liberal license (SISSL)
- Regular courtesy binary updates
- Open development and issue tracking
- Vibrant mailing lists
- Community contributions



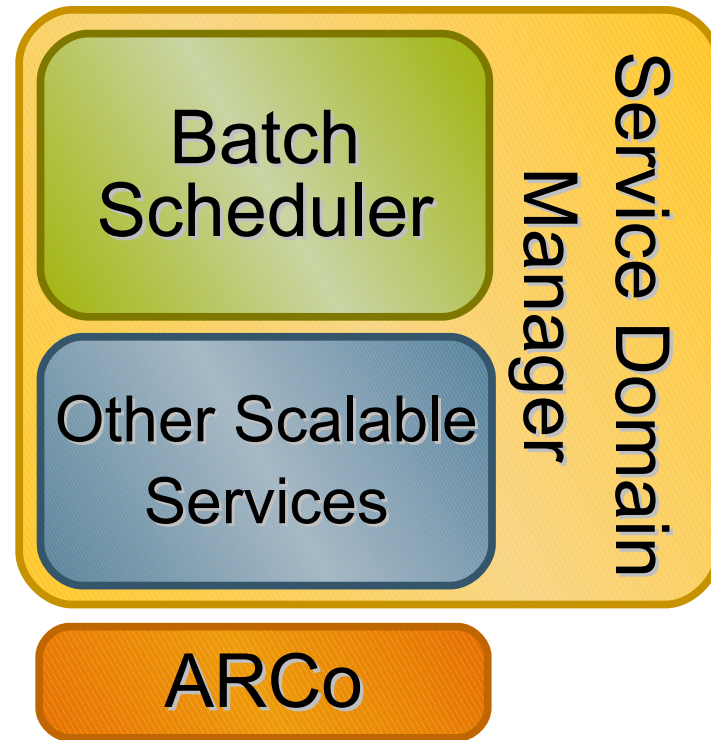

Sun Grid Engine

- Indemnification (also for free download from SDLC)
- Strong Support options; relief patches, on-site diagnosis, etc.
- More “muscle” to influence product direction
- Partner Solutions
- Integrated in Sun's HPC Stack

Future of Sun Grid Engine



Sun Grid Engine 6.x



Sun Grid Engine Next

Grid Engine- Download, Join, Succeed



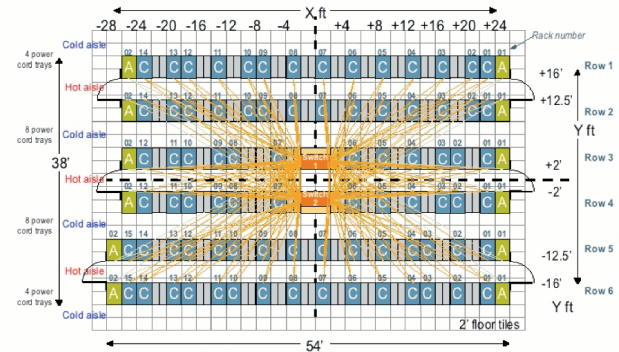
- Download Grid Engine from [Sun Download Center](#) or gridengine.sunsource.net
- Join the Grid Engine community gridengine.sunsource.net
- Join the Hedeby community hedeby.sunsource.net
- Subscribe to the Grid Engine and Hedeby mailing lists
- Get news about Grid Engine at gridengine.info – from the community for the community
- Join Grid.org, a Grid stack community encompassing Grid Engine

Texas Advanced Computing Center

Ranger System


- **First** National Science Foundation Track2 system

- > \$30M acquisition budget
 - \$29M for support over 4 years
- > Awarded September 2006
- > Production December 2007



- TeraGrid member


- > Over 3200 users
- > Over 1000 projects
- > From 48 states
- > Physics, molecular biology, chemistry, astronomy, etc.
- > *Larger than current top 20 TeraGrid systems combined*



Visualization Customer Needs

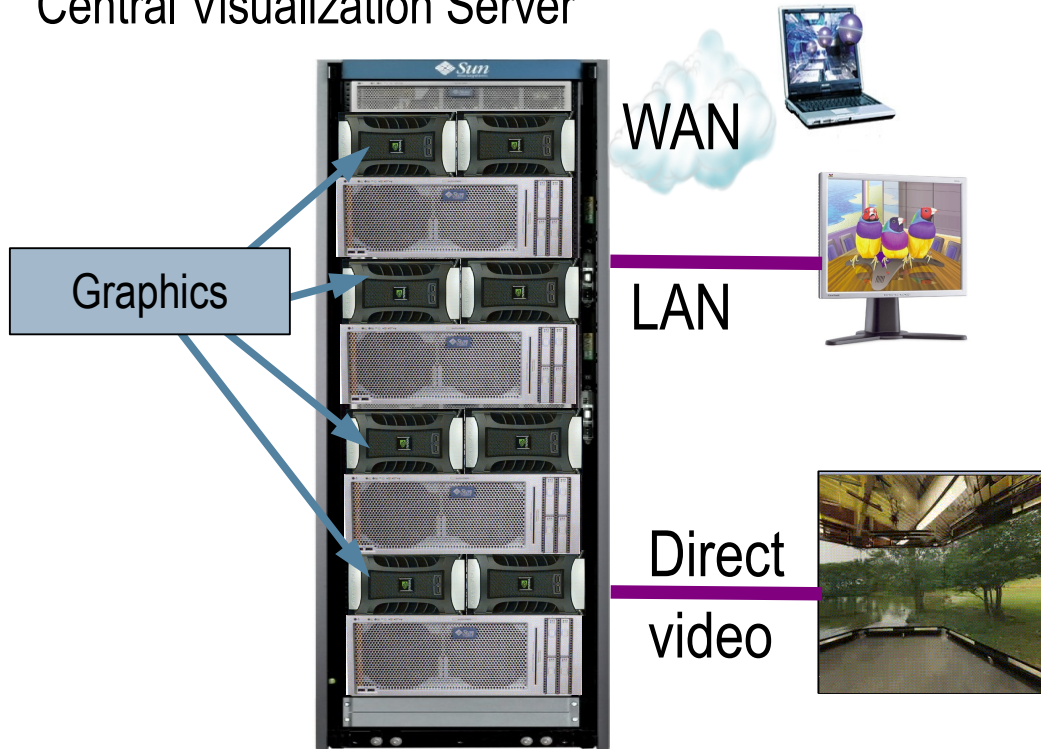
- Understand and act upon increasing quantity and complexity of data
- More and better utilization of graphics resources
- Collaborate with geographically dispersed colleagues
- Secure proprietary data and control access

Share Visualization on the Grid over the Network



Sun HPC Visualization Solutions

Central Visualization Server



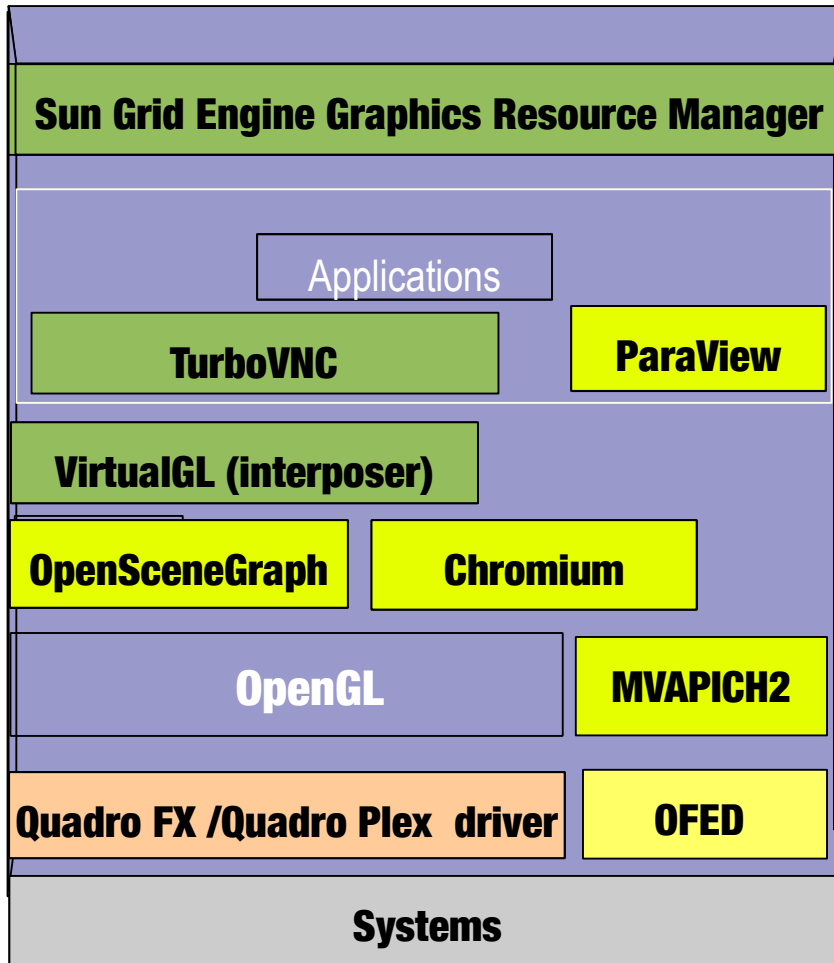
Shared Visualization software provides secure access to 3D apps on a central resource

- > Transparently accessed from a variety of clients
- > Better utilization of resources
- > Access anytime from anywhere
- > Virtualizes Visualization

Scalable Visualization software to combine multiple devices for

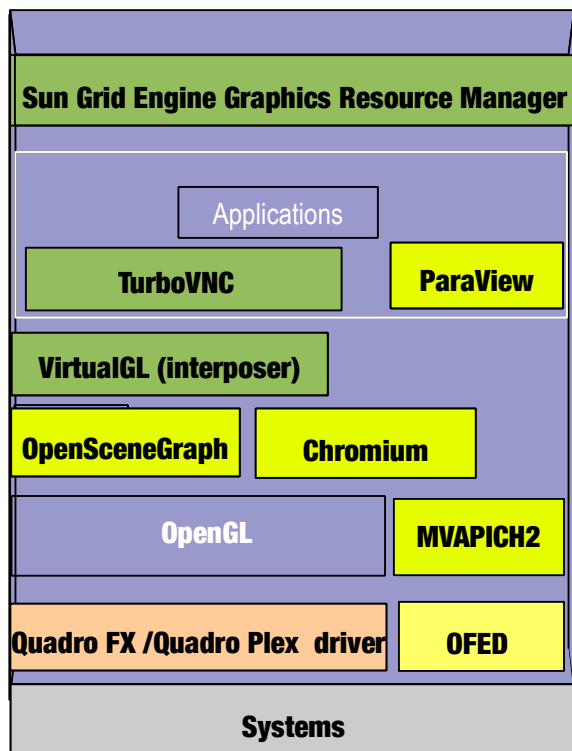
- > Higher Performance
- > More Resolution
- > Virtual Reality

Sun Visualization Stack



- **Shared Visualization software stack** - visualization services to a variety of remote clients
 - **SGE** – open source scheduler extended for graphics resource management
 - **VirtualGL** – open source remote access via any client over standard IP networks
 - **TurboVNC** – open source vncviewer and vncserver that uses TurboJPEG
- **Scalable Visualization software stack**
 - **ParaView** - open-source parallel rendering application optimized for SMPs with multiple graphics.
 - **OpenSceneGraph** - open-source parallel rendering toolkit for building parallel applications.
 - **Chromium** - virtualized graphics devices for Solaris or Linux. Provides transparent parallelization for fill-rate limited applications.
 - **MVAICH2/OpenMPI** - MPI for Linux/Solaris
 - **OFED** – IB for Linux (IB support incl in Solaris)
 - **Quadro Plex** - connects graphics devices to Linux or Solaris servers over a PCI-E cable
- **Systems - Sun Fire x64 & SPARC systems provide scalable platforms**

Viz Stack- Download, Join, Succeed



- Download Visualization Stack from **Sun Download Center**
- Join the Visualization community on **opensolaris.org**
- Join one of the open source communities in the stack
 - > Grid Engine
 - > VirtualGL
 - > TurboVNC
 - > Paraview
 - > OpenSceneGraph
 - > OpenMPI

HPC Development

Sun Studio Software

- C, C++ & Fortran Compilers and Tools
- OpenMP support
- OpenSolaris, Solaris & Linux platforms
- Free license

Sun HPC ClusterTools

- Based on Open MPI
 - Open MPI Message Passing Interface (OMPI)
 - Open Run-Time Environment (ORTE)
- OpenSolaris & Solaris OSs
- Free license

Sun Developer Network

- <http://developers.sun.com>
- Downloads, technical articles, documentation, code samples
- Forums, wikis, events
- CommunityOne, Sun Tech Days

Additional Resources

- Leadership in Programming Language Standards
- NetBeans IDE
- Participation in Industry Conferences
 - SuperComputing
 - International SC

Sun Studio Software Overview

Integrated Toolchain

- Record-setting parallelizing C/C++/Fortran Compilers
- NetBeans-based IDE
- Stable, Scriptable, Multilingual Debugger (dbx)
- Memory Debugger- leak, access, usage (RTC)
- Application Profiling Tools (Performance Analyzer)
- Thread Analysis Tools (Thread Analyzer)
 - > Data race & deadlock detection
- OpenMP v2.5 API Support
- Optimized libraries
 - > LAPACK, BLAS, Sparse-BLAS, SuperLU
 - > FFT, Direct Sparse Solver, Interval BLAS routines



FREE

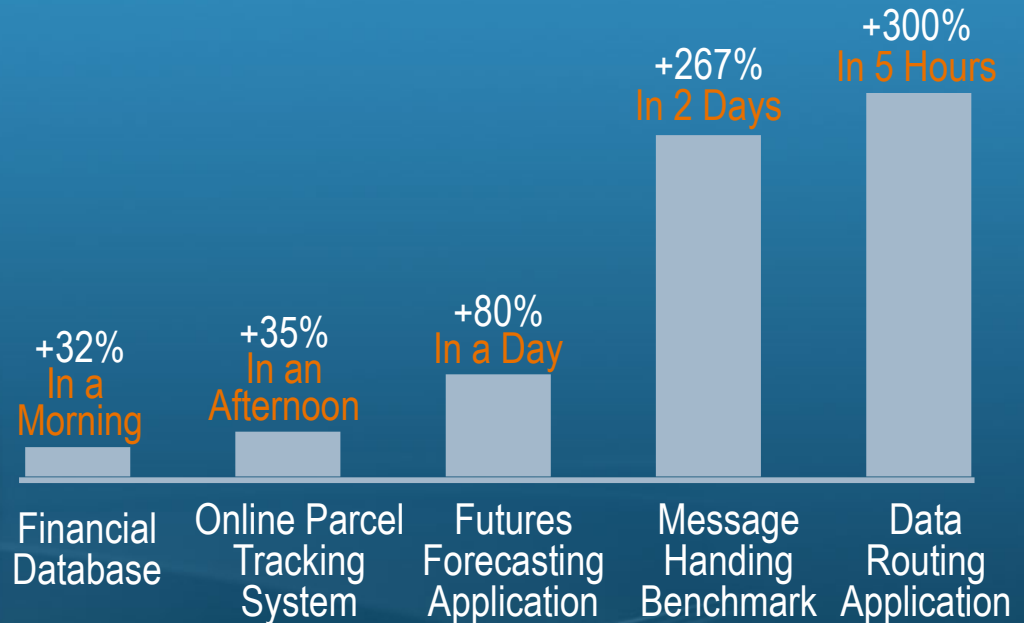


Solaris and OpenSolaris: DTrace



- End-to-end tracing of AJAX/Java/JavaScript Applications
- Monitor JVM internal state
- Diagnose C++ applications with Sun Studio 12 PlugIn
- Easy with NetBeans PlugIn

Increase Performance and Maximize Resource Utilization



Optimize your Applications and your Systems

NetBeans- Download, Join, Succeed



- Download the IDE at netbeans.org
- Join the NetBeans community netbeans.org/community
- Subscribe to the NetBeans mailing lists
- Get news about Netbeans at netbeans.org/community/news
- Participate in the NetBeans Innovators Grants contest (deadline is March 3rd) <http://www.netbeans.org/grant/>
- Learn about and get the Grid Engine and OpenMPI plugins at opensolaris.org/os/project/hpc-stack/

The Open Storage Advantage

Open Source Software with Industry Standard Hardware

Empowers

Developers

to Create Storage
Services Quickly
Creating a Proliferation
of Choice

Offers


Enterprise

Reliability and
Scalability at
1/10th the Cost
of Closed, Proprietary
Storage

The Open Storage Approach

Delivering Breakthrough Economics and Scalability

opensolaris™



General Purpose Components



Storage Server



Network Attached Storage (NAS)



Archive

Virtualization



solaris™



HPC Storage Requirements

Solutions must solve the compute and storage requirements

HPC Data Storage:

- Maximum aggregate at lowest possible costs
- Low latency, accessible by all compute nodes
- Scalable to multiple Petabytes
- Easy to install, deploy and administer

Long-term Retention & Archive:

- Deep repository for very large data multiple Petabytes
- User driven data recall from archive
- Scalable through to move data in and out of archive
- Simple to administer and deploy

Seamless Transfer




TACC Supercomputer Storage

HPC Storage Solutions

Compute Engine Data Cache

- Will scale to over
 - > 72 GB/sec. sustained bandwidth
 - > 1.728 Petabytes of raw capacity
- Configuration includes
 - > 72 SunFire x4500s
 - > Over 3,000 500GB drives
 - > 8 racks



Long-term Retention and Archive

- Will scale to over
 - > 200 Petabytes of near-line
 - > 3.1 Petabytes of on-line
- > Configuration includes
 - > 5 StorageTek SL8500s
 - > 48 StorageTek T10000Bs
 - > 10 StorageTek 6540s
 - > 6 SunFire Metadata servers with SAM-QFS



Seamless Transfer

HPC Storage Solutions

High Bandwidth
Scalable Storage Cluster
with Lustre

Long-Term
Data Retention
with SAM-QFS
Near Line Archive

Compute Cluster

Metadata Servers



Load

Near Line Archive



Archive



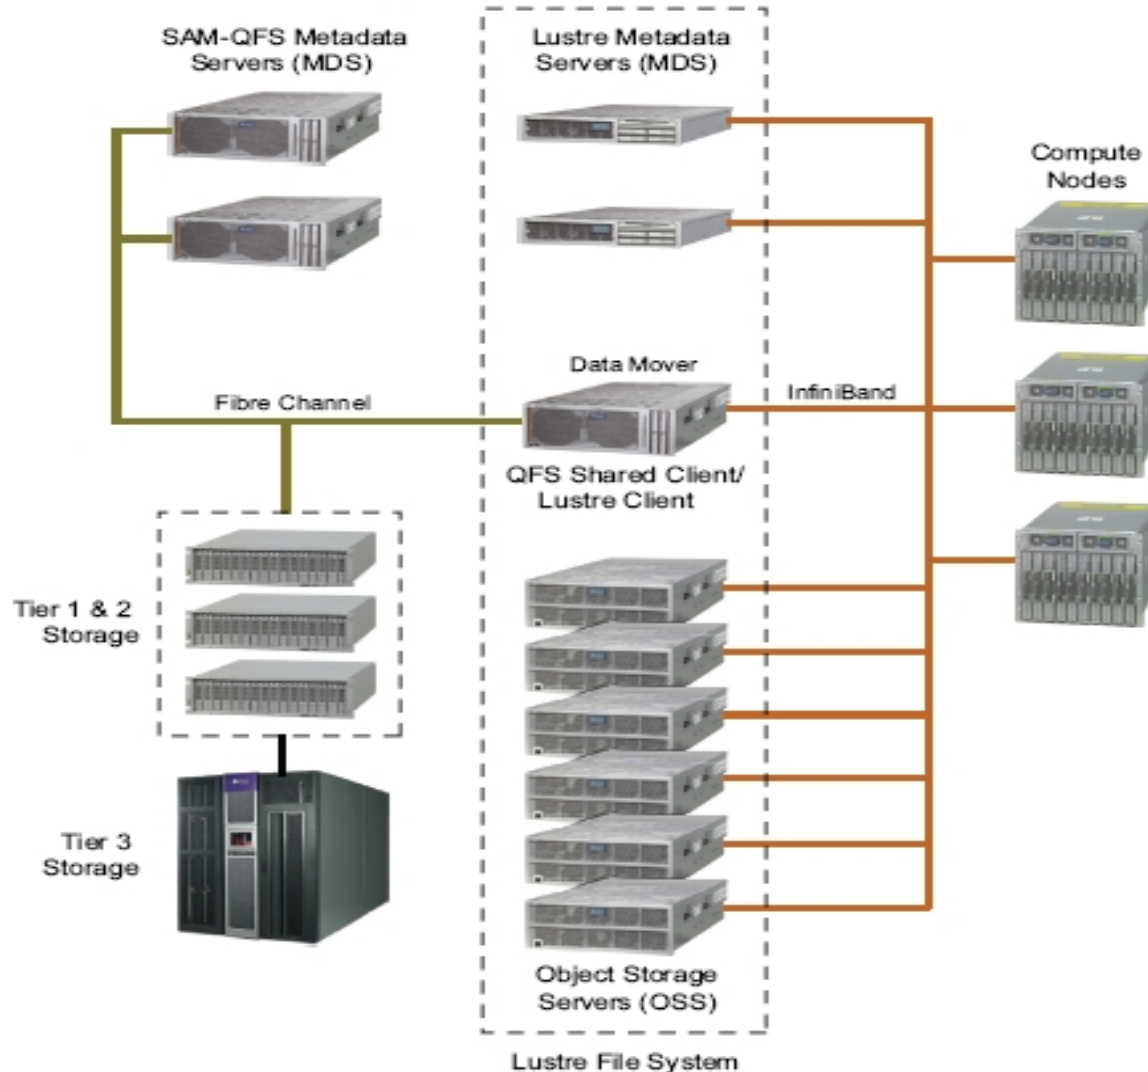
Data Movers



Home Directories
Tier 1 Archive

Object Storage Farm

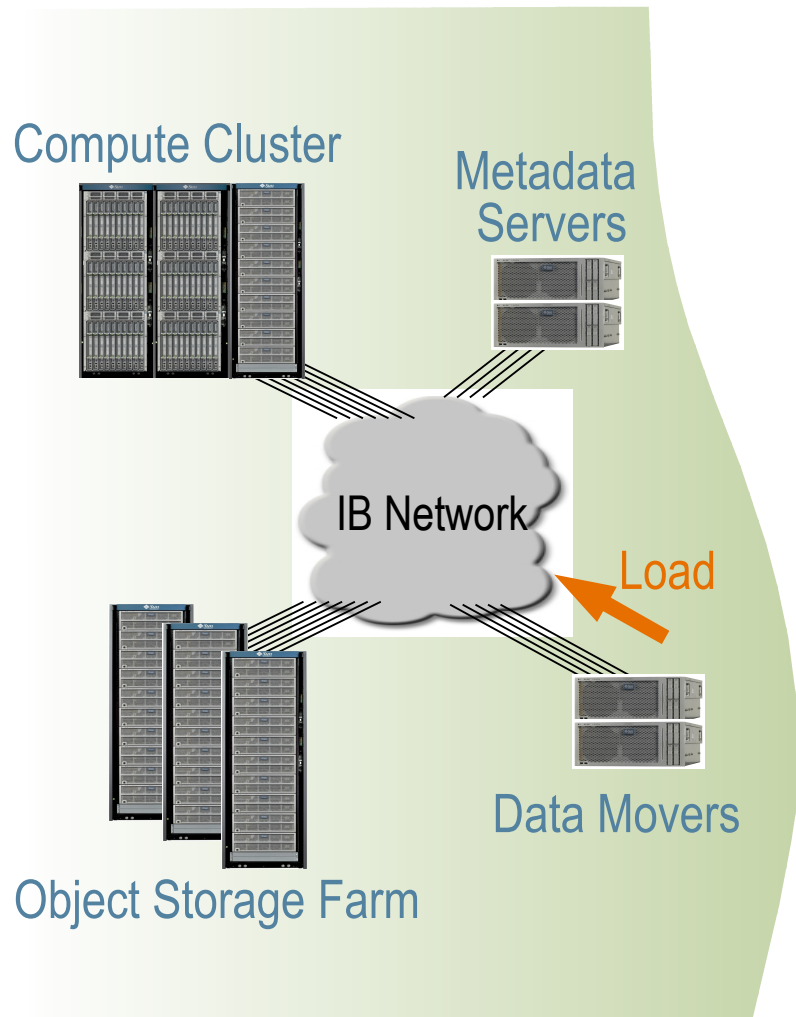
HPC Solution Lustre+QFS+SAM+4500



The combination of Lustre, QFS and SAM, and the Sun Fire X4500 create a truly unique and innovative storage solution that addresses the broadest range of HPC storage requirements in the industry.

Scalable Storage Cluster


High Bandwidth Data Engine



- Achieves extreme scale and aggregate bandwidth required by “Fusion” for a low price/performance
- Leverages both the **SunFire x4500 storage Servers & Lustre** - the defacto standard HPC parallel file system
- Built with industry standard open technologies Infiniband, open source software
- High speed, low-latency interconnects with simplified cabling via Magnum
- Ideal for HPC cache & temp. storage
- Pre-integrate systems – arrive at your site ready to run

Long-Term Retention & Archive

Staging, Storing & Maintaining HPC Data



- Provides a massive on-line/near-line repository to compliment the Scalable Storage Cluster
- Leverages Sun StorageTek Tape Libraries, Modular Arrays and **SAM-QFS**
- Policy driven engine to automate moving data sets in to and results out of the Object Storage Farm
- Enables Tape Libraries as a large near-line repository
- Stores data in open formats (TAR) allowing technology refresh and avoiding vendor lock-in
- Pre-integrate systems – arrive at your site ready to run

Sun Solution Advantages

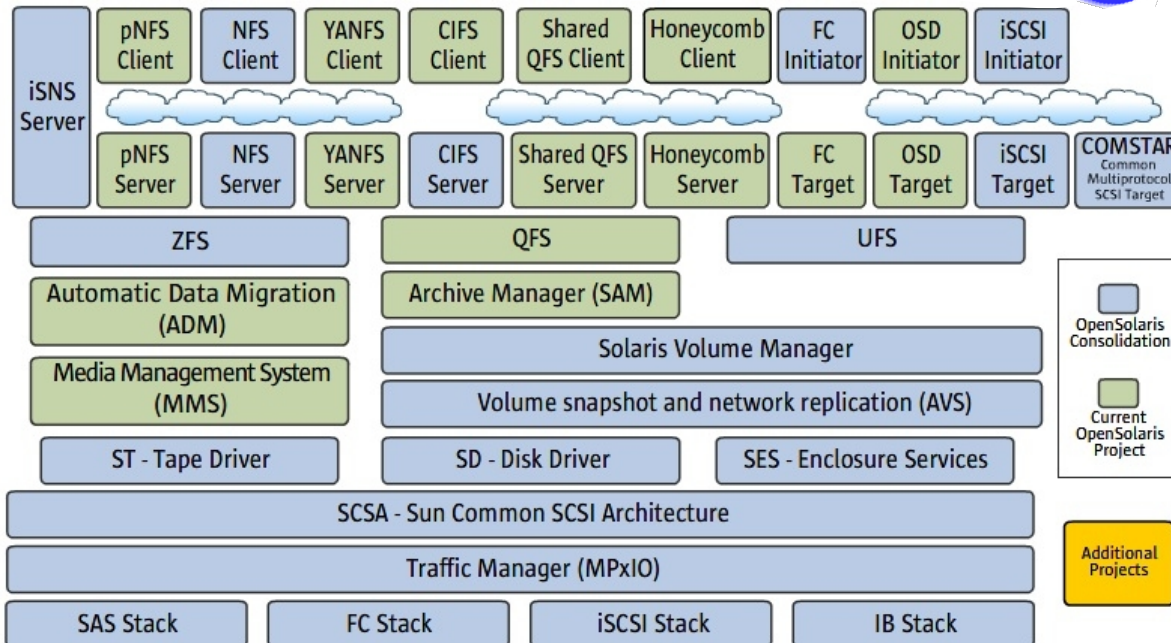
- Lustre, QFS, SAM and ZFS are complimentary.
- X4500 advantages for price performance and density.
- SAM gives Sun automated policy based data migration for data retention.
- Sun has the best Tape Libraries and Tape drives for HPC.
- Sun now has experience delivering these solutions.

Driving Value with Communities and Open Solaris Storage Software

Opensolaris.org/os/storage

Over 3,000 current members!

- 90,000 for opensolaris
- 23 endorsed projects



Participate in the Revolution

Get Informed!

- Watch community town hall, download white papers at www.sun.com/openstorage
- Contact your local Sun representative or authorized reseller

Get Involved!

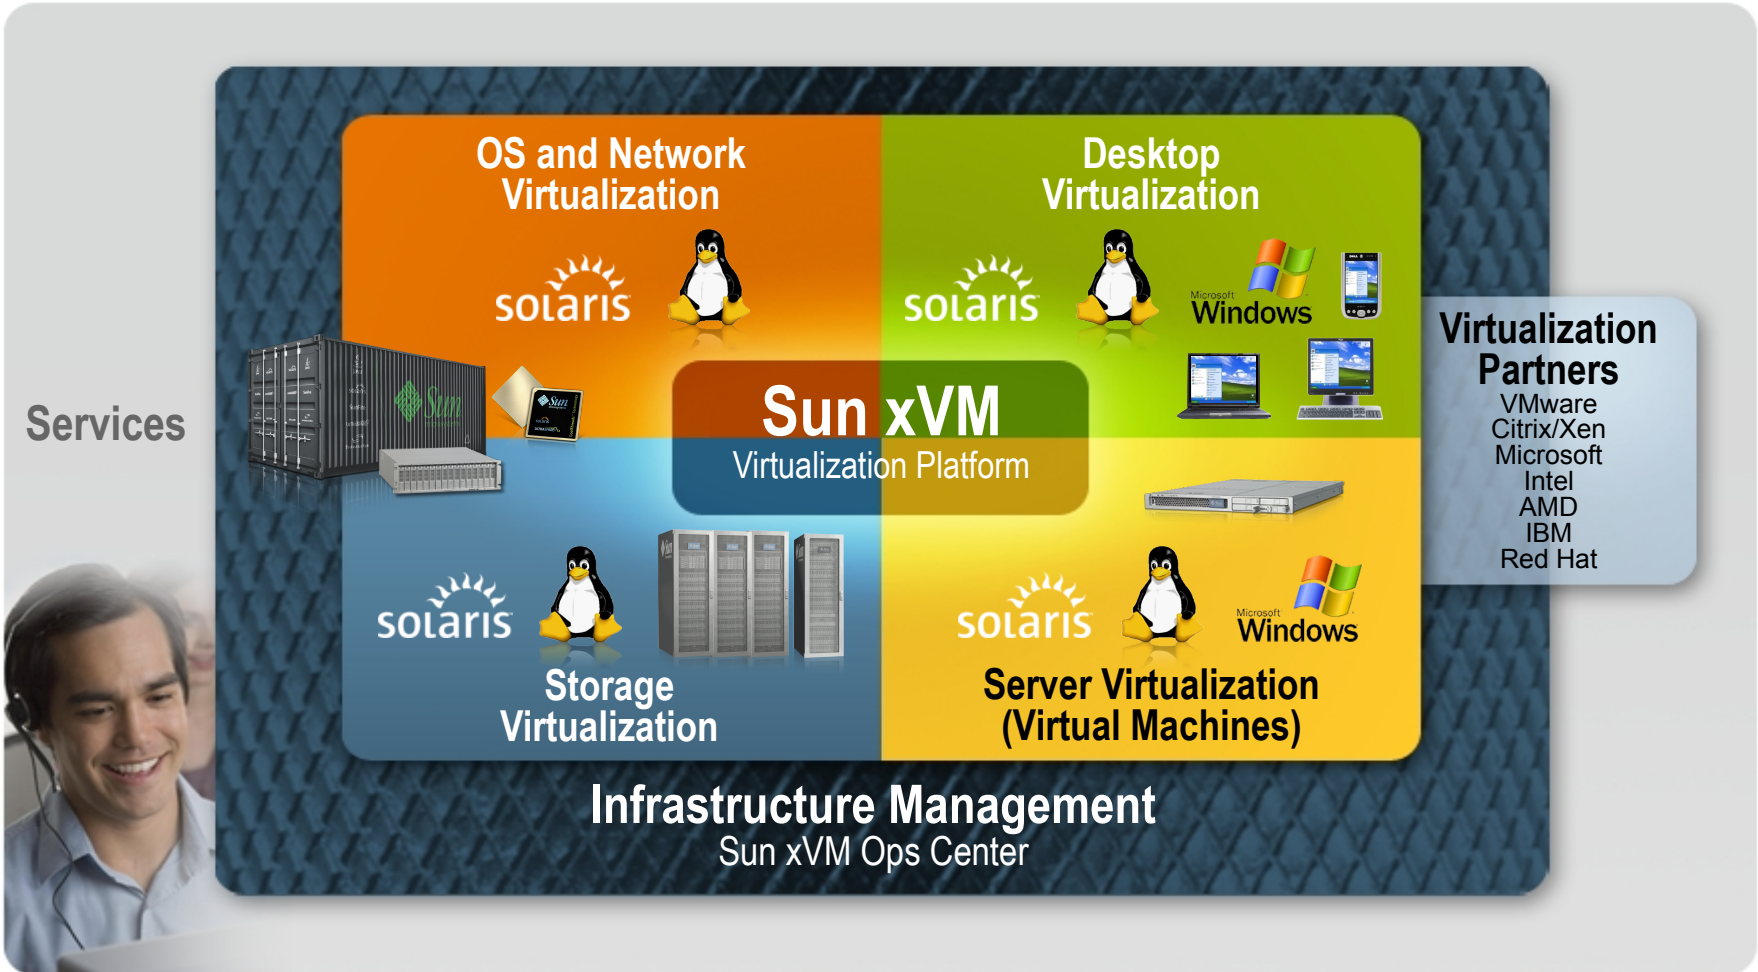
- Join the OpenSolaris Storage Community:
<http://www.opensolaris.org/os/storage/>
- Join the discussion:
<http://www.facebook.com/group.php?gid=12774638>

Get Open!

- Try and Buy Sun products at:
<http://www.sun.com/tryandbuy>
- Engage with Sun Services:
www.sun.com/service/openstorage



End-to-End Virtualization Across the Infrastructure



Sun's xVM Key Components

SUN xVM VIRTUALBOX	SUN xVM SERVER	SUN xVM OPS CENTER
<ul style="list-style-type: none"> • Datacenter automation • Physical and virtual resource management • Manage thousands of hardware and software entities 	<ul style="list-style-type: none"> • Hypervisor family inherits scalability, availability, manageability, observability and security from Solaris • Consolidates Windows, Linux, and Solaris 	<ul style="list-style-type: none"> • Developer desktop virtualization platform • Enables desktop or laptop PCs running Windows, Linux, Mac or Solaris to run multiple operating systems side-by-side • Ease of switching between platforms
<h2>Complete Virtualization and Management Solution</h2>		

2008: Innotek Joined Sun



- Extending Sun's virtualization to the desktop
 - > Ideal for developers to setup multiple virtual machines to develop and test their multi-tier or cross-platforms applications; all on a single laptop or desktop
 - > Inherits scalability, availability, manageability, and observability from Solaris and OpenSolaris
- Expanding Innotek's capacity to deliver global service and support

Open Source Desktop Virtualization

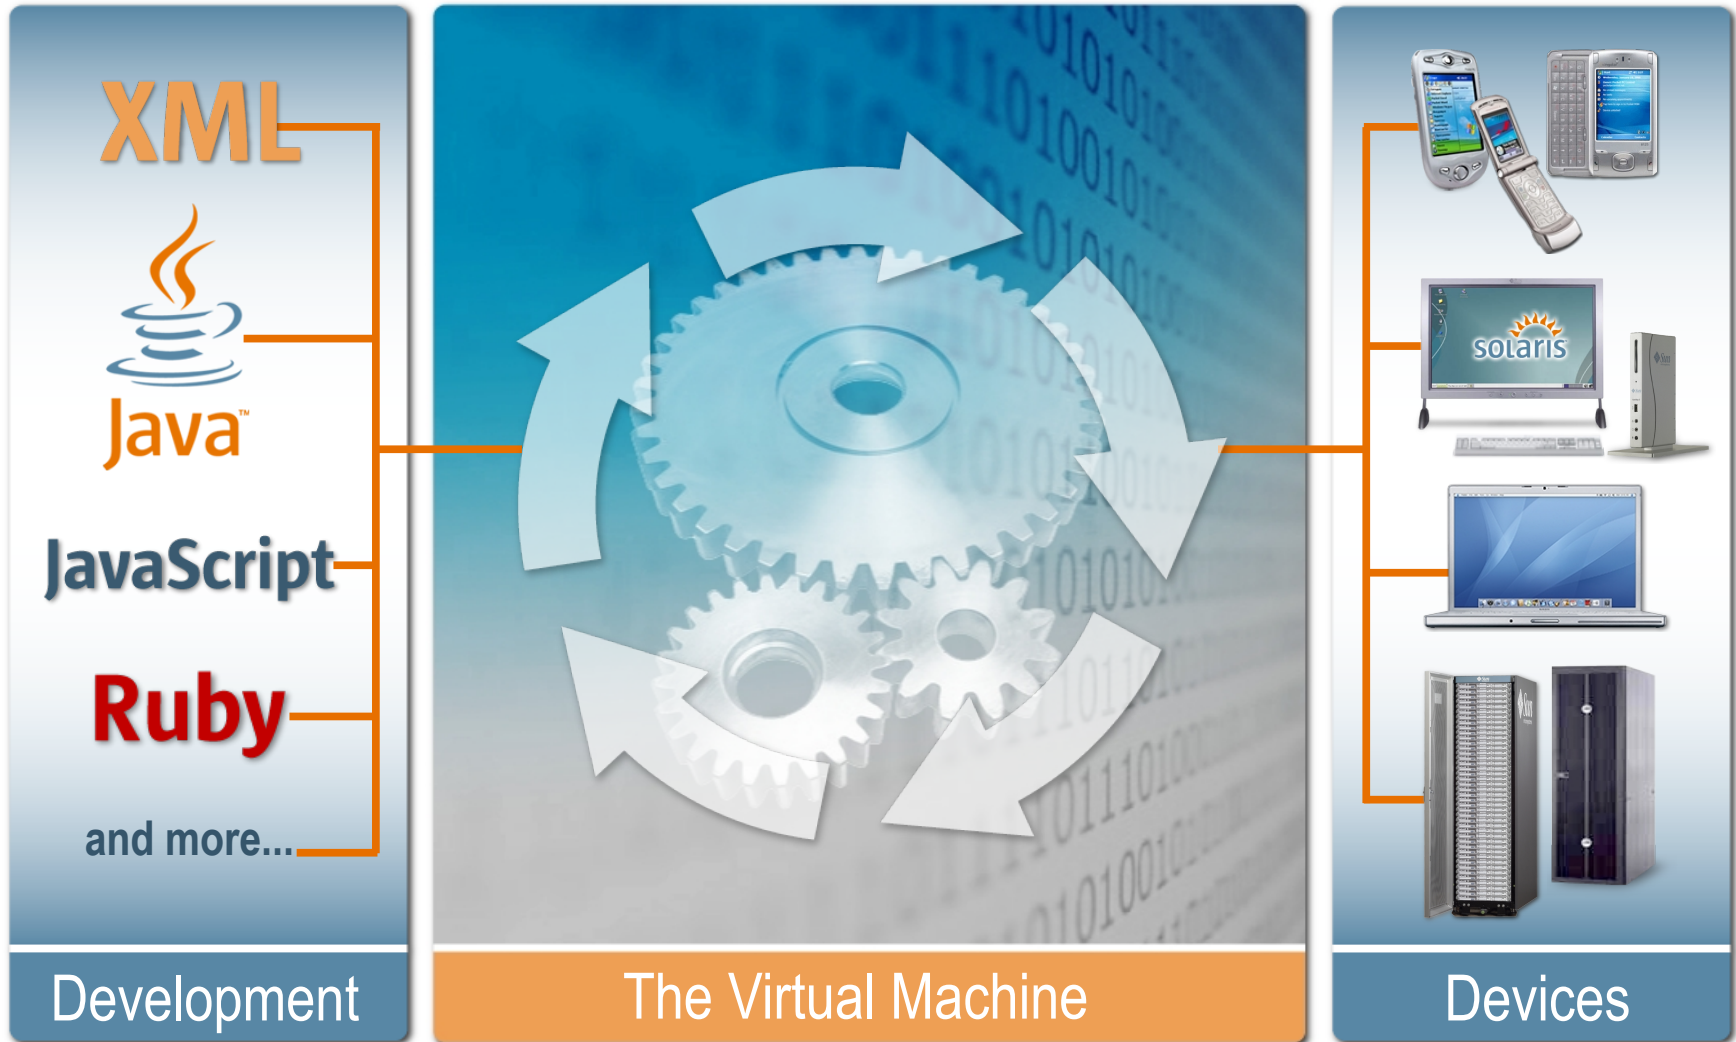
Virtual Box- Download, Join, Succeed

openxVM



- Check out openxvm.org
- Register to join: openxvm.dev.java.net
- Participate in hypervisor development at opensolaris.org
- Download VirtualBox at virtualbox.org/wiki/downloads
- Sign up for mailing lists at openxvm.dev.java.net

Java Virtual Machine: Flexible, Scalable



OpenJDK- Download, Join, Succeed



- Join the conversations at java.net
- Download the JDK at openjdk.java.net
- Join OpenJDK mailing list discussions at mail.openjdk.java.net
- Participate in the OpenJDK Community innovators challenge at openjdk.java.net/challenge
- Check out the Planet JDK blogs at planetjdk.org

Join the community building the World's Apps and Devices

GlassFish: Serving your Applications Faster



- 4 Million downloads since July 2006
 - > Double the downloads of JBoss on a weekly basis
- Dozens of external committers
- Over 7,000 members
- Glassfish(TM) v2 / Sun Application Server set world performance records SpecJAppServer
- Handle business-critical applications and the rigors of production environments

GlassFish- Download, Join, Succeed



- Download GlassFish:
glassfish.dev.java.net
- Register and participate in the GlassFish community
- Enter GlassFish Awards Program
<http://wiki.glassfish.java.net/>
- Subscribe to the GlassFish blog:
<http://blogs.sun.com/theaquarium/>

Sun is Committed to Developer Communities

Building
**Open
and
Free**
Communities



Java

Solaris

SPARC

Building a Vibrant Ecosystem: Sun is the Largest Commercial Contributor to Open Source Communities

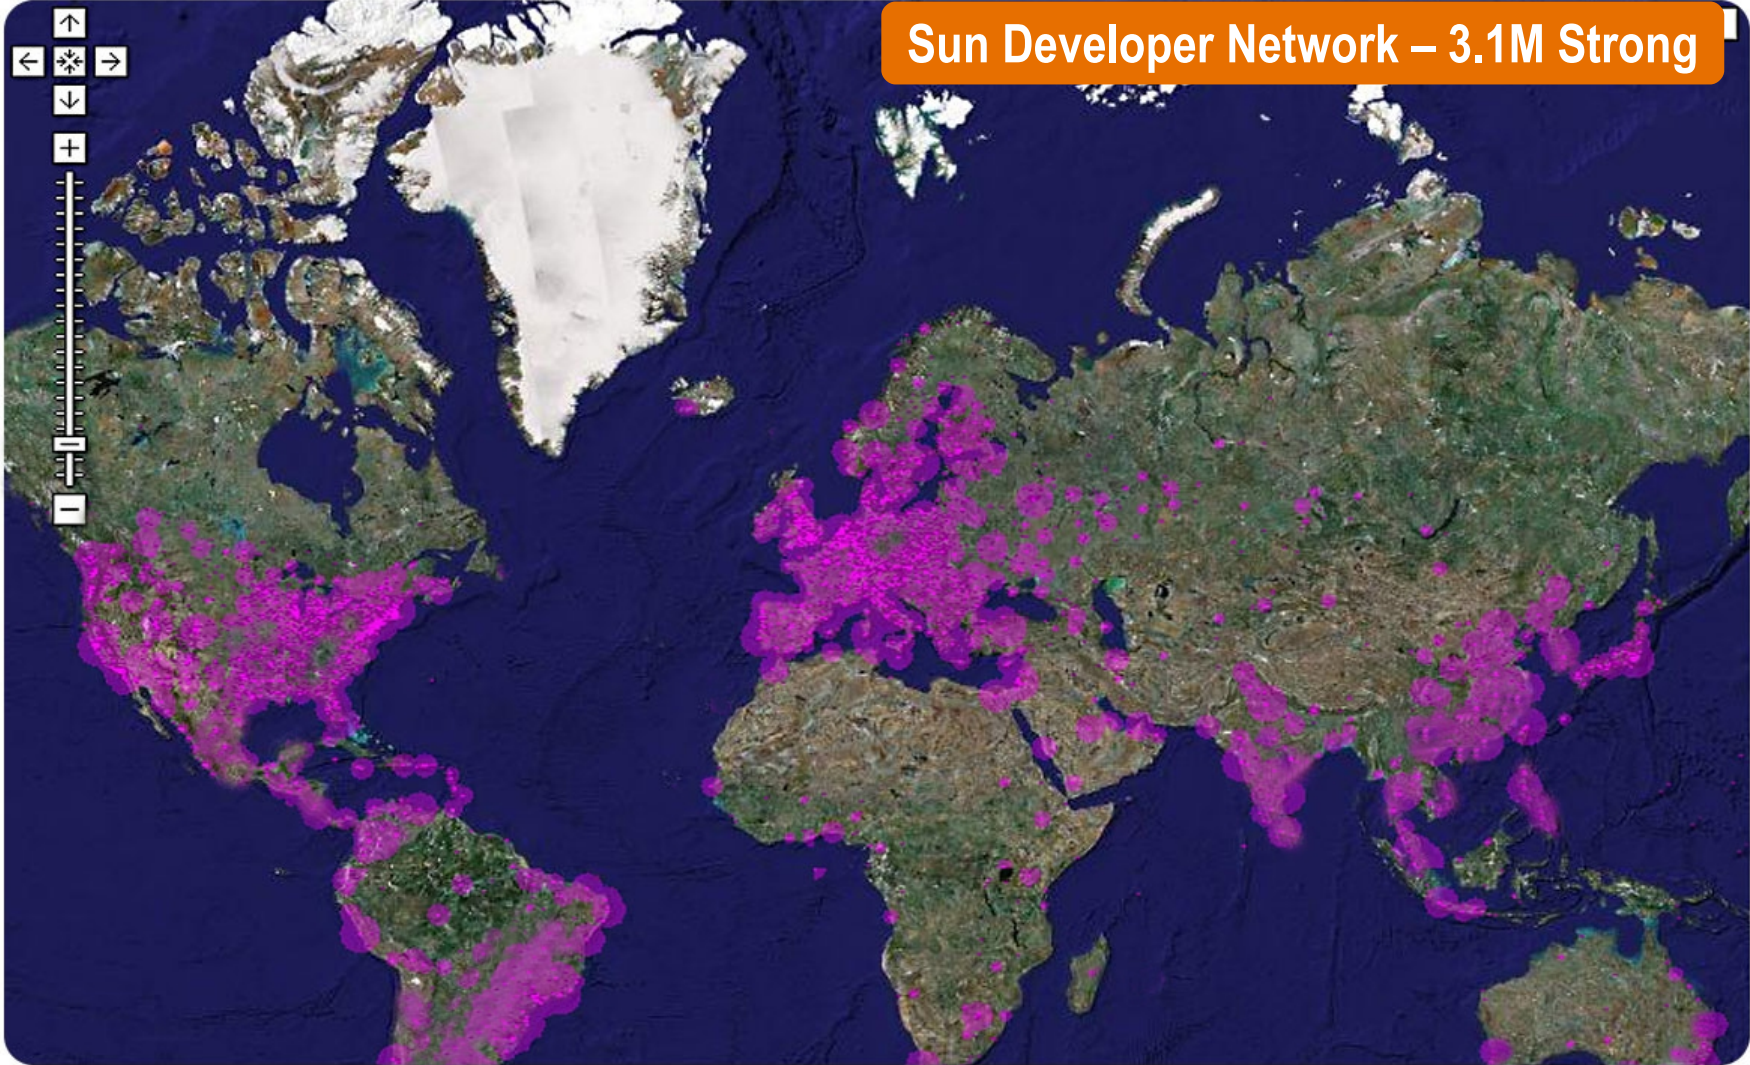
Community

Infrastructure

The Source for Java Technology Collaboration

Ecosystem

The Sun Developer Network Community



Data as of 1/08

Participate, Learn, Thrive



Make the Jump

- **Download** the Tools
- **Leverage** the Technologies
- **Join** SDN developers.sun.com
- **Participate** in the Open Source Communities
- **Build Innovative** Applications
- **Accelerate** Your Career



Thank You.

The Network is the Computer.™

Robert Porras

robert.porras@sun.com

<http://blogs.sun.com/bobp/>

