

Sun Update for IDC

September 9th, 2009

Dave Teszler, Global HPC Sales Manager
Dan Dowling, Ph.D. HPC Sales Specialist

High Performance Computing Division
Sun Microsystems, Inc.

What's up with Oracle?

Questions

What does a growth company need ? - Unsolved problems

1. Are all problems solved in HPC ?
2. Better still, are all Software problems solved ?
3. Are there no Data Intensive problems left ?
4. Would you care about OS Jitter and latency, if you were building a distributed database ?
5. If you were a database and analytics company, would you care about the next grand challenges in data analytics ?
(Say Large Graph Algorithms)
6. Would a database vendor be interested in the faster memory cache hierarchies that Flash and SSD promise ?
7. Would you be interested in a growing product line ?

What's New in HPC at Sun?

Constellation Review: TACC Ranger

- Ranger 579 TFlops



6048 Rack



Pegasus Blade
(4xAMD Barcelona)



DataCenter 3456 IB Switch
(SDR/DDR)



x4500 Storage
(Thumper)

Sun HPC Momentum


- NexGen Constellation Architecture



C48+ Rack
"Glacier" doors
PCIe Gen2 Backplane



Vayu Blade
2x2xNehalem



DataCenter 648 IB Switch
(QDR)
QNEM Leaf Switch





J4400 Storage
(Riverwalk)

Sun Blade 6000 Chassis

10-Rack-Unit Blade Chassis

- Modular Chassis Design
 - > 10RU, 10x SB6000-series servers
 - > Updated midplane supports PCIe Gen2 & SAS2
 - > 4x PCIe Gen2 x8 links per Blade Slot, 2x NEM per chassis
- Power, Cooling and Availability, 2x 5600W PSU (N+N), 1x CMM
 - > All redundant, hot-plug active components



Sun Blade 6048 Unibody Chassis


42-Rack-Unit Blade Chassis

- Unibody Chassis Design
 - > Full rack, 48x SB6000-series Blades, **Up to 96 nodes/768 Cores/Rack ~9 TF (Nehalem 2.93)**
 - > Updated midplane supports PCIe Gen2 & SAS2
 - > 4x PCIe Gen2 links, 2x PCIe EM per Blade Slot, 2x NEM per chassis
- Power, Cooling and Availability, 8x 8400W PSU (N+N)
 - > Glacier Cooling Door option
 - > All redundant, hot-plug active components

Blades, Blades and more Blades



X6275 2x2 Server Module



Sun Confidential - CDA Required



Additional Blade Server Modules

- X6240: AMD Opteron, 2S, 16 DIMMs, 4 HDD, CF 4 PCI-E 8x, REM, FEM, SAS, GigE
- X6270: Intel Nehalem 2S, 18 DIMMs, 4 HDD, CF 4 PCI-E 8x, REM, FEM, SAS, GigE
- X6450: Intel FSB, 4S, 24 DIMMs, CF, 2.5" SSD, 2 PCI-E 8x / 2 PCI-E4x, REM, FEM, SAS, GigE



Sun Confidential - CDA Required



x6270 Constellation Server Module

Dual Socket Enterprise-Class Data Center Compute Engine



- Compute**
 - 2x Intel Nehalem-EP
 - 18x DDR3 DIMM slots - 144 GB using 8GB DIMMs
- I/O**
 - 270 Gbps of throughput (with Gen 2)
 - 2x Hot-Plug PCIe2 ExpressModules
 - 2x Hot-Plug PCIe2 Network Express Modules
 - 1x Gbit Ethernet connection per NEM
 - 2x SAS storage links per NEM (with REM)
 - Optional Fabric Expansion Module (FEM)
 - RAID Expansion Module for external storage (REM)
 - CompactFlash card for boot or storage
- Availability**
 - Four on-board hot-swap disks
 - No Fans or PSU (on blade)
- Management**
 - iLOM 2 Service Processor; IPMI 2.0; remote KVM, Floppy/CDROM - iLOM 3 1Q after RR
 - Solaris, Linux, Windows OS, or VMware
- Density**
 - 2 Sockets/8 cores/16 threads per RU



Sun Confidential - CDA Required



SPARC Blade Server Modules



Sun Blade T6320 Server Module

- One UltraSPARC T2 processor
- 64 threads
- 16 FB-DIMM slots (up to 64GB)
- 2x 10 GbE, 2x GbE
- Four 2.5-inch SATA or SAS HDDs
- Hardware RAID: 0, 1, optional 5 and 6
- Up to 64 LDOMS
- Optional 16GB CF boot media
- iLOM with remote KVMs

Sun Blade T6340 Server Module

- Two UltraSPARC T2+ processors
- 128 threads
- 32 LV FB-DIMM slots (up to 256GB)
- 2x GbE, optional 2x 10GbE
- Two 2.5-inch SATA or SAS HDDs
- Hardware RAID: 0, 1, optional 5 and 6
- Up to 128 LDOMS
- Optional 16GB CF boot media
- iLOM with remote KVMs

Sun Confidential - CDA Required

“Magnum 9”

Densest 2-tier CLOS switch using 36-p chips

- > 648 QDR ports with 72-p CXP Line Card
- > 41 Tbps total capacity
- > 300ns latency (QDR)

Line Cards and Fabric Cards


- > 9 Line Cards
- > 72-p CXP
- > 10G/FCoE/FC gateway
- > 9 Fabric Cards

11RU 19” Enclosure

- > Redundant Power and Cooling
- > 7kW power consumption

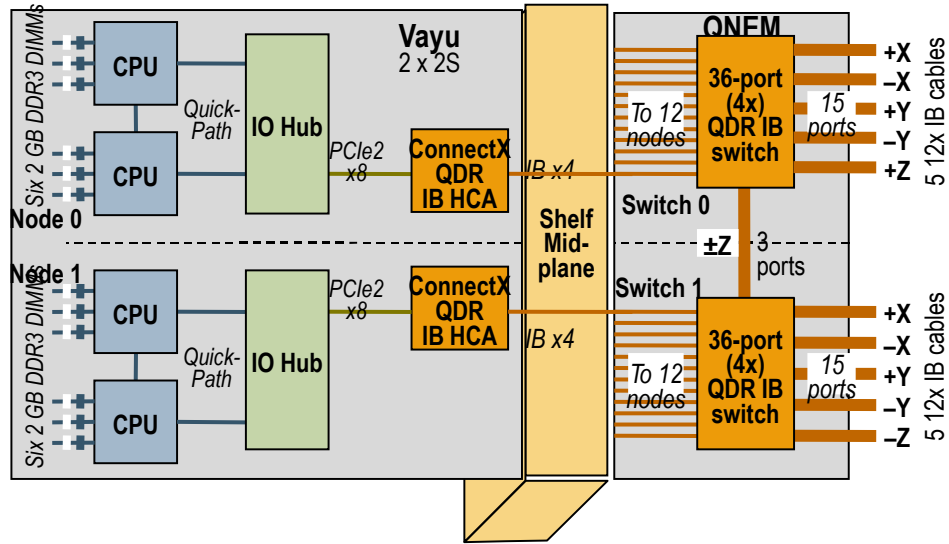
Management

- > Redundant hot-swap service processors
- > External dual redundant subnet managers




3D-Torus Design


1. 12 VAYU Blades + QNEM per Shelf



3. XxYxZ - 4x3x8 Shown (1,152 Nodes) (576 Vayu/48 QNEMs/12 C48)



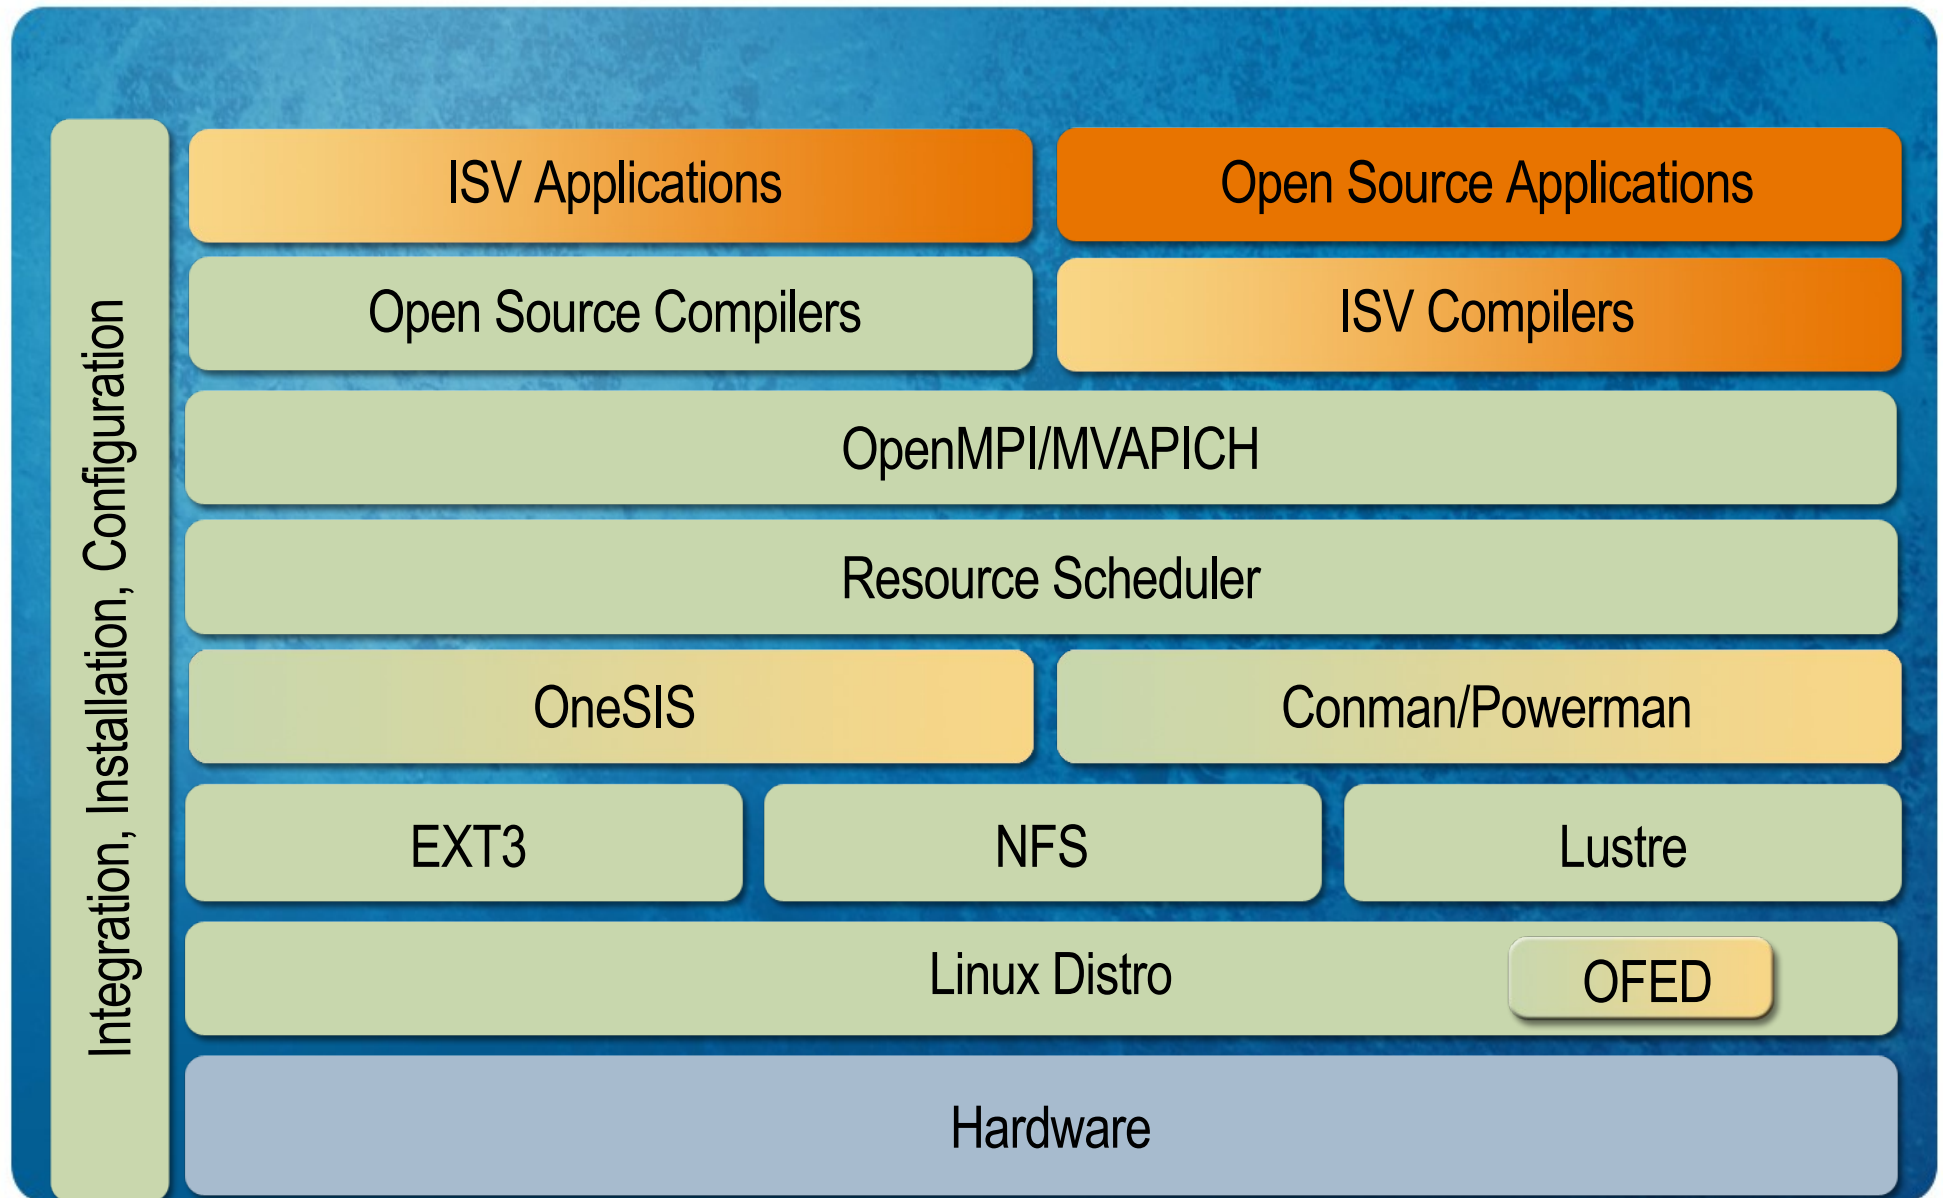
2. C48: 48 Vayu Blades + 4 QNEMs



New Server Memory Hierarchy



Sun HPC Software – Linux Edition



Wins/Momentum

- Major wins in All Geographies (US, APAC, EMEA, EM)
- Mostly Research focused with a few Operational Centers
- Over 20 Systems Contracted at time of Nehalem Launch
- ~200 6048 Racks (Won/Contracted)
 - > ~9600 Vayu Blades
 - > ~38,000 Nehalems
 - > ~1.5 PFlops

200 C48/Vayu Racks (more than 1.5 Pflops)

