

Sun HPC Software and the Sun Constellation System

Linux clusters made simpler



Deploying high performance computing (HPC) infrastructure can be daunting and complex. With the availability of the Sun HPC Software stack, Sun is now able to offer a complete set of infrastructure for HPC environments, from compute servers to scalable InfiniBand interconnects, and even the software stack that manages the cluster. More than individual products, these hardware and software components represent the most successful HPC technologies, all verified to work together. The result is simplified cluster deployments that can scale easily from tens, to hundreds, to thousands of nodes in a space- and energy-efficient fashion.

Highlights

- Sun™ HPC Software, Linux Edition contains all of the components needed to turn bare-metal systems into a running HPC cluster.
- Sun™ x64 servers featuring Intel® Xeon® and AMD Opteron™ processors have produced numerous leading and world-record benchmark results, demonstrating considerable performance and scalability.
- Sun Blade™ modular systems provide new opportunities for dense and scalable HPC clusters, combining speed with large memory capacity and high computational density.
- Sun Customer Ready HPC Clusters offer fast performance and agile deployment, with clusters that are ready-to-deploy and factory integrated for the most demanding HPC applications.
- Sun Datacenter Switches 3456 and 3x24 provide innovative InfiniBand interconnects, scaling from a few blades to the largest supercomputers in existence.

Open and verified HPC software

The software needs of those deploying today's HPC clusters go well beyond the mere cost and performance of technology. To achieve reduced to results, organizations need:

- Simple and scalable provisioning
- Validation of complex configurations
- Integrated management and monitoring
- Community-based technologies

To address these demands, Sun HPC Software, Linux Edition delivers an integrated software stack for Linux-based HPC (Figure 1), with a complete set of highly-scalable tools and processes to configure, manage, monitor, and provision an HPC cluster. The stack includes both Sun and open source components including the operating system, cluster management software, compilers, OpenMPI and OpenMPI development tools, job schedulers, and file systems.

Sun HPC Software is built to be completely modular so that organizations can customize it according to their own preferences and requirements. The distribution also includes a verification suite that determines if the stack and third-party components will install and behave correctly. Downloadable from Sun, Sun HPC Software contains the complete component framework needed to turn bare-metal systems into an HPC cluster.

Key Sun HPC Software components include:

- *Sun Studio software* compilers and tools provide a seamless and integrated development environment for both floating-point and data-intensive computing. Record-breaking C/C++/Fortran 2000 compilers, a native OpenMP API, and state-of-the-art performance analysis tools are included.
- *Sun Grid Engine software* provides policy-based workload management and dynamic provisioning of application workloads. A Sun Grid Engine master can manage a grid of up to ten thousand hosts, meeting the scalability needs of even the largest grids.
- The *Lustre™ file system* provides HPC data storage with unprecedented scalability of capacity and I/O bandwidth. As a parallel file system, Lustre technology can handle very high volumes of data.

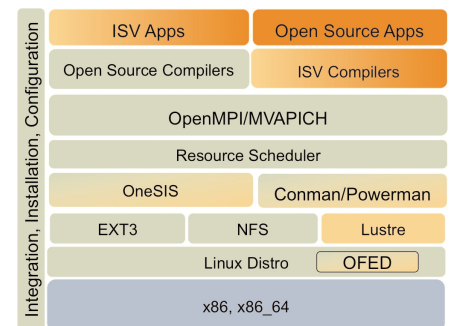


Figure 1: Sun HPC Software, Linux Edition

A choice of robust HPC hardware

Complementing Sun HPC Software, Sun offers a broad family of compute nodes (Table 1), including both x64 rackmount servers and modular systems. These systems offer a some of the largest memory footprints available in the marketplace, with a choice of powerful AMD Opteron™ and Intel® Xeon® processors. Processor choices include both dual-core and quad-core versions so that organizations can select CPUs that are appropriate to their needs.

Innovative modular systems

Sun Blade™ 6000 server modules can be deployed in a choice of available chassis:

- The *Sun Blade 6000 modular system* — a 10-blade chassis that fits into a standard rack
- The *Sun Blade 6048 modular system* — a 48-blade chassis that offers a unibody design

For smaller clusters, the Sun Blade 6000 is ideal as it provides a scalable way to grow from a few nodes to tens of nodes using conventional racks and switches. When more nodes are needed, upgrading to the Sun Blade 6048 chassis is easy, as all the nodes can be simply re-deployed in the larger, more dense chassis.

Table 1: Sun x64 servers provide a choice of processors and options

Sockets/Cores	Memory	AMD Opteron Processors	Intel Xeon Processors
2 / 4 or 8	Up to 64 GB	Sun Fire™ X2200 M2 server Sun Fire X4140 server Sun Fire X4240 server Sun Blade™ X6220 server module	Sun Fire X4150 server Sun Fire X4450 server Sun Blade X6250 server module
4 / 8 or 16	Up to 128 GB	Sun Fire X4440 server	Sun Fire X4450 server Sun Blade X6450 server module
8 / 16 or 32	Up to 256 GB	Sun Fire X4600 M2 server	—

The Sun Constellation System

The Sun Constellation System is the world's first open petascale computing environment, combining ultra-dense high performance computing, interconnect, storage and software. This integrated system delivers massive scalability, dramatically reduced complexity, and breakthrough economics. The Sun Constellation System is based on the highly-dense Sun Blade 6048 chassis.

The InfiniBand fabric for the Sun Constellation System is implemented using Sun components, including the PCI Express Network Expansion Modules (NEMs), innovative cabling, and a selection of switches that scale from a small fabric to terascale or petascale HPC clusters. Sun InfiniBand components include:

- *The Sun InfiniBand Switched NEM* for non-blocking redundant connectivity from 12 compute nodes to the InfiniBand fabric
- *The Sun Datacenter Switch 3x24* for support of up to 72 nodes (288 nodes maximum with four switches deployed)
- *The Sun Datacenter Switch 3456* for support of 3,456 nodes (13,824K maximum with four switches deployed)

Learn More

For more information on the Sun HPC Software stack and the Sun Constellation System, please visit: sun.com/hpc and sun.com/sunconstellationsystem

Integral to Sun Datacenter Switches are innovative 12x cables that utilize the standard Compact Small Form-factor Pluggable (CSFP) connector. These cables provide 3:1 cable consolidation which results in fewer cables to deploy, less weight, and ease of configuration.

The Sun Fire™ X4500 server provides an effective choice for building out clustered storage. Each Sun Fire X4500 server offers the combined capabilities of a two-socket x64 server and up to 48 terabytes of high density storage in a very small (4U) rackmount footprint. When combined with the powerful capabilities of the Lustre file system, this hybrid storage device provides an effective and scalable data cache accessible by all of the cluster nodes.

The Sun Customer Ready HPC Cluster

The Sun Customer Ready HPC Cluster can greatly simplify the design, development, and implementation of HPC infrastructure. With Sun x64 servers, high-performance networking options, and grid-ready resource and system management software, the Sun Customer Ready HPC Cluster provides an integrated system that is ready to deploy. Architected, built, configured, and tested by experts in Sun's award-winning factories, these systems help make clusters simpler than ever before.