

# SUN™ CUSTOMER READY HPC CLUSTER: REFERENCE CONFIGURATIONS WITH SUN FIRE™ X2200 M2 AND X2100 M2 SERVERS

Jeff Lu, Systems Group  
Sun BluePrints™ OnLine — April 2007

Part No 820-0835-11  
Revision 1.1, 4/23/07  
Edition: April 2007

## Table of Contents

<b>Sun Customer Ready HPC Cluster: Reference Configurations with Sun Fire X2200 M2 and X2100 M2 Servers</b> .....	<b>1</b>
Reference Configurations .....	1
Sun Customer Ready HPC Architecture .....	2
Sun Customer Ready HPC Cluster Hardware Components .....	3
Component Configuration Rules .....	5
Reference Architecture Components .....	6
Sun Fire Server Compute Nodes .....	6
Rack System Management Infrastructure .....	6
Data Fabric .....	7
Storage Interface .....	7
Operating System .....	8
System Management Software .....	8
Grid Computing Management .....	8
Reference Configurations .....	8
HPC Entry Level Reference Configuration .....	8
HPC Mid Level Reference Configuration .....	13
HPC InfiniBand Reference Configuration .....	17
HPC InfiniBand Non-Blocking Reference Configuration .....	21
Default Operating System Configuration .....	25
Solaris Operating System .....	25
Red Hat Linux Operating System .....	26
SUSE Linux Operating System .....	27
Default Networking and Access Configuration .....	27
Default Network Assignment .....	27
Default Hostname and IP Address Assignments .....	28
Service Processor Access .....	30
Sun N1 Grid Engine Software Defaults .....	30
Parts List .....	31
Ordering Information .....	33
About the Author .....	33
Acknowledgements .....	33
References .....	33
Ordering Sun Documents .....	34
Accessing Sun Documentation Online .....	34

## Sun Customer Ready HPC Cluster: Reference Configurations with Sun Fire X2200 M2 and X2100 M2 Servers

The Sun™ Customer Ready HPC Cluster includes servers, networking, interconnects, and software installed in a Sun Rack cabinet. The Sun Customer Ready HPC Cluster with Sun Fire™ X2100 M2 and X2200 M2 servers is optimized for High-Performance Computing (HPC) applications such as scientific research, mechanical computer-aided engineering (MCAE), electronic design automation (EDA), financial analysis, and any other compute-intensive application. These configurations focus on delivering computing power with high space density, in a cluster or grid architecture.

All Sun Customer Ready HPC Clusters are configured and tested by the Sun<sup>SM</sup> Customer Ready Systems program, which offers factory-integrated solutions using Sun and third-party products. Such factory integrated systems help enable businesses to deploy IT solutions in a simpler, safer, and swifter manner. As a result, customers gain on productivity and save on deployment costs and risk. (For more information, consult <http://www.sun.com/servers/cr>.)

Warranty and post-sales support for Sun Customer Ready HPC Cluster systems using select, Sun price-listed components is provided through Sun's Service Plans for Systems. For Sun Customer Ready HPC Cluster components that are not on Sun's price list, post-sales support may be contracted through Sun's Managed Services.

### Reference Configurations

The reference configurations described in this paper are starting points for building Sun Customer Ready HPC Clusters configured with Sun Fire X2100 M2 and X2200 M2 servers. The configurations define how Sun Systems Group products can be configured in a typical grid rack deployment. This document describes configurations using Sun Fire X2100 M2 and X2200 M2 servers with a Gigabit Ethernet data fabric, as well as configurations using Sun Fire X2200 M2 servers with a high-speed InfiniBand fabric. These configurations focus on single rack solutions, with external connections through uplink ports of the switches.

These reference configurations have been architected using Sun's expertise gained in actual, real-world installations. Within certain constraints, as described in the later sections, the system can be tailored to the customer needs. Certain system components described in this document are only available through Sun's factory integration. Although the information contained here could be used during an integration on-site, the optimal benefit is achieved when the Sun Customer Ready Systems program integrates the system.

The remaining sections of this document provide the following information:

- High level system architecture
- Interconnect diagrams
- Management node strategy and recommended management software
- Rack limits, size, airflow, weight, and power specifications
- Configuration rules
- List of components
- Default network and disk partition information

## Sun Customer Ready HPC Cluster Architecture

The Sun Customer Ready HPC Cluster consists of a set of components integrated in a Sun Rack 1000-42/38 rack, with interconnections and management utility to deliver compute and system management operations. The Sun Customer Ready HPC Cluster provides an Ethernet management fabric for the management of components, connected to the rack management node. The applications are distributed to run on the Sun Fire servers, delivering results through the data fabric, which consists of either Gigabit Ethernet switches or higher performance InfiniBand switches. An optional provisioning fabric for operating systems and applications can also be configured through Gigabit Ethernet switches in the rack.

Figure 1 illustrates the overall layout of the Sun Customer Ready HPC Cluster components, including the optional Keyboard Monitor and Mouse (KMM) for local interface and the grid master node for scheduling grid computing resources.

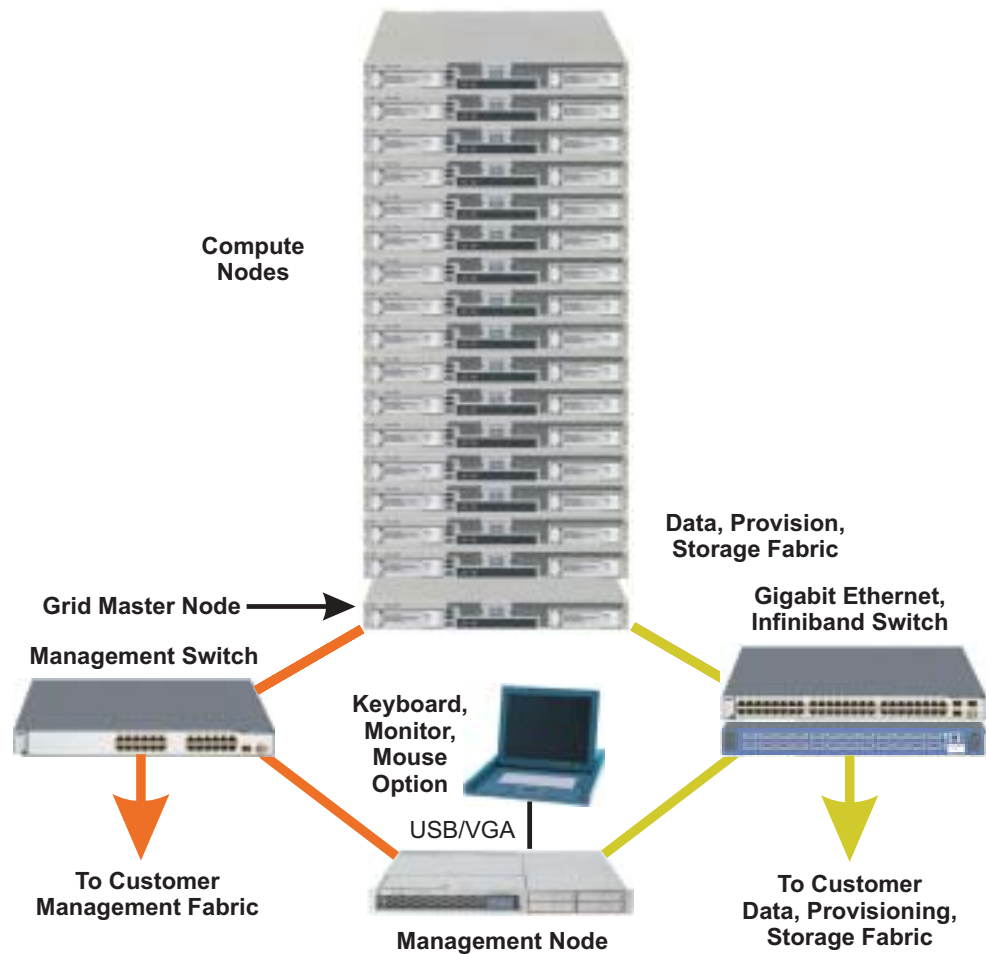


Figure 1. Sun Customer Ready HPC Cluster Architecture.

### Sun Customer Ready HPC Cluster Hardware Components

Each Sun Customer Ready HPC Cluster contains Sun Fire servers configured as compute nodes, a management node, and a grid master node.

- *Compute Nodes* — The compute nodes in the system are the primary engines performing the computation in the rack. Typical usage in HPC involves provisioning of the OS and tasks on to each compute node. Smaller workloads can be loaded on each node's local disk. Larger data sets are delivered through external storage connections.
- *Management Node* — The management node in the Sun Customer Ready HPC Cluster runs the Sun N1™ System Manager software. Connecting through the management port, the management node communicates to each of the server nodes and switches in the rack for setup and monitoring. The management node handles DHCP, PXE booting, OS provisioning, and application provisioning tasks.

- *Grid Master Node* — The grid master node handles task scheduling and communication between compute nodes for a grid environment. The grid master node can also communicate with external compute nodes through uplinks in the switches. Typically, a separate server is configured as the grid master node using the Sun N1™ Grid Engine software; this node communicates with the computing nodes through the data fabric and the provisioning fabric.

In addition, each Sun Customer Ready HPC Cluster contains the following hardware components that provide system interconnects; rack, power, and cabling; and optional KMM and terminal server functionality:

- *Gigabit Ethernet Data Switch* — The rack system uses either a 24-port or 48-port Gigabit Ethernet switch as the data switch. The data switch serves as the primary interconnection for the compute nodes and the external data interface.
- *Management Switch* — The management switch is connected to the management node, the Service Processor port of each of the compute nodes, and the management port of the components in the Sun Customer Ready HPC Cluster. Uplink ports are used to connect to an external management network.
- *InfiniBand Switch* — An InfiniBand switch serves as the main connection between the compute nodes and external InfiniBand fabric. The switch supports 4X InfiniBand, with non-blocking features between ports. A 24-port 1U switch as well as a 96-port 6U switch offer different InfiniBand topologies for the rack, depending on the site requirements.
- *Storage Data Switch* — A separate Gigabit Ethernet storage data path is recommended for servers connecting to the Network Attached Storage (NAS) device. Additional high performance storage architecture options, such as Fibre Channel, or InfiniBand are possible for custom configuration for each server through additional option cards and switching components.
- *Rack and Power* — All components are pre-installed in either a Sun™ Rack 1038 or 1042. The rack is equipped with front and rear doors, and side covers. The Sun Customer Ready HPC Cluster is equipped with a full configuration of power distribution units, which provide sufficient power outlets to cable all installed nodes. Each component has a separate AC power cord. The power cords are routed from the rear, with side-mounted outlet strips to minimize the impact on rear access to the components in the rack.
- *Cabling and Interconnect* — All cable harnesses are pre-cut, routed, and terminated in locations that provide minimal excess cable without impeding component access or replacement. For a Sun Customer Ready HPC Cluster configuration with all component types, the rack can be configured with 3 separate types of harnesses. (Data fabric: yellow, Management fabric: orange, Serial: white.) InfiniBand cables are handled with cable management attachment to ensure proper installation and operation.

- *Keyboard, Monitor, Mouse* — A keyboard, video monitor, and mouse (KMM) unit in the rack system is cabled to the management node by default, and can be used to connect to each of the compute nodes for local interfacing if required.
- *Terminal Server* — All nodes in the rack and the network switches have the option of connecting via their serial console ports to a 48-port terminal server. The terminal server ports use consistent RJ-45 serial port wiring standards. The terminal server can then be used to perform individual configuration functions remotely, if necessary. The terminal server provides normal telnet and secure shell (SSH) access to a serial port using the telnet IP-address port-number protocol.

## Component Configuration Rules

Table 1 lists the minimum and maximum number of components in a single Sun Customer Ready HPC Cluster.

*Table 1. Component Configuration Rules for Sun Customer Ready HPC Clusters.*

<b>Component</b>	<b>Minimum</b>	<b>Maximum</b>
1U compute nodes	2	40
Management node	0	1
Grid master node	0	2
Gigabit Ethernet switch	0	4
InfiniBand 24-port switch	0	2
InfiniBand 96-port switch chassis	0	1
Keyboard, Monitor, Mouse (KMM) unit	0	1
Terminal server	0	1

## Reference Architecture Components

Each Sun Customer Ready HPC Cluster contains integrated Sun Fire server compute nodes, rack system management infrastructure, a Gigabit Ethernet or InfiniBand data fabric, and options for the storage interface. Multiple operating systems, including the Solaris™ Operating System (Solaris OS) and Linux, are supported. In addition, the Sun Customer Ready HPC Cluster is configured with the Sun N1 System Manager and the Sun N1 Grid Engine software.

### Sun Fire Server Compute Nodes

The Sun Fire X2100 M2 and X2200 M2 family of servers are the primary compute nodes in the system. The Sun Fire X2100 M2 server is a single socket Opteron™ platform in a 1U form factor. With a maximum of 8 GB of memory, each Sun Fire X2100 M2 server with a dual-core CPU running at 2.6 GHz can deliver up to 10.4 GFLOPS of theoretical maximum computing. The Sun Fire X2200 M2 server offers two Opteron sockets in a 1U form factor. With a maximum of 64 GB of memory, each Sun Fire X2200 M2 server with two dual-core CPUs running at 2.6 GHz can deliver up to 20.8 GFLOPS of theoretical maximum computing.

Each server has four Gigabit Ethernet connections, and can have up to two SATA or SAS disks for local storage. In addition, there are two PCI Express slots for I/O expansion, allowing a wide array of clustering and storage interface configurations to meet the deployment requirements.

The server management connection to the embedded Lights Out Management (LOM) circuit is through a shared data Ethernet port, allowing flexible configurations. The embedded LOM of each Sun Fire X2100 M2/X2200 M2 server provides control of remote power on/off and status, and browser and command line interfaces through standard industry protocols, including HTTPS, IPMI 2.0, SSH, and SNMP. The embedded LOM also provides full remote keyboard, mouse, video, and storage emulation as a USB device for full feature control from a separate management node.

### Rack System Management Infrastructure

A robust management architecture has been defined to manage all of the components within the rack. All compute and switch components in the rack are connected to the management Ethernet switch through the management port. The management switch also provides uplink ports to the management network outside of the single rack.

The recommended management node within the rack is a Sun Fire X4100 server, running the Sun N1 System Manager software on the Solaris or Red Hat Linux operating system, which provides monitoring and provisioning functions for the rack. A Sun Fire X2100 M2 or X2200 M2 server can also be used as the management node. The management node is connected to the management switch for full access to the

management ports of all components. This management node can also manage additional rack resources, with connections through the management switch. The management fabric can be implemented on the same physical switch unit as the data/provisioning switch, using Virtual Local Area Network (VLAN) functionality to establish different network segments.

A Keyboard, Monitor, and Mouse (KMM) option is available to provide a rack user interface access to all nodes. The KMM is connected to the management node by default.

A terminal server in the rack is not needed, given the KVM (keyboard, video, and mouse) over IP management feature in the Sun Fire server. However, if a customer requires serial connections to all nodes as a standard deployment policy, a terminal server is available.

## Data Fabric

The standard data fabric in Sun Customer Ready HPC Clusters is a Gigabit Ethernet connection to each compute node, with options of multiple Gigabit Ethernet connections for higher bandwidth or redundancy. A separate provisioning data fabric can also be configured using the additional Gigabit Ethernet port of the compute node and a data switch in the rack.

InfiniBand (IB) is a switched data fabric that provides high performance and low latency. InfiniBand 4X Single Data Rate (SDR) connections can achieve a theoretical maximum transmission of 10 Gigabits per second. Each compute node is configured with an InfiniBand host channel adapter (HCA) connected to an InfiniBand switch in the rack. Multiple InfiniBand connections configurations are possible to achieve a wide bandwidth data connection, depending on the deployment requirements.

## Storage Interface

Various data storage architectures can be implemented with a Sun Customer Ready HPC Cluster. The standard storage interface is through a Gigabit Ethernet storage data fabric to each of the servers. A Network Attached Storage (NAS) device can be configured to this storage data fabric as required. Higher performance storage connections through Fibre Channel or InfiniBand to external storage area network (SAN) devices are accomplished through option cards in the server with additional switches.

This document does not address the storage network topology or storage component options, as storage components are not part of the initial Sun Customer Ready HPC Cluster offering.

## Operating System

The computing nodes in a Sun Customer Ready HPC Cluster support the following operating system (or later) revisions:

- Solaris 10 OS Update 2, 64-bit
- Red Hat Enterprise Linux (RHEL) 4.0 U4 64-bit
- SUSE Linux Enterprise Server 9 SP3 64-bit

The recommended operating system for the management node and the grid master node is the Solaris 10 OS.

## System Management Software

The recommended management software for Sun Customer Ready HPC Clusters is the Sun N1 System Manager software, v1.3.2, installed on the management node with either the Solaris or Red Hat Linux OS. The Sun N1 System Manager software performs bare metal hardware discovery for new nodes, provisions the OS and firmware, monitors hardware and software, and integrates with lights-out management features in Sun x64 servers.

## Grid Computing Management

The Sun N1 Grid Engine software running on the Solaris, Red Hat, or SUSE Linux operating system is the recommended grid computing software in a Sun Customer Ready HPC Cluster. The Sun N1 Grid Engine software manages and dispatches jobs onto the compute grid, with policy-based workload management and dynamic provisioning of application workloads.

Typically, a separate Sun Fire server is configured as the master node for the N1 Grid system; this server communicates with the compute nodes through the data fabric and the provisioning fabric. The master node handles task scheduling and communication between compute nodes for a grid environment. The master node can also communicate with external compute nodes through uplink ports in the switches.

## Reference Configurations

The following sections provide details on four distinct reference configurations for Sun Customer Ready HPC Clusters using Sun Fire X2200 M2 and X2100 M2 servers:

- HPC entry level reference configuration
- HPC mid level reference configuration
- HPC InfiniBand reference configuration
- HPC InfiniBand non-blocking reference configuration

### HPC Entry Level Reference Configuration

The HPC Entry Level Reference Configuration emphasizes low cost with single socket Sun Fire X2100 M2 servers. The configuration consists of a single Sun Rack 1000-42 rack

containing 35 Sun Fire X2100 M2 servers used as compute nodes, one 48-port management fabric switch, and one Gigabit Ethernet data fabric switch. This configuration also includes a Sun Fire X4100 server with a KMM for use as the management node, and a Sun Fire X2100 M2 server as the master node for the Sun N1 Grid Engine software. With an entry configuration of 1.8 GHz dual-core Opteron processors, this rack can deliver an estimated 0.25 TFLOPS of theoretical peak computation.

### Entry Level Configuration and Cabling Diagram

Figure 2 shows a high level wiring diagram of the rack for the entry level configuration. The included table details how each component is arranged in the rack, and lists the cable harness mapping to each port of the component. A Gigabit Ethernet port from the compute nodes is connected to the 48-port Gigabit Ethernet data switch, with uplinks for external connections. The same Gigabit Ethernet fabric is used for the provisioning network and storage network in this configuration.

A Sun Fire X4100 server, running the Solaris OS and the Sun N1 System Manager software, is configured as the management node. This management node is not required if an external system management node already exists. The first Sun Fire X2100 M2 server is setup as the master node for the Sun N1 Grid Engine software, providing task scheduling for HPC grid computing.




Figure 2. Rack and Cabling Mapping for HPC Entry Level Configurations.



### Site Planning Considerations for Entry Level Configurations

Table 3 summarizes the size, airflow, power, and weight of the entry level reference configuration, based on preliminary component specification, and compares them against the rack limits.

Table 3. Site Planning Considerations for Entry Level Configurations.

Description	Total Units	Size (RU)		Air Flow (CFM) @35C		Power (W)		Weight (kg, lb)			
		Unit RU	Total RU	Unit CFM	Total CFM	Unit W	Total W	Unit kg	Total kg	Unit lb	Total lb
Sun Fire X2100 M2 compute node, grid master node	36	1	36	40	1440	345	12420	10.7	383.8	23.5	844.3
Sun Fire X4100 management node	1	1	1	96	96	550	550	19.6	19.6	43.1	43.1
1U 24-port 4X IB switch	0	1	0	15	0	96	0	7.7	0.0	16.9	0.0
1U Ethernet switch	2	1	2	15	30	160	320	6.4	12.8	14.1	28.2
KMM (keyboard, monitor, mouse)	1	1	1	0	0	22	22	12.4	12.4	27.3	27.3
Cable harness (GigE), IB cables	2	0	0	0	0	0	0	7.5	15.0	16.5	33.0
Rack cabinet with MPS	1	2	2	0	0	25	25	195.0	195.0	429.0	429.0
<b>Total Reference Config System</b>			<b>42</b>		<b>1566</b>		<b>13337</b>		<b>638.6</b>		<b>1404.8</b>
Max rack limits: 1042MPS60A			42				32448		920.8		2025.7
Max rack limits: 1042MPS32A			42				39936		920.8		2025.7

## HPC Mid Level Reference Configuration

The HPC mid level reference configuration consists of a single Sun Rack 1000-42 rack containing 35 dual-socket Sun Fire X2200 M2 servers used as compute nodes, one 48-port management fabric switch, and one Gigabit Ethernet data fabric switch. This configuration also includes a Sun Fire X4100 server with a KMM for use as the management node, and a Sun Fire X2100 M2 server as the master node for the Sun N1 Grid Engine software. With a configuration of 2.2 GHz dual-core Opteron processors, this rack can deliver an estimated 0.62TFLOPS of theoretical peak computation.

### Mid Level Configuration and Cabling Diagram

Figure 3 shows a high level wiring diagram of the rack for the HPC mid level configuration. The included table details how each component is arranged in the rack, and lists the cable harness mapping to each port of the component. A Gigabit Ethernet port from the compute nodes is connected to the 48-port Gigabit Ethernet data switch, with uplinks for external connections. The same Gigabit Ethernet fabrics is use for provisioning network and storage network in this configuration.

A Sun Fire X4100 server, running the Solaris OS and the Sun N1 System Manager software, is configured as the management node. This management node is not required if an external system management node already exists. The first Sun Fire X2100 M2 server is setup as the master node for the Sun N1 Grid Engine software, providing task scheduling for HPC grid computing.




Figure 3. Rack and Cabling Mapping for HPC Mid-Level Configuration.



### Site Planning Considerations for Mid Level Configurations

Table 5 summarizes the size, airflow, power, and weight of the mid level reference configuration, based on preliminary component specifications, and compares them against the rack limits.

Table 5. Planning Considerations for HPC Mid Level Reference Configurations.

Description	Total Units	Size(RU)		Air Flow (CFM) @35C		Power (W)		Weight (kg, lb)			
		Unit RU	Total RU	Unit CFM	Total CFM	Unit W	Total W	Unit kg	Total kg	Unit lb	Total lb
Sun Fire X2200 M2 compute node, grid master node	36	1	36	47	1678	450	16200	11.2	403.2	24.6	887.0
Sun Fire X4100 management node	1	1	1	96	96	550	550	19.6	19.6	43.1	43.1
1U 24-port 4X IB switch	0	1	0	15	0	96	0	7.7	0.0	16.9	0.0
1U Ethernet switch	2	1	2	15	30	160	320	6.4	12.8	14.1	28.2
KMM (keyboard, monitor, mouse)	1	1	1	0	0	22	22	12.4	12.4	27.3	27.3
Cable harness (GigE), IB cables	2	0	0	0	0	0	0	7.5	15.0	16.5	33.0
Rack cabinet with MPS	1	2	2	0	0	25	25	195.0	195.0	429.0	429.0
<b>Total Reference Config System</b>			<b>42</b>		<b>1804</b>		<b>17117</b>		<b>658.0</b>		<b>1447.6</b>
Max rack limits: 1042MPS60A			42				32448		920.8		2025.7
Max rack limits: 1042MPS32A			42				39936		920.8		2025.7

## HPC InfiniBand Reference Configuration

The HPC InfiniBand reference configuration illustrates how a full rack of Sun Fire X2200 M2 servers can be configured with a high performance data fabric using InfiniBand technology. Given the lower latency and a wider, faster data path, the InfiniBand data fabric is expected to deliver a much higher system efficiency over a Gigabit Ethernet fabric configuration.

Thirty-four dual-socket Sun Fire X2200 M2 servers, each with a 2.6GHz dual-core CPU, are configured in a Sun Rack 1000-1042 rack as compute nodes, with an estimated 0.71 TFLOPS of theoretical peak computational performance. A Sun Fire X4100 server is connected to the management fabric and serves as the management node, and a Sun Fire X2200 M2 server is setup as the master node for the Sun N1 Grid Engine software.

### InfiniBand Configuration and Cabling Diagram

Figure 4 shows a high level wiring diagram of the rack. The included table details how each component is arranged in the rack, and lists the cable harness mapping to each port of the component. Each of the 34 compute nodes is configured with a Voltaire 4X InfiniBand HCA, wired to a Voltaire 24-port InfiniBand switch in the rack. The 4X HCA card is installed in the PCI-E slot in the Sun Fire X2200 M2 server, supporting full bandwidth for InfiniBand SDR mode. With SDR, the 4X HCA offers 10 Gigabits per second of connection bandwidth between nodes. Each of the two 24-port Voltaire InfiniBand switches provides non-blocking connections between 17 compute nodes in the rack. Two InfiniBand switches are connected together with seven InfiniBand ports, giving a bisectional bandwidth ratio of 2.4:1 within the rack between all 34 nodes.

A Gigabit Ethernet port of the compute nodes is connected to a 48-port Gigabit Ethernet switch, configured as the provisioning and the storage data fabric for the rack. The management network is on a separate 48-port switch, with dedicated uplink ports to external management connections.

A Sun Fire X2200 M2 server is setup as the master node for the Sun N1 Grid Engine software, providing task scheduling for grid computing. The grid master node is connected to the Gigabit Ethernet provisioning network, and is not on the InfiniBand network.





### Site Planning Considerations for InfiniBand Reference Configurations

Table 7 summarizes the size, airflow, power, and weight of the InfiniBand reference configuration, based on preliminary component specifications, and compares them against the rack limits.

Table 7. Site Planning Considerations for InfiniBand Reference Configurations.

Description	Total Units	Size (RU)		Air Flow (CFM) @35C		Power (W)		Weight (kg, lb)			
		Unit RU	Total RU	Unit CFM	Total CFM	Unit W	Total W	Unit kg	Total kg	Unit lb	Total lb
Sun Fire X2200 M2 compute node, grid master node	35	1	35	47	1631	450	15750	11.2	392.0	24.6	862.4
Sun Fire X4100 management node	1	1	1	96	96	550	550	19.6	19.6	43.1	43.1
1U 24-port 4X IB switch	2	1	2	15	30	96	192	7.7	15.4	16.9	33.9
1U Ethernet switch	2	1	2	15	30	160	320	6.4	12.8	14.1	28.2
KMM (keyboard, monitor, mouse)	0	1	0	0	0	22	0	12.4	0.0	27.3	0.0
Cable harness (GigE), IB cables	4	0	0	0	0	0	0	7.5	30.0	16.5	66.0
Rack cabinet with MPS	1	2	2	0	0	25	25	195.0	195.0	429.0	429.0
<b>Total Reference Config System</b>			<b>42</b>		<b>1787</b>		<b>16837</b>		<b>664.8</b>		<b>1462.6</b>
Max rack limits: 1042MPS60A			42				32448		920.8		2025.7
Max rack limits: 1042MPS32A			42				39936		920.8		2025.7


## HPC InfiniBand Non-Blocking Reference Configuration

The HPC InfiniBand non-blocking reference configuration illustrates how a full rack of Sun Fire X2200 M2 servers can be configured with a high performance InfiniBand data fabric using a single non-blocking InfiniBand switch. Given the lower latency and a wider, faster data path, the InfiniBand data fabric is expected to deliver a much higher system efficiency over a Gigabit Ethernet fabric configuration.

Twenty-eight dual-socket Sun Fire X2200 M2 servers, each with a 2.6GHz dual-core CPU, are configured in a Sun Rack 1000-1042 rack as compute servers, with an estimated 0.58 TFLOPS of theoretical peak computational performance. A Sun Fire X4100 server is connected to the management fabric and servers as the management node.

### InfiniBand Non-Blocking Configuration and Cabling Diagram

Figure 5 shows a high level wiring diagram of the rack. The included table details how each component is arranged in the rack, and lists the cable harness mapping to each port of the component. Each of the 28 compute nodes is configured with a Voltaire 4X InfiniBand HCA, wired to a Voltaire 96-port InfiniBand switch in the rack. The 4X HCA card is installed in the PCI-E slot in the Sun Fire X2200 M2 server, supporting full bandwidth for Single Data Rate (SDR) mode. With SDR, the 4X HCA offers 10 Gigabits per second of connection bandwidth between nodes. The Voltaire 9096 chassis is populated with two 24-port Voltaire InfiniBand switch modules, providing non-blocking InfiniBand connections between all InfiniBand ports. The InfiniBand switch modules are connected together with a non-blocking backplane between all 28 nodes.



Loc	Device	F	Dev Port	Data Harness Label	SW A	Dev Port	Mgmt Harness Label	SW A	
U42	X2200 M2	C	0	U-42,Rack<->SW-1,Prt-42	42	SP	U-42,Rack<->SW-2,Prt-42	42	
U41	X2200 M2	C	0	U-41,Rack<->SW-1,Prt-41	41	SP	U-41,Rack<->SW-2,Prt-41	41	
U40	X2200 M2	C	0	U-40,Rack<->SW-1,Prt-40	40	SP	U-40,Rack<->SW-2,Prt-40	40	
U39	X2200 M2	C	0	U-39,Rack<->SW-1,Prt-39	39	SP	U-39,Rack<->SW-2,Prt-39	39	
U38	X2200 M2	C	0	U-38,Rack<->SW-1,Prt-38	38	SP	U-38,Rack<->SW-2,Prt-38	38	
U37	X2200 M2	C	0	U-37,Rack<->SW-1,Prt-37	37	SP	U-37,Rack<->SW-2,Prt-37	37	
U36	X2200 M2	C	0	U-36,Rack<->SW-1,Prt-36	36	SP	U-36,Rack<->SW-2,Prt-36	36	
U35	X2200 M2	C	0	U-35,Rack<->SW-1,Prt-35	35	SP	U-35,Rack<->SW-2,Prt-35	35	
U34	X2200 M2	C	0	U-34,Rack<->SW-1,Prt-34	34	SP	U-34,Rack<->SW-2,Prt-34	34	
U33	X2200 M2	C	0	U-33,Rack<->SW-1,Prt-33	33	SP	U-33,Rack<->SW-2,Prt-33	33	
U32	X2200 M2	C	0	U-32,Rack<->SW-1,Prt-32	32	SP	U-32,Rack<->SW-2,Prt-32	32	
U31	X2200 M2	C	0	U-31,Rack<->SW-1,Prt-31	31	SP	U-31,Rack<->SW-2,Prt-31	31	
U30	X2200 M2	C	0	U-30,Rack<->SW-1,Prt-30	30	SP	U-30,Rack<->SW-2,Prt-30	30	
U29	X2200 M2	C	0	U-29,Rack<->SW-1,Prt-29	29	SP	U-29,Rack<->SW-2,Prt-29	29	
U28	X2200 M2	C	0	U-28,Rack<->SW-1,Prt-28	28	SP	U-28,Rack<->SW-2,Prt-28	28	
U27	X2200 M2	C	0	U-27,Rack<->SW-1,Prt-27	27	SP	U-27,Rack<->SW-2,Prt-27	27	
U26	X2200 M2	C	0	U-26,Rack<->SW-1,Prt-26	26	SP	U-26,Rack<->SW-2,Prt-26	26	
U25	X2200 M2	C	0	U-25,Rack<->SW-1,Prt-25	25	SP	U-25,Rack<->SW-2,Prt-25	25	
U24	X2200 M2	C	0	U-24,Rack<->SW-1,Prt-24	24	SP	U-24,Rack<->SW-2,Prt-24	24	
U23	X2200 M2	C	0	U-23,Rack<->SW-1,Prt-23	23	SP	U-23,Rack<->SW-2,Prt-23	23	
U22	X2200 M2	C	0	U-22,Rack<->SW-1,Prt-22	22	SP	U-22,Rack<->SW-2,Prt-22	22	
U21	KMM		0	U-21,Rack<->SW-1,Prt-21	21	SP	U-21,Rack<->SW-2,Prt-21	21	
U20	X2200 M2	C	0	U-20,Rack<->SW-1,Prt-20	20	SP	U-20,Rack<->SW-2,Prt-20	20	
U19	X2200 M2	C	0	U-19,Rack<->SW-1,Prt-19	19	SP	U-19,Rack<->SW-2,Prt-19	19	
U18	X2200 M2	C	0	U-18,Rack<->SW-1,Prt-18	18	SP	U-18,Rack<->SW-2,Prt-18	18	
U17	X2200 M2	C	0	U-17,Rack<->SW-1,Prt-17	17	SP	U-17,Rack<->SW-2,Prt-17	17	
U16	X2200 M2	C	0	U-16,Rack<->SW-1,Prt-16	16	SP	U-16,Rack<->SW-2,Prt-16	16	
U15	X2200 M2	C	0	U-15,Rack<->SW-1,Prt-15	15	SP	U-15,Rack<->SW-2,Prt-15	15	
U14	X2200 M2	C	0	U-14,Rack<->SW-1,Prt-14	14	SP	U-14,Rack<->SW-2,Prt-14	14	
U13	X2200 M2	M	0	U-13,Rack<->SW-1,Prt-13	13	SP	U-13,Rack<->SW-2,Prt-13	13	
U12	X4100	M	0	U-12,Rack<->SW-1,Prt-12	12	SP	U-12,Rack<->SW-2,Prt-12	12	
U11			0	U-11,Rack<->SW-1,Prt-11	11	1	U-11,Rack<->SW-2,Prt-11	11	
U10				U-10,Rack<->SW-1,Prt-10	10		U-10,Rack<->SW-2,Prt-10	10	
U09	Voltaire 96P IB	X		U-09,Rack<->SW-1,Prt-09	9		U-09,Rack<->SW-2,Prt-09	9	
U08				U-08,Rack<->SW-1,Prt-08	8		U-08,Rack<->SW-2,Prt-08	8	
U07				U-07,Rack<->SW-1,Prt-07	7		U-07,Rack<->SW-2,Prt-07	7	
U06				U-06,Rack<->SW-1,Prt-06	6	SP	U-06,Rack<->SW-2,Prt-06	6	
U05				U-05,Rack<->SW-1,Prt-05	5		U-05,Rack<->SW-2,Prt-05	5	
U04	Mgmt 48P switch	X		U-04,Rack<->SW-1,Prt-04	4		U-04,Rack<->SW-2,Prt-04	4	
U03	GigE 48P switch	X		U-03,Rack<->SW-1,Prt-03	3		U-03,Rack<->SW-2,Prt-03	3	
U02	MPS 60A	P		U-02,Rack<->SW-1,Prt-02			U-02,Rack<->SW-2,Prt-02		
U01	MPS 60A	P							

KEY: P MPS T Terminal Server  
X Data Switch M Mgmt Node  
X Mgmt Switch C Compute Node

Figure 5. Rack and Cabling Mapping for InfiniBand Non-Blocking Configurations.



### Site Planning Considerations for InfiniBand Non-Blocking Configurations

Table 9 summarizes the size, airflow, power, and weight of the InfiniBand non-blocking reference configuration, based on preliminary component specifications, and compares them against the rack limits.

Table 9. Site Planning Considerations for the HPC InfiniBand Non-Blocking Reference Configuration.

Description	Total Units	Size (RU)		Air Flow (CFM) @35C		Power (W)		Weight (kg, lb)			
		Unit RU	Total RU	Unit CFM	Total CFM	Unit W	Total W	Unit kg	Total kg	Unit lb	Total lb
Sun Fire X2200 M2 compute node, grid master node	29	1	29	47	1351	450	13050	11.2	324.8	24.6	714.6
Sun Fire X4100 management node	1	1	1	96	96	550	550	19.6	19.6	43.1	43.1
96-port IB switch + 1 RU space	1	7	7	190	190	950	950	38.0	38.0	83.6	83.6
1U Ethernet switch	2	1	2	15	30	160	320	6.4	12.8	14.1	28.2
KMM (keyboard, monitor, mouse)	1	1	1	0	0	22	22	12.4	12.4	27.3	27.3
Cable harness (GigE), IB cables	4	0	0	0	0	0	0	7.5	30.0	16.5	66.0
Rack cabinet with MPS	1	2	2	0	0	25	25	195.0	195.0	429.0	429.0
<b>Total Reference Config System</b>			<b>42</b>		<b>1947</b>		<b>14917</b>		<b>632.6</b>		<b>1391.7</b>
Max rack limits: 1042MPS60A			42				32448		920.8		2025.7
Max rack limits: 1042MPS32A			42				39936		920.8		2025.7

## Default Operating System Configuration

The following section includes information on the default configuration when using the Solaris OS. Similarly, “Red Hat Linux Operating System” on page 26 and “SUSE Linux Operating System” on page 27 contain information on the default configurations using Red Hat and SUSE Linux Operating Systems.

### Solaris Operating System

Table 10 lists the Solaris OS image settings used for the Sun Customer Ready HPC Cluster reference configurations.

Table 10. Solaris OS Image Settings.

Version	Distribution	Locale	Timezone
Solaris 10 U2 64b	SUNWXcall	en_US	EDT
Solaris 10 U3 64b	SUNWXcall	en_US	EDT

### File System Layout for Solaris OS

The default file system layout for a compute node and the Sun N1 Grid Engine master node is shown in Table 11. Free disk space can be mapped to unassigned slices by the user depending on the usage.

Table 11. File System Layout for Compute and Grid Master Nodes.

Slice	Tag	Mount	Size	Logging	Purpose
0	root	/	10 GB	Yes	/ UFS
1	swap		8 GB	NA	SWAP space
2	backup		Entire disk	NA	Entire disk
3	unassigned		0	NA	
4	unassigned		0	NA	
5	var		6 GB	Yes	Variable data
6	unassigned		0	NA	Boot kernels
7	unassigned		0	NA	
8	boot		7.84 MB	NA	Boot kernels
9	alternates		15.69 MB	NA	Use default

Table 12 shows the default file system layout for the management node running the Sun N1 System Management software. Because the Sun N1 System Management software loads the operating system and application images in the `/var` partition, this partition is configured as the remaining disk space.

Table 12. File System Layout for the Management Node.

Slice	Tag	Mount	Size	Logging	Purpose
0	root	/	10 GB	Yes	/ UFS
1	swap		8 GB	NA	SWAP space
2	backup		Entire disk	NA	Entire disk
3	unassigned		0	NA	
4	unassigned		0	NA	
5	var		Free	Yes	OS and application images for provisioning
6	unassigned		0	NA	Boot kernels
7	unassigned		0	NA	
8	boot		7.84 MB	NA	Boot kernels
9	alternates		15.69 MB	NA	Use default

## Red Hat Linux Operating System

Table 13 lists the Red Hat Linux OS image settings used for the Sun Customer Ready HPC Cluster reference configurations.

Table 13. Red Hat Linux OS Image Settings.

Version	Distribution	Locale	Timezone
RHEL AS 4 U4 64b	Red Hat	en_US	EDT

## File System Layout for Red Hat Linux OS

The default file system layout for a compute node and the Sun N1 Grid Engine master node is shown in Table 14.

Table 14. File System Layout for Compute and Grid Master Nodes.

Tag	Mount Point	Size	File System Type	Purpose
root	/	Free	Ext3	Main space and application space
swap		8 GB	swap	SWAP space
var		6 GB	Ext3	Variable data
boot	/boot	100 MB	Ext3	Boot kernels

Table 15 shows the default file system layout for the management node running the Sun N1 System Management software. Because the Sun N1 System Management

software loads the operating system and application images in the `/var` partition, this partition is configured as the remaining disk space.

Table 15. File System Layout for Management Node.

Tag	Mount Point	Size	File System Type	Purpose
root	/	10 GB	Ext3	Root and applications
swap		8 GB	swap	SWAP space
var		Free	Ext3	OS and application images for provisioning
boot	/boot	100 MB	Ext3	Boot kernels

## SUSE Linux Operating System

Table 16 lists the SUSE Linux OS image settings used for the Sun Customer Ready HPC Cluster reference configurations.

Table 16. SUSE Linux OS Image Settings.

Version	Distribution	Locale	Timezone
SUSE 9U3 64b	SUSE	en_US	EDT

## File System Layout for the SUSE Linux OS

The default file system layout for a compute node and the Sun N1 Grid Engine master node is shown in Table 17.

Table 17. File System Layout for Compute and Grid Master Nodes.

Tag	Mount Point	Size	File System Type	Purpose
root	/	Free	Reiserfs3	Main space and root
swap		8 GB	swap	SWAP space
var		6 GB	Reiserfs3	Variable data

## Default Networking and Access Configuration

The following sections include information on the default networking and access configuration, including hostname and IP address assignments.

### Default Network Assignment

The default data network in a Sun Customer Ready HPC Cluster is assigned to the 10.10.0.0/16 segment, while the management network is assigned to the 10.11.0.0/16 segment. The management node and the Sun N1 Grid Engine external connection are left unassigned.

Table 18. Default Network Assignment.

Network	Address	Netmask	Comment
External	<IP Address>	<Netmask>	For connecting to external management network.
Data	10.10.0.0	255.255.0.0	Data network
Management	10.11.0.0	255.255.0.0	Management network

## Default Hostname and IP Address Assignments

While there are numerous hostname and IP address naming conventions for data center equipment, the default scheme for the Sun Customer Ready HPC Cluster is defined for easy tracking of components. For hostnames, the components use a naming convention associated with the function, rack number, the rack unit (RU) location, and the interface:

<function>-<rack number>-<RU location>-<Interface>

The <function> describes the component's role and the <rack number> reflects the rack number in the data center configuration. Starting at the bottom of the rack, the <RU location> is the lowest rack unit location the component occupies in the rack. For example, a 4-RU component mounted in the location 15 to 18 RU will have a 15 in its assignment. The <interface> maps to a port of the component.

The IP Address uses the <rack number> and <RU location> as part of the assignment.

## Management Node Assignment

Table 19 lists the hostname and IP address information for the management node.

Table 19. Management Node Assignment.

Hostname	Address	Interface	Netmask	Comment
manage-1-07-d2	<IP Address>	<int>2	<Netmask>	For connecting to external mgmt network
manage-1-07-d0	10.10.1.7	<int>0	255.255.0.0	Data network
manage-1-07-d1	10.11.1.7	<int>1	255.255.0.0	Management network

### Sun N1 Grid Engine Master Node Assignment

Table 20 lists the hostname and IP address information for the Sun N1 Grid Engine master node.

Table 20. Sun N1 Grid Engine Master Node Assignment.

Hostname	Address	Interface	Netmask	Comment
n1ge-1-08-d2	<IP Address>	<int>2	<Netmask>	For connecting to external mgmt network, submitting jobs
n1ge-1-08-d0	10.10.1.8	<int>0	255.255.0.0	Data network
n1ge-1-08-sp	10.11.1.8	sp	255.255.0.0	Management network

### Compute Node Hostnames and IP Addresses

Table 21 lists the hostname and IP address information for the compute nodes. The data port on each compute node has this naming convention, starting at the bottom of the first rack.

Table 21. Compute Node Assignment.

Hostname	Address	Interface	Netmask	Comment
comp-1-09-d0	10.10.1.9	<int>0	255.255.0.0	Data network default
comp-1-10-d0	10.10.1.10	<int>0	255.255.0.0	Data network default
...	...			
comp-1-42-d0	10.10.1.42	<int>0	255.255.0.0	Data network default

### Service Processor Hostnames and IP Addresses

Table 22 lists the service processor hostname and IP address information. The service processor in each server uses this convention, starting at the bottom of the first rack.

Table 22. Service Processor Assignment.

Hostname	Address	Interface	Netmask	Comment
comp-1-09-sp	10.11.1.9	sp	255.255.0.0	Management network default
comp-1-10-sp	10.11.1.10	sp	255.255.0.0	Management network default
...	...			
comp-1-42-sp	10.11.1.42	sp	255.255.0.0	Management network default

### Switches and Terminal Server Hostnames and IP Addresses

Table 23 lists the hostname and IP address information for the switches and terminal server. The terminal server, network switches, and InfiniBand switches use this naming convention starting at the bottom of the first rack.

Table 23. Switches and Terminal Server Assignment.

Hostname	Address	Interface	Netmask	Comment
datasw-1-03-d	10.11.1.3	Enet port	255.255.0.0	Data switch default
managesw-1-04-m	10.11.1.4	Enet port	255.255.0.0	Management switch default
termserver-1-05-sp	10.11.1.5	sp	255.255.0.0	Terminal server default
ibsw1-1-06-sp	10.11.1.6	sp	255.255.0.0	InfiniBand switch default

### Service Processor Access

Table 24 lists the default root password for the Service Processor.

Table 24. Service Processor Access.

User Name	Password
root	changeme

### Sun N1 Grid Engine Software Defaults

Table 25 lists the default password information for the Sun N1 Grid Engine software.

Table 25. Root Password for the Sun N1 Grid Engine Software.

User	User Name	Password
root	root	admin
admin	sgeadmin	sgeadmin

### Default Values

The Sun N1 Grid Engine software is installed on the head node in the `/opt/N1GE` directory. This directory is shared out to the compute nodes using NFS.

The following default values are used to configure the Sun N1 Grid Engine software on the master and compute nodes:

```
SGE_ROOT="/opt/N1GE"
SGE_QMASTER_PORT="536"
SGE_EXECD_PORT="537"
CELL_NAME="default"
ADMIN_USER=""
QMASTER_SPOOL_DIR="/opt/N1GE/default/spool/qmaster"
EXEC_D_SPOOL_DIR="/opt/N1GE/default/spool"
```

```

GID_RANGE="20000-21000"
SPOOLING_METHOD="berkeleydb"
DB_SPOOLING_SERVER="none"
DB_SPOOLING_DIR="/opt/NlGE/default/spooldb"
ADMIN_HOST_LIST=" "
SUBMIT_HOST_LIST=" "
EXEC_HOST_LIST=" "
EXECD_SPOOL_DIR_LOCAL=
HOSTNAME_RESOLVING="true"
SHELL_NAME="ssh"
DEFAULT_DOMAIN="none"
ADMIN_MAIL="none"
ADD_TO_RC="true"
SET_FILE_PERMS="true"
RESCHEDULE_JOBS="wait"
SCHEDD_CONF="1"
SHADOW_HOST=
EXEC_HOST_LIST_RM=" "
REMOVE_RC="true"
WINDOWS_SUPPORT="false"
WIN_ADMIN_NAME="Administrator"

```

## Parts List

Table 26 lists the major components and parts used in the reference configurations. For an updated and complete list of Sun Customer Ready HPC Cluster supported parts, please refer to the Sun Web site.

Table 26. Parts List.

Description	Configuration Rules	Quantity per HPC Reference Configuration			
		Entry Level	Mid Level	InfiniBand	IB non-blocking
Sun Rack 1000-42, MPS 32 Amps (EMEA, APAC - except Japan and Taiwan)	One rack with MPS per config	1	1	1	1
Sun Rack 1000-42, MPS 60 Amps (Americas, Japan, Taiwan)	One rack with MPS per config	1	1	1	1
Sun Fire X2100 M2 server, Model 1210 (1.8GHz Dual Core, 103W)	38 max with 2 switches	36	0	0	0
Sun Fire X2200 M2 server: 2x2.2GHz, 4x1GB, No Disk/DVD	38 max with 2 switches	0	36	0	0
Sun Fire X2200 M2 server: 2x2.6GHz, 4x2GB, No Disk/DVD	38 max with 2 switches	0	0	35	29

Description	Configuration Rules	Quantity per HPC Reference Configuration			
		Entry Level	Mid Level	InfiniBand	IB non-blocking
Sun Fire X4100 server, 2 Opteron 252 CPUs, 4 GB, 73 GB SAS, DVD, 2 PSU	1 per management, 1 per Sun N1 Grid Engine	1	1	1	1
Slide rail kit for Sun Fire X2100 and X4x00 servers	1 per server	37	37	36	30
Cisco Catalyst 3750 48-port Gigabit-Ethernet switch	1 for management, 1-2 for data ports	2	2	2	2
Keyboard/Monitor/Mouse shelf unit (1RU). Fujitsu model NC14010-B452/US.	1 per rack, optional	1	1	0	1
Management Cable Harness, includes 40 orange Ethernet CAT5e cables	1 per rack	1	1	1	1
Data Cable Harness, includes 43 yellow Ethernet CAT5e cables	Up to 2 per rack	1	1	1	1
Voltaire HCA 410ex PCIe	1 per IB port	0	0	34	28
Voltaire 24-port IB Switch module	4 module max	0	0	0	2
Voltaire ISR power supply unit	2 module max	0	0	0	2
Voltaire ISR9096S-M 96-port IB Switch Chassis	1 per rack	0	0	0	1
Voltaire ISR9024S-M 24-port IB Switch	2 max per rack	0	0	2	0
IB cable, 4X	1 per IB port	0	0	41	28
Solaris 10 OS DVD Multilingual Media Kit, no hard copy doc or license included	1 per server, OS option	1	1	1	1
Sun N1 System Manager license for server with up to 4 sockets	1 per server, management option	37	37	36	30
Sun N1 Grid Engine 6 License for Maximum of 50 CPUs	1 per management cluster, HPC grid option	1	1	1	1
Sun Customer Ready HPC Cluster Ship Kit including documentation for any grid rack.	1 per rack	1	1	1	1
Voltaire sFB-4 Fabric module	4 modules max	0	0	0	4

## Ordering Information

As a starting point to configure and get budgetary references, account teams, customers, and partners can use the external online Sun Customer Ready HPC Cluster configuration tool at <http://www.sun.com/servers/sungridracksystem/configtool/>

All prices and discounts are based on US List Prices and subject to change at any time. They are provided for budgetary reference only, not as a quotation.

To get an official Customer Ready Systems (CRS) quotation using only Sun's price-listed products, contact a Sun Center representative who will use the standard WebDesk or PartnerWebDesk process. They are the first line of support for CRS quoting.

## About the Author

Jeff Lu is a Senior Staff Engineer in the Systems Group and a technical lead for the Sun Customer Ready HPC Cluster program. Jeff has over twenty years of experience in computing systems product design and development with Digital, Dell, and IBM prior to joining Sun in 2005. Jeff holds degrees in MSEE and MBA.

## Acknowledgements

The author would like to thank the Sun Customer Ready HPC Cluster team for reviewing and validating many details in the Reference Configurations. A special thanks to Mark Lamourine for laying out the default network naming convention, and Craig Snow for providing the default OS partitioning template.

## References

1. Sun Customer Ready HPC Cluster.  
<http://www.sun.com/servers/cr/hpccluster>
2. Sun Customer Ready HPC Cluster documentation.  
[http://www.sun.com/products-n-solutions/hardware/docs/Servers/grid\\_rack\\_system/grid\\_rack\\_system/index.html](http://www.sun.com/products-n-solutions/hardware/docs/Servers/grid_rack_system/grid_rack_system/index.html)
3. Documentation for Sun server products, including the Sun Fire X2100 and X2200 families of servers.  
<http://www.sun.com/products-n-solutions/hardware/docs/Servers/>
4. Sun Customer Ready Program.  
<http://www.sun.com/servers/cr>
5. Online Sun Customer Ready HPC Cluster Configuration Tool.  
<http://www.sun.com/servers/sungridracksystem/config-tool/>

## Ordering Sun Documents

The SunDocs<sup>SM</sup> program provides more than 250 manuals from Sun Microsystems, Inc. If you live in the United States, Canada, Europe, or Japan, you can purchase documentation sets or individual manuals through this program.

## Accessing Sun Documentation Online

The `docs.sun.com` web site enables you to access Sun technical documentation online. You can browse the `docs.sun.com` archive or search for a specific book title or subject. The URL is

`http://docs.sun.com/`



**Sun Microsystems, Inc.** 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN (9786) **Web** [sun.com](http://sun.com)

© 2007 Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo, Sun Fire, N1, Solaris, SunDocs, and Sun BluePrints are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. AMD, and the Opteron logo are trademarks or registered trademarks of Advanced Micro Devices, Inc. Information subject to change without notice.

Printed in USA 04/07

