SGI® Altix® 3700 Bx2 Servers and Supercomputers

Highly Scalable Linux Systems for Demanding Technical Applications

Features

- Breakthrough performance and scalability in a standards-based environment
- Large global shared memory supports demanding HPC applications with massive data sets
- Built-in NUMAlink™ interconnect enables sustained performance on HPC applications
- Flexible, high-density packaging with true expand-on-demand upgrade capabilities
- The latest, most powerful Itanium® 2 processors from Intel
- Complete Linux® solution for HPC, including tools, data management, and visualization

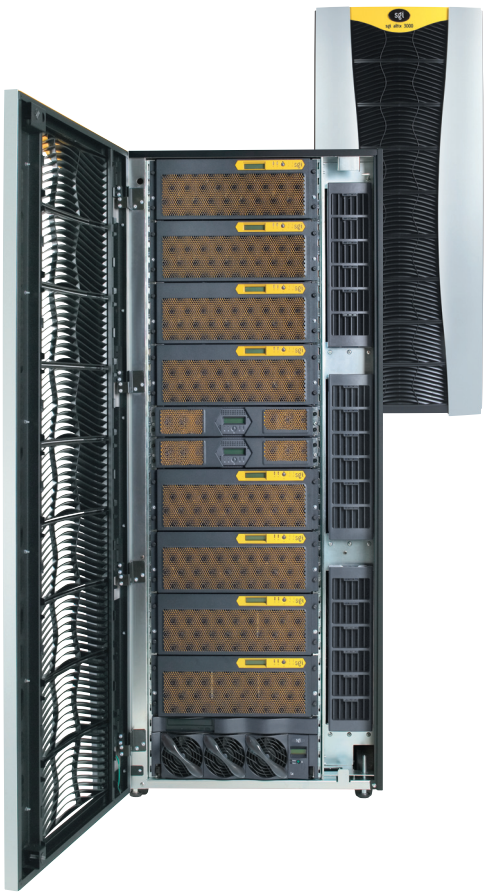

Breakthrough performance and scalability in a standards-based environment

For technical users seeking breakthrough performance with open-source computing, SGI Altix 3700 Bx2 combines the cost-effectiveness of clusters with the scalable performance and big data capabilities of a supercomputer. SGI Altix 3700 Bx2 servers and supercomputers provide significant advantages in performance and capability over traditional Linux clusters, scaling to thousands of Intel Itanium 2 processors and providing up to terabytes of shared memory even across cluster nodes.

Large global shared memory supports demanding HPC applications with big data sets

The SGI Altix 3700 Bx2 supercomputer's high-performance NUMAflex™ architecture was designed to tackle demanding HPC workloads with massive data sets. Individual nodes can scale up to 256 processors and 3TB of memory in a single system image, and SGI Altix supercomputers support shared memory even across cluster nodes. Users can take advantage of up to 24TB of global shared memory space to operate on their multi-terabyte data sets entirely in shared memory. This can provide dramatic application performance improvements, along with substantial cost savings through simplified programming and easy system administration.

Built-in NUMAlink interconnect and new I/O technology enables sustained performance on HPC applications

The high-bandwidth SGI® NUMAlink™ interconnect moves instructions and data between processors and cluster nodes up to 200 times faster than commodity cluster interconnects. Data crosses over an SGI NUMAlink switch, round-trip, in as little as 50 nanoseconds – faster than most clusters' local memory access time. The SGI NUMAlink interconnect enables balanced, sustained application performance on technical workloads. Altix Bx2 also introduces Peer I/O – direct high-speed I/O connection to Altix NUMAlink fabric - enabling breakthrough I/O performance.

Flexible, high-density packaging with expand-on-demand upgrade capabilities

Altix offers flexibility for a wide range of HPC workloads through both shared and distributed programming models. With single system image node sizes from 8 to 256 processors, along with the ability to support shared memory across cluster nodes, the configuration can be scaled for any HPC application. The unique "expand on demand" capability allows I/O and memory to be scaled independently of processors, so that users can precisely tune their system resources for specific application portfolios, and adjust them cost-effectively as their needs change. And Altix Bx2 brings new high density packaging that accommodates up to 64 processors in a rack.

Complete Linux solution for HPC, including tools, data management and visualization

SGI Altix 3700 Bx2 servers and supercomputers run a robust, industry-standard 64-bit Linux environment that is fully optimized to deliver outstanding performance on HPC applications. A complete set of development and run-time tools support superior data handling, application optimization, system administration, and resource management. SGI Altix Bx2 also works with a wide array of storage hardware and software solutions including CXFS™, SGI's heterogenous clustered filesystem. This, combined with SGI's new Linux OS-based visualization solution comprises the industry's most powerful 64 bit Linux HPC solution.

SGI® Altix® 3700 Bx2 Servers and Supercomputers

Configurations
- Processors: 16 to 2048 total; 8 to 256 per node
- Memory: 12GB to 24TB
- XIO I/O Channels: 2 to 512

NUMAflex 8-Port Router Interconnect Module (R-Brick)
- Router brick: Enables large shared-memory configurations above 64 processors, up to thousands of processors
  Note: Integrated router in compute module (CR-Brick) enables shared-memory configurations from 8 to 64 processors

Compute Module with 8P and up to 96GB Memory (CR-Brick)
- Processor options:
  - 1.3 GHz Intel Itanium 2 with 3MB on-chip cache
  - 1.5 GHz Intel Itanium 2 with 6MB on-chip cache
  - 1.5 GHz Intel Itanium 2 with 4MB on-chip cache
  - 1.6 GHz Intel Itanium 2 with 6MB on-chip cache
  - 1.6 GHz Intel Itanium 2 with 9MB on-chip cache
- Memory: up to 96GB DDR ECC
  Memory options:
  - PC2700 166GHz, 512MB DIMMs
  - PC2700 166MHz, 1GB DIMMs
  - PC 2100 133MHz, 2GB DIMMs

  Memory kits: 4GB, 8GB and 16GB
- Memory controller: 5-port crossbar per node board

Memory Expansion Module with up to 96GB memory (M-Brick)
- Memory: up to 96GB DDR ECC
  Memory kits: 4GB, 8GB and 16GB
- Memory controller: 5-port crossbar per node board

Base System I/O Module with PCI-X (IX-Brick)
- Ports:
  - 10/100/1000BaseT Ethernet port, Four RS-232/RS0422 serial ports
- Internal devices:
  - 1 system disk standard, optional redundant system disk, DVD-ROM (LDE)
- Disk interface:
  - Serial ATA
- I/O Interface:
  - 5 64-bit/133 MHz PCI-X buses with 11 available slots, 1 64-bit/66MHz PCI slot

PCI-X Expansion Module (PX-Brick; PA-Brick)
- PX-Brick:
  - Interface: 64-bit/133 MHz PCI-X buses, 3.3V and Universal 64-bit/66 MHz PCI-compatible
  - Number of buses: 6
  - Number of slots: 12 (2/bus) full length; max. 1/bus for 133 MHz PCI-X cards
- PA-Brick: Peer I/O Capability
  - Interface: 64-bit/133 MHz PCI-X buses, 3.3V and Universal 64-bit/66 MHz PCI-compatible
  - Number of buses: 4
  - Number of slots: 6 (2/bus) full length; max. 1/bus for 133 MHz PCI-X cards

JBOD Disk Expansion Module (D-Brick2)
- Interface: 66MHz/2Gb or 133MHz/2Gb Fire Channel
- Drive bays: 16 hot-plug, 3.5" slots with 110/220V power; redundant power supplies standard
- Maximum bandwidth: 200MB/sec peak
- Drive options
  - 36GB (15K RPM)
  - 73GB (15K RPM)
  - 146GB (10K RPM) drive options

External Storage Options
- HBA interfaces: 2Gb Fibre Channel, 200MB/sec peak bandwidth Ultra 320 SCSI, 320 MB/sec peak bandwidth, Gigabit Ethernet copper and optical
- JBOD: SGI® TP900 (Ultra320 SCSI) D-brick2 (2Gb Fibre Channel)
  RAID: SGI® TP9100 (2Gb Fibre Channel)
  SGI® TP9300 (2Gb Fibre Channel)
  SGI® TP9300 (Serial ATA)
  SGI® TP9500 (2Gb Fibre Channel)
  SGI® TP9500S (Serial ATA)
- Data servers:
  - SGI® NAS 2000 & 3000 (Gigabit Ethernet)
  - SGI SAN Server™ 2000 & 3000 (2Gb Fibre Channel)
  - SGI® DLM Server (NAS or SAN attach)
- Tape and libraries:
  - StorageTek® L20, L40, L80, L180, L700, L5500, 9310. ADIC® Scalar® 24, Scalar® 100, Scalar® 1000, and Scalar® 10000 Libraries. StorageTek® T9840B, T9940B, IBM® 3560, LTO, SDLT, and AIT Tape Drives

Power Bay
- Power requirements: 220-240 VAC external source
- Power distribution: 48 VDC internally distributed to all bricks

Dimensions and Weights
- System rack:
  - 75" H, 53" D, 30" W; 40U internal usable space; 1,375 lbs. max.
- I/O rack:
  - 75" H, 53" D, 30" W; 40U internal usable space; 1,225 lbs. max.
- RAID/JBOD rack:
  - 75" H, 31" D, 24" W; 38U internal usable space; 1,265 lbs. max.

Environmental (Operating)
- Temperature: +5°C to +35°C, altitude 5,000 MSL
  +5°C to +30°C, altitude 10,000 MSL
- Humidity: 10% to 90% noncondensing

Environmental (Non-operating)
- Temperature: -40°C to +60°C
- Humidity: 10% to 95% noncondensing
- Altitude: 40,000 MSL

Electrical and Power
- Voltage
  - 180-254 VAC, single phase;
  - 180-254 VAC, 3 phase (North America/Japan);
  - 360-440 VAC, 3 phase (International)

Power and Heat
- System rack (max. per rack):
  12.19 kw/41.59 kBtu
- I/O rack (max. per rack):
  2.14 kw/7.30 kBtu

Software
- System software:
  SGI Advanced Linux™ Environment with SGI ProPack™
- Also available:
  SUSE® Linux Enterprise Server 9 from Novell
- Networking:
  TCP/IP, NFS V2/V3, DHCP, SNMP management, SNMP MIB, NIS/ONC+
- Available server software:
  XFS® 64-bit journaled filesystem, XVM Volume Manager (Plexing & Snap shot options), CXFS shared filesystem, Performance Co-Pilot™ system and network monitoring, SGI® Cluster Manager for Linux, DMF, TMF
- Compilers:
  Intel® Itanium® Processor Family compilers: C/C++, Fortran;
  GNU compilers: C, Fortran 77
- Tools:
  - Libraries: MPT, Array Services, CPU sets, SCSL, FFIO, and Intel® Math Kernel Library
  - Debuggers: Etnus® TotalView®, Intel® idb, GNU gdb (with Fortran extensions)
  - Performance analysis: Intel® VTune™, Intel® Trace Analyzer, Intel® Trace Collector, SGI® Histx
  - System analysis: pfmon , Performance Co-Pilot
- Interoperability:
  Samba® environments for PC
- Partitioning:
  Support for system partitioning for systems up to 512P with maximum partition size 256P

Support and Services
SGI also offers appropriate services to implement and integrate Linux applications in your environment. For more information on available services, please see www.sgi.com/support

*Redhat Enterprise Linux Advanced Server 4 support is limited to 64 CPU cores and 128GB of memory
Redhat Enterprise Linux Advanced Server 5 support is limited to 64 CPU cores and 128GB of memory

Corporate Office
1140 E. Arques Avenue
Sunnyvale, CA 94085
(650) 960-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5195.2860
Europe +44 118.925.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

3709 [05.18.2007]

J14723