Product Guide

SGI™ Origin™ 3000 Series

Modular, High-Performance Servers

The successful deployment of today's high-performance computing solutions is often obstructed by architectural bottlenecks. To enable organizations to adapt their computing assets rapidly to new application environments, SGI™ servers have provided revolutionary architectural flexibility for more than a decade. Now, the pioneer of shared-memory parallel processing systems has delivered a radical breakthrough in flexibility, resiliency, and investment protection: the SGI Origin 3000 series of modular high-performance servers.

Design Your System to Precisely Match Your Application Requirements

SGI Origin 3000 series systems take modularity to the next level. Building on the same modular architecture of award-winning SGI™ 2000 series servers, the SGI Origin 3000 series now provides the flexibility to scale CPU and memory, storage, and I/O components independently within the system. You can design a system down to the level of individual components to meet your exact application requirements—and easily and cost effectively make changes as desired.

Unmatched Flexibility

The unique SGI NUMA system architecture provides a foundation for the most flexible, modular servers in the world. The next-generation SGI Origin 3000 series allows you to deploy, upgrade, service, expand, and redeploy every system component in every possible dimension. As a result, you can confront business changes with unlimited flexibility and astounding agility, making computing infrastructures future-proof for years to come.

A New Snap-Together Approach

This new approach to server architecture allows you to configure—and reconfigure—systems brick by brick. Upgrade CPUs selectively to keep apace of innovation. Isolate and service I/O interfaces on the fly. Pay only for the computation, data processing, visualization, or communication muscle you need. Achieve all this while seamlessly fitting systems into your IT environment with an industry-standard form factor. Brick-by-brick modularity lets you build and maintain your system optimally, with a level of flexibility that makes obsolescence almost obsolete.

SGI™ Origin™ 3200

Two to Eight CPUs

This affordable system scales from two to eight processors in a shared-memory image without routers and can be clustered to leverage the power of hundreds of CPUs managed from a single point of administration. Brick-by-brick expandability makes it an ideal base for high-availability file, Web, or storage serving.

SGI™ Origin™ 3400

4 to 32 CPUs

A price/performance powerhouse, SGI Origin 3400 has two six-port routers to accommodate scaling from 4 to 32 CPU processors in a shared-memory image. For maximum flexibility, the processor and I/O bricks can be configured within the same or separate racks, allowing for extensive I/O expansion.

SGI™ Origin™ 3800

16 to 512 CPUs

Achieve unmatched power and unlimited scalability with the SGI Origin 3800 system. With the largest single-kernel, shared-memory image available, SGI Origin 3800 scales with snap-together modularity from 16 to 512 processors. Built-in Metarouters allow clustering to tens of thousands of CPUs, making supercomputing more accessible than ever.

Introducing NUMAflex™—only from SGI.

The new SGI Origin 3000 series is the first generation of servers to deliver the benefits of NUMAflex, a system concept offered only by SGI. Based upon the award-winning SGI™ NUMA architecture, NUMAflex is a breakthrough design philosophy that rests on three solid pillars:

- Flexibility
- Resiliency
- Future-proof infrastructure

The SGI Origin 3000 series was developed in order to offer server solutions that provide these NUMAflex advantages:

- Lower cost of ownership
- Easier management and administration
- Breakthrough performance and results

NUMAflex delivers flexibility.

The flexibility that is at the heart of SGI Origin system design means that you can build exactly the right system and configuration you need for the job at hand. As your needs change, you can buy, deploy, and redeploy the bricks you need to meet them. Scaling and flexibility are woven into the fabric of the system:

- Tailor the configuration by selecting only the bricks you require
- Upgrade specific features of the server independently over time, as opposed to a forklift replacement
- Service each brick individually without affecting the entire server
- Incorporate the latest technologies as they become available—simply plug them in
- Support your diverse business needs

NUMAflex

Traditionally, scalability has charted a unilinear path. But ideally, scalability should follow the path that best meets the requirements of different applications—a path that you define.

Get the Configuration You Need

The SGI Origin 3000 series allows you to do just that, whether you need heavy-duty computation for weather simulation, massive storage for archiving bioinformatic data, high-performance I/O for media streaming, or integrated visualization of large data sets.

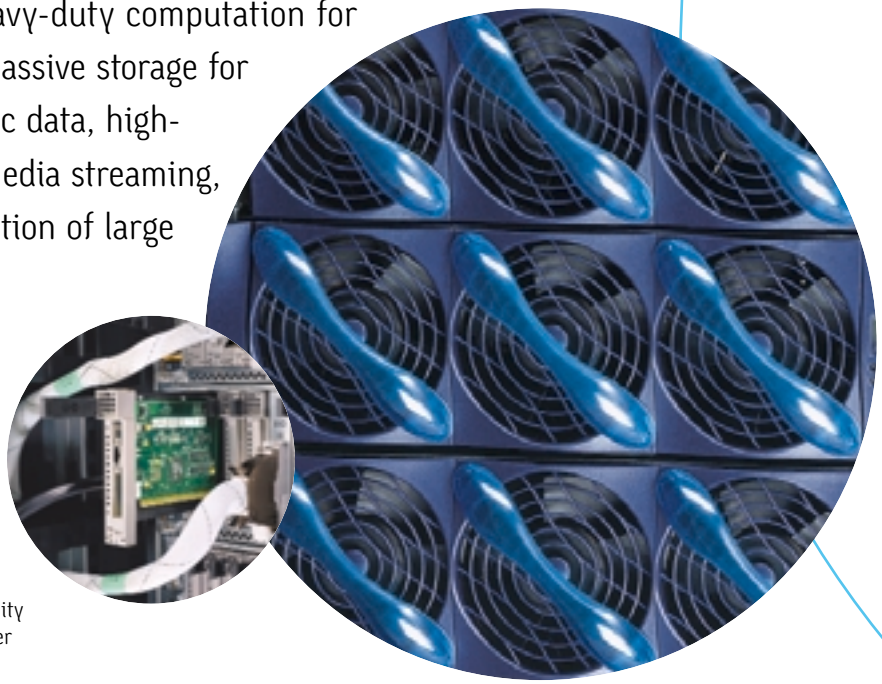

PCI Expansion
Hot-plug PCI maximizes the availability of your data. The patented SGI carrier allows easy access and upgrade or replacement.

NUMAflex

NUMAflex means resiliency.
The resiliency afforded by NUMAflex system design means that you can remove and service only the hardware components that require your attention. Without skipping a beat, the rest of the system continues to produce impressive results. NUMAflex makes it simple to partition a system for failsafe operations.

- Achieve a lower cost of ownership through cluster resiliency
- Leverage the value of hardware investments
- Design cluster solutions that maximize throughput and isolate problems
- Increase availability and uptime

Drawing on the legendary stability of SGI Origin servers, NUMAflex system design delivers the resiliency you need to keep your system up and running no matter what.

Build and Maintain Your System—One Brick at a Time

To scale system performance to meet your needs, choose the bricks your applications require:

R-brick

Router Interconnect

As the structural building block of the system, the R-brick replaces the system bus; it's a high-speed crossbar connecting processors and memory and enabling each system component to be serviced or upgraded individually. Add infrastructure as you need it—from routerless deskside systems to an eight-port router multirack configuration that delivers 512 processors in a single shared-memory environment.

C-brick

CPU Module

The basic C-brick module contains four MIPS® CPUs and local memory. A single crossbar memory controller delivers 200% greater CPU-to-memory bandwidth than previous generations. Now with four CPUs in a C-brick, the system offers a two-fold increase in CPU density, improving memory latency by up to 50% and minimizing the use of valuable floor space.

I-brick

Base I/O Module

The I-brick, standard in all systems, provides base I/O in a module and includes the system disk, CD-ROM, Ethernet, and four available PCI slots. As the system grows, customers have the option to partition the system for greater availability, using additional I-bricks as base I/O for each partition.

P-brick

PCI Expansion

For PCI expansion, a P-brick provides 12 hot-swappable PCI slots distributed over six 64-bit/66 MHz PCI buses. Total peak I/O bandwidth exceeds 3GB per second.

X-brick

XIO Expansion

For high-performance I/O expansion, X-bricks deliver four XIO slots that support HIPPI, GSN, VME, and digital video.

G-brick

Graphics Expansion

Tightly integrated InfiniteReality3 graphics add large-scale visualization capabilities for accelerated insight into complex data sets.

D-brick

Disk Storage

D-bricks provide modular JBOD mass storage for data-intensive applications. D-bricks support up to 12 drives, have dual-power supplies standard, and support drive capacities of 18GB, 36GB, and 73GB.

The SGI Origin 3000 series uses the latest expression of the revolutionary NUMA (nonuniform memory access) architecture. As today's preferred architecture for high-performance, multiprocessor systems, NUMA facilitates access to remote system resources with unmatched efficiency. It's the only way to effectively take advantage of today's high-powered CPUs. Unlike bus-based SMPs, the SGI NUMA architecture allows systems to increase shared memory to meet the growing CPU-to-memory bandwidth demands of additional processors. As a result, memory bandwidth grows proportionately as CPUs are added, making SGI NUMA systems inherently more scalable. It is the unprecedented modularity that makes the superior scalability of the SGI Origin 3000 series possible, allowing you to build large, high-performance systems brick by brick from the same components used to build an entry-level deskside unit.

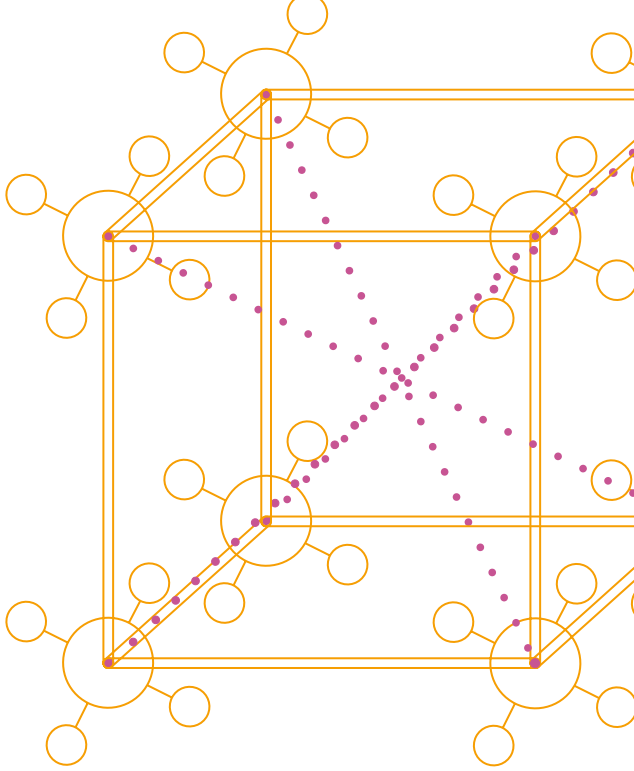

Breakthrough Architecture Enables a New Way of Computing

As the NUMA innovator, SGI has been years ahead of the industry since the introduction of the award-winning SGI Origin servers. Today, IDC calls NUMA "the architecture of the future." But the architecture of the future can be had today in a third-generation implementation at its most robust: the SGI Origin 3000 series.

Performance, Reliability, and Versatility

With their high bandwidth, superior scalability, and efficient distribution of resources, SGI servers are more than just highly modular. They are performance leaders. And now with the brick-by-brick flexibility of the SGI Origin 3000 series, performance gains come even faster—no matter how large the system gets.

- **Connectivity**. The SGI Origin 3000 series provides peak bandwidth for high-speed peripheral connectivity across all leading I/O solutions. And with bottleneck-busting speed, it supports the latest high-performance networking protocols.

- **Storage**. With ample expansion capabilities to accommodate today's ever-growing data sets, the most advanced storage technologies are supported—from 100MB-per-second Fibre Channel to 40MB-per-second Ultra SCSI. Protect your data with RAID storage solutions from SGI, offering support for RAID levels 0, 1, 0+1, 3, and 5, plus global hot sparing across all disk drives. High-end HSM solutions are available for enterprise-wide storage requirements.

- **Availability**. To provide the reliability today's applications demand, the SGI Origin 3000 series can deliver uninterrupted availability through ECC memory, redundant power and cooling, and hot-pluggable disks and PCI, along with IRIS FailSafe™ software.

- **Graphics**. For integrated graphics, the SGI Origin 3000 series supports multiple InfiniteReality3 graphics subsystems, enabling high-performance visualization of large data sets.

A Natural Extension That Protects Your Investment

As the newest generation of SGI servers, the SGI Origin 3000 series allows you to reap new flexibility and performance benefits without sacrificing what has come before. It will leverage and extend your investments in existing shared-memory applications. And since the series uses the powerful UNIX® system-based, 64-bit IRIX® 6.5 operating system, it is fully compatible with other SGI workstations and servers. With the same tools and a familiar operating system, you can integrate the SGI Origin 3000 series without any retraining. Your investment in application software and training is protected, and the availability of open systems software for the future is ensured—facilitating the ongoing development of new applications.

A Fully Supported Solution

The unprecedented flexibility of the SGI Origin 3000 series provides endless possibilities for designing custom solutions for today's competitive-edge applications. To architect and deploy solutions that capitalize on the strengths of the SGI Origin 3000 series, turn to SGI Professional Services. Industries as varied as manufacturing, media, energy, finance, and science rely on SGI consultants for complex data visualization, high-performance data mining and warehousing, e-commerce, and Internet solutions, among others. It's powerful support that leverages resources and maximizes return on investment.

A Breakthrough System for the 21st Century

With the highest level of flexibility, unlimited scalability, and stronger performance than ever before, the new SGI Origin 3000 series delivers an unequaled combination of flexibility and power to meet the changing computing requirements of any organization—today and for years to come.

NUMAflex delivers future-proof infrastructure.
SGI Origin 3000 series servers with NUMAflex mean that you can build a server solution that will be future-proof infrastructure for years to come. With ease and simplicity, you can reconfigure the server to take advantage of upgrades as needed or as they become available with technology advancements. Update only the components you want—you no longer have to replace an entire system to keep up with new technology.

With SGI Origin, you can:
- Keep up with the latest technology
- Add new I/O interfaces by simply plugging them in
- Redeploy or expand systems to meet new application demands
- Protect the value of hardware investments

NUMAflex delivers the best investment protection in the industry. It's a future-proof design approach that will help you get ahead and stay ahead.

NUMAflex

SGI Origin 3000 Series
Technical Specifications

	SGI Origin 3200	SGI Origin 3400	SGI Origin 3800
Processors	2–8	4–32	16–512
System bandwidth	11.2GB/sec	44.8GB/sec max.	716GB/sec max.
Maximum memory	16GB	64GB	1TB
Router type	None	6-port	8-port, Metarouter
Base I/O	One I-brick	One I-brick	One I-brick
Additional system I/O	—	—	One P-brick
System disk	18GB	18GB	18GB
Operating system	IRIX 6.5	IRIX 6.5	IRIX 6.5

R-brick

- 6-port — Supports shared memory system configurations up to 32 CPUs
- 8-port — Supports shared memory system configurations up to 128 CPUs
- Metarouter — Supports shared memory system configurations up to 512 CPUs

C-brick

- Processors — 4 R12000™ or R14000™* CPUs
- Memory — Up to 8GB ECC SDRAM in 4 banks
- Memory kits — 512MB, 1GB, 2GB*
- Memory controller — 5-port crossbar
- Memory bandwidth — 3.2GB/sec peak

I-brick

- Ports — 2-ports USB, 100Base-T, 1-port IEEE 1394, 1-port serial, 1-port Fibre Channel
- Internal devices — 1 system disk standard, CD-ROM drive
- Disk interface — Fibre Channel
- I/O interface — One 64-bit/66 MHz PCI bus, 2 slots; one 64-bit/33 MHz PCI bus, 3 slots

P-brick

- Interface — 64-bit/66 MHz PCI, 3.3 V, and Universal
- Number of buses — 6
- Number of slots — 12 (2/bus) full-length
- Total I/O bandwidth — 3.1GB/sec peak total

X-brick

- Interface — XIO
- Number of slots — 4
- Total I/O bandwidth — 2.4GB/sec peak

D-brick

- Interface — 66 MHz/1GB Fibre Channel, SAN aware
- Drive bays — 12 hot-plug, 3.5" power 110/220 V, redundant power supplies standard
- Maximum bandwidth — 200MB/sec
- Device capacity — 18GB, 36GB, 73GB JBOD

G-brick

- InfiniteReality3 graphics
- 1–2 graphics pipelines per G-brick
- First pipe: 1 or 2 Raster Managers
- Second pipe: 1, 2, or 4 Raster Managers
- 2–8 display channels per graphics pipeline

Processor Data

- Microprocessor — MIPS RISC R12000 at 400 MHz, R14000 at 500 MHz*
- Primary cache — R12000: 2-way set-associative 32KB instruction/32KB data cache; R14000: 2-way set-associative 32KB instruction/32KB data cache
- Secondary cache — R12000: 8MB, SDRAM; R14000: 8MB, DDR full-speed SDRAM*

Power Bay

- Power requirements — 200–240 VAC external source
- Power distribution — 48 VDC internally distributed to all bricks

PCI Adapters

- 1-port Fibre Channel optical
- 1-port Fibre Channel copper
- 1-port ATMOC3
- 1-port ATMOC12
- 1-port Gigabit Ethernet optical
- 1-port Gigabit Ethernet copper
- 2-port serial
- 2-port Ultra SCSI differential
- 2-port Ultra2 SCSI (LVD)
- 8-port digital audio

XIO Adapters

- 1-port FDDI dual attach
- 1-port HIPPI 800 serial
- Digital video
- Digital video with DVCPro
- High-definition video
- 1-port GSN (half bandwidth)
- 1-port GSN (full bandwidth)
- VME 6U
- VME 9U
- 4-port ATMOC3*

Mass Storage

- HBA interfaces — Fibre Channel, Ultra SCSI, Ultra2 SCSI
- RAID controller — Fibre Channel, 128MB cache; 2 controllers per SGI™ TP9100 module
- Internal loops — Two standard per TP9100
- Maximum capacity — 166TB JBOD, 656TB RAID
- RAID storage — TP9100 RAID rack; maximum of 9 TP9100 modules
- Device capacity — 18GB, 36GB, 73GB RAID

Dimensions and Weights

- SGI Origin 3200 — 34" H x 40" D x 24" W; 17U internal usable space; 250 lb max.
- SGI Origin 3400/3800 — 74" H x 50" D x 30" W; 39U internal usable space; 970 lb max.
- I/O rack — 74" H x 50" D x 30" W; 39U internal usable space; 1,050 lb max.
- RAID/JBOD rack — 71" H x 32" D x 24" W; 38U internal usable space; 1,265 lb max.

Environmental (Operating)

- Temperature — +5 to +35°C, altitude 5,000 MSL; +5 to +30°C, altitude 10,000 MSL
- Humidity — 10% to 90% noncondensing

Environmental (Nonoperating)

- Temperature — -20 to +60°C
- Humidity — 10% to 95% noncondensing
- Altitude — 40,000 MSL

Electrical and Power

- Voltage — 200–230 VAC, single-phase and 3-phase
- Heat/power — 4,500 W available per power bay, N+1 (6 x 750 W supplies), 15,100 BTU/hr
- Electrical service/type — NEMA 6-30, 208 VAC @ 30 amp

Software

- System software — IRIX 6.5 Advanced Server Environment, X/OPEN XPG4 BASE 95, IEEE POSIX 1003.2, and 1003.1b, 1003.1c FIPS 151-2, UNIX System 4.4, 4.3 BSD extensions, MIPS ABI, SVID issue 3, X11 R6, Motif Window Manager 1.2, IRIS GL™, OpenGL®
- Networking — TCP/IP, NFS V2/V3, RSVP, DHCP, Bulk Data Service (BDSpro), NetVisualyzer™, SNMP management, SNMP MIB, NIS/ONC+, OS bypass with Schedule Transfer (ST) protocol
- Server software — XFS™ 64-bit journaled filesystem with guaranteed rate I/O, Clustered XFS (CXFS), Networker, HPC Performance Co-Pilot™ system and network monitoring, System MIB (Provision), software distribution (Propel), Enlighten DSM
- Cluster software — MPI Toolkit, LSF, and IRIX Advanced Cluster Environment (ACE) provide centralized administration to support clustered or partitioned servers, job scheduling, accounting, load balancing of batch/interactive jobs, S/W distribution, and user, system, and network management
- Compilers — ANSI C, C++, Fortran 77 and 90, ADA, Pascal, Power C Accelerator (PCA), Power Fortran 77 and 90
- Interoperability — Samba environments for PC
- Security — Trusted IRIX™ B1 security, Commercial Security Pack (CSP)
- Partitioning — CLI interface mkpart in IRIX**

** Available Q1CY01, **Q4CY00*

Corporate Office
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America 1(800) 800-7441
Latin America 1(650) 933-4637
Europe (44) 118.925.75.00
Japan (81) 3.5488.1811
Asia Pacific (65) 771.0290

2775 [8/00]

J11686