

NUMAflex Modular Computing: New Model for Scalable Computing Infrastructure

*Analysts: Earl C. Joseph II, Ph.D., Christopher G. Willard, Ph.D.,
and Debra Goldfarb*

Background

Managers and users of high-performance computers (HPCs) require systems with the performance to address next-generation problems. They require the scalability to both manage large workloads and run large problems, the flexibility to configure systems to meet individual site requirements, the resiliency to be trusted with “bet-your-business” engineering and operations applications, and the investment protection and longevity to incorporate Moore’s Law technology improvements without sacrificing effective components, software, and personnel skills.

Computer vendors have responded to these requirements in part by developing scalable computer architectures. Although great strides in scalability have been made over the last decade, we believe that the full potential of computer scaling has yet to be fulfilled. In designing successful, scalable high-performance computers, manufacturers have to balance such factors as cost, performance, feature richness, form factors, and scalability ranges. The most successful of today’s scalable system designs started with powerful microprocessors, industry standard DRAM or SDRAM memory components, and standard I/O products. The technology for connecting these components is the number-one challenge for computer designers, and it determines the overall effectiveness of the architecture.

A major challenge has been to design systems that can grow in multiple dimensions, such as processing performance, memory and I/O, while maintaining consistent price/performance across all configurations and allowing each dimension to expand independently from the others (e.g., increasing I/O without increasing processors.) The capability to tightly integrate these HPC systems with graphics subsystems and large-scale data storage facilities has also become a requirement for many customers. In summary, the ability for users to “build-to-suit” their scalable systems to meet individual requirements for performance, memory size, I/O, resiliency, and investment protection provides a major opportunity to create a new generation of scalable systems.

This White Paper provides a working definition of what a scalable computer system needs to deliver, a brief overview of scalable system definitions, and market drivers. It presents SGI's next-generation scalable system technology and concludes with IDC's analysis of the market implications of these new technologies.

SGI's "Build-to-Suit" Modular Computing Approach

SGI is a leading provider of scalable systems in the technical market, with a history of designing cost-effective HPCs that receive strong ratings from site managers and end users. SGI is creating its next-generation scalable system by enhancing the system design modularity introduced with its Origin 2000 product into a "build-to-suit" product implementation approach, termed NUMAflex modular computing. With this approach, performance, resiliency, investment protection, and technology evolution become customizable dimensions that customers select and grow according to their needs. The technology that makes this possible is in the memory system design.

SGI is pursuing a dual product strategy with its scalable server products offering both MIPS/Irix, and in the future IA-64/Linux products based on this common highly scalable system architecture.

Scalability Needs

Scalability has become the "Holy Grail" of modern computer design; although scalability has delivered on a number of its promises, it must still meet others. IDC attributes the strong demand for current and future scalable systems to the advantages they bring to end users, including:

- **Ability to solve larger problems.** To the extent that an application can be run in parallel, scalable systems can use multiple processors to either reduce runtime for the application or increase problem size.
- **Investment protection.** Scalable systems offer a wide range of performance while maintaining binary compatibility. Users can expand their systems to match the growth in application and overall organizational requirements while preserving their investments in applications software and personnel training.
- **Capacity management.** Scalability is a tool for fine-tuning system capacity to match overall system requirements. In this case, systems are used as throughput computers, with processors added incrementally to match increases in requirements.

Copyright © 2000 IDC. **Reproduction without written permission is completely forbidden.**

External Publication of IDC Information and Data—Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

*Printed on
recycled
materials*



- **Network consolidation.** A scalable computer can be viewed as a tightly coupled server network, with the system's internal interconnect acting as the LAN. In this case, users can consolidate multiple server functions into a single system by assigning individual server functions to specific nodes. Users gain a simplified systems management, operations, and maintenance environment.
- **Applications isolation.** Scalable systems allow users to set up firewalls between applications by assigning applications to specific sets of processors and specific parts of memory. For distributed memory systems, users can assign applications to nodes; for shared memory systems, a logical partitioning of the system is necessary. This is potentially a highly valuable attribute for the emerging application server providers.
- **High availability.** Distributed memory is a base technology used for configuring high-availability systems, with the physically separate processors, memory, and I/O allowing system functions to be duplicated at a number of levels.

NUMAflex Design Strategy for SGI's Next Generation of Scalable Systems

SGI has a long history of developing high-performance scalable systems. In 1988, the company introduced its first multiprocessor systems with the Power Series, which scaled to four processors using a symmetric multiprocessing model. In 1993, SGI introduced a follow-on SMP system, the Challenge Series, which scaled up to 36 processors over the life of the product. In 1996, the company departed from standard shared memory systems architectures with its Origin product line. These systems implemented a physically distributed but logically shared memory structure that allowed servers to scale up to 256 processors. This distributed/shared memory model is usually referred to as a nonuniform memory access scheme or NUMA. IDC estimates that by the end of 1999, SGI had shipped over 23,000 Origin systems and system level upgrades and had an installed base of about 16,000 systems worldwide.

With the introduction of the SGI Origin 3000 series of servers, the company enters into a new generation of scalable technology and a third generation of NUMA architectures based on its NUMAflex system design strategy.

The NUMAflex design strategy will encompass multiple generations of SGI NUMA-based computer families, all characterized by exceptional modularity, scalability, resilience, high performance, and strong price/performance. NUMAflex systems combine high-speed interconnects and modular "bricks" to create a wide variety of configurations, not only by size but by balance of CPU and memory, I/O, graphics pipes, and storage ["independent resource scalability"]. NUMAflex families often share bricks, which can evolve at their own natural rates

to stay up-to-date [“independent resource evolvability”]. NUMAflex systems offer exceptional system longevity, flexibility, and investment protection by virtue of their modular construction around standard, upgradable system interfaces rather than monolithic construction aimed at one set of technologies.

The Soul of a New Design Strategy: NUMAflex Goals

The NUMAflex design is the blueprint for a new generation of SGI products that will be constantly renewable. Thus, the creation of this strategy provided SGI with the opportunity and challenge to develop an innovative computer architecture, which in turn will allow the company to field uniquely competitive products.

SGI architects first addressed this challenge by adding two new dimensions to the company’s overall approach to scalability:

1. **Independent resource scalability.** In general, scalable systems must be built up from functional nodes, which contain all the components of complete computer systems (i.e., processors, memory, system interconnects, storage, and various I/O interfaces). The NUMAflex strategy breaks apart the node, allowing systems to be built up from more “atomic” components. Thus, a NUMAflex system allows CPU and memory, I/O, graphics, and storage components to scale independently and in varying ratios to meet specific end-user requirements. In addition, the strategy specifies that such scaling should be done without wasting slots, floor space, power, and so forth.
2. **Independent resource evolvability.** Once a node has been broken into its atomic components, the user has the opportunity to incorporate Moore’s Law improvements in the system at this atomic level. Thus, a NUMAflex system allows CPU and memory, I/O interfaces, storage, and system interconnects to be upgraded independently. This independence allows the system to take advantage of different rates of technical improvements for classes of technology.

In addition to extending SGI’s concept of scalable computing, the NUMAflex architecture includes several other basic design goals:

- **Expand the scalability range of SGI systems.** Provide a system design that would support increasing system sizes over several generations of product families.
- **Support future technologies.** Allow the incorporation of future technologies into system architectures. This involves defining an interface between system components that allows for incorporation of technologies that are currently in the design and development phase. For example, faster I/O interfaces, such as PCI-X and

InfiniBand, and new scalable graphics modules using volume graphics processor units (GPUs).

- **Support multiple operations modes.** Allow sites to logically view their systems as operating with large shared memory in single system image mode, with partitioned shared memory, or in distributed memory clustered mode; provide a software-based mechanism to system managers to configure systems and switch between modes.
- **Balance industrial standard components with value-added technology.** Allow system designers to use both SGI-based differentiating technologies and industry-standard components as appropriate.

What Does a NUMAflex Architecture Buy You?

With the addition of resource scalability and resource evolvability, SGI is essentially driving scalable system designs to the next conceptual level. The company is working to expand the overall advantages of computer system scalability in a number of areas:

- **Increased control over systems configurations.** The ability to configure a computer from smaller components allows greater customization of individual system configurations to specific applications and operational requirements.
- **Greater cost efficiency.** From a system cost standpoint, the ability to configure systems from small components is the hardware equivalent to software unbundling — sites do not need to purchase unneeded components to acquire other required components. Money not spent on unneeded infrastructure can, of course, be spent on resources that are more important.
- **Increased investment protection.** Several features of the NUMAflex strategy contribute to investment protection:
 - The ability to configure systems from smaller basic components allows for increased life for components that are not being stressed while other components are upgraded.
 - Sites are less likely to have to sacrifice (i.e., replace) system components in one area to take advantage of technical advances in another area.
 - The ability to incorporate new technology as it becomes available should minimize the risk of lost investment by providing a mechanism to add a new technology to an existing system configuration.
- **Greater system resiliency.** By breaking a system into smaller components the “collateral damage” of a single component failure is reduced (e.g., a processor is not lost if an I/O controller goes down).

- **Enhanced operation control.** The ability to configure and manage large, single-system image configurations, partitioned systems, or clusters from a single point provides the ability to run shared memory, message passing, or OpenMP jobs at different times on the same system.
- **Increased price efficiency.** The use of low-cost standard components in systems works to reduce the overall cost.
- **Response to peak loads.** Quick upgrades and the potential for module rentals to meet temporary demand peaks are possible.

Implementing a NUMAflex System: SGI Origin 3000

SGI's Origin 3000 line of servers and the SGI Onyx 3000 line of visualization systems are the first systems fielded based on the NUMAflex design strategy. They allow users to bridge the gap between the conceptual NUMAflex strategy and real-world, physical systems. To simplify the terminology, we will use the Origin 3000 component and technology names to describe both the NUMAflex strategy and its physical implementation (see Table 1).

Table 1
New SGI Terminology

NUMAflex. Concept conveying the SGI modular system architectural strategy and its features/benefits. The NUMAflex design strategy encompasses multiple generations of NUMA-based computer families, characterized by exceptional modularity, scalability, resilience, high performance and strong price/performance. NUMAflex systems combine high-speed interconnects and modular “bricks.”

NUMAlink. Interconnect technology used to implement SGI's NUMA systems. Multiple generations of NUMAlink are planned. The SGI Origin 2000 uses NUMAlink, and the SGI Origin 3000 uses NUMAlink 3.

SGI NUMA. Non-Uniform Memory Access computer architecture replaces ccNUMA (represents the same concept).

SGI NUMA 3. Third generation of SGI-NUMA architecture based on NUMAlink 3.

Bricks. Basic functional modules (or nodes) in a NUMAflex architecture.

C-brick. CPU module.

R-brick. System router interconnect module.

D-brick. Disk storage module.

P-brick. PCI expansion module.

X-brick. XIO expansion module.

I-brick. Base I/O module.

G-brick. Graphics expansion module for Infinite Reality family graphics.

Partitioned system. Combination of hardware and software technologies that enables NUMAlink-based systems to be configured into a tightly coupled cluster of power-of-two-sized nodes. On the SGI Origin 3000, this complements the large-scale, single-system image capability. In future generations of IA-64-based server products, this technology will be extended in the development of large-scale, Linux-based systems.

Source: IDC, 2000

Functional Modules or Bricks

The Origin 3000 consists of two basic parts: functional modules, or “bricks,” and the NUMAlink interconnect that combines bricks into a single coherent system.

The NUMAflex strategy calls for breaking up nodes into atomic components. In the Origin 3000, the components are called bricks. A brick consists of the circuitry necessary to perform a specific function and an interface to the NUMAlink interconnect. The Origin 3000 can be configured with the following seven brick types:

1. **C-brick.** The compute module that supports up to four processors and up to 8GB of memory. All system memory is contained within the C-bricks. SGI offers C-bricks based on MIPS processor technology. In the future, the company will deliver C-bricks with Intel IA-64 processors.
2. **R-brick.** A crossbar switch designed to interconnect multiple C-bricks within the NUMAlink system interconnect.
3. **D-brick.** The disk storage module holds disk drives and their controllers. By separating the storage technology from the compute module, all bricks can provide equal access to all disks.
4. **I-brick.** The base I/O module contains a minimum level of storage and networking capability to boot a system. It supports PCI cards, Fibre Channel disk drives, and a specialized slot for a DVD/CD-ROM.
5. **P-brick.** The PCI expansion module houses PCI controllers and interfaces for PCI devices.
6. **X-brick.** The XIO expansion module houses XIO controllers and interfaces for XIO devices.
7. **G-brick.** The graphics expansion module holds SGI graphic engine technology. The G-brick is a second-generation (Onyx 3000) graphics subsystem that can be scaled from 1 to 16 pipes (eight G-bricks in a system). Each G-brick has one 2RM (raster manager) port and one 4RM port. The higher the number of RMs per port, the higher the performance of the pipe.

The NUMAlink Interconnect: The Mortar Between the Bricks

The NUMAlink interconnect cements the various bricks into a single coherent system. This technology is based on two SGI proprietary ASIC chips: the Bedrock memory crossbar and the router crossbar chips and high-performance cables.

The Bedrock switch is an 8-input by 6-output crossbar that acts as the memory controller between processors and memory in the system for both local and remote memory accesses. The Bedrock chip has a total aggregate peak bandwidth of 3.2GBps. It also has a channel that

connects processors to system I/O, which allows every processor in a system direct access to every I/O slot in the system. The Router chip is a 6- or 8-port crossbar ASIC found in R-bricks. The Router node channels information between all the bricks in a system. It also connects all Bedrock switches to create a single contiguous memory in the system of up to 1TB.

The crossbar switches and the cabling form the basis of an interconnect fabric that allows Origin 3000 systems to scale from 4 processors to 512 processors. Local and remote memory access times are significantly improved over previous generations of NUMA implementations from SGI. The ratio of remote to local memory access time is also improving, representing a memory performance close to what would be expected from an SMP design. For example, in a 128 SGI Origin 3800 system, the maximum memory access time is less than 525 nanoseconds, and the ratio of local to remote memory access is, at worst, 2.99:1.

What You Can Do with NUMAflex: Origin 3000 Configurations

The NUMAflex design allows a broad range of configurations that are shown in Table 2.

Table 2 SGI Origin 3000 Configurations			
Models	SGI Origin 3200	SGI Origin 3400	SGI Origin 3800
Chassis	Short rack	Single cabinet	Multiple cabinets
CPU range	2–8	4–32	16–512+
Peak GFLOPS	1.6–6.4	3.2–25.6	12.8–409.6
Memory GBs	512MB–16GB	512MB–64GB	2GB–1TB
I/O bricks	1–2	1–8	1–128
Router bricks	None	2 (6 ports each)	2–32 (8 ports each)
Operating system	Irix	Irix	Irix
Key Assumptions:			
<ul style="list-style-type: none"> • Onyx 3000 graphics systems can be added to all cabinet configurations. • Peak GFLOPS are shown in the table, although IDC views peak performance as a poor metric for comparing computers. 			
Source: IDC, 2000			

SGI Origin 3000 Models

The SGI 3000 series family features a total of six models (see Table 3). Configurations are specified as servers or visualization systems.

Table 3
Origin 3000 Models

The SGI Origin 3000 Server Family

SGI Origin 3200 server: 2–8 MIPS processors, half rack, no router

SGI Origin 3400 server: 2–32 MIPS processors, rack, 6-port router

SGI Origin 3800 server: 32–512 MIPS processors, multirack

The SGI Onyx 3000 System Family

SGI Onyx 3200 system: 2–8 MIPS processors, rack, no router

SGI Onyx 3400 system: 2–32 MIPS processors, rack, 6-port router

SGI Onyx 3800 system: 16–512 MIPS processors, multirack

Source: IDC, 2000

No System Is Complete Without Production Software: SGI Origin 3000 Software

High-performance computing requires more than just fast hardware and reliable interconnects. Today's high-performance computing environments demand the highest degree of integrated components: leading-edge hardware platforms, high-bandwidth interconnect networks, and advanced software. The ability to distribute jobs and workloads across an array of systems, ranging from workstations to datacenter servers and supercomputer systems, is a crucial part of these environments. SGI remains committed to delivering and extending this capability with its industry-leading Irix system software.

The Origin 3000 product family supports two newer software capabilities for Irix to dramatically improve system management: partitioned systems and the Advanced Cluster Environment (ACE).

Partitioned System Software

The SGI Origin 3000 series of systems has an innate ability to deliver high levels of resiliency and availability. The very flexible and modular SGI NUMA 3 architecture is designed with reliability, availability, and serviceability in mind, and customers can use it to define a server solution that closely matches their application requirements. With the introduction of SGI NUMA 3, servers from SGI can now be partitioned into separate nodes to create an application environment that emulates a "cluster in a box."

Partitioning is defined as the ability to take a single SGI NUMA system (e.g., any model of the SGI 3000 series) and divide it into a collection of smaller systems. The two primary characteristics of partitioning are:

The ability to run individual partitions, whereby each partition runs in its own protected memory space with its own operating system kernel and behaves as a distinct, standalone system. Partitions can be booted, powered up/down, and rebooted without affecting the normal operation of the other partitions in the system.

The partitions are tightly coupled, through the use of the system's interconnection network (i.e., NUMAlink), as a low-latency, high-bandwidth interconnect. A failure that causes a kernel in one partition to panic will not cause a kernel in another partition to crash.

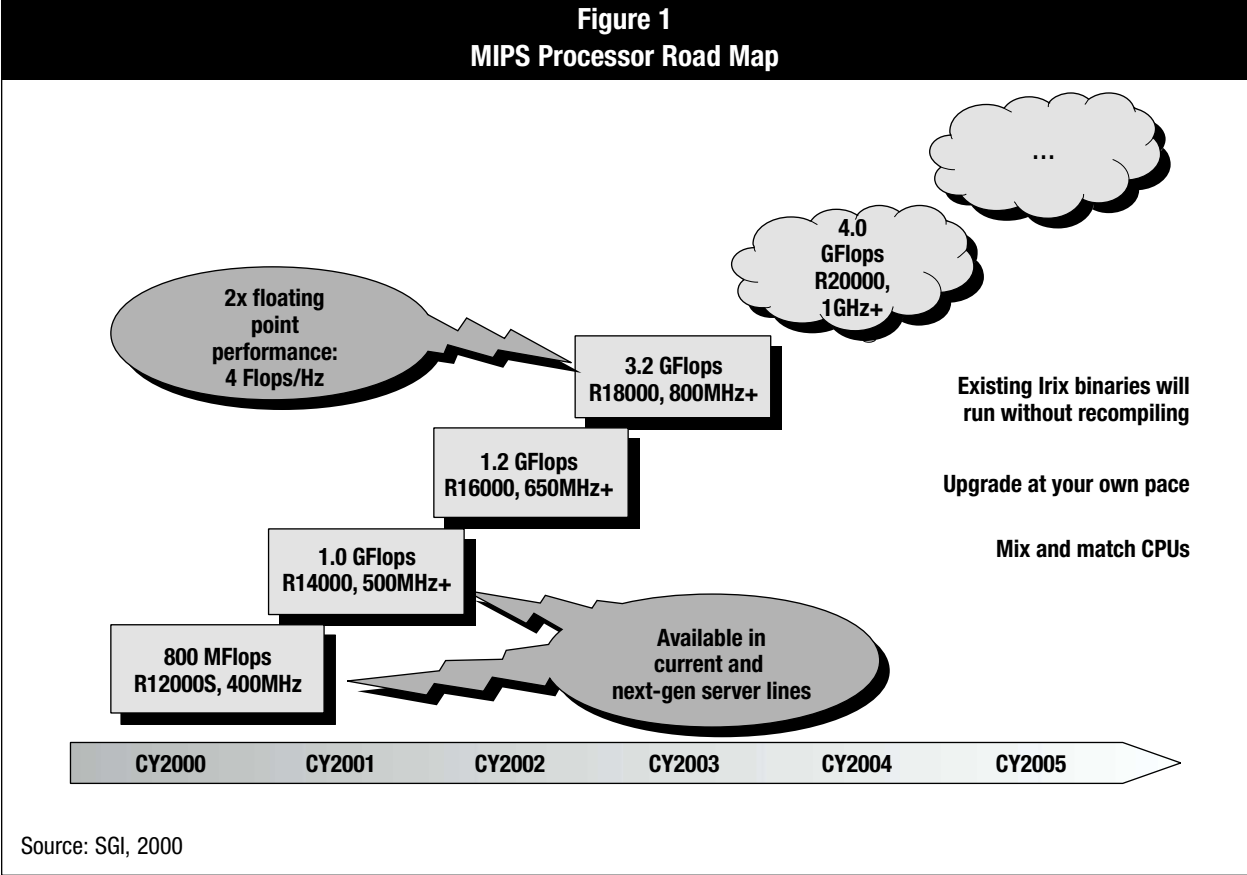
Partitioning can be thought of as a tightly coupled cluster that uses the low-latency and high-bandwidth NUMAlink interconnect instead of a low performance networking interface for the interconnect. The resulting higher level of performance is directly related to the use of NUMAlink to deliver information between partitions.

Irix Advanced Cluster Environment (ACE)

Some application environments benefit from a workflow that involves many small shared-memory servers. The ability to manage multiple servers in a simple but effective manner is the key to enhanced productivity. The Irix ACE line of products is specifically designed to help system managers improve work efficiency by simplifying the effort to manage many servers, both as individual nodes in a cluster or as multiple partitions in a larger shared-memory server. Used in cluster accounts for years, the tools in ACE for Irix 6.5 are fully supported for production environments and deliver a single administrative image and single system view of a cluster or a multipartitioned shared memory system.

What's Next for Origin?

As Figure 1 shows, the MIPS processor road map includes at least five generations of performance enhancements starting with the R12000, which is rated at 0.8 GFLOPS and is expected to grow to over 4 GFLOPS by the year 2004. SGI plans to provide five speedups over the next four years, with each speedup adding approximately 100MHz. SGI is sending a clear message to customers that it intends to keep investing in MIPS RISC processors for many years to come. SGI also announced that it will keep its Irix operating system alive and well over the same period. This should make the SGI Origin 3000 an easy and safe upgrade for current Origin customers.



SGI plans to launch a sister version of the Origin 3000 that uses Intel IA-64 processors and runs the Linux operating system. With NUMAflex, which provides the base for future IA-64 products, the company has addressed the scaling limitations of the Linux operating system by extending the system partitioning software approach to provide a single system view spanning large-scale systems.

Challenges

IDC believes that SGI realizes the need to keep delivering compatible products to their large and loyal MIPS/Irix-based customer base. SGI is directly addressing this requirement with its Origin 3000 product strategy and will need to continue to convince customers that MIPS is here to stay.

The company also needs to communicate clearly to customers how, when, and if they should migrate to IA-64 and Linux. And for customers that decide to migrate, SGI needs to make the path to Linux as painless and as productive as possible.

Conclusion and IDC Opinion

IDC believes that SGI is taking a leadership position in addressing both customer needs and design challenges by enhancing the system

modularity introduced with the Origin 2000 product with NUMAflex-based systems, via a “build-to-suit” product implementation approach. We see this as the first instance of a new architectural class of systems that we expect will gain broad acceptance within technical markets as well as the overall server marketplace over the next few years.

Additionally, we see SGI’s multiple operating system and processor road map as a fundamental cornerstone in its overall strategy of supporting the future requirements of the technical and creative users. How Linux will evolve and which workloads it will support remain uncertain. SGI believes, however, that its NUMAflex platform, coupled with its vision regarding the open source model, will position the company strongly to support both Unix and emerging Linux workloads.

We believe that NUMAflex and its current implementation in the form of the SGI Origin 3000 product line should strongly position SGI to regain customer mindshare and sales. The company should see strong acceptance of this product in its core technical markets as well as in the markets that service creative users. Moreover, we see this as a potentially strong product to support emerging Internet workloads given its flexibility, scalability, and modular attributes.

NORTH AMERICA

Corporate Headquarters

5 Speen Street
Framingham, MA 01701
508-872-8200

IDC Canada

36 Toronto Street, Suite 950
Toronto, Ontario
Canada M5C2C5
416-369-0033

IDC Irvine

18831 Von Karmen Ave, Ste 200
Irvine, CA 92612
949-250-1960

IDC Mtn. View

2131 Landings Drive
Mountain View, CA 94043
650-691-0500

IDC New Jersey

120 Wood Ave South, Suite 509
Iselin, NJ 08830
732-632-9222

IDC New York

2 Park Avenue
Suite 1505
New York, NY 10016
212-726-0900

IDC Texas

100 Congress Ave, Suite 2000
Austin, TX 78701
512-469-6333

IDC Washington

8304 Professional Hill Drive
Fairfax, VA 22031
703-280-5161

ASIA/PACIFIC

IDC Asia/Pacific (Hong Kong)

12/Floor, St. John's Building
33 Garden Road
Central, Hong Kong
852-2530-3831

IDC Asia/Pacific (Singapore)

71 Bencoolen Street, #02-01
Singapore 189643
65-226-0330

IDC Australia

Level 4, 76 Berry Street
North Sydney
NSW 2060, Australia
61-2-9922-5300

IDC China

Rm.610-612 Time Square
No. 88 Xi Chang'an Avenue
Xicheng District
Beijing 100031
China PRC
86-10-6833-1179

IDC India Ltd.

Cyber House
35 (4 Bays)
Echelon Institutional
Sector 32 Gurgaon 122002
Haryana, India
91-124-381673

IDC Japan

10F The Itoyama Tower
3-7-18, Mita Minato-ku
Tokyo 108-0073, Japan
81-3-5440-3400

IDC Korea Ltd

Suite 704, Korea Trade Center
159-1, Samsung-Dong, Kangnam-Ku
Seoul, Korea 135-729
82-2-55-14380

IDC Malaysia

Suite 13-03, Level 13, Wisma KiaPeng
No. 3, Jalan Kia Peng
50450 Kuala Lumpur, Malaysia
6-03-2163-3715

IDC New Zealand

Level 7, 246 Queen Street
Auckland, New Zealand
64-9-309-8252

IDC Philippines

7F, SEDCCO 1Bldg
Rada Street Corner
Legaspi Street
Legaspi Village
Makati City, Philippines
632-750-9477

IDC Taiwan Ltd.

10F, 31
Jen-Ai Rd, Sec 4,
Taipei 106, Taiwan, R.O.C.
886-2-2731-7288

IDC Thailand

27 Soi Charoen Nakorn 14
Charoen Nakorn Road, Klongtongnai
Klongsan Bangkok 10600, Thailand
66-2-439-4591-2

IDC Vietnam

37 Ton Duc Thang Street
Unit 1606
District-1 Hochiminh City Vietnam
84-8-910-1235

EUROPE, MIDDLE EAST, AND AFRICA

IDC Austria

c/o Loisel, Spiel, Zach Consulting
Mayerhofgasse 6
A-1040 Vienna, Austria
43-1-50-50-900

IDC Benelux (Belgium)

29 Avenue Louis Gribaumont
B-1150, Brussels, Belgium
32-2-779-46-04

IDC Benelux (The Netherlands)

A. Fokkerweg 1
1059 CM Amsterdam
The Netherlands
31-20-669-2721

IDC Central Europe (ECE)

Male Namesti 13
Praha 1 110 00, Czech Republic
420-2-2142-3140

IDC Central Europe (Germany)

Westerbachstr. 23A
61476 Kronberg/Ts., Germany
49-6173-7098-0

IDC Central Europe (Switzerland)

WTC, Leutschenbachstrasse 95
CH-8050 Zürich, Switzerland
41-1-308-3619

IDC Egypt

39 Iraq Street
Mohandesseen, Cairo, Egypt
20-2-336-7355

IDC France

Immeuble La Fayette
2, Place des Vosges, Cedex 65
92051 Paris la Defense 5, France
33-14-904-8000

IDC Hungary

Bajcsy-Zsilinszky út. 57
Building 3, Rooms 103-104
H-1065 Budapest, Hungary
36-1-153-0555/ext. 165, 166

IDC Israel

4 Gershon Street
Tel Aviv 67017, Israel
91-124-381-673

IDC Italy

Viale Monza, 14
20127 Milano, Italy
390-2-284-571

IDC Nigeria

House 2, 'C' Close
403 Road, 4th Avenue
New Extension, Festac Town
Lagos, Nigeria
234-1-883585

IDC Nordic (Denmark)

Jagtvej 169B
DK-2100 Copenhagen, Denmark
45-39-162222

IDC Nordic (Finland)

Jarrumiehenkatu 2
FIN-00520
Helsinki, Finland
358-9-8770-466
972-3-561-1660

IDC Nordic (Sweden)

Box 1096 Kistagängen 21
S-164 25 Kista, Sweden
46-8-751-0415

IDC Poland/ProMarket

Wrobla 43
02-736 Warsaw, Poland
48-22-644-4105

IDC Portugal

Av. Antonio Serpa, 36 Piso 9
1050-027 Lisbon Portugal
351-1-796-5487

IDC Russia

c/o PX Post, RDS 186
Ulitsa Zorge 10
Moscow 125525
Russian Federation
7-501-929-9959

IDC South Africa

c/o BMI-TechKnowledge
3rd Floor, 356 Rivonia Blvd.
PO Box 4603, Rivonia, 2128
South Africa
27-11-803-6412

IDC Turkey

Tevfik Erdonmez Sok. 2/1 Gul Apt.
Kat 9D; 46 Esentepe
Istanbul, Turkey
90-212-275-0995

IDC U.K.

6 Dukes Gate, Acton Lane
Chiswick, London W44AE
United Kingdom
44-181-987-7100

IDC U.K.

2 Bath Road
Chiswick, London W4 1LN
United Kingdom
44-181-987-7100

LATIN AMERICA

IDC Miami

Latin America Headquarters
5301 Blue Lagoon Drive, Suite 490
Miami, FL 33126
305-267-2616

IDC Argentina

Trends Consulting
Rivadavia 413, 4th Floor, Suite 6
C1002AAC, Buenos Aires, Argentina
54-11-4343-8899

IDC Brasil

Alameda Ribeirão Preto, 130 cj 41
01331-000 São Paulo
SP Brazil
55-11-253-7869

International Data Corp. Chile

Luis Thayer Ojeda 166 Piso 12
Providencia, Santiago 9, Chile
56-2-231-0111

IDC Colombia

Carrera 40 # 103-78
Bogota, Colombia
571-533-2326

IDC Mexico

Select - IDC
Av. Nuevo Leon No. 54 Desp. 501
Col. Hipodromo, Condesa
C.P. 06100 Mexico, D.F.
52-5-256-1426

IDC Venezuela

Trends Consultores
Av. Francisco de Miranda
Centro Perú, Torre A, Piso 9
Of. 91, Chacao 1060
Caracas, Venezuela
58-2-261-0352

IDC delivers accurate, relevant, and high-impact data and insight on information technology to help organizations make sound business and technology decisions. IDC forecasts worldwide IT markets and adoption and technology trends, and analyzes IT products and vendors, using a combination of rigorous primary research and in-depth competitive analysis. IDC is committed to providing global research with local content through more than 500 analysts in more than 40 countries worldwide. IDC's customers comprise the world's leading IT suppliers, IT organizations, and the financial community. Additional information on IDC can be found on its Web site at <http://www.idc.com>.

IDC is a division of IDG, the world's leading IT media, research, and exposition company.