InfiniBand Architecture

Transforming I/O Technology

An IDC White Paper

Analysts: Jean S. Bozman and Vernon Turner
January 2001

InfiniBand Technology Overview

What is InfiniBand Architecture?

InfiniBand architecture is a new I/O specification that will start to ship in a wide variety of server and interconnect products during 2001. InfiniBand architecture is a well-described technology with a final specification that was prepared by the InfiniBand Trade Association (IBTA) in the fall of 2000. More than 215 member companies participate in the IBTA and will work together to see that InfiniBand architecture is established as an industry standard.

InfiniBand architecture's point-to-point linking technology will be used as the basis for an I/O "fabric" that will increase the aggregate data rate between servers and storage devices. This capability will allow server vendors to replace outmoded system buses and to multiply total I/O traffic by a factor of two to four times or more compared with conventional system buses.

More important, the I/O fabric of InfiniBand architecture will take on a role similar to that of the traditional, mainframe-based "channel architecture," which used point-to-point cabling technology to maximize overall I/O throughput by handling multiple I/O streams simultaneously. That means that I/O need no longer be a bottleneck in improving overall data throughput for server systems.

The Internet Will Drive InfiniBand Architecture Adoption

IDC believes that the Internet's N-tier architectures are being driven by the demands of Internet-driven data processing. By arranging many small servers in Web server farms or in clusters, IT managers have the ability to add more servers as needed. This type of "pay-as-you-grow" build-out of computing infrastructure is necessary because user demand for Internet data services ebbs and peaks unpredictably.

One of the most significant obstacles to overcome in the "N-tier" architectures is that of interoperability created by multiple proprietary architectures. By working from a universal standard, InfiniBand raises

IDC
Analyze the Future

5 Speen Street • Framingham, MA 01701 • Phone (508) 872-8200 • Fax (508) 935-4015

the potential of seamless connectivity across the server landscape, giving rise to improved service levels, product differentiation, and consistent platform management.

InfiniBand technology was designed to meet the emerging requirements of the Internet age in the following ways:

- It acknowledges that there are very fast processors capable of 1GHz or more, which will push more data than most conventional system buses can handle. The result is a technology "mismatch" in which the fast processor has no way to push multiplying data packets to these nearby growing storage resources.
- It will accommodate Internet data types, such as streaming video, multimedia, and high-resolution graphics, and will drive more data traffic than traditional, business-oriented data processing.
- InfiniBand links, combined into an I/O fabric, will help to alleviate potential bottlenecks inside servers.
- InfiniBand architecture will improve link speeds between servers and storage, including storage area networks (SANs).
- InfiniBand architecture can be used as an interconnect technology that will link individual servers into clusters for purposes of high availability, scalability, and improved manageability.

How InfiniBand Technology Works

InfiniBand technology works by connecting host-channel adapters (HCAs) to target channel adapters (TCAs). The HCAs tend to be located near the servers' CPUs and memory, while the TCAs tend to be located near the systems' storage and peripherals. A "switch" is located between the HCAs and the TCAs, directing data "packets" to the correct TCA destination based on information that is bundled into the data packets themselves (see Figure 1).

The glue between the HCA and TCA is the InfiniBand switch, which allows the links to create a uniform fabric environment. One of the key points of this switch is that it will allow packets of information (or data) to be managed based on variables, such as service level agreements and a destination identifier.

InfiniBand architecture has the potential to play a role in transforming or "morphing" computer architectures. Where once there were general-purpose servers, including CPU memory and storage, there can be densely packed application-specific servers or "blades" that can be changed at will. At the same time, storage can be separated from the

Figure 1
The InfiniBand Architecture Model

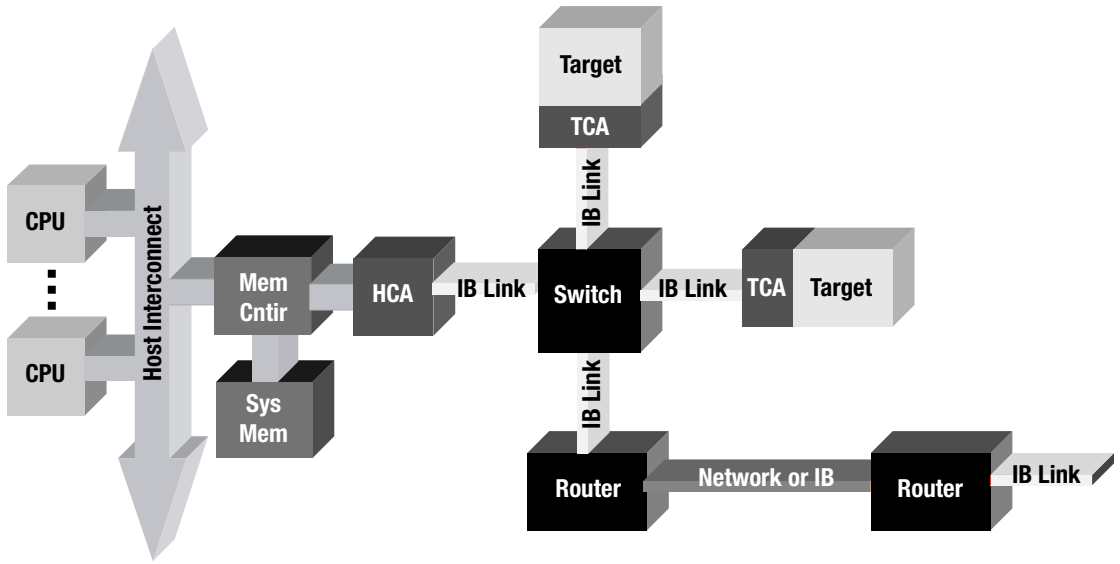

Source: InfiniBand Trade Association

CPU/memory components of these application-specific servers, and extremely large storage resources can be managed centrally by IT staffers.

This transformation can be expected to start in server "farms" that are used as front-end tier 1 and tier 2 servers that handle Web serving and application serving (see Figure 2).

Server vendors will be able to "pack" such servers closer together by moving the I/O subsystems and storage and peripherals to a location "outside" the traditional server box. This will promote high-density server configurations in IT sites and in service provider (SP) data centers.

The result could well be an evolution from today's highly dense rack-optimized servers into a highly dense array of highly modular server appliances or server blades. These blades will be designed for a particular function, or solution, such as caching, email, security, or Web serving. If a blade breaks or fails, it can be rapidly replaced, and workloads can be routed around it until it is repaired.

Flexible configurations of tier 1 and tier 2 servers, linked together (through an InfiniBand fabric) by InfiniBand links, will promote the reallocation of computing resources in fast-growing, Internet computing infrastructure. Further, InfiniBand links can be adapted by using three different link speeds. 1X, 4X, and 12X links are planned for shipment by system and interconnect vendors.

Clustering of servers in all three tiers — tier 1's Web servers, tier 2's application servers, and tier 3's database servers — can be enhanced by InfiniBand links. The InfiniBand architecture I/O fabric, made up of

Figure 2
N-Tier Architecture

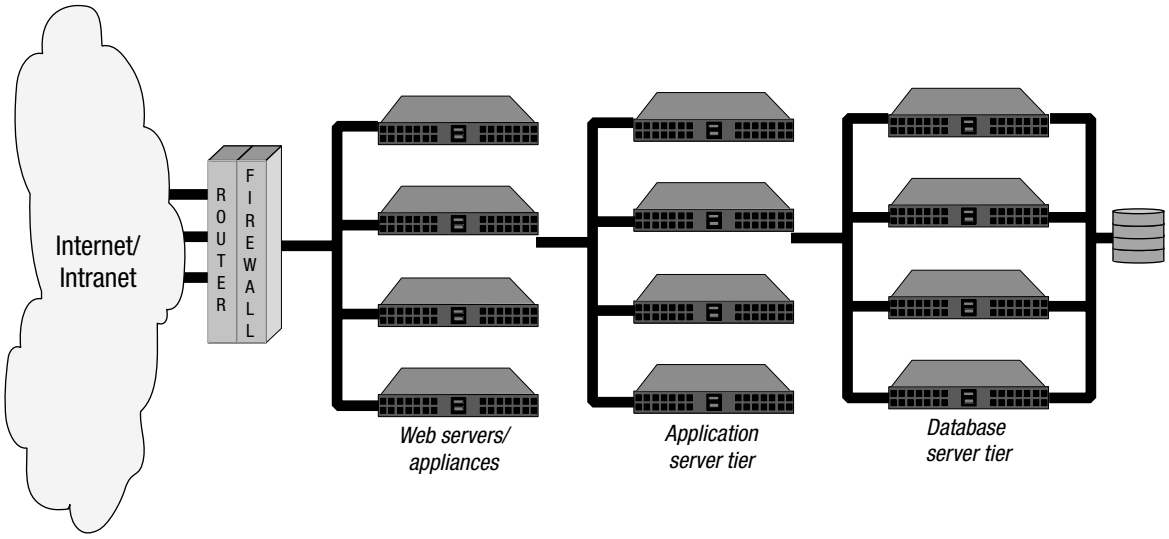

Source: IDC and Dell, 2001

multiple 1X, 4X, or 12X strands, can be used as a new high-speed interconnect in clustered server configurations. These clusters can be deployed in several ways: They can support high-availability failover, scalability of workloads, or improved manageability of multiple servers from a single console.

Finally, scalable server systems will likely use InfiniBand links "inside" the server system itself. However, the design cycle for midrange and high-end server systems is several years. Very likely, InfiniBand architecture adoption will be widespread in entry servers, in clusters, and in connections to SANs by the time that InfiniBand architecture is a major building block for midrange symmetrical multiprocessing (SMP) and cc:NUMA server systems.

InfiniBand Architecture Value Proposition

Why Consider Another Interconnect?

IDC believes InfiniBand fabrics will coexist with a wide variety of existing I/O standards that are already widely deployed in user sites. Often, multiple architectures do and can exist together and help make each other better. Existing architectures include PCI, Ethernet, and Fibre Channel. IDC also expects InfiniBand fabrics to coexist with follow-on I/O standards, including PCI-X, Gbit Ethernet, and 10Gbit Ethernet. Yet, InfiniBand architecture offers a new approach to I/O efficiency by replacing the traditional system bus with an I/O fabric that supports parallel data transfers along multiple I/O links.

However, the aspects of InfiniBand architecture that will likely cause migration away from these well-established I/O standards are the following: InfiniBand architecture offloads CPU cycles for I/O processing, it delivers faster memory pipes and higher aggregate data-transfer rates, and the InfiniBand fabric will reduce management overhead for the server system.

Unmet Needs that InfiniBand Architecture Addresses

In 2000, the server market saw the explosive introduction of appliance servers and rack-optimized servers into the entry server market (servers costing less than $100,000). Many of the general-purpose servers and pedestal-based servers were being placed into racks in every new Internet data center. Earlier we discussed this phenomenon as part of the modularization of the data center. However, as we migrate to smaller and denser servers (1U form factor or smaller), real estate issues on the server motherboard, heat and power dissipation, and consumption combined with sheer connectivity may inhibit the potential of the blade technology. InfiniBand fabrics have the potential to address all of these needs in a simple manner.

InfiniBand Architecture's Performance Characteristics

InfiniBand architecture's role will depend largely on the workload that its server, storage, and network are attempting to support. For example, if the server is simply a front-end device, a Web hosting server, or file and print appliance, then the 1X link should be sufficient to provide both throughput and connectivity. Most of the Internet tier 1 server platforms would be met by this link.

The 4X link may manifest itself in the tier 2 environment. Here servers are likely to be mainly 2-way and 4-way processors that need to run larger workloads where some latency issues can be tolerated. However, many of these workloads begin to take on mission-critical, service-level characteristics.

Finally, the 12X link could be used for transaction processing and database management, where almost no latency is tolerated as features such as two-phase commits are made to keep the data up-to-date. Given these workload characteristics, the 12X links would appear in the higher end of the tier 2 and tier 3 server architectures.

In addition, IDC believes that any data-intensive or heavy transaction processing that is performed outside the bounds of a local data center environment would be ideal candidates for the 12X links. The additional bandwidth and I/O aggregation could be handled by this top-end switch.

Cost Implications

The impact on the cost of interconnect technology may be significant. As with any technology that becomes a standard and ultimately obtains

a ubiquitous role in the market place, InfiniBand architecture probably will drive down the costs of server interconnect, bridges, and switches. IDC believes that vendors that can balance this decline with added value in their products (such as higher-availability solutions, software, and management stacks) will maintain or increase their current margins.

However, it is clear by the number of new and emerging InfiniBand companies that the overall cost of the server technology will continue to decline. IDC can cite previous examples within the server industry where technology disruptions created better value for the consumer (e.g., mainframe technology moving from water-cooled bipolar technology to air-cooled CMOS chipsets).

Market Status and Potential: An Adoption Plan

IDC's InfiniBand Architecture Server Forecast

IDC has attempted to size the server market opportunity for InfiniBand architecture (see Table 1). In doing so, it evaluated the greatest segments of the server market where InfiniBand can be more readily deployed. It has identified that the entry server market (which accounts for more than 90% of all server unit shipments) as being most acceptable to InfiniBand technology. It further constrained this forecast by advising the reader that IDC's appliance server forecast is not integrated into this existing forecast. (Note: IDC forecasts that in 2004, the appliance server market may exceed 2 million additional units, or approximately 22% of the server market).

Table 1
Worldwide InfiniBand-Capable Servers Forecast (2000–2004)

Price Band	2000	2001	2002	2003	2004
< $10K	0	39,744	706,753	1,651,369	3,972,127
$10–24.9K	0	4,362	129,848	290,189	668,548
$25–49.9K	0	493	23,545	72,144	162,015
$50–99.9K	0	82	8,588	28,850	68,914
Total	0	44,681	868,734	2,042,553	4,871,604

Key Assumptions:

Appliance server market is not included in the analysis.

IA32, IA64, RISC, and CISC chip architectures are included.

Messages in the Data:

The greatest InfiniBand architecture opportunity lies in the <$10K server market.

The greatest InfiniBand architecture opportunity lies with IA32-based servers.

Source: IDC, 2001

InfiniBand Architecture and the General-Purpose Entry Server Market

The InfiniBand architecture market is based on workload requirements running on matching server configurations. IDC's InfiniBand server forecast was developed using a constrained model and, as a result, may be viewed as being conservative. IDC used the entry server market as InfiniBand's starting ground because it is made up of more than 90% of all server units. In doing so, the forecast was able to identify the server market capable of supporting InfiniBand architecture with minimal disruption to existing infrastructure because these servers have a high turnover and are widely accepted in new environments where legacy infrastructures do not exist.

InfiniBand Architecture and the Appliance Server Market

Today's all-in-one entry servers are being complemented by a new category of appliance servers that typically support specific workloads and that typically arrive on the customer site with preconfigured, preinstalled, system and application software. This development promotes rapid deployment of workload-specific servers. At the same time, it answers the industry shortage of highly skilled IT professionals and system administrators who are capable of reconfiguring general-purpose servers. IDC research shows that the appliance server market is growing from $740 million in 1999 to more than $11 billion in 2004.

These appliance servers typically cost little, have very specific workload functions, and, therefore, performance requirements are easily measured. In addition, the appliance servers are being readily accepted in new data center environments, where existing fabric infrastructures have not been established. If those fabrics are not in place, then InfiniBand architecture adoption will take place more slowly. Once those fabrics are installed, however, the IT site or SP data center will be able to achieve the following benefits: faster reallocation of server resources, simplified clustering, and better links between servers and storage systems.

IDC believes that workloads will determine the type of server technology to be used. In doing so, the workload environment could determine the success of the InfiniBand fabric. For example, a complex, highly scaled, rack-optimized server environment within a traditional enterprise or service provider offers a great opportunity for InfiniBand fabric management. Here, InfiniBand's fabric management tools will encourage the use of 4X or 12X InfiniBand fabrics in arrays of computers that are being linked in order to tackle ecommerce applications that "tap" the compute power of many individual servers. User requests, for example, could enter the server array via the tier 1 servers, which are cooperating through the use of workload-balancing software and then be passed on to a cluster of servers running a particular application or database. Thus, user requests are navigating through the multiple tiers of computers connected via InfiniBand fabric, which function as a single computing resource from the end users' perspective.

InfiniBand Architecture and the RISC Server Market

IDC believes that vendors of RISC servers, including Unix/RISC servers, the IBM AS/400, and RISC-based proprietary midrange servers, will adopt architecture components more slowly than vendors of entry-based systems. One reason for this is that RISC systems typically use ASIC-controlled I/O subsystems, in which single-vendor technology has been "burned into" silicon. There would be little advantage of using InfiniBand architecture to completely replace these ASIC-based I/O subsystems, which generally perform well and were developed with large R&D investments by the vendors themselves.

However, over time, today's RISC-based systems will likely incorporate InfiniBand technology in some way. They may incorporate "bridge" chip technology to tie InfiniBand architecture-based I/O subsystems to the current RISC server I/O interconnects. In some cases, the RISC servers may be evolving to incorporate IA64 microprocessors, such as those in the SGI Origin series and in the HP N-Class and Superdome series. In those cases, InfiniBand architecture will be used as an I/O subsystem of choice to support I/O from IA64 motherboards to outboard storage resources.

RISC-based entry servers will likely support InfiniBand architecture before midrange and high-end RISC servers support it. Such entry servers are already becoming a staple of SP data centers, known for their high performance and combat-proven RAS characteristics. IDC believes InfiniBand links could link multiple RISC-based entry servers together into clusters or into Web server farms.

InfiniBand Technology and Productization Roadmap

IDC believes that InfiniBand architecture will be phased in over a period of several years. However, its adoption will likely be somewhat more rapid in SP environments than in IT environments, which have considerable investments in legacy computing architectures, including traditional forms of I/O.

IDC believes InfiniBand architecture will likely coexist with a wide variety of existing I/O standards that are already widely deployed in user sites. Among these are PCI, Ethernet, and Fibre Channel. IDC also expects InfiniBand architecture to coexist with follow-on I/O standards, including PCI-X, Gbit Ethernet, and 10Gbit Ethernet. Yet, the InfiniBand architecture offers a new approach to I/O efficiency by replacing the traditional system bus with an I/O fabric that supports parallel data transfers along multiple I/O links.

IT "Conservatism" Will Phase in InfiniBand Adoption

IT managers are, by nature, conservative about making sudden changes in their production computing environments. PCI, for example, is widely installed, having been shipped in production servers since the mid-1990s. For that reason, it is likely that PCI-X is likely to do well

as a replacement/upgrade from PCI. Fibre Channel has a large installed base in the storage arena, in SANs, in clustering of data center servers, and in the WAN-connected, geographically dispersed data centers.

In IT organizations, InfiniBand architecture will be introduced into an environment that is already rich with a variety of I/O technologies. Introduction of InfiniBand architecture will likely take place in those systems that will benefit most — arrays of front-end Web servers and clusters of servers linked to massive amounts of external storage. These elements of the computing infrastructure are the ones that will show the most improvement from the use of a switched fabric for I/O that supports multiple I/O links, working in parallel, to shuttle large packets of data between server engines and system peripherals.

With the development of InfiniBand-Fibre Channel and InfiniBand-Gbit Ethernet routers, switches, and adapters, InfiniBand architecture-based servers can be deployed alongside existing servers and plugged into the existing infrastructure.

Service Provider Adoption is Expected to Accelerate

InfiniBand architecture is likely to have a strong appeal for clustering, for service providers' N-tier networks and for a new generation of Internet-enabled systems for IT data centers. The SP play is a strong one, both for reasons of SPs' openness to modular computer design (dense packing) and their heavy use of clustering (Web clustering or application clustering, or both).

Service provider sites are best positioned to adopt InfiniBand architecture quickly because they have little, if any, legacy computing architecture. These sites are growing rapidly, adding dozens of servers each fiscal quarter. A rapid improvement in I/O between servers and storage resources can only help SPs provide more efficient data services to their end customers.

Bridge Technologies: The First Wave of InfiniBand Architecture Adoption

IDC research has shown that a number of vendors are prepared to ship such bridge technologies, whether at the HCA, TCA, or switch layers of the InfiniBand architecture. This means that IT managers and SP technologists will have a variety of choices in terms of the ways they can tie in InfiniBand-enabled servers to the already-installed servers that were shipped with PCI or Ethernet I/O links, for example.

The advantage of a bridge I/O technology is that it can be immediately adopted on a variety of already-installed server machines. That is, servers with PCI slots or those with fast Ethernet links are already widely used and can become InfiniBand-enabled as soon as the bridging chips ship to customers.

Further, the "bridge chip" technology allows systems vendors to sell servers with several options for I/O ports. This allows vendors to protect their investments in previous I/O technologies, and it allows customers to move to InfiniBand architecture when they are ready to do so.

Clustering: The Next Frontier

Clustering is the next frontier for InfiniBand architecture adoption because InfiniBand links can be "retro-fitted" onto existing server systems to create an industry-standard cluster. InfiniBand architecture could provide high-speed links between multiple servers, boosting total aggregate data bandwidth and minimizing the time for failover of applications and databases from one server to another.

I/O has long been a gating factor in achieving seamless clustering capability. IDC research shows that a wide variety of interconnect links are used to achieve clustering today: Ethernet (10Mbps), fast Ethernet (100MBps), 1Gbit Ethernet, SCI, and Fibre Channel among them. The introduction of InfiniBand fabrics will provide multiple "strands" of I/O between servers. This should help to boost throughput and to reduce latency between the clustered servers.

However, InfiniBand technology will likely compete with two other fast interconnect technologies — PCI-X and 10Gbit Ethernet. It is likely that several interconnect standards will remain in the marketplace, rather than that one will dominate others. Server environment tends to add new servers and new technologies, rather than to discard the older technologies. Thus, the product "mix" evolves over time, with newer server platforms gradually replacing the older ones.

Finally, clustering technology involves a mix of hardware, software, and interconnects. All three elements must be in place in order to forge strong links between individual server nodes. Thus, the need for InfiniBand-enabled software cannot be overlooked: The time to adoption with respect to clustering can be minimized through the shipment of IB-enabled components and IB-enabled software.

Scalable Servers: The Final Frontier

The final stage of the adoption curve, at which InfiniBand architecture has been incorporated into all sizes of servers, is adoption of InfiniBand architecture in future scalable servers and cc:NUMA server designs. The long planning and design cycles inherent in such advanced computer architectures dictates a slow adoption cycle. Thus, SGI has said that it plans to incorporate InfiniBand architecture into the next generation of Origin 3000s that are expected to ship sometime in the next two years.

InfiniBand links are for I/O only — they will not replace internal cache-coherent interconnects between multiple processor boards. Nonetheless, it will take some time for individual server vendors to

phase out the embedded, proprietary I/O links in scalable midrange servers and in high-end servers, in favor of industry-standard InfiniBand technology. Systems vendors have invested many years of time and money in developing these I/O links, and it will likely take some time to redesign systems to the brand-new InfiniBand specification, finalized in the fall of 2000.

The User Equation: IT Sites and SP Data Centers

IDC has long believed that customers are more concerned about an application's performance, or its workload characteristics, than they are about the hardware it is running on. This has been seen clearly in the explosion of rack-optimized servers configured for specific tasks during 2000. Workloads such as email have moved from their role. For example, email, once a communications vehicle that linked PCs scattered across a corporation, has evolved into a mission-critical application running on a single-functioned server, or array of servers. Thus, the characteristics of performance, scalability, and recovery all play increasingly important roles for IT management.

Today's infrastructure models embrace hundreds or even thousands of servers in a multi-tiered computing infrastructure, requiring an extremely dense computing power and cabling and enhanced remote management of these server assets. Further, traditional IT organizations have developed sophisticated remote management and recovery practices that are now being adopted by the new, large SPs, as they "build out" their Internet computer infrastructure. Attempts to replicate this model have been somewhat bumpy, something that is borne out by operational problems with fast-growing computer systems and high-risk server configurations that tend to fail in spectacularly noticeable ways that draw media attention. IDC notes that InfiniBand architecture, with its potential to simplify the server configuration while increasing its reliability, availability, and serviceability (RAS), provides the opportunity for IT managers to significantly lower IT management costs and to improve service levels to end customers.

Challenges

There are significant challenges to InfiniBand architecture from both existing server vendors and their own customers. Given that there have been several I/O standards already under way, many vendors may not have the financial ability or the research capacity to invest in another I/O fabric. Instead, they may want to or need to harvest the investment already made. Likewise, many end users may not want to risk installing another I/O fabric within their data centers.

In addition, the introduction of a "standard" implies that interoperability issues take a back seat. This may be true only after significant cooperation between many of the standard's participants occurs. Historically, the interoperability wild card has inhibited adoption rates. Everyone from the hardware and software vendors to the data center

operations staff have struggled to make a seamless processing from disjointed offerings.

Finally, InfiniBand architecture may have the ability to level the playing field across different operating systems and therefore challenge server vendors to provide additional value through application software and middleware. As a result, InfiniBand architecture may be perceived as a threat to weaker server vendors that may not have strong software portfolios.

All of these challenges are justified, but InfiniBand architecture can be used to help vendors provide their customers with the best I/O fabric choices that their workloads are demanding. InfiniBand architecture should be viewed as complementary to any of the standards and evolutionary to new server farms of the future.

Summary

In the years since the desktop personal computer was elevated to the PC server platform, nearly every element of the server infrastructure has been upgraded and improved. However, the I/O bus that provides connectivity and bandwidth between the server and the storage devices has not kept pace with the speed-up in nearly all of the components surrounding it — the microprocessor, memory, and storage.

Because bandwidth, storage, and network capacities are going to continue to put demands on the server to deliver more content faster and in a sustained manner, IDC believes that InfiniBand architecture is an enabling technology. It will complement the existing I/O delivery vehicles, and it will allow them to better support the growing demands of Internet-computing workloads. IDC believes Internet data requirements will continue to grow as long as the economic model of improved speed and lower costs continue. It is imperative for server vendors to establish a server platform that meets this challenge.

The advantages of a point-to-point, high-speed, switch fabric interconnect architecture are seen every day with existing mainframe-class servers. IDC believes that applying this principle of an I/O fabric that can replace the aging system-bus technology will allow growth opportunities for all computer-technology vendors.

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
Room 611, Beijing Times Square,
88 West Chang'an Avenue, Beijing,
P.R. China, 100031
86-10-8391-3456

IDC (India) Limited
Cyber House
B-35, Sector 32 - Institutional
Gurgaon - 122002, Haryana, India
91-124-6381673 to 80

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
632-894-4808

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, Klongtonsai
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)
Nibelungenplatz 3, 11th Floor
60318 Frankfurt, Germany
49-69-90502-0

IDC Central Europe (Switzerland)
Niederlassung Züerich
WTC, Leutschenbachstrasse 95
CH - 8050 Züerich
Switzerland
41-1-307-1000

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

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

IDC Nordic (Sweden)
Box 1096 Kistagången 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
351-21-796-5487

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 Spain
Ochandiano, 6
Centro Empresarial El Plantio
28023 Madrid
34-91-7080007

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

IDC U.K.
British Standards House
389 Chiswick High Road
London W4 4AE
United Kingdom
44-20-8987-7100

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
Calle Guaicapuro
Edif. Torre Seguros Alianza
Piso 6, Ofc. 6-D, El Rosal
Caracas 1060, Venezuela
58-2-951-3270

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.

IDC
5 Speen Street • Framingham, MA 01701
(508) 872-8200 • Fax (508) 935-4015 • www.idc.com

Sponsored by IBTA
01-072SYSTEM2858