IDC
Analyze the Future

WHITE PAPER

Improving HPC Performance, Scalability, and Availability with Blades: Introducing SGI Altix ICE

Sponsored by: SGI

Earl Joseph, Ph.D. Steve Conway
Jie Wu
June 2007

IDC OPINION

The worldwide market for high-performance computing (HPC) servers has been growing rapidly since 2002. This growth has been driven almost entirely by the emergence of clustered servers ("clusters") that perform well on most but not all technical computing problems and offer substantial price/performance advantages over traditional server architectures. As HPC end users have sought to employ larger cluster installations for scale-up and scale-out workloads, they have run into important barriers. SGI's recently introduced Altix ICE blade-based server line has been explicitly designed to address these issues, especially for HPC end users with larger scale-out workloads. The new Altix ICE product joins the Altix 4700 series to broaden SGI's bladed server product portfolio.

SGI's Altix ICE servers combine cost-effective commodity technologies (e.g., InfiniBand, Xeon processors, and Linux) with SGI value-added technologies to provide HPC-focused differentiation for buyers. The SGI Altix ICE bladed servers are designed specifically to address major "pain points" that are widely experienced by cluster users today, including system manageability, performance scalability, power, cooling, and facilities issues. Aside from providing fine-grained system management capabilities at four distinct levels, this new SGI product has also been designed to address the "memory wall" that limits the performance of many technical applications by optimizing I/O performance using a cost-effective, industry-standard InfiniBand interconnect network. All of this differentiation is highly beneficial in medium- and larger-scale deployments, which is SGI's sweet spot for SGI Altix ICE. Many end users have followed a scale-out strategy, and now that they have hit the limits, they are looking for a scale-up clustered solution that will allow them to address larger and more complex problems.

SGI's strong reputation among industrial HPC users, and the company's demonstrated ability to make large-scale systems work well for government, industrial, and university customers, bodes well for the SGI Altix ICE product's market prospects. HPC buyers' historical penchant for selecting products and services based on high-performance features and related attributes will work in SGI's favor, as will the company's recent refocusing on its established strengths in technical computing. SGI's introduction of the Altix ICE server line is an affirmation of the company's ongoing commitment to serving the HPC market. SGI has disclosed a strong customer interest for the new bladed servers.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

IN THIS WHITE PAPER

This white paper provides an overview of the growing market for HPC clusters and of SGI's new Altix ICE bladed server product that is aimed at this market. It reviews the challenges in acquiring and using clusters today, what end users would like to see in future HPC clusters, and how bladed solutions can meet these needs. The paper also describes the SGI product and assesses the challenges and opportunities SGI faces in the fast-growing cluster market.

SITUATION OVERVIEW

HPC Market Landscape and the Impact of Clusters

The worldwide HPC technical server market has seen very strong growth over the past four years, averaging over 20% yearly growth, and is now a $10 billion a year market. Within this market there has been a major disruption due to the impact of commodity clusters, which have gone from less than 12% of the HPC technical server market to close to 60% within the same four-year period. The strongest growth has been in the middle and lower end of the market, systems that sell for under $1 million. Table 1 shows the market's strong growth and how quickly cluster-based solutions are displacing other types of servers.

TABLE 1

Worldwide HPC Technical Server and Technical Cluster Revenue by Segment, 2002–2006 ($M)

	2002	2003	2004	2005	2006	2002–2006 CAGR (%)
All HPC servers	4,961	5,698	7,393	9,208	10,030	19.2
Clusters	523	1,352	2,425	4,127	5,631	81.1
Cluster ratio (%)	10.5	23.7	32.8	44.8	56.1	

Source: IDC, 2007

Table 2 shows the mix of server and cluster revenue by industry/application area. The largest revenue segments are bio-life sciences, university/academic, defense, and computer-aided engineering (CAE). Clusters made up 46% of the revenue total in 2005. The cluster share of the revenue mix in 2005 varied substantially by segment. Clusters were strongest in digital content creation (DCC) and distribution (64% revenue share) — a segment in which SGI has had considerable historical success — and geosciences (61%). Technical clusters have not penetrated as deeply in the following segments: university/academic (45%), CAE (44%), bio-life sciences (44%), and defense (31%).

CAE is heavily used in the automotive and aerospace industries, both of which have mission-critical, bandwidth-intensive problems that do not scale as well and have not performed adequately on existing cluster offerings. These problems could benefit from improved clusters (bladed or rackmounted) that are designed for better performance on modestly scaling codes. In addition, many of these sites have reached the limits of scaled-out clusters and would like solutions that allow them to scale up better.

TABLE 2

Worldwide Technical Server and Technical Cluster Revenue by Industry/Application Segment, 2005 ($M)

	Technical Servers	Technical Clusters	Cluster Ratio (%)
Bio-life sciences	1,433	630	43.9
CAE	1,108	492	44.4
Chemical engineering	223	125	56.4
DCC and distribution	514	328	63.8
Economics/financial	255	133	52.1
EDA	649	386	59.5
Geosciences	490	300	61.3
Design and drafting	156	37	23.5
Defense	811	254	31.3
Government lab	1,376	606	44.0
Software engineering	20	2	12.0
Technical management	101	26	25.2
University/academic	1,700	773	45.5
Weather	359	165	45.9
Other	3	0	0.0
Total	9,199	4,257	46.3

Source: IDC, 2007

Issues with Clusters Today

At the same time that the market is growing, clusters are still a challenge for most users, especially in midsize and larger configurations. Table 3 shows the top challenges in implementing clusters, as described by end users/buyers:

1. Facility issues in power, cooling, and floor space have gone from a minor nuisance four years ago to the top concern for customers. Part of the reason is that clusters provide a vastly larger number of processors for a given budget, but they result in increased operation expense to power and cool the installed systems in the datacenter.

2. The second major issue cited by customers is managing the cluster system. The tools and system management features available have improved greatly, but not at the rate required by end users. Some sites limit the maximum system size to under 128 processors due to the complexity of system management. As a result, datacenter managers find themselves with separate systems to manage and operate.

3. The third major issue for end users regards the ability to scale user jobs or applications. This applies to sites that write their own applications and even more strongly to sites that use independent software vendor (ISV) applications. More than half of all HPC user jobs today run on a single node due to the complexity of scaling their applications, and many still run on a single processor. The difficulty of creating and implementing parallel algorithms limits the scaling of many user applications, as does interconnect performance. One solution is to follow a scale-up approach to their clusters, but to easily do so requires better products from vendors.

4. The fourth most frequently cited challenge is the unavailability of some applications or, in some cases, their lack of maturity for production technical computing.

5. Storage and data management issues are both large and increasing in importance for customers. Users are generating very large numbers of files, and more users are using clusters, which is causing most sites to have serious storage problems. In addition, users want their data immediately, without having to spend time searching for their files. In many industries, data retention requirements are growing in importance, boosting long-term data storage and data management requirements.

TABLE 3

Top 3 Challenges to Implementing Clusters

Q. What are the most significant challenges in implementing clusters in your organization? (Top 3 choices)

	Number of Mentions	Share (%)
Facility issues — power, cooling	47	16.3
System management capability	35	12.2
Complexity of implementing parallel algorithms	33	11.5
Application availability/lack of maturity of the solution	23	8.0
Complexity of system purchase and deployment	22	7.6
I/O performance	20	6.9
Interconnect latency	19	6.6
Facility issues — space, density	17	5.9
Supported data storage mechanisms (databases, parallel file systems, etc.)	17	5.9
Interconnect bandwidth	14	4.9
Third-party software costs	14	4.9
Interconnect complexity/cable, cards, switches	11	3.8
Facility issues — noise	6	2.1
Other	10	3.5
Total	288	100.0

Source: IDC, 2007

User Requirements for Clusters

The previous section explored the challenges users face with clusters. This section looks at the customer requirements for making truly high-performance clusters available for end-user applications (see Table 4).

1. The most often-cited user requirement for clusters is "total aggregate system performance," a catch-all category that is ultimately related to the cluster's performance characteristics when fully (or at least very heavily) utilized. This category most directly reflects how much value users/buyers are getting from their clusters.

2. I/O performance, that is, the ability of the full cluster system to transfer data internally, is the second most noted user requirement. This requirement is directly related to the "memory wall" issue described earlier. The greater the internal data movement, the busier the data-hungry processors can be kept and the more efficiently (fully) the cluster can be utilized.

3. Processor performance is the third most frequently cited requirement for clusters. Processor performance has grown as an issue in the era of multicore processors, in part because the multiple cores must now share the bandwidth that previously was available to each single-core processor and in part because processor speeds are sometimes deliberately tuned down to reduce power consumption.

4. The final user requirement worth highlighting is the cluster's ability to run specific applications — presumably meaning, in this case, to run them reasonably well. As noted earlier, many applications historically have not run well on clusters, especially bandwidth-intensive codes. As the number of cores per processor escalates through this decade and beyond, this problem will intensify, reducing the breadth of applicability of clusters (rackmounted or bladed) that do not incorporate technologies to alleviate the issues of shrinking bandwidth and ballooning parallelism.

TABLE 4

Top 3 System Performance Requirements for Clusters

Q. What are your top 3 system performance requirements for clusters?

	Number of Mentions	Share (%)
Total aggregate system performance	55	19.2
Aggregate I/O performance	28	9.8
Processor performance	27	9.4
Run specific applications	23	8.0
Node-level memory size	22	7.7
High availability	22	7.7
Node-level memory bandwidth	16	5.6
Compatibility with other installed systems	15	5.2
Total system interconnect bandwidth	14	4.9
Individual node performance	13	4.5
Low power and cooling requirements	13	4.5
Node-to-node latency	11	3.8

TABLE 4

Top 3 System Performance Requirements for Clusters

Q. What are your top 3 system performance requirements for clusters?

	Number of Mentions	Share (%)
Space requirements power/footprint	11	3.8
Total system interconnect latency	9	3.1
Node-to-node bandwidth	7	2.4
Total	286	100.0

Source: IDC, 2007

Application Parallelization Limitations

HPC users have many applications today that do not scale well on clusters (see Table 5), and they are eagerly looking for ways to increase the scale of their application job sizes. There are several major reasons for this lack of scalability. First and foremost is the "memory wall" (i.e., the large and increasing gap between processor speeds and the speeds of the commodity interconnect networks that transfer data between and among the cluster's processors and its physically distributed memory).

The memory wall especially limits the scalability of bandwidth-intensive applications, such as structural analysis codes widely used in the automotive and aerospace sectors. A more systemic barrier to scalability is the erosion, within the past decade, of the business model for ISVs to adapt their application software for use on highly scalable HPC hardware systems. This business barrier is discussed in detail in a comprehensive study done by IDC for the Council on Competitiveness and DARPA, *Study of ISVs Serving the High Performance Computing Market* (downloadable from **www.compete.org**).

SGI historically has offered HPC systems that alleviate the memory wall issue by providing higher-bandwidth (albeit more expensive) SMP architectures with custom SGI NUMAlink interconnect networks. SGI's use of a commodity InfiniBand network reduces the cost of the new SGI Altix ICE bladed servers, while the product's compute density and fine-grained, four-level management hierarchy boost its performance and manageability. (SGI plans to continue offering its Altix 4700 bladed servers with NUMAlink for even more bandwidth-hungry applications.) With SGI Altix ICE's value-added differentiation, the product is well positioned for larger-scale deployments, with which SGI has had considerable experience and success.

SGI has also maintained particularly strong relationships with ISVs that serve the HPC market, which should bode well for the company's continued collaborations aimed at scaling up applications for larger installations.

TABLE 5

Application Parallelization

Q. What percentage of the applications on your cluster are:

	Average (%)	Number of Responses
Running on a single processor	16.8	64
Running on multiple processors within a node	39.7	78
Running on multiple nodes	45.4	76
Running on the entire cluster	14.7	51

Source: IDC, 2007

Integrated Blade Systems for HPC Solutions

Blades are just starting to gain a major position in the HPC market space. In the past, blades were often seen as lagging in technology (i.e., not having the latest processors) and more expensive. Today, these gaps have closed substantially. Table 6 shows the reasons for clustering with blades versus traditional system architectures for HPC solutions, and Table 7 illustrates the advantages of blade-based solutions.

Among HPC customers across all technical application segments, IDC research found substantial representation on both sides of the racks-versus-blades choice (the split was about 50/50). As Tables 6 and 7 show, users see greater density/smaller footprint as the largest attraction of blades. In sum, if interconnect performance is important for respondents' workloads, they tend to steer toward rackmounted clusters (see Table 8); if interconnect performance is less important, blades become considerably more attractive. This implies that certain industry segments and application domains (e.g., oil and gas, bio-life sciences) will be more receptive to blades than other industry segments.

TABLE 6

Reasons for Selecting Blade-Based Systems

Q. Why do you think blades are a good fit for your organization?

	Number of Mentions	Share (%)
Space/footprint/density	14	35.0
Market factors/vendor offerings	6	15.0
Price/performance	5	12.5
Cost/total cost of ownership	5	12.5
Cooling	4	10.0
Power	2	5.0
Other reasons	4	10.0
Total	40	100.0

Source: IDC, 2007

TABLE 7

Advantages of Blade-Based Clusters

Q. What do you see as the advantages of blade-based technical clusters?

	Number of Mentions	% of Respondents
Density	37	67.3
Cost	20	36.4
Interconnect/network cabling	19	34.5
Management	15	27.3
Performance	9	16.4
RAS	8	14.5
Other	8	14.5

Note: Multiple responses were allowed.

Source: IDC, 2007

Blade Limitations and Concerns

The reasons why some HPC users are not yet moving to blades are shown in Table 8. The primary barrier, as mentioned earlier, is interconnect bandwidth and latency (i.e., the system's ability to alleviate the "memory wall" that limits the scalability of many applications today). But users are unwilling to pay a major cost premium for higher-performance interconnects. Hence, vendors that can exploit commodity interconnect technologies (InfiniBand, Gigabit Ethernet) within x86-based architectures designed for performance differentiation are more likely to strike the winning balance between these opposing user requirements.

TABLE 8

Factors in Not Using Blades

Q. If you're not using blades, what has kept you from using them?

	Number of Mentions	% of Respondents
Interconnect performance/availability	16	39.0
Node performance	9	22.0
Cost premium	9	22.0
Cooling concerns	7	17.1
Processor support	1	2.4
Other	19	46.3

Note: Multiple responses were allowed.

Source: IDC, 2007

OVERVIEW OF SGI'S NEW INTEGRATED BLADE PLATFORM: SGI ALTIX ICE

SGI Has Focused on Making Complex HPC Servers Work Well for Users

Deploying large-scale systems in government labs and universities. SGI has a long history of designing, installing, and servicing high-performance large-scale systems. Buyers of large-scale HPC systems are typically in government labs and university HPC centers, although some are also in industry. The key to SGI's success with large installations at sites such as NASA Ames (10,240 processors), Germany's Leibniz Rechenzentrum (4,096 processors), and the University of Dresden (1,500 processors) is the company's deep experience in designing HPC systems that tightly integrate highly scalable, purpose-built hardware and software systems. It is no mean feat to deploy HPC systems of this size that not only fill impressively large rooms but also are highly utilized to run the world's most challenging problems.

Meeting the needs of industry. SGI is well regarded by industrial HPC users/buyers, primarily because the company historically has offered HPC systems that can run a wider variety of optimized ISV applications than those of most competitors, and run them faster. In a study of worldwide HPC users in industry, IDC asked respondents to state their top 3 preferences among HPC systems vendors for their next procurement. SGI has consistently scored well on making it into the top 3 list. The reasons given by the industry end users for their vendor preferences amounted to a checklist of the following desirable traits: leaderlike image, willingness to take risks with the customer, price, performance on applications, discipline-specific knowledge, recommendations from trusted sources, and others.

SGI Altix ICE: A Bladed Solution in the SGI Tradition

With the new Altix ICE product, SGI set out to prove that not all blades are alike. The approach to SGI Altix ICE is squarely in the SGI tradition — marrying the company's differentiated performance technologies with commodity components in an industry-standard environment with high-manageability features. SGI reports that the company talked with more than 250 HPC customers en route to finalizing the product design.

SGI Altix ICE borrows important architectural concepts and performance technologies from the SGI Altix 4700 bladed server system that was introduced in June 2006, such as the 3D torus interconnect design, operating system (OS) synchronization to avoid performance-degrading interruptions ("jitter") from housekeeping tasks, chilled-water doors, and power-saving techniques. But SGI Altix ICE applies these value-added technologies within a more mainstream design that features Xeon-based processors, distributed memory, and an InfiniBand interconnect, as opposed to the Altix 4700's Itanium processors and NUMAlink shared-memory network. (SGI will continue offering the Altix 4000 series server line for bandwidth-intensive workloads in the HPC and high-performance business computing markets.) In addition, Altix ICE bladed servers exploit SGI's lengthy experience with system modularity/flexibility.

The product's four-level management hierarchy (single node, chassis, rack, system) should be especially welcome for system administrators who have been struggling with managing the increasing complexity of existing rackmounted and bladed cluster products. Most HPC sites are growing the size of their clusters each year, in many cases doubling the size of their clusters each year, and will likely welcome these features.

SGI Altix ICE bladed servers are available with either dual-core Woodcrest or quad-core Clovertown Xeon processors from Intel. U.S. list pricing for the new bladed server system will start at about $350,000 for a full rack with 512 processor cores and 1GB of memory per core. The price includes a rack, the power and cooling infrastructure, and the interconnect. SGI expects to bid a wide range of system sizes, including large-scale, multirack configurations.

SGI has clearly stated that today's technical and scientific computing customers will stand for nothing less than genuine HPC performance and scalability — two things that can be difficult to attain from many HPC clusters, especially for large and very large scale-out workloads. The company understands that HPC customers can no longer carry the burden of sprawling small-node cluster deployments that voraciously

consume kilowatts and floor space. It is introducing the Altix ICE as a leading-edge platform that closes the gap between productivity and performance with a densely integrated, rapidly deployable, and reliable solution that delivers all the benefits of high-performance SGI systems combined with groundbreaking energy and cost efficiencies.

Targeting the HPC and Proximity Markets

SGI is targeting SGI Altix ICE primarily at the HPC sector, including traditional HPC industry segments (automotive, aerospace, oil and gas, life sciences, government, and university-based research) as well as proximity markets such as financial services (where SGI has had prior experience). SGI Altix ICE's industry-standard InfiniBand backbone promises to broaden SGI's reach in these target markets.

The system's array of manageability features and tools will be important for HPC users and should also be welcome in certain commercial environments where server virtualization is an important trend. For SGI Altix ICE, SGI leveraged open source cluster manager technology but enhanced it to deal with the large 3D torus interconnect topology and the product's four management levels: single node, chassis, rack, and system. SGI also designed its own fabric manager for the InfiniBand network. Storage can be configured as a separate, virtualized resource. In addition to SGI's own management tool, SGI customers can use Scali Manage and Altair PBS Professional to manage the SGI Altix ICE environment.

Benefits of SGI's Xeon-Based Bladed Servers

IDC believes SGI Altix ICE Xeon-based bladed servers provide the following potential advantages for HPC users:

- **Performance density.** IDC research shows that density is the single most important issue for bladed server customers, and SGI Altix ICE blades excel in this respect. Each 42U rack (footprint: 30" wide x 40" deep) can hold up to 512 quad cores (256 dual cores). Chilled-water doors are optional for large configurations. The blades plug into a backplane, dramatically reducing cabling requirements and improving system reliability.

- **Performance scalability.** SGI Altix ICE bladed servers should do well at scaling application performance, thanks to the product's high-speed, low-latency InfiniBand interconnect (10/20GBps, 2–5μs) and the ability to implement dual InfiniBand networks to improve data management and reduce potential bottlenecks. In addition, OS synchronization minimizes parallel performance degradation from housekeeping tasks, reducing wall clock time on large-scale MPI jobs.

- **Pricing, price/performance.** By using a larger fraction of commodity components (e.g., Xeon processors, InfiniBand, and Linux), SGI is able to deliver more aggressive pricing and price/performance for SGI Altix ICE bladed servers. In addition, SGI has removed components that do not add to performance, such as disk drives on each blade.

- **Power and cooling.** SGI Altix ICE bladed servers leverage high-efficiency power and cooling technologies first developed for the Altix 4700, including the company's third-generation power and chilled-water cooling solutions. The cooling technologies enable scaling to large configurations, while the high-efficiency power supply units reduce electrical consumption and help improve cost of ownership.

- **Manageability.** Clusters, whether rackmounted or bladed, have been difficult to manage, especially in larger system sizes. With SGI Altix ICE blades, SGI provides management capability at four levels: node, chassis, rack, and system.

- **Diskless architecture.** Unlike most bladed servers available today, the Altix ICE product is a diskless architecture. This diskless approach brings with it several important advantages: reduced cost, reduced power and cooling requirements, and, with fewer moving parts, increased reliability and performance.

- **Software stack.** Altix ICE's rack leader controller (one per rack), administrative support node (one per cluster), and other nodes are outfitted with SGI-enhanced, standards-compliant software that draws on SGI's rich Linux heritage. The software supports OS, system, cluster, and InfiniBand management.

Product Overview and Specifications

Foundations for the Performance-Optimized Architecture: Compute Blade and Rack Unit

- **Compute blade.** The SGI Altix ICE compute blade, which was codesigned by SGI and Intel, employs an Intel 5000x chipset (Greencreek) in conjunction with two Xeon processors — either dual-core Woodcrest or quad-core Clovertown. Each blade has eight fully buffered DIMM slots (1, 2, or 4GB DIMMs) and 32GB of memory support, along with two 4x DDR InfiniBand ports. There is no on-board storage (storage can be treated as a separate resource for virtualization).

- **Individual rack unit (IRU).** There are 4 cores (2 dual-core) or 8 cores (2 quad-core) per compute blade, 1 compute blade per node, and 16 nodes per IRU. Each IRU also contains 2 InfiniBand switch blades and a chassis management controller.

Strong Interconnect Between Nodes and Cabinets

The compute and switch blades plug directly into the 4x DDR InfiniBand backplane, avoiding the need for cabling and improving system reliability. Fast Gigabit Ethernet connections link together the leader nodes in different racks, while InfiniBand connects the cluster as a whole.

Pretested and Integrated Software Stack

With SGI Altix ICE bladed servers, SGI provides a factory-tested and integrated software solution stack with the following key elements:

- **Linux OS.** SGI Altix ICE bladed servers are available with the SuSE Linux Enterprise Server 10, and support is planned for the Red Hat® Enterprise Linux® operating system. SGI adds to the mix its SGI ProPack toolkit for workflow optimization, cluster management, and storage administration; the SGI InfiniteStorage data management solution; and comprehensive RAS features. Each rack's leader node provides boot services for all of the diskless nodes (up to 64) in a rack.

- **Support nodes.** As mentioned earlier, the Altix ICE bladed servers come with an administrative support node (one per cluster) that includes management software to provision and manage the cluster, along with leader support nodes (one per rack) that run the fabric management software and monitor, manage, and retrieve data from the rack and its compute nodes. Other node types are as follows: log-in service node, batch service node, gateway service node, storage service node, and fabric management support node. The rationale for this built-in hierarchy is to provide finer-grained, more scalable management functionality — a distinct advantage in the larger-scale deployments SGI is targeting with Altix ICE.

- **System and cluster management.** Aside from node-level management, Altix ICE bladed servers feature an SGI-developed chassis management controller with support for scaling to large systems and system management and monitoring capabilities provided by the cluster management software tool. This feature-rich tool draws certain components from the Open Source Cluster Application Resources (OSCAR) distribution found at OpenClusterGroup.org, with substantial SGI enhancements for scalability to 50+ racks and thousands of diskless compute nodes. This means that customers familiar with OSCAR should be able to transition easily to using SGI's management tool and should also be able to use extensions that they have developed for use with OSCAR.

- **InfiniBand fabric configuration and management.** SGI's solution enhances the OpenFabrics Enterprise Distribution (OFED) software from OpenFabrics.org.

Multidimensional Flexibility

The SGI Altix ICE product is highly flexible and customizable in multiple dimensions:

- Choice of dual-core or quad-core processors to match varying workloads

- Modularity to add compute resources independently

- Redundant, hot-swappable power and cooling

- Blade design that provides rapid serviceability

Part of a Broad Product Portfolio

SGI's server portfolio offers users a wide range of choices in architectures and sizes, in both rackmounted and bladed configurations:

- The SGI Altix 4000 series of high-end servers (including the Altix 4700 bladed Itanium servers)
- SGI Altix 450 midrange servers (based on the Altix 4000 series blade architecture)
- SGI Altix XE x86-64 servers and clusters
- A wide array of storage hardware and software solutions, along with a full set of service offerings

FUTURE OUTLOOK

IDC projects that the overall technical server market will continue to show strong growth over the next five years, but at a more moderate level of 8% to 12% a year. The cluster portion of the HPC market will likely grow at a higher rate during this period, in the 14% to 20% range, and top out at 70% to 80% of the HPC market. Blade-based clusters could exhibit even higher growth if vendors can address the issues raised in this paper. Vendors such as SGI, whose experience prepares them to address these critical performance, scaling, and availability issues, have better opportunities for exploiting the growth of rackmounted and bladed clusters in the years ahead and for moving their HPC customers forward in their missions.

The key areas in which clusters can add value to customers' mission-related work are shown in Table 9. Most important among these is providing the ability to do more science and engineering, closely followed by cost and time savings.

TABLE 9

Value Areas for Clusters

Q. In what areas are clusters providing your organization a major improvement?

	Number of Mentions	% of Respondents
Ability to do more science/engineering (increase workload size)	64	66.7
Reduced overall costs	55	57.3
Reduced research/engineering time (e.g., faster turnaround time on applications)	49	51.0
Ability to do better science/engineering (increase accuracy/precision)	38	39.6
Reduced downtime	28	29.2
Other areas	50	52.0

Note: Multiple responses were allowed.

Source: IDC, 2007

CHALLENGES/OPPORTUNITIES

Challenges

- **SGI is not first to market.** Multiple vendors have already introduced Xeon-based bladed server products into the HPC market. Companies such as IBM and HP have established market footprints in the HPC space using blade-based systems designed for the commercial enterprise market, which creates a barrier to entry for SGI. However, SGI's HPC-focused Altix ICE bladed servers should be an attractive alternative, especially for end users with medium-sized to large scale-out workloads. Nonetheless, SGI will need to educate the market (SGI customers and others) about the product and performance differentiation for its SGI Altix ICE bladed servers and will need to make its case wherever possible with proof points. This will require a clear and concerted effort.
- **SGI's corporate transitions will make some buyers uneasy.** Unlike their counterparts in commercial enterprise computing, HPC users/buyers historically have been willing to switch vendors with relative ease when a better proposition comes along (performance, price/performance). There are repeated examples within the HPC market of vendors whose fortunes have markedly improved due to offering higher-performing products. HP, the current revenue leader in the overall HPC market, is perhaps the best example of a company whose HPC future was in question not long ago. SGI's past corporate history may raise concerns among some buyers that SGI will need to address. However, SGI's recent financial restructuring and refocusing on delivering server and storage solutions to the HPC market, as evidenced by the introduction of the Altix ICE server line, should give customers comfort that the company is back on track and as committed as ever to the HPC market. It will help that throughout this challenging period, users' respect for SGI as a technology powerhouse and a producer of HPC systems that work well at large scale has stayed strong.

Opportunities

- **The market for bladed servers is still new.** This fast-growing market remains fragmented, and no vendor has established dominance yet. For a vendor with SGI's long history and continuing good reputation in HPC, there is ample opportunity to gain a foothold and grow share in the bladed server market.
- **SGI is already having success with SGI Altix ICE bladed servers.** SGI is focusing on a market (HPC) that it knows intimately. SGI reports that it has more than 6,000 technical computing customers, and the company has disclosed an impressive list of initial customer interest for SGI Altix ICE bladed servers.
- **SGI's bladed servers are differentiated.** SGI's greater use of commodity technologies in SGI Altix ICE bladed servers than in the Altix 4700 series provides a cost advantage but limits the degree of differentiation for the newer product. Still, the SGI Altix ICE servers are far from being "vanilla" look-alike offerings. Their design incorporates substantial SGI value-added technologies that are intended to address key customer "pain points" with commodity clusters, whether rackmounted or bladed. The efficacy of the key SGI value-added technologies has been demonstrated in the marketplace on prior SGI products, including on large-scale installations.

- **SGI stands out from many competitors with its customer-intimate style.** Unlike most vendors that sell commodity, "white-box" clusters, SGI has the ability to forge peer-to-peer relationships with customers, based on the company's long-time practice of staffing with experts in the computational disciplines practiced by the customer organizations. This ability provides important advantages for customer acquisition and retention.

- **SGI plans to focus its bladed servers on other markets.** The company also plans to aim the SGI Altix ICE servers in the near term at HPC proximity markets, including financial services (for portfolio pricing/optimization, global risk management, and other back-room activities performed by "quants").

CONCLUSION

IDC expects that the HPC cluster market will continue to see high growth over the next few years and that as users continue to grow their cluster installations, higher-performance features will increase in importance. In addition, we expect that blade-based clusters will provide a higher level of density and should see very strong growth.