IDC Analyze the Future

I D C T E C H N O L O G Y S P O T L I G H T

New Data-Intensive Problems Need Supercomputers with Supersized Memories

May 2011

Adapted from *Data-Intensive Computing in HPC* by Steve Conway, et al., IDC #226231

Sponsored by SGI

Data-intensive ("big data") problems, in fields from cryptography to meteorology, have been as important as their compute-intensive counterparts since the birth of scientific computing in the post–World War II era. In recent years, advances in the capabilities of high-performance computing (HPC) systems have made it feasible to consider using these systems to run much larger, higher-resolution versions of familiar data-intensive problems — as well as a whole new wave of challenging scientific, industrial, and commercial problems related to antiterrorism, fraud detection, energy management, drug discovery, epidemiology, social networking, and many other areas. Some of the new problems involve unstructured (NoSQL) data that does not fit the formats of relational database management systems (RDBMS).

IDC expects data-intensive workloads to become much more prominent in the HPC market in the future. In addition, while many "big data" problems will be run on standard clusters, limitations in the memory sizes and memory architectures of clusters make them ill suited for the most challenging classes of data-intensive problems. Sites that need to run these high-value problems already are seeking out HPC systems with beefed-up memory capabilities. IDC expects this trend to continue.

This paper examines the rise of data-intensive problems that need high-performance computing solutions. It also looks at the role of HPC big-memory systems vendor SGI in this emerging, high-potential market.

Introduction: "Big Data" Problems and Requirements

The emerging wave of HPC-enabled data-intensive solutions can help companies, government organizations, and others tackle a wide range of previously intractable problems that are important for economic competitiveness, scientific advancement, national security, and the quality of human life.

Where HPC is concerned, IDC defines data-intensive ("big data") problems broadly to include tasks involving sufficient data volumes and complexity to require HPC-based modeling/simulation. The problems can employ structured data or unstructured data, or both. They can come from traditional HPC domains in government, industry, and academia — or they can be upward extensions of commercial problems that have grown large and complex enough at the high end to require HPC. In addition, "big data" can accumulate from the multiple results of iterative problem-solving methods in sectors such as manufacturing (parametric modeling) and financial services (stochastic modeling).

Example: Fraud Detection. Fraud detection is a data-intensive problem of great economic importance to both government and commercial organizations. On the commercial side, eBay recently purchased supercomputers to combat fraud within the PayPal system. In the government arena, Italy's social services agency, Istituto Nazionale della Previdenza Sociale (INPS), acquired an HPC system to attack healthcare fraud on a national basis. The U.S. government is considering an HPC solution to attack fraud estimated to top $150 billion a year in federal healthcare programs — Medicare, Medicaid, Veterans Affairs, and others.

Traditional data-intensive problems, some of which require HPC solutions, typically involve "finding a needle in a haystack," that is, searching for discrete items that already exist in a database. In fraud detection, by contrast, the most useful data occurs as a hidden pattern, rather than one or more discrete items. There is a large class of these mission-critical, pattern discovery problems in government and business: fraud detection, financial modeling/risk analysis, antiterrorist analysis, epidemiology, social network analysis, semantic analysis, and a host of others.

These economically and strategically important problems pull from a common list of tasks: pattern matching, scenario development, behavioral prediction, anomaly detection, and analysis of relationships using graphs. They're for things like catching terrorists before they leave the airport, catching bank fraud before the criminal gets away with the money, analyzing patterns in social or biological networks, aiding in drug discovery, optimizing risk in investment portfolios, or protecting a national power grid before it crashes. Some of the powerful algorithms in this domain originated in classified government.

Special Requirements for Data-Intensive Solutions. Performance on compute-intensive problems is generally measured as the speed (ideally, time to solution) for completing a series of calculations. In contrast, data-intensive problem-solving performance typically is gauged by how fast the computer can traverse one or more large data sets, sometimes using special software frameworks such as MapReduce, Hadoop (Linux), or Dryad (Windows). Graph analysis and other methodologies are used to evaluate massive data sets that often have irregular, sparse grid patterns.

High-end, data-intensive methodologies place huge demands on systemwide memory capabilities — memory size, memory bandwidth, and memory latency. Although most data-intensive problems will be run on standard clusters, clusters are not well suited for the most challenging "big data" problems, which also tend to be the most important ones, for these main reasons:

- Systemwide memory on standard clusters is limited in size and is not logically shared. There is no single memory space big enough to hold all the data for a sizable "big data" problem. Instead, the data has to be split into chunks small enough to fit into the system's logically distributed memory locations. This works well enough for compute-intensive problems that exhibit regular (i.e., readily partitionable) data patterns, but some compute-intensive and many data-intensive problems are irregular — they don't allow simple data partitions. Hence, the absence of a large, globally shared memory space can make some challenging "big data" problems intractable and significantly hamper the performance on others by heavily shifting the burden onto the systemwide (especially off-node) communications capabilities of the memory subsystem.
- Standard clusters have limited systemwide communications capabilities. Systemwide (bisection) bandwidth, and especially latency, is inadequate to enable fast solution times on the most demanding "big data" problems — even with the latest InfiniBand or 10 Gigabit Ethernet interconnects. The latencies of standard clusters typically are also too high to support cache coherency across the clusters' distributed memory locations.

For these reasons, sites that need to run the most demanding "big data" problems already are seeking out HPC systems with memory capabilities superior to those available on standard clusters. This is true for the eBay/PayPal, INPS, and U.S. healthcare examples cited earlier.

The Benefits of Data-Intensive Solutions

The benefits of the most demanding, HPC-class "big data" solutions vary according to the problem type, but they are typically large and extensive. Consider the following examples:

- Business fraud detection could save tens of millions of dollars, while government fraud detection could save billions.

- Advanced social network analysis could identify relationships of high economic value to companies that exploit social networks and could lead to the improvement of the hardware/software architectures used by social network providers.
- In the financial sector, data-intensive methods could enable better macroeconomic and microeconomic modeling, including stochastic modeling for risk analysis.
- Data-intensive drug discovery could produce a new generation of high-efficacy cancer-fighting agents with more tolerable side effects, saving or prolonging human lives. Data-intensive epidemiology could improve our ability to control the spread of serious diseases.
- Signal intelligence and entity analysis in the defense sector could identify terrorist threats substantially earlier than is possible today and could improve emergency preparedness plans.
- In the manufacturing sector, data-intensive solutions could speed innovation and help deliver more competitive products to market faster.

Trends Affecting Data-Intensive HPC Solutions

The key trends affecting data-intensive HPC solutions are as follows:

- **The high-end data explosion.** The data explosion affecting the whole IT market is even more pronounced at the high end. To cite a few examples, the Large Hadron Collider at CERN churns out 1 petabyte of data per second when it's running. Multipetabyte storage capacities are no longer uncommon. Manufacturers such as Boeing now use CFD far more extensively than in the past and are required to save engineering data for the 30-year average life of an aircraft.
- **The trend toward increasingly unbalanced HPC systems.** In the past decade, nodal and systemwide memory speeds have fallen behind processor speeds, making it more difficult to feed the processors enough data to keep them busy. This "memory wall" fundamentally limits the capabilities of standard clusters, which have become the dominant species of HPC systems.

Considering SGI Altix UV

The new SGI was formed from the May 2009 acquisition of Silicon Graphics by Rackable Systems Inc. Rackable made a name for itself as a server and storage vendor by landing whale-size Internet-based customers, pioneering cloud and "green grid" offerings, and establishing a footprint in the HPC market. SGI entered the union as an iconic HPC firm known for delivering SMP and other computer systems, along with state-of-the-art visualization, that worked well even at very large scale. SGI's now-integrated dual legacy gives the company strong street credentials in the commercial, government, and university markets for both data-centric (Rackable) and compute-centric (Silicon Graphics) customers.

In the past decade, SGI bucked the trend toward unbalanced HPC systems by continuing to provide servers with big, globally shared, cache-coherent (ccNUMA) memory subsystems. Previously, Silicon Graphics faltered by basing its HPC systems on more costly, nonmainstream (non-COTS) technologies: MIPS, IRIX, and Itanium. Customers loved the capabilities of their SGI systems, but for cost reasons typically bought only a few racks and treated them like crown jewels in their datacenters, limiting access to the centers' most important users. Rackable remedied this by investing in the completion of the Altix UV server family, making these products more cost-competitive while retaining their outstanding memory capabilities.

In recent worldwide HPC studies and site visits, IDC has seen considerable excitement about SGI's Altix products, especially the Altix UV series. This product has already earned strong testimonials from installed customers across the globe, including the University of Tennessee, the North German Supercomputing Alliance (HLRN), France's CALcul en Midi-Pyrénées (CALMIP), and Japan's Hokkaido University.

The SGI Altix UV system is purpose built to deliver fast solution times on demanding data-intensive problems (as well as compute-intensive tasks). The scope of this paper does not permit an in-depth product review, but the following highlights should suffice to show the exceptional capabilities of this product series, especially for challenging "big data" problems and workloads:

- **Standard technology foundation.** The SGI Altix UV product line comes in three series: the Altix UV 10, the Altix UV 100, and the high-end Altix UV 1000. All are based on mainstream COTS technologies: Intel Xeon Processor E7 family of processors and standard operating systems (choice of Novell SUSE Linux, Red Hat Linux, SGI Linux System Software, or Microsoft Windows Server).
- **Turbocharged memory capabilities.** The SGI Altix UV 1000 product's globally shared, cache-coherent memory scales to 16TB — the addressable limit of Intel Xeon 7500 processors. The NUMAlink 5, paired-node 2D torus interconnect provides bandwidth of 120Gb/second (15Gb/second bidirectional) and submicrosecond latency.
- **Range of products and configurations.** The SGI Altix UV 10 is a small-scale, 4U rackmount server with up to four Intel Xeon 7500 processors (quad-, six-, or eight-core), maximum memory of 1TB, and four 1 Gigabit Ethernet channels. The Altix UV 100 and Altix UV 1000 series are bladed servers with two Intel Xeon 7500 E7 processors (quad-, six-, or eight-core) per blade, along with the scalable features noted in the preceding bullet.

In a new customer example, a major global courier company will use a large SGI Altix UV system to perform real-time fraud and error detection on billions of scanned pieces per day, coming from hundreds of processing and distribution centers. This daunting task, which exceeds the bandwidth and latency limitations of RDBMS, could save billions of dollars annually.

The main attractions of the SGI database acceleration solution for this customer are that it's an open Linux platform, it can scale out and scale up using cost-effective standard processors, and it features strong RAS capabilities — especially redundancy, resiliency, and failover protection. In benchmark tests, the SGI-enabled solution was more than 1,000 times faster than the prior method, processing more than 4 billion transactions in a two-hour window.

Challenges

Despite SGI's strong experience in HPC and commercial data-intensive markets, the company faces the following challenges to exploiting new opportunities:

- **Increase sales volumes.** As noted earlier, SGI Altix products historically have received high marks from customers but have not sold in large enough volumes. The move to mainstream base technologies makes the Altix UV line more price competitive. SGI's marketing/sales challenge will be to leverage this new advantage as fully as possible to drive more sales.
- **Communicate the SGI added value to price-conscious cluster buyers.** SGI will need to educate some prospects with differentiated requirements about the value of acquiring SGI's differentiated solutions. Otherwise, these buyers may default to buying standard clusters that are not well suited to their problems.
- **Actively engage in helping customers realize their goals.** IDC research shows that customers increasingly expect HPC vendors to work alongside them to realize the customers' goals for these products. The higher the sales prices, the more participation customers typically expect. Fortunately, SGI is accustomed to engaging in peer-to-peer relationships with customers.

Conclusion

IDC believes that the HPC market for demanding "big data" solutions will become much more prominent in the near-term future. Companies like SGI that provide performance-differentiated, price-competitive solutions and are organized to help customers achieve their goals for these solutions are positioned well to benefit from the anticipated growth of this market.

ABOUT THIS PUBLICATION

This publication was produced by IDC Go-to-Market Services. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Go-to-Market Services makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

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