IDC TECHNOLOGY SPOTLIGHT

Cray Transformation Mines Core Strengths to Address New Markets

October 2013

By Steve Conway, Earl C. Joseph, Ph.D., and Chirag Dekate, Ph.D.

Sponsored by Cray Inc.

The global market for high-performance computing (HPC) servers and related resources has been growing rapidly for more than a decade, spawning economically important new segments with their own technical and pricing requirements. HPC has always lived at the intersection of compute-intensive work ("big compute") and data-intensive work ("big data"). The rise of newer analytics methods is motivating more commercial firms to adopt HPC to support mission-critical advanced searches and pattern discovery that cannot happen today without HPC. The convergence of established HPC markets and high-end commercial analytics markets is already starting to usher in an important new era in high-performance data analysis (HPDA) — that is, big data needing HPC. For several years, Cray Inc. has been transforming itself to benefit more strongly from the expansion and diversification of the worldwide HPC market, especially through the addition of cluster, storage, and advanced analytics products. Because this controlled transformation has happened one step at a time, its cumulative impact may not be apparent. Cray has already created a significantly larger total addressable market (TAM) for itself. IDC believes that Cray is now well positioned to grow to the next level by expanding substantially beyond its established stronghold in high-end supercomputing.

HPC and HPDA Market Drivers

During the past 15 years, HPC has been one of the fastest-growing IT markets. The worldwide market for HPC server systems tripled in size from $3.7 billion to $11.1 billion between 1996 and 2012 — a period including a major economic recession — and is headed toward $15.4 billion in 2017. The broader HPC ecosystem that includes servers, storage, software, and support services swelled to $21.9 billion in 2012 and is expected to reach $30.2 billion in 2017.

There have been two key drivers of HPC market growth in recent years, each of them reflecting trends in user requirements. The main drivers are as follows:

- **Clusters.** The arrival of commercial-grade clusters in 2001–2002 made HPC affordable even for most small and medium-sized businesses (SMBs), small and medium-sized scientific organizations (SMSOs), and start-ups. Thanks to their compelling price/performance, standards-based clusters firmly established themselves as the dominant species of HPC systems by the middle of the decade. Today, cluster pricing extends from under $10,000 to tens of millions of dollars for some of the largest, supercomputer-class clusters. Clusters represent about two-thirds (64–65%) of HPC servers sold today.
- **Supercomputer systems.** The supercomputer segment for HPC systems costing $500,000 and up has grown very rapidly, from $2.7 billion (27.5% of all HPC server revenue) in 2008 to $5.6 billion (51% of HPC server revenue) in 2012. Driven heavily by the ongoing global race for HPC leadership — as a requirement for scientific and industrial leadership — this high-value segment was unfazed by the worldwide economic recession. In 2009, the worst year of the recession, revenue for supercomputers priced at $500,000 and up grew 35%, and revenue for high-end systems selling for $3 million and above jumped a whopping 65%. IDC forecasts that the supercomputer segment will grow to $6.6 billion in 2017.

The supercomputer segment includes both large-scale clusters ("cluster supercomputers") and more tightly coupled HPC systems. Tightly coupled supercomputers have driven high-end market growth more powerfully because of their more prominent role in the global HPC leadership race. But cluster supercomputers have also played an important role, as large government-funded HPC centers and industries such as oil and gas purchase more petascale clusters to meet their users' growing needs for massively parallel, scale-out application performance.

A third important HPC growth driver has emerged more recently:

- **High-performance data analysis.** IDC coined this term to refer to data-intensive ("big data") workloads that require or benefit greatly from HPC resources (see the Definitions section), even though not all HPDA beneficiaries consider themselves HPC users. These workloads include established data-intensive simulations and newer advanced analytics problems. The common denominator for HPDA problems is a degree of algorithmic complexity that is atypical for operational business workloads. The proliferation of newer analytics methods and real-time requirements will propel growth in the HPDA market. IDC believes that graph analytics, in particular, will drive a transition from the current era of static searches for discrete, known items to an emerging era characterized by higher-value, real-time discovery of unknown patterns. IDC forecasts that revenue for HPDA servers will grow robustly (15% CAGR) during the period 2012–2016, increasing from $743.8 million in 2012 to approach $1.3 billion in 2016. HPDA storage revenue will near $800 million by 2016.

HPC/HPDA Market Requirements

No HPC or HPDA system is best for every customer's requirements and budget. That's why a range of products is available in these growing, closely related markets at different price points and with varying features and functions. Many vendors offer a portfolio of products to address a larger portion of the markets' actual and potential revenues (i.e., to increase the sizes of their TAMs).

But not all of the market opportunities are equally attractive. Nearly all HPC and HPDA products today are heavily based on industry-standard technologies. The lowest common denominators are lookalike "white box" computers that do not add much to these standard technologies. IDC research shows that succeeding long term in this white-box commodity market is difficult. Vendors must contend with razor-thin margins and cutthroat competition.

When it comes to capturing and sustaining profitable market share, products that provide meaningful differentiation have a distinct advantage in the HPC and HPDA markets. Differentiation is especially needed for the following important market requirements:

- **Scaling application performance to large system sizes.** On the November 1999 list of the world's most powerful supercomputers (www.top500.org), the number 1 system boasted 9,632 cores and peak performance of 3.2TF. Fast forward to the June 2013 list and the number 1 supercomputer featured 3.2 million processor cores and 77,300TF. That's a 24,156-fold peak performance gain in little more than a dozen years. Not surprisingly, IDC research shows that HPC users consistently report that their number 1 issue is scaling applications performance to exploit substantial fractions (>20%) of leadership-class supercomputer systems. Only a few vendors have demonstrated an ability to enable user applications to perform efficiently at extreme scale.
- **Building on the price/performance of standard technologies.** The compelling price/performance of HPC clusters based on industry-standard technologies such as x86 processors, the Linux operating system, and MPI has made clusters the dominant species of HPC systems. But nearly all cluster vendors today add their own features and functions on top of the standard technologies in order to separate their products from the rest of the crowd. Even more tightly coupled, noncluster HPC systems exploit standard technologies wherever that makes sense, and then differentiate to an even greater degree than clusters.

- **Supporting flops-based and integer-based workloads.** Most data-intensive HPC analysis to date has been flops-based modeling and simulation, but the rising tide of advanced analytics work is primarily integer based. Vendors that can support both types of data analysis will have access to expanded demand among existing HPC users and to commercial firms that adopt HPC technologies and approaches for the first time in order to run advanced analytics.

Definitions

- **Cluster:** IDC defines clusters used in technical markets as a set of independent computers combined into a unified system through systems software and networking technologies. Thus, clusters are not based on new architectural concepts so much as new systems integration strategies. Clusters today are heavily based on standard technologies, such as x86-based processors, the Linux or Windows operating system, and the MPI message-passing protocol.
- **Graph analytics:** IDC defines graph analytics as a computational method for identifying and visualizing relationships among items in a database, and assessing the relative strengths and natures of the relationships based on their connecting lines (edges). Graphs generally exhibit irregular data patterns and cannot be easily subdivided (partitioned) for parceling out to the distributed memory locations of standard clusters. This methodology benefits from computer systems with large shared memories and strong I/O capabilities.
- **HPDA:** High-performance data analysis is the term IDC coined to describe the convergence of established data-intensive HPC markets and the high-end government and commercial analytics markets that require similarly powerful computing resources. In essence, HPDA refers to big data workloads that require HPC resources.

The Cray Transformation

The Cray company name is synonymous with the term supercomputer. Cray Inc. is the linear descendant of Cray Research. Since regaining its independence in April 2000, Cray has added substantially to the company's prodigious record of contributions to HPC technology and practice. In the process, Cray has expanded its revenue from $26.1 million in 2001, the company's first full year of operation, to $421.1 million in 2012, with further growth signaled for 2013.

For years, Cray led the worldwide HPC market when the market consisted overwhelmingly of high-end supercomputers. Today, Cray is again one of the worldwide leaders in the high-end supercomputer segment for HPC systems selling for $3 million and up. This segment has benefited lately from the worldwide race for HPC leadership as a requisite for scientific, industrial, and economic competitiveness. In 2009, the worst year of the global recession, revenue for $3 million-plus supercomputers jumped a whopping 65%. While another leap of that magnitude is not on the near horizon, this segment will remain attractive for the few vendors that are capable of deploying the largest, most powerful HPC systems.

Until recently, Cray's growth was powered almost exclusively by sales and service revenue from the company's flagship supercomputer line, whose current incarnation is the Intel-based Cray XC30 system. Pricing for these products is in the $500,000-and-up range that in 2012 amounted to $5.6 billion, or 50.1% of the $11.1 billion worldwide market for HPC servers. Again, a nice place to compete if you're up to the task.

But in the past few years, Cray has taken bold steps to extend its reach into the marketplace (i.e., TAM) by growing the company's product portfolio while also streamlining its intellectual property (IP). To set the stage for enhanced success, Cray has leveraged existing strengths, acquired complementary capabilities, and divested itself of resources it will no longer need to own outright.

It's accurate to say that the company has been moving into new-to-Cray market segments. It's equally accurate to say, however, that the requirements of these segments have been moving into Cray's competency space. A case in point is advanced analytics ("big data"), where contemporary HPC technologies are finally enabling companies, government agencies, and others to realize their unfulfilled 1990s dreams of turning mountains of data into actionable knowledge — often in near real time.

Elements of the Cray Transformation

Cray's transformation to enhance its market reach and future success is already far along. The strategy is based at heart on further leveraging the R&D investments the company makes in its flagship supercomputer products. But because the ongoing strategy has unfolded one step at a time, its cumulative potential may not be apparent. With that in mind, consider this list of important steps Cray has taken to date:

- **Introduced scaled-down versions of flagship high-end supercomputers.** As part of Cray's broader technical computing strategy, the company has been introducing scaled-down versions of flagship high-end supercomputers. The Cray XE6m series adapted the benefits of its Cray XE6 high-end parent for a price range that straddled the IDC supercomputer ($500,000 plus) and divisional ($250,000–499,000) competitive segments. The "m" series retained the Cray hardware mix (AMD Opteron processors, NVIDIA GPUs, and Cray's Gemini interconnect), but with cost-saving measures such as downshifting the interconnect from a 3-D to a 2-D torus, on the grounds that the 2-D version provides plenty of firepower for midrange workloads. From the Cray XT5m to the current Cray XC30-AC, these scaled-down versions have delivered the key benefits of Cray's flagship supercomputers with features and pricing appropriate for the midrange market.

- **Launched air-cooled versions of high-end systems.** Liquid cooling is inherently more efficient than air cooling, which is why the world's densest, brawniest supercomputers are cooled by water or another liquid. But not all HPC datacenters are equipped to provide liquid cooling. Cray addresses this substantial portion of the global HPC server market with air-cooled versions of the company's flagship line: the Cray X30-AC (Intel Xeon-based) product.

- **Addressed demand for cluster supercomputers.** As noted previously, clusters now reign as the dominant species of HPC systems, with about a two-thirds share of global HPC server revenue. Cray has addressed this demand in two ways. First, the Cluster Compatibility Mode available on Cray's flagship supercomputers and "m" series midrange derivatives allows users to run most standard ISV applications without modification ("out of the box") on these Cray systems. (A stated goal of the company's Adaptive Supercomputing Vision is to provide both extreme-scale computing and cluster compatibility in the same supercomputers.) Second, in 2012 Cray acquired Appro, a successful vendor of large-scale capacity clusters, also called cluster supercomputers. Cray's new CS300 product series, largely derived from the Appro acquisition, enables the company to sell large clusters alone or in conjunction with more tightly coupled Cray supercomputers.

- **Addressed demand for all major processors and coprocessors/accelerators.** In IDC's 2013 HPC worldwide end-user study, more than 90% of all installed processor socket parts were based on the x86 standard. Intel x86 processors predominated, but an important slice of the x86 pie came from AMD. Cray's expanded product portfolio provides choices based on each processor vendor's x86 offerings, along with coprocessors/accelerators from NVIDIA and Intel.

- **Reentered the HPC storage market.** Storage is the fastest-growing part of the global HPC market. IDC forecasts that the HPC storage market will reach $5.6 billion in 2016, growing at an 8.9% CAGR from 2012 to 2016. With its pioneering SSD products, Cray Research was an early pioneer and market leader in high-end storage technology. Cray Inc. has advanced this legacy with its leading-edge Sonexion line of scale-out Lustre storage products. IDC estimates that in 2012, Sonexion was a significant contributor to Cray's reported $50 million in storage-related revenue.

- **Assembled a triple-threat product portfolio for high-performance data analysis.** High-performance data analysis is the term IDC coined to describe the convergence of established data-intensive HPC modeling/simulation markets and newer government and commercial analytics markets that require HPC resources. IDC research shows than in many cases, users are running modeling/simulation and advanced analytics on the same systems. This intersection of big data and big compute is familiar territory for Cray. IDC forecasts that the HPDA server market will grow at a 15% CAGR from 2012 to 2016 to approach $1.3 billion in 2016, by which time the HPDA storage market (18.1% CAGR from 2012 to 2016) will be worth about $789 million. Cray is addressing this opportunity with a trifecta of products that includes the flagship supercomputers, the CS300 cluster supercomputers (configured with the Intel Apache Hadoop distribution), and the Urika graph analytics appliance from the company's YarcData business unit. The Sonexion product should also benefit substantially from the fast-growing HPDA storage market.

- **Developed and deployed differentiated system software.** Deploying hardware systems with stratospheric peak performance is not so difficult, but scaling systems software and applications to exploit a large fraction of the peak performance efficiently is another matter entirely. Cray has continued its long history of enabling users to work at large scale. In particular, the Cray Linux Environment has allowed Cray supercomputer users to scale the Linux operating system (SUSE Linux) to new heights, as a prerequisite for scaling up the performance of their codes. Other Cray systems software isolates and synchronizes housekeeping tasks to minimize interruptions to computing jobs. In addition, software baked into the Sonexion storage products optimizes disk pools to enhance resiliency and recovery times in the event of failed disks.

- **Sold selected interconnect assets to Intel.** In 2012, Cray sold its interconnect hardware assets to Intel for $140 million and transferred some 74 employees who were working on this technology. IDC said then that it was a smart move to monetize these assets because it would become increasingly challenging for proprietary interconnects to stay far enough ahead of the advancing InfiniBand and Ethernet standards. Cray didn't dispose of all its interconnect-related capabilities, however. The company retained personnel with expertise in scaling interconnect software. IDC believes that Intel and Cray may be collaborating in ways that will benefit Cray's interconnect offerings in the future, starting in about 2016.

The Cray initiatives described here are highly synergistic. The company is not venturing into new businesses that have little to do with each other. On the contrary, these initiatives fit together well and can be pursued together efficiently and economically. An example is the recent win at Japan's Railway Technical Research Institute, which put into production a Cray CS300 cluster supercomputer, a Cray XC30-AC supercomputer, and a Cray Sonexion storage system.

Opportunities and Challenges

IDC believes that the transformation under way at Cray is creating substantial new market opportunities, along with some related challenges.

Opportunities

- **Grow top and bottom lines substantially over time.** By breaking out of the high-end supercomputing box where Cray continues to excel, the company is significantly expanding the portion of the worldwide HPC market it addresses. Along with the growing market demand for tightly coupled, supercomputer-class systems and support, Cray can now exploit demand for cluster supercomputers, advanced storage solutions, and platforms for high-performance data analysis. Cray is already benefiting from its expanded product portfolio, as the recent multiproduct sale to Japan's Railway Technical Research Institute demonstrates. IDC expects Cray to pursue a controlled growth strategy that increasingly exploits the company's larger TAM over time. This strategy should substantially benefit both the company's revenue and its bottom line, especially because the expanded portfolio of related products creates increased opportunities for R&D and operating efficiencies.

- **Alleviate the historic "lumpiness" of the business.** Investors like companies to show predictable growth, but the high-end supercomputer market is inherently "lumpy." Annual revenue in this segment depends heavily on a limited number of large transactions. It is not unusual for the customer acceptance of a high-end supercomputer to move to an earlier quarter than expected or to slip to a later quarter. Either occurrence reduces the predictability of the company's quarterly financial results. Cray's expanded product portfolio should make the company's quarterly results more predictable over time by increasing the number of quarterly transactions and quarterly revenue amounts. Any "surprise" acceptance timing would have less impact when the rest of the quarterly revenue is larger.

- **Accelerate the convergence of the HPC and commercial markets.** Cray has long operated at the intersection of big compute and big data. HPC is arguably the original home of big data. Until recently, HPC big data generally meant flops-based, data-intensive modeling and simulation — although a few customers, notably some government agencies, have relied heavily on integer-based advanced analytics. As noted previously, the market for HPDA has been growing quickly in recent years. Both established HPC users in government, industry, and academia and first-time HPC adopters in commercial markets are employing MapReduce/Hadoop, graph analytics, and other advanced analytics methods on HPC platforms to gain new insights into mission-critical problems. Cray's triple-threat product portfolio for HPDA positions the company well to benefit from this important trend. In the process, Cray's advanced solutions can help accelerate the convergence of the HPC and commercial HPDA markets.

Challenges

- **Transform into a multimarket company.** Although all of Cray's initiatives benefit from the same set of core competencies, mapping these competencies to different markets and market requirements is no easy matter. The challenge for Cray and any company entering new markets is to compete on the terms of each market, with the long-term goal of becoming one of the leaders in every targeted market. This strategy is far more sustainable than entering markets with the narrower goal of augmenting company revenue. IDC believes that Cray is pursuing the path toward long-term success in the new markets and segments the company has entered. For example, the company is allowing its YarcData business unit to pursue a strategy of controlled exploration and growth that IDC considers prudent in the formative HPDA marketplace.

- **Update buyer perceptions.** The fact that the Cray name is synonymous with supercomputing is a major advantage, but it means that Cray needs to make sure buyers in existing and new markets are well informed about the company's expanded product portfolio and plans. This is largely a marketing/sales challenge, and IDC observes that the company has been focused on this task, assisted by employees in the YarcData business unit, the CS300 group, and the Sonexion team who are intimately familiar with the requirements and preferences of their respective target markets.

Conclusion

The worldwide HPC ecosystem has skyrocketed in size, especially in the past decade, and is now worth about $22 billion. Important growth drivers have been the price/performance of standards-based clusters and the global race for HPC leadership as a requirement for national and regional competitiveness. HPDA is emerging as a third important market driver. It is a focal point for the convergence of "big compute" and "big data" demand, including from new HPC adopters in commercial markets.

HPC began as the nearly exclusive province of scientists and engineers pursuing advanced research on high-end supercomputers. As the HPC market has grown, it has spawned important new segments, with new technical and pricing requirements. The increased exploitation of standard technologies has made it more difficult for vendors to differentiate, especially in the most price-sensitive segments. Key differentiators include the abilities to scale application performance to large system sizes, to add unique features and functions to systems based on standard technologies, and to support both flops-based and integer-based workloads.

Cray played a dominant role in creating and growing the supercomputer market during the market's first decades, and the company has remained one of the leaders in high-end supercomputers. In recent years, Cray has transformed itself to address a larger portion of today's HPC market. Among other things, the company has introduced scaled-down and air-cooled versions of its flagship high-end supercomputers, addressed demand for cluster supercomputers and for all major processors and coprocessors/accelerators, reentered the HPC storage market, assembled a strong product portfolio for high-performance data analysis, developed and deployed differentiated system software, and divested itself of selected assets it does not need going forward.

IDC believes that Cray's ongoing transformation positions the company to benefit more strongly from the growth of the HPC and HPDA markets.
