IDC
Analyze the Future

IDC VENDOR SPOTLIGHT

Green HPC: Performance Meets TCO

September 2008

Adapted from Worldwide Technical Computing Server Market 2008 Top 10 Predictions, by Steve Conway and Earl Joseph; IDC doc # 210271

Sponsored by Cray Inc.

Seymour Cray, the father of supercomputing, bent the 1976-era Cray-1 into a "C" shape to shorten the electron paths and gave it a new type of liquid cooling to allow the system's single processor to run at a blazingly fast clock speed for that time. From then until recently, advances in high-performance computing (HPC) system density, and power and cooling efficiency, were all about boosting performance. But skyrocketing energy prices and HPC system sizes in recent years have expanded the rationale for these developments.

Today, HPC system designers improve density and power/cooling efficiency not only to drive performance, but also to address power consumption and spatial requirements. In addition, systems with multi-generational cabinets can lower purchase pricing and reduce equipment disposal issues. As a result, performance and total cost of ownership (TCO) considerations have joined forces as key drivers of HPC system design.

This paper examines the interplay of these factors and the trend toward more environmentally friendly, "green HPC." It also looks at the role of Cray Inc. and its approach to addressing these requirements.

Introduction

High-performance computing, once known as supercomputing, was born in an era of cheap electricity. When industry pioneer Seymour Cray and his colleagues fashioned the Cray-1 supercomputer into a "C" shape, it was not to conserve electricity but to speed the transit times of electrons by shortening their pathways within the computer. The name of the game was higher performance — period. Performance was also the impetus for using Freon in conjunction with heat-conducting cold copper plates in this supercomputer.

This early example of HPC-specialized liquid cooling permitted the heat-generating circuits to be packed more densely and to run at higher clock speeds because Freon, like water and most other liquids, can absorb and dissipate more heat than air can. Liquid cooling in a computer functions much as it does in an automobile, where the liquid-filled radiator system allows the engine to perform harder without overheating. As HPC users prepare to enter the petascale computing era, the challenges of cooling large-scale HPC systems — and paying for the associated energy costs — have never been greater.

Performance Drives Cooling Strategies

To prevent overheating in large-scale HPC systems designed for ever-higher performance, HPC vendors have employed liquids including not only Freon, but also 3M's dielectric Fluorinert solution (which is environmentally safe) and water itself, although water is electrically conductive and components such as processors need to be carefully insulated from it. HPC vendors' strategies for deploying liquid cooling have varied, from immersing the entire computer in the liquid, to circulating the liquid through hollow cold plates, to spray evaporative cooling that relies on the substantial heat-dissipating capacity of a liquid-to-gas phase change.

Fan-based air-cooling suffices for many small and midrange HPC systems. There are multiple approaches to air-cooling HPC systems today. Most fans blow air horizontally over the system nodes, from cabinet back to front, but some propel the air vertically from the bottom to the top of the cabinet. As processing densities for even smaller systems have outgrown the cooling capacities of standard fans, HPC system cabinets have increasingly employed a combination of air and liquid cooling, especially high-performance fans in combination with water-chilled doors.

In addition to cooling the compute cabinets themselves, air-cooled systems must deal with the airflow dynamics of the datacenter, which can differ for every site. The risk of inefficiencies or unacceptable hotspots increases if the strategy for cooling the datacenter is not approached holistically. The most efficient cooling strategies take a datacenter rather than system perspective, combining facility-wide air conditioning, system placement to minimize the impact of hot spots, and customized heat removal. In the HPC center at Imperial College London, a low-profile liquid-chilled CO2 wall (built by TROX ATICS) augments the fan-based air-cooling ability of a row of adjoining system cabinets. This approach allows the college's datacenter to use a wide selection of system cabinets from a wide set of vendors.

Another "cooling" strategy, of course, is to minimize the amount of heat that’s generated in the first place. This can be accomplished using low-power processors or by tuning down the frequencies of standard processors. In either case, the downside can be having to apply more processors to achieve a given level of sustained performance, and some application codes do not scale high enough to exploit the larger processor counts. The low-power processor strategy may reduce an HPC system's breadth-of-applicability, making it a more limited-purpose machine.

Performance Drives Compute Density

In an HPC context, IDC defines compute density as the amount of peak processing power (floating point operations per second, in most cases) that can be fitted into a given volume of space. Historically, the main advantage of greater density has been the potential for higher performance on applications and workloads. When circuits are packed more tightly together, electron paths are shorter and transit times (related to solution times) are correspondingly faster, i.e., latencies are lower, assuming the comparison is for connections of equal bandwidth.

A related advantage of higher compute density is the increased potential for running problems faster "locally" (local processor, local node, or local system), rather than having to traverse slower, performance-degrading connections to less-local levels of the system architecture hierarchy. In addition, storage density, especially access density, is also important for performance and has become a major issue of late. The rise of HPC clusters, which more loosely link together commodity servers, reversed the long-standing trend toward higher system-level compute densities. Despite this disadvantage, cluster market share grew to 65% of all HPC server revenue in 2007, aided by distributed programming models, notably MPI.

Green HPC: Environmental Factors Change the Game

In recent years, environmental factors have begun altering the rules of engagement in the HPC market. The most important environmental factors are the following:

- **Power and Cooling.** Escalating system sizes have caused electricity bills to skyrocket. A recent IDC study showed that the size of the average HPC cluster grew six-fold in a two-year period, from 683 to 4,148 processors. Already, the electricity needed to power and cool the largest HPC systems is enough to run a small city. Future extrapolations of system sizes and power requirements are even more alarming. The U.S. Department of Energy's NERSC facility has estimated that petascale supercomputers slated for delivery in the 2010 time frame under the DARPA High Productivity Computing Systems (HPCS) program may require as much as 20MW of power, 16,000 sq ft of space, and $12 million per year in electricity costs. NERSC projects this could go to 60MW in the 2015–2017 time frame. And in some cases where the infrastructure is lacking, adequate power is simply not available to an HPC site at any price. When the sharply rising costs of oil and gas, and the proliferation of multicore processors are factored in, it is no surprise that power and cooling has become one of the top few issues for HPC users.

- **Facility Space.** Ballooning HPC system sizes are stressing the physical capacity of technical datacenters, as users continue on a high growth path of increasing the performance levels in their datacenters. To upgrade their HPC systems, many sites have had to undertake expensive new facility construction or expansion. For example, with affordable space no longer available in Manhattan, financial services firms have been relocating their datacenters to Long Island or New Jersey.
- **Equipment Disposal.** Less publicized but also important is the cost of equipment disposal. All too often, buying the next-generation supercomputer from a vendor requires a forklift upgrade replacement of the cabinets. The HPC buyer bears the substantial added cost of the new cabinets, and the “old” cabinets often have to be trucked off at some cost to a landfill, with negative environmental consequences.

Environmental issues have inaugurated a new era of “green HPC.” While a true concern for the environment may not always be the top motivating factor for vendors and users who attempt to address these issues, performance improvement is almost always a key motivation. Today, however, the desire for cost savings in power and cooling, as well as facility space, is running a close second.

Definitions

High-Performance Computing

IDC uses the term *HPC* to cover all servers (non-desktop systems), as well as related software, networking, storage, and services used for technical computing tasks. HPC servers can be priced from a few thousand dollars to hundreds of millions of dollars each. IDC uses the terms *HPC server* and *technical server* synonymously.

Compute Density

In an HPC context, IDC defines *compute density* as the amount of peak processing power (floating point operations per second, in most cases) that can be fitted into a given volume of space. It can be measured in FLOPS per square feet, or FLOPS per cubic feet.

Technical Servers and Clusters

In the context of HPC, IDC uses the term *technical server* to refer to a technical computer larger than a desktop system. IDC defines a *cluster* as a set of independent computers (typically servers) combined into a unified system through systems software and networking technologies.

Benefits

The major benefits of "green" HPC systems with lower power and cooling requirements, higher compute densities, and lower equipment disposal frequency and costs are as follows:

- ***Higher Performance.*** Higher compute densities provide more peak processing power per square foot/meter (floor space) or cubic foot/meter (datacenter volume). A related advantage of higher compute density is the increased potential for running problems faster "locally" (local processor, local node, or local system), rather than having to traverse slower, performance-degrading connections to less-local levels of the system architecture hierarchy. Reducing power and cooling costs, and avoiding the need to purchase new cabinets with each system upgrade, can also contribute to higher performance by freeing up budgetary funds to purchase additional processing power.

- ***Higher Productivity.*** Higher performance in turn sets the stage for higher productivity. Completing more work in a given timeframe is important for the efficient use of highly paid research and engineering personnel. Being able to run more problems at larger sizes and higher resolutions, and running more iterations, such as in parametric studies, can also produce superior, more-competitive research solutions and commercial products.

- ***Lower Total Cost of Ownership (TCO).*** Because power and cooling costs are a large and increasing portion of operating expenses, reducing these costs can have a major beneficial impact on TCO. Minimizing the frequency of cabinet purchases and disposals can also benefit TCO.

- ***Smaller Carbon Footprint.*** Using less electricity for power and cooling, and less landfill space for equipment disposal, reduces an organization's carbon footprint.

Trends

- ***Return to Higher Densities.*** Pursuing the twin goals of performance improvement and environmental cost savings, in recent months a number of hardware system vendors have introduced higher-density versions of cluster products. HPC clusters, whose distributed, loosely coupled architectures diverged from the historical HPC design trend toward higher compute densities, are at least in some cases aggressively pursuing higher densities today under the pressure to increase performance and decrease environmental costs for users. Meanwhile, leading vendors of non-cluster HPC products have continued their uninterrupted pursuit of higher densities. These products remain the density leaders.

- ***Heterogeneous Processing.*** HPC accelerators and other alternative processors are an increasingly important strategy for augmenting the stalled single-threaded performance of standard microprocessors. But accelerators (some more than others) can also save on electricity for certain applications by exploiting more parallelism at lower clock periods to deliver higher performance. Examples of alternative processors include FPGAs, GPGPUs, vector and Cell processors. Tuning down the clock speeds of standard microprocessors can also conserve electricity, but more processors and more parallel decomposition may then be needed to achieve a given performance level.

- ***Increasing Carbon Footprints.*** Initiatives to boost power and cooling efficiency will dampen the growth of HPC carbon footprints, but these footprints will nevertheless continue to trend upward. This is because of the HPC community's insatiable desire for more performance, in combination with the ability to buy more and more peak performance within fixed budgets.

- ***Migration to Areas with Cheaper Energy.*** With energy prices increasing and no lasting relief on the horizon, there is a growing trend toward locating the most powerful computing facilities in areas with comparatively cheap energy (commercial example: Google; HPC example: Oak Ridge National Laboratory). This trend could eventually cause top scientific talent to concentrate in the same areas to take advantage of the best computational tools.

Cray Raises The Green Ante

Cray Inc. is a linear descendent of Cray Research, the iconic company often described during its lifetime (1972-1996) as the leading provider of supercomputer packaging and cooling. Since its founding in 2000, Cray Inc. ("Cray") has maintained and advanced those historic competencies. In particular, the company has worked to provide fast proprietary interconnects in order to offer systems that are more balanced (bytes/flops ratios) than commodity clusters and competing products that are outfitted with interconnects available in the open market.

Higher-bandwidth, lower-latency interconnects command a price premium, but compensate for that by enabling communications-bound problems to be run with fewer processors. Compute density is not an absolute term. Although the compute densities of Cray systems may differ little from those of competing products in relation to *theoretical peak performance*, Cray systems exhibit strong compute density where *sustained performance* on communications-intensive codes is concerned. Most datacenter workloads include a broad mix of applications, some of them communications-bound. This means that in some fraction of HPC sites, datacenters could handle their workloads with fewer Cray processors and cabinets. Fewer Cray processors and cabinets can translate into significant dollar savings in some cases, and in nearly all cases it will translate into savings in costly facility space.

With its Cray XT multi-generational product line, the company is providing multi-generational cabinets. The progenitor of this product line, the “Cray Red Storm” system at Sandia National Laboratories, has already undergone multiple processor upgrades without a forklift cabinet replacement, as has the large Cray supercomputer system at Oak Ridge National Laboratory. Minimizing the frequency of cabinet replacements makes system upgrades considerably less expensive over time (i.e., TCO) and reduces the environmental impact of sending old equipment to landfills.

Cray High-Efficiency (HE) Cabinet

The next-generation Cray XT5 supercomputer series will feature the new HE cabinet with the company's ECOphlex technology, with general-availability shipments starting in late 2008 and testing at customer sites occurring before then. One early installation will be the DOE petascale system headed for Oak Ridge National Laboratory.

The HE cabinet with ECOphlex technology is not just another blade enclosure. With it, Cray aims to substantially advance the company's Green heritage and capabilities. Here are the particulars:

- Within the same dimensions as the existing Cray XT4 cabinet (24" wide by 36" deep by standard 42U high), Cray says the HE cabinet will immediately provide multiple times the Cray XT4 compute density: up to 24 blades, 192 sockets, and several thousand cores. The HE cabinet and its ECOphlex technology are designed to handle additional Moore's Law-driven growth in the number of cores per socket as the multicore/many-core trend advances.
- The cabinet is designed to fit into a wide range of air-cooled and liquid-cooled environments. For air cooling, a single high-efficiency axial turbofan that draws cooling air from the floor space and exhausts it from the top of the cabinet is sufficient to support the HEUs. The HE cabinet fan is an improved version of the one that debuted with the Cray XT3 system and boasts a minimum maintenance interval of 7.5 years.
- For clients needing more than air cooling, the HE cabinet supports a flexible approach to liquid cooling. Cray claims that the cabinet's phase change evaporative cooling, which uses the inert gas r134a, is more than 10 times as efficient in removing heat as a water coil of similar size (see Figure 1). Cray also says its Green HE cabinet can greatly reduce the number of air handlers a facility needs and alleviate or eliminate the need for datacenter chillers, resulting in meaningful energy and TCO savings over time.

Figure 1

Phase Change Evaporative Cooling in Cray HE Cabinet

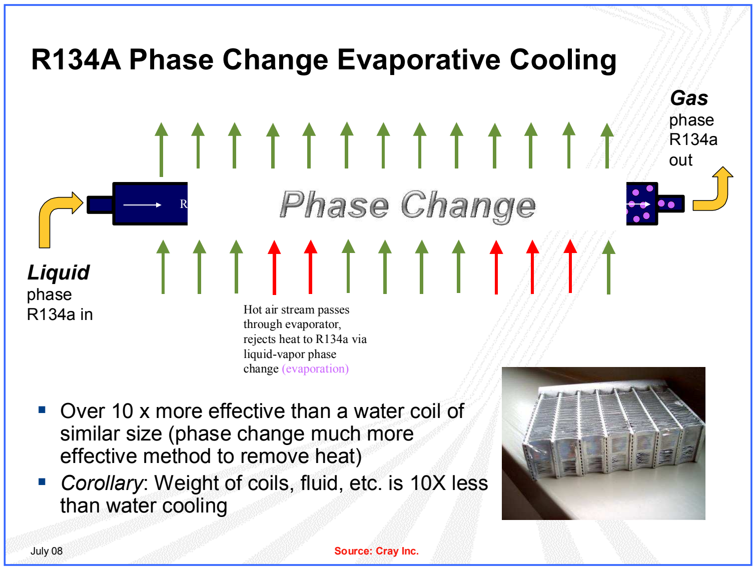

R134A Phase Change Evaporative Cooling

Gas
phase
R134a
out

Liquid
phase
R134a in

Hot air stream passes
through evaporator,
rejects heat to R134a via
liquid-vapor phase
change (evaporation)

- Over 10 x more effective than a water coil of similar size (phase change much more effective method to remove heat)
- *Corollary*: Weight of coils, fluid, etc. is 10X less than water cooling

July 08

Source: Cray Inc.

Source: Cray Inc., 2008

- The HE cabinet has noise reduction components to alleviate another environmental concern in datacenters.
- Cray has also invested in AC/DC power rectification with a 90%-plus efficiency rating. The HE cabinet provides power distribution for less power loss (52V instead of 12V). These power supplies are capable of supporting much higher-power processors than any Cray is using today.

Opportunities

- ***Increase Cray XT5 Product Sales.*** Power and cooling has become one of the top issues for HPC users. Cray's traditional “green” strengths in compute density ("packaging") and cooling, augmented by the technical advances and environmental flexibility embodied in the new HE cabinet, create the potential for additional sales of the HE-equipped Cray XT5 product line beyond the Cray XT4 install base starting in late 2008. In addition, this new approach may remove a potential sales roadblock that some sites could have with Cray, viewing Cray as not “green” enough.
- ***Further Boost Sales by Adding Key ISV Applications.*** The Cray XT product line has achieved notable success at large-scale government sites and university HPC centers, but has not been adopted by industry because key ISV applications have not yet been optimized for this product line. By optimizing a short list of key ISV codes, Cray can establish the basis for selling the Cray XT5 product into select industries.

Challenges

- ***Increasing the TAM for Cray Products.*** Cray's "Green" advances (e.g., the new HE cabinet) may boost sales of the next-generation Cray XT5 product line, especially if key ISV applications are also made available. But this product will still be constrained to the IDC "supercomputers" category for HPC systems priced at $500,000 or more. To accelerate revenue growth, Cray should consider making systems available below this price range.

- ***Marketing and Maintaining Cray's Differentiation.*** HPC vendors have caught the Green fever and are rushing to position their environmental attributes in the market place. In the midst of this marketing noise, where fluff and substance can be hard to distinguish, it will take a concerted effort to hammer home Cray's Green differentiation. And because leading competitors are also striving to make real environmental progress by improving their compute densities, and their power and cooling efficiencies, Cray will need to continue its Green R&D efforts to maintain differentiation.

Conclusion

High-performance computing was born in an era of cheap electricity. HPC system vendors pursued higher compute densities, and better power and cooling efficiencies, in order to boost performance. Clusters reversed the trend toward higher compute densities with their more loosely coupled architectures while vendors of non-cluster HPC systems, including Cray, continued to advance densities.

The recent sharp spike in energy and facility costs has combined with growing environmental concerns to launch a new era of "green HPC." Today, vendors pursue improvements in compute density, and in power and cooling efficiency, not only to boost performance but also to reduce energy and facility costs. The Green HPC trend plays to Cray's historic strengths in packaging and cooling, and to Cray's more recent practice of providing multi-generational system cabinets that minimize the frequency and cost of equipment disposal.

Later in 2008, Cray will raise the ante on Green HPC by introducing its high efficiency (HE) cabinet with ECOphlex technology, in conjunction with the firm's next-generation Cray XT5 product line. The HE cabinet, with its phase change cooling, has the potential to be a truly "green" system by reducing or eliminating the need for datacenter air handlers and chilled water. IDC believes that the HPC market's twin goals of greater sustained (as opposed to peak) performance and lower environmental costs constitute long-term, increasingly important trends.

To the extent that Cray can continue advancing compute density, power and cooling efficiency, and its multi-generational approach to cabinets, the company has an opportunity for expanded success. Cray can further extend this success into industrial markets by making key ISV codes available on its systems.

ABOUT THIS PUBLICATION

This publication was produced by IDC Go-to-Market Services. 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.

IDC Asia/Pacific, 80 Anson Road, #38-00, Singapore 079970 P.65.6226.0330 F.65.6220.6116 www.idc.com.