

Could Power & Cooling Costs Spur a Scientific Migration?

Recent research shows HPC sites plan expansion despite growing concerns

Steve Conway



A decade ago, power and cooling didn't make it onto the top 10 list of issues HPC data centers said they were facing. Today, power and cooling consistently ranks among data center managers' top two or three challenges. What's changed?

- The rapid growth in HPC system sizes (including memory subsystems) has elevated energy requirements. Today's largest HPC data centers consume as much electricity as a small city, and their multi-petascale and exascale plans promise to devour even more.
- Energy prices have risen substantially above historic levels, although prices have moderated from their 2008 highs.
- A third element in this "perfect storm" is the challenge of making HPC processors and data movement more energy-efficient without overly compromising performance — the holy grail of HPC. As Intel Labs' John Gustafson has said, "HPC users are not tree huggers." They tend to use energy savings not to reduce overall energy use or enhance the bottom line, but to buy more gear.
- HPC data center power and cooling developments are occurring at a time of growing sensitivity toward carbon footprints and global climate change.
- Finally, some of the biggest HPC data centers worry that their local power companies may balk at fully supplying their future demands. One such site, with a 250-megawatt data center on the drawing board, may need to go off the grid and build a small nuclear reactor.

A worldwide study IDC conducted a year ago for DICE, Avetec's HPC Research Division, surveyed managers representing more than 200 HPC data centers on the topic of power, cooling and facility space (www.diceprogram.org/reports/request_power_and_cooling.shtml).

And a companion IDC study for Avetec found that, when it comes to power and cooling, HPC and commercial data centers are more alike than different.

Here are the key findings of the HPC data center study:

1. Nearly all the HPC sites (96 percent) considered "green" design criteria important for their HPC system and data center planning process. They described steps they took to make their HPC resources and operations "greener," including data center flow analysis, hot-aisle cold-aisle containment, moving to higher voltage distribution systems, more regular maintenance schedules, the use of "free" cooling, and the purchase of liquid-cooled systems.
2. Two-thirds of the sites were planning to expand or build new HPC data centers. Most also had budgets in place to upgrade their power and cooling capabilities, with the average budget for power and cooling upgrades amounting to \$6.87 million.
3. Half of the sites planned, or had already begun, to distribute their HPC resources to multiple locations. The majority (63 percent) had begun distributing their resources among multiple buildings at a single site, while a smaller number (37 percent) were distributing their resources either regionally or nationally.
4. Power and cooling infrastructure limitations were the biggest barriers to increasing available HPC resources. The most important constraints impeding the sites from expanding their HPC resources were supplying enough additional cooling and providing sufficient additional power. At government sites, funding limitations were equally important.
5. Liquid cooling was the alternative approach being considered most often by the user sites. A fair number of sites expected to maintain existing air-based cooling methods, but departures from status quo tended toward increased adoption of water and other liquid-cooling technologies.
6. Most sites (61 percent) had analyzed their data centers' heat flow and power consumption. Their HPC server systems consumed the vast majority of their power and cooling (90 percent), followed by their HPC storage systems (nine percent).
7. Only about half of the sites (48 percent) paid for power and cooling costs out of their own budgets. Of the government sites, 75 percent paid for power and cooling directly out of their own budgets, whereas only 50 percent of academic sites, and 14 percent of industry sites, did.

8. Most HPC vendors saw power and cooling efficiency as a brake on compute performance. The majority of the vendors agreed that the trend toward more cores and away from greater single-core performance means that the pure pursuit of HPC performance will be tempered by the need for greater power/cooling cost-efficiency.
9. Most sites cited tradeoffs between performance and power efficiency. These included considering non-mainstream processing technologies, such as GPUs and FPGAs, that offer more efficient performance per watt on appropriate codes, but are harder to program; accepting more service disruptions because of shorter upgrade cycles; and staffing to address poorly scaling applications and harder-to-manage equipment.
10. HPC users and vendors differed sharply on the likelihood of game-changing cooling technologies. Just over one-third (36 percent) of the user sites expected game-changing cooling technologies in the next five years. The vendors were much more optimistic than the user sites, with 62 percent of them foreseeing game-changing cooling technologies. The difference implies that users have more aggressive expectations for what constitutes a game-changing technology.

The companion study that included both HPC and commercial data centers revealed that, while many data centers monitor power usage, few actively track power efficiency. Only 31 percent were using specific metrics and/or statistics (e.g., PUE). Most sites (80 percent) were operating under relatively weak internal or external mandates to reduce energy consumption. The management attitude often is that “the cost of electricity is simply the cost of doing business.” For most data center sites, power capacity and fitting into a facility’s power envelope were more important than power consumption or energy efficiency per se. A just-finished IDC worldwide study of HPC user sites reconfirmed these findings.

What’s ahead —a power-based scientific migration?

An important goal of exascale initiatives is dramatically increasing the energy efficiency of HPC systems, ideally so that these systems will consume no more power than today’s petascale computers. That’s a daunting goal to achieve by the latter part of this decade.

Meanwhile, power usage grows apace, and a pattern is already forming in which the largest data centers are being located or expanded in geographical areas where power is comparatively cheap and plentiful. Google set the tone five years ago, by building a vast new data center along the Columbia River near Oregon’s Dalles Dam, with its 1.8-gigawatt power station and relatively inexpensive hydroelectric power. Another prominent example is Oak Ridge National Laboratory, whose considerable power appetite is fed by the Tennessee Valley Authority. The lab’s 50-megawatt data center hosts petascale systems for the Department of Energy and the National Science Foundation, with a NOAA system slated for a petascale upgrade soon and defense agencies reportedly scouting out the lab as a potential HPC system host.

If the pattern of skyrocketing power usage and costs continues, it could spur the migration of top scientific talent to HPC sites with access to enough power to support the most capable supercomputers. The scientific migration patterns could become global, rather than merely national or regional in scope. In 2002, even before the era of expensive power began, the U.S. Department of Energy was concerned about the potential for an intercontinental scientific migration. Reacting to the Sputnik-like shock of Japan’s “Earth Simulator” supercomputer, a DOE

Office of Science paper, “Reasserting Leadership in Scientific Computing,” warned that “The U.S. has lost the lead in climate science research. Other areas of computational science critical to the DOE’s mission are at risk ... U.S. climate scientists have no recourse other than to travel to Japan and form partnerships with Japanese scientists to use the [Earth Simulator].”

The subsequent rise of the power supply and cost issue has elevated the scientific migration issue from a single instance in 2002 (the Earth Simulator) to a more general concern today. In either case, migration decisions will ultimately be driven by the need of top scientists to remain at the forefront of their fields by exploiting the most powerful supercomputers available — and by the need of the host sites to power and cool these mega-systems.

Steve Conway is Research Vice President, HPC at IDC. He may be reached at editor@ScientificComputing.com.

Copyright 2011 Scientific Computing, Advantage Business Media