

Money for Research, Not Energy Bills: Finding Energy and Cost Savings in High Performance Computer Facility Designs

Dale Sartor, Lawrence Berkeley National Laboratory

Mark Wilson, Drewmark Communications

November 2010

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Money for Research, Not for Energy Bills: Finding Energy and Cost Savings in High Performance Computer Facility Designs

By Dale Sartor (Lawrence Berkeley National Laboratory) and Mark Wilson (Drewmark Communications)

Abstract

High-performance computing facilities in the United States consume an enormous amount of electricity, cutting into research budgets and challenging public- and private-sector efforts to reduce energy consumption and meet environmental goals. However, these facilities can greatly reduce their energy demand simply through the energy-efficient design of the facility itself. Using a case study of a facility under design, this article discusses strategies and technologies that can be used to help achieve energy-reduction.

Keywords: High performance computing, HPC, supercomputer, energy, energy-efficiency, Computational Research and Theory Facility, CRTF, University of California, facility design.

The need for high-performance scientific computing (HPC), or “supercomputing” continues to grow, to meet national security, materials design, climate protection, and energy goals, among others. As the need for faster, higher-performance computers rises, the energy needs and costs associated with those machines also rise—sometimes to startling levels. For example, the U.S. Department of Energy (DOE) supercomputer system uses an aggregate \$100 million of energy annually, and that number is rising rapidly.¹ With electricity needs for some future supercomputer designs topping over 100 megawatts (MW) each, the challenge to rein in costs is immense.

This burgeoning energy demand will not only strain already limited research budgets, but also force energy needs to rise at a time when research institutions are focused on reducing energy consumption. Numerous efforts are underway to redesign supercomputers to work more efficiently, and enormous technological advances continue to be made. Technologies that focus on parallel processing and the adoption of design principles from the low-power embedded computer industry can greatly reduce supercomputer energy use, with little compromise in computing performance. For example, proposed specialized Tensilica architecture would achieve a next-generation climate-modeling tool for a task such as 1.5-kilometer resolution climate modeling requiring 2.5 MW (200 petaflop) for the computing infrastructure at an estimated cost of \$75 million, as compared a more conventional x86-based system requiring about 180 MW (5 petaflop) at an estimated cost of \$1.8 billion.²

However, machine efficiency is only part of the solution. Just as residential and commercial buildings are now being designed to improve energy efficiency and comfort, the buildings that house these supercomputers must be designed to provide the cooling they need as efficiently as possible. Moreover, they must do so while ensuring that the design is flexible enough to accommodate generations of future supercomputer designs.

One building that could be setting the standard is the Computational Research and Theory Facility (CRTF) currently being designed for the University of California. The university’s goal is

¹ Mills, Evan. “Sustainable Scientists.” *Environ. Sci. Technol.* 2009. 43, 979.


² Wehner, M, L. Olikar, and J. Shalf. 2008. “Towards Ultra-High Resolution Models of Climate and Weather” *International Journal of High Performance Computing Applications* 22 (2). May. 149–165.

to create a center that is both a model of high-performance computing and a showcase for energy efficiency.

The project is not simply a theoretical exercise, but one of necessity. As one of the premier research entities in the world, the university needs high-performance computing facilities that are capable of meeting formidable computing challenges. However, the university estimates that this one 126,000 square-foot building with offices, a 32,000 square foot computer space, and infrastructure will *more than double* energy use for the entire 4000-person campus. This both complicates the university's vision to reduce its greenhouse gas emissions below 1990 levels by 2014 and potentially siphons money from other priorities to cooling computers.

Setting Goals and Identifying Strategies

Because it is impossible to know the number and type of computers that will inhabit the building over its lifespan, project planners sought to design a structure that incorporates energy efficiency while providing flexibility for a range of future supercomputers and cooling strategies. They established a design target for data center infrastructure efficiency (DCiE, the ratio of IT energy to energy use of the whole facility) at 83%. This target represents a considerable improvement over data centers previously benchmarked by LBNL, with the most efficient previously around 75%, and the current average about 50% (Figure 1). The Power Usage Effectiveness (PUE), another metric for looking at data center energy efficiency, was set at 1.2.³ An arbitrary power budget of 7.5 MW was established for the initial build-out and 17 MW was allotted for the completed build-out.



Of course, setting a target energy efficiency is easy enough; but how to get there? Although efficiencies can be gained from both IT loads and infrastructure loads, the focus here is on infrastructure, including cooling and air flow. The designers looked at a variety of features to identify areas where efficiencies could be found: air versus liquid cooling, air handler size and air management, modular design, cooling options, heat recovery, power distribution, and the local climate. For all of these design points, they recognized that flexibility was essential. As a result,

³ For the PUE metric, “1” represents the benchmark use of the IT equipment, so the closer to that number, the better. With DCiE, the most efficient facility would be 100%, where 100% of the ideal building’s energy use went to the computers, which, of course, would not be feasible.

space and other design considerations were configured so that additional capacity can be added as the facility load is increased. Not only did this approach reduce the facility's first costs, it also allowed components to be sized to better match the load requirements.

Natural Cooling a Big Advantage

The design approach to the cooling system is a perfect example. Located on San Francisco Bay, Berkeley benefits from a natural air conditioner—the cool air that sweeps in from the ocean through the Golden Gate toward the campus. Average temperatures throughout the year range from the mid-forties to the mid-seventies. So the designers chose to take advantage of the area's natural cooling. However, a key design concept for the CRTF is to accommodate many generations of supercomputers over the course of several decades, so air cooling could not be the only design option available. Because there is a general trend in the computer industry toward liquid cooling as power densities increase, it was necessary that they also accounted for future equipment that could be largely or completely liquid-cooled. In addition, the lower-density areas in the facility that house memory and network equipment will likely be air-cooled for a longer period than the scientific computing machines, so the design must enable different types of computer systems in different areas of the building to be cooled with various methods—air, liquid, or a combination of both—and to offer that flexibility with the highest energy efficiency and lowest up-front costs.

Cool Air In, Hot Air Out

The airflow and temperature levels of any building are critical contributors to occupant comfort and productivity. Employees in more than one office have come to blows over the heating, ventilating, and air conditioning (HVAC) settings. However, this situation is even more intense when a large number of your occupants are finicky supercomputers that will simply go on strike if the temperatures get too hot.

Keeping this in mind, the design team examined how air flow and cooling could best be accomplished. The power and energy requirements of the air-handling equipment are determined by two factors: (1) the fan and motor efficiencies, and (2) the system's total pressure drop. These issues are both addressed in the CRTF design by the choice of central air handler units (AHUs). Simply put, the AHUs move air through the HVAC system. With a multi-story, 126,000 square-foot building, the trick is to move the air at the desired rate throughout all parts of the system, without experiencing pressure drop that would reduce that airflow to a rate that cannot meet the diverse needs of the equipment and building's occupants.

As shown in Figure 2, the modular AHU configuration at the facility is at the lowest portion of the building, leaving the expensive raised-floor space above for optimal IT equipment placement flexibility. Because a modular design allows flexibility, each 20-foot-wide bay can be equipped with one or two AHUs. If airflow demand requires that a bay house two AHUs, they will be vertically stacked in the basement area. Each unit is sized at 100,000 cubic feet per minute (cfm), maximum. To enhance energy efficiency, flow is adjusted using variable-speed fans that enable the system to match the airflow with that needed at any given time.

The AHUs pull the cool San Francisco Bay air from the west side of the building, through the AHUs, and into the ductwork that feeds supply air into the four-foot-high raised floor plenum (the chamber that holds the pressurized air) in multiple locations. The air is delivered to the IT

equipment either through a typical cold aisle arrangement⁴ or directly into the bottoms of the racks, depending on the computer design.

The hot air discharged from the equipment is exhausted through exhaust fans located high on the east wall of the high-performance computing (HPC) area, returned to the AHUs through ductwork down the HPC's west wall, or most commonly, through a combination of the two. The modular AHU scheme enables the system to supply the right amount of air to the right place with a minimum of excess capacity.

Of course, given that over 80% of the 17 MW entering the CRTF will be converted to heat by the IT equipment, and with discharge air temperatures of about 100°F, there is significant opportunity to recover and use this heat. The office floors of the building are being designed to high standards of energy efficiency, but they will still require some heating, making heat recovery within the building one obvious opportunity. However, there will be so much waste heat available from the CRTF that it could heat a whole cluster of nearby buildings, replacing the heat now provided by their natural gas-fired boilers. These facilities include laboratories that need high volumes of outside air on a continuous basis, so that cooler Berkeley replacement air must also be heated continuously. The feasibility of these opportunities, which include using such technologies as run-around air-to-water coils and a local district heating system, is being investigated.

Because the cross-sectional area of the AHUs is large, peak face velocities of the filters, cooling coils, and evaporative cooling media is approximately 500 feet per minute (fpm), which is relatively low. *Face velocity* is the air stream velocity as it enters a filter, and pressure drop increases with airflow velocity, so a low face velocity is good. It will be even lower over most conditions at times when the fan speed is reduced.

Supply air from the air handlers is supplied to the plenum via short ductwork, which is designed at 1500 fpm maximum at full design flows. These velocities, combined with careful attention to other internal AHU pressure drops and low pressure drops in the ductwork and air distribution, result in a low total initial static pressure of 1.5" at design flow. This low static pressure saves fan energy while providing sufficient air to move through the system without experiencing pressure drop that would adversely affect performance. However, because of the low pressure, it is essential that filters and evaporative media pads be replaced regularly, before they become clogged, so that they do not contribute to excessive pressure drop.

CFD Modeling Eliminates Guesswork

To fine-tune the design and confirm that the it would perform as intended, Computational Fluid Dynamics (CFD) models were used from the beginning and throughout the process. Using different design scenarios for input, a CFD model can graphically show how the air will move through the system and whether the design will achieve established performance goals. Because of the difficulty designing for high heat loads, it was decided that CFD would be a valuable tool in designing the CRTF. Even with a limited budget, designers can use CFD modeling to evaluate multiple scenarios, enabling them to predict performance capabilities of each design and optimize them, rather than just taking educated guesses. The technology helps designers move away from overdesigning cooling capacity to make sure that it can to cool the higher heat loads to a far more nuanced approach that saves both space and initial costs for equipment that most likely won't be used.

⁴ A cold-aisle arrangement consists of open tiles that allow air to cool equipment without allowing it to mix with warm air in the facility, greatly increasing the system's efficiency.

The CFD modeling started out with no preconceived assumptions or specific design, and the team went through a number of permutations as they searched for the best option. One early design was ruled out because the modeling results showed that it would provide poor pressure distribution. To keep construction costs and complexities to a minimum, the team converged upon an off-the-shelf design solution using a conventional raised floor. Modeling showed that it would be able to meet the project's cooling targets without excessive pressure drops.

Figure 3 illustrates one configuration evaluated in the CFD modeling.




Figure 2. Building section showing AHU located below computer floor, exhaust air path (return is to left and down to AHUs), and office floors above.




Figure 3. Computational Fluid Dynamics study was performed to optimize under floor plenum design. (Source: ANCIS Incorporated)

Zoning Fine-Tunes the Cooling

For really efficient cooling, however, it's necessary to combine efficient airflow with directed airflow. To reduce fan energy and ensure adequate cooling for the high-intensity computing equipment, the hot and cold air streams must be separated. To accomplish this, the design provides for physical separation of hot and cold sides of the IT equipment. Of course, the detailed design cannot be finalized until the IT equipment set is defined, but some general design features can be incorporated from the outset. For example, by providing a high ceiling area, hot exhaust from the IT equipment is collected and flows to the building exhaust or is recirculated for cooling or mixing with outside air.

However, separating the hot and cold sides of the equipment is only part of the strategy. Since the environmental condition requirements and heat intensity of the IT equipment may vary widely, it is also necessary to zone the HPC area, both above and below the raised floor. Zoning allows the airflow to match the environmental requirements of each area's particular equipment. For example, a central data storage area is planned that will operate at significantly lower load density compared to the supercomputers. Zoning will allow this area to be operated at much different conditions than the main computing areas. Also mixing liquid- and air-cooling solutions could require varying amounts of air, so providing zoning capability will optimize cooling for those various conditions.

Liquid Cooling Looks to the Future

Although it is not anticipated that liquid cooling will be used in the initial CRTF configuration, it is designed to accommodate the distribution of cooling water for direct or indirect use at or in the racks or computer itself. A four-pipe distribution scheme will include chilled water from the chiller plant (using water-cooled, electrically driven centrifugal chillers, which are more efficient than air-cooled chillers) and closed-loop, treated cooling water from the cooling towers (via plate-and-frame heat exchangers, which have a high heat transfer efficiency, and therefore provide efficient cooling). Mixing valves will allow individual computing systems within the facility to use 100% chilled water, 100% treated water, or anything in between, as needed to satisfy the temperature requirement for water entering the system. Chilled water and treated water temperature setpoints and reset schedules will be established to optimize energy-efficient operation. Lawrence Berkeley National Laboratory researchers are working with computer manufacturers to increase the required water temperature for cooling. The warmer the water, the

more “free” cooling can be derived from the cooling tower only, without the use of the chiller. Further, chiller efficiency is increased when the chilled water set point is warmer, and there is reduced change of condensation in the data center.

Until the treated water is needed, it will not be fully installed, but will be accommodated by headers, valves, blank-off plates, and space for pipe runs. The chilled water system will initially run only to the AHUs, but taps with valves and blank-off plates will be installed also for future water cooling requirements.

Meeting Today’s Load

The CRTF’s anticipated electricity load growth from 7.5 MW to 17 MW over the life of the facility will more than double its energy use, so the building’s support systems need to grow with it. However, load growth throughout the life of the facility is not always a given; the demand for electricity will vary both upward and downward as computing systems are added, changed, turned on for use, and shut down for maintenance. Weather variations can also result in diurnal and seasonal load changes. For the most efficient operation, all of the facility’s systems must be able to meet load variations in a way that provides uninterrupted service while modulating efficiently.

To make that happen, the cooling plant will be modular, and all of the significant loads in the plant and system (such as tower fans; chiller compressors; chilled, tower, and treated water pumps; and AHU and exhaust fans) are all designed with variable-frequency drives. Part-load curves will be integrated into the building automation system so that overall energy and power use are minimized at any combination of cooling load and outdoor conditions.

How Cool is Cool Enough?

When designing energy-efficient cooling for an HPC center, choosing the environmental conditions (temperature and humidity) acceptable for supercomputer operations can be tricky. Some computers are likely to run optimally at lower temperatures than others, so it is necessary to be able to provide that lower temperature, at least for that equipment. On the other side, however, providing even a degree of cooling more than necessary throughout the entire life of the facility could result in enormous energy and economic costs.

The project team debated whether the environmental conditions recommended by the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) could be used as a design basis for the facility, since some of the supercomputers on the market required more stringent conditions. To resolve whether ASHRAE-recommended ranges could be specified, LBNL held a workshop with all of the major supercomputer vendors. Ultimately, it was decided that all of the vendors would agree to using the recommended ranges. Subsequent to this meeting, the ASHRAE data center committee voted to broaden the recommended ranges even further. With these assurances, the team agreed to use a maximum of 77°F as the design inlet temperature of the IT equipment. This temperature suited the Berkeley climate well, since the facility will be able to use outside air for cooling to temperatures lower than 77°F during most of the year. A broad design humidity range was also established at 30% to 60% relative humidity at the inlet to the IT equipment.

Natural Cooling and Moist Air Yield Big Savings

After evaluating the data, it was decided that the cooling system could meet temperature and humidity operational requirements through four different modes: (1) a mix of outside and return air, (2) direct evaporative cooling, (3) direct evaporative cooling and chilled water coil, and (4) chilled water coil alone.

For over 90% of the year, the indoor conditions can be met by mixing outside and return air. The psychrometric process of this mode is shown by the arrows in Figure 4. Direct evaporative cooling (with a mix of return air as needed) brings the humidity into the proper range when outdoor conditions are too dry, as does the condition where a combined use of direct evaporative cooling and the chilled water cooling coil is indicated. Approximately 500 hours per year require the chilled water coil alone.

Evaporative cooling combines water evaporation from a wet material with airflow to cool the air. By using a wetted media and the heat in either the outside air or the return air to evaporate the water, the CRTF cools the air and avoids the energy use that would be necessary to power steam or infra-red humidifiers. Use of onsite evaporative cooling improves efficiency.

Another cooling method—a direct-spray system—was considered for lower pressure drop, but the extra pressure drop caused by the wetted media did not justify the cost of the reverse-osmosis or deionization system that would be required for the direct spray system’s make-up water.

Other alternate humidification strategies were explored, such as floor-mounted and plenum-located humidifiers, but these were rejected due to concerns that they would not provide uniform humidity distribution and because there was a preference to keep the plenum and floor clear for maintenance, accessibility, and flexibility. It must be noted that the facility will install multiple supercomputers (with a variety of rack configurations) at a given time which will be replaced by new-generation super computers every five years or so.

When water-based IT cooling is implemented at the CRTF, close-approach cooling towers and plate-and-frame heat exchangers will be used to supply as much of the cooling as possible without operating the chillers. *Close-approach* means that the temperature of the water cooled by the tower gets close to the wet bulb temperature of the outside air. It is anticipated that most of the cooling will be provided without using the chillers.

Further Savings Through the Electricity Supply

Beyond the HVAC savings, data center energy and power savings can also be found in the electrical distribution system. The CRTF is fortunate in that most of the IT equipment will be used for scientific computing, so unlike IT equipment that services bank transactions or other time-sensitive services, it does not need to be supported by uninterruptible power supplies (UPS) and standby generation. Of the initial computer load of over 6 MW, only 500 kilowatts (kW) will be on UPS power. The topology of the UPS at the center is still being decided, but high efficiency across a wide range of loading is a prime selection criterion.

The CRTF will also be able to minimize distribution losses through the use of IT equipment that takes 480-volt power directly at the racks. Thus, for the main computing load, there is only one voltage transformation (from the 12-kilovolt site distribution system to 480 volts). Using 480 volts rather than the normal 208-volt distribution means further savings due to lower currents in the building wiring.

Commissioning Helps to Ensure Success

Commissioning is often conducted after a building is completed, to test and verify that key building systems perform to their designed level. For this project, a commissioning agent was on the design team, helping to choose a controls vendor, providing review of proposed control sequences, helping to develop the measurement and verification plan, and providing general design review, such as conceptualizing the chiller selection and associated turn-down strategies.

Bringing a commissioning agent on-board early in the process greatly increases the chances that they system will perform as designed once it is up and running.

Energy Efficient Data Centers Begins to Emerge

The University of California's CRT facility is designed to be a model for energy efficiency as well as a leading scientific computing facility. By employing best practices identified in previous research, the project team was able to design a very energy-efficient building with a DCiE above the design goal of 83%. And by combining use of the higher end of the ASHRAE-recommended temperature range with a favorably cool climate, very little cooling will need to be provided by chillers. And in the event that liquid cooling becomes the preferred cooling solution, even further energy and power efficiency gains are possible.

Annual Psychrometric Chart of Oakland, CA
(relative humidity lines are stepped by 10%,
wetbulb lines by 10 degrees F)




Figure 4. Psychrometric data for Oakland, California (adjacent to Berkeley). The dashed line of 60°F–77°F drybulb temperature and 30%-60% relative humidity is the designed supply air condition. The arrows represent mixing of outside air and return air.

However, the CRTF is not the only HPC center exploiting better building design to reduce energy and costs; other facilities are being designed nationwide. For example, the U.S. Department of Energy's National Renewable Energy Laboratory is also designing a new supercomputer facility with these goals in mind. And because they're working in a colder, drier climate than that of Berkeley, they will be able to cool it without a single chiller. In fact, the request for proposals for the facility's construction specified that mechanical cooling would not be allowed and required a peak PUE under 1.1. The building will be a combination lab, office, and data center, so waste heat will be used for preheating ventilation air and heating the office.

Construction on the University of California's Computational Research and Theory Facility, which is registered under the LEED NC 2.2 Silver designation, is slated to begin in 2010.

Acknowledgments

The authors would like to thank the following organizations for information on the work described in this article: Perkins+Will (architects), Arup (mechanical and electrical engineering), ANCIS (CFD modeling), and the National Energy Research Scientific Computing Center (NERSC) and Environmental Energy Technologies Division (EETD) at Lawrence Berkeley National Laboratory (LBNL). Special thanks to LBNL's Steve Greenberg, Amit Khanna, William Tschudi, and Evan Mills. This work was supported by the Assistant Secretary of Energy Efficiency, Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Program, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.