Case Study

Energy Efficiency Opportunities in Federal High Performance Computing Centers

Summary

The energy use at six Department of Defense High Performance Computing (HPC) data centers was assessed in 2011. Table 1 provides a breakdown of the potential energy and cost savings and average payback periods for each data center. The total energy saved was estimated at more than 8,000 Megawatt-hour (MWh) with an annual greenhouse gas (GHG) emission reduction of 7,500 tons. Energy cost savings of approximately $1,000,000/yr were possible through energy efficiency measures (EEM) that had an average payback period of less than two years.

The individual data centers contained a variety of IT systems and cooling systems. Server rack densities were measured from 1kW (non HPC racks) to 25kW. In a conventional data center with rows of sever racks, separation of the rows into hot and cold aisles followed by containment of the hot or the cold aisle can result in substantial reductions of cooling energy. This arrangement does not fit well with some of the HPC data centers because of the integrated local cooling and different cooling air path. The most effective EEM was increasing the data center temperature which made it possible to save cooling energy by raising the chilled water supply temperature.

Sites	Payback years	Annual Energy Saving MWh	Annual GHG Emission Reduction Ton
Site 1	2.9	1,090	940
Site 2A	2.5	3, 060	2,750
Site 2B	2	2,520	2,300
Site 3	2	937	860
Site 4	2.7	520	480
Site 5	0.5	1,000	950

Table 1- HPC Sites Potential Energy/GHG Savings

Overview

The DoD Office of High Performance Computing Management Program and the Federal Energy Management Program (FEMP) jointly funded assessments by personnel from the Lawrence Berkeley National Laboratory (LBNL) because of an earlier success in identifying EEM at an HPC data center in Hawaii. This case study includes the results of the assessments, lessons learned, and recommended EEMs. Table 2 shows the IT equipment for each site with their corresponding cooling systems. High power density is the most common feature of these equipments. Power used by each rack averaged up to 10-25kW. Another common feature is the use of 480V power without intervening transformation. Since with each transformation there is power loss and heat rejection, eliminating several power transformations, reduced the need to remove the generated heat. Many of the systems also have self-contained cooling using local heat exchangers in the form of refrigerant cooling or rear door heat exchangers. The advantage of this feature is that the heat removed is very close to the source enabling heat transfer to be done at higher temperatures. This by itself can make the cooling very efficient since very cold chilled water will not be needed. The result is increased hours of air side or water side economizer operation and reduced compressor operating hours. Most of the HPC server racks are different from the usual server racks in the way cooling air flows through them. An example is the Cray unit with cooling air flowing from the bottom to the top, rather than horizontally.

Assessment Process

The first step of the assessment was to baseline the environmental conditions in the HPC data centers vs. the ASHRAE recommended and allowable thermal guidelines. LBNL installed a wireless monitoring system. The following environmental sensor points were installed throughout the data center: temperature sensors at all computer room air handlers (CRAH) or/and computer room air conditioners (CRAC) supply and return. Supply thermal nodes were installed under the raised floor just in front of the unit. Return nodes were placed over the intake filter or were strapped to the beam on the top of the intake chimney. For CRAH/CRAC with their chimneys extended to the data center ceiling, thermal nodes were placed over the ceiling tile very close to the CRAH/CRAC intake. Humidity was measured by the same thermal node. Temperature sensors were placed on front of the rack near the top, at the middle and near the base. Other temperature sensors were installed on the back of the rack, in the sub-floor, and on every third rack in a row. Pressure sensors were located throughout the data center to measure the air pressure differential between sub-floor supply plenum and the room. The same approach was used for data centers with no raised floor by installing the sensors in the ductwork. The wireless sensor network continuously sampled the data center environmental conditions and reported at five minute intervals. Power measurements (where it was possible), power readings from equipment (Switch gear, UPS, PDU,…), and finally power usage estimation otherwise, facilitated the calculation of power usage effectiveness. Generally, Power use information is critical to the assessment. IT power should be measured as close as possible to IT equipment. More recent power supply units measure the power and can communicate that to a central monitoring system. If this is not possible then communication or manual reading of PDU or at a higher power chain level, UPS power output can be assumed as IT power usage. Loss in the electrical power chain can be measured or can be estimated while considering how efficient the UPS units are loaded. Lighting power usage can be done by counting the fixtures or by an estimation based on the wattage per square foot. Cooling load measurement including chiller power usage can be a challenge. If a reading from the control panel is not possible, then estimation based on theoretical plant efficiency can work. While spot measures at CRAH units with constant speed fans provided sufficient information, it is more advantageous to continuously monitoring CRAC units because of the variable power usage of the compressors.

The next step was to analyze the data and present the result of the assessment. The collected data provided empirical measures of recirculation and by-pass air mixing, and cooling system efficiency. The assessment established the data center's baseline energy utilization and identified the EEMs and their potential energy savings benefit. In some of the centers, a few low cost EEM were completed and their impacts were observed in real time.

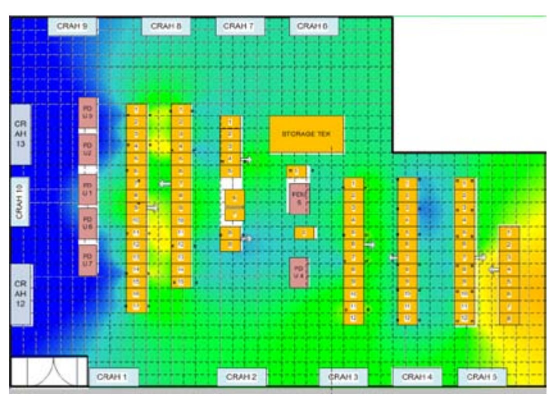

Figure 1: Data Center Thermal Map, After Containment

For instance, Figure 1 is the thermal map of one the data centers that was assessed. The map shows the impact of partially contained hot aisle. The major power usage comes from two rows on the left. The containment helped to isolate hot air to some extent. Figure 2 shows the impact of raised CRAH supply air temperature. The result is a more uniform and higher temperature within the data center space. Chilled water temperature then was increased which reduced energy usage by the chiller plant. In this case, the IT equipment cooling air intake was in front and exhaust at the back of the rack.

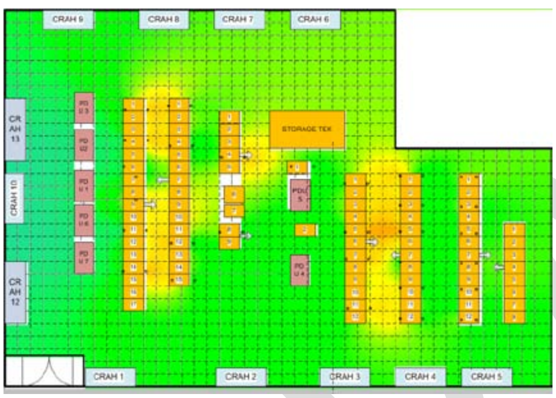

Figure 2: Data Center Thermal Map, After Raised Temp

Challenges

The LBNL technical team was challenged by these energy assessments because many of the HPC data centers employed different cooling types and configurations in the same room. For example, while one area used building HVAC, another system had rack based air cooling using refrigerant coils and a third system had rear door cooling heat exchangers using central chilled water system. The air intake to the IT equipment posed another challenge. While in conventional server racks air enters horizontally at the face of the rack and exits from the back, in some HPC systems (e.g. Cray) air enters from the bottom (under the raised floor) and exits at the top. With vertical air displacement the containment of aisles by itself does not impact the cooling energy use. The same holds true with in-row and in-rack cooling systems, or rear door heat exchangers, since the exhaust air can be as cold the intake air. This cold air mixes with the hot air. The temperature difference across the air handler's coils is reduced which results in inefficiencies in the cooling systems.

Site	IT System	Cooling Type
Site 1	Cray XE6	Refrigerant Cooled using local water cooled refrigerant pumps
	SGI Altix 4700	Rear Door Heat Exchanger using central chilled water
Site 2A	SGI Altix Ice 8200	Air cooled
	Cray XT5	Air cooled with integrated fan, bottom air intake, top exhaust
Site 2B	SGI Altix Ice 8200	Air cooled, ready for rear door heat exchanger using central chilled water
	Linux Woodcrest	Air cooled
Site 3	Cray XE6	Refrigerant Cooled using local water cooled refrigerant pumps
	Cray XT3	Air cooled with integrated fan, bottom air intake, top exhaust
	Cray XT4	Air cooled with integrated fan, bottom air intake, top exhaust
	SGI Altix Ice 8200	Rear Door Heat Exchanger using central chilled water
Site 4	IBM Cluster 1600 P5	Air cooled
	Cray XT5	Air cooled with integrated fan, bottom air intake, top exhaust
	IBM Power 6	Air cooled
Site 5	Dell Power Edge M610	Air cooled

Table 2- Computers and Cooling Types in HPC sites

EEMs for HPC Data Centers

For the presentation of results, LBNL personnel created packages of EEMs categorized based on their cost and simple payback. Typical EEMs covered in each package with its rough cost estimate are described as follows:

Air Management Adjustment Package

- Seal all floor leaks
- Rearrange the perforated floor tiles locating them only in cold aisles for conventional front to back airflow; solid everywhere else
- Contain hot air to avoid mixing with cold air
- Seal spaces between and within racks
- Raise the supply air temperature (SAT)
- Disable humidity controls and humidifiers
- Control humidity only on makeup air only
- Turn off unneeded CRAH units

Typical Cost of this package is $100/kW of IT power. A typical simple payback is ~1 year.

Cooling Retrofit Package

- Install variable speed drives for air handler fan and control fan speed by air plenum differential pressure
- Install ducting from the air handler units to the ceiling to allow hot aisle exhaust to travel through the ceiling space back to the computer room air handling units
- Convert computer room air handler air temperature control to rack inlet air temperature control
- Raise the chilled water supply temperature thus saving energy through better chiller efficiency

Typical Cost of this package is $180/kW of IT power. Typical simple payback is 2.5years.

Generator Block Heater Modification Package

- Equip generator's block heater with thermostat control
- Seal all floor leaks
- Reduce the temperature set point for the block heater

Typical Cost of this package is $10/kW of IT power. A typical simple payback is 2.5 years.

Full Lighting Retrofit Package

- Reposition light fixtures from above racks to above aisles.
- Reduce lighting
- Install occupancy sensors to control fixtures

Typical Cost of this package is $15/kW of IT power. A typical simple payback is 3 years.

Chilled Water Plant Package

- Install water side economizer
- Investigate with HPC equipment manufacturer whether chilled water supply temperature to their heat exchanger can be increased
- Run two chilled water loops, one for those equipment with lower chilled water temperature requirement and the other for those with rear door heat exchangers
- Install VFD on pumps
- Purchase high efficiency chillers and motors if renovation or capacity increase is planned

Typical Cost of this package is $400/kW of IT power. A typical simple payback is 3 years.

Additional Opportunities

1. The use of water cooling can lead to major reductions in power use due to the higher energy carrying capacity of liquids in comparison to air. Higher temperature water can also be used for cooling which saves additional energy. This strategy can work very well with rear door heat exchangers. This strategy can also be integrated with the use of cooling towers or dry coolers to directly provide cooling water thus bypassing the compressor cooling.
2. Refrigerant cooling systems can be specified to operate using higher supply water temperatures than the typical 45-50^{o}F. There is already equipment that requires 65^{o}F. This will increase compressor-less cooling hours.

Site	Recommended Packages	Implementation
Site 1	Air Management Adjustment	Partially done
	Cooling Retrofit	Potential
	Lighting Retrofit	Partially done
	Chilled Water Plant	potential
Site 2A	Air Management Adjustment	Potential
	Cooling Retrofit	Potential
	EG Block Heater	Potential
	Lighting Retrofit	Potential
	Chilled Water Plant	Potential
Site 2B	Air Management Adjustment	Potential
	Cooling Retrofit	Potential
	EG Block Heater	Potential
	Lighting Retrofit	Potential
	Chilled Water Plant	Potential
Site 3	Air Management Adjustment	Partially done
	Cooling Retrofit	Partially done
	EG Block Heater	Potential
	Lighting Retrofit	Potential
	Chilled Water Plant	Potential
Site 4	Air Management Adjustment	Partially done
	Cooling Retrofit	Partially done
	EG Block Heater	Potential
	Lighting Retrofit	Implemented
	Chilled Water Plant	Potential
Site 5	Air Management Adjustment	Partially done
	Cooling Retrofit	Partially done
	EG Block Heater	Potential
	Lighting Retrofit	Implemented
	Chilled Water Plant	Potential

Table 3- EEM Packages status for HPC data centers

Lessons Learned

- The main barrier for the increasing supply air temperature was the IT equipment and refrigerant pumping system maximum temperature requirements. In one of the sites, the refrigerant pumping system required chilled water supply temperatures of as low as 45°F.
- Granular monitoring enables temperature measurements at the server level allowing the implementation of some low cost simpler EEMs during the assessment without interrupting the data center operation.
- Chilled water supply temperature setpoint optimization can result in large energy savings especially for the cooling systems utilizing air cooled chillers.
- Some conventional approaches (e.g. sealing the floor) are applicable to HPC data centers but some commonly applied EEMs to enterprise data centers such as hot/cold aisle isolation are not suitable to HPC data centers where the racks are cooled internally by refrigerant coils or rear door heat exchangers.
- Hoods are installed as is shown in Figure 3, to direct cold air to Cray power supply units, an example of a site engineer's remedy to prevent mixing of cold and hot air. This mixing is a universal problem and the use of hoods will work in similar situations where IT equipment has an unusual configuration. In this case, the airflow from the bottom of the racks did not reach the power supply unit so additional air flow was provided from the back of the cabinet. Installing just the perforated tiles would have caused mixing of cold and hot air without sufficient cooling. With the installation of the hoods on top of the perforated tiles the cold air was supplied from under the raised floor directly to the cabinet. To avoid cold air supply release into the aisle the remainder of the perforated tiles that were exposed to the room were blanked off.

- Openings within the HPC racks created problems with overheating of the cores when the data center temperature was increased to save energy. Installation of internal air dams to prevent recirculation of air within the racks helped to address this problem.

Site	Current IT Load W/sqft	Current IT Load kW	Elec Dist. Loss kW	Cooling Load kW	Fan Load kW	Other users kW	Current PUE	Potential PUE
Site 1	120	2,000	150	750	200	260	1.68	1.64
Site 2A	180	1,050	170	450	195	150	1.92	1.57
Site 2B	240	810	170	370	160	95	1.98	1.63
Site 3	260	1,670	100	700	125	120	1.63	1.56
Site 4	130	550	158	180	47	65	1.82	1.71
Site 5	130	510	73	265	80	33	1.88	1.65

Table 4- Summary of Power Losses and PUE in DOD HPC Data Centers

Next Steps

In order to implement the remaining EEMS, FEMP recommends an investment grade assessment by a firm experienced in data center efficiency improvements. If agency funds are not available for implementing the EEMs, then private sector financing mechanisms such as energy savings performance contracts (ESPC) or utilities energy savings contracts (UESC) may be appropriate, considering the attractive payback periods and the magnitude of the savings. FEMP can assist in exploring such opportunities.

In addition, the LBNL experts recommend installing metering and monitoring systems, especially with a comprehensive dashboard to present the environmental data and the energy efficiency related data including the power use by different components and systems. The dashboard will allow the operators to make real time changes to optimize the energy efficiency.

The LBNL experts also recommend that future purchases of IT equipment include a preference for water cooled systems. At a minimum, the IT equipment should be capable of operating at more than a 90^0F air intake temperature, operate at a high voltage (480V is preferred), and contain a variable speed server fan controlled by the server core temperature.

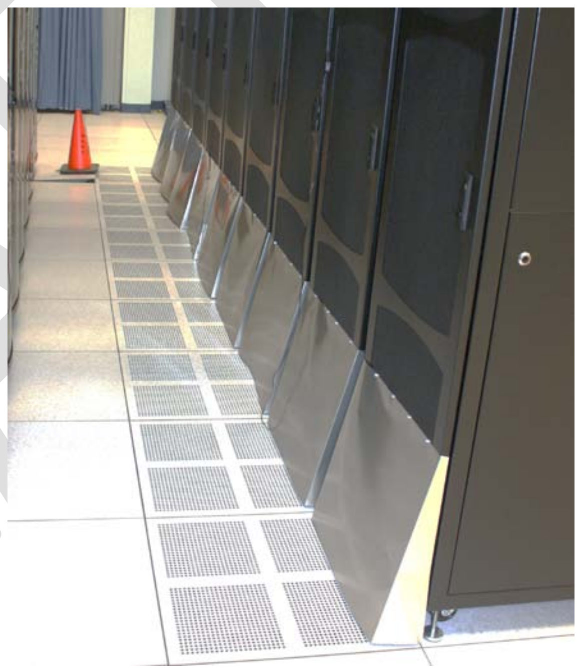

Figure 3- Hoods for Power Supply Unit air intake

For more information on FEMP:
Will Lintner, P.E.
Federal Energy Management Program
U.S. Department of Energy
202.586.3120
william.lintner@ee.doe.gov

Case Study Primary Author:
Rod Mahdavi, P.E. LEED AP
Lawrence Berkeley National Laboratory
510.495.2259
rmahdavi@lbl.gov