High Performance Computing Data Center Metering Protocol

Prepared for:
U.S. Department of Energy
Office of Energy Efficiency and Renewable Energy
Federal Energy Management Program

Prepared by:
Thomas Wenning
Michael MacDonald
Oak Ridge National Laboratory

September 2010

FEMP
Federal Energy Management Program

Introduction

Data centers in general are continually using more compact and energy intensive central processing units, but the total number and size of data centers continues to increase to meet progressive computing requirements. In addition, efforts are underway to consolidate smaller data centers across the country. This consolidation is resulting in a growth of high-performance computing facilities (i.e. - supercomputers) which consume large amounts of energy to support the numerically intensive calculations they perform. The growth in electricity demand at individual data centers, coupled with the increasing number of data centers nationwide, are causing a large increase in electricity demand nationwide.

In the EPA's *Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431,* 2007, the report indicated that US data centers consumed about 61 billion kilowatt-hours (kWh) in 2006, which equates to about 1.5% of all electricity used in the US at that time. The report then suggested that the overall consumption would rise to about 100 billion kWh by 2011 or about 2.9% of total US consumption. With this anticipated rapid increase in energy consumption, the U.S. Department of Energy (DOE) is pursuing means of increasing energy efficiency in this rapidly transforming information technology sector.

This report is part of the DOE effort to develop methods for measurement in High Performance Computing (HPC) data center facilities and document system strategies that have been used in DOE data centers to increase data center energy efficiency.

Contacts

Oak Ridge National Laboratory
Thomas Wenning
R&D Staff
One Bethel Valley Road
Oak Ridge, TN 37831
865-241-8676
wenningtj@ornl.gov

U.S. Department of Energy Federal Energy Management Program
Will Lintner
Federal Energy Management Program
1000 Independence Ave., S.W.
Washington, D.C. 20585
202-586-3120
william.lintner@ee.doe.gov

Abbreviations and Acronyms

ASHRAE	American Society of Heating, Refrigerating, and Air Conditioning Engineers
BAS	Building automation system
CPU	Central processing unit (computer)
CRAC	Computer room air conditioner
CSB	Computer Science Building — Building 5600
DCeP	Data center energy productivity
DCiE	Data center infrastructure efficiency
DOE	US Department of Energy
EERE	DOE Office of Energy Efficiency and Renewable Energy
EPA	US Environmental Protection Agency
FEMP	Federal Energy Management Program
FLOP	Floating-point operation (computer calculation)
HVAC	Heating, ventilating, and air conditioning
IT	Information technology
kW	kilo-Watt
kWh	kilo-Watt-hour
LEED	Leadership in Energy and Environmental Design
MFLOP	Mega-FLOP
MRF	Multiprogram Research Facility
NSF	National Science Foundation
ORNL	Oak Ridge National Laboratory
PDU or RDU	Power distribution unit
PUE	Power usage effectiveness
TGG	The Green Grid
TVA	Tennessee Valley Authority
UPS	Uninterruptible power supply
VFD	Variable frequency drive (or variable flow drive)
W	Watt

Contents

1 Metering Background

Significant efforts to further understanding and benchmarking of data center energy efficiency have occurred over the past few years. Numerous potential regulatory and institutional initiatives have driven these efforts. The U.S. Environmental Protection Agency (EPA) "Report to Congress on Server and Data Center Energy Efficiency" (2007) and the European Commission's "Code of Conduct on Data Centres Energy Efficiency, Version 1" (2008) are only two examples of regulatory interest in data center efficiency.

Future rules and regulations regarding data center power consumption will be developed primarily by the federal government. DOE is currently pursuing multiple requirements and guidelines for DOE facilities also. DOE and its facility contractors invest more than $2 billion/yr in information technology resources, a large component of which includes desktop and laptop computers utilized by end users, and servers and storage media maintained in data centers. The Energy Independence and Security Act of 2007 directed agencies to improve energy efficiency, reduce energy costs, and reduce greenhouse gas emissions. In addition, the Department issued Order 450.1A in 2008 that required programs and sites to implement a number of environmental stewardship practices, including enabling power management features on computers and other electronic equipment. Further goals have been spelled out in the Department of Energy's *Strategic Sustainability Performance Plan* recently released in 2010 to address the requirements of Executive Order 13514, *Federal Leadership in Environmental, Energy and Economic Performance,* signed by the President on 10/5/2009.

DOE's Office of Energy Efficiency and Renewable Energy (EERE) is working to evaluate energy efficiency opportunities in data centers. As part of this process, energy assessments of data center facilities will be conducted under the *Save Energy Now* program. These assessments are designed to help data center professionals identify energy saving measures that are most likely to yield the greatest energy savings. The assessments are not intended to be a complete energy audit, but rather, the process is meant to educate data center staff and managers on an approach that can be used to identify potential energy saving opportunities that can be further investigated. The intent is that sites will continue to track improvement in their energy performance to document their energy performance metrics and actions implemented over time. This performance tracking will be done by the sites.

Practical performance metrics have been developed by The Green Grid (TGG), an industry consortium active in developing metrics and standards for the IT industry. DOE and TGG have a Memorandum of Understanding signed in 2007 to cooperate on multiple fronts to address improving data center energy and water efficiency. In 2009, TGG and ASHRAE published a book entitle, "Real-Time Energy Consumption Measurements in Data Centers." The book is a comprehensive resource discussing various data center measurements.

2 Metering Purpose

Metering projects are undertaken for a variety of reasons. Meters allow for a better understanding of a facilities power and cooling systems. Metered data can be used to determine benchmarks for the current state of operations. Benchmarking and baselining enables a facility to track performance and efficiency improvements over time. In addition, benchmarking allows for management and operations to take a proactive approach to identifying performance improvement projects and technologies. With the use of meters, new equipment, retrofits, and system development (HVAC, lighting, controls, etc) performance can easily be tracked and documented over time. Meters provide reliable measurement and verification for the expected energy benefits of system improvement projects.

Metering can be utilized as a diagnostic tool to continuously monitor, track, and improve facility performance to ensure long-term efficient operation. If properly set up, meters can also enable a facility to consistently track, report, and communicate various metrics.

From a facility management perspective, installing meters can help determine system benchmarks and use them to compare against other facilities. This information can also be used to set future performance goals, determine specific improvement targets, and verify that targets have been met. This could enable management to incentivize goals for meeting and exceeding established targets. Analysis of metered data can provide improved planning for future utilities.

Metering projects typically carry a burdensome cost, thus site managers typically struggle to justify projects from solely a capital cost-savings perspective. Installing meters does not directly save energy; however, intelligent integration of meters into a system can allow for improved system performance, which in turn can save energy. Building management systems often use the information from various meters to directly influence and dictate operations and control of key systems. In high performance computing facilities, intelligent integration of meters can play a vital role in efficient system operation.

In high performance computing facilities, there are a number of potential measurements to be taken. These measurements may include any of the following:

- Power demand (kW)
- Power consumption (kWh)
- Amperage per phase
- Percent amperage
- Voltage per phase
- Power factor
- Line harmonics
- Head / suction pressure
- Flow rates
- Flow velocity
- Humidity
- Entering / Leaving Chilled Water Temperature
- Entering / leaving condenser water temperature
- Condenser / Evaporator refrigerant temperature
- Instantaneous equipment capacity
- Air quality

With all these possible measurements, it is easy to imagine how a metering project can quickly become a costly and complex undertaking. Thus, it is important to establish the purpose at the onset of a metering project and to understand what data needs to be acquired. Site managers must be able to set clear goals for their metering project. This often requires resolving differences between desired and realistic metering expectations. Several main considerations to account for during the planning stages of a metering project include (ASHRAE):

- Project Goals – Establish the project goals and data requirements before selecting hardware.
- Project Cost and Resources – Determine the feasibility of the project given the available resources.
- Data Products – Establish the desired final output data type and format before selecting data measurement points.
- Data Management – Identify proper computer and personnel resources to handle data collection needs.
- Data Quality Control – Identify the system to be implemented to check and validate data quality.
- Commitment – Projects often require long-term commitments of personnel and resources.
- Accuracy Requirements – Determine the required accuracy of the final data early in the project.

Developing a metering project often becomes an iterative process, revolving between budget constraints, equipment costs, and metering goals. The following figure from *ASHRAE Handbook – HVAC Applications, 2007* shows a nine-step flowchart for the iterative planning process.

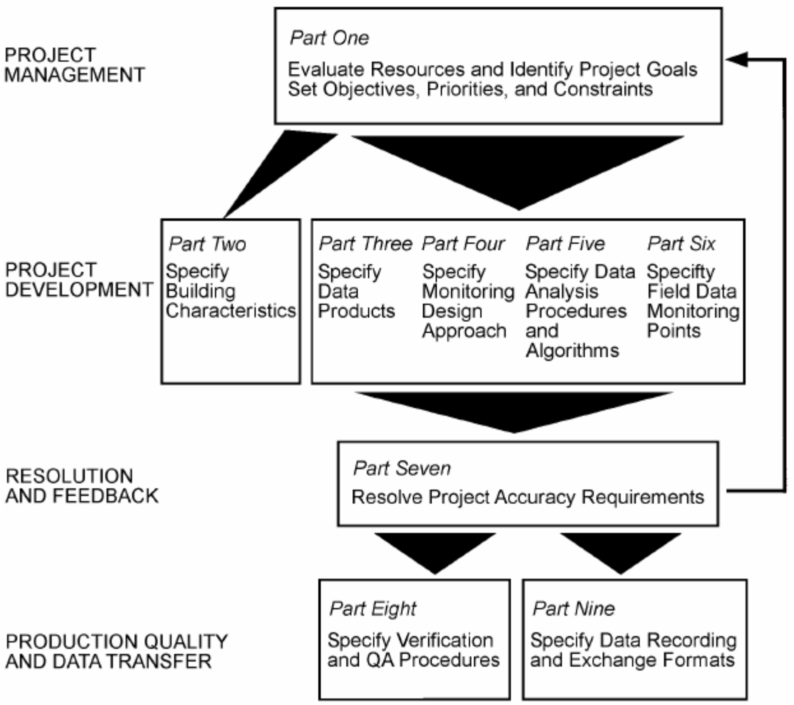

Methodology for Designing Field Monitoring Projects

3 Levels of Metering

There are four general levels of utility data that can be captured with varying degrees of metering. The levels start off with broad site consumption metering and it narrows down to specific end use measurements (ASTM E 1465). The four levels consist of:

1. Site Level
    - General utility coming into a site. The site may have several buildings and other end use equipment.
2. Building Level
    - Metering at the utility feed to an individual building. This encompasses all energy used within a given building.
3. System Level
    - Sub-metering at the system level. This may include whole systems such as, chiller plants, lighting, computer room air conditioners (CRACs), etc.
4. Component Level
    - Sub-metering within systems. This may include flow and temperatures within a chiller system.

Detailed sub-metering at the component level will provide the greatest resolution of energy consumption within a facility and will provide for more control option. It will also provide the most feedback for the user to analyze; however, this is often the most expensive option when considering permanent metering installations. Metering only at the site and building level is often the cheapest option, however, it is generally insufficient when trying to determine system and facility performance.

One method for minimizing metering costs is to monitor energy as high in the distribution system as possible. Doing so minimizes the number of monitoring nodes and therefore reduces equipment needs. However, measuring too high in a system will lead to a poor understanding of end-use consumption and results in difficulty trying to assess system performance. A rule-of-thumb is to not separately meter an end-use if its expected consumption is less than 10% of the higher nodes consumption (ASTM E 1465). Multiple levels of nodes allow for some redundant metering which can be used to help identify installation problems and can be used to facilitate a comparison between end-use and utility meter data.

4 Performance Metrics

The primary metrics reported are related to energy and are the Power Usage Effectiveness (PUE) and Datacenter Infrastructure Efficiency (DCiE), metrics from The Green Grid (TGG, www.thegreeengrid.org). These metrics are defined in several white papers by TGG (see White Paper #6 from the website).

The PUE is defined as follows:
PUE = Total Facility Power/IT Equipment Power

PUE Benchmarking Values -

Standard	Good	Better
2.0	1.4	1.1

and its reciprocal, the DCiE is defined as:
DCiE = 1/PUE = IT Equipment Power/Total Facility Power x 100%

DCiE Benchmarking Values -

Standard	Good	Better
0.5	0.7	0.9

For the PUE and DCiE equations, the Total Facility Power is defined as the power measured at the utility meter — the power dedicated solely to the datacenter (this is important in mixed-use buildings that house datacenters as one of a number of consumers of power). The IT Equipment Power is defined as the equipment that is used to manage, process, store, or route data within the data center. It is important to understand the components for the loads in the metrics, which can be described as follows:

IT EQUIPMENT POWER. This includes the load associated with all of the IT equipment, such as compute, storage, and network equipment, along with supplemental equipment such as KVM switches, monitors, and workstations/laptops used to monitor or otherwise control the datacenter.

TOTAL FACILITY POWER. This includes everything that supports the IT equipment load such as:

- Power delivery components such as UPS, switch gear, generators, PDUs, batteries, and distribution losses external to the IT equipment.
- Cooling system components such as chillers, computer room air conditioning units (CRACs), direct expansion air handler units, pumps, and cooling towers.
- Compute, network, and storage nodes.
- Other miscellaneous component loads such as datacenter lighting.

One so-called “green” metric used for the “Green500” list of computers is MFLOP/W. This “green” metric has some limited value for understanding how “green” a computer is, but the metric is unduly influenced by running with limited memory, and thus reduced ability to handle certain types of important tasks.

A better data center productivity (DCP) metric has been identified by TGG as DCeP. DCeP is envisioned as one of a family of DCP metrics, designated generically as DCxP. The energy productivity metric, DCeP, is defined by TGG (White Paper #18) as:

$$\text{DCeP} = \frac{\text{useful work produced in a data center}}{\text{total energy consumed in the data center to produce that work}}$$

The major issue facing this productivity metric at this time is developing a meaningful measure of the numerator. Several proposals have been made, but no good solution has yet emerged (see TGG White Paper #24).

5 Metering Equipment

In the broadest sense, there are two types of data, time dependant and time independent data. Time dependant data includes weather and energy consumption data. Time independent data includes facility descriptive data and project cost data. Various time dependant data can be measured using an array of different measurement devices. Most devices are capable of being used to capture data at various time intervals. Capturing data on the smallest possible time intervals provides the most resolution; however, it also provides exceedingly large quantities of data. This can lead to issues of having too much information to sort through. If data is to be captured on increments of less than an hour, it is very beneficial to have an automated collection and processing system. Using outdoor air temperature for example, the minimum time interval of data collection should be once daily. Shorter time increments, such as hourly, can provide more clarity when comparing it to equipment’s electricity consumption. Smaller increments, less than hourly, can be used to continuously calculate near-instantaneous system performance for systems such as a chiller plant; however, this can only be achieved by using an automatic collection system.

The following tables describe a cursory breakdown of various metering technologies in the marketplace. Included in the table are the measurement type, sensor, application, and relative accuracy of each technology. For expanded information on the applications and limitations of the sensors listed in the table below, please refer to *ASHRAE Handbook – Fundamentals, 2009* and *Real-Time Energy Consumption Measurements in Data Centers, 2009.*

Thermodynamic Measurements

Measurement	Sensor	Application	Accuracy
Temperature	Thermocouples	Any	1.0 - 5.0%
	Thermistors	Any	0.1 - 2.0%
	Resistance Temperature Detectors	Any	0.01 - 1%
Pressure	Bourdon Tube	Pressure in pipe	0.25 - 1.5%
	Strain Gage	Pressure in pipe	0.1 - 1%
Humidity	Psychrometer	Above freezing temperatures	3 - 7%
	Hygrometer	Any	Varies

Flow Rate Measurements

Measurement	Sensor	Application	Accuracy
Liquid	Ultrasonic Flow Meter	Flow in Pipes	1.0 - 5.0%
	Variable Area / Orifice Plate	Flow in Pipes/Ducts	0.5 - 5.0%
	Turbine Wheel	Flow in Pipes/Ducts	0.3 - 2.0%
	Paddle Wheel	Flow in Pipes/Ducts	0.5 - 5.0%
	Shedding Vortices	Flow in Pipes	1%
Gas	Pitot Tube and Manometer	Any	1.0 - 4.0%
	Anemometer	Any	1.0 - 5.0%

Electrical Measurements

Measurement	Sensor	Application	Accuracy
Current	Solid Core	Permanent Installations	Varies
	Split Core	Permanent Installations	Varies
	Clamp-on / Flex	Temporary	Varies
Voltage	Pressure Transducer	Any	Varies
	Voltage Divider	Low Voltage AC or DC	Varies
Power	Portable Meter	Temporary	Varies
	Panel Meter	Permanent Installations	Varies
	Revenue Meter	Permanent Installations	Varies
	Power Transducer	Monitoring	Varies

The following table from *ASHRAE Handbook – HVAC Applications, 2007* explicitly calls out the accuracy and reliability issues of some of the most used metering instrumentation. One important issue that is often overlooked in metering installations is the need to periodically re-calibrate sensors. If sensors are not in calibration and are being used in system controls, they may be causing substantial inefficiencies in the system. Sensors that are only used for data collection and not control still need to be recalibrate to ensure accuracy of the calculated metrics and benchmarks.

Instrumentation Accuracy and Reliability

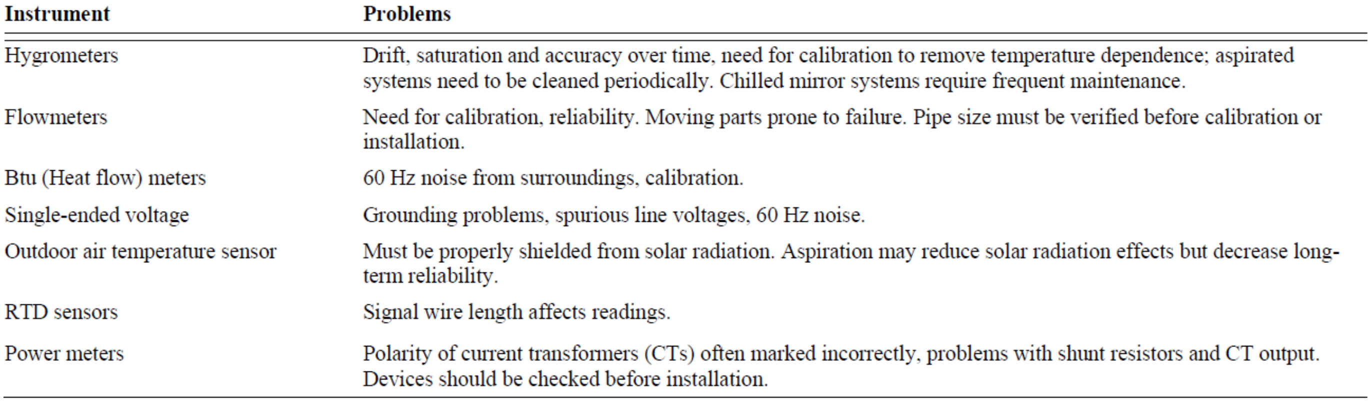

Instrument	Problems
Hygrometers	Drift, saturation and accuracy over time, need for calibration to remove temperature dependence; aspirated systems need to be cleaned periodically. Chilled mirror systems require frequent maintenance.
Flowmeters	Need for calibration, reliability. Moving parts prone to failure. Pipe size must be verified before calibration or installation.
Btu (Heat flow) meters	60 Hz noise from surroundings, calibration.
Single-ended voltage	Grounding problems, spurious line voltages, 60 Hz noise.
Outdoor air temperature sensor	Must be properly shielded from solar radiation. Aspiration may reduce solar radiation effects but decrease long-term reliability.
RTD sensors	Signal wire length affects readings.
Power meters	Polarity of current transformers (CTs) often marked incorrectly, problems with shunt resistors and CT output. Devices should be checked before installation.

Source: 2007 ASHRAE Handbook – HVAC Applications

When it comes to capturing data from all the various sources in a HPC facility, it is best to do so using an automated system. Though some measurements can be captured manually, this method often proves to be cumbersome over time and ultimately becomes an unsustainable practice. The best method to capture, store, and analyze the measurement information is through the means of a data acquisition system, also known as building automation systems, building management systems, energy monitoring and control systems, and supervisory control and data acquisition systems. These systems often serve multiple purposes. They can monitor, trend, record system status, record energy consumption and demand, record hours of operation, control subsystem functions, produce summary reports, and print alarms when systems do not operate within specified limits. One example of the usefulness of controlling subsystems is in room humidity control. DAS allows for the central processing and control of maintaining a computer rooms relative humidity instead of having numerous humidification and dehumidification units fighting one another in an effort to meet their localized sensor requests. Numerous platforms are already used in the marketplace; examples include PowerNet and Metasys. A breakdown of various data acquisition systems and there general purposes are described in the table below.

General Characteristics of Data Acquisition Systems (DAS)

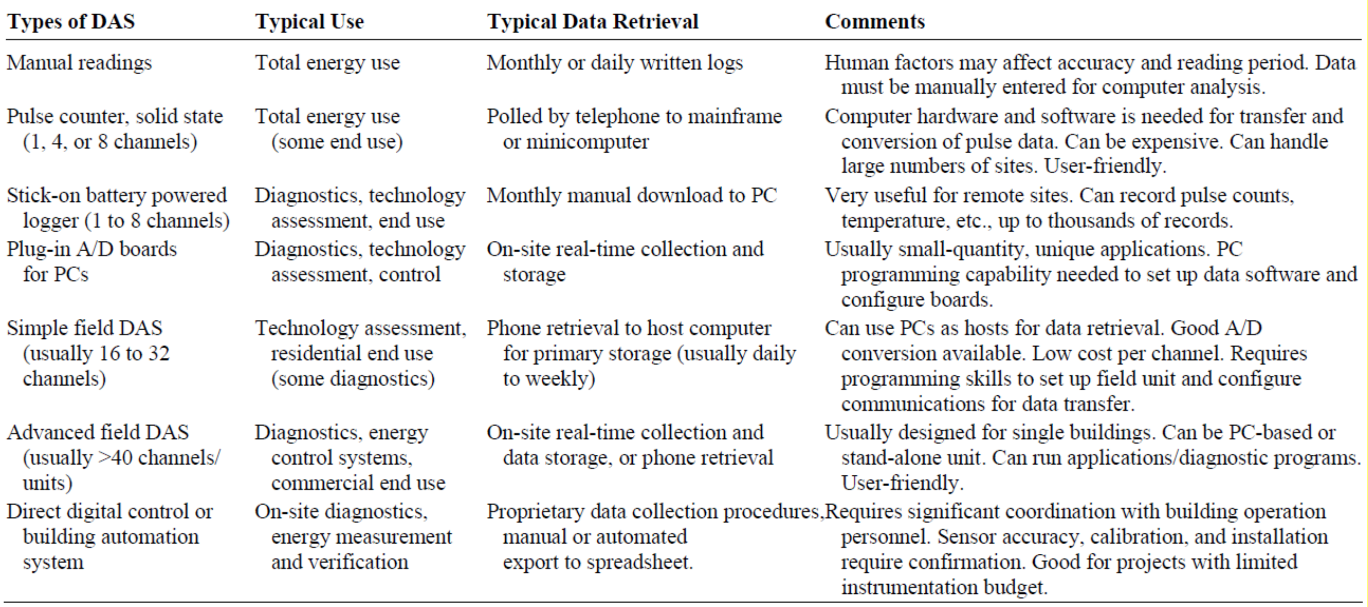

Types of DAS	Typical Use	Typical Data Retrieval	Comments
Manual readings	Total energy use	Monthly or daily written logs	Human factors may affect accuracy and reading period. Data must be manually entered for computer analysis.
Pulse counter, solid state (1, 4, or 8 channels)	Total energy use (some end use)	Polled by telephone to mainframe or minicomputer	Computer hardware and software is needed for transfer and conversion of pulse data. Can be expensive. Can handle large numbers of sites. User-friendly.
Stick-on battery powered logger (1 to 8 channels)	Diagnostics, technology assessment, end use	Monthly manual download to PC	Very useful for remote sites. Can record pulse counts, temperature, etc., up to thousands of records.
Plug-in A/D boards for PCs	Diagnostics, technology assessment, control	On-site real-time collection and storage	Usually small-quantity, unique applications. PC programming capability needed to set up data software and configure boards.
Simple field DAS (usually 16 to 32 channels)	Technology assessment, residential end use (some diagnostics)	Phone retrieval to host computer for primary storage (usually daily to weekly)	Can use PCs as hosts for data retrieval. Good A/D conversion available. Low cost per channel. Requires programming skills to set up field unit and configure communications for data transfer.
Advanced field DAS (usually >40 channels/units)	Diagnostics, energy control systems, commercial end use	On-site real-time collection and data storage, or phone retrieval	Usually designed for single buildings. Can be PC-based or stand-alone unit. Can run applications/diagnostic programs. User-friendly.
Direct digital control or building automation system	On-site diagnostics, energy measurement and verification	Proprietary data collection procedures, manual or automated export to spreadsheet.	Requires significant coordination with building operation personnel. Sensor accuracy, calibration, and installation require confirmation. Good for projects with limited instrumentation budget.

Source: 2007 ASHRAE Handbook – HVAC Applications

The following table from *ASHRAE Handbook – HVAC Applications, 2007* lists details about some of the concerns associated with using and installing data acquisition hardware in systems.

Practical Concerns for Selecting and Using Data Acquisition Hardware

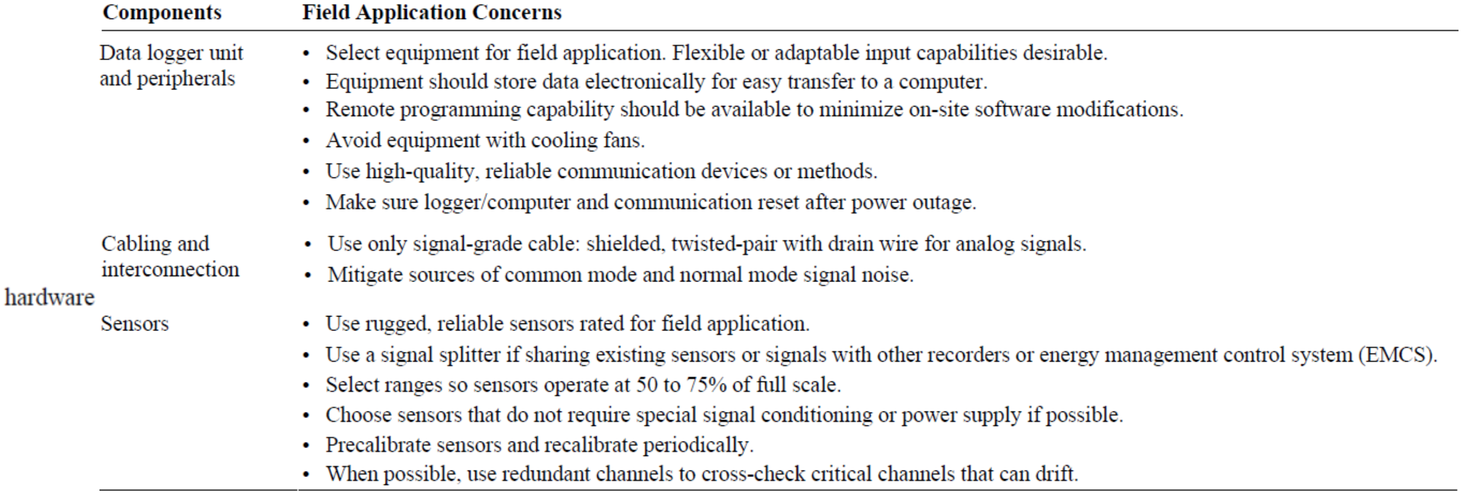

Components	Field Application Concerns
Data logger unit and peripherals	• Select equipment for field application. Flexible or adaptable input capabilities desirable. • Equipment should store data electronically for easy transfer to a computer. • Remote programming capability should be available to minimize on-site software modifications. • Avoid equipment with cooling fans. • Use high-quality, reliable communication devices or methods. • Make sure logger/computer and communication reset after power outage.
Cabling and interconnection	• Use only signal-grade cable: shielded, twisted-pair with drain wire for analog signals. • Mitigate sources of common mode and normal mode signal noise.
hardware	
Sensors	• Use rugged, reliable sensors rated for field application. • Use a signal splitter if sharing existing sensors or signals with other recorders or energy management control system (EMCS). • Select ranges so sensors operate at 50 to 75% of full scale. • Choose sensors that do not require special signal conditioning or power supply if possible. • Precalibrate sensors and recalibrate periodically. • When possible, use redundant channels to cross-check critical channels that can drift.

Source: 2007 ASHRAE Handbook – HVAC Application

6 Key Equipment in HPC Facilities

High performance computing (HPC) facilities contain a large number and various types of equipment. With each type of equipment, there is a different need when it comes to metering. The table below shows a generic breakdown of the various systems found within an HPC facility. The table also lists the potential components within each system and the key measurements to be obtained. The measurements for each system can be used to calculate various performance metrics for the system and the facility. More detailed information on the subsystems and on the minimum practical, best practical, and state-of-the-art measurements can be found in *Real-Time Energy Consumption Measurements in Data Centers, 2009.*

Measurable Equipment and Key Measurements in HPC Facilities

System	Components	Key Measurements
General Measurements		Outdoor Temperature, Outdoor Relative Humidity, Indoor Temperatures, Indoor Relative Humidites
Servers / Storage / Networking	Internal Fans	Current, Voltage, Power, Air Intake Temperature
Uninterruptible Power Supplies / Power Distribution Units		Power, Current, Voltage
Transformers		Current, Voltage
Automatic Transfer Switches		Power, Current, Voltage
Computer Room Air Conditioner / Computer Room Air Handling Units	Compressors Blowers/Fans Pumps Humidifiers Reheaters	Temperature, Flow Rate, Power, Voltage, Current, Power
Chillers	Compressor Heat Exchangers	Temperature, Flow Rate, Power, Voltage, Current
Cooling Towers	Blowers/Fans Pumps	Current, Voltage, Power, Flow Rate, Pressure
Pumps / Fans / Blowers		Current, Voltage, Power, Flow Rate, Pressure
Heat Exchangers		Temperature, Flow Rate
Lighting		Current, Voltage
Distributed Energy Systems / Combined Heat & Power Systems	Varies	Varies upon Equipment: Temperature, Flow Rate, Power, Voltage, Current

7 Oak Ridge National Laboratory Metering Case Study

Oak Ridge National Laboratory (ORNL) has been concerned with data center efficiency since the design and construction of the new East Campus facilities began. One major interest is improving data center efficiency as new centers are built, taking what is learned from previous centers or center upgrades to make each new one more efficient. After the Computer Sciences Building - Building 5600 (CSB) was completed, DOE approved construction of the Multiprogram Research Facility (MRF), work on which began in February 2005. The MRF has several areas that are dedicated to computer applications, and lessons learned from Building 5600 on energy efficiency and LEED certification were applied to these areas as well as other

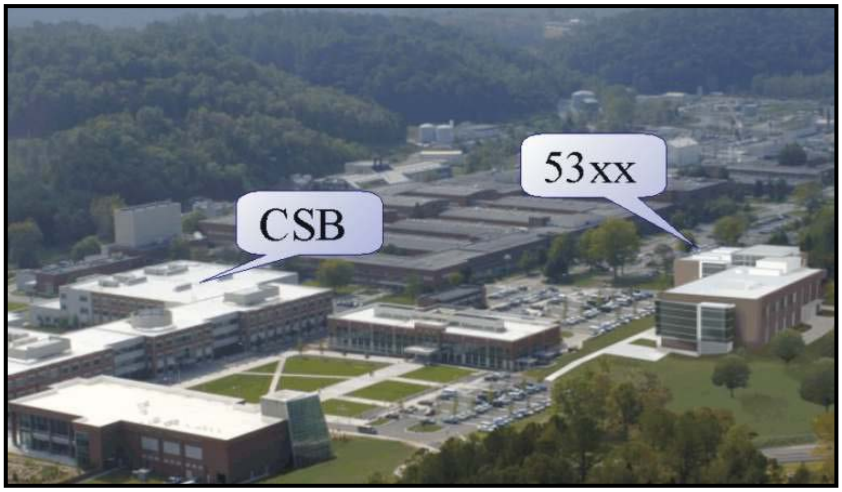

areas in the facility. Subsequent upgrades to data centers in either of these buildings have incorporated the latest energy efficiency ideas that were implemented in the other building.

7.1 Site Background

Originally known as Clinton Laboratories, Oak Ridge National Laboratory (ORNL) was established in 1943 to carry out a single, well-defined mission: the pilot-scale production and separation of plutonium for the World War II Manhattan Project. From this foundation, the Laboratory has evolved into a unique resource for addressing important national and global energy and environmental issues.

With the creation of DOE in the 1970s, ORNL's mission broadened to include a variety of energy technologies and strategies. Today the laboratory supports the nation with a peacetime science and technology mission that is just as important as, but very different from, its role during the Manhattan Project. ORNL is DOE's largest science and energy laboratory.

Case Study Data Center Identification
Location: Oak Ridge, TN
5600 Computational Sciences Building (CSB) - 137,000 sq ft (3-story) & Central Plant - 7,650 sq ft
Completed August 2003
LEED Rating: NC, v2.0--Level: Certified

8 Metering Protocol

Determination of the energy parameters for the data centers in Building 5600 requires an extensive array of meters. The original data center's electrical and cooling infrastructure received major upgrades in the last few years to allow installation of major new supercomputers. As of June 30, 2010, Building 5600 houses the #1, #4, #20, and #36 most powerful supercomputers in the world. In addition, ORNL is currently experiencing the installation and start-up of a new supercomputer to support the National Oceanic and Atmospheric Administration (NOAA).

The three-story CSB contains offices and a 2nd floor computer center that houses a typical data center in one half of the area and a supercomputer in the other half. A large raised-floor computer center currently houses two supercomputers on the first floor. The NOAA supercomputer will be the third supercomputer on the first floor once completed. The CSB data centers are part of the larger facility that includes other functions.

In the continuous upgrading process, extensive electric metering has been installed in the building (about 100 in all). Of these meters, 56 were required to determine the energy benchmarking metrics. Chilled water meters were also installed to measure chilled water flows to the first floor and second floor centers. The electrical data metering network is Eaton (Cutler-Hammer) PowerNet, and the chilled water meters are handled by the Johnson Controls Metasys building automation system (BAS). The two metering and control systems are currently not integrated; however, infrastructure is expected to evolve to a type of higher level integrative enterprise solution, such as one based on OLE Process Control.

8.1 Electric System and Metering

Power is supplied by the Tennessee Valley Authority (TVA) from a substation that is supplied with three independent 161 kV transmission lines. TVA power provides four to five 9's of uptime, so expected interruptions are minor. Electrical power is delivered to the CSB at plant distribution voltage of 13.8 kV. Multiple transformers at Building 5600 step power down to 480V or 208V. Currently, 13.8kV-

480Y/277V transformers are located inside the building close to the computer room to reduce distribution losses. ORNL power is measured every half-hour by TVA and is aggregated to the hourly level in this figure. Building 5600 electric meters collect data at different intervals but is aggregated to hourly data.

The uninterruptible power supply (UPS) requirements were minimized to only required support for dual corded systems such as disk drives, networking and communications equipment, and business applications. UPS ride-through power is supplied until the backup generator can start and come on line and provide power in the event of a power outage. One UPS unit has battery energy storage and one has flywheel energy storage, with both of these units being double conversion types.

The PowerNet power distribution metering and control system can be used to manage energy cost, analyze harmonics, view waveforms for transient events power quality, trending meter equipment usage, maximize use of available capacity, etc.. The system has a total of 567 monitored points and provides integrated metering from the 161 kV distribution level down to the 208V end user level. ORNL has one of the largest PowerNet installations in the Southeast.

Overall, PowerNet is used as an engineering design / operation tool, for real time data monitoring of the electrical infrastructure, for internal power billing, for power quality analysis and system monitoring, and during medium voltage switching to verify system operation. Networking is via an ethernet interface to several dedicated data servers in Building 5600.

8.2 Electric Data Measurement

ORNL power is fed into the main substation at 161 kV, where it is stepped down to 13.8 kV for distribution. Building 5600 power is on the 480V side of the 13.8kV/480V transformers at the building. All line losses and all transformer losses are included in the ORNL value, but no losses down to the 480V level are included in the 5600 power. There are 13 electric meters measuring electric use from the main 480V feeds. These meters do not isolate the supercomputer data center electric use. To perform efficiency calculations on the data centers, a total of 56 meters are required for the computer center. The following summary information is an example breakout for the meters on the main 480V power supply.

Building 5600 Electric Use Simple Breakout Example				
Winter Month Total kWh				
Meter ID	**kWh**	**480V Panel ID**	**Breakout Item**	**Breakout kWh**
132	345,194	1A	mixed	
138	669,813	1B	mixed	
570	354,349	2A	1st floor center other	354,349
582	285,924	2B	2nd floor center other	285,924
594	2,496	3A	Mixed, mainly (2) 1200-ton chillers when running	
606	1,503	3B	Mixed, mainly (1) 1200-ton chiller when running	
183	1,367,119	4	2nd floor Cray	1,367,119
501	1,281,000	7	1500-ton chillers and plant	1,281,000
512	1,150,963	8	Kraken computer	1,150,963
524	1,133,532	9	Jaguar computer	1,133,532
536	1,097,951	10	Jaguar computer	1,097,951
548	1,156,162	11	Jaguar computer	1,156,162
554	478,791	12	Kraken computer	478,791
Total	9,324,798			

Given the high level of metering, estimates are made to determine equipment consumption and losses. Power losses from transformers are measured with on a spot-check basis and then estimates are made for continual operation. Uninterruptable power supply (UPS) losses are estimated using load information and manufacturers specification sheets. Likewise, lighting energy consumption throughout the computer rooms is estimated based on fixture counts, power per fixture, and operating control. Light operation includes dimming them to half power at night. The lighting power estimate is expected to be very close to actual.

The 56 meters throughout the system along with system estimates are used to calculate the PUE for the datacenters. The metering plan point list and the calculation methodology for the first and second floor supercomputer centers in Building 5600 at ORNL is shown below. Dxxx stands for PowerNet device number ‘xxx,’ where all the devices in this list are electric meters.

PUE (or DCiE) Calculation for CSB Electric Data Meters	
Chillers 1, 2, 3	
	D427-5600 MSB-3A MTR
+	D430-5600 MSB-3 MTR
+	D406-5600 MCC14 MTR
+	D403-5600 CT2
+	D404-5600 CT3
+	D395-5600 CHWP2
+	D396-5600 CHWP3
-	D401-5600 ATS-7
-	D402-5600 ATS-8
+	1% xfmr loss
Chillers 4, 5	
	D253-5600 MSB-7 MTR
+	D405-5600 MCC13 MTR
+	1% xfmr loss
Lighting, both floors	
	Estimated value. Load is small and constant
CRU's first floor (AC units)	
	D398-5600 ATS-4
+	D399-5600 ATS-5
+	D400-5600 ATS-6
+	1% xfmr loss
CRU's second floor (AC units)	
	D397-5600 ATS-3
+	D401-5600 ATS-7
+	D402-5600 ATS-8
+	1% xfmr loss
Supercomputer Center, first floor	
LCF main computer (Jaguar XT5)	
	D299-5600 MSB-9 MTR
+	D322-5600 MSB-10 MTR
+	D345-5600 MSB-11 MTR

+	1% xfmr loss
NSF main computer (Kraken)	
	D276-5600 MSB-8 MTR
+	D368-5600 MSB-12 MTR
+	1% xfmr loss
ERP	**Local power units on the UPS**
	D68-5600 ERP1
+	D69-5600 ERP2
+	D70-5600 ERP3
+	D71-5600 ERP4
+	**D79-5600 ERP14 -- not used any longer**
+	7% xfmr / UPS loss
RDUs	**Local power units / panels**
	D81-5600 RPA2
+	D82-5600 RPA3
+	D86-5600 RPA7
+	D89-5600 RPB3
+	2% xfmr loss
RDF	**Data communication facility for all network traffic**
	D72-5600 ERP5
+	7% xfmr / UPS loss
	subtotal
	D84-5600 RPA5
+	2% xfmr loss
	subtotal
	add subtotals
Disk drives	
	D111-5600 PDU-A3A MTR
+	D110-5600 PDU-A2A MTR
+	2% xfmr loss
	subtotal
	D137-5600 PDU-UPS2-3A MTR
+	D138-5600 PDU-UPS2-4A MTR
+	D136-5600 PDU-UPS2-2A MTR
+	7% xfmr / UPS loss
	Subtotal
	add subtotals
Computer Center, second floor	
Cray XT4 Supercomputer	
	D1-5600 MSB-4 MTR
+	1% xfmr loss
	Subtotal
Local power units / panels	
+	D109-5600 PDUB4 MTR
+	**D92-5600 RPB6 -- not used any longer**
+	D85-5600 RPA6
+	D91-5600 RPB5

+	D97-5600 RPD1
+	D98-5600 RPD2
+	D99-5600 RPD3
+	D100-5600 RPD4
+	D101-5600 RPD6
+	D105-5600 RPD7
+	D106-5600 RPD8
+	D107-5600 RPD9
+	D108-5600 RPD10
+	2% xfmr loss
	subtotal
	add subtotals
ERP	**Local power units on the UPS**
	D74-5600 ERP7
+	D75-5600 ERP8
+	D76-5600 ERP9
+	D77-5600 ERP10
+	D78-5600 ERP13
+	7% xfmr / UPS loss
Disk drives	
	D135-5600 PDU-UPS2-1A MTR
+	7% xfmr / UPS loss
Cray XT fan power has to be measured separately	

The method of calculating electric power for the chiller plants is described under the Chiller Plant Measurements section of this report. Cooling unit (AC unit) electricity is metered via the local power units.

The electric meters are electronic programmable meters with extensive capabilities, but after initial analysis of data variability and chiller plant performance variations, a decision was made to calculate the performance metrics on a daily basis. Currently, PUE is determined monthly, based on the metering protocol here. All the required meters were programmed to log daily kWh readings, and the daily data are used to calculate the daily values of PUE for each month.

8.3 Data Center Room Cooling and Measurement

The Jaguar and Kraken supercomputers use the Cray ECOphlex (phase change liquid exchange) system that pumps refrigerant to each cabinet, where the phase-change refrigerant heat transfer system provides ebullient cooling. Liquid refrigerant is pumped to heat exchangers located at different levels in each computer cabinet, and an axial fan blows air up through the cabinet to absorb heat from electronics and then transfer it to heat exchangers. The refrigerant vapor from the heat exchangers then flows back to another heat exchanger, where it transfers heat to the chilled water system, is condensed to liquid , and is pumped back to the computer. This means that the first floor data center only requires minimal cooling from underfloor air. This approach reduces system size and enhances stability of temperature control by operating in a constant-temperature phase-change condition that stays above air dewpoint. In addition, use of liquid cooling can lead to major reductions in fluid transport power due to the higher energy density potentials of liquids in comparison to air. Liquid refrigerant pumping power is not directly measured. However, it is included as part of the “technical power” total.

Data storage and networking systems do receive underfloor air cooling, but minimal underfloor air is provided to the Jaguar and Kraken systems with the exception of 16 air cooled cabinets in the center of the Jaguar system rows.

In the second floor center, underfloor air is provided to the previous-generation Cray systems, but the other systems in the room do not have special cooling systems, so they are cooled by air from the computer room AC units. The 2nd floor Cray cooling system utilizes supply air temperature control instead of return air temperature control to ensure a constant temperature is delivered to the Cray systems. This has reduced fan horsepower requirements and nearly eliminated issues with hot spots.

To acquire a better understanding of the affects that the Cray XT5 fan power has on the PUE, manual frequency readings are taken on each of the 200 cabinets. The fan power consumption is characterized as a function of frequency with the fan motor variable speed frequencies being measured at each computer cabinet. The fan power relationship is shown in the table to the right. Linear interpolation is used to calculate between points.

XT5 fan power	
Hz	**kW**
40	1.2
45	1.6
50	2.2
55	2.8
60	3.7
65	4.7

Cray indicates the fan frequencies are controlled by the inlet air temperature. Since the fans are controlled by inlet air temperature, and since most cabinets are cooled by the XDP system, it is assumed that the level of computing work does not impact the frequencies much. Total XT5 fan power for the supercomputers is obtained by ratioing up the representative measured Jaguar fan power for its 200 cabinets, to 288 to include Kraken.

8.4 Chilled Water System and Metering

One chiller system is installed in the complex where Building 5600 is located. This system consists of three 1200 ton chillers and two 1500 ton chillers. This system is cross-connected with the chilled water plants in the MRF complex. This cross-connection will eventually be phased out with the addition of new equipment. A new chiller plant for Building 5800 is in progress to allow complete separation of building comfort cooling and data center cooling. The two 1500 ton chillers provide most of the cooling to Building 5600 data centers, with the three 1200 ton chillers being run as little as possible at this time. Chilled water is piped directly to the data center rooms.

The chilled water system currently has an interconnection with the MRF complex. Only a limited amount of cooling can be supplied from the MRF complex systems, but the overall backup is important and the chiller plants in the MRF complex are newer and more efficient. The chiller plants are controlled by the Building Automation System (BAS), which also controls the data center cooling overall. The next graphic shows the BAS main chiller plant display for Building 5600 as an example.

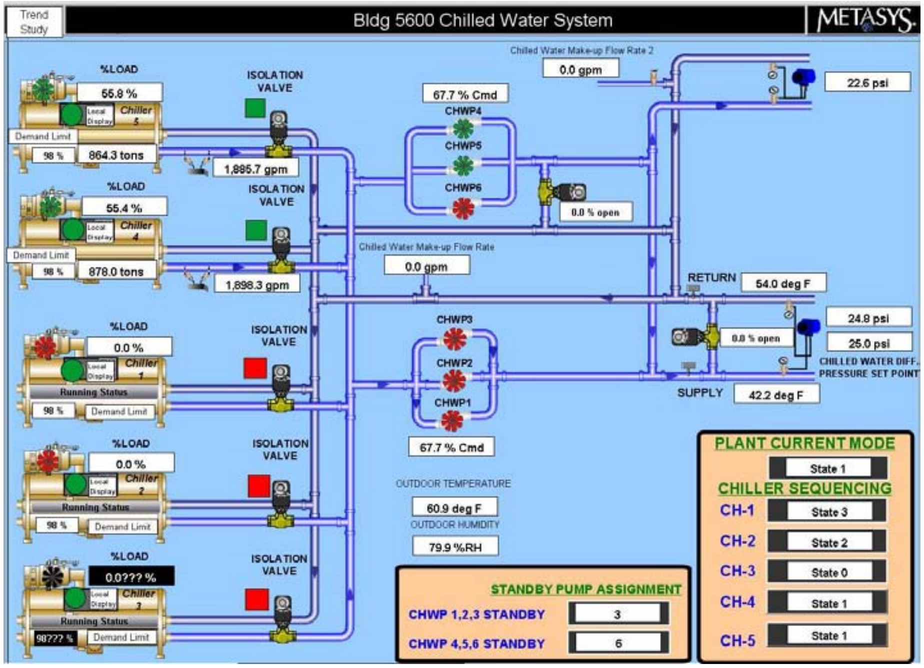

Chillers 1, 2, and 3 are the original 1200-ton chillers for the building (2002 vintage). Chillers 4 and 5 were added as part of the 2003–2008 upgrade (2006 vintage), and since they are more efficient, they are the primary units to run and are run almost full out most of the year. In the winter, the balance of the 5600-5700-5800 complex cooling load can be provided by the chiller input from the MRF complex plants, which are newer and more efficient than Chillers 4 and 5. Thus, during the winter, Chillers 1, 2, and 3 are not run, and total cooling load for the data centers and the balance of the 5600-5700-5800 complex is provided by chillers 4 and 5 and the MRF complex interconnection. A diagram of data center chilled water flows and meters is shown below.

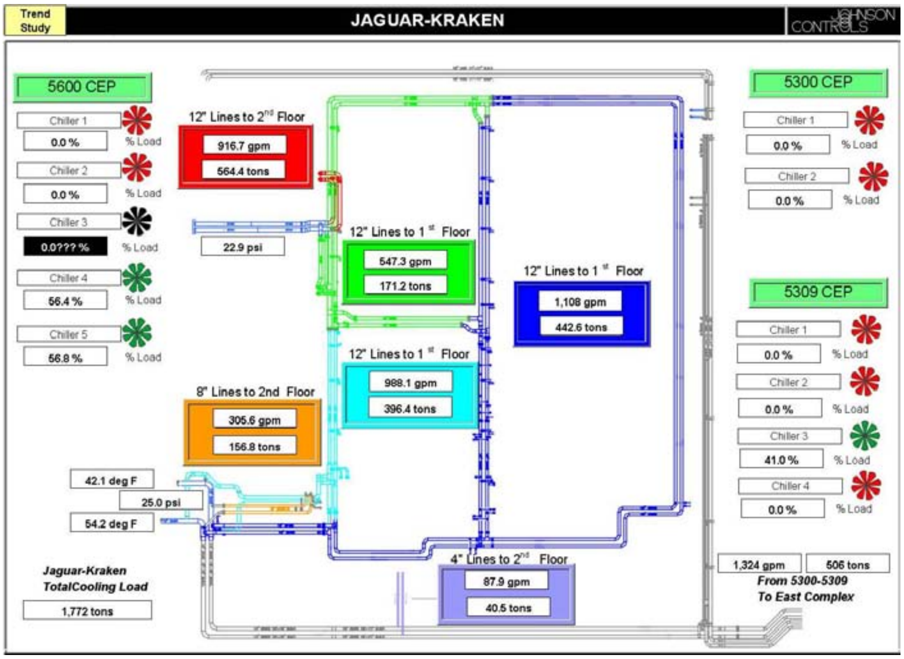

The cumulative ton-hr are totalized for the key values needed to calculate chilled water electricity, and the chilled water electricity data are used in calculating energy metrics.

8.5 Chiller Plant Measurements

The total chiller plant complex serves several buildings, including, the Computational Sciences Building, Research Office Building, and Engineering Technology Facility. These three buildings are all part of a single facility served mainly by the CSB chiller plant.

Chillers 1, 2, and 3 have the lowest operating efficiency and thus are run as little as possible, typically only in the summer when to meet overall building loads. Chillers 4 and 5 run almost continuously to meet the data center loads. An interconnection to the MRF complex provides some chilled water that typically is about equal to the general building loads much of the year. Chillers 4 and 5 are more efficient than chillers 1, 2, and 3, and the MRF Chiller Complex is more efficient than any of the CSB chillers.

Chilled water energy metering currently provides the following data points:

- Ton-hr produced by chiller 4
- Ton-hr produced by chiller 5
- Ton-hr delivered to first-floor data center
- Ton-hr delivered to 2nd-floor data center

Chilled water energy from chillers 1, 2, and 3 is not metered, but future plans include installation of these energy meters. The protocol for calculating chilled water energy uses chillers 4 and 5, together with the balance of plant installed with chillers 4 and 5 as representative of the total plant energy use for all chilled water delivered to the data centers. The chilled water interconnection from MRF is expected to be closed when the 5800 chiller plant comes on line and the current CSB chiller plant serves only data center loads.

Electricity use for chillers 4 and 5 (include related tower, pump, and peripheral electricity) is measured separately. The daily kW/ton for chillers 4 and 5 and peripherals is calculated as total daily kWh divided by total daily ton-hr delivered.

Chillers 4 and 5 and peripherals in the CSB at ORNL have a daily kW/ton of 0.62–0.75 in cold to mild weather, and 0.75–0.85 in hotter weather. The annual average appears to be around 0.75. Climate adjustment might make the climate-normalized value about 0.72. Daily chiller plant electricity for each of the data centers is calculated as: kW/ton x ton-hr/day.

9 First-Floor Performance Measurement Results

Measured results on total electricity for the first floor data center using the methodology presented previously are presented below. PUE values are presented for both the case with XT5 fan power included in total computer loads and the case without fan power included. The rationale for this presentation is that similar supercomputer CPU cabinets may have 32 small fans in each cabinet which may use nearly the same power as the one large fan in each XT5 cabinet uses. [Example: 32 fans at 70 W each would be 2.2 kW / cabinet]

Pumping power of the refrigerant cooling system is not measured and is included in the “technical power” total. This approach puts a high and low value on the performance metrics spectrum. Inclusion of the fans in the “technical power” suggests higher performance than actual, while removal of the fan energy completely suggests lower than actual performance.

Transformer losses have been measured on a spot basis and estimates should be close to actual. UPS losses are estimated at 6% of UPS load. Lighting energy is also estimated, based on fixture counts, power per fixture, and operating control. Estimated lighting power is expected to be close to actual.

A breakout of the July 2009 energy use is shown in the pie chart. Energy for Jaguar and Kraken includes refrigerant cooling system energy and XT5 fan energy. The notation in the figure means:

Jaguar = Jaguar supercomputer
Kraken = Kraken supercomputer
ERP = local power supplies on UPS power
RDUs = local power supplies
RDF = CSB main networking / comm facility
Disks etc = disk storage units not on ERP
CRUs = room AC units
Chiller = chiller plant energy
Lighting = room lighting
xfmr est = estimated losses from UPS and transformers

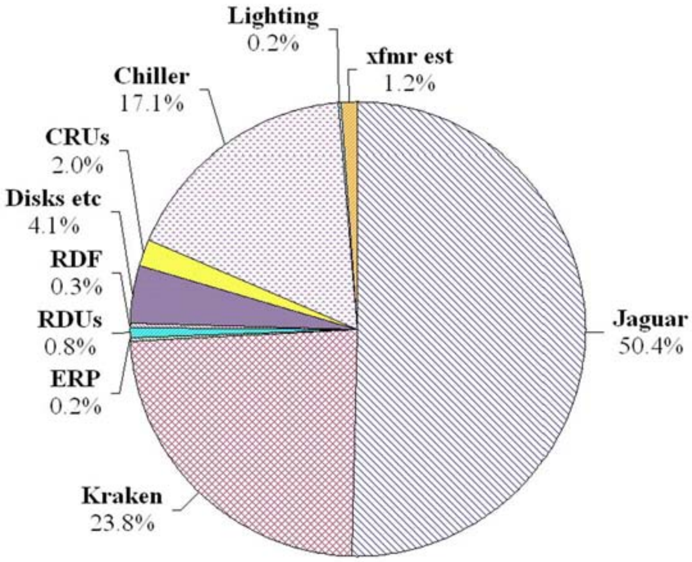

Breakout of 1st-floor data center power

The average power density in the CSB has increased from 150 W/sq-ft in 2007 to 250 W/sq-ft currently. Power density is expected to increase as the new National Oceanic and Atmospheric Administration supercomputer comes online.

The key performance metrics are summarized in the table below for the three cases of: XT5 fan power included in total technical power, XT5 fan power excluded completely, and a reasonable middle ground. The XT5 fan power is about 7.5% of total electricity use. Since all supercomputer cabinets have fan power that is regularly included in the technical power when calculating PUE, the “reasonable” case is bracketed by the other two cases in the table below for the first-floor data center in Building 5600. Since the XDP cooling system power is still included in the technical power total, a PUE = 1.33 appears most appropriate. PUE is expected to decrease in cold weather as cooling tower related energy consumption drops.

CSB First-Floor Metrics			
Metric	**Reasonable**	**w/ XT5 fans**	**w/o XT5 fans**
PUE	1.33	1.26	1.39
DCiE	75%	80%	72%

10 Second-Floor Performance Measurement Results

Measured results on total electricity for the second floor data center using the methodology presented previously are presented below. The second floor houses multiple systems, including an older Cray XT4 cabinet system that is still the #20 supercomputer in the world as of June 2010. The second floor houses the enterprise data systems for ORNL.

All computer systems are air-cooled. XT4 fan energy is included in the technical power for the breakdown below. Lighting energy is estimated for this data center also, based on fixture counts, power per fixture, and operating control. The lighting power estimate is expected to be close to actual. Transformer losses and UPS losses are also estimated, similar to the estimates for the first-floor data center.

A breakout of the July 2009 energy use is shown in the pie chart.

The notation is similar to the first floor breakout:

Cray cpu = older supercomputers
Other cpu = all other computers
ERP = local power supplies on UPS power
Disks etc = disk storage units and other local power not on ERP
CRUs = room AC units
Chiller = chiller plant energy
Lighting = room lighting
xfmr est = estimated losses from UPS and transformers

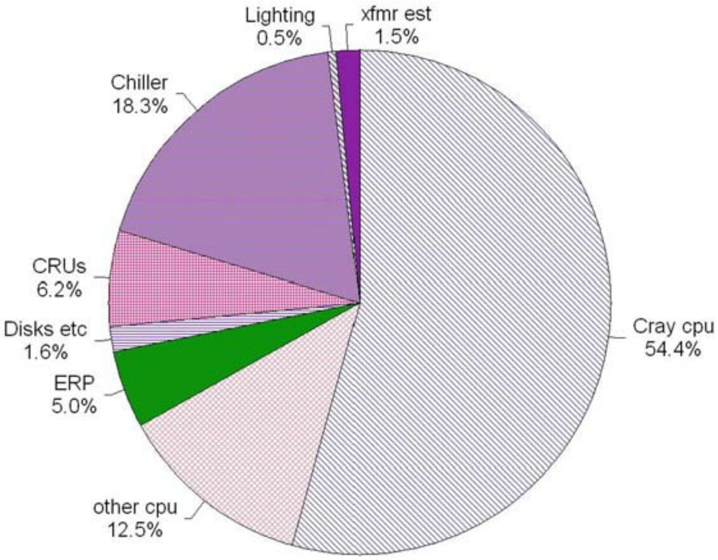

Breakout of 2nd-floor data center power

The key performance metrics for the entire 2nd-floor center are summarized in the table below for the three cases of: XT4 fan power included in total technical power, XT4 fan power excluded completely, and a reasonable middle ground. The XT4 fan power is estimated based on the measurements for the XT5 cabinets at 4.6% of total electricity use for the 2nd-floor datacenter. Similar to the first floor data center, when calculating PUE, the “reasonable” case is bracketed by the other two cases in the table below. A PUE = 1.41 appears most appropriate for the second floor. PUE is expected to decrease in cold weather as cooling tower related energy consumption drops.

CSB Entire Second-Floor Metrics			
Metric	**Reasonable**	**w/ XT4 fans**	**w/o XT4 fans**
PUE	1.41	1.36	1.45
DCiE	71%	73%	69%

11 Future Plans for the Data Centers

The data center at ORNL see a lot of change over time, and more changes are coming. A new large computer for the National Oceanic and Atmospheric Administration is being installed and ramped up in the first floor data center. Chiller plant reconfigurations and additional metering are in progress and will continue over the next year. The chilled water connection with the MRF complex likely will be completely closed after the new chiller plant in Building 5800 is fully functional.

Performance metrics will continue to be calculated monthly and reported based on daily measurements of electric use for over 50 electric meters and multiple chilled water energy meters. The current ability to report PUE on a monthly basis is deemed less than desirable by internal management. There is a need to have the capability to generate these values on a more real-time basis and then have these results displayed on an electronic dashboard. Even if instant values are not an option, the possibility of a time-delayed display would still be beneficial. A lag of a couple hours would still be helpful for diagnosing system performance.

The ability to report real time results will require major improvements in data base configurations that allow consistent data queries at regular intervals. Such queries are not possible with current data base configurations for the electric metering and building automation systems. ORNL is actively seeking out control products to allow for the integration of electric (PowerNet) and thermal (Metasys) data into one common source. It is predicted that this would allow for better control of the various datacenter support systems.

Efforts are being made to evaluate various energy technologies. For hot, humid climates like Tennessee, it would be desirable if computer cabinets could be made to function with only cooling tower water to cool the computers (or cool an intermediate ebullient system). Some room air cooling would still be needed to keep room dewpoints at acceptable levels. Thus far, the design of such an approach remains challenging. LED lighting technology will be tested in the near future at ORNL. The small pilot testing will take place before deciding on large scale implementation in the datacenters. Deployment of the technology will reduce energy consumption and alleviate maintenance issues.

Another large push is being made to find a control system that will optimize chiller plant efficiency. It is believed that there are large potential savings in being able to properly control and stage various cooling equipment on and off depending upon IT load and outdoor weather. ORNL is actively searching out various options to allow for continuous optimization of chiller plant operation.

In addition, ORNL is studying the benefits of utilizing non-OEM installed metering to deal with issues experienced in standard meters installed in OEM equipment. Facilities are generally plagued with issues resulting from standard metering that comes with purchased equipment, including: inaccessibility, poor performance, calibration difficulties, etc. Future studies will highlight issues and lessons learned at ORNL during the installation and pilot testing of two new meters being installed on power distribution panels.

The computing industry continues to experience high rates of change, and DOE supercomputer data centers also see high rates of change. The information in this report is intended to help others consider possible means of handling data center design, to document the energy performance metrics for the two data centers in Building 5600, and also to understand how DOE’s largest data center operates.

12 References

ASHRAE, *2007 ASHRAE HVAC Applications.* American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., 2007. Chapter 40 – Building Energy Monitoring.

ASHRAE, *2009 ASHRAE Handbook – Fundamentals.* American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., 2009. Chapter 36 – Measurement and Instrumentation.

ASHRAE & The Green Grid, *Real-Time Energy Consumption Measurements in Data Centers*. American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., 2009.

ASTM E 1465: *Standard Guide for Developing Energy Monitoring Protocols for Commercial and Institutional Buildings or Facilities.* ASTM International, 2005.

European Commission, *European Union Code of Conduct for Data Centres, Version 1.0,* 2008. http://re.jrc.ec.europa.eu/energyefficiency/html/standby_initiative_data%20centers.htm

MacDonald, M., *Energy Performance of ORNL Supercomputer Data Centers in Building 5600.* FEMP. December 2009.

MacDonald, J.M., Sharp, T.R., Getting, M.B., *A protocol for monitoring energy efficiency improvements in commercial and related buildings.* ORNL/Con-291, 1989.

NREL, *Best Practices Guide for Energy-Efficient Data Center Design.* FEMP. February 2010.

The Green Grid. www.thegreeengrid.org

United States Department of Energy, *Strategic Sustainability Performance Plan.* September 2010

United States Environmental Protection Agency, *Report to Congress on Server and Data Center Energy Efficiency, Public Law 109-431*, ENERGY STAR Program, 2007. http://www1.eere.energy.gov/femp/pdfs/epa_dc_report_congress.pdf

Appendix A: Other Data Center Metrics and Benchmarking

TGG has published several white papers on performance metrics for data centers. These and related metrics are discussed below. Further metrics information can be found in the FEMP publication, *Best Practices Guide for Energy-Efficient Data Center Design.*

Cooling System Efficiency
Overall chiller plant efficiency at many plants was studied by Ben Erpelding in California and an efficiency scale was developed, as shown below. This scale is one of the best ways to consider overall chiller plant efficiency (wire-to-water efficiency) and also one of the easiest. This chart applies to annual average kW/ton. This chart has been published in many forms in several places (e.g., *HPAC Engineering*, May 2007).

Cooling System Efficiency = Average Cooling System Power (kW) / Average Cooling Load (ton)

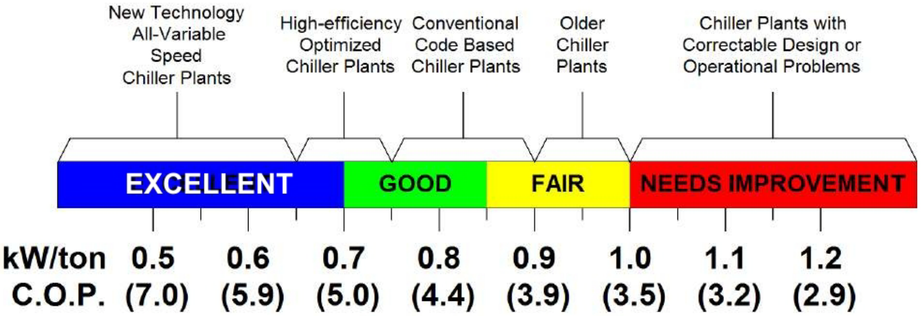

AVERAGE ANNUAL CHILLER PLANT EFFICIENCY IN KW/TON (C.O.P.)
(Input energy includes chillers, tower fans, and condenser & chilled water pumping)

Based on electrically driven centrifugal chiller plants in comfort conditioning applications with 42F (5.6C) nominal chilled water supply temperature and open cooling towers sized for 85F (29.4C) maximum entering condenser water temperature.
Local Climate adjustment for North American climates is +/- 0.05 kW/ton

Benchmark Values -

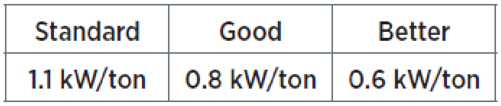

Standard	Good	Better
1.1 kW/ton	0.8 kW/ton	0.6 kW/ton

Airflow Efficiency
This metric provides an understanding of how efficiently air is moved through a data center.

Airflow Efficiency = Total Fan Power (W) / Total Fan Airflow (cfm)

Benchmark Values -

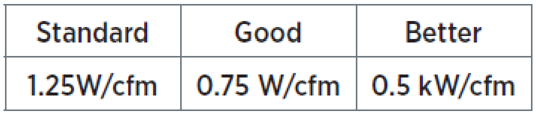

Standard	Good	Better
1.25W/cfm	0.75 W/cfm	0.5 kW/cfm

Heating, Ventilation and Air-Conditioning (HVAC) System Effectiveness
This metric is simply the ratio of the annual IT equipment energy consumption to the annual HVAC energy consumption.

Effectiveness = $[\text{kWh/yr}]_{IT} / [\text{kWh/yr}]_{HVAC}$

Benchmark Values –

Standard	Good	Better
0.7	1.4	2.5

Rack Cooling Index (RCI)
The rack cooling index is a measure of compliance with ASHRAE/NEBS temperature specifications. It effectively gives a numerical representation of how well equipment racks are cooled based on equipment intake temperatures. The equations below reflect ASHRAE Class 1 (2008) conditions.

$$RCI_{HI} = \left[1 - \frac{\sum_{T_x > 80} (T_x - 80)}{(90 - 80)n}\right] \times 100\ [\%] \qquad RCI_{LO} = \left[1 - \frac{\sum_{T_x < 65} (65 - T_x)}{(65 - 59)n}\right] \times 100\ [\%]$$

Where,
T_x = Mean temperature at equipment intake x
n = Total number of intakes

Benchmark Values -

RCI_{HI} = 100%	No temperature above max recommended
RCI_{LO} = 100%	No temperature below min recommended
$RCI_{HI/LO}$ < 90%	Often considered poor operation

Return Temperature Index (RTI)
The return temperature index is a measure of the net by-pass or net recirculation of air in a data center.

$$RTI = \frac{\Delta T_{AHU}}{\Delta T_{EQUIP}} \times 100\%$$

Where,
ΔT_{AHU} = Typical air handler temperature drop (airflow weighted)
ΔT_{EQUIP} = Typical IT equipment temperature rise (airflow weighted)

Benchmark Values -

RTI = 100%	Balanced airflows
RTI > 100%	Recirculating air
RTI < 100%	By-passing air

EERE Information Center
1-877-EERE-INFO (1-877-337-3463)
www.eere.energy.gov/informationcenter

U.S. DEPARTMENT OF ENERGY | Energy Efficiency & Renewable Energy

ORNL/TM-2011/49 September 2010

Prepared by the Oak Ridge National Laboratory (ORNL). ORNL is a national laboratory of the U.S. Department of Energy operated by UT-Battelle.