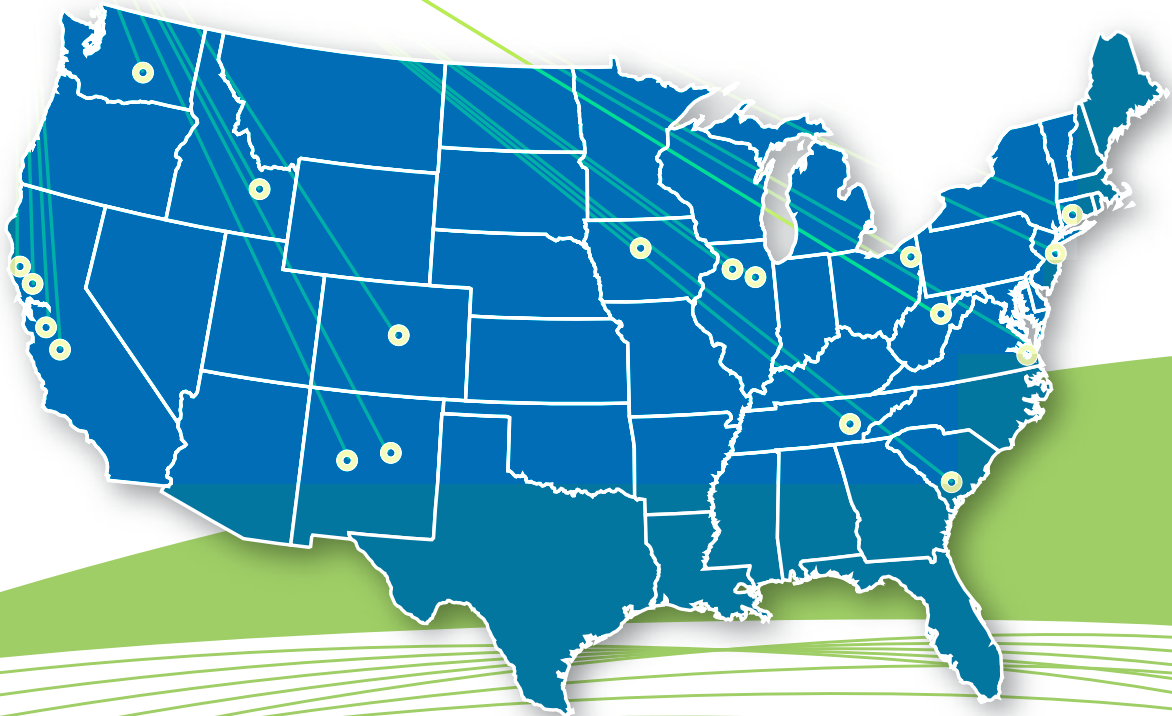


DEPARTMENT OF ENERGY LABORATORIES

Leadership in Green IT



U.S. DEPARTMENT OF
ENERGY

Prepared by the National Renewable Energy Laboratory (NREL), a national laboratory of the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy; NREL is operated by the Alliance for Sustainable Energy, LLC.

Prepared by: Chuck Powers, IT Strategist
National Renewable Energy Laboratory
Email: chuck.powers@nrel.gov
Phone: 303-275-4182 ■ Website: www.nrel.gov



TABLE OF CONTENTS

Table of Tables	iv
Table of Figures	iv
Introduction	1
Department of Energy Laboratories	2
Showcase of Accomplishments, Leadership, and Innovation	7
Sustainable Management of Desktops, Laptops, and Associated Peripherals	7
Cloud Computing and Server Virtualization	12
Energy Efficient Data Centers	18
DOE Labs Redefine World-Class PUE	19
Data Center Efficiencies	19
Cooled with “Free” Cooling	22
Powered by Renewable Energy	25
Waste Heat Reuse	26
Green IT Leadership Provided by Labs	27
High Performance Computing	31
Conclusion	38
References	39
Resources for Green IT	40



TABLE OF TABLES

Table 1. FEC partner definition levels	8
Table 2. 2010 FEC winners	10
Table 3. 2011 FEC applicants	10
Table 4. PNNL Comparison of hardware and electricity costs between standalone and virtualized servers	18
Table 5. DOE Lab World Class PUEs	19
Table 6. Sandia cold aisle containment system before and after comparison	21
Table 7. INL HPC data center PUE	23
Table 8. DOE labs in Green 500 List top 100	32

TABLE OF FIGURES

Figure 1. NETL’s LEED Gold Technology Support Facility (TSF)	10
Figure 2. ORNL dashboard illustrating substantial improvements in six key measures	11
Figure 3. Sandia Labs’ HPC clusters	12
Figure 4. Server equipment avoided through virtualization	13
Figure 5. FNAL’s Grid Computing Center interior	13
Figure 6. Energy reduction from green IT initiative	14
Figure 7. LANL achieves return on investment	14
Figure 8. The Magellan management and network control racks at NERSC	16
Figure 9. Cisco UCS	16
Figure 10. Server virtualization effects—20 servers on one blade server	17
Figure 11. PNNL ISB2/1 data center power capacity	17
Figure 12. PPPL’s computing environment	18
Figure 13. Ames energy monitoring results	19
Figure 14. LBNL wireless environmental monitoring system	20
Figure 15. Effect of humidification control set to “off”	20

TABLE OF FIGURES (CONTINUED)

Figure 16. Effect of adding blanking panel to rack of high-performance servers	20
Figure 17. Sandia Labs' cold aisle containment curtain	21
Figure 18. Savannah River power consumption results	22
Figure 19. ALCF's Blue Gene/P supercomputer	23
Figure 20. RSF data center cooling and reuse of heat from hot aisle containment system	24
Figure 21. NREL's Data center air intake	24
Figure 22. PNNL aquifer cooling system	24
Figure 23. PPPL's hot aisle containment enclosure	25
Figure 24. NREL's roof-installed photovoltaic panels	25
Figure 25. Sandia Labs' DETL existing array w 9300 SW array	26
Figure 26. Hot aisle containment system used in NREL's RSF data center	26
Figure 27. PNNL CSF/BSF building entry area	27
Figure 28. IBM Rear Door Heat eXchanger	27
Figure 29. A high-level view of the Energy Sciences Network	28
Figure 30. Sample DC Pro Profiler output for a data center	29
Figure 31. Best practices guide for energy-efficient data center design	29
Figure 32. National Snow and Ice Data Center CRAC upgrade results	30
Figure 33. SLAC and Stanford University modular data center design concept	31
Figure 34. ALCF's "Intrepid" memory modules	33
Figure 35. INL data center chilled water metrics	34
Figure 36. INL HPC data center PUE results	34
Figure 37. Jefferson Lab's Hadron Supercomputer	34
Figure 38. Jefferson lab's 2010 high performance computing clusters	35
Figure 39. LLNL Terascale Simulation Facility	35
Figure 40. ORNL's Jaguar HPC	37
Figure 41. ORNL's Jaguar facilitates design of an energy-efficient UnderTray System for semis	37
Figure 42. Sandia's Red Sky HPC	38

Information technology (IT) is the driving force behind managing the billions of bytes of data processed globally. IT equipment—including servers, storage, networks, and end-user computing devices—is estimated to use a massive 9.4% of the total amount of energy produced in the United States. Globally, these devices consume 5.3% of all of the energy produced.¹ Of this consumption, 3% of all energy produced is used in U.S. data center facilities alone; and, demands are increasing. Today, public utilities and organizational infrastructure are challenged to meet the increasing demand. This growing concern opens the door for significant opportunities for the implementation of energy efficiencies.

In response to these challenges, the federal government has issued Executive Orders 13514 and 13423, which require all government agencies to meet certain energy reduction and environmental performance requirements by 2020. The Department of Energy (DOE) has set specific goals that their national facilities must meet to comply with these requirements. DOE laboratories have risen to the challenge and are well ahead of schedule with demonstrated leadership and innovative practices in IT sustainability in their 26 data centers.

Join the 17 DOE laboratories as we recap stories about our efforts to make each laboratory's energy-intensive IT capabilities “green” and supportive of DOE sustainability efforts.

The report highlights examples of practices, plans, and successes in the following areas.

- Sustainable management of IT devices
- Virtual and high performance computing environments
- Planning and implementation best practices
- Partnering with nature for “free cooling”
- Power using renewable resources
- Waste reuse and recycling
- Energy efficient data center practices
- Leadership through partnering with industry and each other.

We would like to thank the national laboratory chief information officers for making this report possible and allowing us to present our accomplishments to the world. We hope this report offers ideas about how other IT organizations can benefit from our planning and implementation in their efforts toward IT sustainability.

The DOE laboratories

¹ Sarokin, D. “Question: Energy Use of Internet.” *Uclue* online, <http://uclue.com/index.php?xq=724>. Accessed May 26, 2011.

DEPARTMENT OF ENERGY LABORATORIES

The 17 DOE laboratories that comprise the nation's federal scientific research system provide strategic scientific and technological capabilities. Their collective goal is to meet the nation's challenges and priorities in these areas, which often reach beyond the scope of academia and private industry; and also to ensure that our government has access to these crosscutting discoveries and innovations.²

The following are brief overviews of the missions of each lab and the IT professionals who contributed to this report.

AMES LABORATORY

The Ames Laboratory (Ames) is dedicated to energy research and is a national center for the synthesis, analysis, and engineering of rare-earth metals and their compounds. Its mission is to create “innovative materials, technologies and energy solutions” and to leverage the laboratory's capabilities to devise solutions to global challenges. Ames conducts fundamental research in the physical, chemical, and mathematical sciences associated with energy generation and storage.

Mark Grootveld, Ames Laboratory, USDOE ■ 111 TASF ■ Ames, IA 50011
Phone: 515.294.7895 ■ E-mail: grootveld@ameslab.gov ■ Web: <http://www.ameslab.gov/>

ARGONNE NATIONAL LABORATORY


Argonne National Laboratory (ANL) is one of the DOE's oldest and largest laboratories dedicated to science and engineering research. This multidisciplinary research center provides innovative research in basic science, scientific facilities, energy resources programs, environmental management and national security. The lab strives to “create new knowledge that addresses the most important scientific and societal needs of our nation.”

Paul Domagala, Argonne National Laboratory ■ 9700 S. Cass Avenue ■ Argonne, IL 60439
Phone: 630.252.5197 ■ E-mail: domagala@anl.gov ■ Web: <http://www.anl.gov/>

BROOKHAVEN NATIONAL LABORATORY

Brookhaven National Laboratory (BNL) was established in 1947. It is a multi-program lab conducting research in physical, biomedical, and environmental sciences; energy technologies; and national security. BNL has received seven Nobel Prizes for discoveries made at the lab.

Frank Quarant, Brookhaven National Laboratory ■ P.O. Box 5000, NY 11973
Phone: 631.344.2438 ■ E-mail: quarant@bnl.gov ■ Web: <http://www.bnl.gov/world/>



² U.S. Department of Energy, Office of Science. “Laboratories.” *Office of Science* Online, <http://science.energy.gov/laboratories/> Accessed April 19, 2011.



FERMI NATIONAL ACCELERATOR LABORATORY

The Fermi National Accelerator Laboratory (FNAL) advances understanding of the fundamental nature of matter and energy. FNAL's world-class scientific research facility allows qualified researchers from around the world to conduct fundamental research at the frontiers of high-energy physics and related disciplines.

Eric Mieland, Fermi National Accelerator Laboratory ■ P.O. Box 500 ■ Batavia, IL 60510-5011
 Phone: 630.840.2248 ■ E-mail: mieland@fnal.gov ■ Web: <http://www.fnal.gov/>



IDAHO NATIONAL LABORATORY

The Idaho National Laboratory (INL) has been in operation since 1949. The lab is a science-based, applied engineering national laboratory dedicated to ensuring the nation's energy security with safe, competitive, and sustainable energy systems and unique national and homeland security capabilities.

Rebecca Ristrem, Idaho National Laboratory ■ P.O. Box 1625 ■ Idaho Falls, ID 83415
 Phone: 208.526.1588 ■ E-mail: rebecca.ristrem@inl.gov
 Web: <https://inlportal.inl.gov/portal/server.pt/community/home>



LAWRENCE BERKELEY NATIONAL LABORATORY

The Lawrence Berkeley National Laboratory (LBNL) conducts unclassified research across a range of scientific disciplines. Its key efforts are in fundamental studies of the universe, quantitative biology, nanoscience, new energy systems, and environmental solutions; and the use of integrated computing as a tool for discovery. Founded in 1931, the laboratory boasts 11 scientists who have won the Nobel Prize, and has many other distinguished awards to its credit.

Adam Stone, Lawrence Berkeley National Laboratory ■ 1 Cyclotron Road Mail Stop 65-0113
 Berkeley, CA 94720- 8105 ■ Phone: 510.486.4650 ■ E-mail: adstone@lbl.gov
 Web: <http://www.lbl.gov/>



LAWRENCE LIVERMORE NATIONAL LABORATORY

The Lawrence Livermore National Laboratory (LLNL) was founded on 1952. LLNL is dedicated to ensuring the safety and security of the nation through applied science and technology in nuclear security, international and domestic security, and energy and environmental security.

Mike Payne, Lawrence Livermore National Laboratory ■ P.O. Box 808, Livermore, CA 94551-0808
 Phone: 925.422.0661 ■ E-mail: payne29@llnl.gov ■ Web: <https://www.llnl.gov/>

DEPARTMENT OF ENERGY LABORATORIES

LOS ALAMOS NATIONAL LABORATORY

Los Alamos National Laboratory (LANL) is a premier national security research institution. Since 1978, the lab has delivered scientific and engineering solutions for the nation's most crucial and complex problems. LANL's primary responsibility is ensuring the safety, security, and reliability of the nation's nuclear deterrent. The lab also advances bioscience, chemistry, computer science, earth and environmental sciences, materials science, and physics disciplines.

Virginia Romero, Los Alamos National Laboratory ■ P.O. Box 1663 ■ Los Alamos, NM 87545
Phone: 505.667.9682 ■ E-mail: vir@lanl.gov ■ Web: <http://www.lanl.gov/>



NATIONAL ENERGY TECHNOLOGY LABORATORY

The National Energy Technology Laboratory (NETL) was designated as a national laboratory in 1999 after a long history of supporting the nation's fossil energy resource needs. NETL implements a broad spectrum of energy and environmental research and development programs that enable domestic coal, natural gas, and oil to economically power our nation, while protecting our environment and enhancing our energy independence.

Mark Cerullo, National Energy Technology Laboratory ■ 626 Cochrans Mill Road
Pittsburgh, PA 15236-0940 ■ Phone: 304.285.4345 ■ E-mail: mark.cerullo@netl.doe.gov
Web: <http://www.netl.doe.gov/>



NATIONAL RENEWABLE ENERGY LABORATORY

In operation since 1977, The National Renewable Energy Laboratory (NREL) is the nation's only laboratory dedicated solely to renewable energy and energy efficiencies research and development. NREL develops renewable energy and energy efficiency technologies and practices, advances related science and engineering, and transfers knowledge and innovations to address the nation's energy and environmental goals. These areas span from understanding renewable resources for energy, to the conversion of these resources to renewable electricity and fuels, and ultimately to the use of renewable electricity and fuels in homes, commercial buildings, and vehicles.

Chuck Powers, National Renewable Energy Laboratory ■ 1617 Cole Blvd
Golden, CO 80401-3305 ■ Phone: 303.275.4182 ■ E-mail: chuck.powers@nrel.gov
Web: <http://www.nrel.gov/>



DEPARTMENT OF ENERGY LABORATORIES



OAK RIDGE NATIONAL LABORATORY

The Oak Ridge National Laboratory (ORNL) is a multi-program science and technology laboratory established in 1943. The lab conducts basic and applied research and development to create scientific knowledge and technological solutions that strengthen the nation's leadership in key areas of science. These include increasing the availability of clean and abundant energy, restoring and protecting the environment, and contributing to national security.


Sheila Causby, Oak Ridge National Laboratory ■ P.O. Box 2008 ■ Oak Ridge, TN 37831
Phone: 865.241.9786 ■ E-mail: causbysl@ornl.gov ■ Web: <http://www.ornl.gov/>



PACIFIC NORTHWEST NATIONAL LABORATORY

The Pacific Northwest National Laboratory (PNNL) has delivered leadership and advancements in science, energy, national security, and the environment since 1965. The lab conducts applied research in information analysis, cyber security, and the nonproliferation of weapons of mass destruction; research in hydrogen and biomass-based fuels to reduce U.S. dependence on oil; and works to reduce the effects of energy generation and use on the environment.

Bary Elison, Pacific Northwest National Laboratory ■ P.O. Box 999 ■ Richland, WA 99352
Phone: 509.372.4595 ■ E-mail: bary.elison@pnl.gov ■ Web: <http://www.pnl.gov/>



PRINCETON PLASMA PHYSICS LABORATORY

The Princeton Plasma Physics Laboratory (PPPL) is a national center dedicated to plasma and fusion science with a leading international role in developing the theoretical, experimental, and technology innovations needed to make fusion practical and affordable. Since 1951, PPPL has worked with collaborators across the globe to develop fusion as an energy source for the world, and conduct research along the broad frontier of plasma science and technology.

Robert Sheneman, Princeton Plasma Physics Laboratory ■ P.O. Box 451 ■ Princeton, NJ 08543-0451 ■ Phone: 609.243.3392 ■ Email: rsheneman@pppl.gov
Web: <http://www.pppl.gov/>



SANDIA NATIONAL LABORATORIES

Sandia National Laboratories (Sandia Labs) has developed science-based technologies that support national security since 1949. The lab develops technologies to sustain, modernize, and protect the U.S. nuclear arsenal; prevent the spread of weapons of mass destruction; defend against terrorism; protect national infrastructures; ensure stable energy and water supplies; and provide new capabilities to the U.S. armed forces.

David Martinez ■ Sandia National Laboratories, New Mexico ■ PO Box 5800 Albuquerque, NM 87185-0165 ■ Phone: 505.844.6531 ■ E-mail: davmart@sandia.gov
Web: <http://www.sandia.gov/>

DEPARTMENT OF ENERGY LABORATORIES

SAVANNAH RIVER NATIONAL LABORATORY

The Savannah River National Laboratory (Savannah River) was founded in 2004 and is the applied research and development laboratory at the Savannah River Site (SRS). The lab is dedicated to solving complex national defense, homeland security, and nuclear material problems. They also provide applied research in environmental management, energy security, and technologies.

*David Wolfe, Savannah River National Laboratory ■ Savannah River Site
Aiken, SC 29808 ■ Phone: 803.557.9911 ■ E-mail: david02.wolfe@srs.gov
Web: <http://Savannah.River.doe.gov/>*

SLAC NATIONAL ACCELERATOR LABORATORY

The SLAC National Accelerator Laboratory (SLAC) is dedicated to the design, construction and operation of state-of-the-art electron accelerators and related experimental facilities for use in high-energy physics and synchrotron radiation research. Founded in 1962, SLAC is a multipurpose laboratory for astrophysics, photon science, accelerator and particle physics research. The lab boasts six Nobel Prize winning scientists.

*Ted Schnipper, SLAC National Accelerator Laboratory ■ 2575 Sand Hill Road. Mail Stop 58
Menlo Park, CA 94025-7015 ■ Phone: 650.926.7438 ■ E-mail: teds@slac.stanford.edu
Web: <http://www.slac.stanford.edu/>*

THOMAS JEFFERSON NATIONAL ACCELERATOR FACILITY

The Thomas Jefferson National Accelerator Facility (JLab) began operation in 1995. The lab provides forefront scientific facilities, opportunities, and leadership essential for discovering the fundamental nature of nuclear matter, to partner with industry to apply its advanced technology, and to serve the nation and its communities through education and public outreach. Scientists from around the world use the laboratory's facilities to conduct their research.

*Andy Kowalski, Jefferson Laboratory ■ 12000 Jefferson Avenue ■ Newport News, VA
23606 ■ Phone: 757.269.6224 ■ E-mail: kowalski@jlab.org ■ Web: <http://www.jlab.org/>*



SHOWCASE OF ACCOMPLISHMENTS, LEADERSHIP, AND INNOVATION

The DOE national laboratories are tasked with finding ways to support the nation's energy goals through identifying and implementing sustainability practices across their sites. The IT groups have embraced this challenge through innovations and leadership in Green IT. This section highlights the accomplishments of the labs in five key areas:

- Sustainable management of desktops, laptops, and associated peripherals
- Cloud computing and server virtualization
- Energy-efficient data centers
- Green IT leadership provided by labs
- Sustainable use of high-performance computing (HPC) resources.

■ SUSTAINABLE MANAGEMENT OF DESKTOPS, LAPTOPS, AND ASSOCIATED PERIPHERALS

IT end-user computing devices consume a considerable amount of energy. DOE labs have implemented several practices to optimize power use at the desktop. This section highlights achievements at DOE laboratories in sustainable management of these IT resources.

The electronic equipment used in offices have a significant impact on the amount of waste going into landfills and CO₂ production.

Full life-cycle management

The electronic equipment used in offices not only use a lot of power, but also have a significant impact on the amount of waste that goes into landfills and the carbon produced in their manufacturing and delivery.

Presidential Executive Order 13514 focuses on “Federal Leadership in Environmental, Energy, and Economic

Performance.” All federal facilities must meet specific requirements under this order, but many DOE labs have implemented practices that far exceed these expectations.

The Federal Electronics Challenge (FEC) was created to help implement the sustainable practices that the order was intended to achieve.¹ The FEC partners with federal agencies to help them meet and exceed the order. The FEC also

rewards federal facilities for their initiative in documenting how they met the requirements of the executive order, and to exceed expectations by going above and beyond those requirements.

The FEC awards are coveted achievements and designate leadership in the following phases of the electronic equipment life cycle:

- Acquisition and Procurement— Participant procurement processes must require that both leased and purchased electronic equipment be Electronic Product Environmental Assessment Tool (EPEAT™) registered
- Operation and Maintenance— The environmental impact of electronic equipment must be reduced during operation, and its useful life must be no longer than four years
- End-of-Life Management— Obsolete electronics must be managed in a sustainable manner, including environmentally sound disposal practices and recycling.

There are four possible FEC award levels—Platinum (new in 2011), Gold, Silver, and Bronze. These levels require that a certain number of mandatory and optional activities be completed to apply for a certain level, and that the level below is achieved prior to applying for the next highest level.

¹ Federal Electronics Challenge. <http://www.federalelectronicschallenge.net/>. Accessed May 02, 2011.

SUSTAINABLE MANAGEMENT OF DESKTOPS, LAPTOPS, AND ASSOCIATED PERIPHERALS

Table 1 details the mandatory and optional activities required to be awarded FEC partnership.

What the labs are doing to manage the full life cycle and obtain FEC standards

The DOE labs are doing several things to manage the full life cycle of electronic equipment, including hiring dedicated staff to manage the process. The labs have implemented procurement processes that require the acquisition of equipment that is EPEAT and Energy Star compliant, have instituted programs that recycle or reuse equipment, are modifying print practices to reduce energy and paper use, and are right-sizing employee computer processing needs by moving toward thin clients.

Procurement

The labs are embracing EPEAT and Energy Star standards, which support the full life cycle of the electronic product itself. The DOE labs are also collectively working toward making smart equipment decisions, even before the product reaches the labs.

The Energy Star program was developed jointly by the Environmental Protection Agency (EPA) and DOE to identify electronic products (from computers to refrigerators) that reduce greenhouse gas (GHG) emissions and save money for both businesses and consumers. The DOE labs have a commitment to buy only electronics with

Table 1. FEC partner definition levels

AWARD LEVEL	MANDATORY ACTIVITIES	OPTIONAL ACTIVITIES
● Platinum	Complete all mandatory activities in all three of the life-cycle phases, and complete one mentoring activity	Complete three optional site-selected activities in any life-cycle phase, or one that is not selected from the mandatory list, but supports FEC goals
● Gold	Complete all mandatory activities in all three of the life-cycle phases, and complete two mentoring activities	Complete nine optional site-selected activities in any life-cycle phase, or one that is not selected from the mandatory list, but supports FEC goals
● Silver	Complete all mandatory activities in any two of the three life-cycle phases	Complete six optional site-selected activities in any life-cycle phase, or one that is not selected from the mandatory list, but supports FEC goals
● Bronze	Complete all mandatory activities in any one of the three life-cycle phases	Complete three optional site-selected activities in any life-cycle phase, or one that is not selected from the mandatory list, but supports FEC goals

the Energy Star certification, which is a start to reducing the labs' carbon footprint and energy cost savings even before the equipment ships.

EPEAT is a global registry tool that evaluates electronic equipment on 23 required environmental performance criteria and a percentage of 28 optional criteria. These criteria follow the entire manufacturing life cycle of the product, which includes design, production, the amount of energy it uses during operation, and the amount of recycled materials it includes. EPEAT also independently verifies manufacturer claims. Based on how they measure up in the tool, businesses can qualify as a Bronze, Silver, or Gold partner. The DOE labs strive to purchase products from manufacturers that are certified EPEAT partners.

The DOE labs have committed to purchasing only electronics that are Energy Star & EPEAT compliant.

The labs are taking equipment purchasing even further to reduce the carbon footprint caused by desktop printing and personal computer use.

SUSTAINABLE MANAGEMENT OF DESKTOPS, LAPTOPS, AND ASSOCIATED PERIPHERALS

All labs have already implemented or plan to implement thin clients in place of laptop and desktop computers for employees who do not need extensive computing power. Thin clients offer several advantages above and beyond energy footprint reduction, including the following:



- Increased data security with all data living on data center servers
- Smaller grain size at the desktop
- Greater agility of where and how work can be done.

Printing

Most labs have, or have plans to, replace copiers, fax machines, and individual printers with centralized multifunction printers. These devices reduce not only the energy required for printing, but also support programs aimed toward reducing printing as a whole.

Multifunction print devices require that employees travel farther to retrieve their print jobs, which prompts them to review documents on their computers, rather than sending to print. These devices are Energy Star compliant and have power management built in. Additionally, the labs are purchasing paper that is composed of at least 30% post-consumer recycled materials for these printers, in line with the goals of E.O. 13514.

To further support printing goals, the labs are instituting automatic duplex printing, black-and-white print unless color is absolutely necessary, and the use of solid ink toner cartridges.

Equipment end-of-life management

Another important area of full life cycle management is the disposition of electronic equipment once it has fulfilled its useful life. The labs have

implemented several practices for end-of-life management, all targeted toward reuse or recycling in a secure and environmentally sound manner.

The following are some of the ways the DOE labs manage equipment end of life:

- Reuse within the lab environment—can the equipment be repurposed in a way that still meets sustainability goals?
- Resale—can the equipment be purged of lab data and resold outside the lab as office equipment?
- Donation—can the equipment be purged of lab data and donated?
- Recycling—can the equipment be sold for recycled parts?
- Return to vendor—can the equipment be returned to the vendor at end of lease?
- Disposed—can the equipment be disposed of in an environmentally responsible manner?

DOE Lab Achievements in FEC

DOE labs have met the sustainable electronics life cycle management challenge by applying for and

achieving FEC status in one, two, or all of the life-cycle phases. DOE labs are showing leadership in full life cycle management by achieving one of the highest participation rates in FEC. Each year, more DOE labs apply for and achieve FEC awards; and keep coming back the next year to reach even higher!

Tables 2 and **3** list awards for DOE labs in 2010 and applications for 2011. It is important to note that the labs who do not participate in FEC have met the requirements of the executive order and may have plans to participate in FEC in the future.

NETL implements Best Practices to Manage Energy Consumption

NETL has implemented energy efficiency best practices for the management of their end-user computing devices and monitoring of their IT environments (**Figure 1**). They have developed environmental management plans (EMP) that are tracked quarterly and detail the lab's strategies for reducing environmental impact.

The lab achieved its goal to buy at least 95% EPEAT and Energy Star-designated electronic equipment in FY 2010. It implemented a monitoring and tracking system in FY 2011

SUSTAINABLE MANAGEMENT OF DESKTOPS, LAPTOPS, AND ASSOCIATED PERIPHERALS

Table 2. 2010 FEC winners

Lab	Bronze	Silver	Gold	Platinum
Idaho National Laboratory		●		
Lawrence Livermore National Laboratory	●			
National Energy Technology Laboratory		●		
National Renewable Energy Laboratory			●	
SLAC National Accelerator Laboratory	●			

to monitor the full life cycle of that equipment and will continue to replace energy-inefficient equipment with products that meet the efficiency standards. Power management, duplex printing, and other management practices are now standard on all IT devices.

NETL began monitoring power consumption in its data centers in 2011 to determine the baseline for the laboratory’s energy efficiency metrics. NETL began using DC Pro—a data center energy efficiency monitoring tool developed in part by LBNL—to identify efficiency opportunities and track performance. With the baseline metrics in place, NETL will continue to use DC Pro to evaluate its progress in

Table 3. 2011 FEC applicants

Lab	Bronze	Silver	Gold	Platinum
Idaho National Laboratory	●			
Lawrence Livermore National Laboratory	●			
National Renewable Energy Laboratory				●
Princeton Plasma Physics Laboratory	●			
Sandia National Laboratories			●	

comparison to the baseline. Recently, the lab installed new computer room air conditioner (CRAC) units to replace inefficient units, and installed an efficient uninterrupted power supply (UPS). The lab is also working on a plan to rearrange data center racks in a hot aisle/cold aisle fashion, with the intent to install more efficient HVAC systems to address requirements.

The lab is installing sub-metering to allow data center managers to view energy consumption real-time and to optimize management practices immediately.

ORNL IT Sustainability Dashboard

In fiscal year (FY) 2009, ORNL Information Technology Services Division (ITSD) piloted computer power management on Windows-based desktop computers. A primary consideration during this testing was the reduction of energy consumption without creating difficulties or associated issues for users. The lab deployed Verdiem® Surveyor, a software application system that enforces a set of energy-management policies through the central IT system, and piloted it on a small number of Windows® desktop computers. Following a successful 30% reduction in power use for the pilot group, power management software was approved to roll out lab-wide.



Figure 1. NETL's LEED Gold Technology Support Facility (TSF) .
Photo courtesy of NETL

SUSTAINABLE MANAGEMENT OF DESKTOPS, LAPTOPS, AND ASSOCIATED PERIPHERALS

The concept behind Verdiem Surveyor is the ability to “sleep” a computer after a standard work day has been identified, and then to initiate a “wake” to start a remote desktop connection to systems and monitors after nonworking hours. This initiative was achieved, but it was far from a simple “out-of-the-box” answer. The ITSD modified the application significantly to overcome reporting limitations, management concerns, and power management issues with various assets. IT personnel addressed desktop computers of various ages, laptops, research-oriented systems, and various nuances within the deployment effort.

Benefits of this initiative included:

- The ability to model energy-saving opportunities in an area of operations common to most organizations
- Energy savings through standardized processes and procedures
- The ability to monitor power consumption of IT equipment in real time.

This customized power management application is deployed on all 8,000 Windows desktop operating systems. Exemptions were given for systems that must run 24 hours a day to support research activities. **Figure 2** shows the sustainability improvements in six key measures, including Money, Energy, CO₂, Fuel, Trees, and Cars.

As of August 13, 2010, Verdiem is part of the standard configuration for ORNL desktop computer systems. Implementing this project is one of the important steps ORNL has taken—and will continue to take—as part of its effort to reduce the lab’s energy consumption and greenhouse gas emissions.

Sandia Labs’ Red Sky Foam Recycling Pilot Project

Sandia Labs’ Red Sky high-performance computer (HPC) cluster installation team participated in a pilot diversion effort to keep packaging foam NO. 4 LDPE (low-density polyethylene) from being disposed in the solid waste. Thirty cubic yards of foam material, weighing between 200 and 300 pounds, was diverted.

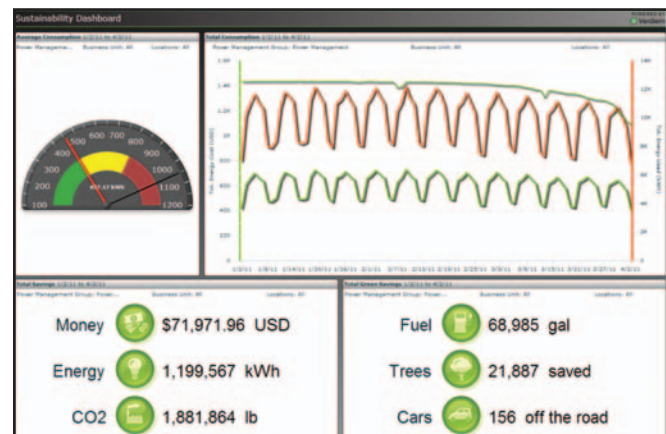
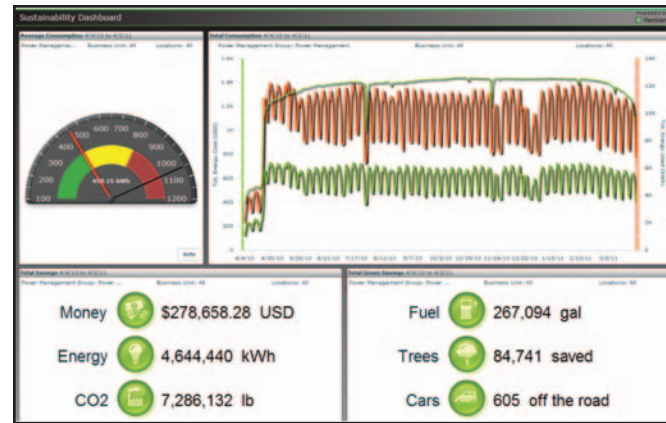


Figure 2. ORNL dashboard illustrating substantial improvements in six key measures. Image courtesy of ORNL

The Annex of Building 880 houses many HPC clusters, including the new Red Sky, installed in spring 2009 (**Figure 3**). These high-dollar pieces of equipment are packaged very thoroughly to minimize the chance of breakage during shipment from the manufacturer. The packaging is then disposed of after the equipment is opened.

The installation of the Red Sky HPC generated large quantities of metal, cardboard, and wood for recycling. Low-density polyethylene (LDPE) foam was also generated and is also labeled as a Plastic Resin No. 4 recyclable. Knowing that the project would generate a huge surge in

CLOUD COMPUTING AND SERVER VIRTUALIZATION



Figure 3. Sandia Labs' HPC clusters. Photo courtesy of Sandia National Laboratories

waste, the Red Sky installation team contacted the Solid Waste Transfer Facility (SWTF) and Pollution Prevention (P2) in advance. The timing of the installation project coincided with P2's effort to increase waste-plastic diversion for recycle. The team decided that the LDPE from the Red Sky project would be the material used to pilot recycling of packing foam.

A 30-cubic-yard roll-off container was placed next to the receiving area of the Annex, and the installation team sorted all of the packing materials. Wooden pallet boxes were stacked to one side, loose cardboard went into two recycle dumpsters, the LDPE foam went into the roll-off, scrap metal was neatly stacked in collection containers, and other waste material (such as bags and polyurethane foam) went into the solid waste containers.

■ CLOUD COMPUTING AND SERVER VIRTUALIZATION

Cloud computing and server virtualization are important strategies to reduce data center power consumption. Cloud computing provides IT with the ability to increase capacity or add capabilities as needed without investing in new infrastructure—which results in no added energy requirements. Virtualization enables data centers to consolidate their physical server infrastructure by hosting multiple virtual servers on significantly fewer energy-efficient physical servers.

Both of these strategies reduce the demands on data center infrastructure, and they require fewer physical servers. This results in significant financial savings for data center infra-

structure expansion, electricity, server replacement, and technical support. Virtualization is what makes today's cloud computing technology agile, sustainable, and affordable.

BNL virtualizes with a Shift to Blade Technology

BNL's Information Technology Division (ITD) manages nearly 400 servers that operate and consume electricity around the clock for thousands of hours each year. As part of the BNL sustainability plan, ITD reviewed the hardware configurations used in its data centers to determine whether there were any opportunities to conserve energy.

ITD decided to migrate into blade-server technology, which is a data center technology that typically uses 25% less electricity than traditional stand-alone servers. Individual blades are comparable in cost and can easily handle the same jobs that traditional servers handle. Because they share resources, blades do the same thing using less electricity. For example, as many as 14 blades can sit in the same chassis and share cooling systems, power supplies, and network cards.

ITD realized significant cost savings just by using more energy-efficient blade technology hardware. They took this a step further with the implementation of server virtualization using VMware® software. Virtualization allowed ITD to replace more servers with fewer blades, which allowed a single physical blade server to be divided into a number of virtual servers. This technology uses almost all of the server's processing power, a significant improvement over blades alone.

A considerable number of the servers managed by ITD were due for replacement in 2010 and 2011. Many of these servers supported the main administrative business applications used at BNL, such as Microsoft® Exchange®, Microsoft SharePoint®, Maximo®, and PeopleSoft® systems. The servers also provided support for scientific research. Because of the critical nature of the supported applications, each of the servers and applications were reviewed to determine whether they would be good candidates for virtualization.

The end result was that 164 stand-alone servers, which alone consumed almost 85kW of power, have now been replaced by 37 blade servers (Figure 4).

FNAL Grid Computing Center Virtualizes Global Physics Research

FNAL's Grid Computing Center (GCC) is a 16,000-square-foot data center containing computers, networking, and data storage robotics with related power and cooling infrastructure. The GCC serves the computing needs of the lab's research programs, and supports the collection, archiving, processing, simulation, and analysis of data from global scientific programs.

The GCC provides the ability to conduct research using grid and cloud computing, which is connected by high-speed networking to other FNAL facilities and globally (Figure 5).¹ Grid computing is a form of distributed computing in which multiple clusters of nodes work together



Figure 5. Grid Computing Center interior.
Photo courtesy of FNAL



Figure 4. Server equipment avoided through virtualization. Photo courtesy of BNL

to complete tasks. Physicists submit “jobs”—or computer programs that physicists use to extract physics results from data—to the grid. To maximize resource utilization, the grid determines which resources are free and uses those nodes to process the job.

All of FNAL's grid resources are organized into a single, local architecture called the FermiGrid Gateway. FNAL physicists first submit jobs to the gateway, and they may run on FNAL computers or they may be sent out to computers at other institutions through the Open Science Grid network. Similarly, jobs submitted by researchers at other institutions through Open Science Grid could be routed to run on FNAL computers.

FNAL also works with two other grids, TeraGrid and EGEE, to support physicists' needs. FNAL is also a regular user of LBNL's Energy Sciences Network (ESNet).

The GCC has 255 rack spaces for high-density computers and provides power and cooling for computers that consume between 10 and 15 kW per rack. More than 7,000 computers in the GCC use 1.5 MW of power and this number is grow-

ing. There is more than 1,000 tons of air conditioning available to remove the heat generated by computers, and it takes 0.7 MW to run the air conditioners plus lighting and other support infrastructure.

To help maximize energy efficiency, FNAL incorporated energy-conservation measures during the engineering and construction phases. These measures include managing warm and cold unobstructed air flows, applying high efficiency power components, using room designs influenced by modeling and prototyping tools, and implementing other measures that together contribute to efficient energy use.

FNAL's GCC received the agency's Energy Star award for 2010 for ranking in the Top 25 percentile of data centers registered with the EPA.²

Jefferson Lab combines Green IT efforts to reduce GHGs and energy use

Jefferson lab reduced energy consumption in its core server infrastructure by migrating from physical to virtual machines. As of mid-2011, 200 servers had been converted to run virtually on seven physical computer systems. The average server at Jefferson lab required 308 W to operate (220 W plus 88 W for cooling). Virtualization represented an annual savings of 539,000 kWh of electricity. Concurrently with virtualization, Jefferson lab implemented power management on 750 desktops by putting

¹ FermiLab. “Grid Computing.” *FermiLab* online, <http://www.fnal.gov/pub/science/computing/grid.html>. Accessed May 10, 2011.

² Bellendir, G. “Fermilab data center receives EPA Energy Star award.” *FermiLab Today* online, http://www.fnal.gov/pub/today/archive_2011/today11-02-04.html. Accessed May 10, 2011.

CLOUD COMPUTING AND SERVER VIRTUALIZATION

desktop computers and monitors to sleep when not in use. The implemented power management system places the lab's desktops in sleep mode 63% of the time. The average desktop uses 120 W while awake, and 12 W while in sleep mode. This has resulted in an annual savings of 447,000 kWh of electricity.

The combined effort of server virtualization efforts and desktop power management has reduced the lab's GHGs by 1,344,000 pounds a year and has saved enough energy to power 90 average U.S. homes each year (Figure 6).

LANL Developed Infrastructure on Demand to Support Green IT and Enhance Customer Experience

LANL pioneered an infrastructure-as-a-service (IaaS) cloud computing environment that saves energy and costs, while actually improving the customer experience. Called infrastructure on demand (IoD), this virtualized computing portal keeps LANL's data centers running efficiently while saving the taxpayers \$1.4 million a year. IoD also results in the reduction in carbon footprint inherent in virtualization.

So what is IoD and how does it support Green IT? IoD is a virtualized server environment that revolutionizes the way service requesters ask for and receive computing time. In the past, LANL researchers and employees requested a physical server for their specific computing needs. It could then take up to 30 days to get this server, and once installed, the cost in dollars and energy was high. With IoD, requesters log in to a self-service Web portal and request

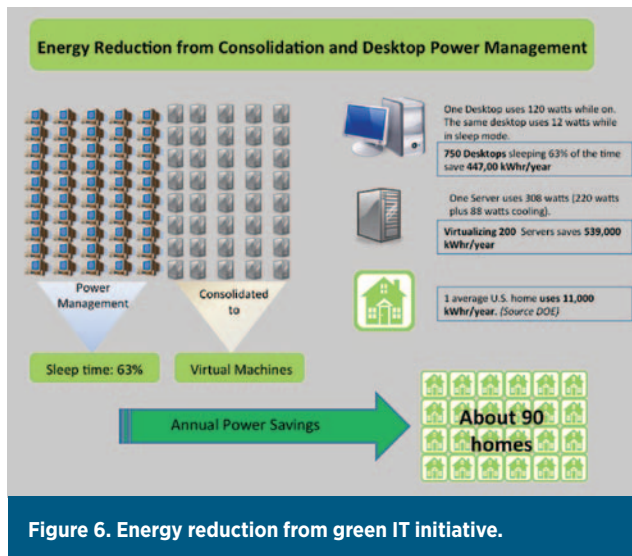


Figure 6. Energy reduction from green IT initiative.

energy savings estimated at more than 1.8 MWh/yr. This translates to a reduction of 50%-70% of physical data center footprint without any loss in computing capacity or security.

LANL realized a full return on investment in IoD within nine months from the project's inception (Figure 7), which was a full year ahead of expectations.

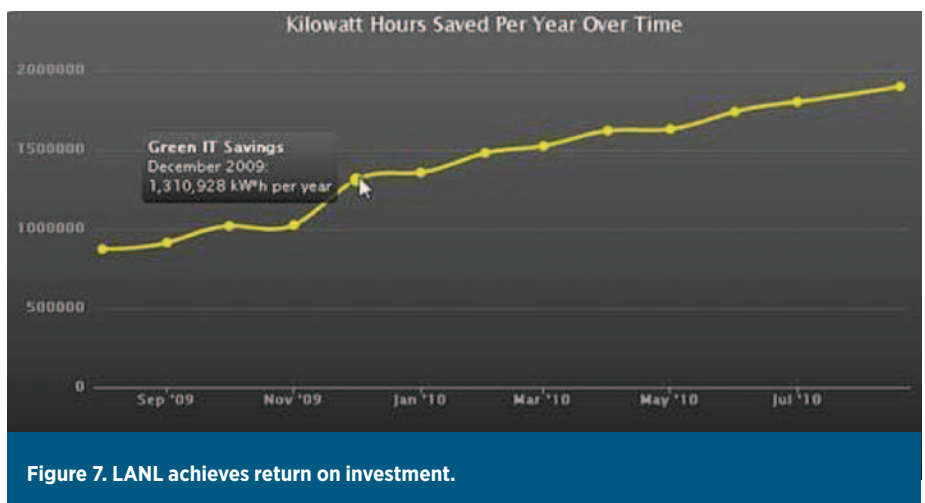


Figure 7. LANL achieves return on investment.

IoD supports Green IT real-time metrics by offering a dashboard that computes the amount of energy saved per year by using virtual servers, rather than physical systems. System owners can track energy savings over time as related to the number of homes that can be powered annually by the energy saved. The annual energy savings enabled by using IoD provides enough electricity to power 216 houses!

LANL's remarkable success in saving energy, costs, and enabling great science is just one of many ways in which the labs are supporting how DOE is finding innovative solutions to some of America's biggest energy efficiency challenges.

The IoD portal provides many advantages over a physical server, even beyond the cost and energy savings.

- Users can quickly request the computing resources they need, and know they have made the correct fiscal and computing choices, directly through the IoD portal.
- A dynamic cost calculator shows the requester the cost of his/her environment up front, dynamically identifying the exact amount of computing needed before purchasing the virtual environment.
- Once requested, the IoD tool handles the workflow to gain the proper approvals for the requested environment.
- Once approved, the consumer is charged back automatically through the system, eliminating the need for paper-based accounting.
- The tool provides a "Single Pane of Glass," allowing the customer to view and manage all of the virtual machines in their cloud.
- IoD supports full life cycle management by proactively notifying virtual machine owners when it is time to renew their environment. IoD archives obsolete servers, freeing up space for a newly provisioned environment without the need for waste disposal.

infrastructure for scientific computing. When sourced from large vendors, these products can provide substantial environmental benefits. For example, LBNL migrated to Google Apps for email, calendar, and collaboration tools in 2010. Google's data centers exceed the EPA's most optimistic scenarios for "state-of-the-art" data centers, with a power usage effectiveness (PUE) substantially below 1.2. For the most recent quarter, Google's average weighted PUE for its data centers was 1.13, with some data centers scoring as low as 1.09. The lab's use of Google Apps not only reduces its physical data center footprint, but also reduces the total energy footprint.

LBNL has implemented additional cloud computing vendor services to further reduce the energy footprint made by IT resources. The lab uses a wide range of software-as-a-service (SaaS) and infrastructure-as-a-service (IaaS) offerings, which permits reduced impact on its data centers and increased capabilities for the research staff.

The next major opportunity for green computing at LBNL is HPC. While HPC in the cloud is still relatively expensive compared to on-premises services, some uses of HPC are immediately attractive. For example, occasional-use, highly parallel jobs can be run in the cloud, avoiding the need to purchase additional hardware for occasional use. The lab expects commercial cloud HPC services to continue to become more competitive with on-premises HPC.

In many ways, DOE lab IT centers already have many of the best characteristics of cloud computing. They support very large resources optimized for running scientific jobs from a wide array of scientific customers, which minimizes the costs to manage the systems and idle time. To further extend this model to mid-range computing, the DOE Office of Science funded the Magellan Project housed at LBNL's National Energy Scientific Research Computing Center (NERSC), a joint partnership with LBNL and Argonne lab.¹

LBNL Early Adopter of Green Clouds and Gmail

LBNL was an early adopter of a variety of cloud technologies, from productivity and collaboration tools to

¹ Simon, H.; Yelick, K.; Broughton, J.; Draney, B.; Bashor, J.; Paul, D.; Vu, L.; Beckman, P.; Coghlan, S.; Taylor, E. "Exploring CLOUD Computing for DOE's Scientific Mission." *SciDAC Review* online, <http://www.scidacreview.org/1002/html/hardware.html>. Accessed May 5, 2011

CLOUD COMPUTING AND SERVER VIRTUALIZATION

The funding allowed the labs to evaluate the suitability of various cloud approaches to scientific workloads (Figure 8). Magellan was awarded the HPCwire Readers' Choice Award 2010 for "best use of HPC in the cloud."

LLNL Adopts "Virtualization First"

Gartner deemed virtualization as a key Fortune 100 best practice. LLNL commissioned and expanded its virtualization service, offering the business side virtual servers in lieu of power-draining, resource-hogging physical servers.

LLNL commissioned its virtualization cluster in spring 2010 using the Cisco Unified Computing System (UCS) platform and the VMWare virtualization software suite (Figure 9). The virtualization cluster hosts nearly 260 servers virtually, which achieves dramatic energy savings with an investment payback period of less than one year. Savings are realized through reductions in energy consumption, the facilities footprint, and system/network administration support. Further savings can be achieved through longer system life cycles and more uniform system management practices.



Figure 9. Cisco UCS. Photo courtesy of Cisco Systems Inc.

LLNL adopted a "virtualization first" policy, which requires that any system under consideration for consolidation or collocation into the enterprise data center is examined as a virtualization candidate. A programmatic virtual machine offering, or infrastructure-on-demand, is planned as a service expansion in FY12 with the expansion of the virtualization cluster to support up to 800 servers in progress.



Figure 8. The Magellan management and network control racks at NERSC. Photo courtesy of LBNL

NREL Server Virtualization

NREL's Research Support Facility (RSF) was designed to use half of the energy specified in the Denver building code. To support these aggressive goals to reduce energy consumption in the RSF, it was imperative for NREL's IT organization to replace high energy-consuming legacy IT infrastructure with energy-efficient equipment. The team needed to architect IT infrastructure to meet current requirements, while ensuring the ability to appropriately scale capabilities with changes in demand. The data center also had to extensively leverage server virtualization technologies.

NREL replaced 90% of its legacy server environment with HP Blade servers that use variable speed fans and energy-efficient power supplies. Additionally, a virtualized server environment was introduced to decrease the required number of physical servers.


Nearly 85% of NREL's server environment is "virtualized." The original goal was to reach a 20:1 ratio for server virtualization, which would convert the workload that used to run on 20 physical servers to only one single-blade server. However, in some environments, NREL has experienced as much as a 29:1 ratio. Dell Equallogic storage area networks (SANs) are used to pool storage resources, which can reduce the amount of hardware that would typically be required for storage dedicated to server resources.

Figure 10 illustrates the server virtualization effects of running a workload that used to require 20 servers on one blade server. In this example, the energy footprint is reduced by more than 96% for each server.

NREL is planning to implement 500 thin clients in its expansion to the RSF, which will significantly reduce the energy footprint of employee computers and provision computing power where and as needed. Additional individual computers will be replaced as employee computing habits are monitored.

PNNL institutes Cash for Clunkers server recycling program, replaced with virtualization

Today, there are servers of many types (rogue servers) operating within normal office spaces. The extra heat load requires the building operator to compensate by over-cooling a zone that includes many normally occupied offices, which often leads occupants to resort to space heaters. Additionally, building management traditionally



used to relax their cooling requirements in the evenings and weekends (called setbacks) to save energy and wear and tear on equipment. PNNL has implemented an amnesty program, called Cash for Clunkers, where rogue servers are traded in for leased virtual server space.

There are many benefits to this trade-in program. Virtual servers are more secure and are managed by professional system administrators, which frees the customer from performing mundane update tasks. These systems are much more energy efficient, and have built-in high availability via migration. Building setbacks can be reinstated to save energy and equipment costs. Finally, energy-consuming space heaters can be eliminated by eliminating over-cooling.

A campus-wide advertising campaign ensured that all facilities knew about and could take advantage of the program. The virtual machines were offered at no cost to the user to encourage participation, and setbacks were reinstated to further motivate users to embrace this energy-saving program.

Server Virtualization and Energy Conservation at PNNL’s ISB2/1 Business Data Center

PNNL’s ISB2/1 Business Data Center customer use required the lab to triple its available business server applications and increase their onsite disk storage by 200% (**Figure 11**). The data center needed to meet this large increase in capacity while decreasing occupied space and maintaining a reasonably flat rate of power consumption.

The implementation of a converged network (Ethernet and disk data) onto a mirrored 10Gb backbone was key in reducing network complexity and increasing bandwidth, which allowed implementation of more capable virtual host hardware servers.




Figure 11. PNNL ISB2/1 data center power capacity

ENERGY EFFICIENT DATA CENTERS

PNNL is about 75% virtualized, which results in three out of every four business applications operating in the virtual environment.

Table 4 compares a traditional stand-alone server solution to a virtualized solution using the server count and a PUE of 1.7. Virtualization represents a substantial cost savings for both server hardware and electricity.

Table 4. PNNL comparison of hardware and electricity costs between standalone and virtualized servers

	Standalone	Virtualized	Savings
Hardware Cost	\$2,848K	\$550K	\$2,298K
Electricity Cost	\$180K/yr	\$8K/yr	\$172K/yr

PPPL's Cloud Computing and Thin-client Initiatives

PPPL is pilot testing both cloud-based computing applications (e.g., Google Apps) and thin-client desktop computing solutions. The objective of these pilot tests is to prove the suitability and performance of state-of-the-art IT applications in preparation for large-scale adoption. Such applications will significantly reduce the overall environmental footprint of PPPL's IT operations through improved energy efficiency, reduce materials use, and reduced waste generation, while providing better user experiences (**Figure 12**).

Additionally, PPPL's Environment Services Division has used the cloud-based environmental management information system Enviance® to track, manage, and document environmental compliance activities for more than nine years.

ENERGY-EFFICIENT DATA CENTERS

Data centers are the operational heart of Green IT systems. The U.S. Environmental Protection Agency Energy Star Program estimates that servers and data centers alone use approximately 100 billion kWh of electricity, which repre-

sents an annual cost of about \$7.4 billion.¹ They also estimate that without the implementation of sustainability measures in data centers, the United States may need to add 10 additional power plants to the grid just to keep up with the energy demands of these facilities. Given these trends, the federal government electricity costs could reach \$740 million in 2011, with a peak energy load of about 1.2 GW.

Gartner estimates that for every watt of power used to run data center equipment, more than 2 watts is required just for power management and cooling systems. These staggering statistics illustrate the need for federal data centers to become more sustainable. There are significant opportunities to reduce power consumption by implementing simple best practices.

DOE labs have done so; and, in many cases, used innovative, crosscutting techniques to further reduce the energy footprint in data centers and be good stewards of our energy resources.



Figure 12. PPPL's computing environment. Photo courtesy of PPPL.

¹ U.S. Environmental Protection Agency Energy Star Program. "Report to Congress on Server and Data Center Energy Efficiency, Public Law 109-431." *Energy Star* online, http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf. Accessed May 26, 2011.

This section highlights the industry-leading work that DOE laboratories are accomplishing in data center sustainability.

DOE LABS REDEFINE WORLD-CLASS PUE

PUE is an industry standard metric used to evaluate data center performance. It is defined as the ratio of the total power required to run a data center, divided by the power used by all of the IT equipment in the data center. A PUE of between 2.0 and 3.0 is considered standard, and Gartner has identified a PUE of less than 1.3 as world class. Several DOE labs are approaching or have exceeded a world-class designation in PUE, as shown in **Table 5**.

Table 5. DOE Lab World Class PUEs

Lab	PUE
NREL	-1.12
LBNL (via Google)	-1.13
NERSC (ANL & LBNL HPC)	-1.15
INL	-1.2
SLAC (SRCF target)	-1.3
PPPL	-1.308
LLNL (TSF)	-1.34
PNNL	-1.5

DATA CENTER EFFICIENCIES

Data center practices, technologies, and innovations have opened the door for a number of sustainable implementations

in the data center environment. From simple modifications to how air is handled, to inventing new ways to track energy use, the DOE labs are leading the way in green data center efficiencies. This section highlights practices and innovations in the efficiency arena.

Ames Cluster Energy Usage Monitoring

Ames lab is located on the Iowa State University campus. Because Ames buys its electricity from the city rather than its host university, the Division of Materials Sciences and Engineering (DMSE) subscribed to email notifications from City of Ames Electric Power Plant to learn when city power use increased with heat-related uses during the summer of 2010. The goal was to determine if DMSE could possibly help reduce energy use, CO₂ generated, and IT-related costs, especially for electricity bought during peak hours at a premium rate.

DMSE enlisted the help of undergraduate interns to monitor cluster computer use. The goal was to determine how IT infrastructure consumes less power by shutting down unused nodes and turning them on as needed, or after peak alert hours. The interns wrote computer scripts to monitor power use.

Following the trial period, DMSE showed that it could significantly reduce costs, kW hours used, and CO₂ released just by using nodes only when needed (**Figure 13**). The lab calculated savings of 17,700 kWh and roughly 7 metric tons of CO₂ on just one of its server rack clusters.

The program has since been expanded and incorporated into the data center’s power management protocols for all clusters.

Greening and Monitoring LBNL’s Data Centers

LBNL’s information technology (IT) Division has been working to improve the energy efficiency of its primary data center since 2007. In collaboration with researchers in the Environmental Energy and Technologies Division, IT developed a multiyear strategy to reduce waste. To date, IT has:

- Converted the overhead cold air supply plenum to a hot air return
- Added ducts connecting data center air conditioners to the new return plenum
- Eliminated unnecessary and wasteful humidification and dehumidification systems




Figure 13. Ames energy monitoring results. Illustration courtesy of Ames

DATA CENTER EFFICIENCIES

- Installed plumbing and hardware to bring water cooling to individual computer racks
- Virtualized servers with the aim of reducing space and electrical requirements
- Installed a wireless environmental monitoring system to measure temperature, humidity, current, and under-floor air pressure.

The wireless monitoring system allows data center operators to get real-time visual feedback in response to physical changes in the room (Figure 14).




Figure 14. LBNL wireless environmental monitoring system.

Illustration courtesy of LBNL

Sometimes, simple modifications in the data center produce dramatic results. For example, Figure 15 shows the effect of turning off humidification control on a single air-conditioning unit, and Figure 16 shows the effect of adding a “blanking panel” (a piece of sheet metal designed to block air flow) to a rack of high-performance servers.

As these figures show, investing five minutes for installation of a \$10 component can reduce air temperature at the top of a rack of computers by 20°F or more.

With the information provided by the monitoring system, operators have removed and relocated dozens of air-permeable floor tiles to eliminate hot spots, reduce over-cooling, and increase under-floor air pressure. These efficiency efforts allowed IT to deploy a new 100kW computing cluster in 2008 without purchasing additional air-conditioning equipment.

LLNL Enterprise Data Center

Sustainability in the LLNL Enterprise Data Center (EDC) is at the forefront of the lab’s strategic roadmap and a key feature in all service offerings. The EDC achieves this through a variety of initiatives, including server consolidation, energy monitoring, virtualization services, service standardization, multi-modular data center design, increased EDC temperature, and additional cloud computing. In FY12, the lab plans to incorporate free cooling when the outside ambient air temperature is cool enough (estimated to be 68% of the year in Livermore), and hot aisle/cold aisle containment.

Traditionally, data centers have maintained a cool 68°F temperature. Studies have shown that this amount of cooling is excessive. LLNL’s EDC increased its temperature to 74°F with no loss in server performance, and is studying ways to increase the temperature to 75°F in FY12.




Figure 15. Effect of humidification control set to “off.”



Figure 16. Effect of adding blanking panel to rack of high-performance servers.

Facility consolidation across the LLNL landscape continues as programs shut down server rooms and computer closets and receive solid, high-availability, reliable tenant hosting in the data center. Facility costs have dropped substantially as more than 430 servers from 30 facilities have been migrated to the EDC.

The EDC has been reconfigured for multi-modular design, rather than a one-size-must-fit-all approach. Programs now have the capability to select an availability

zone that meets their server needs. For example, a development environment will likely just need single-sourced power—anything above that (e.g., dual power, high availability) would waste efficiency. The EDC has a variety of configurations available in its tenant housing service offering, driving efficiencies by matching available power with program requirements.

Power metering and monitoring facilitate measurement of the EDC’s power utilization efficiency, and the data center PUE rating of 2.5 is right at industry standards.

The laboratory plans to request funding to implement free cooling into the EDC in FY12. A major energy consumer in the data center is powering the chill water plant, which requires a cooling system that provides 74°F air to the computing floor.

However, in Livermore, the outside air temperature is less than 74°F for 68% of the year. This project will implement a bypass of the chill water system and air-conditioning plant, and will bring filtered outside air directly to the EDC floor. The lab also plans to venture into cloud computing, beginning with the storage of existing public-facing (green network) content into the cloud through a storage-hosting service such as Amazon S3. The EDC expects to realize direct savings by decommissioning storage infrastructure on-site. Finally, LLNL plans to implement hot- and cold- aisle containment in the EDC, keeping cool air segregated from hot exhaust air and, in turn, making cooling units substantially more efficient.

Sandia Labs’ High Performance Computing Containment Curtain

In anticipation of a future high-density, indirect-cooled computer cluster, Sandia labs is testing the advantages of a cold-aisle containment system. A containment curtain was installed around the Phase I cluster in Sandia labs’ data center to see what energy efficiencies would result (Figure 17). The Phase I system is colocated in a shared environment that requires nine computer room air conditioners (CRAC) operating at 100% utilization and an under-floor temperature of 54°F. With the installation of the containment curtain, Sandia labs was able to maintain the same environmental parameters with only six CRAC units running at 80% and also has been able to raise the under-floor temperature to 58°F.

Sandia labs tested the air flow, ambient temperature, and under-floor temperature for optimal energy-efficient operations in its specific computer environment. Testing was conducted in a live environment and will be replicable in the Phase II design and layout for the most energy-efficient operations.

Table 6 compares before and after operation of the cold aisle containment system.

Table 6. Sandia cold aisle containment system before and after comparison

Before*		After*	
9 CRACs @ 100% for 678,024 kWh/year	\$61,022*	6 CRACs @ 80% for 346,896 kWh/year	\$31,220*
Annual operating cost savings of \$29,802			
Cost of containment curtain \$5,500**			
Return on Investment			
29,802 annual savings = \$2,483.50 monthly savings		\$5,500 / \$2,483 = 2.2 months	
*Based current kWh rate of \$0.09 **Time and Material			
This project’s Return on Investment (ROI) was two and a half months.			



Figure 17. Sandia Labs’ cold aisle containment curtain. Photo courtesy of Sandia National Laboratories

DATA CENTER EFFICIENCIES

This project's return on investment (ROI) was two-and-a-half months. The key benefits of this project were:

- Testing the predicted energy efficiency rewards in a live environment
- Achieving ROI rapidly
- Applying results to other computer environments
- Improving energy savings, costs, and efficiencies.

Savannah River Increases Data Center Operational and Energy Efficiency with Simple, Low-Cost Measures

Savannah River strived to meet DOE goals for Green IT at a minimal cost. With a partnership between Savannah River, Berkeley, and DOE, Savannah River brought its PUE down from 4.00 to 2.77. This more than 30% reduction occurred by implementing simple, low-cost modifications to the lab's existing stand-alone central computing facility (CCF) over the space of three days using Retro-Commissioning (Retro-Cx).¹

Savannah River split the Retro-Cx process into three phases. The first day was dedicated to the initial project setup and collecting baseline metrics. The selected improvements—such as turning some air handlers off and rearranging floor tiles—were implemented on the second day. The third

day was spent taking final readings and analyzing the changes in energy consumption (Figure 18).

These fast, inexpensive changes dropped the data center's initial PUE of 4.00 to 2.77%, which increased the amount of electricity powering the IT server racks from 25% to 36%. The improvements to air handling and the modifications to floor tile arrangement combined to drop the total system cooling load from 402.9 kWh to 229.8 kWh. This equates to a decrease of 173.1 kWh, or almost 43%. These changes equated to a total estimated energy savings of approximately 1.44 million kWh/year for a total estimated cost of \$25,000. The cost was recouped within two-and-a-half months of operation.

COOLED WITH "FREE" COOLING

Data center facilities require a significant amount of energy to cool the equipment they house. Many DOE labs are located in climates that are conducive to sharing the work of cooling data centers with the region's weather patterns. This section details what the DOE labs are doing to partner with nature for free cooling.

Initial Power Consumption		Final Power Consumption	
Total CCF CRAH Units	133.2 kW	Total CCF CRAH Units	58.40 kW
All Others (Chillers, Pumps etc.)	309.7 kW	All Others (Chillers, Pumps, etc.)	211.4 kW
Telecom deduction*	(-40 kW)	Telecom deduction*	(-40 kW)
Total Cooling System Load	402.9 kW	Total Cooling System Load	229.8 kW
Total UPS-1 Input - for IT Load	96.3 kW	Total UPS-1 Input - for IT Load	95.2 kW
Total UPS-2 Input - for IT Load	77.3 kW	Total UPS-2 Input - for IT Load	76.7 kW
Total UPS Load - Gross Input	173.6 kW	Total UPS Load - Gross Input	171.9 kW
TOTAL LOAD	576.5 kW	TOTAL LOAD	401.7 kW
Total UPS-1 Output - net IT Load	77 kW	Total UPS-1 Output - net IT Load	78 kW
Total UPS-2 Output - net IT Load	67kW	Total UPS-2 Output - net IT Load	67 kW
Total IT Load - Net UPS Output	144 kW	Total IT Load - Net UPS Output	145 kW
Initial PUE	4.00	Final PUE	2.77

Figure 18. Savannah River power consumption results

¹ Bell, G. C. "Retro-commissioning increases data center efficiency at low-cost: Success at Savannah River Site (SRS) at low-cost. Success at Savannah River Site (SRS) at low-Cost." *Federal Energy Management Program* online, http://www1.eere.energy.gov/femp/pdfs/datacenter_savannah.pdf. Accessed May 4, 2011.

Argonne Lab's Free Cooling for Super Hot High Performance Computer

The Argonne Leadership Computing Facility (ALCF) supercomputer, Blue Gene/P, is a highly energy-efficient computer system (Figure 19). However, a machine of this type requires a significant amount of cooling, which translates to high energy costs. In fact, if left uncooled, the room that houses the Blue Gene/P would



Figure 19. ALCF's Blue Gene/P supercomputer. Photo courtesy of Argonne National Laboratories

heat to 100°F in just 10 minutes.¹ By partnering with the climate, an ultra energy-efficient machine is cost-effective when teamed with an ultra energy-efficient facility.

The ALCF design includes a chilled water system that uses cooling towers. When the Chicago temperature drops to 35°F, the chilled water system temperature is maintained solely by the cooling towers. Partnering with the

weather saves Argonne lab nearly \$25,000 a month in cooling costs.

The ALCF's innovative, energy-saving approach to cooling was awarded the DOE EStar award for green computing in 2010.

INL's HPC Relies on Free Cooling to Support Energy Efficiency

In 2007, INL completed the construction of a 3,700 square-foot data center to support its HPC resources, and also ensured it could be expanded to 10,000 square feet in the near future. Several practices were incorporated to assist with energy efficiency goals for the lab.

The data center space was right-sized to minimize the associated operating energy costs.

Cooling for the data center uses a green solution called "free cooling." As long as the outside temperature ranges between 40°F and -31°F, and the temperature of the water leaving the data center is not too high, the chillers do not operate. This means that the water used

to cool the heat transfer plate and subsequently to extract the heat from the data center water can be cooled solely through the evaporative process of the cooling towers. This condition of "free cooling" saved a significant amount of electrical costs by not having to run the expensive chillers.

INL's large compute clusters include water-cooled doors to improve the overall cooling efficiency of the data center. The exhausted warm air from the compute nodes is immediately cooled as it passes through the rear cooling doors on the racks and reenters the room at temperatures near those of the open air in the data center. The CRAC units are cooled by the chilled water and are required to do less work, which reduces electricity consumption.

Finally, the last HPC cluster procured was designed using the latest technologies in high-density processors from AMD®. Having four 8-core processors in each node (versus traditional configurations) greatly reduced the requirements for space, power, and cooling of the entire system. The configuration includes fewer racks, fewer nodes, less network infrastructure, and fewer power supplies, all resulting in less consumed power.

Table 7. INL HPC data center PUE

Instantaneous Power (5/19/11)	Consumption (KW)	3-Year Power (2/2008-5/2011)	Consumption (KW)
Compute	420	Compute	7018
Cooling	140	Cooling	2750
PUE	1.3	PUE	1.4

As a result of these efforts, the current PUE rating for the INL HPC data center ranges from 1.1 to 1.3, depending on system load and outside weather conditions (Table 7).

¹ West, J. "Baby, It's Cold Outside...Let's Calculate Something." *HPCwire* Online, http://www.hpcwire.com/hpcwire/2008-12-18/baby_its_cold_outside_lets_calculate_something.html. Accessed May 9, 2011.

COOLED WITH "FREE" COOLING




Figure 20. RSF data center cooling and reuse of heat from hot aisle containment system. *Illustration courtesy of NREL*

NREL Partners with Climate to Cool Data Center

NREL is located in a climate that is favorable for “free cooling,” which pulls outside air through an innovative air intake to provide the majority of the data center’s cooling needs. The cooling system was designed to use direct air pulled from the external environment and evaporative cooling

methods for all but the hottest and most humid days of the year; free cooling is unavailable for only 10 days a year.

By using passive methods, data center cooling temperatures range from 65° F–80° F, with humidity ranging from 20%–60%. **Figure 20** shows the cool outside air intake system and its interaction with the data center’s hot and cold containment aisles. **Figure 21** shows NREL’s air intake, integrated into the employee lunchroom’s courtyard in the RSF.

PNNL Data Center Uses Aquifer for Year Round Data Center Cooling

PNNL has access to a natural aquifer that provides free cooling to the data center in its Computational Sciences Facility (CSF) and Biological Sciences Facility (BSF). The aquifer is used as a geothermal heat sink for these buildings’ chilled water facilities, supporting the sustainable design feature. Chilled water is supplied to the CSF/BSF data center via efficient centrifugal pumps, and is capable of displacing 4.5 MW of heat.

The CSF/BSF geothermal system is capable of providing 65° F water year-round to rear-door heat exchangers in the data center without the use of chillers (**Figure 22**). At a flow rate



Figure 21. Data center air intake.
Photo by Dennis Schroeder, NREL/PIX 19220




Figure 22. PNNL aquifer cooling system. *Illustration courtesy of Graham Whitmore, Motivair*

of 13 gallons per minute each, they can dissipate 20kW of heat generated by rack-mounted servers, and more (if needed) at higher flow rates.

PPPL's Data Center Energy Improvements

Installation of a hot aisle containment enclosure and enhanced air conditioning enabled PPPL to raise the data center's ambient temperature from 67°F to 74°F, while maintaining appropriate cooling for critical computer equipment (Figure 23). These combined actions provided significant immediate energy savings by reducing cooling load on the existing chillers. During July 2010, the average day-



Figure 23. PPPL's hot aisle containment enclosure.
Photo courtesy of PPPL

time high temperature was 92°F. An evaluation of chiller loads for this period revealed a significant reduction in cooling from the 65 tons noted for April 2010 to only 37.5 tons in July 2010. This represents an energy savings of approximately 42%.

Adding an electrical substation, emergency generator, and sub-metering will allow a better understanding of energy use (IT and non-IT energy) in the data center, which currently operates with an average PUE of 1.308.

POWERED BY RENEWABLE ENERGY

Many DOE lab data centers are within facilities that use renewable energy resources to power their systems. This section highlights the practices and innovations that the labs have implemented to save energy, decrease costs and reduce their carbon footprint by using sustainable resources to power their data centers.

NREL Takes Energy from Net Zero Energy Building Design

NREL's data center delivers its administrative IT products and services from its Leadership in Energy and Environmental Design (LEED®) Platinum Research Support Facility (RSF), one of the most energy-efficient office buildings in the world. To help meet the building's energy goals, the RSF and

the data center it houses are powered by solar panels installed on the roof (Figure 24). The data center contains approximately 100 kW of power, all of which is powered by the sun.

The center was designed from the ground up to minimize its energy footprint without compromising

needed service quality. After a year in operation, NREL's data center operates at between 1.1–1.15 PUE, which is significantly lower than the legacy data center that operated at a 3.3 PUE. The data center supports efficient use of its solar power by enforcing strict power management settings for all devices connected to the network, and by using only EPEAT-registered and Energy Star-compliant devices. The data center itself requires 81% less energy to operate than the legacy facility NREL operated before moving to the RSF. This helps support the building's net-zero goal.

Sandia Labs' Distributed Energy Technologies Laboratory Partners to Power Data Center with Solar

Sandia Labs' Distributed Energy Technologies Laboratory (DETL) studied ways to reduce energy consumption in its data centers. The lab determined that an integral component of the plan would be the capability to offset consumption with energy production



Figure 24. NREL's roof-installed photovoltaic panels. Photo by Dennis Schroeder, NREL/PIX 18224

through a renewable energy source. Buying photovoltaic panels was one step toward achieving this goal, but getting the panels up and generating

POWERED BY RENEWABLE ENERGY

power was thwarted by the lack of a suitable installation site. Sandia Labs' potential sites either did not have any of the required infrastructures or were pending review for environmental impact.

The solution was born through an unusual alliance on the Sandia Labs' campus. The Computing and Network Services Center (CNSC) owned photovoltaic arrays, and DETL could provide a site to install them (Figure 25). The DETL site would use these panels to support programmatic research projects, while CNSC would reap the renewable energy credits (RECs) generated by their arrays and all other existing arrays on the site.

This partnership presented an opportunity to take advantage of readily available infrastructure, minimize environmental impact for the installation, and ensure that the project was aligned with the objectives of the Sandia Strategic Energy Water and Transportation Resource Management Plan. CNSC manages the corporate computing facilities and is using the energy credits from the DETL site to offset its energy consumption.

This project links to two of the federal government's renewable energy goals. The first is the DOE's goal of having 7.5% of its electricity consumption come from renewable energy (RE) sources by FY 2013. The second is the Energy Policy Act (EPA) 2005 statute that provides a double bonus if the RE is produced on-site and if the RECs are retained.

This collaboration represents two organizations with different missions coming together on a project that allows both to achieve fundamental mission goals. Each organization is able to demonstrate a commitment to environmental stewardship and as-



Figure 25. Sandia Labs' DETL Existing Array w 9300 SW Array. Photo courtesy of Sandia National Laboratories

simulate renewable energy production with a direct impact on Sandia Labs' day-to-day operations. Additionally, CNSC was able to maximize use of an existing site (no addition to Sandia Labs footprint) and infrastructure (inverters, combiner boxes, and conduit runs) to install the solar arrays.

■ WASTE HEAT REUSE

In traditional data center operations, managers used air-conditioned cooling to cancel out the heat produced by equipment. Innovations in air-handling technology allow the heat produced in a data center to be used as a resource to heat facilities outside of the data center.

The following are examples of how the DOE labs minimize their energy footprint to cool a data center, while using the waste heat productively to heat adjacent facilities.

NREL Reuses Data Center Waste to Heat 800-Employee Office Building

The equipment racks in NREL's data center are arranged in a hot aisle configuration with hot aisle containment. Air from the hot aisle is extracted from

the data center for reuse in the RSF building. Figure 26 shows the configuration of the hot aisle containment system in the NREL data center. Air ducts pull the heated air into the RSF's basement labyrinth thermal storage area, where it is distributed to workspaces through a heating and cooling system that runs beneath the floors.

Efficient Building Design



Figure 26. Hot aisle containment system used in NREL's RSF data center. Photo by Dennis Schroeder, NREL/PIX 18789

Enables Waste Heat Reuse and Energy Savings at PNNL

PNNL opened its LEED Gold Computational Sciences Facility (CSF) and Biological Sciences Facility (BSF) in early 2010 (Figure 27) with Green IT in mind. These office and lab spaces house 300 staff and the lab's new 10,000-square-foot data center in the CSF.



Figure 27. PNNL CSF/BSF building entry area.
Photo courtesy of PNNL

The design of the CSF/BSF targeted a significant reduction in energy use by capturing the waste heat from data center operations and using it to heat the BSF building and the CSF/BSF entry area in the colder months. The lab also implemented a strategy to recover heat using heat pipes with a summertime evaporative spray cooling enhancement, which transferred energy between the BSF outside and exhaust air streams. The design decision to install a central system was made because the BSF requires 100% outside air for its research, and the CSF houses a computational research laboratory that requires 24-hour-a-day cooling.

The inclusion of the data center into the sustainable design of the CSF/BSF LEED gold building supported the exceptional accomplishment of an energy savings of greater than 30% from American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) 90.1 in the first year of operation. The data center's design earned PNNL a DOE EStar award in 2010.

GREEN IT LEADERSHIP PROVIDED BY LABS

The DOE labs provide global leadership in Green IT not only through individual innovations, but also through partnerships with industry and each other. The labs have worked with technical innovators to develop technologies that significantly reduce a data center's carbon footprint and energy consumption. They have also partnered with businesses to help them improve their data centers' efficiencies and developed tools available to the world to identify energy efficiency opportunities. This section highlights the DOE labs' leadership in Green IT.

Argonne and LBNL's NERSC Partnering with Industry on Center's Rear-Door Cooling

The National Energy Research Scientific Computing Center (NERSC), a joint partnership between Argonne lab and LBNL, is a high-performance scientific-computing laboratory focused on increasing the world's ability to perform science through computer modeling. NERSC needed to replace two super-computing clusters without immedi-

ately decommissioning the old clusters, which resulted in a significant power and cooling strain on the center.

To achieve power and energy goals without compromising computation quality, NERSC selected IBM® System x iDataPlex dx360 M3 servers and iDataPlex racks with an innovative water-cooling solution from IBM and Vette Corp®.¹ To tackle the huge cooling needs for the center's server racks, the center installed the IBM Rear Door Heat eXchanger (Figure 28).



Figure 28. IBM Rear Door Heat eXchanger.
Photo Image reprinted from IBM Rear Door Heat eXchanger Fact Sheet².

NERSC partnered with Vette Corp, the licensor and supplier of the doors, to customize a multi-pass coil that first runs cold water through the outside of the door as hot air leaves the racks, then passes it back through the inside of the door. The water returning to

¹ Mellanox® Technologies. "NERSC Creates an Ultra-Efficient Supercomputer." *Mellanox Technologies* online, "http://www.voltaire.com/ResourceCenter/customer_success/nersc_creates_an_ultra-efficient_supercomputer. Accessed May 11, 2011.

GREEN IT LEADERSHIP PROVIDED BY LABS

the chillers is actually warmer than the air exiting the door, which helps the chillers run more efficiently.

The doors are extremely large, which gives the air time to exchange heat with the water running through them. The slow-moving water allows the fans to move more slowly as well and reduces energy wasted as noise. Additionally, the server's half-depth form factor reduces the air flow that is required across the rack components, which further lessens the power needed to cool. Power management and more efficient rack supplies also contribute to minimizing power requirements.

Another innovation in NERSC air flow management is that the center does not have the alternating hot aisle/cold aisle configuration that is a standard best practice for green data centers. Server racks are configured to pull exhaust air from one row into the next row as input air. This means that air exhausted from one row, traditionally very hot, is actually at a lower temperature than the 75 degrees required for input to the next iDataPlex rack. This solution allows the center to add racks without adding additional air handlers.

NERSC's PUE rating was previously 1.35, meaning that for every watt consumed by the computing resources, a further 0.35 watts are required for cooling. Following installation of the iDataPlex cluster with Rear Door Heat eXchanger, the center's PUE dropped to 1.15.

INL Implements Sustainability Community of Practice for Innovation

INL has implemented a sustainability community of practice (CoP) to lead the innovation and implementation of sustainable science by conducting business in a manner that benefits customers, employees, and the communities in which the lab operates. The CoP develops and implements sustainable business practices, promotes sustainable behavior, and provides for integration and leadership opportunities for members. The CoP also supports institutionalization of sustainability as a way of doing business at INL.

2 IBM®. IBM Rear Door Heat eXchanger. *IBM System X* online. <http://public.dhe.ibm.com/common/ssi/ecm/en/xsd03029usen/XSD03029USEN.PDF>. Accessed June 13, 2011.

The sustainability CoP is tasked with meeting several objectives. These include cultivating an awareness of and communicating sustainability principles and practices across the lab, defining program implementation strategy, generating a positive stewardship philosophy to build sustainability into the lab's culture, providing a framework for issue resolution and program improvements, formulating and instituting sustainability metrics, and rewarding innovation in sustainability.

LBLN Provides Green Wide Area Networks through ESnet

The Energy Sciences Network (ESnet) at LBNL is a high-speed network serving thousands of DOE scientists and collaborators, and is a worldwide leader in high-performance network services (Figure 29). With \$62 million in American




Figure 29. A high-level view of the Energy Sciences Network.

Photo courtesy of NERSC

Recovery and Reinvestment Act (ARRA) funds, ESnet is building a new network infrastructure capable of moving science data at 100 Gigabit per second speeds and beyond. As link speeds increase to accommodate the need for bandwidth (averaging 70% year-over-year growth in ESnet's case), it is important that energy consumption not increase proportionally. Otherwise, the critical role networks play in reducing the carbon impact of information communications technology (ICT) will not be realized. To address this risk, ESnet is working with vendors and partners to characterize and quantify the problem by instrumenting its network for power consumption. When this project is complete, ESnet will be

the first national-scale network to collect and release measured power data to the research community and industry.

FNAL is one of the largest users of ESnet, sharing large amounts of data to collaborate with the scientific community. Without high-speed networks like ESnet, this type of collaboration would not be possible.¹

In addition, ESnet is working with research partners and the larger scientific community to explore architectural solutions to the problem of network energy consumption by moving long-lived data flows down to the optical layer, for example, or by rerouting traffic during periods of low utilization. Finally, ESnet is reducing its own carbon footprint by aggressively virtualizing servers, migrating to more efficient blade technology, promoting video conferencing as an alternative to travel, and improving data center infrastructure.

LBNL Develops Data Center Profiler Software Tool Suite (DC Pro) to Aid Data Center Efficiency

LBNL partnered with DOE's Energy Efficiency and Renewable Energy (EERE) Industrial Technologies Program to develop a suite of tools that help data center administrators determine how to most efficiently and sustainably manage their assets (Figure 30). This free software, called Datacenter Pro (DC Pro), helps data center

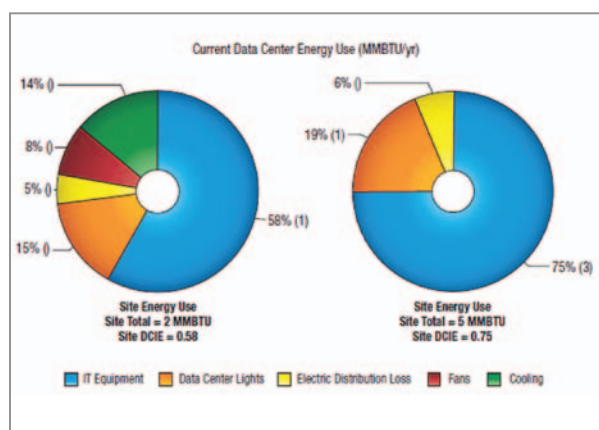


Figure 30. Sample DC Pro Profiler output for a data center. Image courtesy of LBNL

managers determine the efficiency of their centers and evaluate different energy saving opportunities.² The software includes a profiling tool and a set of system assessment tools that allow administrators to assess energy use in specific areas of their data center. LBNL and other labs not only developed the software, but they also test it operationally and inform its direction.

LBNL and NREL Partner to Develop Globally Available Tools to Support Green Data Centers

In conjunction with the Federal Energy Management Program (FEMP), LBNL and NREL published the “FEMP Best Practices Guide for Energy-Efficient Data Center Design (Figure 31).”³ This guidebook details planning and implementation practices, tools, and metrics designed to help any data center operate sustainably. It includes tips for purchasing desktop and data center equipment, configuring a data center for optimum operation, managing air flow and power, and applying full life cycle management practices.



Figure 31. Best practices guide for energy-efficient data center design. Image courtesy of NREL

1 FermiLab. “Networking.” FermiLab Computing, <http://www.fnal.gov/pub/science/computing/networking.html>. Accessed May 10, 211.

2 Department of Energy Industrial Technologies Program. “Data Center Profiler Software Tool Suite.” *Saving Energy in Data Centers* online, <http://www1.eere.energy.gov/industry/datacenters/software.html>. Accessed May 10, 2011.

3 Lintner, W.; Tschudi, B.; VanGeet, O. “FEMP Best Practices Guide for Energy-Efficient Data Center Design.” *Federal Energy Management Program* online, <http://www1.eere.energy.gov/femp/pdfs/eedatacenterbestpractices.pdf>. Accessed May 5, 2011.

GREEN IT LEADERSHIP PROVIDED BY LABS

NREL and LBNL also partnered with the Green Grid and FEMP to develop a new metric targeted toward measuring the reuse of energy from data centers. This Energy Reuse Effectiveness (ERE) metric measures the reuse of heat energy captured from the data center and distributed to energy-consuming areas external to the data center. In the past, the PUE measurement was the only metric available to data center administrators to measure energy efficiency. This metric measures power management efficiency in the data center, but not heat reuse in other areas. The ERE is targeted specifically toward how well heat energy is collected from the data center and used to heat non-data center spaces. Green Grid authors and the labs are presenting a paper on this metric at the 2011 ASHRAE summer conference.

NREL Supports Design and Operations of Low Energy Usage Data Centers

NREL has significant expertise and experience in low-energy data center design and operation. The lab has brought this expertise to the market by partnering with government and industry to plan and implement new and retrofitted data centers sustainably.

NREL has developed innovative approaches to data center cooling using very low cooling energy economizers and evaporative cooling. The lab has also developed a cooling energy analysis tool to quickly calculate the energy-saving potential for low-energy cooling compared to computer room air conditioning (CRAC) units using typical meteorological year (TMP) data.

The following are just two of the agencies NREL has partnered with to help them improve data center efficiency.

National Snow and Ice Data Center (NSIDC) at University of Colorado, Boulder:

The legacy NSIDC is being reconfigured to virtualize and consolidate data center equipment and rearrange equipment in a “hot-aisle/cold-aisle” configuration with curtains. Additionally, CRAC units are being replaced with an economizer and Coolerado® multipass indirect evaporative coolers (Figure 32). The new cooling system came online May 2011 and is expected to reduce annual cooling energy by 96%.

NREL partnered with Coolerado, a product division of IDALEX® Technologies USA, to develop its innovative cooling core and to provide testing and analysis of early




Figure 32. National Snow and Ice Data Center CRAC upgrade results. Illustration courtesy of NREL

prototypes as well as the initial product offering. This is the first major application of the Coolerado technology in a data center.

Environmental Protection Agency (EPA) Data Center located in Denver, Colorado:

The three-year-old EPA data center was designed for energy efficiency. An energy audit by NREL in February 2011 identified 14 energy conservation measures that provide additional energy efficiency opportunities. These were consolidated into five general categories, including IT environment improvements, airflow management optimization, equipment replacement, cooling, system retrofit, and lighting upgrades.

When the EPA completes implementation of these measures, its data center's PUE metric is expected to be reduced from the current 3.2 to approximately 1.5. A PUE of 2.0 is considered standard for data centers.

SLAC and Stanford Partner on New Modular Data Center

SLAC and Stanford University are collaborating on a joint venture to design a new Data Center Building—the Scientific Research Computing Facility (SRCF). This new building is designed to LEED Gold standards and sports a new design concept; its design is modular, allowing for incremental phasing as required by the lab and university (Figure 33). The SRCF is anticipated to consist of two modules. The initial phase is planned for approximately 19,000 square feet, and the second phase will be about 26,000 square feet.

The SRCF's modular concept allows for a flexible building design and systems that can support various IT equipment, including water cooled racks.

The concept also contains a goal of high density—the power and cooling systems must be flexible and able to support energy-efficient racks. The facility will be energy efficient, requiring that the concept incorporate the latest in energy-efficient design strategies to achieve a PUE target of less than 1.3.

The SRCF's design concept also incorporates operational economy, working toward lower operational costs as compared to traditional IT facilities.

■ HIGH PERFORMANCE COMPUTING

High-performance computing (HPC) uses considerably more power than an average data center uses for business applications.¹ DOE has 89 data centers overall, including 26 HPC data centers. Where HPCs make up 30% of DOE

data centers, they account for 90% of power consumption across the entire department.

Despite this huge demand for energy, the sustainable value that DOE laboratory HPC data centers provide is much greater than all other DOE laboratory efforts in IT sustainability combined. HPC data centers typically achieve much better PUEs compared to administrative computing data centers.

The HPC systems housed in HPC data centers are also models of sustainability. Many times, the utilization for these systems far exceeds 90%, while many business systems operate at less than 10%. HPC systems allow for fewer physical laboratories by enabling computational research. The research outcomes for tens of thousands of projects that are run on computational science resources annually result in an immeasurable amount of environmentally sustainable benefits. Finally, HPC system owners are very focused on maximizing the flops-per-watt performance, ensuring the highest performance for the least amount of energy use.

HPC systems have made important contributions in DOE's IT sustainability. This section highlights the accomplishments of the labs in HPC.




Figure 33. SLAC and Stanford University modular data center design concept.

Illustration courtesy of SLAC National Laboratory

20 DOE Labs High Performance Computing Facilities Ranked in Top 100 on the Green 500 List; Other DOE Lab HPCs fall in the Top 500

The supercomputing industry has traditionally measured super computers by their speed. From a green IT

¹ Dunlap, J.; Powers, C.; Sartor, D.; Wooley, J. "Lean Green (Data Center) Machine - DOE's Holistic Approach: Expert Panel from DOE OCIO and National Laboratories to Discuss Data Center Efficiency Best Practices." *GITEC Summit presentation; March 6-9, 2011, Orlando, Florida.*

HIGH PERFORMANCE COMPUTING

standpoint, the problem with rating super computers using this metric is this—with great speed comes great energy consumption.

HPC provides significantly increased ability to do science through modeling problems that would be very expensive to characterize in the laboratory. LLNL reports that the energy costs to run just one of their HPC labs is roughly \$14 million annually for power consumption alone, with more than half of that going into cooling the very fast computers.² It is evident that HPCs must be more than just fast—they need to be both fast and energy efficient. And they need to be both without impeding reliability.

The industry uses the Green 500 List to evaluate HPC systems based on their speed, reliability, and energy efficiency. Companies that implement the most efficient HPCs are then ranked on the list based on the number of megaflops per watt the system can produce. This means that the more computational operations that can be performed for every watt of energy, the more energy efficient the system. The DOE has proven a leader on the Green 500 List, evidencing 20 DOE lab HPCs in the Top 100, with many other lab facilities included in the Top 500.³

The DOE has more supercomputers than any other institution or govern-

ment in the world on the Green 500 List. Those in the Top 100 are included in **Table 8.**⁴

Table 8. DOE labs in Green 500 List top 100

Rating Points		Owner	Computer
1	1684.20	Argonne/LBNL	Blue Gene/Q Prototype
12	628.13	Lawrence Livermore National Laboratory	Appro GreenBlade Cluster Xeon X5660 2.8Ghz, nVIDIA M2050, Infiniband 160.00
15	458.33	DOE/ LANL	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Infiniband
17	444.25	DOE/NNSA/LANL	BladeCenter QS22/LS21 Cluster PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband
28	366.58	DOE/NNSA/LLNL	Blue Gene/P Solution
29	363.98	DOE/SC/Argonne National Laboratory	Blue Gene/P Solution
31	362.20	DOE/SC/LBNL/NERSC	Cray XE6 12-core 2.1 GHz
43	318.32	NOAA/Oak Ridge National Laboratory	Cray XT6-HE, Opteron 6100 12C 2.1GHz
57	292.21	DOE/SC/Argonne National Laboratory	iDataPlex, Xeon X55xx QC 2.66 GHz, Infiniband
59	288.48	Lawrence Livermore National Laboratory	Dell Xanadu 3 Cluster, Xeon X5660 2.8 Ghz, QLogic InfiniBand QDR
60	285.16	DOE/SC/LBNL/NERSC	iDataPlex, Xeon X55xx QC 2.66 GHz, Infiniband 129.25
60	285.16	DOE/SC/LBNL/NERSC	iDataPlex, Xeon X55xx QC 2.66 GHz, Infiniband

Argonne Lab's Leadership Computing Facility Supercomputer

Argonne's Leadership Computing Facility (ALCF) is a supercomputing facility that enables breakthrough science across the globe. Everything from modeling the Earth's climates to modeling tools for energy research is run through this facility. The ALCF is one of only two such leadership computing facilities in the country that provides both

world-class supercomputing resources, and dedicated teams of computational scientists and engineers to support

research efforts across a wide spectrum of scientific disciplines.

The ALCF meets the demands of the advanced simulation and modeling community without compromising the need for environmental sustainability. The facility's HPC (the IBM® Blue Gene/P) is one of world's fastest supercomputers, ranking in the Top 20 of the Green 500 list.⁵

2 Sharma, S.; Chung-Hsing, H.; Wu-chun, F. "Making a Case for a Green500 List." *Green 500 List* online, <http://www.green500.org/docs/pubs/hp-pac2006.pdf>. Accessed May 5, 2011.

3 Eagan, J. Personal e-mail. U.S. Department of Energy, Washington, DC, April 28, 2011.

4 Green 500 List. "The Green Lists." *Green 500 List* online, <http://www.green500.org/lists.php>. Accessed April 28, 2011.

5 West, J. "Baby, It's Cold Outside...Let's Calculate Something." *HPCwire* Online, http://www.hpcwire.com/hpcwire/2008-12-18/baby_its_cold_outside_lets_calculate_something.html. Accessed May 9, 2011.

The “Intrepid” features 40,960 quad-core compute nodes (163,840 processors), 80 terabytes of memory, and a peak performance of 557 teraflops (**Figure 34**).¹ With all of this power at its call, Intrepid still uses only one-third of the electricity of comparable HPCs.² The facility’s “Surveyor” Blue Gene/P system sports 1,024 quad-core nodes (4,096 processors), and two terabytes of memory that the site uses for testing and development tasks. The Intrepid and Surveyor use just over 1 MW of power, rather than the several required for other such facilities. This saves taxpayers over \$1 million annually.³

The ALCF takes advantage of the cold Chicago winters to chill the water used to cool the machine room free for several months of the year. ANL also implemented smart power management functions in the ALCF, including turning off chips and storage systems when they are not in use, and scheduling intensive compute jobs to run at night when the power grid has more capacity and temperatures are lower.

The Blue Gene/P supercomputer is high on the Green 500 list of supercomputers; however, the ALCF is not stopping there. The facility worked with IBM to develop the next generation Blue Gene/Q supercomputer, which is four times faster than today’s fastest supercomputer and is in line to be the No. 1 green supercomputer in the world.⁴ The ALCF will implement “Mira,” a Blue Gene/Q supercomputer, in 2012. Mira will feature 10-petaflops of performance, with 48K 16-way compute nodes (768K processors), and 768 terabytes of memory.

INL Constructs Mission Appropriate HPC Designed to Accommodate Future Requirements

In 2007, INL completed the construction of a 3,700 square-foot data center to support its HPC resources. The data center was also designed with the ability to readily expand to 10,000 square feet in the near future.

Several practices were incorporated to assist with meeting the lab’s energy efficiency goals. The data center space was right-sized for initial needs, and by implementing a mission appropriate data center size, the lab minimized the associated operating energy costs.




Figure 34. ALCF’s “Intrepid” memory modules.
Photo courtesy of Argonne National Laboratories

The data center was also designed to take advantage of “free cooling.” As long as the outside temperature is between 40° and -31°F, and the temperature of the water leaving the data center is not too high, the chillers do not operate. The water used to cool the heat-transfer plate, and subsequently to extract the heat from the data center water, can be solely cooled through the evaporative process of the cooling towers. This condition of “free cooling” saves an enormous amount of electrical costs by not having to run the expensive chillers.

The data center’s compute clusters were designed with water-cooled doors to improve the overall cooling efficiency of the facility. The exhausted warm air from the compute nodes is immediately cooled as it passes through the rear cooling doors on the racks and reenters the room at temperatures near those of the open air in the data center. The CRAC units are

1 Argonne National Laboratory. “Argonne Leadership Computing Facility Leadership-Class Computing for breakthrough science and engineering.” *Argonne National Laboratory online*, http://www.alcf.anl.gov/news/media_files/alcf-fctsht-0411_r6.pdf. Accessed May 5, 2011.

2 Argonne National Laboratory. “Argonne Leadership Computing Facility: Producing more science per watt for a greener America.” *Argonne National Laboratory online*, http://www.alcf.anl.gov/news/media_files/alcf-greenpstr-0410.pdf. Accessed May 5, 2011.

3 Argonne National Laboratory. “Argonne Leadership Computing Facility Leadership-Class Computing for breakthrough science and engineering.” *Argonne National Laboratory online*, http://www.alcf.anl.gov/news/media_files/alcf-fctsht-0411_r6.pdf. Accessed May 5, 2011.

4 Argonne National Laboratory. “Argonne Leadership Computing Facility Leadership-Class Computing for breakthrough science and engineering.” *Argonne National Laboratory online*, http://www.alcf.anl.gov/news/media_files/alcf-fctsht-0411_r6.pdf. Accessed May 5, 2011.

HIGH PERFORMANCE COMPUTING

cooled by the chilled water, which requires less work and reduces electricity consumption (**Figure 35**).

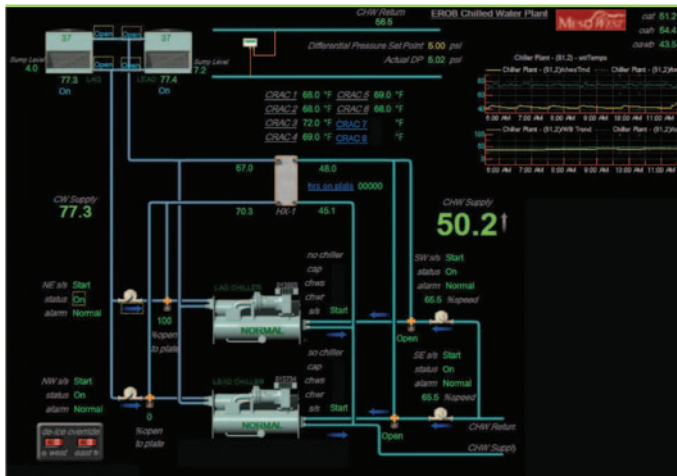


Figure 35. INL data center chilled water metrics. Image courtesy of INL

Finally, the last HPC cluster procured was designed using the latest technologies in high-density processors from AMD®. There are four 8-core processors in each node (versus other configurations), which greatly reduced the requirements for space, power, and cooling of the entire system. The data center’s configuration includes fewer racks, fewer nodes, less network infrastructure, and fewer power supplies—all resulting in less consumed power.

As a result of these efforts, the PUE rating for the INL HPC data center ranges from 1.1 to 1.3, depending on system load and outside weather conditions (**Figure 36**).

Instantaneous Power (5/19/11)	Consumption (KW)	3-Year Power (2/2008-5/2011)	Consumption (MWh)
Compute	420	Compute	7018
Cooling	140	Cooling	2750
PUE	1.3	PUE	1.4

Figure 36. INL HPC data center PUE.

Jefferson Lab Works with NVIDIA to Pioneer the Use of GPUs for HPC

Jefferson lab pioneered the use of graphics processing units (GPU)—the chips that allow processing-power-hungry video games to be so realistic—in supercomputers. GPUs must be able to compute and draw images that are too small for the eye to see at very high speeds, making them uniquely suited to depicting subatomic particles.⁵ The lab combined standard computer processors (CPU) with the video and gaming processors found in game consoles like the Sony Playstation® to bring 100 Teraflops of processing power to scientific applications.

Jefferson lab’s supercomputer, called Hadron, includes 480 GPUs and 266 CPUs (**Figure 37**). The GPU processing power is “used to compute how the building blocks of matter, quarks, build the protons, neutrons and other particles that makeup everyday matter.”⁶ Scientists then analyze the results on the CPU side of the system. The calculations needed



Figure 37. Jefferson Lab’s Hadron Supercomputer.

Photo courtesy of Jefferson Lab

to describe the properties of these subatomic particles were almost impossible to carry out with ordinary CPUs. Jefferson lab’s breakthrough using GPUs for these calculations made the science possible.

A quad GPU compute node (95% of the lab’s capacity) has an effective performance of 232 Intel cores, or 29 dual socket quad core nodes, that otherwise would have needed to be bought. The quad GPU node directly uses 1 kW, whereas the dual socket CPU node uses 280 W. Each four-GPU server saves 7.1 kW. Jefferson lab has 119 such nodes, which saves 845 kW compared to achieving the same computing power with conventional nodes at that time. Power conditioning and cooling savings pushes this number to more than 1 MW.

5 Karter, C. Hot Graphics Cards Fuel Supercomputing. *On Target* online, <http://www.jlab.org/news/OnTarget/2010/2010-06/June2010.pdf>. Accessed June 1, 2011.

6 Karter, C. JLab Cluster Tops 100 Teraflops. *Jefferson Lab News Releases* online, <http://www.jlab.org/news/releases/2010/2010-100Teraflops.html> Accessed June 1, 2011

In addition to the power savings, Jefferson lab also significantly reduced costs using GPUs. To gain the processing power of Hadron using traditional CPUs, Jefferson lab would have spent more than \$100,000. The implementation of GPUs cost the lab about \$6,000.¹

Jefferson Lab's Sustainable Technology Refresh

According to Moore's Law, the ability to place an increasingly larger number of transistors on a computer chip without increasing the size of the chip enables the performance of computing technology to double every 18-24 months. The energy required to realize this doubling of performance will increase only slightly, if at all.

Jefferson lab has developed a sustainable technology refresh policy for its HPCs that follows Moore's Law. The



Figure 38. Jefferson lab's 2010 high performance computing clusters. Photo courtesy of Jefferson Lab

policy states that if the energy savings that will be realized by upgrading a compute node exceeds one-half of the purchase price of upgrading the device, the upgrade will be performed. If the energy savings does not meet this criteria, the node upgrade will wait until the energy savings parameter is met (**Figure 38**).

The UC Berkeley Computational Research and Theory Supercomputing Building

The proposed University of California Computational Research and Theory Building (CRT) represents a best-in-class green data center optimized for scientific supercomputing workloads. The building incorporates numerous sustainable optimization approaches, which are driven by research conducted by LBNL scientists and engineers.

The CRT building will use free-cooling, flexible cooling approaches that are highly load adjustable; water-side economization for maximum efficiency; and heat recovery for building water and air management. The proposed building is expected to have a Data Center Infrastructure Efficiency (DCIE) above 0.83. This metric is calculated by dividing IT equipment power by total facility power.

The mild climate in Berkeley provides an ideal opportunity to remove massive amounts of heat from data center equipment using free cooling throughout most of the year. Through aggressive use of the upper end of the ASHRAE recommended temperature range, coupled with direct use of outside air, chiller energy demand was minimized in the design. If liquid cooling emerges as the preferred vendor cooling solution, then additional energy and power efficiency gains are possible; however, it will require added water use. Best practices identified through previous investigations and benchmarking were incorporated into the design, which resulted in a very energy-efficient design.

LLNL Realizes Big Cost Savings with Temperature Adjustments in HPC

A four-year effort by the Computation Directorate to conserve energy by raising the temperature and managing air flow in the Terascale Simulation Facility (TSF) computer



Figure 39. LLNL Terascale Simulation Facility. Photo courtesy of LLNL

rooms is yielding big dividends—an estimated \$3.5 million in energy savings annually so far (**Figure 39**).

Called “Megawatts to PetaFLOPs,” the program earned a 2009 Department of Energy/National Nuclear Security Administration FEMP award. The savings derived from a

multi-track approach. This included increasing the temperature in the computer rooms, implementing a chilled water

¹ Karter, C. Hot Graphics Cards Fuel Supercomputing. *On Target* online, <http://www.jlab.org/news/OnTarget/2010/2010-06/June2010.pdf>. Accessed June 1, 2011.

HIGH PERFORMANCE COMPUTING

system that cools the computers, and finding efficiencies in the configuration of high-performance computing (HPC) systems.

As it turns out, raising the temperature in the TSF's computer rooms was not simply a matter of adjusting the thermostat. For starters, the TSF east and west computer rooms (which total 48,000 square feet), house multiple HPC systems, each with different cooling requirements. To address these varied requirements, LLNL used a self-benchmarking tool available from DOE Federal Energy Management Program (FEMP) and LBNL. The lab also modeled the air flow required to adequately cool machines using computational fluid dynamics.

Starting in October 2007, the temperature in the computer rooms was raised about 1°F each month, with impacts on the systems carefully monitored to ensure the smooth 24/7 operation of this high-performance computing facility. Since the program began, the temperature has been raised 12°F to 64°F in the east room and 10°F to 62°F in the west room. In addition, the temperature of the cooling water from the chillers also has been raised 8 degrees from 42° to 50°F. The temperature was raised again in the west room in June 2011 to accommodate the installation of the first indirect liquid cooling Petascale machine. The air supply temperature for this machine must be a minimum of 70°F.

Computing systems in the TSF are cooled by blowing air from the ground level up through the computer room floor. The air is then moved up through the machines and vented out the top or the sides, depending on system design. The air is also recirculated out the ceiling and back down to the ground level through the walls — the reverse of a clean room. Air flow in the computer rooms is controlled by using solid, grated, or perforated floor tiles, which allow adjustment of the temperature and the creation of “hot and cold aisles.” All of this must be done while equalizing the pressure in the room.

The computational fluid dynamics model of the computer rooms also revealed 450,000 cubic feet per minute of air leakage from louvered perforated tiles on the computer

room floor. By replacing the leaking tiles, the team was able to shut down two 100 horsepower 80,000 cubic-foot-per-minute air handlers.

The TSF has “excellent power efficiency,” noting the building has an industry standard PUE rating of 1.34 (where the most efficient is 1.0 and most inefficient is 3.0).

New energy conservation initiatives are being designed for the lab's Building 453. One option under consideration is a “free cooling” air economization project. The goal is to modify the existing building air-handling system. The upgrade may include the installation of louvers, which are dampers that can bring outside air in or send exhaust air out. It may also include implementing humidifiers, exhaust fans, and associated controls to provide free cooling to the HPC space for those times when the outside temperature is suitable for computer-room cooling. In this way, the TSF will be able to take advantage of “free” seasonal and nighttime outside air variations to help cool the computer rooms. This free-cooling, outside air-handling system will be built into a new wall structure of the building, with the majority of the system being on the north side of Building 453 and a small portion on the south side. The upgrade is estimated to result in an annual savings of \$600,000 per year.

LLNL has approached updates to additional HPCs in the lab through the lens of lessons learned and successes gained from the updates to the TSF and Building 453. Most notably, Buildings 451, 439, 115, and 117 are next on the radar for Green IT updates.

LLNL is actively pursuing LEED certifications for buildings containing HPCs. B-453 was LEED Gold Certified in December 2010 and B-451 was LEED Silver Certified in April 2011.

[ORNL Works with SmartTruck Systems to Design Energy-Efficient Semi Trucks](#)

HPCs help the scientific community perform calculations that would never be possible without access to these large-scale systems. Due to a partnership between ORNL and SmartTruck Systems, green HPCs are now also helping the nation meet its goal of fuel independence. SmartTruck



Figure 40. ORNL's Jaguar HPC. Photo courtesy of ORNL

Systems leveraged the processing power of ORNL's Jaguar HPC (Top 100 on the Green 500 list) to help create add-on parts that are increasing the fuel efficiency of semi trucks. Retrofitting these 18-wheelers will save billions of gallons of fuel annually over conventional trucks. Jaguar (Figure 40) is at ORNL's Leadership Computing Facility (OLCF).

Semi trucks average 6 mpg or less and add 423 million pounds of CO₂ to the atmosphere annually.¹ To address the need to retrofit existing trucks to be environmentally sustainable, SmartTruck Systems first turned to its in-house

cluster to model semi aerodynamics. This system did not have the processing power to run the complex simulations of airflow around the truck and verify that the company's

add-on part designs would meet its fuel efficiency goals. SmartTruck Systems turned to ORNL for access to Jaguar, a 2.33 petaflop HPC, to model the air flow simulations and help design an under tray system to increase the aerodynamics for semis. With access to Jaguar, SmartTruck Systems ran simulations that broke the truck into hundreds of pieces so they could calculate air drag accurately, and ran those simulations 10 times faster than was possible on their in-house system.² The resulting UnderTray System works by compressing and accelerating high-energy incoming air and attached airflow from the top of the trailer down into the trailer's wake (Figure 41).

SmartTruck Systems thought that the innovation would take three-and-a-half years to go to market – with Jaguar, they went to market in just 18 months.³

The new technology achieves fuel savings of between 7% and 12% over a conventional semi, translating to a significant




Figure 41. ORNL's Jaguar facilitates design of an energy-efficient UnderTray System for semis. Simulations courtesy of SmartTruck Systems

1 Oak Ridge National Laboratory. "Designing a Smart Truck with the Power of Jaguar." Industrial HPC Case Study.

2 Oak Ridge National Laboratory. "Designing a Smart Truck with the Power of Jaguar." Industrial HPC Case Study.

3 Schwartz, A. "Supercomputer Designs Energy-Efficient Trucks, Potentially Saving Billions in Fuel Costs." *Fast Company* online, <http://www.fastcompany.com/1725255/jaguar-supercomputer-designs-energy-efficient-trucks?partner=rss>. Accessed June 20, 2011.

CONCLUSION

decrease in the amount of fossil fuels needed to transport goods. Several truck companies have already installed the SmartTruck UnderTray System in their fleets,

benefiting from savings in operating costs and contributing to a greener planet. Fast Company states that “if all 1.3 million semi trucks in the U.S. used the SmartTruck UnderTray, it would save 1.5 billion gallons of diesel fuel (\$5 billion in costs) and cut CO₂ emissions by 16.4 million tons.” These results would not have been possible without this high-performance green machine.

Sandia Labs Upgrades Super Computer for Ultra Efficiency

Sandia Labs’ Computing and Network Services Center (CNSC) built the Red Sky super computer in response to the nation’s critical need for HPC (Figure 42). The project involved the decommissioning and replacement of the Thunderbird HPC system and the deployment of the Red Sky HPC system.



Figure 42. Sandia’s Red Sky HPC. Photo courtesy of Sandia National Laboratories

One of the goals of Red Sky was to have an HPC system that performed faster than its predecessor—Thunderbird. In addition to raw horsepower, the lab set a goal that its newest supercomputer would be designed to maximize its eco-efficiency by using cutting-edge technological innovations.

Following implementation of HPC facility best practices, Red Sky needed 40% less water for cooling, resulting in more than 5 million gallons saved annually. Cooling efficiency was increased tenfold, from 70% to 97%. The HPC’s carbon footprint was reduced to 203 tonnes of carbon dioxide equivalents, or CO₂E and Footprint (global hectares or gha) of 46. The replaced Thunderbird HPC had a carbon footprint of 912 tonnes CO₂E and Footprint (gha) of 205.

Each year, the National Nuclear Security Administration’s (NNSA) awards national laboratories one of 21 Pollution Prevention (P2) Awards for innovative initiatives across the enterprise.⁴ Sandia was awarded the Best in Class “Comp. Energy &/or Fleet Management: High Performance Computing Water Reduction and Energy Efficient Cooling” in 2011 for the Red Sky HPC project.

CONCLUSION

The DOE labs are showing exceptional leadership in Green IT. Whether in full life cycle management, energy-efficient data centers, high-performance computing, or managing the end-user computing environment, the labs are constantly looking for ways to save taxpayer money and support sustainability through best practice implementation, innovation, and partnerships.

We hope this guide offers ideas to IT professionals looking for ways to make their IT more sustainable.

⁴ National Nuclear Security Administration. “NNSA Recognizes Labs and Sites for Achievements in Environmental Stewardship.” *National Nuclear Security Administration* online, <http://nnsa.energy.gov/mediaroom/pressreleases/p2awards32311>. Accessed May 11, 2011.

■ REFERENCES

- Argonne National Laboratory. "Argonne Leadership Computing Facility Leadership-Class Computing for breakthrough science and engineering." *Argonne National Laboratory* online, http://www.alcf.anl.gov/news/media_files/alcf-fctsht-0411_r6.pdf. Accessed May 5, 2011.
- Argonne National Laboratory. "Argonne Leadership Computing Facility: Producing more science per watt for a greener America." *Argonne National Laboratory* online, http://www.alcf.anl.gov/news/media_files/alcf-greenpstr-0410.pdf. Accessed May 5, 2011.
- Bell, G. C. "Retro-commissioning increases data center efficiency at low-cost: Success at Savannah River Site (SRS) at low-cost. Success at Savannah River Site (SRS) at low-Cost." *Federal Energy Management Program* online, http://www1.eere.energy.gov/femp/pdfs/datacenter_savannah.pdf. Accessed May 4, 2011.
- Bellendir, G. "Fermilab data center receives EPA Energy Star award." *FermiLab Today* online, http://www.fnal.gov/pub/today/archive_2011/today11-02-04.html. Accessed May 10, 2011.
- Dunlap, J.; Powers, C.; Sartor, D.; Wooley, J. "Lean Green (Data Center) Machine - DOE's Holistic Approach: Expert Panel from DOE OClO and National Laboratories to Discuss Data Center Efficiency Best Practices." *GITEC Summit presentation*; March 6-9, 2011, Orlando, Florida.
- Eagan, J. Personal e-mail. U.S. Department of Energy, Washington, DC, April 28, 2011.
- Federal Electronics Challenge. <http://www.federalelectronicschallenge.net/>. Accessed May 02, 2011.
- FermiLab. "Grid Computing." *FermiLab* online, <http://www.fnal.gov/pub/science/computing/grid.html>. Accessed May 10, 2011.
- FermiLab. "Networking." *FermiLab Computing* online, <http://www.fnal.gov/pub/science/computing/networking.html>. Accessed May 10, 2011.
- Karter, C. Hot Graphics Cards Fuel Supercomputing. *On Target online*, <http://www.jlab.org/news/OnTarget/2010/2010-06/June2010.pdf>. Accessed June 1, 2011.
- Karter, C. JLab Cluster Tops 100 Teraflops. *Jefferson Lab News Releases* online, <http://www.jlab.org/news/releases/2010/2010-100Teraflops.html>. Accessed June 1, 2011.
- Green 500 List. "The Green Lists." *Green 500 List* online, <http://www.green500.org/lists.php>. Accessed April 28, 2011.
- Lintner, W.; Tschudi, B.; VanGeet, O. "FEMP Best Practices Guide for Energy-Efficient Data Center Design." *Federal Energy Management Program* online, <http://www1.eere.energy.gov/femp/pdfs/eedatacenterbest-practices.pdf>. Accessed May 5, 2011.
- Mellanox® Technologies. "NERSC Creates an Ultra-Efficient Supercomputer." *Mellanox Technologies* online, http://www.voltaire.com/ResourceCenter/customer_success/nersc_creates_an_ultra-efficient_supercomputer. Accessed May 11, 2011.
- National Nuclear Security Administration. "NNSA Recognizes Labs and Sites for Achievements in Environmental Stewardship." *National Nuclear Security Administration* online, <http://nnsa.energy.gov/mediaroom/press-releases/p2awards32311>. Accessed May 11, 2011.
- Sarokin, D. "Question: Energy Use of Internet." *Uclue* online, <http://uclue.com/index.php?q=724>. Accessed May 26, 2011.
- Schwartz, A. "Supercomputer Designs Energy-Efficient Trucks, Potentially Saving Billions in Fuel Costs." *Fast Company* online, <http://www.fastcompany.com/1725255/jaguar-supercomputer-designs-energy-efficient-trucks?partner=rss>. Accessed June 20, 2011.
- Sharma, S.; Chung-Hsing, H.; Wu-chun, F. "Making a Case for a Green500 List." *Green 500 List* online, <http://www.green500.org/docs/pubs/hp-pac2006.pdf>. Accessed May 5, 2011.
- Simon, H.; Yelick, K.; Broughton, J.; Draney, B.; Bashor, J.; Paul, D.; Vu, L.; Beckman, P.; Coghlan, S.; Taylor, E. "Exploring CLOUD Computing for DOE's Scientific Mission." *SciDAC Review* online, <http://www.scidacreview.org/1002/html/hardware.html>
- U.S. Department of Energy Industrial Technologies Program. "Data Center Profiler Software Tool Suite." *Saving Energy in Data Centers* online, <http://www1.eere.energy.gov/industry/datacenters/software.html>. Accessed May 10, 2011.
- U.S. Department of Energy, Office of Science. "Laboratories." *Office of Science* Online, <http://science.energy.gov/laboratories/>. Accessed April 19, 2011.
- U.S. Environmental Protection Agency Energy Star Program. "Report to Congress on Server and Data Center Energy Efficiency, Public Law 109-431." *Energy Star* online, http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final.pdf. Accessed May 26, 2011.
- West, J. "Baby, It's Cold Outside...Let's Calculate Something." *HPCwire* Online, http://www.hpcwire.com/hpcwire/2008-12-18/baby_its_cold_outside_lets_calculate_something.html. Accessed May 9, 2011.

RESOURCES

RESOURCES FOR GREEN IT

- American Society of Heating, Refrigerating and Air-Conditioning Engineers <http://ashrae.org/>
- Data Center Pro (DC Pro) <http://www1.eere.energy.gov/industry/datacenters/software.html>
- Department of Energy DOE O 430.2B, Departmental Energy, Renewable Energy and Transportation Management <https://www.directives.doe.gov/directives/current-directives/430.2-BOrder-b/view?searchterm=None>
- Department of Energy National Laboratories and Technology Centers <http://www.energy.gov/organization/labs-techcenters.htm>
- Department of Energy Office of Science <http://science.energy.gov/laboratories/>
- Department of Energy Strategic Sustainability Performance Plan http://www.energy.gov/media/DOE_Sustainability_Plan_2010.PDF
- Department of Energy Office of the Chief Information Officer <http://www.cio.gov/>
- Energy Sciences Network <http://www.es.net/>
- Energy Star <http://www.energystar.gov/>
- Environmental Protection Agency <http://www.epa.gov/>
- ERE: A Metric for Measuring the Benefit of Reuse Energy from a Data Center <http://www.thegreengrid.org/en/Global/Content/white-papers/ERE>
- Executive Order 13423 <http://www1.eere.energy.gov/femp/regulations/eo13423.html>
- Executive Order 13514 <http://www1.eere.energy.gov/femp/regulations/eo13514.html>
- Federal Electronics Challenge (FEC) <http://www.federalelectronicschallenge.net/>
- Federal Energy Management Program (FEMP) Best Practices Guide for Energy-Efficient Data Center Design <http://www1.eere.energy.gov/femp/pdfs/eedatacenterbestpractices.pdf>
- Federal Energy Management Program Data Centers http://www1.eere.energy.gov/femp/program/data_center.html
- Federal Energy Management Program <http://www1.eere.energy.gov/femp/>
- Federal IT Dashboard <http://it.usaspending.gov/>
- The Green 500 <http://www.green500.org/home.php>
- The Green Grid <http://www.thegreengrid.org/> U.S. Green Building Council <http://www.usgbc.org/>
- Industrial Programs Data Centers <http://www1.eere.energy.gov/industry/datacenters/>
- The National Laboratories Chief Information Officers <http://nlcio.nationalallabs.org/>
- NREL Renewable Resource Data Center <http://www.nrel.gov/rredc/>

DEPARTMENT OF ENERGY LABORATORIES LEADERSHIP IN GREEN IT



National Renewable Energy Laboratory

1617 Cole Boulevard, Golden, Colorado 80401
303-275-3000 • www.nrel.gov

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency and Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

DOE/GO-102011-3295 • August 2011

Printed with a renewable-source ink on paper containing at least
50% wastepaper, including 10% post consumer waste.