

Vision and Roadmap

**Routing Telecom and
Data Centers Toward
Efficient Energy Use**

DRAFT

Sponsored by

Emerson Network Power

Silicon Valley Leadership Group

**Telecommunications
Industry Association**

Yahoo! Inc.

In cooperation with the

U.S. Department of Energy

March 18, 2009

This document reflects the ideas and solutions expressed by participants in the “Vision and Roadmap Workshop: Routing Telecom and Data Centers Toward Efficient Energy Use” held on October 15 and 16, 2008, at the Yahoo! Headquarters in Sunnyvale, California. The workshop participants developed industry visions for the energy-efficient future of data centers and telecommunication central offices and identified some of the key challenges and potential solutions to significantly reduce energy requirements and costs to these growing and vital facilities.

The workshop was jointly sponsored by Emerson Network Power, the Silicon Valley Leadership Group, the Telecommunications Industry Association, and Yahoo! Inc. in collaboration with the U.S. Department of Energy’s Industrial Technologies Program in the Office of Energy Efficiency and Renewable Energy. Yahoo! Inc. graciously provided all meeting facilities, audiovisual services, and amenities.

DRAFT

**Vision and Roadmap
for
Routing Telecom and Data Centers
Toward Efficient Energy Use**

DRAFT

March 18, 2009

Contents

Foreword

- 1. Escalating Need for ICT Energy Efficiency 1**
 - Accelerating Energy Demand 2**
- 2. Technology R&D Challenges and Solutions 3**
 - A. Equipment and Software 7**
 - Chip-level and Automated Power Management 8
 - Dynamic Energy Consumption Management in Network Devices 9
 - Data Storage Technologies 10
 - Hardened ICT Equipment 12
 - Novel Computing Architectures 13
 - Nano-electronic Circuitry 14
 - All-Optical Networks 15
 - Superconductive Components 16
 - B. Power Supply Chain 17**
 - Eliminate Conversion Steps/Losses 19
 - High-Efficiency Power Conversion Circuits 20
 - Efficiency-Optimized Control Systems 21
 - Transition to DC Operation 22
 - On-Site DC Generation and Micro-Grid 23
 - C. Cooling 25**
 - Liquid Cooling of Components 26
 - Advanced Cooling at Chip/Server-level 27
 - Efficiency-Optimized Control Systems 28
 - Free Cooling & Standards Development 29

Appendices

- Supporting Technology Needs 31**
 - Metrics 31
- Workshop Breakout Results 31**

DRAFT

Foreword

Routing Telecom and Data Centers Toward Efficient Energy Use identifies promising technologies for solving the energy and reliability challenges facing our essential information and communication technology (ICT) sector. The document synthesizes expert input from facility operators, equipment and software developers, university and laboratory researchers, commercial vendors, industry associations, and others in the ICT community. These experts came together at the “Vision and Roadmap Workshop: Routing Telecomm and Data Centers Toward Efficient Energy Use” in October 2008.

These dedicated individuals recognized the importance of looking beyond incremental efficiency improvements to envision far more energy-efficient ICT facilities. Workshop participants developed industry visions, explored key challenges, and identified solutions to critical ICT energy issues in each of three energy-intensive areas: Equipment and Software, Power Supply Chain, and Cooling.

Members of the ICT Energy Efficiency Roadmap Organizing Committee would like to thank everyone who attended the workshop for their active participation and valuable contributions. The committee particularly appreciates all efforts by the many industry and laboratory experts who wrote descriptions for the advanced technology solutions identified in the workshop. Many thanks also go to the workshop sponsors and co-sponsors: Yahoo! Inc., Emerson Network Power, the Telecommunications Industry Association, the Silicon Valley Leadership Group, and other leading ICT companies. The Industrial Technologies Program within the U.S. Department of Energy’s Office of Energy Efficiency and Renewable Energy will consider the contents of this roadmap in shaping its R&D portfolio to achieve critical DOE missions and national goals.

ICT Energy Efficiency Roadmap Organizing Committee

David Bellandi, ICT Green Initiative,
Telecommunications Industry Association
Dennis DelCampo, Emerson Network Power
Bob Hines, Silicon Valley Leadership Group
Kevin P. Kenny, Sprint Nextel
Moe Khaleel and Andres Marquez,
Pacific Northwest National Laboratory
KC Mares, MegaWatt Consulting, *Chair*
Chris Page, Yahoo! Inc.
Larry Plumb, Verizon
Dale Sartor and William Tschudi,
Lawrence Berkeley National Laboratory
Gideon Varga, U.S. Department of Energy

Contributing Writers

Whit Allen, Liebman & Associates, Inc.
Kerry Bernstein, IBM T.J. Watson Research Center
Ken Brill, Uptime Institute
Brian Fortenbery, EPRI
Patrick Giangrosso, Eaton
Phil Hughes, Clustered Systems
Yogendra Joshi, Georgia Institute of Technology
Mark Johnson, Lineage Power
Doug Kelley, Cray
Kevin P. Kenny, Sprint Nextel
Moe Khaleel, Pacific Northwest National Laboratory
Daniel Kharitonov, Juniper Networks
Teja Kuruganti, Oak Ridge National Laboratory
Kalyan Perumalla, Oak Ridge National Laboratory
Chris Page, Yahoo! Inc.
Rick Ridgley, National Reconnaissance Office
John Shalf, Lawrence Berkeley National Laboratory
John Stanley, Uptime Institute
William Tschudi, Lawrence Berkeley National
Laboratory
Ben Yoo, University of California, Davis

Workshop Participants

Whit Allen, Liebman & Associates, Inc.
Ken Bikelar, Alcatel-Lucent
Haris Basit, Multigig, Inc.
Jonathan Biederer, Sprint/Nextel
Mark Bramfitt, Pacific Gas & Electric
Rich Brown, LBNL
Van Carey, University of California
Richard Craig, Verizon Wireless
Thomas G. Croda, CSI Communications
Joyce Dickerson, Stanford University
Andrew Dugan, Level 3 Communications
David Filo, Yahoo! Inc.
Brian Fortenbery, EPRI
Gary Garcia, Network Appliance
Ludwig Graff, Verizon Corporate
Shaun Harris, Microsoft
Jeff Holt, Fujitsu Network Communications
Nick Holt, Yahoo! Inc.
Leland Horstmann, AT&T
Phil Hughes, Clustered Systems
Tim Jeffries, ATIS
Mark Johnson, Lineage Power
Mike Jump, CBRE 7x24 Exchange
Aurangzeb Khan, Multigig Board of Directors
Moe Khaleel, PNNL
Daniel Kharitonov, Juniper Networks
Wison Korol, Nortel
Teja Kuruganti, Oak Ridge National Laboratory
Alexis Kwasinski, University of Texas
Ken Lear, AT&T
Davis Lewis, Nokia Siemens Networks
Frank Lukas, Qwest Communications
Chris Malone, Google
KC Mares, MegaWatt Consulting
Ed Mardiat, Burns & McDonnell
Bob Marinik, Emerson
Andes Marquez, Pacific Northwest National Lab.
David Mastrandrea, Jacobs
Mary Medeiros McEnroe, Silicon Valley Power
Ed Mikoski, Telecommunications Industry Assoc.
Ayaz Mundarris, Yahoo! Inc.
Kieran Nolan, AT&T
Tom Okrasinski, Alcatel Lucent
Chris Page, Yahoo! Inc.
Rakesh Patel, Altera Corporation
Michael K. Patterson, Intel
Jeremy Rodriguez, VMware
Paul Roggensack, CA Energy Commission
Donna Sadowy, AMD
Dale Sartor, Lawrence Berkeley National Lab.
Ken Schneebeil, IBM
Eddie Schutter, AT&T
Bob Seese, Advanced Data Centers
Steve Smith, Oak Ridge National Laboratory
Eric Soladay, Rumsey Engineers
John Stanley, Uptime Institute
Ron Thompson, Yahoo! Inc.
Roger Tiple, Hewlett-Packard
Bill Tschudi, Lawrence Berkeley National Lab.
Lou Warner, Verari System
Henry Wong, Intel

1. Escalating Need for ICT Energy Efficiency

Exponential growth in the demand for data processing and storage continues to stimulate rapid growth in the U.S. data center industry. Over the past six years, energy use by these centers and their supporting infrastructure is estimated to have increased by nearly 100 percent. If data center energy demand continues to grow at the current rate, the nation will need to build approximately two large power plants *per year* to meet the demand. Just one 500-MW plant can cost \$1.5 billion or more, depending on the technology.^{1,2} This power demand raises energy costs for business and government, strains the existing power grid, and can increase greenhouse gas emissions.

The telecommunications industry, which increasingly relies upon similar hardware and systems to provide signal processing and switching services, faces similar challenges to reduce energy use in processing equipment and cooling systems. These two rapidly growing industries, which together have been designated a “critical infrastructure” to the American economy, account for about three percent of total U.S. electricity use, and that share is projected to continue growing at a rapid rate.

As information and communication technology (ICT) services continue to converge, the data center and telecommunications industries face increasingly similar challenges to control the power usage of their ICT equipment and supporting power and cooling systems. In the face of growing global energy demand, uncertain energy supplies, and volatile energy prices, innovative solutions are needed to radically advance the energy efficiency of these systems, which represent the engine of the American economy today. Enhanced energy efficiency in the central offices and data centers supporting our ICT systems will enhance U.S. energy and economic security.

Computing technology has proven a pervasive, driving force in the U.S. economy over the past two decades. It affects nearly every aspect of life—from education, entertainment, transportation, and personal communication to the basic infrastructure of our economy, medicine, engineering, and science. Society has come to depend not just on computing but on the increases in computing capability that have been made available each year at a given cost and power budget. Today’s powerful search engines, sophisticated

Vital ICT Functions

Increased energy efficiency in data and telecom centers can improve power reliability and security for vital functions and services:

- Financial services (e.g., on-line banking and electronic trading)
- Emergency, health, and safety services and disaster recovery
- Retention of electronic records (e.g., medical and government records)
- Global commerce and telecommunications
- Internet communication and entertainment
- Satellite navigation and electronic shipment tracking
- Information security and national security
- High-performance scientific computing
- E-commerce and telecommuting to reduce freight and passenger transport energy use

¹ EPA 2007 *Report to Congress on Server and Data Center Efficiency*, p.7-8. A large power plant is 500 MW.

² Power plant costs based on overnight capital cost estimates of \$3,000 per kW for a new Integrated Coal-Gasification Combined Cycle (IGCC) plant, without carbon capture. (The cost is higher if carbon capture is included). Estimates from: Energy and Environmental Economics, Inc. (E3). 13 May 2008. "Greenhouse Gas Calculator for the California Electricity Sector." Energy and Environmental Economics, Inc. (E3). Available at www.ethree.com/GHG/GHG%20Calculator%20v2b.zip. Accessed 6 Feb 2009. Another source is: Energy Information Administration (EIA). June 2008. "EIA - Assumptions to the Annual Energy Outlook 2008." EIA. Available at www.eia.doe.gov/oiaf/aeo/assumption/pdf/tbl38.pdf. Accessed 6 Feb 2009.

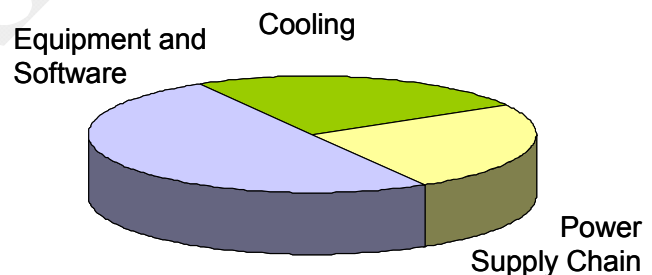
personal electronics, and other advances (e.g., medical breakthroughs, climate models, etc.) would not have been possible without these major increases in data processing and computing performance. As demand for ICT services continues to grow, ICT facilities will need radical technology advances to enable continued support and expansion of the societal and economic benefits they deliver.

Accelerating Energy Demand

The three percent of U.S. electricity consumed by ICT facilities today equates to approximately 120 billion kWh annually. Of that amount, data centers account for roughly 60 billion kWh, while large communications centers and network trunk lines account for another 20 billion kWh. Local equipment, private exchanges, and mobile phone towers use the remaining 40 billion kWh. The burgeoning demand for ICT services is steadily increasing electricity demand by these facilities. According to some projections, data center energy use may double between 2006 and 2011, while strong demand for information processing, broadband, and mobile communications is rapidly outpacing efficiency improvements.

Most ICT energy use occurs in large data centers and telecommunications facilities that consume an average of about 80 billion kWh annually. Energy consumption patterns within these large ICT facilities vary considerably with the type of center (supercomputing to switching), the complexity of the backup power and cooling systems, and age of the facility infrastructure. A simplified distribution of overall facility energy use is shown in Figure 1. As depicted, rack-mounted IT equipment consumes approximately half of all energy used by these facilities.³ The sophisticated power conditioning and delivery systems, which supply “clean,” reliable power to the IT equipment through the use of backup power devices, system redundancies, and multiple power conversion steps, consume approximately 25 percent of all facility power when actual system loading is considered. All of this power use generates heat loads that can be 10 to 100 times greater than those of typical commercial buildings on a square foot basis. Cooling systems therefore use about 25 percent of facility energy use to move large volumes of chilled air to the hardware to maintain acceptable equipment temperatures (with today’s prevalent cooling systems).

Figure 1. Equipment and software may account for half of the energy used by ICT facilities



The balance of ICT energy consumption (40 billion kWh) is in small installations such as mobile phone towers, server closets, and neighborhood facilities. These smaller facilities use about the same amount of power as a single-family home (20 to 200 amp service). The IT equipment consumes approximately 75 percent of all energy, as power systems are simpler and have fewer redundancies than larger facilities. Low heat loads and hardened IT equipment typically eliminate the need for additional cooling equipment. Yet major technology advances that benefit the larger ICT facilities are likely to yield power efficiency benefits in many of these smaller installations as well.

³ Typically, all power consumed by a server is considered as equipment power use, yet up to 25 percent of that power may be used for internal equipment cooling (fans) and power functions. Although a comprehensive energy analysis would allocate a portion of this energy use to these other functions, the simpler division of power use is used here because the proposed hardware technologies may radically alter both the internal cooling and internal power needs of the equipment.

2. Technology R&D Challenges and Solutions

Efforts to reduce ICT energy demand are appropriately focused on the larger ICT facilities that today represent a major engine and critical infrastructure for the entire U.S. economy and standard of living. Based on the main areas of energy use or loss within these facilities, energy-efficiency improvements naturally center on three areas: (a) the equipment and software, (b) power supply chain, and (c) cooling. In many respects, these areas are highly interdependent. For example, as shown in Figure 2, the power is initially provided to the equipment and software so that those components can perform the real computing work of these centers, and the heat given off by these components generates the need for cooling (as all energy used is converted into heat).

Ultimately, major energy efficiency gains that may be achieved in the equipment and software will decrease the demand for power, and therefore also the demand for cooling. Another promising solution is the development of equipment that can withstand a wider range of temperatures, humidity levels, and other environmental conditions. Success in this area would obviate research investments to improve the efficiency of facility cooling technologies.

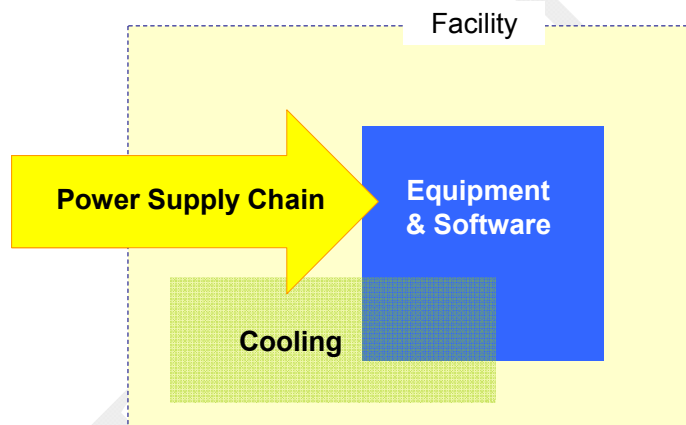


Figure 2. Equipment and software drive the need for power and cooling in ICT centers

The ICT industry and many of the industry’s leading organizations are actively engaged in a range of efforts to foster wider adoption of the diverse technologies and practices that are available today to reduce energy use in ICT facilities. While these activities are critically important and necessary, most industry observers understand the growing need for transformational technology innovations that will radically alter the energy outlook for these essential facilities and components. In the current economic climate, an aggressive, targeted, public-private program of collaborative research and development (R&D) offers the potential to effectively accelerate delivery of these sophisticated breakthroughs in the required timeframe.

Sixty forward-thinking representatives of the ICT industry stepped forward to create industry visions and paths forward at the “Vision and Roadmapping Workshop: Routing Telecomm and Data Centers Toward Efficient Energy Use” in October 2008. The success of this effort is due to the active involvement by facility owners and operators, software specialists, equipment manufacturers, suppliers, other experts, university researchers, and scientists from the national laboratories. For each of the three energy areas (noted above), they developed a clear vision, identified key challenges, and proposed potential solutions. This document highlights a number of the promising solutions or technology opportunities proposed.

Figure 3 summarizes the key workshop results, which reflect industry input on future directions for ICT research and development by collaborative public-private partnerships.

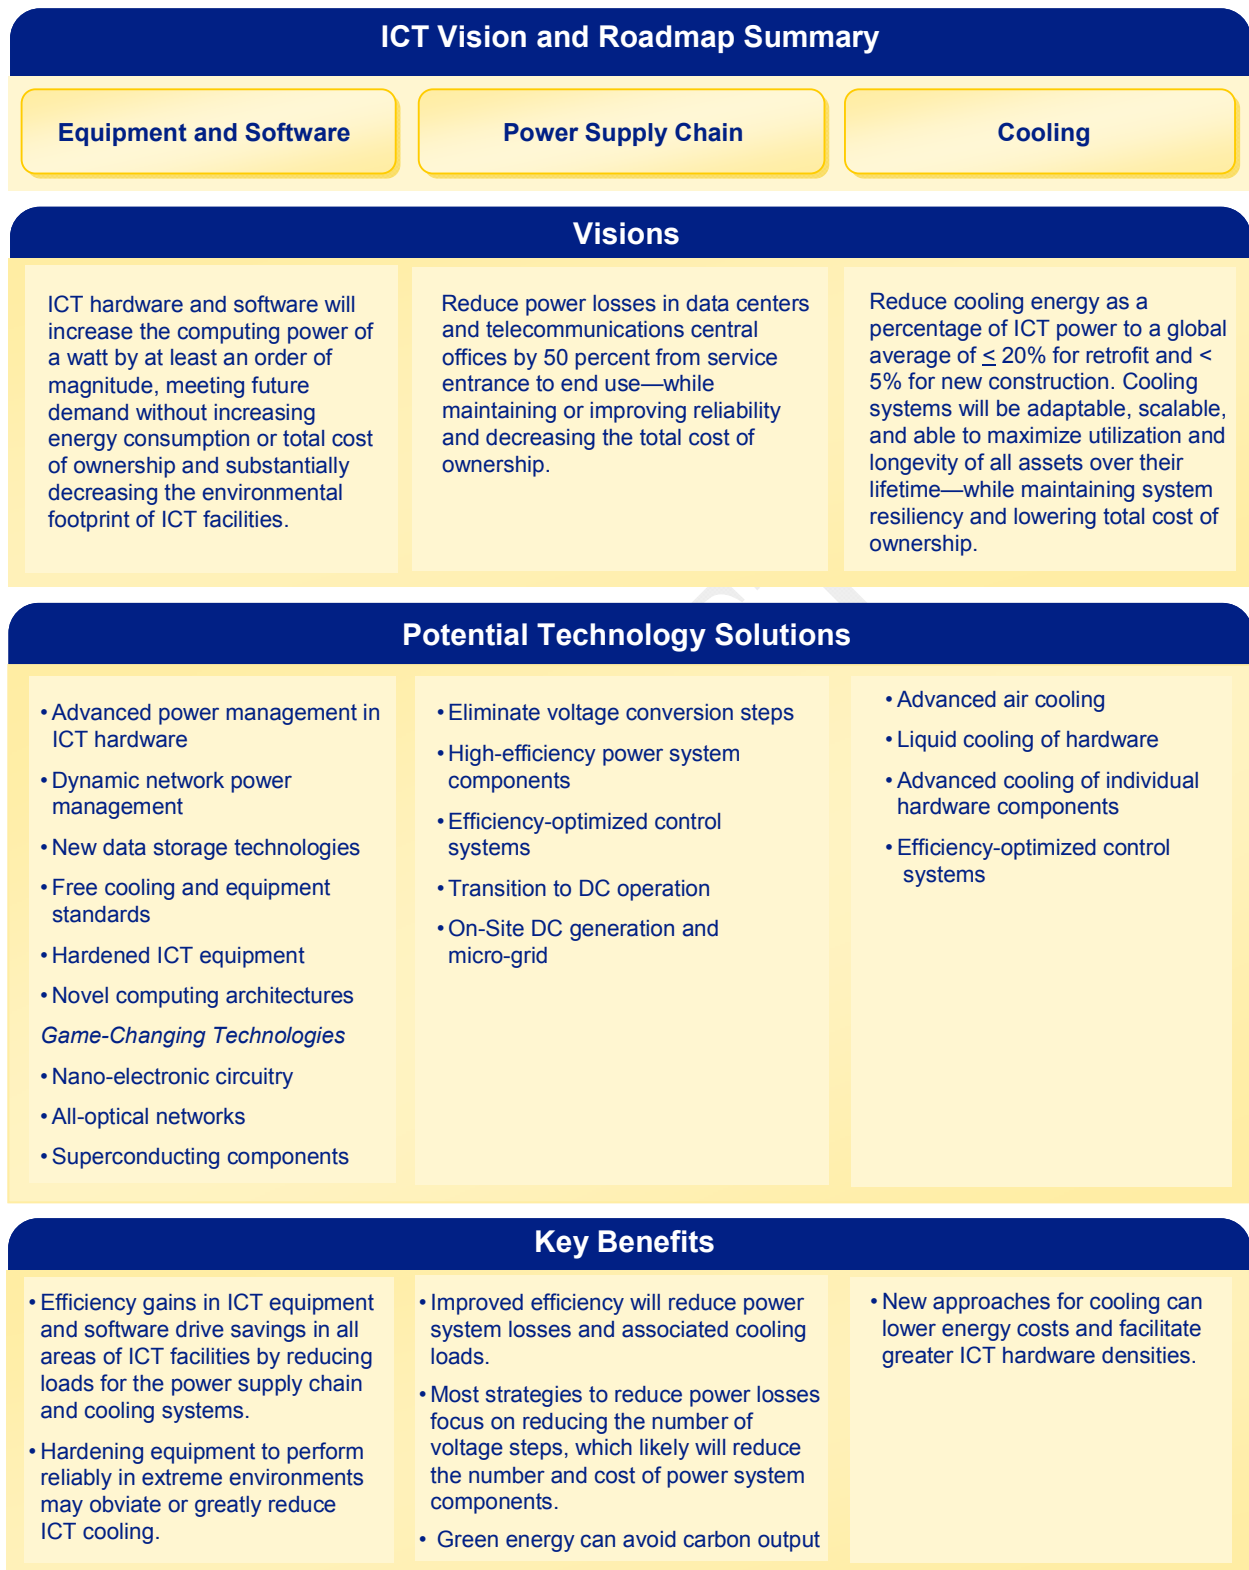


Figure 3. Summary of Key Workshop Results

The technology opportunities identified at the workshop are a small but generally representative sampling of the best potential avenues for R&D. For each of these energy-saving technology opportunities, Table 1 indicates the estimated potential energy savings in three ways: (1) percentage savings within the area of consumption, (2) percentage savings in the overall facility, and (3) the potential energy savings at the national level in billions of kilowatt hours. Note that values shown in the National Impact column reflect potential impacts of technology advances in data centers, whereas the estimated impacts for telecommunications centers are less certain. The potential savings estimates are independent and use current technology as a basis. While some opportunities can be combined, the savings are not necessarily additive.

Table 1. Summary of Proposed ICT Technology Solutions and Estimates of Potential Savings

Area of Energy Consumption		Opportunity	Area Savings	Facility Savings	National Impact (in billion kWh)
			(% of IT energy)	(% of facility total energy)	(2006 baseline: 60 billion kWh)
Equipment and Software	50% of Energy Use	Chip-level and Automated Power Management	50%	50%	30
		Dynamic Energy Consumption Management in Network Devices	40%	40%	24
		Data Storage Technologies	10%	10%	6
		Free Cooling & Equipment Standards	0%	15%	9
		Hardened ICT Equipment	0%	25%	15
		Novel Computing Architectures	50%	50%	30
		Nano-electronic Circuitry	75%	75%	45
		All-Optical Networks	75%	75%	45
		Superconductive Components	50%	50%	30
			(% of power or cooling energy)	(% of facility total energy)	(2006 baseline: 60 billion kWh)
Power Supply Chain	25% of energy use	Eliminate Conversion Steps/ Losses	20%	5%	3
		High-Efficiency Power Conversion Circuits	20%	5%	3
		Efficiency-Optimized Control Systems	16%	4%	2.4
		Transition to DC Operation	40%	10%	6
		On-Site DC Generation & Micro-Grid	15%	15%	9
Cooling	25% of energy use	Liquid Cooling of Components	70%	17%	10
		Advanced Cooling at Chip/Server-Level	20%	5%	3
		Energy-Optimized Control Systems	30%	30%	18

The remainder of this document is divided into three sections, each of which addresses one of the three main areas of ICT energy use shown above. Each section provides an overview of the area, the industry's vision for the future of that area, and brief descriptions of key technology opportunities to reduce energy demand in the area. Workshop participants and other members of the ICT industry with specialized knowledge in the technology topics graciously prepared these technology descriptions. Most of the technology opportunities were identified during the workshop, as indicated in the results of the workshop sessions presented in the Appendix. In accordance with the wishes of those who attended the workshop, additional experts who were unable to attend the event were invited to submit additional technology opportunities.

DRAFT

A. Equipment and Software

Despite ongoing efficiency improvements in ICT hardware and software, power demand continues to rise in these core functioning components of a data or telecommunications center. Growing demand for ICT services is forcing processor manufacturers to search for new ways to balance processor performance, heat dissipation, and power consumption. In the new generation of servers, server power consumption has been shown to increase linearly with utilization of the central processing unit (CPU). Achieving high levels of energy performance will require novel approaches to the design and management of these hardware and software systems. Reducing ICT hardware power consumption will mean that less heat will need to be removed by the cooling system.

INDUSTRY VISION STATEMENT

In the next 15 yrs, ICT hardware and software will increase the computing power of a watt by at least an order of magnitude, substantially decrease the environmental footprint of ICT facilities, and meet future demand without increasing energy consumption. Improvements in ICT computing efficiency will actively enable productivity and efficiency improvements throughout the American economy.



Figure 4. Improving the energy efficiency of computing hardware will affect energy use throughout ICT centers.

Chip-level and Automated Power Management

Current Issues or Challenges

Power consumption within a processor core is directly proportional to the voltage and clock frequency. The microprocessor manufacturers have started including advanced power management options that adaptively scale the core voltage or the clock frequency during run-time (examples include Intel's SpeedStep Technology and AMD's Cool'n'Quiet). Clock frequency refers to a processor's computing speed in terms of cycles per second. One way to increase a processor's performance is to increase the clock frequency, but this also increases heat generation. Higher clock frequencies and voltages thus result in greater heat dissipation from the chip, requiring better heat removal techniques. Techniques for efficient chip-level power management can reduce the computing power budget and HVAC power budget.

Clocking in electronics has been a major bottleneck to the continued improvement of electronic systems. In computing systems, clocking is a major source of power consumption and a principal obstacle to the improved processor performance. Clock and clock distribution technology have not changed significantly in several decades and still contribute to significant part of the energy consumed. In current microprocessors the power required for distributing the clock across the chip could be as high as 30% of the total power consumed by the chip.

Key Idea/Approach to Solution

Research in advanced clock distribution is in the early stages. Widespread deployment requires resources for developing novel design techniques (near term), industry standard design tools, and fabrication techniques (long term). This research will require about three years of effort for staged development, testing, and integration of all components.

Expected Benefits/Outcomes

Advanced clock and clock distribution techniques have the potential to save 30-75 percent of the energy consumed in various chips used in the computing systems (e.g., storage networks, optical networks, peripheral communication chips etc.). These energy savings will also reduce the energy needed for heat removal, therefore significantly lowering the energy consumption of the data center.



Figure 5. Distributed clocking technology can reduce heat generation in microprocessor chips.

Dynamic Energy Consumption Management in Network Devices

Current Issues or Challenges

Virtually every data communication transaction—from credit card purchase to sending an e-mail—generates data packets traversing data links with the help of network devices (routers, switches, firewalls, etc.). These networks are typically designed to serve humans—who don't like to wait—so every network device must be “always ready” for full and reliable operation. This explains why computer networks are planned and built for peak usage and why average utilization levels are typically very low. The peak usage profile typically defines the energy consumption pattern because energy consumption by most network devices varies little in response to the load level. Strict rules governing real-time operation do not allow these devices to enter low-power modes in the way many personal computers do to save energy. The percentage of energy consumed by network devices is rising steadily worldwide, placing greater importance on dynamic energy consumption management. To date, however, no commercial technology allows high-speed networking devices to consume energy in a manner proportionate to their instantaneous load.

Key Idea/Approach to Solution

Time pattern analysis is a critical first step in dissecting and solving this problem. Dramatically different solutions are likely to be needed to address idle timeouts at different scales—ranging from milliseconds to minutes. Longer idle periods may be expected to require software development, through which various system components could be slowed down or disabled and then enabled dynamically, within the window of the real-time tolerance. Shorter idle periods would require novel approaches to the embedded hardware and software structures so that they would be able to detect, utilize, and propagate knowledge on “windows of opportunity” for entering the low-power states. This research, development, and deployment (RD&D) activity would require close cooperation with universities and national laboratories and would take about three years.

Expected Benefits/Outcomes

Development of this technology is likely to generate a significant amount of intellectual property and should offer a significant competitive advantage to the economies capable of producing silicon designs of such complexity. Usage pattern analysis from routers and switches assumes no less than 40 percent of today's energy consumption could be saved in the aggregation and no less than 80 percent in the access layers. Other network types offer less dramatic energy savings potential, yet all topologies can benefit from the new technology. Other benefits would include increased density, higher speed and lower cost of the network devices. The technology could also spark another wave of system miniaturization with embedded network agents.

Data Storage Technologies

Current Issues or Challenges

Large data centers can contain multiple, heterogeneous storage technologies distributed over the data center network. These storage systems are accessed by a distributed network of processing nodes that service data requests. In complex data processing applications, data requests can require significant processing time due to the complexity of the requests (e.g., database SQL queries). Energy use by the data center directly depends on the amount of time the storage and processing nodes actively engage in processing the queries. A number of solutions to this challenge are possible. Three are presented here.

1. Data Layout Optimization

Key Idea/Approach to Solution

Significant reduction in the processing time can be achieved if the storage is optimized to improve locality with computation. Since both the data as well as the requests are constantly changing, it becomes necessary to periodically and dynamically re-evaluate the optimality of data layout. A re-layout and data movement can help reduce unnecessary delays due to network congestion stemming from suboptimal data-task assignment. Research is needed to study this data layout optimization problem and arrive at the best load-balancing algorithms that minimize the total data processing work performed in the data center.

Expected Benefits/Outcomes

Optimized layout can reduce processing times by as much as 50 percent (by improving the speed of request service by a factor of two), enabling energy savings if the systems are then switched to lower power mode as soon as the request has been serviced. Some of the existing research in locality enhancement algorithms, load-balancing methods, and distributed database optimization can be leveraged as a starting point, but the research on these topics needs to be revisited in the context of the massive scale of modern data centers. Research and development for dynamic optimization of data organization for energy use reduction will take at least three years.

2. Carbon Nanotube-based Memory

Key Idea/Approach to Solution

Carbon nanotube (CNT) based memories have been successfully demonstrated and are scheduled for space flight tests on the upcoming Hubble Telescope repair mission. These memories are unique in that the use of a CNT nano-electro-mechanical switch to hold the bit information does not require any standby electric power to retain the information. Additionally, CNT memories offer the potential to rapidly transition from the current 4Mbit memory size to 512Mb memories for space applications. However, the speed of CNT memories, the low power consumption, the ability to rapidly move to large memory sizes (512 Mb), and the potential low costs in high-volume manufacturing make them candidate replacements for commercial flash memory applications.

Expected Benefits/Outcomes

The time and technology progress for CNT-based memories are such that an investment in establishing a domestic CNT-memory manufacturing capability would provide the U.S. a dramatic and profitable lead in the highly competitive memory markets. Additionally, a commercial “pull” for this memory would rapidly advance this technology and enable heretofore inaccessible capabilities for strategic satellites.

R&D to create a successful, low-cost, high-volume CNT memory manufacturing capability will provide the U.S. a strategic lead for memory applications, such as flash memory, on-board data processing, low-power solid state recorders, and lower power consumption in servers.

Establishing a manufacturing capability should create several thousand new U.S. jobs. The work is expected to require significant public-private investment over three years.

3. Holographic Data Storage

Key Idea/Approach to Solution

Holographic technologies offer improvements in the performance and cost curves of storage thereby improving the ability to make massive amounts of data accessible to users while reducing the total costs incurred to store their exponentially growing quantity of data. Holographic storage performance is demonstrated at data density of 515 Gb/in², random access (~250 ms), and very high transfer rates capable of exceeding 120 MB/s. The next technology node of 1.6TB per platter (DVD size) requires investment in three key technologies a micro electro-mechanical-system (MEMS) scanner, spatial light modulator (SLM), and photo-diode camera system.

Expected Benefits/Outcomes

Moving this technology to the next generation of capacity will provide the U.S. the ability to regain worldwide leadership in optical disk technologies for uses such as,

- Ultra-low cost, low-power, and long-term data storage, as needed to meet National Archiving and Records Administration requirements for health care, financial transactions, etc.
- HD & 3-D movie and television recording/playback; software and game content distribution
- Personal recording and archival storage

Hardened ICT Equipment

Current Issues or Challenges

Currently, about one-third to half of the energy in a typical data center is used to run air conditioners and other cooling equipment—instead of performing useful IT work. Making the equipment less sensitive to temperature will require rethinking server circuitry. Current silicon technology is approaching a performance plateau; enhancements on the horizon are low impact, and costs are becoming prohibitively high. At the same time, next-generation and emerging applications are continuing to demand substantial improvements in power conversion efficiency. To meet these new requirements and challenges, novel materials and transistor structures are needed.

Key Idea/Approach to Solution

Developing ICT equipment that can function reliably over all temperature and humidity ranges would allow more data centers to operate with unconditioned outdoor air, even near the equator. Military applications already use shock-resistant and environmentally hardened equipment. Non-classified military and aerospace (NASA) standards may help guide development of robust electronic devices capable of operating in truly hostile environments (extreme heat, extreme cold, high/low humidity, etc.). Novel materials will be needed.

GaN-based power devices are expected to offer significant improvement over currently available options. The use of a GaN device provides many advantages for the user, including reduced losses, elimination of active components, increased efficiency, and improved temperature performance. Gallium nitride (GaN) high-power modules can be implemented in available packaging schemes. The reduction in losses in GaN devices can be applied in a number of ways to optimize circuit design by increasing efficiency, reducing heat sink requirements, and reducing the power consumption of the transistor.

Developing circuit design methodologies on these novel, high-temperature materials can significantly reduce the need for heat removal. Advanced device research is needed to understand the performance limitations of components made with these materials. Circuit design and fabrication techniques should be enhanced to include design rules and techniques for using these materials.

Expected Benefits/Outcomes

If hardened IT equipment can reduce consumption by 25 percent, it would save \$7.5 billion in avoided power plant construction over 10 years, and all associated permitting and lifecycle greenhouse gas emissions. The Federal government has a clear role in sponsoring this research into the feasibility of hardening, as individual ICT manufacturers would gain no direct cost benefit. Without feasibility demonstration to users and industry, adoption of hardening is highly unlikely. The reward-to-risk balance is highly favorable. The underlying materials science, which relates reliability with temperature and humidity, has not been updated for consumer-grade electronics for a long time. Research in circuit and integrated circuit design using novel materials is at a critical stage. Although fabrication of high-density chips is at least 5 to 10 years away, a critical need exists to develop circuit design techniques for these materials. The research will broadly impact communication circuits, industrial wireless sensors, and data center equipment. These chips can significantly increase the operating temperature of the servers and reduce the cooling budget. Primary research activities would include:

- Study existing materials for developing temperature and humidity hardened devices
- Develop automated circuit design tools and techniques using these novel materials
- Develop packaging and fabrication techniques for low-cost mass production and utilization of integrated circuits using the novel materials

Novel Computing Architectures

Current Issues or Challenges

In accordance with Moore's law, the number of transistors packed on each computing chip continues to double approximately every two years, yet the efficiency of CMOS logic on silicon is no longer improving at that rate. This power density issue has become the dominant constraint in design of new processing elements and ultimately limits clock-frequency growth for future microprocessors. Under these conditions, the cost of power will exceed the procurement costs of such systems, and will ultimately limit the practicality of future state-of-the-art computing platforms—from mobile phones to the largest data centers. These trends are leading to a crisis for large-scale data centers in the not-too-distant future unless power-efficient solutions are developed.

Approach to Solution

New approaches to software writing and organization of computing functions offer many opportunities for improving computing efficiency. The development of applications-specific machines is one approach that is described below. Large-scale data centers and supercomputing facilities rely on general-purpose machines, able to handle many types of algorithms and applications—some better than others. As a result, the efficiency at which many applications use large parallel systems is often rather poor. Until now, designing machines targeted to specific applications has not been practical, except for the few companies that use very large quantities of homogeneous hardware devices. The few attempts to do so were rapidly overtaken by clock frequency improvements in the commodity markets. However, the rapid pace of clock-frequency scaling has hit a plateau, power is now the leading design constraint, and computational efficiency again plays a central role in computer architecture design. In addition, the vibrant market for power-efficient processors for mobile electronic devices has yielded new tools for chip designers, enabling more rapid and cost-effective turn-around for power-efficient semi-custom designs. With power now a leading design constraint and marketing advantage for ICT equipment, the energy-efficient embedded processors that enable these modern conveniences could be an attractive technology for large-scale ICT facilities.

This approach differs from past forays into semi-custom computer system design by leveraging mainstream tools, commodity IP, and design processes already successful in the consumer electronics industry. Data centers and high-performance computing applications can tap into this approach if they can rigorously define their requirements. Since application codes vary greatly, optimal machine characteristics can be expected to vary widely from one code to another. However, the flexibility of the embedded processor approach converts this from problem to opportunity. Different exascale machines could be affordably designed and manufactured for each discipline requiring this kind of compute power. The underlying components of such a machine would be largely the same, but organized in a manner optimized for each application.

Demonstrating such technology and the software infrastructure to make it useful will require sustained investment over a period of three years, but some of the groundwork has already been laid by existing research in this area [Green Flash www.lbl.gov/CS/html/greenflash.html]. Commercialization of the technology will require public-private partnerships with industry to create point-designs for large-scale systems.

Expected Benefits/Outcomes

Embedded technology has routinely demonstrated 50x power efficiency benefit over conventional desktop processing chips for delivered performance on targeted applications [see Rowen: www.tensilica.com/news_events/presentations.htm]. Even higher power-efficiency benefits can be realized with more attention to tailoring the design for the application. Rather than wringing 20 percent better efficiency out of datacenters with incremental improvements, we can enable 20x improvements by tailoring datacenter designs to their workload. Such an approach would transform the industry.

Nano-electronic Circuitry

Current Issues or Challenges

The practice of continually improving computer chips in accordance with the well-known “Moore’s Law,” is confronting hard limits. For the first time in decades, advances in computing technology are now threatened because, while transistor density has continued to increase, the energy efficiency of silicon has not. Scaling down of CMOS transistors is already generating more heat than we have the ability to remove. Ironically, as these transistors shrink in size, chips become harder to cool as more and more are integrated onto chips for increased capability.

Development of an effective “post-CMOS” replacement switch is an industry research priority. The CMOS transistor switch uses a charge potential or “thermionic” barrier to turn on and off the flow of current. Alone, it uses a small amount of energy; but on chips of perhaps 2 billion transistors switching on and off at the rate of 3-5 billion cycles per second, power quickly becomes a big problem. In addition, the interconnections between these devices dissipate power and limit performance.

Approach to Solution

Nanoscale devices may have the ability to push back the power barrier. A number of proposed new switches use non-thermionic barriers to gate logic, mitigating an intrinsic source of power dissipation. Some of these structures do not shuttle electron-based tokens between devices to retire transactions – rather, proposed new state variables include electron spin, magnetic field, excitons, molecular position or even “qubits,” rather than charge. Current technologies spend as much energy moving information as creating it; these new tokens have the potential to dramatically improve communication speed and power. Nanophotonics, nanoelectronics, nanomagnetics, nanoelectromechanical systems (NEMS), and nanostructured materials are promising building blocks for energy-efficient, scalable, and miniature computing systems of the future. These nanoscale devices can form processors, memory, nano-circuits (interconnects), and potentially realize a future computing system with unprecedented energy-efficiency, scalability, and miniaturization that would fundamentally change the electronics industry and re-enable the performance scaling that society has come to expect. Molecular nanoelectronics and nanophotonics offer a new opportunity to incorporate very high-performance nano-FPGA (field-programmable gate arrays) or optical logics with unprecedented programmability and performance.

Development and deployment of micromechanical and nano-electronic devices are high-risk, high-reward technologies. Because they are so far from current design practices, the lead-time for maturing such technologies is 5-8 years and requires sustained research investment. The benefits from these investments will become increasingly important as we move beyond CMOS technology.

Expected Benefits/Outcomes

Nanotechnology advances promise numerous benefits to energy efficiency and programmability. First, efficient and parallel interconnections by nanophotonics and nanoelectronics can bring the computing systems to more balanced and optimized operation so that far more computing and data processing tasks can be completed for a given amount of energy consumption. Nano-optical structures provide ultra-high throughput, minimal access latencies, and low power dissipation that remains independent of capacity and distance. Multi-wavelength operation can bring massive parallelism to the computing system. Massively parallel nanoelectronic interconnection offers low-power, short-distance interconnection between many cores. The new generation of memory and data storage technologies based on nanomagnetics, nanoelectronics, and NEMS offer extremely high-density, low-power, and high-bandwidth access such that highly integrated large-scale data storage is possible at high-throughput.

All-Optical Networks

Current Issues or Challenges

With trends clearly toward further multiplication of the on-chip processing cores, chip multiprocessors (CMPs) have begun to resemble highly parallel computing systems integrated onto a single chip. In this context, the role of the interconnect and associated global communication infrastructure is becoming central to chip performance. As with highly parallel systems, performance is increasingly tied to how efficiently information is exchanged and how well the growing number of computational resources are used. The realization of a scalable, on-chip communication infrastructure faces critical challenges in meeting the large bandwidth capacities and stringent latency requirements demanded by CMPs in a power-efficient fashion. With vastly increasing on-chip and off-chip communication bandwidths, the interconnect power consumption is an acute and growing problem. Conventional CMOS scaling of electronic interconnects and networks-on-chip (NoCs) do not appear capable of satisfying future bandwidths and latency requirements within the CMP power budget.

Key Idea/Solution

Technologies like silicon nanophotonics or gallium nitride (GaN) optical circuits could lift the limitations of conventional electronic networks by using optics to deliver much higher bandwidth within the same power budget and reduce heat removal requirements. The insertion of *photonics* in the on-chip global interconnect structures for CMP can potentially leverage their unique advantages in optical communication and capitalize on their capacity, transparency, and low energy consumption. Construction of photonic NoC could deliver performance-per-watt scaling that would be impossible with all-electronic inter-connects. Unlike prior photonic technologies, nano-scale silicon photonics offer the possibility of creating highly integrated photonic platforms for generating and receiving optical signals with superior power efficiencies. These tremendous gains in power efficiencies for optical modulators and receivers are driven by the nano-scale device footprints and corresponding capacitances, as well as by the tight proximity of electronic drivers enabled by the monolithic CMOS platform integration.

Evaluating the tradeoffs between electrical and photonic network designs will require extensive cycle-accurate simulations using custom software that is calibrated against prototype silicon devices. It will require development of a comprehensive, event-driven simulation to model the low-level electronic and photonic details of the evaluated interconnect configurations. As an early application, photonic ring resonators could be used to modulate laser light for existing Infiniband and Ethernet implementations. This would yield a ten-fold improvement in power and another eight-fold improvement in bandwidth because of the faster modulation rate.

Research is needed to explore innovative, on-chip micro-architectures for massively parallel chip multiprocessors that leverage nano-scale silicon photonics, GaN-based optical interconnects, and complementary devices developed in synergy with standard CMOS electronics. It may take up to five years just to produce test devices that can directly assess the manufacturability of such products. The integration of electrical and optical devices along with the passive devices and optical fibers on the same substrate will provide efficient, compact modules that will not only increase the power efficiency of the data centers, but will provide substantial area savings for the servers.

Expected Benefits/Outcomes

Photonic NoCs can deliver a dramatic reduction in power expended on intra-chip global communications while satisfying the high bandwidths requirements of CMPs. Photonic NoCs change the rules of power scaling: as a result of low loss optical waveguides, once a photonic path is established, the data is transmitted end-to-end without the need for repeating, regeneration, or buffering. Power consumption by optical switching elements is independent of the bit rate, so, once generated, high-bandwidth messages do not consume additional dynamic power when routed.

Superconductive Components

Current Issues and Challenges

The advent of higher operating frequencies in high-performance CMOS processors has increased dynamic power consumption dramatically. Static power numbers have not fared better, with ballooning consumption due to sub-threshold, junction and gate quantum leakage effects introduced by small circuit feature sizes ranging below 100 nm. Emerging semiconductor technologies/processes as well as architecture designs have tried to stem the adverse power consumption requirements for processors that are quickly approaching multi-billion transistors on a single die. Nevertheless, barring an unforeseen, disruptive technological breakthrough, the “thermal design power (TDP)”⁴ for this new class of high-performance microprocessor is expected to approach 200W in the near future.

Key Idea/Approach to Solution

Instead of squeezing incrementally better performance-power ratios out of CMOS technology, superconductive (SC) components hold the promise of better performance-power ratios. Improved ratios are achieved by avoiding heat dissipation as the SC circuits switch at temperatures below the critical temperature, enabling “reversible computing”—if cooling inefficiencies are disregarded. Cooling is provided by liquid nitrogen. This system would require SC components with higher critical temperatures and some energy would be needed to maintain the equipment in a liquid nitrogen environment, but the overall energy demand is anticipated to be lower than for traditional circuit technology.

SC technology operates at the quantum level. Whereas quantum effects in CMOS in the form of tunneling currents are highly undesirable, the binary states encoded as magnetic flux quanta in SC are valuable. Rapid Single Flux Quantum (RSFQ) and Quantum Flux Parametron (QFP) technologies today overcome switching shortfalls (hysteresis) by better controlling the operation point (current biasing). Architecture and programming remain a challenge because the technology imposes bit-sequential processing, albeit at frequencies of multiple-100 GHz. A multi pronged approach is needed to advance SC processing:

- Explore suitable SC components for microprocessor technology applying high critical temperature components if industrially feasible;
- Push programs to accelerate SC architecture and programming development;
- Evaluate the impact and operations of state-of-the-art cryo-cooling in the ICT center.

This research, development and deployment (RD&D) activity requires at least a four-year window that can begin immediately.

SC technology does not require <0.1 micron semiconductor technologies/processes to provide high processing performance, although SC would doubtless benefit from more advanced, high-density semiconductor fabrication processes. In fact, SC could reuse “low performance” fabrication plants usually targeted for low-performance processing (e.g., Ethernet, USB) that are one or two generations past state-of-the-art.

Expected Benefits/Outcomes

Initial RSFQ prototype work for “Hybrid Technology Multi Threading” (HTMT) and others show great promise. A major thrust in high-temperature superconductivity, architecture, and software research could make this technology viable in the near future for large ITC centers with extreme cooling capability.

⁴ Note: TDP is determined by the manufacturer based on a typical work mix in conjunction with a desired thermal management solution. Maximal power consumption data is usually withheld by the manufacturer and can substantially exceed TDP.

B. Power Supply Chain

Data and telecommunications centers require large quantities of electricity to be conditioned, converted, and delivered to the diverse components, including servers, switches, routers, and hard drives. The power supply chain can include electricity purchased from the grid, backup power, onsite-power generation, switchgear, UPSs, power distribution systems, rack-level and unit-level power supplies, and power management technology. Traditionally, telecommunications centers have used direct current (DC) power, and data centers have used alternating current (AC) power distribution systems.

Although modern data centers typically distribute electric power in the form of AC, the end loads require DC. As a result, the power chain requires some power conversion stages, from AC to DC. Because of reliability concerns and varying voltage level standards, most power distribution systems in data centers employ multiple conversion steps, as shown in Figure 6. The most widely used power conditioning equipment is the double-conversion uninterruptible power supply (UPS), which converts AC to DC, where some DC energy storage element can be connected, followed by conversion back to AC. (Double conversion UPS requires all the power to flow through the two conversion steps, all the time). A step-down transformer usually takes the voltage level from 480V AC to 208V AC. The end loads (ICT equipment) are powered with AC, and immediately convert this to DC inside their own power supplies. This DC voltage is then converted (DC to DC) to lower voltages as appropriate for use with the computer chips found on the motherboard. At a minimum, this equates to five conversion steps, but could be even more, depending on the voltage levels chosen inside the server. Each of these conversions introduces an attendant power loss, and the more steps used, the less efficient the overall power chain.

INDUSTRY VISION STATEMENT

In 15 years, reduce power losses in data centers and telecommunications central offices by 50 percent from service entrance to end use—while maintaining or improving reliability and decreasing the total cost of ownership.

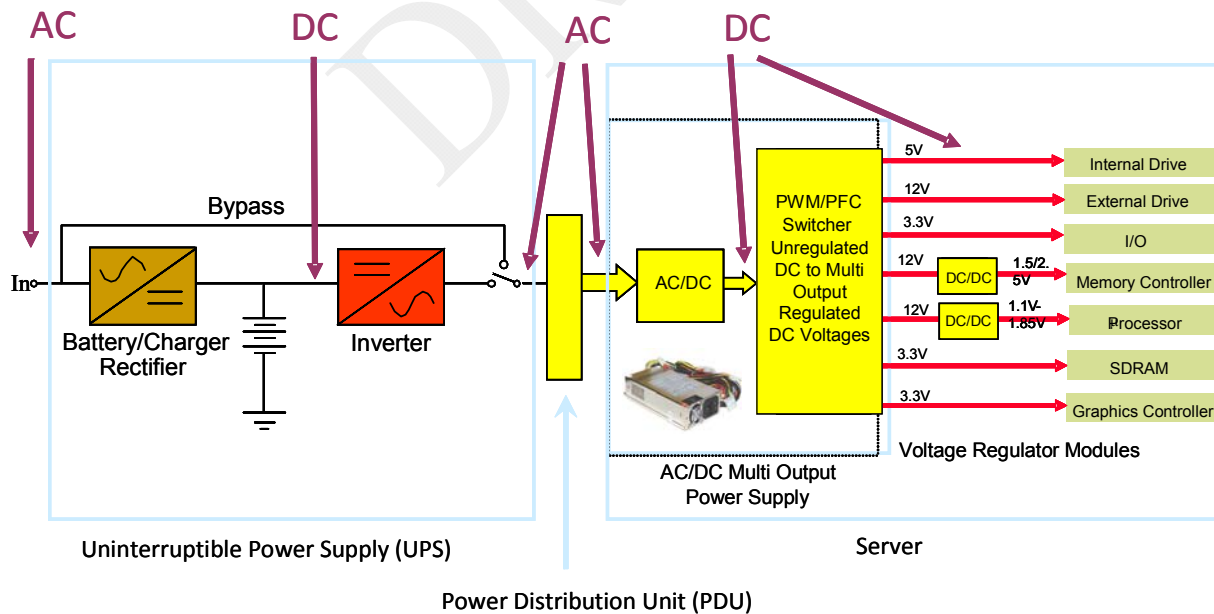


Figure 6. From utility power to the chip, every conversion step generates an inefficiency and produces heat that must be removed

THE CASE FOR ALTERNATIVE ENERGY

As electricity demand from ICT continues to rise, there is growing scientific consensus that global carbon emissions must drastically decrease, and that reducing the carbon intensity of electricity is a crucial piece of the puzzle. To achieve climate stabilization, scientists estimate that we must reduce carbon emissions by 25 to 40 percent by 2020. As the ICT sector continues to grow, it is imperative that we encourage clean and reliable supply solutions to complement ICT advances in efficiency.

For data centers, key challenges will include the following:

- Development of cost-effective, regionally compatible sources of onsite generation for power for data centers (including but not limited to solar, wind, tidal and wave power).
- Research and development of cost-effective and more reliable/safe storage technology (including improvement of batteries, flywheels and fuel cells) to enable use of intermittent energy sources such as wind and solar.
- Increasing both real and perceived reliability of intermittent alternative energy sources for onsite generation and as an expanded part of the grid, via storage and other technologies.
- Continued development and implementation of robust and cost-effective CCHP technologies. This requires both better technology, and a better understanding of the economics of combining power and cooling on the part of data center design and operation. This may also include collaborative use of municipal/industrial uses of power to optimize efficient use of waste heat (e.g., for wastewater treatment).

Key opportunities will include:

- New business opportunities presented by the use of ICT technology to both expand the use of intermittent clean power, and more efficiently utilize existing clean and alternative energy capacity (e.g., real time communication for dispatch of hydropower and wind).
- Facilitating efficiency and use of alternative energy in developing countries, in order to foster both climate stabilization and continued level of global competitiveness for the US-based ICT industry.
- Most renewable energy sources generate DC electricity, which is the form used by ICT equipment and power storage devices. Renewable sources are thus likely to integrate well with DC power distribution in ICT facilities.

Eliminate Conversion Steps/Losses

Current Issues or Challenges

Typical power distribution systems can have losses as high as 40 percent just from the conversion steps. In light of these facts, it makes sense to re-evaluate the power distribution system and consider methods to eliminate conversion steps or reduce their losses.

Key Idea/Approach to Solution

Several approaches can be taken to achieve this goal. These approaches fall into two main categories: AC distribution or DC distribution.

DC Distribution: The simplest way to eliminate conversion steps is to do just that – by distributing DC throughout the building, so that only one AC-to-DC conversion is necessary. If the voltage level is carefully chosen, the number of DC-to-DC conversions can be minimized. The main barriers to DC distribution seem to be that the majority of installed base in the industry is the more traditional AC architecture, and facility managers are reluctant to try new approaches. Transition to DC would require capital investment for retrofits but could save capital investment for new construction. This would represent a paradigm shift for designers and facility managers.

AC Distribution: With this scheme, the approach must be to reduce conversion losses, since some conversions will be necessary. Depending on the design approach, some designers use 400V AC with no step-down transformer and supply the load directly with this voltage at the racks. The UPS losses, however, are still there. Efforts to improve individual component efficiency are ongoing, with improvements in UPS and transformer losses reported by vendors. New semiconductor materials and novel electronics design approaches hold much promise, however, third-party measurement and verification of any approach is advisable.

A highly effective measure to increase UPS efficiency is to use a line-interactive topology instead of double conversion. Line interactive UPSs supply the load directly from the source, and the converters only come into play during a power disturbance, whereas the double conversion type require power flow through the converters at all times. These disturbances occur infrequently, so the line interactive UPS is highly efficient most of the time (97-99 percent). Unfortunately, the data center industry is highly risk-averse; any potential for unplanned shutdowns is avoided. Test results from typical servers will be needed to help shift mindsets and enable widespread use of line interactive UPSs.

This research, development and deployment (RD&D) activity requires federal support supplemented with industry cost-share. National Laboratory support is also desired for analysis and demonstration assistance. The activity requires a three-year time window that can begin immediately.

Expected Benefits/Outcomes

Initial analysis in the DC architecture area has shown a potential for up to 7 percent input power savings over the best possible AC designs, along with increases in reliability and reduction in footprint. The majority of the installed base today, however, is not representative of the best possible AC design. Efficiency gains in voltage regulator module (VRM) topologies of 6 percent and PSU topologies of 15 percent are also possible, though initial costs are a significant development and adoption barrier.

The AC options may achieve lower savings, but could still be effective through wider adoption. Third-party, independent engineering analysis is needed to validate the data to educate industry stakeholders. Key needs include (a) Analysis, measurement and verification, (b) Education, and (c) Standardization of components, voltages, and system design.

High-Efficiency Power Conversion Circuits

Current Issues/Challenges

Microprocessors run on very low-voltage (even sub-volt) DC power, but the electric grid provides AC power, and battery backups are maintained at 48 volts (V) or higher to reduce distribution losses. Multiple conversion stages are necessary to accommodate these different parameters. Unfortunately, every conversion stage provides a double energy efficiency problem to the ICT industry and causes further inefficiencies. Every watt (W) of power that goes into an ICT system turns into heat that must be removed from the facility via air conditioning. Inefficient power conversion circuits require more power to operate the system and give off more heat. Inefficiencies along the power conversion chain multiply both effects.

Traditional power conversion circuit designs have reached their limits in efficiency and power density due to a combination of distribution bus losses and fundamental restrictions in topology performance (as processor voltages reach sub-volt levels). New architecture and converter technologies are needed to increase efficiency while optimizing processing real estate.

Approach to Solution

Solutions to this challenge must address the issues at two levels: (1) the overall conversion architecture from AC to point of load and (2) the topologies of each individual conversion step. From an architecture standpoint, three high-level approaches should be evaluated for development: various AC architectures, various rack-level DC architectures, and various facility-level DC architectures. Each of approaches presents multiple options, including different voltages and other considerations (such as the amount of copper needed). Focus should be placed on the power path to reduce conversion watts, eliminate redundant regulation stages, and utilize more efficient unregulated converters, where possible. The evaluation must also consider efficiency during light loads.

From a topology standpoint, higher computation speeds from higher clock frequencies will continue to increase power loss. To decrease these losses, the operating voltage should be reduced, but microprocessors operating at significantly lower voltages demand much higher currents, requiring voltage-regulating modules (VRM) to provide the desired regulated power supply. In addition, due to increased clock frequencies and higher currents, these microprocessors present high slew rates (maximum rate of change in a signal at any point in a circuit) during transient loads. To meet these demands, high power density, heavier load, and tight voltage regulating modules are needed to provide high efficiency over a wide range of loads with good transient response. Further upstream, traditional power supply units (PSU) that convert AC to bus-level DC (e.g., 12V) have operated at 75 percent efficiency, though EPA Energy Star and other programs have called for 80+ percent efficiencies, and over 90-percent efficient power supplies are available. This activity requires a two-year time window that can begin immediately.

Expected Benefits/Outcomes

Initial analysis in the architecture area has shown a potential for up to 7 percent input power savings. Efficiency gains in VRM topologies of 6 percent and PSU topologies of 15 percent are also considered possible, though initial costs have been a significant barrier to development and adoption to date.

Conversion Loss Example

A typical AC to 48 V DC to 1.x V system has an overall 67-percent efficiency from AC to point-of-load (POL),* even though individual conversion stages are 85-95 percent efficient. This means a 1kW load has to draw 1.5 kW from the AC grid, and the difference of 500W is lost as heat.

* "Thoughts on Server Metrics," C.Belady, Hewlett Packard, Enterprise Servers and Data Centers: Opportunities for Energy Efficiency, 2006.

Efficiency-Optimized Control Systems

Current Issues/Challenges

One route to greater energy efficiency in ICT is to optimize the efficiency of AC/DC rectifiers, which convert utility power to the form of electricity used by electronic devices (DC). Rectifiers are the standard power conversion technology used by the telecommunications industry and are increasingly common in data centers to take advantage of the enhanced reliability and energy efficiency opportunities afforded by DC.

Rectifier-based power systems in telecom and data center applications use multiple rectifiers for reliability purposes. Some facilities use an N+1 configuration, in which there is always at least one more rectifier in the system than is needed to carry the load and recharge the battery backup system at the same time. That leaves room for one rectifier to fail and the system will still be able to provide the needed power. Since battery re-charging is only needed after the end of a utility failure, most N+1 systems are running at closer to N+2. Other facilities use an A/B configuration, in which two separate sets of rectifiers (or rectifier systems) are available to power the load. If the A system fails, then all of the load can be picked up by system B.

In both types of facilities, standard system configurations have all rectifiers ON at all times. For reliability purposes, a standard rule of thumb for N+1 facilities is that even for N rectifiers, they are commonly operated at less than 80 percent of their capacity (at which point additional rectifiers would have already been added). For this and many other reasons, including intentional over-sizing and engineering safety factors, many systems (especially large systems) are operated at light loads. By design, rectifiers in an A/B facility must be running at less than 50 percent capacity, and most are running significantly lower than that. In other words, when all rectifiers are sharing load, each rectifier may only be operating at a fraction of its total capacity. Unfortunately, rectifier efficiency decreases at low utilization.

Key Idea/Approach to Solution

A new efficiency-optimized control technology is needed to both ascertain the peak rectifier operating points and to maintain operation at those points. This control system will focus on achieving the maximum efficiency of the entire DC power system while still maintaining required redundancy and security. The approach includes algorithmically “experimenting” to find the peak point, utilizing all available control features of the rectifiers.

This research, development, and deployment (RD&D) activity requires a public-private partnership. National Laboratory support is needed for analysis and demonstration assistance. The activity requires a nine-month time window that can begin immediately.

Expected Benefits/Outcomes

A conservative estimate of just 4 percent savings (increasing system efficiency from 85 to 89 percent), extrapolated against earlier estimates of energy consumption for the ICT industry, suggests a total savings opportunity of approximately 4.88 billion kWh per year, or the electricity consumed by nearly half a million U.S. households.

Transition to DC Operation

Current Issues/Challenges

Telecommunications facilities have traditionally operated very reliably with 48-volt DC power. Distribution of 48V DC power is not currently a good option for data centers because of their higher power needs, which would require very large conductors. Today's data centers use various electrical power distribution schemes, including primarily AC distribution backed up with batteries—which require DC. These schemes typically involve multiple conversions from AC to DC and from DC to AC at the various voltages required by devices. Other distribution schemes distribute DC power at various voltages. The industry organization Green Grid has described these power distribution schemes in a series of white papers.

A proof-of-concept demonstration of a DC powering scheme using commercially available equipment concluded that a DC distribution scheme using 380V DC could be safely deployed to lower capital costs while improving reliability. Such a scheme could also efficiently use renewable DC energy sources, such as fuel cells, photovoltaics, or wind power. Interest in this technology in Europe and Asia suggests a unique opportunity to form a global standard for the voltages and connections. Deployment of this technology is hampered by industry forces resistant to change.

Transitioning to DC power today is feasible, yet requires market pull and push. Many barriers must be overcome for this type of market transformation to occur. Progress will also require leadership in interfacing with codes and standards organizations.

Key Idea/Approach to Solution

Near-term approaches to foster a move to DC power include measuring and documenting the range of efficiency achieved with current electrical distribution equipment, including uninterruptible power supply (UPS) systems, power distribution units (PDUs), transformers, rectifiers, and IT equipment power supplies. This information can then be used to validate a range of system-level efficiencies. Work will also include demonstration of an energy-efficient DC system in an operating facility. Collaboration with European and Asian organizations implementing DC solutions will help to ensure a common approach. Agreement on appropriate connectors and distribution voltages could even enable development of a global standard.

In the longer term, end-to-end efficiency can be studied to optimize conversion voltages throughout the power delivery scheme. This study would be accomplished in collaboration with the Electric Power Research Institute (EPRI), the Power Sources Manufacturers Association (PSMA), Green Grid, IEEE, and other leading organizations. Further demonstrations would feature use of renewable on-site generation. The activity requires a three year time window that can begin immediately.

Expected Benefits/Outcomes

By taking advantage of all possible DC use, overall facility-level energy savings could be in the range of 10 percent or more. Additional benefits would include reduced first cost, improved reliability, potential worldwide adoption, improved power quality, and compatibility with renewable DC sources.

On-Site DC Generation and Micro-Grid

Current Issues/Challenges

A variety of distributed power generation technologies are in use today, including highly efficient combined heat and power (CHP) systems. Yet distributed generation (DG) technology is extremely difficult to apply in mission-critical environments that requires 24/7 computer-grade electricity. The challenges are numerous. When power is generated by a single, primary power unit, availability of power may vary more than when it comes from the grid (e.g., variable availability of wind or sunlight). In addition, using multiple units in parallel often presents synchronization and interdependence issues. Fuel cells, which might be used to provide energy, possess limited step load capability as well as poor fault clearing and transient overload capacity. Some of the biggest obstacles come from utility distribution systems and protective elements. In view of these challenges, many server-based facilities have been reluctant to widely embrace DG technology.

Key Idea/Approach to Solution

Using any DG technology, high-availability microgrids can combine power sources and continuously distribute power to multiple critical loads. This approach avoids the failure points created when paralleling rotating equipment in the AC environment (DC power sources can be combined without invoking the synchronization and interdependence issues that frequently cause downtime in AC systems). As a result, the system can couple all of the power sources to all of the loads at all times without any switching. This ultra-high availability is based on a redundant array of independent devices. The independence and redundancy of each element in the system means that, with appropriate application, multiple devices in the system could fail without affecting the critical load. Any component that fails can be cut off from the interconnections without propagating sympathetic failures to other components. Each prime mover control is semi-autonomous, responding to commands from a central control when appropriate, but capable of taking independent action if system conditions dictate.

The three technologies that appear best suited to critical applications are reciprocating engines, gas turbines or microturbines, and high-temperature, solid-oxide fuel cells, which should become commercially available in a few years. The grid can provide backup power on an interruptible basis or may provide prime power economically for a portion of the total load—an improvement over traditional DG systems that require backup utility for the entire system. The system architecture allows equipment to be serviced and repaired, and components replaced, without interrupting or affecting power output. DG heat recovery enables delivery of diverse thermal outputs, including economical cooling for the servers. To achieve high availability, the system cannot rely upon a single source of energy. Depending upon the DG technology selected, geographic location of the facility, and the capability to store fuel onsite, options for backup fuel include biofuels, diesel, propane, and liquefied natural gas (LNG). A variety of power conversion devices are available as system outputs, depending upon the type of load served.

Analysis, research, development, demonstration, and deployment are needed on multiple aspects of a DC microgrid for data centers and telecom central offices, including DC architectures, safety and protection components, controls, and system integration. The activity requires a two-year time window that can begin immediately.

Expected Benefits/Outcomes

High-availability DC microgrids can provide primary power to critical facilities on a continuous basis with a 1-percent chance of failure over a typical 20-year lifespan⁵. Incorporation of CHP can increase input energy efficiency to 70 percent. Biofuels may also be incorporated to diversify fuel usage.

⁵ Whit Allen, Derek W. Fletcher and Kevin J. Fellhoelter, "Securing Critical Information and Communication Infrastructures Through Electric Power Grid Independence," Proceedings of the 25th Annual International Telecommunications Energy Conference, INTELEC '03, Yokohama, Japan, October 2003, pp. 170 – 177.

DRAFT

C. Cooling

Cooling by air may account for 25 percent or more of all power consumed by major telecommunications central offices and data centers. This proportion could increase further as the density of servers and switching equipment increases.

In most air cooling, air picks up heat from a solid, such as a CPU, and carries it to another solid, such as the cooling coil of an air conditioner, where the heat is removed from the air. Liquid or refrigerant passes through the cooling coil and carries the heat to a cooling system, where the heat is dissipated to the environment. While air may be convenient, it is highly inefficient as a heat mover compared to liquids such as water—which has a latent heat capacity 24 times greater than air. Considerable energy must be used to move the large amounts of air required to provide adequate cooling. In addition, large temperature differences are required at the air-solid interfaces to transfer heat effectively. In many cases, air movement is not well controlled, and the temperature differences have to be made even larger to ensure adequate cooling under worst-case conditions. In most cases, mechanical cooling systems are used to create these differences, requiring even more energy.

While most ICT facilities today rely heavily on air cooling, more efficient options may include close-coupled liquid cooling, spray cooling, or a combination of systems.

INDUSTRY VISION STATEMENT

In 15 years, we should reduce all cooling energy as a percentage of ICT power to a global average not to exceed 20 percent for retrofit and 5 percent for new construction. Cooling systems should be adaptable over time, scalable, and able to maximize utilization and longevity of all assets over their lifetime—while maintaining system resiliency and lowering total cost of ownership (TCO).

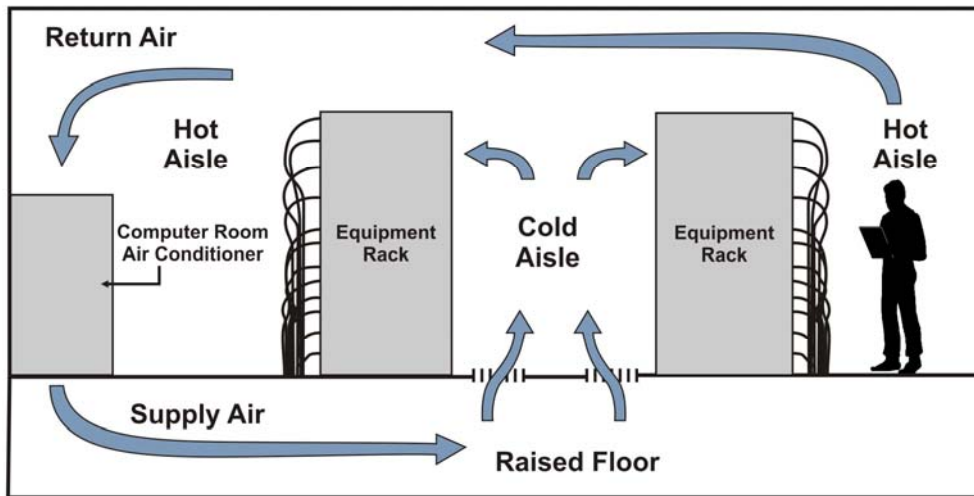


Figure 5. Most data centers today rely on air cooling

Liquid Cooling of Components

Current Issues/Challenges

Most data centers are cooled with chilled air, yet liquid cooling is far more efficient at transferring heat. Circulating a liquid also requires less than a tenth of the energy required to circulate air, and the required temperature difference at solid-to-liquid interfaces is about one fifth that required for air-to-solid interfaces for the same cooling effect. In many cases, the need for mechanical cooling can be sharply reduced or even eliminated. On the down side, liquid cooling suffers from an early history of low reliability and higher capital costs. Most operators of ICT equipment are extremely resistant to running liquid anywhere around their electronic equipment due to concerns over leakage. While many reliability issues have been resolved, the costs associated with purchasing and maintaining pipes and hose couplings to circulate the liquid remain. This issue arises from the current strategy of providing each component with its own cold plate, through which the coolant is circulated. This expensive and difficult-to-service design presents cost and logistical hurdles that force facilities to rely on air for at least some hardware cooling, thus limiting the energy savings potential.

Although some developers have investigated use of a single cold plate to cool all components on a circuit board, such systems have encountered technical difficulties. Intimate contact between the cold plate and the cooled component is essential for efficient heat transfer, yet such contact is prevented by variations in component heights that occur even with the most modern electronic assembly methods. Nevertheless, a single cold plate system would offer advantages in reliability, initial cost, and serviceability.

Key Idea/Approach to Solution

One approach would be to design a system incorporating an enclosure containing many components to be cooled by a single, liquid-cooled cold plate. In addition to enabling devices, further work is needed in infrastructure redesign to take advantage of the delivered benefits. Specific needs are as follows:

- An elastic interface that can maintain a highly conductive heat path between each individual integrated circuit component in an enclosure and the cold plate;
- Other heat movers for components other than integrated circuits, such as hard disk drives;
- Cold plate optimization, including flow rates, sizing etc. as well as serviceability. This implies that the interface between the cold plate and the various heat path assemblies must be easy to break so that the enclosure may be removed for repair, component changes, or additions.

The activity requires a two-year time window that can begin immediately.

Expected Benefits/Outcomes

Such systems will reduce cooling energy by 70 percent, lower capital expenditures on new builds by up to 50 percent, extend the life of existing facilities, and improve the work environment. A standard volume server's heat removal limitation of 600-700W can be increased to 1500W or more, and racks can be fully populated to handle 60KW, up from 6KW in some cases. These benefits could cut infrastructure capital costs by up to 50 percent due to savings in HVAC, electrical capacity, and floor space. The freed-up energy from cooling could power more IT equipment, postponing new facility construction. Finally, liquid-cooled systems are totally silent. Overall energy consumption for cooling can thus be reduced to 5 to 10 percent of total IT power.

Advanced Cooling at Chip/Server-level

Current Issues or Challenges

Two opposing densification trends at data centers are (1) the emerging distributed, parallel applications that benefit from low power, low density, scale-out, distributed IT, and (2) closely coupled modeling, knowledge discovery, and visualization applications requiring high-density, high-performance processing technology. Requirements of the latter class of applications, running on state-of-the-art CMOS semiconductor technology, demand advanced thermal management solutions that avoid performance throttling and would still find wide acceptance in the ITC community.

High-performance microprocessors currently dissipate nearly 100 W/cm^2 of background heat flux, with localized heat fluxes that can be $1,000 \text{ W/cm}^2$ or higher. This is currently handled through over-provisioning of cooling resources. Novel, site-specific, on-demand cooling approaches are needed to handle these challenges.

Key Idea/Approach to Solution

Microfluidic cooling using liquids for next-generation chip architectures is a promising approach. While the background heat fluxes can be handled with the high heat transfer coefficients in single-phase, microfluidic liquid cooling, localized hot spots can be handled using local heat transfer enhancement via phase change. A key challenge is to develop miniature solutions that will fit in commercial servers. One approach is to incorporate thin, liquid-cooled heat spreaders, coupled with active and passive enhancements on the air-side. The active enhancements would include synthetic jets and vibrating reeds to improve localized heat transfer. The liquid-cooled heat spreaders will use enhanced boiling from micro-fabricated structures.

A major hurdle to the introduction of advanced cooling technologies is the lack of a clear understanding on how these technologies will perform in a production environment. Another hurdle is the reluctance of data center managers towards liquid cooling solutions at the server or chip level. The evaluation of these technologies would have to be performed in the setting of pilot projects that measure energy efficiency, reliability, manageability and total cost of ownership. The activity requires a two-year time window that can begin immediately.

Expected Benefits/Outcomes

Not only will advanced cooling solutions enable further chip densification and IT performance improvements, but they could reduce by half the energy consumed by cooling hardware by easing, improving, or potentially eliminating forced-air cooling systems.

Efficiency-Optimized Control Systems

Current Issues or Challenges

A number of studies show that the use of state-of-the-art technology could reduce energy consumption by up to 40 percent, with a big part of this savings potential (26 percent) coming from use of high-efficiency cooling and air movement (chillers and fans). Additional gains (5-10 percent) can be derived by dynamic allocation of computational resources. A 30 percent reduction of energy use in data centers will result in 20 billion kWh/year in electricity savings nationwide. Further reductions (10-20 percent) were reported as possible through the use of advanced cooling systems, such as independent dehumidification, driven by waste heat from onsite power generation equipment; direct liquid cooling in the racks; and cascade cooling systems. Central to achieving these savings is the optimization of the temperature at which air is delivered and the airflow and localized heat generation management in the computer rooms.

One barrier to implementing these solutions is the lack of understanding of the coupling between the various subsystems (HVAC, air flow, computing allocation). These couplings result from the interaction between the refrigerant and cooling loops, air movement, computational load distribution, and heat generation from the computational racks. For example, it is highly desirable to raise the temperature of the air entering the room as this increases the efficiency of the chiller and/or combined heat and power (CHP) plant. However, this, in turn, increases the demand for ventilation and fan power. The resulting trade-off should be actively controlled given rack temperature limits. A challenge is that the thermal environment in data centers varies both temporally and spatially. The temporal variation is the result of varying load demands within short time scales. There are also changes of computing systems and physical configurations at longer time scales.

Key Idea/Approach to Solution

The key to achieving the above-mentioned energy savings is to develop an integrated, holistic, real-time, system-level approach to optimize the energy usage. This approach requires rapid, high-resolution sensing and control of the key operational parameters and electrical and cooling demands of the chiller/CHP and air delivery systems to maximize energy efficiency. This research effort will require approximately three years for developing an overall framework and the underlying technologies that can be seamlessly deployed across various sites.

Expected Benefits/Outcomes

Up to 30 percent achievable energy gains can be expected in data centers through proper system integration and real-time control (both computing and HVAC).

Free Cooling & Standards Development

Current Issues or Challenges

Today's ICT equipment is, for the most part, assembled using components and processors that are designed to operate best in an environmentally conditioned space (i.e., controlled humidity and temperature). These mandatory operating parameters force the operator to expend scarce capital dollars to procure the heating, ventilation, and air conditioning (HVAC) systems necessary to maintain this operational environment. On top of the initial cost, facilities must pay for the associated monthly electricity use (approximately 40% of the utility bill is attributed to powering the HVAC systems) plus the cost of ongoing maintenance / repair / replacement activities as needed to keep the systems operating as designed. Not only do these HVAC systems impose a burden on an already taxed power grid, they add an incremental hidden cost to the price of all products and services sold.

Key Idea/Approach to Solution

In cooler climates, it may be possible to simply filter (and possibly condition, but not cool) outside air for use in the computer room of an ICT facility.⁶ Modern servers, which have protective coatings on circuit boards and hermetically sealed hard drives, appear to be far more robust than their vendors claim. A number of enterprises, including the Lawrence Berkeley National Laboratory, have been experimenting with this concept and detected surprisingly little impact on server performance. Microsoft's well-know "data center in a tent" showed no adverse impacts on reliability, even when water dripped on a server. Other data center design groups are capitalizing on the cooler air available at night or during certain seasons of the year to drastically cut their cooling costs.

Server vendors may be unwilling to expand their operating environments because there is no satisfactory data on reliability. Research is needed to reexamine the underlying science of reliability versus temperature and dew point, to identify the acceptable range of operating conditions for current equipment, and to define acceptable standards. In the telecommunications world, specifications (Telcordia GR-487, UL/NEBS) exist today that serve as guidelines for the manufacture of more robust equipment and associated outdoor enclosures—devices that can operate in a less carefully controlled operating environment (higher heat, greater humidity).

Expected Benefits/Outcomes

If revised design guidelines or standards were incorporated into the manufacture of data center mainframes, servers, storage, etc., the operating environment controls could be relaxed without adversely affecting the performance of the electronics. The difference can mean significant data center operations savings, which may be viable from a net cost perspective (higher front end cost offset by reduced power / maintenance costs over the life of the equipment). Beyond the ICT facility, use of free cooling can reduce demands on the utility grid, cut power plant fuel consumption, and lower greenhouse gas emissions.

⁶ It is currently possible to use "free" cooling without allowing outside air into the data center, if outside air is passed through a heat exchanger to cool off warm indoor air. However, allowing cold outside air directly into the data center could potentially save more energy—if the problems with humidity and particles in outside air could be solved or reduced.

DRAFT

Appendices

Supporting Technology Needs

Although technology improvements are essential in improving data center efficiency, many of the most important problems to be addressed in the industry are *non-technical*. Often times, energy efficient choices and best practices are not implemented because of organizational barriers, outdated management practices, or insufficient executive-level attention to data center efficiency issues.

For example, in many organizations the “IT group” that makes decisions regarding servers and IT equipment does not communicate well with the “facilities group” that pays for the data center’s power equipment, cooling equipment, and electricity. Even in the presence of better technology, it is difficult to make data centers efficient if these two groups do not work together.

Executive-level decision-makers in the ICT industry need to be better informed about the *financial and business impacts* of data center energy use. Executive-level decision makers are key, because typically only they have the authority to drive the cross-departmental changes in their organizations that are necessary to achieve large data center efficiency improvements. Financial and business impacts are also important, because these arguments are necessary to get the attention of “bottom-line focused” executives.

Metrics

The industry must address efficiency metrics both at the “whole data center” level and at the component level. Whole system metrics are essential for capturing trade-offs. Imagine a new data center technology that saves energy by removing the cooling fans from each server and consolidating them into a more efficient facility-level cooling unit. This would be an overall efficiency gain, even though it makes power utilization efficiency (PUE) look worse by reducing the apparent server energy use.

Component-level metrics are just as important. Data center purchasing managers need easy-to-use metrics for comparing the efficiency of one server versus another, one external storage device versus another, one UPS versus another, etc., without having to have a complicated “whole data center” model to evaluate every individual equipment decision.

Workshop Breakout Results

The charts on the following pages represent the ideas articulated by the participants in the breakout sessions at the workshop. The orange circles represent votes for the concepts with the highest priority. While many of the ideas refer to worthwhile activities, this document selectively focuses on technologies that require collaborative R&D.

EQUIPMENT AND SOFTWARE: WHAT ARE THE POTENTIAL SOLUTIONS?

ECONOMIC	VISIBILITY EDUCATION & MANAGEMENT	TOOLS	BEHAVIOR
<ul style="list-style-type: none"> • “Space Program” for energy innovation ● • Taxation on energy costs to fund DOE project pool● • Corporate tax advantage to buy/develop energy-efficient equipment • DOE and DOD dollars for breakthrough technical innovation on efficiency • Use military efficiency R&D to pull civilian products to market • Eliminate disincentives (or conflicting incentives) to innovation/efficiency • Price on resources to increase consideration (e.g., CO₂ tax) • Adequate federal funding for industry R&D 	<ul style="list-style-type: none"> • CEO/CTO/CIO/CFO forums to drive top-down leadership - partnership with Federal government ●●●●●● • DOE/Industry should fund think tank to develop green solutions • Knowledge sharing at all levels • Green technology and solutions sharing (public patents) • Professional accountability of energy usage/consumption - energy mgrs and execs. • <i>Holistic approach to reliability to avoid overkill</i> 	<ul style="list-style-type: none"> • Invest in tools to capture and report on integrated “datacenter” efficiency statistics (severs, switches, facilities, network, etc.) ●●●●●● • Tools to track utilization and method to correlate output to use/utilization ●●● • Develop a framework with optional monitoring on control parameters for high-fidelity real-time closed-loop control of facilities and IT ● • Develop technical and business tools that equip decisions in: <ul style="list-style-type: none"> - Life-cycle predictions - Energy recommendations - Right sizing to risk • Instrumentation for continued energy efficiency certification monitoring (as the facility gets old, it is still green) • Simplified tools for life-cycle analysis of product environmental impact 	<ul style="list-style-type: none"> • Investigate usage model – how is ICT currently operating in terms of power management, etc., and develop a feedback loop to developers ●●● • Research, report and forecast user trends annually

EQUIPMENT AND SOFTWARE: WHAT ARE THE POTENTIAL SOLUTIONS? (CONTINUED)

TECHNOLOGY		STANDARDS & REGULATORY	METRICS
<ul style="list-style-type: none"> • Improved compute technology (Gflop/watt) e.g., accelerated computing ●● • Investigate energy impacts of alternative network/information architectures (distributed vs. centralized) ●● • Develop low-latency ICT power management technologies and strategies ●● • Develop flexible power systems in ICT that are source agnostic ● • Evaluate whether it makes sense to harden equipment and expand its temp/humidity tolerances, rather than trying to implement better cooling ● • Develop applications to use ICT for energy savings in other sectors ● • Develop ICT cooling solutions that are cost effective, low environmental impact, and low reliability risk ● • Advanced data security as an enabler to adopt energy-efficient technology (if we can't do it securely, we can't do it) 	<ul style="list-style-type: none"> • Nano-electronic circuitry capable of morphing (self-modifying hardware; adapt hardware to fit the program, not vice versa) ● • Easier reuse and recycling of equipment components • Develop ICT equipment that can operate in existing worker spaces and infrastructures without impact to the average worker space, environmental requirements • Ultra low power circuits like multi-phase clock, asynchronous circuits within ICT equipment • Develop all-optical systems to improve latency and increase energy efficiency • Government/DOE labs proof of concept data centers for new technology • Parallel computing/cloud computing (distributed vs. centralized) • Automated heuristic controls for ICT & facilities and equipment (self intelligence) 	<ul style="list-style-type: none"> • Open network and communication standards for energy management of ICT and facilities equipment ●●●●● • Create national mandatory⁷ standards on green ●●● • Deregulation - increase incentives for enabling change/ adaptation ●●● • Regulated energy and environmental reporting (10K) ●● • Legislative change to remove regulatory barriers to efficiency innovation • Building codes specific to energy-efficient data centers • Association advocacy to deliver industry side of issues to governments and compilation of industry progress and trends into market intelligence • Total cost of ownership (TCO) measurement standard testing required for all new products • Global standards development instead of regional regulations through use of industry associations and alliances • Agreed-upon, trustworthy, vendor-neutral testing body for validating vendor energy performance claims 	<ul style="list-style-type: none"> • Acceptable cross-industry metrics for measuring performance ●●●●●●●●●● • End-customer-initiated metric definitions (ex: TEER - Telecom Energy Efficiency Ratio) • Create an energy efficiency metric for CFO's (e.g., GAAP) • Develop metrics & tools for efficient code (minimize instructions to do a given function) • Baselining ICT Green characteristics in order to measure improvement • Improve ICT standards to require efficiency and sustainability targets; develop key benchmarks that can be applied • As roadmap goes forward, focus on understanding and accounting for different needs of 1) telecomm, 2) "cost-center" DCs, 3) "profit center" DCs, and 4) others • Mandated infrastructure sharing

⁷ Support for mandatory nature of standards not universal among participants. Cap and trade may be preferred.

POWER SUPPLY CHAIN - WHAT ARE THE POTENTIAL SOLUTIONS?

POWER SYSTEM ARCHITECTURE	COMPONENTS AND EQUIPMENT	
<ul style="list-style-type: none"> • Transition to full DC operation at uniform voltage ●●●●●●●●●● <ul style="list-style-type: none"> – Develop < 600 volt DC bus architectures that enable distributed energy resources (DER) and efficiency simultaneously – Develop, design, and install a modular, on-site DC generation system and micro grid • Develop new utility DC distribution • <i>Develop system using several low-voltage dc buses and rectifiers</i> • Eliminate isolation requirement from 48V to 12 V PC to 1X VDC to other (telcomm) • Eliminate AC to AC conversion stages where possible. i.e., 480 VAC 30 to the server/rack ●●●●●●●● • Take “system” look at power conversion • Provide passive DC power generation (solar/wind) ●●●●● <ul style="list-style-type: none"> – Wind/solar generator and hydrogen storage • Develop/promote alternative storage devices ●● • Introduce nuclear power (≥ 65% efficient) into available energy mix) ● • Develop alternate power grid/<i>on-site, self-generated power</i> • Migrate from energy-based to power-based systems <ul style="list-style-type: none"> – Go from backup energy storage to onsite generation • Research lossless power transmission (fiber? wireless?) ● <ul style="list-style-type: none"> – Move from signal over copper or power over copper to wireless signal or wireless power and optical signal and optical power • Investigate energy issues of IT ↔ telecom, e.g., VOIP 	<ul style="list-style-type: none"> • Research and develop high-efficiency power conversion circuits ●●●●●●●● <ul style="list-style-type: none"> – Optimize server/telcomm equipment, i.e., consider wide range DC/DC to silicon coupled with fixed duty cycle front end ● • Control systems built to optimize efficiency ●● • Develop special purpose chips, multiphase clocking, ternary/other processing modes, lower-power chips (Intel, AMD) ● • Promote use of DC connectors in early 2009 ● • Commercialize direct carbon fuel cell ● • Research issues affecting conversion loss (e.g., transformer materials and construction) ● • Research the use of optical switching to eliminate many conversion steps & losses ● • Hi efficiency converters (SiC, GaN) <ul style="list-style-type: none"> – Better field effect transistors (FETs) • Optimize UPS to maintain full operations and space conditioning for long-duration power outages 	<ul style="list-style-type: none"> • Eliminate ‘lossy’ availability circuits, or-ing, N+N, other • Conduct R&D of superconducting components • New processing paradigms/ algorithms, Neural -3D arrays in servers (specialized co-processors), programmable architectures • Power of software: Embrace digital controlled conversion to combine stage, i.e., power factor correction (PFC) and output into one stage • Develop more energy efficient breakers, fuses, shunts • Research use of piezoelectrics <ul style="list-style-type: none"> – Develop micro mechanical air conditioning for point of load cooling

POWER SUPPLY CHAIN - WHAT ARE THE POTENTIAL SOLUTIONS? (CONTINUED)

MIND SETS	STANDARDS	BUSINESS CASE	TRAINING
<ul style="list-style-type: none"> • Reduce or eliminate the criteria for UPS backup ●●●●● • Make technology leaders responsible and accountable for data center facilities including decisions, capacity, cost and environmental impact ●● • Demonstrate more efficient “systems” for all to see ● • Increase demands and expectations from clients and consumers ● • Develop demand response techniques (electrical and thermal) that meet reliability requirements of ICT industry • Change framework of research funding, e.g., direct DOE funding to systems research centers (like DOE’s basic science centers) 	<ul style="list-style-type: none"> • Develop standards (e.g., IEEE) for data center power systems (including recommended practices) [operators, vendors, stakeholders, A&E’s, investors] ●● <ul style="list-style-type: none"> – Develop IEEE/IEC standard on high-voltage DC systems • Encourage high level early adopter (with power) champions ●● • Encourage development of missing components (e.g., 400V connectors) ● • Organize for standards development ● <ul style="list-style-type: none"> – Establish V levels and standards, which reduces risk to manufacturers and leads to production development, which drives component development • Achieve _____ TBD _____ emissions standards 	<ul style="list-style-type: none"> • Conduct studies and demonstration projects to compare/contrast AC vs. DC systems ●●● • Establish clear return on investment (ROI) ● • Achieve economy of scale (cost effective) ● • Develop a federal program dedicated specifically to invest in research on data/telecomm energy issues and distributed generation (DG) ● • Establish federal tax incentives to reduce energy consumption ● <ul style="list-style-type: none"> – Government incentives to help drive ROI intervals beyond 2 years • Challenge industry groups 	<ul style="list-style-type: none"> • Stimulate interest in power engineering education ●● <ul style="list-style-type: none"> – Set up power electric engineering scholarship programs • Foster training on safety issues

COOLING - WHAT ARE THE POTENTIAL SOLUTIONS?

PARTNERSHIPS	EDUCATION	STANDARDS POLICY MEASUREMENT	BEST PRACTICES
<ul style="list-style-type: none"> • Fund test facility (public/private) and publicize ●●● • Cross functional teams, e.g., cooling -IT-Power • Inter utility cooperation (incentives, standards, interconnection) • Joint effort among several companies to trial/model element of vision 	<ul style="list-style-type: none"> • Education - sharing of best practices ●●●●● • Provide actual college/high school courses on data centers for engineering students • Academic/university research • Change customer perceptions • K-12 education on science 	<ul style="list-style-type: none"> • Leadership directive/ government mandate ● • Develop world-wide database of power use for ICT systems ● • Economic/financial modeling ● • Set specific goals (metrics) for component manufacturers ● • LEED standard for ICT • Develop plug & play standards for tower-cooled water cooling of ICT equipment • Publicize independent performance data of cooling equipment, e.g., chillers and computer room air conditioners (CRACs) • Corporate focus more on environmental responsibilities in addition to bottom line • Lobby for global standards (aggressive standards) • Provide for either rewards or penalties for standards 	<ul style="list-style-type: none"> • Wide adoption of metering system that generate cost signals ●●● • Develop first in Japan and Europe and bring back the technology to the U.S. ● • Industry leaders fully committed reference designs • Low pressure drop design • Commissioning standards • Inclusion of server/cabinet fans in building cooling strategy

COOLING - WHAT ARE THE POTENTIAL SOLUTIONS? (CONTINUED)

R&D	R&D AND IMPLEMENTATION INCENTIVES	SITING
<ul style="list-style-type: none"> • Liquid cooling of components (some or all) ●●●●●● • Advanced cooling techniques at the component level ●●●●●● • Develop environmental-tolerant ICT equipment for outside air or tower-cooled water cooling ●●●●●● • Develop mitigation techniques to reduce failures associated with “free” cooling ●● • Better use of low-quality waste heat ● • Biomimicry, fan efficiency and noise improvement • Develop free cooling retrofit for data center CRAC units • Cooling integrated onto circuit packs (new methods/research) • Separation of data processes (compute, memory, storage) for specialized conditioning) 	<ul style="list-style-type: none"> • Incentive for innovation with grants, “X” prize, national competitions, monetary rewards, rebates ●●●●●● • Market pull from government organizations and large industries ●●●● • Legislation ●● <ul style="list-style-type: none"> – Strong carbon tax – Tax breaks for early adopters – R&D credits • Utility funding for R&D demonstrations in facilities 	<ul style="list-style-type: none"> • Maps OA/water economizer zones ● • Identification of regional DC “Power Parks” (allows for integration with utility at utility site) • Company locations of data centers and power generation