



THE GREEN GRID OPPORTUNITY

DECREASING DATACENTER AND OTHER IT ENERGY USAGE PATTERNS



With energy shortages and energy costs taking top rank among the concerns of our times, a movement to decrease datacenter and other IT energy usage patterns is gathering momentum. A key step has been the formation of The Green Grid. Founded in February 2007, this growing industry organization seeks to define and propagate the best energy-efficiency practices in datacenter operation, construction, and design. In this paper, we'll examine why this is so important and the initial steps being taken.

THE NEED FOR DATACENTERS TO GO GREEN

Until recently, electricity usage wasn't a major IT concern. In fact, the electricity costs for running and cooling computers generally wasn't even considered part of the IT budget. Servers didn't draw much power and their density in datacenters wasn't significant enough to create any concern. All that has changed with the nonstop growth of the Internet and the world's businesses, governments, and other institutions filling up rack after rack with more powerful servers to keep up with growth.

At Gartner's 25th Annual Datacenter Conference (November 2006), electrical power consumption was a headline issue, demonstrating the IT industry's growing energy consciousness and concern about unsustainable levels of power usage. Gartner analysts at the conference predicted that half of the world's datacenters will run out of power by the end of 2008.¹

This doesn't mean that these datacenters will go dark. It means they won't have enough energy capacity to meet the power and cooling requirements of the latest high-density computing equipment. "With the advent of high-density computer equipment such

as blade servers, many datacenters have maxed out their power and cooling capacity," explained Michael Bell, a vice president for research at Gartner. "It's now possible to pack racks with equipment requiring 30,000 watts-per-rack or more in connected load. This compares to only 2,000 to 3,000 watts-per-rack a few years ago."² Richard Edwards, a senior research analyst at Butler Group, points out that the density of IT equipment per cubic meter is 20 to 50 times what it was 50 years ago and that IT's hunger for power still remains unquenched.³

Of particular concern is that power consumption is rising faster than the average datacenter can accommodate. It's not just a matter of building out of space to house new servers. IT also needs to add additional electrical capacity and cooling systems.

POWER CONSUMPTION: AN ISSUE THROUGHOUT THE DATACENTER

Electrical power needed to run today's high-performance datacenter servers is only part of the problem. Non-IT devices that consume datacenter power include such things as transformers, uninterruptible power supplies (UPS), power wiring, fans, air conditioners, pumps, humidifiers, and lighting. Some of these devices, like UPS and transformers, are in series (they're part of the power path) with IT loads, while others, like lighting and fans, are in parallel with IT loads because they perform other support functions in the datacenter. Virtually all the electrical power feeding the datacenter ultimately ends up as heat. Unfortunately, according to Gartner, the vast majority of hardware devices have been designed to provide maximum functionality and performance with little

Gartner analysts advise enterprises to make the most efficient use of existing systems with technologies such as workload balancing and virtualization, and to plan for greener datacenter designs in the years ahead.

With energy rates also rising, companies have tried to save power where they can. Both Microsoft and Yahoo Inc., for example, are putting server farms in eastern Washington state, near less costly hydroelectric power supplies. Google Inc. has put a datacenter near hydro supplies in Oregon.

A HARDWARE COMPANY MIGHT IMPROVE THE ENERGY EFFICIENCY OF A BLADE SERVER AND ANOTHER COMPANY MIGHT IMPROVE THE ENERGY EFFICIENCY OF A TRANSFORMER TO CONVERT AC TO DC. RIGHT NOW, VERY FEW PEOPLE ARE DOING RESEARCH ON HOW TO INTEGRATE THOSE TWO TOGETHER.

—David Rodgers, acting deputy assistant secretary of the U.S. Department of Energy’s Efficiency and Renewable Energy (EERE) office



regard for wider environmental issues. For instance, Gartner research vice president Rakesh Kumar says traditional datacenters typically waste more than 60% of the energy they use to cool equipment.⁴

Other inefficiencies are also easy to find. For example, vendors commonly put inefficient power supplies in high-volume servers because they don’t see a competitive advantage in putting in more efficient components. Since IT typically doesn’t pay the electrical bill, their buyers have shown greater interest in the initial cost of a server rather than




Figure 1: Datacenter Cost Factors

on energy savings over the server’s lifetime. Jon Koomey, a consulting professor at Stanford University, calls this a “perverse incentive that pervades the design and operation of datacenters.”⁵ For instance, commonly used power supplies have a typical efficiency of 65–70% and are a huge generator of waste heat, while units with efficiencies of 90% or better are available and pay for themselves over the life of the equipment. Yet, IT buyers still go for the less-expensive servers with the inefficient power supplies because the benefits of specifying higher-efficiency power supplies are not obvious.

Power conversion is also a potential area of significant improvement when it comes to datacenters. From the time power enters into a datacenter until the time it reaches a server’s microprocessors, power is converted many times. A study by a major microprocessor manufacturer found that IT datacenters typically burn more power in power conversion and cooling at light loads (zero to 25% platform utilization) than the computer systems themselves are using. That means there are significant opportunities for energy saving in the design of power conversion and cooling systems that scale better with the load.

Many customers are unaware or unable to convert all of the available power and cooling into compute cycles. Given that unlimited access to energy is unlikely and the cost of building or retrofitting datacenters is high, it makes sense that customers are provided the tools to increase the efficiencies of today’s datacenters. Measuring and improving these facilities using the proper performance and energy metrics, along with a set of best practices, will help companies plan a more efficient datacenter.

THE NEED FOR BETTER DATACENTER METRICS

To cope with the coming power crisis for datacenters, companies will need to retrofit existing facilities, build new ones, or risk running out of capacity for growth. The power and cooling issues facing large IT organizations are being driven by the following factors:

- The increase in computing demand.
- The increase in IT equipment power density.

FROM AN END-USER STANDPOINT, A STANDARDIZED METHOD OF ESTABLISHING AND REPORTING SERVER ENERGY CONSUMPTION IS ESSENTIAL.

—Rick Salmon, vice president of the mission-critical facilities group at Lehman Brothers Inc., quoted in a Computerworld article, October 30, 2006



- The increase in energy cost and its subsequent importance as a significant part of the total cost of ownership (Gartner estimates that energy bills traditionally have accounted for less than 10 percent of an overall IT budget but soon could account for more than half).⁶
- The availability of power to meet computing demands.

Fair comparisons for datacenter energy efficiency (and the components that comprise it) require a standardized set of performance and energy efficiency metrics. Much like miles-per-gallon (MPG) comparisons on vehicles allow the buyer to compare how well a car converts fuel (in gallons) into work (in miles), datacenter managers need a standard set of metrics to understand the efficiency of their datacenters, improve the performance-per-watt of their IT equipment, and make smarter IT purchases. In addition, these metrics will provide a common way to benchmark against existing or legacy datacenters as well as other organizations and industry norms.

Conventional models for estimating the electrical efficiency of datacenters are grossly inaccurate for real-world installations. Currently, many manufacturers provide efficiency data for power and cooling equipment. For power equipment, efficiency is typically expressed as the percent of power out to power in. For cooling equipment, efficiency is typically expressed as the ratio of heat removed to electrical input power (coefficient of performance). Unfortunately, these individual values of efficiency often lead people to think that the efficiency losses of a datacenter can be determined by simply adding up the inefficiencies of various components.

Unfortunately, this approach does not provide accurate results in the case of real datacenters. In fact, it can cause datacenter designers to overestimate efficiency, and consequently underestimate the losses or over-provision for these losses.

GOING GREEN THROUGH THE GREEN GRID

The Green Grid (www.thegreengrid.org) is a global consortium dedicated to advancing energy efficiency in datacenters and business computing ecosystems. A non-profit trade organization of IT professionals, The Green Grid's charter is to address power and cooling throughout datacenters and the information service delivery ecosystem. The Green Grid does not endorse any vendor-specific products or solutions, but will provide industry-wide recommendations on best practices, metrics, and technologies that will improve overall datacenter energy efficiencies. Green Grid members include industry-leading technology companies (AMD, APC, Dell, HP, IBM, Intel, Microsoft, Rackable Systems, SprayCool, Sun, and VMware). These companies all realize that a standards-based industry solution is needed to coordinate efforts in developing specifications and metrics for power efficiencies.

The scope of The Green Grid is broad in the context of energy and power efficiency. It includes the definition of meaningful, end-user-centric models and metrics; development of standards, measurement methods, processes and technologies to improve performance against the defined efficiency metrics; and the promotion of energy-efficient standards, processes, measurements, and technologies.



INITIAL GOALS: DEVELOPING BETTER METRICS

The Green Grid has identified both short-term and long-term objectives to increase the energy efficiency of datacenters and IT infrastructure equipment around the globe. In the short term, the organization seeks to define a way for facilities organizations, IT managers, CIOs, regional power utilities, and governmental regulatory agencies to evaluate the performance-per-watt of the datacenter and their components (including islands or sub-components of the datacenter). While there are metrics today used to gauge the performance of the datacenter, their usefulness falls short when measuring datacenter performance-per-watt.

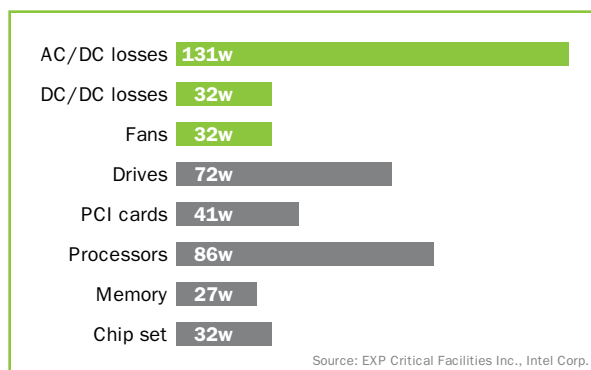


Figure 3: How 160 Watts Are Lost Per Server

Based on a typical dual-processor 450W 2U server approximately 160W (35%) are losses in the power conversion process.

One such metric that has been used for years is datacenter density (DCD):

$$\text{DCD} = \frac{\text{Power of All Equipment on Raised Floor}}{\text{Area of Raised Floor}}$$

The primary use for this benchmark is to determine if the deployment is low-, medium-, or high-density. While this is a great metric to determine the

absolute performance of a datacenter relative to other datacenters, it fails to capture whether the deployment is being done effectively. When taking a holistic view of the datacenter, it becomes clear that current metrics such as DCD, measured in kilowatts-per-square-foot, are not useful metrics for datacenter efficiency since the power to cool and convert power for the IT infrastructure in a datacenter is now greater than the power used by the IT equipment itself. The question of datacenter energy efficiency really starts at the electrical meter coming in from the utility and needs to look at how much of that power coming in is actually being used to by the IT computing infrastructure to do work. Once that can be determined, it will be easier to pursue three Green Grid goals:

- Minimize the power needs of the datacenter.
- Maximize the percentage of that power coming in that is used for IT computing work.
- Minimize the amount of power spent on non-IT computing equipment (i.e., power conversion, cooling, etc.).

Other short-term objectives of The Green Grid include the publication of best-practices white papers and checklists, definition of standard workload and deployment models, engagement with other relevant standards organizations, and extensive engagements with datacenter operators to define key energy-efficiency requirements going forward. In addition to data publication and consultative efforts, The Green Grid will also define a long-term power efficiency roadmap for the datacenter that charts relevant technology and practices to a timeline of continuous improvement.

STANDARDS HELP US ALL COMPETE BETTER.

—Paul Perez, HP vice president for storage, networks and infrastructure
for industry-standard servers at HP



A NEW DATACENTER POWER EFFICIENCY ARCHITECTURE

Long-term objectives for The Green Grid include defining a new datacenter power efficiency architecture to implement energy-efficiency policies natively through the instrumentation of devices. This architecture will include automatic control of datacenter components via policy-based management geared toward IT objectives for power efficiency. Included in this effort will be specifications, as well as compliance and interoperability testing linked to a Green Grid logo program.

To develop this architecture and set of specifications, much groundwork needs to be done in proposing, aligning, and defining each element. To do this, The Green Grid has organized a governance body (board of directors), communications committee, and technical committee. The technical committee is further segmented into workgroups chartered to manage a specific discipline. Each workgroup consists of industry and end-user experts in its particular area. The initial four workgroups in the technical committee are Data Collection and Analysis, Metrics and Measurement, Data Center Operations, and Data Center Technology and Strategy.

PROGRESS WILL REQUIRE AN INDUSTRY-WIDE EFFORT

The expertise for Green Grid initiatives will come from the extensive proficiency and knowledge of its members. In order to achieve its long-term objectives, The Green Grid needs industry participation. Membership in The Green Grid is open to IT professionals who are concerned or chartered with datacenter operations and facilities management.

For more information, visit The Green Grid at www.thegreengrid.org.

The Green Grid is devoted to working closely with both the EPA and the Alliance to Save Energy. Both organizations are in full support of The Green Grid and will play a key role in the development of the organization's charter as well as collaborating to identify innovative solutions that cut energy costs as well as benefit the environment.

REFERENCES

- 1 "Gartner Says 50 Percent of Data Centers Will Have Insufficient Power and Cooling Capacity by 2008," Gartner Inc. Press Release, November 29, 2006.
- 2 Ibid.
- 3 "The Future of the IT Organisation: Do You Run a Green Machine," Computer Weekly, December 12, 2006.
- 4 "Gartner Says Look Beyond Power Issue As Pressure Mounts for 'Greener' IT," Gartner Inc. Press Release, November 7, 2006.
- 5 "Sidebar: The Big Drain: Where the Most Power is Wasted," Computerworld, April 3, 2006
- 6 "These Energy Efforts are Hot, Hot, Hot," Network World, October 6, 2006.