Chapter 8

Choosing Greener Gear

This chapter offers strategies for making a Data Center greener through IT hardware choices, outlines server energy efficiency standards and metrics, and discusses both hardware utilization and how to reduce hardware power consumption.

Environmental Impact of IT Hardware

In and among greening the facilities-related technologies of your Data Center—power, cooling, and fire suppression—don't overlook opportunities to green the systems that drive your Data Center power consumption in the first place—your IT hardware. Because servers, networking devices, and storage units are the final destination of your Data Center's power and cooling delivery chains, anything you do to reduce their power consumption and heat generation has a cumulative benefit.

As mentioned in Chapter 4, "Powering Your Way to a Greener Data Center," every watt of power conserved at the server level actually saves nearly three. Although energy consumption is the most obvious green factor of Data Center hardware, other elements such as what materials are consumed to manufacture a device and their footprint—smaller systems use fewer materials than larger ones—are also valid to consider.

Aside from the environmental benefits of choosing energy-efficient hardware are the cost savings. Exactly how much money is saved depends upon the cost of power where your Data Center is located, but in the mid-2000s, several Data Center industry groups began predicting that energy costs over a server's lifetime would soon exceed the initial price tag of the device.

Note Electrical costs associated with IT hardware have become so significant that a British manufacturer of energy-efficient servers now offers to provide its machines to companies at no capital cost in exchange for a portion of the Data Center utility savings its hardware provides over other models.

Under VeryPC's Free Green Server Initiative, customers can trade in more than 135 types of servers—a variety of Dell, Fujitsu, Hewlett Packard, and Sun systems—in exchange for leasing the company's GreenServer Janus 2 hardware. A customer's monthly lease payments are calculated based upon anticipated electrical cost savings from replacing the old hardware with the new servers.

You can find more information about the Free Green Server Initiative at http://www.freegreenservers.co.uk/.

When evaluating IT equipment for your Data Center, keep in mind that your overall goals are to obtain hardware that can meet your processing needs for a long time and are efficient not just as individual machines, but as a complete system. That means rather than simply defaulting to whatever hardware draws less power, focus on processing per kilowatt. Strategically choose more powerful devices that have greater capabilities and can perhaps enable other systems to be consolidated. For example, a small economy car consumes much less fuel than a large school bus, yet the greater seating capacity of the bus makes it more efficient for transporting a large number of people. The bus, in essence, takes the place of multiple small cars.

If your Data Center is equipped with servers that draw less electricity than other models yet don't provide adequate performance, it will only be a matter of time before you feel compelled to add more hardware or upgrade.

Note I've encountered some businesses over the years that learned the impact of hardware choices upon their Data Center resources the hard way. Having run critically low on electrical capacity in their Data Center, they were forced to ration what little power capacity remained until a new server environment could be brought online. New hardware purchases and installations were heavily scrutinized for their power consumption, and the default policy was that no additional hardware could be installed until other systems were removed that accounted for at least as much power usage.

Not surprisingly, the experience left a lasting impression, causing the companies to focus on power consumption when choosing new hardware long after they had resolved the capacity shortage.

Hardware Energy Efficiency Targets

As the role of power has increased in importance in the Data Center—power is both the defining factor of a server environment's true capacity and its greatest operational expense these days—so too has interest grown concerning the energy efficiency of IT hardware. Understanding power-efficiency targets for servers and networking devices that have been set by various agencies and choosing hardware that meet them can help make your Data Center greener.

Energy Star Specifications

The U.S. Environmental Protection Agency is issuing a series of specifications for Data Center-related hardware to receive its Energy Star label.

The first standard, issued in 2009, applies to individual servers that include one to four processing slots. A second tier of the standard, scheduled to come out in 2010, is to encompass systems with more than four sockets, blade servers, fully fault tolerant servers, server appliances, and multinode servers.

Additional Energy Star specification are intended to be developed in the future for storage devices and networking hardware. Initial work on the storage specification began shortly before publication of this book.

To qualify for the Energy Star designation, servers are required to meet several criteria, including

- **Efficiency standards for power supplies:** Server power supplies are required to meet minimum efficiency and power factor requirements under various loads. Distinct performance minimums are called out for a variety of power supplies: AC-DC and DC-DC, varying wattages, and multiple and single output power supplies.
- **Idle power limits:** Servers must not exceed certain power thresholds while idle. The maximums vary based upon the configuration of the server, including the quantities of processors, installed memory, and hard drives.
- **Reporting requirements:** Manufacturers are to provide data sheets on their website detailing the Energy Star-qualified hardware model, including maximum, minimum, and typical configurations for the hardware.
- **Data measurements:** Most of the server configurations eligible for the Energy Star designation must additionally be capable of providing input power consumption, inlet air temperature, and processor utilization data during normal operation.

Table 8.1 shows the minimum power supply energy-efficiency requirements for servers to qualify for the Energy Star designation.

Server power supplies are additionally required to meet the minimum power factor thresholds outlined in Table 8.2.

Table 8.1 *Energy Star Efficiency Requirements for Computer Server Power Supplies (Tier 1)*

Power Supply Type	Rated Output Power	10-Percent Load	20-Percent Load	50-Percent Load	100-Percent Load
Multi-Output (AC-DC and DC-DC)	All Output Levels	N/A	82 percent	85 percent	82 percent
Single-Output (AC-DC and DC-DC)	500 watts	70 percent	82 percent	89 percent	85 percent
	>500 to 1,000 watts	75 percent	85 percent	89 percent	85 percent
	>1,000 watts	80 percent	88 percent	92 percent	88 percent

Various governments, energy agencies, and public utility companies provided comments during the 2 1/2-year development period of the specification as did hardware manufacturers including Dell, EMC, Fujitsu, HP, IBM, Intel, NetApp, Sun Microsystems, and VMware. The agency estimates that servers with the Energy Star label are 30 percent more efficient than conventional systems.

The Energy Star Computer Server Specification is available online at http://tinyurl.com/qwxc8a.

Table 8.2 *Energy Star Power Factor Requirements for Computer Server Power Supplies (Tier 1)*

Power Supply Type	Rated Output Power	10-Percent Load	20-Percent Load	50-Percent Load	100-Percent Load
DC-DC (All)	All Output Levels	—	—	—	—
AC-DC Multi-Output	All Output Levels	—	0.80	0.90	0.95
AC-DC Single-Output	500 watts	N/A	0.80	0.90	0.95
	>500 to 1,000 watts	0.65	0.80	0.90	0.95
	>1,000 watts	0.80	0.90	0.90	0.95

Power factor requirements do not apply where output power is less than 75 watts.

Note Energy efficiency and power factor are common benchmarks for defining power supply efficiency, but what do they mean?

An energy-efficiency rating equals how much power a power supply provides divided by how much energy is input into it. So, a server with a power supply that is 80 percent efficient uses 8 watts of every 10 watts drawn from the power source to which it is connected (8 watts / 10 watts = 80 percent). To put it another way, a computing device that requires 400 watts to function that has a power supply that is 80 percent efficient actually draws 500 watts. The less efficient a power supply is, the more power wasted—that same device equipped with a 60-percent efficient power supply would draw 667 watts to reach its needed 400 watts.

Power factor, meanwhile, is the ratio of active power to apparent power of an alternating current circuit. Expressed as a value between zero and one, it indicates how closely current and voltage are in phase with one another. A score of one indicates they are in phase; the nearer a score is to zero, the closer to 90 degrees that the current leads or lags the voltage. The lower the power factor, then, the more amperage required to provide a given amount of useful power.

Climate Savers Computing Initiative

In 2007, Google and Intel began the Climate Savers Computing Initiative, inviting makers and users of personal computers (PCs) and servers to commit to manufacture and purchase, respectively, systems that meet energy-efficiency standards of 80 percent and higher at various power load levels. More than 300 companies and organizations are, as of this writing, members of the initiative.

Organizers projected that by 2010 the program could, with sufficient participation, reduce the energy use of computers shipped that year more than 50 percent compared to 2007 and over that 3-year period save 62 billion kWh of energy and avoid 54 million tons (49 million metric tons) of carbon dioxide emissions. At 8.85 cents per kWh, that totals more than $5.5 billion in avoided energy costs.

The Climate Savers Computing Initiative introduces higher efficiency targets year over year, reaching their highest mark—deemed gold—in 2010. Table 8.3 shows efficiency targets for single-output power supply units typically used in Data Center hardware. Table 8.4 shows efficiency targets for multi-output units typically used in desktop PCs and workstations.

The Climate Savers Computing Initiative additionally calls for participating members to include an increasing proportion of energy-efficient servers and desktops among their hardware purchases and commit to using power management features on their computers.

An online database of hardware that meets the Climate Saver Server Initiative standards is available at http://tinyurl.com/56eenh.

Table 8.3 *Climate Savers Computing Initiative Power Efficiency Targets for Single-Output Power Supply Units (Servers)*

Power Load Condition	Bronze (2007)		Silver (2008)		Gold (2010)	
	Efficiency	Power Factor	Efficiency	Power Factor	Efficiency	Power Factor
20 percent	81 percent	—	85 percent	—	88 percent	—
50 percent	85 percent	—	89 percent	0.9	92 percent	0.9
100 percent	81 percent	0.9	85 percent	—	88 percent	—

Note The desktop computer efficiency levels named by both the Climate Savers Computing Initiative and the U.S. Environmental Protection Agency's Energy Star program are based upon those of the 80 Plus Program, which was launched in 2004 in the wake of a study of power supply energy efficiency and called for computer systems to be at least 80-percent efficient at various load levels.

Table 8.4 *Climate Savers Computing Initiative Power Efficiency Targets for Multi-Output Power Supply Units (Desktop Computers)*

Power Load Condition	Base (2007)		Bronze (2008)		Silver (2009)		Gold (2010)	
	Efficiency	Power Factor	Efficiency	Power Factor	Efficiency	Power Factor	Efficiency	Power Factor
20 percent	80 percent	—	82 percent	—	85 percent	—	87 percent	—
50 percent	80 percent	—	85 percent	0.9	88 percent	0.9	90 percent	0.9
100 percent	80 percent	0.9	82 percent	—	85 percent	—	87 percent	—

Efficiency Metrics for Hardware

As discussed in Chapter 2, "Measuring Green Data Centers," there are several metrics used in varying degrees across the Data Center industry to gauge energy efficiency and environmental impact of server environments. As valuable as those metrics are, however, they give only a partial view into how efficiently a Data Center's resources are used and, consequently, how green the facility is functioning.

For further insight, it's useful to also determine the relative efficiency of the specific hardware deployed in the Data Center. This is especially relevant considering that Data Center metrics universally include IT energy consumption as a factor. If you know precisely the ratio of power consumed in your Data Center by IT hardware versus facilities equipment (that is, Power Usage Effectiveness [PUE] or Data Center Infrastructure Efficiency [DCIE]), for instance) yet have no idea what sort of performance is accomplished by that IT hardware energy consumption, it's hard to draw complete conclusions about the efficiency of your company's computing activities.

A handful of metrics have been suggested to assess the energy efficiency of Data Center hardware. Although none have reached the level of discussion and usage as PUE and DCIE, consider whether they can provide insight into your Data Center's computing activities.

Energy Consumption Rating (ECR)

One proposed metric, Energy Consumption Rating (ECR), tallies the amount of energy required to move data across a networking device; specifically, the energy consumed to move one gigabit worth of line-level data per second:

$$\text{Energy Consumption Rating (ECR)} = \frac{\text{E (Energy Consumption)}}{\text{T (System Throughput)}}$$

Normally expressed in watts/gigabits per second, ECR can be calculated at a networking device's peak load or weighted to include a device's energy-saving idle mode as well.

Because networking hardware models can vary considerably, various product classes are suggested to be used in conjunction with ECR, including

- **Class 1:** Routers. These include core, edge, and multipurpose routing platforms.
- **Class 2:** WAN/Broadband aggregation devices.

- **Class 3:** Ethernet Layer 2/Layer 3 switches. Carrier-grade Ethernet switching platforms, including Data Center/large enterprise switches and desktop/generic Ethernet platforms.
- **Class 4:** Experimental. (This class is a placeholder for equipment that does not fit another class.)
- **Class 5:** Security appliances. Various security platforms including deep packet inspection (DPI), firewalls, virtual private network (VPN) gateways, and more.
- **Class 6:** Application gateways. Variable application platforms including load balancers, accelerators, and compressors.

Energy Efficiency Rating (EER)

Just as the factors used to calculate PUE can be flipped to provide the inverse DCIE metric, so too can the factors of ECR be flipped to provide the inverse known as Energy Efficiency Rating (EER):

$$\text{Energy Efficiency Rating (EER)} = \frac{\text{T (System Throughput)}}{\text{E (Energy Consumption)}}$$

ECR and EER were developed by IP performance test system provider Ixia, Juniper Networks, and Lawrence Berkeley National Laboratory and introduced in 2008.

You can find more information about ECR and EER at http://www.ecrinitiative.org/.

Space, Watts, and Performance (SWaP)

Another metric, known as Space, Watts, and Performance (SWaP), evaluates server efficiency by juxtaposing the computing performance of the hardware against its physical footprint and energy consumption. The higher a machine's SWaP rating, the greater productivity it presumably provides relative to your Data Center resources:

$$\text{Space, Watts and Performance (SWaP)} = \frac{\text{Performance}}{\text{Space} \times \text{Power Consumption}}$$

For example, if one server in your Data Center is 2U high and performs 250 operations during a given period of time while consuming 200 watts, while another server is 4U high, performs 350 operations, and consumes 400 watts, which machine is more efficient? According to SWaP, the first machine is about three times as efficient, achieving a score of 0.63 ($250 / (2 \times 200) = 0.63$) compared to the second machine's score of 0.22 ($350 / (4 \times 400) = 0.22$).

SWaP was developed by Sun Microsystems and introduced in 2005. The company offers an online calculator to compare SWaP ratings between servers, which you can find at http://www.sun.com/servers/coolthreads/swap/index.jsp#how.

Note SWaP can be an effective tool for making Data Center users think about how servers with various physical characteristics provide computing capability versus consume Data Center resources. Looking at the formula does make me wonder, though, if the relative weights of power consumption and physical footprint should be adjustable in certain circumstances.

For example, if two servers provide the same performance and one is twice the height, yet consumes half of the power of a second system, both receive the same SWaP score. Because more Data Centers these days seem to have power constraints rather than space limitations—plus power usage impacts operational costs—I consider power savings more valuable than space savings. Given the choice of the two sample machines, I would rather have the bigger, more energy-efficient server every time. Anyone managing a Data Center where cabinet space is limited but ample power is available thinks the opposite.

Applying a small multiplier to either the power consumption or space value would increase its importance in the calculation in favor of whichever resource that you consider most valuable within your Data Center.

Hardware Utilization

Just as important as choosing energy-efficient hardware to make your Data Center green is using them efficiently after they are installed. The more of a hardware device's capacity that you use, known as its *utilization*, the better. Having the most energy-efficient servers and storage devices available on the market provides little green value if each of those systems are lightly utilized, and you end up operating—and therefore powering and cooling—several of them.

Think of utilization and capacity in terms of a motor vehicle. The more seats that are filled with passengers the fewer trips you have to make, the less gas that you consume, and ultimately the fewer automobiles you need to provide the necessary transportation. As with Data Center hardware, the most fuel-efficient passenger car isn't particularly green if it is used only to transport one person and its several other seats are vacant.

Industry approximations of average Data Center server and storage utilization vary significantly—optimistic estimates suggest 40 percent utilization, others estimate as low as just 5 percent. Chapter 9, "Greening Your Data Center Through Consolidation, Virtualization, and Automation," offers methods to improve the utilization rates of your server and storage devices, potentially to as high as 80 percent.

Beyond Energy Consumption and Utilization

Although energy efficiency is arguably the most important factor for evaluating how green a server, networking device, or storage unit is, it's not the only one. Other elements contribute to the environmental impact of Data Center hardware and are relevant to consider when making purchasing decisions:

- **Cooling efficiency:** Aside from the power that they consume directly, Data Center machines indirectly use additional energy due to their need to be kept cool. Servers that are optimized for cooling—generating less heat and oriented with their air intake in the front and exhaust venting in the back, thereby matching the hot- and cold-aisle designs prevalent in modern Data Centers—reduce that secondary energy consumption.
- **Materials:** Limit buying hardware made of materials that are bad for the environment or require large quantities of resources to produce. The European Union has restricted the use of six materials—lead, cadmium, mercury, hexavalent chromium, and flame retardants polybrominated biphenals and polybrominated diphenyl ether—in electrical and electronic equipment since mid-2006 and it's likely that similar regulations will ultimately be adopted in other countries in the future. The Restriction of Hazardous Substances [RoHS] Directive does allow lead-based solders through 2010 for servers and storage arrays and indefinitely for networking infrastructure equipment.

Note Material Declaration Data Sheets (MDDS) that outline what substances a product is made of—and often what hazardous materials the item does not include—can be obtained from hardware manufacturers. Many offer them online on their company website.

- **Capability to upgrade:** What happens when your Data Center hardware ages? Systems that enable you to replace key components with new ones are greener because you can extend their useful life using fewer materials than if you fully decommissioned and replaced them.
- **Capability to recycle:** The more components that can be kept out of the landfill when a piece of Data Center hardware does reach the end of its useful life, the better. Some manufacturers facilitate this process by promoting the return of their old systems. (Chapter 10, "Greening Other Business Practices," discusses green considerations and strategies for dealing with e-waste.)
- **Small form factor:** Smaller hardware involves fewer manufacturing materials and requires less Data Center supporting infrastructure. If your company's processing needs can be supported by a given number of 1U servers instead of the same number of 4U servers, for instance, your Data Center could theoretically be built at one fourth the size, reducing everything from the quantities of server cabinets to the lengths of structured cabling to the overall amount of building materials. This

assumes your Data Center has sufficient power and cooling capacity to support significant equipment density and, therefore, take advantage of the small form factor, of course. Be sure to consider this in the context of groups of hardware as well—a larger machine can be a greener choice when it takes the place of multiple smaller devices that, taken together, have a larger footprint (and perhaps consume more energy).

Data Centers typically house hundreds or even thousands of pieces of hardware, and many of those machines come from a relatively small number of vendors. If you are going to purchase multiple systems from a manufacturer within a short period of time, ask whether it can consolidate how the items are packaged. If everyone did this, manufacturers would consume fewer resources (and spend less money) for packaging and accessories, and customers would end up with fewer materials that are often discarded.

If you order 10 of a given server model, it's doubtful that you need more than one copy of its installation manual, for instance. Cables and adapters that are packed with many hardware models are also superfluous if those items are already stocked in your server environment—a common practice in rooms where color-coding schemes are used.

Note It's also a good idea to ask manufacturers of the consumable items used in your Data Center for consolidated packaging options. I once placed a bulk order for a few hundred patch cords and when they arrived, I was startled to discover each cable had been individually wrapped in plastic—a complete waste of material, not to mention of my time to open every wrapping.

The most wasteful packaging of Data Center items I have ever seen involved shelving for server cabinets. During my first week working for Cisco, I was asked to help clean up and organize a caged portion of a receiving dock that contained a jumble of Data Center items and miscellaneous items that were more suitable for an office supply closet. The biggest objects in the cage were a dozen or so boxes, each about the size of a kitchen refrigerator and extremely heavy, that contained shelves for use with four-post server cabinets.

To my surprise, each box contained only about 15 shelves. Despite being made of sturdy steel, each shelf was thoroughly wrapped in packing foam, capped on all four corners with cardboard, nestled within an individual box and then all those boxes were in turn stacked inside an outer box. The shelves seemed quite damage resistant, yet had been wrapped as if constructed of flimsy, breakable material. After I removed all the packaging, the shelves occupied maybe one-third of the space of the outer box, freeing up considerable room in the cage.

How Budget Process Impacts Hardware Choices

Although it might not be obvious, the manner in which funds are budgeted at your company could be hindering how green your Data Center is.

In many businesses, the IT department has a defined budget for the purchase of computing hardware. The funding is typically available for a finite period of time, such as a fiscal quarter or perhaps an entire year. For a large company, smaller groups within the IT organization might each have their own budget, but otherwise the conditions are the same. Because one of IT's main measures of success is how well it provides computing services for the company and because IT personnel know their funding will go away if they don't spend it—and perhaps even resulting in a smaller budget allocation next period—there is a natural tendency to buy the highest-performing machines possible. This approach can seem harmless. If some of the hardware capabilities exceed what is needed, who cares?

Hardware performance does come with a price, though. The faster processing a server can do, typically the more power it consumes and more heat it generates. If your company continuously buys more powerful servers than it requires, it's putting excess power and cooling demand upon the Data Center. It's much like buying a high-performance sports car even though you need to drive only in town where speed limits are low. You rarely see the benefit of the car's excellent acceleration (the server's greater processing) over that of an average vehicle, but you suffer with its poorer gas mileage (excess power consumption) every day.

The ongoing cost of higher-performing hardware can be more difficult to detect at companies where utility bills are paid by the facilities organization. With no view into those costs, the IT department has no idea of how much impact it has upon Data Center power consumption and no incentive to change its buying habits.

Communication between the IT and facilities departments can improve this situation. At minimum, have IT management see the monthly power bills associated with the Data Center. Give them visibility into the energy consumption associated with their hardware choices and an understanding of the finite resources of your server environment.

The ultimate extension of this is to establish a chargeback model for Data Center resources. Under this approach, you operate your server environment much like a colocation facility, charging clients for using Data Center capacity. Although you're not looking to make a profit off of the chargeback model the way that an external colocation business does, by assigning a monetary value to your Data Center resources and requiring departmental groups to pay when they use them can influence behavior. Groups are less likely to overstate their capacity needs "just in case" they need it, and with a chargeback policy in place, they even have a direct incentive to consume fewer Data Center resources.

Note I vividly recall the year that Cisco began to charge a small fee to internal groups for the building floor space they occupied. Managers who had absolutely required storerooms and to own various functional work areas suddenly didn't feel they needed the space so badly anymore. Whereas my team had once had to compete to obtain sufficient space in which to keep Data Center-related equipment, plenty became available.

Years ago Data Center chargeback models involved rack units or floor space because physical dimensions, and available cabinets were a server environment's leading finite capacity. Today, chargeback models focus on different critical Data Center resources. For Data Centers housing conventional servers, it's power. For Data Centers in which systems are virtualized, they're increments of CPU usage, RAM, and storage (see Chapter 9 for details).

Consider the user behavior that is promoted by the different models. Charging based upon cabinet space encourages Data Center clients to buy hardware with the smallest footprint and to install those systems as tightly as possible. Such high-density machine installations can lead to hot spots and exceed the power budget of cabinet locations. Charging based upon power usage puts the value on maximizing the use of energy, leading to reduced consumption and emphasizing greater processing efficiency per amp. Charging based upon processing and storage encourages clients to maximize the use of those resources.

During the 2000s, many colocation facilities changed their pricing structure so that customers began to be charged for energy usage rather than cabinet space. Some continue to charge for overall floor space but also have a hard cap on how much power is provided to the floor space that a client leases. In that case, obtaining more power requires leasing more space.

Note I worked at Syracuse University before coming to Cisco, and for several months some of the on-campus cafeterias had a salad bar in which patrons could buy a salad for a set price based upon the size of the bowl. I became very good at constructing a tall, heaping salad—not unlike a Data Center manager occupying every rack unit of cabinet space with hardware.

One day when I went to pay for a salad that I had so carefully constructed, the cashier weighed it and charged me on a cost-per-ounce basis. Now calculated based upon the resources that I was consuming, the price doubled!

I still ate salads at the cafeteria, but nowhere near as frequently, and when I did I focused more on which ingredients I truly wanted.

Idling Servers

Data Centers are massive consumers of electricity, not just because they power and cool thousands of pieces of high-performing computing hardware, but also because they power and cool those systems constantly—24 hours a day, 7 days a week, 365 days a year, year after year. Imagine how much higher the power bills for your home would be if you kept all your electronic appliances and lights on around the clock.

Even when Data Center machines are idle, they draw a significant amount of energy—various estimates place that consumption as 30 percent, 50 percent, or nearly 70 percent of when a system is at peak load. Because hardware serves no useful purpose while idling, how about turning them off and avoiding that energy usage?

Several studies have been conducted to demonstrate the feasibility and energy savings of deactivating idle Data Center hardware and then reactivating them when demand warrants:

- Microsoft tracked usage patterns such as login rates, connection counts, and connection failures on its instant messaging service, Windows Live Messenger, over a 45-day period along with server performance data such as CPU usage, memory usage, and power consumption. The company then developed server provisioning and load dispatching algorithms, adjusting parameters such as how many logins are routed to various servers, in hopes to find a balance point where energy usage could be reduced without notable impact upon connection services for users. The study, *Energy-Aware Server Provisioning and Load Dispatching for Connection-Intensive Internet Services*, published in 2008, reported energy savings of 20 percent to 30 percent with only minor impacts to user experience.
- In a separate study, Microsoft determined that spinning down storage disks when they are idle can save 28 percent to 36 percent in energy consumption. By redirecting blocks normally written to one volume to other storage systems, a technique Microsoft calls *write off-loading*, idle time can be increased and raise energy saving to 45 percent to 60 percent. The study, *Write Off-Loading: Practical Power Management for Enterprise Storage*, published in 2008, analyzed 36 volumes containing 179 disks on 13 servers during a 1-week period and then conducted write off-loading techniques on a test bed of 56 disks on four servers.
- A company that sold software that can automatically power servers off or on according to various preset conditions, performed a 6-week study in which 89 servers were power-cycled a cumulative 3,500 times, simulating 6.3 hours per day of idle time. Having the systems offline saved approximately 26 percent of the power typically used by those system, a total about 2,400 kWh, according the documented study, *Resource Optimization through Active Power Management*. The study was performed by Cassatt, whose assets have since been purchased by IT software firm CA.

Despite the obvious potential to reduce energy consumption by powering down hardware, many Data Center managers are hesitant to do so because they are concerned about possible system malfunctions due to repeated power-cycling. Energy savings are desirable for a Data Center but not as important as reliable uptime. The study was performed by Cassatt, whose assets have since been purchased by IT software firm CA.

Note I attended the Data Center Energy Summit sponsored by the Silicon Valley Leadership Group in 2008, at which the results of the aforementioned Cassatt case study were first presented. Although audience interest in the results was strong, there were some skeptical comments based on the relatively limited times that machines were powered off and on—3,500 cycles of 89 machines equals about 40 times apiece.

Although not documented in a published paper, former Cassatt vice president of product engineering Vinay Pai blogged in 2007 about a study performed in one of the company's engineering labs in which 123 servers were each power-cycled about once a day for a 5-month period—for a total of 18,826 times—without a single power supply or disk drive failure in the ensuing 2 years.

More than half of the servers were HP machines, and the remainder were a mix of Dell, IBM, and Sun systems according to Pai's blog entry, titled "Yes It's Still Safe to Power Off and Power On That Server."

"So if you're still afraid to power down that server, don't worry!" Pai wrote. "Power supplies and hard drives are very reliable these days. From several different studies we've seen that power supplies hold up quite well from (and are even designed for) power cycling."

Eliminating Less Efficient Hardware

As you focus on purchasing new, greener hardware, don't overlook the systems already in your Data Center. Companies often have ample processes in place concerning obtaining and installing new servers, but few concerning the decommissioning or replacement of older ones. Avoid allowing outdated hardware that is poorly utilized and not energy-efficient to consume your Data Center's valuable resources.

It's a good idea to regularly check the machines in your server environment, especially legacy machines, to determine their utilization. If a piece of equipment isn't used, turn it off, and remove it. If the equipment is used but only lightly, determine whether the machine can and should be upgraded to a newer configuration that is more energy-efficient.

Newer systems are generally more energy-efficient, so it can sometimes make financial and environmental sense to upgrade to new models—much like trading in an old vehicle for an economy car when fuel prices are high.

Cisco Product Efficiency Calculator Don't be shy about asking hardware manufacturers for information about how energy-efficient their models are that you use (or are considering using) in your Data Center. This can help you choose greener systems that won't need to be upgraded for an extended period of time.

As part of its Efficiency Assurance Program, Cisco offers an online tool that enables you to determine the power usage, electrical efficiency, and energy costs of its devices. You can even set certain variables in the tool to reflect specific conditions within your particular Data Center.

Users enter local electric rates (per kWh), the type and quantity of various networking devices, an estimate of how many watts are consumed cooling the Data Center relative to those consumed by IT hardware, and an estimate of the network devices' overall utilization. In return, the Product Efficiency Calculator shows the nominal power draw of the networking gear, typical thermal energy loss, and the total annual electrical costs. The tool also provides a chart of the electrical efficiency of the specified networking device at various power loads.

Figure 8.1 shows the Cisco Product Efficiency Calculator, profiling a single Nexus 7000 model.

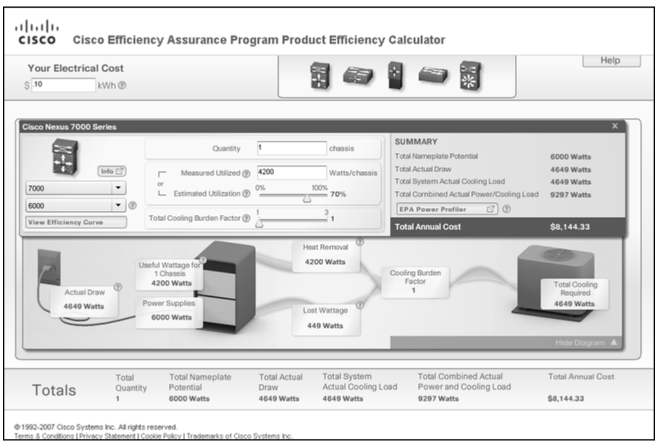

Figure 8.1 *Cisco Product Efficiency Calculator*

Figure 8.2 shows the power efficiency of a Nexus 7000 power supply, at various load levels, according to the Cisco online Product Efficiency Calculator.

The Cisco Efficiency Assurance Program is located at www.cisco.com/go/efficiency.

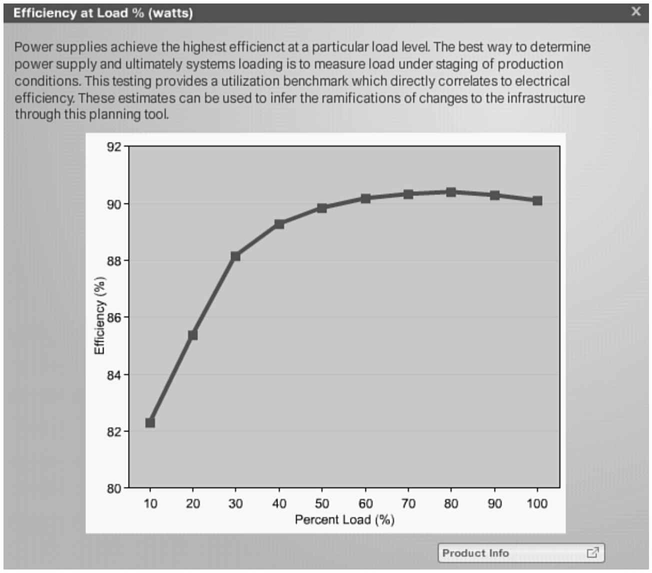

Figure 8.2 *Power Efficiency of a Cisco Nexus 7000 Power Supply*

Summary

Hardware within your Data Center impacts how green your facility is as well its operational costs. When considering equipment purchases, look for an overall configuration that meets your computing needs, today and in the foreseeable future, and provides the most processing per kilowatt.

The U.S. Environmental Protection Agency's Energy Star Specification for Servers calls for hardware to meet minimum energy-efficiency ratings with their power supplies, to stay within certain power consumption limits when idle, and to facilitate data gathering about operational conditions such as power usage, inlet air temperature, and processor utilization.

The Climate Savers Computing Initiative, which calls for energy-efficiency standards of 80 percent and higher for desktop PCs and servers, has the goal to cut energy usage of new computing equipment in half between 2007 and 2010, saving 62 billion kWh of energy, 54 million tons (49 million metric tons) of carbon dioxide, and more than $5.5 million in energy costs.

Various metrics have been proposed to capture IT hardware efficiency, including ECR (dividing energy consumption by system throughput), EER (dividing system throughput by energy consumption), and SWaP (dividing performance by space times power consumption).

Utilization, that is how much of a server's, networking device's, or storage unit's capacity is used, is another indicator of efficiency within your server environment.

Other factors to consider when determining how green a piece of Data Center hardware is include cooling efficiency, what materials the machine is made of, how easy it is upgrade with new components, how easy it is to recycle the machine, and the machine's size. Further limit the environmental impact of your Data Center hardware by requesting that packaging and other associated materials be consolidated by the manufacturer.

The budget cycle used by many businesses in which the IT department is given money for Data Center hardware on a quarterly, use-or-lose basis promotes the acquisition of more powerful machines that consume more energy than other models. This is especially problematic in businesses where IT has no visibility into company power bills whereas facilities, which pays those bills, has no input into hardware selection. Implementing a chargeback model based upon power consumption provides visibility for IT and rewards energy reduction and efficiency within the Data Center.

Idle servers consume 30 percent to nearly 70 percent of what they do at peak load. Various software applications can shut down Data Center hardware when not in operation, reactivating them at a later time. This practice can provide significant power savings if you are comfortable with regularly power-cycling such equipment.

Remove old Data Center hardware that is less energy-efficient and poorly utilized to optimize energy consumption.