Chapter 2

Energy-Efficient and Ecologically Friendly Data Centers

Pay your speeding tickets and parking fines or spend your savings elsewhere.

In this chapter you will learn:

- How to identify issues that affect the availability of power, cooling, and floor space
- Why achieving energy efficiency is important to sustain growth and business productivity
- How electrical power is generated, transmitted, and used in typical data centers
- How electrical power is measured and charges determined

By understanding fundamentals and background information about electricity usage as well as options and alternatives including rebates or incentives, IT data centers can deploy strategies to become more energy-efficient without degrading service delivery. Reducing carbon footprint is a popular and trendy topic, but addressing energy efficiency—that is, doing more work with less energy—addresses both environmental and business economic issues. The importance of this chapter is that near-term economic as well as environmental gains can be realized by making more efficient use of energy. By reducing energy consumption or shifting to a more energy-efficient IT model, businesses can reduce their operating expenses and enable more useful work to be done per dollar spent while improving service delivery. This chapter looks at challenges with electrical power for IT data centers as well as background information to help formulate effective strategies to become energy-efficient.

2.1 Electric Power and Cooling Challenges

Asking the right questions can help you to close the "green gap" and address **power, cooling, floor space, and environmental (PCFE)** issues. That is, insight into how infrastructure resources are being used to meet delivery and service levels is critical. For example, instead of asking whether there is a green mandate or initiative, try asking the following questions:

- Does the data center have a power issue or anticipate one in the next 18–24 months?
- Does the IT data center have enough primary and backup power capacity?
- Is there enough cooling capacity and floor space to support near-term growth?
- How much power does the data center consume?
- How much of that power goes for cooling, lighting, and other facility overhead items?
- How much power is used by servers, storage, and networking components?
- Is power constrained by facility, local substation, or generating capability limits?
- What floor space constraints exist, and is there adequate cooling capabilities for growth?
- Can energy usage be aligned with the level of service delivered or amount of data stored?
- What hazardous substances and materials exist in the data center?

Closing the green gap is important in that core IT PCFE issues can be addressed with positive environmental and economic benefits. For example, building on the previous questions, common PCFE-related pain points for many IT data centers include:

- A growing awareness of green and environmentally friendly issues and topics

- The need to remove heat from IT equipment and the power required for this cooling
- Excessive power consumption by older, less energy-efficient technology
- Insufficient primary or standby power
- Rising energy costs and insufficient availability of power
- Lack of sufficient floor space to support growth and use of heavier and denser equipment
- Aging and limited HVAC (heating, ventilating, and air conditioning) capabilities
- Disposing of older technology in compliance with recycling regulations
- Complying with environmental health and safety mandates
- Improving infrastructure and application service delivery and enhancing productivity
- Doing more with less—less budget, less head count, and more IT equipment to support
- Support applications and changing workloads with adaptive capabilities

The available supply of electricity is being impacted by aging and limited generating and transmission capabilities as well as rising fuel costs. While industries such as manufacturing consume ever more electrical power, IT data centers and the IT equipment housed in those habitats require continued and reliable power.

IT data centers rely on available power and transmission capabilities, which are being affected by rising fuel costs and increasing demands as shown in Figure 2.1. The U.S. Environmental Protection Agency (EPA) estimates that with no changes, U.S. IT data center electric power consumption will jump to 3% of the U.S. total by 2010–2012. IT data centers require ever more power, cooling, and physical floor space to accommodate the servers, storage, and network components necessary to support growing application demands. In an era of growing environmental awareness, IT data centers, information factories of all sizes, and enterprise data centers in particular have issues and challenges pertaining to

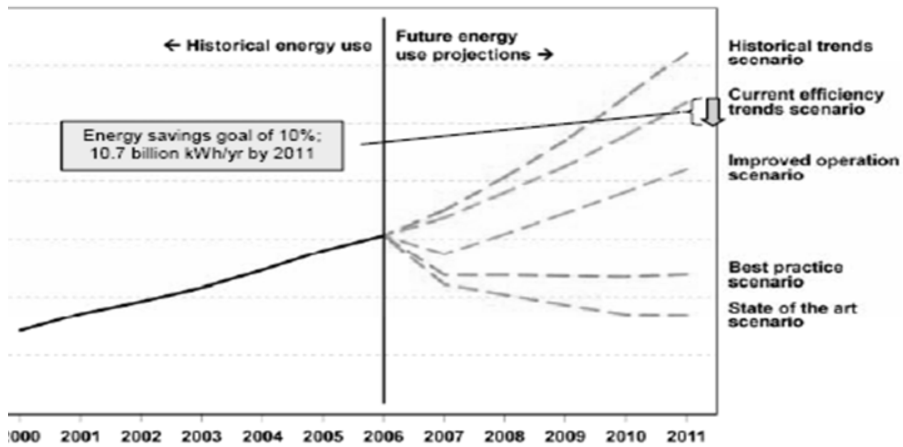

Figure 2.1 Projected Electricity Use, 2007–2011 (Source: U.S. Environmental Protection Agency)

power, cooling, floor space, and greenhouse gas emissions as well as "clean" disposal of IT equipment.

Data center demand for electrical power is also in competition with other power consumers, leading to shortages and outages during peak usage periods. There are also increasing physical requirements for growing data centers in the form of more servers, storage, and network components to support more IT and related services for business needs. Other pressing issues for IT data centers are cooling and floor space to support more performance and storage capacity without compromising availability and data protection.

Thus, if a data center is at its limit of power, and if the data center needs to increase processing and storage capabilities by 10% per year, a corresponding improvement in efficiency of at least the same amount is required. Over the past decade or so, capacity planning has been eliminated in many organizations because of the lowering cost of hardware; however, there is an opportunity to resurrect the art and science of capacity planning to tie power and cooling needs with hardware growth and to implement data center power demand-side management to ensure sustained growth.

Figure 2.2 shows typical power consumption and energy usage of typical data center components. With a current focus on boosting performance

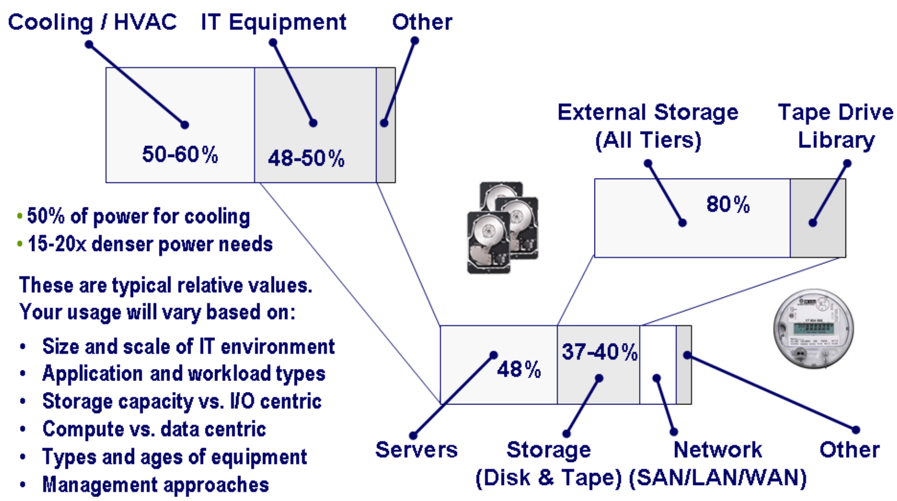

Figure 2.2 Average IT Data PCFE Consußmption (Source: www.storageio.com)

and reducing power consumption for servers and their subsequent cooling requirements, the PCFE focus will shift to storage. Even with denser equipment that can do more work and store more information in a given footprint, continued demand for more computing, networking, and storage capability will keep pressure on available PCFE resources. Consequently, addressing PCFE issues will remain an ongoing issue, and, thus, performance and capacity considerations for servers, storage, and networks need to include PCFE aspects and vice versa.

2.2 Electrical Power—Supply and Demand Distribution

Adequate electrical power is often cited as an important IT data center issue. The reasons for lack of available electrical power can vary greatly. Like data networks, electrical power transmission networks, also known as the power grid, can bottleneck. For example, there may be sufficient power or generating capabilities in your area, but transmission and substation bottlenecks may prevent available power from getting where it is needed. For example, consider the following scenarios:

- Power is available to the facility, but the facility's power distribution infrastructure is constrained.
- Power is available to the facility, but standby or backup power is insufficient for growth.
- Power is available in the general area, but utility constraints prevent delivery to the facility.
- Power costs are excessive in the region in which the IT equipment and facilities are located.

General factors that affect PCFE resource consumption include:

- Performance, availability, capacity, and energy efficiency (PACE) of IT resources
- Efficiency of HVAC and power distribution technologies
- The general age of the equipment—older items are usually less efficient
- The balance between high resource utilization and required response time
- Number and type of servers, type of storage, and disk drives being used
- Server and storage configuration to meet PACE service-level requirements

There is a correlation between how IT organizations balance server, storage, and networking resources with performance and capacity planning and what electrical **generating and transmission (G&T)** utilities do. That correlation is capacity and demand management. For G&T utilities, building or expanding existing generating and transmission facilities are cost- and time-consuming activities, so G&Ts need to manage the efficient use of resources just as IT managers must. G&Ts rely on supervisory control and data acquisition management systems to collect data on G&T components and to enable real-time management of these resources.

Thus, there are many similarities between how IT centers manage resources using simple network management protocol traps and alerts together with capacity planning and how the G&T industry manages its

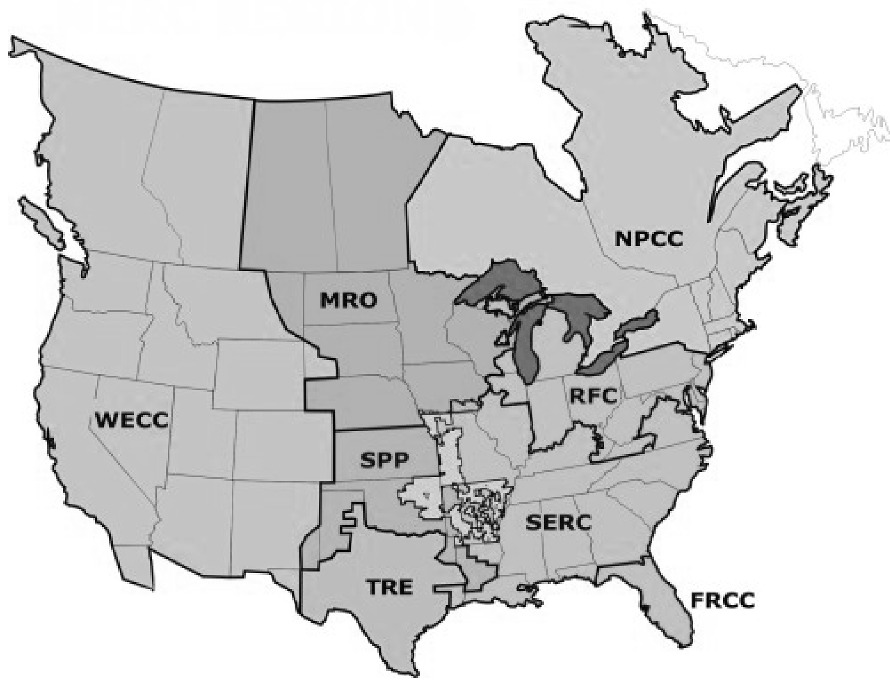

Figure 2.3 North American Electrical Power Grid (Source: www.nerc.com)

resources. Think of the power plants as servers and the transmission grid as a network. Like IT data centers, which have historically used performance and capacity planning to maximize their resources versus the expense of buying new technologies, the power companies do so on an even larger scale. Power companies have a finite ceiling to the amount of power they can provide, and data centers have a ceiling on the amount of power they can consume based on available supply.

Figure 2.3 shows how electrical power is managed in different regions of the United States as part of the North American electrical power grid. The U.S. power grid is, as is the case in other parts of the world, a network of different G&T facilities that is able to shift available power to different areas within a region as needed.

Figure 2.4 shows a simplified example of how the G&T power grid works, including generation facilities (power plants), high-voltage distribution lines, and lower-voltage distribution from local substations. IT data centers typically receive electric power via local distribution from one or more substation power feeds. As in an IT data network, there are many

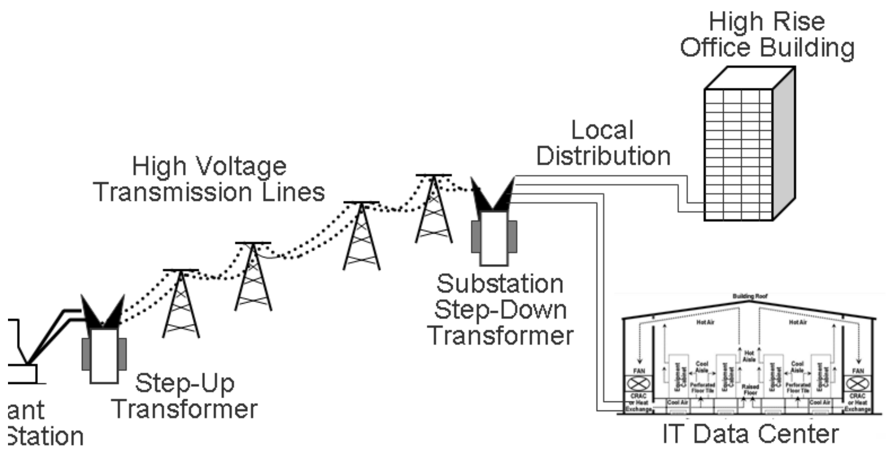

Figure 2.4 Electrical Power G&T Distribution

points for possible bottlenecks and contention within the G&T infrastructure. As a result, additional power may not be available at a secondary customer location; however, there may be power available at a primary substation or at a different substation on the G&T grid.

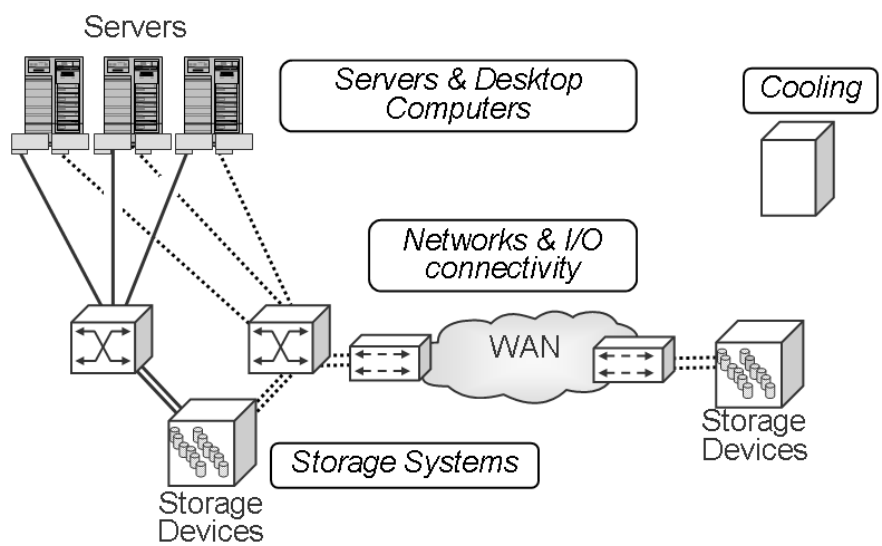

Figure 2.5 IT Data Center Consumers of Power

Once electricity is supplied to an IT data center, various devices, some of which are shown in Figure 2.5, consume the power. Additional items not shown in Figure 2.5 that consume power include HVAC or computer room air conditioning, power conversion, and distribution, along with general facility components such as battery chargers and lights.

2.3 Determining Your Energy Usage

When was the last time you looked at your business energy bill or your home electric bill? If you have not done so recently, look at or ask someone about what your facilities energy bill looks like. Also, look at your home energy bill and see how many kilowatt-hours you used, what the base energy rates are, what surcharges and other fees are assessed, and other information. Once you have reviewed your energy bill, can you determine where, when, and how electrical power is being used?

There are different approaches to determining energy usage. One is to take a macro view, looking at how much energy a data center consumes in total and working down to individual devices. The opposite approach is to work backwards from a micro view of components in a device and total the results across the various devices. Measured electricity usage can be obtained from utility bills, meters, or facilities management software tools, depending on what is available. Sources for estimating power consumption on a device or component level are vendor-supplied information, including equipment faceplates, and site planning guides. Other sources for determining energy usage are power meters and analyzers as well as power distribution devices that can also measure and report on power usage.

Electrical power is typically charged by utilities based on the number of kilowatt-hours used or the number of 1,000 Watt (W) of electricity used. For example, a server that draws 500 W consumes 0.5 kWh, or a storage device that consumes 3,500 W when being actively used consumes 3.5 kWh. Note that while energy usage varies over time and is cumulative, energy from a utility billing standpoint is based on what is used as of an hour. That is, if a server draws 500 W, its hourly energy bill will be 500 W or 0.5 kWh, as opposed to 500 × 60 seconds × 60 minutes. Likewise, electrical power generation is quoted in terms of kilowatt-hours or mega (million) watt-hours (MWh) as of a given point in time. For example, an 11-MW power plant is capable of producing 11,000 kWh at a given point in time, and if usage is constant, the energy is billed as 11,000 kWh.

Typically, energy usage is based on metered readings, either someone from the utility company physically reading the meter, remote reading of the meter, or, perhaps, estimated usage based on historical usage patterns, or some combination. Electric power is charged at a base rate (which may vary by location, supply, and demand fuel sources, among other factors) per kilowatt-hour plus fuel surcharges, peak demand usage surcharges, special fees, applicable commercial volume peak usage minus any applicable energy saver discounts. For example, for voluntarily reducing power consumption or switching to standby power generation during peak usage periods, utilities may offer incentives, rebates, or discounts.

IT technology manufacturers provide information about electrical energy consumption and/or heat (Btu per hour) generated under a given scenario. Some vendors provide more information, including worst-case and best-case consumption information, while others provide only basic maximum breaker size information. Examples of metrics published by vendors and that should be visible on equipment include kilowatts, kilovolts, amperage, volts AC or Btu. Chapter 5 discusses various PCFE-related metrics and measurements, including how to convert from watts to Btu and from amperes to watts.

To calculate simple energy usage, use the values in Table 2.1, selecting the energy costs for your location and the number of kilowatt-hours required to power the device for one hour. For example, if a server or storage device consumes 100 kWh of power and the average energy cost is 8¢/kWh, energy cost is $70,100 annually. As another example, a base rate for 1 kWh might be 12¢/kWh but 20¢/kWh for usage over 1,000 kWh per month. Note that this simple model does not take into consideration regional differences in cost, demand, or availability, nor does it include surcharges, peak demand differentials, or other factors. The model also does not differentiate between energy usage for IT equipment operation and power required for cooling. The annual kWh is calculated as the number of kWh × 24 × 365.

A more thorough analysis can be done in conjunction with a vendor environment assessment service, with a consultant, or with your energy provider. As a start, if you are not in the habit of reading your monthly home energy bill, look to see how much energy you use and the associated costs, including surcharges and fees.

Table 2.1 Example Annual Costs for Various Levels of Energy Consumption

Hourly Power Use (kWh)	**5¢/ kWh**	**8¢/ kWh**	**10¢/ kWh**	**12¢/kWh**	**15¢/kWh**	**20¢/kWh**
1 kWh	$438	$701	$806	$1,051	$1,314	$1,752
10 kWh	$4,380	$7,010	$8,060	$10,510	$13,140	$17,520
100 kWh	$43,800	$70,100	$80,600	$105,100	$131,400	$175,200

2.4 From Energy Avoidance to Efficiency

Given specific goals, requirements, or objectives, shifting to an energy-efficient model can either reduce costs or enable new IT resources to be installed within an existing PCFE footprint. Cost reductions can be in the form of reducing the number of new servers and associated power and cooling costs. An enabling growth and productivity example is to increase the performance and capacity, or the ability to do more work faster and store more information in the same PCFE footprint. Depending on current or anticipated future power and/or cooling challenges, several approaches can be used to maximize what is currently in place for short-term or possibly even long-term relief. Three general approaches are usually applied to meet the various objectives of data center power, cooling, floor space, and environmental aims:

- Improve power usage via energy efficiency or power avoidance.
- Maximize the use of current power—do more with already available resources.
- Add additional power, build new facilities, and shift application workload.

Other approaches can also be used or combined with short-term solutions to enable longer-term relief, including:

- Establish new facilities or obtain additional power and cooling capacity.

- Apply technology refresh and automated provisioning tools.
- Use virtualization to consolidate servers and storage, including thin provisioning.
- Assess and enhance HVAC, cooling, and general facility requirements.
- Reduce your data footprint using archiving, real-time compression and de-duplication.
- Follow best practices for storage and data management, including reducing data sprawl.
- Leverage intelligent power management such as MAID 2.0-enabled data storage.
- Use servers with adaptive power management and 80% Plus efficient power supplies.

Virtualization is a popular means of consolidating and eliminating underutilized servers and storage to reduce cost, electricity consumption, and cooling requirements. In their place, power-efficient and enhanced-performance servers and storage, including blade centers, are being deployed to support consolidated workloads; this is similar to what has historically been done in enterprise environments with IBM mainframe systems. However, for a variety of reasons, not all servers, storage, or networking devices lend themselves to being consolidated.

Some servers and storage as well as network devices need to be kept separate to isolate different clients or customers, different applications or types of data, development and test from production, online customer-facing systems from back-end office systems, or for political and financial reasons. For example, if a certain group or department bought an application and the associated hardware, that may prevent those items from being consolidated. Department turf wars can also preclude servers and storage from being consolidated.

Two other factors that can impede consolidation are security and performance. Security can be tied to the examples previously given, while application performance and size can have requirements that conflict with those of applications and servers being consolidated. Typically, servers with applications that do not fully utilize a server are candidates for consolidation. However, applications that are growing beyond the limits of a single dual-, quad-, or multi-core processor or even cluster of servers do not lend themselves to consolidation. Instead, this latter category of servers and applications need to scale up and out to support growth.

Industry estimates and consensus vary from as low as 15% to over 85% in terms of actual typical storage space allocation and usage for open systems or non-mainframe-based storage, depending on environment, application, storage systems, and customer service-level requirements. Low storage space capacity usage is typically the result of one or more factors, including the need to maintain a given level of performance to avoid performance bottlenecks, over-allocation to support dynamic data growth, and sparse data placement because of the need to isolate applications, users, or customers from each other on the same storage device. Limited or no insight as to where and how storage is being used, not knowing where orphaned or unallocated storage is stranded, and buying storage based on low cost per capacity also contribute to low storage space capacity utilization.

The next phase of server virtualization will be to enhance productivity and application agility in order to scale on a massive basis. Combined with clustering and other technologies, server virtualization is evolving to support scaling beyond the limits of a single server—the opposite of the server consolidation value proposition. Similarly, server virtualization is also extending to the desktop to facilitate productivity and ease of management. In both of these latter cases, transparency, emulation, and abstraction for improved management and productivity are the benefits of virtualization.

2.5 Energy Efficiency Incentives, Rebates, and Alternative Energy Sources

Carbon offsets and emissions taxes have their place, particularly in regions where legislation or regulations require meeting certain footprints. In such locations, a business decision can be to do an analysis of paying the emissions tax fee to comply near term versus cost to comply long term. In other words, pay carbon offsets or get money back and achieve efficiency.

Some U.S. energy utilities provide incentives and rebates for energy efficiency and/or use of renewable energy. The programs vary by utility, with some being more advanced than others, some more defined, and some more customer oriented. Some programs provide rebates or energy savings, while others provide grants or funding to conduct energy efficiency assessments or

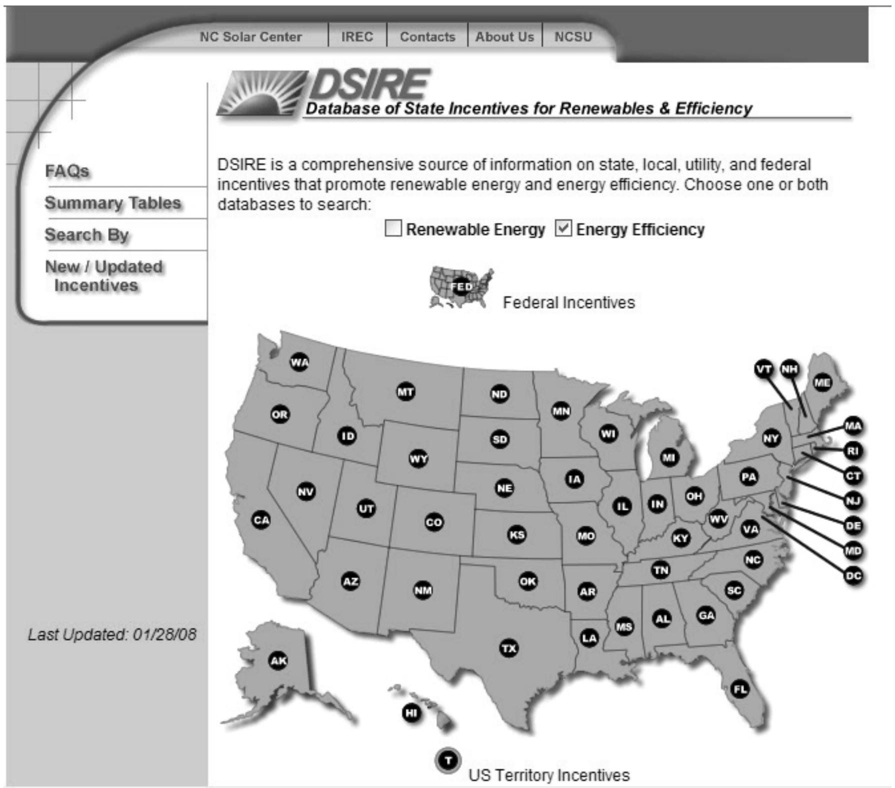

Figure 2.6 Database of State Incentives for Renewables & Efficiency (DSIRE)

make infrastructure and facilities changes. In general, utilities do not differentiate between an IT data center and vendor development and testing lab facilities or consider the size of the data center.

Pacific Gas and Electric (PG&E) has been a pioneer in energy rebates incentives for IT data centers. PG&E has programs targeted toward energy demand-side management for various localities and industry sectors. Other energy utilities also leverage demand-side management as part of their capacity planning and performance management of their resources: energy generation via their G&T facilities. A group of utilities led by PG&E has created a Consortium for Energy Efficiency (CEE) to exchange and coordinate ideas and to further develop specific programs to address IT data center power issues. Another venue is DSIRE (Database of State Incentives for Renewables & Efficiency), which is a portal that provides information on various available energy efficiency as well as

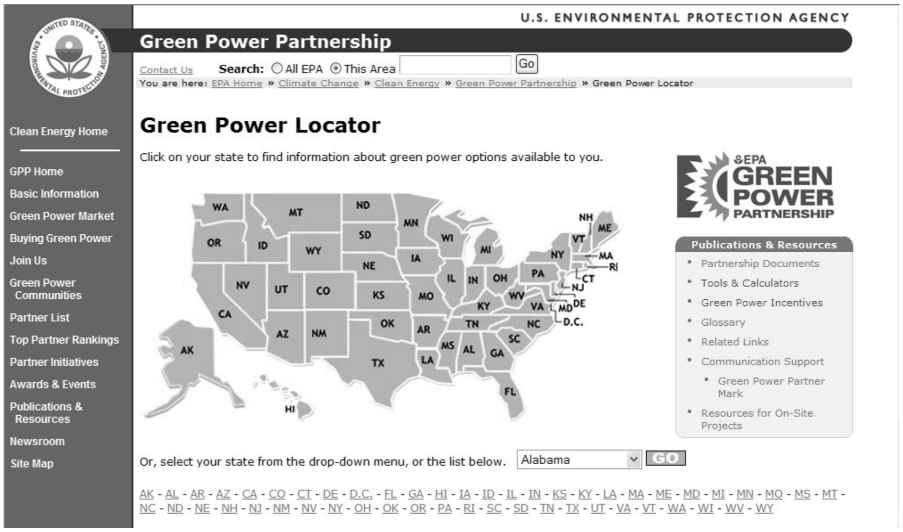

Figure 2.7 The EPA Power Portal (Source: U.S. Environmental Protection Agency)

renewable energy programs across the United States. As shown in Figure 2.6, additional information about such incentives on a state-by-state basis can be found at www.dsireusa.org.

The EPA has many programs associated with power and energy that combine an environmental viewpoint with a perspective on sustaining supply to meet demand. Examples of EPA programs include Energy Star, Green Power (www.epa.gov/grnpower), and others. The EPA Green Power portal shown in Figure 2.7 provides information on various programs, including alternative and green power sources, on a state-by-state basis. Other agencies in different countries also have programs and sources of information, for example, the Department for Environment, Food and Rural Affairs (DEFRA; www.defra.gov.uk), in the United Kingdom. In Canada, Bullfrog Power has a portal (www.bullfrogpower.com) that provides information on green and alternative power for homes and businesses.

Fossil fuels for primary and secondary electric power generation are coal, oil and gas (natural gas, liquefied propane [LP] gas, gasoline or aviation fuel, diesel). Alternative and renewable sources for electricity generation include biomass (burning of waste material), geothermal, hydro, nuclear, solar, wave and tidal action, and wind. Part of creating an energy-efficient and environmentally friendly data center involves leveraging different energy sources for electricity. For example, a local power utility can

Table 2.2 Some Relevant Standards, Regulations, and Initiatives

Abbreviation or Acronym	Description
DOE FEMP	U.S. Department of Energy Federal Energy Management Program
ECCJ	Energy Conservation Center Japan (www.eccj.or.jp)
ELV	End of Life Vehicle Directive (European Union)
Energy Star	U.S. EPA Energy Star program (www.energystar.gov/datacenters)
EPEAT	Electronic Product Environmental Assessment Tool (www.epeat.net)
ISO 14001	Environmental management standards
JEDEC	Joint Electronic Device Engineering Council (www.jedec.org)
JEITA	Japan Electronics Information Technology Industry Association (www.jeita.or.jp)
J-MOSS	Japanese program for removal of hazardous substances
LEED	Leadership in Energy Efficiency Design
MSDS	Material Safety Data Sheet for products
NRDC	Natural Resources Defense Council (www.nrdc.org)
REACH	Registration, Evaluation, Authorization and Restriction of Chemicals
RoHS	Restriction of Hazardous Substances (www.rohsguide.com)
SB20/50	California Electronics Waste Recycling Act of 2003
USGBC	U.S. Green Building Council (www.usgbc.org)
WEEE	Waste from Electrical and Electronic Equipment
WGBC	World Green Building Council (www.worldgbc.org)

provide a primary source for electric power, leveraging the lowest-cost, most effectively available power currently available in the power grid. As a standby power source, backup generators fueled by diesel, propane, or LP gas can be used.

From an economic standpoint, working with local and regional utilities to improve electrical efficiency and obtain rebates and other incentives should all be considered. For example, during nonpeak hours, electrical power from the local power grid can be used; during peak demand periods, backup standby generators can be used in exchange for reduced energy fees and avoiding peak demand surcharges. Another economic consideration, however, is the cost to run on standby generator power sources. These costs, including fuel and generator wear and tear, should be analyzed with respect to peak-demand utility surcharges and any incentives for saving energy. For organizations that have surplus self-generated power, whether from solar, wind, or generators, some utilities or other organizations will buy excess power for distribution to others, providing cash, rebates, or discounts on regular energy consumption. Learn more about electrical energy fuel sources, usage, and related statics for the United States and other countries at the energy information administration website www.eia.doe.gov/emeu/aer/elect.html.

2.6 PCFE and Environmental Health and Safety Standards

The green supply chain consists of product design, manufacture, distribution, and retirement, along with energy production, deliver, and consumption. Table 2.2 provides a sampling of initiatives relating to PCFE and environmental health and safety.

Later chapters in this book will present additional technologies and techniques to boost efficiency and productivity to address PCFE issues while balancing performance, availability, capacity, and energy efficiency to meet various application service requirements.

2.7 Summary

Action items suggested in this chapter include:

- Gain insight into how electrical power is used to determine an energy efficiency baseline.
- Investigate rebates and incentives available from utilities and other sources.
- Explore incentives for conducting data center energy efficiency assessments.

- Understand where PCFE issues and bottlenecks exist and how to address them.
- Investigate alternative energy options, balancing economic and environmental concerns.
- Review your home and business electric utility bills to learn about power usage and costs.
- Learn more about the various regulations related to environmental health and safety.

Other takeaways from this chapter include:

- Energy avoidance may involve powering down equipment
- Energy efficiency equals more useful work and storing more data per unit of energy.
- Virtualization today is for consolidation
- Virtualization will be used tomorrow to enhance productivity.