The Greening of IT

How Companies Can Make a Difference for the Environment

John Lamb

The Greening of IT – by John Lamb – IBM Press – 0137150830/9780137150830 – May 2009

1

The Importance of Green IT

"As more and more people understand what's at stake, they become a part of the solution, and share both in the challenges and opportunities presented by the climate crises."

—Al Gore on global warming

Information Technology (IT) is at the heart of every successful modern business. Without it, success is impossible. Yet, the pervasive deployment of IT has had significant, unintended side effects, namely as a significant contributor to the economically unsustainable worldwide dependence on fossil fuels. The awareness of these side effects, though somewhat late in coming, has led some successful companies to turn to a sustainable practice known as "IT greening." IT greening is about using IT more efficiently to achieve reductions in energy consumption, and therefore, considering the acquisition of energy-efficient IT solutions. Within this book, you can find details on the environmental impact of IT, including data centers' consumption of fossil fuel-based electric energy. In addition, we examine many case studies, extracting lessons learned and best practices for implementing green IT.

IT is so pervasive that energy efficiency through the implementation of green IT has moved to center stage for many companies in their pursuit of environmentally helpful practices. This book provides details on the importance of implementing green IT; the significant and growing role of IT and data centers in the world's consumption of electric energy and carbon footprint; and especially the case studies for "lessons learned" and the best-practice approaches for implementing green IT.

As I mentioned in the Preface, green IT is an ideal way for most companies to make a significant step in reducing their carbon footprint for several reasons. First, for competitive reasons, most companies already refresh their computer hardware—laptops, desktops, servers, and storage devices—every three to four years. That refresh cycle provides a recurring opportunity to buy increasingly energy-efficient technology, such as virtual servers, virtual networks, and virtual data storage. Such virtualization can easily reduce IT power consumption for the replaced equipment by up to 50 percent. (For examples, refer to the Environmental Protection Agency's {EPA} "Report to Congress on Server and Data Center Energy Efficiency" or the reports by Jonathan Koomey listed in the Bibliography.) A second compelling reason to move to green IT is that virtualization technology enables you to reduce equipment and system management costs for your data center. Data center green technology is based on a solid business case—even before we consider the savings due to reduced energy costs. A third reason for moving to green IT is that all large companies are moving to such implementation improvements (in IT virtualization, cloud computing, and so on). In addition to information on IT virtualization, this book also includes information on new energy-efficient cooling technologies that support IT, and the impact of electric utility-rate case incentives and government incentives and regulations on promoting IT energy efficiency.

Green IT has many different aspects. In this book, we use the terms **green IT**, **green computing**, and **green data centers**. Green IT—as used here—is the most comprehensive because it includes all computing, inside and outside the data center. The emphasis of our discussion is on the business aspects of green IT, so the focus is on what to do, rather than the details of how to do it. However, several chapters, especially the case studies, do give details on how to implement green IT, using best practices based on recent experience and lessons learned through dealing with many companies and organizations throughout the world.

In the following chapters, we look at the benefits and roadblocks in moving to green IT, including the following:

1. Organizational issues in addressing the problem (for example, CIO doesn't pay the electricity bill).
2. The future of regulations as external factors for change.
3. Overall motivation for executives to move to green IT.
4. Evaluation of product end of life and asset disposal, procurement policies, and supply-chain issues (solutions to avoid climate impact, and such).

Executives have one significant aspect of motivation to move to green IT that is not covered in any depth in this book, and that is the area of corporate social responsibility. A growing body of evidence shows that companies can do well by doing good. In fact, books have been written about this corporate motivation for going green. (See the reference to *Green to Gold* in the Bibliography.) The Internet's ubiquitous connectivity has created new relationships among businesses, customers, employees, and partners. People now have access to massive amounts of information and opinions about products and company practices. This information is available in every part of the globe, every minute of every day. Collaboration over the Internet is taking place during a time of increased visibility of corporate actions, a time when customers' perceptions of companies—and their consequent purchasing behaviors—are fundamentally changing. Thus, having your company become part of the **green wave** (to use a term used in the *Green to Gold* book) should be an additional motivation (besides the standard business case) for companies to pursue green IT.

Although this book emphasizes the business aspects, rather than the technical aspects, of green IT, several chapters give technical details, including the case study chapters (Chapters 9, 10, and 11). I'm an engineer by training, and I'm fascinated by the technical aspects of green IT. The energy used for green IT and green data centers is electricity, so I would like to give a brief review of the familiar electricity concepts and relationship of volts, amps, and watts, which are fundamental in our quest to reduce energy used for green IT. The relationship between volts, amps, and watts is this: watts = volts × amps. A watt is the measure of electrical power. Energy is power over a unit of time. We pay for electricity in terms of energy used with a measure of kilowatt hours or KWH. One kilo watt hour (KWH) of electrical energy is the energy used by having 10 one-hundred-watt light bulbs on for one hour. In the New York City area, one KWH costs about 20 cents, whereas in West Virginia, a KWH costs only about 5 cents. The big difference in cost is due to the big difference in generation costs. With the high cost of electricity, the region of the country is becoming a significant factor in deciding where to locate a new data center.

Besides the cost per KWH, another aspect of basic electricity to consider for your data center is the voltage level. In the United States, we typically have two voltages to use in our homes, offices, and data centers: 110 volts or 220 volts. The actual volts can fluctuate somewhat during the day (as you can discover using a simple voltmeter), and electrical engineers often give the two voltages available as 120V or 208V. If you have an electric range, an electric clothes dryer, or a large electric air conditioner in your home, they'll be connected to the higher 208V service. The reason is that they need more power; using a higher voltage gives more power and also saves energy. Higher voltage saves energy because the formula for electric losses due to transmission over a wire is $I^2 R$, where I represents amps, and R is the fixed resistance of the wire. Because watts = volts × amps, we can double the power (watts) by doubling the volts or doubling the amps. However, doubling the amps would increase the losses by four times. Thus, to transmit electricity over long distances, the practice is to increase the voltage as much as possible. In high-tension towers, the voltage is often as high as 120,000 volts, and even in the power lines outside our houses, the voltage is often 4,000 volts. That voltage drops to 110V for use in our houses to reduce danger of electrocution. As noted, 220V power is used only for electric ranges, clothes dryers, and so on, although in Europe, the base power is 220V. One easy way to reduce energy transmission losses at data centers is to use 220V (208V) service. Higher input voltage results in more-efficient operation. Most servers (just like our laptops or hair dryers) are capable of either 110V or 220V service. Older data centers often have 110V power sources for servers, but switching to 220V (208V) would provide significant savings. This change in voltage at data centers to reduce energy use is mentioned in several places throughout the book.

The Growing Significance of Green IT and Green Data Centers

Much of the emphasis in the following chapters is on data centers because they almost always represent the starting point for green IT initiatives for companies. Data centers—the facilities that primarily contain electronic equipment used for data processing, data storage, and communications networking—have become common and essential to the functioning of business, communications, academic, and governmental systems. Data centers have been growing and expanding quickly as our economy continues to shift from paper-based to digital information management. The U.S. EPA's 2007 "Report to Congress on Server and Data Center Energy Efficiency" estimated that the energy use of the nation's servers and data centers doubled from 2000 to 2006 to approximately 61 billion KWH. Under current efficiency trends, national energy consumption by servers and data centers could nearly double again by 2011 to more than 100 billion KWH, representing a $7.4 billion annual electricity cost.

Data centers are found in nearly every sector of the economy, including financial services, media, high-tech, universities, and government institutions. Dramatic server growth at data centers is indicated by well-known web services such as Google, Amazon, and eBay. Estimates indicate that Google maintains more than 450,000 servers, arranged in racks located in clusters in cities around the world. Google has major data centers in California, Virginia, Georgia, and Ireland, and new facilities in Oregon and Belgium. In 2009, Google is planning to open one of its first sites in the upper Midwest in Council Bluffs, Iowa, close to abundant wind power resources for fulfilling green energy objectives and proximate to fiber optic communications links. For additional information on the positioning of new data centers close to abundant renewable electric power sources, see Appendix C, "Comparison of Different Power-Generation Methods." Amazon.com and eBay also have thousands of servers. It is estimated that the Second Life Internet-based virtual world launched in 2003 has more than 9,000 servers. Even with these large numbers of current servers, IBM consultants estimates that in the next decade, server shipments will grow by six times and data storage by an amazing 69-fold.

Green energy-efficient data centers can help us reduce greenhouse gases—which, in turn, can help reduce global warming. The recent UN and White House sessions on climate change emphasize the environmental importance of green projects. Although the extent of the global warming danger might continue to be open to debate, implementing green data centers presents a significant opportunity for all of us to help reduce greenhouse gasses.

In many instances—such as building alternative energy sources by implementing solar panels, or wind turbines—going green has not been economical and can be justified only by government or energy utility rebates. Yet implementing green data centers can be quite financially rewarding—especially when you go first after the low-hanging fruit. As I've indicated throughout this book, going to green IT is a win/win for all parties involved. Energy expenditures for IT keep increasing. Figures mentioned previously bear repeating. According to the research firm IDC: By 2010, for every $1 spent on hardware, 70 cents will be spent on power and cooling, and by 2012, for every $1 spent on hardware, $1 will be spent on power and cooling. Green IT has generated significant customer interest throughout the world. Much of the interest comes from the financial return on green data center investment.

Here is a general definition of a green data center: A repository for the storage, management, and dissemination of data in which the mechanical, lighting, electrical, and computer systems are designed for maximum energy efficiency and minimum environmental impact. The construction and operation of a green data center involve use of advanced technologies and strategies. The strategies and goals include the following:

- Minimizing the footprints of the buildings
- Using low-emission building materials, carpets, and paints
- Creating sustainable landscaping
- Initiating waste recycling
- Installing catalytic converters on backup generators
- Using alternative energy technologies such as **photovoltaics** (PVs) and fuel cells
- Increasing the efficiency of heat pumps, variable speed fans, and free-cooling technology

However, in the following chapters, we concentrate on the ways data centers can become more energy efficient by first exploring the low-hanging fruit. The basic technologies that we should first examine for existing data centers range from the use of efficient cooling towers and variable speed blowers to the use of energy-efficient IT systems, such as virtual servers, blade centers, and virtual data storage. Server consolidation—although initially undertaken to save server hardware capital—is also an excellent way to reduce server energy use. A step way beyond server consolidation is data center consolidation—also done to reduce facility and personnel resource cost; however, a significant side effect is reduced data center energy use. Most data centers have already started to employ newer IT technology such as virtual servers or server consolidation, so this book first explores the technologies that have already started to be employed at your data center for capital cost-saving—and discuss the ways this same technology can significantly reduce energy use. The experiences described in the case studies presented in later chapters are a way to leverage those lessons learned for your data center.

Many consultant reports indicate that data centers are at a "tipping point." Some well-publicized issues supplying adequate electrical power to data centers include Canary Wharf in London and the area south of 14^{th} Street in New York City. In 2006, the financial institutions at Canary Wharf were told that the power infrastructure could not supply power for additional servers at their data centers. In recent years, financial organizations have been adding significant server power, often with racks of blade servers. The racks of blade servers can greatly increase the power required per square foot in the data center. Each blade server requires about the same energy as larger, older servers, and the data center needs similar levels of electricity to cope with the heat generated. Canary Wharf didn't have the power infrastructure to support the increased demands. A similar limit of the power structure occurred during 2008 for data centers south of 14^{th} Street in Manhattan. Power restrictions to data centers based on inadequate power infrastructure is only a part of the problem. Data center floor space has also become a significant concern for data centers, especially in large cities. Often, a company runs out of data center floor space with no easy capability to expand.

The green IT techniques described in later chapters (such as server and data storage virtualization, and server consolidation), in addition to cutting power requirements by 50 percent, also reduce data center floor space requirements. Using virtual server techniques to replace ten stand-alone physical servers with one large physical box that includes ten virtual servers can easily reduce the data center floor space required by 80 percent. Practicing green IT promotes a "win-win" situation for all aspects of your data center: electric-power reduction, server cost, data center floor space, and management of the physical boxes.

Although building and certifying a green data center or other facility can often be expensive upfront, substantial long-term cost savings can be realized on operations and maintenance. The green data center technologies described later can all be based on the typical business case, where a significant **return on investment** (ROI) would be required before proceeding with a project. Of course, there are also significant nonfinancial returns to consider—because green facilities (including green data centers) offer employees a healthy, comfortable work environment. In addition, we cannot ignore the fact that green facilities enhance relations with local communities.

We are all aware of the growing pressure from environmentalists and, increasingly, the general public for governments to offer green incentives: monetary support for the creation and maintenance of ecologically responsible technologies. Server refresh offers data centers a convenient opportunity to go green, which always makes economic (as well as environmental) sense. IBM estimates that a typical 25,000 SF data center with electrical costs at 12 cents per KWH will cost a company $2.5 million a year in electrical energy costs for IT power and cooling. IBM also estimates that the typical data center can reduce its annual electricity cost by up to 50 percent by going green. Of course, as energy costs continue to climb, so will the savings due to the installation of energy-efficient IT equipment and optimization of data center cooling techniques.

Recent EPA reports stress that the U.S. data center industry is in the midst of a major growth period stimulated by increasing demand for data processing and storage. This demand is driven by several factors, including the following:

- Increase in electronic transactions in financial services, such as online banking and electronic trading
- Growth of Internet communication and entertainment use
- Increase in online shopping and related transactions
- Shift to electronic medical records for healthcare
- Growth in global commerce and services
- Adoption of satellite navigation and electronic shipment tracking in transportation

Other important trends contributing to data center growth in the government sector include the following:

- Use of the Internet to publish government information
- Government regulations requiring digital records retention
- Enhanced disaster recovery requirements
- Emergency, health, and safety services
- Information security and national security
- Digital provision of government services (for example, e-filing of taxes and U.S .Postal Service online tracking)
- High-performance scientific computing

During the past five years, increasing demand for computer resources has led to significant growth in the number of data center servers, along with an estimated doubling in the energy used by these servers and the power and cooling infrastructure that supports them. This increase in energy use has a number of important implications:

- Increased energy costs for business and government
- Increased emissions, including greenhouse gases, from electricity generation

- Increased strain on the existing power grid to meet the increased electricity demand
- Increased capital costs for expansion of data center capacity and construction of new data centers

For these reasons, there has been mounting interest in opportunities for energy efficiency in this sector. To its credit, the Information Technology (IT) industry is actively investigating and developing solutions, such as power-managed servers and adaptive cooling.

The direct energy use of IT and infrastructure equipment is not, however, the only way that data centers affect energy use. The data-processing and communication services provided by data centers can also lead to indirect reductions in energy use in the broader economy, which can exceed the incremental data center energy expenditures in some cases. For instance, e-commerce and telecommuting reduce both freight and passenger transportation energy use. When we use an electronic bookstore such as Amazon.com, that use of e-commerce can save us from driving to the local bookstore, and, hence, save energy. We can attend a "virtual" conference using a web-conferencing service such as Webex or Microsoft® Live Meeting and thus save the energy expenditure of an airline flight to the conference, use of a rental car, and all the other energy use that travel entails.

The pursuit of energy efficiency opportunities in data centers is especially important because of the estimated continued rapid growth of direct energy use in data centers and the resulting impact on both the power grid and U.S. industries.

To repeat the theme: We're all aware of rising energy costs in today's data centers and the growing concerns over global warming and other environmental issues. These problems have made green IT one of the hottest topics in the IT area. But what exactly is green IT and green computing and how does it affect IT infrastructures? This book provides an outline on the concepts, benefits, and business value of green computing, such as the following:

- A definition/analysis of green computing and its benefits
- An overview of green computing solutions
- The business case for going green
- Implementation of an energy management solution
- Why energy efficiency is so important

All Companies Can Take Basic Steps Toward Green IT

According to Gartner research firm, the green wave has only begun to rise. The research company predicts that in 2009, more than one-third of all IT organizations will place environmental concerns among their top six buying criteria. By 2010, Gartner says, three-quarters of companies will use carbon-footprint considerations in calculating their hardware-buying strategy, and by 2011, large enterprises will develop policies requiring their suppliers to prove their green credentials through an auditing process.

Most companies are talking a good game but not actually going green where it counts. According to a survey of 124 IT operations by Forrester Research in May 2007, some 85 percent of respondents said environmental factors are important in planning IT operations. But only one-fourth of survey respondents have actually written green criteria into their company's purchasing processes. Enterprises that have started the green journey, however, have found that reducing total energy requirements can be accomplished through some fairly straightforward improvements that don't take years to implement or to bring return. The following six tasks are applicable to all green IT projects. Chapter 2, "The Basics of Green IT," gives details on the five steps used by IBM to implement green data centers. Those five green data center steps include the virtualize, cooling, and measure tasks in the following list. Also, Chapter 9, "Green IT Case Studies for Energy Utilities," and Chapter 10, "Green IT Case Studies for Universities and a Large Company," give details on the five steps used for case studies.

1. Communicate Green IT Plans and Appoint an Energy Czar

Measuring the current state of affairs, energy wise, is one of the first steps to take. A baseline on which to start measuring the impact of an organization's energy-saving initiatives in the green IT area is needed. Of course, you must also communicate your proposed energy-efficiency initiatives right away. Inform all employees about the plans and goals to save energy via green IT. Besides communicating with your employees, set up an organization to drive the effort. You may start by making one person responsible; give that person a title (like "Energy Czar"). Details on the importance of communication and collaboration for green IT is the subject of Chapter 3, "Collaboration Is Key for Green IT."

2. Consolidate and Virtualize

Consolidating IT operations, and using virtualization to reduce server footprint and energy use, are the most well-recognized and most-often-implemented efficiency strategies of the past few years. Some of the largest technology organizations in the world—including Advanced Micro Devices®, Hewlett-Packard®, Intel®, IBM, and Sun Microsystems®—have recently (2008) completed major data center consolidation projects. The projects also included server consolidation and virtualization. Details on the significance of virtualization for your IT systems in going to green data centers is the subject of Chapter 6, "A Most-Significant Step—'Virtualizing' Your IT Systems."

3. Install Energy-Efficient Cooling Units

In most cases, traditional data center design called for bulky **computer room air conditioners (CRAC)** units that are placed on the perimeter of the floor to move large amounts of air around the data center. However, in-row or supplemental cooling units have been shown to save energy. The in-row units typically enclose a row or two of servers, and the backs of all the servers are pointed into a single "hot" aisle. Heat in the aisle is contained by a roof and end-row doors, allowing cooling to be applied directly to the heat source, rather than trying to cool after the heat is dispersed into the general data center floor. Details on data center cooling strategies for green data centers are given in Chapter 8, "What About Chillers, Cooling Tower Fans, and All That Cooling Equipment Usually Ignored by IT?"

4. Measure and Optimize

In 2009, several groups (including the The Green Grid) are expected to release important deliverables in the form of metrics that businesses can use to measure the power-usage effectiveness of facilities infrastructure equipment. Most businesses can already readily identify areas where infrastructure optimization can achieve increased efficiency by simply monitoring and measuring their existing infrastructure equipment. The EPA is also working to create metrics. About 100 companies have indicated that they will provide raw power data and other information to the EPA for use in developing its new benchmark. The EPA indicated that the results of the benchmark should be available by 2010.

Until widely accepted metrics become available, businesses should make sure that the utility costs associated with their data center operations are broken out separately from those for other corporate facilities. In addition, metering specific equipment racks or types of equipment such as servers can provide valuable insight into which specific consolidation, virtualization, and optimization projects will yield the best ROI going forward. The status of energy-use metrics is the subject of Chapter 7, "The Need for Standard IT Energy-Use Metrics."

5. Implement Efficient Applications and Deduplicate Data

Software and application efficiency can be significant for green IT. The author has had a recent experience where the procedure for creating a data warehouse report was reduced from eight hours to eight minutes merely by changing the Oracle data warehouse search procedure. (For example, don't search the entire database each time when only a much smaller search is required.) During the eight hours required to create the report, the large server was running at near peak capacity. Sure, that type of significant application inefficiency has been created and fixed many times over the history of programming. But what about the cases where a few application efficiencies can make an application run 20 percent faster? That 20 percent more-efficient application can also result in 20 percent lower energy use. The steps required to improve application efficiency by a few percent are often not easy to determine; however, the added incentive of saving energy—while making the application run faster—is a significant plus.

Data-storage efficiency, such as the use of tiered storage, is also significant. **Data deduplication** (often called **intelligent compression** or **single-instance storage**) is a method of reducing storage needs by eliminating redundant data. Only one unique instance of the datum is actually retained on storage media, such as disk or tape. Redundant data are replaced with a pointer to the unique data copy. For example, a typical email system might contain 100 instances of the same one-megabyte (MB) file attachment. If the email platform is backed up or archived, all 100 instances are saved, requiring 100MB storage space. With data deduplication, only one instance of the attachment is actually stored; each subsequent instance is just referenced back to the single saved copy. In this example, a 100MB storage demand can be reduced to only one MB.

Data deduplication offers other benefits. Lower storage space requirements can save money on disk expenditures. The more efficient use of disk space also allows for longer disk-retention periods, which provides better **recovery time objectives** (RTO) for a longer time and reduces the need for tape backups. Data deduplication also reduces the data that must be sent across a WAN for remote backups, replication, and disaster recovery.

Data deduplication uses algorithms to dramatically compress the amount of storage space needed. Many organizations deal with increased scrutiny of electronically stored information because of various regulations; this need to preserve records is driving significant growth in demand for storing large sets of data. Depending on the type of information compressed, deduplication can enable a compression rate of between 3:1 and 10:1, allowing businesses to reduce their need for additional storage equipment and associated tapes and disks. Many businesses are already using the technology. Application efficiency as part of green IT strategy is discussed in Chapter 2.

6. Make Use of Rebates and Incentives

More utility providers offer rebates or other incentives that encourage businesses to update equipment and adopt efficient operational practices that can help reduce peak and total power demands. Companies doing this include Pacific Gas and Electric in San Francisco and Austin Energy in Austin, Texas.

New electric power-generation stations are very expensive, and power companies are more than willing to avoid building new capacity. Thus, the power companies encourage data center efficiency through rebates and other incentives. Also, the organization's facilities team doesn't have to build as much new data center space. The IT organization and engineering groups get new equipment that is smaller, cooler and faster than before—and everyone ends up happy. The roles of government and energy utility rebates and incentives are the subjects of Chapter 4, "The Government's Role—Regulation and EPA Activity," and Chapter 5, "The Magic of 'Incentive'—The Role of Electric Utilities."

What This Book Covers

This book includes the following topics to help you understand green data centers and your potential role in creating and maintaining them:

- The significant role data centers have in the world's consumption of electric energy and carbon footprint.
- How companies are offering services and products to help reduce data center energy use—for example, IBM's Big Green $1 billion annual investment in green data centers.
- How IT employees (for example, corporate CIOs (chief information officers), IT architects, IT specialists, and IT project managers) can help drive the implementation of green data centers.
- Case studies of organizations implementing green data centers.
- Details on the best ways to measure data center energy use and report to your executives. Because "You can't manage what you can't measure," the first step is to start the measurement process and understand the need to continually improve your measurement process. This is necessary to better quantify the savings due to your energy initiatives.
- Study of the different ways to measure server utilization and look at trends. You need to answer the question: How has customer server virtualization increased server CPU utilization?
- Continuing follow-up on the literature on green data centers because technology is progressing at a fast pace. The U.S. EPA is key to following the U.S. government recommendations and incentives for data center energy efficiency.
- Survey of emerging technology for server and storage enhancement to reduce data center energy use. This includes the following:
  - **Information Lifecycle Management** (ILM), overall storage management, tiered storage
  - Server virtualization enhancements such as PowerVM®, VMware enhancements, and such
  - Active energy management
  - Enhanced cooling technology
- Analysis of emerging technology for server and storage enhancement to reduce data center energy use.