


The Green Data Center: Steps for the Journey

Going green: Improving energy
efficiency and costs

Green IT products

Green tips for site and
facilities



Mike Ebbers
Alvin Galea
Michael Schaefer
Marc Tu Duy Khiem



International Technical Support Organization

The Green Data Center: Steps for the Journey

August 2008

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (August 2008)

Contents

Notices	vii
Trademarks	viii
Preface	ix
The team that wrote this paper	ix
Become a published author	x
Comments welcome	x
Chapter 1. The benefits of a green data center	1
1.1 Overview	2
1.1.1 Managing the increasing cost of the energy	2
1.1.2 Running out of power capacity	3
1.1.3 Running out of cooling capacity	3
1.1.4 Running out of space	3
1.2 How energy is used in a data center	3
1.3 Environmental laws and the company image	4
1.4 What is happening around the world	5
1.5 The benefits	5
1.6 Next steps	6
Chapter 2. Developing a strategy	7
2.1 Assess the greenness of your data center	8
2.1.1 Calculate data center infrastructure efficiency	8
2.1.2 Important questions to consider	9
2.2 Strategy recommendations	12
2.3 A summary of best green practices	12
Chapter 3. Energy optimization with IT equipment	17
3.1 Energy flow in a computer system	18
3.1.1 How the electric power is used	18
3.1.2 What to do with the heat	19
3.2 Working with energy efficiency	21
3.3 System instrumentation for power management	22
3.3.1 IBM systems with embedded instrumentation	22
3.3.2 IBM intelligent power distribution units	23
3.4 Power management: The hardware side	23
3.4.1 IBM POWER6 EnergyScale	23
3.4.2 IBM BladeCenter features	24
3.4.3 System x features	25
3.4.4 ACPI and friendly power saving features	25
3.4.5 High input voltages	25
3.5 Power management: The software side	26
3.6 Consolidation and virtualization overview	27
3.6.1 Consolidation: A key in energy efficiency	27
3.6.2 Virtualization: The greenest of technologies	28
3.7 Server virtualization	29
3.7.1 Partitioning	29
3.7.2 Special virtualization features of IBM systems	30
3.7.3 Other virtualization techniques	30

3.8	Storage virtualization	31
3.8.1	IBM SAN Volume Controller	32
3.8.2	Virtual tapes	33
3.9	Client virtualization	33
3.10	Integration of energy and systems management	35
3.11	Where to start	36
Chapter 4. Site and facilities		37
4.1	Tips to start moving your data center towards green	38
4.1.1	Reducing power consumption with innovative technologies	38
4.1.2	Reducing cooling requirements	38
4.1.3	Improving physical infrastructure	39
4.2	Optimize your data center cooling	39
4.2.1	Manage airflow	40
4.2.2	Structured cable management	42
4.2.3	Recommendations for raised floor height	43
4.2.4	Data center insulation	43
4.2.5	Hot aisle and cold aisle configuration	43
4.3	Localized cooling equipment options	46
4.3.1	IBM Cool Blue rear door heat exchanger	47
4.3.2	Modular water unit	48
4.3.3	Enclosed rack cooling	49
4.3.4	Sidecar heat exchanger	50
4.4	Heating ventilation and air conditioning (HVAC)	51
4.4.1	Types of chillers	51
4.4.2	Variable-speed drive pumps	52
4.4.3	Air handling unit	52
4.4.4	Economizers to enable free cooling	52
4.5	Cool less, save more	52
4.6	Uninterruptible power supply	53
4.6.1	Components and types	53
4.6.2	Flywheel technology	54
4.7	Power	54
4.7.1	Utility power supply	54
4.7.2	Power factor correction	55
4.7.3	Power distribution and resistance	55
4.7.4	Intelligent power distribution unit	55
4.7.5	DC versus AC	55
4.7.6	On-site generation	55
4.8	Generators	55
4.8.1	Standby generator	55
4.8.2	On-site electrical generation	56
4.9	Recommendations for existing data centers	57
4.9.1	Investigate, consolidate, and replace	57
4.9.2	Investigate industry initiatives	57
4.10	Summary	58
Chapter 5. Who can help: IBM services and Business Partners		59
5.1	Services IBM can deliver	60
5.1.1	IBM Global Technology Services	60
5.1.2	IBM Asset Recovery Solutions	63
5.1.3	IBM financing solutions	65
5.1.4	IBM green data centers	65

5.2 IBM Business Partners	65
Chapter 6. Conclusion: Green is a journey, act now	67
Appendix A. Commitment to green from IBM: The past, present, and future	69
A.1 A history of leadership in helping the environment.	70
A.2 Project Big Green	70
A.3 IBM internal efficiency results	71
A.4 Future directions	71
Related publications	73
IBM Redbooks publications	73
Online resources	73
Other publications	75
ASHRAE publications	76
How to get IBM Redbooks publications	76
Help from IBM	76

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®	IBM®	System p5™
BladeCenter®	iDataPlex™	System x™
Cool Blue™	POWER™	System z™
DB2®	POWER6™	System z10™
DFSMS™	PR/SM™	System/360™
DPI®	Processor Resource/Systems Manager™	Tivoli®
EnergyScale™	Redbooks®	Virtualization Engine™
FICON®	Redbooks (logo)  ®	X-Architecture®
HiperSockets™	System p™	z/OS®
IBM Project Financing™		z10™

The following terms are trademarks of other companies:

AMD, AMD PowerNow!, the AMD Arrow logo, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

InfiniBand, and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

Power Management, Solaris, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel SpeedStep, Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

The information technology (IT) industry is a big user of resources as well as an enabler of efficiencies that reduce CO₂ emissions. However, as companies continue to grow to meet the demands of their customers and as environmental concerns continue to be an issue, organizations are looking for ways to reduce corporate energy consumption and to become more environmentally responsible—to become *green*.

This IBM® Redpaper publication can help your IT organization as it begins the journey to becoming a *green data center*. IBM wants to help others, particularly our clients, to chart a course to reap the benefits of lower costs and improved sustainability that running a green data center can provide. Understanding what is possible can speed your journey to an optimized green data center with sustainability designed into both the IT and facilities infrastructures. Although this paper is not all inclusive, it provides a quick start for going green in data centers. It also provides additional pointers and information. You can use this paper as a guide to becoming more energy efficient.

The team that wrote this paper

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Raleigh and Poughkeepsie Centers.

Mike Ebbers is a Consulting IT Specialist and ITSO Project Leader, working for IBM for 34 years. He has published a number of books on various IT solutions. Before joining the ITSO 13 years ago, Mike worked in as an Instructor and Develop for IBM U. S. Technical Education.

Alvin Galea is a data center and facility infrastructure specialist for the Australia Post Corporation in Australia. He has over 20 years experience in IT and has worked the last eight years directly with data center and facility infrastructure. He has been involved in a number of large data center and facility upgrades, including utility substation, uninterruptible power supply, main switch board, data center layout reconfigurations, standby generator, chilled water loops, and the replacement of three chillers.

Michael Schaefer is an IT specialist for systems management in Germany. He has over 20 years of experience in the systems management field. He holds a masters degree (Diplom-Informatiker) in computer science from the University of Karlsruhe, Germany. His areas of expertise include systems, database, network, and intranet management on many platforms, including AIX® and Solaris™, Oracle®, and DB2®. He has been involved in numerous projects with customers all over Germany, where his focus was on integration of various systems management tools into overall seamless management applications. Currently, he works for the Chip Design and Simulation Center in IBM Germany Research and Development in Boeblingen, Germany. He is actively involved in the Boeblingen Energy Efficiency Initiative.

Marc Tu Duy Khiem is an IT Specialist in France. He is currently leading the transformation of Montpellier Products and Solutions Support Center (PSSC) Data Centers towards energy efficient data centers. He is also involved in the project of building a *Green Data Center of the Future* showcase in Montpellier to demonstrate IBM and Partners solutions that answer current data center customer pain points. He has worked at IBM for five years. His areas of expertise also include Oracle products, N series, System x™, BladeCenter®, and Linux®.

Thanks to the following people for their contributions to this project:

David F Anderson PE is a Green Architect from the IBM Briefing Center in Poughkeepsie supporting IBM worldwide. He is a licensed Professional Engineer with a B.S. from the United States Military Academy (West Point), an MBA from the University of Puget Sound, and a Masters in Engineering Science from Rensselaer Polytechnic Institute.

Don Roy is a Senior IT Consultant and Server Presenter for the IBM RTP Executive Briefing Center. Prior to joining the RTP Executive Briefing Center, Don was the World Wide Product Marketing Manager and Team Lead for System x high volume servers. In the Executive Briefing Center, Don provides specialized coverage for System x servers and enterprise solutions.

Dr. Roger Schmidt is a Distinguished Engineer, National Academy of Engineering Member, IBM Academy of Technology Member, and ASME Fellow. He has over 25 years experience in engineering and engineering management in the thermal design of IBM's large scale computers. He has led development teams in cooling mainframes, client/servers, parallel processors, and test equipment utilizing such cooling mediums as air, water, and refrigerants. He currently leads IBM's lab services team, providing customer support for power and cooling issues in data centers. Dr. Schmidt is a past Chair of the ASHRAE Technical Committee 9.9 (TC 9.9) on Mission Critical Facilities, Technology Spaces, and Electronic Equipment.

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks® in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- ▶ Send your comments in an e-mail to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099, 2455 South Road
Poughkeepsie, NY 12601-5400



The benefits of a green data center

Many data centers are running out of power and space. They cannot add more servers because they have reached either their power or space limits, or perhaps they have reached the limit of their cooling capacity.

In addition, environmental concerns are becoming priorities because they too can impede a company's ability to grow.

Sustainability is a business imperative. Thus, many data center clients are looking at ways to save energy and cut costs so that the company can continue to grow.

In this chapter, we discuss how running a *green* data center can help clients move toward a more sustainable business model.

1.1 Overview

The journey to a *green* data center has cost and sustainability rewards. Energy consumption and environmental concerns are becoming priorities because they can impede the ability to grow and respond to organizational IT needs.

In September 2007, IBM and the Economist Intelligence Unit issued a report entitled *IT and the environment: A new item on the CIO agenda?*¹ This report found that although most organizations say they are green organizations, many of them are not actually doing as much as they could. Two-thirds of the two hundred or more executives polled said that their organizations have a board-level executive responsible for energy and the environment; however, only 45% of firms had a program in place to reduce their carbon footprint. In addition, of those that did have a carbon reduction strategy, the majority (52%) had no specific targets for it, although a small hard core (9%) aimed to be carbon-neutral by 2012.

Green awareness can become *green action*. Using green technologies to extend capabilities while reducing costs and risks is more than smart business. Planning for energy efficiency and corporate responsibility with positive results go together. The global economy has made the world more competitive and cooperative. As organizations strive to thrive in a constantly changing world, adapting for sustainability is a business imperative. Figure 1-1 lists multiple concerns that CIOs have as they look toward the future.

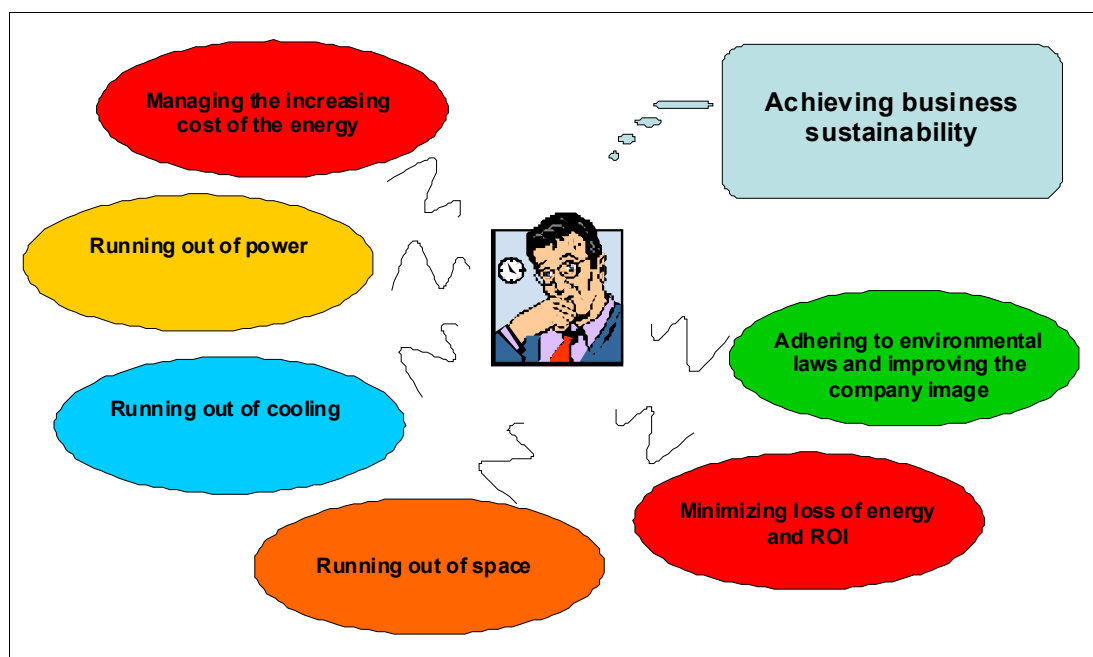


Figure 1-1 Concerns that CIOs have about IT needs

1.1.1 Managing the increasing cost of the energy

The cost of a kilowatt of electricity is constantly rising. According to the Uptime Institute, the three-year cost of powering and cooling typical servers is currently one-and-a-half times the cost of purchasing server hardware.² With the growing demand for less expensive and more powerful high-performance computer clusters, the problem is not just paying for the computers but also determining whether the budget is available to pay for power and cooling.

¹ http://www-05.ibm.com/no/ibm/environment/pdf/grennit_oktober2007.pdf

² Kenneth G. Brill, "Data Center Energy Efficiency and Productivity", The Uptime Institute, 2007

1.1.2 Running out of power capacity

Some companies cannot deploy more servers because additional electrical power is not available. Many utilities, especially those in crowded urban areas, are telling customers that power feeds are at capacity and that they simply have no more power to sell. New server, storage, and networking products give better performance at lower prices but can also be power hungry. The effort to overcome a power supply threshold is a huge investment.

1.1.3 Running out of cooling capacity

Many customer data centers are now 10 to 15 years old, and the cooling facilities are not adapted to the present needs. Traditional cooling methods allowed for 2-3 kW of cooling per rack. Today's requirements are 20-30 kW per rack. Heat density is many times past the design point of the data center.

1.1.4 Running out of space

Each time a new project or application comes online, new servers or storage subsystems are added. Thus, space utilization is growing exponentially because of business requirements. When servers and storage cannot be added, except by building another data center, this growth becomes expensive.

1.2 How energy is used in a data center

To understand how to reduce energy consumption, we need to understand where and how energy is used. We can study energy use in a data center by taking three different views:

- ▶ How energy is distributed between IT equipment (servers, storage, network equipment) and supporting facilities (power, cooling and lighting).
- ▶ How energy is distributed between the different components of the IT equipment (processor, memory, disk, and so forth).
- ▶ How the energy allocated to IT resources is really used to produce business results (Are idle resources powered on, using energy without productive results?)

Figure 1-2 on page 4 shows how energy is used in several components of a typical non-optimized data center. Each component is divided into two portions:

- ▶ IT equipment (servers, storage, and network) uses about 45% of the energy.
- ▶ The infrastructure that supports this equipment—such as chillers, humidifiers, computer room air conditioners (CRAC), power distribution units (PDU), uninterruptible power supplies (UPS), lights, and power distribution—uses the other 55% of the energy.

Thus, IT equipment does not use 55% of the energy that is brought into the data center, so this portion of the energy is not producing calculations, data storage, and so forth. Fortunately, reducing this inefficiency is possible. We discuss energy savings and efficiencies for non-IT equipment in Chapter 4, “Site and facilities” on page 37.

Companies must also consider the energy consumption of the components at the IT equipment level. For example, in a typical server, the processor uses only 30% of the energy and the remainder of the system uses 70%. Therefore, efficient hardware design is very important also. We discuss IT hardware energy savings in Chapter 3, “Energy optimization with IT equipment” on page 17.

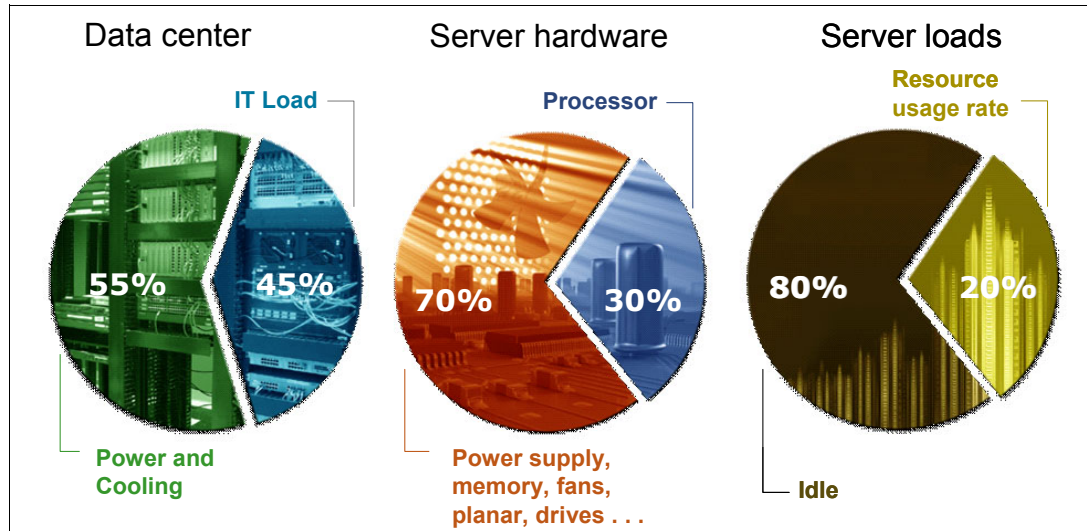


Figure 1-2 How energy is used in a typical data center

Finally, companies should consider the use of IT resources in the data center. Commonly, servers are underutilized, yet they consume the same amount of energy as though they were running at 100%. A typical server utilization rate is 20%. Underutilized systems can be a big issue because a lot of energy is expended on non-business purposes, thus wasting a major investment. Virtualization and consolidation help utilize the entire capacity of your IT equipment.

System p™ and System z™ offer integrated virtualization capabilities that are often more optimized than other servers. By using your IT equipment at its full capacity with consolidation and virtualization, you can achieve energy efficiency and maximize the return on investment. We discuss this in section 3.6, “Consolidation and virtualization overview” on page 27.

1.3 Environmental laws and the company image

Earth’s resources are limited. Oil and coal supplies are finite. Fossil fuel use has been linked to CO₂ increases. The air, water, and mineral resources of the Earth require good stewardship for future generations. Excess CO₂ can impact the climate. Information technology currently accounts for 2% of CO₂ production. At the current rate, energy consumed by servers will grow to 76% of the 2005 levels by 2010.³

Today, climate change is a major environmental and political issue. Environmental laws have affected the IT industry. For example, in the U. S. a recent law authorized the Environmental Protection Agency (EPA) to analyze the growth of energy consumption in data centers. The European Union (EU) has established a directive to drive 20% reduction in energy usage by 2020. The EU, Canada, Japan, and other countries have signed the *Kyoto treaty*. Australia requires all companies using more than 150 000 MWh of electricity per year to prepare an assessment and an action plan.

Beyond respecting environmental laws, being green also enables a company to fulfill its objectives of social responsibility as a *green company*. Being green improves the company’s image for its customers and employees as a good company with which to do business and for which to work.⁴

³ J. Koomey, “Estimating Total Power Consumption by Servers by the US and the World”, 15 Feb 15 2007

1.4 What is happening around the world

A quick survey trip around the world shows what companies are doing to address environmental concerns in the following countries:

- ▶ **North America:** A survey shows that North American companies are beginning to realize the climate challenge, often entrusting responsibility for this issue to their CIO. However, companies are slow to take steps to reduce their overall carbon impact. Today, more than half of North American companies (53%) have someone at the board level who is responsible for energy and environmental issues.⁵ According to recent surveys, climate change, energy, and water are the top three environmental issues.⁶
- ▶ **The U.K.:** While U.K. companies have been quick to assign responsibility at the board level to address the climate challenge, they have been slower to take action, with few U.K. companies actively putting plans in place to reduce their overall carbon impact. Almost two thirds of companies in the U.K. have someone at board level responsible for energy and environmental issues, more than in Western Europe (62%), in North America (53%), or in Asia Pacific (37%). In addition, 16% of U.K. companies have a director with a role that is focused specifically on energy or environmental issues, 7% leave it to the CIO or IT director, and 38% leave it to another director on the board. However, 57% of U.K. companies have no program in place to reduce their overall carbon impact.
- ▶ **Western Europe:** Over 60% of companies in Western Europe (62%) have someone at board level who is responsible for energy and environmental issues, more than in North America (53%) and far more than the in Asia Pacific (37%). In addition, 17% of Western European companies have a director with a role that is focused specifically on energy or environmental issues, 7% leave it to the CIO or IT director, and 38% leave it to another director on the board. However, more than half of companies in Western Europe (57%) have no program in place to reduce their overall carbon impact, compared to 55% in North American and 54% in Asia Pacific.
- ▶ **Asia Pacific:** A survey shows that companies in the Asia Pacific (AP) region have been slower than their counterparts in either Western Europe or North America to address the climate challenge at board level and that over half do not have active plans in place to take steps to reduce their overall carbon impact. Less than 40% of AP companies (37%) have someone at board level who is responsible for energy and environmental issues, less than in North America (53%) and far less than the in Western Europe (62%). In addition, 15% of AP companies have a director with a role that is focused specifically on energy or environmental issues, 7% leave it to the CIO or IT director, and 15% leave it to another director on the board. More than half of AP companies (54%) have no program in place to reduce their overall carbon impact, compared to 55% in North American and 57% in Western Europe.

1.5 The benefits

Figure 1-3 on page 6 outlines the benefits of a green data center.

⁴ For more information, see the information provided by the U.S. EPA:
<http://www.epa.gov/climatechange/index.html>

⁵ *IT and the environment: A new item on the CIO's agenda?*; IBM and The Economist Intelligence Unit Business Study Findings at:
http://www-05.ibm.com/no/ibm/environment/pdf/grennit_oktober2007.pdf

⁶ "Green to Gold," p33. Daniel C. Esty and Andrew S. Winston, Yale University Press, 2006.

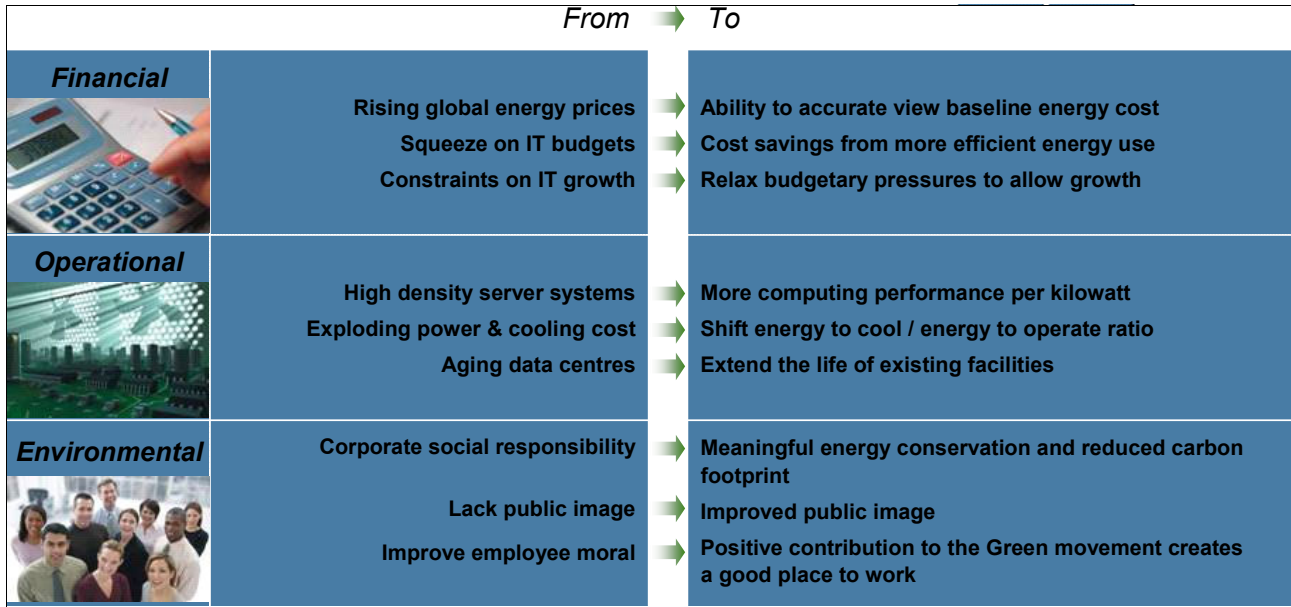


Figure 1-3 Benefits of going green

1.6 Next steps

The remainder of this paper can help you as you move towards a green data center. In the next chapter, we discuss how to assess your data center and provide some strategy recommendations. We also give an overview of best practices. In Chapter 3, “Energy optimization with IT equipment” on page 17 and Chapter 4, “Site and facilities” on page 37, we discuss the infrastructure of the IT systems and data center facilities, and we suggest recommendations. In Chapter 5, “Who can help: IBM services and Business Partners” on page 59, we list ways that IBM and IBM Business Partners can help. Finally, in Appendix A, “Commitment to green from IBM: The past, present, and future” on page 69, we list examples of how IBM is committed to *going green*.



Developing a strategy

A thoughtful green strategy can reduce operational costs and risks. Being green by design can also improve sustainability and corporate image. In this chapter, we discuss how to assess your data center and we provide strategy recommendations. We also provide an overview of best practices at the end of this chapter, and then develop these practices in the chapters that follow.

2.1 Assess the greenness of your data center

Before you can move to a green data center, you must assess the current situation of your data center with as many measurements as are available, so that you can track your progress. You should also research what is available in the following areas:

- ▶ Facilities
- ▶ IT equipment
- ▶ Utilization techniques
- ▶ Use of best practices

2.1.1 Calculate data center infrastructure efficiency

For an indicator of the greenness of a data center, you can use two metrics:

- ▶ The data center infrastructure efficiency (DCiE)
- ▶ The power usage effectiveness (PUE)

PUE is the original metric, while *DCiE* was created to understand more easily where a data center stands in terms of efficiency.

Formula definitions of DCiE and PUE

This section offers the formulas for calculating DCiE and PUE measurements. Both formulas require that you know the total power that your data center uses and the portion of total power that your IT equipment uses.

Note: DCiE is the *reciprocal* of PUE as shown in Figure 2-1 on page 9.

The formulas are:

$$\text{DCiE} = (\text{IT equipment power} / \text{total facility power}) \times 100\%$$

$$\text{PUE} = \text{Total facility power} / \text{IT equipment power}$$

IT equipment power includes the load that is associated with all of your IT equipment (such as servers, storage, and network equipment) and supplemental equipment (such as keyboard, video, and mouse switches; monitors; and workstations or mobile computers) that are used to monitor or otherwise control the data center.

Total facility power includes IT equipment and everything that supports the IT equipment load, such as:

- ▶ Power delivery components such as uninterruptible power supply (UPS), switch gear, generators, power distribution units (PDUs), batteries, and distribution losses external to the IT equipment
- ▶ Cooling system components such as chillers, computer room air conditioning units (CRACs), direct expansion air handler (DX) units, pumps, and cooling towers
- ▶ Computer, network, and storage nodes
- ▶ The decreased efficiency of uninterruptible power supply (UPS) equipment when run at low loads
- ▶ Other miscellaneous component loads such as data center lighting

Figure 2-1 summarizes the relationship of DCiE to PUE.

$$DCiE = \frac{1}{PUE} = \frac{IT\ equipment\ power}{Total\ facility\ power}$$

Figure 2-1 DCiE relationship to PUE

A DCiE value of 33% (equivalent to a PUE of 3.0) suggests that the IT equipment consumes 33% of the power in the data center. Thus, for 100 dollars spent in energy, only 33 dollars are really used by IT equipment. Improvements in energy efficiency result in movement towards 100% DCiE, the ideal number.

IBM is now using the DCiE metric instead of PUE. However, PUE is still in use.

Using DCiE to gauge your situation

A recent study initiated by IBM showed that the average DCiE is 44%. An excellent DCiE is more than 60%, thus many data centers can be optimized. Tips to improve DCiE are discussed in Chapter 4, “Site and facilities” on page 37.

Figure 2-2 shows a scale to illustrate relative measurements of the DCiE and the PUE.

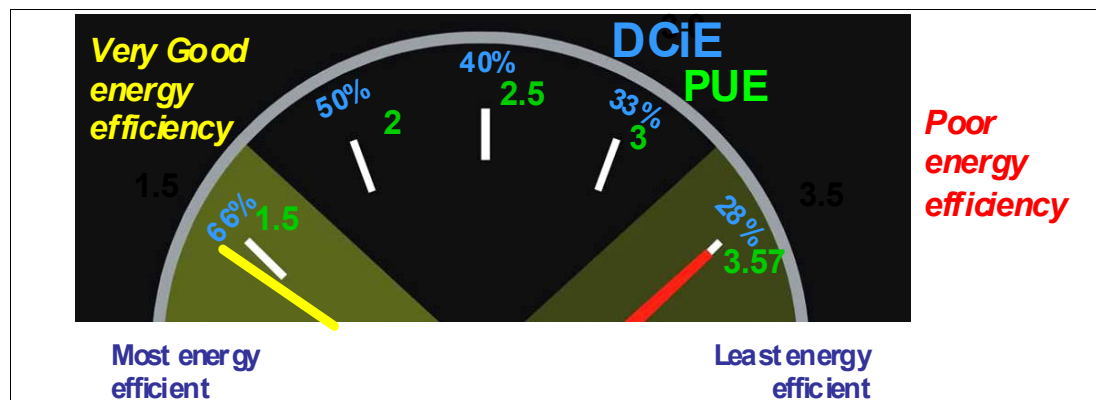


Figure 2-2 Relative measurements of the DCiE and PUE

Important: DCiE indicates the percentage of the energy that is used by IT equipment compared to the total energy drawn. It shows the efficiency of the facility infrastructure that supports your IT equipment. It does *not* provide information about the energy efficiency of the IT equipment itself (servers, storage, and network equipment) or whether the return on investment (ROI) of this equipment is maximized (few idle resources).

2.1.2 Important questions to consider

To assess the energy efficiency of your data center, ask yourself important questions about the following areas:

- ▶ Facilities
- ▶ IT equipment
- ▶ Utilization rate

The facilities

The following set of questions pertains to your site and facilities:

- ▶ How and where is energy being used?
- ▶ What is the facility's current DCiE or PUE?
- ▶ Am I power- and performance-oriented or only performance-oriented?
- ▶ Do I invest in a new data center or can I invest in the evolution of my current data center?
- ▶ Is the physical site of my data center adaptable to changes?
- ▶ Is my desired level of reliability driving the facility's energy consumption? How much idle capacity for redundancy or resilience exist? Is it enough or too much? Could I eliminate any equipment?
- ▶ What support equipment should I choose (uninterruptible power supply, flywheel, generators, power distribution, chillers, CRACs, and so forth)? What are the future trends? Is the facility's infrastructure adaptable to the power and cooling requirements of the next generation hardware? For example, more IT equipment will be water-cooled in the future.
- ▶ Does the facility have problems of overheating? Of humidity?
- ▶ Can I use free cooling? (Refer to section 4.4.4, "Economizers to enable free cooling" on page 52.)
- ▶ Does power, cooling, or space impact operations today? Which will impact business growth in the future? Can I add future compute capacity within my energy envelope?
- ▶ Is my site infrastructure optimized, with regard to the following categories?
 - Airflow and heat removal
 - Power distribution
 - Cooling
 - Lighting
 - Monitoring and management
- ▶ Am I able to exploit water for cooling if needed?
- ▶ Who can help?

The IT equipment

The following set of questions pertains to IT equipment. It includes design of the hardware but also the options for cooling, powering, and monitoring that exist at the rack level:

- ▶ Does the equipment use energy-efficient hardware? Does it use power-reduction features?
- ▶ Should I choose power and cooling options at the site, facilities, or at the rack level?
- ▶ Does the hardware provide options for power, thermal, and usage resource monitoring? Do I monitor and control power consumption?
- ▶ How is billing of the power usage done?
- ▶ Who can help?

The utilization rate

The following set of questions covers the server and storage utilization rate:

- ▶ Is the utilization of my infrastructure optimized?
- ▶ Do unneeded redundancies exist?
- ▶ Can I consolidate and virtualize?
- ▶ How do I go from multiple silos or islands of computing to a *shared model*?

- ▶ Do I monitor utilization rates of my resources? What about real-time and trends?
- ▶ How is billing done for the services that the infrastructure is providing?
- ▶ Who can help?

Summary of questions and additional assessment help

Figure 2-3 shows typical questions that you should ask to evaluate the situation of your data center using the three areas: facilities, IT equipment, and utilization rate.

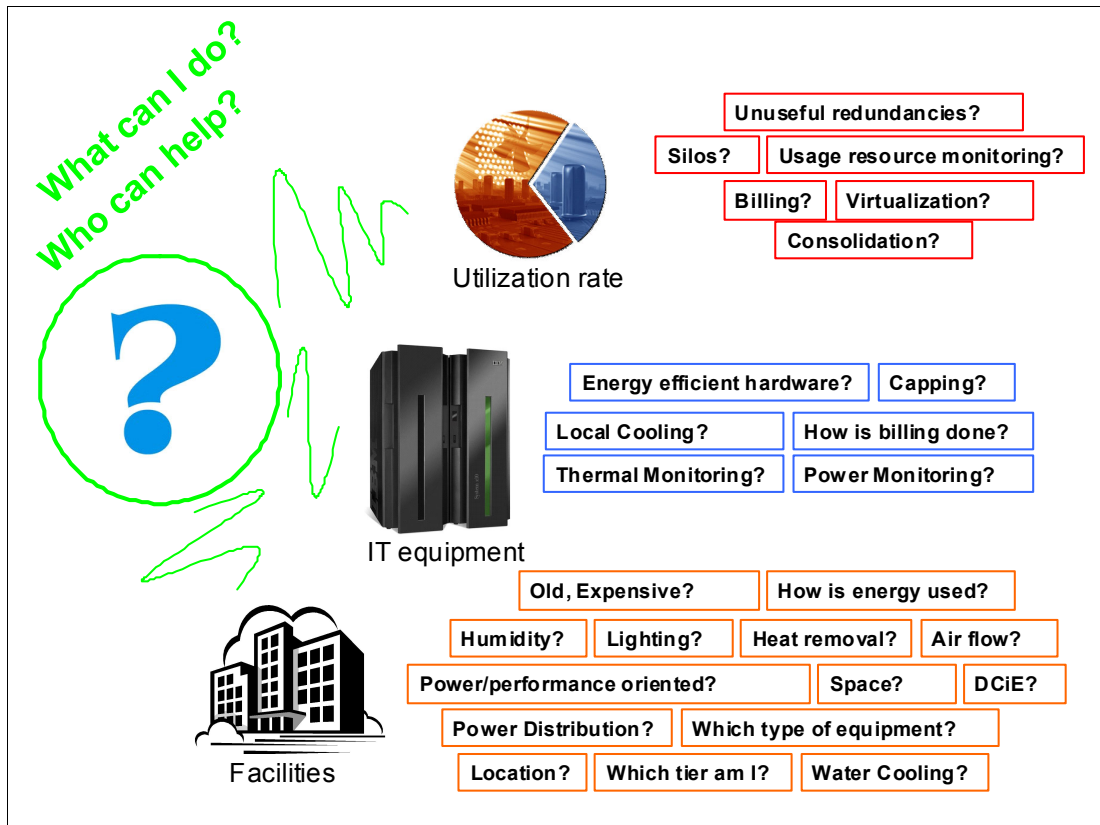


Figure 2-3 Data center assessment: Ask effective questions

Note: To continue the assessment on energy efficiency, access the free Web questionnaire from IBM at the following location:

http://www.ibm.com/systems/optimizeit/cost_efficiency/energy_efficiency/services.html

The output from the assessment is a listing of areas that deserve additional focus for better energy management. It also recommends IBM services that can be deployed to help improve the energy efficiency of your data center.

IBM Global Technology Services also offers assessment services. You can find further details in Section 5.1, "Services IBM can deliver" on page 60.

2.2 Strategy recommendations

After the assessment, you can take actions, but first you should consider the following guidelines in your endeavor to creating an energy efficient data center:

- ▶ Become informed about private and public energy-efficiency initiatives. With the importance of environmental issues, these initiatives are more and more numerous. For example, in the U.S., over 80 local utility and state energy efficiency programs are offering rebates for increasing energy efficiency. To find incentives in the U.S. or Canada, go to: <http://www.energycommons.com>
- ▶ Obtain executive management approval, have a dedicated team, and involve everybody. Becoming energy efficient is a total and complex transformation. It is not a quick fix. You should review your high-availability strategy, study how hardware and resources are allocated to projects, and look at how accounting and billing are done. Having executive management approval is the only way to be sure that everybody is involved. Having approval avoids a situation in which the site and facilities team does something while the IT team does something else. Becoming energy efficient concerns all parts of the IT department. Having a dedicated team is the only way to put energy-efficiency in the list of the business priorities.
- ▶ Consider IT, site, and facilities together. Infrastructure transformation goes together with IT transformation. Removing 1 kW of load from the IT equipment will reduce an additional 1.35 kW of load from the site and facility infrastructure.
- ▶ Begin with an assessment of the current situation.
- ▶ Make a plan, take quick win and long term goals into account, balance the costs against the budget, and execute the plan.
- ▶ Review information in the New Enterprise Data Center (NEDC).¹ The New Enterprise Data Center is an evolutionary new model for efficient IT delivery that helps provide the freedom to drive business innovation. Through a service-oriented model, IT can better manage costs, improve operational performance and resiliency, and more quickly respond to business needs. Becoming energy-efficient is one part of the model. Other parts are, of course, being compliant with the Information Technology Infrastructure Library (ITIL®) and for customers, being dynamically services-oriented and not resources-oriented within a shared and simplified data center. Energy efficiency has to go along with service management, processes, and service level agreements (SLAs). Notice that applying consolidation and virtualization, which are part of energy-efficiency, helps you progress in the direction of a unified infrastructure that has a pool of resources.
- ▶ More than ever, monitoring, accounting, and billing solutions are necessary. Share the costs of energy among all users. *You cannot manage what you do not measure.*

2.3 A summary of best green practices

Table 2-1 summarizes important points to consider when creating a green data center, in what section you can find more information, complexity level, cost, time frame, and payback.

Table 2-1 Best practices

Strategy best practices	Section	Complexity	Cost	Timeframe	Payback
Research private and public energy-efficiency initiatives.	2.2, "Strategy recommendations"	Low	Low	<1yr	High

¹ <http://www.ibm.com/systems/optimizeit/datacenter>

Have executive management approval, have a dedicated team, and involve everybody.	2.2, "Strategy recommendations" on page 12	Low	Low	<1yr	High
Make plans that consider IT and site/facilities together.	2.2, "Strategy recommendations" on page 12	Low	Low	<1yr	High
Begin with an assessment of the current situation.	2.2, "Strategy recommendations" on page 12	Medium	Low	<1yr	High
Make energy costs part of every business case	2.2, "Strategy recommendations" on page 12	Low	Low	<1yr	High
Consider the benefits offered by the New Enterprise Data Center.	2.2, "Strategy recommendations" on page 12	Medium	Low	1-3yrs	High
Share the costs with everybody	2.2, "Strategy recommendations" on page 12	High	Med	1-3yrs	High
IT equipment best practices	Section	Complexity	Cost	Timeframe	Payback
Only buy servers with virtualization and power management features. Lay the foundation for maximum flexibility and a sustained investment into the future.	3.7, "Server virtualization" on page 29	Low	Low	<1yr	High
Enable share everything architectures, such as the mainframe, to drive up utilizations and reduce the need for systems.	3.7, "Server virtualization" on page 29	Low	Low to High	<1yr	High
Install new servers as virtual instances for flexibility.	3.7, "Server virtualization" on page 29	Low	Low	<1yr	High
Move from less efficient to more efficient hardware. Even better, consolidate old servers to virtual servers on efficient hardware. This reduces energy, CO ₂ , and even the space footprint of the data center.	3.7, "Server virtualization" on page 29	Medium	Low	<1yr	High
Identify systems with complementary loads and consolidate them. Partially utilized systems can often make one fully utilized system. Reduce energy, CO ₂ , and space footprint.	3.10, "Integration of energy and systems management" on page 35	Medium	Low	<1yr	High
Manage power consumption of your IT systems. Take advantage of the energy consumption and heat load according to the workload. Reduce the energy and CO ₂ footprint (carbon dioxide emitted) of the data center	3.5, "Power management: The software side" on page 26	Medium	Low	<1yr	High

Measure thermal and power loads of your IT systems. Take control of the energy and CO ₂ footprint.	3.3, "System instrumentation for power management" on page 22	Medium	Med	<1yr	Low
Virtualize your storage if possible. This can have a high initial cost, but it is the foundation for flexibility and energy savings in the future. A sustained investment.	3.8, "Storage virtualization" on page 31	High	Med	<1yr	High
Virtualize your desktops. High initial cost but reduces your total cost of ownership dramatically. Reduces the energy and CO ₂ footprint of the site, simplifies management. A sustained investment.	3.9, "Client virtualization" on page 33	High	High	1-3yr	High
Site and facilities best practices	Section	Complexity	Cost	Time frame	Pay back
Manage airflow. To increase airflow efficiency, have a clear path for the cool air to travel under the raised floor and to get to the loaded areas. Above the raised floor, allow a path for the hot air to return back to the CRAC units.	4.2.1, "Manage airflow" on page 40	Low	Low	<1yr	High
Arrange hot and cold aisles. The hot aisle and cold aisle configuration enables much better airflow management on the raised floor, for both hot and cold air.	4.2.5, "Hot aisle and cold aisle configuration" on page 43	Medium	Low	<1yr	High
Localize cooling. Locate heat exchangers at the heat source, which lessens the need for a CRAC unit. This increases the efficiency of the remaining CRAC units and the capacity of the available cooling within the data center. These heat exchangers are all scalable.	4.3, "Localized cooling equipment options" on page 46	Medium	Med	<1yr	High
Plan for water-cooling in your data center.	4.3, "Localized cooling equipment options" on page 46	Medium	Med	<1yr	High
Newer infrastructure products are more energy efficient.	4.1.3, "Improving physical infrastructure" on page 39	High	High	1-3yr	Medium
Replace the oldest infrastructure equipment first, because it is likely to fail next and is less energy efficient.	4.9, "Recommendations for existing data centers" on page 57	High	High	1-3yr	Medium

Figure 2-4 shows the strategy of moving towards having a green data center, with several recommended steps. It shows the need to coordinate all actions simultaneously with the IT infrastructure and with the people involved in the process.

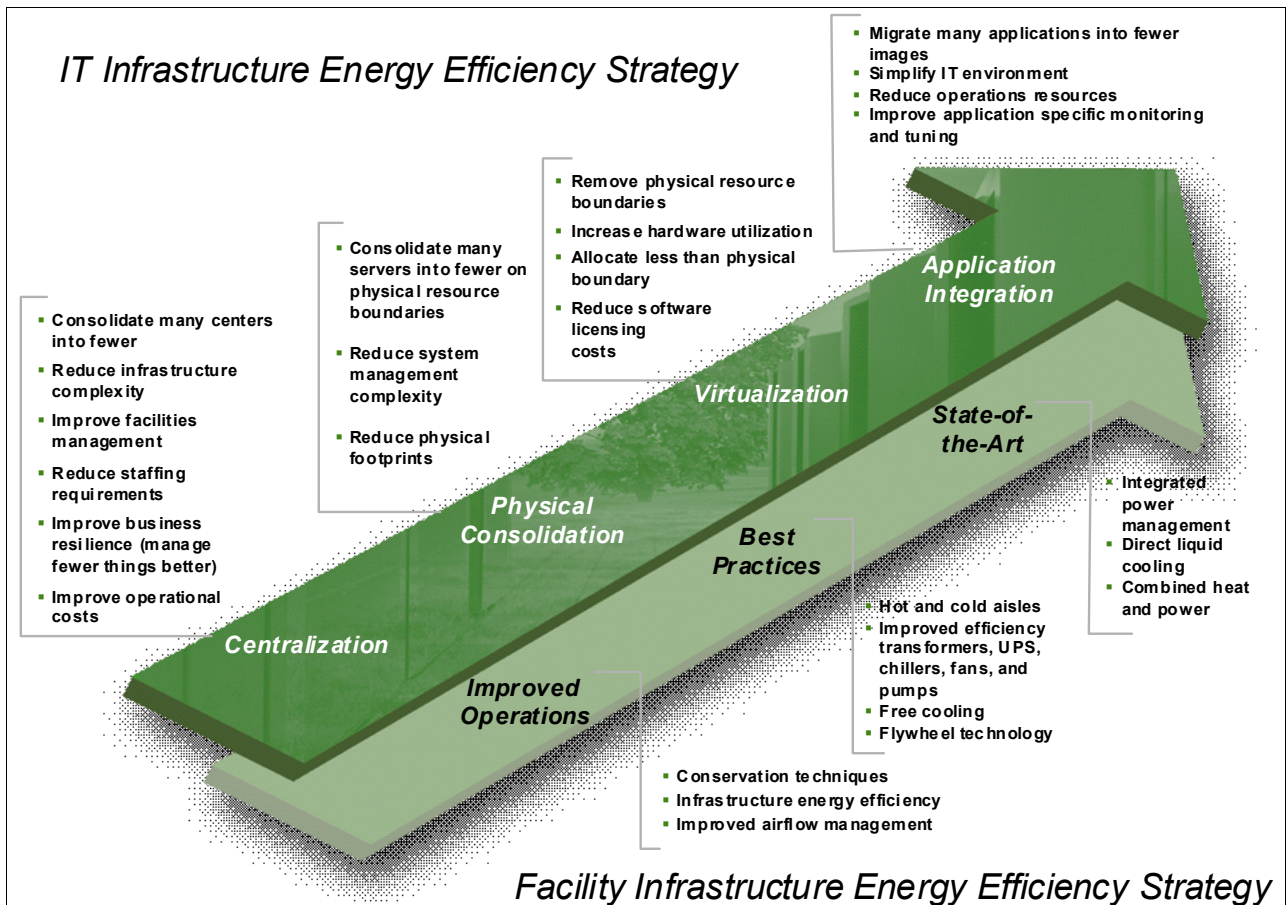


Figure 2-4 Journey towards the green data center



Energy optimization with IT equipment

Electrical energy, when used for computing, gets transformed into heat. Optimizing IT equipment by reducing power and heat at the source has a direct impact on facility efficiencies.

This chapter discusses the IT equipment aspect of a green data center. We describe two approaches to improve IT energy efficiency, show how energy efficiency management can fit into an overall systems management environment, and provide tips for moving your IT systems to a green operation.

3.1 Energy flow in a computer system

Understanding how energy flows through IT equipment can help improve its efficiency. Computers transform information from one form to another. However, the energy that we put into the computer does not actually go into the information itself. Instead, the energy flows to the devices that hold the information. The physical work—energy consumption—results in a change to the state of the devices over time.

Eventually, all electrical energy that goes into the computer is converted to *heat*. For particular appliances, such as a steam iron, this heat generation is the desired effect when we want to remove wrinkles from clothing. However, in computer information processing, the heat is an unneeded by-product. In fact, we spend additional energy to remove the heat from the data center.

3.1.1 How the electric power is used

After power is brought into the data center site (see Chapter 4, “Site and facilities” on page 37), it is fed to the various IT components. The power allocation in a system depends on its architecture and purpose. Figure 3-1 illustrates the relative power allocation to the components of several IBM systems. The total amount of power supplied and the distribution of power to the components varies with the workload that is processed. Each system is built with a certain purpose in mind.

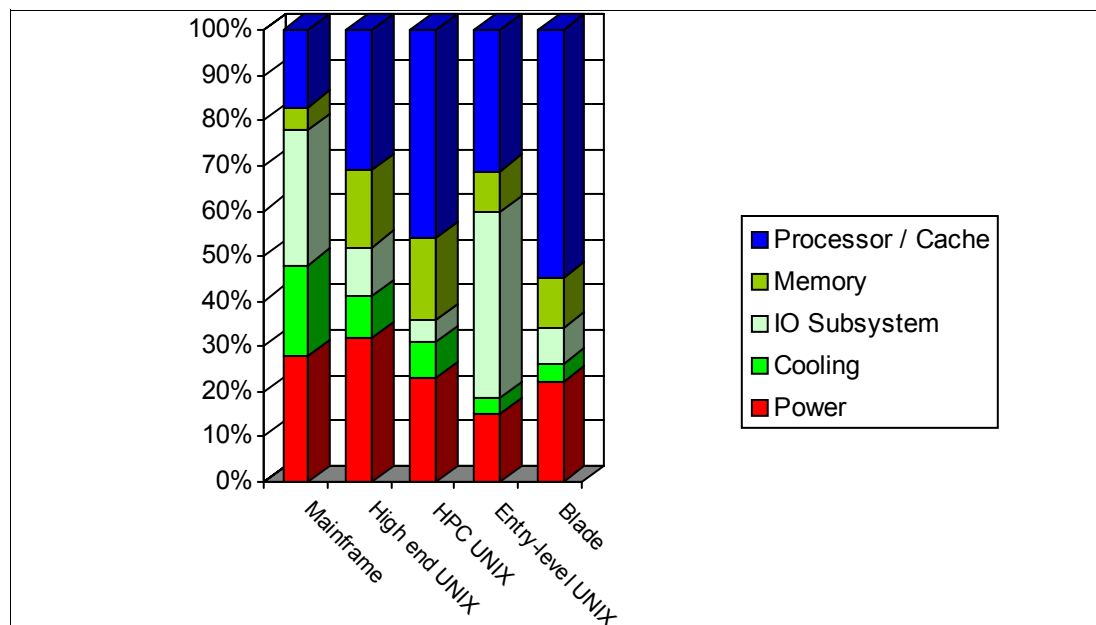


Figure 3-1 Typical relative power consumption by component for typical systems

From left to right, Figure 3-1 shows typical power allocations for mainframe systems (such as the IBM System z10™), high-end UNIX® servers (such as the IBM POWER™ 595), high performance computing (HPC) servers (such as the IBM POWER 575), entry-level UNIX systems (such as the IBM POWER 520), and blade servers, which represent a power-efficient, typical 1U server replacement. Because the power consumed by the processor makes up approximately 20-30% for the mainframe while the blade is over 50%, we use different strategies for each when optimizing their energy effectiveness.

The red portions at the bottom of the vertical bars show the energy required to provide power to each system. Transforming AC power into DC loses some energy. The efficiency of a transformer (power supply) depends on the load, and it is non-linear. The most efficient load is 50-75%. Efficiency drops dramatically below a 50% load, while it does not improve significantly, with higher loads. The challenge is to balance the system so that each component can operate in the most efficient way.

3.1.2 What to do with the heat

The power that is input is transformed into heat output. To prevent the destruction of the computer chips, the heat must be removed from the chip and then from the system. Figure 3-2 illustrates the removal path from the transistor device to the data center.

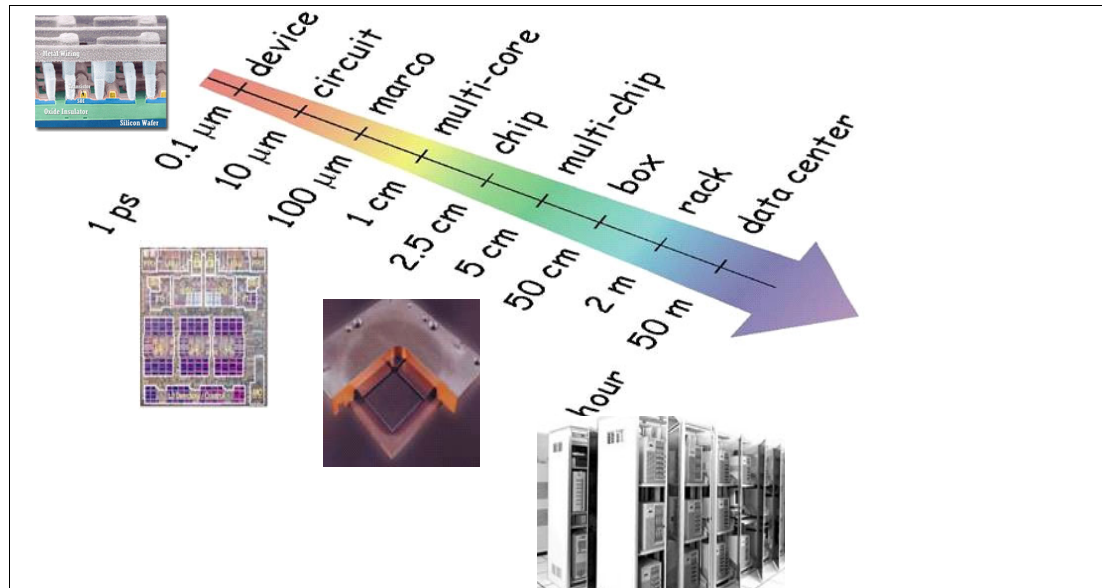


Figure 3-2 Heat flow path from transistor device to data center

Each watt saved in system power results is approximately another watt saved in heat load. These savings also have an effect on the uninterruptible power supply (UPS) and cooling. Therefore, *reducing system power consumption pays back more than double*, which is a big benefit when moving to a green data center.

Air is a very inefficient cooling medium, therefore liquid cooling is increasing in popularity. Water is currently the most common liquid used for cooling. One liter of water can absorb about 4000 times more heat than the same volume of air. As more heat is generated in a smaller space, water cooled systems seem inevitable in the near future. When planning new data centers or revamping current centers, consider that new IT equipment will undoubtedly require liquid cooling for efficiencies.

IBM hardware solution examples

Let us begin with the component having the greatest heat density—the *chip*. Usually the heat is removed from the back of the chip by using massive copper plates. IBM actively researches this field and has current solutions that allow for direct water cooling at the chip's backside. The basic principle is to have water stream through microscopic channels in metal plates. After the heat is transferred to water, it can be handled efficiently.

The IBM z10 processor

The IBM z10 processor modules are cooled primarily with a redundant pair of in-rack refrigerators. This special liquid-cooling is transparent to the owner, because the machine does not have requirements other than power and fresh cool air. The system itself ensures proper cooling by using the internal refrigerators. It is also designed to be completely air-cooled. If the entire redundant liquid cooling system fails, a controlled slowdown of the system will maintain its heat at a level that can solely be managed with its built-in redundant air moving devices.

The POWER 575

The POWER 575 system for high performance computing has water-cooled processors. A single 32-way node operates at 4.7 GHz, and up to 14 nodes can be housed in one frame. In contrast to the System z local refrigerator, the 575 requires a chilled water infrastructure in the data center. However, looking at the rack as a total entity, it manages its cooling autonomously. As Figure 3-3 shows, in addition to the processor cooling system, there are also connections for an optional rear door heat exchanger (RDHX). Chilled water is fed to the redundant modular water units (MWU) and then distributed to the nodes and the RDHX.

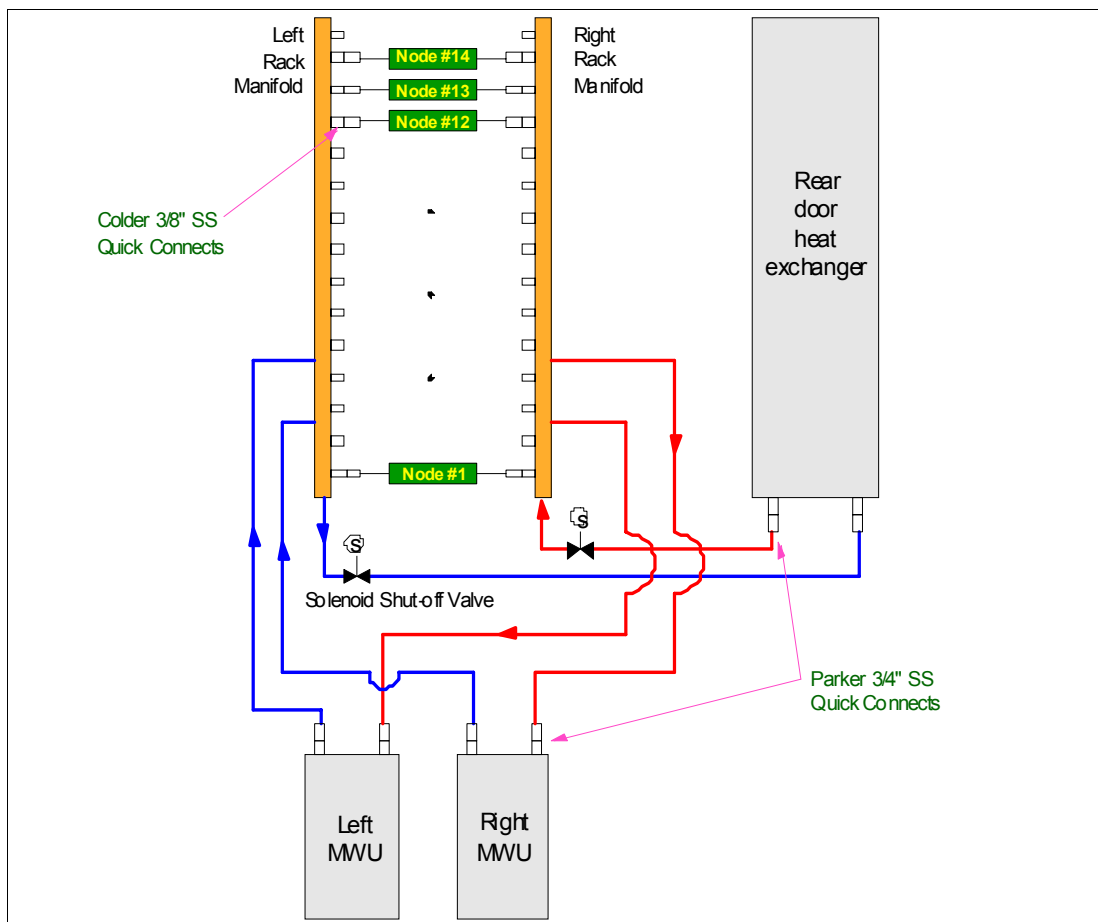


Figure 3-3 Schematic cooling system overview of the POWER 575 system

The rear door heat exchanger (RDHX)

The rear door heat exchanger (RDHX) is described in Section 4.3.1, “IBM Cool Blue rear door heat exchanger” on page 47 in more detail. It is another device to help reduce the heat load in your data center. Any of these devices are helpful when there are hotspot problems, or if the air-based cooling of the data center is at its limit. When the chilled water infrastructure is

in place, the RDHX is a very attractive solution because the dissipated heat now bypasses the CRAC unit and can be more efficiently absorbed by the chillers.

The pictures in Figure 3-4 on page 21 were taken in a production data center. The rack houses nine System p5™ 550Q servers. The RDHX reduces the interior temperature of 46.9° Celsius (114.8° Fahrenheit) to 25.1° Celsius (77°F) on its outside.



Figure 3-4 Heat exchanger in action

iDataPlex

System x takes a new approach to solving data center challenges through its latest innovation, iDataPlex™. This is a data center solution for Web 2.0, HPC cluster, and corporate batch processing clients who are experiencing limitations of power, cooling, or physical space. iDataPlex servers help pack more processors into the same power and cooling envelope, better utilize floor space, and *right size* data-center design. With the iDataPlex solution, less power per processor means more processing capacity per kilowatt. The iDataPlex can run cooler to deliver greater reliability.

iDataPlex offers flexibility at the rack level. It can be cabled either through the bottom, if it is set on a raised floor, or from the ceiling. Front-access cabling and Direct Dock Power allow you to quickly and easily make changes in networking, power connections, and storage. The rack supports multiple networking topologies including Ethernet, InfiniBand®, and Fibre Channel. It is an excellent option for a green data center because it offers the following benefits:

- ▶ Efficient power and cooling
- ▶ Flexible, integrated configurations
- ▶ New levels of density
- ▶ Cost-optimized solution
- ▶ Single-point management
- ▶ Data center planning and power and cooling assessment

3.2 Working with energy efficiency

In this section, we introduce two IT manipulation handles, power and workload, that we then discuss in later sections.

Key questions from the previous section are:

- ▶ What makes IT equipment energy efficient?
- ▶ How do systems compare to each other?

One approach to answering the two questions is to examine the relationship between the power that is used and the workload that is processed. This comparison provides a handy concept of an energy efficient IT system. Choose equipment that is efficient, and when options exist, go with the system that is more efficient.

Note: The more energy efficient a system is, the more workload it can process with a certain power input, or the less power it needs for processing a certain workload. In other words, being energy efficient means that you can either decrease power and do the same work or increase workload with the same power.

Most efficiency improvements result in a combination of these two solutions. For example, we could assign the workload of computer A to another, more power-efficient computer B, then switch off computer A. This means we choose from both alternatives at a time.

Note: For IT equipment energy efficiency, we will primarily look at the power to workload relationship. But there is another workload-related characteristic we do not want to miss: the workload to time relationship, better known as performance. Using physics, we can see our power-to-workload relationship as an analogy to mechanical work, so we view performance as analogous to mechanical power.

After each change, we want to know that we actually improved our overall energy balance. Obviously, this can only be achieved with proper before-after evaluation. The very basic starting point and first step in all our efforts must therefore be to *enable proper measurement*.

3.3 System instrumentation for power management

Remember, *you cannot improve what you do not measure*.

As you move toward an energy efficient system, knowing how the system actually behaves is essential. You must understand about the power consumption of the system and its heat dissipation.¹ Although there are many systems in the marketplace, in this section, we first look at the instrumentation that IBM systems provide.

3.3.1 IBM systems with embedded instrumentation

Virtually all newer IBM systems have built-in instrumentation for power usage measurement and thermal sensors. Although system-specific tools are also available, it is primarily the IBM Active Energy Manager tool that can read the information provided by these sensors, regardless of the particular system platform (Figure 3-5 on page 23). The tool allows us to display the current value and to initiate actions based on the overall system status.

¹ The instrumentation for workload measurement is out of scope for this document. Tracking the operating system's utilization and using accounting measurement tools is (or should be) an established practice for all data centers, whether going green or not.

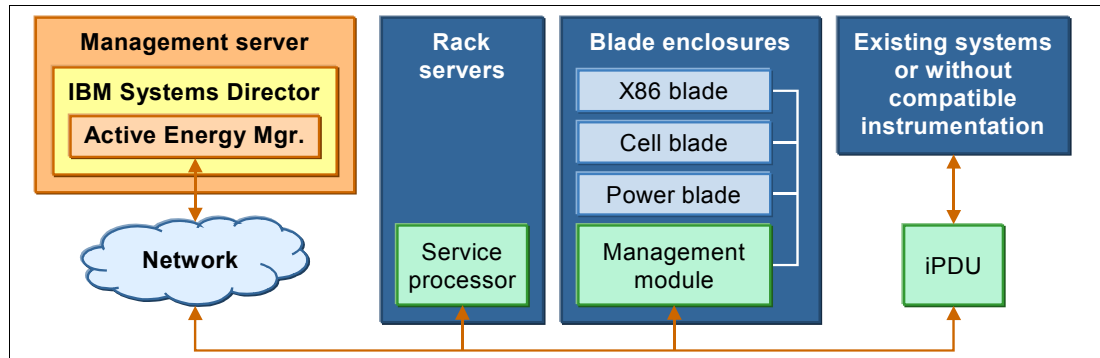


Figure 3-5 How system instrumentation and Active Energy Manager work together

The following overview describes instrumentation in IBM systems:

- ▶ System z Gas Gauge: The System Activity Display within the Hardware Management Console shows the machine's current power and temperature information.
- ▶ EnergyScale™ technology in POWER6™ processors, which are built into POWER blades and the Power System computers. Described in greater detail later, EnergyScale and Power Trending monitors the actual power and thermal load of the processor.
- ▶ System x servers and BladeCenter chassis provide their power and monitoring information to the Active Energy Manager product, also described later in more detail.

3.3.2 IBM intelligent power distribution units

For systems that do not have embedded or manageable instruments on board, intelligent Power Distribution Units (iPDUs) are available. The IBM DPI® C13 PDU+ and IBM DPI C19 PDU+ iPDUs contain versatile sensors that provide power consumption information of the attached devices, and environmental information such as temperature and humidity. The iPDU's serial and LAN interfaces allow for remote monitoring and management through a Web browser, any SNMP based Network Management System, Telnet, or a console over a serial line. Events can be notified by SNMP traps or e-mail and it is possible to send out daily history reports, also by e-mail. Not surprisingly, Active Energy Manager is also capable of managing the iPDUs.

3.4 Power management: The hardware side

While we described the tools for monitoring a system in a previous section, we now concentrate on the possibilities that IBM systems provide for actively managing their power consumption. (In the next section, 3.5, "Power management: The software side" on page 26, we show how the IBM systems can be managed from one point.)

3.4.1 IBM POWER6 EnergyScale

The IBM POWER6 processor-based systems offer the *IBM EnergyScale* technology. It provides several power management functions:

- ▶ Power Trending: Enables the computer to collect power usage data and store the data internally. The data can then be displayed with IBM Active Energy Manager.
- ▶ Power Saver Mode: Drops processor voltage and frequency by a pre-determined percentage. While Power Saver Mode still allows proper and safe operation, it helps

reduce peak energy consumption. For example during low CPU utilization cycles in the night, the processor may run in Power Saver Mode. This feature resembles the Intel® SpeedStep or AMD™ PowerNow!™ technologies.

- ▶ **Power Capping:** Enforces a specified power usage limit. This feature is handy when there are general power limitations such as maximum power available for a set of systems. However, this should not be used as a power saving feature like Power Saver Mode, as it has a significant impact on performance.
- ▶ **Processor Core Nap:** Uses the IBM POWER6 processor low-power mode (called Nap) for reduction of power consumption by turning off the clock for the core. Signalled by the operating system, the hypervisor controls cores to enter or leave Nap mode.
- ▶ **EnergyScale for I/O:** Enables Power System models to automatically power off pluggable PCI adapter slots. Slots are automatically powered off if they are empty, are not assigned to a partition, or their assigned partition is powered off. This saves approximately 14 W per slot.

The power management functions can be configured using the Advanced System Management Interface (ASMI), the Hardware Management Console (HMC), or Active Energy Manager. However, only Active Energy Manager allows configuration of Power Capping and evaluation of Power Trending. Note, that all POWER6 based systems employ this technology.

3.4.2 IBM BladeCenter features

IBM BladeCenter servers offer a high-density solution for the data center. Up to 14 blade servers can be held in a single 7U high BladeCenter. The BladeCenter provides its infrastructure services—such as power, networking, and cooling—to the installed blade servers. Because power and cooling is concentrated in one device for up to 14 servers, the BladeCenter solution is far more efficient than a comparable installation with 14 individual machines. Especially noteworthy are the very efficient redundant power supplies (up to 91%, already 85% efficient at 20% load).

Another point to consider is the cooling concept. A BladeCenter provides central redundant blowers that use less total power than individual fans, compared to an appropriate number of conventional servers. Blade servers are available on different platforms including Intel, AMD, POWER, and cell processors. So, simply replacing a single conventional server with an equivalent blade server would be a step in power reduction.

Power management for both BladeCenters and System x is already supported in the planning phase, where a power configurator² tool can estimate the power requirements for a specific configuration, which usually is below the machine's label rating.

BladeCenters also provide power capping in combination with Active Energy Manager. In conjunction with the very efficient redundant power supplies, BladeCenters are an excellent platform for power management. Besides, dependent on the respective blade server platform, the power management options of the particular server, for example AMD PowerNow!, Intel SpeedStep®, and POWER6 EnergyScale are available.

² The IBM System x and BladeCenter Power Configurator can be downloaded from:
<http://www.ibm.com/systems/bladecenter/resources/powerconfig/>

3.4.3 System x features

System x has many of the same key features as the IBM BladeCenter, such as:

- ▶ Design elements that lead to reduce power consumption include the option of using lower power 50 W processors, the use of lower power 2.5 inch disk drives, efficient cooling design called Calibrated Vector Cooling
- ▶ Enterprise servers featuring the fourth-generation of X-Architecture® uses DDR II memory which uses up to 37% less power than the Fully Buffered DIMMs used by the competition
- ▶ Power Configurator for preinstallation planning
- ▶ Power and thermal monitoring and trending by Active Energy Manager
- ▶ Power capping with Active Energy Manager (see section 3.5, “Power management: The software side” on page 26)
- ▶ Extensive support of virtualization environments to enable server consolidation (VMware, Xen, and Microsoft® Virtual Server)

3.4.4 ACPI and friendly power saving features

The Advanced Configuration and Power Interface (ACPI) allows the operating system to take power-saving actions, such as putting the processor to sleep, spinning the disks down, and others. While primarily used for mobile and desktop computers, servers offer these features too. Note, however, that your system’s mainboard, chipset, timer, BIOS, CPU, and the operating system itself have to be ACPI-enabled. ACPI is broadly introduced in the marketplace and is well-supported. However, Linux often runs into poorly documented manufacturer-specific control tables and their features.

Processor-specific features like intel SpeedStep, AMD PowerNow!, or Cool’n’Quiet can also help reduce overall system power consumption. By reducing the processor clock frequency and core voltage, its power requirements decrease. Obviously, this carries with it some thermal relief. (Figure 3-1 on page 18 shows the candidates for this type of power management, which are those systems whose processor plays a dominant role in power consumption.)

3.4.5 High input voltages

IBM has designed all System z and some POWER6 systems (575 and 595) with a 200-480 VAC 3-phase power input so that they can be installed in any location in the world. In some instances, the customer can choose to run a system from a high voltage (380-480 VAC) or a lower voltage (200-240 VAC). When this choice is available, the higher voltage group provides the following advantages:

- ▶ Smaller line cords, line cord connectors, and building wiring, which all reduce installation cost and complexity.
- ▶ Step-down transformers on the computer floor (such as 480V-to-208V) are not required, which can reduce the installation cost significantly.
- ▶ The 595 and typically the building infrastructure that feeds power to it offer higher energy efficiency, reducing energy costs.
- ▶ As energy prices increase, the savings become more significant.
- ▶ Using the UPS system in front of the system at the building level, eliminating the building transformer, results in even greater savings.

3.5 Power management: The software side

The primary tool for power management of IBM systems is Active Energy Manager (AEM), which is an extension to the IBM Director product. It provides a common systems management environment for all the IBM platforms. Active Energy Manager offers the following basic features:

- ▶ Measures and displays current power and temperature data from the managed systems.
- ▶ Provides trend data over selected periods of time.
- ▶ Sets power capping for IBM systems if their firmware supports it, and manages the Power Saver mode for Power Systems.

Figure 3-6 shows a scenario that demonstrates how AEM can help in optimizing your rack or BladeCenter layout, as follows.

1. Assume that power trending indicates that the server’s actual power consumption remains below its labeled rating. Find a certain power usage that the server never exceeds.
2. Employ AEM for power capping. Cap the power for the systems to the maximum observed level. Remember, power capping sets the limit to something that should not happen under normal circumstances, because of the significant performance impact it implies. But obviously, the allocated power for the rack as a whole and the power actually consumed are now under control. The power that previously over-allocated for each single system can now be managed at a rack-level of over-allocation.
3. Use the previously over-allocated power effectively by adding more systems to the rack. Ideally, install new systems with a better energy footprint, replacing less efficient machines. See 3.6, “Consolidation and virtualization overview” on page 27 for more information.

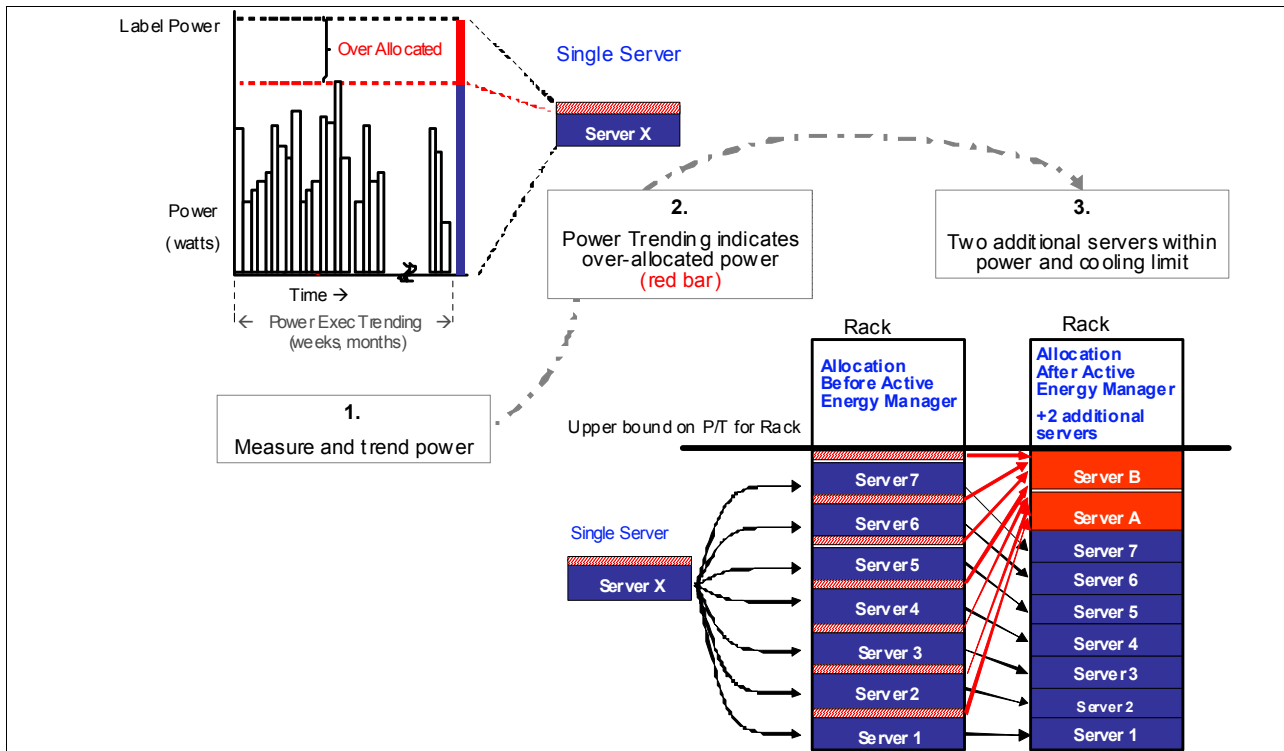


Figure 3-6 Optimize rack layout with the help of Active Energy Manager

When you have a power management infrastructure in place, you can take additional steps. For example, you can locate hot spots in your data center, perhaps supported by physical location information about your servers. Alternatively, after you know the hot spots, you can prevent them by provisioning your applications to *cool* servers. If your utility provider offers load dependent rates, or day and night rates, why not optimize your power consumption to these? Many options are available. However the first step is to build the infrastructure.

3.6 Consolidation and virtualization overview

While we discussed the concept of power management in preceding sections, we now look at the second key energy efficiency technique for workload management: *consolidation*. We concentrate our workload together, to have it processed by as few machines as possible, and have these machines work at their most efficient energy level. A powerful tool to help us do so is *virtualization*.

Note: Workload is not limited to the applications a computer has to run. We also want to understand workload as storage usage or network usage.

3.6.1 Consolidation: A key in energy efficiency

Figure 3-7 illustrates the idea of *consolidation*. Let us assume we have four systems, each running two applications (APP). Also, each machine consumes 2 kW power, 8 kW in total. However, as is often the case for small x86 servers, they are utilized at only 10%.

If we are able to consolidate these eight applications to a single, more powerful server and run this at a utilization of 70% with a power usage of 4 kW, this single server can operate more energy efficiently, according to our definition described in 3.2, “Working with energy efficiency” on page 21. In addition, if we perform a simple power management technique of switching off the previous four systems, the result in a total power consumption is 4 kW and a 70% utilized system.

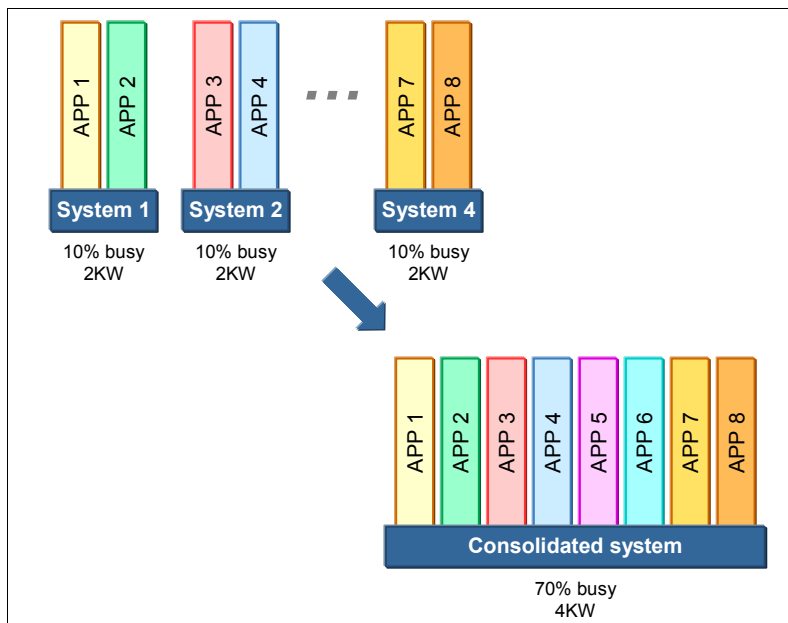


Figure 3-7 Consolidation of applications from under utilized servers to a single, more efficient server

As we indicated earlier, a decrease in overall power consumption is not the only factor. Hand-in-hand with the power reduction goes the same amount of heat load reduction and another add-on for the infrastructure. This double reduction is the reason why consolidation is an enormous lever to moving to a green data center.

However, a particular drawback of consolidation is that none of systems 1 through 4 is allowed to be down during the time that the respective applications are moving to the consolidated system. So, during that migration time, higher demands on resources might occur *temporarily*.

3.6.2 Virtualization: The greenest of technologies

In addition to consolidation is *virtualization*, the concept of dealing with abstract systems. It can dramatically reduce the amount of IT equipment needed in a data center. Virtualization eliminates the physical bonds that applications have to servers, storage, or networking equipment. A dedicated server for each application is inefficient and results in low utilization. Virtualization enables “car pooling” of applications on servers. The physical car (server) may be fixed but the riders (applications) can change, be very diverse (size and type), and come and go as needed.

The example in Figure 3-7 on page 27 shows how specific applications were moved to another system with a better energy footprint. In our simple case, we assume all systems are running at the same operating system level. However, what if the applications require different operating system levels or even completely different operating systems? That is where virtualization comes into play.

Important: Virtualization enables us to take the old system *as a whole* and reimplement it on the target machine.

The term virtualization is widely used and has several definitions, as follows:

- ▶ Can create logical instances of a computer system consisting of CPU, memory, and I/O capabilities.
- ▶ Can be put together from other virtual components.
- ▶ Can consist of a virtual CPU or virtual memory and disk.
- ▶ Can be a virtual network between a virtual computer and the outside world.

To have real work done by a virtual system, it must run on a real system. Obviously additional intelligence is required to do this. There are pure software solutions, a system’s firmware may offer virtualization features or such features may be hardwired into the system. Many of the current processor architectures have virtualization features integrated, which can be taken advantage of by software solutions such as the IBM System z and p machines. In the field, various other solutions are available, such as VMware Server, VMware ESX, Microsoft Virtual Server, and Xen.

To continue with our example, using virtualization gives a slightly different picture, as shown in Figure 3-8 on page 29. Instead of moving the applications to the consolidated server, we now virtualize the existing systems 1 through 4 on our consolidation target. The effect is clear: not only is the application moving, but its complete operating environment has moved with it.

Taking a closer look, we find other attractive opportunities:

- ▶ Consider the three separate systems. To communicate, they require a network infrastructure such as NICs, cables, and switches. If our virtualization system supports network virtualization, this infrastructure is no longer needed. The virtualized systems can

communicate using the virtualization system's capabilities, often transferring in-memory data at enormous speed. Performance and energy efficiency increase because the network components are dropped. Once again, this reduces the need for site and facilities resources.

- ▶ Each of the separate systems has its own storage system, namely disks. The virtualized systems can now share the disks available to the virtualization system. By virtualizing its storage, the virtualization system can provide optimal disk capacity—in terms of energy efficiency—to the virtualized systems.

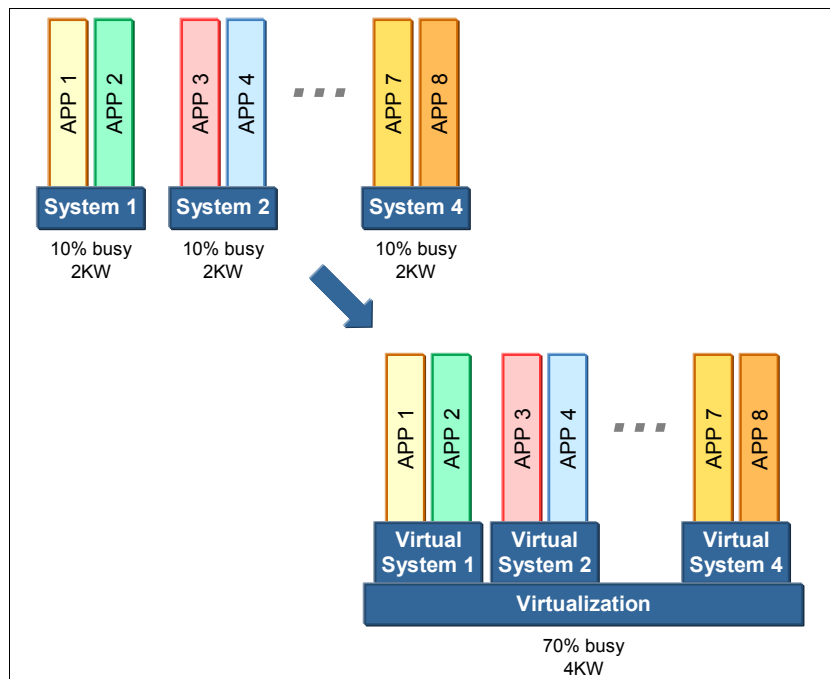


Figure 3-8 Virtualization allows us to consolidate systems the way they are

3.7 Server virtualization

This section discusses the IBM techniques that are available for server virtualization, the most attractive approach to consolidation. In many cases, it is the easiest and most effective way to transfer workload from inefficient, underutilized systems to efficient, well-utilized equipment.

It is worth noting that since 1964, the mainframe, with its shared architecture and built-in virtualization, has been the gold standard for server virtualization. System z is highly regarded as a consolidation platform in terms of diversity of workloads and number of instances virtualized while maintaining high qualities of service.

3.7.1 Partitioning

Partitioning is sometimes confused with virtualization, but the partitioning feature is a tool that supports virtualization. Partitioning is the ability of a computer system to connect its pool of resources (CPU, memory, and I/O) together to form a single instance of a working computer or *logical partition* (LPAR). Many of these LPARs can be defined on a single machine, as long as resources are available. Of course, other restrictions apply, such as the total number of LPARs a machine can support. The power supplied to the existing physical

computer system is now used for all these logical systems, yet these logical systems operate completely independently from each other.

Important: LPARs each work independently at the maximum performance of the underlying system. All partitions share the energy provided to the overall system.

LPARs have been available on the IBM System z since the late 1980s and on System p since approximately 2000. Although the System z and System p partitioning features differ in their technical implementations, they both provide a way to divide up a physical system into several independent logical systems.

3.7.2 Special virtualization features of IBM systems

Besides providing logical computers, each running an independent operating system instance, IBM systems offer further virtualization features, such as:

- ▶ **Micropartitioning:** Allows assignment only of fractions of CPUs to virtual machines. In uncapped mode, this assignment may become dynamic.
- ▶ **HiperSockets™ on System z:** A network virtualization technique that allows LPARs to communicate through memory instead of using dedicated networking peripherals. This allows for a very large communication bandwidth compared to normal networks.
- ▶ **Live Partition Mobility on System p:** Allows you to move active LPARs from one physical machine to an other without disrupting the respective applications.
- ▶ **Virtual Ethernet on System p:** Is a similar technique to HiperSockets. LPARs can communicate using the Hypervisor. Instead of using traditional networking peripherals, the communication takes place within the machine.
- ▶ **Virtual I/O on System p:** The use of a virtual I/O server (VIO server) on a System p machine allows for further I/O virtualization. The VIO server can, for example, divide up physical disks in the machine and offer them as single virtual disks to the LPARs.
- ▶ **Multiple image facility (MIF) on System z:** Enables channel sharing among Processor Resource/Systems Manager™ (PR/SM™) logical partitions. With MIF, multiple logical partitions can directly share channels and optionally share any of the control units and associated I/O devices configured to these shared channels. MIF also provides a way to limit the logical partitions that can access a reconfigurable channel or a shared channel to enhance security. Dynamic Channel Management takes MIF a step further to dynamically adjust the channel configuration in response to shifting workload patterns. It is a function of the Intelligent Resource Director and works together with WLM to virtualize I/O paths out to the controllers.

The partitioning technique is based on a *hypervisor*. The hypervisor or *virtual machine monitor* (VMM) establishes a virtual machine to the operating system. In theory, the operating system is unaware of this intermediate layer between itself and the hardware (full virtualization). That is the case for System z. However, for Power Systems, their operating system must be aware of the Hypervisor (paravirtualization). Therefore, AIX and Linux for POWER architecture include the necessary support.

3.7.3 Other virtualization techniques

Other virtualization techniques are available, in addition to partitioning. Popular in the market are the VMware products, Xen, and Microsoft Virtual Server. Also, hardware manufacturers extend their products to support virtualization.

VMware ESX Server and Microsoft Virtual Server come with a hypervisor that is transparent to the virtual machine's operating system. These products fall into the full virtualization category. Their advantage is their transparency to the virtualized system. An application stack bound to a certain operating system can easily be virtualized, as long as the operating system is supported by the product.

VMware offers a technology for moving servers called VMotion. By completely virtualizing servers, storage, and networking, an entire running virtual machine can be moved instantaneously from one server to another. VMware's VMFS cluster file system allows both the source and the target server to access the virtual machine files concurrently. The memory and execution state of a virtual machine can then be transmitted over a high speed network. The network is also virtualized by VMware ESX, so the virtual machine retains its network identity and connections, ensuring a seamless migration process.³ System p Live Partition Mobility offers a similar concept.

Xen uses either the paravirtualization approach (mentioned in 3.7.2, "Special virtualization features of IBM systems" on page 30), as the POWER architecture does, or full virtualization. In the partial approach (paravirtualization), virtualized operating systems should be virtual-aware. Xen for example, requires virtual Linux systems to run a modified Linux kernel. Such an approach establishes restrictions to the usable operating systems. However, while they are hypervisor-aware, different operating systems with their application stacks can be active on one machine. In the full approach, the hardware must be virtual-aware, such as Intel's Vanderpool or AMD's Pacifica technology. In this case, running unmodified guests on top of the Xen hypervisor is possible, gaining the speed of the hardware.

Another technique is operating system level virtualization. One operating system on a machine is capable of making virtual instances of itself available as a virtual system. Solaris containers (or zones) are an example of this technique. In contrast to the other techniques, all virtualized systems are running on the same operating system level, which is the only operating system the machine provides. As mentioned in the introductory section, this can become a very limiting restriction, especially when consolidating different server generations. Often the application stack is heavily dependent on the particular operating system. We reach a dead end when we want to consolidate servers running different operating systems such Windows® and Linux.

Attention: In addition to the network virtualization products mentioned in this paper, a popular virtualization technique is to combine related applications on one central server or complex. This allows the networking between them to be done internally at computer speeds rather than network speeds, and it saves the cost of networking hardware and software.

3.8 Storage virtualization

Computer systems are not the only candidates for virtualizing; storage can be virtualized too. This section describes IBM SAN Volume Controller, which provides a virtual pool of storage consisting of SAN-attached physical storage devices. In addition, the IBM System z platform includes storage virtualization capabilities in its hardware and its DFSMS™ software under the z/OS® operating system.

These devices can be part of an overall Information Lifecycle Management (ILM) solution being offered by IBM. This approach starts from a business perspective and aligns the cost for storing information with the information's value. Its foundation is tiered or virtualized

³ Quoted from: <http://www.vmware.com/products/vi/vc/vmotion.html>

storage. ILM is not the scope of this paper, however, keep in mind that optimizing your storage landscape by adapting it to your actual needs can be a very green strategy.

Dynamic address translation (DAT) is the process of translating a virtual address during a storage reference into the corresponding real address. If the virtual address is already in main storage, the DAT process can be accelerated through the use of a translation look-alike buffer. If the virtual address is not in main storage, a page fault interrupt occurs, z/OS is notified, and the page is brought in from auxiliary storage. Looking at this process more closely reveals that the machine can present any one of a number of different types of faults. A type, region, segment, or page fault is presented, depending at which point in the DAT structure invalid entries are found. The faults repeat down the DAT structure until, ultimately, a page fault is presented and the virtual page is brought into main storage either for the first time (there is no copy on auxiliary storage) or by bringing the page in from auxiliary storage. DAT is implemented by both hardware and software through the use of page tables, segment tables, region tables and translation look-alike buffers. DAT allows different address spaces to share the same program or other data that is for read only. This is because virtual addresses in different address spaces can be made to translate to the same frame of main storage. Otherwise, there would have to be many copies of the program or data, one for each address space.

3.8.1 IBM SAN Volume Controller

The SAN Volume Controller (SVC) is a hardware device that brings storage devices in a SAN together in a virtual pool. This makes your storage appear as one logical device to manage. To the connected computers, SVC offers virtual disks as ordinary SCSI devices. On the SAN side, SVC integrates various storage subsystems, even multivendor, and takes care of the correct block mapping between the SAN devices and the virtual disks for the computers. Figure 3-9 illustrates how it works.

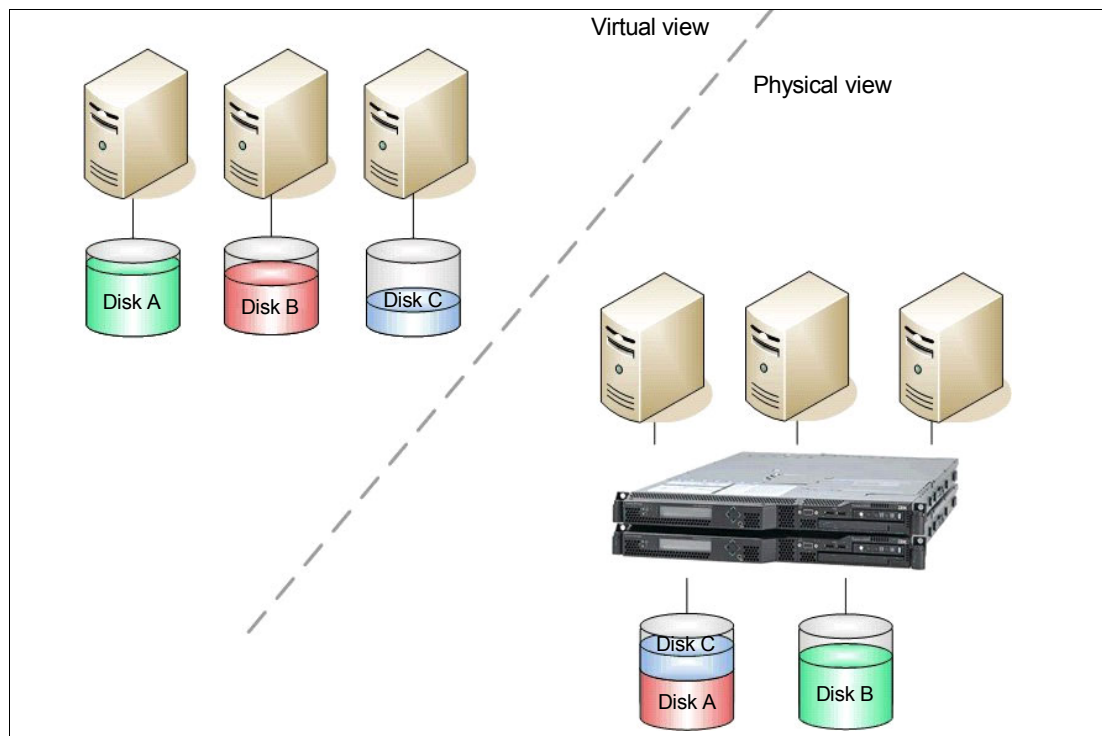


Figure 3-9 Storage virtualization: virtual view and physical view

Obviously, storage virtualization is another tool for consolidation. If underutilized, disks can be virtualized and consolidated in a SAN, the data can reside in more efficient storage pools.

The SVC supports migration of data among the connected device and to remote sites for redundancy or backup purposes. And it can help manage storage hierarchies, where low-activity or inactive data can be migrated to cheaper storage. The integrated cache on the other hand is able to improve the performance of lower tier storage. The data migrations are managed transparently, so they do not interrupt the applications.

The following points make the SVC an attractive tool for an energy efficient storage strategy.

- ▶ Data migration from older to newer, more efficient systems can happen transparently
- ▶ Tiered storage enables you to use media with a smaller energy footprint, while the SVC cache improves its performance.
- ▶ Consolidation of the system individual storage devices to virtual storage has the same effect, of increasing storage utilization, as is shown for server virtualization.

Storage virtualization requires more effort than server virtualization, often requiring us to rethink the existing storage landscape. During consolidation, large amounts of data must be moved from the old systems to the consolidated storage system. This can become a long task that requires detailed planning. However, when done, the effect can be enormous because now storage can be assigned to systems in the most flexible way.

3.8.2 Virtual tapes

Looking at the cost of storage, tapes are the cheapest media on which to store data. They offer the largest storage volume at the lowest cost, which is the reason that they are the optimal backup media. Tapes have a long latency compared to disks. This is not always a benefit. Currently, data centers are faced with a time limitation for backing up and probably restoring their data, because the time frames for backups shrink, while the amount of data to back up expands. For this reason, many sites prefer large disk-based backup systems instead of tapes.

Tape virtualization may be a solution to this. A virtual tape server behaves just like a tape library, but a very fast one. This is made possible with internal disk arrays and a migration strategy to export to and import from real tape libraries.

The IBM Virtualization Engine™ TS7700 is a mainframe virtual-tape solution that optimizes tape processing. It has a fully integrated tiered storage hierarchy of disk and tape, and can leverage the benefits of both technologies. Deploying this innovative subsystem can help reduce batch processing time, total cost of ownership and management overhead. For example, if the batch process is constrained, the TS7700 Virtualization Engine's processing power, cache capacity, and support for 4 GBps IBM FICON® attachment may help alleviate bottlenecks and reduce the batch window. Another system, the TS7520 Virtualization Engine, is a general purpose system for non-mainframes.

3.9 Client virtualization

A great potential in energy savings is client, or desktop, virtualization. IBM has estimated an energy savings of more than 60% by using client virtualization. In a typical workplace, the installed PCs show marginal usage rates. Except when launching an application, an office PC spends most of its time waiting for the user to press a key or click a mouse. However, the PC continues to need a considerable amount of energy to operate, heats up its surrounding

environment, and produces noise. Desktop virtualization can improve the situation dramatically.

The underlying principal of client virtualization is to replace the office workstation with a box having a much smaller energy footprint. The needed computing power is moved into the data center. This does not sound excitingly new, remembering terminal-to-host applications or the Xwindows system, but those configurations also had their advantages and today's virtualization techniques make this approach even more attractive. Benefits are many, and not only to the energy balance.

Software deployment, for example, can become a mess if the desktop machine contains many historically grown software stacks. If we do not want to bother users by running updates during the day, machines can run overnight. An erroneously switched off machine can make a whole deployment fail. Central machines reduce the risk and cost.

The three strategies for using virtualization in desktop consolidation are:

- ▶ **Shared services:** The client PC continues to run the full operating system and has access to remote service providers for special applications. Examples of this strategy are Pure Windows Terminal Service (WTS), Exceed on Demand, and Citrix Metaframe and Presentation Server.
- ▶ **Virtual clients:** The client PC is replaced by a *thin client*. The user's PC runs now as a virtual machine on a central server, while the thin client supports only appropriate protocols such as WTS, Windows, and so forth, but cannot run programs locally.

The *IBM virtual client solution* is an example of the virtual client strategy. It is based on BladeCenter or System x servers. VMware virtualizes these physical servers into multiple virtual servers. Each of these servers can be assigned to a user (thin client). Connections between clients and servers are managed with Leostream Virtual Desktop Connection Broker, which itself runs in a virtual machine. Provisioning of the virtual machines to the respective servers is supported with the IBM Virtualization Manager.

- ▶ **Workstation blades:** The client PC is replaced by a special device because the advanced graphics requirements must be handled with special compression features or special hardware. There is a 1:1 assignment between user machine and central workstation blade.

An IBM solution with workstation blades is the HC10 Blade Workstation and the CP20 Workstation Connection Device approach. The HC10 Blade Workstation serves as a replacement for desktop workstations, which can take advantage of all BladeCenter features (as discussed in 3.4.2, "IBM BladeCenter features" on page 24), such as optimized power and cooling efficiency, or the ability of power management using AEM. On the desktop side, the CP20 Workstation Connection Device is the user interface. This device has no moving parts and serves as the remote KVM device for the HC10. It provides connections for a keyboard, video display, and mouse (KVM) and optionally other USB devices. To support advanced graphic applications, the video output from the HC10 is compressed before it is sent to the CP20 device.

Connections and assignments between the CP20 and HC10 devices are maintained by the connection management software, also known as connection broker. The connection broker assigns the CP20 to its corresponding HC10 using different modes:

- Fixed seating assigns a distinct CP20 to a specific HC10
- Free seating assigns a distinct user to a specific HC10
- Pooling as a subset of free seating assigns users to a pool of HC10s to which they are connected.

In contrast to the virtual client solution, the HC10 and CP20 solution aims at environments using advanced graphics applications such as CAD, geographic info systems or trading

floor systems. In these environments, the workstation is used for higher computing and graphics requirements than the typical office PC.

Note: When installing blade systems, pay attention to additional cooling requirements (hot spots) that are generated when placing multiple blades in the same area.

3.10 Integration of energy and systems management

In Section 3.5, “Power management: The software side” on page 26, we introduced Active Energy Manager as tool of choice for all IBM platforms. However, in an overall systems management environment, power management is only one aspect of many. This section discusses how to integrate AEM into the IBM Tivoli® systems management platform.

IBM Tivoli Monitoring is the tool to monitor the status of your IT environment. It allows you to monitor your operating systems, databases, and servers throughout distributed environments through a flexible, customizable portal. A monitoring agent sits in place to tap the desired data from the monitored system.

From the Open Process Automation Library (OPAL), the IBM Tivoli Monitoring Power Management™ Agent available. This agent constitutes the interface between AEM and IBM Tivoli Monitoring. With this agent in place, IBM Tivoli Monitoring can monitor the power, temperature, and CPU usage of the respective systems. Figure 3-10 shows how the components interact.

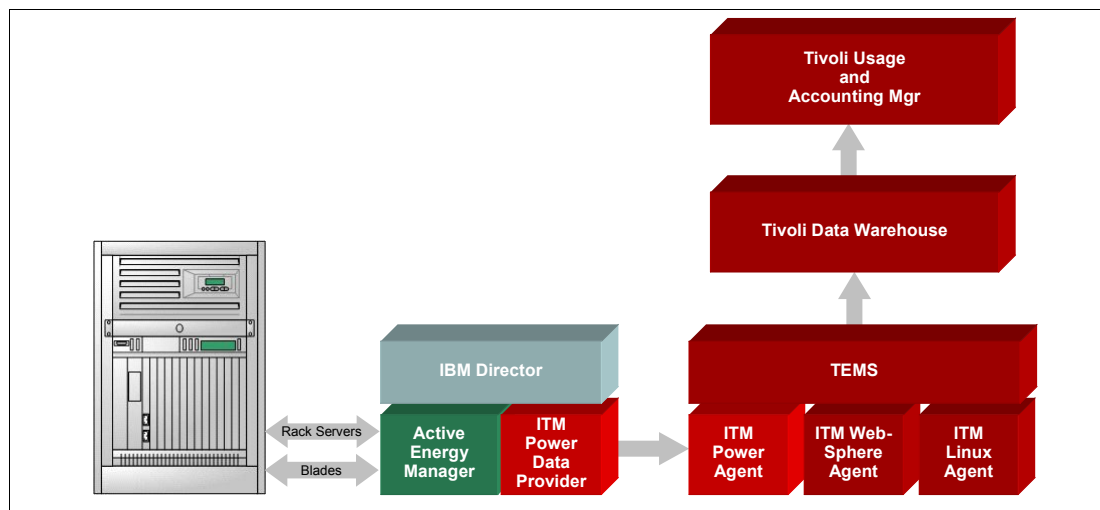


Figure 3-10 Integration of AEM with Tivoli systems management tools

Having this entry point into the Tivoli environment enables you to employ all the well-known features of IBM Tivoli Monitoring and other tools with which it interacts. You can also add the performance aspect discussed in section 3.2, “Working with energy efficiency” on page 21. Optimizing for power and performance might include the following scenarios:

- ▶ Reprovisioning a server based on the machine’s environmental temperature or overall rack power consumption to another rack in a cooler area of your data center. On a temperature alert in ITM, you would trigger the reprovisioning in IBM Tivoli Provisioning Manager.
- ▶ Power capping a single server having a temperature problem, perhaps because of an obstructed airflow, until the problem is solved on site.

- ▶ Feeding power, temperature, and CPU usage data into the IBM Tivoli Monitoring Warehouse. Using IBM Tivoli Usage and Accounting Manager this data can be correlated with accounting data. Charge the IT users according to their CPU and correlated power usage.

The opportunities are many, after the AEM data is available to the Tivoli environment. As energy management begins to play an important role, additional integration products from Tivoli are evolving. Due to the flexible nature of the Tivoli toolset, user setup might be complex. IBM services can help you find the best-fit solution.

3.11 Where to start

First measure energy, cooling, and inlet temperatures where possible. Trending analysis can help identify efficiency needs and opportunities. Find out what tools your systems provide. The objective is to make your energy consumption measurable. This is the most important step for all future decisions.

Aggressively drive consolidation of servers and storage using virtualization. When you are prepared for a virtualized infrastructure, consolidation comes easy. Ask the following questions:

- ▶ Which servers should be replaced in the near future? Which of the new, energy efficient systems would be able to take on this workload because they are underutilized?
- ▶ Which systems, if any, are complementary loaded? Consider the following points:
 - Low CPU and high memory usage can nicely coexist with high CPU and low memory usage systems in a partitioned environment.
 - Look for non-overlapping workloads such as day versus night jobs.
 - Is any system completely filled with non-business-critical work (low performance or response time needs)? If yes, virtualize it. Pair it with another virtual system on which you run occasional peak loads, such as weekly or monthly balancing. By setting appropriate entitlements, the peak load can displace the non-critical load, while the overall system usage stays close to 100%.
- ▶ Are there individual *infrastructure servers* such as LDAP, DNS, and licensing? These servers are often built on a single system that then has a low utilization. If so, virtualize and consolidate to get rid of these inefficiencies. Virtualization in this case also helps in reducing the cost of redundancy. Move the virtualized image to another machine instead of providing a hot standby.
- ▶ Use virtualized storage for new and virtualized servers. Although this can require start-up costs, it pays back quickly. Instead of having terabytes of unused disk space being turned around in individual machines, consolidate and *right-size* your storage. Moving a virtual operating system from one machine to another is as simple as connecting the target machine with the virtualized disks.

Consider opportunities for client consolidation. The cost of managing a distributed environment of fully equipped user PCs is nearly always more than you think, so the potential for big savings exists.



Site and facilities

This chapter describes the various components of the site and facility infrastructure, the types of infrastructures that are available, and why they are an integral part of your data center.

To be green, the data center utilizes a high-efficiency infrastructure and best practice initiatives. The life span of the infrastructure is three to five times longer than most IT hardware, and requires a much larger investment. Therefore, it is important to understand the timing of updating your infrastructure. This chapter also offers tips for utilizing best practice initiatives to increase efficiencies of your infrastructure.

The site and facility infrastructure can be divided into following five sections, each of which is dependent on the other:

- ▶ Data center cooling
- ▶ Heating, ventilation, and air conditioning (HVAC)
- ▶ Uninterruptible power supply (UPS)
- ▶ Power
- ▶ Standby generator or alternative power sources

IBM Business Partners, such as APC Schneider, Emerson Liebert, Eaton, GE, and Anixter, have a variety of green solutions that we discuss in this chapter and the next one.

4.1 Tips to start moving your data center towards green

Moving your data center towards green is all about saving energy. What infrastructure can you change? What best practices initiatives should you employ to become more energy efficient and reduce your energy consumption?

4.1.1 Reducing power consumption with innovative technologies

Applying innovative technologies within the data center can yield more computing power per kilowatt. IT equipment continues to become more energy efficient. With technology evolution and innovation outpacing the life expectancy of data center equipment, many companies are finding that replacing older IT equipment with newer models can significantly reduce overall power and cooling requirements and free up valuable floor space. For example, IBM studies have demonstrated that blade servers reduce power and cooling requirements by 25-40% over 1U technologies. While replacing equipment before it is fully depreciated might seem unwise, the advantages that new models can offer (lower energy consumption and two to three times more computing power than older models) combined with potential space, power, and cooling recoveries, are usually enough to offset any lost asset value.

IBM Business Partners such as APC¹ and Eaton² offer innovative green solutions in the power and UPS area. Companies such as GE³ design data center lighting solutions to save energy.

4.1.2 Reducing cooling requirements

A number of factors should be considered when you develop a plan for improving power and cooling efficiency by reducing the heat generated in the data center. Improvements in rack and room layout can increase energy efficiency with a relatively low initial investment. Consider the following improvement opportunities:

- ▶ Organizing IT equipment into a hot aisle and cold aisle configuration.
- ▶ Positioning the equipment so you can control the airflow between the hot and cold aisles and prevent hot air from recirculating back to the IT equipment cooling intakes.
- ▶ Taking advantage of supplemental cooling options, such as water or refrigerant heat exchangers.
- ▶ Improving rack cooling efficiency by employing a rear door heat exchanger or an enclosed racking system to dissipate heat from high-density computer systems before it enters the room. Similarly, relatively simple airflow management improvements can boost energy efficiency. For example, you can:
 - Take advantage of the current capacity by clearing under-floor blockages and implementing effective cable management.
 - Ensure that floor openings match the equipment thermal load by adding or removing vented tiles at the equipment air intakes.
 - Consider adding ducted returns. Ultimately, companies should consider organizing their data centers into thermal zones, assigning a defined set of IT equipment and floor space to specific HVAC or CRAC units. This type of space and thermal planning will eliminate hotter areas of the room (hot spots) that challenge cooling systems and will enhance system reliability by helping to avoid heat-related hardware failures. Also avoid cool spots.

¹ <http://www.apc.com>

² <http://www.eaton.com/EatonCom/ProductsServices/index.htm>

³ http://www.ge.com/products_services/lighting.html

Companies such as APC and Emerson Network Power⁴ have worked as business partners with IBM to devise cooling solutions for energy efficient data centers.

4.1.3 Improving physical infrastructure

Energy efficiency for infrastructure equipment has significantly improved in recent years. Replacing chiller or uninterruptible power supply (UPS) systems that have been in service for 15 years or more can result in substantial savings. New best-in-class UPS systems can experience as much as 70% less loss than existing UPS equipment. New chiller systems can improve efficiency by up to 50%. New chiller plants also can be installed with variable-speed drives, reducing pumping system energy usage and allowing better integration of the liquid cooling system into the chilled water infrastructure. Water-side economizers, which utilize outside air to directly cool the chilled water, can further reduce the energy required to cool the data center.

The capacity and efficiency of chilled water systems can be augmented with thermal storage systems that store energy generated at night, when chillers typically operate more efficiently, and then release this energy during the day, when energy costs are higher.

Air delivery to the data center also can be made more efficient, either through central HVAC systems or through CRAC units with variable speed drives. Central HVAC tends to be more efficient, because the systems are larger and more amenable to taking advantage of no-cost cooling when outside air temperatures are sufficiently low to provide some or all of the cooling requirements. CRAC units, alternatively, provide greater flexibility in managing the data center.

Even without upgrading facilities equipment, companies can save energy and gain cooling capacity by relaxing stringent relative humidity and temperature requirements for their data centers. Because these specifications are usually driven by the presence of hot spots, removing those hot spots (and cold spots) will allow temperature and relative humidity requirements to be relaxed, helping to reduce the energy required to operate the data center. In addition to cutting back on power usage inside its data center, a company can also reduce its carbon footprint⁵ by taking advantage of options for more eco-friendly sources of power. Integrating renewable energy into the power supply—including solar, wind, hydro and bio-mass generated energy—is a good way to reduce dependency on fossil fuels. Companies with the flexibility to relocate or open new data centers are even choosing locations that are rich in renewable energy sources as part of their corporate environmental strategy.

IBM Business Partners such as Anixter⁶ can supply cabling solutions for a new or reorganized data center.

4.2 Optimize your data center cooling

By addressing some of the most common cooling issues first, you will regain cooling capacity with a relatively low up-front investment. There are a number of opportunities for improving cooling efficiency within the data center. Improvements to airflow management, rack and

⁴ <http://www.liebert.com/>

⁵ The amount of carbon dioxide emitted. See <http://www.carbonfootprint.com/>

⁶ <http://www.anixter.com/>

room layout, and localized cooling all can increase energy efficiency with relatively low up front investment.

Note: Thanks to Dr. Roger Schmidt. Parts of this information came from his presentation, CIO's guide to *The green data center*. This is located at:

http://www-07.ibm.com/systems/includes/content/optimiseit/pdf/CIO_Guide_to_Green_Data_Center.pdf

4.2.1 Manage airflow

Where possible, avoid mixing hot and cold air. To increase airflow efficiency, you must have a clear path for the cool air to travel under the raised floor to get to the loaded areas. Above the raised floor, allow a path for the hot air to return back to the CRAC units. The following tips can help improve the airflow management:

- ▶ Add or remove vented tiles.
 - Remove vented tiles from hot aisle and open areas.
 - Add tiles to heavy heat load areas.
 - Adjust down the vented tiles with dampers to deliver smaller amounts of air in low heat areas, and open them fully for high heat areas.
 - Replace any tiles that have unused cutouts with full tiles.
- ▶ Improve airflow through racks.
 - Where possible, prevent the mixing of cold and hot air. Install blanking plates in racks that have empty space.
 - Provide large openings in racks for the cool air to bypass the heat load generated by the servers, and for cool air to cycle back to the CRAC unit. Air always takes the path of least resistance. This leaves hot air in the rack to be drawn through the server again.
 - Populate unused module bays in the rear of BladeCenter chassis with appropriate blanking plates/fillers.
- ▶ Seal cable cutouts and penetrations.

Cutouts in raised floors affect the air distribution and reduce the static air pressure under the floor. Block the openings by using brushes, foam, or pillows. This will allow more air to get to where it is needed.

Figure 4-1 on page 41 shows an open cable cutout. In this situation, air is being released, thus reducing under-floor static pressure and efficiency.

Anixter has a Web page of case studies involving cabling solutions.⁷

⁷ <http://www.anixter.com/AXECOM/US.NSF/ProductsTechnology/SolutionsCaseStudiesOverview>




Figure 4-1 Open tile cutout

Figure 4-2 shows pillows used to seal the cable cutout. This solution dramatically reduces cool air escaping into unnecessary areas.



Figure 4-2 Sealed tile with pillows

► Clear under-floor obstructions.

Excessive under-floor obstructions can lead to an increase in static pressure. High static pressure can have a reverse impact on the airflow under and above the raised floor.

Remove under-floor obstructions such as:

- Unused cables and wiring
- Unused under-floor equipment or communication boxes

Figure 4-3 on page 42 shows excessive cables obstructing airflow.




Figure 4-3 Under-floor blockages

4.2.2 Structured cable management

Using overhead and vertical cable management trays helps reduce the number of cables under the raised floor. As new cables are installed, they should be placed in overhead trays and obsolete cables should be removed. Figure 4-4 shows a vertical cable management frame being used to manage cables above the raised floor rather than being left under the floor, restricting airflow.



Figure 4-4 vertical cable management frame

4.2.3 Recommendations for raised floor height

Current raised floor modeling suggests a height to support at least 600 mm (24 in) of unobstructed space, enabling a clear path for cool air to travel. Some new raised floors are 900 mm (36 in) high, which allows additional quantities of air for extreme cooling requirements. See Figure 4-5.

For low raised floors, such as 150-300 mm (6-12 in), do not place equipment close to CRAC units, because low airflow or reverse flow can occur from the perforated tiles.

Under-floor partitions can be placed under the raised floor to direct air into the desired areas.




Figure 4-5 Under-floor view with chilled water pipes

4.2.4 Data center insulation

Consider the following alternatives for data center insulation:

- ▶ Insulate the data center walls and ceiling.
- ▶ Seal penetrations at the perimeter of the data center.
- ▶ Replace windows with double glazed windows.
- ▶ Install weather strips on doors.

These simple steps can assist in maintaining the temperature and humidity set-points of the data center and increase thermal efficiencies.

4.2.5 Hot aisle and cold aisle configuration

The hot aisle and cold aisle configuration enables much better airflow management on the raised floor, for both hot and cold air. This type of configuration offers a number of opportunities to route the airflow for both hot and cold using independent air corridors that minimize air mixing and improve efficiency.

Note: It is not always practical to move existing equipment; however, great efficiencies can be gained if doing so is possible.

Figure 4-6 shows a hot and cold aisle configuration, as recommended by the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE).

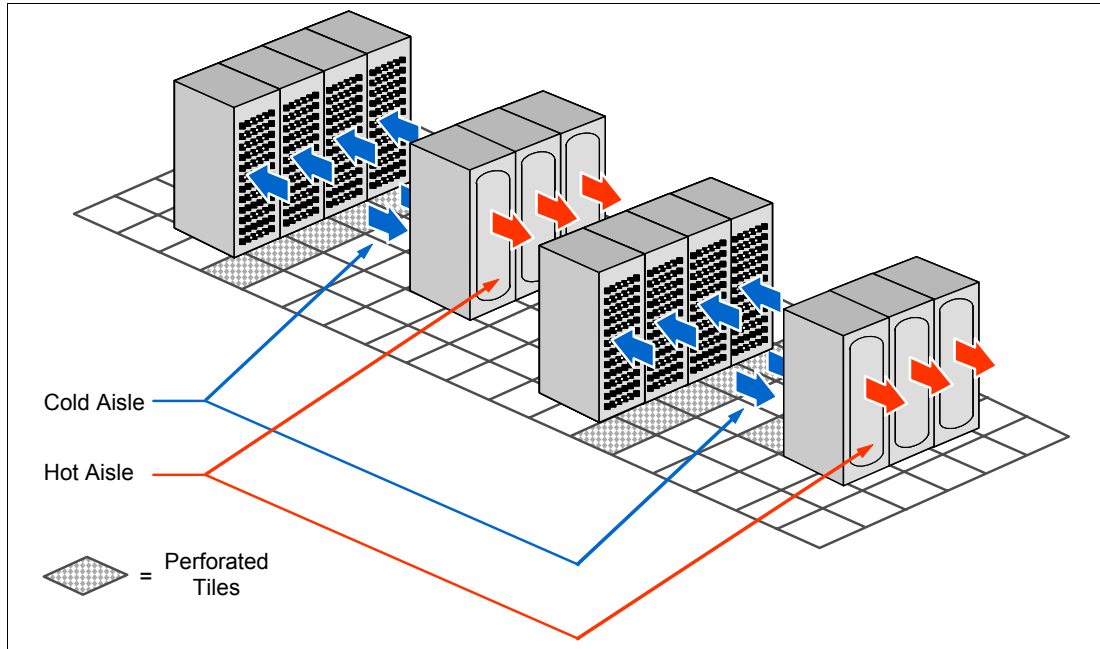


Figure 4-6 Hot and cold aisle configuration

Figure 4-7 shows the thermal flow benefits achieved by hot and cold aisles. The hotter the air returning back to the CRAC unit, the greater heat transfer that is achieved, increasing the efficiency of the CRAC unit.




Figure 4-7 Thermal flow of hot and cold air

If flexibility exists, locate CRAC units facing the hot aisle rather than cold aisles. The under-floor velocity pressure will be maximized in the cold aisles for this layout versus locating the CRAC at the end of the cold aisle.

Strategically placed ceiling partitions can help prevent recirculation of the hot exhaust air into the inlets of the servers.

Important: Be aware of the local fire codes before placing these partitions.




Figure 4-8 Cold aisle with perforated tiles

Figure 4-9 on page 46 shows simple, effective options to create corridors for hot and cold air.

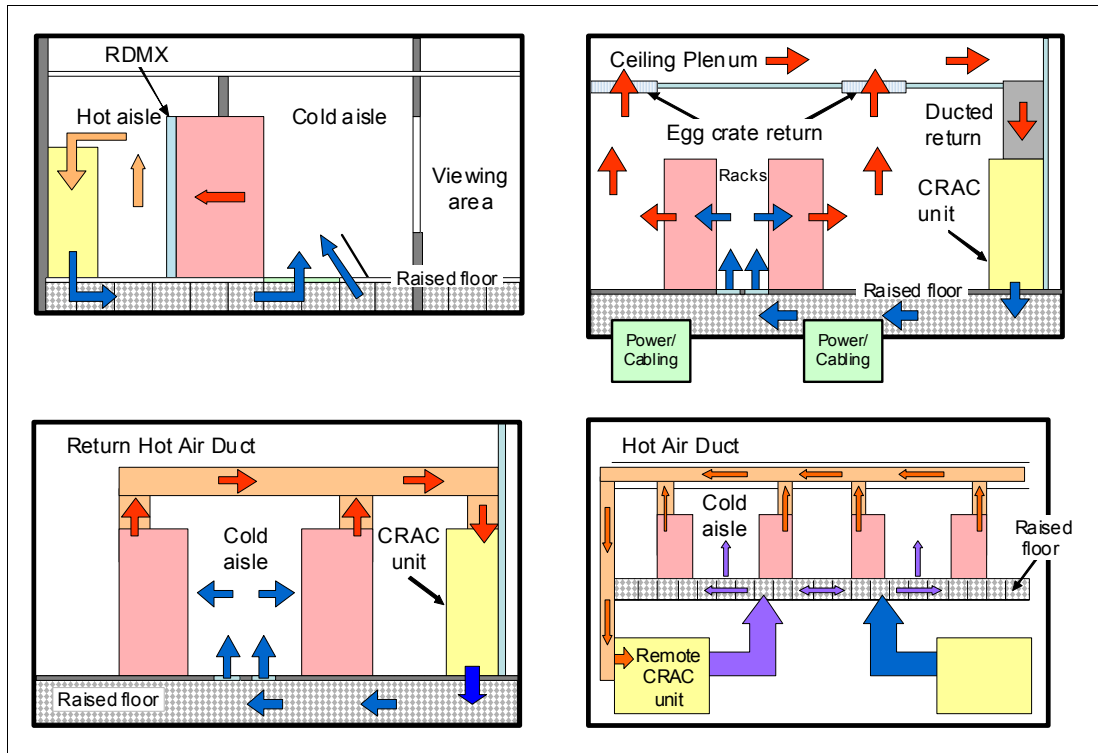


Figure 4-9 Data center ventilation schemes

4.3 Localized cooling equipment options

Traditional IT facility cooling designs utilize multiple CRAC units that are located around the perimeter of the data center. The cold supply air from the CRAC units should travel the distance from the unit to the server racks and then return back to the CRAC units.

To resolve hot spot issues raised by high density racks in data centers, cooling equipment manufacturers have begun to offer alternative cooling solutions to remove the heat load in the data center. The objective is to provide localized liquid cooling by locating the heat exchangers closer to the problem. By locating heat exchangers directly at the heat source, lessens the need for the CRAC unit. This increases the efficiency of the remaining CRAC units and the capacity of the available cooling within the data center. These heat exchangers are all scalable.

The Liebert XD system⁸ from Emerson Network Power offers a number of components for localized cooling. APC's InfraStruXure solution⁹ provides integrated rack-optimized packages.

To take advantage of using localized cooling, there is a requirement for chilled water in the data center. Chilled water in data centers is not new. It was used in the 1980s for mainframes. The need has returned as air cooling is unable to keep up with the thermal demand produced by high density servers. Several viable options based on this strategy have been developed for racks:

- ▶ Front or rear mounted fin and tube heat exchangers
- ▶ Internal fin and tube heat exchangers either at the bottom or mounted on the side of a rack

⁸ http://www.liebert.com/product_pages/SecondaryCategory.aspx?id=33&hz=60

⁹ http://www.apc.com/go/promo/isxglobal_registration/index.cfm?tsk=h423x

- ▶ Overhead fin and tube heat exchangers
- ▶ Internal server cooling

4.3.1 IBM Cool Blue rear door heat exchanger

The IBM Cool Blue™ rear door heat exchanger (RDHX) is a water-cooled device that serves as the rear door of an IBM enterprise rack. The RDHX is designed to directly cool 15 kW of heat from high density rack. The hot air is exhausted from the rear of the server. The hot air passes through the RDHX and is cooled. The heat is then transferred back to the chilled water loop, bypassing the CRAC units. Correct water temperatures are necessary for the successful operation of this device. Figure 4-10 shows how the RDHX works.

RDHX case study

In a case study of rack level testing with high power density servers, IBM BladeCenter servers were used in the rack. Functional 1U servers were also used to determine heat extraction. The tests were conducted in a room where the inlet temperature could be maintained between 23-25 degrees Celsius. The water flow rate was maintained at 31 LPM in the test cases and the water inlet temperature was kept at 17 degrees Celsius.

The results for the BladeCenter and 1U rack servers showed a 50-60% heat extraction of the input heat load. The increase in airflow impedance to the servers resulted in a negligible 1 degree Celsius increase in the CPU temperatures, due to a slight airflow loss to the servers. The exit air temperature from the rear door heat exchanger was reduced by an average of 25 degrees Celsius from the discharge end of the servers.

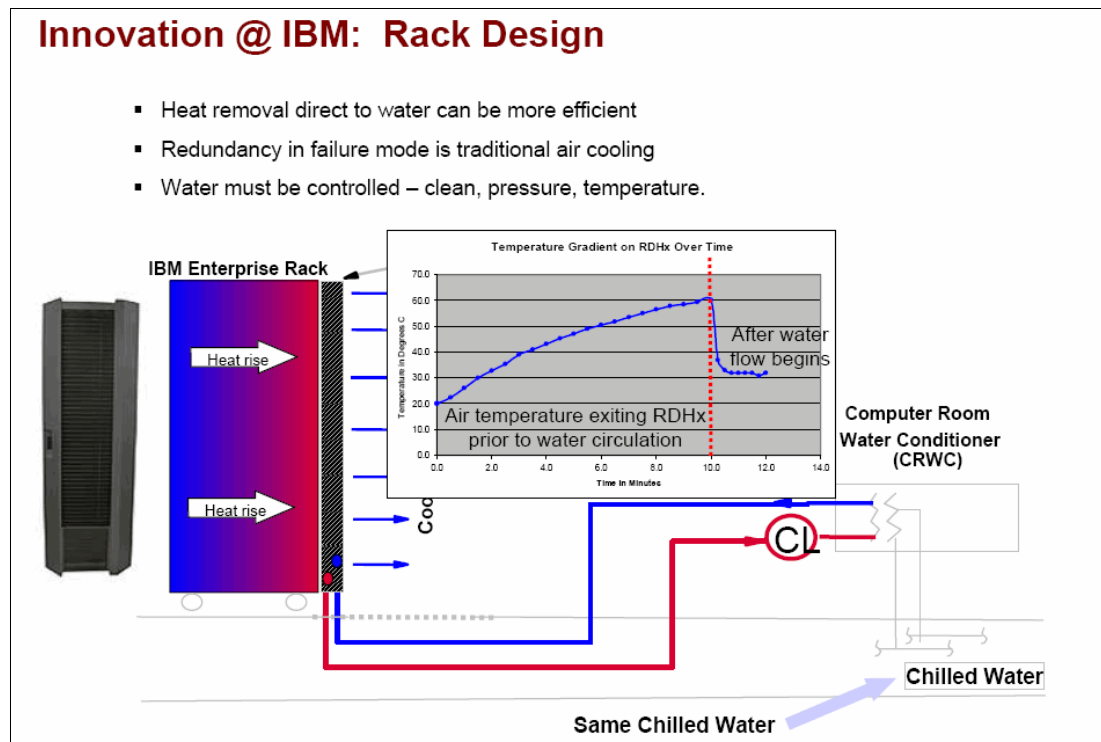


Figure 4-10 IBM Cool Blue RDHX

A graph of these results is illustrated in Figure 4-11 on page 48.

Comparison – with and without Cool Blue RDHX

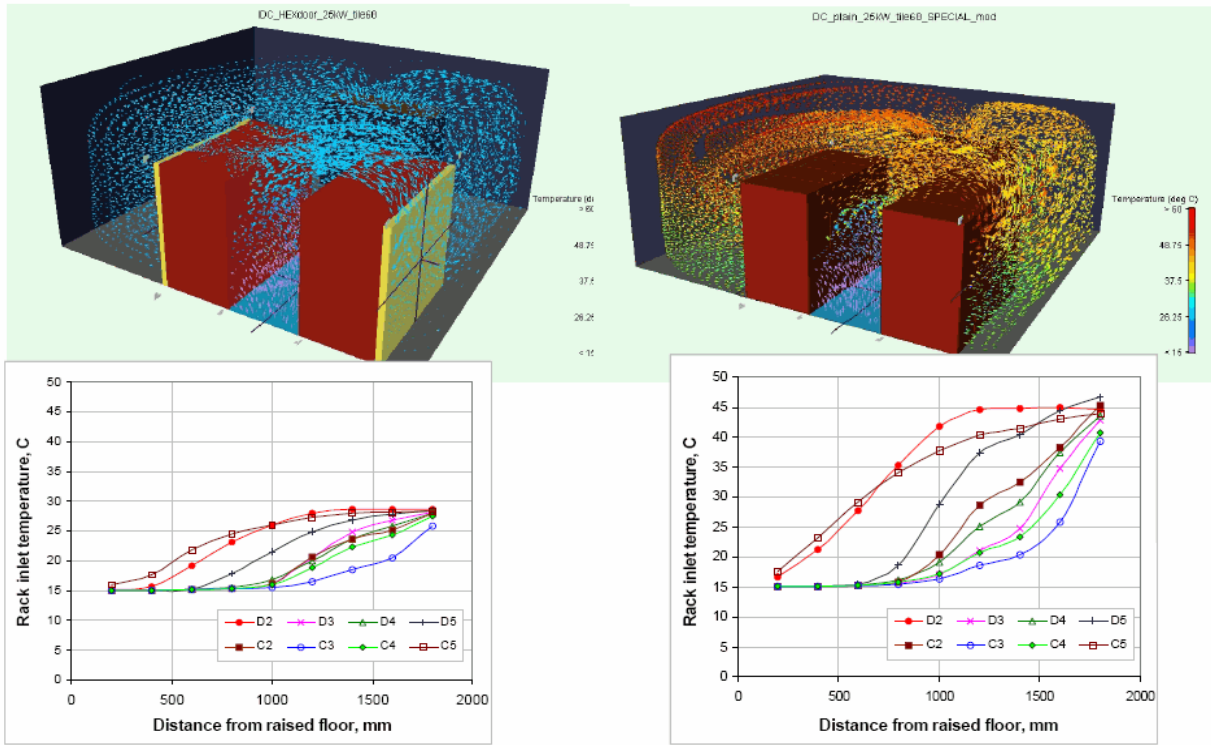


Figure 4-11 IBM cool blue rear door heat exchanger comparison

4.3.2 Modular water unit

Figure 4-12 on page 49 shows the modular water unit (MWU) located in the POWER 575 rack. The MWU supplies water internally to the server. The MWU acts as a heat exchanger transferring heat collected from the servers via a closed loop to the facility water loop. This increases efficiency by transferring heat back to the chilled water loop directly and bypassing the CRAC units.

Note: The size is significantly smaller than the technology that was available 15 years ago.




Figure 4-12 POWER 575 with MWU

4.3.3 Enclosed rack cooling

Enclosed rack cooling or self contained rack cooling provides increased capacity for high density racks. These environments can be located on both slab and raised floor. The self contained environment, using APC¹⁰ technology, allows the hot air from the rear of the servers to be ducted through the in-row cooler and then out to the front of the rack to start the cycle again. Increased efficiencies are gained by directly ducting the hot air over the condenser coil within the enclosed rack. The hot air cannot mix with the cold air, thus allowing maximum heat transfer to take place. The amount of heat rejection is dependent on the size of enclosure. Correct sizing is important for this to work. Figure 4-13 on page 50 shows a scalable modular data center.

¹⁰ <http://www.apc-mge.com/>



Figure 4-13 IBM GTS scalable modular data center

4.3.4 Sidecar heat exchanger

A sidecar heat exchanger is a water cooled heat exchanger located on the side of a rack to remove 100% of the heat generated by a rack of servers (see Figure 4-14).




Figure 4-14 Sidecar heat exchanger

4.4 Heating ventilation and air conditioning (HVAC)

There is no doubt that HVAC is the highest consumer of energy for the whole data center. Yet energy efficiency for chillers has significantly improved in recent years. Replacing chillers and pumps that have been in service for 10-15 years or more can result in substantial savings. New chillers can reduce energy consumption by up to 50%. Each type of chiller operates differently depending on environmental conditions. It is important to understand what is the correct type of chiller for your environment. A number of components make up HVAC:

Chiller	Removes heat from a data center. Sometimes used to describe the device that chills water that is used as the transfer medium for heat removal.
Pump	Circulates water through the data center via computer room air conditioning (CRAC) units, in-row cooling units, and rear door heat exchangers.
Air handling unit (AHU)	Cools or heats outside air depending on the ambient temperature and the required internal air temperature.

4.4.1 Types of chillers

Air cooled and water cooled chillers each can use various types of compressors such as reciprocating, centrifugal, and screw driven.

Air cooled chillers reject the heat with a condenser coil. Air is blown over the condenser coils with a fan so heat is transferred into the atmosphere. These units use more power than water cooled systems because the heat transfer works more effectively on a wet surface than a dry surface.

To further increase the efficiency of an air-cooled chiller, a pre-cooler evaporator can be placed on the air intake side. This pre-cools the air up to 10 degrees Celsius (17 degrees Fahrenheit) before the air passes over the condenser coil, which allows a greater transfer of heat.

Water-cooled chillers are generally used for very large capacity applications. They reject heat into a separate water source typically a cooling tower that uses the evaporation process to reject the heat into the atmosphere. However, if you are close to a river or lake, this water source can be used as a cheap and energy efficient option to reject this heat via immersion cooling. The cooling tower requires more maintenance, because condenser water tends to build up mineral deposits and requires chemical dosage to eliminate the chance of Legionnaires' disease. A water-cooled chiller requires much less power than an air-cooled chiller. The heated water by-product can also be used for heating buildings or swimming pools. The efficiency differences between the three compressors are:

- ▶ *Reciprocating* compressors are least efficient. They are usually set to run at three stages, which are 33%, 66%, and 100%. When the load requirement is 25%, the additional 8% is manufactured and then discarded.
- ▶ *Centrifugal* compressors are more efficient because they have fewer moving parts.
- ▶ *Screw-driven* compressors are the most efficient. They are able to stage to the exact load that is required from as little as 10% to 100%.

Note: Chillers that continually operate above 75% lose efficiency. So, if you are able to reduce the heat load within the data center, this will benefit the chiller load. For every watt you save in the data center, you will save 1.25 watts at the chiller.

A new technology chiller system can improve efficiency by up to 50%. New chiller plants also can be installed with variable-speed drives, reducing pumping system energy usage and allowing better integration of the liquid cooling system into the chilled water infrastructure. Water-side economizers, which use outside air to directly cool the chilled water, can further reduce the energy required to run the chillers.

4.4.2 Variable-speed drive pumps

Variable-speed drive (VSD) pumps self-adjust the output pressure depending on the demand. A standard pump operates at 100% load regardless of the demand, making it very inefficient. To increase your pump's efficiency, VSD control units can be added to your existing pumps, increasing their efficiency and reducing the electrical load.

4.4.3 Air handling unit

The air handling unit (AHU) cools or heats outside air, depending on the outside ambient temperature and the required internal air temperature. Energy efficient AHUs can reuse a percentage of conditioned air, hot or cold, for energy savings and increased efficiency.

4.4.4 Economizers to enable free cooling

Two types of economizers are air-side and water-side.

Air-side economizers

Air-side economizers can be used as a free cooling unit. However, depending on your location, these units work best where there is a constant supply of cool clean air. The consistency can be maintained by running these units overnight. The outside air economizer directly draws outside air for use by the data center.

Attention: Be aware of the following items with inside economizers:

- ▶ Temperature and humidity control is pivotal
- ▶ Gaseous contamination that can enter your data center and that can destroy equipment
- ▶ Particle contamination can enter your data center

Water-side economizers

Water-side economizers use cool outdoor air to generate chilled condenser water that is used to partially or fully meet the cooling demands of the facility. When the outside air is cool enough, the water-side economizer will take over part or full load of the chiller. This can result in a number of free cooling hours per day.

4.5 Cool less, save more

Increasing the data center temperature set points by just one degree can save energy costs dramatically. It reduces the load on CRACs and enables more free cooling.

Active energy manager (AEM) can monitor, manage, and provide trend analysis for power and heat within the various IBM series of systems. AEM can provide substantial cost savings by controlling energy consumption and improving energy use.

The American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) has a Technical Committee 9.9 (TC 9.9). The members of this committee are currently evaluating moving the upper and lower temperature set points for data centers beyond the previous recommendations. Moving these set points, if feasible, would have multiple benefits:

- ▶ Allow better utilization of outside air.
- ▶ Reduce the direct load placed on the chillers by increasing chilled water temperature.

IBM has increased the temperature in its own data centers from 20 degrees Celsius to 22 degrees Celsius (68 degrees Fahrenheit to 72 degrees Fahrenheit).

4.6 Uninterruptible power supply

An uninterruptible power supply (UPS) is a device that continues to supply electricity during a utility power failure. UPS units are often used to support critical computing environments.

Energy efficiency for UPS systems has significantly improved in recent years. Replacing a UPS system that has been in service for 15 years or more can result in substantial savings. New best-in-class UPS systems are approximately 70% more efficient than older ones.

4.6.1 Components and types

Various types of UPS units are available, and all have a varying design and performance characteristics. The major components of a UPS unit are as follows:

- Inverter:** The inverter converts direct current (DC) to alternating current (AC). If the power fails, the inverter provides 115 or 240 volts AC from a 12 volt DC power store.
- Rectifier:** The rectifier converts AC to DC. It provides DC power to ensure batteries remain charged. In larger UPS units, when the rectifier is used in conjunction with an inverter, acts as a premium power filter. This is achieved by the rectifier converting AC to DC, then the inverter converting DC back to AC. During this process, power spikes and dips are removed and clean power is delivered to the IT equipment.
- Power Store:** Although batteries are the most common power store, fuel cells and fly wheels are also widely used. The DC electricity is stored to be used by the inverter in case of a power failure.

Each UPS type caters to different applications. Understanding your equipment is important because one size does not fit all. The following types are the most common UPS units:

- ▶ **Standby UPS**
This type is commonly used for personal computers or single servers. The inverter starts only when the power fails. This is very efficient and cheap.
- ▶ **Line Interactive UPS**
This type is commonly used for 5-15 servers. When the power fails, power flows from the battery through the UPS to the servers. The inverter is always connected to the output. This provides filtering compared with the standby UPS.
- ▶ **Double Conversion UPS**
This type is used for larger data centers. It recreates the AC power by using a rectifier to change from AC to DC and then using an inverter to change back from DC to AC. This provides the best quality power and protection and is usually the most expensive option.

Although UPS units are typically rated in kilovolt-amperes (kVA), for example 330 kVA, the usable power is measured in watts. The efficiency of the UPS unit depends on the power factor (PF) of the unit, as shown in the following examples:

- ▶ A 330 kVA UPS unit operating at 0.8 = 264 kW of usable energy
- ▶ A 330 kVA UPS unit operating at 0.9 = 297 kW of usable energy

As you can see, the unit with the higher PF rating provides more usable energy.

4.6.2 Flywheel technology

As new flywheel technology emerges, the trend is to replace old UPS batteries with a new flywheel. UPS batteries have a limited life span. The quality of the battery and the number of times it is discharged can limit the life to 10 years or even less in some cases. The flywheel can operate in a greater temperature range than batteries, have a much smaller footprint than batteries, and does not have to convert AC to DC and back to AC.

The advantages of the flywheel are its efficiency and size. The disadvantages are its limited hold-up time and it does not condition power. Have at least one flywheel in your UPS bank.

Figure 4-15 is a cutaway drawing of a flywheel.




Figure 4-15 Flywheel example

4.7 Power

This section discusses terms to understand when planning for more efficient power.

4.7.1 Utility power supply

Utility companies sell power at a rate of X cents per kWh. How much of this do you get to use? Inefficiencies from the utility pole to your equipment can be as great as 20%. For example, 1 kW of power into your facility can equate to 0.8 kW usable power for your equipment.

4.7.2 Power factor correction

Power factor correction (PFC) on your site's electricity supply can regain some of that lost power. With PFC, 1 kW of power-in equates to 0.95 kW usable power for your equipment. For sites that use 2500 kW to 3000 kW, the pay-back period is three to four years.

4.7.3 Power distribution and resistance

Undersized wiring and lengthy cable runs from the power distribution point to the rack can lead to energy loss due to resistance. Ensure that the correct size of wiring is installed and limit the distance. This can improve the efficiency of the power delivered to the server.

In-row UPS units and power distribution boards save on both wiring and power loss due to resistance.

4.7.4 Intelligent power distribution unit

The intelligent power distribution unit (iPDU) interfaces with Active Energy Manager (AEM). The AEM collects power information to present a complete view of the energy consumption of the servers. AEM can provide substantial cost savings by controlling energy consumption and capping to the connected servers. This improves the energy efficiency.

4.7.5 DC versus AC

Power supplies are now offered in a choice of direct current (DC) or alternating current (AC). As AC power supplies are developed and the technology moves to switch-mode power supplies, the efficiency gains in this area have been astounding. The result has been a marginal difference between DC and AC, where DC is up by only 5% to 7%. AC parts are widely available and cost considerably less than that of DC. Consider the following questions:

- ▶ Should we change our infrastructure for this improvement?
- ▶ Will AC computer power supplies continue to improve?

4.7.6 On-site generation

For on-site power generation, see Section 4.8.2, "On-site electrical generation" on page 56.

4.8 Generators

The last infrastructure component that we discuss is generators.

4.8.1 Standby generator

Today's standby generators are designed to have greater fuel efficiency than older style generators and to be much greener in releasing CO₂ emissions into the atmosphere. They also have a much faster start and transfer time—less than 30 seconds—which supports the use of flywheel and fuel cell technologies to replace UPS batteries.

The standby generator is a key piece of the site infrastructure for high availability data centers. UPS systems can energize the data center for a few minutes or even hours. However, without a standby generator, the HVAC systems that supply the cooling to the data

center will not operate and the likelihood of overheating the data center is a certainty. High-power availability is a key in achieving a high-availability data center.

A standby generator supplies power when the utility power supply is not available. Both the standby generator and the utility power supply are connected to an automatic transfer switch (ATS). When the utility power supply is not available, the generator automatically starts, taking approximately 40 seconds to assume the load. During this time, the UPS supports the data center load. Without the UPS, the data facility loses power and IT systems power down.

Standby generators typically run on diesel fuel or natural gas. These units have a life span of 15 to 20 years. In areas that have a very good utility power supply, these units might not get very much run-time, so it is important to maintain and test them.

When sizing these units, be sure to account for all the infrastructure that maintains the data center, including chillers, pumps, CRACs, UPS, AHUs, and other site infrastructure.

4.8.2 On-site electrical generation

On-site power generation can augment or replace utility power as follows:

- ▶ Sites that have limited utility power can use on-site generation during peak periods.
- ▶ Sites that have access to cheap fuels, where the cost of electrical generation is cheaper than the cost of utility power, can use on-site generation.
- ▶ Utility power might be unreliable so on-site power generation is useful.

Co-generation or combined heat and power (CHP) is an option. Refer to Figure 4-16. Conventional generators produce electricity, and the heat by-product can be harnessed to assist with domestic and industrial heating. Consider these factors when building a new site.

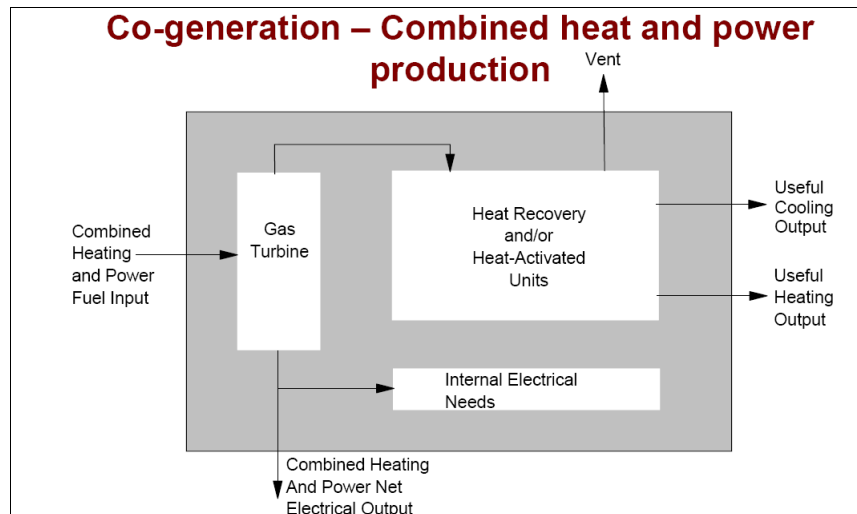


Figure 4-16 Co-generation schematic

New technologies evolving for on-site generation

The following new technologies are evolving to provide for on-site generation:

- ▶ Fuel cells, which could eliminate the need for UPS because reliability is 99.99% or better
- ▶ Nuclear, which could eliminate need for UPS because reliability is 99.99% or better
- ▶ Wind, which is not as developed as other technologies
- ▶ Solar, which is not as developed as other technologies

4.9 Recommendations for existing data centers

You should understand how much energy your data center consumes, and how efficient it is, beginning at the sever and tracing back to the utility company. As you begin to understand where the inefficiencies are, you can then start resolving them. By employing best practice initiatives in airflow management and data center cooling as described in this paper you can have a positive impact on the site and facility infrastructure. Begin with easy steps that have a relatively low up-front investment.

4.9.1 Investigate, consolidate, and replace

Consolidating and replacing old IT equipment can have a direct impact on the site and facility infrastructure. For every 1 kW of energy saved in the data center, an additional 1.35 kW is saved on the site and facility infrastructure, which results in a 2.35 kW saving from the utility provider.

As parts of your infrastructure progressively reach the end of serviceable life, investigate all the options available and their impact on other components.

Adding more energy-improving components to your existing infrastructure can increase the energy efficiency and capacity of your infrastructure.

The following recommendations are for the data center:

- ▶ Assess the data center environment.
- ▶ Implement no-cost and low-cost initiatives.
- ▶ Replace the oldest systems first; they are more likely to fail and are less energy efficient.
- ▶ Invest in systems that save the most energy.

The following recommendations are for the site and facilities:

- ▶ Assess the site and facility infrastructure.
- ▶ Replace the oldest systems first; they are more likely to fail and are less energy efficient.
- ▶ Invest in infrastructure that saves the most energy.

IBM services can assist in all this areas. See Chapter 5, “Who can help: IBM services and Business Partners” on page 59 for more details.

4.9.2 Investigate industry initiatives

We recommend that you also investigate the following industry activities and guidelines:

- ▶ Green Grid, which is a consortium of IT and data center equipment companies to help:
 - Establish metrics for data centers.
 - Develop standards and practices to improve energy efficiency in the data center.

Green Grid is located on the Web at the following location:

<http://www.greengrid.com/>

- ▶ American Society of Heating, Refrigerating, and Air Conditioning Engineers (ASHRAE) offers the following ways to help:
 - Has various subcommittees that provide information.
 - Develops standards for the design of efficient data center cooling.

ASHRAE is located on the Web at the following location:

<http://www.ashrae.com/>

- ▶ Uptime Institute, offers information about, but is not limited to the following areas:
 - Has data center availability and reliability information.
 - Provides operating criteria for adding energy efficiency to a data center.

Uptime Institute is located on the Web at the following location:

<http://www.uptimeinstitute.org/>

4.10 Summary

You might believe you have outgrown your data center, but you probably have not. The inefficiencies in older design principles and technologies are the main things holding you back from improving the data center. As you move towards implementing best practice initiatives and upgrading your IT equipment, you will generate additional capacity for space, power, and cooling from your existing infrastructure.

Advances in technology of IT systems are delivering more computing power for your dollar, but they are also stressing your power and cooling infrastructures and, in some cases, back to your local utility grid.

With virtualization offering increased performance per watt of power and advanced thermal diagnostics delivering pinpoint control for your cooling infrastructure, you can regain control of your energy and cooling requirements.

The introduction of best practice initiatives to improve data center cooling can start overnight, with minimal up-front costs.

The introduction of new equipment does not happen overnight. Server migration adds a short-term increase to the data center power and cooling demands. So, implement the new and remove the old, and the benefits will start to show.

As the power and cooling demands reduce in the data center, the cascading affect on the site and facility infrastructure will immediately take place.

As an example mentioned previously, If you remove 1 kW of power load from the data center, you will reduce an additional 1.35 kW of load from the site and facility infrastructure. Total savings is 2.35 kW.

Planning for site and facility infrastructure upgrades takes time, sometimes between 12 to 24 months to implement. So when designing a solution, do not simply look at what is on the market today, but ask the vendors for their roadmaps, and then plan for the best, newest, technology available.

In summary, review the following suggestions:

- ▶ Assess your environment, including the data center and facility.
- ▶ Improve cooling efficiencies by implementing best practice initiatives.
- ▶ Use less space through consolidation. Consider consolidating multiple data centers.
- ▶ Decrease the amount of servers to reduce the heat load.
- ▶ Optimize your data center in conjunction with the site and facility infrastructure.
- ▶ Work towards energy efficiency for CO₂ reductions.



Who can help: IBM services and Business Partners

This chapter discusses services and assistance that IBM and IBM Business Partners can deliver.

5.1 Services IBM can deliver

IBM professionals can help at the following levels by providing information for hardware and software solutions, and recommendations for designing and building highly efficient data centers:

- ▶ System and technology
- ▶ Management systems
- ▶ Data center

IBM offers over 500 services of all types. We highlight several of the relevant services in this section. They are also described, along with other services, on the IBM Web site, located at:

<http://www.ibm.com/services/us/index.wss/allservices/>

To find them, enter relevant keywords in the search field of the Web site.

5.1.1 IBM Global Technology Services

IBM Global Technology Services can help with the following categories in your move toward having a green data center:

- ▶ Diagnose
- ▶ Build
- ▶ Cool
- ▶ Virtualize and simplify
- ▶ Manage, measure, and enhance

Diagnose

IBM can help assess energy efficiency of your data center by using power management analysis and thermal analytics. The options of server and storage virtualization and consolidation are also studied. A fact-based business case is provided and major opportunities for improvement can be detected with reduction of energy costs up to 40%. Actions at the thermal level can help eliminate hot spots (regions of high power density) and undesired intermixing of hot and cold air. Payback on the investment can be achieved in as little as two years, thereby covering the cost of the assessment in the first year.

For example, a major U.S. utility provider needed to support IT growth with its existing 400 square meters (4400 square foot) data center as well as demonstrate to its customer base that the company is a leader in energy efficiency. It asked IBM to conduct a comprehensive, fact-based analysis of its IT infrastructure. The IBM analysis evaluated cooling system components, electrical systems, and other building systems. It was then possible to provide a baseline metric for data center energy efficiency and deliver a roadmap of cost-justified recommendations. IBM established that if the company spent U.S. \$18 000 once, it could save U.S. \$23 000 annually by reducing its energy consumption by 46%, providing a payback in less than two years. It could go further and realize U.S. \$100 000 in annual energy savings with upgrades to its UPS systems.

Table 5-1 shows examples of the Diagnose category offerings.

Table 5-1 Examples of Diagnose category offerings from IBM

Data center energy efficiency assessment	Provides a comprehensive, fact-based analysis that prioritizes tactical plans across the data center to help improve efficiency and reduce costs. Provides from 15% to 40% annual energy savings and includes an industry standard comparison (MPG for your data center), identifies low performance areas and actions for improvement, and a business case to prioritize investment.
Server consolidation efficiency study	As an attractively priced streamlined assessment, provides you with the facts base required to justify initiatives for optimizing your server environments. Identifies solutions that can help to reduce energy consumption in transitioning to a greener data center, and quantifies savings to reduce TCO and energy costs by up to 72%.
Storage optimization and integration services	Helps you reduce complexity, optimize performance, and manage growth by creating cost-effective, highly utilized, scalable and resilient storage infrastructures. Helps maximize energy efficiency through consolidation, virtualization, and storage optimization leading to a greener infrastructure.
Data mobility services	As a suite of technologies and services, enables the movement of stored data—typically included in energy assessment recommendations—in a nondestructive manner regardless of server platform or storage array vendor.

Build

IBM can provide expertise to customers, based on the experience of building 3 million square meters of data centers for clients worldwide. Planning, building, or upgrading a new data center provides the perfect opportunity to rationalize the data center strategy as a way for the customer to gain substantial savings on capital, operations, and energy efficiency.

For example, Bryant University¹ in Rhode Island needed to reduce costs and grow the capacity of its IT infrastructure to meet rising student enrollments and student expectations for IT services. The University's previous decentralized IT infrastructure was costly, inefficient, and increasingly unable to scale up to meet these growing demands. The University worked with IBM to consolidate and upgrade its IT operations. IBM helped Bryant design and build a centralized, 500 square-foot data center. The project consolidated 75 servers in three data centers down to three IBM BladeCenter platforms holding a total of 40 IBM System x and IBM System p servers. The new data center was implemented in half the total space required by the previous three data centers. This smaller footprint, coupled with energy efficient components, significantly reduced the university's energy costs—contributing to a 40% reduction in overhead costs.²

Table 5-2 shows examples of Build category offerings.

Table 5-2 Examples of Build category offerings from IBM

IT facilities assessment, design and construction services	Helps you create stable, security-rich, energy efficient, future-ready data centers and enterprise command-center facilities. The end-to-end services include a review of your existing data center's reliability, points of failure, growth, floor space, and power and cooling needs.
Server optimization and integration services	Helps you create a cost-effective, scalable, flexible and resilient server infrastructure to support business applications using industry-leading practices based on IBM experience and intellectual capital.

¹ See the Build section in *The green data center: cutting energy costs to gain competitive advantage*, April 2008, at: <http://www-935.ibm.com/services/us/cio/outsourcing/gdc-wp-gtw03020-usen-00-041508.pdf>

² See *The art of the possible: Rapidly deploying cost-effective, energy-efficient data centers*, February 2008, at: <http://www.ibm.com/services/us/its/pdf/smdc-eb-sfe03001-usen-00-022708.pdf>

Storage optimization and integration services	Helps you reduce complexity, optimize performance and manage growth by creating cost-effective, highly utilized, scalable, and resilient storage infrastructures. Helps maximize energy efficiency through consolidation, virtualization and storage optimization leading to a greener infrastructure.
Data mobility services	As a suite of technologies and services, enables the movement of stored data—typically included in energy assessment recommendations—in a non disruptive manner regardless of server platform or storage array vendor.
Intelligent and green building, especially in Asia Pacific region	Through specialized facilities services, helps enable the integration of all building subsystems so that they can operate in a safe, efficient, and ecologically friendly environment. The system manages the air-conditioning, fire and security systems of the building and can lead to significant energy savings.

Cool

One example of deploying cooling technology from IBM is at Georgia Institute of Technology's Center for the Study of Systems Biology³. The Center required supercomputing capabilities for protein structure simulations and other techniques supporting research of new drugs. Its supercomputer demanded the highest possible computational performance while generating significant heat output from its ultradense blade servers. By implementing a mix of advanced IBM cooling technologies, including an innovative rear door heat exchanger, the university was able to maintain computing performance while reducing air-conditioning requirements by 55%. The resulting energy savings helped cut operational costs by 10-15% and helped save an estimated U.S. \$780 000 in data center costs.⁴

Table 5-3 shows examples of Cool category offerings.

Table 5-3 Examples of Cool category offerings from IBM

Installation of Cool Blue Rear Door Heat eXchanger (RDHX)	Provides a simple, cost effective, energy efficient solution to solve hot spot problems within the data center. The overall approach is to provide and oversee a simple step-by-step process for implementation of the RDHX units.
Data center infrastructure energy efficiency optimization	Allows clients to work one-on-one with IBM Data Center Services Power / Thermal Development Engineers to formulate a balanced plan to improve efficiency, reduce total cost of ownership, and maximize the aggregate IT equipment supported by an existing data center infrastructure.
Server and storage power / cooling trends and data center best practices	Helps you understand the current and future power, and cooling and I/O demands that IT equipment places on your existing or planned data center infrastructure.

Note: IBM has four types of services to address facilities problems you might have:

- ▶ Data Center and Facilities Strategy Services
- ▶ IT Facilities Consolidation and Relocation Services
- ▶ IT Facilities Assessment, Design and Construction Services
- ▶ Specialized Facilities Services

Virtualize and simplify

Virtualizing servers and storage devices can increase processing loads and boost individual utilization rates.

³ See the Cool section in *The green data center: cutting energy costs to gain competitive advantage*, April 2008, at: <http://www-935.ibm.com/services/us/cio/outsourcing/gdc-wp-gtw03020-usen-00-041508.pdf>

⁴ "Georgia Tech implements a cool solution for green HPC with IBM," October 2007

An example of that comes from the University of Pittsburgh Medical Center (UPMC)⁵, which is seeking to become a truly integrated, self-regulating health care system, using evidence-based medicine to produce superb clinical outcomes and lower costs. To support this goal, UPMC has been undergoing an IT service transformation program with help from IBM. Health Industry Insights, an IDC Company, has been tracking the project and reports that, “UPMC is facing multiple challenges that translate into doing more with less.”

IDC noted such challenges as pressures “to improve customer service, patient safety, and service quality while reducing care delivery costs.”

“In order to deliver highly integrated, efficient care in the face of this rapid growth and industry pressures, UPMC recognized that enterprise-wide IT systems, data integration, and platform standardization were crucial for its quality and business integration goals and to achieve the economies of scale expected to accrue from these acquisitions,” IDC said.

UPMC worked with IBM in virtualizing its Wintel and UNIX systems and consolidating 1,000 physical servers down to 300 IBM servers. Storage also was reduced from 40 databases down to two centralized storage area network (SAN) arrays. IDC reports, “Our initial estimate that UPMC would avoid almost \$20 million in server costs has grown to \$30 million and is likely to exceed \$40 million by the conclusion of the project in 2008.”⁶

IBM services can help with assessments, methods, technology and knowledge.

Manage, measure, and enhance

Tracking energy usage helps the customer control his billing but also will aid in allocating energy where it is needed.

For example, Bryant University is working with IBM and APC to deploy software that automatically manages server clock speed to lower power consumption. In addition, software helps monitor and control fan speed, power level used at each outlet, cooling capacity, temperature and humidity, and it distributes power to server blades as needed.⁷

IBM services can help with assessments, methods, technology, and knowledge.

Note: The following links are to information about site and facilities, servers, and storage:

- ▶ <http://ibm.com/services/siteandfacilities>
- ▶ <http://ibm.com/services/server>
- ▶ <http://ibm.com/services/storage>

5.1.2 IBM Asset Recovery Solutions

The disposition of obsolete information technology has become a major issue for corporations. International Data Corporation estimates that most corporations store old computer equipment for up to three years at a cost of U.S. \$360 per machine, and then pay an additional U.S. \$218 for its eventual disposal. In the face of these statistics, it is clear that the customer faces real challenges every time he upgrades its IT equipment. The customer will want to maximize recovery costs at every opportunity. In addition, disposal methods must

⁵ See the Virtualize and simplify section in *The green data center: cutting energy costs to gain competitive advantage*, April 2008, at:

<http://www-935.ibm.com/services/us/cio/outsourcing/gdc-wp-gtw03020-usen-00-041508.pdf>

⁶ Health Industry Insights, an IDC Company, “Virtualization: Healthcare’s Cure for the Common Cost,” Part 2, Doc #HI209705, December 2007 at:

<http://www-03.ibm.com/industries/healthcare/doc/content/resource/insight/3721514105.html>

⁷ J. Nicholas Hoover, “Data Center Best Practices,” InformationWeek, March 3, 2008,

<http://www.informationweek.com/management/showArticle.jhtml?articleID=206900660&pgno=1&queryText=>

comply with applicable environmental regulations. Also, the customer will want one single-source provider instead of using a patchwork of different services.

IBM Global Financing (IGF) can meet these challenges and more. Asset Recovery Solutions, an offering from our Global Asset Recovery Services division, offers a suite of highly competitive solutions to dispose of IT equipment and equipment related to IT from IBM and other sources. The equipment can include hardware, fax machines, and printers. IBM methods provide for safe and proper processing of used equipment. IBM has worldwide product remarketing and logistics capabilities complemented by state-of-the-art systems, and brings these best-of-breed capabilities to the customer, with the reach, resources, and knowledge to serve its customers anywhere in the country, providing real peace of mind. Figure 5-1 provides an overview of Asset Recovery Solutions.

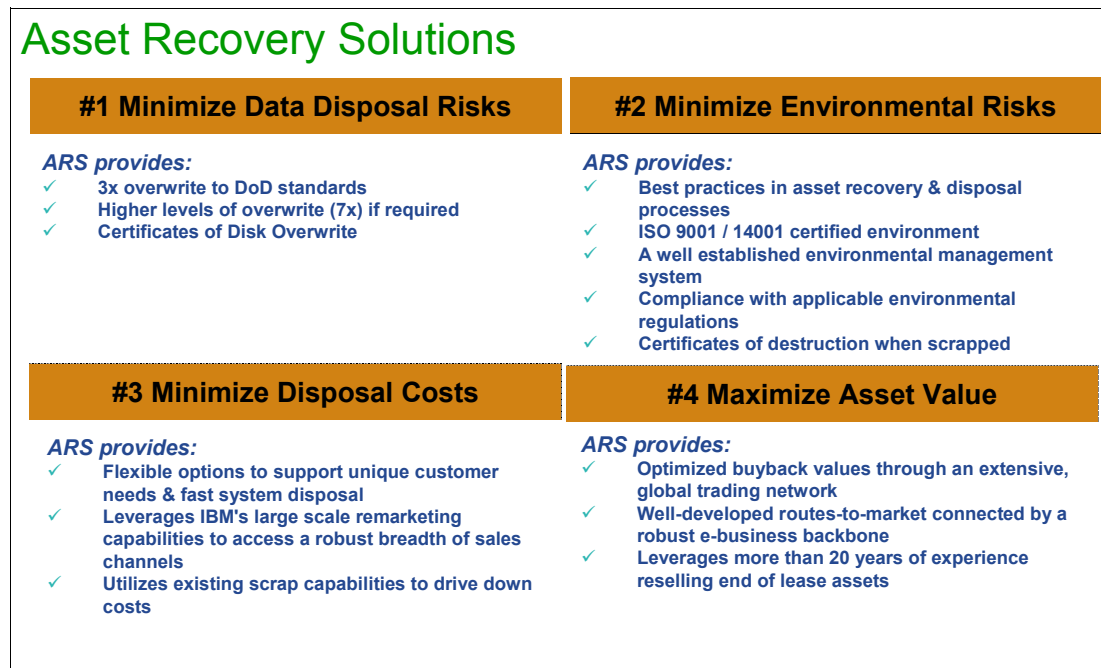


Figure 5-1 Asset Recovery Solutions

To give you an idea of the magnitude of what can be achieved, review the following accomplishments from IBM:

- ▶ Processes an estimated 40 000 machines per week across 22 centers around the world
- ▶ Reused or resold almost 87% of the assets returned to its Remanufacturing Centers
- ▶ Processed over 108 million pounds of end-of-life products and product waste in 2006 and sent only 78% of that to landfill
- ▶ During a three-year period, took in and reused more than 1.9 million machines; generated billions in revenue; processed over 147 000 metric tons of equipment, parts, and waste from materials and products, such as:
 - Steel: Over three times the amount used in the Eiffel Tower
 - Plastic: Over 22 railroad cars of condensed plastic
 - Paper: Enough bales to span the Golden Gate Bridge 23 times
 - Non-ferrous material: Equivalent to 150 18-wheeler trucks

For more information, see the following Web site:

<http://ibm.com/financing/us/recovery>

5.1.3 IBM financing solutions

In 2007, IBM Project Financing™ announced, with *Project Big Green*, a way to help gain the financial agility to seize new opportunities through business transformation and renewal of critical infrastructure of data centers. Solutions are offered from design to construction to outfitting and relocation. Refer to Appendix A.2, “Project Big Green” on page 70.

IBM Project Financing can structure financing for all of the components that make up a data center—from the real estate and construction to the Hardware and Software, as well as services. We provide industry knowledge to ensure payment structures in line with the useful life of various asset classes. IBM Global Financing (IGF) is uniquely positioned as part of Project Big Green to provide a *green wrapper* of financing solutions.

It offers the possibility to control cash flow and to let the customer decide when to draw funds and how much to draw. Easy one-time negotiations cover the entire project and let you finance IT and data center equipment and all related services for design, construction, and relocation. It is a one-stop shop for a single, integrated IBM solution. IBM can provide consistent terms globally for all leases and loans, enabling worldwide competitive pricing. You can be prepared for current and future geographical expansion. Financing is available for planned expenditures, and the flexibility to address unforeseen challenges as they arise. For more information, see:

<http://www.ibm.com/financing>

5.1.4 IBM green data centers

IBM is continuing to evolve its own data centers with greenness and sustainability in mind. You can find more information about green data centers at:

http://www.ibm.com/systems/optimizeit/cost_efficiency/energy_efficiency/

Alternatively, you may search the IBM Web site for *green data center* at:

<http://www.ibm.com/us/>

5.2 IBM Business Partners

IBM has a number of Business Partners that can assist in offering the right green solution for your data center. With strategic alliances and offering best innovative technologies with best practice solutions, these IBM Business Partners can help you achieve a state-of-the-art green data center.

The IBM Executive Briefing Center in Triangle Park, NC was designed and built with the following IBM Business Partners:

- ▶ APC provides power and cooling solutions. The company often partners with IBM Global Technology Services (GTS) to support scalable modular data centers. See their green solutions announcement.⁸
- ▶ Emerson Network Power has provided cooling products since the 1960s, when the IBM System/360™ mainframe was introduced. See the company’s Web pages on energy efficiency.⁹
- ▶ Eaton is known for power delivery products. See its products and services Web pages.¹⁰

⁸ <http://www.apc.com/solutions/display.cfm?id=39E81F1D-5056-AE36-FE2C8E82CEB07885>

⁹ http://www.liebert.com/topissue_pages/top_issues.aspx?x=efficiency

¹⁰ <http://www.eaton.com/EatonCom/ProductsServices/index.htm>

- ▶ GE supplied the lighting, using energy-efficient compact fluorescent bulbs. See the company's Web pages about "Ecomagination¹¹", containing products and practices that are both economical and ecological.

IBM has joined forces with Neuwing Energy Ventures to enable organizations to receive Energy Efficiency Certificates (EEC) for projects that reduce energy.¹²

ADC has an online brochure on green data center solutions.¹³

Other organizations and concepts to watch include:

- ▶ Leadership in Energy and Environmental Design (LEED) is a voluntary consensus-based national rating system for developing high-performance, sustainable buildings. Developed by U.S. Green Building Council (USGBC), LEED addresses all building types and emphasizes state-of-the-art strategies for sustainable site development, water savings, energy efficiency, materials and resources selection, and indoor environmental quality. LEED is a practical rating tool for green building design and construction that provides immediate and measurable results for building owners and occupants. The Web site is: <http://www.usgbc.org/>
- ▶ Energy Star¹⁴ is a joint program of the U.S. Environmental Protection Agency (EPA) and the U.S. Department of Energy (DOE). Its purpose is to encourage efficient use of energy, thereby protecting our environment and saving us money. With the help of Energy Star, Americans have saved considerable amounts of energy to reduce greenhouse gas emissions and to save significantly on their energy bills.

For business, Energy Star offers help with energy management, which can produce a considerable savings for a company's bottom line and for the environment. Partnering with Energy Star helps a company set goals, track savings, and reward improvements. Top performing buildings can earn the ENERGY STAR, which is a mark of excellence for energy efficiency that is recognized nationally.

¹¹ <http://www.ge.com/innovation/eco/index.html>

¹² For more information regarding EECs in data centers, contact David F. Anderson PE at dfa@us.ibm.com

¹³ <http://www.adc.com/Library/Literature/106057AE.pdf>

¹⁴ <http://www.energystar.gov/>



Conclusion: Green is a journey, act now

Today's CIOs face many concerns in their data centers to achieve business sustainability. The cost of the energy is growing, as is public awareness about data centers impacting our environment. Data center power, cooling, and space capacities are at or near the limit. With new environmental laws coming and company images at stake, going green makes sense, from cost and sustainability perspectives and also as a way to reduce organizational risk.

This paper has highlighted green concepts, technologies, products, and services that can assist sustainability as data centers evolve by simplifying, sharing, and dynamically managing and using IT.

Going green can enable organizational and competitive advantages. Many solutions and actions are available to assist data centers in their endeavors to become more green. By working simultaneously at the site, facilities, and IT equipment levels, you can achieve new green ways to run your data centers.

Figure 6-1 illustrates the vision of the future from IBM, the *New Enterprise Data Center*. This vision supports being green by design, with energy efficiency as a strength.

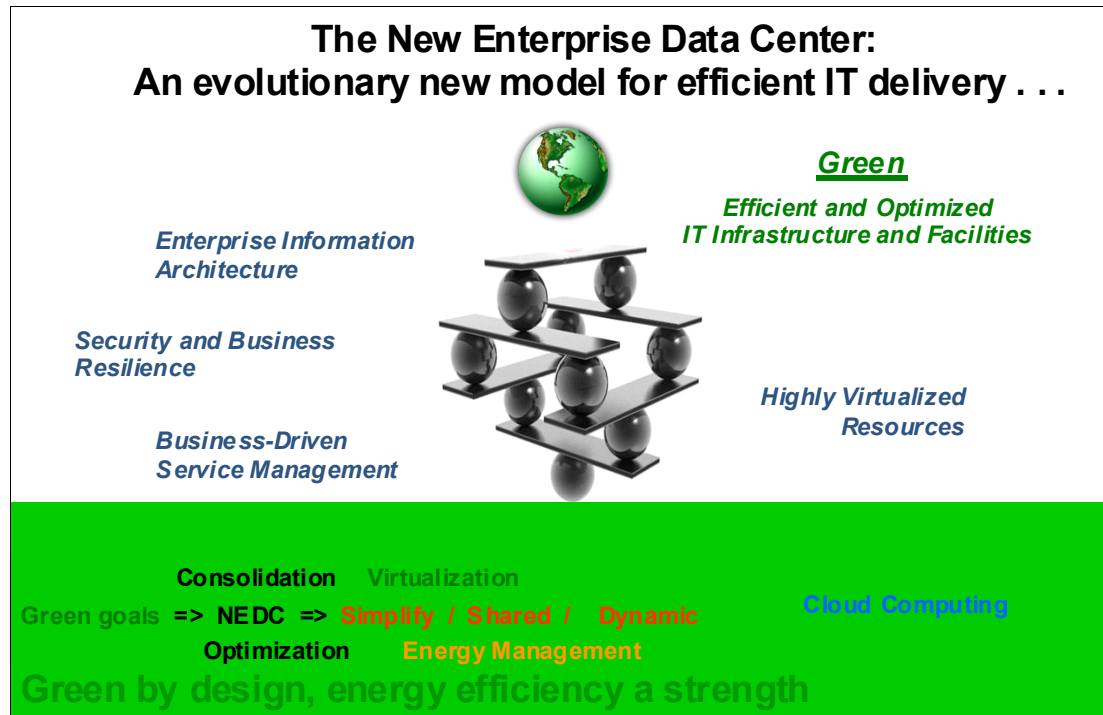


Figure 6-1 The New Enterprise Data Center

Going green is a collaborative experience. IBM welcomes ideas and suggestions on this topic. For information about how to contact IBM, see “Comments welcome” on page x.

Becoming green is part of the journey towards a better future. Each step in the process brings more benefits. Develop a green strategy. Enable IT equipment to become more energy efficient with green technologies such as virtualization. Optimize data center power and cooling. Manage for sustainability using new measuring and management tools. Use IBM Services and partners to boost your green capabilities.

Keep in mind that a green data center is not a finite project with a destination at the end. Rather, a green data center is an evolving program and should be considered a *journey* that provides dividends of lower costs, improved sustainability, and a better public image.



A

Commitment to green from IBM: The past, present, and future

For many years, IBM has been committed to energy efficiency and to the protection of the environment. Today, this involvement continues with new activities. This appendix describes these actions.

A.1 A history of leadership in helping the environment.

To establish an exhaustive list of accomplishments attributed to IBM for the protection of the environment would be difficult. Several significant accomplishments are:

- ▶ In 1998, U.S. EPA Climate Protection Award and U.S. EPA ENERGY STAR Computer Partner of the Year and Alliance to Save Energy's Star of Energy Efficiency Award
- ▶ In 1999, U.S. EPA ENERGY STAR Computer Partner of the Year, U.S. EPA Climate Wise Partner Achievement Award, U.S. EPA ENERGY STAR Excellence in Corporate Commitment Award, and U.S. EPA Top 20 Best Workplaces for Commuters
- ▶ In 2005 and 2006, U.S. EPA Top 20 Best Workplaces for Commuters (first 2-time winner)
- ▶ In 2007, U.S. EPA's SmartWay Excellence Award
- ▶ A founding member of *The Green Grid*

IBM has a policy of transparency, and its efforts are public information, including:

- ▶ Production of corporate environmental reports annually since 1990
- ▶ Disclosure of greenhouse gas emissions that are verified by several institutions

A.2 Project Big Green

On 10 May 2007, IBM committed to allocate \$1 billion each year toward the following achievements:

- ▶ Accelerate *green* technologies and services.
- ▶ Offer a roadmap for clients to address the IT energy crisis while taking advantage of IBM hardware, software, services, research, and financing teams.
- ▶ Create a *global green team* of approximately 1000 energy efficiency specialists from throughout IBM.
- ▶ Double the computing capacity by 2010 without increasing power consumption or carbon footprint saving 5 billion kilowatt hours per year.
- ▶ Reduce energy costs by half for a typical 25 000 square-foot data center that spends U.S. \$2.6 million in power annually. The resulting reduction of emissions can be compared with taking 1300 passenger cars off the road.

Press release for Project Big Green

NEW YORK, NY—10 May 2007: IBM announced that it is redirecting U.S. \$1 billion per year across its businesses, mobilizing the company's resources to increase dramatically the level of energy efficiency in IT. The plan includes new products and services for IBM and its clients to reduce sharply data center energy consumption, transforming the world's business and public technology infrastructures into *green* data centers.¹

¹ IBM Press Releases, <http://www.ibm.com/press/us/en/pressrelease/21524.wss>

Recognition of the Project Big Green by the market

The Project Big Green initiative aligns with past IBM environmental programs. IBM leadership has been recognized by the IT marketplace. In February 2008, IDG's Computerworld selected IBM as the top Green IT Company for 2008.² Computerworld is the leading source of technology news and information for IT influencers worldwide.

A.3 IBM internal efficiency results

Let us take a brief look at the IBM internal IT account that runs our manufacturing, research, and general business operations. Table 6-1 represents a decade of improvement in how IBM manages and operates its centers. IBM has several thousand servers, and over the last 10 years has implemented its own recommendations.

IBM moved from over 200 data centers to about a dozen. Fewer centers to manage allows IBM to focus on doing a few things well. IBM did server and storage consolidation and virtualization. It then consolidated the workload by merging and reducing applications. This application virtualization allows IBM to reduce the number of applications running significantly, by more than two thirds.

Table 6-1 IBM internal computing metrics in 1997 versus today

IBM Metrics	1997	Today
CIOs	128	1
Host data centers	155	7
Web hosting centers	80	5
Network	31	1
Applications	15,000	4,700

A.4 Future directions

IBM continues to work toward a greener working environment and provides regular announcements. For example, in November 2007, IBM launched the world's first corporate-led energy efficiency certificate program in an effort to help clients benchmark and improve efficiency of their IT operations and to reduce environmental impact. The certificates, based on energy use reduction verified by a third-party, provide a way for businesses to attain a certified measurement of their energy use reduction. The certificates can be either traded for cash on the growing energy efficiency certificate market or retained to demonstrate reductions in energy use and associated CO₂ emissions.³

² IBM Press Releases, <http://www.ibm.com/press/us/en/pressrelease/23563.wss>

³ IBM Press Releases, <http://www.ibm.com/press/us/en/pressrelease/22513.wss>

Related publications

We consider the publications listed in this section particularly suitable for a more detailed discussion of the topics covered in this paper.

IBM Redbooks publications

For information about ordering these publications, see “How to get IBM Redbooks publications” on page 76. Note that some of the documents referenced here might be available in softcopy only.

- ▶ *IBM Systems Virtualization: Servers, Storage, and Software*, REDP-4396
- ▶ *IBM System z10 Enterprise Class Capacity on Demand*, SG24-7504
- ▶ *PowerVM Virtualization on IBM System p Managing and Monitoring*, SG24-7590
- ▶ *PowerVM Virtualization on IBM System p Introduction and Configuration*, SG24-7940
- ▶ *Virtualizing an Infrastructure with System p and Linux*, SG24-7499
- ▶ *IBM System x3950 M2 and x3850 M2 Technical Introduction*, REDP-4362
- ▶ *Virtualization in a SAN*, REDP-3633
- ▶ *Going Green with IBM Systems Director Active Energy Manager*, REDP-4361, draft available at:

<http://www.ibm.com/redbooks>

Online resources

These Web sites are also relevant as further information sources:

- ▶ APC and APC MGE
<http://www.apc.com/>
<http://www.apc-mge.com/>
- ▶ American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE)
<http://www.ashrae.com/>
- ▶ Anixter and Anixter cabling solution case studies
<http://www.anixter.com>
<http://www.anixter.com/AXECOM/US.NSF/ProductsTechnology/SolutionsCaseStudiesOverview>
- ▶ Assessment questionnaire from IBM
http://www.ibm.com/systems/optimizeit/cost_efficiency/energy_efficiency/services.html
- ▶ Carbon Footprint to determine the amount of carbon dioxide emitted.
<http://www.carbonfootprint.com/>
- ▶ Centrinet launches UK's first operational zero carbon data center with help from IBM

- http://www.ibm.com/software/success/cssdb.nsf/CS/JGIL-7BRTQ3?OpenDocument&Site=default&cty=en_us/
- ▶ Creating the Green Data Center, Simple Measures to Reduce Energy Consumption
[http://www.anixter.com/.AXECOM/AXEDocLib.nsf/\(UnID\)/9D5FF4F709EB68E9862573D7007788D2/\\$file/Energy_Awareness.pdf](http://www.anixter.com/.AXECOM/AXEDocLib.nsf/(UnID)/9D5FF4F709EB68E9862573D7007788D2/$file/Energy_Awareness.pdf)
 - ▶ Data Center Best Practices. J. Nicholas Hoover, InformationWeek
<http://www.informationweek.com>
 - ▶ Eaton
<http://www.eaton.com/EatonCom/ProductsServices/index.htm>
 - ▶ Emerson Network Power
<http://www.liebert.com/>
 - ▶ Energy Star
<http://www.energystar.gov/>
 - ▶ Energy tips for small and medium businesses
http://www-304.ibm.com/jct03004c/businesscenter/smb/us/en/contenttemplate/!!/gc1_xmlid=80670/
 - ▶ Environmental Protection Agency
<http://www.epa.gov/climatechange/index.html>
 - ▶ GE lighting and Ecomagination
http://www.ge.com/products_services/lighting.html
<http://www.ge.com/innovation/eco/index.html>
 - ▶ Georgia Tech implements a cool solution for green HPC with IBM; case study
http://www-01.ibm.com/software/success/cssdb.nsf/cs/STRD-788DC6?OpenDocument&Site=gicss67educ&cty=en_us%20for%20green%20HPC%20with%20IBM
 - ▶ Green Grid
<http://www.greengrid.com/>
 - ▶ How going green can improve day-to-day operations
http://www.ibm.com/industries/government/doc/content/landing/3727822109.html?g_type=pspo
 - ▶ IBM Asset Recovery Solutions
<http://ibm.com/financing/us/recovery>
 - ▶ IBM Global Financing
<http://www.ibm.com/financing>
 - ▶ IBM green data centers
http://www.ibm.com/systems/optimizeit/cost_efficiency/energy_efficiency/
 - ▶ IT and the environment: A new item on the CIO agenda?
http://www-05.ibm.com/no/ibm/environment/pdf/grennit_oktober2007.pdf
 - ▶ Kalbe Collaborates With IBM to Build Green Data Center and Reduce Energy Consumption
<http://www.ibm.com/press/us/en/pressrelease/23467.wss>

- ▶ New Enterprise Data Center (NEDC)
<http://www.ibm.com/systems/optimizeit/datacenter>
- ▶ SearchDataCenter.com from Uptime Institute, Inc.
<http://searchdatacenter.techtarget.com/?int=off&offer=DCregtui208>
- ▶ Service offerings from IBM
<http://www.ibm.com/services/us/index.wss/allservices/>
- ▶ Site and facilities, servers, and storage
<http://www-935.ibm.com/services/us/index.wss>
- ▶ The IBM System x and BladeCenter Power Configurator can be downloaded from:
<http://www.ibm.com/systems/bladecenter/resources/powerconfig/>
- ▶ The art of the possible: Rapidly deploying cost-effective, energy-efficient data centers, February 2008
<http://www.ibm.com/services/us/its/pdf/smdc-eb-sfe03001-usen-00-022708.pdf>
- ▶ The green data center: cutting energy costs for a powerful competitive advantage.
<http://www-935.ibm.com/services/us/cio/outsourcing/gdc-wp-gtw03020-usen-00-041508.pdf>
- ▶ U.S. Green Building Council (USGBC)
<http://www.usgbc.org/>
- ▶ Uptime Institute
<http://www.uptimeinstitute.org/>
- ▶ VMware VMotion technology
<http://www.vmware.com/products/vi/vc/vmotion.html>
- ▶ Virtualization: Healthcare's Cure for the Common Cost? Part 2
<http://www-03.ibm.com/industries/healthcare/doc/content/resource/insight/3721514105.html>
- ▶ White paper: The green data center
<http://www.ibm.com/industries/education/doc/content/resource/thought/2794854110.html>

Other publications

These publications are also relevant as further information sources:

- ▶ Esty, Daniel, and Winston, Andrew S. *Green to Gold: How Smart Companies Use Environmental Strategy to Innovate, Create Value, and Build Competitive Advantage*. Yale University Press, 2006. ISBN-13: 9780300119978
- ▶ Hoover, J. Nicholas. *Data Center Best Practices*. InformationWeek, March 3, 2008
- ▶ Koomey, Jonathan. *Estimating total power consumption by servers in the U.S. and the world*. Oakland, CA: Analytics Press, 2007.

ASHRAE publications

The American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) has a committee named Technical Committee 9.9 (TC 9.9) that is concerned with mission-critical facilities, technology spaces, and electronic equipment. TC 9.9 has a number of books available to help with data energy efficiency and best practices.

ASHRAE Datacom Series of books by TC 9.9

Books that ASHRAE TC 9.9 has published include:

- ▶ Thermal Guidelines for Data Processing Environments (2004), ISBN/ISSN: 1-931862-43-5
- ▶ Datacom Equipment Power Trends and Cooling Applications (2005), ISBN/ISSN: 1-931862-65-6
- ▶ Design Considerations for Datacom Equipment Centers (2006), ISBN/ISSN: 1-931862-94-X
- ▶ Liquid Cooling Design Considerations for Datacom Equipment Centers (2006), ISBN/ISSN: ISBN-10: 1-933742-05-4; ISBN-13: 978-1-933742-05-2
- ▶ Structural and Vibration Guidelines for Datacom Equipment Centers (2008), ISBN/ISSN: 978-1-933742-20-5
- ▶ Best Practices for Datacom Facility Energy Efficiency (2007), ISBN/ISSN: 978-1-933742-27-4
- ▶ High Density Data Centers—Case Studies and Best Practices (2008), ISBN/ISSN: 6331-2-403-000000-43
- ▶ Contamination and Economizers in Datacom Facilities (not yet available)

How to get IBM Redbooks publications

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



The Green Data Center: Steps for the Journey

**Going green:
Improving energy
efficiency and costs**

Green IT products

**Green tips for site and
facilities**

The information technology (IT) industry is a big user of resources as well as an enabler of efficiencies that reduce CO₂ emissions. However, as companies continue to grow to meet the demands of their customers and as environmental concerns continue to be an issue, organizations are looking for ways to reduce corporate energy consumption and to become more environmentally responsible—to become *green*.

This IBM Redpaper publication can help your IT organization as it begins the journey to becoming a *green data center*. IBM wants to help others, particularly our clients, to chart a course to reap the benefits of lower costs and improved sustainability that running a green data center can provide. Understanding what is possible can speed your journey to an optimized green data center with sustainability designed into both the IT and facilities infrastructures. Although this paper is not all inclusive, it provides a quick start for going green in data centers. It also provides additional pointers and information. You can use this paper as a guide to becoming more energy efficient.

**INTERNATIONAL
TECHNICAL
SUPPORT
ORGANIZATION**

**BUILDING TECHNICAL
INFORMATION BASED ON
PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:
ibm.com/redbooks**