

Greening the Data Center

Opportunities for Improving Data Center Energy Efficiency

George Spafford



Greening the Data Center

Opportunities for Improving Data Center Energy Efficiency

Greening the Data Center

Opportunities for
Improving Data Center
Energy Efficiency

GEORGE SPAFFORD



IT Governance Publishing

Every possible effort has been made to ensure that the information contained in this book is accurate at the time of going to press, and the publishers and the author cannot accept responsibility for any errors or omissions, however caused. No responsibility for loss or damage occasioned to any person acting, or refraining from action, as a result of the material in this publication can be accepted by the publisher or the author.

Apart from any fair dealing for the purposes of research or private study, or criticism or review, as permitted under the Copyright, Designs and Patents Act 1988, this publication may only be reproduced, stored or transmitted, in any form, or by any means, with the prior permission in writing of the publisher or, in the case of reprographic reproduction, in accordance with the terms of licenses issued by the Copyright Licensing Agency. Enquiries concerning reproduction outside those terms should be sent to the publishers at the following address:

IT Governance Publishing
IT Governance Limited
Unit 3, Clive Court
Bartholomew's Walk
Cambridgeshire Business Park
Ely
Cambridgeshire
CB7 4EH
United Kingdom

www.itgovernance.co.uk

© George Spafford 2009

The author has asserted the rights of the author under the Copyright, Designs and Patents Act, 1988, to be identified as the author of this work.

First published in the United Kingdom in 2009 by IT Governance Publishing.

ISBN 978-1-84928-009-9

PREFACE

IT organizations are under intense pressure to manage the power consumed by data centers and the resulting cooling demands. To address these needs, IT must properly blend people, processes and technology to create solutions.

In my first Green IT pocket guide, we reviewed governance and processes for Green IT. In this guide we will provide a sample of technical improvement opportunities at a high level. Our intent is to foster discussions in your own organization around what should be done first, second, third and so on. The result of this should be a technical implementation roadmap that supports the objectives of Green IT.

Indeed, most data centers will not suffer from a lack of opportunities to improve. What is needed is careful planning and deliberate execution to manage power and cooling while continuing to create and protect value for the organization.

ABOUT THE AUTHOR

George Spafford is a Principal Consultant with Pepperweed Consulting, LLC and is an experienced practitioner in business and IT operations. He is a prolific author and speaker, has consulted and conducted training on regulatory compliance, IT Governance and process improvement in the US, Australia, New Zealand and China, and has co-authored *The Visible Ops Handbook* and *Visible Ops Security*.

George holds an MBA from the University of Notre Dame, a BA in Materials and Logistics Management from Michigan State University, and an honorary degree from Konan Daigaku in Japan. He is an ITIL[®] Service Manager, TOCICO Jonah and a Certified Information Systems Auditor (CISA). George is a current member of the IIA, ISACA, ITPI, ITSMF, and the TOCICO.

ACKNOWLEDGEMENTS

As always, I'd like to thank my family for all their support. If it wasn't for them, I would never have survived writing two books back to back.

CONTENTS

Introduction	10
Chapter 1: What is “Energy Efficiency”?	13
Energy versus power.....	13
Ratio of inputs to outputs.....	13
Energy consumption and productivity	14
Chapter 2: Processes and Planning	15
Develop situational awareness	15
Implement a Green IT process	15
IT asset management (ITAM).....	17
Capacity management.....	17
Project management.....	17
Identify available assistance	18
Employee involvement	18
Involve the right stakeholders	19
Chapter 3: Applications and Data.....	20
Reduce application variation.....	20
Utilize cloud computing.....	21
Software as a service (SaaS), platform as a service (PaaS) and storage as a service	22
Application design	23
Data management	23
Chapter 4: Broad Themes	25
Maximize utilization.....	25
Newer tends to be more efficient	25
Modular approach.....	25
Utilize zones	26
Looser environmental demands	27
Chapter 5: IT Hardware	28
Identify and decommission ghosts.....	28
Server consolidation and virtualization.....	29
Blade computers	30
Power management.....	31
Compression	32

Contents

Power supplies	32
Other server components	33
Reassess fault tolerance	34
Reduce hardware variation	34
“Turn it off!” campaign	35
Chapter 6: Facilities – Electrical	36
Scalable modular power.....	36
Power back-up systems.....	36
Power zones	37
Power distribution.....	38
Distributed generation	39
Chapter 7: Facilities – Cooling	41
The Arrhenius equation and data center cooling	41
Input temperature: ASHRAE standards, vendors and reality	42
Cooling zones	43
Hot aisle/cold aisle.....	43
Correcting air flow.....	44
Leverage experts.....	47
Water cooling	47
Insulation	48
Economizers	48
Chapter 8: Selecting a Data Center Location	50
Climate	50
Cost of power.....	50
Energy availability.....	51
Tax incentives.....	52
Bonds/government financing.....	52
Political climate	52
Natural disasters	52
Chapter 9: Monitoring and Reporting	54
Power usage effectiveness (PUE)	54
Data center infrastructure efficiency (DCiE)	55
Monthly data center energy costs	55
Monthly energy costs per IT service.....	55

Contents

Baseline now!	56
Chapter 10: Conclusion	57
Appendix: Additional Resources	58
ITG Resources	59

INTRODUCTION

Organizations around the world rely on data centers to house their IT services in a manner that optimizes confidentiality, integrity and availability in support of the entity's goals. The growth in the number of servers and data centers in the past five years has been nothing short of amazing.

An accelerating demand for power is coinciding with an increasing demand for computing and storage capacity. In 2006, data centers consumed 61 billion kilowatt hours (kWh) of electricity, with growth projected at 12% per year.¹ Unfortunately, only 30% of the power that enters a data center is actually consumed by IT hardware.² Of that, an even smaller fraction of that power is used to do anything productive.

Not only must the IT hardware be powered but it must also be cooled. One emerging phenomenon is the increasing density of high-performance computing in terms of computing capacity, power demanded and cooling needs in a given volume of space. In many respects, we can view cooling as a function of power consumption. As power consumption rises, so too do cooling demands and

¹ Paul Scheihing, "DOE Data Center Energy Efficiency Program" (U.S. Department of Energy, Energy Efficiency and Renewable Energy, May 2008), at www1.eere.energy.gov/industry/saveenergynow/pdfs/doe_data_centers_presentation.pdf.

² The Green Grid, "Guidelines for Energy-Efficient Datacenters," available, after registering, at <http://whitepapers.techrepublic.com.com/abstract.aspx?docid=353987>.

Introduction

the power consumption associated with that cooling.

While 3,000 watts per rack was considered a significant draw until approximately the early part of this decade, today's high-density blade racks are 10,000–30,000 watts with a potential 60,000–70,000 watts being discussed. Of course, that power demand will trigger a demand for approximately the same amount of power for cooling.

The rise in power and cooling demands is creating a number of risks:

- Operating expenses are increasing rapidly, with power being a sizable component.
- Utilities are informing data centers that they cannot provide any additional power.
- Data centers lack the power and/or cooling infrastructure necessary and cannot add servers, thus constraining their ability to support their respective organizations. Indeed, Gartner's Rakesh Kumar estimates that 70% of the Global 1,000 currently face growth challenges due to power constraints.³
- The projected growth in power demands would put data centers ahead of airlines in terms of green house gas (GHG) emissions.⁴ Given concerns over global climate change and today's political climate, it is very likely

³ Rakesh Kumar, "Green IT: Immediate Issues for Users to Focus On" (Gartner, 7 August 2008).

⁴ Steve Lohr, "Data Centers Are Becoming Big Polluters, Study Finds" in *The New York Times*, 1 May 2008, available at <http://bits.blogs.nytimes.com/2008/05/01/data-centers-are-becoming-big-polluters-study-finds>.

Introduction

that organizations will see increased regulation around GHG emissions at the overall entity level and potentially even at the data center level.

- Brand reputation can be affected by what customers, investors and employees learn about what the organization is doing to protect the environment.

Addressing these risks requires a systemic approach that blends people, processes and technology. In my first book⁵, the emphasis was on Green IT governance and processes. The current book will provide an overview of processes for planning, but will mainly focus on technical improvement opportunities at a high level.

Each chapter deals with a specific aspect of data center power consumption. Rather than take a linear “cookbook” approach, we have organized the chapters in such a manner that the reader can skip to the topic of interest.

⁵ George Spafford, *The Governance of Green IT* (ITGP, 2008). See www.itgovernance.co.uk/products/2106.

CHAPTER 1: WHAT IS “ENERGY EFFICIENCY”?

The term “energy efficiency” is currently being bandied around a lot in the trade press and in meetings. Before we progress, this term needs to be reviewed, as it can give an insight into the work that needs to be done.

Energy *versus* power

Firstly, when energy is discussed we are talking about power (watts) over a period of time. Thus, energy is typically reported in kilowatt hours, (kWh), bearing in mind the following equation:

$$\text{kWh} = (\text{Watts} \times \text{Duration of use in hours}) / 1000.$$

This is an important point – energy is a combination of power and time. All things being equal, if we decrease either, then we decrease the energy used. If we can cut both the wattage and the duration of use, then the results will be better still.

Ratio of inputs to outputs

Secondly, we need to recognize that efficiency is defined as the ratio of inputs to outputs. If the outputs increase faster than the inputs, then efficiency is improving. In terms of data center power, energy inputs are in kWh and outputs are some degree of operation of the intended IT hardware.

1: What is "Energy Efficiency"?

With data centers, only 30% of the power gets to the IT hardware.⁶ The balance is used by coolers, power systems, and so forth. If we can decrease the power used by supporting systems and increase the ratio going to IT hardware, we will have improved efficiency. The challenge is that this view of efficiency is not sufficient. Outputs only matter if they move the organization towards its goals (thus improving productivity) or serve to mitigate risks (thus protecting value). Anything else is wasteful.

Energy consumption and productivity

Efficiency alone does not solve the problem. We need the outputs to support IT's mission of creating and protecting value for the organization.

Improving efficiency is just the first step. The next step is to make sure that maximum useful work is being accomplished with the power that does make it to the IT services.

IT must be able to identify non-value add (NVA) or low-value add services and work with the business to make the proper decisions as to whether to decommission, consolidate with another service, etc.

⁶ The Green Grid, "Guidelines for Energy-Efficient Datacenters," available, after registering, at <http://whitepapers.techrepublic.com.com/abstract.aspx?docid=353987>.

CHAPTER 2: PROCESSES AND PLANNING

Action without a plan is a recipe for suboptimal results if not outright failure. IT needs to develop an approach that supports the goals of the organization while also providing a means of lowering risks and costs.

The following is an overview of processes that can assist with the planning and execution of data center energy efficiency improvement initiatives.⁷

Develop situational awareness

The first step is to develop an understanding of what the corporation is doing in terms of a green strategy and what senior management expects of IT in terms of supporting that strategy.

Implement a Green IT process

The best way to manage an ongoing effort is formally to define a process that enables the achievement of desired objectives while also mitigating risks.

The goal of the Green IT process is to support the green objectives of the organization as defined by the board and/or a corporate social responsibility group.

⁷ For more information on the Green IT process, see George Spafford, *The Governance of Green IT* (ITGP, 2008). See www.itgovernance.co.uk/products/2106.

2: Processes and Planning

The process is comprised of four sub-processes:

- 1 Green IT policy management – tasked with monitoring the organization's position on Green IT and developing IT policies accordingly. One input is ongoing analysis from the Green IT functional area objective management sub-process (step 4 below).
- 2 Green IT architecture and standards – takes the policies of the organization and develops architecture standards that will guide the procurement, development, implementation and operation of IT services. This sub-process also returns feasibility and status information to the Green IT policy management sub-process (step 4 below) in regard to proposed new or changed policies.
- 3 Green IT functional area objective management – helps explicitly maintain ongoing alignment between the objectives of IT and the organization's green strategy.
- 4 Green IT assessments – provides the means to validate progress through the use of ongoing assessments. This is a necessary part of the plan-do-check-act (PDCA) process improvement cycle.

Furthermore, the process can be used to define roles and responsibilities, metrics, data exchanges and so forth.⁸

⁸ For more information on the Green IT process, see George Spafford, *The Governance of Green IT* (ITGP, 2008). See www.itgovernance.co.uk/products/2106.

2: Processes and Planning

IT asset management (ITAM)

To improve the value derived from assets while also managing the total cost of ownership (TCO), organizations need to implement a formal ITAM program.

ITAM is best thought of as a management discipline that spans multiple functional areas in both IT and the business. For green data center initiatives, ITAM can both help improve the utilization of IT assets and enforce standards around the procurement of IT assets and services.

Capacity management

IT's ability formally to plan for future capacity is critical. The traditional approach of ballpark estimates and over-specification of capacity has contributed to the creation of data centers that are under-utilized and woefully inefficient.

Project management

It is critical that the individual improvement efforts be treated as projects. Informal efforts that are not managed properly are likely to have a high failure rate, which can have a number of negative internal and external ramifications.

IT needs to use project and portfolio management (PPM) to maintain proper command and control over the various efforts. The project methodology used should include phase gates for review and approval of milestones as well as proper requirements definition, solution building or purchasing, testing, deployment and overall management reporting.

2: Processes and Planning

If PPM is not in place, then a basic process should be implemented to support the data center power and cooling initiatives. This should be evolved over time rather than inordinately delaying the improvements.

Identify available assistance

When planning your approach, it is useful to understand what forms of external assistance are available. Assistance can be of a financial nature, such as tax credits at various levels, state and municipal bonds to spur property development, and incentive programs from utilities. For example, PG&E recently paid a customer a record \$1.4 million to fund its improvement efforts.⁹

Indeed, utilities may offer a number of incentive programs, including offers of technical assistance. The truth is that they will achieve a better return on investment by helping customers become more efficient than they will on the massive capital investment required to build another power generation facility.

Employee involvement

Employee involvement is more of a need than a process; employees must be involved in the effort. This is not only because they will be the people who actually implement the solutions, but also because they may have been thinking about

⁹ Sara Stroud, "Data center earns record rebate" in *Sustainable Industries*, 11 December 2008, available at www.sustainableindustries.com/breakingnews/35933274.html.

2: Processes and Planning

various energy consumption and cooling problems and their solutions for a long time.

There must be a method of collecting their suggestions and rewarding their ideas. Rewards can vary dramatically and include monetary bonuses, recognition in front of peers and management, the chance to work on ideas as fully funded approved projects, and so forth.

Involve the right stakeholders

To effectively and efficiently address the issue of power consumption in the data center, both the IT and the facilities management teams must be involved. Traditionally these two groups have been siloed and not inclined to work together.

These two groups must work together. It is critical that the facilities team understands capacity requirements and expected growth in order to provide the appropriate power and cooling systems. At the same time, the IT team needs to understand what is feasible given the current infrastructure and how to maximize value. Otherwise, it is apt to pursue changes that the power and cooling infrastructure cannot support.

Doubtlessly, if the barriers can be broken down, both groups have much to share. The first step is to appeal to mutual interests and build informal relationships. Ultimately, the two groups must be tasked to work together in order continuously to improve.

CHAPTER 3: APPLICATIONS AND DATA

The reason businesses have data centers is, of course, to provide their IT services. Integral to those services are applications and the data that they generate. Changing some of the perceptions of how applications are designed and hosted can lead to a dramatic reduction in power and cooling requirements.

Reduce application variation

The first step is to reduce the number of applications. Organizations that have grown their IT services without centralized planning, or through mergers and acquisitions, will often find themselves with a number of applications performing similar operations. It makes business sense to reduce the number of instances of each class of application to as few as possible. The ideal would be one, but the reduction needs to be pragmatic and satisfy cost *versus* benefit rationale.¹⁰

The reason for this is that these applications require hardware to operate – including servers, storage, etc. The perspective of services is useful because no application is provided in a vacuum. The service view recognizes that it takes hardware, software, people, facilities, documentation and so forth to actually meet the needs of the business. Reducing the number of redundant services means

¹⁰ It must also better manage organizational risks such as regulatory compliance, information security, etc.

3: Applications and Data

that the hardware enabling these services can be decommissioned, thus reducing power and cooling loads.

Utilize cloud computing

Another approach is to shift the entire service-hosting concern to cloud computing providers.¹¹ These vendors provide extremely scalable services to customers at a far lower cost to the organization than that of trying to build a similar infrastructure itself.

Cloud computing vendors are investing hundreds of millions of dollars in their data centers. In some cases, such as Microsoft and Google, the sums are in the billions of dollars, with individual data center costs ranging from \$200–600 million for the facility alone.¹² Most firms can not remotely approach this level of investment.

Shifting the server burden to the provider leaves the customer free to build out and maintain the balance of the IT service, which includes the network infrastructure necessary to provide reliable and secure access to the cloud computing vendor. This shifts the bulk of the power and cooling needs to the provider, who hopefully

¹¹ “Cloud computing” refers to some level of IT-related services being hosted in a network, most commonly the Internet. A common attribute of cloud computing is massive scalability.

¹² “Facility” here means the building, land, power and cooling infrastructure, but not the IT equipment. The investment level is stunning.

3: Applications and Data

makes more efficient use of natural resources than any one isolated organization could.¹³

Software as a service (SaaS), platform as a service (PaaS) and storage as a service

SaaS enables organizations to use a vendor to provide virtually the entire application via the Internet. This avoids many of the costs associated with the traditional development and/or acquisition of the system and its ongoing hosting and maintenance.

PaaS could allow developers to have access to a platform that enables them to build scalable services that are hosted elsewhere. This allows for creation of custom applications but avoids the ongoing costs associated with hosting the application internally.

Storage as a service could be a tremendous boon for healthcare organizations and other groups struggling with massive storage requirements. There are many potential approaches to this, such as using a hierarchical storage model that provides fast local access to very active data and pushes less-frequently accessed data out to the cloud.

As with cloud computing, the hardware burden – with its underlying concerns regarding power, cooling and overall management – is shifted to a provider. The SaaS and PaaS solutions could be considered cloud-based if the offering includes the ability to scale on a massive level.

¹³ Due diligence is mandatory to verify that the provider has adequate security, business continuity plans, strategic direction, financial viability, etc.

3: Applications and Data

Application design

Development and engineering resources will need to collaborate to ensure that there is a proper balance between performance and the needs of the organization to reduce power consumption.

By collaborating, these groups can leverage concepts such as tiered resourcing, rather than assuming that all resources are fully available at all times. This may mean adding additional checks to allow a storage subsystem time to spool up to speed or other near-line subsystem time to power up, initialize and come online.

Data management

The next layer down, the data layer, must also be optimized in order to minimize power and cooling demands.

There are a number of potential improvement areas, including:

- Normalization – data should exist in a system of records and not be needlessly duplicated. The less normalized the data, the greater the amount of storage space required, drawing more power and generating more heat.
- Compression – data should be compressed to use as little space as possible.
- Tiered storage – contrary to what some may think, storage is not cheap. Organizations must investigate storing data on a platform commensurate with the likelihood of access. For example, data that is accessed frequently should be stored on hard drives. Data accessed less frequently will be stored on optical media.

3: Applications and Data

Data rarely accessed will be stored on tape and archived.

- Retention policy – it is not feasible to try and store everything forever. Not only does this require an incredible amount of storage space and high operating expenses, but it exposes the organization to potential lawsuits and discovery actions. IT, legal counsel, compliance and other stakeholders must develop and implement a retention policy that identifies how long and by what method different types of data will be retained. Not only will this lower risks for the organization, but it will also lower operating expenses and ongoing capital investment.

CHAPTER 4: BROAD THEMES

There are some broad concepts that can guide efforts to improve power efficiency in the data center for both IT equipment and facilities infrastructure.

Maximize utilization

Generally speaking, IT hardware and facilities equipment make more efficient use of electricity the closer they operate to their rated loads. For example, a server running at only 10% will be less efficient than one running at 90%. Likewise, a chiller running at 95% is better than one running at 40% of rated capacity. As a result, all things being equal, efforts to improve utilization of a given asset will result in improved power efficiency.

Newer tends to be more efficient

In general, newer hardware tends to be more efficient than older equipment. Vendors are moving rapidly not only to make the hardware more efficient through new technologies and designs, but also to add an increasing array of power management capabilities.

Modular approach

Engineers and those tasked with data center growth need to plan for modular implementations that allow for capacity to be added in phases. This ensures that only the resources that are needed are in use, and thus avoids having IT equipment,

4: Broad Themes

uninterruptable power supplies (UPSes), coolers and so forth needlessly underutilized, consuming power and not delivering any significant benefits.

Utilize zones

The historical approach has been to build a homogenous data center, by applying averages. Essentially, designers estimated the types of computers that would be installed, and then calculated the average watts-per-square-foot and cooling-per-square-foot values, and built accordingly.

Today, the data center cannot be viewed in this way, and cannot be built with equal power and cooling capabilities. This approach can lead to inordinate building expenses and low utilization. There is no point in IT and facilities trying to accommodate the potential of every rack drawing 30,000 kW when, in fact, only a minority of racks actually contain blade servers.

Instead, the data center needs to be designed, built and operated with zones in mind. In this way, high-density systems with their heavy power and cooling needs can be in one area supported by the appropriate infrastructure, and other zones can accommodate various other power and cooling demands accordingly. This approach would allow for more effective and efficient systems.¹⁴

¹⁴ Pods, or shipping container systems, present an opportunity to combine zones and a modular phased implementation approach to reduce capital investment in facilities as well as ongoing operating costs.

Looser environmental demands

Data centers need to invest in servers and other IT equipment that can run in warmer conditions and with broader temperature and humidity variations. This will then encourage vendors to improve the efficiency of the components that they provide. In its 4th generation data centers, Microsoft is actively pursuing servers that are more tolerant of warmer and more varied temperature and humidity levels. This move will lower the need for compensating environmental controls and result in more energy efficient data centers.¹⁵

¹⁵ See Mike Manos's fascinating blog post about Microsoft's new approach: "Our Vision for Generation 4 Modular Data Centers – One way of Getting it just right...", LooseBolts, entry posted 2 December 2008, at <http://loosebolts.wordpress.com/2008/12/02/our-vision-for-generation-4-modular-data-centers-one-way-of-getting-it-just-right>.

CHAPTER 5: IT HARDWARE

The increase in demand for IT services creates a corresponding increase in the number of servers in the data center. If not properly managed, the demands for power and cooling will rise in an unsustainable manner. This chapter reviews opportunities to improve the management of the demands from IT hardware.

Identify and decommission ghosts

Data centers may be running hardware that is generating operating expenses but no longer adding value. Essentially, these systems were put into production for some business purpose that is no longer relevant and nobody told IT to decommission the service and its underlying hardware. As a result, these “ghosts” continue to run, consuming power, generating heat and incurring support costs needlessly.

Eliminating ghosts is a low-cost, high-reward initiative. IT needs to monitor hardware and/or review activity logs to identify ghost servers, confirm with the business that they are no longer needed, and then decommission them.

There are a number of areas that can be reviewed to identify ghosts, including processor activity, network adapter activity, switch port activity, analyses from a network behavior analysis (NBA) tool, or power consumption from a metering power distribution unit (PDU). A combination of these can be used to confirm that a given server is not being used.

5: IT Hardware

A procedure should be implemented to ensure that hardware is appropriately decommissioned when it is no longer needed. If this is not done, ghosts can begin to reappear even while the clean-up effort is in progress. In fact, the routine appearance of ghosts during assessments might indicate a process weakness.

Server consolidation and virtualization

Servers that are running with low utilization need to have their services consolidated and any unused or redundant hardware removed.¹⁶ In some cases, services can be readily consolidated by combining those that are needed on a fewer number of hosts; for example, reducing the number of SQL servers by using one larger server that can host all of the required databases.

Many data centers are aggressively pursuing virtualization strategies to reduce the number of discrete servers. The intention is to create multiple virtual environments on a given hardware platform and therefore make more efficient use of that hardware than if the independent physical servers were to exist.

It is vital to have an understanding of how this technology will affect the data centre. For example, virtualization software that can put unneeded blades into sleep mode or simply turn them off and on as appropriate will have an impact

¹⁶ In an e-mail exchange with the author, James Hamilton, of Amazon, made the very sound observation that if the business benefit justifies the use of the newly available equipment's costs, then it should be used rather than removed from the data center.

5: IT Hardware

on power and cooling demands, and this impact needs to be understood. This new variability in demand can bring significant savings if holistically planned and correctly executed, in collaboration with the facilities department.¹⁷

There is one important thing to remember – there will only be savings if the old servers are powered off. In much the same way that work expands to fill all available time, IT services can expand to use all available resources – efficient or not – if not properly managed.

Blade computers

The use of blade servers is central to many of the high-density, high-performance computing plans that data centers are setting forth. The deployment of these servers must be carefully planned if they are to be supported by existing data centers.

As mentioned previously, the potential power and thermal densities could rapidly overwhelm a legacy data center's capacities. There are methods of decreasing the number of blades per chassis, the number of chassis per rack, and so forth, that can balance supply and demand.¹⁸ This is an example of where the Capacity Management process that is

¹⁷ The best way to tackle Green IT and improve data center efficiency is for IT and facilities groups to work together and not attempt isolated solutions that may be problematic for other stakeholders.

¹⁸ APC has a very good white paper which discusses this at a high level. See Neil Rasmussen, "Strategies for Deploying Blade Servers in Existing Data Centers" (APC, 2006), available, following free registration, at www.apc.com/go/promo/whitepapers/index.cfm?tsk.

codified in the Information Technology Infrastructure Library[®] (ITIL[®]) can play a critical supporting role by helping with the planning for future IT resources.¹⁹

Power management

Some IT services are very user-driven in terms of resource utilization. Services that are used primarily during an 8am–5pm business day will see usage spike during that time and then taper off outside of “prime time.” The problem is that many servers and other equipment are still running for prime time, consuming unnecessary power and needlessly generating heat.

In situations like this, power management options need to be configured to balance any performance needs against the need to cut power consumption and minimize heat generation. Newer servers have many options to power down processors, drives and other subsystems to reduce power consumption.

This recommendation is appearing a lot in the trade press, but there are two best practices that still need to be followed.

First, always test power management to verify compatibility with your technical production environment. Not all applications and peripherals can handle power management and some may require special considerations.

¹⁹ Capacity Management is codified in the book *ITIL Service Design* (Office of Government Commerce, 2007).

5: IT Hardware

Second, pre-production and production systems must match in terms of what power management applications are loaded and how they are configured. If these two areas are allowed to differ, then the risk of application failures or security breaches will increase, because testing in pre-production will not truly reflect the production environment.

Compression

Compression was mentioned in *Chapter 3: Applications and Data*, and the same benefit applies here – if data can be compressed and use less space, then less storage is needed and thus power and cooling demands decrease. The reason this topic is covered in two places is that compression can be written in software or embedded in hardware.

All things being equal, hardware-based compression tends to be faster and can compress control structures that the software may not be able to access, such as the database table structures, indices, and so forth.

Power supplies

Traditional servers have a power supply designed to convert AC power to DC. In addition, some servers may have an additional power supply for fault tolerance that operates if the primary unit fails.

The problem is that each and every power supply is inefficient during operation and as utilization drops, so does the efficiency.

5: IT Hardware

One way to address this is to use more efficient power supplies. Vendors are responding to market demands and producing more efficient supplies. The Environmental Protection Agency's 80 PLUS website²⁰ lists power supplies that have attained at least 80% efficiency. However, this is at full load, a state that most supplies will never actually enter. So while this is good information to have, power consumption must still be reviewed, specified in contracts, and then verified during acceptance testing.²¹

Another approach is to move towards a centralized power conversion facility that feeds DC power to the racks, thus removing the burden of having distributed power supplies in each server.

There are further options, such as running 240 volts to the servers, using AC-only devices that do not need DC power conversion, and reducing the power conversions that do happen inside devices to only one – such as ± 12 volts DC only. There are pros and cons to each that need to be evaluated.

Other server components

The power supply is only one component. The processor, memory, video card, motherboard and so on all have power and, thus, cooling requirements. Not all organizations build their own servers; instead, many rely on their vendors to build the systems and meet certain requirements – of which those for power consumption and cooling

²⁰ www.80plus.org.

²¹ The old adage of “trust but verify” is still appropriate.

should be formally written into the contract or purchase order. Compliance with the identified specification(s) should be verified during acceptance testing.

Reassess fault tolerance

With a desire to boost the availability of services, IT has often made fault tolerance a requirement when purchasing or building systems. As a result, we see $N+1$ or $2(N+1)$ fault tolerance²² being applied to all services delivered, regardless of value to the business.

To reduce power and cooling requirements as well as ongoing operational costs, IT organizations need to shift to a tiered fault-tolerance offering that is commensurate with business risks. In this model, only IT services that require $N+1$ or $2(N+1)$ fault tolerance are built to have these, and thus only those systems incur the costs, consume the power and generate the heat.

Reduce hardware variation

Where possible, reducing variation will help improve IT's ability to manage the various systems and, as a result, power efficiency. Imagine trying to manage 500 different types of file server; then imagine managing just five. A deeper

²² “ $N+1$ ” refers to device-level fault tolerance, where one server gets a second back-up server. “ $2(N+1)$ ” applies to a fault-tolerance approach where there is both device-level and geographic fault tolerance, for example, a cluster in a data center in one state and another cluster, synchronized to the primary cluster, located in a data center in another state.

5: IT Hardware

understanding of the system can be developed. Training becomes much more feasible. Best practices can be readily transferred from one like device to the next. There are a lot of benefits to this. While it may not be feasible to get down to only one type of system per category, “the fewer the better” is a good mantra to follow.

“Turn it off!” campaign

We will close this chapter with an approach that is extremely easy to understand but takes cultural change to enact. Simply put, people need to turn off unnecessary equipment.

This can apply to workstations, test networks and servers, development platforms, back-up systems that do not require hot failover, and so on. This is perhaps the easiest concept in the book to understand in principle, but it may be the most challenging to get people to follow.

One group made the message very clear by tying power costs to employee bonuses and putting a sign over light switches that read “remember your bonus.” If senior management supports the initiative and employees can clearly understand “what’s in it for me?” (WIIFM), there is greater likelihood that the initiative will achieve its desired objectives. This is true of everything in this book.

CHAPTER 6: FACILITIES – ELECTRICAL

Electrical infrastructure in a data center needs to be carefully planned. In today's era of high-performance, high-density computing systems, the design needs to accommodate significant differences in power requirements, both today and into the future.

Scalable modular power

Data centers need to be designed such that power can be added in phases. Otherwise, utilization is negatively impacted.

A modular approach ensures that as demand increases and the data center grows, so does the internal power grid, including distribution and back-up systems. This could take the form of zones in the data center that are provisioned once demand hits a certain threshold, planned expansions of the physical data center, or – the newest approach – using shipping containers, or “pods,” to add infrastructure and IT equipment in an incremental manner.

Power back-up systems

The ubiquitous UPS can use up to 18% of data center power and ranks third in terms of consumption.²³ The challenge, of course, is that IT

²³ The Green Grid, “Guidelines for Energy-Efficient Datacenters,” available, after registering, at <http://whitepapers.techrepublic.com.com/abstract.aspx?docid=353987>.

6: Facilities – Electrical

equipment needs clean, continuous power, and the UPS has been deployed in response.

Essentially, a UPS is a stop-gap measure meant to provide power during the interval between initial power loss and generators coming online and providing stable power. Historically, correctly sized and installed UPSes did just that.

Today, we recognize that UPSes create a number of challenges in addition to power consumption:

- Online UPSes can generate tremendous heat and should not be located in the environmentally controlled data center area.
- Batteries have a limited lifespan of 5–7 years and present a hazardous waste disposal challenge.
- UPSes have traditionally been significantly oversized in an attempt to provide sufficient current and future growth capacity.

In response, firms are investigating a number of new technologies, including flywheels, fuel cells and capacitors.

Another possible approach is to supply different zones of the data center with different levels of power back-up commensurate with the criticality of the particular services that they deliver.

Power zones

Running massive amounts of power to each rack is not a very realistic approach. One solution to this is to create different zones in the data center. High

6: Facilities – Electrical

performance systems drawing considerable power are put in one area, along with the requisite infrastructure, systems with another level of demand are put in another area, and so on. In this manner, power and cooling capacity can be focused on the demands of the types of systems present.

Power distribution

Transformers and distribution systems generate heat, either by stepping down high-voltage AC to branch circuits, or by converting AC power to DC. This not only wastes the power at that point, but potentially affects cooling as well. For this reason, power distribution units (PDUs) should be located outside the data centre wherever possible.²⁴

At the IT hardware level, US data centers that are running at lower voltages should review their equipment and eliminate 120 volts where possible, and move to the higher 208-volt rating, which is more efficient. Most IT equipment in the US is sold with switched-mode power supplies, which either automatically change to or require that a physical switch on the power supply be toggled to the higher voltage.

In addition, there are opportunities to leverage the use of the data center's three-phase 230-volt supply (line to neutral) to further boost efficiency over 208 volts (line to line), and reduce the

²⁴ This is assuming that a design cannot be developed that pragmatically reduces the number of PDUs and/or voltage conversions.

6: Facilities – Electrical

number of wires in the data center.²⁵

Newer PDUs offer additional power management capabilities, including the ability to turn ports off and on and capture actual power utilization data. This ability to track power usage is instrumental in developing a true understanding of consumption and enabling effective management.

Organizations should investigate how to distribute power in a manner that makes the most sense given their current infrastructure, project growth, and available resources. There are many options that can be explored.²⁶

Distributed generation

Data centers have always used generators as a method of augmenting utility power in the event of disruptions to supply. While effective, it has never been particularly efficient. Organizations today are investigating a number of renewable energy alternatives as means not only of dealing with some forms of supply disruption, but also of supplementing utility power and thus driving down energy costs.

²⁵ See Neil Rasmussen, “Increasing Data Center Efficiency by Using Improved High Density Power Distribution” (APC, 2006), available at http://www.apcmedia.com/salestools/NRAN-6CN8PK_R0_EN.pdf.

²⁶ See The Green Grid, “Quantitative Efficiency Analysis of Power Distribution Configurations for Data Centers,” posted 1 December 2008, available at www.thegreengrid.org/en/sitecore/content/Global/Content/white-papers/Quantitative-Efficiency-Analysis.aspx.

6: Facilities – Electrical

Distributed power generation through renewable energy sources is a fast-growing industry. Some examples of technologies that firms are leveraging include solar, wind, hydroelectric, geothermal and wave power.

CHAPTER 7: FACILITIES – COOLING

Cooling is driven by power consumption and itself consumes power. Moreover, not only are some data centers constrained by power, but some are limited by their cooling capacity. The good news is that there are many ways to improve the use of existing services before building a new data center is required.

The Arrhenius equation and data center cooling

There is a very simple reason why data centers must be cooled. As the temperature rises, so too does the component failure rate. Svante Arrhenius was a Swedish chemist who noted that chemical reaction rates double for every additional 10°C.²⁷ We can apply this to IT equipment in the following way: for each 10°C rise in temperature, the component failure rate will double.²⁸ As this occurs, IT service availability will suffer and, thus, negatively impact the organization. As a result, some level of cooling is mandatory in the data center, but the real question is “how much is sufficient?”

²⁷ See http://en.wikipedia.org/wiki/Arrhenius_equation.

²⁸ See Kirk W. Cameron, Rong Ge and Xizhou Feng, “High-Performance, Power-Aware Distributed Computing for Scientific Applications” in *Computer*, November 2005, available at www.computer.org/portal/site/computer/menuitem.5d61c1d591162e4b0ef1bd108bcd45f3/index.jsp?&pName=computer_level1_article&TheCat=1005&path=computer/homepage/1105&file=highperf.xml&xsl=article.xsl&.

Input temperature: ASHRAE standards, vendors and reality

An early action is to verify that thermostat settings are as warm as possible, based on inlet air temperatures and not on the hot exhaust. It is a common misconception that thermostats should be located by the exhaust. Vendors and standards bodies are setting temperature and humidity values based on inlet air.

If the thermostat is located near the hot exhaust air, then the data center will be colder than necessary. This additional cooling results in unnecessary energy consumption and wear on mechanical systems.

There are guidelines that data centers can follow. The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) publishes recommendations on data centers. Recently, they recommended expanding the cooling and humidity requirements to 18–27°C (68–77°F) dry bulb temperature and 40–55% relative humidity.²⁹

While it is useful to know the ASHRAE standards, it is important to bear in mind that they are not always followed by vendors. Hardware warranties often have clauses stipulating that a warranty will be void if the equipment in question is operated outside the environmental tolerances specified. Organizations need to be aware of these clauses

²⁹ See ASHRAE, “2008 ASHRAE Environmental Guidelines for Datacom Equipment” (ASHRAE, 2008), available, after registration, at www.ashrae.org/publications/page/1900.

7: Facilities – Cooling

and act accordingly to avoid unintentionally voiding warranties – especially on critical systems.

Cooling zones

Just as the data center should be zoned for different tiers of power, it should also be zoned for different cooling requirements. Given that cooling is a function of power, the need to improve cooling to support the higher power systems is a given. In practice, the cooling zones will likely mirror the power zones.

An important characteristic of high-performance systems is that if cooling fails they will overheat very quickly. As a result, it is critical not only to have the correct cooling capacity in production, but also to put in place the necessary redundancy and/or contingency plans should the primary cooling system fail.

Hot aisle/cold aisle

This is a critical design concept. IT equipment needs to be positioned in aisles such that intakes all face one aisle and exhausts dump into another aisle, as shown in Figure 1.³⁰

³⁰ CRAC = Computer Room Air Conditioner.




Figure 1: Hot aisle/cold aisle

This approach allows for the segregation of hot and cold air, which dramatically improves efficiency. Preventing the uncontrolled mixing of cool and warm air means that the exhaust air can be as hot as possible without degrading the inlet air. This provides the further opportunity of using economizers to create free cooling due to the potential for a larger differential between the exhaust and external air.³¹

Correcting air flow

Data centers that have not paid sufficient attention to air flow management can reap significant benefits by ensuring that:

- the volume of air supplied matches the volume of air demanded for each device in the data center.
- dampeners and airflow control mechanisms are in place and properly set.
- a hot aisle/cold aisle layout with partitions is used to segregate exhaust and inlet air.

³¹ Economizers are discussed later in this chapter.

7: Facilities – Cooling

- server cases are closed correctly – do not leave them open when running. If chassis have been left open so that a component could be changed, then close them. Servers need their cabinets closed for air to flow correctly.
- server face blanks are installed – racks that are in a hot aisle/cold aisle arrangement need blanking plates installed so that cold air is only drawn through IT equipment in an appropriate manner.
- cable access points are sealed – if there are air movement methods in use that have openings for cables to pass through, such as a raised floor, these openings should be sealed to prevent unmanaged air flow.
- non-standard air inlets/exhausts are corrected – if overhead cool air ducts are used, they should dump straight down in front of the servers. General-purpose four way diffusers such as those seen in office spaces should not be used.
- long cool-air runs are avoided – the computer room air conditioner (CRAC) should be located as close to the equipment as possible, to reduce demands on fans to push the air.
- variable speed fans (VSFs) are used – by replacing constant speed fans with variable speed units that can increase or decrease the flow of air as needed, efficiency can be improved.
- ducting is optimized – flexible air ducts should be as straight as possible, avoiding kinking and any unnecessary openings. For example, someone may have performed a service and knocked a duct off a connector, so that the system is now either pushing cool air into the

7: Facilities – Cooling

ceiling or drawing air from the ceiling – both are costly.

- raised floors are of sufficient height – IBM, for example, recommends no less than two feet.³²
- raised floor spaces are clear – the crawl space can accumulate a lot of debris and dust over time, and this can hamper air flow. The area should be cleaned of debris and any abandoned cables or equipment removed.
- cooling tiles are located correctly – raised floors should have cooling tiles located directly in front of populated racks only. They should not be located haphazardly, wasting cool air.
- hot air returns are correctly located – in general, they should be above the hot aisle and not located haphazardly.
- wiring is organized – ideally, wiring should follow a management plan and be organized to minimize impacts on air flow. If wiring is located under the raised floor, review the cost-benefit of moving to overhead trays and clearing out the area under the raised floor.
- partitions are used to avoid uncontrolled mixing of air – if a hot aisle/cold aisle or zoned system is in place, use partitions to avoid unmanaged mixing of air between dedicated hot and cold areas.

³² See Mike Ebberts, Alvin Galea, Marc Tu Duy Khiem and Michael Schaefer, “The Green Data Center – Steps for the Journey,” IBM RedPaper (August 2008), at www.redbooks.ibm.com/abstracts/redp4413.html.

Leverage experts

In some instances, utilities are offering free professional services to help improve efficiency and can leverage techniques such as computational fluid dynamics modeling to predict air flow and cooling in the data center.³³ There are also experts in the data center and cooling community who can provide objective recommendations.

Lastly, don't make air flow improvements a one-time event. All of the above should be reviewed at regular intervals. This might be once a quarter, on a rolling weekly basis, etc.

Water cooling

In another example of the old becoming new, water cooling is gaining interest again, in part because it is extremely efficient at reducing heat. In fact, a liter of water can absorb about 4,000 times more heat than the same volume of air.³⁴ Opportunities exist to use water cooling not only in the data center but also within racks, chassis and even at the component level. While implementation costs may be higher, the savings occur through the reduction of ongoing cooling costs and increased data center capacity.

³³ See http://en.wikipedia.org/wiki/Computational_fluid_dynamics.

³⁴ See Mike Ebbers, Alvin Galea, Marc Tu Duy Khiem and Michael Schaefer, "The Green Data Center – Steps for the Journey," IBM RedPaper (August 2008), at www.redbooks.ibm.com/abstracts/redp4413.html.

Insulation

Another basic area to review is the insulation in the data center. This can help protect the data center from changes in the external environment.

Examples of insulation include:

- Insulation in external walls and ceilings
- Insulation in internal walls and ceilings where appropriate
- Insulated windows
- Door seals to prevent uncontrolled air exchange.

Economizers

Data centers are increasingly deriving significant benefit from economizers. These devices leverage the external climate to make use of “free cooling” when possible. In other words, rather than powering a chiller, the environment is tapped to provide the cooling needed. For some data centers, the outside temperatures are cool enough at night, at certain times of year or even year-round to make this appealing.

There are two basic types of economizer:

- Air side – mixes cooler external air with the warm exhaust air from the data center. The air must be filtered to remove particulates, and humidity adjustments may be necessary.
- Liquid side – uses a water, glycol or some other refrigerant cooling loop, which then has one or more heat exchangers to remove the heat from the exhaust air before it is returned to the data center. This method is sometimes

7: Facilities – Cooling

identified as “closed,” as it does not directly expose the controlled data center environment to the uncontrolled external environment.

The use of economizers is catching on quickly, and is one of the reasons why locating data centers in cooler climates is appealing. Organizations are reviewing locations in northern parts of the United States, Canada, Iceland, and so forth to, in part, make use of free cooling.

A related method is making use of the relatively cool ground by creating a well field, with wells potentially extending down 800 feet to cool the water. This water is then used to cool IT equipment, thus forgoing the use of the chiller, which of course consumes power.³⁵

This variation helps illustrate that there are many ways to improve the efficiency of cooling the data center. Doubtlessly new approaches will continue to appear and serve to highlight why training is needed. Exposure to new ideas through trade shows and industry groups can help organizations stay abreast of this rapidly advancing field.³⁶

³⁵ See John Peterson, “Using a Geothermal System to Cool Your Data Center,” in *The Data Center Journal*, 11 December 2008, available at http://datacenterjournal.com/index.php?option=com_content&task=view&id=2318.

³⁶ Another interesting example is IBM’s Zurich lab, which uses waste heat to warm a swimming pool and nearby buildings.

CHAPTER 8: SELECTING A DATA CENTER LOCATION

Gartner estimates that 70% of the Global 1000 will need to replace their data centers in the next five years.³⁷ The primary drivers of this need are cooling, energy costs, and energy availability, but there are other dimensions that must also be considered.

The following sections describe data center site-selection criteria that could affect cooling and thus power consumption.

Climate

Because cooling is such a concern, there are benefits to picking locations that are cooler. This allows the use of economizers a greater percentage of the time to make use of “free cooling” from the environment, and thus reduce the need for mechanized cooling that requires power. Spending money just to reduce the ambient temperature is a non value add (NVA) expenditure.³⁸

Cost of power

This has become an important criterion. Energy costs can vary dramatically and thus affect ongoing operations costs. In January 2008,

³⁷ Rakesh Kumar, “Green IT: Immediate Issues for Users to Focus On” (Gartner, 7 August 2008).

³⁸ Would you rather cool a data center in a 100-degree desert or one that is located in a cool climate, wherein waste heat is actually used to warm the building?

8: Selecting a Data Center Location

commercial power in Connecticut cost 15.43 cents/kWh. In contrast, the cost in Idaho in the same month was 5.12 cents/kWh.³⁹ As well as examining this data at a single point in time, it is also important to trend the data over time to try and understand where the prices are headed. To underscore the differences even further, one data center operator expected to pay 1.85 cents/kWh by locating in East Wenatchee, Washington.⁴⁰ Clearly, different locations can yield very different costs.

Energy availability

Research must be conducted in order to understand how robust the power grid is and identify any availability issues. This may include talking to the utility, reviewing news sources, and interviewing local companies, particularly energy-intensive ones such as foundries, glass producers, forestry, etc.⁴¹

³⁹ “Average Retail Price of Electricity to Ultimate Customers by End-Use Sector, by State,” US Energy Information Administration, available at www.eia.doe.gov/cneaf/electricity/epm/table5_6_a.html.

(Note: this link yields the most recent data. The author has used the January table for this comparison.)

⁴⁰ Mark Fontecchio, “Data center operations flock to central Washington” in *Data Center News*, 22 May 2007, available at http://searchdatacenter.techtarget.com/news/article/0,289142,sid80_gci1255876,00.html.

⁴¹ The US Department of Energy provides a list of energy-intensive industries in the US. See www1.eere.energy.gov/industry/program_areas/industries.html.

8: Selecting a Data Center Location

Tax incentives

State and local governments are offering tax incentives to entice data centers to be built in their areas. These can help to defray some of the construction and ongoing operating costs.

Bonds/government financing

As with taxes, state and local governments are looking at how they can provide capital to lure data centers into their areas.

Political climate

The way in which government perceives data centers should be factored in, not to mention political stability – especially if potential sites are located in other countries. For example, Iceland and Siberia are considerations given their cold climate but Iceland's political climate would appear to be more favorable at this time.

Natural disasters

Consideration must be given to the types of natural disasters a given site may face. For example, depending on the design of the power grid, hurricanes, snowstorms, ice storms, thunderstorms or tornadoes could disrupt the flow of power – not to mention disrupt data center activities in other ways.

An organization can create a site selection worksheet and rate various potential locales on the above, plus any other criteria to aid in decision making. It is important to look at the larger

8: Selecting a Data Center Location

picture, including total implementation costs, ongoing support costs, risks associated with the site, and so on. The more informed the site location decision, the better.

CHAPTER 9: MONITORING AND REPORTING

A large part of understanding what needs to be done lies in understanding the current state of the data center. To control power and cooling costs and maximize utilization, we must be able to measure and report on this data, trended over time.

Monitoring is becoming increasingly inexpensive and granular. Examples include:

- The use of metering, or intelligent, power distribution units that can give visibility down to the port level.
- Cooling systems that can track and report on their own activity.
- Servers that report on internal chassis temperature.
- The use of a sensor grid to track temperature and humidity throughout the facility.

As with any endeavor, there must be agreement on what to track, how to analyze the data, who needs what reports, and so forth.

The following are some examples of data elements that can be reported on and analyzed.

Power usage effectiveness (PUE)

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

Total facility power is all power going into the data center, whereas IT equipment power is only IT systems such as servers, switches, firewalls, etc.

9: Monitoring and Reporting

This metric is useful because the resulting ratio can help identify how much overhead exists when planning for future needs. For example, if PUE is 1.3 and a new server is being looked at that will need 1,000 watts, then the estimated power demand will be $1,000 \times 1.3 = 1,300$ watts.

Data center infrastructure efficiency (DCiE)

This is the reciprocal of PUE and can illustrate what percentage of the power in the data center is going towards IT equipment.⁴²

Monthly data center energy costs

Surprisingly few CIOs know how much electricity the data center consumes and what it costs the organization on a monthly and annual basis. This needs to change.

Monthly energy costs per IT service

To help drive accountability, IT must become increasingly granular with its measurements and ideally be able to tie energy and cooling costs to services that IT is providing to the business.

⁴² Both PUE and DCiE are metrics being recommended by the Green Grid industry group. See The Green Grid, “Green Grid Metrics: Describing Data Center Power Efficiency,” 2007, at <http://www.thegreengrid.org/en/sitecore/content/Global/Content/white-papers/The-Green-Grid-Data-Center-Power-Efficiency-Metrics-PUE-and-DCiE.aspx>

9: Monitoring and Reporting

Baseline now!

There are not any absolute standard metrics at this point in time. Instead, organizations need to identify what metrics are important to them and begin collecting baseline data today. Without baselines, IT will have a difficult time showing what progress has been made, as well as understanding cause and effect.

CHAPTER 10: CONCLUSION

Improving data center power and cooling to protect the environment, while reducing ongoing operating expenses and protecting the brand, is a great example of how environmental benefits and business benefits are not mutually exclusive.⁴³

Green IT is also an area in which solutions must properly address people, processes and technology. In terms of people, there must be the right culture, training, awareness, incentives and support from senior management. Processes must be designed, implemented and properly monitored to achieve their stated objectives of creating and protecting value. Lastly, we must understand Green IT requirements and both acquire and develop technical solutions accordingly.

This pocket guide has given a brief overview of people and process considerations, and then mainly focused on high-level technical opportunities.

To leverage the ideas in this book, a roadmap should be assembled to identify what can be done today, with a view to both the short term and the long term, to help support the organization while reducing the negative environmental impacts of the power consumed by the data center. Fortunately, IT organizations have many opportunities before them to make this a reality.

⁴³ Indeed, some organizations are using environmental achievements and certifications as a marketing opportunity to promote their brands.

APPENDIX: ADDITIONAL RESOURCES

There are many resources available to help organizations to green their data centers, and new resources are appearing all the time. Rather than provide a static list of resources in this book, the author has provided a dynamic resource list, which is available at:

www.spaffordconsulting.com/GreenITResources.html.

To stay abreast of news relating to data centers and process improvement, subscribe to the author's newsletter via email or RSS:

www.spaffordconsulting.com/dailynews.html.

ITG RESOURCES

IT Governance Ltd sources, creates and delivers products and services to meet the real-world, evolving IT governance needs of today's organizations, directors, managers and practitioners. The ITG website (www.itgovernance.co.uk) is the international one-stop-shop for corporate and IT governance information, advice, guidance, books, tools, training and consultancy.

www.itgovernance.co.uk/green-it.aspx provides a growing range of Green IT books and tools for use by corporations and organizations who are tackling this 21st-century critical business area.

www.27001.com is our information security-focused website that deals specifically with information security issues in a North American context.

Pocket guides

For details of the entire range of pocket guides, which includes a number of free Green IT titles, follow the links at www.itgovernance.co.uk/publishing.aspx.

Toolkits

ITG's unique range of toolkits includes the IT Governance Framework Toolkit, which contains all the tools and guidance that you will need in order to develop and implement an appropriate IT governance framework for your organization. Full details can be found at www.itgovernance.co.uk/products/519.

For a free paper on how to use the proprietary Calder-Moir IT Governance Framework, and for a free trial

version of the IT Governance Framework toolkit, see www.itgovernance.co.uk/calder_moir.aspx.

Best Practice Reports

ITG's new range of Best Practice Reports – including one on Green IT – can now be found at www.itgovernance.co.uk/best-practice-reports.aspx.

These offer you essential, pertinent, expertly researched information on an increasing number of key issues.

Training and consultancy

IT Governance also offers training and consultancy services across the entire spectrum of disciplines in the information governance arena. Details of our range of training courses can be accessed at www.itgovernance.co.uk/training.aspx and descriptions of our consultancy services can be found at www.itgovernance.co.uk/consulting.aspx.

Why not contact us to see how we could help you and your organization?

Newsletter

IT governance is one of the hottest topics in business today, not least because it is also the fastest-moving, so what better way to keep up than by subscribing to ITG's free monthly newsletter *Sentinel*? It provides monthly updates and resources across the whole spectrum of IT governance subject matter, including risk management, information security, ITIL® and IT service management, project governance, compliance and so much more. Subscribe for your free copy at www.itgovernance.co.uk/newsletter.aspx.