Realtime publishers

The Shortcut Guide™ To

Data Center Energy Efficiency

sponsored by

APC by Schneider Electric

David Chernicoff

What we can do for your data center efficiency is unlimited, however the FREE trial of our tools is not.

How efficient is your current data center?

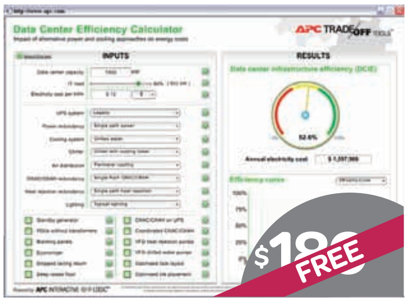

Data Center Efficiency Calculator

What's your data center's actual carbon footprint?

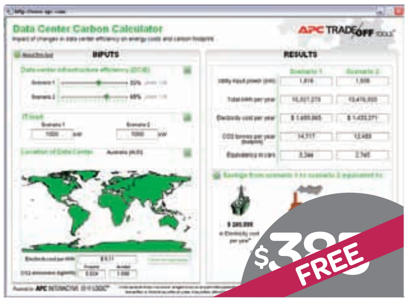

Data Center Carbon Calculator

Know how much you can save by re-designing?

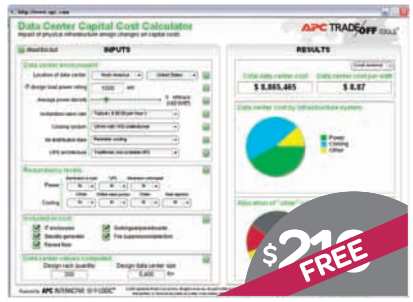

Data Center Capital Cost Calculator

APC's web-based TradeOff Tools™ are the practical application of data center best practices. The tool's decision-making criteria and recommendations stem from customer feedback about actual installations and are founded upon the logic and research determined by the APC Science Center, the internationally renowned research team responsible for more than 100 white papers on the subject.

How much will virtualizing save you?

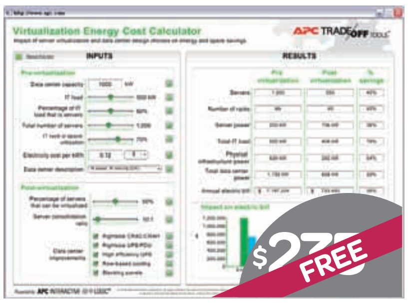

Virtualization Energy Cost Calculator

Is your power balanced with what you need?

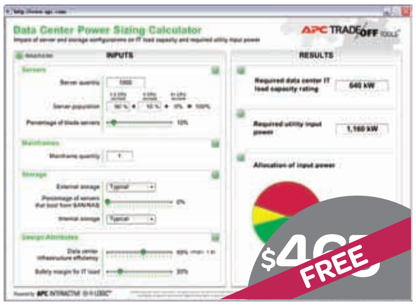

Data Center Power Sizing Calculator

Which cooling solution will work for you?

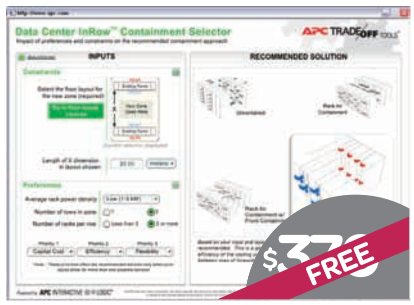

Data Center InRow™ Containment Selector

Introducing APC interactive TradeOff Tools,™ web-based calculators for predicting efficiency and availability in your data center planning.

The time has come to shift the efficiency debate from efficiency as a conceptual ideal to implementation - the measuring of and planning for efficiency in your installation. Decisions you make today will certainly affect your efficiency tomorrow. Want to virtualize but unsure of how it will affect your power and cooling efficiency? Considering changes to your data center architecture approach or heat containment strategy, but are reluctant to rock the boat? Balancing the pros and cons of data center planning can be a real challenge. The effects of a bad decision are further complicated by rising energy costs and the need for green.

Introducing APC's latest innovation for these critical decisions: the TradeOff Tools suite. These interactive, online, and (most important) precisely formulated calculators deliver the actual implications of decisions in your data center design or planning. With TradeOff Tools, you can accurately calculate the impact that new equipment, server virtualization, design changes, and heat containment strategies will have on your facility. Because the data plugged in is your own data, and not industry averages, the results are your own, as well.

More than a widget

The benefit of TradeOff Tools is easy to calculate. Now you can know, ahead of time, without question, how changes will affect your facility. There has never been a tool this easy to use, that provides information this actionable. And did we mention these great tools are FREE for a limited time? Saving money, energy and time has never been so simple. Try out TradeOff today and know before you go.

APC TRADEOFF TOOLS™

Try online for the next 30 days for FREE!

Visit www.apc.com/promo Key Code i202w
Call 888-289-APCC x6029 • Fax 401-788-2797

APC® by Schneider Electric

e-mail: esupport@apc.com • 132 Fairgrounds Road, West Kingston, RI 02892 USA • 998-1434

Introduction to Realtime Publishers

by Don Jones, Series Editor

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We've made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book's production expenses for the benefit of our readers.

Although we've always offered our publications to you for free, don't think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you $40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the "realtime" aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We're an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I'm proud that we've produced so many quality books over the past years.

I want to extend an invitation to visit us at http://nexus.realtimepublishers.com, especially if you've received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you're sure to find something that's of interest to you—and it won't cost you a thing. We hope you'll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Copyright Statement

Chapter 1: Impact of Virtualization on Data Center Power & Cooling

Virtualization is one of the base technologies driving IT planning. With the clear benefits that virtualization technologies bring to the data center, the impact that virtualized environments have on the overall power consumption and cooling capacity currently present in most data centers is often overlooked. Optimizing the physical infrastructure that provides power and cooling to the data center, to work in concert with the deployment of virtualized environments, is critical to allowing the business to reap full benefit from virtualization in the data center.

Introduction

As virtualization becomes more widespread in your data center, you will likely find that the structure of your data center is changing. Although some servers may be removed, many servers will simply be repurposed, while additional, more powerful servers are installed. These new servers will have increased power and cooling needs when compared with lower-powered servers that previously occupied the same space and will impact the efficiency of your data center power delivery and cooling. This is especially true if you are in a growing organization where the level of systems density maintained in the data center was fairly low. You may find that additional energy is now required by your data center, both as power and cooling.

In an environment already heavily dependent on dense technologies such as blade servers, the trend towards virtualization may result in the opposite case: a surfeit of power and cooling capability. In both cases, the result is similar—a poorly optimized energy delivery model that results in a decrease in efficiency.

The second situation is the more common one, as early technology adopters are likely to have already built a high-density data center environment. Properly optimizing power and cooling will often result in a significant reduction in the overall power consumption of the data center, with its attendant cost savings.

Properly designed, the power and cooling systems deployed in your data center will provide the support needed to deploy your current virtualized environment in the most cost-efficient fashion while allowing those same systems to scale to meet the needs of a growing business. Potentially, this may require changes to your power and cooling infrastructure. As one of the hallmarks of virtualization is the ability to quickly provision and scale your computing environment, it is only natural to deploy power and cooling systems that can meet the challenge presented by the non-static data center environment.

Data Center Infrastructure Efficiency

The basic, primary requirement for implementing an energy efficient data center is to be able to measure the relative efficiency of the current, or planned, power and cooling systems deployment within that data center. As with any technology, the trick becomes finding an effective measuring technology that is vendor-independent and applicable to a wide variety of conditions.

To that end, The Green Grid (www.thegreengrid.org), an industry alliance made up of more than 100 member companies with products or technologies focused on the data center, has been launched to develop a standard set of measurements of data center efficiency to be used by the industry. To date, their first two benchmarks, Power Usage Efficiency (PUE) and Data Center infrastructure Efficiency (DCiE), have received relatively wide acceptance. These benchmarks can be used to build a model of your data center's energy efficiency that you can employ when determining which products and technologies, from a wide range of vendors, will be best utilized in your environment.

Understanding the Concepts Behind PUE and DCiE

The basic concept between PUE and DCiE is the ability to establish a metric for evaluating the efficiency of your data center. Making use of these metrics assists you in determining whether you can meet your needs with your existing systems, as well as whether those needs are being met efficiently. PUE and DCiE are complementary measurements. PUE is defined as:

Total Facility Power/IT Equipment Power

The reciprocal measurement, DCiE, is defined as:

1/PUE = IT Equipment Power/Total Facility Power x 100%

Total Facility Power is the power measured at the utility meter. This can be an issue in a mixed-use facility due to the need to establish the amount of power used solely by the data center facilities within the building without including power requirements of other organizations housed within. **IT Equipment Power** is the power delivered to devices within the data center used specifically for the management, processing, storage, or routing of data (servers, storage devices, network components, and so on). This measurement would most likely be taken as a value of the output of all power distribution units within the data center.

Using the PUE and DCiE Model in Your Environment

Clearly there is no high end limit to a PUE score, but a PUE value of 1 would indicate the utilization of 100% of the delivered power by IT equipment only. For comparison purposes, the ability to determine how efficiently power is being utilized within your data center gives you a useful metric for evaluating the impact of changes to the efficiency of your data center energy utilization. Reports on data center surveys provided by the Green Grid show common values in the 3.0 range, with the conclusion that, with proper redesign, those PUE values can be cut by as much as one third to one half.

The DCiE measurement gives you the reciprocal information to that provided by the PUE measurement. For example, if the data center was seeing power delivery from the grid equaling 5000 watts and the measurement of the total IT power consumption was 2000 watts, the DCiE metric would indicate 40% efficiency.

Both PUE and DCiE metrics are in use by vendors offering power efficiency services to business users; although the metrics present the same information, in different forms, you will likely find that PUE is in more common use. Determining your own PUE calculations is a valuable task before beginning any data center-related projects.

Understanding the Benefits of Virtualization

One of the reasons that data centers have proliferated is due to the explosive increase in the number of servers used by large enterprises and the need to consolidate their location in order to improve management and support. But the cost of supporting individual physical servers, in terms of both IT support and facilities costs, has also continued to grow. This increase reduces the effectiveness of the management efficiencies that have been gained by consolidating the physical location of your enterprise servers.

A large percentage of the servers in a data center spend the bulk of their time in an idle state. Although some applications place continual loads on the server, more often than not, the average data center server is simply sitting there, drawing power as much as 80% of the time, according to some analyst estimates. Even at idle, these servers are drawing approximately 30% of the power that they would draw at peak load, while they sit and wait for work. These costs have to be absorbed somewhere in the enterprise, and optimizations that can reduce these costs are important to IT budgets.

An IDC report in 2007, "Enterprise Class Virtualization 2.0," estimated that 50 cents of every dollar spent on servers was spent to provide power and cooling. The same report estimated that by 2010, that amount will have risen to 70 cents. With the potential lifetime costs of providing power and cooling to individual physical servers being 70% of the total expense, server consolidation using virtualization becomes the logical progression for upgrading any server running suitable applications.

The changes in the data center wrought by the move to virtualized environments can be very obvious; a smaller number of more powerful physical servers replacing a large number of less powerful standalone servers. Thus, the overall power and cooling requirements of the data center are likely to go down, resulting in reduced expenditures for those needs, at least initially. This results in a degraded PUE value, as it increases due to the reduction in IT load while the physical infrastructure remains the same, unable to adjust for the drop in demand.

However, there are a number of issues that will crop up, some of them unexpected, that provide an opportunity to more carefully optimize the energy efficiency of your data center and may impact how you manage your data center facility:

- The behavior of your servers will change—Virtualized environments allow for the dynamic implementation of servers and applications. This means sudden changes on the load on the physical server hosting multiple virtual servers and the need to deliver efficient power and cooling to those physical servers in order to support the server state changes without over-sizing the availability of resources provided to your servers.
- Server support priorities will change—When a single server was hosting and supporting a single, low-priority application, it received the level of support suitable for a server that wouldn't have a severe or immediate impact on the business process if it was offline temporarily. When you concatenate multiple servers with those types of applications to a single physical server, that host server suddenly acquires a much more mission-critical role, as its health can now affect a much larger pool of users. Thus, the smaller number of physical servers that are the hosts for those low-priority servers need to be considered high-value servers and get the same level of power and cooling support offered to your mission-critical servers.
- Matching the power and cooling model to the leaner IT model—By providing an optimized power and cooling environment to match the newly streamlined server environment IT can see a much greater overall cost reduction than would be achieved by simply reducing the number of physical servers.

The Benefits of Virtualized Environments and Efficient Energy Design Implementation

Adding virtualization to your data center will not only allow you to enjoy the benefits that virtualization brings to your computing environment but also allow IT to reap the benefits of the knowledge that already exists, in-house and in the industry, on the technologies and techniques for building, deploying, and managing power and cooling for high-density rack mounted servers. Because the physical server consolidation that virtualization brings reduces the number of servers necessary to deploy in your data center, the most obvious benefit is the reduction in power required to maintain the same level of service to the business consumer of IT resources. But additional savings may be possible due to the reduced needs of the supporting power and cooling infrastructure. Maximizing the potential savings should be part of the overall migration plan for the data center virtualization implementation.

The obvious way in which this would work would be the need for fewer power delivery systems and the reduction in cooling capacity due to the lower number of servers. However, it's not quite that simple; the layout or design of your data center may preclude major changes to the delivery of cooling services or fixed losses may make your reduction in power and cooling needs effectively less than might be expected.

Impact of Virtualization on Power Consumption and Efficiency

Seen from the outside, the implementation of a virtualized computing environment will always deliver better efficiency when seen in terms of computing power and optimization. What isn't as clear is that the conversion to virtual systems will always reduce the energy efficiency of the physical infrastructure of the data center due to the fixed losses inherent in the existing power and cooling systems. Not taking this opportunity to match the power and cooling delivery to the new conditions in the data center is effectively leaving money on the table; the savings potential in matching the power and cooling needs to the new reduced load is considerable.

We've already discussed the importance of understanding the PUE of your data center; grasping the total power aspects of the PUE will help you comprehend the impact of fixed losses on your data and cooling infrastructure. All power consumed in the data center, by both IT workloads and support workloads, is your total power consumption. The power consumed by those non-IT workloads—such as the cooling systems, power system inefficiencies, and other data center physical infrastructure systems—is the loss we are addressing here.

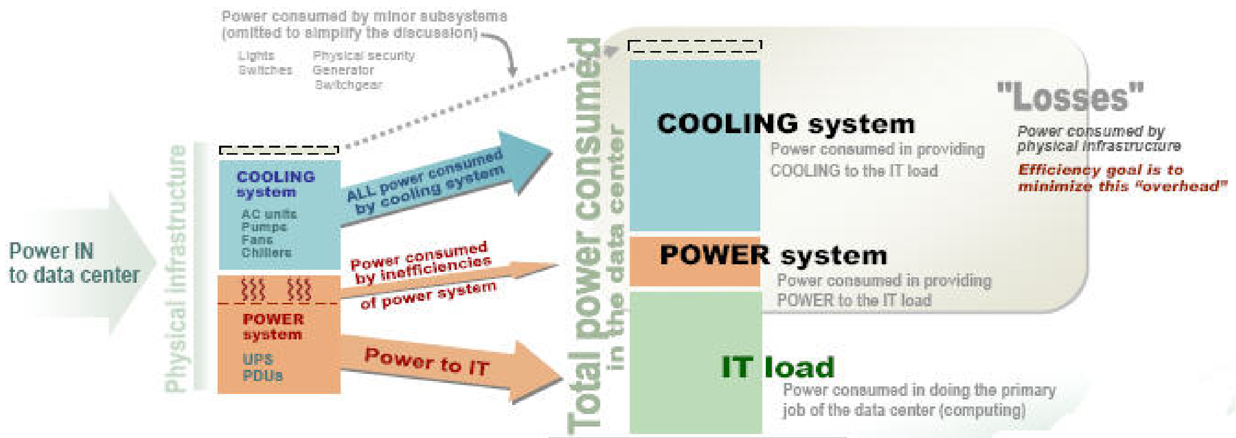

Figure 1.1: Power consumed by the physical infrastructure defines losses.

Losses that are experienced are of two types; fixed and proportional. Some equipment will draw the same level of power regardless of load; others will vary based on workload, giving you the two types of infrastructure loss.

- Fixed loss—This loss remains the same regardless of the workload. The system or device consumes a fixed amount of power regardless of the task at hand. When there is a high load present, as might have been the case before the data center virtualization was accomplished, these fixed losses are a small percentage of the overall power requirement. As the power need of these devices doesn't change as the overall load is reduced, these fixed losses, and their attendant cost, become a noticeable percentage of the overall power consumption of the data center.
- Proportional loss—The amount of loss for a device is directly proportional to the workload on the device: increase the workload 50%, the loss increases 50%; decrease the workload 25%, the loss decreases 25%.

The bottom line here is that fixed losses will limit the savings that can be achieved relative to power and cooling. The fixed costs must be reduced in order to see the maximum savings from a data center consolidation project. There are three options for reducing fixed costs:

- Elimination—Elimination means removing some of the fixed-loss devices from the infrastructure. An update of your infrastructure support systems might make certain devices redundant or unnecessary; removing them from the data center takes the fixed losses of these devices out of the equation.
- Reduction—Reduction of fixed costs can be accomplished by switching to more efficient power delivery or backup systems. Looking for devices with reduced parasitic drain or more efficient power conversion to update existing systems will result in reduced fixed losses.
- Conversion—Conversion means taking some of the fixed-loss devices, such as cooling fans and pumping systems, and replacing the existing fixed speed fans and pumps with variable speed units, changing this portion of the overall power consumption from a fixed loss to a variable loss.

By reducing the power consumed by fixed losses within the data center, IT is able to maximize the cost savings benefits that the virtualization process enables. Remember that your losses are the aggregate of all the devices within the data center that are not supporting the IT load directly. Reducing and converting necessary supporting loads from fixed to proportional loads will improve the overall efficiency of the data center.

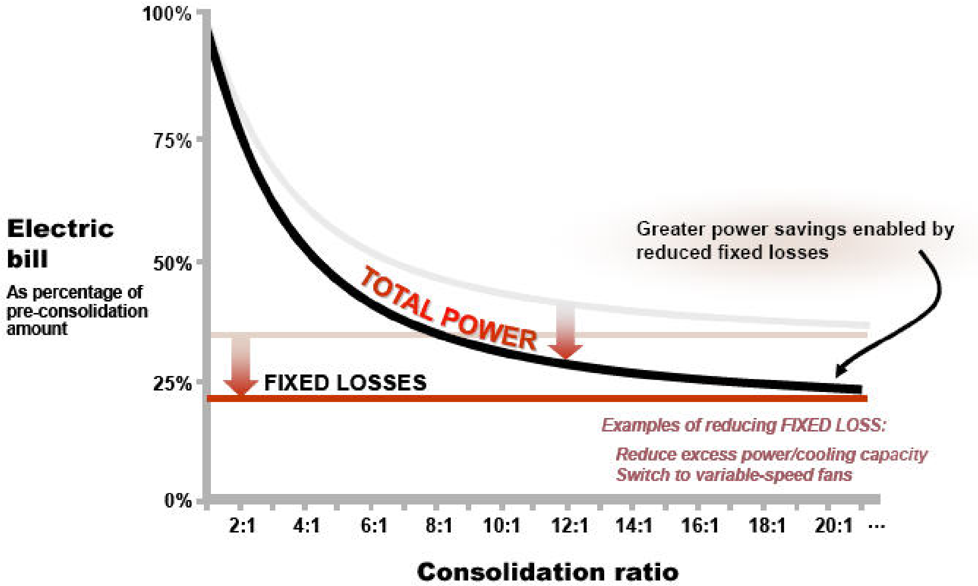

Figure 1.2: With the typical fixed loss averaging 35% pre-consolidation, reducing or removing fixed consumption devices and optimizing for the virtualized environment results in significant savings.

Remember that the move to virtualization and consolidation will always reduce the power consumption of the data center. Optimizing the energy efficiency of the data center requires changing the infrastructure, as necessary, to match the new load demands. Without making these changes, IT will simply be wasting power and reducing the overall data center efficiency.

Making these changes is rarely a simple matter of removing devices that are no longer necessary. Reconfiguring your power delivery and the design and layout of your data center as well as adjusting the cooling infrastructure are all just pieces of the puzzle. It is almost always safe to say that the solutions that were optimized for a standalone server environment will be oversized and inefficient for your new virtualized data center.

Addressing the Challenges

Many institutions that are building data centers using server consolidation and virtualization look at the project and realize that power and cooling that is sufficient for their existing environment should be good enough for their new environment, with its smaller number of physical servers doing the same work. The problem with this attitude is that “good enough” rarely is, and if IT is going to the trouble of building a new data infrastructure, it is reasonable to build an optimized power and cooling infrastructure to maximize the return on investment (ROI) on the investment in virtualization.

The proper way to approach these issues is to take the whole-system approach; power and cooling are simply another component of the complete data center model and need to be fully integrated as part of the system to be deployed. This approach requires laying out the data center IT hardware in such a way that the dynamic issues of the environment can be addressed, implementing a scalable power and cooling infrastructure, and deploying the software tools to mange power and cooling as an integral part of the entire data center system.

The Differences Between Physical and Virtualized Environments

The primary difference between physical and virtualized environments is characterized by the flexibility of the virtualized model versus the traditional static model. In the traditional data center with individual servers, the role of those servers is defined when the server is installed. An estimate can be made of its planned workload, and appropriate power and cooling can be configured for the physical location of the server. That process has served IT well, but the dynamic nature of virtualized environments has rendered it obsolete.

The virtualized environment can be changed at a moment's notice; servers can be deployed and provisioned as new virtual machines in minutes, running compute- and power-intensive applications immediately. The reverse is equally true, and as the load on the virtual hosts changes to meet the demands of IT, the load on the power and cooling infrastructure needs to change with it to provide the most optimized and efficient environment.

The very nature of a data center using virtualization to its maximum business benefit is one of change. Because of this, the power and cooling systems designed for the static model will be incredibly inefficient in providing services in the virtualized space. For example, the thermal characteristic of the data center server room was determined by the simple expedient of taking temperature readings in the operating data center. Once the hot spots were identified, the appropriate measures were taken to ensure that adequate cooling was available to the locations that needed it.

Sizing Your Virtualized Environment

With the virtual model, it's no longer that simple. Because the workload on the servers can shift dramatically, without any physical server changes, the thermal profile of the room will change with the dynamic software changes to the physical host servers. The hot spots in the server room will change as the workloads change, and the cooling system used to maintain the optimum temperature profile must change accordingly.

Thus, cooling needs to be provided only when and where it is needed, ideally by a solution that is able to sense the changes in environment and react appropriately. In addition to this dynamic response to load changes, the solution needs to have a short air path between cooling and workload.

This need has brought forth the row-based cooling model. In this design, cooling units are located within the rows of server racks and are capable of responding to temperature changes they detect. In this way, cooling is delivered when and where it is needed, and the amount of cooling provided is matched to the needs of the environment. Placing the computer room air conditioners or air handlers (CRAC) directly in the rows that host the physical servers meets this need.

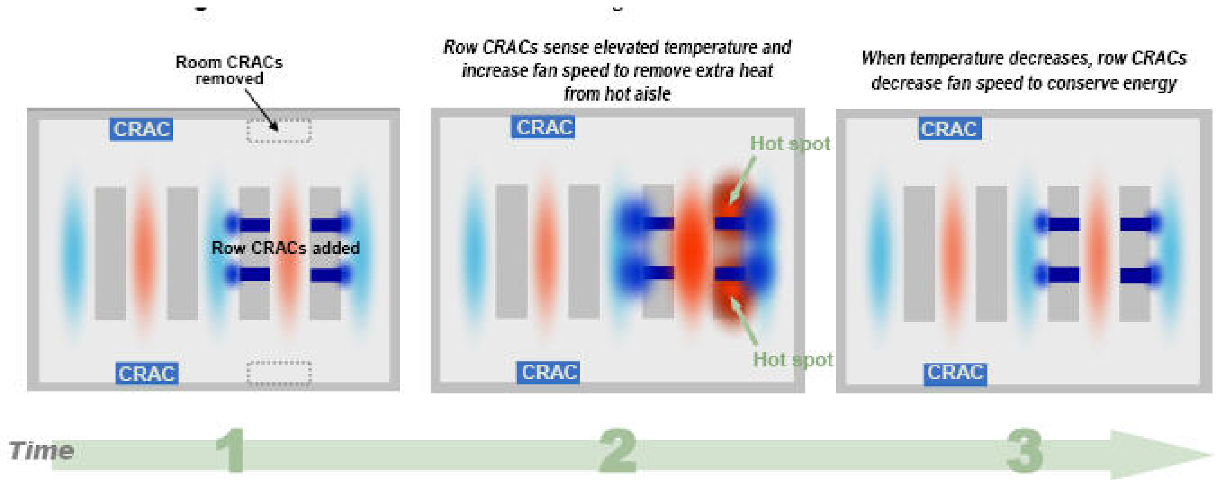

Figure 1.3: Row-based cooling systems are in the right place to most efficiently deal with hot spots.

This short air path cooling model brings many advantages and benefits to data center cooling efficiencies and is especially effective when adding virtual server hosts into existing data centers.

Cross Reference

We will cover this technology in greater depth in Chapter 2.

The implementation of a row-based cooling methodology provides an additional technique when deploying high-density servers into existing low-density data centers: the high-density zone. With this technique, an area of the existing data is designated as the "high-density zone" where a set of racks containing these high-density servers is effectively isolated from the rest of the data center.

This isolation may be achieved by using a curtain system to define its limits; it can also be done strictly with row-based cooling that is self-contained in the high-density zone and is effectively a neutral presence with no negative thermal impact on the rest of the data center. There is even a potential for positive impact, as additional cooling could be provided to devices outside the cooled rows. The high-density zone can be cooled and managed independently of the rest of the data center.

Matching Power and Cooling Delivery to the Virtualized Model

One of the most significant issues related to power and cooling when bringing virtualization to existing data centers is the under-loading of the power and cooling systems. Most data centers don't currently have a scalable power and cooling system. The primary selling point of such systems has been the ability to start small and grow with the data center. By doing so, a new data center didn't have to over-invest in their power and cooling infrastructure, purchasing equipment that might never be used.

Virtualization technology moving into existing data centers has reversed the problem. The goal in this situation is to be able to scale down to meet the new needs supporting the physical servers while retaining the ability to increase capacity to meet the growth needs of the organization. Matching need to demand provides the most efficient solution to the problem and avoids potential problems that can result from an over-sized solution.

Both issues focus on right-sizing the power and cooling requirements to meet the existing needs of the organization while retaining the flexibility to scale to whatever the data center demands become. Failure to act on this issue takes us back to the problem of inefficient use of the power and cooling infrastructure due to the power that is used regardless of the load—the fixed loss we discussed earlier.

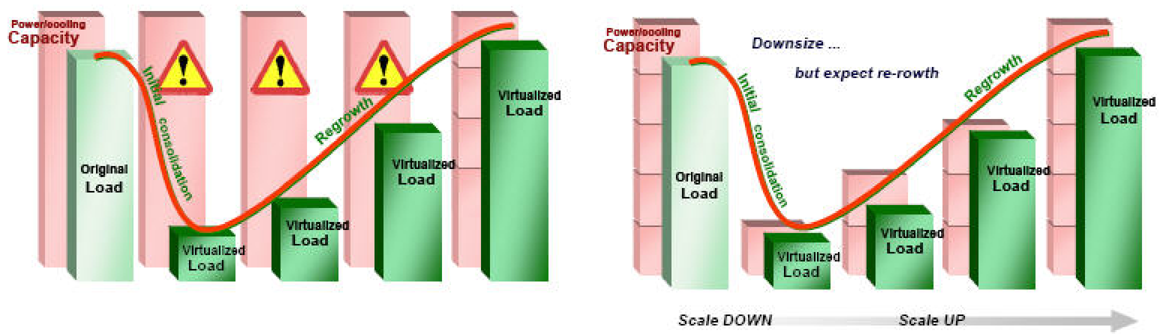

Figure 1.4: Comparing the efficiency of scalable power and cooling delivery to a legacy static over-sized system.

When looking at the two charts in Figure 1.4, you see that the overall capacity remains the same between the scalable and over-sized examples. However, with the scalable model shown on the right, you see that the unavoidable inefficiencies in the data center that cause the losses are effectively proportional to the amount of resources being used. The converse is evident when the capacity exists at the same maximum level found in the over-sized solution that results from leaving an existing data center power and cooling implementation in place after the move to virtualization.

Remember that we are looking at the data center as a single entity, with multiple interlocking systems. In addition, the power and cooling infrastructure is, itself, made up of many different underlying systems, all of which will be affected by being severely oversized, if that turns out to be the case after the move to virtual machines. Everything from cooling systems to backup generators can be negatively affected by being underutilized. Systems designed for providing power and cooling to the data center are designed for a specific operational range, and if the load drops sufficiently due to virtualization, it is possible to fall below the lowest recommended operational range on your cooling equipment—which can have results ranging from voided warranties to shortened compressor life due to repeated short-cycling.

It is important not to get caught up in the obvious when considering the effect of virtualization on your data center. You will see a reduction in overall electric costs. You won't see the increase in wasted energy because the percentage of power that is being used to drive the IT equipment significantly drops as the fixed losses take up more of the actual power and cooling expense.

Utilizing a scalable architecture gives you some measure of future-proofing your environment. If higher-density virtualization is where the data center goes, you will be in a position to effectively deliver power and cooling services to a smaller, more efficient data center. If the data center needs to scale out the physical server presence, your scalable solution is right there with it. Flexibility and agility is critical to the delivery of cost-effective IT services.

Assuring Availability

The arena of ensuring that power and cooling services are available falls into two categories: hardware and software. Hardware, such as backup generators, redundant cooling systems, and power protection systems, are not within the purview of this discussion. For this guide, we will be looking at the concepts of capacity management tools, which combine hardware instrumentation and software management to deliver the capability to provide dynamic management and capacity planning. Optimizing power and cooling delivery to the data center goes a long way to improving the delivery of IT services as it prevents potential problems with servers that are often the result of power and cooling problems.

Designing to Address Potential Availability Issues

There are three central challenges in capacity management that allow IT to address potential availability issues. Addressing these three issues gives IT the ability to quickly respond to changes in the data center and is crucial to getting optimal benefit in the virtualized data center. The need to know, and be able to manage, what is going on in real time is the driving force behind getting control of these three tasks:

- Power
- Cooling
- Physical Space

An effective capacity management solution will have the ability to monitor power and cooling at the row, rack, and server level and allow for the remote management of the power and cooling supplies in order to ensure the most efficient operation of the data center. If any of the three challenges aren't met, the solution can't be delivered. Avoiding this situation is important to an effective deployment, but the prevention of *stranded capacity* is also an aspect of this problem.

Stranded capacity is what you have when one or two of the three required resources are available but you have no way to deliver the remaining resources necessary for a successful deployment. These stranded resources are incredibly inefficient and can range from having over-sized cooling at a specific location, to having racks with power and space available but insufficient cooling capability to allow the servers to be used, to having power and cooling available but not having the space to utilize it. This issue is more common than you might think and one of the goals of capacity management is to prevent this situation from happening.

Anyone who has worked with network or IT systems management tools will find the concepts behind power and cooling capacity management tools familiar. "What if" analysis will play a large role in the selection, configuration, and administration of capacity management tools, especially with the use of virtual servers. By making use of data that is continually collected, your capacity management tool will be able to give IT the proper locations for adding physical servers—in terms of space, power, and cooling requirements—as well as define the impact of potential changes to the data center infrastructure.

A quality tool will also allow for input from IT regarding issues such as the need to group servers with other specific servers to meet business or networking needs and whether the servers will require additional support for features such as redundancy. In addition, a tool will help IT understand the difference in impact on the data center environment between virtual server hosts and standalone servers.

Planning isn't the only role of capacity management tools, however. The ongoing management of the delivery of power and cooling is central to the implementation of an efficient data center. Being able to respond to changes in the cooling requirements is only the start, albeit an important one.

By building heuristic data from the management of the existing environment, the management tool will be able to predict where potential problems can occur and suggest or perform corrective actions to prevent problems from impacting operations. These technologies are old hat in the network and systems management world. Not taking this opportunity to apply them to the data center infrastructure means that a chance to improve the overall efficiency of the data center has been wasted. And don't think of it only as a way to make the power and cooling delivery more effective; consider also the issue of reducing the need for IT manpower by reducing the amount of time IT personnel are physically trying to resolve the problems we are addressing here in the data center.

We've talked a lot about efficiency of the data center; capacity management is a critical tool for attaining maximum efficiency (on any measurement, such as PUE or DCiE). As we discussed earlier, the benefits of measuring the efficiency of the data center will appear on the bottom line as an improvement in the ROI of your data center projects.

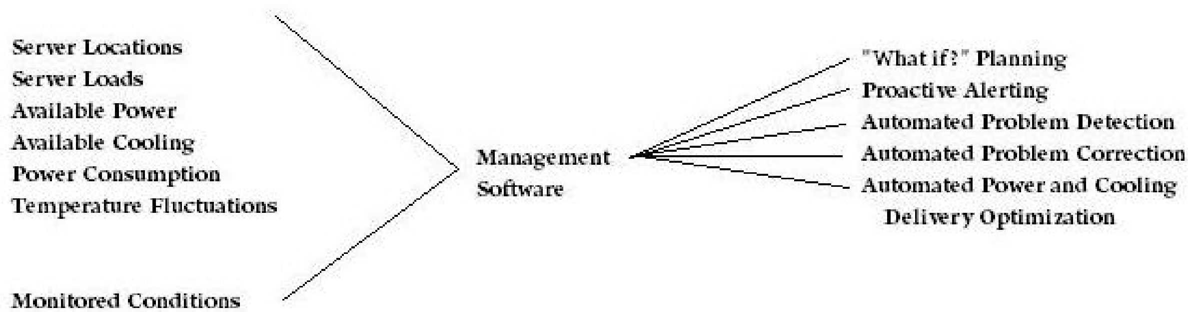

Figure 1.5: Effective capacity management software is able to concatenate a large amount of data to provide an optimized data center power and cooling environment.

Much, if not all, of the data that can be collected by the capacity management system is already being made available by devices in place in the data center. Actions are being taken manually, by IT personnel, to address problems as they occur, and the value of the information being provided is wasted due to the lack of a comprehensive proactive management system.

With the implementation of virtual servers in the data center, the adoption of an effective capacity management and planning tool is critical to maximizing the success of the new virtual infrastructure. We have already discussed some of the changes that virtualization will make in your data center power and cooling needs, so let's take a look at how the implementation and deployment of virtualization within the data center will affect the demands on the capacity management tools.

Top Virtualization Issues Affecting Availability

The very nature of the virtualized data center has an impact on the availability and delivery of power and cooling services. To fully utilize a data center with a combination of standalone servers and virtual machine hosts, capacity management tools need to be able to adapt to the changing conditions, especially those introduced by the virtualization deployment. The following list will aid the IT reader in identifying the issues that will most be encountered when optimizing cooling and power solutions for the virtualized environment.:

- Changing density—Compared with the state of the data center prior to the implementation of virtualization, there will often be an increase in density in one set of racks, where the servers acting as physical hosts to the virtual servers will be placed. Conversely, the overall density of the data center will go down due to the reduction in the number of physical servers. Thus, the thermal conditions of the overall data center will change and the power and cooling requirements will need to be reevaluated to meet the changing environment.
- Changing loads—The workload that was once spread out over a large number of servers has now been confined to a much smaller number. As server loads migrate from standalone servers to virtual machines, the dynamics of the data center will continually change.
- Increased rate of change—With one of the virtues of virtualization being the ability to quickly provision new servers and services, businesses will take advantage of this flexibility to test and adopt new server-based technologies that might gain them a competitive edge. In the past, the issues involved in adding a server to the data center would have acted as a brake on rapid changes. With virtualized environments, these changes no longer have the negative business impact they once had. Thus, business units will jump on that process, forcing IT to make sure all the support pieces, including the power and cooling infrastructure for the virtual server hosts, is in place to allow these rapid changes.
- Unforeseen changes—With virtualization, changes can be made to existing servers that significantly impact the cooling needs of those servers, and as the resulting hot spot is created, it will impact cooling on servers in the same physical area. As these changes can be made without any physical access to the server, IT needs reliable tools that monitor for hot spots and have the ability to dynamically adjust for these changing conditions.

- Interdependencies—The data center is a complex organism. Changes made to the behavior of the data center—such as moving servers, virtualization, changing cooling patterns, and modifying space requirements—can have unexpected impacts due to the close relationships between the component entities that comprise the data center.
- Lean provisioning—The move to a lean provisioning model for power and cooling allows maximum optimization of the data center infrastructure. However, unexpected changes can have unexpected results in the behavior of the power and cooling needs of the data center.

Planning to address these issues as part of the data center migration to virtualized environments will give IT a hand up in implementing and deploying next-generation technologies.

Summary

Virtualization is coming soon to your data center, if it isn't already there. And as the drive for server consolidation continues to increase, the performance of your data center in supporting high-density, variable-load systems will become critical.

The optimal way to address these needs is to take the whole-system approach to building, or re-building, your data center power and cooling infrastructure. These issues need to be addressed while the changes in the data center are being planned, with the impact that virtualization will have on the data center infrastructure being considered as part of the implementation and deployment process. With this in mind, building a highly-optimized and electrically-efficient data center will not be an afterthought; the design and capabilities of the infrastructure can be tailored directly to the business needs of the organization, resulting in the most cost-effective solution possible.