PART V

The Greening Process

CHAPTER 11
Datacenter Design and Redesign
CHAPTER 12
Virtualization
CHAPTER 13
Greening Your Information Systems
CHAPTER 14
Staying Green

CHAPTER 11

Datacenter Design and Redesign

Chances are you aren't sitting in an empty building, thinking about installing computers, and wondering, "If I do install computers, how can make the process as ecologically friendly as possible?" The reality is that there's probably already a computer on every desk, a datacenter humming away, and a throbbing electrical bill.

In this chapter we'll talk about issues that pertain to redesigning your datacenter. When it is time to replace equipment, how you can downsize your datacenter and replace equipment with energy efficiency in mind.

We'll talk about power consumption, how you can tweak your datacenter with that in mind, along with designing for optimal cooling. We'll also talk about the hardware in your datacenter and how that can be optimized.

Energy Consumption

Every 5 years computing capacity doubles. And when capacity doubles, so does everything that goes with it. For instance, between 2000 and 2005 the amount of energy consumed by datacenters and other infrastructure in the United States also doubled.

What accounts for this gorging? It's going on both in the datacenter and in the cubicle farm. Users are taking advantage of everything the World Wide Web has to offer: YouTube, music downloads, Internet telephony, and whatever else is popular this week. But in the server room, power consumption has spiked because of other trends, including the bloat of low-end servers that cost less than US$25,000.

Also, many IT managers are deciding to move to Linux and other free operating systems, rather than pay a charge for each server. As such, there's even more demand for a large number of low-end servers.

Growth

As Figure 11-1 shows, in U.S. datacenters in 2000 about 5.6 million servers were installed. Of those, 4.9 million were low-end, 663,000 were midrange, and 23,000 were high-end servers, according to a survey by Lawrence Berkeley National Labs in 2007.

By 2005, U.S. datacenters had 10.3 million servers. Of those, 9.9 million were low-end, 387,000 were midrange, and 22,200 were high-end servers.

The same study showed that it's not just the U.S. that spiked in its server growth. Around the world server demand expanded. In 2000, as Figure 11-2 shows, there were 14.1 million servers. Of those, 12.2 million were low-end servers, 1.8 million were midrange, and 66,000 were high-end servers.

By 2005, that number had swelled to 27.3 million servers. Of those, 26 million were low-end servers, 1.2 million were mid-range servers, and 66,000 were high-end servers. To feed that growth, in the United States, the amount of power necessary was equivalent to about five 1000-megawatt power plants. Worldwide, it was the same as 14 power plants.

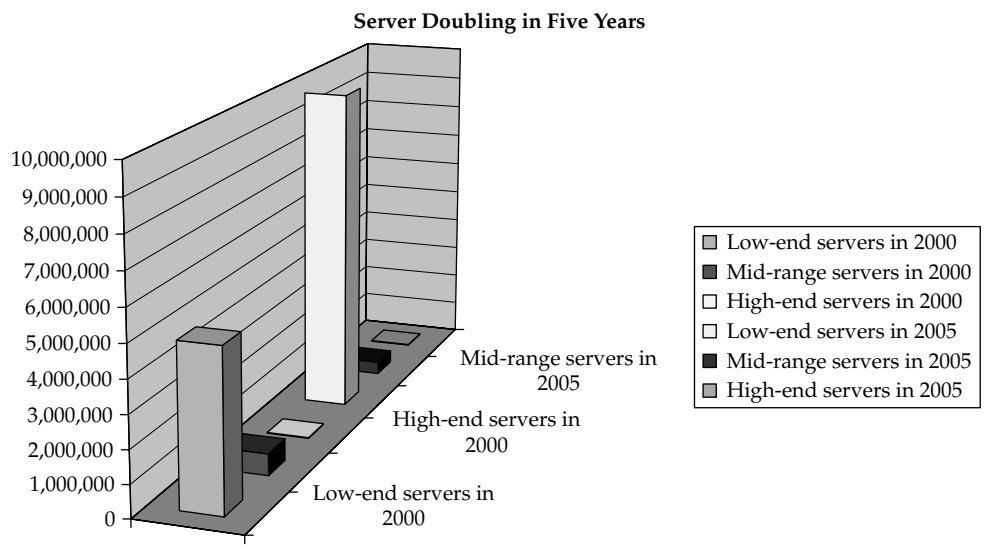

Figure 11-1 The total number of American servers nearly doubled between 2000 and 2005—and so did power consumption.

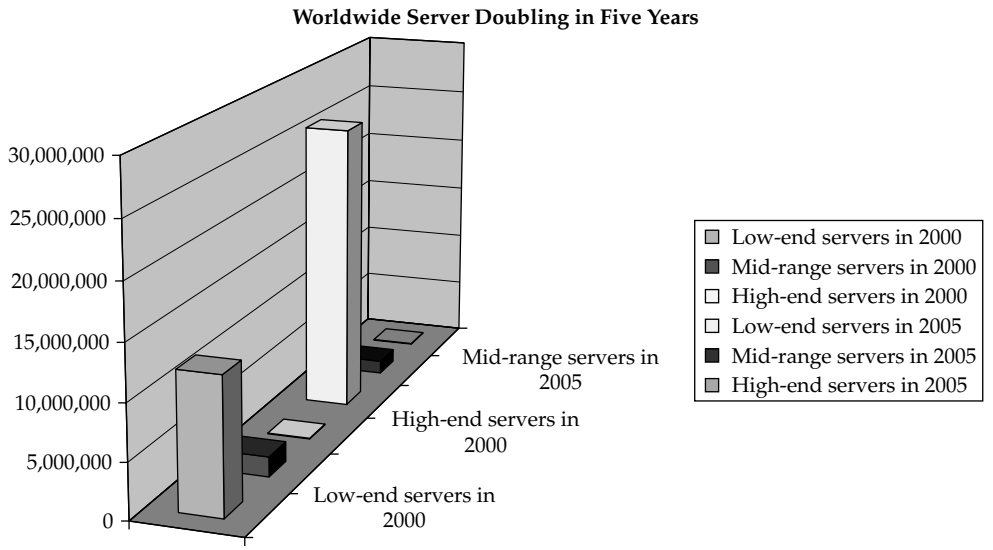

Figure 11-2 Around the world, server usage doubled.

Other Costs

Naturally, when we talk about energy consumption, we tend to immediately think of the power meter and just how fast it's spinning. But you might also want to think of your organization holistically. That is, how a change in one area can help another area.

For example, Isothermal Systems Research (ISR) in Liberty Lake, Washington offers a technology called SprayCool. SprayCool uses a fluid called Fluoinert, which is nonconductive. The coolant is sprayed directly on a chip or processor (the procedure is called "chip-based" cooling), and then cools the chip as the coolant is converted into a gas. This and other cooling technologies for the datacenter can be found at www.spraycool.com and at Link 11-1.

As Figure 11-3 shows, that gas is then taken down to a heat exchanger at the bottom of the unit, where the heat is transferred to another fluid, such as water. Because the water has been heated up, it can be used to regulate the building's temperature.

In a system with 10 racks of equipment, there is a 150-kW load just for cooling. Using the spray coolant removes a quarter of that cooling load.

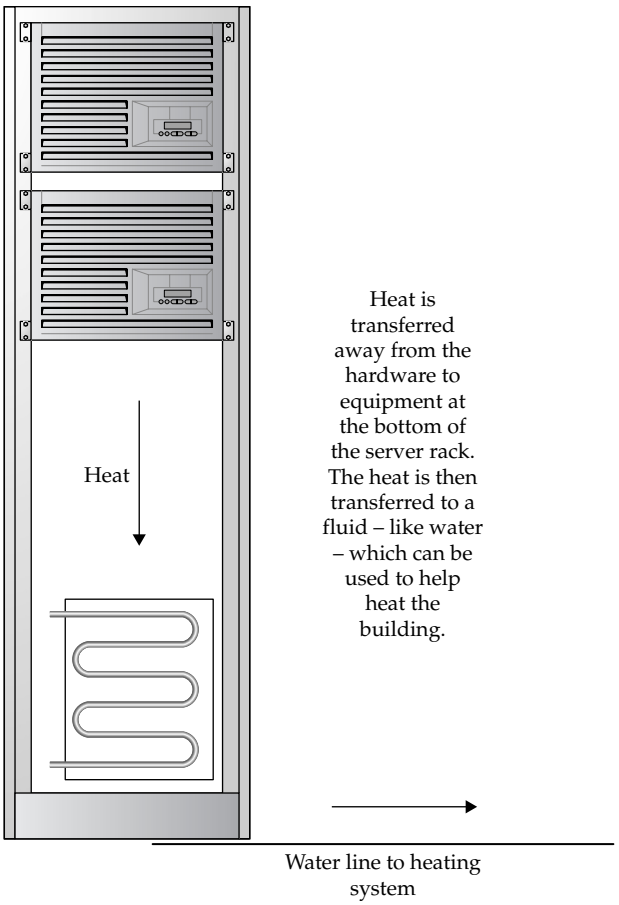

Figure 11-3 The heat can be taken from servers and used to help heat the facility.

Design

When you design or redesign your datacenter with green in mind, approach it from the perspective of rightsizing the physical infrastructure. Use efficient devices, and design with energy efficiency in mind. The following sections will examine how you can make the most efficient decisions and plans for your datacenter.

Efficiency

You can ensure better efficiency in your datacenter by looking at some common sources of problems:

- Power distribution units operating well below their full load capacities
- Air conditioners forced to consume extra power to drive air at high pressures over long distances

- N+1 or 2N redundant designs, which result in underutilization of components
- Oversized UPSs to avoid operating near capacity limits
- Decreased efficiency of UPS equipment when run at low loads
- Under-floor blockages adding to inefficiency by forcing cooling devices to work harder

Floor Layout

As we discussed in Chapter 4, the way you lay out your datacenter has a huge impact on the efficiency of the air conditioning system. If airflow is blocked, it is not as efficient. If a hot-aisle/cool-aisle layout is not adhered to, you lose efficiency. When you design your floor layout, you want to do it with hot air and cold air segregation in mind.

Also, be mindful of where cooling systems overlap in your datacenter. This is the ideal location for equipment that runs especially hot or is especially important.

Server Configuration

It's not enough just to get the servers up and running; be sure that you enable any power-saving features on the servers. A lot of times IT managers don't set power-saving modes, rather opting for high availability. Setting power-saving modes will help save money by reducing power consumption and cooling costs.

Floor Vent Tiles

Vented floor tiles don't just look cool, they're designed to facilitate airflow. In many datacenters, however, vented tiles are either incorrectly placed or an insufficient or excessive number of vented tiles are installed.

If you use Computational Fluid Dynamics (CFD)—an analysis of airflow patterns—in your datacenter, you can optimize datacenter cool-airflow. This allows you to fine-tune the center's cooling by placing vented floor tiles in their optimal locations. Some vendors will even help you with CFD. By simply fine-tuning tile locations, some datacenters have achieved a 25 percent reduction in cooling costs.

Rightsizing

Datacenter design and redesign involves a lot of components. For instance, a seemingly small task such as the placement of vented floor tiles can lead to big savings. However, the biggest thing you can do to optimize your datacenter is to design with rightsizing in mind. That simply means building for what you need.

Rightsizing has the most impact on your datacenter's power consumption. Fixed losses in your power and cooling systems are present whether or not you have a datacenter in place. Organizations that have light IT loads might not even see an impact by its IT department. However, as the IT load gets larger, organizations will see power and cooling

expenses climbing. Rightsizing can be accomplished through server consolidation and virtualization.

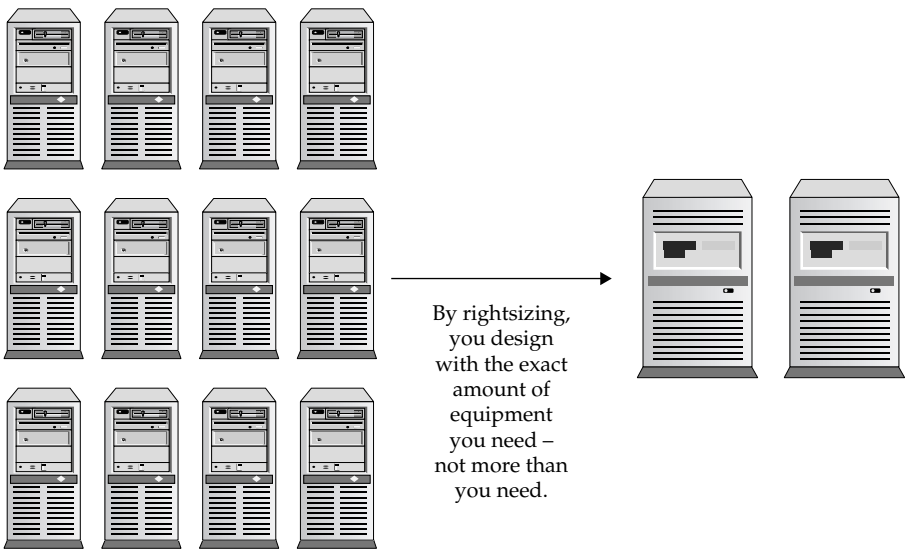

NOTE *We talk about consolidation later in this chapter and show you how to do it in Chapter 12.*

When you redesign with rightsizing in mind, you must also be cognizant of the fact you will need additional capacity in the years to come. However, if you spend the time and money you need to predict the datacenter power and cooling load, it will pay for itself in both reduced capital and operational expense.

Upgrading to Energy-Efficient Servers

When you plan to green your datacenter, it would be nice if you could do it in one fell swoop. But, realistically, that's not going to happen. You can't really take down your entire system and replace it from scratch.

It won't work for a couple reasons. First, the cost would be too prohibitive. The budget just probably won't be there for so much expensive equipment. Second, there's the issue of workflow. Can anyone really envision being dark for the amount of time it would take to turn everything off and install new equipment? In most cases, when a new server is installed, it has to be hot-swapped in, so that there isn't even a fraction of a second of downtime.

So what you'll wind up doing is replacing your equipment with green alternatives as you would normally replace servers and storage units. When you do design your green datacenter, there are some issues to keep in mind and some practices you should follow.

Consolidation

The first place to start with any datacenter greening project is to consolidate servers. In some cases, you might find servers that you don't even need and you can just turn them off. Between 10 and 30 percent of servers aren't even being used and could be deactivated.

NOTE *If you take one physical server out of service, you'll save about US$560 in yearly electrical costs (that's assuming 8 cents per kWh).*

As soon as unused servers have been taken out, the next step is to consider consolidation. Many servers are underused, in terms of their processing power, as Figure 11-4 shows. Most physical servers run at a maximum of 15 percent.

Move as many server-based applications as you can to a virtual machine. A virtual machine is simply a software server installed on a physical server. You can put a number of virtual machines onto a physical server. This will result in a sharp drop in the number of physical servers in your datacenter. The servers that take on the additional roles will be more efficient in terms of their utilization levels.

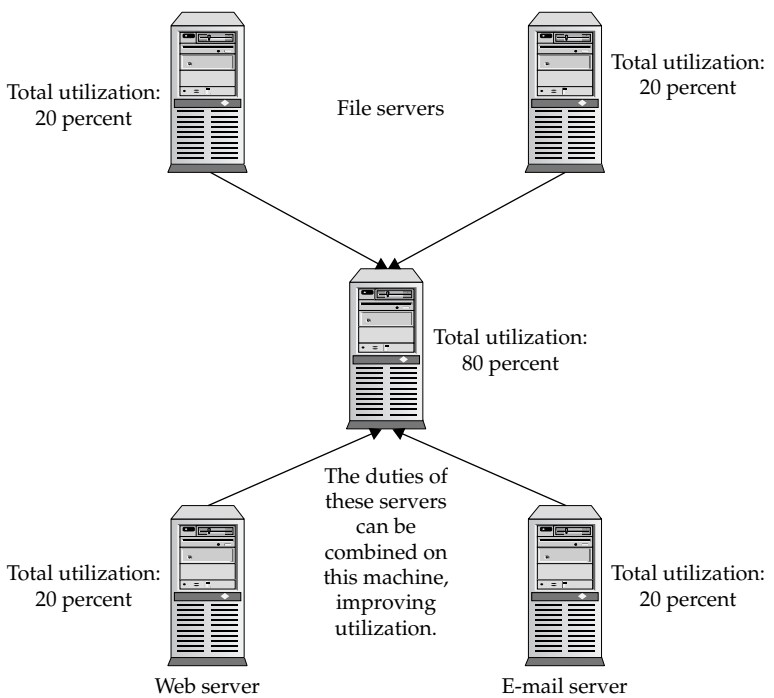

FIGURE 11-4 By consolidating servers, you reduce the amount of servers you're paying to power, plus you are increasing the utilization of the tasked machines.

New Replacement Policies

The biggest place where you will see cost savings, power reduction, and cooling reduction is in server consolidation. It's a no-brainer—the less hardware you have, the less consumption will occur. When you replace servers, keep these guidelines in mind:

- Replace two or more old servers with a two-way server or dual-core, single-processor server.
- Replace old servers with a blade based on a low-voltage or mid-voltage processor.
- Replace dual-processor servers with a single, dual-core processor.
- Replace four-way servers with two-way, dual-core servers.

Use Power Management Features

It seems obvious, but some IT managers don't get it: Use your equipment's power management tools. Although electricity usage is reasonably constant, the IT load swings up and down according to the needs of the system. It can be as much as a factor of three.

By using your equipment's power management functions and turning off unused servers, you can cut datacenter energy usage by 20 percent. But it's not because IT managers are dimwits—rather, they're focusing on uptime and performance. In reality, however, turning on power management features can improve overall reliability and uptime because it reduces the stress on datacenter power and cooling systems.

NOTE *Experiment with your system, however, before turning on any power management features. In some systems running Linux, for example, waking up an idle server can cause unexpected behavior.*

Get Energy-Efficient Servers

As you upgrade to new equipment, look for energy-efficient servers. They're using less power than before, and they'll probably only get more efficient as new products are introduced.

The first generation of multicore chips resulted in huge reductions in power consumption. For instance, the Intel Xeon 5100 delivered twice as much performance with a 40 percent drop in power usage.

Use Energy-Efficient Power Supplies

Probably the biggest wastes of power in the datacenter are the power supplies. Inefficient units tend to ship with servers, and they aren't really selected based on how the server is configured.

Inefficient power supplies can waste almost half of the power before it gets to the server. Further, every wasted watt is turned into heat, which then requires at least another watt of power to cool.

Look for power supplies that achieve 80 percent or higher efficiency, even at 20 percent load. They cost more, as much of this type of equipment does, but they cost less to use. We'll talk about power supplies later in this chapter.

Talk to Your Facilities Manager

The IT department might keep scrupulous records and track the performance of its equipment, but most IT managers don't see the electrical bill. The facilities group does.

In order to truly understand how much power you're using (and costing the company), pick up the phone and punch in the facilities manager's extension. You might make his day. At the very least, open some communications with the facilities manager so that he knows that you are trying to address the power usage problem.

NOTE *This disconnect is also affecting how vendors develop equipment. Often, vendors just work on their one piece of equipment and don't worry about how their power needs will affect the rest of the datacenter.*

Refer to the Standards

If you're not sure which equipment to buy and you don't have the time, patience, or energy to sit down with product brochures and data sheets, you can still make wise procurement decisions.

A number of standards are out there that you can refer to. For instance, the 80 Plus certification program was initiated by electric utilities, and it names power supplies that achieve 80 percent efficiency at load levels of 20, 50, and 100 percent.

Energy Star is most commonly affiliated with desktop machines, but there is an effort underway to extend certifications to servers. Certainly, if you purchase a server in Europe, it will have to be RoHS and IEEE compliant. It's likely that vendors would tout that feature in other parts of the world, should their servers be so certified.

Ask for It

Remember 15 years ago when fat-free food was all the rage? To make it taste better it was jacked up with extra sugar. Remember 5 years ago when the Atkins diet made it forbidden to eat carbs, but you could eat steak wrapped in bacon and smothered with cheese? In each case, industry responded to demand. Fifteen years ago store shelves were stocked with fat-free foods. Five years ago store shelves were stocked with low-carb food.

Industry only responded because customers wanted it. It's the same way for Green IT equipment. If you don't ask your vendors for it, they won't see the need to make it.

Server Consolidation

Chances are you aren't going to shut down your datacenter for a weekend, wheel out the old servers, and install brand new, more powerful ones. It's more likely that you're going to bring in new equipment as your replacement plan dictates. However, when you do start replacing equipment, it's a smart idea to consolidate your servers.

By consolidating servers you can:

- Increase utilization of existing hardware from 10 to 15 percent, up to 80 percent
- Reduce servers at a 10-to-1 ratio
- Reduce hardware and operating costs by as much as 50 percent

Traditionally, servers were approached with a "one workload, one box" philosophy. That is, if you needed an e-mail server, you bought a new server. If you needed a file server, you bought a new server. However, each server is not used often enough to truly justify its location on a single machine. As mentioned earlier, most servers only operate at about 10 to 15 percent of their total load capacity.

The result of having a server for each network application is server sprawl. In this section, we'll talk about the merits of consolidating your servers into fewer machines.

Utilization

Server sprawl can be eliminated by consolidating and virtualizing your physical machines into virtual machines. Applications such as VMware run independently from the underlying hardware and are supported on a range of physical servers.

And if you have different operating systems for your different applications, that's not a problem. A virtual machine represents a complete system—processors, memory, networking, storage, and BIOS. This allows you to run Windows, Linux, Solaris, and NetWare operating systems and applications on the same server. This is illustrated in Figure 11-5.

Normally, a server running a lone application experiences total utilization of 15 percent, tops. That means all that processing power you paid for is being wasted—and all the power you're paying the electric company for to run that machine is also being wasted.

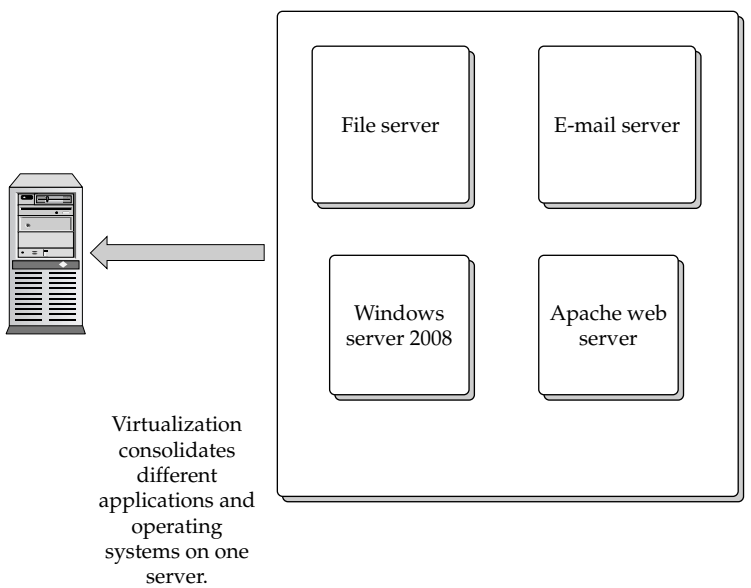

Figure 11-5 Consolidation and virtualization allows you to run multiple virtual servers on one machine, and each server can be a different operating system.

By running multiple virtual servers on one physical server, you're increasing your utilization of that machine from 15 percent up to 80 percent.

Hardware Reduction

Consolidation doesn't just mean that one server can do double-duty. It's more like a 10-to-1 ratio. For example, the VMware ESX Server can support more than 100 virtual machines.

You can also use the same strategy on storage and network consolidation, thus reducing more equipment from your server room.

Reducing Operating Costs

According to a study performed by Lawrence Berkeley National Laboratory, the servers in a datacenter account for about 55 percent of the electricity costs. The remaining power is spent to support that equipment—Lawrence Berkeley National Laboratory, 2007. So when you consolidate your equipment, you also need fewer pieces of equipment to play support roles, so there's more cost savings there.

Consolidating can save you about US$560 per year per server. So let's use the simple math example of consolidating 10 servers into 1. Consolidating 10 servers down to 1 results in a US$5040 savings per year. And that's at a smallish company. If you have thousands of servers, you'll recognize some big savings.

Consolidation also cuts down on the amount of heat generated in your datacenter. But, like power consumption, it isn't just the servers that generate the heat, it's also the equipment supporting the servers that add to the heat generation. Figure 11-6 shows this.

Repurposing Servers

As part of your green datacenter design process, you should also look at the equipment you have and evaluate its worth. You might have a systems life cycle that dictates that

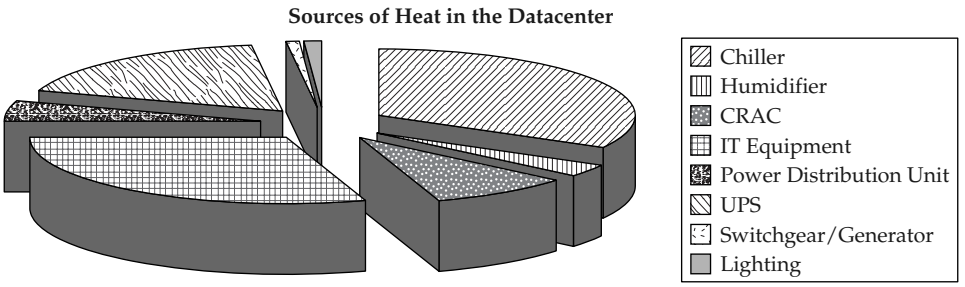

Figure 11-6 Servers generate a lot of heat, but it's the equipment supporting them that adds up to create more heat.

the equipment needs to be phased out and replaced with new stuff. However, you might be able to save a few thousand dollars by simply repurposing the equipment.

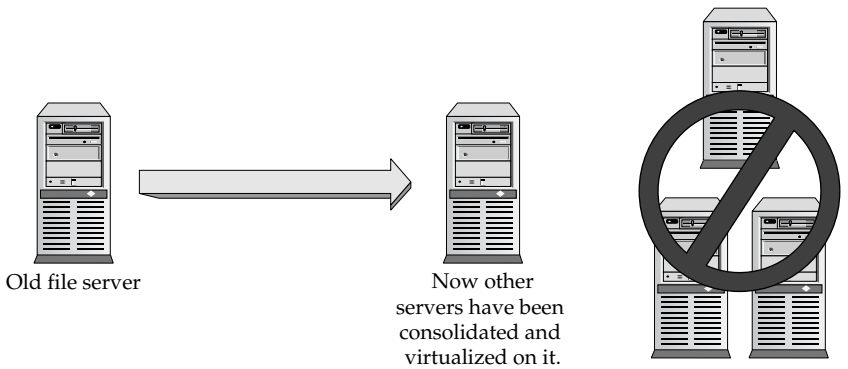

For instance, if you have a server that has sufficient processing power and memory, why not keep it, but consolidate and virtualize the duties of other servers? If it is equipment you already own, by repurposing the server you don't need to buy a new one (thus saving you the expense of buying a new one and recycling the old one), plus you maximize the server's utilization.

Cabling Considerations

Think of datacenter design as a Hollywood movie. The stars are the placement of servers and overall cooling. But the supporting actors in datacenter design are the small things you might not think about, such as cabling. True, cabling doesn't generate heat, and there are no toxic emissions that come out of your cabling. However, you need to make certain considerations when redesigning the cabling placement in your datacenter.

Largely, cooling is affected by your cabling choices. For the most efficient cooling, cabling has to be properly designed, remediated, and routed so that air flows optimally.

TIA-942

The Telecommunications Industry Association (TIA)—the same people who came up with CAT standards for unshielded twisted pair cabling—have published a definitive document encompassing best practices and design considerations for the modern datacenter.

The standard, titled "Telecommunications Infrastructure Standard for Data Centers, TIA-942," is 148 pages long and covers all sorts of issues, from site selection to rack-mounting methods.

TIA-942 is a best-practices document that aims to standardize a lot of requirements throughout a datacenter. One of the upshots is that by standardizing requirements for contractors, companies can save on one of IT's biggest budget items—IT management salaries.

You can order a copy of the standard from TIA at www.tiaonline.org (Link 11-2). They're not giving it away, though. The standard costs US$278, and you can buy either paper or PDF format.

Cable Routing

The key consideration in cabling your datacenter is keeping airflow in mind. TIA-942 and other datacenter standards advise horizontal and vertical cabling to be run accommodating growth so that you don't have to revisit cabling when your datacenter does grow (and it will). There are a number of reasons to do this:

- Eliminating the adverse affects of removing floor tiles and decreasing static pressure under raised floors during moves, adds, and changes (MAC) work
- Ensuring that pathways allow the flow of cold air in cold aisles to be unobstructed by cabling
- Installing cabling to provide a baffle of sorts, channeling cool air into cold aisles

Unless the issue of cabling is addressed every time a datacenter is changed, cabling can add to your cooling woes. A lot of older datacenters suffer because abandoned cabling channels are left behind, creating an air dam that can obstruct airflow.

In everyday use, cabling in and of itself isn't toxic, but given extraordinary circumstances it can be. Old cabling jackets might not meet current RoHS requirements. Older cabling also carries a significant fuel load, which can pose fire threats and can release toxins—such as halogens—if it catches fire.

Safety issues are important, but if you're removing old cabling, you also have to consider the disposal and recycling of cabling. Like with other IT equipment, you can't just chuck the cabling into the dumpster; it must be disposed of properly.

When you design your datacenter, you can save future problems and reduce the amount of abandoned channels through proper management. Infrastructure management systems (such as MapIT) can provide a detailed monitoring of any moves or changes in your datacenter. This gives you an up-to-date diagram of the physical layer connections so that channels can be managed and fully utilized before they become a problem. More information on MapIT and infrastructure management systems can be found at Link 11-3.

Bigger Bandwidth

When designing or redesigning a datacenter, there's a balance to be struck between what you need and what you can spend. If cost were no issue, you'd get the biggest, fastest equipment you could afford. However, there is a financial reality you need to keep in mind. That said, it is in your best interest to get the highest capacity equipment you can afford.

Currently, CAT 7/Class F cabling is the highest performing cabling on the market. CAT 7A/Class FA is due soon, offering 1GHz per channel. This provides a high amount of bandwidth above the latest 10Gbps speeds for copper. The new cabling is backward compatible with older technology.

But what if you don't need CAT 7 cabling right now? Well, you will. And when you do, you'll have to take out the CAT 5 or CAT 6 cabling you've got and rewire with CAT 7. If you install the highest performing cable now, you will pay more upfront. However, it will cost significantly more if you have to keep upgrading every time your equipment gets a performance boost.

For instance, the installation of a CAT 5e system would need replacing in a few years as 10GBASE-T is implemented to the desktop. Once your organization outgrows CAT 6, you'll have to replace it with CAT 7, and so forth.

Installing high-performance cable not only saves you money, but copper, aluminum, and other natural resources are conserved. Further, the discarded cable doesn't have to be recycled or disposed of, thus saving some environmental impact.

Other Savings

Although optimization of airflow with proper cable planning and routing will save energy, there are other energy saving opportunities within the cable infrastructure.

For example, if you move to a higher performing class of cabling, you will reduce noise on the cabling channel, thus resulting in a significant power savings in the equipment by eliminating Digital Signal Processing (DSP) complexity used to reduce noise levels. In fact, using fully shielded cabling could save 20 percent in the overall power budget in 10BASE-T deployments.

Alien crosstalk (shown in Figure 11-7) happens in 10GBASE-T deployments when signals from neighboring cables bleed into other cables.

LEED Considerations

If you are concerned with LEED, good cabling practices can help propel you toward your certification goal. Buildings registered under the LEED system are awarded "credits" for reaching certain environmental goals, and are assigned a rating within the LEED system based on the following scale:

- **Platinum** 52–69 credits
- **Gold** 39–51 credits
- **Silver** 33–38 credits
- **Certified** 26–32 credits

Table 11-1 lists some design considerations, what they accomplish, and which LEED credits they can help you earn.

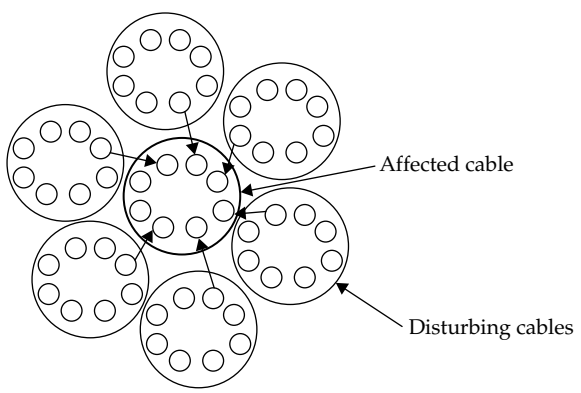

Figure 11-7 Alien crosstalk occurs when signals bleed from one cable to another.

Design Consideration	LEED Credit	Explanation
Using intelligent infrastructure management software	**MR 2.1** Construction Waste Management (50 percent)	Reduction of unnecessary channels due to undocumented or poorly managed movement, addition, or change work.
	MR 2.2 Construction Waste Management (75 percent)	
	MR 3.1 Resource Reuse (5 percent)	Identification and utilization of unused cabling channels to limit installation of new channels.
	MR 3.2 Resource Reuse (10 percent)	
	EA 1 Optimize Energy Performance	Maximization of active port usage to limit the installation of unnecessary active equipment. Identification and utilization or elimination of abandoned channels to maximize pathway space or increase airflow for energy-efficient cooling.
CAT 7 cabling and cable sharing	**MR 2.1** Construction Waste Management (50 percent)	Cable sharing as a means to reduce the number of installed cabling channels. Future-proof performance extends the life cycle of the cabling, decreasing the frequency of cable removal/disposal and installation of additional cabling.
	MR 2.2 Construction Waste Management (75 percent)	
	EA 1 Optimize Energy Performance	Shielded construction may limit noise sufficiently to reduce active equipment power consumption through elimination of DSP.
Using trunking cable	**MR 2.1** Construction Waste Management (50 percent)	Factory termination eliminates onsite waste created by field terminations.
	MR 2.2 Construction Waste Management (75 percent)	Efficient installation of trunk cables requires fewer contractor visits and smaller crews.
	MR 3.1 Resource Reuse (5 percent)	Modular design of trunks allows for onsite reuse.
	MR 3.2 Resource Reuse (10 percent)	
	EA 1 Optimize Energy Performance	Well-organized channels eliminate air dams in pathways caused by poorly managed individual channels to maximize airflow for energy-efficient cooling.

Table 11-1 Some Cabling Considerations Can Help Earn LEED Certification

Power Supplies

A couple of years ago, Google decided to build its new facility on the banks of the Columbia River in The Dalles, Oregon. The facility is two football fields in size with two cooling towers extending four stories into the sky.

But Google didn't move to this location on the Oregon/Washington border to get closer to Microsoft. In fact, Microsoft and Yahoo! are following suit and have plans to build their own massive datacenters upstream in Wenatchee and Quincy, Washington, 130 miles to the north.

Is there something gloriously beautiful about the area? Actually there is—the flowing rivers are harnessed to feed the power that these power-hungry datacenters need. It is the cheapest power in the country. This location underscores the needs of datacenters—they need lots of power.

Although you may not be able to uproot everything and move your organization next to a river, you still need to manage the power in your datacenter. That is a key consideration when designing or redesigning.

Too Much Power

When systems engineers design power supply requirements for enterprise servers, they approach it the same way a civil engineer designs a tunnel—they build it for the largest possible traffic that will ever go into it. Even though the vast majority of the traffic in that tunnel will be commuter vehicles, they still have to make sure it will be large enough to accommodate semis and tanker trucks.

It's the same way for servers. It's common to specify power supply requirements based on the maximum system configuration and load requirements. But systems engineers don't need to hit the gas so hard. They can specify different power requirements for different configurations. Naturally, this would cost manufacturers more money, so they tend to just put the biggest power supply needed in all of a line's models, regardless of how they will be used.

For instance, consider the servers in Figure 11-8. Both are from the same line—WasteTech's Escalenté line. The WasteTech Escalenté 5000SUX, shown on the left, is configured with enough hardware to only require 60 W of power. The 6000SUX, on the other hand, is fully maxed out and contains enough hardware to consume 540 W. Both servers use the same power supply. Although the power supply feeds the 6000SUX model with enough power, it provides more than enough power for the smaller server. So much so that it wastes energy.

But the issue isn't limited to servers in the same line. Many times manufacturers select a power supply that can fit all of their equipment. Again, cost is the issue. It's cheaper for the manufacturer to use the same supply in all their equipment—you're the one who ultimately pays for power inefficiency. Further, the planet pays a price, because of the carbon emissions spent to inefficiently power that device.

Thankfully, many companies such as Dell offer different power supply configurations so that the customer can choose the right size for the job a particular server is going to do. This can lower power consumption significantly, especially when factored over the life of the server.

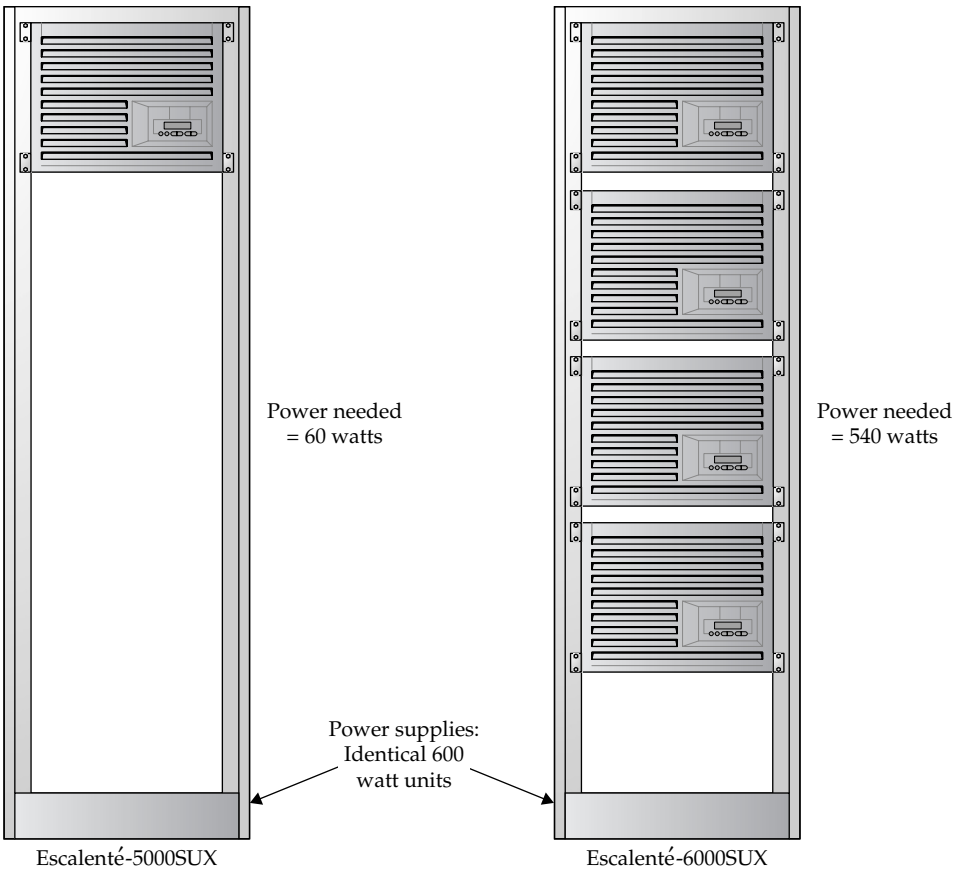

Figure 11-8 Servers with different power needs are often fitted with the same power supply, leading to inefficiency.

Efficiency

A power supply's efficiency is calculated by taking the DC output of the power supply and dividing it by the AC input. If a power supply draws less AC to produce more DC, then the power supply is more efficient than one that uses more AC power than the DC power it generates. Here's the formula:

Efficiency = (DC output)/(AC input)

For example, if we have a power supply that's putting out 200 W and it's using 300 W from the wall, simply divide 200 by 300:

200 / 300 = 66 percent

The higher the efficiency, the better. In this case, 66 percent isn't so great. Typically, the efficiency of a quality power supply is between 75 and 85 percent. Using our calculation,

the missing 34 percent doesn't just disappear. It turns into heat—that same heat you have to pay to cool off.

When power supply efficiency and system cooling are taken into account, the actual useful work done by the computer can be less than 50 percent of the total power draw of the system. That means over half of the power consumed is simply a drain on the ROI, an increased operating cost with no benefit to the client.

There are various power supply calculators available for different machine types and uses. At Link 11-4, we have provided you with a list of several that you may find useful.

Loads

In terms of efficiency, server power supplies are the least efficient when they are just idling. That doesn't mean they are using more power, they're just using power inefficiently. It's like an automobile. When you're parked with the engine running, you're still consuming gas, but not going anywhere. When you pull out of the driveway, you're finally going somewhere, but now wasted energy is a smaller percentage than it was when you were just idling. Now, consider a server that is configured minimally, but with a large power supply. The power supply is churning out a lot of unused power.

Manufacturers are trying to address this issue. However, as you've seen before, there will be more upfront costs, especially in the form of more expensive circuit materials.

Redundancy

Quite often, datacenters are configured with redundant power supplies. The reason is obvious—if one power supply fails, the other is there to keep the server online. In fact, in large datacenters, the power supplies are on separate power grids, so if the AC power goes out on one grid, the other supplies power and the servers keep working.

Redundancy is great for the sake of ensuring uptime, but it's tough on energy efficiency. For example, if a server needs 200 W to operate and it has a single 800 W power supply, then that server is using just 25 percent of the supply's total capacity. If a redundant power supply is added, power use is split between the two supplies and the draw on each is only 100 W. The efficiency then drops from 83 percent for a 25 percent load to 65 percent for a 12.5 percent load.

As with many issues in the datacenter, this comes down to a balancing act and some negotiations. How important is system availability versus reduced energy usage (and cost)?

The biggest thing you should take from this chapter is the importance of consolidating and virtualizing your servers. The technology now exists to run multiple servers virtually. In the next chapter we'll talk about server consolidation and show you how you can do it.