InformationWeek
analytics

December 2008

How To Ride The Coming Data Center Efficiency Wave

Analytics Report

Big changes are afoot, driven by the winds of green awareness. Communication is key to steering a steady course between IT and facilities while making the most of your CEO's newfound interest in data center operations. Here's how to keep an even keel in the face of transformational forces.

By Kenneth Miller

InformationWeek
BUSINESS TECHNOLOGY
NETWORK

TABLE OF CONTENTS

TABLE OF CONTENTS

Kenneth Miller has more than 15 years of executive-level experience implementing complex IT systems and is adept at developing and driving strategic technological direction for a wide range of enterprises. Currently, Mr. Miller is data center architect with the IT Infrastructure and Operation Services division of Midwest ISO, developing mission-critical facilities. Mr. Miller also leads the IT and business consulting practice of Thinx, including leading feasibility studies and developing large-scale IT capital projects for executive management.

Prior to Thinx, Mr. Miller was chief technology officer of n|Frame, a provider of IT business continuity services, including high-availability and high-performance IT infrastructure and e-commerce systems. As CTO, he was responsible for development of the next generation of n|Frame products and services, as well the long-term technical vision and roadmap for the company. Prior to joining n|Frame, Mr. Miller was president and founder of Data Communications Consulting, an IT consulting practice. Mr. Miller directed the company's enterprise communications practice, focusing on applying new data and telecommunications technologies to solve high-profile business challenges.

Executive Summary

A massive wave of efficiency will sweep through data centers large and small over the next few years. We've survived many changes over the last several decades, of course, but the current situation has one element that makes it more transformational than anything we've seen before: Scrutiny from the highest echelons of business management are moving data centers from business necessity to strategic advantage.

Virtualization. Cloud computing. Escalating power densities and equipment weight. Environmental awareness. All of these trends, along with explosive data center growth at odds with emerging efficiency mandates, demand fundamental change in the way data centers are designed and operated—and not just because of energy or ecological impact, but because the cost of traditional data center models is escalating out of control.

Between rapid obsolescence of the physical infrastructure and the rise of more sophisticated monitoring and control capabilities, building electrical and mechanical systems can no longer operate solely in the domain of facilities. Rather, management of these systems must play in integral role the IT infrastructure, forcing facilities and IT leaders to work much more closely. We'll examine this and other trends and discuss some of the newest technologies to improve efficiency.

Research Synopsis

Survey Name: *InformationWeek* Data Center Survey
Survey Date: November 2008
Region: North America
Number of Respondents: 279

Purpose:
To determine the state of data center architecture in the enterprise

Methodology:
InformationWeek Analytics surveyed business technology decision makers at North American companies. The survey was conducted online, and respondents were recruited via an e-mail invitation containing an embedded link to the survey. The e-mail invitation was sent to qualified *InformationWeek* subscribers.

State Of The Data Center

Those responsible for operating their organizations' data centers are feeling the heat, and not just from the airflow in overloaded facilities. While more than half of the 279 business technology professionals we surveyed for this *InformationWeek* Analytics report expect demands on their data centers to increase in 2009, only 25% will see their facilities budgets grow. Meanwhile, thanks in large part to the green movement, IT and line-of-business leaders alike are inundated with stats, drawn from across the industry, on how data centers are operating.

They're taking notice.

While benchmarking performance against peers has value, it's also fraught with pitfalls. For example, Gartner and other pundits estimate the average operating cost for a Tier 3/Tier 4 facility at $1,200 to $1,300 per square foot. However, this figure is highly variable based on location, regional utility costs, and the fundamental intent and use of the data center. These single-number metrics fail to account for these differences. Without more insightful analysis, such simple comparisons can result very wrong conclusions.

What isn't variable is the pressure on capital and operating budgets, and the fact that IT is driving considerable growth in energy use. In fact, according to the EPA, the energy consumed by U.S. commercial and industrial buildings is responsible for nearly 50 percent of our national greenhouse gas emissions, and within these buildings, the power used by data centers has doubled over the past five years. Meanwhile the national average rate for electricity has jumped 44% since 2004. Yet, our 2008 *InformationWeek* 500 report revealed that the stagnant economy is

Figure 1

Resource Demand Expected To Increase In 2009

Your data center facilities resource demands for 2009 are projected to be:

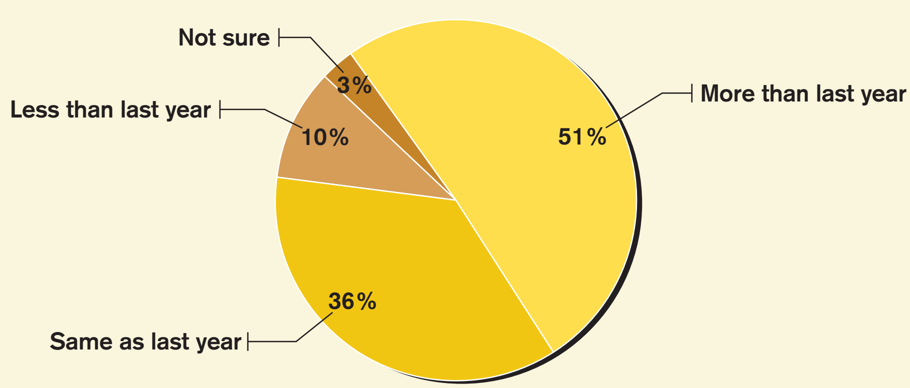

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

keeping overall IT spending relatively flat over last year, even at large companies, with the average IT budget stuck at 2.8% percent of revenue. However, with revenues down, average dollars spent actually dropped, from $435 million in 2007 to $354 million in 2008.

For some small companies, the answer is to outsource their data center challenges. In our recent *InformationWeek* Analytics report on cloud computing, we found that virtualization is enabling large service providers to offer shared platforms at a fraction of the cost of physical platforms—significant because 68% in that poll say cost is the most important factor in moving toward a hosted model.

Clearly, cloud computing, virtualization, the green movement, roller-coaster utility costs, budget pressures, and other factors are running smack into an increasing business dependence on software and computing resources. This is all helping drive the apparently supercharged growth of data centers, with Microsoft spending billions to build new sites in Chicago, San Antonio, and Dublin, with the Chicago location using Redmond's much-touted "Generation 4 Modular Data Center" model. Not to be outdone, Google's new $600 million high-tech center in Council Bluffs, Iowa, is fast taking shape, while IBM says it will construct a massive new data center in North Carolina's Research Triangle and open a "cloud computing center" in Tokyo, at a cost of nearly $400 million. All are placing great emphasis on energy-efficient construction methods.

For those with more modest resources, the top challenges in data center operations over the next several years will be meeting increasing demand in the face of continuing budget pressure, finally bridging the gap between IT and facilities groups, and learning to speak the language of

Figure 2

Insufficient Cooling, Space Top Concerns

Please rate following concerns for your organization's data center.

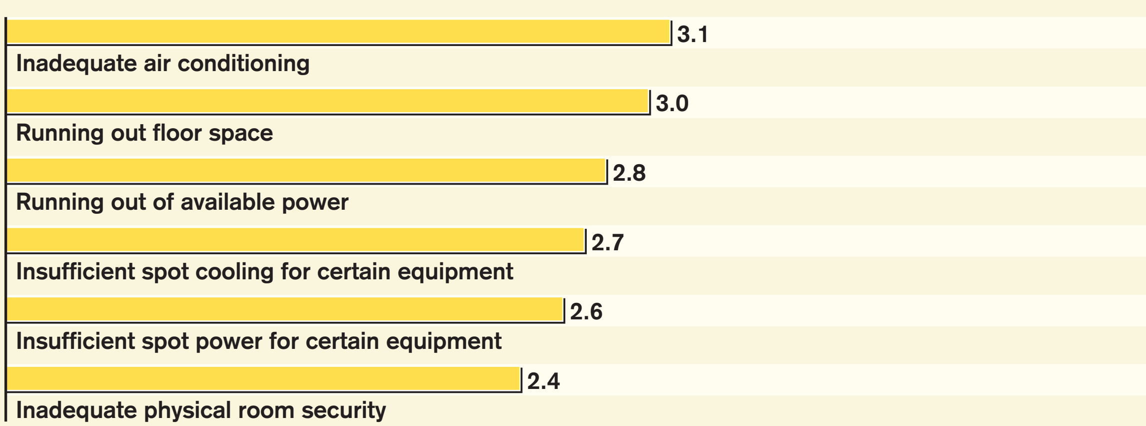

Note: Mean averages based on a 5-point scale where 1 is "not a concern" and 5 is "primary concern"
Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

business executives, all while preparing for the coming wave of change brought on by energy-efficiency mandates.

IF YOU CAN'T BUILD, THEN REBUILD

Data center managers facing escalating resource demands have only two options: Add floor space, power, and/or cooling, or make more efficient use of those resources. With just 25% of respondents saying their budgets will increase in 2009 and close to 60% of respondents cost-justifying each project as it comes along, we don't expect a large number of green-field data center construction projects to launch next year.

The primary way for most of us to meet growing demand will be to incrementally increase the efficiency of older facilities that likely weren't built to handle modern loads. The challenge here is that it's very difficult to achieve the efficiency gains necessary to meet future demand by taking small nibbles out of the problem.

Yet, is making radical changes to the way our data centers operate even feasible?

In many cases, it will be the only option. But "radical change" doesn't have to mean breaking ground for a new site with all the green bells and whistles. For many CIOs, bridging a divide with facilities groups that's reminiscent of the old days of IT vs. telecom could yield surprising economic advantages, as can new technologies and techniques that we'll discuss later.

Mounting business demand on the data center isn't all we need to worry about. Major IT industry trends are adding unique challenges as well; these include increasing power and weight

Figure 3

Few Will Increase Data Center Budgets In 2009

How will your data center facilities budget for 2009 compare with your 2008 budget?

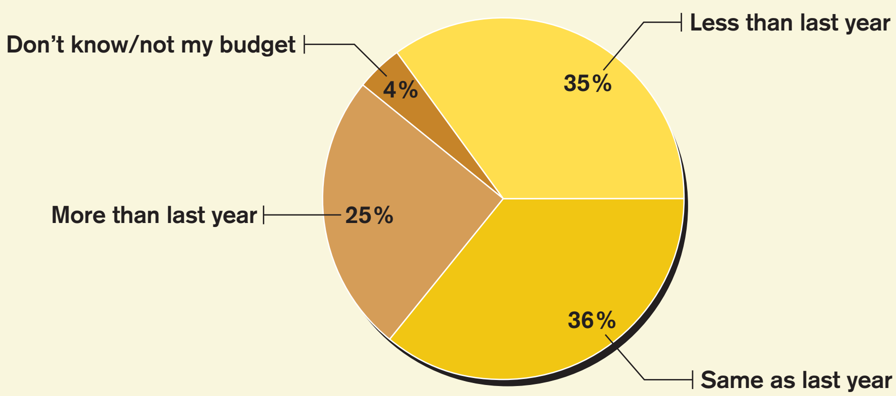

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

density of equipment, a proliferation of multicore processors, virtualization, the spiraling rate of storage growth, a scarcity of technically qualified staff, high construction costs associated with even modest data center infrastructure upgrades, and increasing demand for all forms of communications and associated equipment.

Even more frightening is the number of projects that we've seen recently where anticipated upgrades were stopped in mid-stride due to unanticipated structural, power, and cooling system constraints outside of the data center. In one case, the local utility refused to deliver more power unless the company paid for an entire substation upgrade. In another project we consulted on, the data center expansion was stopped cold because the building could not structurally support the new infrastructure and projected data center weight loads.

Time for plan B—as in, "back to the drawing board."

WHY EFFICIENCY WILL DRIVE REAL CHANGE

Efficiency has a multiplier effect in terms of capital and operating costs. Small gains produce not only a large payback over time but may delay the need for capital investment. As the economics of this multiplier become understood at the executive level, efficiency will enter the realm of strategic economic advantage.

In terms of environmental impact, data center efficiency is gaining real traction because it brings benefits for many parts of the organization, including those outside IT. Growing social awareness of green initiatives, for example, is driving large enterprises to evaluate their environmental stances in an effort to strengthen their positions with customers and employees.

Figure 4

Ad Hoc Funding For Capacity Increases

How do you allocate funding for data center capacity increase

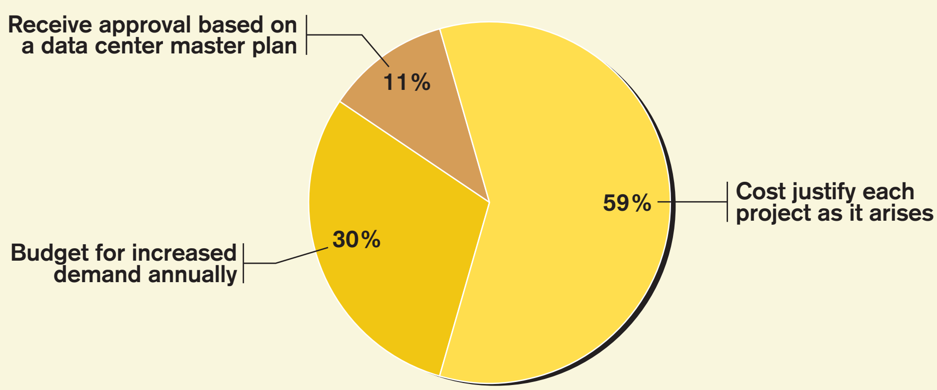

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

When assessments are done, depending on the industry, data centers frequently top the list of energy consumers—and therefore, carbon producers. Organizations seriously looking to minimize their environmental footprints will inevitably end up on the data center doorstep.

Still, as we discuss in our upcoming *InformationWeek* Analytics Green IT report, you shouldn't bet your carbon credits on IT reducing energy consumption overnight: A mere 12% of the 419 business technology professionals we surveyed for that report say they'd be willing to pay more for energy-efficient gear. The shift toward green awareness must gain traction with an organization as a strategic effort before major data center renovations, or even individual purchases, can take place.

In addition to driving a significant portion of an organization's energy costs, the capital investment in data center IT infrastructure (networks and storage) and physical plant infrastructure (CRACs, UPSes, and PDUs) has been growing at a more rapid pace than the spend on IT servers and software, according to a study by McKinsey. For many of the organizations we've worked with, this mounting IT budget item doesn't even include the costs for mechanical and electrical systems, which typically come out of the facilities budget. In fact, in our reader poll for this report, just 16% of respondents say it's IT that gets the tab for data center electricity and takes responsibility for managing it. A discouraging 23% say no one is accountable for governing power usage.

Efficiency isn't just about energy. In the data center it's also about squeezing the last dollar out of what's been invested and extending a resource's life for as long as humanly possible. Much of the disproportionate growth in data center investment is being driven by inefficient mechanical and electrical infrastructure design. We are encouraged that 58% of respondents to our poll favor a modular approach to data center design. Perhaps the next generation will dodge some of the problems we're struggling with. For now, though, we'll discuss how IT can improve these systems without scrapping the data center and starting over.

"The first rule of any technology used in a business is that automation applied to an efficient operation will magnify the efficiency. The second is that automation applied to an inefficient operation will magnify the inefficiency." —Bill Gates

NEW TECHNOLOGICAL SCALE

Something is broken in the way data centers are scaling up. The canaries in our data center coalmine are the giants—Google, Microsoft, IBM, Verizon, and various collocation providers. These entities are building data centers on scales that few outside the industry can grasp.

Those of us inside the business are sitting back and watching with great interest.

What these titans are discovering is that the traditional design of raised-floor cooling with perimeter CRAC units, bulky 480/208 volt in-room PDUs, and traditional mechanical systems does not scale. In fact, if IT demand is projected into the future at historical efficiencies, there are very hard limits to what can be achieved for a reasonable cost. And when you move outside "standard" equipment and facility sizes, prices escalate even faster. High-end power consumption can require new or custom substations. High-density racks call for custom PDUs. Large sets of parallel generators or UPSes may require custom parallel equipment. As we all know, anything custom is rarely cost-efficient.

In searching for solutions to these problems of scale, we've found enormous benefits in designing efficiently. It sounds obvious, but these new designs mean giving up entrenched data center stereotypes. Who wants to build a mission-critical facility that breaks with the tradition of a raised floor, or runs at 80 degrees? Yet, that's exactly what these leaders are doing, at great economic advantage. Their lessons can be used in our legacy data centers to hold back the rising tide.

TECHNOLOGICAL INFILTRATION OF PHYSICAL SYSTEMS

A major factor in the coming tsunami of efficiency is that technological progress is—finally—beginning to transform traditional building electrical and mechanical systems. For those coming from a technology background, the concept seems simple, but in our practice we've found that it's considered cutting edge for a facility to monitor itself as a total system and dynamically respond based on sensing resource demand, rather than on preprogrammed schedules or manually adjusted set points.

For example, building components, such as air handlers and humidifiers, have long had individual control points that react to their immediate surroundings. However, this does not allow IT to make globally efficient decisions—for example, operating only the minimum number of air handlers necessary to meet facility demand, or implementing facility-wide humidity control vs.

Figure 5

Planned Projects, ROI Impact Gear Upgrades

In the following situations, what is the likelihood you will upgrade non-IT equipment (CRAC units, PDUs, UPSes) for purposes of improving efficiency?

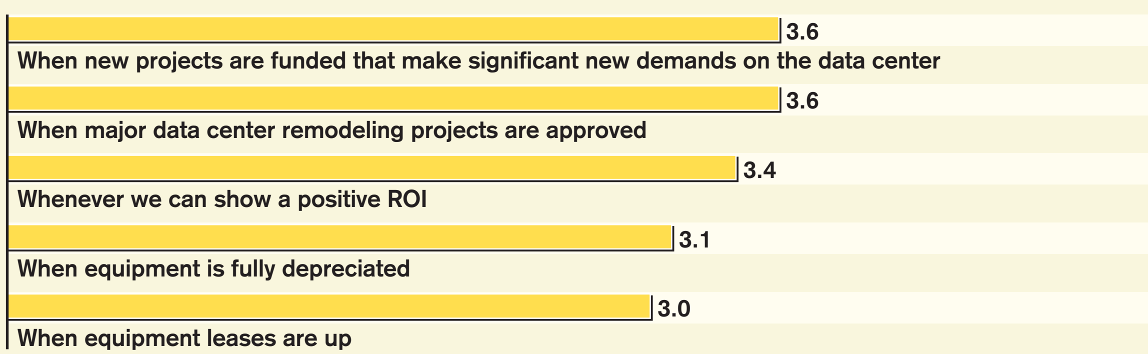

Note: Mean averages based on a 5-point scale where 1 is "won't do it" and 5 is "planned/already doing it"
Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

localized humidification and dehumidification. For that, we need centralized monitoring, control, and intelligent software, but this has been hard to come by because it was difficult to justify the operational gains.

In our experience, in fact, it's not uncommon for data center facilities to have minimal or even no monitoring of energy consumption outside the UPS and critical IT equipment loads. "Keep it simple, reliable, and inexpensive" has long been the mantra for mechanical and electrical engineers. That means power consumption/distribution has been viewed as a static system—install a breaker panel and PDU and you shouldn't have to worry for years, right?

But such thinking means many breakers, and much power, become stranded in odd places, and without real data on electrical system performance, a global decision to reallocate resources cannot be made.

This approach is finally being challenged, thanks to the cost of monitoring and control systems being driven down by commoditization. In fact, a range of inexpensive digital instruments for monitoring mechanical and electrical system performance, loads, and faults over time in legacy facilities is available now and will allow for an informed approach to optimization, updates, or redesign.

A top barrier to justifying even these lower costs is that the evolution of technology in a building infrastructure is typically on a 20- to 30-year cycle, due to the view that a structural investment is depreciated over this length of time. But data centers are accelerating facilities investments and making infrastructure elements obsolete long before they're scheduled for replace-

Figure 6

Data Center Depreciation Under 10 Years

What is your organization's depreciation timeline for data center facilities?

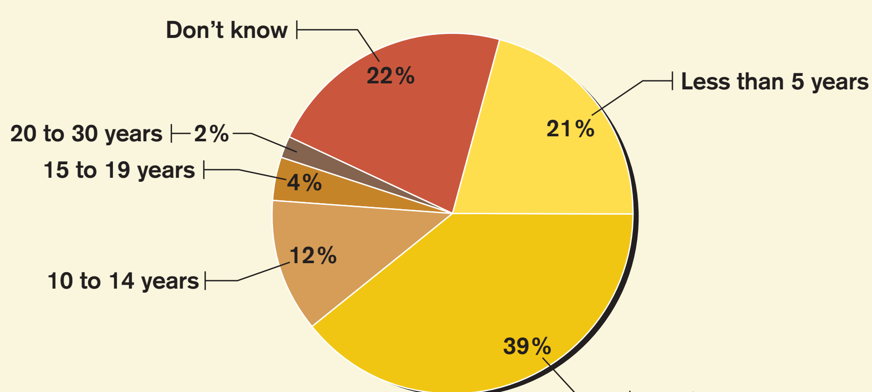

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

ment. At today's rate of change, a data center can become unmanageable, or at least cost much more to operate than newer models, in just seven to 10 years. For a company to simply keep on par with efficiency advances, facility infrastructure upgrades must occur on a flexible timescale previously unheard of.

So what's a good place to begin? Products that can help with unified management include variable-speed pump and fan controls based on variable-frequency drives (VFDs). Although VFDs have been available for years as optional equipment, vendors have now begun to include these capabilities in mainstream products and designs. With the price of energy rising, the savings realized by operating only what is needed easily offset upfront design and instrumentation costs. Retrofit paybacks can be more complex to calculate but are still worth the investment, in our view, if your organization's ROI horizon allows.

Between rapid obsolescence of the building infrastructure and the rise of more-sophisticated monitoring, facility controls, software, and programming must become integrated with IT monitoring and control software.

We believe coming changes will drive this détente.

ORGANIZATIONAL CHANGE

In most companies, the building is managed by the facilities group, data center equipment is managed by IT, and while the dividing lines vary from organization to organization, these camps are rarely organizationally integrated. The classic example is an IT group that shows up to install 50 kW of new servers without having discussed whether there's available capacity on the UPSes or PDUs.

This operating model presents challenges, but they haven't generally been insurmountable, so the separation continues. In fact, the false dichotomy between computing platforms in the data center and the facilities infrastructure of the data center itself is reminiscent of the gulf we saw between the telecom and the data communications worlds a decade ago—a state some would argue is still latently present. It's about distinct experiences, budgets, and domains of expertise—and about who reports to which office in the executive suite.

The irony is that IT and facilities can learn a great deal from one another. Understanding of power and cooling isn't typically a strength of technologists, yet IT decision makers are increasingly pressured to make equipment choices with a poor conceptual grasp of key issues. And complex communications and dynamic software control systems have not typically been a competency of facilities staffs, yet with the rapid evolution of these control systems, they're increasingly making decisions about software, communication, and security architectures of new control systems.

So why is the wall crumbling now? Simply, the rate of data center change and the growing cost consequences to an organization when there are problems. To operate a data center efficiently, business decisions that drive IT investments must include a complete cost picture, which involves a significant facilities component. Unless facilities is integrated into this process, the business will not understand that a $1 million blade server investment could drive an additional $5 million in facilities capital costs and a 5% annual increase in facilities operating expenses. So, a $3 million decision to rewrite software to operate more efficiently, instead of increasing hardware performance, won't likely be made.

True efficiency in mechanical and electrical use also requires coordination between IT and facilities around the confluence of demand, projected growth, and planned replacements. Instead of facilities reacting to IT, which can drive an inefficient "fix it quick" approach, teams should work together so that long lead-time changes to building systems can be made with good planning. In one organization we worked with, for example, the justification for a larger and more efficient UPS as a replacement for aging UPSes was made based on a clear IT project schedule covering a 24-month period. In this case, facilities worked with IT to develop power requirements for projected IT equipment installs, and IT helped facilities cost-justify the new UPS.

What this organizational change will look like will be highly dependent on your organization and industry. A cooperative process will help each be successful and drive value for the organization instead of struggling through long learning curves on their own.

But make no mistake, change *is* coming. If both camps don't get ahead of the situation with an effective and cooperative approach, we're confident that a solution will be handed down from above soon enough.

Figure 7

Most Favor Modular Design

Which best describes your approach to data center design?

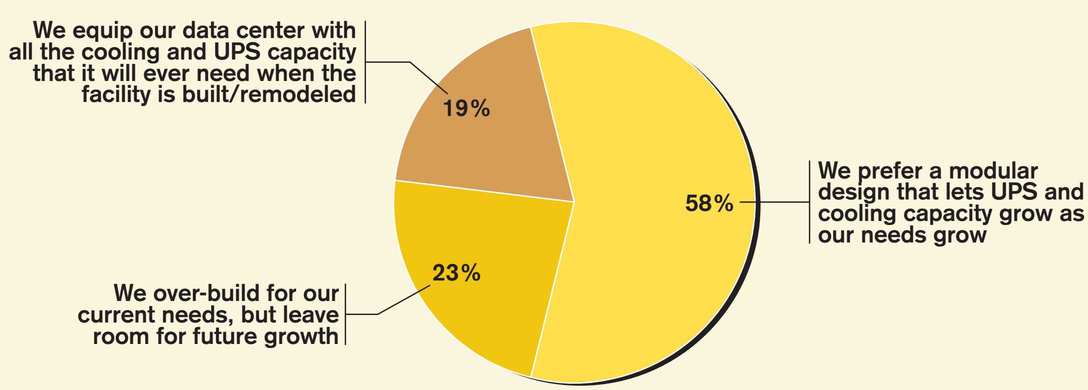

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

PRAEMONITUS, PRAEMUNITUS

As we've discussed, the data center is about to become an intense focus for your organizational leadership, if it hasn't already. Forewarned is forearmed, so get prepared. In fact, where there's a divide of culture, trust, and cooperation between IT and facilities, you can use the impending scrutiny as a way to bridge the gap. If the transition is handled well, instead of a revolution, a gradual evolution can be managed that's much less threatening to staff and professional careers than emergency board meetings and unexpected large-scale capital investments.

Once you have communication with facilities down, how do prepare for the scrutiny of executive leadership? Measure, cover the basics, then educate.

FIRST, MEASURE

The first step in preparation is to measure. It may sound simple, but providing a historical view of your space, power, and cooling resources is key in telling the story of efficient management. We know many data center professionals who've made truly miraculous capacity improvements in their data center operations by picking off the low-hanging fruit, but instead of appreciation, the next time they request capital to support another business initiative, they're grilled as to why there are no more magic rabbits in the hat.

Figure 8

Efficiency Evaluation A Collaborative Effort

Who do you look to for data center efficiency evaluation and change recommendations?

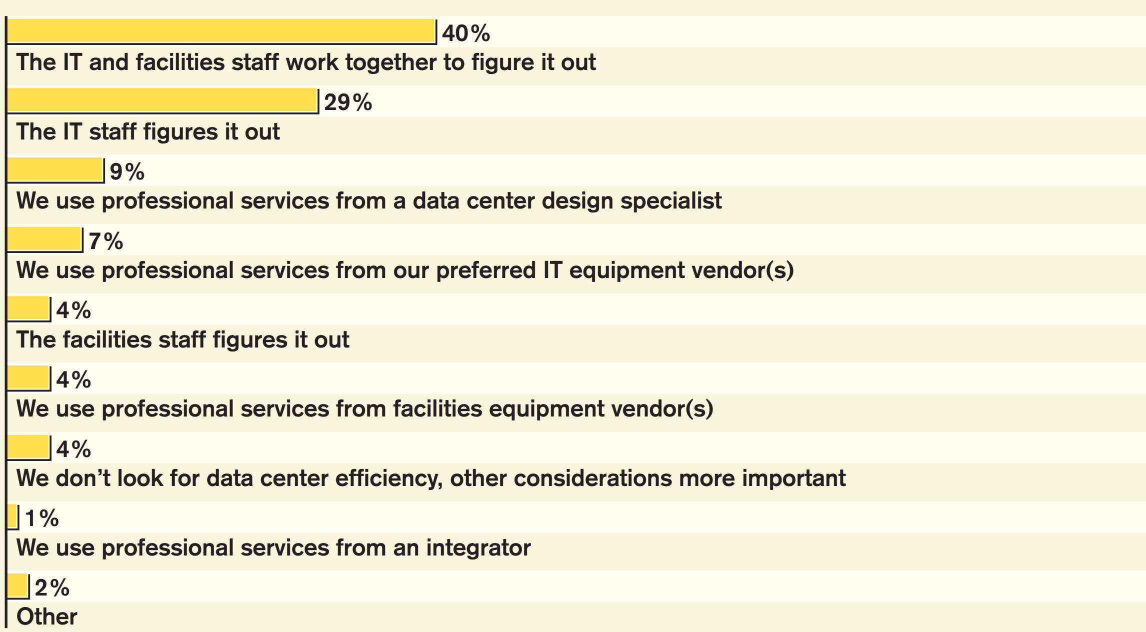

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

While measuring is extremely important to demonstrating that every last ounce of capacity has been squeezed out of existing investments, it can also help you contrast your data center's performance with that of other organizations. There are many factors that might influence your choice of an efficiency metric for your data center, however, once you've picked a standard, stick with it.

And, make a decision on your metric *before* initiating improvements—if you can't document enhancements, in our experience, it will be as if they never happened.

As a general guideline, you should measure and trend at minimum the following data points on at least a monthly basis:

- Energy consumption of all IT equipment;
- Energy consumption of all supporting mechanical and electrical systems;
- Cost for the energy consumption of the two items above;
- IT data center room thermal load, which may include loads from non-IT equipment such as PDUs;
- Thermal loads for all mechanical and electrical systems outside the data center supporting the IT load; and
- Available—and useable—IT rack space. In your documentation, note space that cannot be used for environmental or technical reasons and so should not be counted as available capacity.

Compare the above items with installed, and designed upgradeable, capacity. Now, if yours is like most organizations, you don't have the systems in place to automatically track the above infor-

Figure 9

Room For Improvement

Which best describes your company regarding its use of data center power?

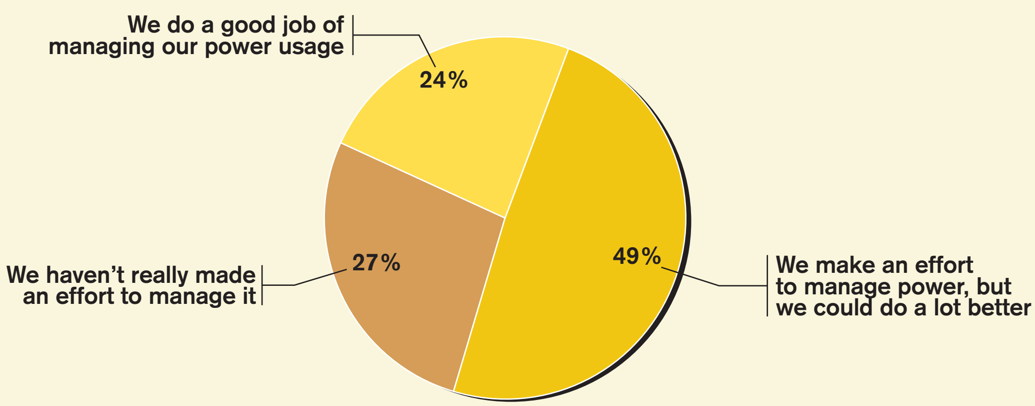

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

mation. Still, do the best you can to collect this data, even if it's just a monthly manual "amp clamp" snapshot of key electrical circuits or a calculated thermal load based on the theoretical heat generation of your IT equipment and infrastructure. Even a quick rack-space count done on a monthly basis is better than no data at all.

It's also very important to carefully examine your capacity metrics. We've worked on many projects where the real capacities of mechanical and electrical systems were misunderstood by both IT and facilities. Wrong assumptions have led to improper operation, translating into a risk to data center reliability—and large, unexpected capital expenditures to fix problems.

For example, the capacity of the power system must properly factor in the 80% nominal ratings required by code. Another common mistake is the "ton" rating on CRAC units. Talk with your vendor or manufacturer to understand the nominal rating *for your equipment at your operating conditions*. You may find that the actual thermal capacity of your CRAC units is far less than you think. Both electrical and mechanical system capacities should be evaluated from end to end. Just because you have room for another 30-ton CRAC unit on the data center floor does not mean that the cooling tower will be able to support the additional load.

COVER THE BASICS

If you examine the "best practices" lists that data center vendors, consultants, and industry associations provide, you'll find quite a few things that can be done to improve efficiency without rebuilding the entire facility. Don't reinvent the wheel—taking advantage of the wisdom of your peers is smart, and following these practices really can help improve your data center's performance. When you implement best practices, make sure that they're documented and, where possible, correlated to your measurements. Being able to demonstrate that the basics have been covered shows that IT has made significant efforts to prevent, or at least delay, requests for additional data center resources. We'll discuss some best practices later.

There are also quite a few high-level white papers from organizations including Gartner, Uptime Institute, McKinsey, Hewlett-Packard, and IBM circulating within executive circles that tout the vast improvements that can be made in your data center without having to add capacity. These should be viewed with some skepticism.

Measuring your data center metrics and covering the basics according to a best-practices framework will help you dodge a painful death by executive white paper.

EDUCATE

More than a few great ideas have languished on the drawing board because of a lack of understanding by decision makers. In our experience, the larger the capital costs and the more strategic the change, the longer education takes. In some data centers, we could be talking years. Continually apprising management of resource constraints and growth rates in advance of the coming revolution will help you build your case and avoid this fate.

Given an extended education timeline, it's important that your message be consistent. You need a clear idea of where you're going as a business and what kind of data center you'll need when you get there. This analysis will also help you know which metrics you should be measuring.

What business drivers or functions are growing the fastest in the data center? Is the sales department requesting a new server every day, or is your compliance department archiving every single e-mail and instant message? Do you primarily purchase entry-level server platforms by the dozen, or are you bringing in full-rack midrange or mainframe systems? Breaking your resource consumption down in a way that the CEO can understand ties resource requests back to the business. Saying you want to buy a new 250 kW UPS and associated electrical infrastructure will make management's collective eyes glaze over. But saying you want to add data center capacity to support the new SAS 70 requirements being driven by corporate compliance requirements means something to the business.

Education should include regular circulation of data center status reports, or quarterly resource briefings, detailing your measurements/metrics, best practices efforts (for example, 50% complete with floor-penetration sealing), and your educational campaign (for example, legal

Figure 10

Data Center Consolidation Primary Power-Saving Method

Please rate the likelihood that your organization will take the following measures to save power.

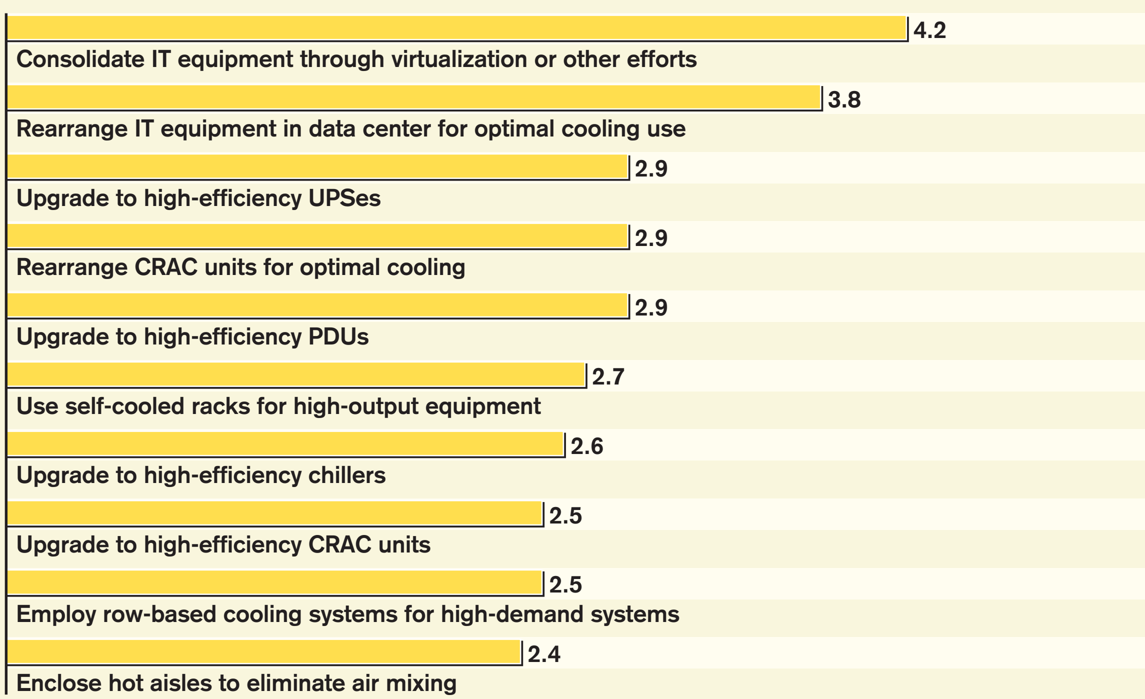

Note: Mean averages based on a 5-point scale where 1 is "won't do it" and 5 is "planned/already doing it"
Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

installed 20 new servers this month). Circulate industry information about data center performance and how your organization compares with these numbers. Compare norms against your metrics, best practices, and business-specific drivers. Show the capacities of and loads on your systems. Start communicating in advance of when these resources will be constrained. Work to get this key information included in IT project planning.

SPEAK THE LANGUAGE OF ENERGY AND HEAT

It's important to explain what's driving the massive growth in resource consumption in the data center, beyond the obvious installation of more equipment. There are a few general areas where educational materials will enlighten business leaders.

First, one result of Moore's Law—packing double the transistors into chips every 18 months—is increased server power consumption, despite improvements in chip energy efficiency. This is not just a function of the processor but also the motherboard, memory, and storage systems.

Discuss growth in the total amount of data within an organization. Storage in particular is growing more rapidly today than other computing resources. Finally, the energy consumed by mechanical and electrical systems is directly related to the energy consumed by IT equipment. Yet mechanical and electrical systems are typically not designed for modular growth or efficient operation at suboptimal load.

These issues combine to represent a multiplier effect. And, seeing as just 24% of our poll respondents say they believe they're doing a good job managing data center power usage—compared with 27% who aren't expending any effort—we need all the help we can get here.

Figure 11

Expected Data Center Rating

If audited today, what tier rating do you believe your data center would achieve?

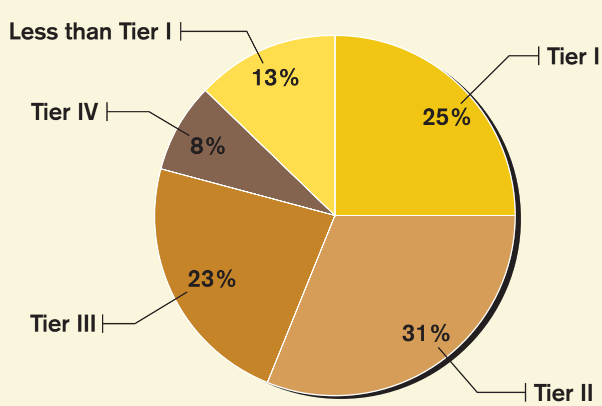

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

The more energy IT equipment draws, the more additional energy is required to remove the heat generated, and the more capital must be invested in expensive mechanical and electrical systems. As discussed above, these large investments are typically absorbed when a building is constructed, so the sheer magnitude of the numbers involved are rarely understood.

ENERGY GUZZLERS

One factor driving the data center to center stage in the board room is energy consumption. Of course, the cost of power and the cost to deliver it are part of the equation. But the real source of the attention is embarrassment over the amount of energy that's wasted by outdated equipment and old designs. For example, older UPSes at lower loads can be less than 50% efficient—meaning we may lose 50% of the power going to them to wasted heat. Additionally, business and IT project discussions typically do not factor in energy consumption on the front end. Instead, it tends to fall into unexpected operating increases on the back end of a project.

We're always interested in talking with CIOs about the amount of power drawn by the current set of servers in their data centers. Don't get us wrong—most know the numbers, but they don't always hit home in terms of real-world equivalence. To communicate effectively, it's helpful to show people just how much energy is involved, in terms they can understand. Say you buy a server today that, in round numbers, continuously consumes 1 kW. Over a year, it will draw 8,766 kWh of power. When we take into consideration that many data centers consume twice a

Environmental Overkill

The American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) has published guidelines for its mechanical-engineering members. Historically, ASHRAE has classified data centers as Class 1 and Class 2 spaces and provided a recommended environment range of 20 to 25 degrees Celsius (68 to 77 degrees Fahrenheit, or dry bulb temperature) and a relative humidity range of 40% to 55%. The group also publishes a much wider "allowable" operating range.

Many industry professionals have questioned whether these tight requirements are appropriate, and there's been much recent discussion about wider recommended and allowable ranges. After some debate, ASHRAE published new guidelines in 2008; find them at: tc99.ashraetcs.org/documents/ASHRAE_Extended_Environmental_Envelope_Final_Aug_1_2008.pdf

Intel, Microsoft, and Google have all published proof-of-concept data center tests this year showing that operating ranges well outside of traditional environmental guidelines can work. Widening these ranges translates directly into energy savings and more attractive ROI calculations. Warmer allowable room temperatures reduce the energy required to remove heat from the data center and increase the effective capacity of your cooling units by permitting a larger temperature drop across the cooling coils. In addition, managing humidity can require a surprising amount of energy, which can be saved if tight control is not required.

Finally, if higher temperatures are recommended or allowable, "free cooling" or "economizer" equipment can operate for longer periods of the year at wider geographic latitudes, making investments in these new assets produce a faster ROI.

server's power draw to support its operation, that device drives a whopping 26,298 kWh, excluding linked infrastructure elements, such as communications and storage. This is equivalent to the energy released when burning approximately 800 gallons of gasoline. The average American uses approximately 500 gallons of gasoline per year, along with 11,571 kWh of electricity every year in their households, according to the U.S. Energy Information Administration. Simply put, a single server running for a year consumes approximately enough energy to power the average household and one automobile for an entire year.

The following rough scenario explains the multiplier effect. If you add 1 kW of IT load in the data center, this drives additional power consumption. Given the inefficiency of getting power to IT equipment resulting from UPS, PDU, and other types of electrical system losses, some estimates indicate a load increase at the meter of another 1 kW consumed in the process of delivering power from the utility to your data center IT equipment, with most of it converted to waste heat. In addition, any increase at the meter leads to an upsurge in mechanical system power consumption to remove that heat from the building. Based on other industry studies, that means, on average, an additional 1.3 kW in mechanical system power consumption.

In total, an increase in IT load of 1 kW can drive an additional 2 kW to 3 kW, for a total of 3 kW to 4 kW of consumption growth.

Industry groups including the Green Grid, Uptime Institute, and McKinsey are pushing for standards around energy efficiency calculations for data centers. The specs gaining the most traction today are Power Usage Effectiveness (PUE), Data Center Infrastructure Efficiency (DCiE), Site Infrastructure Energy Efficiency Ratio (SI-EER), and Corporate Average Data Center Efficiency

Figure 12

Power Management Responsibilities Split

Which of the following best describes the management of power usage at your organization?

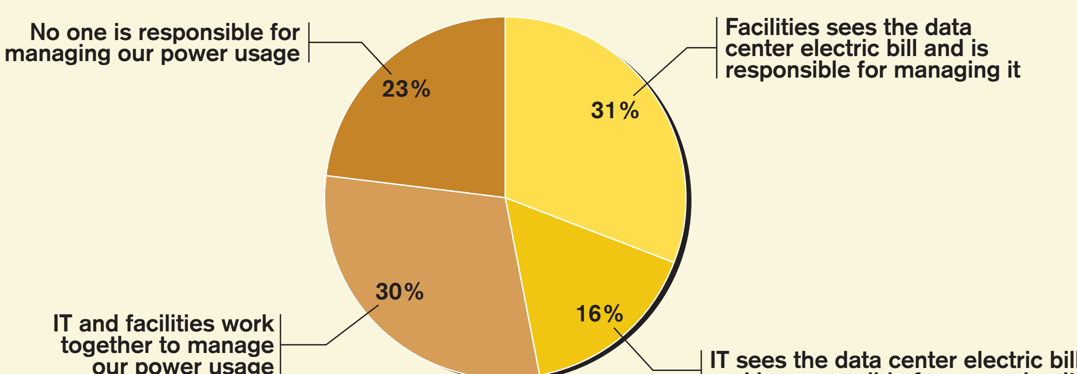

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

(CADE). It's important to understand how these are calculated, and to be aware that there are still substantive differences of opinion around how they should be used, and which is most accurate (see p. 23 for links).

Although IT may sound smarter tossing around sophisticated acronyms, all these standards are just different ways of looking at the energy overhead associated with operating IT equipment. Most of our work has been centered around the use of PUE, but depending on your industry and location, you may find one of the other efficiency calculations more effective in communicating with your management and making comparisons with peers. What's important is to understand the multiplier effect, choose one metric, and then apply it consistently.

For a typical data center today, most studies indicate an average PUE of 2.0 to 3.0. This means that for every kW of IT load, an additional 2 kW to 3kW of other system power consumption will be added.

The EPA is also getting into the act with its August 2007 *Report to Congress on Server and Data Center Energy Efficiency*, which showed a sample of 22 West Coast data centers running at an average PUE of 2.0. It's predictive of things to come that the EPA is running these studies and trying to understand why data centers are one of the fastest-growing energy consumers. Official practices governing energy efficiency, specifically in the data center, are already present in Europe; for example, the European Commission in October issued its *Code of Conduct on Data Centres Energy Efficiency, Version 1.0*. Similar regulations in the United States are likely inevitable, and it's clear from historical precedent that rules will tighten over time.

Figure 13

Power Monitoring Strategies

How often do you currently use the following measures to gauge power consumption in your data center?

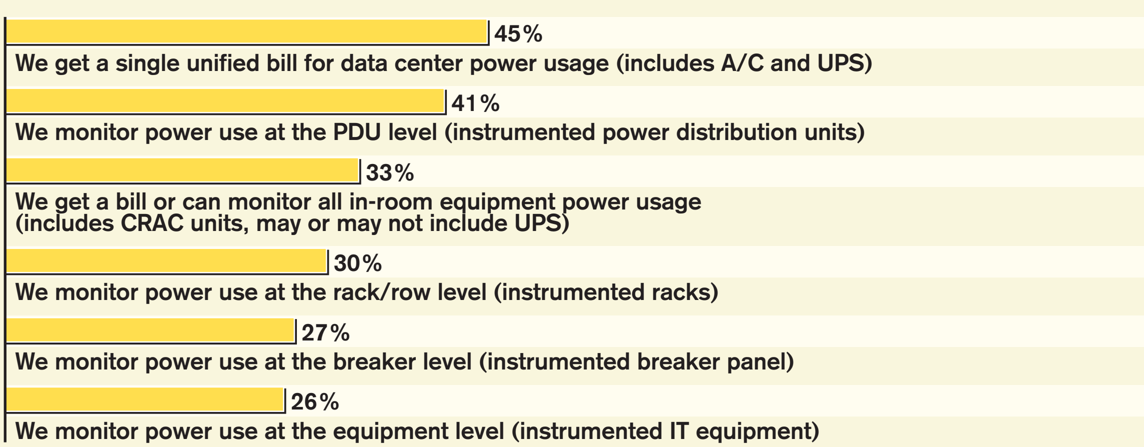

Note: Percentages reflect those responding "use often" or "always use"

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

So what can IT do now to improve? We've heard various opinions about what's possible to achieve from an efficiency perspective. Some big operators—notably Google—claim a PUE of 1.2. Depending on how usage is calculated, this is a bit extreme given today's technology, but it does demonstrate that the energy consumption of mechanical and electrical systems *can* be vastly improved, given the will and budget. These improvements translate into significant operating cost reductions, even postponement of major infrastructure upgrades. For organizations that are resource-bound in their current locations, extreme efficiency boosts can significantly extend a facility's operating life.

DON'T BE DENSE

Adding to energy-efficiency challenges is the growing power density of IT gear. From a computer equipment manufacturer's perspective, we can vastly increase the energy efficiency of servers through physical consolidation. This makes sense on an individual server or rack basis, but it really just transfers pressure to data center cooling systems.

We've seen many data center groups stretch the original design intent over the last 10 years to support racks of approximately 10 kW, but more than 70% of our poll respondents say that their data centers have not yet reached this level of power density. The good news is, there's still room for improvement in many facilities to reach these power densities. The bad news? It can be expensive and time consuming to push the original design intent to these levels, and it's very difficult to get above 10 kW without a major change in the data center operation—translation: significant renovation, rebuild, or green-field project.

In addition, pushing the original design intent almost always means less efficiency compared with data centers engineered from the ground up for higher power densities. The twin drivers of

Figure 14

Rack Power Demand

What is your target average rack power demand?

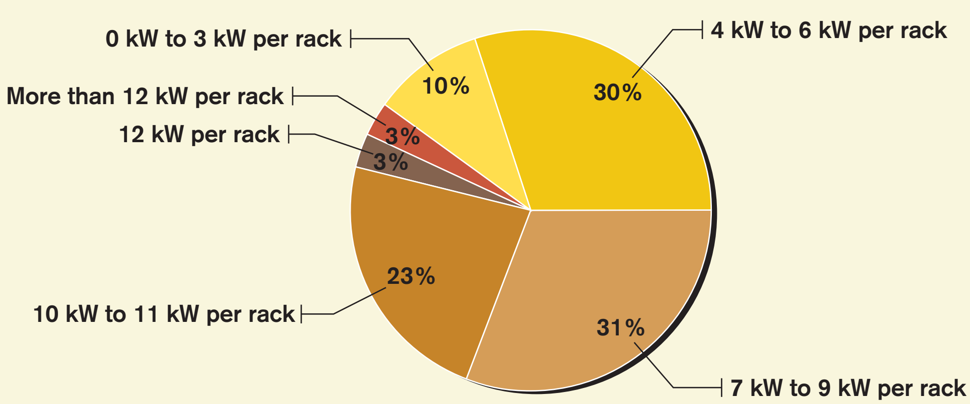

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

power density and energy efficiency demand a new approach to power delivery and to the cooling systems that support the IT equipment.

COOLING 101

A large portion of the energy consumed in operating a data center goes to systems that have one simple job: Remove heat. It's a shocking statistic that for the average data center today with a PUE of 2.0+, more energy is needed to deliver power to, and remove heat from, a server than is required to actually run the server.

It's been more than 100 years since Willis Carrier outlined the first compressor refrigeration system for office and industrial applications. The use of chilled water for cooling systems reached widespread use in the 1950s, and under-floor air distribution in the data center has remained virtually unchanged since the concept was "perfected" in the 1960s. There have been improvements in controls, tolerances, and other details leading to better efficiency, but fundamentally, these systems have not changed much since that time.

Our point is that our modern computing industry, also born in the 1960s, has evolved in performance and efficiency by several orders of magnitude. During the same time frame, our data center cooling systems have remained almost unchanged. This design and operating philosophy has been so effective for 40 years, that today approximately 60% of our survey respondents still use perimeter CRAC units with under-floor air distribution to cool their data centers. There's nothing inherently wrong with this approach—it's reliable, and it works.

If it's not broken, don't fix it, right?

Figure 15

Under-Floor Cooling Prevalent

What kind of design do you use for cooling your data center?
Please select all that apply to the cooling design used at your largest data center.

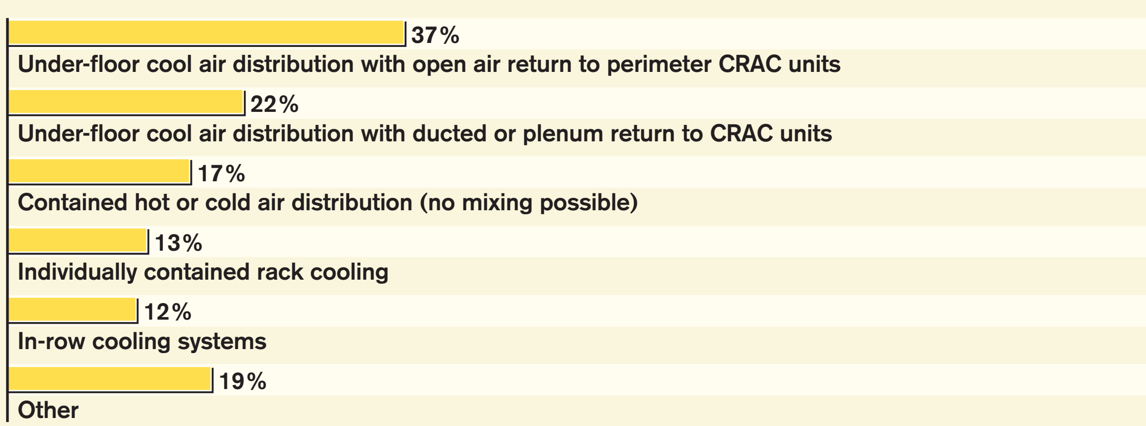

Note: Multiple responses allowed
Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

The problem is that the current model won't survive the next generation of IT gear. Just "working" isn't good enough. High-density blade servers—and the big multicore boxes needed for virtualization—really require fine-grained control of air flow and the resultant exhaust. If more than a few racks of blades are deployed, hot spots are almost inevitable with under-floor air and open returns. In fact, we've seen a few cases where traditional perimeter CRAC and under-floor air distribution wastes 80% to 90% of the energy consumed for airflow to keep the room within target temperatures.

These figures will vary depending on best practices being followed in the data center, of course, and often don't reflect efficiencies (or inefficiencies) in mechanical infrastructure outside the data center walls. Still, there's room for efficiency gains using perimeter CRAC unit design, apparent from the proliferating lists of best practices. Most guidance provided by vendors and consultants contains similar recommendations. When you consider that most data center designs face the same challenges, this makes sense. Here's our take on the basics.

CONTROL THE FLOW

- Plug all unnecessary floor penetrations to focus cooling to carefully placed vent tiles, and evenly pressurize the under flow; some studies claim a 10% reduction in energy consumption for the same cooling loads.
- Remove all unneeded infrastructure from under the floor to allow good air flow.
- Use hot and cold aisles.
- Install blanking panels to prevent in-rack wrap-back of hot air to the server intake.
- Use a plenum return.
- Use chimney or ducted hot-air return from the rack.
- In advanced installations, use Computation Fluid Dynamic (CFD) models to properly place CRAC units.

Links To Specifications

Green Grid Power Usage Effectiveness (PUE), Data Center Infrastructure Efficiency (DCiE)
thegreengrid.org/gg_content/TGG_Data_Center_Power_Efficiency_Metrics_PUE_and_DCiE.pdf

Uptime Institute Site Infrastructure Energy Efficiency Ratio (SI-EER)
uptimeinstitute.org/wp_pdf/(TUI3009F)FourMetricsDefineDataCenter.pdf

McKinsey Corporate Average Data Center Efficiency (CADE)
mckinsey.com/clientservice/bto/pointofview/pdf/Revolutionizing_Data_Center_Efficiency.pdf

European Commission *Code of Conduct on Data Centres Energy Efficiency Version 1.0*
re.jrc.ec.europa.eu/energyefficiency/pdf/CoC%20data%20centres%20nov2008/CoC%20DC%20v%201.0%20FINAL.pdf

COORDINATE THE FLOW

- Ensure CRAC unit air flow and cooling are coordinated among logical groups. We've seen an old CRAC unit begin heating the air placed under the floor because the unit next to it was making the return temperature too cold.
- Coordinate humidification so multiple units aren't performing opposite functions.
- Operate only the number of CRAC units that you need to cool the current room load and maintain reliability. Run all CRACs as close to maximum capacity as possible because a single unit operating at 100% rated load is significantly more energy efficient than three CRAC units each operating at 33% load. Admittedly, this can be difficult if you need the static pressure in a wide-open air-return design.
- In advanced installations, instrument and analyze the room as a complete system—this can help stamp out simple problems and, along with CFD models, enable changes that allow cooling systems to operate at the same performance with less energy—think working smarter, not harder. Again, this might be a difficult strategy for smaller data centers to implement, but typically, instrumenting a room and finding problems is far less costly than adding an unneeded CRAC.

CHALLENGE ASSUMPTIONS

- Operate the room at higher air temperatures. This is hard to manage if you're still using wide-open air circulation because higher air temps will only cause your hot spots to get hotter. If you have all your hot air contained, however, it's much easier to raise the room temperature without issues. Rule of thumb: Each degree below outdoor ambient temperature becomes increasingly more expensive to achieve.
- Widen acceptable humidity operating ranges to reduce or eliminate energy costs for adding or subtracting humidity from the data center.

WAYS TO INNOVATE

These suggestions aren't appropriate for everyone, but don't dismiss them out of hand. All can yield large savings via reduced energy consumption. For example, historically, chilled-water systems have been designed around supplying 45-degree water to cooling systems and expecting a 55-degree return water temperature. However, depending on CRAC equipment, target room temperature, and mechanical system design, running chilled water at 55 degrees, or even 60 degrees, is very possible. Running at higher loop temperature saves significant energy in cooling and increases the effective capacity (tons) of your chillers without additional equipment.

Also investigate new fan and pumping technologies. Electrically Commutated (EC) fans consume significantly less energy to move exactly the same amount of air as traditional fan technologies. In one recent project, we measured the energy consumption of a roomful of new EC fans supporting a 240 kW loaded facility to 100% airflow capacity as consuming the same energy as a single 20-ton perimeter CRAC unit fan motor.

Consider variable-speed and soft-start controls. This is a much larger and complex system design issue, but worth investigating for improvements in mechanical system energy efficiency. Varying fluid flows to match demand eliminates the energy costs for moving excess amounts of thermal medium (water), while soft starts can reduce demand on electrical systems and lower electrical system component sizing.

Finally, look into a "closed cool" design. Many manufacturers have products to bring cooling fluids to racks, or even into servers themselves. This approach significantly reduces energy consumption by reducing air circulation requirements.

ECONOMIZERS

Water Economizer: Depending on your risk tolerance, geographic latitude, and ambient environmental conditions, water returned from data centers can be circulated to a water-side economizer—in other words, a giant radiator—that uses outside air to cool the water instead of expensive compressor-based refrigerant cycles.

Air Economizer: Again depending on your risk tolerance, geographic latitude, and ambient environmental conditions, outside air can be brought into the data center to cool equipment. While there is debate regarding the risk of negatively impacting computing systems, this idea is compelling enough from an energy conservation perspective that proof-of-concept projects have been conducted by major vendors; see more in "Environmental Overkill," p. 21.

Figure 16

Air-Cooled Chiller Most Common

Which of the following does the primary mechanical system design include?

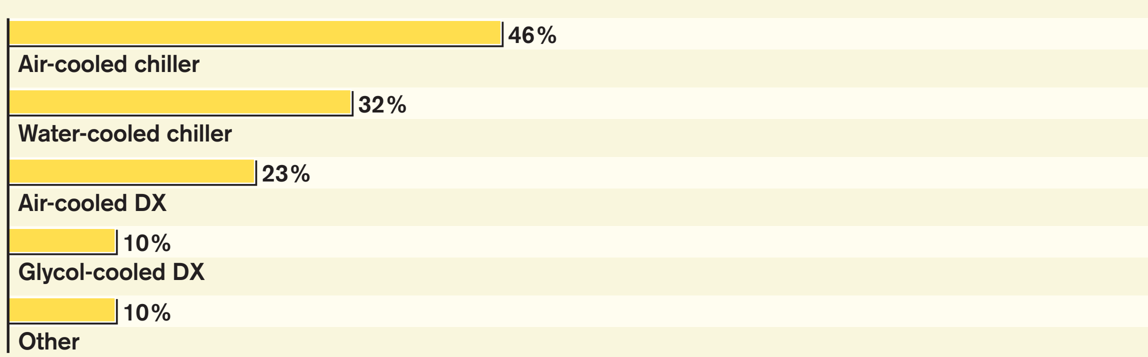

Note: Multiple responses allowed
Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

INDUSTRY AT THE CROSSROADS

Data center professionals have weathered many changes to typical operations since the 1960s, but we haven't gone far enough, and we're fast approaching the point where the decision will be taken out of our hands if we don't act. We no longer have the protection of the "little understood but business necessity" mantle. With collocation providers and readily available industry data showing CEOs the potential for cost reduction and strategic advantage in outsourcing compute capacity, we can either get ahead of the coming wave or be swamped by more nimble alternatives.

Think of this time not as stressful, but as an incredible opportunity to take a large and business-critical function and turn it into an agile and strategic asset, while at the same time contributing to your organization's green environmental position.

Appendix

Figure 17

Most Own Their Data Centers

Which best describes your company's data center strategy?

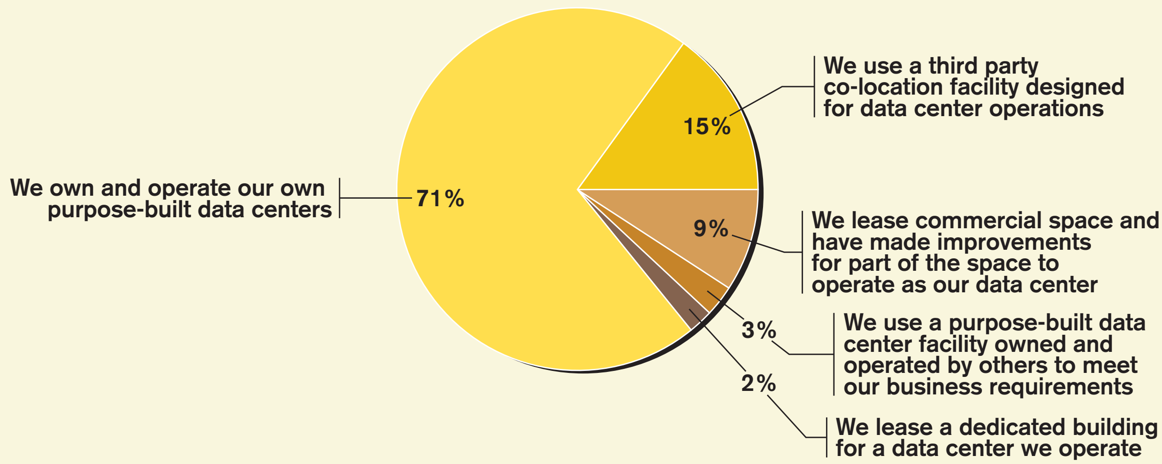

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

Figure 18

Data Center Size Matters

What is the approximate size of the data center for which you have the most management or decision-making responsibility? If more than one, indicate the size of the largest.

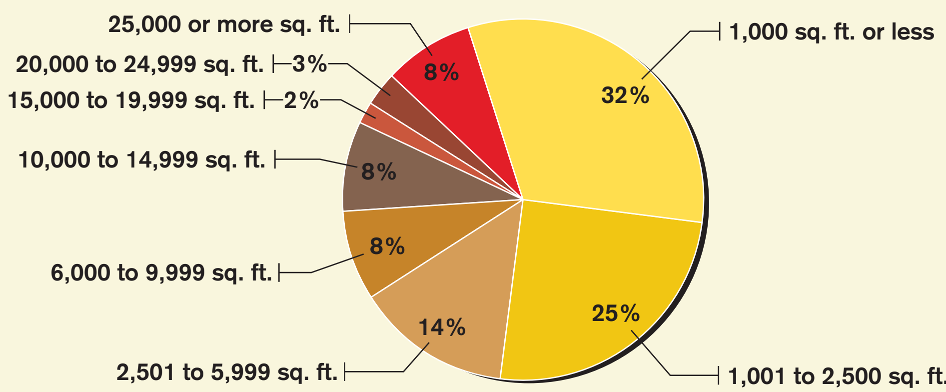

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

Figure 19

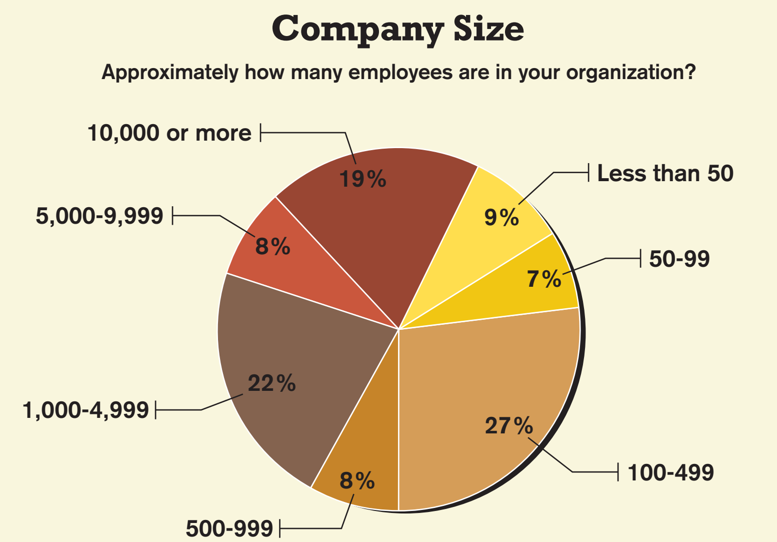

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

Figure 20

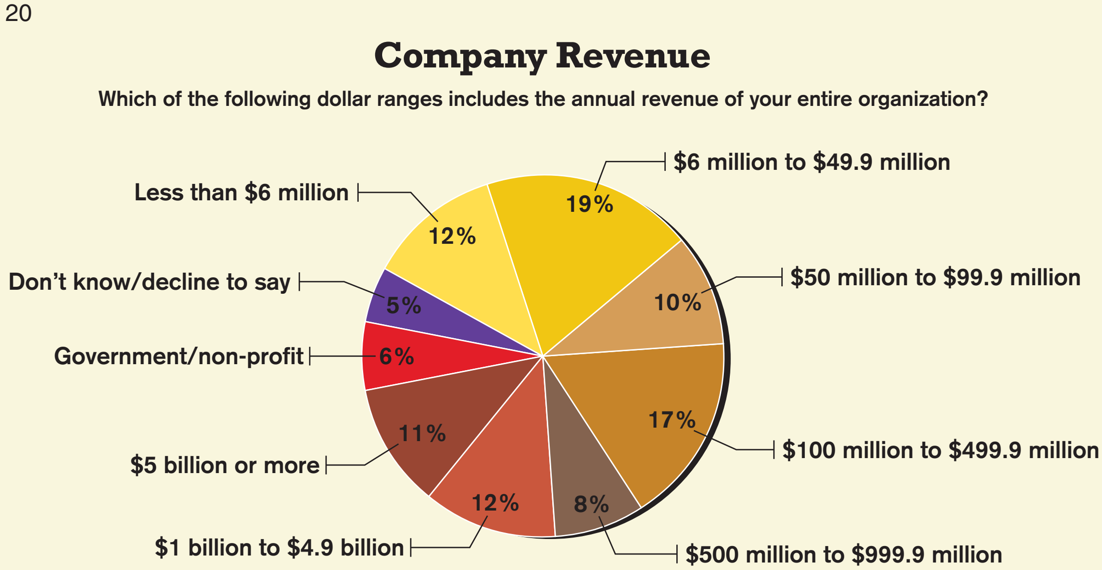

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

Figure 21

Industry

What is your organization's primary industry?

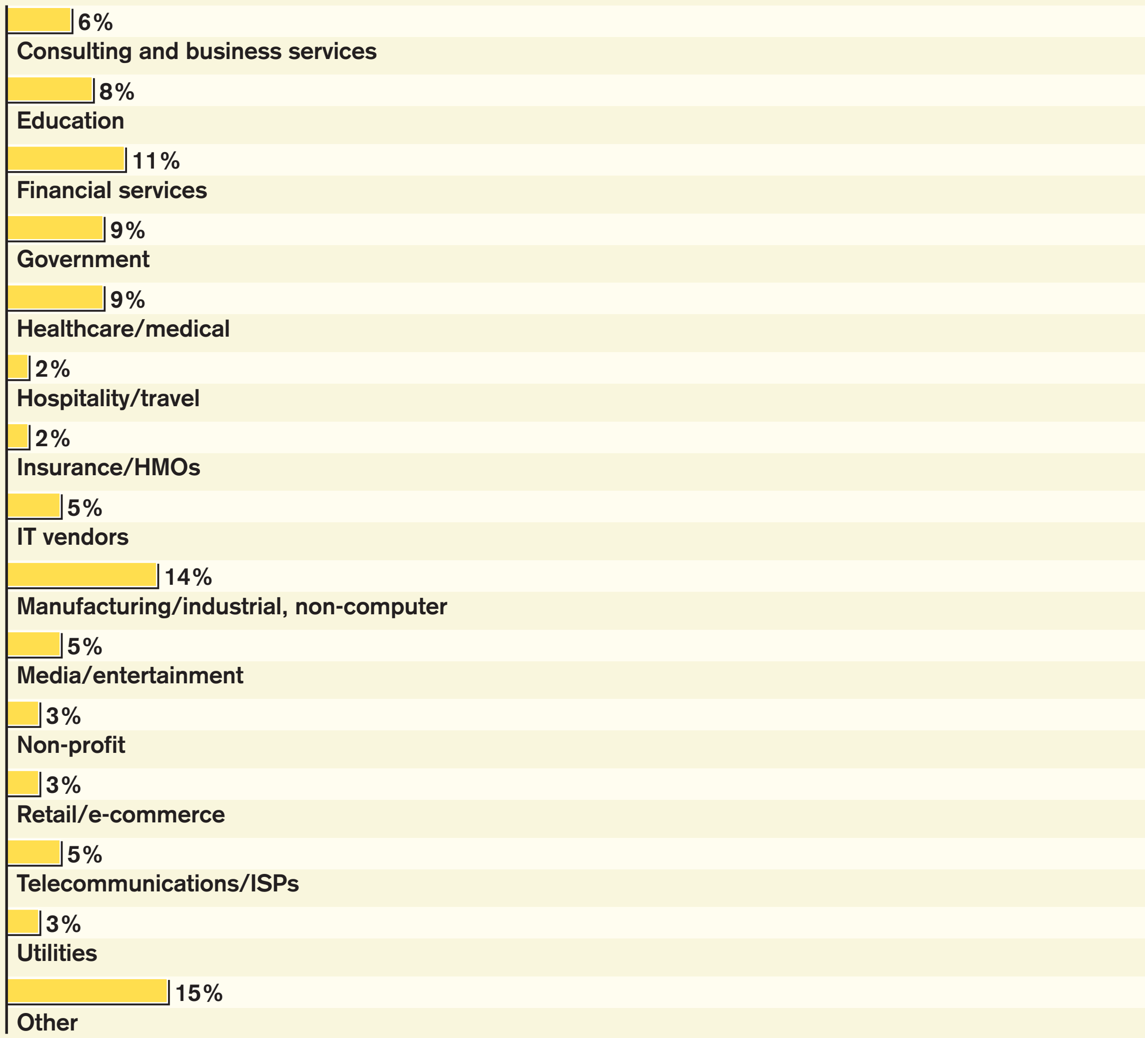

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals

Figure 22

Job Title

Which of the following best describes your job title?

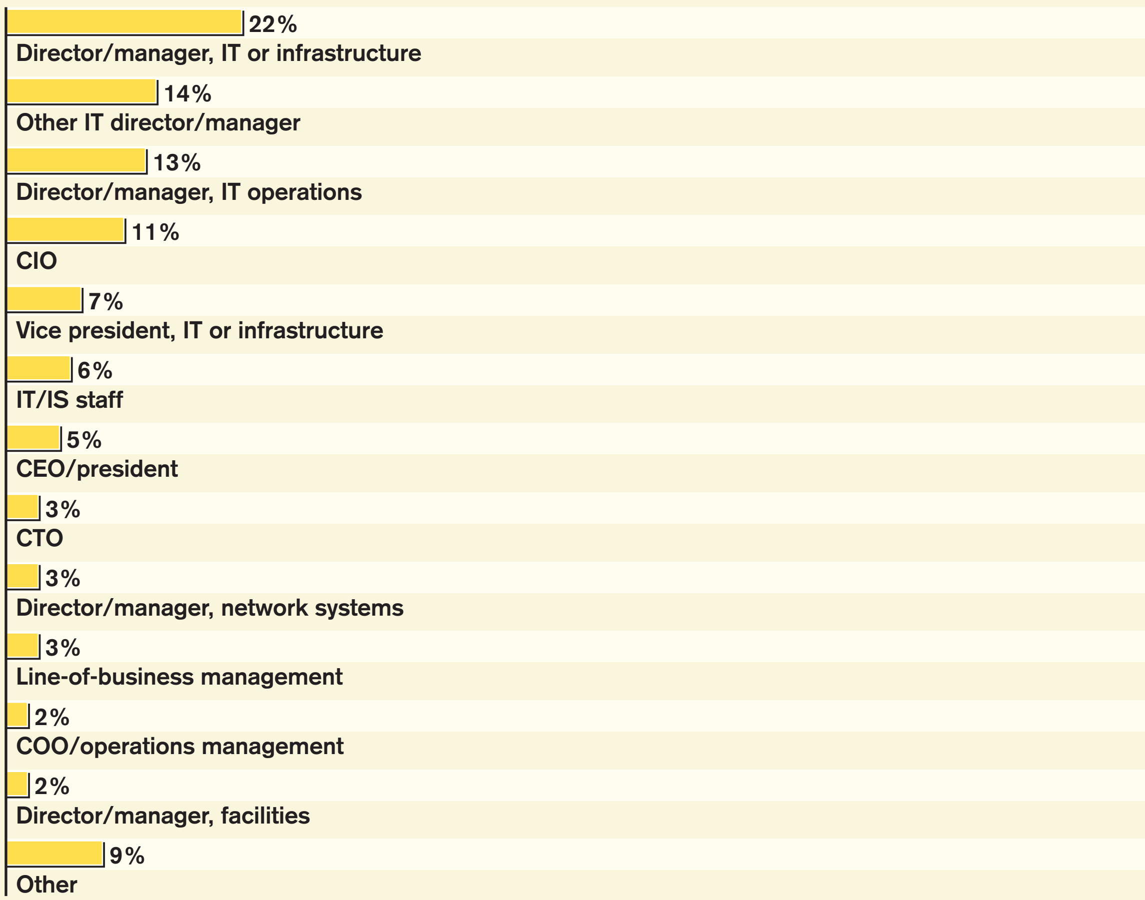

Data: *InformationWeek* Analytics Data Center Survey of 279 business technology professionals