ISSUE BRIEF

AUGUST 2014
IB:14-08-A

America's Data Centers Are Wasting Huge Amounts of Energy

Critical Action Needed to Save Billions of Dollars and Kilowatts

Data centers have become the backbone of the U.S. economy, powering businesses, communications, and online consumer services and helping make our society more productive and efficient. All of our online activity—including email, social media, and conducting business—is delivered through data centers. Ranging from small computer server rooms to mammoth "server farms," these data centers house an estimated 12 million computer servers for data processing, data storage devices, and networking machines.

Data centers also are among the nation's largest and fastest-growing population of consumers of electricity. In 2013 alone, nearly 3 million computer rooms used enough electricity to power all of the households in New York City for two years, equivalent to the annual output of 34 large coal-fired power plants. While the data center industry has made progress in cutting energy waste, an analysis by the Natural Resources Defense Council in partnership with Anthesis finds up to 30 percent of servers are "comatose" and no longer needed, other machines are grossly underutilized, and a number of strategic and tactical barriers still remain. This report recommends steps to accelerate the pace and scale of energy savings across the industry, including the use and public disclosure of energy efficiency metrics.

If just *half* of the savings potential from adopting energy efficiency best practices were realized, America's data centers could slash their electricity consumption by as much as 40 percent.[1] In 2014, this represents a **savings of $3.8 billion and 39 billion kilowatt-hours**, equivalent to the annual electricity consumption of all the households in the state of Michigan.

NRDC
THE EARTH'S BEST DEFENSE

Anthesis

For more information, go to our full report at
www.nrdc.org/energy/data-center-efficiency-assessment.asp

DATA CENTERS ON TRACK TO CONSUME EVEN MORE ELECTRICITY

From the server rooms supporting small- to medium-sized organizations, to the data centers running large corporations, to the server farms hosting Amazon, Facebook, Google and others' Internet-based "cloud" computing services, nearly 3 million U.S. data centers help power our economy, make our lives easier, and render our buildings and electricity grid "smarter."[2]

However, the continuing explosion of digital content, big data, e-commerce, and Internet traffic puts data centers among the fastest-growing users of electricity in the United States and makes them a key driver in the construction of new power plants. U.S. data centers are on track to consume roughly 140 billion kilowatt-hours of electricity annually by 2020, equivalent to the output of 50 large power plants (each with 500 megawatts capacity). Already, if worldwide data centers were a country, they would be the globe's 12th-largest consumer of electricity, ranking somewhere between Spain and Italy.[3]

It is therefore crucial that we clearly understand the current energy efficiency of data centers as a whole. NRDC and Anthesis surveyed more than 30 industry stakeholders and experts, and reviewed the latest industry literature to find the answers to these key questions:

- How efficient are typical data centers in the United States?
- What key opportunities exist to save energy?
- What are the barriers to capturing these opportunities?

Our analysis, *Data Center Efficiency Assessment—Scaling Up Energy Efficiency Across the Data Center Industry: Evaluating Key Drivers and Barriers*, found that much of the energy efficiency progress over the past five years has been achieved in the buildings that house the IT (information technology) equipment and their cooling systems, and that substantially more savings are possible in the operation of the IT equipment itself. We identified two key opportunities for U.S. data centers:

Increasing server utilization. The more work a server performs, the more energy efficient it is—just as a bus uses much less gasoline per passenger when ferrying 50 people than when carrying just a handful. However, the average server operates at no more than 12 to 18 percent of its capacity while still drawing 30 to 60 percent of maximum power. Even sitting virtually idle, servers draw power 24/7, which adds up to a substantial amount of energy use. To put this in perspective, **much of the energy consumed by U.S. data centers is used to power more than 12 million servers that do little or no work most of the time**. Our analysis suggests increased server utilization remains one of the industry's largest energy-saving opportunities, but several barriers hinder progress.

Aligning incentives between those who make decisions affecting efficiency and those who pay the energy bills. In most organizations, the department responsible for data center management is separate from the one paying the electric bills, and they have different agendas. This issue of split incentives is spreading outside organizations with the rapid growth of multi-tenant data centers where customers lease space and power to run computing equipment rather than manage their own centers. Aligning incentives for data center efficiency both within and between organizations would remove many obstacles to implementing commonsense efficiency measures.

SERVER FARMS NOT THE WHOLE STORY

U.S. data centers devoured an estimated 91 billion kilowatt-hours of electricity in 2013, and annual consumption is projected to increase by roughly 47 billion kilowatt-hours by 2020. This would cost businesses $13 billion annually and be equivalent to the amount of electricity generated by 50 large coal-fired power plants emitting nearly 150 million metric tons of carbon pollution.

Table 1: Estimated U.S. data center electricity consumption by market segment (2011)

Segment	Number of Servers (million)	Electricity Share	Total U.S. Data Center Electricity Use (billion kWh/y)
Small and Medium Server Rooms	4.9	49%	37.5
Enterprise/Corporate Data Centers	3.7	27%	20.5
Multi-Tenant Data Centers	2.7	19%	14.1
Hyper-Scale Cloud Computing	0.9	4%	3.3
High-Performance Computing	0.1	1%	1.0
Total (rounded)	**12.2**	**100%**	**76.4**

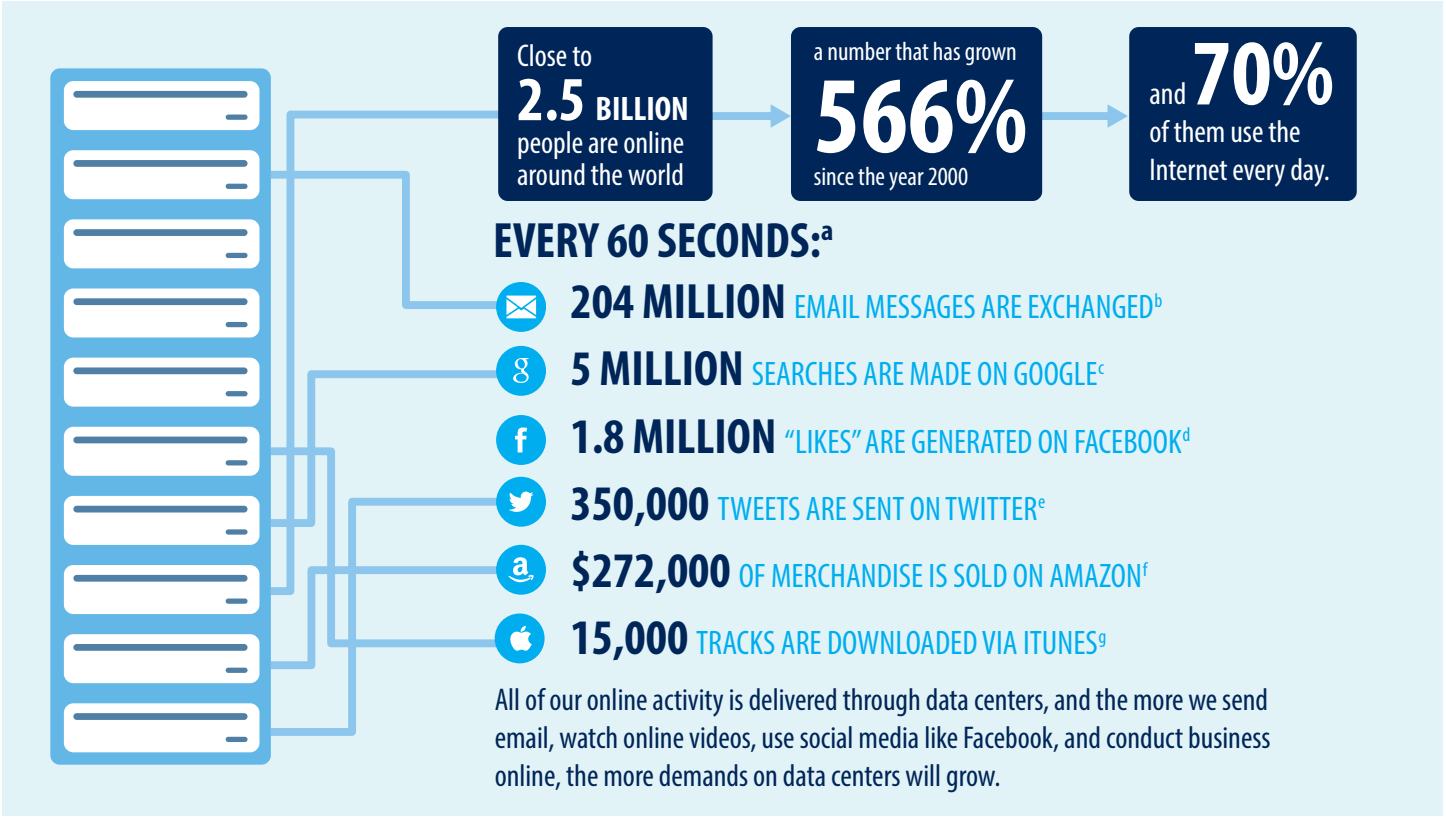

a As of 2013
b http://www.domo.com/learn/infographic-data-never-sleeps
c http://www.statisticbrain.com/google-searches/
d http://gizmodo.com/5937143/what-facebook-deals-with-everyday-27-billion-likes-300-million-photos-uploaded-and-500-terabytes-of-data
e http://www.internetlivestats.com/twitter-statistics/
f http://www.domo.com/learn/infographic-data-never-sleeps
g http://www.billboard.com/biz/articles/news/1538108/itunes-crosses-25-billion-songs-sold-now-sells-21-million-songs-a-day

There has been significant progress in data center energy efficiency in the past decade, with server farms operated by companies like Google and Facebook leading the way. But these hyper-scale cloud computing companies that have moved computing applications to an Internet-based platform represent only approximately 5 percent of all data center servers.[4] The less-visible enterprise (corporate-owned), small- and medium-sized organizations (SMO), and multi-tenant data centers lag far behind in efficiency. As NRDC found in its 2012 report *Is Cloud Computing Always Greener? Finding the Most Energy and Carbon Efficient Information Technology Solutions for Small- and Medium-Sized Organizations*, smaller server rooms and closets are responsible for about half of U.S. server electricity consumption, with much of that wasted due to lack of awareness and inducements for efficiency. There is a critical need for action, including developing utility incentive programs to reduce waste in the massive amounts of electricity used by data centers of all sizes.

PERSISTENT ISSUES SLOW PROGRESS ON EFFICIENCY

NRDC's latest assessment finds that beyond the largest companies providing cloud computing services, data center energy efficiency progress has slowed. The key persisting issues include:

- **"Comatose" servers.** An estimated 20 to 30 percent of servers in large data centers are obsolete or unused—because projects have ended or business processes have changed—but are still plugged in and consuming electricity 24/7. Operators may not realize these devices are no longer needed or may not want to decommission them and risk negative impacts on business operations.
- **Peak provisioning.** Organizations install sufficient equipment to handle their peak annual load, and then some, resulting in vast underutilization of this equipment for the majority of the time.
- **Limited deployment of virtualization technology.** It used to be common practice for IT operations to run each computing application on a dedicated server. Over the past few years, however, virtualization technology has allowed multiple applications to be run, each in its own virtual environment, on a single physical server. This utilizes more capacity on each server and greatly reduces the number

Figure 1: Estimated U.S. data center electricity consumption by market segment (2011)

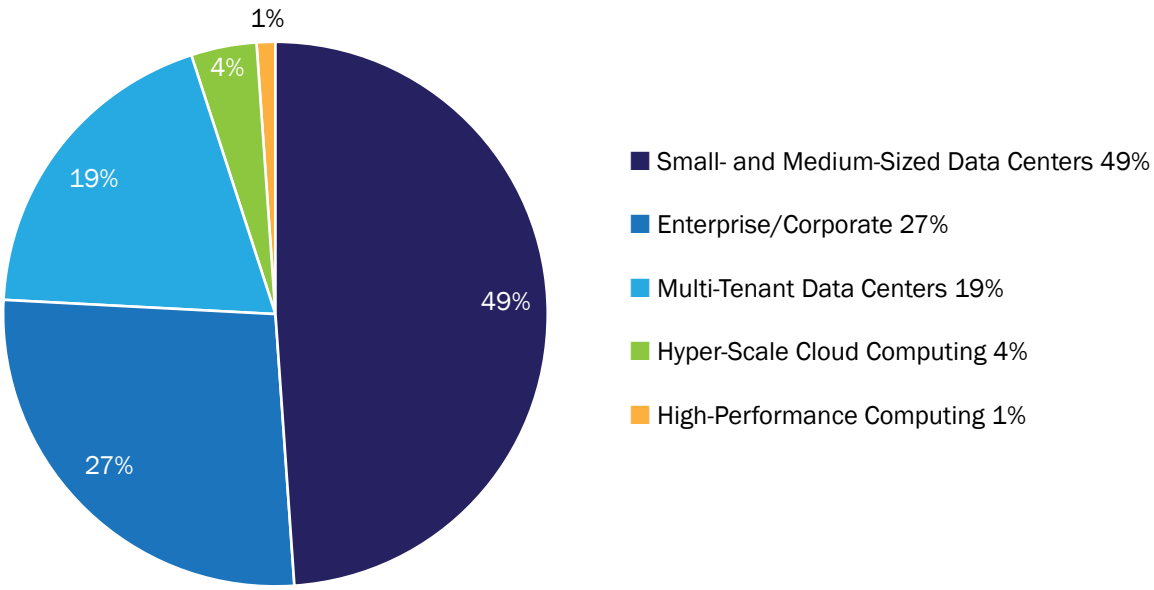

of servers—and the amount of energy—required to run a given number of applications. While most organizations are now familiar with this technology and have deployed it at least partially, few have adopted it at scale across most of their server fleet, particularly in production environments.

- **Failure to power down unused servers.** Many data center operators leave unused servers up and running 24/7, fearing that if they are powered down, they might not restart or might take too long to respond when needed. However, variations in demand can now be managed safely through planning and technology, as the operators of the electrical grid have shown. The data center industry should follow the lead of the utility industry, which ramps its power plants up and down depending on demand.
- **Shortsighted procurement practices.** When taking into account the full cost of operating servers and other IT equipment, including purchase price and lifetime energy costs, it can be significantly cheaper to buy the more efficient models, such as those certified by the Environmental Protection Agency's ENERGY STAR™ program, even if the initial price is higher. Unfortunately, most procurement personnel consider only the upfront purchase price rather than the total cost of ownership.

MULTI-TENANT DATA CENTERS POSE NEW CHALLENGES

The multi-tenant data center business model has become increasingly popular over the past decade. In fact, this outsourcing approach has become the "new normal" for many companies looking to save money in their computing operations.[5] Such data centers are seeing rapid year-to-year growth rates of 18 to 20 percent, leading to projections that they will make up more than one-fourth of all data center capacity by 2016. However, they have not received the attention they deserve by stakeholders concerned with energy consumption. Along with the challenges faced by nearly all data centers, multi-tenant operations are subject to unique challenges that make deployment of energy efficiency more difficult.[6,7] They include:

- **Competing priorities.** Multi-tenant data center operators focus primarily on keeping costs low and maintaining high levels of security, reliability, and uptime for their clients' mission-critical applications, and this focus undermines interest in energy efficiency.[8] Many interviewees noted that the ability to offer an energy efficient facility was a low priority for customers, if it was a priority at all.
- **Exacerbated split incentives.** In multi-tenant data centers, the people paying the power bills and those managing the IT equipment work for different companies. Without contracts whose pricing reflects the costs of powering and cooling the equipment, there is little motivation for data center customers to invest in more-efficient equipment.[9] A retail provider of multi-tenant services commented that there is a perception that it would cost extra time and money to monitor equipment for energy efficiency, but consolidating space and power capacity to allow for more customers would save money for everyone.
- **Challenges with efficiency incentive programs.** Utilities are often restricted to offering efficiency incentives to their customers of record, in this case the multi-tenant service provider. This means they cannot directly incentivize the service provider's customers to improve the efficiency of their IT operations.

IT IS TIME TO ALIGN INCENTIVES FOR ALL DATA CENTERS

Only 20 percent of IT departments pay the data center power bill, a statistic unchanged in more than five years, and one that represents a major barrier to implementing efficiency.[10] In addition, IT and facilities departments are often situated in different parts of an organization, creating a division of accountability and incentives for using data center energy as efficiently as possible. The problem is even more acute with multi-tenant data centers, with the data center owner paying the electricity bill, the tenants paying for blocks of power regardless of use, and IT purchasers separately specifying which equipment tenants should buy.

RECOMMENDATIONS

A series of critical actions is necessary for faster adoption of efficiency-focused technologies, metrics, and initiatives in all data centers, including:

- **Adoption of a simple server utilization metric.** The data center industry should adopt a simple metric, such as the average utilization of the server central processing unit(s) (CPU), to help resolve one of the biggest efficiency issues in data centers: underutilization of servers. Measuring and reporting CPU utilization is a simple, affordable, and adequate way of gauging data center efficiency that could be used immediately to drive greater IT energy savings in data centers.
- **Alignment of incentives between decision-makers.** Data center operators, service providers, and multi-tenant customers should review their internal organizational structure and external contractual arrangements and ensure that incentives are aligned to provide financial rewards for efficiency best practices. Multi-tenant data center stakeholders should develop a "green lease" contract template to make it easier for all customers to establish contracts that incentivize, rather than stand in the way of, energy savings. While the standard contract does not promote efficiency today, nearly one-third of all leased data center space in the United States will come up for renewal over the next year, with nearly all of those leases being based on square footage.[11] If data center providers and their customers are going to recognize together the cost savings of better efficiency, the lease turnover provides an optimum opportunity to do so.
- **Disclosure of data center energy and carbon performance.** Public disclosure is a powerful mechanism by which to demonstrate leadership and drive behavior change across an entire sector. Industry leaders in data center efficiency should voluntarily disclose in their corporate and social responsibility reports a broad range of information and metrics on their entire data center operations, including fleet-wide utilization levels and ways in which they address split incentive issues internally and externally.

CONCLUSION

The Internet has created myriad new opportunities for society and the economy, but the data center industry supporting it has become one of America's fastest-growing sectors for electricity consumption. While the industry at large continues to make progress with many of the challenges to efficiency identified in this report, there remain a number of strategic and tactical barriers that need to be prioritized to curb data center energy growth, particularly in server utilization and within the multi-tenant data center market.

If the recommendations resulting from this survey were adopted and just half of the technical savings potential for data center efficiency were realized (to take into account the market barriers discussed in this report), it could save U.S. businesses $3.8 billion while avoiding the need to generate 39 billion kilowatt-hours of electricity in 2014—enough to power all of Michigan's homes for a year.

Endnotes

1 E.R. Masanet et al., "Estimating the Energy Use and Efficiency Potential of U.S. Data Centers," *Proceedings of the IEEE* 99, No. 8 (August 2011): 1440-1453.

2 Michelle Bailey et al., Data Center of the Future, IDC Special Study, April 2006, IDC #06C4799.

3 DCD Intelligence Report, *Powering the Data Center*, April 2013, http://www.dcd-intelligence.com/Products-Services/Powering-the-Datacenter.

4 Mirko Lorenz, *How Many Servers Worldwide?,* Vision Cloud, August 2011, www.visioncloud.eu/content.php?s=191,324.

5 Uptime Institute, *2013 Uptime Institute Data Center Industry Survey*, www.uptimeinstitute.com/images/stories/DataCenterSurvey_assets/uptime-institute-2013-data-center-survey.pdf (accessed November 2013).

6 David Freeland, *Colocation and Managed Hosting*, Focus Investment Banking, 2012, http://www.focusbankers.com/telecom/pdfs/FOCUS_Colocation_Winter12.pdf.

7 451 Research, *Multi-Tenant Data Center Global Providers—2012*, November 2012.

8 Gary Cullen et al., *Evaluation, Verification, and Measurement Study FY 2008/2009 Program for Silicon Valley Power*, Summit Blue Consulting, 2009, www.ncpa.com/images/stories/LegReg/SVP% percent20-% percent20EM&V% percent20Report% percent203-2-10% percent20final.pdf.

9 Ibid.

10 Mirko Lorenz, *How Many Servers Worldwide?*

11 Rich Miller, *Data Center Leasing: It's All About the Megawatts*, Data Center Knowledge, November 11, 2009, www.datacenterknowledge.com/archives/2009/11/11/data-center-leasing-its-all-about-the-megawatts.

For more information contact **Pierre Delforge at pdelforge@nrdc.org**

Natural Resources Defense Council
40 West 20th Street
New York, NY 10011
212 727-2700
Fax 212 727-1773

Beijing

Chicago

Los Angeles

Bozeman

San Francisco

Washington, D.C.

www.nrdc.org

www.nrdc.org/energy
www.twitter.com/nrdcEnergy
www.twitter.com/nrdc
www.facebook.com/nrdc.org

Printed on recycled paper