

GROWTH IN DATA CENTER ELECTRICITY USE 2005 TO 2010

Jonathan G. Koomey, Ph.D.

Consulting Professor, Stanford University

Contact: JGKoomey@stanford.edu, <http://www.koomey.com>

A report by Analytics Press, completed at the request of *The New York Times*

<http://www.analyticspress.com/datacenters.html>

August 1, 2011

EXECUTIVE SUMMARY

This study assesses growth in data center electricity use from 2005 to 2010 for the US and the world. It uses data and methods similar to earlier analyses to create a consistent time series of data center electricity use. The key results for the analysis are as follows:

- Growth in the installed base of servers in data centers had already begun to slow by early 2007 because of virtualization and other factors.
- The 2008 financial crisis, the associated economic slowdown, and further improvements in virtualization led to a significant reduction in actual server installed base by 2010 compared to the IDC installed base forecast published in 2007.
- Growth in electricity used per server probably accounted for a larger share of demand growth from 2005 to 2010 than it did in 2000 to 2005.
- Assuming that the midpoint between the Upper and Lower bound cases accurately reflects the history, electricity used by data centers worldwide increased by about 56% from 2005 to 2010 instead of doubling (as it did from 2000 to 2005), while in the US it increased by about 36% instead of doubling.
- Electricity used in global data centers in 2010 likely accounted for between 1.1% and 1.5% of total electricity use, respectively. For the US that number was between 1.7 and 2.2%.
- Electricity used in US data centers in 2010 was significantly lower than predicted by the EPA's 2007 report to Congress on data centers. That result reflected this study's reduced electricity growth rates compared to earlier estimates (see **Figure ES-1**), which were driven mainly by a lower server installed base than was earlier predicted rather than the efficiency improvements anticipated in the report to Congress.
- While Google is a high profile user of computer servers, less than 1% of electricity used by data centers worldwide was attributable to that company's data center operations.

In summary, the rapid rates of growth in data center electricity use that prevailed from 2000 to 2005 slowed significantly from 2005 to 2010, yielding total electricity use by data centers in 2010 of about 1.3% of all electricity use for the world, and 2% of all electricity use for the US.




Figure ES-1: Predicted US electricity use for data centers from the EPA report to Congress (EPA 2007) and the range estimated in this study

INTRODUCTION

Revolutionary technologies always capture the popular imagination. At the dawn of the twentieth century it was the telephone, the airplane, the automobile, radioactivity, and the electric power system. Then came antibiotics, radio, television, nuclear weapons, nuclear power, and space travel.

By the end of the twentieth century, the most pervasive world-changing technology was the computer, which revolutionized modern society. Computers land our airplanes, run our businesses, control our industrial processes, and help us keep us in constant touch. Information is now easier to collect, find, and use than ever before. It is also more plentiful than at any time in human history (Hilbert and López 2011).

The importance of computers has prompted many observers to investigate the unintended consequences of their use. Some have assessed the environmental impacts of computer manufacturing and recycling, while others have analyzed ergonomic issues associated with computer use or the ways that our brains change when we interact with computers. The side effect generating some of the deepest popular interest, however, is the amount of electricity used by computers and related information technology (IT) equipment.

The most recent detailed work on electricity used by all IT equipment was created in the year 2000 and soon afterwards in response to erroneous claims of huge IT electricity demand put forth in an article in *Forbes* (Huber and Mills 1999). Those claims were erroneous in every important respect, overestimating IT electricity use by substantial margins, but they prompted the creation of research studies that are still useful to this day (Kawamoto et al. 2002, Koomey 2003, Koomey et al. 2002, Koomey et al. 2004, Mitchell-Jackson et al. 2003, Roth et al. 2002).

More recently, the electricity used by the high-density computing facilities that power the Internet (i.e., data centers) has generated the most intense interest, in part because of the importance of these facilities to the broader economy, in part because the power used by individual data centers can rival some industrial facilities, and in part because the total electricity used by these facilities has been growing rapidly (Koomey 2008a, Masanet et al. 2011).

So just how fast has data center power use been growing? Koomey (2007a, 2007b, 2008b) analyzed trends in worldwide and regional data center power use from 2000 to 2005, showing roughly a doubling in power used by these facilities over this period. Those papers also showed that this growth was mostly driven by growth in the numbers of so-called “volume servers”, which dominate server sales, and that growth in power use per server was a secondary contributor to total growth in power use.

This short report assesses what’s happened since then. It focuses on total electricity used by data centers from 2005 to 2010, and creates those estimates in a way that is consistent with the earlier work. While not as detailed as the previous work, it uses some of the same data sources to make the 2010 estimates comparable to the earlier data.

DATA AND METHODS

To estimate total electricity used by data centers¹ in 2010, the most straightforward approach would have been to use the same methods documented in the earlier studies by Koomey (2007a, 2007b, 2008b) with updated data from IDC on the installed base of servers in 2010. Since those studies came out, IDC has changed some of its data and reporting assumptions so I was not able to easily create a consistent time series in trends in power use per server. I'm working on how to address that issue, but in the meantime, I felt it was important to understand what has happened in the data center market since 2005 as best I could, so I choose a simplified approach for this analysis, anticipating that more detailed analyses will follow the publication of this report.

Servers are still the largest and most important electricity consumer in data centers, so I begin there (as I did in the previous analysis). First, I calculated the total electricity used by servers as the product of the installed base of servers by server class (from new IDC data) and the electricity used per server per year for each scenario I considered (see below for definitions). To the server electricity use total I added electricity used for data storage, communications, and infrastructure equipment (cooling, fans, pumps, and losses in the backup and power distribution systems), based on aggregate estimates of the electricity intensity of these end-uses.

Trends in installed base

Table 1 shows data on the installed base of servers from IDC's 2007 and 2010 forecasts. Because of various data issues I needed to define several different categories of installed base and shipment statistics. The terms "old" and "new" in Table 1 refer to IDC installed base estimates from February 2007 and January 2010, respectively. The term "adjusted" refers to the "new" installed base estimates for 2010 adjusted to reflect updated historical shipment statistics for 2010 (the original "new" estimates were a forecast for 2010). It is unclear how the year 2000 data would be affected by the changes made by IDC in the new installed base forecast since that analysis does not extend back to 2000, but the installed base numbers for 2005 differ by up to 8% between the old and new installed base figures (depending on server class).

Defining scenarios for 2010

For 2010, I created four scenarios to help illustrate the trends occurring in the data center market. First, I show what would have happened if trends in server installed base and electricity used per server continued to grow at the rates prevailing from 2000 to 2005

¹ There are many different ways to define data centers. For this analysis I follow IDC (Bailey et al. 2007) and EPA (2007) by including as data centers any space whose main function is to house computer servers, including data closets and server rooms. This definition does not include the parts of telecommunications central offices whose primary purpose is to house telephone switches, routers, and other large-scale network equipment (the parts of those facilities devoted to housing standard servers for voice over Internet protocol applications would be included in this definition).

(All Trends Continued). Then I show an estimate of 2010 data center electricity use based on the 2007 IDC installed base forecast and the 2000 to 2005 trends in electricity used per server (**Best Guess 2007**). Then I show two scenarios that bound the likely range in 2010. The **Upper Bound Case** is based on the 2010 adjusted installed base from the latest IDC installed base forecast and assumes that the 2000 to 2005 trends in electricity used per server continue to 2010. Finally, the **Lower Bound Case** uses the same installed base numbers for 2010 as in the Upper Bound Case but assumes that there is no growth in electricity used per server from 2005 to 2010 (which is an attempt to reflect the industry's intense focus on server efficiency starting in 2006).

Keeping the analysis consistent with the previous work

The discrepancy in the 2005 installed base figures forced me to simplify the analysis approach compared to the previous work. For the Lower and Upper bound cases I applied the 2005 to 2010 *growth rates* in installed base taken from the new (Jan. 2010) adjusted forecast to the installed base figures for 2005 from the old (Feb. 2007) forecast. This approach allowed me to assess trends in a way consistent with my earlier analysis of data center electricity use (Kooimey 2008b).

Storage and communications electricity use

In my 2008 study that analyzed data center power use in 2000 and 2005 (Kooimey 2008b), I wrote:

Assuming that the IDC forecast of worldwide server installed base is correct, that trends in electricity use per server continue to 2010 as they did from 2000 to 2005, and that the networking and storage electricity use figures continue to track server electricity use as assumed for the historical data, the total growth in data center electricity use would be 76%.

This simple scenario did not include a change in the trends in electricity used by storage and communications equipment that was embodied in the forecast for 2010 from the EPA Report to Congress (US EPA 2007). In that forecast, both storage and communications were projected to grow more rapidly than historical trends. To capture that change, I used the EPA's projected ratio of storage (or communications) electricity use to total server electricity use in 2010 in their "Current Efficiency Trends" case to estimate storage or communications electricity use for my four 2010 scenarios.²

² The details of this calculation are described in Tables 2 and 3, below.

Changes in infrastructure efficiency

For the All Trends Continued and the Best Guess 2007 cases (as in the previous studies) I assumed an average Power Usage Effectiveness (PUE)³ of 2.0, which implies that for every kWh of IT electricity use there is another kWh of infrastructure electricity use. This assumption may result in a slight overestimate of electricity used because of the increasing prevalence of data centers performing cloud computing (because those providers have been able drastically reduce infrastructure electricity use compared to in-house data centers), but the in-house facilities still dominate the industry in terms of numbers of servers. That assumption will need to be revisited as cloud computing becomes more widely used. For the Upper Bound case in 2010 I assumed an average PUE of 1.92, and for the Lower Bound case I assumed an average PUE of 1.83, reflecting progress in this parameter since the earlier analyses were conducted.⁴

RESULTS

I summarize and discuss key findings from the analysis below.

Installed base

The installed base numbers in **Table 1** show some interesting trends. The dominant driver of electricity demand from 2000 to 2005 was growth in the installed base of volume servers, which doubled over that five-year period for the US and the world. The IDC installed base forecast from 2007 showed slower growth in volume servers to 2010 compared to the 2000 to 2005 period, and the newer forecast reduced installed base growth still further. Considering the latest installed base estimates for 2010, the growth in volume servers slowed considerably in the 2005 to 2010 period, growing only about 20% in the US and about one-third in the world. The installed base of mid-range servers fell even faster than expected in the earlier forecast, but high-end servers grew rapidly instead of declining as had been projected earlier.

³ PUE = Total data center power use/Information technology equipment power use

⁴ I based my choices for this parameter on measured data, although the question of how well they represent actual historical developments is ultimately an empirical one that can only be answered with further research. The EPA's Energy Star program set up a measurement protocol for PUE <<http://www.energystar.gov/datacenters>> and found a mean PUE for 61 data centers in 2010 of 1.92, and a range from 1.36 to 3.6. The First Annual Uptime Institute Data Center Industry Survey <<http://uptimeinstitute.org/content/view/302/281/>>, conducted in 2011 showed an average PUE of about 1.83. In the case of Energy Star, the sample is self-selected, so it is likely that there's some bias in the reported values (operators with less favorable PUEs are less likely to participate in the survey). The Uptime Institute Network membership, which responded to the second survey, represents operators with above average sophistication, which is why I chose that value as the lower bound number.

Electricity used by data centers

Table 2 shows estimated worldwide electricity used for data centers for 2000, 2005, and for the four scenarios in 2010, and **Table 3** shows the same results for the US. The All Trends Continued case shows more than a doubling of data center electricity use 2005 to 2010, while the Best Guess 2007 case shows an increase over 2005 levels of 93 to 98%. The Upper and Lower Bound cases, which are based on historical installed base statistics (estimated as described in Table 1), show lower electricity use for 2010 than projected in 2007, partly because of installed base in 2010 being lower than expected (in both cases) and partly because of reductions in growth in electricity used per server (in the Lower Bound case).

The actual electricity used by data centers in 2010 will likely fall in between the Upper and Lower Bound cases because growth in electricity used per server has almost certainly slowed since 2005. The data center industry really started to focus on efficiency in early 2006, which is when the US EPA sponsored two public meetings on efficiency with representatives of all the major players in the industry in attendance.⁵ There are also some anecdotal trends data indicating that growth in electricity used per server slowed since 2005 (Kooimey et al. 2009), but further research is needed to accurately quantify and generalize these trends.

Comparing projected to actual electricity use, 2005 to 2010

The 2007 EPA report to Congress on data centers (US EPA 2007) predicted a little less than a doubling in total data center electricity use from 2005 to 2010 if historical trends continued. **Figure 3** plots the graph from that study with the range spanned by the Lower and Upper Bound estimates from this study for the US. That range coincidentally overlaps with the EPA report's "Improved operations" scenario, but the main reason for the lower estimates in this study is the much lower IDC installed base estimates, not the significant operational improvements and installed base reductions from virtualization assumed in that scenario. Of course, some operational improvements are captured in this study's new data (in the form of PUEs less than 2.0) but they are not as important as the installed base estimates to the results.

The case for 2010 labeled "Best Guess 2007" approximates the expectations embodied in the EPA report to Congress, and it shows 98% growth for the world compared to 2005 (Table 2). If we just correct the worldwide installed base numbers for 2010 using the latest IDC forecast (Table 1) and keep the trends in electricity use per server the same as for 2000 to 2005, we end up with the "Upper Bound" scenario, which shows about a 78% increase in electricity used by data centers worldwide from 2005 to 2010.

The effect of the revised 2010 installed base data is larger for the US, which moves from 93% growth (Best Guess 2007) to 53% growth (Upper Bound Case). It would appear

⁵ <http://www.energystar.gov/datacenters>

that the financial crisis had a larger effect on the US data center market compared to earlier expectations than it did on the world market.

The trends in electricity use per server are likely to be more important to the results in 2010 than they were during the period 2000 to 2005. In the earlier study (Kooimey 2008b) I found that most of the growth in data center electricity use was associated with growth in the installed base of volume servers from 2000 to 2005, with the growth in electricity use per server contributing modestly. Depending on what happened with electricity use per server from 2005 to 2010, that parameter could be as or more important than growth in server installed base over the 2005 to 2010 period. That result implies that more research is needed to assess what the trends in electricity use per server actually were over this period.

Rough estimate of electricity used by Google's data centers

Google's servers are not included in the IDC installed base data because the company assembles its own custom servers (and has been doing so for a long time). A few other companies have recently started down this path, but Google is likely the biggest one to do so. How big an omission is this?

Table 4 makes some educated guesses about Google's servers to estimate electricity used for that company's data centers over time. While there is substantial uncertainty in these estimates (because of the lack of data on the installed base and other characteristics of Google's servers), the calculations (and an explicit upper bound on Google's data center electricity use supplied by that company)⁶ show that Google's data center electricity use is about 0.01% of total worldwide electricity use and less than 1 percent of global data center electricity use in 2010. This result is in part a function of the higher infrastructure efficiency of Google's facilities compared to in-house data centers, which is consistent with efficiencies of other cloud computing installations, but it also reflects lower electricity use per server for Google's highly optimized servers.

Summary of results

Figures 1 and 2 summarize the key results from the analysis. In 2005, electricity used by data centers totaled about 1% of world electricity use and 1.5% of US electricity use. By 2010, world data center electricity use represented between 1.1 and 1.5% of world electricity use, while in the US data center electricity use represented between 1.7 and 2.2% of the total. These ranges represent significant reductions from expectations of a few years ago, and reflect the effect of the economic crisis of 2008 and the efforts the

⁶ To my knowledge this is the first time Google has made any specific statements about how much power their data center operations use. David Jacobowitz of Google told me (on the record in a phone conversation on May 13, 2011) that Google's total data center electricity use was less than 1% of the 198.8 billion kWh that the review draft of this report gave as total worldwide data center electricity use in 2010 in the lower bound case. That absolute number for the lower bound case has since changed but I've cited and used the exact statement that Jacobowitz gave to me when I created the calculations in Table 4.

industry has made to improve data center efficiency since 2005. The lower bound figures represent 20 to 33% growth in data center electricity use compared to 2005—the increase in the numbers and per unit power use of servers, storage, and communications were offset somewhat by the reduction in PUE from 2.0 in the lower and upper bound cases.

FUTURE WORK

Server installed base

There is still some uncertainty in the installed base of servers for 2010, for which I've attempted to correct in Table 1 (using the latest shipments for 2010). Once the IDC installed base data is updated to reflect later data for 2010, the adjusted numbers may change further, but probably not by much.

The bigger issue will be how to create a consistent time series going back to 2000. The revised data set from IDC starts in 2003, but as Table 1 shows, there are some differences between the old and new IDC data for 2005. It is not known if similar differences would show up for 2000 if the new database were to be extended backwards to that date, but determining how to adjust the old 2000 data to make them consistent with the new data series should be a high priority for further research.

In addition, the question of “comatose” or orphaned servers has yet to be convincingly addressed. Anecdotal evidence indicates that 10-30% of servers in many data centers are using electricity but no longer delivering computing services. These servers have not yet been decommissioned and are probably not counted in installed base statistics. In many facilities nobody even knows these servers exist, so it is likely at least some of them are not included in the installed base statistics. Actual server electricity use could therefore be higher than estimated here. This issue is one that needs further detailed investigation.

Trends in electricity use per server

As discussed above, trends in electricity use per server could be as or more important than changes in the installed base of servers in determining total data center electricity use in 2010, which suggests that additional analysis of this parameter is warranted. In the earlier analysis I estimated or measured electricity use for the five most widely owned server models for each class and used those data to calculate the average electricity use for servers in each class. Because of the importance of this parameter to the results it makes sense to revisit this method, expanding it to the top ten or fifteen server models and redoing the analysis for 2000, 2005, and 2010. The 2000 installed base by model can be done using the old IDC installed base database, while the 2005 and 2010 numbers could be calculated using the new database. It makes sense to wait until the new IDC database is revised to reflect actual 2010 shipments before undertaking this analysis.

This new approach has the disadvantage of affecting the total energy use estimates from the earlier work in unpredictable ways, but it's the best way I can think of to address the change in the IDC data assumptions and generate a consistent time series. It has the advantage of using a larger number of models to estimate electricity use per server, which will likely result in greater accuracy in assessing this important parameter. One wildcard

in this analysis would be the increasing prevalence of custom designed (“self assembled”) servers that do not appear in the IDC database and that use less electricity than standard servers—that issue will presumably need to be addressed explicitly in future analyses.

Servers have also been changing over time, which complicates matters. Servers in 2010 have much higher processing power, more memory, faster network connections, more components, and bigger power supplies. They also have power management and other clever technologies to reduce electricity consumption. In my conversations with various industry experts for this paper, it became clear that peak power per server (which seems to be increasing) may be diverging from annual electricity use per server (which is growing more slowly than peak power and may actually be declining). This divergence is an important area for further research.

Data storage and communications electricity use

Data storage (principally disk drives) and communications electricity use are treated at a very high level in these calculations, based on US EPA (2007). More analysis and data collection are needed to assess and update the assumptions for future studies. The most urgent need is for measurements of power use for such equipment paired with data on the storage capacity and data access speeds (for disk drives) and on network throughput speeds (for communications technologies). More detail on the installed base of such equipment is also needed, as these devices have not been studied in as much detail as servers.

There is anecdotal evidence that the need for data storage has been growing more rapidly than the need for computing power in data centers, but the data density of storage devices has also been growing rapidly (doubling every year or so according to Grochowski and Halem (2003)). In addition, the total power used by these devices is primarily related to the number of drive spindles, not to the amount of data they hold, so the relationship between total data storage capacity and total energy use is not a simple one. This analysis uses the projected growth in storage electricity use relative to server electricity use from EPA (2007) but more research is needed to confirm if the history actually matched that projection.

Prevalence of cloud computing

Because cloud computing installations typically have much higher server utilization levels and infrastructure efficiencies than do in-house data centers (with PUEs for some specific facilities lower than 1.1) increased adoption of cloud architectures will result in lower electricity use than if the same computing services were delivered using more conventional approaches. While the US EPA has made strides in compiling available data on data center infrastructure efficiencies for in-house facilities⁷ and some cloud

⁷ See the benchmarking tool at <<http://www.energystar.gov/datacenters>>.

providers have released data on their own PUEs,⁸ I'm aware of no publicly available data on what fraction of total data center floor area or servers is associated with cloud computing. Once such data become available, a more accurate analysis of infrastructure electricity use will be possible, since the cloud computing facilities can then be segregated from the in-house facilities in the calculations.

Trends after 2010

The January 2010 IDC forecast shows virtually no growth in installed base from 2010 to 2013, which presages continued slower growth in data center electricity use. The installed base forecast is driven principally by assumptions about virtualization becoming more and more prevalent, thus reducing the need to install more physical servers. That improves energy efficiency by reducing the number of physical servers needed (thus eliminating the “fixed” energy cost of keeping some servers available to deliver computations) and driving up utilization of the servers that remain (which spreads energy use and other costs over more computations per server). Of course, it's important to recall Neils Bohr's dictum that “Predictions are very difficult, especially about the future”, and to take this forecast (and all such forecasts) with a grain of salt.

CONCLUSIONS

World data center electricity use doubled from 2000 to 2005, but that rate of growth slowed significantly from 2005 to 2010. This slowing was the result of the 2008-9 economic crisis, the increased prevalence of virtualization in data centers, and the industry's efforts to improve efficiency of these facilities since 2005. If we take the midpoint between the Low and High cases in this analysis, US and world data center electricity use grew by about 36% and 56%, respectively, from 2005 to 2010, totaling about 1.3% of World electricity use and 2% of US electricity use in 2010. The US data center market appears to have been hit harder than the world market by the economic crisis, with growth slowing more noticeably in that market than in the world as a whole. Finally, while Google is a large user of data centers, that company's facilities represent less than 1 percent of all data center electricity use worldwide.

⁸ See, for example, the data released recently on Facebook's Prineville, OR facility <<https://www.facebook.com/notes/facebook-engineering/building-efficient-data-centers-with-the-open-compute-project/10150144039563920>>, the specifications for Yahoo's Lockport, NY facility <<http://www.greentechmedia.com/articles/read/yahoo-nears-perfection-with-chicken-coop-data-center/>> and the aggregate data released by Google <<http://www.google.com/corporate/datacenter/efficiency-measurements.html>>.

ACKNOWLEDGMENTS

This research was created as an exclusive for *The New York Times*, at the request of *NY Times* reporters Miguel Helft, James Glanz, and John Markoff (Helft has since left the *NY Times* for *Fortune*).

I would especially like to thank Lloyd Cohen and Vernon Turner of IDC for making their data available for this research and for invariably answering my innumerable questions promptly and courteously.

I would also like to thank the technical reviewers for their insights and comments. The reviewers included (in alphabetical order by company):

AT&T: Michele Blazek
Dell: David Moss, John Pflueger
Ecos Consulting: Chris Calwell
EPA: Katharine Kaplan
Facebook: Jason Taylor
Google: David Jacobowitz
Green Grid: Mark Monroe
HP: Paul Perez
IBM: Steve Kinney, Joe Prisco, Roger Schmidt
ICF: Arthur Howard, Rebecca Duff, Cody Taylor
IDC: Lloyd Cohen, Vernon Turner
Intel: Henry Wong, Scott Shull, Winston Saunders
LBNL: Rich Brown, Bruce Nordman, Dale Sartor, Bill Tschudi
Lumina Systems: Max Henrion
McKinsey: William Forrest
MegaWatt Consulting, Inc.: KC Mares
Microsoft: Mark Aggar, Christian Belady, Stephen Berard, Rob Bernard
Rocky Mountain Institute: Amory Lovins
Stanford University: Richard Mount, Phil Reese
Sun Microsystems (former VP): Subodh Bapat
Uptime Institute: Ken Brill, John Stanley
Virginia Tech: Kirk W. Cameron, Wu-chun Feng
Yahoo: Christina Page

REFERENCES

Bailey, M., M. Eastwood, T Grieser, L. Borovick, V. Turner, and R.C. Gray. 2007. *Special Study: Data Center of the Future*. New York, NY: IDC. IDC #06C4799. April.

Grochowski, E., and R. D. Halem. 2003. "Technological impact of magnetic hard disk drives on storage systems." *IBM Systems Journal*. vol. 42, no. 2. pp. 338-346.

Hilbert, Martin, and Priscila López. 2011. "The World's Technological Capacity to Store, Communicate, and Compute Information." *Science*. vol. 332, no. 6025. April 1. pp. 60-65.

Huber, Peter, and Mark P. Mills. 1999. "Dig more coal—the PCs are coming." In *Forbes*. May 31. pp. 70-72.

Kawamoto, Kaoru, Jonathan Koomey, Bruce Nordman, Richard E. Brown, Maryann Piette, Michael Ting, and Alan Meier. 2002. "Electricity Used by Office Equipment and Network Equipment in the U.S." *Energy—The International Journal (also LBNL-45917)*. vol. 27, no. 3. March. pp. 255-269.

Koomey, Jonathan. 2003. "Sorry, Wrong Number: Separating Fact from Fiction in the Information Age." In *IEEE Spectrum*. June. pp. 11-12.

Koomey, Jonathan. 2007a. *Estimating regional power consumption by servers: A technical note*. Oakland, CA: Analytics Press. December 5. <<https://files.me.com/jgkoomey/atrxl>>

Koomey, Jonathan. 2007b. *Estimating total power consumption by servers in the U.S. and the world*. Oakland, CA: Analytics Press. February 15. <<https://files.me.com/jgkoomey/98ygy0>>

Koomey, Jonathan. 2008a. *Turning Numbers into Knowledge: Mastering the Art of Problem Solving*. 2nd ed. Oakland, CA: Analytics Press. <<http://www.analyticspress.com>>

Koomey, Jonathan. 2008b. "Worldwide electricity used in data centers." *Environmental Research Letters*. vol. 3, no. 034008. September 23. <<http://stacks.iop.org/1748-9326/3/034008>>

Koomey, Jonathan, Chris Calwell, Skip Laitner, Jane Thornton, Richard E. Brown, Joe Eto, Carrie Webber, and Cathy Cullicott. 2002. "Sorry, wrong number: The use and misuse of numerical facts in analysis and media reporting of energy issues." In *Annual Review of Energy and the Environment 2002*. Edited by R. H. Socolow, D. Anderson and J. Harte. Palo Alto, CA: Annual Reviews, Inc. (also LBNL-50499). pp. 119-158.

Koomey, Jonathan, Huimin Chong, Woonsien Loh, Bruce Nordman, and Michele Blazek. 2004. "Network electricity use associated with wireless personal digital

assistants." *The ASCE Journal of Infrastructure Systems (also LBNL-54105)*. vol. 10, no. 3. September. pp. 131-137.

Koomey, Jonathan G., Christian Belady, Michael Patterson, Anthony Santos, and Klaus-Dieter Lange. 2009. *Assessing trends over time in performance, costs, and energy use for servers*. Oakland, CA: Analytics Press. August 17. <<http://www.intel.com/pressroom/kits/ecotech>>

Masanet, Eric R., Richard E. Brown, Arman Shehabi, Jonathan G. Koomey, and Bruce Nordman. 2011. "Estimating the Energy Use and Efficiency Potential of U.S. Data Centers." *Proceedings of the IEEE*. vol. 99, no. 8. August.

Mitchell-Jackson, Jennifer, Jonathan Koomey, Bruce Nordman, and Michele Blazek. 2003. "Data Center Power Requirements: Measurements From Silicon Valley." *Energy--The International Journal (also LBNL-48554)*. vol. 28, no. 8. June. pp. 837 - 850.

Roth, Kurt, Fred Goldstein, and Jonathan Kleinman. 2002. *Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings--Volume I: Energy Consumption Baseline*. Washington, DC: Prepared by Arthur D. Little for the U.S. Department of Energy. A.D. Little Reference no. 72895-00. January. <<http://www.eere.energy.gov/>>

US EPA. 2007. *Report to Congress on Server and Data Center Energy Efficiency, Public Law 109-431*. Prepared for the U.S. Environmental Protection Agency, ENERGY STAR Program, by Lawrence Berkeley National Laboratory. LBNL-363E. August 2. <<http://www.energystar.gov/datacenters>>

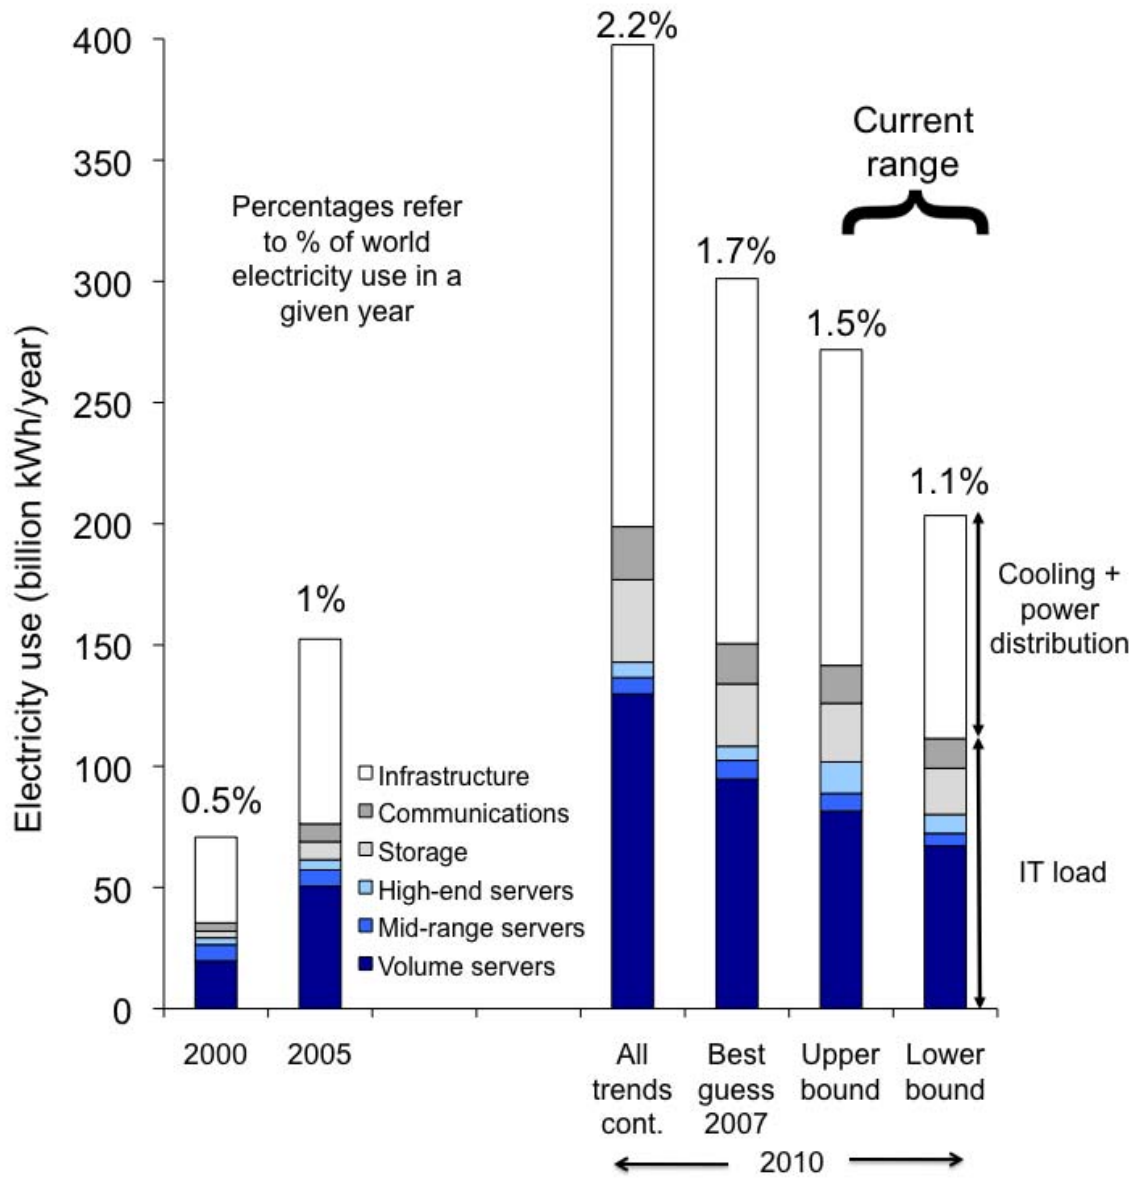


Figure 1: Worldwide electricity use for data centers (2000, 2005, and 2010)




Figure 2: US electricity use for data centers (2000, 2005, and 2010)




Figure 3: Predicted US electricity use for data centers from the EPA report to congress (EPA 2007) and the range estimated in this study

Table 1: IDC server installed base and shipment estimates for the US and the world (thousands)

	Year	US			World			Notes
		Volume	Mid-range	High-end	Volume	Mid-range	High-end	
Installed base old (IDC Feb. 2007)	2000	4,927	663	23.0	12,240	1,808	65.6	1
	2005	9,897	387	22.2	25,959	1,264	59.4	1
	2010	15,434	326	15.2	40,144	1,013	49.6	1
Installed base new (IDC Jan 2010)	2005	9,700	416	23.5	23,765	1,303	63.3	2
	2010	11,325	314	34.8	31,048	982	114.4	2
Shipments								
From IDC January 2010	2010	2,499	35	2.9	6,933	116	10.1	3
From IDC March 2011	2010	2,670	47	4.6	7,504	126	13.0	4
Difference	2010	171	12	1.6	572	10	2.8	
Installed base, adjusted	2010	11,496	326	36.5	31,620	991	117.2	5
<i>Ratio 2005 old to 2000 old</i>		<i>2.01</i>	<i>0.58</i>	<i>0.96</i>	<i>2.12</i>	<i>0.70</i>	<i>0.91</i>	
<i>Ratio 2010 old to 2005 old</i>		<i>1.56</i>	<i>0.84</i>	<i>0.69</i>	<i>1.55</i>	<i>0.80</i>	<i>0.84</i>	
<i>Ratio 2005 new to 2005 old</i>		<i>0.98</i>	<i>1.07</i>	<i>1.06</i>	<i>0.92</i>	<i>1.03</i>	<i>1.07</i>	
<i>Ratio 2010 adjusted to 2010 new</i>		<i>1.02</i>	<i>1.04</i>	<i>1.05</i>	<i>1.02</i>	<i>1.01</i>	<i>1.02</i>	
<i>Ratio 2010 adjusted to 2005 new</i>		<i>1.19</i>	<i>0.78</i>	<i>1.55</i>	<i>1.33</i>	<i>0.76</i>	<i>1.85</i>	

- 1) "Old" installed base numbers for 2000 and 2005 taken from Koomey, Jonathan. 2008. "Worldwide electricity used in data centers." Environmental Research Letters. vol. 3, no. 034008. September 23. <<http://stacks.iop.org/1748-9326/3/034008>>, based on the IDC installed base forecast from February 2007 (which also includes a forecast for installed base in 2010).
- 2) "New" installed base statistics for 2005 and 2010 taken from the IDC January 2010 installed base/shipments forecast, delivered to Jonathan Koomey by Lloyd Cohen of IDC in email on 26 April 2010. The 2010 number represents a forecast for installed base projected to exist at the end of 2010.
- 3) Forecast shipments in 2010, taken from the IDC January 2010 installed base/shipments forecast.
- 4) Actual shipments in 2010, taken from the 15 March 2011 IDC shipments forecast, delivered in email to Jonathan Koomey by Lloyd Cohen of IDC on 22 March 2011.
- 5) The 2010 IDC "new" installed base estimates were adjusted to reflect higher than originally projected 2010 shipments by calculating the difference between the actual shipments in 2010 and the January 2010 estimate, then adding that difference to the estimated 2010 installed base.

Table 2: Total electricity used by data centers worldwide (2000, 2005, and 2010)

<i>Scenario name</i>		<i>All trends cont. Best guess 2007 Upper bound Lower bound</i>							<i>Notes</i>
		<i>Installed server base trend</i>		<i>2000-2005</i>	<i>IDC 2007</i>	<i>IDC adj.</i>	<i>IDC adj.</i>		
<i>W/server trend</i>		<i>Units</i>	<i>2000</i>	<i>2005</i>	<i>2000-2005</i>	<i>2000-2005</i>	<i>2000-2005</i>	<i>No growth</i>	
					<i>2010</i>	<i>2010</i>	<i>2010</i>	<i>2010</i>	
Volume servers	BkWh	19.7	50.5		129.8	94.7	81.5	67.2	1, 2
Mid-range servers	BkWh	6.7	6.7		6.7	7.7	7.3	5.1	1, 2
High-end servers	BkWh	2.8	4.2		6.4	5.9	13.0	7.8	1, 2
Storage	BkWh	2.8	7.5		34.0	25.7	24.2	19.1	1, 3
Communications	BkWh	3.4	7.3		21.9	16.6	15.6	12.3	1, 3
Infrastructure	BkWh	35.4	76.2		198.8	150.6	130.2	92.0	1, 4
Total	BkWh	70.8	152.5		397.6	301.1	271.8	203.4	5
<i>Index relative to 2005</i>		<i>0.46</i>	<i>1.00</i>		<i>2.61</i>	<i>1.98</i>	<i>1.78</i>	<i>1.33</i>	
Average power	Gigawatts	8.1	17.4		45.4	34.4	31.0	23.2	6
World consumption, all uses	BkWh	13238	15747		18118	18118	18118	18118	7
<i>Index relative to 2005</i>		<i>0.84</i>	<i>1.00</i>		<i>1.15</i>	<i>1.15</i>	<i>1.15</i>	<i>1.15</i>	
<i>% of world total</i>	<i>%</i>	<i>0.53%</i>	<i>0.97%</i>		<i>2.19%</i>	<i>1.66%</i>	<i>1.50%</i>	<i>1.12%</i>	8

1) Server, storage, communications and infrastructure electricity use for 2000 and 2005 taken from Koomey, Jonathan. 2008. "Worldwide electricity used in data centers." Environmental Research Letters. vol. 3, no. 034008. September 23. <<http://stacks.iop.org/1748-9326/3/034008>>.

2) I estimate total electricity use for servers by class in 2010 by applying trends in server installed base from 2005 to 2010 (from Table 1) and in electricity use per server for 2000 to 2005 (from Koomey, op cit note 1) to the 2005 electricity use for servers by class from Koomey (op cit note 1). The annual average worldwide growth in electricity use per server from 2000 to 2005 from this source was 3.9%, 7.5%, and 10.7% for volume, mid-range, and high-end servers, respectively. This approach avoids the discrepancy in 2005 installed base documented in Table 1 and allows for the creation of a consistent time series. In the Lower Bound Case, I assume zero growth in electricity use per server, to reflect the industry's increased focus on efficiency in the post 2005 period.

3) Storage and communications/networking electricity use in 2010 is calculated using storage and communications electricity as a fraction of server electricity from the 2010 Current Efficiency Trends forecast for 2010 in EPA (2007). This approach assumes that storage and communications electricity use in 2010 matches that in the EPA forecast as a fraction of total server electricity use. The ratio of storage to server power in 2010 is 24% and for communications it is 15%.

4) Infrastructure electricity use = IT load x (PUE - 1), using PUE of 2.0 for 2000 and 2005 and for the All Trends Continued and Best Guess 2007 cases. For the Upper and Lower Bound cases I used a PUE of 1.92 and 1.83, respectively (see text for details). PUE is the ratio of total data center load to IT load. Infrastructure includes cooling, fans, pumps, and distribution losses.

5) Total electricity use is the sum of servers, storage, communications, and infrastructure.

6) Average power calculated as BkWh/hours per year assuming 100% load factor and 8784 hours in 2000 and 8760 hours in 2005 and 2010.

7) World total electricity use estimated for 2010 using EIA's *International Energy Outlook 2010* generation statistics for 2007 and 2015, interpolated to 2010 and adjusted to reflect 9.8% T&D losses. T&D losses derived using 2007 IEO generation data and 2007 electricity consumption data from <<http://tonto.eia.doe.gov/cfapps/ipdbproject/IEDIndex3.cfm>>.

8) % of world total = total data center BkWh /World total BkWh.

Table 3: Total electricity used by US data centers (2000, 2005, and 2010)

<i>Scenario name</i>				<i>All trends cont.</i>	<i>Best guess 2007</i>	<i>Upper bound</i>	<i>Lower bound</i>	<i>Notes</i>	
		<i>Installed server base trend</i>		<i>2000-2005</i>	<i>IDC 2007</i>	<i>IDC adj.</i>	<i>IDC adj.</i>		
<i>W/server trend</i>		<i>Units</i>	<i>2000</i>	<i>2005</i>	<i>2000-2005</i>	<i>2000-2005</i>	<i>2000-2005</i>	<i>No growth</i>	
					<i>2010</i>	<i>2010</i>	<i>2010</i>	<i>2010</i>	
Volume servers	BkWh	8.0	18.9		44.8	34.8	26.4	22.5	1, 2
Mid-range servers	BkWh	2.5	2.1		1.8	2.6	2.5	1.7	1, 2
High-end servers	BkWh	1.1	1.5		2.0	1.4	3.2	2.3	1, 2
Storage	BkWh	1.1	2.7		11.6	9.2	7.6	6.3	1, 3
Communications	BkWh	1.4	2.7		7.4	5.9	4.9	4.0	1, 3
Infrastructure	BkWh	14.1	28.0		67.6	54.0	41.0	30.3	1, 4
Total	BkWh	28.2	56.0		135.1	107.9	85.6	67.1	5
<i>Index relative to 2005</i>		<i>0.50</i>	<i>1.00</i>		<i>2.41</i>	<i>1.93</i>	<i>1.53</i>	<i>1.20</i>	
Average power	Gigawatts	3.2	6.4		15.4	12.3	9.8	7.7	6
US consumption, all uses	BkWh	3421	3661		3884	3884	3884	3884	7
<i>Index relative to 2005</i>		<i>0.93</i>	<i>1.00</i>		<i>1.06</i>	<i>1.06</i>	<i>1.06</i>	<i>1.06</i>	
<i>% of USA total</i>	<i>%</i>	<i>0.82%</i>	<i>1.53%</i>		<i>3.48%</i>	<i>2.78%</i>	<i>2.20%</i>	<i>1.73%</i>	8

1) Server, storage, communications and infrastructure electricity use for 2000 and 2005 taken from Koomey, Jonathan. 2008. "Worldwide electricity used in data centers." Environmental Research Letters. vol. 3, no. 034008. September 23. <<http://stacks.iop.org/1748-9326/3/034008>>.

2) I estimate total electricity use for servers by class in 2010 by applying trends in server installed base from 2005 to 2010 (from Table 1) and in electricity use per server for 2000 to 2005 (from Koomey, op cit note 1) to the 2005 electricity use for servers by class from Koomey (op cit note 1). The annual average US growth in electricity use per server from 2000 to 2005 from this source was 3.3%, 8.1%, and 6.7% for volume, mid-range, and high-end servers, respectively. This approach avoids the discrepancy in 2005 installed base documented in Table 1 and allows for the creation of a consistent time series. In the Lower Bound Case, I assume zero growth in electricity use per server, to reflect the industry's increased focus on efficiency in the post 2005 period.

3) Storage and communications/networking electricity use in 2010 is calculated using storage and communications electricity as a fraction of server electricity from the 2010 Current Efficiency Trends forecast for 2010 in EPA (2007).

This approach assumes that storage and communications electricity use in 2010 matches that in the EPA forecast as a fraction of total server electricity use. The ratio of storage to server power is 24% and for communications it is 15%.

4) Infrastructure electricity use = IT load x (PUE - 1), using PUE of 2.0 for 2000 and 2005 and for the All Trends Continued and Best Guess 2007 cases. For the Upper and Lower Bound cases I used a PUE of 1.92 and 1.83, respectively (see text for details). PUE is the ratio of total data center load to IT load. Infrastructure includes cooling, fans, pumps, and distribution losses.

5) Total electricity use is the sum of servers, storage, communications, and infrastructure.

6) Average power calculated as BkWh/hours per year assuming 100% load factor and 8784 hours in 2000 and 8760 hours in 2005 and 2010.

7) Year 2000 and 2005 US total electricity use from Koomey (2008), op cit. Note 1. US total electricity use for 2010 from the EIA *Monthly Energy Review* <<http://www.eia.doe.gov/totalenergy/data/monthly/>>. Total use includes internal electricity consumption for industries that self generate, not just retail sales by utilities.

8) % of US total = total data center BkWh /US total BkWh.

Table 4: Educated guesswork about electricity used by Google's servers (1)

	Units	2000	2005	2010	Notes
Google servers	Estimated # of servers	25000	350000	900000	2
	% of total volume servers	0.2%	1.3%	2.8%	3
Volume servers	BkWh	0.028	0.477	1.339	4
Mid-range servers	BkWh	NA	NA	NA	5
High-end servers	BkWh	NA	NA	NA	5
Storage	BkWh	0.003	0.058	0.162	6
Communications	BkWh	0.003	0.057	0.160	6
Infrastructure	BkWh	0.027	0.148	0.266	7
Total	BkWh	0.1	0.7	1.9	8
Average power	Gigawatts	0.01	0.08	0.22	9
World total, all uses	BkWh	13238	15747	18118	10
% of world data centers	%	0.1%	0.5%	0.8%	11
% of world total BkWh	%	0.000%	0.005%	0.011%	12

1) Because Google creates its own custom servers, the company is treated as a server "self assembler" so its servers are not included in the IDC world totals in Table 1.

2) Google does not reveal the number of servers it has. The year 2000 estimate of 25k is a guess corresponding to about 30k sf of electrically active data center floor space (40 1U servers/rack, 50 sf/rack). The 2005 estimate of 350k is a guess based on 450k estimate for 2006 from Markoff, John, and Saul Hansell. 2006. "Hiding in Plain Sight, Google Seeks an Expansion of Power." The New York Times. New York, NY. June 14. The best guess for 2010 is 1,000,000 servers, from sources for *NY Times* reporter Miguel Helft. I reduced this estimate to calibrate the total power use to be under the upper limit electricity use reported by Google's David Jacobowitz (see Note 8 below).

3) Google servers as a % of total volume servers. Year 2000 and 2005 totals from Koomey, Jonathan. 2008. "Worldwide electricity used in data centers." Environmental Research Letters. vol. 3, no. 034008. September 23. <<http://stacks.iop.org/1748-9326/3/034008>>. Year 2010 totals based on adjusted 2010 data from Table 1.

4) Power used by Google servers equals % of total volume servers times total BkWh for all volume servers from Table 2 times (1 - 30% savings for Google servers compared to conventional ones on a per server basis).

5) The number of mid-range and high-end servers at Google is unknown but assumed to be negligible.

6) Storage and communications/networking electricity use is calculated using storage and communications electricity as a fraction of server electricity from Koomey 2008 (op cit note 3) for 2000 and 2005 (the fractions for 2010 are assumed to be the same as 2010). This approach assumes that Google uses storage and communications in amounts comparable to those in other data centers.

7) Infrastructure electricity use = IT load x (PUE - 1), using PUEs of 1.8 in 2000, 1.25 in 2005, and 1.16 in 2010. PUE is the ratio of total data center load to IT load. The 2010 number is from <http://www.google.com/corporate/datacenter/efficiency-measurements.html> (lagging 12 month average from Q4 2010). 2005 is a bit higher than the earliest available data (2008, taken from the same google blog post), and 2000 is an educated guess, assuming that Google's infrastructure was about 10% more efficient than others in 2000.

8) Total electricity use is the sum of servers, storage, communications, and infrastructure. In 2010, total electricity use is calibrated to be less than 1% of 198.8 billion kWh, based on a personal communication about the upper limit of Google data center power use with David Jacobowitz of Google, May 13, 2011.

9) Average power calculated as BkWh/hours per year assuming 100% load factor and 8784 hours in 2000 and 8760 hours in 2005 and 2010.

10) World total electricity use taken from Table 2.

11) % of world data centers= total Google BkWh /World data center BkWh (average of upper and lower bound cases from Table 2).

12) % of world total = total google BkWh /World total BkWh.

Copyright notice

© 2011 by Analytics Press. This report was published in July 2011 by Analytics Press and is licensed under a Creative Commons Attribution-Noncommercial-No Derivative Works 3.0 United States License <<http://creativecommons.org/licenses/by-nc-nd/3.0/us/>>.

Permissions beyond the scope of this license may be available at <<http://www.analyticspress.com>>.

Suggested citation:

Jonathan Koomey. 2011. *Growth in data center electricity use 2005 to 2010*. Oakland, CA: Analytics Press. July. <<http://www.analyticspress.com/datacenters.html>>