Editor: Eoin Woods
Endava
eoin.woods@endava.com

Data Center Energy Demand

What Got Us Here Won't Get Us There

Rabih Bashroush, Eoin Woods, and Adel Noureddine

WE'VE ENTERED A new era of IT activity. In 2008, the number of devices connected to the Internet exceeded the number of people.[1] We're witnessing explosive data growth driven by more affordable storage systems and the proliferation of mobile devices, the Internet of Things, social media, and smart cities. And don't forget the software applications being created around this data, from business analytics to agriculture apps that can monitor a cow's pregnancy via a sensor on its tail. Most apps are cloud based, meaning that demand is constantly increasing on the datacenters large tech companies use to expand their capacity.

Datacenters are believed to emit more greenhouse gases than the entire aviation sector. US datacenters alone consumed an estimated 91 billion kWh in 2013 and are expected to consume 140 billion kWh by 2020.[2] Unsurprisingly, a recent survey of datacenter facility managers showed that power density and energy efficiency were among their top current and future concerns.[3] Such concerns are particularly important for colocation and managed-service datacenter providers in or near large cities with limited grid power.

Will datacenters be able to cope with these energy demands and support the exponential growth of big data and the cloud? Although we can't answer this yet, we know that researchers widely recognize the challenges. They're investigating numerous related areas, from new cooling technologies, to more energy-efficient servers and building designs, to runtime workload consolidation and management techniques.

However, the software architecture community has been slower to recognize its role in energy efficiency and mobilize to meet the challenges. Addressing energy efficiency at the architecture level is still far from mainstream. Architects must ask themselves whether they can continue designing systems without considering energy and power efficiency. An important issue is whether energy efficiency should be a bolted-on system property or a quality attribute addressed during design.

Why Software Architects Aren't Helping Yet

Software architects might not prioritize energy efficiency for three main reasons. First, we have little understanding of how design decisions affect energy efficiency or other system qualities such as user experience, reliability, and performance. Without this knowledge, analyzing tradeoffs to elucidate the benefits or costs of improving energy efficiency is difficult. Minor system design changes could yield substantial benefits, such as avoiding unnecessary polling or eliminating redundant housekeeping tasks that prevent equipment from entering lower power states. However, a lack of relevant design tools and frameworks makes it difficult for architects to achieve more sophisticated optimizations that consider contextual information about the runtime environment.

Second, to achieve the next order of magnitude in energy efficiency, architects must think beyond traditional design boundaries. This will require that people from different specializations and departments work together. Such collaboration is challenging given current organizational software governance structures, wherein teams might have competing objectives, and human dynamics and political barriers. Moreover, existing technologies provide few mechanisms to allow communication across different technology layers (the application software, middleware, hardware, network, cooling, power infrastructure, and so on), which would enable cross-layer optimization.

Finally, end users rarely require or worry about energy efficiency. On one hand are split incentives. System operators such as administrators or datacenter managers don't pick up the energy bill—the money tends to come out of the facilities budget. Accordingly, they would see little return from any energy savings. On the other hand, given current energy prices, information and communications technology energy costs constitute less than 3 percent of a typical organization's budget. So, when an organization pursues energy efficiency, it often does so by addressing areas with a larger budget share (such as payroll!). Exacerbating this problem is the lack of benchmarks, metrics, and reliable data that would allow realistic comparisons of different energy efficiency opportunities and their returns.

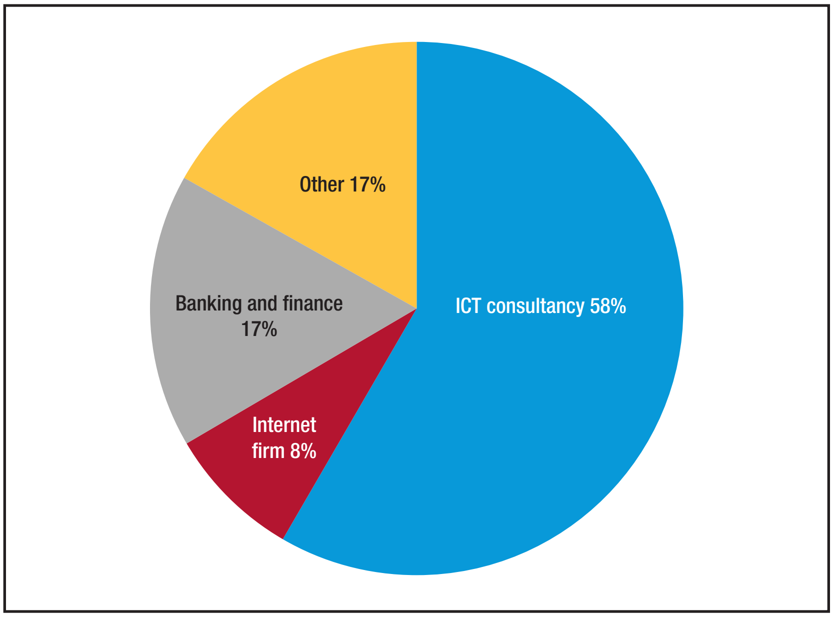

FIGURE 1. Participating architects by organization type. Although the survey was small, the respondents were representative of the IT sector. ICT stands for information and communications technology.

What Software Architects Think

To understand practicing architects' perspectives on energy efficiency, we surveyed 12 representative, experienced architects from various organization types. We asked them whether they had encountered energy-efficiency-related challenges in the last five years and whether they had the right tools to address such challenges. We also asked them whether they believed energy efficiency would be a major architectural concern over the next five years.

Figure 1 breaks down the participants by organization type. Although the survey was small (a confidence level of 80 percent and an error margin of 20 percent), the pie chart confirms that the respondents were representative of the IT sector.

As Figure 2a shows, 83 percent of the participants hadn't dealt with energy efficiency concerns over the last five years. However, 67 percent thought that energy efficiency would be a major concern over the next five years (see Figure 2b). Yet only 25 percent agreed or strongly agreed that they had the right tools to address these challenges (see Figure 3).

The survey results confirmed our view of the challenges that software architects face. They lack the tool support required to address energy at an architectural level, and

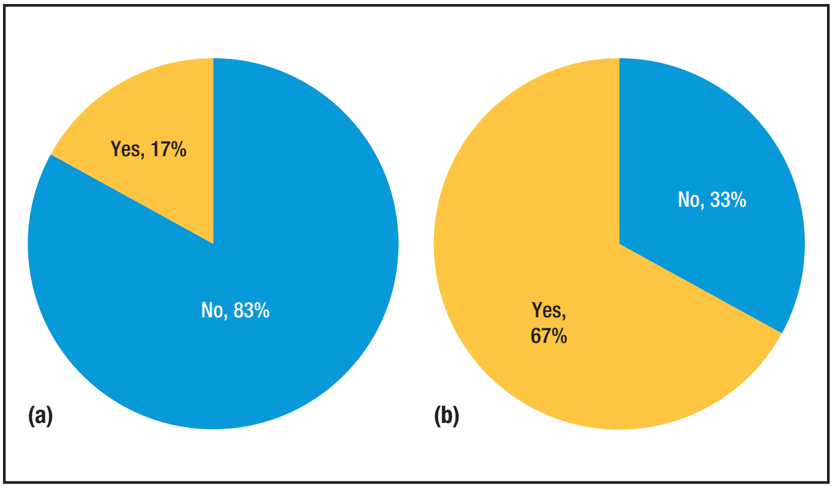

FIGURE 2. Participant responses to whether (a) they had addressed energy efficiency over the last five years and (b) energy would be a major architectural concern in the next five years.

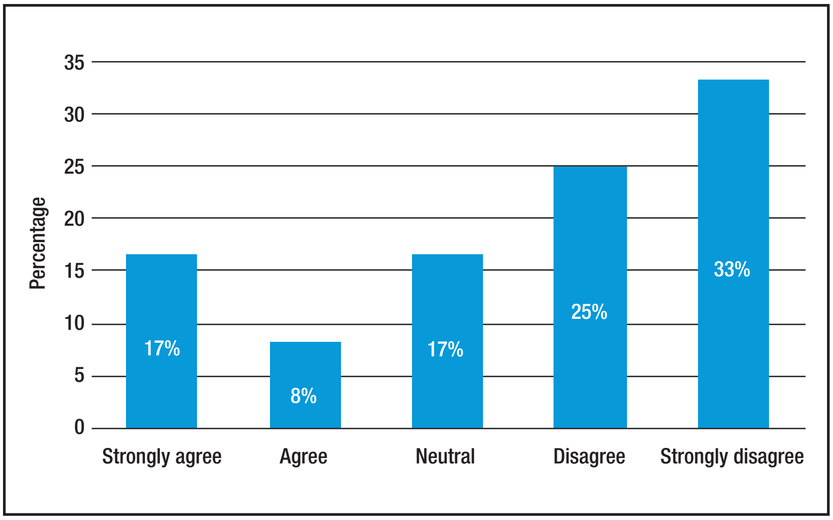

FIGURE 3. Participant responses to whether they had the right tools to address energy efficiency at the design level.

stakeholders place a low priority on energy concerns in solution design.

Opportunities and Future Directions

Although energy efficiency's situation appears gloomy, potential solutions and energy-saving opportunities exist—we just need to recognize and exploit them.

From the system deployment viewpoint, when we consider workload across the datacenter, we could base allocation on the cooling profile rather than on the power or performance profiles often used today. The cooling profile consumes on average 40 percent of total datacenter energy.

When designing applications, we could move beyond system quality requirements based on service-level agreements (SLAs). Most throughput, availability, and performance SLAs are rigid, time-based measures that don't factor in energy usage. Specifying these requirements more flexibly, on the basis of outcomes over time, would let developers incorporate energy efficiency and pricing into their applications. By taking into account the time of day, real-time energy prices, and other environmental factors, some large applications could manage their processing in real time and thus minimize their energy costs.

As application designers, we should also consider how we match workloads to processing environments to minimize energy costs. We tend to focus on ease of application construction and the performance of different application platforms. But we also need to keep in mind the energy consumption available for different platform types. For example, vector-based hardware architectures can offer considerable energy savings for some data-intensive applications. Also, some mainstream hardware—if run at lower clock speeds—can reduce energy consumption while having little effect on perceived application performance.

One inevitable trend is the energy rating of software products—similarly to how industrial equipment and consumer domestic appliances must clearly state their energy efficiency. Today's IT hardware is energy rated, but software isn't. Given the continuing rise of environmentalism, energy prices, and IT workloads, software also seems likely to become energy rated.

However, we lack a metric that quantifies how much work a piece of software performs for a certain amount of energy consumed.

Researchers are investigating various approaches. The basic idea is to

- identify a representative characteristic operation for a particular piece of software (perhaps sending or receiving a million messages for a message bus or processing a thousand single-item orders via a Web shop), and
- measure the energy consumed to process that workload.

We could express application work energy efficiency as work performed divided by energy consumed[4] and thereby compare similar applications' energy efficiency.

Software architects have few tools they can use off the shelf to monitor and minimize software energy consumption. However, only a few years ago, the same could have been said of our colleagues in datacenter and infrastructure architecture. Today, they must comply with standards such as the European Code of Conduct for Data Centres[5] and the Green Grid Data Centre Maturity Model.[6] A similar process for software seems inevitable. So now is the time for us, as software architects, to seriously consider energy efficiency and ensure that our contributions and concerns are heard.

References

1. D. Evans, "The Internet of Things," *Cisco Blogs*, 15 July 2011; http://blogs.cisco.com/diversity/the-internet-of-things-infographic.
2. P. Delforge, *America's Data Centers Consuming and Wasting Growing Amounts of Energy*, Natural Resources Defense Council, 2015; www.nrdc.org/energy/data-center-efficiency-assessment.asp.
3. "Data Centre Manager Overview," Intel Corporation, 2011; https://software.intel.com/sites/datacentermanager/kr/DCM_Overview_041312.pdf.
4. "Description of Energy Metrics for Data Centres D7.1," DC4Cities, 2014; www.dc4cities.eu/en/?smd_process_download=1&download_id=3430.
5. *Data Centres Energy Efficiency*, Inst. for Energy and Transport, European Commission Joint Research Centre, 2015; http://iet.jrc.ec.europa.eu/energyefficiency/ict-codes-conduct/data-centres-energy-efficiency.
6. *WP#36—Data Centre Maturity Model*, white paper, Green Grid, 21 Feb. 2011; www.thegreengrid.org/en/Global/Content/white-papers/DataCenterMaturityModel.

RABIH BASHROUSH is a faculty member and director of the Enterprise Computing Research Group at the University of East London. Contact him at r.bashroush@qub.ac.uk.

EOIN WOODS is the chief technology officer at Endava. Contact him at eoin.woods@endava.com.

ADEL NOUREDDINE is a researcher in software engineering and green computing in the School of Architecture Computing and Engineering at the University of East London. Contact him at a.noureddine@uel.ac.uk.