Uptime Institute INTELLIGENCE

PLANNING & STRATEGY

UI Intelligence report 34

Beyond PUE: Tackling IT's wasted terawatts

After a decade of work on data center energy efficiency, a big shift in focus is now needed.

Authors
Dr. Rabih Bashroush, Research Director, Uptime Institute
Andy Lawrence, Executive Director of Research, Uptime Institute

Analysis of data center energy efficiency trends over the past decade, supported by detailed power usage data from over 300 data centers, shows that significant improvements have been made. However, Uptime Institute's analysis also shows that very substantial energy reduction opportunities still remain untapped. While gains in mechanical and electrical efficiency have stalled over the past few years, it remains that over 65% of the power used by IT in data centers is used to process just 7% of the work, due to aging equipment inefficiencies. But if it is time for operators to pay more attention to IT energy efficiency, they need to do their analysis carefully: the slowdown in Moore's law is creating new complications.

30-45 minutes to read

UII-34 v1.0 published January 6, 2020 last updated January 6, 2020

This Uptime Institute Intelligence report includes:

ABOUT UPTIME INSTITUTE INTELLIGENCE

Uptime Institute Intelligence is an independent unit of Uptime Institute dedicated to identifying, analyzing and explaining the trends, technologies, operational practices and changing business models of the mission-critical infrastructure industry. For more about Uptime Institute Intelligence, visit uptimeinstitute.com/ui-intelligence.

KEY FINDINGS

- **Efforts to improve the energy efficiency of the mechanical and electrical infrastructure of the data center are now producing only marginal improvements. The focus needs to move to IT.**
- **The most-implemented practices for energy reduction are those that do not require substantial investment but do require process, discipline or relatively minor and incremental investments.**
- **Initiatives that span IT, that involve most cultural and multi-disciplinary changes and that require major strategic operational changes are the least-implemented energy efficiency practices.**
- **Major energy efficiency opportunities involving IT remain untapped — partly due to a misplaced management focus on infrastructure.**
- **Energy-saving opportunities on the IT side are so great that if fully addressed, they would significantly reduce data center energy use and carbon footprint, would slash energy bills and would likely lead to reduced demand for cooling and critical power equipment.**
- **In a study of 300 data centers, aging IT kit (older than five years) accounted for 66% of IT energy use but contributed just 7% of the compute capacity.**
- **All these issues are well-known and can only be resolved by senior management, which is empowered to make decisions that cross the IT/facilities boundary or drive behavior among suppliers and clients. An understanding of the sheer scale of the energy savings should encourage executives to address the issues more directly.**
- **Over the past few years, while processor lithography has stagnated at 14 nanometers, the increase in performance per watt has been accompanied by a steady increase in idle power consumption (perhaps due to the increase in core count to achieve performance gains). This is one reason why the case for hardware refresh for more recent kit has become weaker: Servers in real-life deployments tend to spend a substantial part of their time in idle. As such, the increase in idle power may overall offset energy gains from performance.**
- **If a server spends a disproportionate amount of time in active idle mode, the focus should be on active idle efficiency (i.e., choosing servers with lower core count) rather than just on higher server performance efficiency, while satisfying overall compute capacity requirements.**

Introduction/Context

It is widely known that the aggregated energy consumption of the infrastructure of global IT (data centers, servers, networks, devices) has been rising steadily for many years — even if the scale of overall energy consumption is a matter of debate and requires further research (see **Appendix, Note 1**). It seems likely that the annual consumption of energy by data centers is somewhere between 400 terawatt-hours (TWh) and 500 TWh, depending on what is counted as a data center. To put things in perspective in terms of demand, research by Uptime Institute Intelligence shows that every time an image is posted on Instagram by the Portuguese soccer star Cristiano Ronaldo (who at the time of writing has the highest number of followers on the platform), his more than 195 million followers consume nearly 30 megawatt-hours (MWh) of energy to view it.

Some forecasters (most notably the International Energy Administration, or IEA) have predicted a flattening or even downturn in overall energy use by data centers — the result of improving efficiency. Certainly, the explosion in IT demand in the past decade (2010-2020) did not translate directly into the same rate of growth for infrastructure energy consumption. However, Uptime Institute expects that demand for IT services and data centers will substantially outpace the gains from efficiency practices over the next five years, resulting in steadily increasing energy use.

Increasing energy demand by IT — especially if that energy is used inefficiently — has both direct and indirect consequences for data center operators and major customers. First, power is a major cost component of all IT services; inefficiency, therefore, increases costs and, very likely, reduces overall margins. Second, a shortage of power in some geographies increases power prices; and third, most power sources produce carbon dioxide (CO_2), and all polluters will come under increasing political and economic pressure to reduce emissions in the coming years.

PUE and efficiency

Over the past decade, the data center industry has stepped up its efforts to address the issue of energy inefficiency. Several indicators to help measure and track infrastructure efficiency have been developed, of which power usage effectiveness (PUE) is the most widely used. PUE, adopted and promoted by The Green Grid in 2006, is the ratio of the total facility energy consumption to IT equipment energy consumption. The ISO/IEC 30134-2:2016 (ISO, International Organization for Standardization; IEC, International Electrotechnical Commission) Standard now defines what PUE is and how it is measured.

Since its introduction, PUE has helped focus attention on mechanical and electrical (M&E) infrastructure efficiency — a major cause of energy waste and inefficiency in data centers. The metric has helped drive uptake of many power and cooling best practices.

Uptime Institute data, based on global surveys, shows a substantial drop in PUE from 2007 to 2014, after which the law of diminishing returns seems to have begun to limit the impact of energy savings at the infrastructure level (see Figure 1).

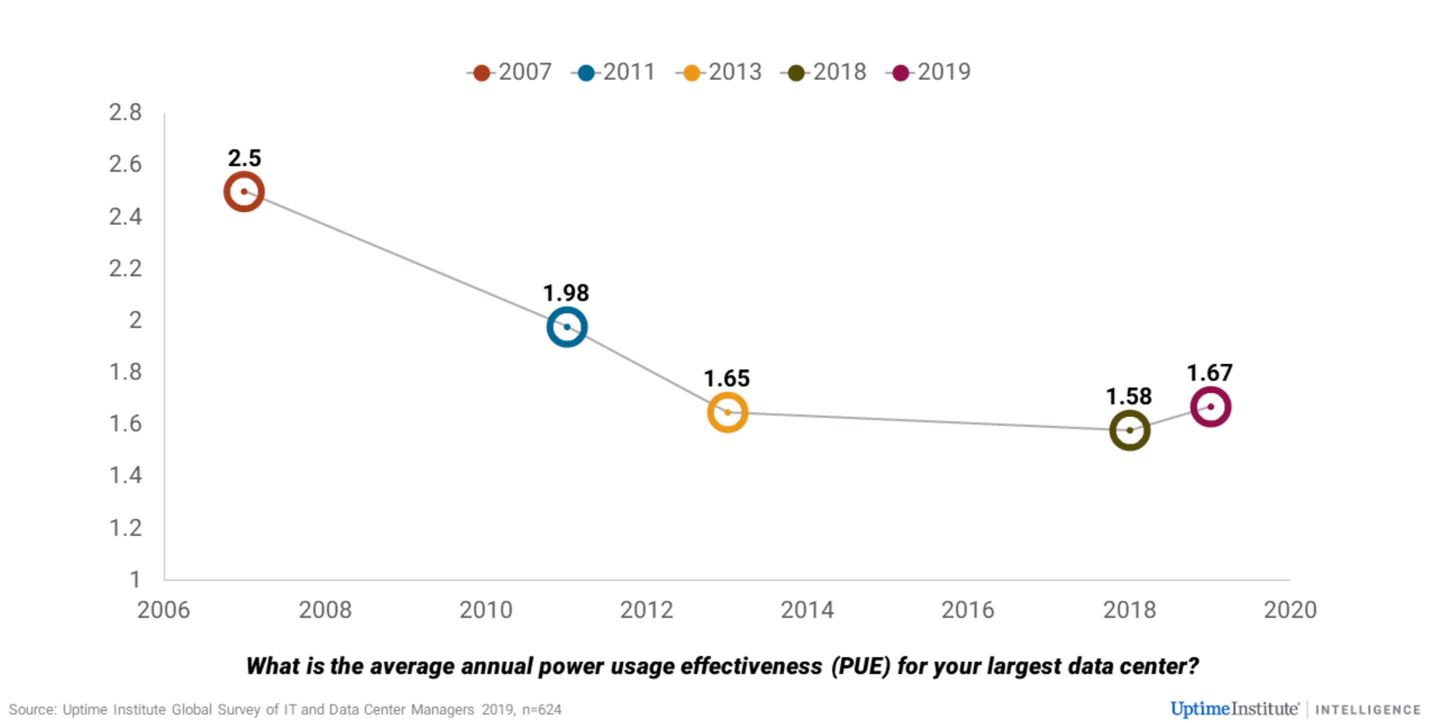

Figure 1. **Data center efficiency gains have stalled**

These improvements are the result of several major steps taken by owners and operators of data centers over the period. Interventions range from implementing simple best practices (e.g., separating hot and cold air) and using energy-efficient technologies (e.g., indirect and direct free cooling) to building entire data centers engineered to minimize power inefficiencies and energy waste.

It seems likely that the long trend in PUE will be marginal decreases, partially the result of newer data centers replacing older, less-efficient ones (the small uptick in 2019 is likely the result of temporary factors). In other words, most of the large, short gains from a focus on M&E efficiency have been made. Further gains need larger investments, a bigger focus on overall demand and, most importantly, greater involvement on the IT equipment and IT management side.

Tough measures delayed?

Data from the adoption of measures in the European Code of Conduct (CoC) for Data Centre Energy Efficiency supports the point. The program was introduced in 2008 as a voluntary initiative to help increase energy efficiency in data centers. (Although it is a European initiative, adoption of the CoC is not limited to Europe, especially after being published by ISO/IEC as a technical report.)

The CoC provides a list of infrastructure efficiency best practices that is updated annually. The lists below show the most and least-implemented best practices according to analysis published by the European Union (EU) Commission Joint Research Centre, based on data submitted by over 350 participant data centers.

The first list shows that the most-implemented best practices seem to be the ones that do not require substantial investment but do require process, discipline or relatively minor and incremental investments.

MOST-IMPLEMENTED EUROPEAN COMMISSION CODE OF CONDUCT BEST PRACTICES

- Encouraging group involvement (i.e., multi-team focus)
- Building resiliency to business requirements
- Following lean provisioning principles
- Engineering to increase energy efficiency under partial load conditions
- Designing hot/cold aisle containment
- Installing blanking plates
- Separating hot aisles from cold aisles
- Using perforated doors on server cabinets
- Installing high-efficiency uninterruptible power supplies
- Turning off lights
- Installing low-energy lighting
- Using IT energy consumption meters
- Performing periodic manual readings of entry-level energy, temperature and humidity

Source: Joint Research Centre (Ispra), European Commission

The second list shows the least-implemented practices. These practices include those that span IT, involve most cultural and multi-disciplinary changes and require major strategic operational changes (with the possible exception of uninterruptible power supply [UPS] operating modes, which may be an investment issue or one that involves risk).

LEAST-IMPLEMENTED EUROPEAN COMMISSION CODE OF CONDUCT BEST PRACTICES

- Using server power management tools
- Selecting energy-efficient software
- Developing energy-efficient software
- Reviewing cooling requirements
- Revising cooling strategy
- Considering more energy-efficient uninterruptible power supply operating modes

Source: Joint Research Centre (Ispra), European Commission

IT energy efficiency

If energy demand growth in the past decade has lagged the growth in IT demand, it is only partially the result of data center efficiency initiatives. The biggest improvements are the result of increased hardware efficiency (from processors, memory/storage and more), virtualization and the rise in consolidation activities. Despite this progress, major energy efficiency opportunities remain untapped. This may, in part, be due to a misplaced management focus on infrastructure and on key performance indicators (KPIs) such as PUE.

Server efficiency — the data speaks dollars

It is hardly news to the industry that more work needs to be done "on the 1.0" — the denominator side of the PUE ratio. This represents the IT load, the work done, and the energy used by the equipment. The challenge is to initiate effective action.

In one sense, it is already a very positive story. Servers have seen consistent increases in energy efficiency over the past several decades — largely but not entirely the result of Moore's law (see below). This enables infrastructure operators to steadily increase compute capacity, while reducing energy consumption, through hardware refresh.

However, gauging the remaining untapped energy-saving opportunity has been problematic for most operators, given the lack of appropriate KPIs (IT efficiency is not captured by PUE). Good data on the scale and scope of potential improvements, especially relative to reductions in PUE, is rare.

Uptime Intelligence has been able to scope these improvements by analyzing the data from over 300 data center facilities, assessed as a part of EURECA (EU Resource Efficiency Coordination Action), an EU-funded project. EURECA provided a unique snapshot of the data center industry in general, rather than just large-scale facilities. The project ended in 2018, but the data is still representative of the industry, given the known pace of change (investment lifecycles) in the sector.

The findings suggest that the energy-saving opportunities on the IT side are so great that if fully implemented, they would significantly reduce the overall global data center energy use and carbon footprint, would slash energy bills, and would likely lead to a significant reduction in demand for cooling and critical power equipment. For some data centers, it would help managers cut operational power costs and would provide long-term headroom for expansion. It would also help enterprise data centers narrow the wide lead in operational efficiency that hyperscale cloud providers are ravenously exploiting.

A key finding of the data (in the EURECA sample) is that aging IT kit (older than five years) represented 40% of deployed servers, consumed 66% of the energy, but contributed just 7% of the compute capacity (see Figure 2).

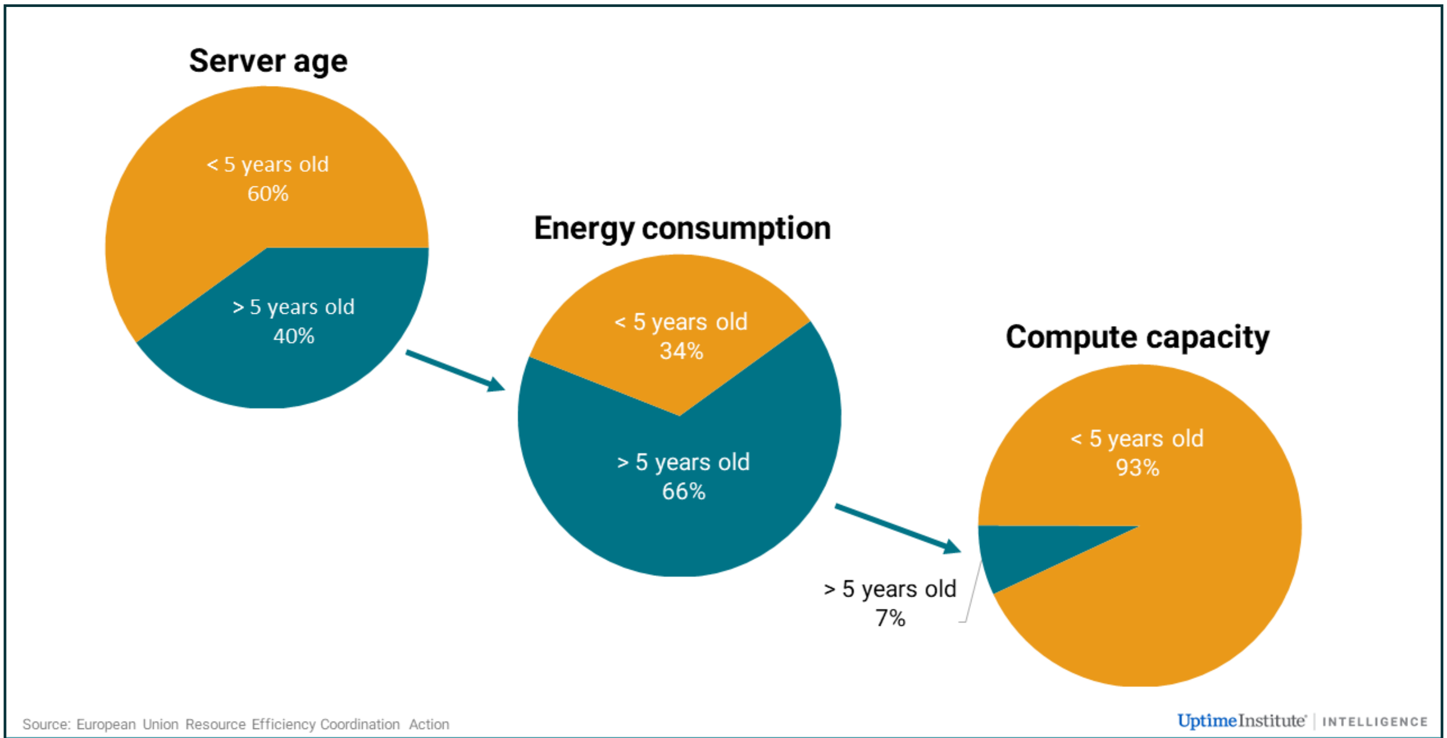

Figure 2. **Older servers use more energy than newer servers yet have far less compute capacity**

The message from this is clear: updating older servers and eliminating many through consolidation will slash the big 66% number and reduce the overall energy consumption significantly. While there may be some servers that cannot be easily replaced or some workloads that require proprietary and possibly older machines, there is a still a significant opportunity for updating.

Uptime Intelligence also examined the energy consumption of a fixed workload (equivalent to the need of 200 million server-side Java operations — ssj_ops — see **Appendix, Note 2**) in different operating environments, using servers of varying age from current-generation equipment to 9-year-old kit. The number of servers required to support the workload varies according to compute capacity — from 48 to 4,473.

Various scenarios were analyzed — from using old, under-utilized servers that were not virtualized, to using modern, highly utilized servers with full virtualization. The results are shown in Table 1.

Table 1. **Workload energy consumption under different conditions**

	PUE*	Utilization	Energy consumption of a 200-million ssj_ops** workload (in megawatt-hours)					
			9 years old	7.5 years old	6 years old	4.5 years old	3 years old	up to 1.5 years old
On-premises, not virtualized								
Worst case	3	5%	10,349	5,846	2,747	1,746	1,534	1,459
Average	2	10%	3,778	2,202	1,049	697	629	580
Best practice	1.5	25%	1,428	889	435	313	294	258
Colocation, not virtualized								
Worst case	2.5	5%	8,624	4,872	2,289	1,455	1,279	1,216
Average	1.9	10%	3,400	1,982	944	627	566	522
Best practice	1.3	25%	1,238	770	377	271	255	224
On-premises, with virtualization								
Worst case	3	6%	8,788	4,998	2,356	1,512	1,337	1,263
Average	2	30%	1,696	1072	528	386	366	318
Best practice	1.5	60%	882	592	298	231	226	189
Private cloud								
Worst case	2.5	7%	6,394	3,661	1,730	1,121	997	935
Average	1.8	30%	1,527	965	475	347	330	286
Best practice	1.3	60%	764	513	258	200	195	164
Public cloud								
Worst case	2	7%	5,115	2,929	1,384	897	798	748
Average	1.5	40%	1,077	698	347	260	250	214
Best practice	1.1	70%	606	412	208	163	160	134

**PUE – Power usage effectiveness*
***ssj_ops - server-side Java operations*

Note. *The table uses server performance data from the Standard Performance Evaluation Corporation (SPEC). These SPECpower results (see* ***Appendix, Note 2****) were submitted over a nine-year period ending August 2019. The servers are dual-socket, industry-standard volume servers running a normalized server-side Java workload.*

The five scenarios are not presented as alternatives but rather as examples using likely environments. For example, a colocation company may have a better worst-case power usage effectiveness than many enterprise data centers but a lower average utilization than many dedicated private operators. However, there will be many exceptions, and often the data will vary widely at the rack level.

Looking at the first row in Table 1, note that running the selected workload on 9-year-old servers consumes over 10 gigawatt-hours (GWh) of electricity at a cost of over $1 million (assuming an energy price of $0.10 per kilowatt-hour, or kWh). Running the same workload within the same environment at the same utilization level but using the latest servers drops the energy consumption to 1.5 GWh, at a mere cost of $150,000. This represents a staggering reduction of 85% in energy consumption, all for the same PUE (see **Appendix, Note 3**) — and an operational energy savings of $850,000.

These savings are vast. But what about the investment needed? The server refresh in this scenario requires 667 new servers, maintaining the same utilization level (to approximate return on investment, or ROI, from hardware refresh only, without increasing utilization level). At a cost of $3,000 per server (assuming a volume dual-socket server) and 15% procurement cost/overhead, the refresh in this case would cost $2.3 million and would pay for itself within three years (due to the reduced energy consumption).

Financial reasons are good enough for most managers. But the environmental benefit is significant, too: The refresh will also save over 2,623 metric tons of CO_2 emissions annually in Europe (assuming a grid emission factor of 295 grams per kWh), or over 5,067 metric tons of CO_2 emissions in the United States (assuming a grid emission factor of 570 grams per kWh). (For context, 5,000 metric tons is equivalent to 12 million miles driven by an average passenger vehicle.)

The analysis above does not factor in any return from recycling existing equipment; initially this would be low, but it would rise as a regular refresh program is instituted, because the equipment being replaced would likely be newer. It also does not include returns attributed to the reduction in server count (from over 4,000 old servers to 667 new ones), leading to substantially reduced maintenance costs, risk factors, floor space, and so on. Nor does it consider associated potential savings in UPS, batteries and generator capacity, or in cooling. In addition, for some at least, the space freed by a refresh would enable a capacity increase, the potential for more consolidation from other data centers, or for enterprises, a reduced use of colocation services.

For all these reasons, the business case for a full-on refresh and utilization/efficiency program is likely to include a ROI of better than three years — especially in geographies where the power prices are high or are likely to become higher.

Utilization as an energy saver

At the workload level, the rise of the cloud and virtualization over the past decade has helped increase server utilization levels from around 5%[1] at the beginning of the decade to 25% more recently, as per EURECA data and Shehabi (2016).[2]

[1] A. Beloglazov, R. Buyya, Y. C. Lee, and A. Zomaya. 2011. "A taxonomy and survey of energy-efficient data centers and cloud computing systems" in **Advances in computers,** vol. 82, pp. 47-111. Elsevier. http://www.cloudbus.org/papers/GreenCloudTaxonomy2011.pdf

[2] Arman Shehabi, Sarah Smith, Dale Sartor, Richard Brown, Magnus Herrlin, Jonathan Koomey, Eric Masanet, Nathaniel Horner, Inês Azevedo, and William Lintner. 2016. **United States data center energy usage report.** https://eta.lbl.gov/publications/united-states-data-center-energy

To better understand the opportunity, consider the workload discussed in Table 1 but analyzed in a few example environments at different utilization levels against a ***fixed*** PUE, as shown in Table 2.

Table 2. **Workload energy consumption at different utilization levels and fixed PUE**

Utilization	Energy consumption of a 200-million ssj_ops** workload (in megawatt-hours)					
	9 years old	7.5 years old	6 years old	4.5 years old	3 years old	up to 1.5 years old
PUE* = 2 (On-premises)						
5%	6,899	3,897	1,831	1,164	1,023	973
10%	3,778	2,202	1,049	697	629	580
25%	1,904	1,185	580	417	393	344
PUE* = 1.8 (Colocation and private cloud)						
5%	6,209	3,508	1,648	1,048	921	875
10%	3,400	1,982	944	627	566	522
25%	1,714	1,066	522	375	353	310
PUE* = 1.5 (Public cloud)						
5%	4,394	2,499	1,178	756	669	631
30%	1,272	804	396	289	275	238
60%	882	592	298	231	226	189

**PUE – Power usage effectiveness*
***ssj_ops - server-side Java operations*

Note. *The table uses server performance data from the Standard Performance Evaluation Corporation (SPEC). These SPECpower results (see **Appendix, Note 2**) were submitted over a nine-year period ending August 2019. The servers are dual-socket, industry-standard volume servers running a normalized server-side Java workload.*

The scenarios are not presented as alternatives but rather as examples.

The data in this table (fixed PUE, variable utilization) shows the major opportunity for energy reduction that can be achieved by increasing utilization levels (reducing the number of servers needed to run the workload). For example, running the workload on 3-year-old kit, with a PUE of 2 but increasing the utilization from 5% to 25%, reduces the energy consumption by 62%. At the same time, reducing the PUE from 2 to 1.5, which might require some substantial investment, yields only 25% savings.

It should be noted that these benefits may not always be available — a careful analysis of workloads and capacity is needed. As with the server refresh (see below), there are complications in increasing utilization levels. These include:

- The need to keep utilization below 50% to allow for failovers in active/active environments.
- The fact that some processor manufacturers do not guarantee performance beyond certain utilization levels (due to performance degradation).
- The need to reserve compute capacity for peak demands in some deployments (when it is not possible to dynamically scale-out across servers).
- The possibility that the server may not be configured appropriately for the workload (e.g., memory lookup-heavy workloads). This means the memory or input/output (I/O) can become the bottleneck, leaving the central processing unit (CPU) underutilized.

In spite of this, the opportunity for cutting IT energy waste remains considerable. Evidence suggests most of the above complications — which apply only in some situations — occur only when utilization levels go above 35% or 40%. Given that industry-wide utilization levels are still (we believe) around 25%, there is significant opportunity left.

Focus on IT energy use?

The clarity of the business case for refreshing servers and raising utilization (for efficiency reasons) raises the question: Why don't more organizations have active policies and processes (and investment programs) to decrease their IT energy use? Our evidence suggests that large commercial (hyperscale) operators do this rigorously, as do a few leading-edge enterprises — but the majority do not.

There are many reasons for this: a lack of awareness of a clear business case/good data; budgetary issues (lack of capital versus operational monies); split incentives and interdepartmental cultural issues (the IT management are not interested in/concerned with energy consumption); a misconception that reducing IT energy use will usually result in less IT power; and misunderstandings about environmental impact.[3]

Of all these, the issue of interdepartmental responsibilities and split incentives is probably the biggest. Facilities managers are often very concerned with reducing power consumption, but their colleagues in IT may have little interest or no real incentives. The use of colocation services can further entrench this difference, since colocation companies/managers usually have no influence on IT energy use by their clients (and may even be incentivized to allow energy waste to continue).

[3] Rabih Bashroush. 2018. "A comprehensive reasoning framework for hardware refresh in data centers." **IEEE Transactions on Sustainable Computing** vol. 3(4): 209-220. https://ieeexplore.ieee.org/document/8263130

All these issues are well-known and can only be resolved by senior management, which is empowered to make decisions that cross the IT/facilities boundary or drive behavior among suppliers and clients. An understanding of the sheer scale of the energy savings should encourage executives to address the issues more directly.

Moore's law and beyond

There is a complication to the refresh/upgrade policy that is relatively recent: the slowing down of Moore's law. As stated earlier, the most dramatic examples of savings in energy are achieved when replacing old servers — up to nine years old. However, if we consider refreshing more recent servers (e.g., 3-year-old servers), the picture may far less clear than it once was. This is due to the stagnation witnessed in Moore's law over the past few years (see below).

Moore's law refers to the observation made by Gordon Moore (co-founder of Intel) that the transistor count on microchips would double every two years. This implied that transistors would become smaller and faster, while drawing less energy. Over time, the doubling in performance per watt was observed to happen around every 18 months. (This was first observed by Moore's colleague David House but is known by some in the data center industry as "Koomey's law" after analyst Jon Koomey, who published a 2010 paper[4] showing the trend.)

It is this doubling in performance per watt that underpins the major opportunity for increasing compute capacity while increasing efficiency through hardware refresh. But in the past five years, it has been harder for Intel (and immediate rivals AMD) to maintain the pace of improvement. This raises the question: Are we still seeing these gains from recent and forthcoming generation of CPUs? If not, the hardware refresh case will be undermined … and suppliers are unlikely to be making that point too loudly.

To answer this question, Uptime Intelligence analyzed the SPECpower dataset containing energy performance results from hundreds of servers, based on the SPECpower benchmark. To be able to track trends and eliminate potential outlier bias in reported servers (e.g., high-end servers versus volume servers), only dual-socket servers were considered for trend consistency. These were then broken down into 18-month intervals (based on the published date of server release in SPECpower) and the performance averaged for each period. The results are shown in Table 3.

[4] Jonathan Koomey, Stephen Berard, Maria Sanchez and Henry Wong. 2010. "Implications of historical trends in the electrical efficiency of computing." **IEEE Annals of the History of Computing** vol. 33(3): 46-54. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.323.9505&rep=rep1&type=pdf

Table 3. **Analysis of dual-socket volume servers (2007-2019)**

Time intervals (1.5 years)	Average watts @ 100% of target load	Average watts @ active idle	Performance/ power @ 100% of target load	Dynamic range	Server count (n)
2007/09 - 2009/02	255	156	1,352	1.64	61
2009/03 - 2010/08	227	76	2,928	3.00	57
2010/09 - 2012/02	247	80	3,648	3.10	30
2012/03 - 2013/08	253	63	5,277	3.99	73
2013/09 - 2015/02	245	55	9,791	4.41	18
2015/03 - 2016/08	270	46	12,710	5.93	13
2016/09 - 2018/02	398	57	12,754	6.96	21
2018/03 - 2019/08	387	67	15,335	5.76	34

Note. *The table uses server performance data from the Standard Performance Evaluation Corporation (SPEC). The processing power and energy consumption data in the table is based on servers submitted — not on the processors.*

Time intervals *are 18 months (column 1), roughly in line with Moore's law. Date format = year/month.*

Target load *is the maximum throughput of a server (determined by the SPEC application and considered to represent 100% load).*

Active idle *describes a steady state of processing readiness — the server is on but is not processing any workload; it is not in sleep state.*

Performance/power *shows the average performance (operations per watt) of servers, which increases (but not evenly) for each new period.*

Dynamic range *is the ratio of power consumption at maximum work to power consumption at active idle (indicates the power proportionality of the server).*

Server count *is the number of servers submitted and analyzed for each period. (In total, n=307.)*

Figure 3 — which shows server performance per watt (based on Table 2, column 4), along with the trend line (polynomial, order 3) — demonstrates how performance increases have started to plateau, particularly over the past two periods. The data suggests upgrading a 2015 server in 2019 might provide only a 20% boost in processing power for the same number of watts. In contrast, upgrading a 2008/2009 server in 2012 might have given a boost of 200% to 300%.

To further understand the reason behind this, we charted the way CPU technology (lithography) has evolved over time, along with performance and idle power consumption (see Figure 4).

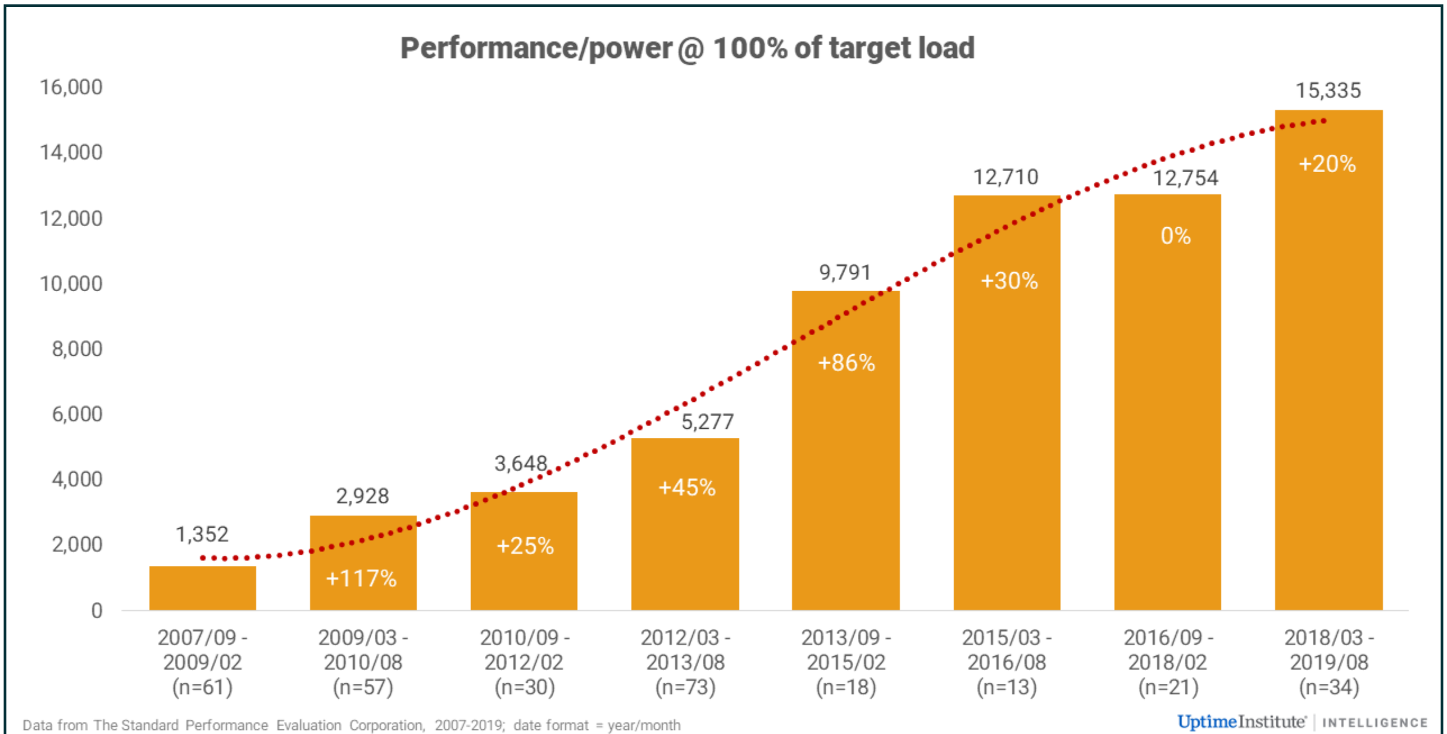

Figure 3. **Server performance over time**

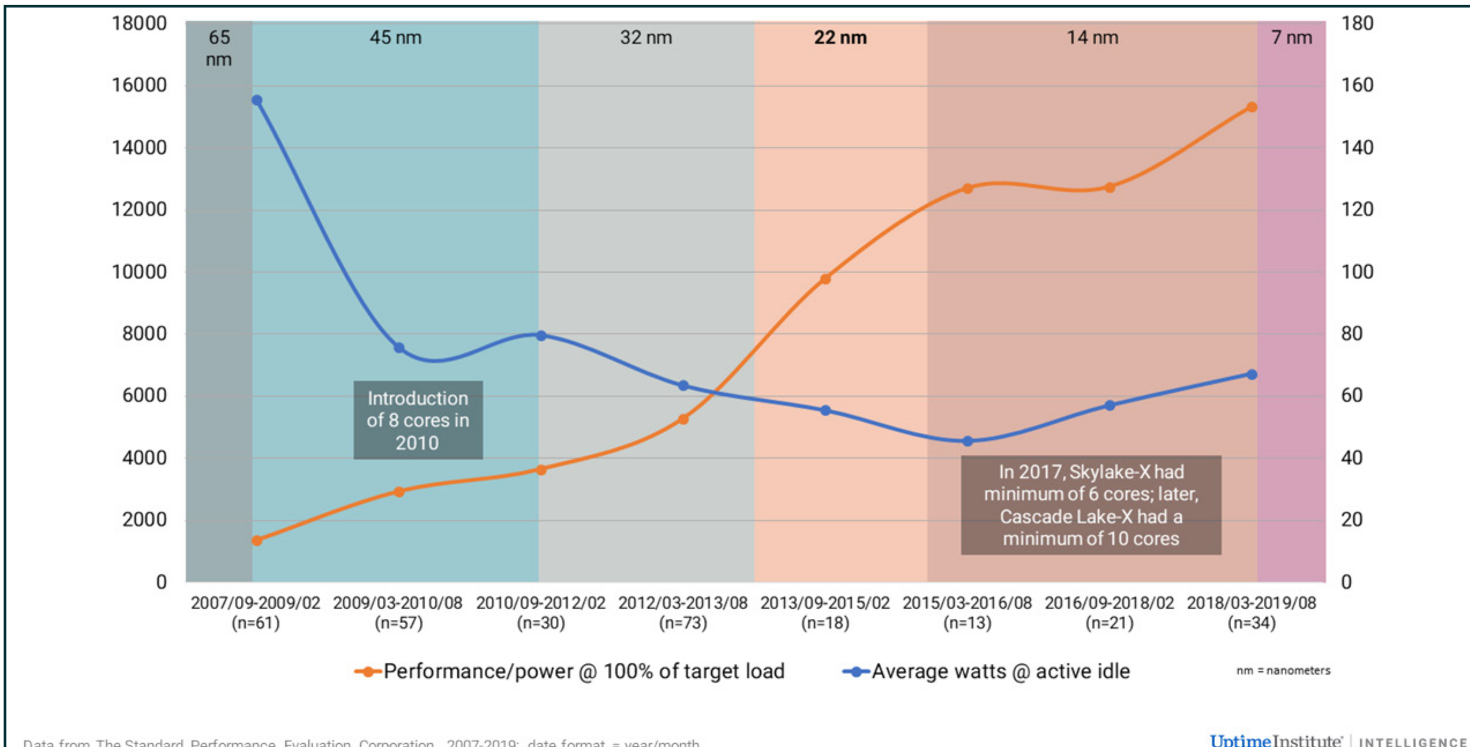

Figure 4. **Performance vs. idle power vs. lithography**

Note. *The color-coded vertical bars represent generations — lithography — of processor technology (usually, Intel). For each generation of around three to four years, hundreds of servers are released. The steeper the rise of the orange line (compute performance per watt), the better. For the blue line — power consumption at idle — the steeper the decine, the better.*

Figure 4 reveals some interesting insights. During the beginning of the decade, the move from one CPU lithography to another — for example, 65 nanometer (nm) to 45 nm, 45 nm to 32 nm, etc. — presented major performance-per-watt gains (orange line), as well as a substantial reduction in idle power consumption (blue line), thanks to the reduction in transistor size and voltage.

However, it is also interesting to see that the introduction of a larger number of cores to maintain performance gains produced a negative impact on idle power consumption. This can be seen briefly during the 45 nm lithography and very clearly in recent years with 14 nm.

Over the past few years, while lithography stagnated at 14 nm, the increase in performance per watt (when working with a full load) has been accompanied by a steady increase in idle power consumption (perhaps due to the increase in core count to achieve performance gains). This is one reason why the case for hardware refresh for more recent kit has become weaker: Servers in real-life deployments tend to spend a substantial part of their time in idle (discussed in **Power proportionality**). As such, the increase in idle power may offset energy gains from performance.

This is an important point that will likely have escaped many buyers and operators: If a server spends a disproportionate amount of time in active idle mode — as is the case for most — the focus should be on active idle efficiency (e.g., choosing servers with lower core count) rather than just on higher server performance efficiency, while satisfying overall compute capacity requirements.

It is, of course, a constantly moving picture. The more recent introduction of the 7 nm lithography by AMD (Intel's main competitor) should give Moore's law a new lease on life for the next couple of years. However, it has become clear that we are starting to reach the limits of the existing approach to CPU design. Innovation and efficiency improvements will need to be based on new architectures, entirely new technologies and more energy-aware software design practices.

Power proportionality

Another way to assess server efficiency is to look at how the dynamic range has evolved over time. The dynamic range is the ratio of the energy consumption at 100% target load over idle. The higher the dynamic range is, the more energy proportionate a server is. Ideally, a server uses 1% of power when doing next to no work, and 100% when fully active, giving a ratio of 100. More realistically, servers use 17% to 25% of their power at active idle.

Figure 5 shows how the dynamic range evolved over the past decade and confirms the issues highlighted above. While dynamic range has increased over the years, it has now started to dip, with performance gained through increased core count rather than a reduction in lithography.

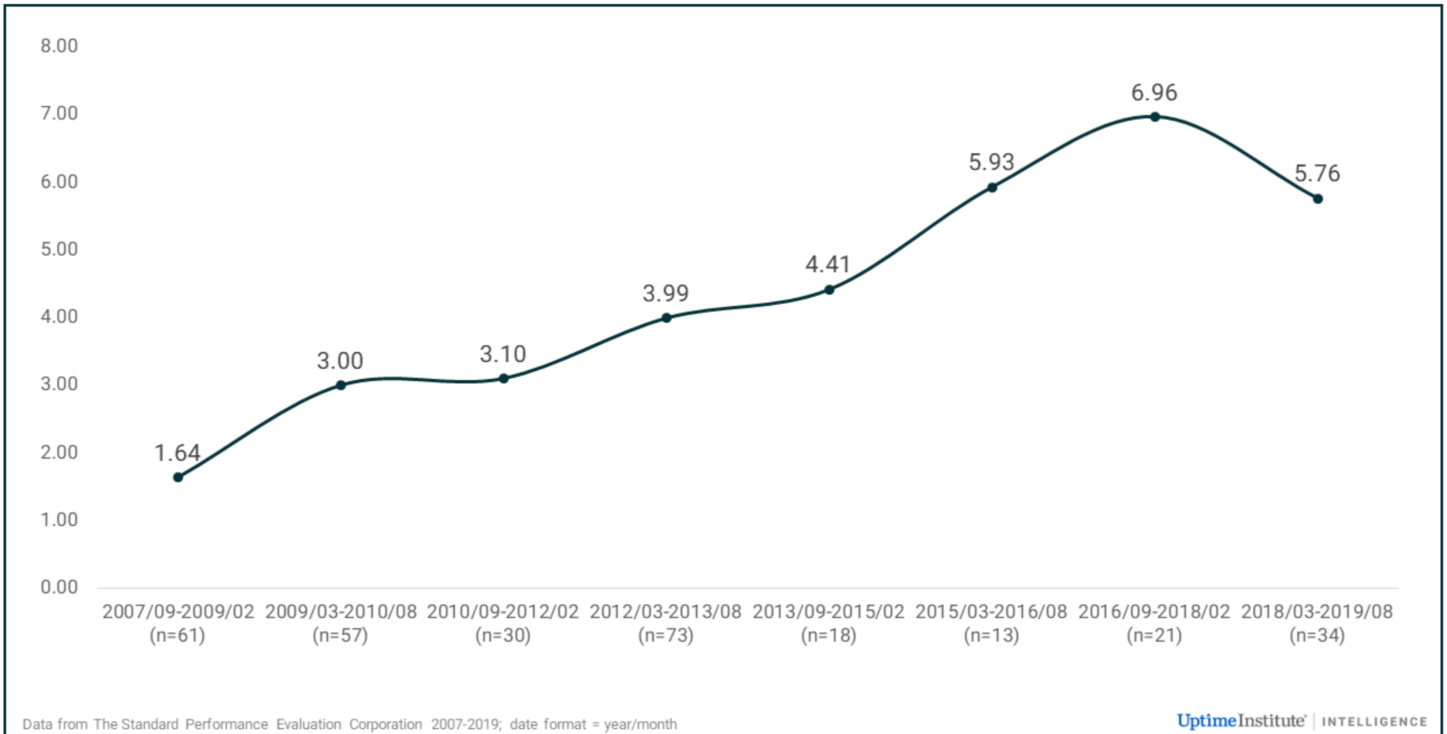

Figure 5. **Dynamic range over time**

The implications of this, for operators and managers, are twofold: First, for situations and workloads where the work is dynamic and often drops to near zero, it may pay to use models with higher dynamic range (i.e., fewer cores) — possibly older ones — or to manage workloads to avoid periods of low activity. And second, it may also pay to use power management states when activity drops. The use of caching, virtualization and more advanced workload management software can reduce or eliminate any performance impact resulting from the recent drop in dynamic range.

IT consolidation

The third main area that has contributed to increasing energy efficiency in IT — and one that still presents a major opportunity for operators — is consolidation, at the facility as well as the workload level. Larger facilities are more efficient in terms of shared infrastructure and operational costs. Hyperscale operators, large colocation companies and some enterprises are benefiting from consolidation and building/managing at scale.

But the data from the EURECA project shows that 80% of the data centers studied contained fewer than 25 racks; 17% contained between 25 and 125 racks, and only 3% had more than 125 racks (see Figure 6). Small (and often inefficient) data centers remain a large, and largely unseen, proportion of global critical infrastructure.

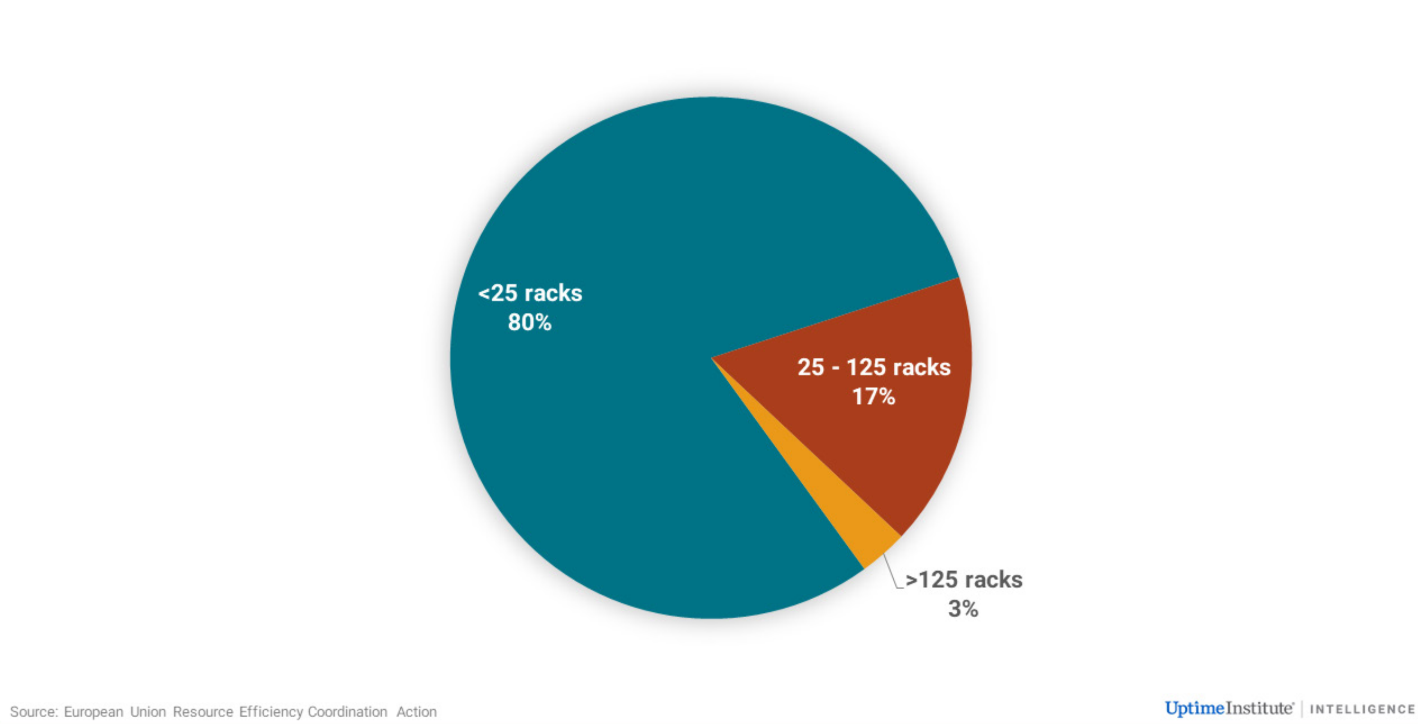

Figure 6. **Rack count per data center (European dataset)**

While a growing proportion of demand for new digital services over the past decade was met by hyperscalers (a fact reflected in the growth enjoyed by these big operators over that period), small facilities are still very common in the public sector, as well as in highly regulated industries such as banking and finance.

Owing to their known inefficiencies (much like the high proportion of aging servers), smaller data centers are likely to be using more energy (proportionately) than larger facilities and doing a much smaller proportion of the work. There still remains a large opportunity for consolidation, whether into large enterprise data centers, colos or the cloud.

Summary, conclusions and opportunities

Data shows that the past decade has seen substantial gains in IT and data center energy efficiency. This can largely be put to three factors: A big focus on M&E infrastructure, improved designs and equipment, and optimized management and processes. This is highlighted and emphasized by the way PUE data has evolved over time. However, for IT, efficiency gains were largely driven by steadily increased equipment throughput relative to power consumption, thanks to technological advancement driven by the likes of Moore's law. There has been little active adoption of energy conservation best practices by operators (as demonstrated by the EU CoC data above).

With PUE gains plateauing, the major energy-saving opportunities in the future will have to come from shifting the focus to IT optimization best practices. The five main opportunities that we believe will present the best savings opportunities are:

1. **Optimize the server refresh lifecycle.** With 40% of deployed servers in the sector older than five years doing just 7% of the work and consuming more than 66% of energy, the optimization opportunity is enormous.

2. **Increase server utilization.** Uptime Institute's analysis shows that increasing server utilization can yield much more savings than reducing PUE, for much less upfront investment. With average sector utilization levels at around 25%, there is plenty of low-hanging fruit to be had, at least until we've reached the 40% average utilization mark.

3. **Right-size redundancy based on workload requirements.** While it seems easier to design for the highest common denominator in terms of workload resiliency requirements, best practices show that right-sizing redundancy levels (and tracking the appropriate KPIs to reflect it) can yield energy savings of up to 90%.[5]

4. **Consolidate infrastructure to benefit from economies of scale.** The industry has come a long way in terms of consolidation, but in a recent study, 80% of data centers still contained fewer than 25 racks. This represents a significant consolidation opportunity that could be achieved through internal organizational infrastructure rationalization or a move to colo or the cloud.

5. **Address energy consumption across traditional design boundaries.** Work over the past decade has shown that the next order of magnitude in energy efficiency will only be achieved by working across traditional design boundaries and organizational silos, engaging business, IT and infrastructure teams.[3]

[5] Rabih Bashroush and Eoin Woods. 2017. "Architectural principles for energy-aware internet-scale applications." **IEEE Software** vol. 34(3): 14-17. https://ieeexplore.ieee.org/document/7927928

Appendix

Note 1. Information and communications technology energy demand

Several reports have been published in recent years on IT energy consumption and its predicted growth rates. An IEA report published in 2019 noted that workloads and internet traffic will double, but it also forecast that data center energy demand will remain flat to 2021, due to efficiency trends. It cited various references for the basic research.

But Uptime Institute Intelligence is very wary of this prediction and is collaborating with various parties to research this further. There are very strong factors driving up IT energy consumption, and some of the existing data on IT energy use contradicts the IEA figures. The IEA report, for example, stated that global data center energy consumption was 195 TWh in 2017 and is expected to drop slightly by 2021. However, research by the EURECA Project found that European data centers consumed 130 TWh in 2017, and Greenpeace put energy consumption by the Chinese data center industry at 160 TWh in 2018. This suggests an annual total for China and Europe alone in the neighborhood of 290 TWh, far higher than the IEA global figures.

Note 2. Server-side Java operations as a benchmark

Server-side Java operations were adopted in 2007 by the SPECpower benchmark as a way to model the energy efficiency of servers by measuring ssj_ops/watt at various server utilization levels. Since then, the results from hundreds of servers have been assessed and posted on the SPECpower database, making it the most comprehensive and publicly available dataset.

There are other workload types that could be used for benchmarking (e.g., LINPACK), but these are not representative of commercial scenarios, as most operators do not run their servers at peak loads. Other workload types, such as Floating Point Operations Per Second, or FLOPS (used by the Green 500 list), are mostly suited to scientific or high performance workloads.

A more recently introduced SPEC benchmark, the server efficiency rating tool, or SERT, uses different worklets to simulate load on CPU, memory, and I/O. However, at this stage, there is very limited server performance data available using the SERT benchmark. Additionally, SERT allocates different fixed weightings for CPU, memory and I/O, which might not be reflective of the different types of workloads.

The SPECpower dataset used in this study was the most up-to-date version available as of the date of this report (with latest performance data uploaded in August 2019).

Finally, to calculate the energy consumption needed to run a fixed workload at various utilization levels, server age and PUE, we used the models described in Bashroush (2018).[3]

Note 3. Changing IT load and PUE

For the purposes of the models discussed, we have assumed that the PUE will remain constant even when the overall IT load energy consumption is reduced, either through server updating or increasing utilization.

This assumption enables us to make clear comparisons and to measure the impact of hardware refresh and utilization levels on IT energy use. In practice, of course, it will not be so simple: When inefficient servers are replaced with more efficient ones, the IT energy consumption, and the overall energy consumption, will drop; however, the M&E energy will usually not drop proportionately — and may not even drop at all. The extent of the drop will depend on the responsiveness of the M&E equipment and the design and configuration of the electrical and cooling systems.

In practice, the PUE is likely to increase significantly when the IT energy load reduces — even though overall efficiency will have improved. This higher PUE signals a new opportunity to increase IT capacity or improve the efficiency/responsiveness of the M&E equipment and is another reason why IT and M&E teams should work closely together.

Uptime Institute | INTELLIGENCE

ABOUT THE AUTHORS

Dr. Rabih Bashroush is Uptime Institute's Research Director. An international authority in the field of information and communication technology digital infrastructure and consultant for various European governments (including the European Commission), Dr. Bashroush has led numerous research programs designed to provide a better understanding of digital infrastructure, its impact, and how it can operate more effectively. Contact: rbashroush@uptimeinstitute.com

Andy Lawrence is Uptime Institute's Executive Director of Research. Mr. Lawrence has built his career focusing on innovative new solutions, emerging technologies, and opportunities found at the intersection of IT and infrastructure. Contact: alawrence@uptimeinstitute.com

ABOUT UPTIME INSTITUTE

Uptime Institute is an advisory organization focused on improving the performance, efficiency and reliability of business critical infrastructure through innovation, collaboration and independent certifications. Uptime Institute serves all stakeholders responsible for IT service availability through industry leading standards, education, peer-to-peer networking, consulting and award programs delivered to enterprise organizations and third-party operators, manufacturers and providers. Uptime Institute is recognized globally for the creation and administration of the Tier Standards and Certifications for Data Center Design, Construction and Operations, along with its Management & Operations (M&O) Stamp of Approval, FORCSS® methodology and Efficient IT Stamp of Approval.

Uptime Institute – The Global Data Center Authority®, a division of The 451 Group, has office locations in the US, Mexico, Costa Rica, Brazil, UK, Spain, UAE, Russia, Taiwan, Singapore and Malaysia. Visit uptimeinstitute.com for more information.

All general queries:
Uptime Institute
5470 Shilshole Avenue NW, Suite 500
Seattle, WA 98107 USA
+1 206 783 0510
info@uptimeinstitute.com