Uptime Institute® INTELLIGENCE

PLANNING & STRATEGY

UI Intelligence report 29

Uptime Institute data center supply-side survey 2019

Lead Author
Dr. Rabih Bashroush, Research Director, Uptime Institute

Contributing Authors
Rhonda Ascierto, Vice President of Research, Uptime Institute
Andy Lawrence, Executive Director of Research, Uptime Institute

This report examines the supply-side response to Uptime Institute's 2019 global data center survey, providing vendor and influencer insights into industry trends. While data center spending is on the rise, forecasting capacity requirements remains a top challenge, followed by the need to maintain cost-competitiveness compared with cloud. Efficiency technologies are becoming more common, as are limited deployments of lithium-ion batteries.

30-45 minutes to read

UII-29 v1.0 published July 26, 2019 last updated July 26, 2019

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 https://uptimeinstitute.com/ui-intelligence.

KEY FINDINGS

• Spending is rising, and not just for hyperscale operators — smaller facilities (under 20 MW) are also spending more, with some investing in higher levels of redundancy at primary sites. The spending is strategic and in line with the ongoing trend toward consolidation; 17% of suppliers say that customers expect to shut down at least one major data center.

• Forecasting data center capacity and maintaining cost-efficiency are top challenges for operators, their suppliers say.

• The edge of tomorrow is not here just yet. Although anticipation of a surge in edge capacity is still strong, data center engineers, consultants and vendors are becoming more cautious about the rate of adoption.

• Lithium-ion batteries are becoming more accepted in data centers. As predicted in our 2018 supplier survey, more than half of respondents now say that at least some of their customers are deploying lithium-ion batteries in their facilities. This may open new applications for intelligent use of energy storage in the future.

• Adoption of efficiency technologies — for example, free air economizers, data center infrastructure management software, prefab data centers and smart power distribution units — is now high and the technologies have become mainstream.

• The staffing crisis will likely worsen: artificial intelligence and robots won't solve the ongoing skills shortage.

Introduction

Uptime Institute has been conducting a large global survey of the data center industry for nearly a decade. In our 2018 survey, we separated out suppliers and influencers in the industry from the wider group of operators and asked this group different questions.

We continued this practice in 2019. The results of the operator survey were published earlier this year (**Uptime Institute global data center survey 2019**).

Our 2018 supplier findings suggested rapid growth in free cooling and smart power, with artificial intelligence (AI) and micro sites coming soon. The 2019 supplier survey confirmed many of the results of our initial 2018 survey and, in fact, suggested accelerated adoption of prefab, smart power distribution units (PDUs), data center infrastructure management (DCIM), free air cooling economization and lithium-ion (Li-ion) batteries. The findings suggest that the industry is willing to adopt new technologies that improve efficiency and resiliency, even if the pace of adoption is slow in some cases.

Respondents to this year's survey also say forecasting future capacity requirements continues to be a top challenge for their customers,

followed by the need to maintain competitive and cost-efficient operations compared with cloud/colocation venues.

End-user data and supplier data are not always consistent, as vendors (which in our survey includes consultants, design engineers and product providers) can hold different opinions from IT managers and data center operators. For example, vendors were more than twice as likely to say that AI will help address the severity of the widely discussed staffing and skills shortage in the next five years.

A full description of the 2019 supply-side survey demographics is provided in the **Appendix**.

Data center spending is on the rise

Suppliers say they have witnessed higher than normal data center spending patterns (39% of respondents), with a further 45% reporting consistent spending, as shown in Figure 1. This is in line with general market trends, where the demand is driven by continued growth in data and digital services.

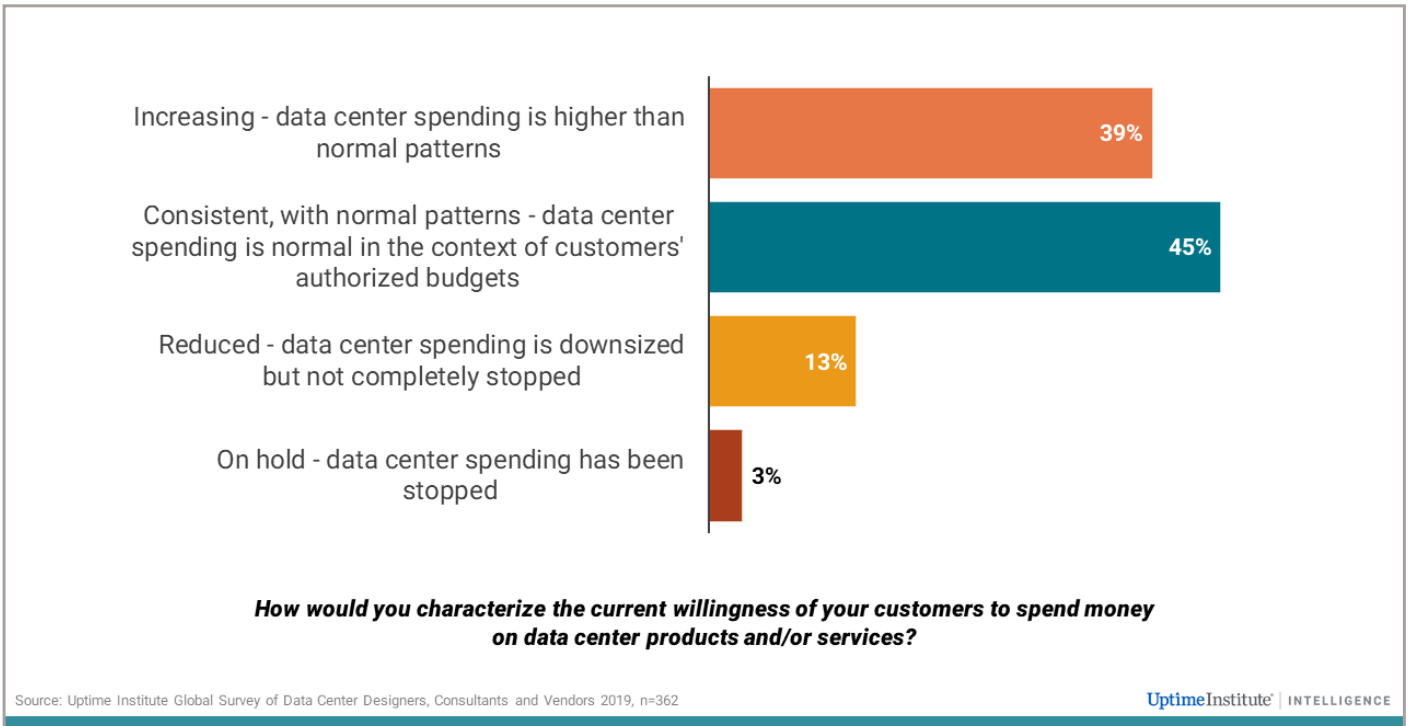

Figure 1. **Respondents report consistent to higher than normal data center spending patterns**

The increase in spending was not only by hyperscalers, which are known to be spending heavily and to be designing their new infrastructure for 100x scalability and building for 10x growth. Smaller facilities (under 20 MW) also saw continued investment, including in high levels of redundancy at primary sites, with overall redundancy levels remaining surprisingly high: 53% of respondents report an increase in power redundancy and 42%, an increase in cooling redundancy (see Figure 2).

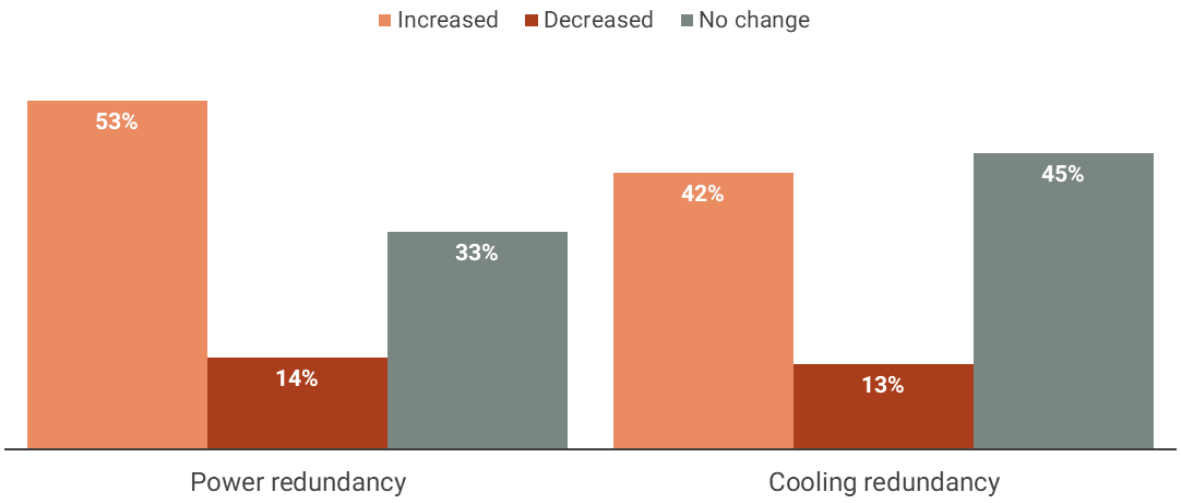

Thinking about the majority of your customers, how have redundancy levels changed in the past 3-5 years in their primary data centers? Please pick one per power and cooling.

Source: Uptime Institute Global Survey of Data Center Designers, Consultants and Vendors 2019, n=359

Uptime Institute | INTELLIGENCE

Figure 2. **Power and cooling redundancy in customers' primary data centers have increased**

Comparing power and cooling redundancy levels with last year's survey results, we notice that N+1 remains dominant for both, with a slight increase in N+2 for power and 2N for cooling, as shown in Figure 3. Note that these trends are being witnessed at customers' primary data center sites; across the industry, the picture might be slightly different. For

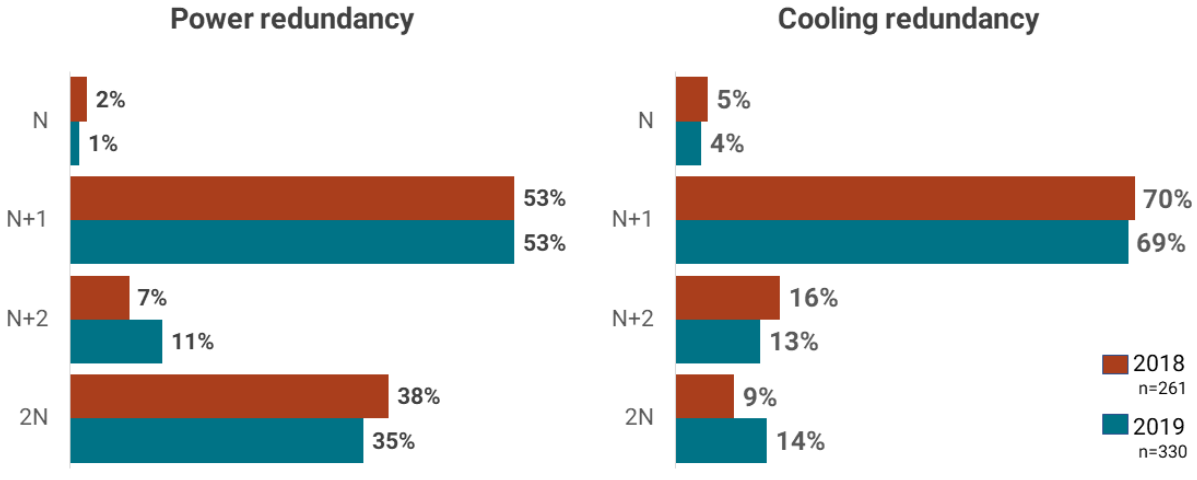

For your customers' primary data centers, which of the following best describes the most common level of redundant power and/or cooling equipment? Please choose only one option in the power column and one in the cooling column.

Source: Uptime Institute Global Survey of Data Center Designers, Consultants and Vendors 2019

Uptime Institute | INTELLIGENCE

Figure 3. **N+1 power and cooling redundancy configurations remain dominant**

example, in hybrid/distributed environments, the move is more toward application-level resiliency, as per the findings from our 2019 operator survey.

There is a clear trend toward the consolidation of regional and smaller facilities into centralized, premium facilities for enterprises. This is reflected in closure figures, as 17% of respondents indicated that most of their customers plan to close at least one of their major data centers, and a further 24% saying that, within five years, most of their customers will own a data center larger than 20 MW. The latter could also reflect the growth of large and hyperscale facilities supporting cloud computing and various other third-party compute and data center services.

With this backdrop, almost half of suppliers (49%) reported that large customers (with data centers 20 MW or larger) often seek projects to be delivered on timelines, budgets, or at a scale that prove challenging for them. This seems to be driven by the speed and scale of growth requirements, as well as the continued need to increase cost-efficiency. However, the majority of suppliers (more than half) still reported that they believe they work with large customers successfully. If this is to be believed, it suggests that the supplier market is currently keeping pace with the surge in demand — at least for now.

Planning and cost challenges persist

The two top challenges identified by suppliers this year were forecasting future data center capacity requirements (28%, similar to last year's supplier survey) and the need to maintain competitive and cost-efficient operations compared with cloud/colocation providers (18%). Managing different data center environments dropped to fourth place (14%), after coming second in last year's supplier survey. This finding is in line with the results of our 2019 operator survey, which attributed the change to the advancement in tools and models, as well as market maturity.

Figure 4 shows the top challenges operators faced in 2018 and 2019, as reported by their suppliers.

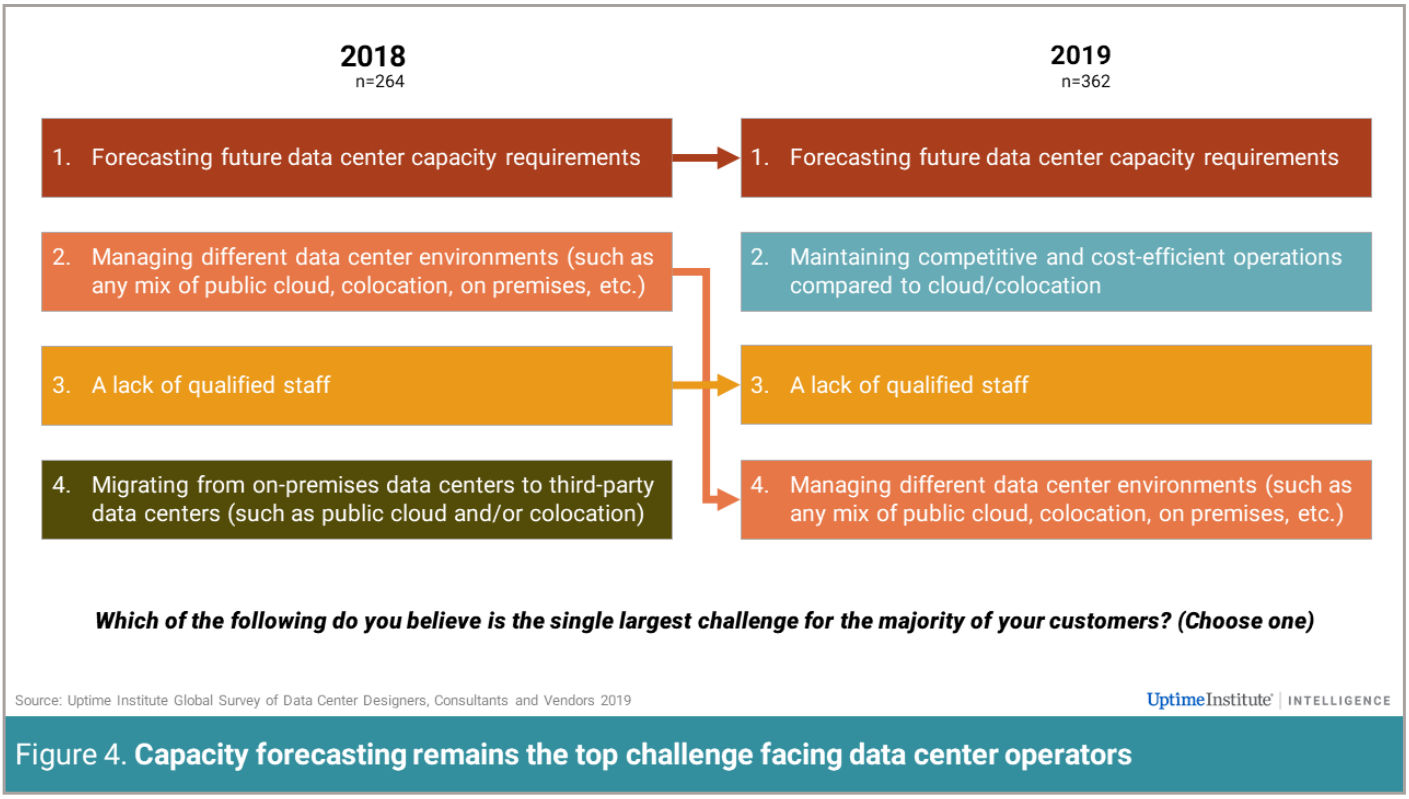

Figure 4. **Capacity forecasting remains the top challenge facing data center operators**

Forecasting data center capacity is a long-standing challenge, not the least due to rapid changes in technology, as well as the difficulty of anticipating future workload growth and the corresponding IT equipment needs. Overprovisioning capacity, the most commonly adopted strategy, leads to inefficiencies in operations (and unnecessary upfront investment), while underprovisioning capacity is an operational risk and could also mean facilities reach their limit before their planned investment lifecycle.

Depending on the sector and type of workload, many organizations currently resort to modular data center designs to alleviate the expense of overprovisioning and may move highly unpredictable workloads to public cloud environments where appropriate. These strategies, plus various other factors driving the uptake of mixed IT infrastructures, mean more organizations are accumulating expertise in managing hybrid environments, which may explain why the challenge of managing different data center environments dropped to fourth place in our survey this year. Additionally, cloud computing suppliers are offering more effective tools to help customers better manage their costs when running a high volume of workloads using cloud services.

Due to the adoption — or at least the aspiration — of cloud-first policies in many organizations over the past few years, managers are having to demonstrate more than ever the cost-effectiveness of maintaining private in-house data centers compared with public cloud, managed service or colocation environments.

This is also reflected in the 2019 Uptime Institute operator survey, where 40% of participants indicated that they are not confident in their organization's ability to compare costs between provisioning workloads to owned sites, colocation facilities and cloud venues. Indeed, this is not a straightforward exercise. On the one hand, calculating the cost of running owned sites is tricky given how enterprises are structured (e.g., energy, staffing and maintenance costs are often split among different budgets). On the other hand, calculating the true cost of moving to the cloud is also not straightforward — there may be costs inherent in the transition related to application re-engineering, potential repatriation or network upgrades, for example. This, in addition to other risks, such as vendor lock, is now driving many organizations and public bodies to change their policies to be more about cloud appropriateness, rather than cloud-first.

The edge of tomorrow is not here yet

The development and adoption of mobile edge computing and the Internet of Things is leading to a proliferation of connected devices. The high volume of data produced and the latency and resiliency requirements of many applications and services has led to a resurgence in distributed IT, with an increased need for processing, storage and, hence, data center capacity at the edge.

The Uptime Institute supplier survey finds that, while there is much hype surrounding edge computing, a surge in demand has yet to materialize. In this year's survey, 41% of data center engineers, consultants and vendors say that most of their customers will own very small edge data centers (less than 150 kW per data center) within the next five years. While this is a significant proportion, it is substantially down compared with last year's survey, where the figure was 63%, as per Figure 5. This drop may indicate that suppliers are becoming more realistic about the adoption rate and more cautious of inflated expectations surrounding edge computing, which is in agreement with the findings from our end-user survey. That survey found the proportion of overall work performed at the edge is expected to be flat or fall slightly from 2019 to 2021.

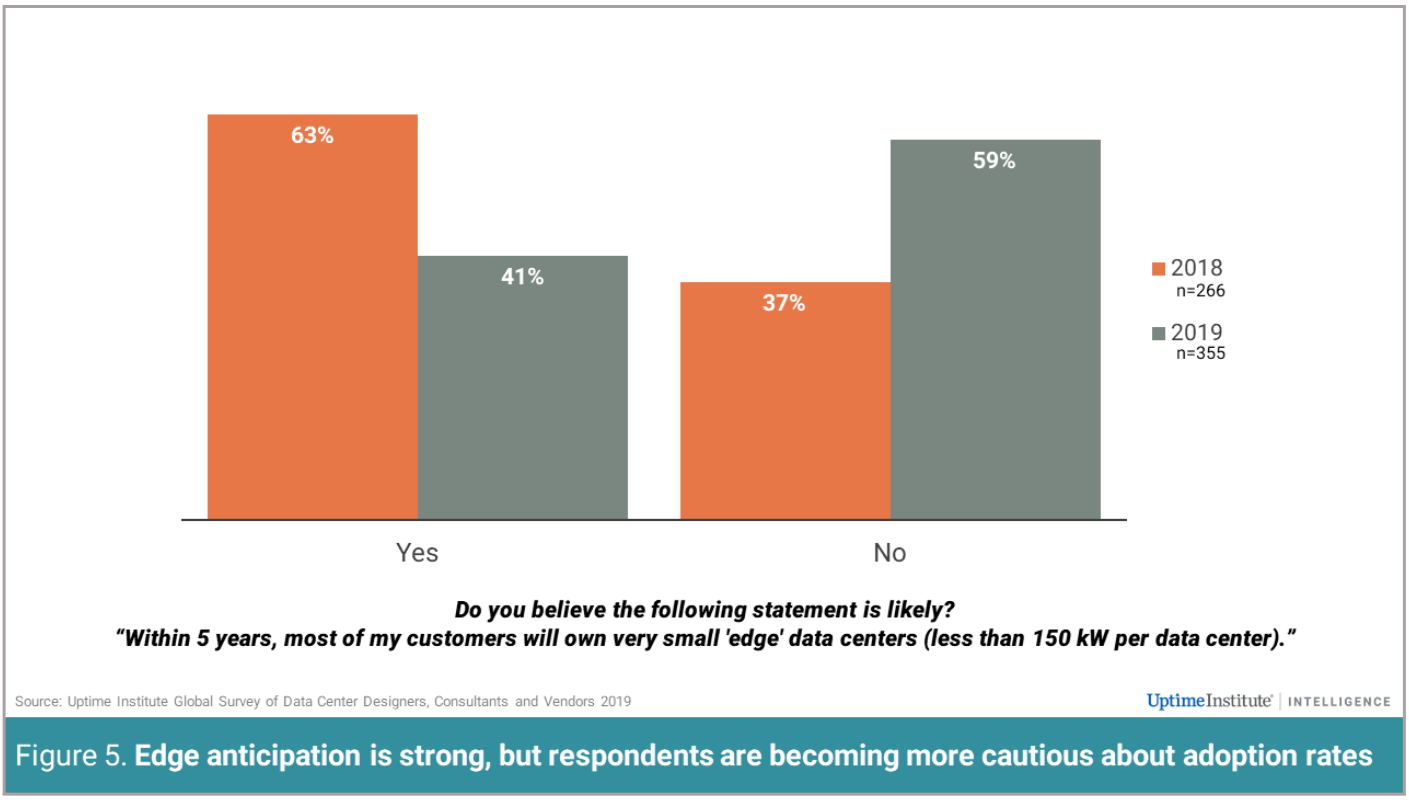

Figure 5. **Edge anticipation is strong, but respondents are becoming more cautious about adoption rates**

This also agrees with the analysis published in our report **Ten Data Center Industry Trends in 2019**, which predicted that 2019 will not be the breakthrough year for edge adoption at scale due to business and technology issues that remain outstanding. (For more forward-looking analysis, read the full report here.)

Technology trends

Data center engineers, consultants and vendors were asked about the adoption of various technologies by their customers. The results are discussed below.

Lithium-ion is becoming more accepted in data centers

Lithium-ion batteries continue to mature in terms of density, safety profile and cost (thanks primarily to the electric car industry) compared with valve-regulated lead-acid batteries, which have not changed much in recent years. Li-ion, although currently more expensive, can be charged and discharged thousands of times, opening up the possibility of using energy storage more dynamically and under intelligent control. It also has a higher power density, enabling operators to free up space.

Li-ion adoption is becoming more accepted in data centers, with more than half of respondents in this year's supplier-side survey indicating that at least some of their customers are deploying Li-ion batteries (see Figure 6). This is a substantial increase that was predicted by our survey respondents last year, with 30% saying that at least some of their

customers deployed Li-ion batteries at scale or in limited projects, and a further 23% saying that at least some of their customers were planning limited projects in 2018.

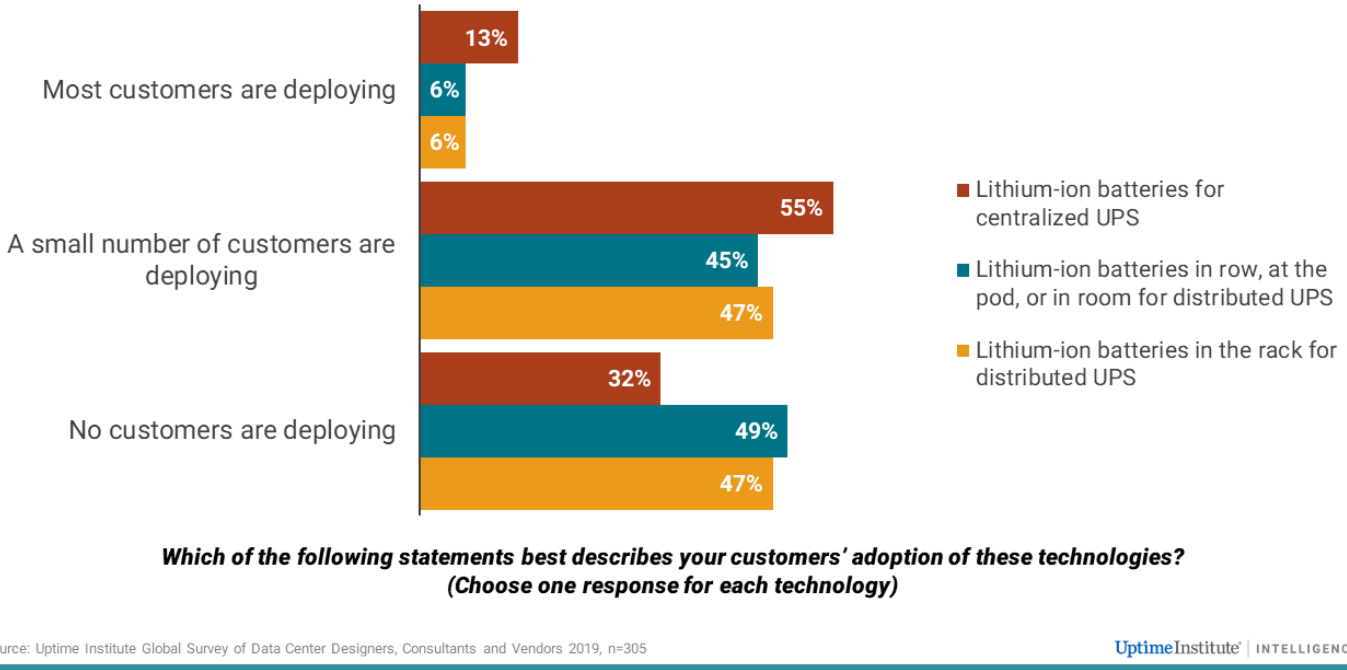

Figure 6. **More than half of respondents say that some of their customers are deploying Li-ion batteries**

According to participants, the most common Li-ion deployments are in centralized uninterruptible power supplies (UPSs), followed by distributed UPS at rack level, and then distributed UPS at the row, pod or room level. This is in line with deployed UPS architectures in the industry and corresponds with the findings from our 2019 operator survey. Overall, no single architecture for Li-ion adoption is dominant, and it appears that, over time, the use of batteries in the white space will become commonplace.

One recent development that could potentially impact Li-ion adoption rate and deployment architecture is the introduction of the National Fire Protection Association's (NFPA's) new standard 855 — Standard for the Installation of Stationary Energy Storage — that passed the voting stage in June 2019. Although it is a US-based standard, it could impact any jurisdiction that follows NFPA. The Standard was introduced to provide safety guidance on the use of batteries, including the design, deployment and testing of UPS systems for Li-ion deployments over 20 kWh (70 kWh for lead-acid). Unfortunately, the data center sector made almost no contribution to the Standard's creation and reacted slowly when it was being reviewed, although it could have a profound impact on the industry.

It is worth noting that NFPA 855 could potentially apply retrospectively to existing facilities. Some of the main points when considering the Standard are its relevance to the facility location; the cap on deployed capacity for certain battery groupings; minimum separation between

groupings and surrounding walls (which impacts space requirements); and fire suppression requirements for all battery deployments deemed within scope — given current limits, this applies to the vast majority of data centers. Additionally, there are other requirements around ventilation, detection, emergency plans and more, so any organization considering a new design/deployment should check with the authority having jurisdiction or the relevant local authority to see if the new standard (or a derivation of it) might apply.

DCIM and smart PDUs are now common

DCIM adoption has been famously slow, but our survey data suggests deployment has seen another increase year-on-year — it may now be said that the technology is in common or mainstream use. This year, 90% of respondents say that at least some of their customers are using/deploying DCIM, with 42% indicating that most of their customers are using/deploying DCIM. Last year (when slightly different questions were asked), 41% of participants indicated that most of their clients were deploying DICM at scale or in limited projects, with a further 34% saying that only large clients were deploying DCIM at scale or in limited projects, as per Figure 7.

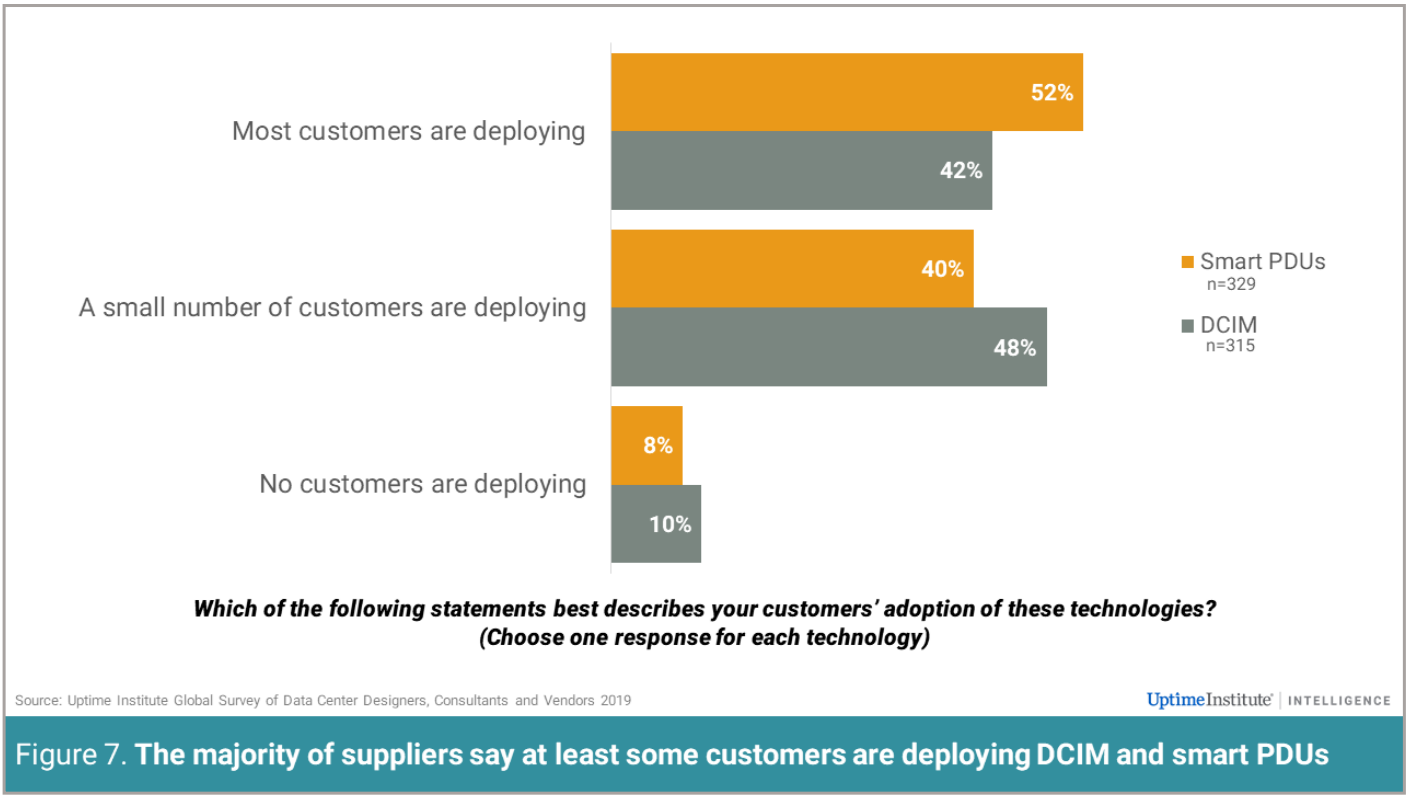

Figure 7. **The majority of suppliers say at least some customers are deploying DCIM and smart PDUs**

Uptime Institute expects that adoption rates will likely continue to rise as operators develop a better understanding of what DCIM can and cannot offer, the level of investment required for value realization and the minimum facility size above which DCIM economics start adding up (note that when we refer to "DCIM," we include cloud-based data center management as a service). The continued introduction of new DCIM services, which may offer AI-based solutions, further strengthens the

case for DCIM, allowing operators to draw more value from the large amount of data readily collected.

A similar increase in adoption patterns is being witnessed with smart PDUs. “Smart PDU” is an umbrella term used to refer to networked power distribution units capable of providing a number of features such as power measurement and monitoring, environmental monitoring, power switching and integration with DCIM tools.

This year, 92% of respondents say that at least some of their customers are deploying smart PDUs, with 52% indicating that most of their customers are deploying the technology (see Figure 7). The strongest smart PDU deployment figures came from the Middle East and Africa, where 74% of participants say that most of their customers are deploying smart PDUs, followed by Asia-Pacific (including China) at 68%.

Last year, 48% of participants in our supplier survey indicated that most of their clients were deploying smart PDUs at scale or in limited projects, with a further 33% saying that only large clients were deploying the technology at scale or in limited projects.

Smart PDU adoption is becoming the norm, with only 8% of suppliers saying that customers are not deploying the technology this year, down from 20% last year. Indeed, participants in 2018 supplier-side survey predicted this year’s increase, as 11% stated that their clients were planning to deploy the technology by the end of the year.

Free cooling economization is on the rise

Free air cooling economization projects continue to gain traction, with indirect free air cooling being slightly more popular than direct air. In our survey, 84% say that at least some of their customers are deploying indirect air cooling (74% for direct air). Only 16% of participants say that no customers are deploying indirect free air cooling (26% for direct air), as per Figure 8.

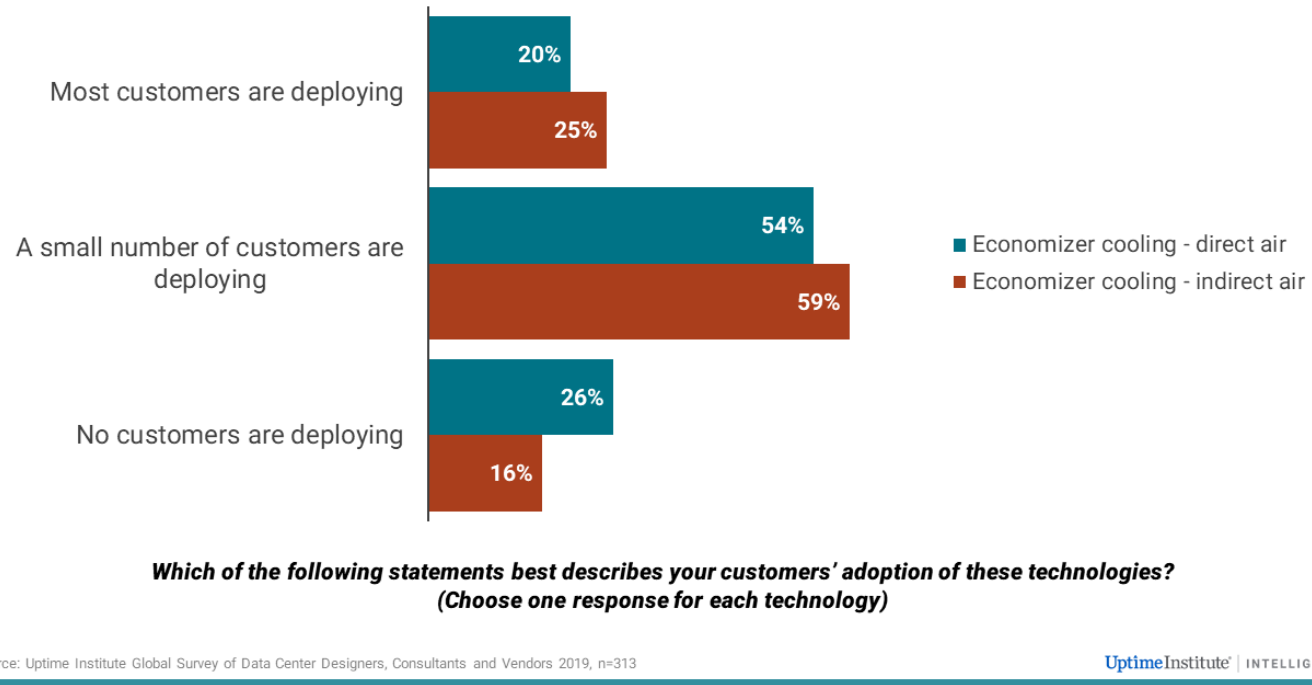

Figure 8. **Free air economization projects continue to gain traction**

The data suggests that there is more momentum behind direct free air cooling in North America than in other parts of the world. Among North American respondents, 70% indicated that some of their customers are deploying direct air cooling (compared with 63% indirect air). As shown in Figure 9, this was not the case in Europe or Asia-Pacific (including China), where suppliers reported that more customers were deploying indirect air. This perhaps could be linked to the way internet giants represent a bigger data center market share in North America when compared with other parts of the world — internet giants are known to favor direct free air cooling when deploying at scale.

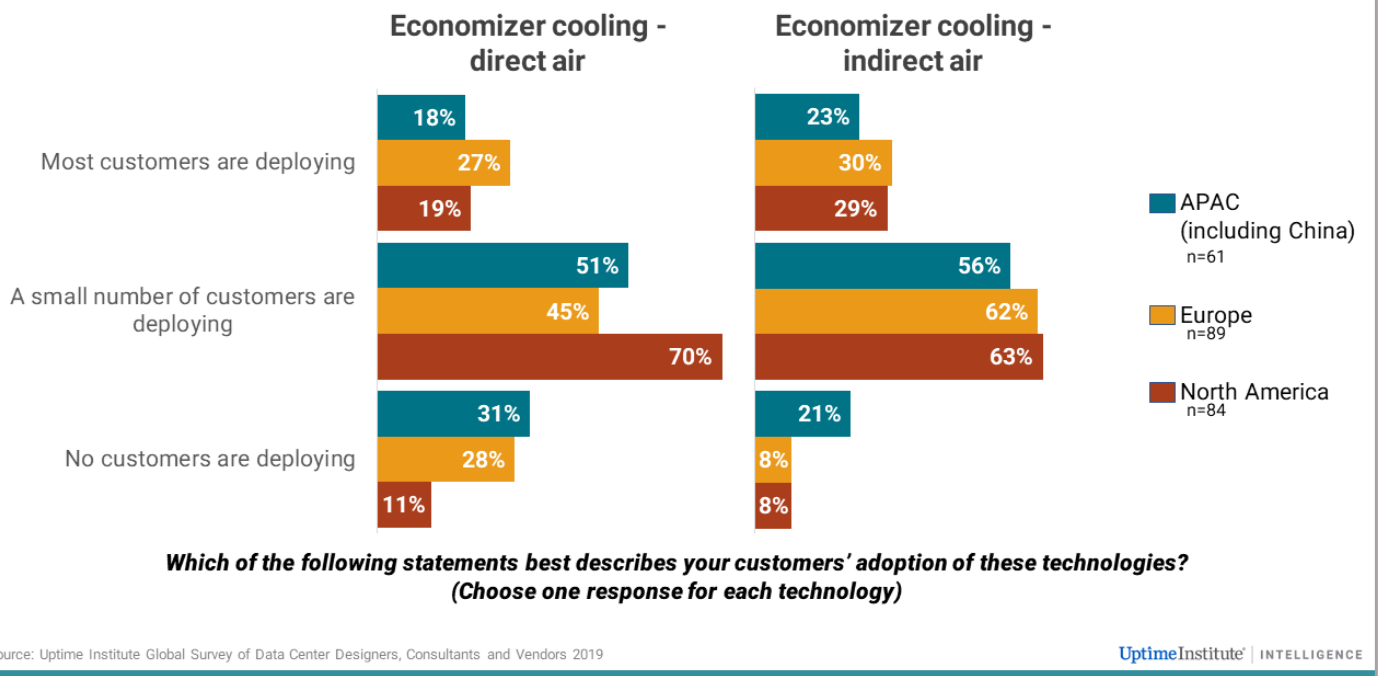

Figure 9. **Respondents say direct free air cooling has greater momentum in North America**

The continued pressure to increase cost-efficiency, as well as the rising awareness and interest in environmental impact, is likely to continue driving uptake of free air cooling. Compared with traditional compressor-based cooling systems, free air cooling requires less upfront capital investment and involves lower operational expenses , while having a lower environmental impact (e.g., no refrigerants, low embedded carbon and a higher proportion of recyclable components).

Yet, some issues hampering free air cooling uptake will likely continue in the short term. These include the upfront retrofit investment required for existing facilities; humidity and air quality constraints (which are less of a problem for indirect air cooling); lack of reliable weather models in some areas (and the potential impact of climate change); and restrictive service level agreements, particularly in the colocation sector.

Moreover, a lack of understanding of the American Society of Heating, Refrigerating and Air-Conditioning Engineers' (ASHRAE) standards and clarity around IT equipment needs is driving some operators to design to the highest common denominator, particularly when hosting legacy or mixed IT systems. As such, the opportunity to take advantage of free air cooling is missed, due to the perceived need to adopt lower operating temperatures. At least in Europe, this problem might be partially addressed by the introduction of the new European EcoDesign legislation for servers and online storage devices, which will take effect from March 2020. The new legislation will require IT manufacturers to declare the operating condition classes and thermal performance of their equipment. This, in turn, will help enterprise data

centers better optimize their operations by segregating IT equipment based on ambient operating requirements.

Is prefab winning mainstream market uptake?

A prefabricated modular or prefab data center is a technology built using one or more standardized blocks within a factory environment and transported for final on-site integration. A previous study by 451 Research (the sister company of Uptime Institute) found that prefab has turned a corner toward mass adoption. The study evaluated the benefits of prefab over traditional construction approaches, identifying advantages such as better quality control, reduced risk, speed of deployment and modular scalability. (The full report is available here: https://clients.451research.com/reports/95342.)

In this year's Uptime supply-side survey, 72% of our participants say that their customers are either already using prefab (including electrical skids, packaged mechanical plants, etc.) or plan to deploy in a planned future project, with a further 8% indicating that their customers already use prefab but are not likely to do so in a future project (see Figure 10). Of those considering or using prefab, the majority (more than 57%) say that it made a difference if the prefab products had been assessed as Tier-Ready by Uptime Institute. (For clarity, prefab can refer to significant components, not just complete data centers.)

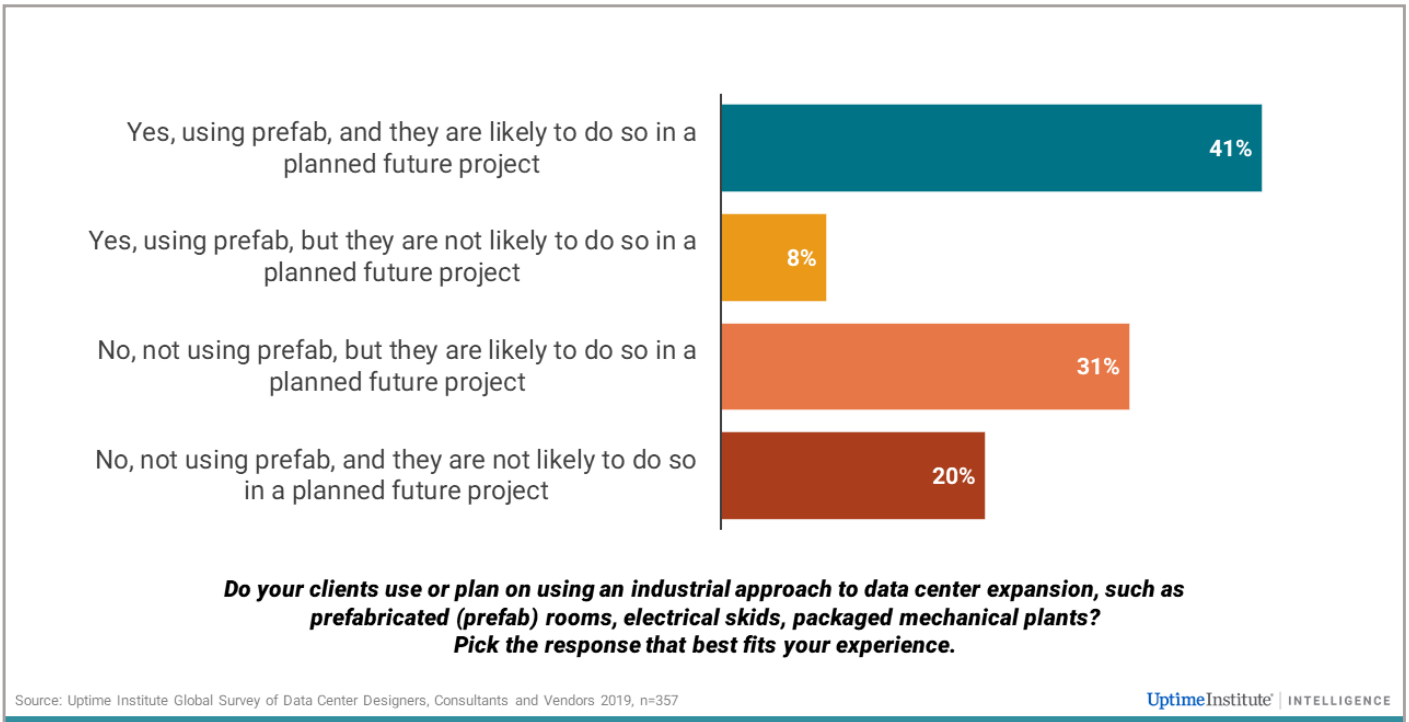

Figure 10. **Prefab data centers seem set to become mainstream**

Analyzing the survey data based on the geographical location of participants, we notice a slightly larger uptake in North America, where 47% of participants indicated that their customers are using prefab and are likely to do so in a planned future project, compared with 44% in the Asia-Pacific region (including China) and 42% in

Europe, the Middle East and Africa (EMEA). However, when asked about customers who do not use prefab but plan to do so in a future project, 24% agreed with this question in North America, compared with 30% in Asia-Pacific and 29% in EMEA. This confirms the findings from the previously referenced 451 Research report that North America is more advanced in its use of prefab products, but the growth rate in Asia-Pacific and EMEA is faster.

Staffing crises

With decades of strong growth behind it and no slowdown in sight, the data center sector is struggling with staffing shortages. In our 2019 survey of IT and data center managers, 61% of respondents said they had difficulty retaining or recruiting staff, up from 55% a year ago. When asked about finding qualified candidates for open jobs, 41% of operators and IT managers said they were having difficulty, up slightly from 37% last year.

We also asked our supply-side respondents about staffing issues and related topics in the future, as shown in Figure 11. Most were optimistic, with just 28% saying that they believe a shortage of data center facilities staff will limit the data center industry's growth in the coming five years.

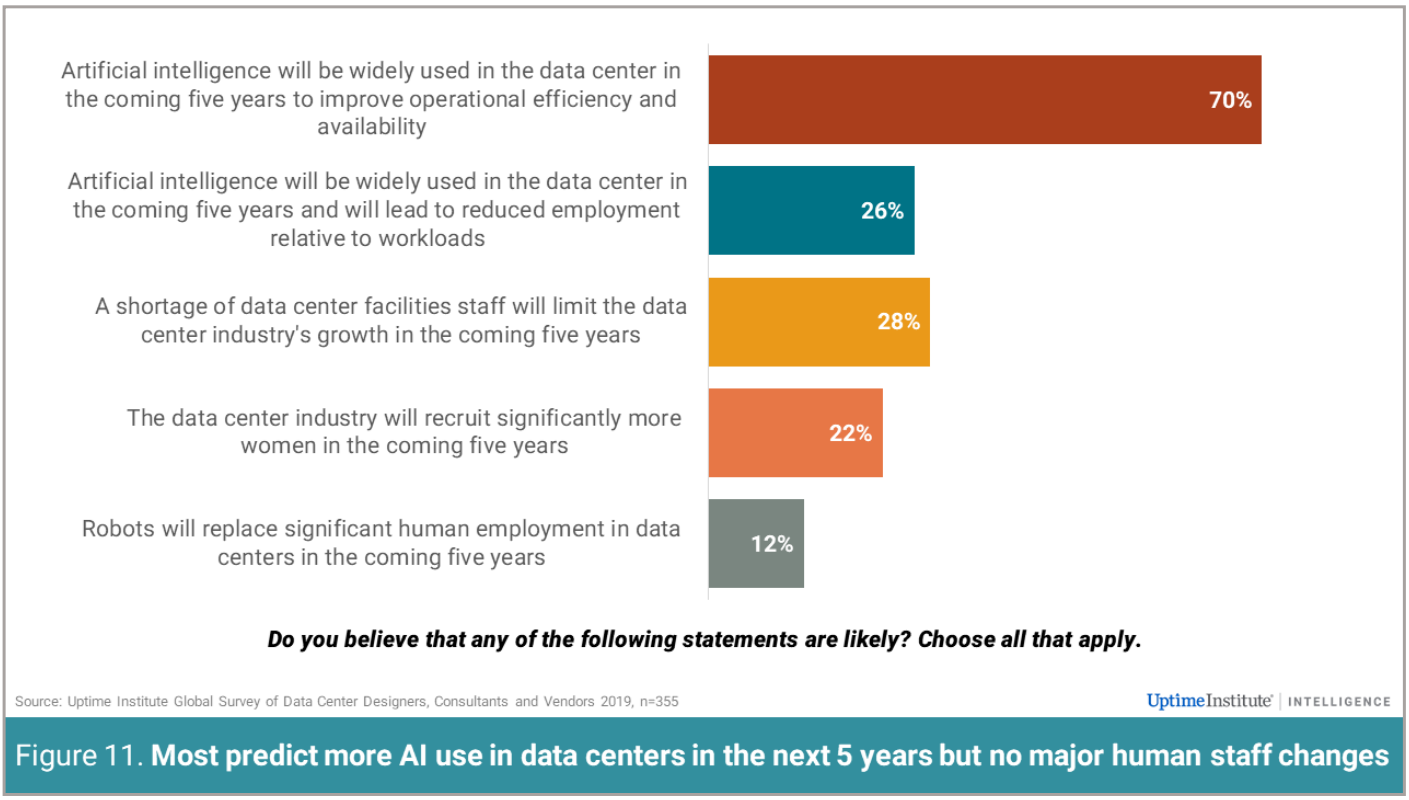

Figure 11. **Most predict more AI use in data centers in the next 5 years but no major human staff changes**

Will technology help alleviate the staffing issues? The majority of supply-side respondents (70%) believe AI will be widely used in data centers in the coming five years to improve operational efficiency and

availability. This echoes similar results in last year's supplier survey: the same proportion of respondents (70%) said AI will be widely used in data centers in the future.

In 2019 we took this a step further, asking whether AI could help fulfill staffing requirements. Just 26% of respondents said they believe that AI will be widely used in the data center in the next five years and will lead to reduced employment relative to workloads. While AI will almost certainly play a key role in future data center robotics, few respondents (just 12%) believe that robots will replace significant human employment in data centers during the same timeframe.

One of the most widely discussed approaches to bolstering the sector's talent pipeline is to address its gender imbalance. Just 5% of operators and managers in our 2019 end-user survey said women represented 50% or more of staff, while one-quarter had no women at all among their build, design or operations staff. While efforts are underway across the sector to attract and retain a wider pool of talent, including women, there is no data to suggest early success.

In our supply-side survey, just 22% say they believed that the data center sector will recruit significantly more women in the coming five years. This is a large decline from last year's supplier survey, when 65% of respondents said the sector would recruit significantly more women in the next three to five years. There could be several explanations for this apparent change in sentiment: While more in the sector are aware of and are discussing gender imbalance, relatively few have changed their hiring processes. Conversely, organizations that have implemented new hiring approaches may be struggling for early wins. Experience suggests that there is no short or even medium-term fix to the gender imbalance, which is culturally ingrained in engineering industries and in the education system.

Each of the questions in Figure 11 was asked separately; they were not grouped under the guise of potential staffing problems or solutions. Yet taken together they may be reasonably interpreted as a snapshot view — from the supply-side — of the sector's near-term future state of employment: staffing shortages will largely be met, but AI/robotic technologies and concerted efforts to attract more women will play limited roles.

Summary

As with last year's survey, there were divergences between some of the results of the Uptime Institute's 2019 supply-side and operator surveys, primarily around data center skills shortages and its threat to the sector.

For the most part, however, the thinking was aligned. The main findings include:

• Data center spending is rising for large as well as small facilities. At the same time, closures continue as the drive for more data center consolidation advances.

• Capacity forecasting and planning continue to top the challenges faced by customers, followed by the need to maintain cost-efficiency compared with cloud environments.

• Edge anticipation is still strong, though suppliers now have more realistic expectations around adoption rates.

• Lithium-ion battery use is becoming more accepted in data centers.

• Efficiency technologies such as free air cooling, DCIM, prefab, and smart PDUs are now mainstream.

• The staffing crisis will likely worsen, with AI and robots falling short of providing a solution, at least in the short term.

Appendix: Survey demographics

The Uptime Institute Annual Survey, now in its ninth year, is conducted online and by email. The 2019 survey was conducted in March and April 2019. Participants were split (in the survey) into two groups: data center operators (managers, engineers, designers, etc.) and data center suppliers (vendors, designers and consultants). This report includes responses from 504 data center service and equipment suppliers — people responsible for designing and building data centers, as well as suppliers of data center equipment. As shown in Figure A1, the majority are data center consultants (47% of respondents), followed by data center product vendors (27%) and data center design engineers (26%). This is a slight change in job function of participants compared with last year, when the majority were data center design engineers (40%), followed by consultants (36%) and equipment product vendors (24%).

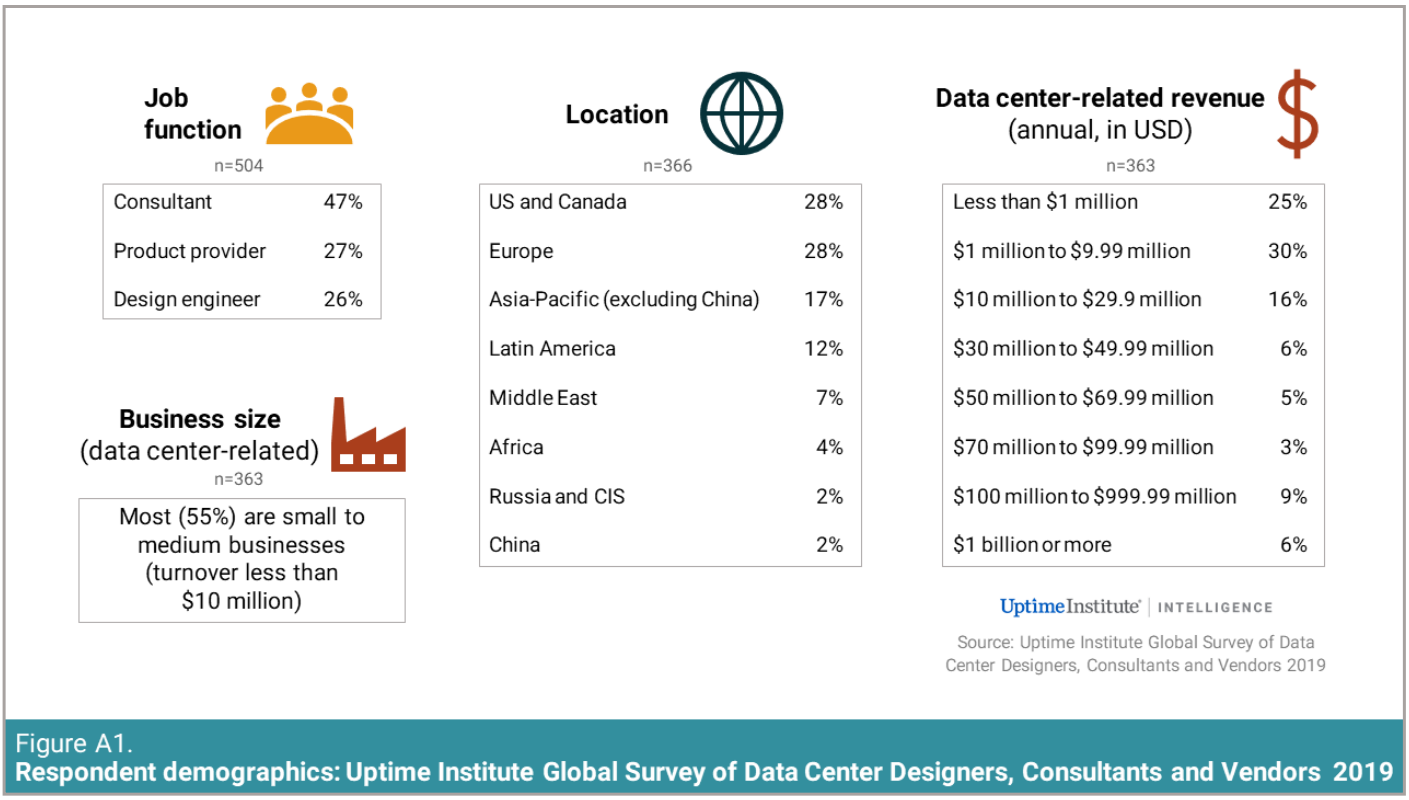

Figure A1.
Respondent demographics: Uptime Institute Global Survey of Data Center Designers, Consultants and Vendors 2019

As for location, more than half of participants were from Europe and North America (Europe, 28%; US and Canada, 28%). This was followed by Asia-Pacific (including China, 19%) and Latin America (12%).

In terms of data center-related revenue (reported in US dollars, for their most recent fiscal year), the breakdown is consistent with the data center market demographics: 30% had $1 million to $9.99 million in revenue, followed by 25% with less than $1 million in revenue. Most respondents were from smaller organizations — just 6% of respondents worked for organizations with more than $1 billion in revenue.

Previous survey findings are available on the Uptime Institute Network member website, Inside Track.

For more information on our surveys or Uptime Intelligence, contact Rhonda Ascierto, Vice President of Research (rascierto@uptimeinstitute.com) or Brenda South, Vice President Communications (bsouth@uptimeinstitute.com).

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

Rhonda Ascierto is Uptime Institute's Vice President of Research. She has spent nearly two decades at the crossroads of IT and business as an analyst, speaker, adviser, and editor covering the technology and competitive forces that shape the global IT industry. Contact: rascierto@uptimeinstitute.com

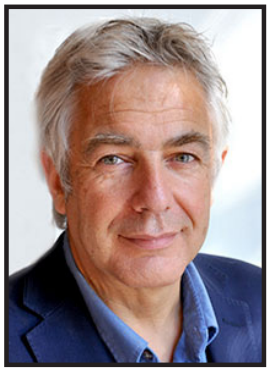

Andy Lawrence is Uptime Institute's Executive Director of Research. Mr. Lawrence has built his career focusing on innovative new solutions, emerging technologies, and opportunities found at the intersection of IT and infrastructure. Contact: alawrence@uptimeinstitute.com

ABOUT UPTIME INSTITUTE

Uptime Institute is an unbiased 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 U.S., Mexico, Costa Rica, Brazil, U.K., Spain, U.A.E., Russia, Taiwan, Singapore, and Malaysia.

Visit www.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

