Uptime Institute®
INTELLIGENCE

PLANNING
& STRATEGY

UI Intelligence report 26

Uptime Institute global data center survey 2019

Lead Author
Andy Lawrence, Executive Director of Research, Uptime Institute

Contributing Authors
Rhonda Ascierto, Vice President of Research, Uptime Institute
Kevin Heslin, Chief Editor, Uptime Institute

The Uptime Institute 2019 survey is the largest and most comprehensive in the data center sector. The findings discussed in this report reveal what operators around the world are thinking, doing and planning in the areas of efficiency, resiliency, workload placement, climate change, staffing and new technology adoption.

30-45 minutes to read

UII-26 v1.6 published April 23, 2019 last updated August 1, 2019

This Uptime Institute Intelligence report covers:

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.

Introduction

The Ninth Annual Uptime Institute Data Center Survey provides an overview of the shape, practices and major trends driving the mission-critical digital infrastructure of today. This survey, the most comprehensive research survey of its kind, was conducted online during March and April 2019. For more details, see the **Appendix**.

Our survey results find that the sector is large, diverse, complex and adapting to change in many different ways. In almost every area under discussion — whether outages, resiliency, staffing, placement of workloads, use of innovation or use of cloud — there is considerable diversity in the strategies being employed. Overall, the industry is open to adopting new approaches and technologies but is doing so cautiously.

Some findings stand out:

• A high proportion of organizations expressed concern at the lack of visibility into the design and operations of public cloud services on which they could be or are dependent. This has hindered adoption of the public cloud.

• There is still strong adherence to the large, privately owned enterprise data center, which accounts for half of all IT workloads currently and is predicted to continue doing so in the near future. Meanwhile, a high proportion of respondents expect to own and run their own edge data centers, perhaps in conjunction with third parties.

• The number of outages reported mapped very closely onto our 2018 survey results. This is not a good thing — a third of those surveyed have suffered some form of outage or serious service degradation in the past year. A number of these had serious financial consequences.

• Meanwhile, the industry faces a staffing crisis, with most reporting difficulty recruiting or retaining staff. And while the lack of women working in data centers is well-known, the extent of the imbalance is notable: One-quarter of operators surveyed had no women at all among their design, build or operations staff.

The findings of the 2019 IT and data center manager survey are discussed in this report. Analysis of previous years' surveys are available on the Uptime Institute Network's site, Inside Track, along with presentations. The results from the designer/supplier/advisor survey 2019 are presented in report 29, **Uptime Institute data center supply-side survey 2019**.

KEY FINDINGS

• The large, privately owned enterprise data center facility still forms the bedrock of corporate IT and is expected to be running half of all workloads in 2021.

• The staffing problem affecting most of the data center sector has become a crisis. Sixty-one percent (61%) of respondents said they had difficulty retaining or recruiting staff — up from 55% a year earlier.

• Outages continue to cause significant problems for operators. Just over a third (34%) of all respondents had an outage or severe IT service degradation in the past year, while half (50%) had an outage or severe IT service degradation in the past three years.

• A lack of visibility, transparency and accountability of public cloud services is a major concern for enterprises with mission-critical applications. A fifth of operators surveyed said they would be more likely to put workloads in a public cloud if there were more visibility. Half of those using public cloud for mission-critical applications also said they do not have adequate visibility.

• Improvements in data center facility energy efficiency have flattened out and even deteriorated slightly in the past two years. The average PUE for 2019 is 1.67.

• Kilowatt (kW) rack density is rising, following a long period of flat or minor increases, causing many to rethink cooling strategies. Uptime Intelligence regards this as a medium- to long-term trend.

• Ten percent (10%) of all respondents said that their most recent significant outage cost more than $1 million ("most recent" could have been at any time in the past).

• The data center sector continues to be dominated by men. Only five percent (5%) of respondents said women represented 50% or more of staff, while a quarter had no women at all among their build, design or operations staff.

OTHER NOTABLE FINDINGS

• Power loss was the single biggest cause of outages, accounting for one-third, as in 2018. Networking issues were close behind, at 31%.

• Distributed resiliency using active-active data centers is becoming more common. Forty percent (40%) of those surveyed said they use availability zones for resiliency — a strategy that requires at least two active data centers replicating data to each other.

• Sixty percent (60%) of respondents said their data center's outage could have been prevented with better management/processes or configuration.

• Most operators surveyed have a hybrid strategy. IT workloads are being spread across a range of services and data centers, with about a third of all workloads expected to be contracted to cloud, colocation, hosting and Software as a Service (SaaS) suppliers by 2021.

• Ownership and management of edge capacity will be diverse, depending on the function and the applications. But a majority of enterprises expect to mostly operate their own edge data centers or to use a mix of their own and third-party operator facilities.

• Confidence in costing best-execution venues is growing. Sixty percent (60%) of respondents said they are confident in their organization's ability to compare the costs of provisioning workloads in the cloud or at their owned and leased colocation sites.

• Climate change is not causing the data center industry a lot of concern. As in 2018, half of operators said they are not currently preparing for climate change, with more than half of that cohort stating that their existing plans are sufficient (30% of total respondents). There were small increases in the number re-evaluating flood risk, site selection and their ability to deal with rising temperatures, compared with last year.

• While most think the industry would benefit from more women, they also think that it is easy for women to pursue a data center career — at least in their organization — and that the lack of women presents no threat to the industry.

• Automation and artificial intelligence (AI) will not reduce data center staffing requirements in the next five years, according to the majority of respondents. After that, however, most think automation will reduce staff requirements.

• Awareness and adoption of both the Open Compute Project (OCP) and Open19 as lower-cost architectures for racks and power distribution is still low, years after their introduction, as is evaluation and deployment. Among data center staff, senior IT management are the most likely to be aware of these initiatives. Nearly half of the data centers represented in our survey now have some lithium-ion (Li-ion) batteries in use or contracted to be installed.

Capacity and workloads: Where does IT live?

It is sometimes convenient to speak of data center capacity as if it were a generic or standardized commodity — perhaps measured in megawatts (MW), square feet or even compute or storage units. By MW of uninterruptible power supply (UPS) power, for example, 451 Research (the sister company of Uptime Institute) has calculated that roughly 60% of global data center constructed capacity today is enterprise; about 20%, colocation; and 20%, cloud.

But where are mission-critical and enterprise loads run, and in what types of data centers? This is a different question. In the 2019 Uptime Institute survey, we asked respondents to estimate, by percentage, how much of their workload/data is processed/stored in different types of data centers. The results, shown in Figure 1, are consistent with separate studies by 451 Research and Uptime Institute.

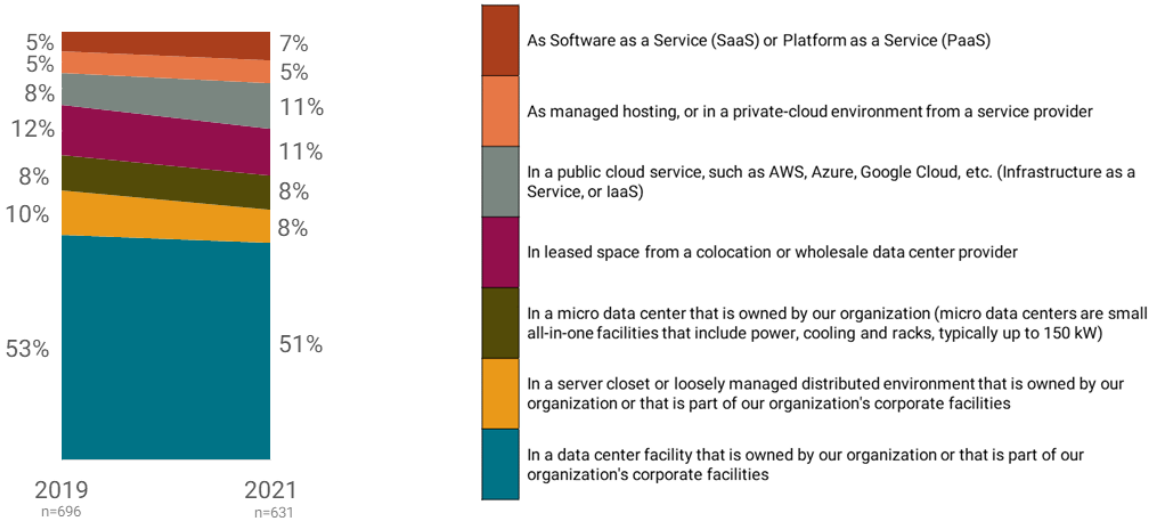

Figure 1. **Only partly cloudy: Corporate IT has many venues**

These findings confirm Uptime Institute's view that, while the enterprise-owned data center sector may not be the most free-spending or the most innovative, it will still form the bedrock of mission-critical IT for the next decade. (Our separate study, reported in **Capacity planning in a complex, hybrid world**, also shows that capacity demand for the enterprise sector is growing despite the cloud and colocation data center build-out.) IT workloads are being spread across a range of services and data centers, with about two-thirds running in privately owned environments (large facilities, server closets or micro data centers) and a third contracted to external suppliers by 2021. While the enterprise data center sector is falling as a percentage of the whole, it is actually growing — slowly — in real terms.

There are clearly many factors determining where workloads (or applications) are run. Software as a Service (i.e., complete applications , such as Salesforce) is growing, as is managed hosting and public cloud — showing that businesses are keen to offload the burden and complexity of managing the full stack of modern IT and applications.

From 2019 to 2021, colocation is expected to fall very slightly as a percentage of direct enterprise workload placement (although growing slowly in real terms). At first pass this is surprising, in that many colocation providers are enjoying strong growth. However, this anomaly is likely due to "secondary" placement — the work is first given to a managed service provider, a SaaS or cloud company, which

then places its workloads in colocation data centers. This can be invisible to the enterprise.

What about small data centers and edge facilities? This sector has been much hyped by vendors and analysts, but the predicted surge in demand has yet to materialize. In our study, 16% to 18% of workloads are reported to be in either privately owned small, standalone data centers (below 150 kW) or in server rooms and closets. Again, the proportion of work is expected to drop slightly, while the overall capacity may, in reality, increase.

Demand for small data centers is a mixed picture. Some smaller data centers are being eliminated by public cloud services; others may be hardened, consolidated and upgraded. Many server rooms are being or will be closed, but more small facilities (perhaps specialist micro data centers) are yet to be widely introduced. Demand drivers — for example, edge analytics, 5G and IoT (Internet of Things) applications — are not easy to accurately forecast, and many depend on innovations and business cases not yet fully imagined. As new services and applications come on line, we expect edge capacity to rise; many ad hoc server rooms and small data centers will eventually be supplanted by remotely managed, optimized purpose-built edge data centers.
As Figure 2 shows, most enterprises expect to meet edge demands either with their own capacity or in conjunction with a third party. When comparing our survey data from 2018 to 2019, there has been a slight shift toward more organizations expecting to meet edge demands using their own private data centers and edge facilities. Overall, ownership and management of edge capacity will be diverse and will depend on the function and specific applications.

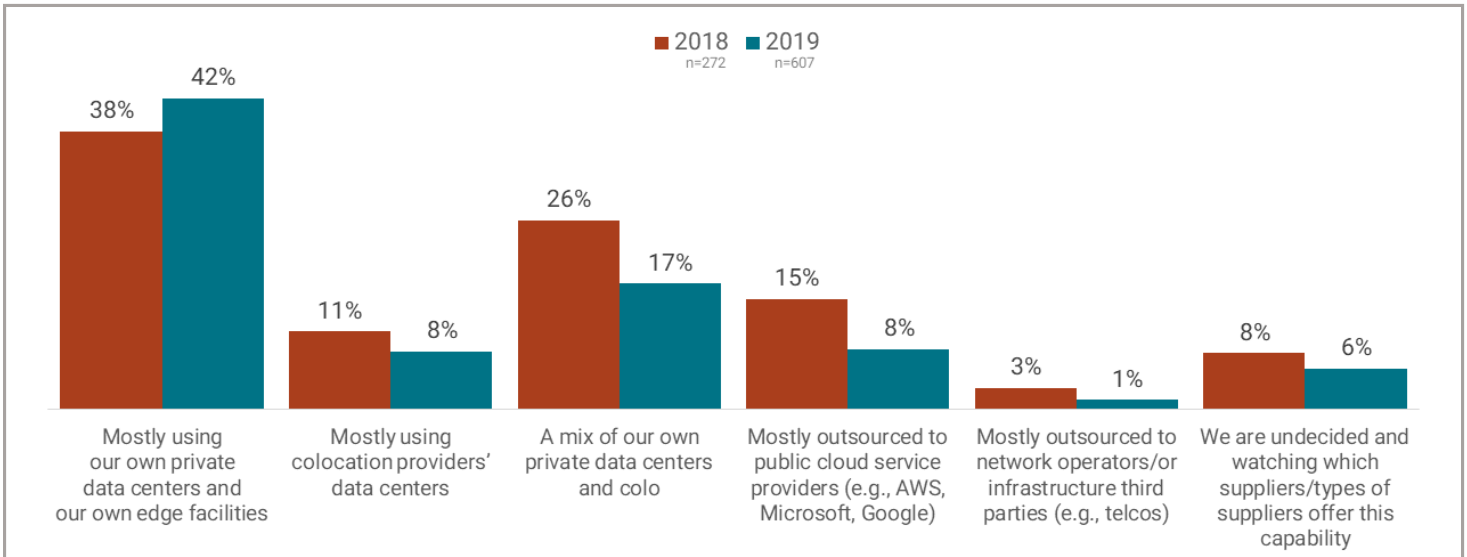

Figure 2. **Ownership and management of edge data centers will be diverse**

While many new edge data centers will be deployed by telecom companies and other commercial operators, it is clear that many enterprises (retailers, banks and manufacturers are examples) expect to own and manage their own small edge facilities.

Colocation and cloud placement

A lot of research has been conducted on how operators choose colocation and cloud companies and what drives their decision. In this survey, we identified two factors of interest: First, 77% view good connectivity (connections to other clouds, services or partners/customers, etc.) as critically or very important in their choice of a colocation partner; and second, 61% considered Uptime Institute Tier Certification status as part of the evaluation process.*

One of the barriers in placing workloads in colocation data centers or in public cloud services has been the difficulty of assessing true costs (visibility and resiliency is also an issue — see below). But confidence is growing: sixty percent (60%) said they are confident in their organization's ability to compare costs between provisioning workloads to owned sites, colocation facilities and cloud options. This is likely due to better tools and models, market maturity and greater attention to costs as the work placement is managed more formally.

Slight fall in energy efficiency

The power usage effectiveness (PUE) metric developed in 2007 has become the de facto standard for measuring the energy efficiency of a data center facility (excluding the efficiency of the IT equipment). PUE was not intended to be a comparative metric (data centers are too dissimilar for that), nor was it intended to be used for tracking the overall energy efficiency of the data center industry. Even so, if the sample is large enough and values are tracked over time, the data can be useful to show how the industry is progressing. Uptime Institute has tracked the numbers, at intervals, over 12 years (see Figure 3).

* In this survey, we asked a small number of questions related to areas in which Uptime Institute has a direct interest. Primarily for control and analysis purposes, we also asked whether the respondent's data center has Uptime Institute Tier or Operations Certification or whether their organization is a member of the Uptime Institute Network. Sixty percent (60%) of those surveyed responded "No" to these questions.

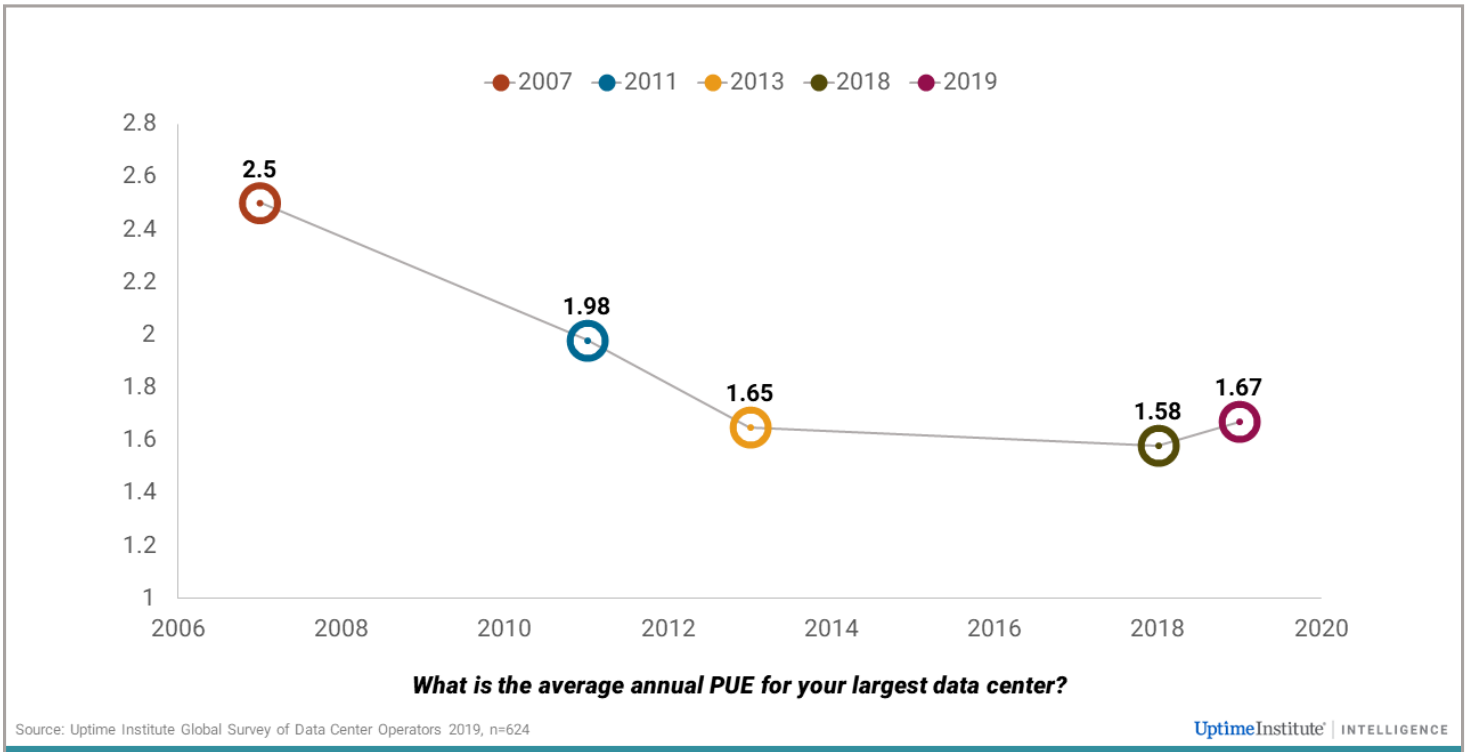

Figure 3. **Data center efficiency gains have stalled**

The latest figure provides a slight surprise: For the first time, there was no recorded improvement in facility energy efficiency. In fact, energy efficiency deteriorated slightly, from an average PUE of 1.58 in 2018 to 1.67 in 2019 (lower is better). The consistency of our data shows, once again, that most operators cannot compete with the finely tuned, aggressively efficient hyperscale data centers in energy efficiency, nor indeed with newer, highly efficient colocation sites. In these sectors, PUE values of 1.1 to 1.4 are frequently claimed.

This marginal PUE increase may be seen in the long-term context of a trend of diminishing returns in energy efficiency, at least in pre-existing facilities. The most dramatic increases in energy efficiency were achieved between 2007 and 2013, often by taking steps such as hot/cold air separation, raising temperatures, or applying more control on cooling, fans and power distribution. The widespread adoption of free air cooling (direct and indirect) in newer builds has also helped to bring the overall level of energy use down. But it is clear that the easiest steps have largely been taken.

Could efficiency improvements have actually gone into reverse? It is speculation, but we think that several factors could have caused a slight, and probably temporary, halt in PUE improvements. For example, the higher and extreme temperatures experienced in many parts of the world where data centers are situated during 2019 could account for increased use of cooling. Another factor is that utilization in many data centers — although certainly not in all — has fallen as certain workloads are moved to public cloud services (see the Uptime

Intelligence report **Capacity planning in a complex, hybrid world**). Low utilization can lead to a reduction in efficiency.

Uptime Institute advises that data center operators closely watch and understand the total energy consumption of their data centers, with the goal of improving both IT and facility energy efficiency.

Density is rising

The power density per rack (kW per cabinet) is a critical number in data center design, capacity planning and cooling and power provisioning. Miscalculations or unexpected changes in density demand can create technical challenges for operators and can have major financial implications. Most operators anticipate this by over-provisioning power and cooling, which can be an expensive and wasteful precaution.

In the past, experts predicted an explosion in rack density that in most cases failed to materialize. In our 2017 survey, we found the average density to be 6 kW per rack, a figure that had increased only a little over the previous decade.

In recent surveys (2018 and 2019), we have asked about highest density racks, rather than average (average can be misleading and is skewed by utilization). In 2018, a third of respondents said their highest density rack was less than 10 kW, and 28% said it was 10-19 kW. When rack densities go higher than this, direct liquid cooling and precision air cooling becomes more economical and efficient.

In 2019, the density of the highest racks is similar to a year earlier (see Figure 4), and perhaps has even decreased in some levels. The explanation for this is not clear. But more say their highest density racks are in 10-19 kW range, suggesting that the density of many fairly standard air-cooled racks, rather than those with very power-hungry specialist servers, is creeping up.

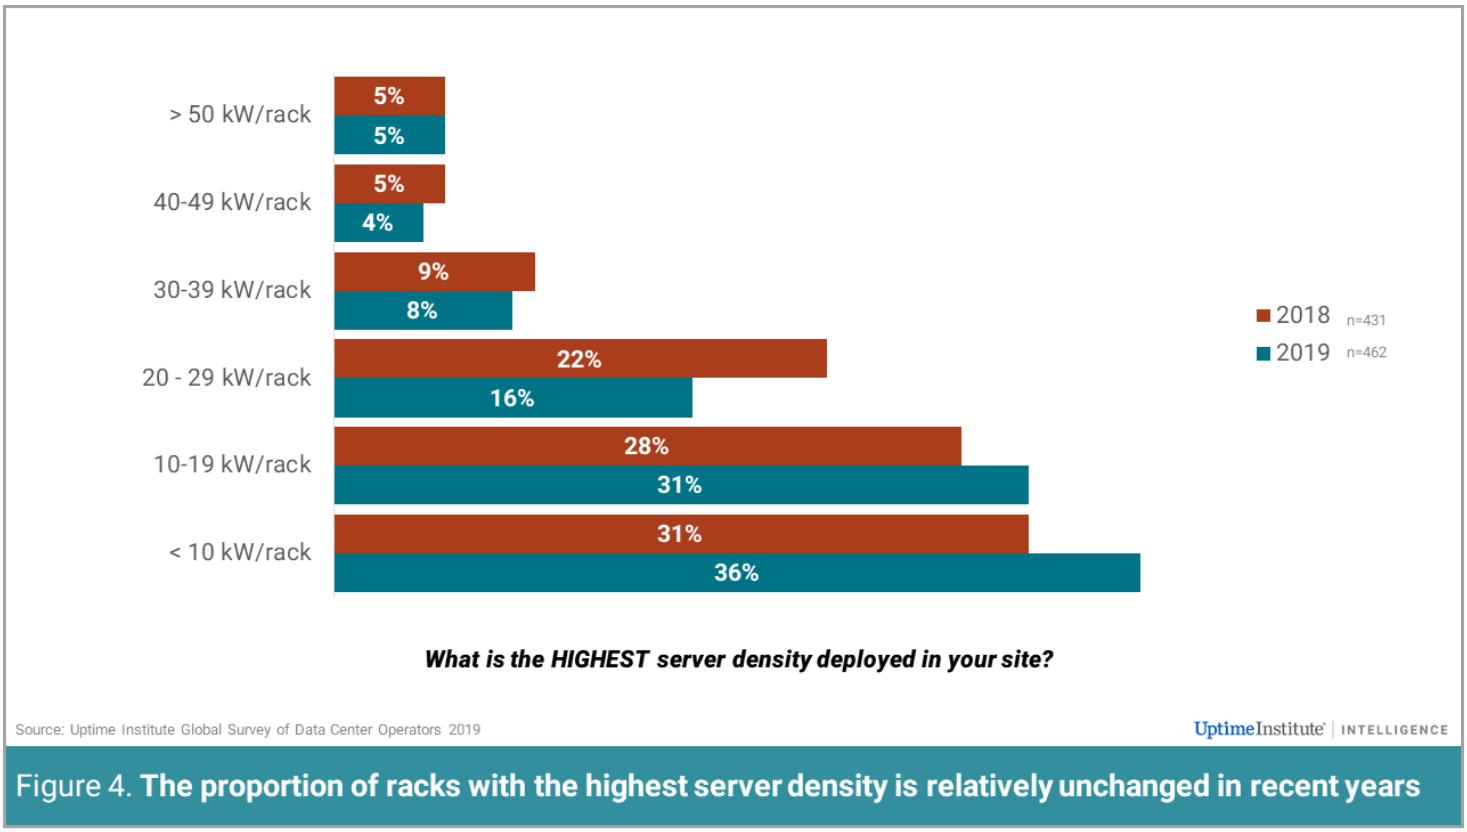

Figure 4. **The proportion of racks with the highest server density is relatively unchanged in recent years**

However, there is evidence that rack density is rising faster than in the past. As Figure 5 shows, a fifth (21%) of data center operators said density is rising rapidly, while another 48% said it is rising slowly — a total of 69% said it is rising; only six percent (6%) said it is shrinking.

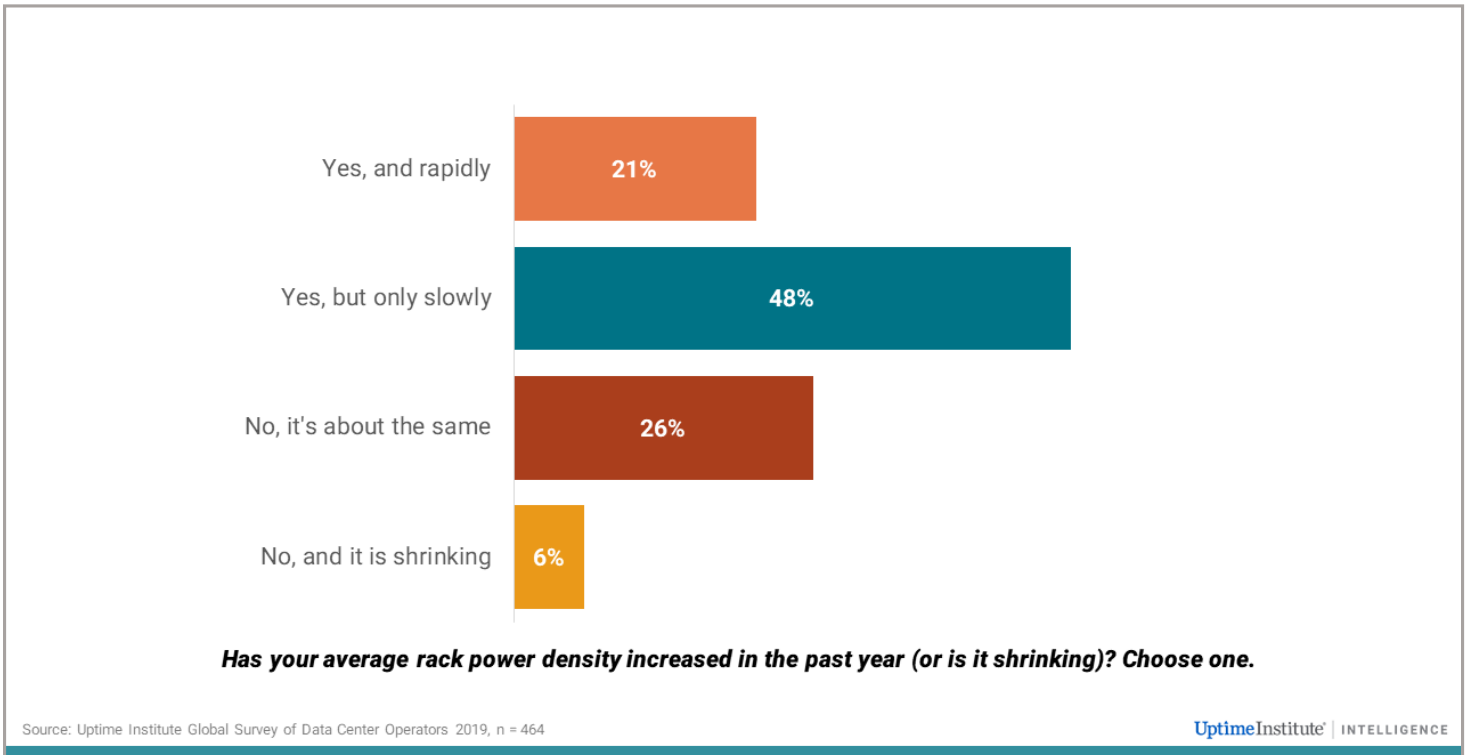

Figure 5. **Most say their average rack density is increasing**

This is in line with our expectations: Our research shows that the use of virtualization and software containers pushes IT utilization up, in turn requiring more power/cooling. Similarly, AI and analytics require a lot of compute power, sometimes enabled using co-processors — again pushing up power use, rack density and heat. With Moore's Law and processor energy use slowing down, this can be regarded as a long-term trend.

Outages haunt the sector

For the past three years, Uptime Institute Intelligence has been tracking outages (or service degradation) closely, including (as far as is possible and sensible) the costs/impact of outages. This information can be useful in budgeting and in understanding the value of investments in resiliency.

Data center outages is a complex topic, and data is easily misread. Different people in the same organization can have different perceptions of what has happened, even when they both work in digital infrastructure. A customer can have a different view altogether, while a service provider will clearly prefer not to define a service degradation as an outage if it means paying out for a service level agreement breach.

The 2019 Uptime Institute survey shows, as in our 2018 findings, that outages continue to be commonplace, and cause significant problems for operators; and once again, there was a small increase over earlier years. Just over a third (34%) of all respondents had an outage or severe IT service degradation in the past year, while half (50%) had an outage or severe IT service degradation in the past three years — very marginally up over last year (48%).

Causes of outages

What is causing these outages? Again, the data is remarkably consistent with 2018 (see Figure 6). Power was the single biggest cause, accounting for one-third, as in 2018 (respondents were asked about their most recent significant outage). Networking issues were close behind, at 31%; other IT issues (e.g., software and systems failures) fell significantly, from 28% in 2018 to 19% in 2019. Problems at third-party suppliers (e.g., colocation companies, hosting and cloud companies) jointly accounted for 24%.

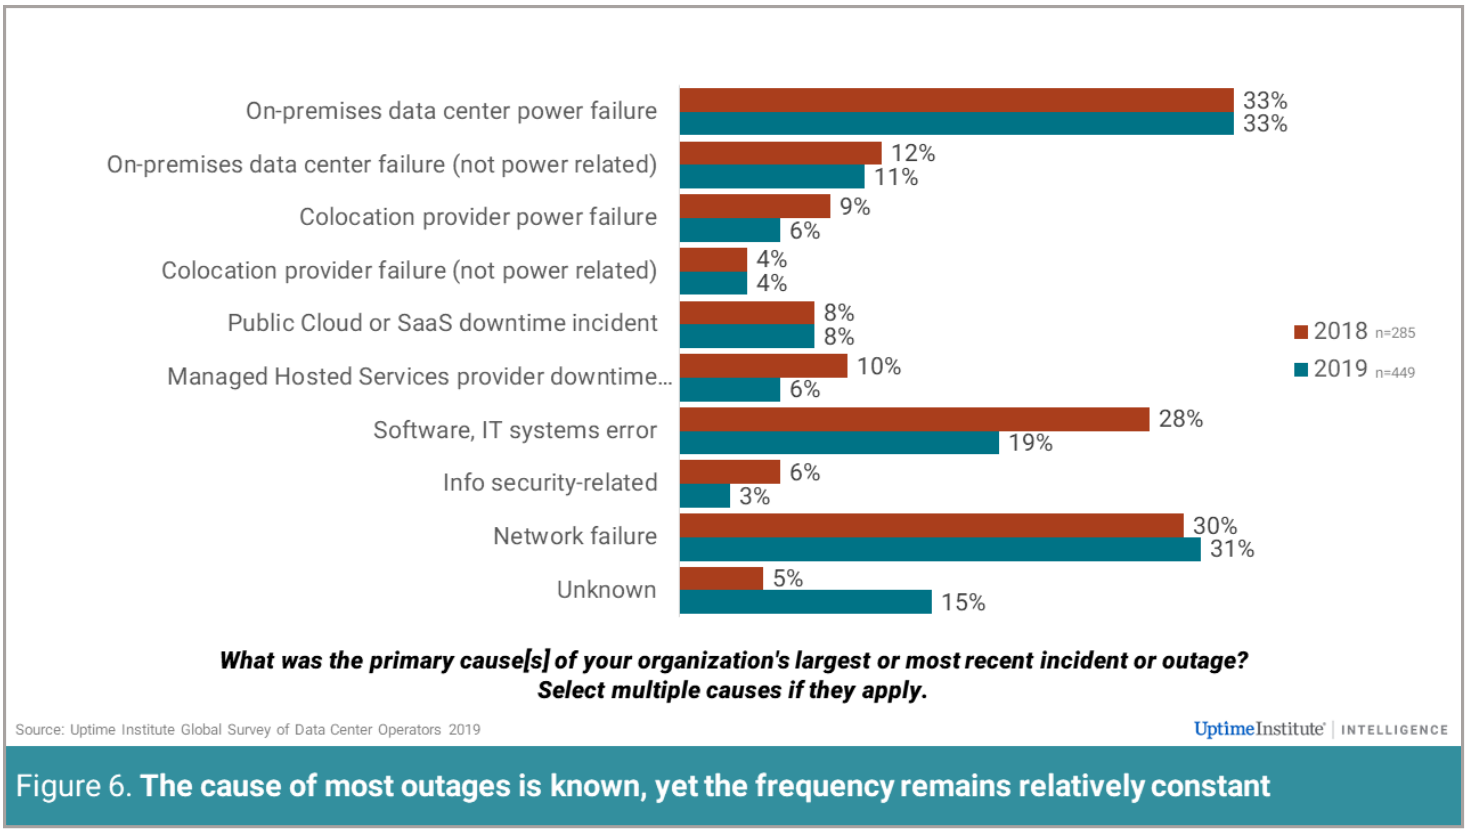

Figure 6. **The cause of most outages is known, yet the frequency remains relatively constant**

Overall, in this survey and other research, Uptime Institute has seen a clear increase in failures due to networking incidents, third-party problems and other IT concerns. Power or power-related failures continue to be a major issue.

Uptime Institute is frequently asked: What proportion of outages is the result of human error? This of course depends on interpretation — is a factory fault human error? The number is clearly high, but difficult to determine.

In this study, we asked, "Would your organization's most recent significant downtime incident have been preventable with better management/processes or configuration?" Sixty percent (60%) said "Yes." This supports the view that even the best-designed data centers can be run less than optimally and suffer more incidents — while data centers with less-resilient designs can exceed expectations if they are well-managed.

Damaging disruption

Uptime Institute also asked about the impact of the outage or disruption on the business, using the Uptime Institute Outage Severity Rating. Eight percent (8%) of outages were classed as having a "Severe" Category 5 impact, and a further 10% having a "Serious" Category 4 impact. As Figure 7 shows, both of these categories are likely to involve a damaging interruption of service, financial losses, compliance breaches and possible safety issues. In terms of numbers

of incidents, this represents 45 Category 4 outages and 36 Category 5 outages. Although the period was not defined (we asked, "most recent major incident,") the majority were in the last three years.

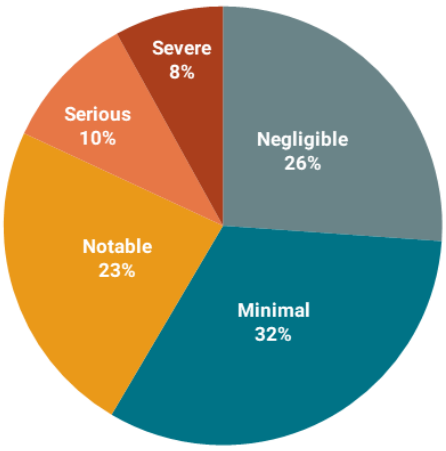

Category 1 Negligible	Recordable outage, but little or no obvious impact on services.
Category 2 Minimal	Services disrupted. Minimal effect on users/customers/reputation.
Category 3 Notable	Customer/user service disruptions, mostly of limited scope, duration or effect. Minimal or no financial effect. Some reputational or compliance impact(s).
Category 4 Serious	Disruption of service and/or operation. Ramifications include some financial losses, compliance breaches, reputation damages, possibly safety concerns. Customer losses possible.
Category 5 Severe	Major and damaging disruption of services and/or operations with ramifications including large financial losses, possible safety issues, compliance breaches, customer losses, reputational damage.

Thinking of your most recent service outage, how would you categorize it? Choose one.

Source: Uptime Institute Global Survey of Data Center Operators 2019, n=448

UptimeInstitute | INTELLIGENCE

These are significant numbers and demonstrate just how serious and damaging major outages can be, and further, that big costly outages should perhaps be viewed not as extremely rare events, but rather as "quite likely or possible" without proper attention/prevention. Vigilance, training, design review and investment are required.

This is all the more important given the financial impact of outages. While most incidents (60%) reported in our survey had a relatively negligible cost impact, just above 10% cost over $1 million (Figure 8). There were six recorded incidents where the cost of the outage exceeded $40 million. Although this sounds remarkable, Uptime Institute is aware of between five and 10 incidents with this order of losses over the past three years. As in earlier years, a high percentage of this year's respondents — 60% — admitted that their organization did not formally calculate the business cost of downtime incidents.

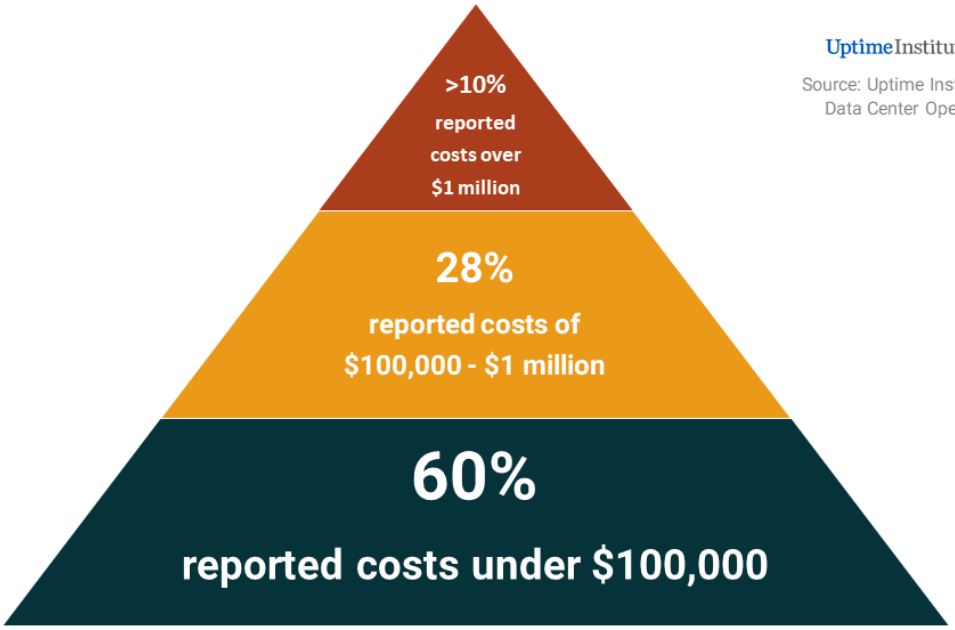

Uptime Institute | INTELLIGENCE

Source: Uptime Institute Global Survey of Data Center Operators 2019, n=420

Please estimate the total cost of your most recent downtime incident (from outage to full recovery) for your organization, including direct, opportunity and reputation.

Figure 8. **One in ten major outages cost over $1 million**

With so many services, workloads and applications spanning multiple sites, it is becoming more difficult to contain outages to a single site. In 2019, almost a fifth of outages affected multiple sites (Figure 9).

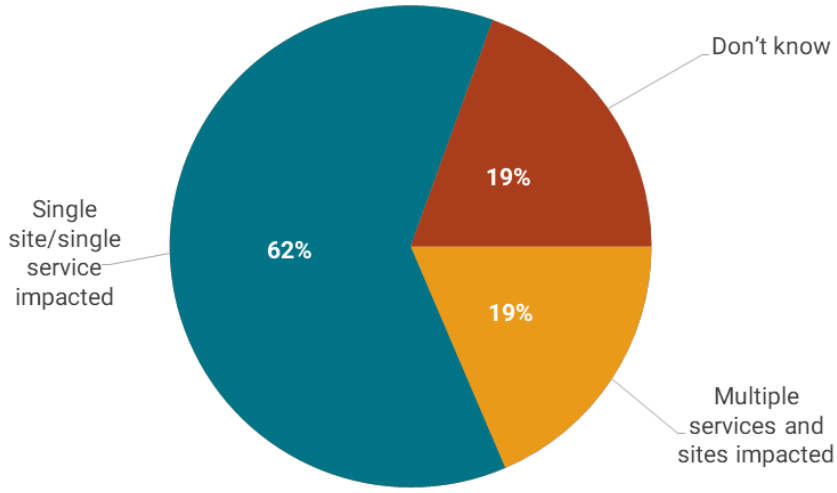

Did your most recent downtime or degradation incident impact services delivered from multiple data centers or was the downtime incident's impact confined to a single site?

Source: Uptime Institute Global Survey of Data Center Operators 2019, n=447

Uptime Institute | INTELLIGENCE

Figure 9. **Many outages affect multiple sites/services**

Distributed resiliency spreads risk

For many data center operators, the biggest single focus of their efforts to ensure uptime and resiliency has been to maintain availability at the primary data center, through rigorous attention to power, infrastructure, connectivity and on-site IT replication. If there is a failure, disaster recovery procedures kick in, which likely means some availability and even short-term data will be lost.

For most, this has been effective strategy — until now. First, even temporary unavailability causes serious business problems; second, many applications and business processes, if not underlying systems management, span several or many data centers; and third, modern architectures and advances in networking have made it possible to spread work more reliably and cost effectively across multiple sites, reducing risk. While this strategy may in some cases mean that failures or problems may not be contained to a single site, overall resiliency improved.

Over time, it is likely that more operators will adopt a distributed strategy — but doing so requires careful planning and possibly a major change in applications and databases. Our research suggests that this move is gathering momentum. Forty percent (40%) say they use availability zones for resiliency — a strategy that effectively requires two or more active data centers replicating data to each other (usually, three or more is most effective). One-quarter have not considered this strategy, and about 14% say it is too expensive or is not appropriate for their workloads (Figure 10).

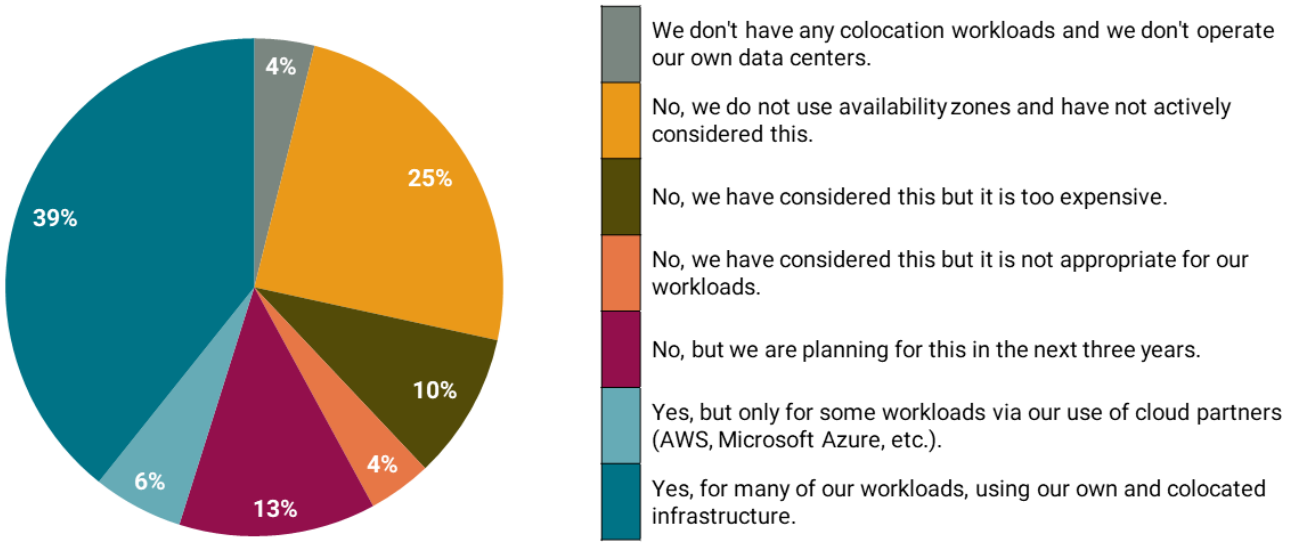

In thinking about your organization's workloads in on-premises and/or colocation data centers, are they organized into availability zones, so that data is automatically replicated and workloads diverted when there is a failure of a single site within a region? (Exclude traditional disaster recovery and back up.)

Source: Uptime Institute Global Survey of Data Center Operators 2019, n = 461

Uptime Institute | INTELLIGENCE

Figure 10. **Over half of respondents either use or plan to use availability zones**

The most committed users of distributed resiliency and availability zones are the public cloud operators. Their scale, immense resources and technical expertise enable the largest operators to create availability zones across the globe, promising very high levels of performance and availability and, in theory, a high tolerance of infrastructure failure. (Separate Uptime Institute research based on public outages, as well as our outage data above, shows that public cloud operators, however expert and well resourced, still commonly suffer degradations in service, and occasional outages, caused by infrastructure and other failures. For more information, see our **Annual outage analysis** report.)

Despite this, our research shows that a majority of organizations — around three-quarters in our survey — are not placing their mission-critical applications into a public cloud (Figure 11). There may be many reasons for this — not all workloads are suited to the cloud.

Does your organization have adequate visibility into the resiliency of public cloud operations (e.g., AWS, Azure, Google) in terms of architecture, availability record, management processes and full understanding of options? Choose one.

Source: Uptime Institute Global Survey of Data Center Operators 2019, n = 464

Uptime Institute | INTELLIGENCE

Figure 11. **More would use public cloud if providers gave better visibility into options and operations**

But a lack of visibility, transparency and accountability of public cloud services is a major issue for enterprises with mission-critical applications. A fifth of operators said they would be more likely to place workloads in a public cloud if there were more visibility into resiliency. Even among those that do use the cloud for mission-critical applications, almost half said they do not think they have adequate visibility.

Climate change denial?

Climate change is, unfortunately, going to be a major issue for all industries for as far ahead as we can see and model. In the data center industry, this manifests itself primarily in the form of extreme weather events and gradual changes that can quietly undermine data planning. Much of this was detailed in the Uptime Intelligence report **A mission-critical industry unprepared for climate change**.

In 2019, as in 2018, Uptime Institute found that about half of respondents are not currently preparing for climate change; over half of that cohort reported that their existing plans are sufficient (30% of total respondents; see Figure 12). This approach is, in some cases, likely to prove complacent — Uptime Institute suggests that frequent reviews are necessary.

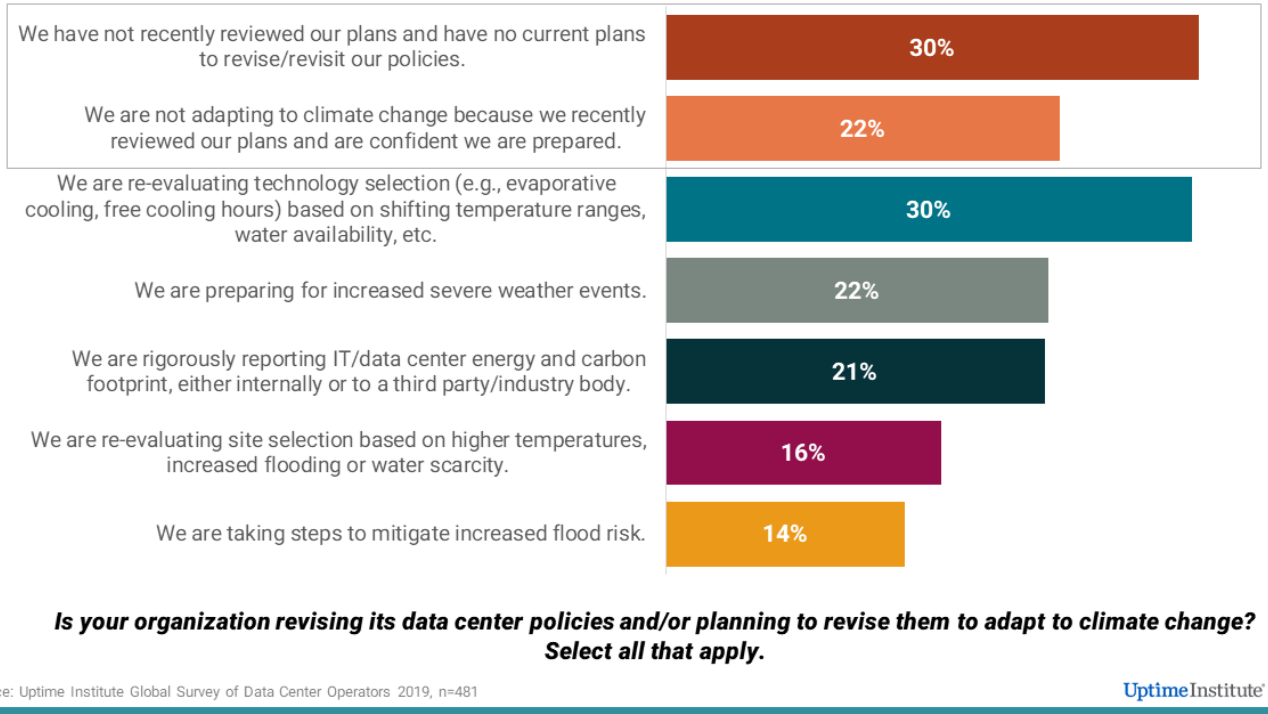

Figure 12. **Few data center operators are preparing for climate change**

The effects of climate change may be incremental, with air temperature, flood plains or maximum storm surges increasing inexorably until some threshold is exceeded and disaster follows. The threshold might be set by the height of a levee or dike, or simply by the maximum capacity of a data center cooling system. Water shortages are also a major concern in many geographies, and in some cases prevent the use of certain cooling technologies.

Data center operations are particularly vulnerable to the localized effects of climate change, as floods can interrupt utility services and connections; water can block roads and in doing so, prevent fuel deliveries and strand employees. For this reason, Uptime Institute advises data center owners/operators to review their resiliency plans regularly — at least annually.

The 2019 survey found only small increases in the number of respondents taking steps to mitigate flood risk (14%), re-evaluate

site selection (16%) and prepare for increased severe weather events (22%). The survey found decreases in the number of respondents re-evaluating cooling technology (30%) and reporting energy and carbon (21%). Necessity, and government action, may ultimately lead to a rethink.

Uptime Institute also notes that energy and carbon reporting can drive improved efficiencies that result in energy savings and reduced stranded capacity. Committing to environmental sustainability and efficiency goals through annual reports and the Global Reporting Initiative or the Carbon Disclosure Project can help create company-wide accountability and drive change.

The people problem

In the past two years, the staffing problem that is affecting most of the data center sector has been elevated to the level of a crisis; many operators are struggling to effectively staff their data centers, with all the risks this entails. It is not one problem, but many — and most but not all of these can be attributed to the rapid growth of the industry. The entry of the hyperscales into the market has had a dramatic effect in some areas.

The problem is shown in Figure 13. Sixty-one percent (61%) of respondents said they had difficulty retaining or recruiting staff — up from 55% a year earlier. Only a small number (12%) are expecting staff cuts. Given the slowdown in enterprise data centers and growth of automation, this is a very low percentage. The extent of the staffing problem has been widely discussed in the industry, with many operators now recruiting outside the industry and retraining staff much more actively.

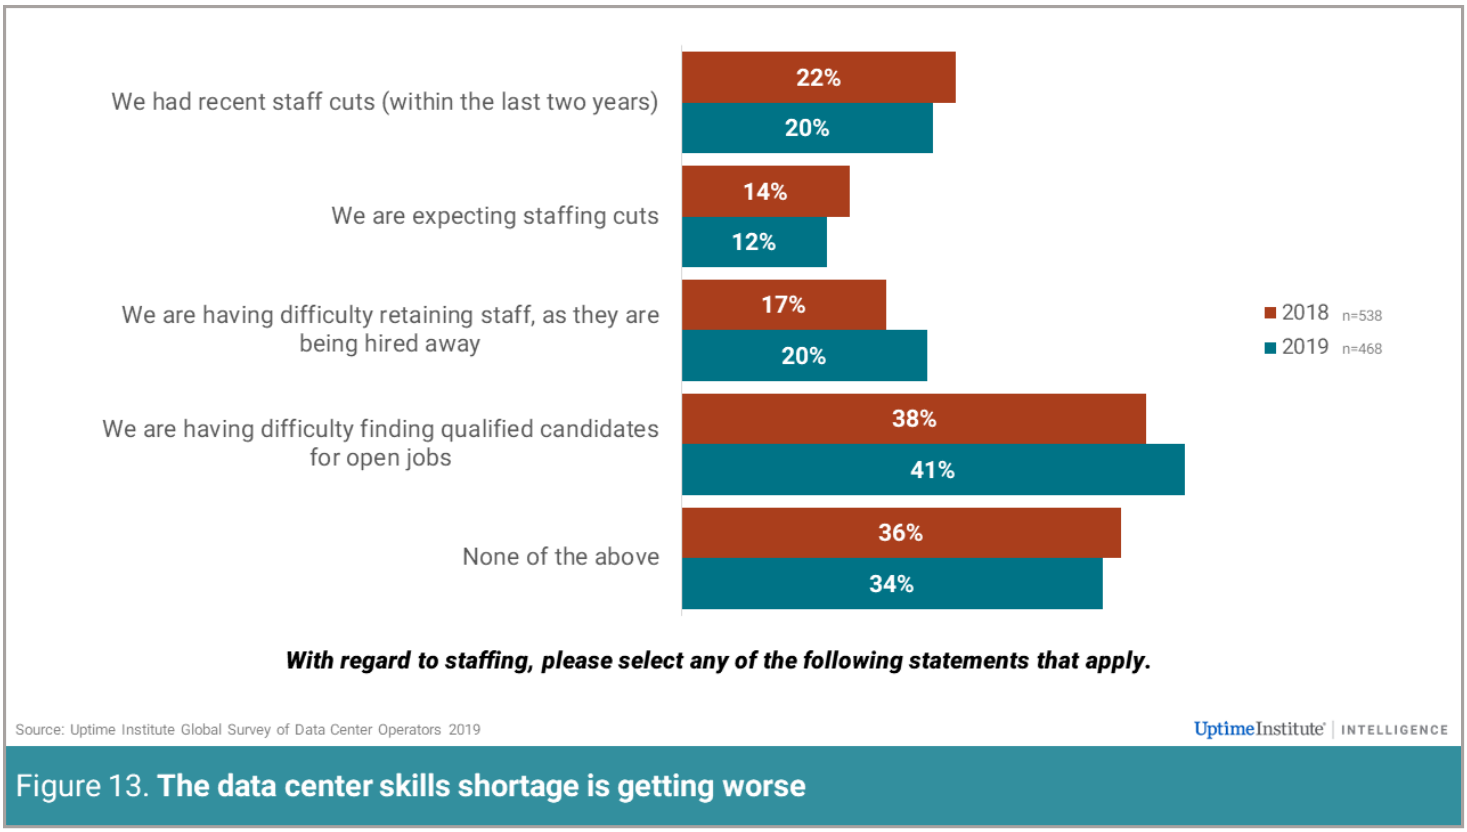

Figure 13. **The data center skills shortage is getting worse**

One known problem is aging: many management and operational staff are over 50 years old and, in many cases, over 60. (Although this was reflected in the survey responses, we have discounted this data because younger or more junior staff may not be qualified to respond to surveys.) Many experienced staff are retiring, leaving a shortage of qualified management, engineers and operators.

The lack of women working in all roles in the data center industry has been widely highlighted. Only five percent (5%) of respondents said women represented 50% or more of staff — while one-quarter had no women at all among their build, design or operations staff (Figure 14). Whatever the causes — and they are deep, wide and long-standing — there is a clearly a general, if not urgent, desire across the industry to redress the gender imbalance. For some, this is primarily a commercial issue: A more inclusive approach would bring more people, and more skills, into the sector.

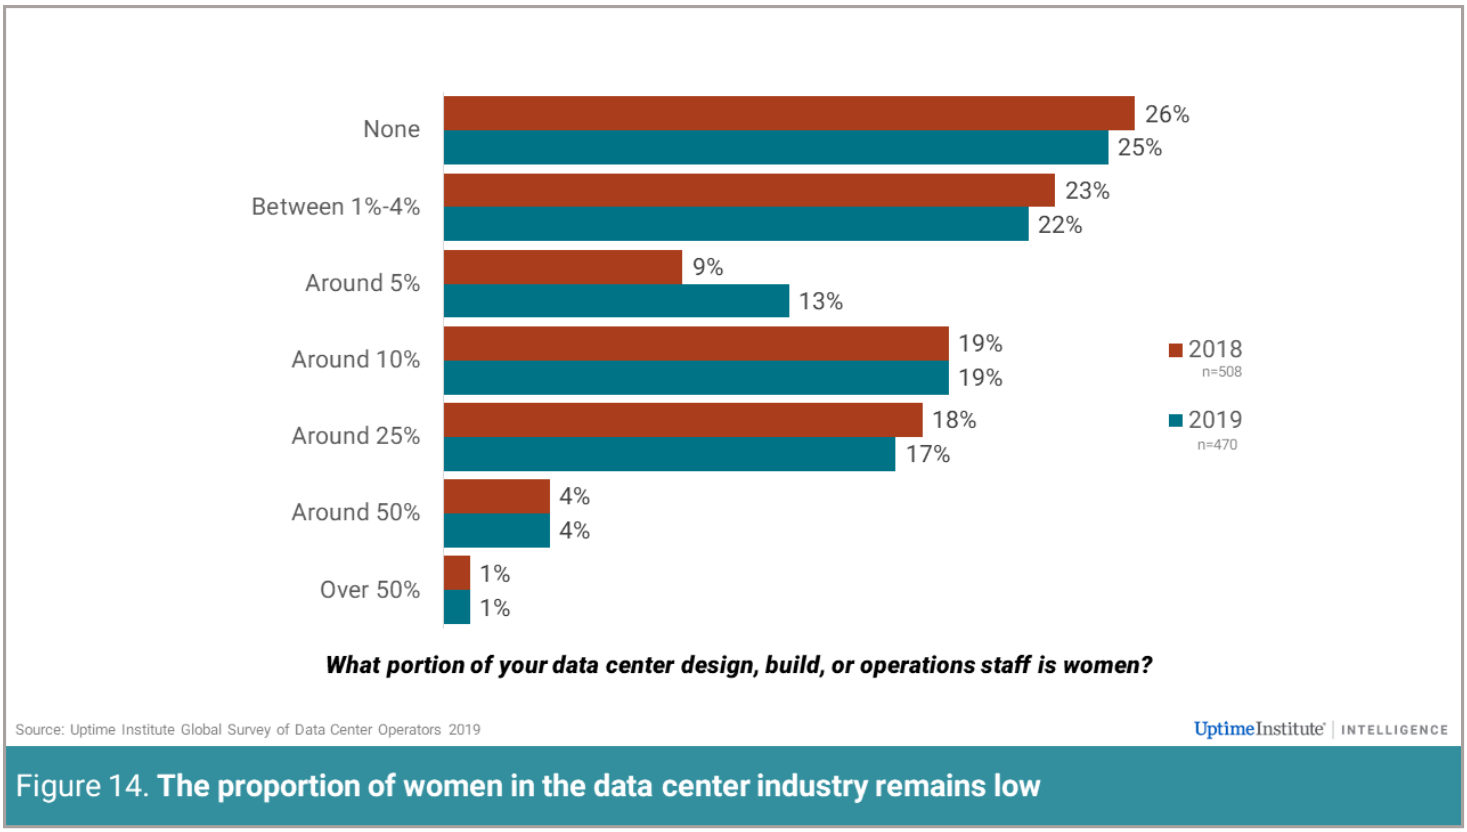

Figure 14. **The proportion of women in the data center industry remains low**

While most (primarily men) in the industry (73%) do think the industry would benefit from employing more women, most do not appear to view this is a serious problem or one that needs urgent attention, at least as far as their organization is concerned. As shown in Figure 15, just over half (53%) of those surveyed think it is easy for women to pursue a career in data centers in general, against 47% who thought it difficult. But 85% thought it was easy for women to pursue a career in data centers in their own department (just 15% thought it would be difficult). Fifty-five percent (55%) of respondents think the lack of women presents no threat to the industry.

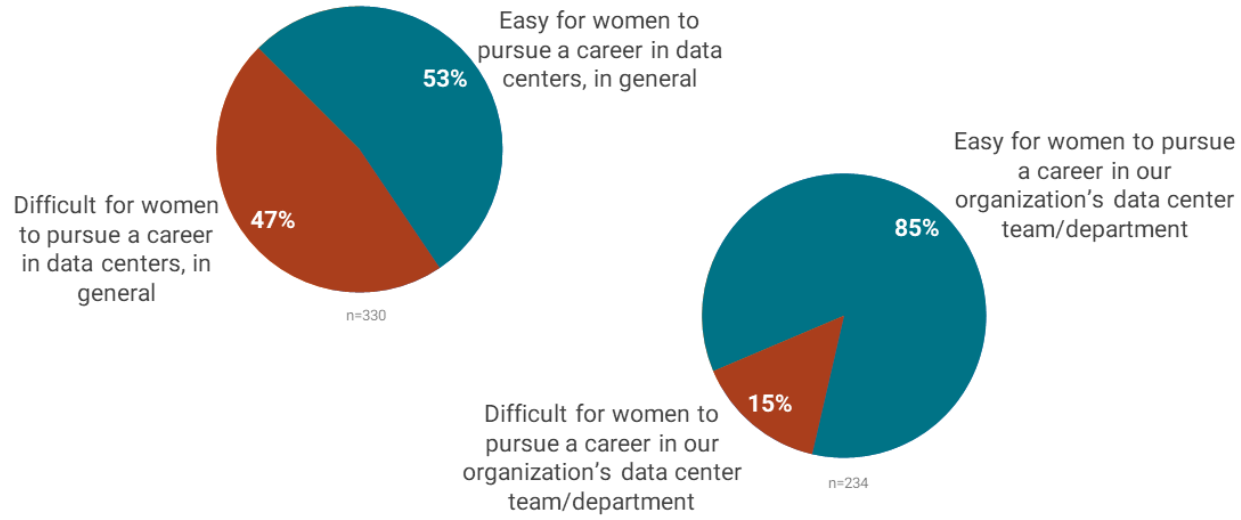

Figure 15. **Women can do well here, most operators say**

Given the attention around a lack of gender diversity in the industry, it is perhaps surprising to see that slightly fewer organizations have initiatives in place to recruit more women (see 2017 vs. 2018 comparison in Figure 16). This may be an anomaly — it is certainly not because the initiatives were successful — or perhaps this is because some initiatives simply failed to yield any results.

More see the lack of women in the industry as a problem

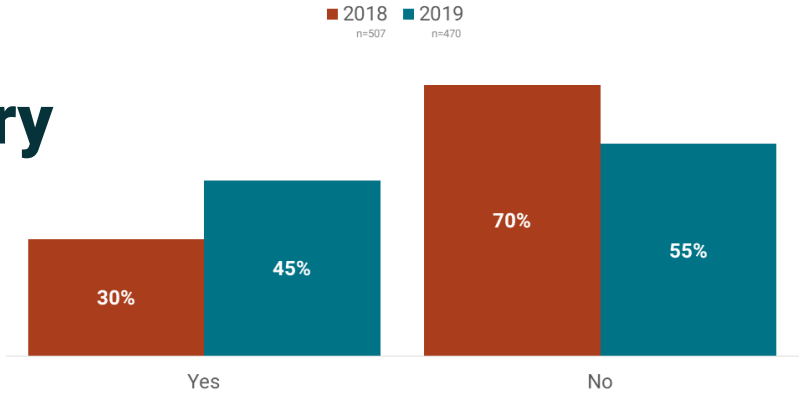

Do you believe a lack of women in the data center profession is a threat to the industry at large (lack of pipeline for hiring, technical stagnation, or other)?

Source: Uptime Institute Global Survey of Data Center Operators 2019

UptimeInstitute | INTELLIGENCE

and most feel the industry would benefit if more women were in the workforce

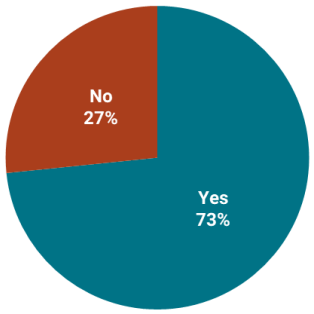

Do you think the data center industry would benefit from having a higher proportion of women employees than it does today?

Source: Uptime Institute Global Survey of Data Center Operators 2019, n = 466

UptimeInstitute | INTELLIGENCE

yet fewer are doing anything to increase the proportion of women in their data center

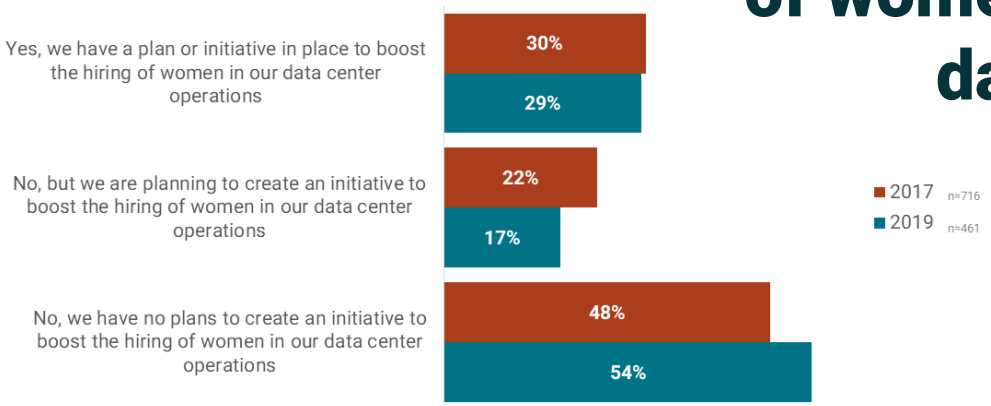

Does your organization have a plan or initiative in place to boost the hiring or participation of women in your data center operations?

Source: Uptime Institute Global Survey of Data Center Operators 2019

UptimeInstitute | INTELLIGENCE

If not women, then perhaps robots and AI will help manage and operate the data center? Uptime's own research does suggest that certain tasks can be effectively automated and that smaller data centers can, of course, operate without any staff at all. But in larger data centers, automation clearly has its limits and is unproven in many areas. Most of those operating data centers today agree; nearly two-thirds think automation will not reduce staffing requirements in the data center in the next five years, if at all. After that, however, most believe automation will reduce staff (Figure 17).

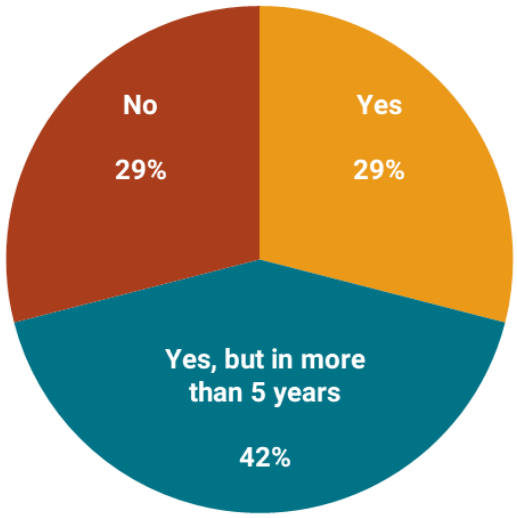

Do you believe artificial intelligence (AI) will reduce your data center operations staffing levels in the next 5 years?

Source: Uptime Institute Global Survey of Data Center Operators 2019, n = 466

Uptime Institute | INTELLIGENCE

Figure 17. **Most say that AI will not reduce staffing levels within the next five years**

OCP and Open 19: Hidden in open view?

In 2011, a group of data center designers and pioneers led by Facebook introduced the Open Compute Project. The OCP architecture involves some major data center design changes, especially in rack design and power distribution, resulting in dramatically reduced costs for large-scale operations. In 2016, another group led by LinkedIn operators developed a less revolutionary and less comprehensive (but more accessible) alternative, Open19, which would also reduce data center costs. Both groups hoped to have a major impact on the industry.

So far, the impact has been limited. Uptime Institute's Data Center Survey finds that awareness of the Open Compute Project has stalled, with 39% of respondents to the 2019 survey saying they were not aware of it (see Figure 18). This is the same result as the 2017 survey. In addition, the percentage of respondents who reported evaluating or deploying OCP remained at about 20%.

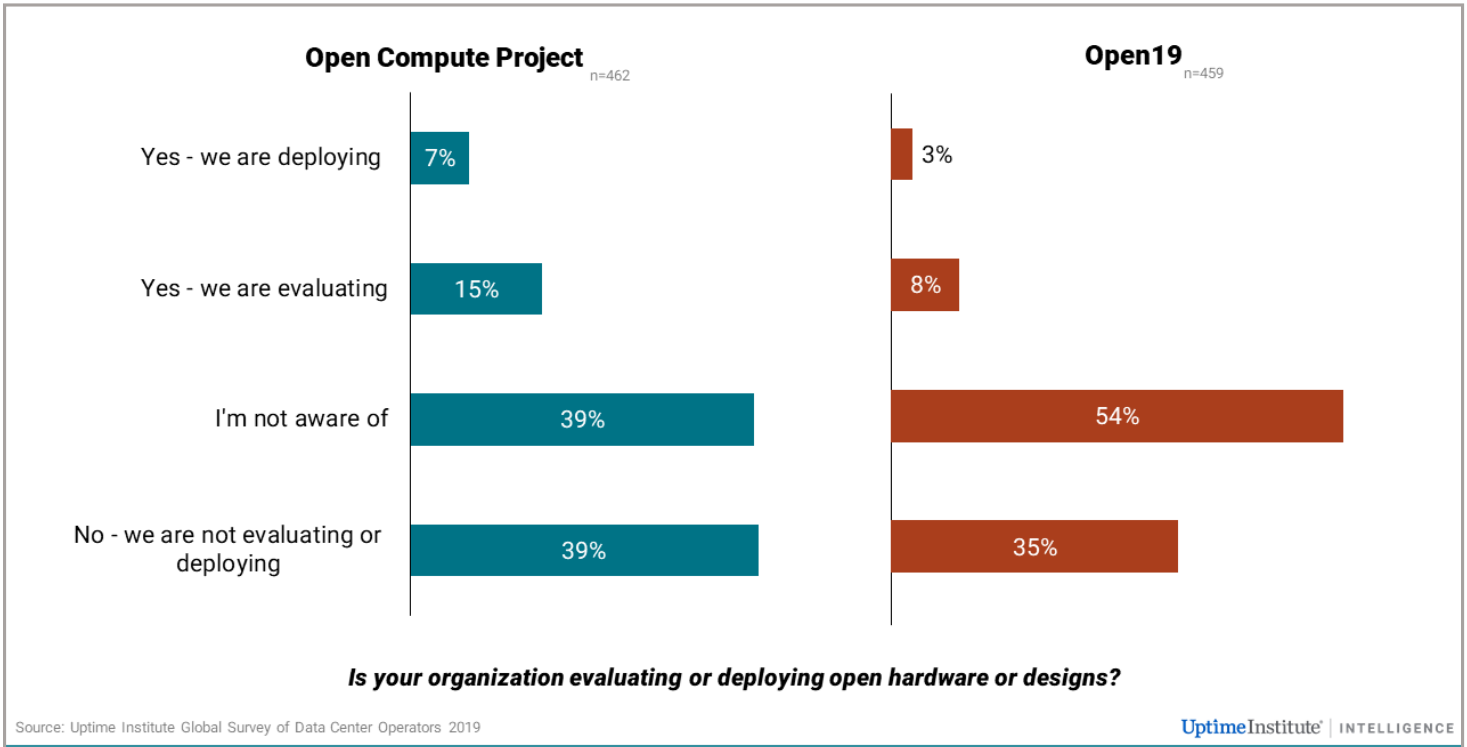

Figure 18. **Few data center operators are adopting open architecture standards**

Awareness of Open19 is even lower, with 89% of respondents reporting that they either were not aware of it (54%) or were aware but not evaluating or deploying it (35%). Although this is not surprising (by comparison, OCP has more and bigger commercial backers), awareness of Open19 is still far lower than OCP awareness (39%) was back in 2017 — despite the fact that, in the enterprise sector, Open19 is probably more appropriate than OCP for many operators.

Even so, regardless of low awareness levels, Uptime Institute finds that both architectures are finding some acceptance in the market, with seven percent (7%) of respondents having deployed OCP and three percent (3%) now having deployed Open19. Although these numbers seem small, they are growing, especially as these standards are adopted in the larger data centers. In addition, conditions may be right for an increase in OCP and Open19 adoption. Senior executives are far more likely (22%) than IT management (11%) or critical facilities management (11%) to report evaluating OCP hardware or designs. This senior group was also more aware of OCP (75%), with all other groups having 60% or less awareness of the program.

The Open Compute Project Foundation is a rapidly growing, global community whose mission is to design, use, and enable mainstream delivery of the most efficient infrastructure designs for scalable computing. https://www.opencompute.org/about

The Open19 Foundation aims to create project-based open hardware and software solutions for the data center industry. The main goal of Open19 is to create a set of community-driven standards and designs that are more customizable, flexible and economical. https://www.open19.org/foundation/

The same pattern held true for Open19, with senior executives (56%) showing far more awareness than IT management (41%) and critical facilities management (40%) staff. Design engineers (seven percent [7%]) were more likely to say they were deploying Open19 designs than the other three groups combined.

Li-ion batteries

Our survey showed that Li-ion batteries are becoming well-established in the data center, although it is still early in the adoption cycle. As shown in Figure 19, almost half of those surveyed said that they have distributed Li-ion batteries installed or contracted. Although this is higher than many would have expected, Uptime Institute is seeing many operators using small amounts of Li-ion tactically to increase power or resiliency in certain racks (and in major projects, of course). The number saying they have deployed Li-ion (around 20%) has not varied much from 2018 — but the number considering using distributed batteries jumped to one in five. Meanwhile, over half (56%) are either using Li-ion batteries centrally or are considering doing so.

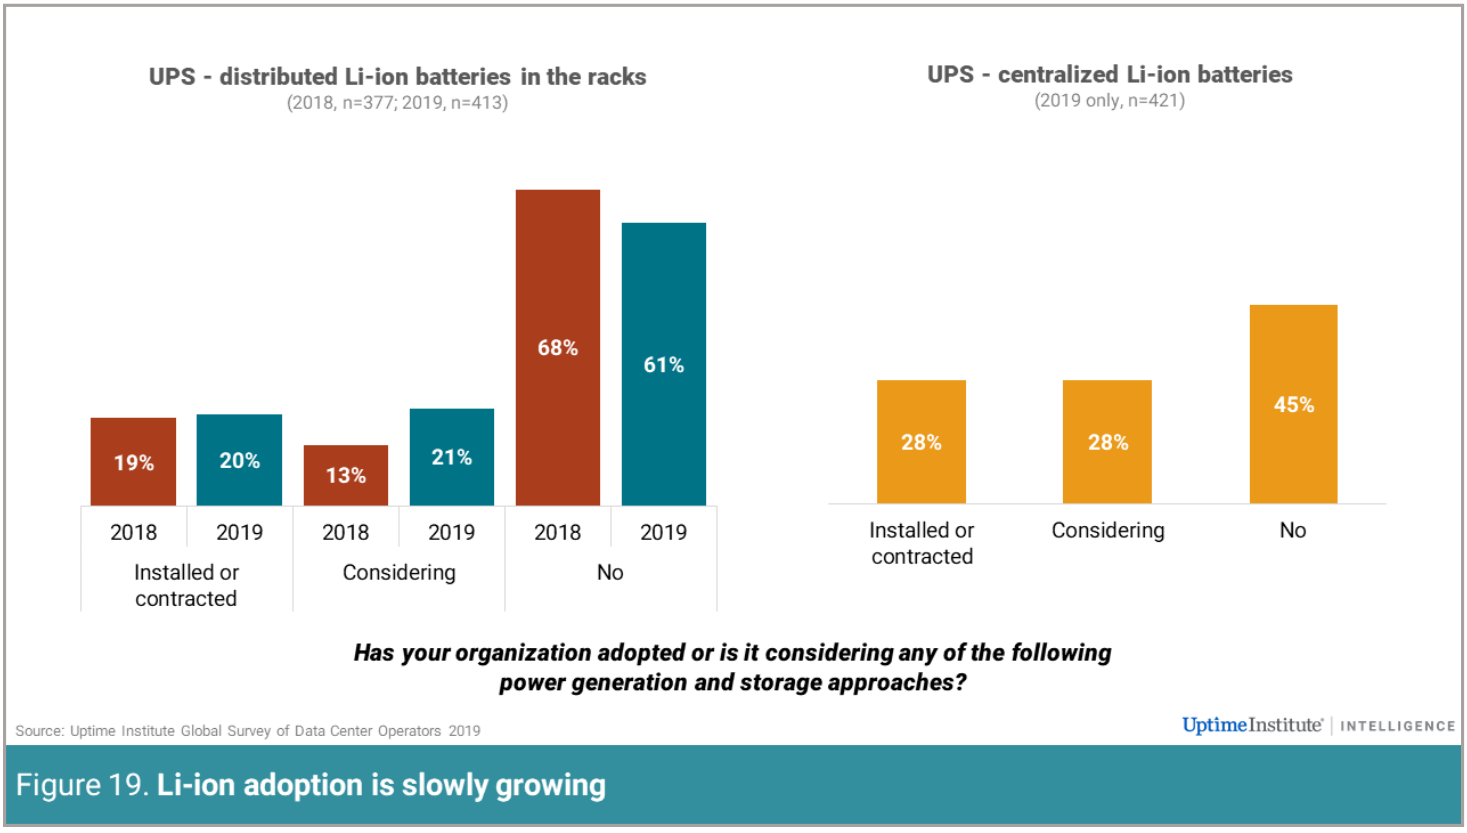

Figure 19. **Li-ion adoption is slowly growing**

Organizations considering OCP or Open19 are far more likely to have installed or to be considering Li-ion batteries than other organizations. The same holds true for fuel cells. This is unsurprising, as OCP is designed to support distributed storage, and because larger operators considering OCP are more likely to be willing to adopt newer technologies.

Uptime Institute considers this point significant because Li-ion batteries and fuel cells, as well as renewable energy, microgrids and even engine-generator sets, play an important role as energy storage devices in emerging "smart energy" configurations. These energy storage and generation technologies (including fuel cells), along with automated management systems and smart meters that monitor, report and optimize energy use and demand, can form the basis for transactive utility arrangements and dynamic energy optimization, which can improve reliability, reduce costs and increase the ability to maintain or improve resiliency without high levels of redundant equipment (see Uptime Intelligence's recent report **Smart energy in the data center**.)

Conclusions and recommendations

Based on our survey findings, Uptime Institute Intelligence recommends that managers/operators consider the following:

• Public cloud and other third-party data center services are not immune to outages or service degradations. The vetting process for third-party services should go beyond a simple review of the service level agreement and include visibility, resiliency, accountability and true costs.

• The biggest infrastructure efficiency gains happened five to six years ago, according to our analysis of the industry average PUE. Further improvements will require significant investment and effort, with increasingly diminishing returns. While managers and operators should remain vigilant and seek to maintain high facility efficiency levels, higher gains may be found by focusing on IT efficiency.

• Outages continue to be commonplace and costly, and incidents increasingly span multiple data centers. Uptime Institute recommends comprehensive and ongoing resiliency reviews of digital infrastructure that include company-owned data centers and third-party service providers, and that also take into account the effects of climate change (at a regional level).

• Data center skill shortages will intensify: In this aging and overwhelmingly male sector, most operators are struggling with staffing issues. While most do not believe a lack of diversity in their ranks is an issue to be concerned about, Uptime believes actively recruiting women and other underrepresented populations into data centers will help alleviate the staffing crisis.

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. Respondents are separated into two groups: owner/operators of data centers; and suppliers, designers and advisors. This report focuses on the findings from the owner/operator survey — people responsible for managing infrastructure at the world's largest IT organizations. Job titles include senior executive, IT management, critical facilities management and design engineer.

As Figure A1 shows, the participants represent a wide range of industries in multiple countries, with just over half coming from North America and Europe. Almost half of the respondents work for professional IT/data center service providers (i.e., staff with operational or executive responsibilities for a third-party data center, such as those offering colocation, wholesale, software or cloud computing services). This is closely aligned with the percentage of data center build-out (by MW) attributed to this sector.

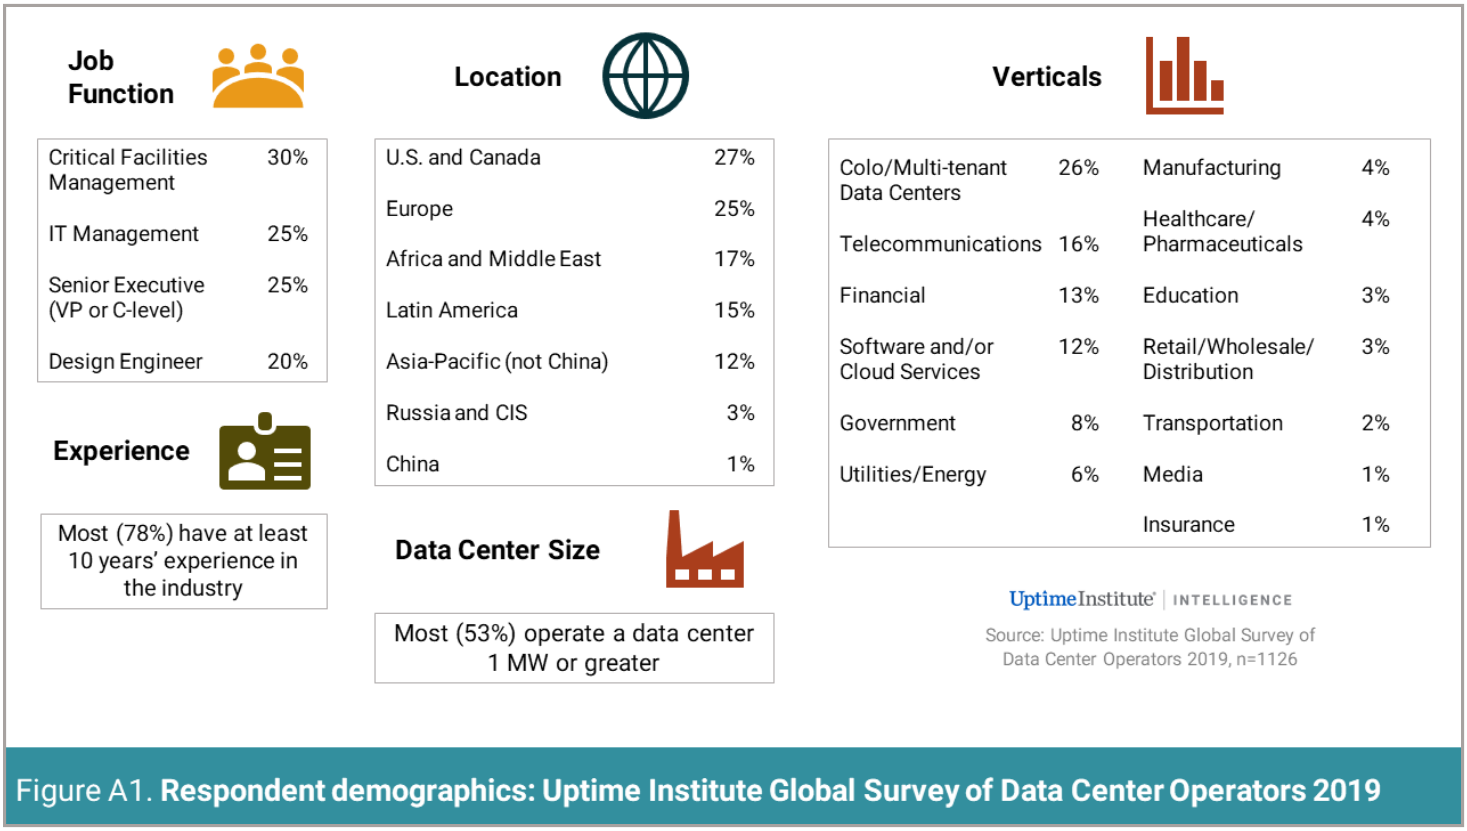

Figure A1. **Respondent demographics: Uptime Institute Global Survey of Data Center Operators 2019**

Approximately 1,100 end users registered for the survey — this means they answered at least one question. Because respondents were not required to answer all questions, the number of respondents for individual questions ('n') varies widely.

Previous survey findings are available on the Uptime Institute Network member website, Inside Track. The results of the designer/supplier/advisor survey are available in report 29, **Uptime Institute data center supply-side survey 2019**. We will also publish some particular slices ("cuts") of findings (e.g., by region, by industry) in the form of **Uptime Institute Intelligence Notes**.

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

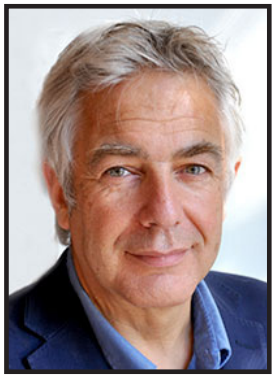

Andy Lawrence is Uptime Institute's Executive Director of Research. Mr. Lawrence has built his career focusing on innovative new solutions, emerging technologies, and opportunities found at the intersection of IT and infrastructure. Contact: alawrence@uptimeinstitute.com

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

Kevin Heslin is Uptime Institute's Chief Editor. He has over two decades of experience serving as an editor at New York Construction News, Sutton Publishing, IESNA and BNP Media, where he founded Mission Critical. In addition to editorial responsibilities, he serves as community manager for Uptime Institute Network's member community, Inside Track. Contact: kheslin@uptimeinstitute.com

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

