Uptime Institute Global Data Center Survey 2022

Resiliency remains critical in a volatile world

Uptime Institute® | INTELLIGENCE

The Uptime Institute Global Data Center Survey, now in its 12th release, is the most comprehensive and longest running of its kind. The findings in this report reveal the practices and experiences of data center owners and operators in the areas of performance, resiliency, efficiency and sustainability, staffing and innovative technologies.

Authors
Jacqueline Davis, Research Analyst
Daniel Bizo, Research Director
Andy Lawrence, Executive Director of Research
Owen Rogers, Research Director for Cloud Computing
Max Smolaks, Research Analyst

Contributing authors
Lenny Simon, Senior Research Associate
Douglas Donnellan, Research Associate

20-30 MINUTES TO READ

Synopsis

The Uptime Institute Global Data Center Survey 2022 reveals an industry that is growing, dynamic and increasingly resilient. Spending on data centers and related services is strong, despite persistent staffing shortages, supply chain delays and other obstacles. Creating a more environmentally sustainable footprint is a major and growing focus as operators brace for heightened scrutiny, new regulations and reporting requirements.

- Sustainability pressures are growing. Most operators expect carbon emissions reporting requirements soon — yet many are unprepared.
- Outages are becoming more expensive and are still too frequent, but the proportion of operators reporting a significant outage is down slightly on last year.
- Many cloud applications are vulnerable to outages, despite growing confidence in cloud's suitability for mission-critical workloads.
- Future efficiency gains will need to focus on IT power, requiring additional metrics to supplement power usage effectiveness (PUE).
- Data center equipment vendors are caught between high demand and lingering supply chain problems. Most vendors deliver with delays.
- Attracting and retaining qualified staff remains highly problematic for many operators.

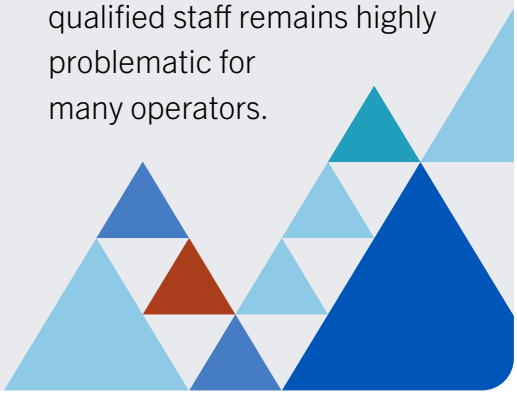

Contents

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 or contact research@uptimeinstitute.com.

Illustrations

Introduction

The 12th annual Uptime Institute Global Data Center Survey presents a snapshot of the practices, trends and challenges shaping the mission-critical digital infrastructure industry. The survey is the most comprehensive and longest running of its kind.

The survey examines the state of the industry in terms of performance, resiliency, efficiency and sustainability, staffing and innovative technologies.

This report focuses primarily on the practices and experiences of IT and data center managers globally. For a more rounded view, Uptime has also canvassed equipment and engineering suppliers in certain areas. The survey was conducted online during the first half of 2022. For more details, including demographics, see **Appendix A**.

Industry benchmarks

Power usage effectiveness (PUE) and rack power density are two major metrics the data center industry uses to track progress toward greater facility efficiency. The simplicity and wide adoption of these metrics are their strengths, but their scope does not extend to the IT infrastructure. Typically, IT accounts for the largest proportion of a data center's energy use and other environmental impacts, yet IT operational efficiency is largely not quantified.

PUE is in stasis — for now

Since its introduction in 2007, PUE has become the de facto standard metric for data center efficiency. Despite long-acknowledged limitations to its scope (such as IT efficiency and trade-offs with resiliency and water), PUE helps operators track the energy efficiency of their facilities over time. When backed up by a large sample size, PUE is also the most useful proxy to illustrate trending in facility efficiencies of the global data center sector. Uptime has been tracking PUE for 15 years.

Survey respondents' average annual PUE in 2022 was 1.55 (see **Figure 1**), which means that, in aggregate, their data centers expend 55% as much energy on cooling, power distribution and ancillary facility functions as on IT. This is consistent with the average PUE trend observed by Uptime in recent years — improvement slowed markedly in 2014, with only marginal gains since then.

Figure 1 **PUE progress has stalled**

What is the average annual PUE for your largest data center? (n=669)

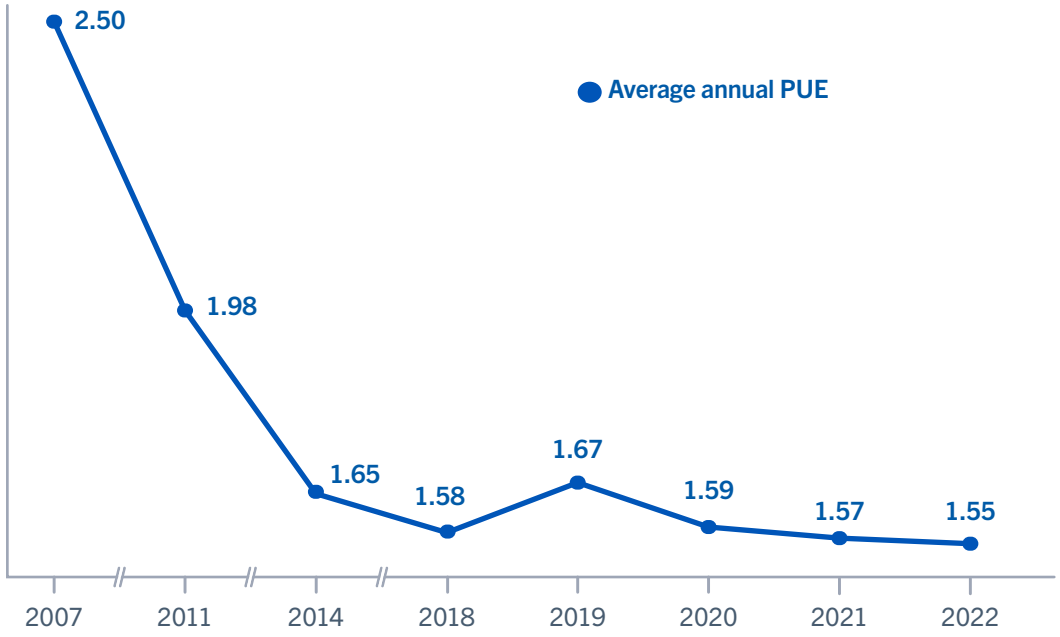

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2007-2022

Uptime Institute | INTELLIGENCE

If hotter chips become common within a few years, the industry average PUE may rise before it falls

The rapid decrease in PUE from when Uptime started tracking it in 2007 up until 2014 largely reflects the broad adoption of inexpensive efficiency measures, such as hot / cold air containment, optimized cooling control and increased air supply temperatures. Often, efficiency improvements beyond these were not economically or technically feasible for many older data centers, which contributed to the industry average PUE stabilizing.

This is not to suggest the sector has reached an efficiency limit. New data center builds routinely outperform the average, achieving PUEs of 1.3 and below using facility designs and more advanced equipment that are optimized for lower energy use.

In contrast to other industries with intensive cooling needs, such as some manufacturing and power generation, air cooling still dominates in data centers — even in new facilities. Upcoming server processors with high thermal power may pose a new challenge to air cooling. In a reversal of the historical trend toward more relaxed temperature ranges, a growing number of operators will need to restrict supply air temperatures more tightly, with a resulting penalty in cooling efficiency. Although PUE has been trending downward over the long term, if these hotter chips become common within a few years, the industry average PUE may rise before it falls.

Performance requirements and expectations around efficiency will likely push a growing number of operators (and their IT tenants) toward direct liquid cooling (DLC). Greater adoption of DLC could contribute to greater efficiency gains through the 2020s and beyond, both in new builds and retrofits. Commercial DLC offerings are likely to mature and standardize, especially as more large operators deploy at scale — as expected (see our report *The coming era of direct liquid cooling: it's when not if*).

For a more comprehensive view of infrastructure energy efficiency, technical organizations will need additional metrics that address IT energy performance. This is a considerably more complex task than measuring PUE. A major challenge is establishing a baseline because there is no clear method to quantify the amount of work IT performs. It is also difficult to isolate the energy used for performing work from parasitic losses (such as fan power, IT system internal power distribution losses and silicon power losses) within IT's power consumption. The difficulty of adopting composite IT efficiency metrics makes it likely that most operators will adopt a collection of relatively simpler indicators of efficiency to promote IT efficiency gains, such as server utilization, hardware age or the application of power management tools — all outside the scope of PUE. Relying too heavily on PUE as the industry's key efficiency metric may reduce operators' motivation to pursue IT efficiency improvements.

Rise of rack densities accelerates

Uptime has been closely tracking data center power densities as a key design and operational metric for many years. Increases in modal (most common) rack power have been slow for several years, with no major shift from the typical 4 kilowatt (kW) to 6 kW cabinets, even at flagship data centers. Emerging compute-hungry workloads, such as the application of artificial intelligence (AI) in analytics and engineering simulations, have up to now failed to significantly push up modal rack power. Uptime does not track the innumerable racks spread across all small computer rooms; if we did, modal rack power would likely be considerably lower still.

Population averages tend to reveal trends slowly, hiding underlying dynamics. But in a marked uptick to previous readings, more than a third of data center operators surveyed say their densities have rapidly increased in the past three years, with only 4% decreasing. There are some, albeit small, differences between enterprises, colocation operators and IT services providers, such as cloud, hosting and software as a service (SaaS), as shown in **Figure 2**.

Figure 2 **Rack power density rising fast across major data center segments**

Over the past three years, how has the most common (modal average) rack power density deployed in your organization / colocation data centers changed? (n=611)

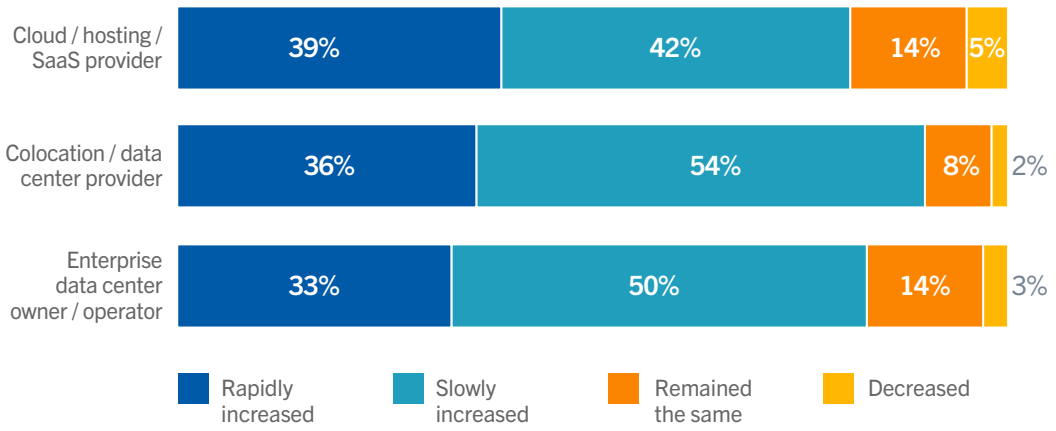

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2022 Uptime Institute | INTELLIGENCE

The larger a data center, the more likely it has seen a fast rise in typical rack power: about 40% of organizations that operate facilities with capacities above 5 megawatts (MW) say their densities are increasing rapidly — a larger proportion than the 30% with smaller facilities. Larger data centers typically run higher-powered racks to begin with, which suggests an increasingly widening density spread in the future.

The same pattern holds when Uptime asks operators about their most power-hungry IT racks. Excluding dedicated high-performance computing facilities, only one-quarter of respondents reported that they have any racks above 20 kW, a slight decrease from 2021 and few have racks above 30 kW (see **Figure 3**). This may be attributed to a larger, more diverse survey sample in 2022 than in previous years.

Figure 3 **Few data centers have racks above 30 kW**

What is the highest server rack density deployed in your site?

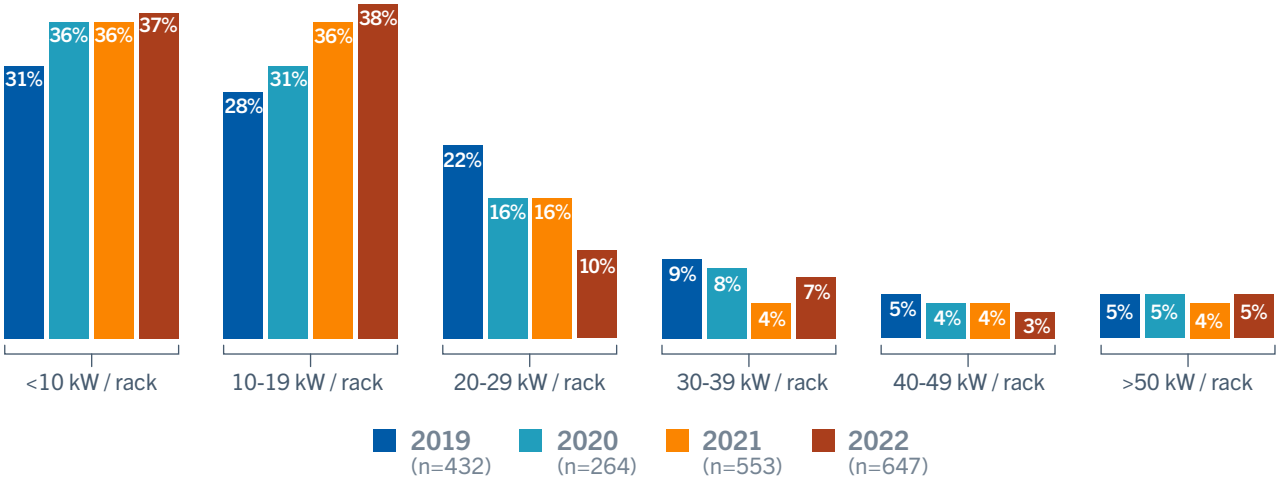

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2018-2022 Uptime Institute | INTELLIGENCE

However, when it comes to the largest facilities at 10 MW and above, nearly half have cabinets above 20 kW, and almost one in five run some racks over 40 kW power. There is also a small but growing number of data centers that house some cabinets above 70 kW power, concentrated mostly in the largest facilities.

Our survey data suggests that these larger facilities also tend to have the lowest PUEs, which means they will likely be running increasingly denser, more efficient infrastructures than smaller operations. Larger organizations generally have more complex business and IT needs, the benefit of economies of scale and larger budgets to install advanced infrastructure and power hungry, denser computing.

Operators refresh servers less often

Regular replacement of older IT hardware usually improves data center efficiency, with more recent equipment typically capable of delivering more compute capacity per watt. And yet, our survey data suggests that server lifespans are becoming longer, often exceeding vendors' recommended three to five years. In 2015, when we first asked this question, only 34% of survey respondents said they kept their servers in operation for five years or longer; in 2022, this proportion has grown to 52% (see **Figure 4**).

Figure 4 **Server refresh cycles are slowing down**

How often does your organization typically refresh its servers?

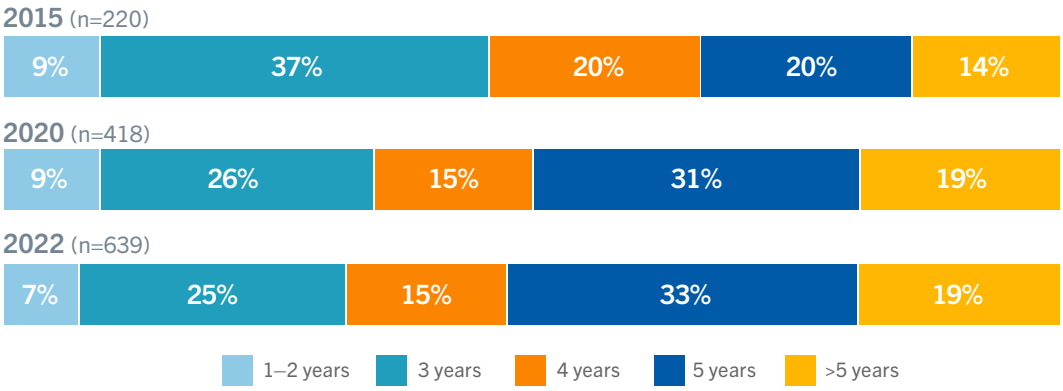

(All figures rounded)

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2015-2022 Uptime Institute | INTELLIGENCE

There are multiple factors behind this increase. One of them is ongoing semiconductor shortages beginning in 2020, resulting in higher prices and increased delivery times for some IT hardware. Smaller organizations with less buying power were often required to hold off on nonessential upgrades.

The trend may also reflect a slowdown in server power efficiency gains. Generational changes, particularly in Intel-powered servers, which make up most of the market, are delivering much lower performance and energy improvements than before. Supply of more efficient servers using alternative (AMD and ARM-based) processors is still limited.

Sustainability and measurement

Sustainability has been a serious consideration for infrastructure operators for many years, but only since 2020 has the environmental footprint of a data center rivalled resiliency as a major concern. In the coming years, legislators and other authorities will force operators to report significantly more data and to demonstrate a commitment to good environmental stewardship.

Uptime's surveys in 2021 and 2022 show that data center operators have a long way to go in this area. Most operators collect data that relates to power efficiency, which is as much about saving money as it is about reducing environmental impact. In 2022, 85% say they report their overall data center power use and 73% report PUE (for either internal or external use), as shown in **Figure 5**.

But when it comes to carbon emissions, the proportion of those who collect this data is still low (37%). This may quickly become an area of concern for businesses, as most organizations and / or their customers will be required to report this data under new laws, initiatives and rules that are being implemented around the world. Worse still, when asked if they collected only Scope 1 and 2 carbon emissions (direct and supply energy) or if they collect Scope 1, 2 and 3 (including emissions from supply chain partners), the numbers were low at 17% and 12%, respectively.

Legislators and other authorities will force operators to report on significantly greater levels of data and to demonstrate a commitment to good environmental stewardship

Uptime recommends that all operators should put plans in place to report all carbon emissions associated with their data centers, regardless of whether there is an immediate legal requirement. Uptime also recommends that water data is collected — even in relatively non-water-stressed environments, water use is becoming a concern.

Only 39% of respondents currently report their water use, which is a 12-percentage point drop from 2021. This drop does not indicate that water efficiency is becoming less important; the change may be attributed to a larger, more diverse survey sample in 2022 than in previous years. Other Uptime survey data shows that most operators that don't track water use say it is because there is no business justification, which suggests a low priority for management (in terms of cost, risk or environmental considerations). However, a growing number of municipalities will permit data center developments only if they are designed for minimal or near-zero direct water consumption. These types of rules will heavily influence facility design and product choices in the future, mandating cooling equipment that uses water sparingly (or not at all).

Figure 5 **Most operators don't track and report key environmental data**

Which IT or data center metrics do you compile and report for corporate sustainability purposes? Choose all that apply.

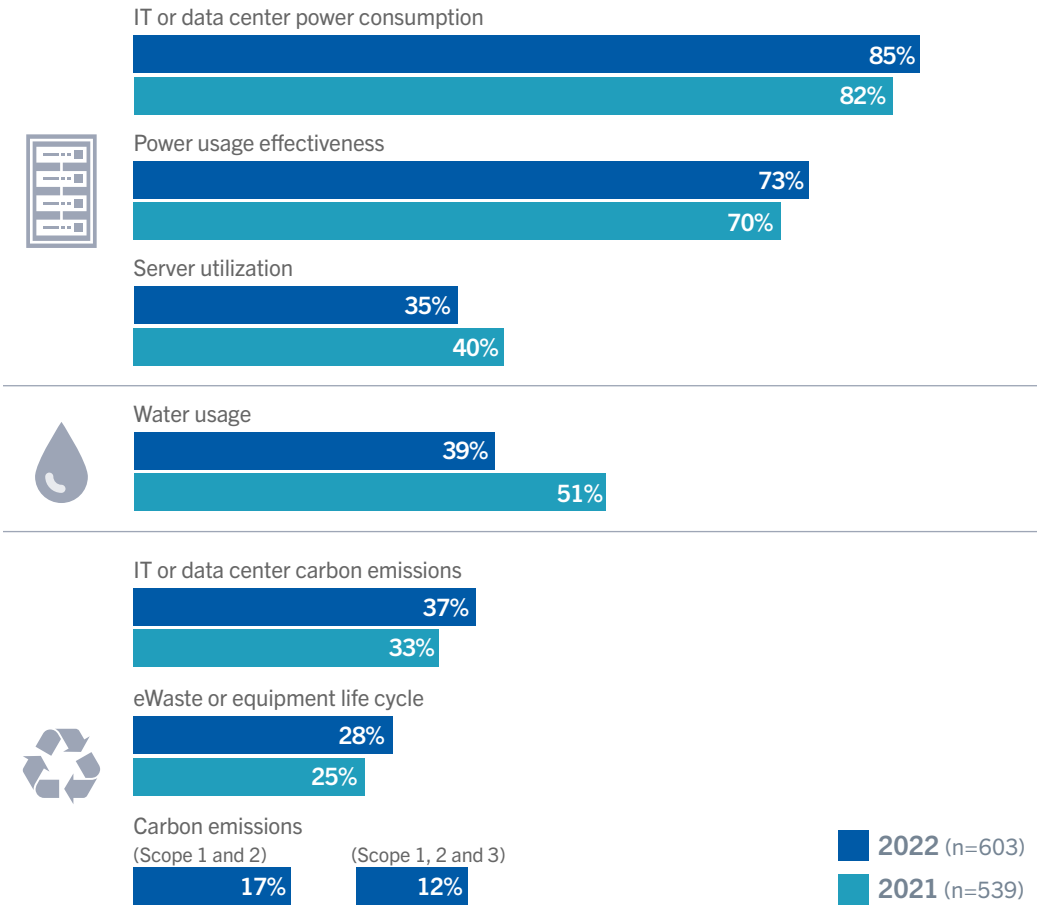

Questions relating to carbon emissions were not included in the 2021 Uptime Institute data center survey

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2021, 2022 UptimeInstitute | INTELLIGENCE

For a detailed explanation of the key areas a data center sustainability strategy should cover and the actions necessary to implement a successful plan, see our report series 'Digital infrastructure sustainability – A manager's guide'.

Operators expect sustainability legislation

Although many are ill-prepared, most data center operators surveyed (63%) think that authorities in their region will require data centers to publicly report environmental data in the next five years, as shown in **Figure 6**. Given that much of this legislation is already in motion (even in less regulated countries, public financial reporting is likely to require mandatory sustainability reporting), many of the remaining 37% will be forced to revise their opinion and instigate preparations. Notably, in North America only half of respondents expect mandatory reporting for sustainability metrics in the next five years.

Figure 6 **Operators expect mandatory sustainability reporting**

In the next five years, do you think that data centers in your region will be required to publicly report on environmental sustainability metrics? ("Yes" responses) (n=675)

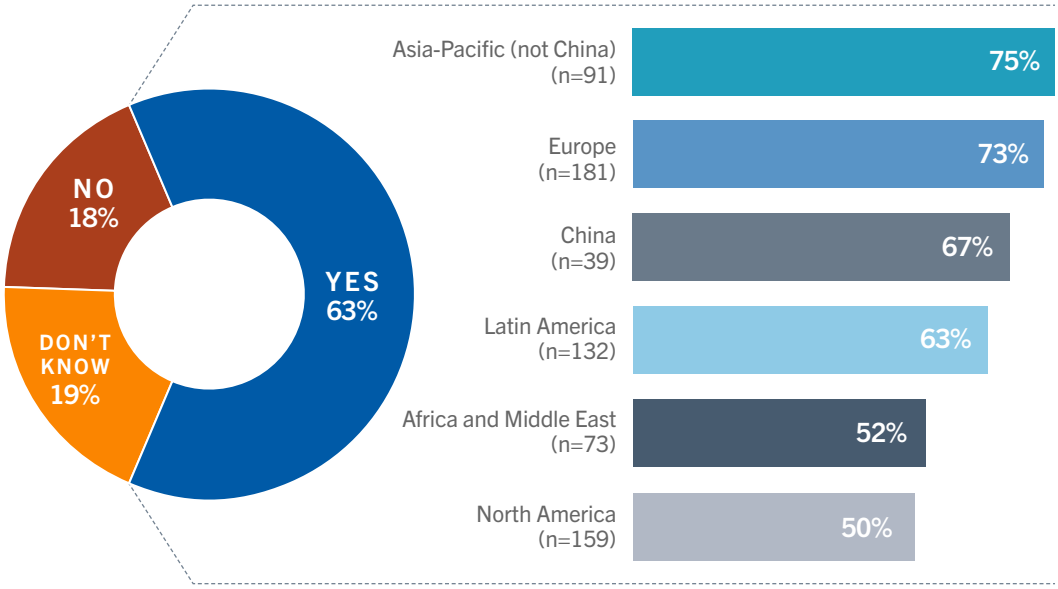

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2022 Uptime Institute | INTELLIGENCE

Data centers "part of the solution"

The data center industry is often singled out as a major consumer of energy and, by extension, a large emitter of carbon emissions. TV documentaries, protests at data conferences and some harsh planning requirements have stoked up the contention. Many in the industry, however, think this finger-pointing is unfair: data centers are estimated to only account for 2% to 3% of electricity use, which represents about 0.4% to 0.75% of global CO2 emissions (electricity accounts for 20% to 25% of total global CO2 equivalent emissions). Some argue that, even allowing for frivolous online services and energy-hogging applications, such as some blockchain technologies, IT enables significant efficiencies and fuel savings elsewhere (for example, more advanced engineering and less business travel) that more than offset its consumption.

Uptime's data shows that two-thirds of data center operators think that data centers (and the services they offer) are part of the solution, not the problem (see **Figure 7**). This viewpoint will, however, be difficult to maintain if increased reporting requirements show that IT and data centers are wasteful.

Figure 7 **We are part of the solution, not the problem, say operators**

In terms of reducing carbon emissions, would you say the data center industry is more likely to be part of the solution or part of the problem in the next five years? (n=664)

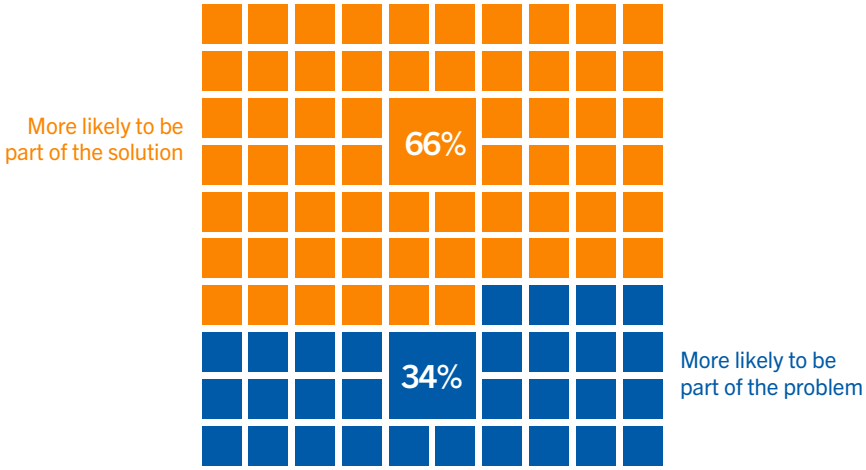

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2022 Uptime Institute | INTELLIGENCE

Making data centers more sustainable

Over several decades, efforts to improve the sustainability of data centers have spanned a large number of technologies, directives, strategies and initiatives. Many of these have played a role. But is there a single driver that would make the industry more sustainable in the next three to five years? Operators were asked to select one option from a range of drivers. Response levels for two of these stood out: More options for buying renewable energy (which primarily requires investment by generators and utilities); and improved cooling (see **Figure 8**). Improved IT utilization — an issue that Uptime argues needs more attention — is only selected by one in six.

Figure 8 Renewables, cooling are biggest drivers for sustainability gains

In your opinion, which of the following will have the biggest impact in making the data center industry more environmentally sustainable in the next three to five years? Choose one. (n=667)

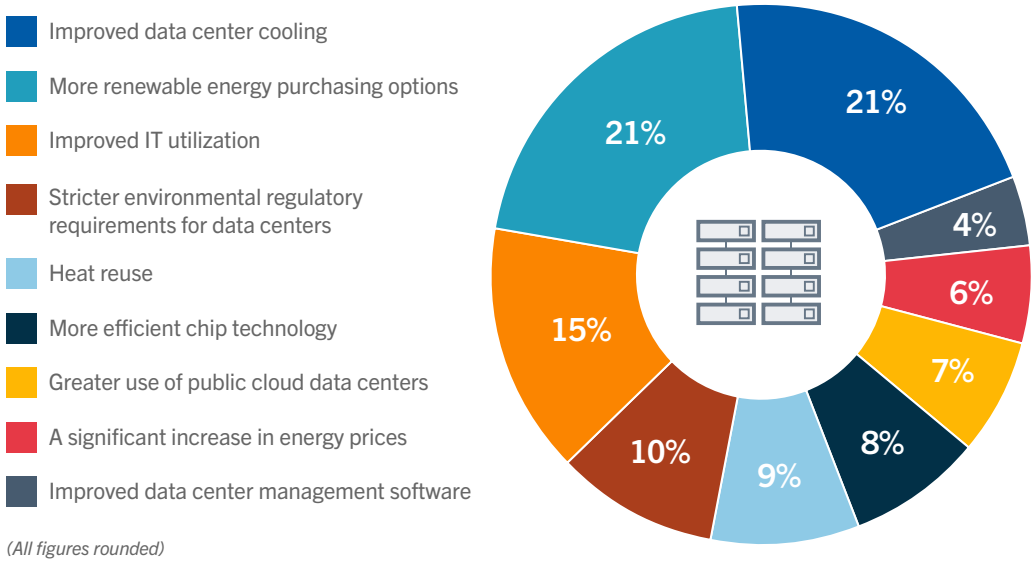

(All figures rounded)

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2022

Uptime Institute | INTELLIGENCE

In practice, none of these operational techniques and energy factors will be transformational on their own and almost all will likely play a role. Higher energy prices and more regulations can be considered near certainties — as is more renewable energy on the grid, although the transition to a clean energy-powered global electrical grid could take several decades. Other technologies will emerge, evolve and be adopted gradually, such as green hydrogen (generated from excess renewable capacity), 10-hour to 10-day storage technologies, including various batteries, compressed air and stored approaches, and other developing technologies. Public cloud operators have invested heavily in attempting to persuade the industry that they are more sustainable, but currently only a small number of respondents see this as a major part of the solution.

Industry cautiously supports nuclear

At the beginning of 2022, Uptime forecast that data center operators, in their search for low carbon, firm (non-intermittent) power sources, would increasingly favor (and even lobby for) nuclear power. Uptime's 2022 annual survey shows that data center operators / owners in major data center economies around the world are cautiously in favor of nuclear power. There are, however, significant regional differences (see **Figure 9**).

Figure 9 **Nuclear is needed, say operators in most regions**

Society is divided on providing nuclear energy to the power grid for lowering carbon footprints. Which of the following options is closest to your view?

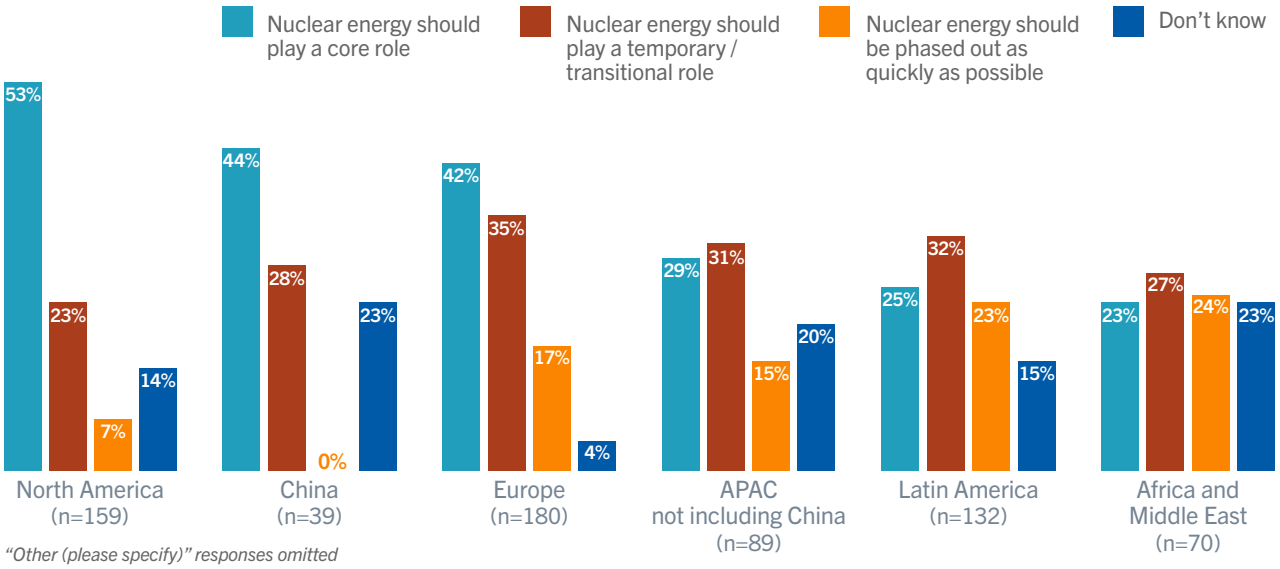

"Other (please specify)" responses omitted from graph (all figures rounded)

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2022

Uptime Institute | INTELLIGENCE

In both North America and Europe, about three-quarters of data center operators believe nuclear should either play a core long-term role in providing grid power or is necessary for a period of transition. However, Europeans are more wary, with 35% saying nuclear should only play a temporary or transitional role (compared with just 23% in North America).

In Europe, attitudes to nuclear power are complex and politicized. Following the Chernobyl and Fukushima nuclear accidents, green parties in Europe lobbied strongly against nuclear power, with Germany eventually deciding to close all its nuclear power plants at the expense of burning more fossil fuel. More recently, the Russian invasion of Ukraine has exposed Germany's over-reliance on energy imports from Russia and many have called for a halt to this nuclear shutdown.

In the US, there is greater skepticism among the general population about climate change being caused by humans, and consequently surveys record lower levels of concern about carbon emissions. There also appears to be lower levels of concern about nuclear safety in the US, too. As the issues of climate change and energy security intensify across the world, the gap in opinion between the US and Europe is likely to close in the years ahead.

In China, not a single respondent thought nuclear power should be phased out — perhaps reflecting both its government's stance and a favorable view of nuclear technology. China, more than most countries, faces major challenges in both meeting energy needs and simultaneously reducing carbon emissions.

In Latin America, and Africa and the Middle East, significantly lower proportions of data center operators think nuclear power should play a key role. This may reflect political reality: there is far less nuclear power already in use in those regions, and concerns about political stability and nuclear proliferation (and cost) will likely limit even peaceful nuclear use.

Realistically, data center operators will not have a major impact on the use (or non-use) of nuclear power. Decisions will primarily be made by grid-scale investors and operators and be steered by government policy. However, large-scale energy buyers can make investments more feasible — and existing plants more economic — if they choose to class nuclear power as a renewable (zero-carbon) energy source and include nuclear in power purchase agreements. They can also benefit by siting their data centers in regions where nuclear is a major energy source. Early-stage discussions around the use of small modular reactors for large data center campuses (see *Data center operators ponder the nuclear option*) are, at present, just that — exploratory discussions.

Resiliency and outages

Operators striving to deliver services with resilient data centers continue to invest in redundancy but outages remain an issue.

Operators report fewer disruptive outages

The frequency of outages does not grow as fast as the global data center footprint

Uptime has been tracking the frequency, causes and impacts of outages in the data center / critical infrastructure industry for many years. Nondisclosure is still the first instinct of most operators and IT outages themselves have become less binary: failures are now often partial, distributed and dependent on user configurations. Knowledge of when an outage occurs, and its severity and cause, may depend on the role and viewpoint of the observer.

Overall, Uptime's data suggests that the number of outages globally increases year on year, as the industry itself expands. This increase, along with the obvious public impact of some of the outages, inevitably attracts attention and headlines. It can also lead to the wrong conclusions: the frequency of outages does not grow as fast as the global data center footprint. Uptime has tracked a steady improvement in the outage rate per site (or per survey respondent). In 2022, 60% of operators surveyed say they had an outage in the past three years (see **Figure 10**) — down from 69% in 2021 and 78% in 2020.

Uptime is reluctant to herald this apparent improvement for two reasons. First, the impact of COVID-19 on business levels and on IT and data center operations has made recent year-on-year comparisons difficult and it may be too early to call this a strong trend. Second, the level of outages is still high, even if seen to be improving.

Figure 10 Most operators had no or negligible outages in the past 3 years

On a scale of 1 (negligible) to 5 (severe), how would you classify the most impactful outage your organization has had in the past three years, either in your own facility or because of a third-party service provider? (n=730)

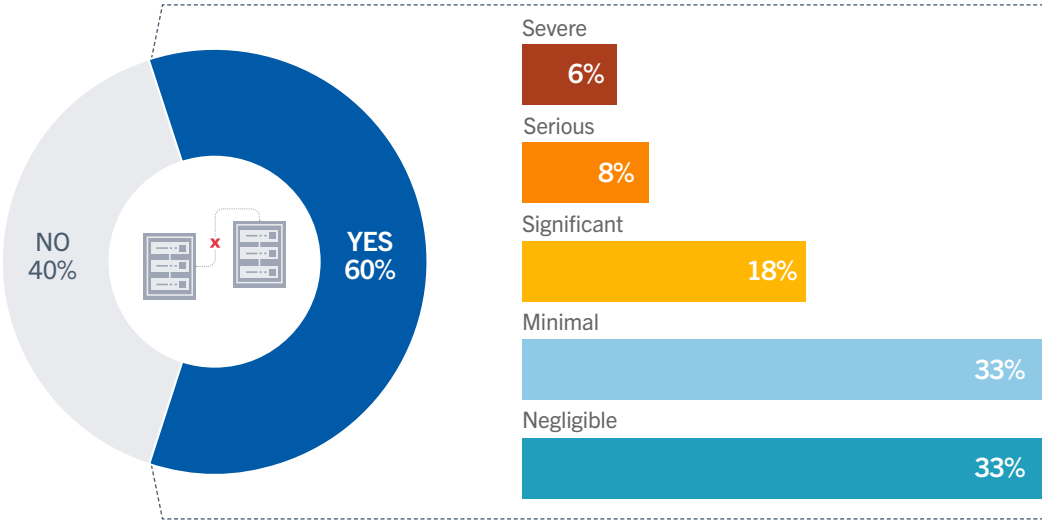

(All figures rounded)

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2022 Uptime Institute | INTELLIGENCE

There are also signs that the impact of at least some outages is decreasing. Uptime classes outages on a scale of 1 to 5, and those in the two most serious / severe categories historically account for about one in five of all outages (among survey respondents who had an outage in the past three years). In 2022, outages in the serious / severe outages categories fell to 14%, or one in six (among those who had an outage in the past three years) (see **Figure 11**). Also, many outages are increasingly caused by partial failures of systems or equipment, rather than total failures — which may also help to lessen the impact.

Figure 11 **Fewer data center outages are significant, serious or severe**

On a scale of 1 (negligible) to 5 (severe), how would you classify the most impactful outage your organization has had in the past three years, either in your own facility or because of a third-party service provider?

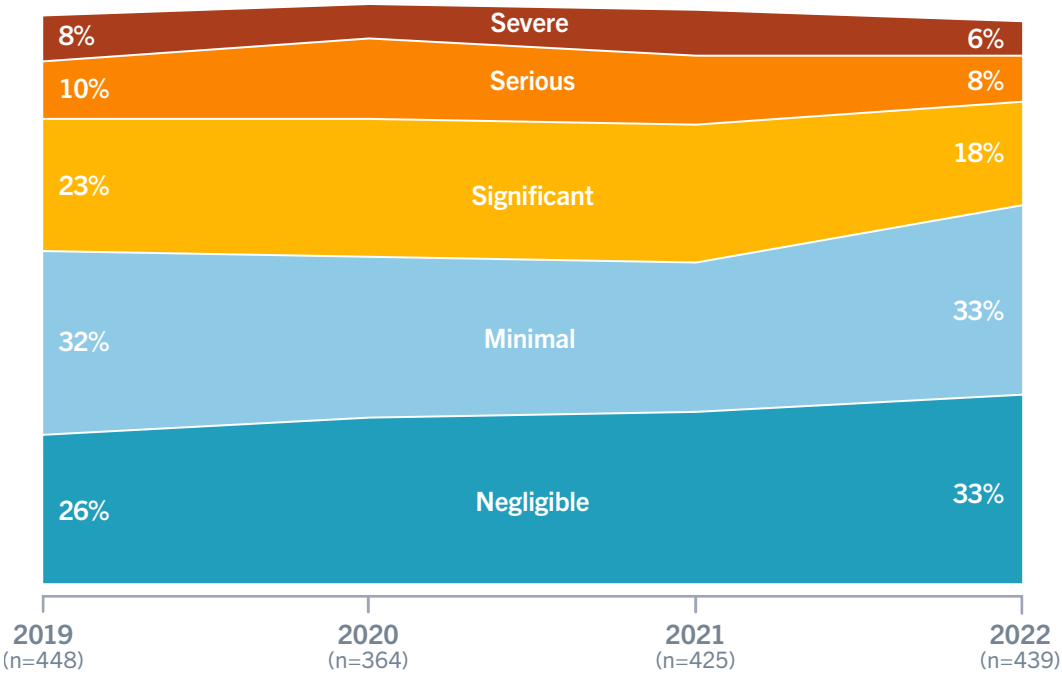

(All figures rounded)

UPTIME INSTITUTE GLOBAL SURVEY OF
IT AND DATA CENTER MANAGERS 2019-2022 Uptime Institute | INTELLIGENCE

Outages become more expensive

A falling proportion of managers may be reporting serious / severe outages, but our data also shows a strong counterpoint: outages are becoming more expensive. When asked about the cost of their most recent outage, a quarter of respondents say the outage had cost more than $1 million in both direct and indirect costs, a significant increase from 2021 and continuing a clear trend (see **Figure 12**). A further 45% say their most recent outage cost between $100,000 and $1 million.

With more than two-thirds of all outages costing more than $100,000, the business case for investing more in resiliency — and training — is becoming ever stronger. As organizations make this investment, the overall number of outages can be expected to fall.

Why is the cost of outages increasing? This can be attributed to a variety of factors, ranging from inflation, fines, service level agreement breaches and the cost of labor, call outs and replacement parts — but the biggest single reason is the growing dependency of corporate economic activity on digital services and on the data center. The loss of a critical IT service often translates directly and immediately into disrupted business and lost revenue.

Figure 12 **Outages costing over $1 million are increasing**

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

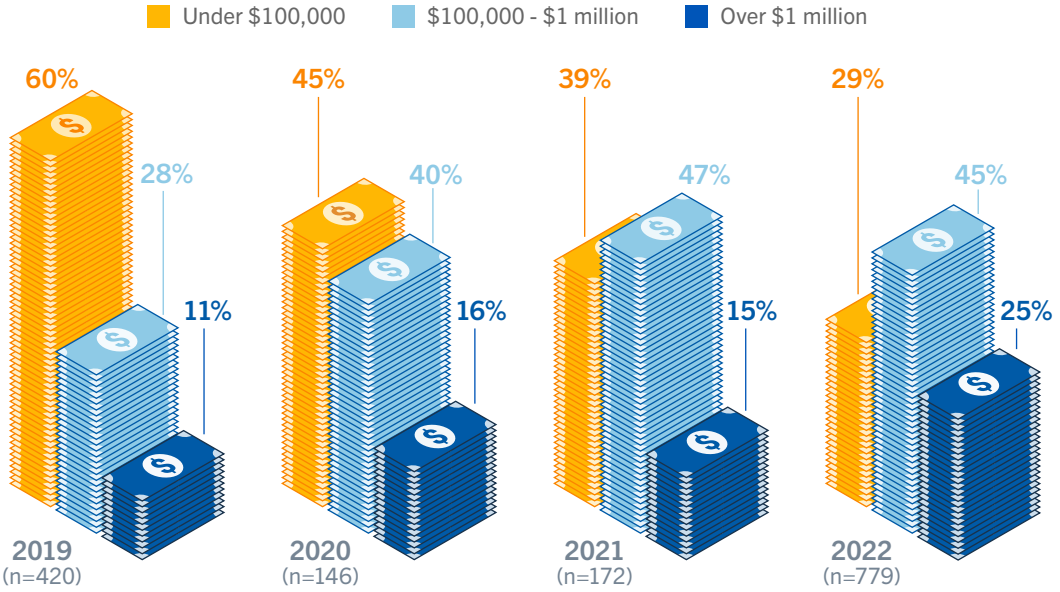

(All figures rounded)

UPTIME INSTITUTE GLOBAL SURVEY OF
IT AND DATA CENTER MANAGERS 2019-2022 Uptime Institute | INTELLIGENCE

Power is still main cause of outages

Understanding the causes of outages is critical to preventing them and to knowing where any investment is necessary. This is not always as simple as it sounds. Most outages have several causes, and, as we have noted, knowledge and understanding of outages may depend on the individual questioned.

Uptime's annual global survey provides a robust, consistent data set from a group of respondents who have detailed knowledge of data center operations — but not all respondents from this group have detailed knowledge of all their organization's IT outages and causes. Deeper analysis of the different perspectives is provided in Uptime's report *Annual outage analysis 2022: The causes and impacts of data center outages*.

Uptime's 2022 annual survey findings are remarkably consistent with previous years. They show that on-site power problems remain the single biggest cause of significant site outages by a large margin (see **Figure 13**). As in previous years, all other outage causes are far less common. Three other common causes stand out as particularly troubling: cooling failures, software / IT system errors and network issues. The frequency of problems at third-party providers (e.g., SaaS, hosting and cloud providers) is creeping up, reflecting a greater use of cloud, SaaS and colocation.

Deeper analysis in separate Uptime research identifies the biggest causes of power-related outages to be uninterruptible power supply failures, followed less commonly by transfer switch (generator / grid) and generator failures. While utility grid failures are never attributed by Uptime as a primary cause of outages, the slight increase in power-related failures in recent years may correlate with degrading grid reliability that lay bare substandard maintenance and training at some data center sites.

Figure 13 **Power is still main cause of outages**

What was the primary cause of your organization's most recent impactful incident or outage?

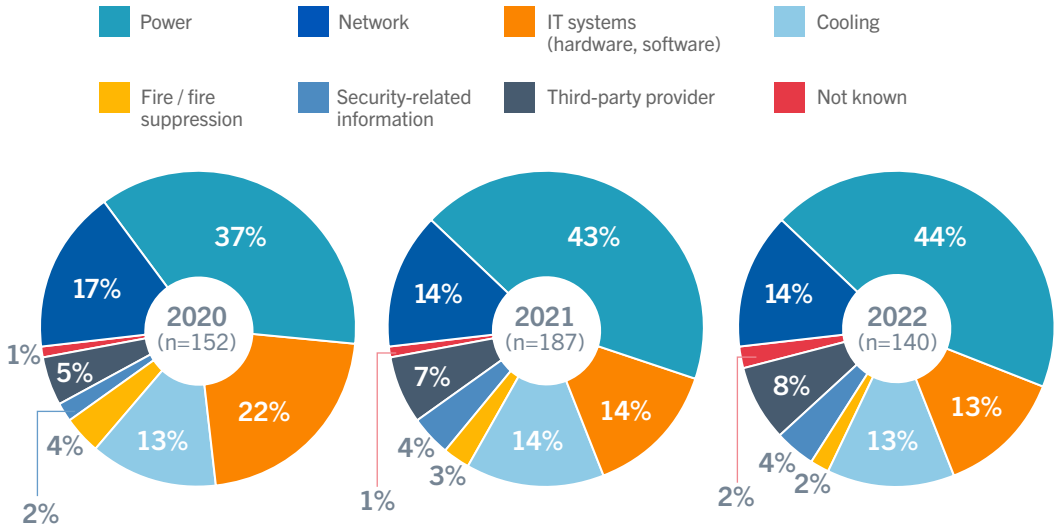

(All figures rounded)

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2020-2022 Uptime Institute | INTELLIGENCE

Most outages are preventable

Uptime does not class human error as a technical cause of failure — but rather as a contributory factor that has a role in a large percentage of outages (usually between 60% and 80% of outages, depending on the data source and how errors are categorized). This is analyzed in more detail in *Uptime's Annual outage analysis 2022: The causes and impacts of data center outages*.

In our 2022 annual survey, we also ask those who have experienced an outage whether they think their most recent impactful outage would have been preventable with better management, processes or configuration. As **Figure 14** shows, nearly four in five say "yes" — a number that has risen gradually in recent years. This supports Uptime's frequently made point: the most impactful and cost-effective way to reduce outage occurrences is to improve management, planning and training.

Figure 14 **Most operators still view downtime as preventable**

Would your organization's most recent impactful downtime incident have been preventable with better management / processes or configuration?

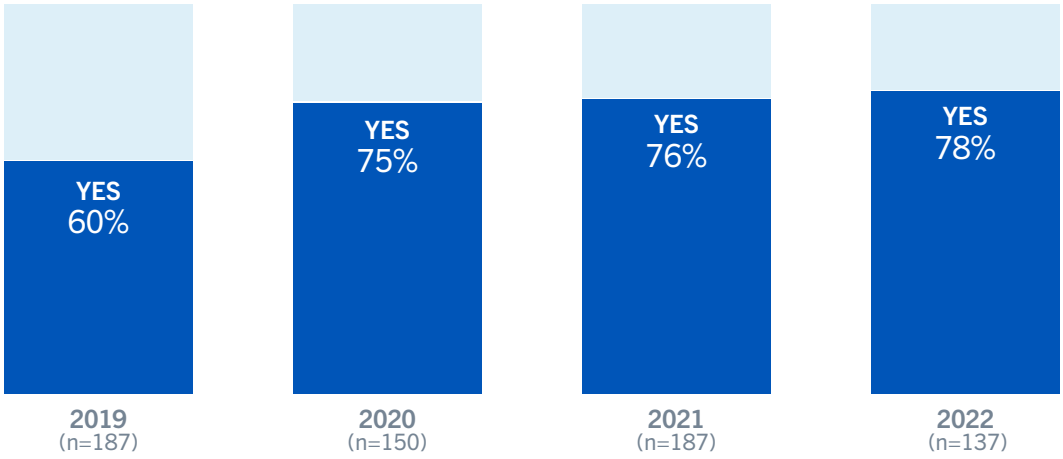

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2019-2022

Uptime Institute | INTELLIGENCE

More are increasing data center resiliency

The availability of IT services continues to lean primarily on the resiliency of data center facilities, with no apparent move toward lower physical infrastructure redundancies. This is perhaps unsurprising, although not necessarily a given. The growing use of public cloud infrastructures and cloud-style enterprise IT has been accompanied by the broader adoption of multisite resiliency, which was once exclusive to mission-critical applications. New software development techniques help make distributed redundancy easier because network traffic and workloads can be dynamically diverted to, and lost service easily recovered at, other sites.

Yet, operators continue to invest in increasing the resiliency of their physical infrastructure. About 40% of respondents say that they have increased the redundancy levels of their primary data centers in the past three to five years. Power and cooling systems have received similar attention, with about a third of operators surveyed upgrading either or both. IT services providers (cloud / hosting / SaaS) were more likely to add redundancy to their infrastructure, as shown in **Figure 15**. Only about 4% of all respondents reduced physical redundancy levels (power and / or cooling).

Figure 15 **Levels of physical infrastructure redundancy continue to rise**

How have redundancy levels changed in the past three to five years in your primary data center? (n=619)

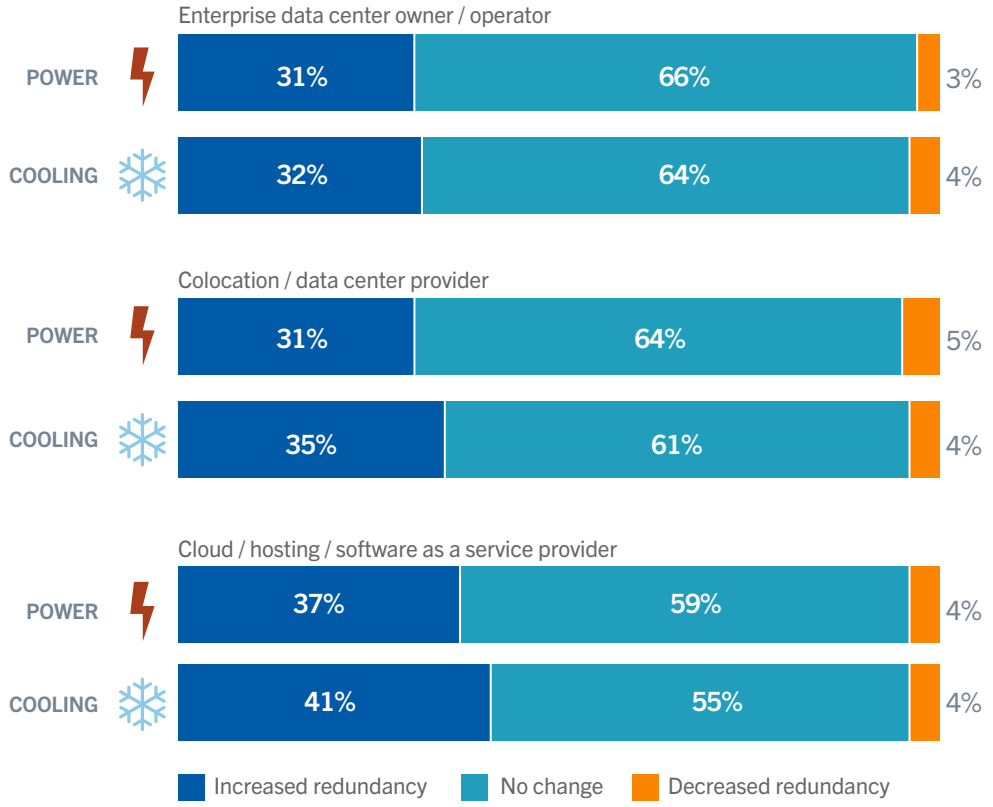

(All figures rounded)

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2022 Uptime Institute | INTELLIGENCE

The growing business criticality of many IT services shapes the risk profiles of digital infrastructure and, in turn, redundancy upgrades. While Uptime expects more operators to make greater use of distributed resiliency in the future, there is no sign of a diminishing need for site-level resiliency. More resiliency at every level is not only the least risky approach but also a justified expense for the business, according to the consensus.

Users unprepared for inevitable cloud outages

Organizations are becoming more confident in using the cloud for mission-critical workloads, partly due to a perception of improved visibility into operational resiliency. The proportion of respondents not placing mission-critical workloads into a public cloud today dropped from 74% (2019) to 63% (2022), while those saying they have adequate visibility into the resiliency of the service provided by a public cloud rose from 14% to 21% (see **Figure 16**).

Figure 16 **Improved visibility drives greater trust in public cloud**

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

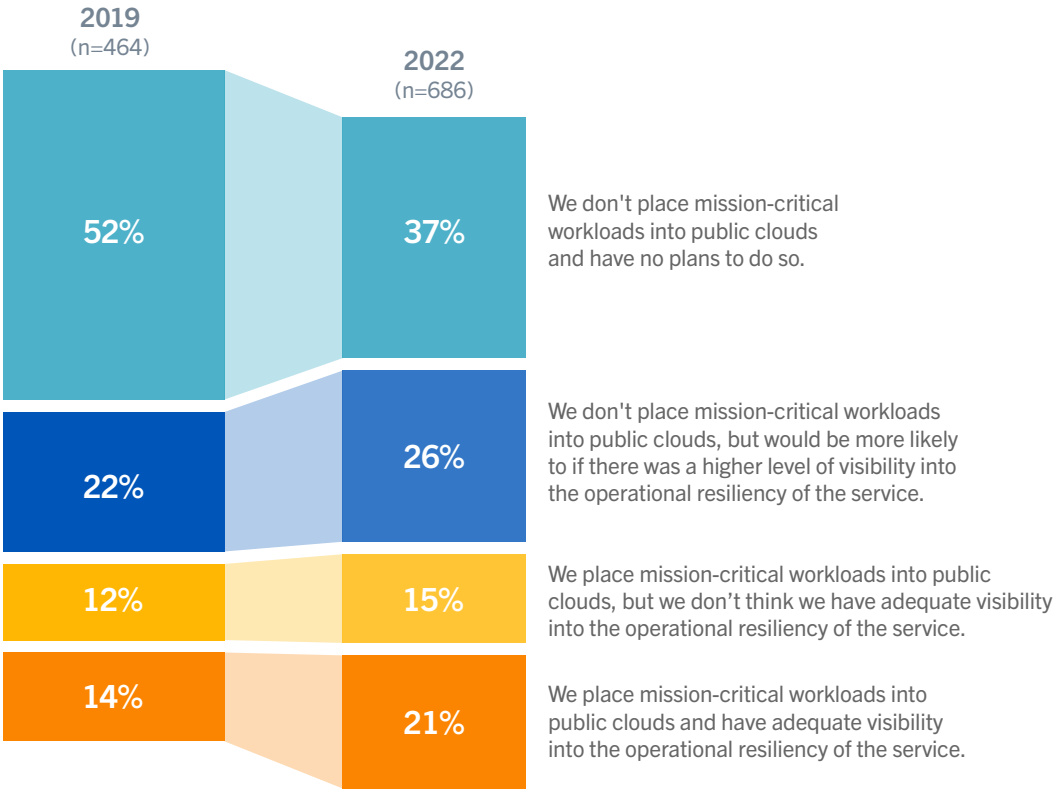

(All figures rounded)

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2019, 2022 Uptime Institute | INTELLIGENCE

However, other data suggests cloud users' confidence may be misplaced. Cloud providers recommend that users distribute their workloads across multiple availability zones. An availability zone is a logical data center, often understood to have redundant and separate power and networking. Cloud providers are upfront that zones will suffer outages occasionally. Their position is that the user must architect their applications to handle the loss of an availability zone.

Zone outages are relatively common and more than a third of respondents say that the loss of an availability zone would result in significant performance issues, as shown in **Figure 17**.

Figure 17 **Many cloud applications vulnerable to availability zone outages**

If your primary cloud provider were to experience an outage across a single availability zone, what would be the likely impact on most of your cloud applications running in that availability zone? (n=237)

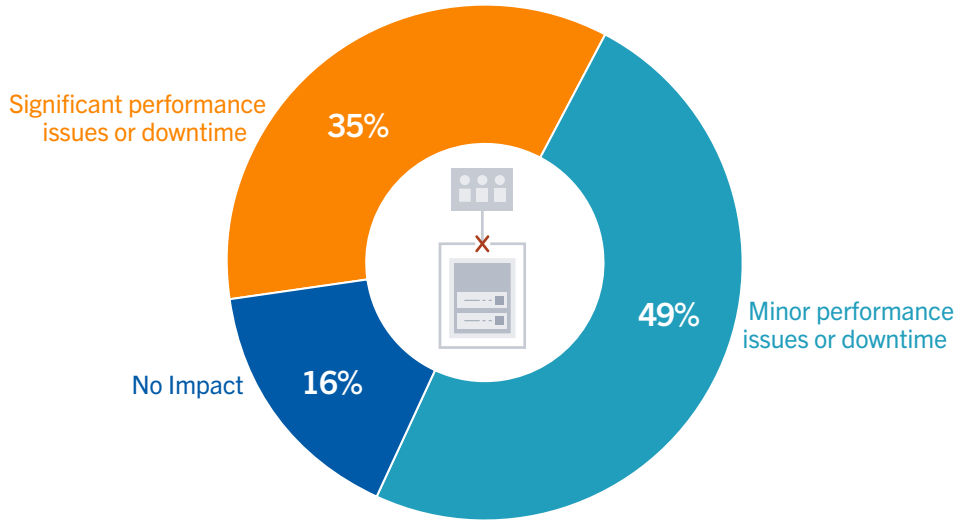

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2022

Uptime Institute | INTELLIGENCE

This presents a clear contradiction. Users appear more confident that the cloud can handle mission-critical workloads, yet over a third of users are architecting applications vulnerable to relatively common availability zone outages. This contradiction is due to a lack of clarity on roles and responsibilities between provider and user.

As shown in **Figure 18**, half of respondents believe that if a single availability zone fails and an application becomes unavailable, it is primarily the cloud provider's fault.

Figure 18 **Ambiguous accountability threatens cloud application resiliency**

Who would you say is most at fault when an application fails due to a public cloud outage across a single availability zone and why? Choose one. (n=328)

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2022 Uptime Institute | INTELLIGENCE

Organizations that want to achieve high availability in the cloud must architect their applications to endure frequent outages of single zones

The provider is, of course, responsible for the operational resiliency of its data centers. But cloud providers neither state nor guarantee that availability zones will be highly available. So, why are users assuming that a single zone will provide the resiliency their application requires?

Some of this misunderstanding is likely due to the simplified view that the cloud is just someone else's computer in someone else's data center — but this is not the case. A cloud service is a complex combination of data center, hardware, software and people. Services will fail from time to time due to unexpected emergent behavior from the complexity of interacting systems and people.

As a result, organizations that want to achieve high availability in the cloud must architect their applications to endure frequent outages of single zones. Lifting and shifting a workload from an on-premises server to a cloud virtual machine might reduce resiliency if the workload is not rearchitected to work across cloud zones.

As cloud adoption increases, the impact of outages is likely to grow as a significantly higher number of organizations rely on cloud computing for their applications. While many will architect their applications to weather occasional outages, many are not yet fully prepared for inevitable cloud service failures and the subsequent impact to their applications.

Vendors and supply chains

Thrown off-balance by the pandemic, supply chains remain stretched by continued demand for new data center capacity and facility upgrades. Three-quarters of vendors surveyed, including equipment makers, engineering services firms and consultants, project their revenues to increase in 2022, compared with 2021. Vendors of all sizes and in every geography expect to perform well, with North America leading the charge. Chinese vendors were the least upbeat, likely due to some prolonged COVID 19-related lockdowns early in 2022. Globally, more than half of vendors surveyed say spending on data centers is above the normal trend, which is higher than when we asked in 2021.

This strength of demand is against a backdrop of major supply chain difficulties. Most vendors say their revenues have been negatively affected by supply chain issues. Nearly half of those respondents involved with data center construction have had at least some major delays (or other events) in their supply chains, with an additional one-third reporting moderate issues. Fewer than one in five say they have had only minor or no disruptions.

Prolonged troubles in supply chains mean vendors have little means to shield their customers against the resulting impacts. Only about a quarter of data center equipment vendors can meet all orders in a timely fashion (see **Figure 19**). A slightly higher number have prioritized supply to key customers. Supply in North America appears to be considerably tighter than the global average, which is most likely due to unusually high demand for capacity after a period of high absorption rates.

Figure 19 **Most vendors delay deliveries to customers**

Which of the following options best describes the timeliness of your product supply to your data center customers? (n=467)

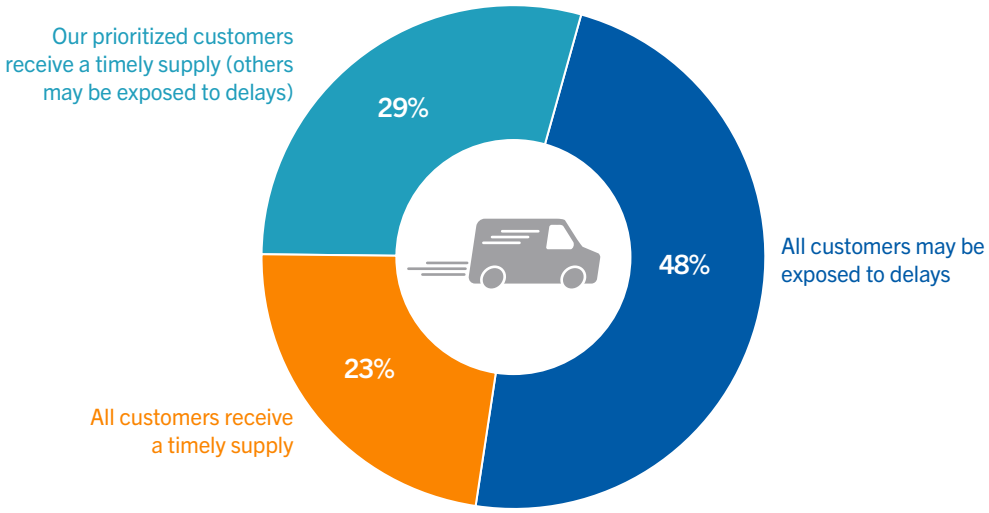

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2022

Uptime Institute | INTELLIGENCE

The outlook is similarly mixed. Vendors continue to be bullish about the data center industry's continued growth, which they expect will be further boosted by demand for micro-edge capacity. Conversely, two-thirds expect data center projects to be affected by supply chain delays in the next two years. These views predate the recent slowdown in semiconductor demand, which may offer some respite for supply chains to regain their balance, but any sign of a significant improvement will take months to transpire. Furthermore, half the vendors surveyed say it is not only supply chains, but also staffing shortages that will hinder the data center sector's expansion (see **Staffing shortfalls**).

Staffing shortfalls

Attracting and retaining qualified data center staff has been a challenge for operators for more than a decade. Continuously escalating demand for data center capacity has driven an increase in the number and size of facilities, and a proliferation of job openings that still outpaces recruitment. In 2021, Uptime published the first forecast of data center workforce needs, reported by region, data center type and education requirements. We estimate staff requirements will grow globally from about 2.0 million full-time employee equivalents in 2019 to nearly 2.3 million in 2025. (See our report *The people challenge: Global data center staffing forecast 2021-2025*.)

The staff shortage affects almost all data center job roles globally. In mature data center markets, such as North America and Western Europe, much of the existing workforce is aging and many professionals expect to retire around the same time, leaving data centers with a shortfall on both headcount and experience. Hiring efforts are often offset by jobseekers' poor visibility of the sector. Efforts to bolster talent pipelines by attracting career-changers to the data center industry are still nascent.

Problems with attracting and retaining staff appear to be worsening. As shown in **Figure 20**, over half (53%) of operators surveyed report difficulty finding qualified candidates for open jobs — up from 47% in 2021, and 38% in 2018. Operators also face difficulties with employee retention — 42% report staff being hired away, which is more than double the 2018 figure of 17%. Most staff who leave are hired by data center competitors and few are leaving the industry.

Roughly one in three (31%) respondents say they are not having difficulty attracting or retaining staff.

Figure 20 **More operators struggle with attracting, retaining staff**

Please select any of the following statements that apply.

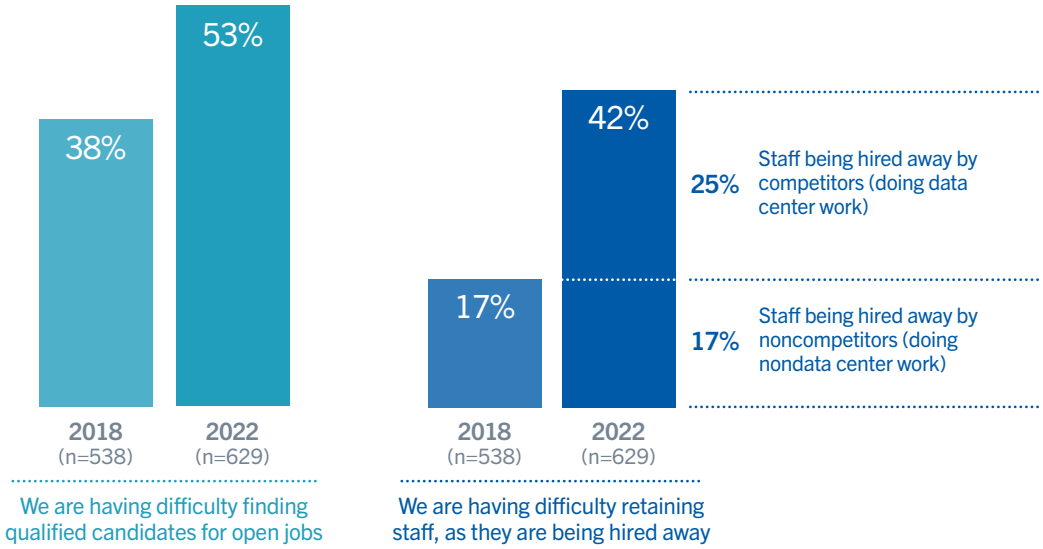

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2018, 2022 Uptime Institute | INTELLIGENCE

The most successful employers in the sector take multiple steps to attract and retain top talent, including revisiting advertised job requirements, implementing training and mentoring programs and ensuring adequate diversity efforts. Gender and other demographic imbalances can deter some candidates from pursuing a career in the data center sector if they are not well represented, reducing the size of the labor pool.

On average, data center design, build and operations teams employ few women (and often other underrepresented groups in certain geographies). Uptime has collected data on the gender demographics of the sector since 2018 and the imbalance has not changed materially. As shown in **Figure 21**, only 4% of operators say that about half of their data center staff are women. More than three-quarters report that their data center workforce is around 10% women or less. One in five data center teams do not employ any women at all.

Figure 21 **Women remain underrepresented in the data center industry**

What portion of your organization's data center design, build or operations staff is women?

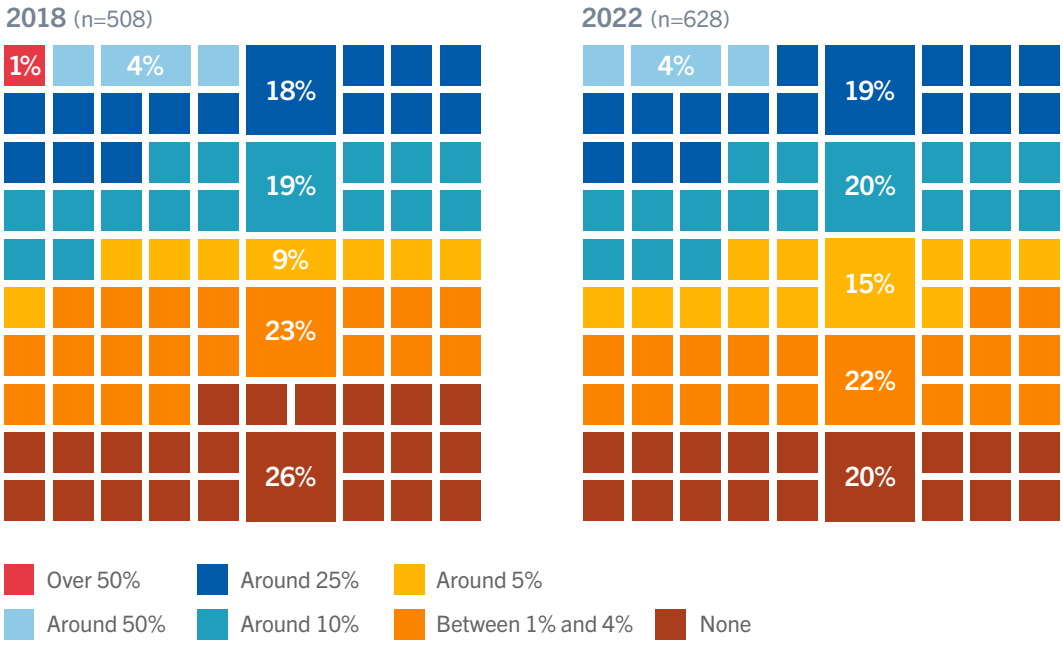

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2018, 2022

Uptime Institute | INTELLIGENCE

The growing number of unfilled job positions relative to the low and stagnating proportion of women workers suggests that the industry has much work ahead to leverage the untapped potential of qualified women and other underrepresented candidates.

Innovation and impact

The industry continues to look for ways to cut back on capital needs of new builds and to improve energy efficiency. Several technologies have been proposed as the next step in data center evolution and these are currently undergoing widespread testing or proof of concept deployments.

Half of data center operators surveyed say software-defined power is one of the technological innovations they thought most likely to deliver significant improvements in energy efficiency in the next five years, as shown in **Figure 22**. This category includes various capabilities that rely on software interacting with electrical systems, including load shedding, load balancing and server throttling.

Artificial intelligence (AI) for data center operations, the development of multisite resiliency as an alternative to equipment redundancy and the adoption of DLC are among other efficiency technologies being closely observed. Which technologies are least likely to impact data center efficiency? Few respondents expect metaverses (for example, augmented reality and digital twins) and the practice of rejecting heat into water to drive the most significant efficiency improvements across the sector.

Figure 22 **Operators expect power and cooling to deliver better efficiency**

Thinking about the next five years, which of these innovations is likely to deliver the most significant improvements in the efficiency of the data centers? Choose no more than three. (n=744)

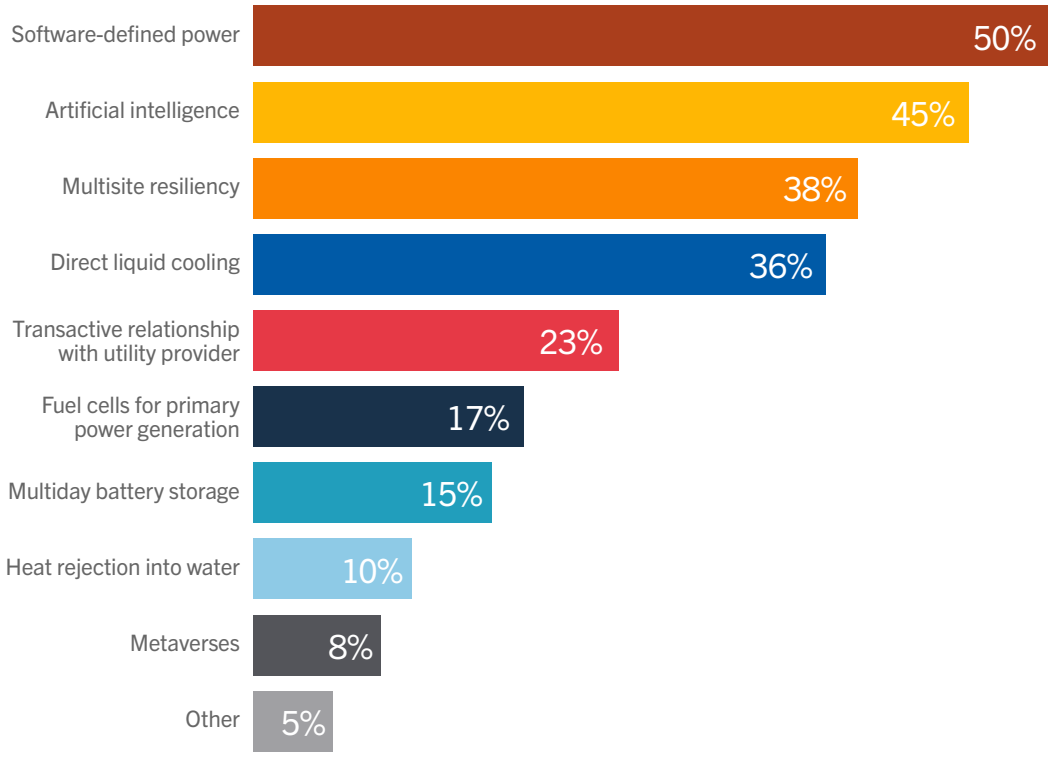

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2022 Uptime Institute | INTELLIGENCE

AI is not replacing operations staff — yet

Owner and operator confidence in AI as a tool for data center management is rising, with 57% of respondents saying they would trust an adequately trained machine learning model to make operational decisions, up from 49% in 2021.

AI-based components are already being incorporated into readily available data center power and cooling systems. It remains unclear whether AI will be used to replace employees. Most operators still expect AI to reduce staff numbers over the longer term, but fewer than one in five believe this will happen in the next five years (see **Figure 23**). The proportion of respondents hoping that AI can provide a near-term fix to their staffing challenges has dropped steadily since 2020, hinting at lowered expectations that are more aligned with the current and near-term capabilities of AI.

Figure 23 Fewer operators expect AI to reduce staffing in the near term

Do you believe artificial intelligence will reduce your data center operations staffing levels in the next five years?

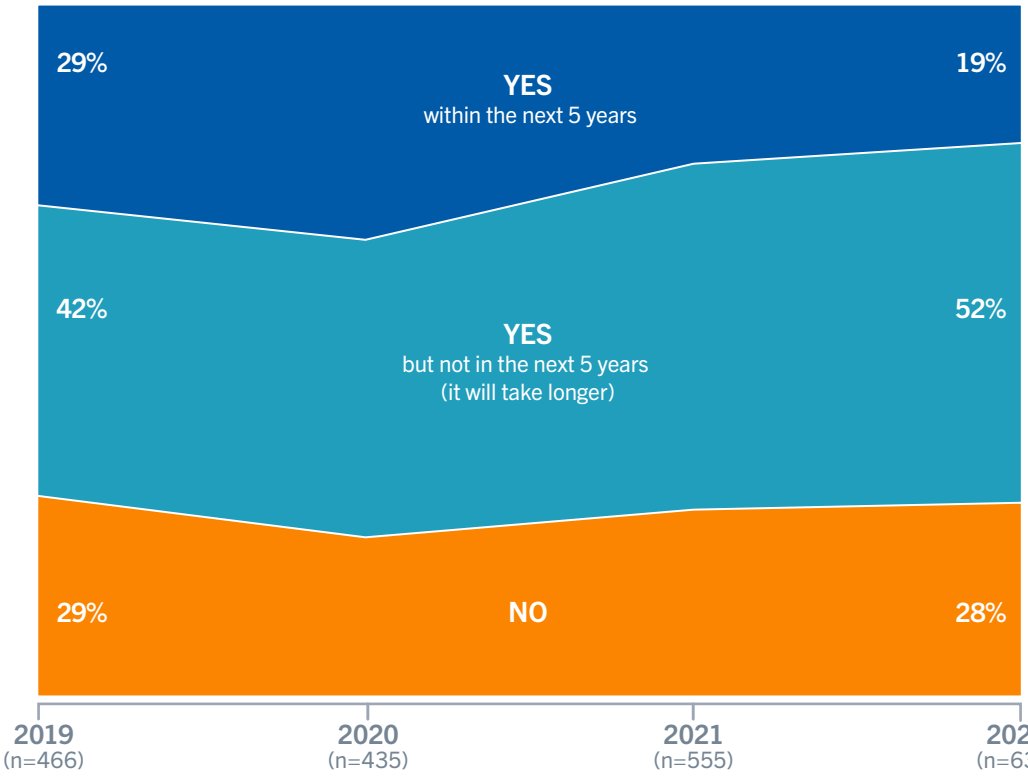

(All figures rounded)

UPTIME INSTITUTE GLOBAL SURVEY OF IT AND DATA CENTER MANAGERS 2019-2022

Uptime Institute | INTELLIGENCE

Appendix

Survey methodology and demographics

Uptime Institute's Global Data Center Survey, now in its 12th year, is conducted annually online and by email. The 2022 survey was conducted in the first half of the year.

Respondents are separated into two groups: data center owners and operators (Uptime Institute Global Survey of IT and Data Center Managers; 2022 demographics shown in **Figure A1**) and data center suppliers, designers and advisors (Uptime Institute Global Survey of Data Center Designers, Consultants and Vendors; 2022 demographics shown in **Figure A2**).

This report focuses on responses from the owners and operators of data centers, including those responsible for managing infrastructure at the world's largest IT organizations. Job titles include senior executive, IT manager, IT operations staff, critical facilities manager, critical facilities operations staff, design engineer and consultant.

Figure A1 **Uptime Institute Global Data Center Survey 2022: End user demographics**

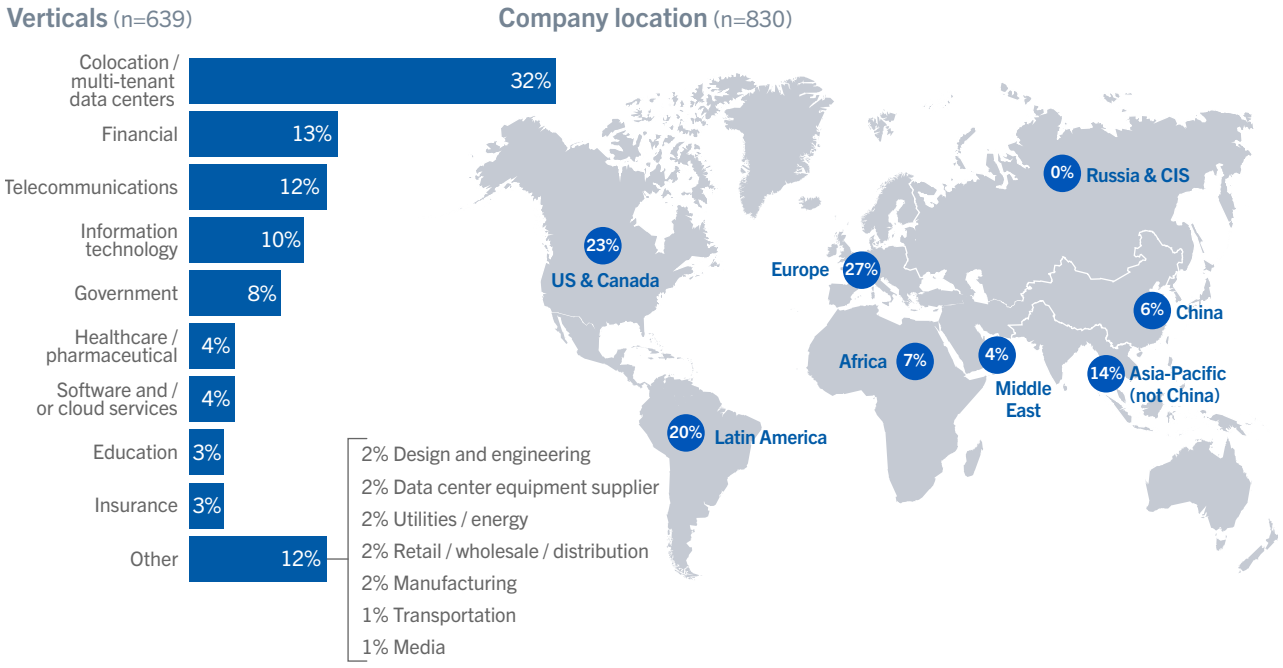

Job function (n=829)

24% Critical facilities manager	**20%** Senior executive (VP or C-level)	**13%** Critical facilities operations staff	**11%** IT manager	**7%** Design engineer	**6%** IT operations staff	**5%** Consultant	**15%** Other

(All figures rounded)

Appendix

Figure A2 Uptime Institute Global Data Center Survey 2022: Supplier demographics

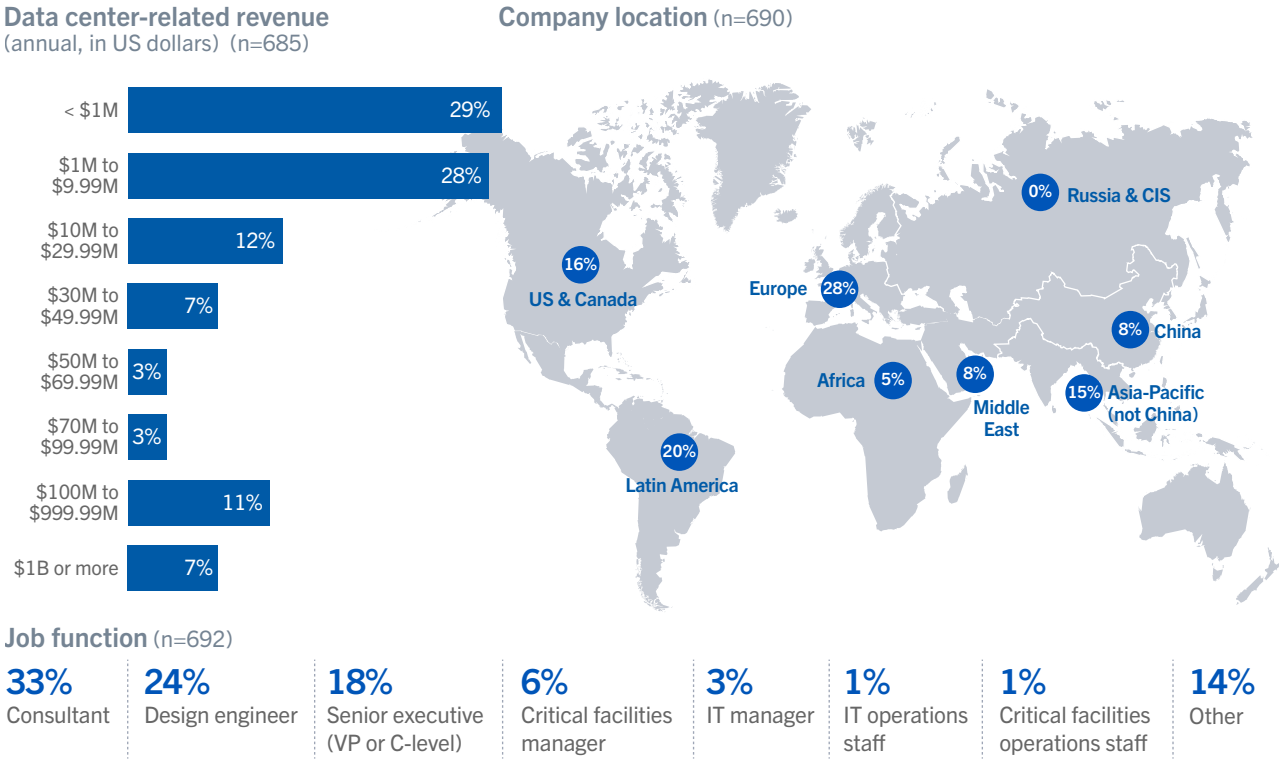

(All figures rounded)

UPTIME INSTITUTE GLOBAL DATA CENTER SURVEY 2022 Uptime Institute | INTELLIGENCE

The participants represent a wide range of industry verticals in multiple countries. Half are in North America and Europe. Approximately 40% of respondents work for professional IT / data center service providers — that is, staff with operational or executive responsibilities for a third-party data center, such as those offering colocation, wholesale, software or cloud computing services.

A total of 830 end users registered for the survey and answered at least one question. The number of respondents ("n") varies between individual questions because respondents are not required to answer every question.

Findings of previous surveys are available **here**.

If you have questions, comments or seek further insights, please contact **research@uptimeinstitute.com**.

About the authors

Jacqueline Davis

Jacqueline Davis is a Research Analyst at Uptime Institute covering global trends and technologies that underpin critical digital infrastructure. Her background includes environmental monitoring and data interpretation in the environmental compliance and health and safety fields.

jdavis@uptimeinstitute.com

Daniel Bizo

Daniel Bizo is Uptime Institute's Research Director. He has been covering the business and technology of enterprise IT and infrastructure in various roles for 15 years, including a decade as an industry analyst and advisor.

dbizo@uptimeinstitute.com

Andy Lawrence

Andy Lawrence is Uptime Institute's Executive Director of Research. He is Uptime Institute's Executive Director of Research and has spent three decades analyzing developments in IT, emerging technologies, data centers and infrastructure. He also advises companies on their technical and business strategy.

alawrence@uptimeinstitute.com

Dr. Owen Rogers

Dr. Owen Rogers is Uptime Institute's Research Director of Cloud Computing. He has been analyzing the economics of cloud for over a decade as a product manager, a PhD candidate and an industry analyst. Rogers covers all areas of cloud, including economics, sustainability, hybrid infrastructure, quantum computing and edge.

orogers@uptimeinstitute.com

Max Smolaks

Max Smolaks is a Research Analyst at Uptime Institute. His expertise spans digital infrastructure management software, power and cooling equipment, and regulations and standards. He has 10 years' experience as a technology journalist, reporting on innovation in IT and data center infrastructure.

msmolaks@uptimeinstitute.com

Contributing authors

Lenny Simon

Lenny Simon is a Senior Research Associate at Uptime Institute covering global trends, sustainability and technologies in critical digital infrastructure. Before joining Uptime, he was an analyst at GreenMax Capital Advisors and worked in the Office of the Mayor of New York City.

lbrandes-simon@uptimeinstitute.com

Douglas Donnellan

Douglas Donnellan is a Research Associate at Uptime Institute covering sustainability in data centers. His background includes environmental research and communications, with a strong focus on education.

ddonnellan@uptimeinstitute.com

About Uptime Institute

Uptime Institute is the Global Digital Infrastructure Authority. Its Tier Standard is the IT industry's most trusted and adopted global standard for the proper design, construction, and operation of data centers – the backbone of the digital economy. For over 25 years, the company has served as the standard for data center reliability, sustainability, and efficiency, providing customers assurance that their digital infrastructure can perform at a level that is consistent with their business needs across a wide array of operating conditions. With its data center Tier Standard & Certifications, Management & Operations reviews, broad range of related risk and performance assessments, and accredited educational curriculum completed by over 10,000 data center professionals, Uptime Institute has helped thousands of companies, in over 100 countries to optimize critical IT assets while managing costs, resources, and efficiency.

Uptime Institute is headquartered in New York, NY, with offices in Seattle, London, Sao Paulo, Dubai, Singapore, and Taipei.

For more information, please visit www.uptimeinstitute.com

All general queries:

Uptime Institute
405 Lexington Avenue,
9th Floor, New York,
NY 10174, USA
+1 212 505 3030
info@uptimeinstitute.com