

DATA CENTER

2025

Closer To The Edge



The Journey Continues

Two major developments are driving the future of the data center industry: more distinct segmentation in data center types and the growth of edge computing.

This is the context for a new survey by Vertiv that builds on research conducted in 2014. That research asked data center professionals across the globe to look 10 years into the future and share their perspectives on how market and technology trends would shape the future of the industry.

Five years later, we now have answers to some of the questions posed by the original Data Center 2025 report, “Data Center 2025: Exploring the Possibilities,” but new questions have emerged.

This report, “Data Center 2025: Closer to the Edge,” serves as a midpoint check-in on the original report. Comprising insights from more than 800 data center professionals, as well as Vertiv experts, it reviews the progress made in key areas while providing new perspective on the trends that have emerged since the original report was published.





Contents

- 4 Five Years of Innovation and Disruption**
How the industry is evolving since the original Data Center 2025 report and why a single vision of the 2025 data center is no longer possible.
- 5 Powering the Data Center of the Future**
How expectations around electrical energy sources have changed since the original Data Center 2025 research and the role renewables will play in the future.
- 7 Managing the Growing Demand for Compute**
The impact growing compute and storage demand will have on IT utilization rates, rack density and cloud dependency.
- 10 The Rise of Edge Computing**
A look at how networks are expected to grow by 2025 and the new applications most likely to benefit from 5G.
- 13 Critical Infrastructure Technology**
The thermal management, backup power and data center management strategies most likely to support data centers in 2025.
- 17 A Revised View of 2025**
A summary of key findings and takeaways from the Data Center 2025 initiative.
- 17 Participant Profile**
A breakdown by region, industry, job function, data center dependence of 2019 Data Center 2025 survey participants.

Five Years of Innovation and Disruption

Tracking expectations against progress can highlight ways in which technology is advancing unpredictably or, more significantly, the impact of unanticipated disruptive forces on the trajectory of an industry.

In the case of Data Center 2025, we see more of the latter than the former. While technologies such as servers, infrastructure systems and management tools have advanced somewhat predictably, the nature of what constitutes a data center has changed dramatically.

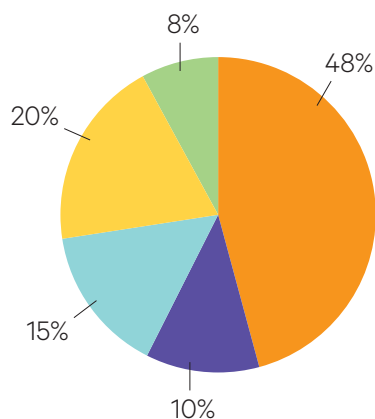
Specialization and Segmentation

In 2014, the primary compute platform for most businesses was still a discrete, multi-purpose enterprise data center that was just beginning to shift workloads to the cloud and colocation in a meaningful way.

Cloud computing was certainly building momentum, as two-thirds of original Data Center 2025 participants (67%) expected at least 60% of compute

How would you best describe the type of data center facility you are most involved with?

Figure 1: Participant profile by data center type for 2019 Data Center 2025 survey



- On-premise core enterprise/private cloud
- Hyperscale/public cloud
- High-performance computing
- Colocation
- Remote or edge

capacity to be performed in the cloud by 2025. That may seem low today compared to some analyst projections but at least demonstrates an understanding of the impact cloud computing would have.

Yet despite this growing presence, data centers across the cloud, colocation and enterprise segments included in the original survey were more similar than different.

That's no longer the case.

Today, when we say "data center" it's necessary to specify exactly what type of facility we are talking about. Is it the high-performance computing facility with rack densities approaching 50 kW? The hybrid enterprise data center that is increasingly managing resources across cloud, colocation and multiple distributed sites? Or is it those distributed sites themselves, which are becoming more critical as services and applications move closer to users?

All have different physical characteristics and different roles in an increasingly integrated and interconnected ecosystem that has evolved to meet the capacity and service demands of today's digital world.

This was reflected in the participant profile for the 2019 survey (Figure 1). Less than half (48%) of participants identified the type of data center they are most involved with as "on-premise core enterprise/private cloud." The remaining identified the type of data center they work in, manage or own as colocation (20%), High Performance Computing (15%), hyperscale/public cloud (10%) and remote or edge (8%).

Emerging Trend: 5G and Edge Computing

451 Research says 5G will be “the most impactful and difficult network upgrade ever faced by the telecom industry.” But the firm also emphasizes that telecom operators that successfully navigate this transition will create huge opportunities for themselves and their customers. According to a [global study by the firm](#), 98% of global telecom decision makers expect their organizations to be supporting 5G by 2021. With its high bandwidths and ultra-low latencies, 5G has the potential to accelerate the development of a host of digitally enabled innovations that increase the demand for, and amplify the benefits of, edge computing.

Centralization and Distribution

One of the questions posed in the original Data Center 2025 report was, “Will data centers cluster in regions with low energy costs and cool climates, or will proximity to users drive location decisions?” In other words, would compute be centralized or distributed?

The answer, of course, is turning out to be both. Cisco [projected a tripling in data volume](#) between 2016 and 2021, and no single approach could possibly handle that growth.

Since 2014, we’ve seen larger and larger cloud facilities being developed, creating a class of hyperscale facilities with distinct and innovative architectures. At the same time, more data is being generated and consumed at the network edge, forcing compute and storage closer to users and devices in the form of mini and micro data centers.

This trend toward edge computing will be accelerated by what may turn out to be the most significant technological advance of the second five years of the Data Center 2025 timeline: 5G (see sidebar).

Looking Ahead to 2025

Today we have two evolutions occurring simultaneously: one in technology itself, and one in the increasing segmentation of the industry, which has evolved from primarily on-premise, core-focused data centers to increasingly distributed and dynamic data center networks.

There clearly isn’t a single vision of Data Center 2025 that will apply across segments. Each of the main segments covered in this report—hyperscale, HPC, colocation, enterprise and edge—will evolve somewhat independently. But they will also become increasingly interdependent as they function together to meet the demands of the future.

While technologies such as servers, infrastructure systems and management tools have advanced somewhat predictably, the nature of what constitutes a data center has changed dramatically.

Powering the Data Center of the Future

More data ultimately means more power. In 2014, when the original Data Center 2025 survey was deployed, the industry was coming off an extended period in which energy efficiency was a major focus and significant progress had been made on that front.

According to the [Uptime Institute 2018 Global Data Center Survey](#), average data center PUEs went from 2.5 in 2007 to slightly above 1.5 in 2014. This essentially allowed more compute to be performed with less energy.

However, since 2014 PUEs have plateaued and the demand for capacity has continued to rise. As the 2019 survey was deployed, the industry was several years into a major building boom among colocation and hyperscale operators to bring new capacity online, driving overall industry energy consumption even higher. In the telecom industry, the transition to 5G could increase network energy consumption by up to 170%. Managing these growing energy

requirements will continue to be a challenge for the industry.

Renewables to the Rescue?

There were some significant shifts in expectations for some energy sources this year compared to the original Data Center 2025 survey. Compared to 2014, participants projected a lower percent of total data center energy to come from solar, natural gas, nuclear and wind (Figure 2).

Participants in the original survey projected 22% of data center power would come from solar and an additional 12% from wind by 2025. That’s a little more than one-third of data center power from these two renewable sources, which seemed like an unrealistic projection at the time.

This year’s numbers for solar and wind (13% and 8% respectively) seem more realistic, although still ambitious. We are likely not close to advancing at a pace that would put us on course to meet even those reduced projections, although precise numbers on the use of renewable energy in data centers are hard to

find. While solar and wind technologies continue to advance in terms of capacity and cost competitiveness, distribution and reliability challenges have limited their application in the data center.

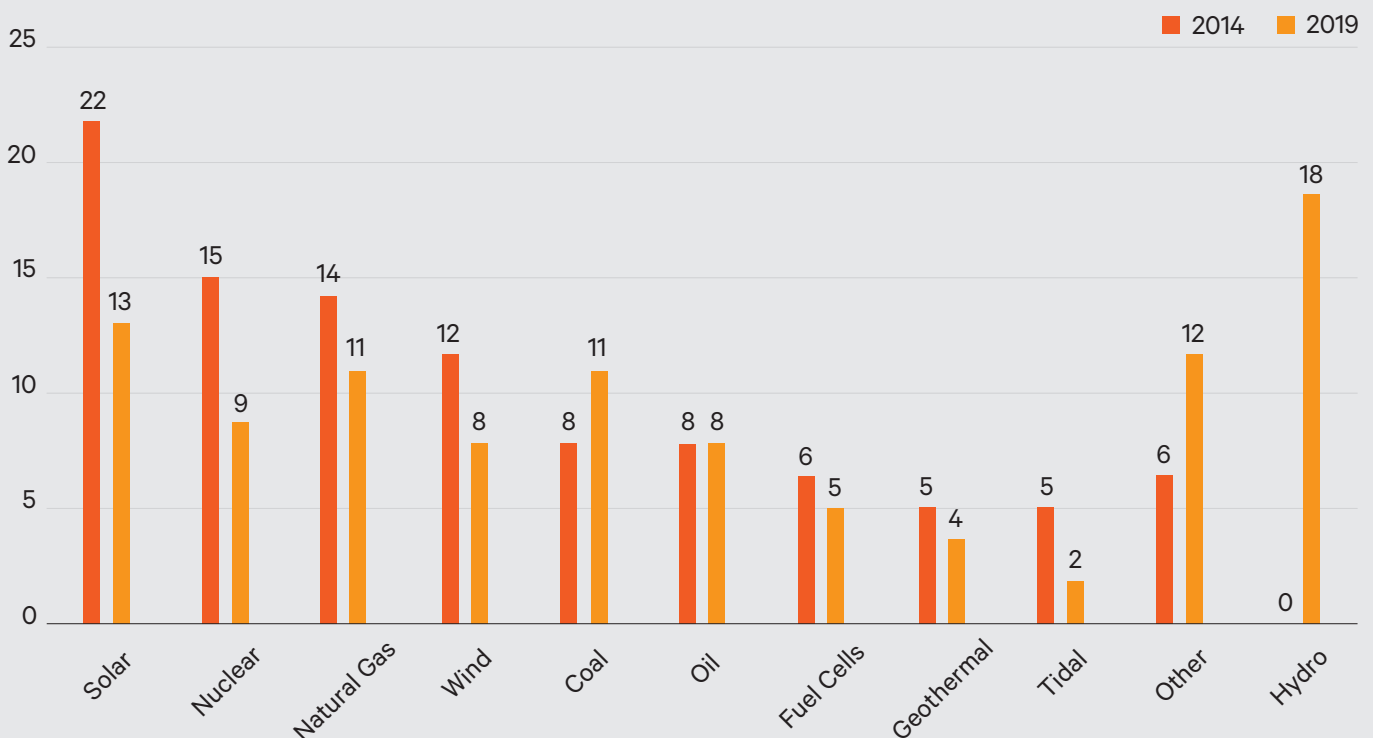
The more modest outlook for wind and solar was countered by increased expectations around hydroelectric power, another form of renewable energy. In fact, participants in this year’s survey expect hydro power to be the largest energy source for data centers in 2025.

Projections for energy sources were fairly consistent across regions, with a few exceptions. Expectations for solar were higher in Latin America (17%) than other regions. Latin America was also high on hydroelectric power, projecting 29% of data center power from this source in 2025.

Combining average responses for solar, wind and hydro power from Latin American participants creates the expectation of 56% of data center power coming from these three renewable sources in 2025. The next highest percentages for these three sources were

Electrical Energy Sources

Figure 2: Comparison of 2014 and 2019 Data Center 2025 survey results for the question, “In 2025, what percent of the electrical energy used by data centers do you expect to come from each of these sources?”



EMEA at 37% and China at 36%. APAC had the lowest projections for the three renewables at 31%.

Beginning the Transition

Large hyperscale and colocation providers are increasingly committing to transition to renewable energy, using renewable energy purchases and credits as a path to achieve their objectives. Equinix, for example, says it “covered approximately 90% of our global electricity consumption in 2018 with equivalent renewable energy purchases.” Likewise, Digital Realty “procured more than 1,100 GWh of above-baseline utility renewables in 2018.”

“While direct use of renewables such as wind and solar may be limited due to capacity and reliability concerns, we are seeing more data center operators entering into power purchasing agreements that include high percentages of renewables,” said Emiliano Cevenini, VP sales mobility & critical energy verticals for Vertiv in Europe, Middle East and Africa. “This shifts the reliability challenge to the distributor who is then responsible for meeting the agreed-to SLA. These financial incentives could ultimately drive greater reliability of renewables and lead to lower cost per kilowatt hour—and increased usage—as the costs of not meeting SLAs is minimized.”

Managing the Growing Demand for Compute

The primary challenge the industry faces as it moves closer to 2025 is meeting the growing demand for compute and storage. As noted previously, there is no one solution. From higher density equipment racks to continued investments in new hyperscale and colocation facilities to increased edge computing, a multi-faceted approach is required.

Large hyperscale and colocation providers are increasingly committing to transition to renewable energy, using renewable energy purchases and credits as a path to achieve their objectives.

Rack Density

Many data center professionals have been hearing about, but not experiencing, rising rack densities for years.

The impact of all of those warnings was exhibited in the original Data Center 2025 survey. Despite rack densities at a relatively stable 5-6 kW at the time, participants in the original survey expected densities to rise to an average of 55 kW by 2025. We are clearly not on a path that will take us anywhere near that projection.

But there are signs that we are approaching that point in some segments. This isn't showing up in industry averages because broad averages don't accurately reflect what is happening in these segments.

As the 2018 Uptime Institute Global Data Center Survey noted: “The high level of consolidation and the movement of workloads to public cloud has rendered the metric of average rack density less meaningful than it used to be.” The report goes on to point out that rack density has become more about extremes than averages, and here the move to higher density racks is unmistakable. In the Institute's 2017 survey, 9% of participants had average densities of 10 kW per rack or higher. In 2018, about one fifth had racks 30 kW or higher.

As Vertiv expert, Tony Gaunt, senior director for the colocation, cloud and financial services markets in Asia and India, noted, “The growth in AI, machine learning and gaming is driving the demand for high-density pods within many industries. These pods typically feature 3-8 racks with densities from 30-60 kW and will place new demands on power and cooling infrastructure that was sized to support a much lower average rack density across the facility.”

IT Utilization

If there is a great untapped resource within the current data center ecosystem, it is IT asset utilization. While utilization rates are difficult to determine without a detailed analysis, the best studies typically put the utilization rate in enterprise data centers at around 20%.

Yet, precisely because it's so hard to measure and

Emerging Trend:

Active Rear Door Cooling

Active rear door cooling has emerged as a high-efficiency solution for racks up to 50 kW. This approach uses the equipment rack as a containment system with rear-door chilled-water systems removing heat before the air leaves the rack. This results in a room-neutral design that does not require a hot-aisle/cold-aisle configuration. In addition to high efficiency, these systems offer easy installation and maintenance and take up very little floor space, adding just six inches to the footprint of the rack.



participants may have different definitions for utilization, many data center professionals may be unaware of exactly how low their actual utilization rates are. Many may think of IT utilization rates in the same way as UPS utilization, which generally ranges from 30% to 75%.

That's one theory for explaining the results from the original Data Center 2025 survey, where 72% of participants expected IT utilization rates to be at least 60% in 2025 (Figure 3). Now, five years later, with little apparent progress made outside of the growth of higher-utilization hyperscale and cloud facilities,

expectations have shifted down. Today, 57% expect IT utilization rates to reach at least 60% by 2025.

If this projection is to be realized it will almost certainly come via continued growth in hyperscale and cloud operators, which typically achieve significantly higher utilization rates based on the varying needs of multiple users.

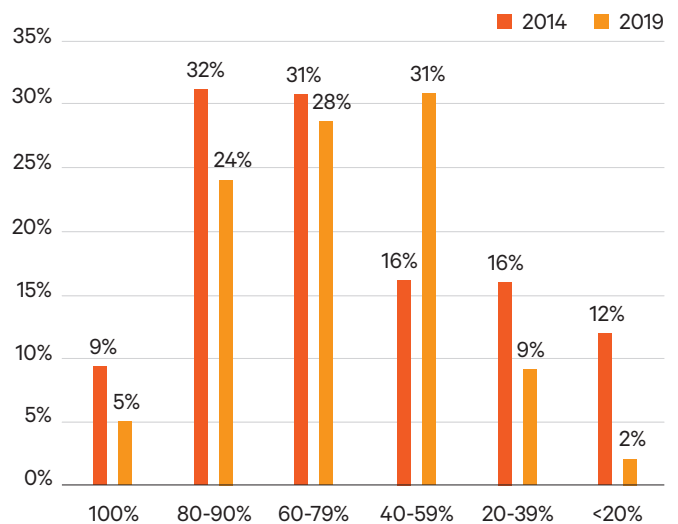
Across all data center types, survey participants almost unanimously expect IT utilization at the core to be above 20% (98%)—a significant increase from the 2014 survey (88%).

Participants in the hyperscale/private cloud and colocation segments were more pessimistic than those in the HPC or enterprise/private cloud segments, perhaps indicating greater awareness of current utilization rates. Slightly more than half of participants who identified their data center type as hyperscale/public cloud (51%) or colocation (54%) expect IT utilization rates to reach 60% by 2025. Sixty-four percent of participants who identified their data center type as HPC and 60% of those who identified as enterprise/private cloud expect it to reach that level.

Looking at the data regionally, participants in China

IT Utilization Rates (Core)

Figure 3: Comparison of 2014 and 2019 survey results for the question, "In 2025, what do you expect the average IT resource utilization rate to be at the network core?"



and Asia Pacific were more optimistic than other parts of the world with 58% and 61% respectively expecting utilization rates to reach at least 60%. The U.S./Canada was the most pessimistic at 50%.

“Higher utilization rates are baked into the value proposition of public cloud providers, but my experience is that utilization is a major concern across the industry,” said Peter Panfil, VP of global power sales for Vertiv. “Multi-tenant data centers, in particular, are challenged in this area because they don’t control the IT resources in their facilities. Nevertheless, there are specific actions they can take to increase utilization by 50-100%.”

Cloud Dependence

Just as in the previous two sections, participants in the new Data Center 2025 survey lowered their expectations when it comes to the percent of computing that would be in the cloud as we move closer to 2025, although in this case the trend seems less warranted.

Sixty-seven percent of participants in the 2014 survey expected at least 60% of compute capacity to be performed in the cloud by 2025 (Figure 4). In 2019, that number dropped to 60% potentially due to the expected increase in edge computing sites discussed later in this report.

This indicates the industry now has a better sense of which applications work best in the cloud. Overall, there is the expectation that the cloud will play a dominant but not exclusive role in the future data center ecosystem.

Not surprisingly, participants who identified their data center type as hyperscale/public cloud were the most optimistic, projecting an average of 69% of computing to be performed in the cloud by 2025.

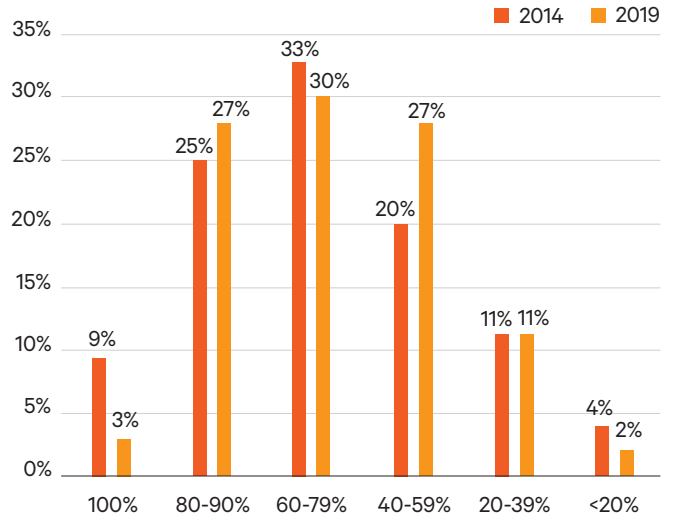
This was followed by colocation (64%), HPC (63%), and enterprise/private cloud (62%). Regionally, Asia Pacific had the highest expectations at 67% while EMEA had the lowest at 59%.

Managing the Workforce

The 2014 Data Center 2025 survey uncovered the potential for a significant brain drain in the industry.

Percent of Computing in the Cloud

Figure 4: Comparison of 2014 and 2019 survey results for the question, “In 2025, what percent of data center computing do you expect will be done in the cloud, rather than by in-house data centers?”



Only 56% of survey participants expected to be working in the industry in 2025, with the largest percentage leaving because of retirement (23%).

Five years in, it’s no surprise that those numbers are better in this version of the survey as we are dealing with a shorter window of time. Nevertheless, there are some trends that may raise concern.

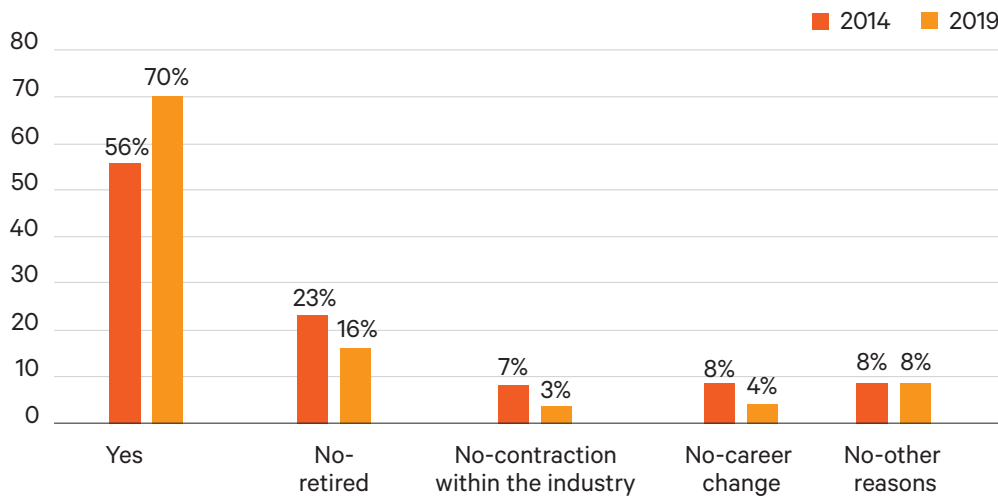
Seventy percent of participants expect to be working in the industry in 2025 with 16% expecting to be retired in five years (Figure 5). Only 3% expect contraction to affect their employment in this year’s survey compared to 7% in 2014, reflecting the strong growth the industry has experienced in the last five years.

With the labor market already tight in some key markets, the potential to lose 16% of the workforce to retirements could impede the ability of organizations to adapt to changing requirements. On the other hand, it is also driving more operators to normalized designs and the application of rapid deployment configurations that require less intellectual capital to deploy and support.

Impacts of labor shortages could be most severe in

Data Center Employment in 2025

Figure 5: Comparison of 2014 and 2019 survey results for the question, “Do you plan to be employed in the data center industry in 2025?”



the U.S. and Canada, which showed a markedly higher percentage of data center professionals expecting to be retired by 2025 at 33%. China, reflecting the relative maturity of the data center industry in that country, had the lowest projected retirement rates at 8%.

An unexpected benefit of workforce challenges is the drive to normalized designs and the application of rapid deployment configurations that require less intellectual capital to deploy and support.

“Shifting workloads to the cloud has helped mitigate the impact of labor challenges in the short term, but as hybrid IT and edge computing continue to grow the skills gap is becoming a more serious issue,” said Robert Linsdell, managing director for Vertiv in Australia and New Zealand. “These changes require new skills that may not exist in the legacy workforce and this could hamper the ability of some IT organizations to support their businesses as they continue to evolve.”

The Rise of Edge Computing

The network edge is not new, but it is being re-purposed and expanded. Over the last several years, “edge computing” has become one of the most talked about trends in IT, and for good reason. Nearly every industry is recognizing the limitations of

supporting users and emerging technologies through centralized IT infrastructures and is pushing storage and computing closer to users and devices.

Much of the increase in data generation will come from mobile sensors and must be transmitted on wireless or mobile networks rather than wired Internet connections, putting a strain on mobile network infrastructure. Mobile IP traffic was projected to increase seven-fold from 2016 to 2021, double the pace of growth in fixed IP traffic. The changes in the compute and storage infrastructure required to support the smart and connected future, particularly at the local level, will be profound.

The magnitude of this impact is exhibited in response to the 2019 Data Center 2025 question, “How many computing sites is your company supporting today, and how many do you expect by 2025?” Of participants who have edge sites today or expect to have edge sites in 2025, more than half (53%) expect the number of edge sites they support to grow by at least 100% with 20% expecting a 400% or more increase (Figure 6). Yet, even this doesn’t fully capture the magnitude of the change.

The challenge the industry faces may be put in better perspective when considering the total number of edge sites today and in 2025. For the 494 qualified respondents to this question, the total number

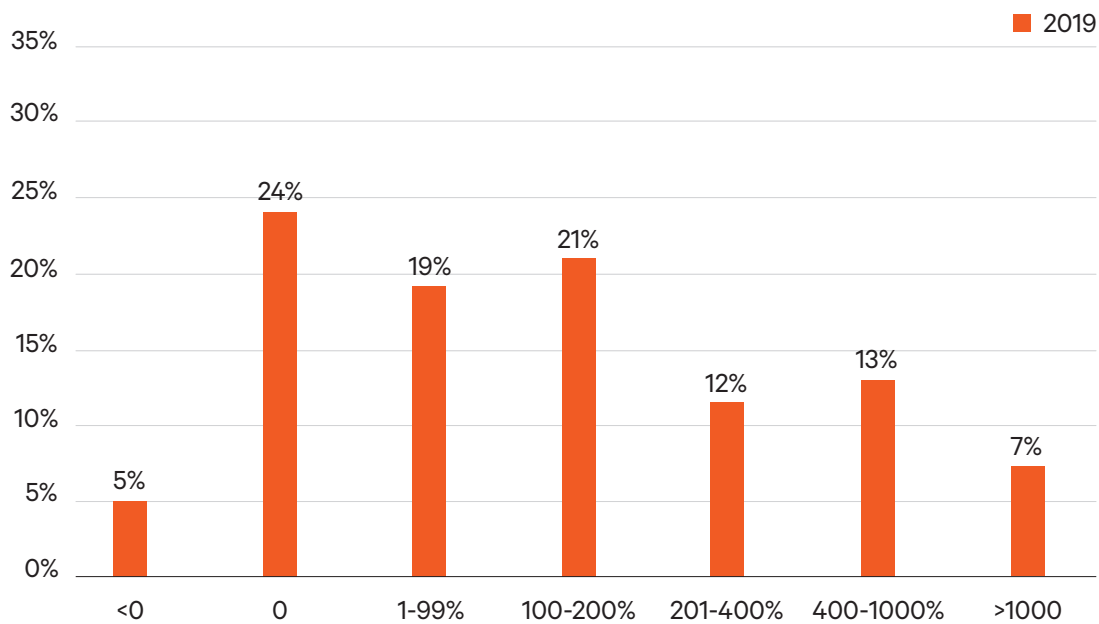
of edge sites supported is expected to grow from 128,233 today to 418,803 in 2025—a 226% increase. The challenge of configuring, deploying and managing this growing network of sites has the potential to strain IT organizations beyond their limit unless standardized configuration options and remote management tools are employed to streamline processes and minimize the need for on-site technical support.

Understanding Current Edge Use Cases

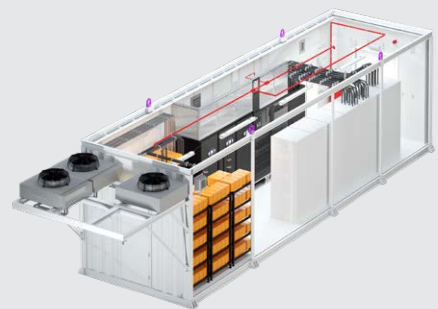
One of the challenges organizations face as they scale up the edge of their networks is the wide variety of potential emerging use cases. How can they create infrastructure tailored to their needs while streamlining deployment and enabling standardization?

Growth in Edge Computing Sites

Figure 6: Percentage growth in edge sites for participants with edge sites today or those who plan to have edge sites in 2025.



Emerging Trend: Modular Prefabricated Data Centers



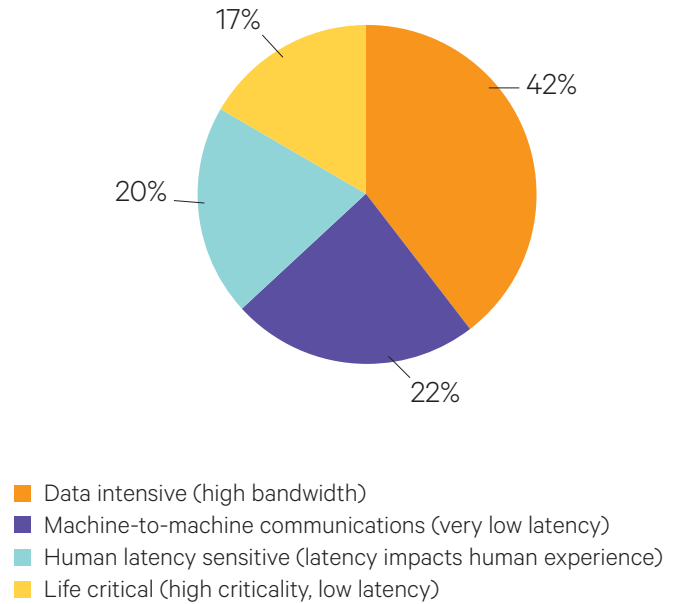
Tightly integrated and pre-manufactured data center systems have long been used to speed the deployment and improve the management of small and remote data centers. In the last several years, this same approach has been applied to large, freestanding data centers. By designing and integrating all components, including the data center shell, in the factory and then shipping the facility in modules that are assembled on site, the traditional stick-build process has been streamlined. This allows operators to meet rising capacity demands faster while achieving enhanced scalability and efficiency.

To simplify and accelerate the deployment of edge infrastructure, Vertiv classified the most impactful edge use cases into four categories:

- Data Intensive:** Use cases where the amount of data makes it impractical to transfer over the network directly to the cloud, or from the cloud to the point-of-use, because of data volume, cost or bandwidth issues. Examples of Data Intensive use cases include smart factories, smart cities, high-definition content delivery and virtual reality. In the Data Center 2025 survey, Data Intensive was identified as the primary data requirement of participants edge applications in 2025 by 42% of respondents (Figure 7). Expectations for supporting Data Intensive edge applications were highest in Latin America (54%) followed by the U.S./Canada and Asia Pacific (45%). They were lowest in China (26%).
- Human-Latency Sensitive:** This category includes use cases where services are optimized for human consumption or to improve the human experience with technology-enabled services. Examples include augmented reality, smart retail and natural language processing. Twenty percent of Data Center 2025 participants identified Human-Latency Sensitive as their primary edge data requirement in 2025, with China and the U.S./Canada having the highest expectations at 28% followed closely by EMEA at 25%. Asia Pacific and Latin America had significantly lower expectations for Human-Latency Sensitive edge applications at 13% and 12% respectively.
- Machine-to-Machine Latency Sensitive:** This category covers use cases where services are optimized for machine-to-machine consumption. Because machines can process data so fast low latency communication is required to support these use cases, which include arbitrage, smart security and smart grid. Twenty-two percent of Data Center 2025 participants identified Machine-to-Machine Latency Sensitive as their

Among respondents, the total number of edge sites supported is expected to grow by 226%.

Figure 7: 2019 Data Center 2025 response to the question, “What will be the primary data requirement for your edge applications in 2025?”



primary edge data requirement in 2025. EMEA and China had the highest expectations for machine-to-machine data requirements (24% and 25%) while the U.S./Canada had the lowest expectations (19%).

- Life Critical:** This category encompasses use cases that directly impact human health and safety. Probably the best examples of the Life Critical Archetype are autonomous vehicles and digital healthcare. Seventeen percent of Data Center 2025 participants identified Life Critical as their primary edge data requirement in 2025. Asia Pacific and China had the highest expectations at 21%, followed by EMEA at 18%. The U.S./Canada had the lowest expectations, with just 8% of participants expecting their primary edge data requirement to support Life Critical use cases.

For more information on edge use cases, see the Vertiv White Paper, [Defining Four Edge Archetypes and Their Technology Requirements](#).

Emerging Use Cases and 5G

5G will play an important role in providing the high bandwidth and low latency required to support many

emerging edge use cases. Data Center 2025 participants felt 5G would be most impactful in enabling smart cities, smart security, smart transportation and connected/autonomous vehicles (Figure 8).

Smart cities received the highest percentage of responses in China (78%) and Latin America (72%) while smart security received the highest percentage of responses in Latin America (71%) and the U.S./Canada (68%). China also had the strongest support for smart transportation (80%), virtual reality (57%) and augmented reality (46%). EMEA had below-average expectations for all applications presented. “It’s hard to underestimate the impact of 5G and edge computing,” said Martin Olsen, global VP of edge and integrated solutions at Vertiv. “Broadly it will enable some of the biggest innovations coming in the next five years. Specific to our industry, it will require data center operators to rethink their fundamental network architecture as they transition from core-heavy architecture to architectures that are more balanced between edge and core.”

Critical Infrastructure Technology and Management

Critical infrastructure technology—the power and thermal management systems that enable data centers of all sizes to deliver uninterrupted services to

users—in many ways looks much like it did in 2014.

However, these systems have benefited from significant innovations that have enabled them to effectively adapt to the changing needs of today’s facilities, from the largest hyperscale data centers to remote edge sites.

Most notably, they have added intelligence to enable machine-to-machine communication and simplify remote management; become more scalable to adapt to changing capacity demands; and are increasingly integrated off-site to speed deployment.

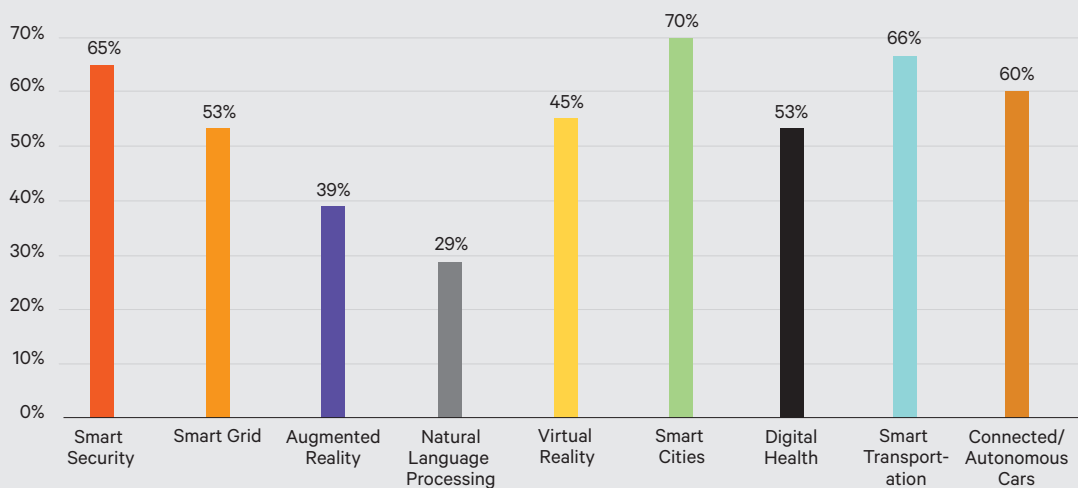
As a result of these advances, participants in the 2019 Data Center 2025 survey continue to show confidence in the ability of these core technologies to meet the needs of the data center ecosystem of the future.

Thermal Management

Perhaps no data center system has changed as dramatically in the five years since 2014 as thermal management. The industry has seen a large-scale shift to economization driven by hyperscale operators and colocation providers, while simultaneously driving heat removal closer to servers through rear door and liquid cooling systems designed to support the high-density racks common in HPC facilities.

Applications Enabled by 5G

Figure 8: 2019 Data Center 2025 responses to the question, “Thinking broadly of the world in 2025, which applications do you believe will require 5G technology to be fully functional? Your responses can include, but are not limited to, the specific requirements of your job or the business you work for.”



In our 2014 survey, participants expected cooled air, delivered by precision cooling systems, to account for 41% of data center cooling. This was followed by ambient or outside air at 20% and liquid or immersive cooling at 20%.

For 2019, we have re-labeled the responses to this question to better reflect the current state of the technology. One of the major developments in this category occurred right around the time our original survey was launched—the integration of economization into precision cooling systems, blurring the traditional line between free cooling and precision cooling.

These integrated perimeter cooling systems have received broad market acceptance and likely contributed to the strong confidence participants had in the future of mechanical cooling to carry more of the cooling load in the future. They expect 42% of future cooling requirements to be met by mechanical cooling systems (Figure 9). Liquid cooling and outside air also saw slight growth from 20% in 2014 to 22% in 2019, likely driven by the more extreme rack densities being observed today, as discussed earlier.

However, it's also important to note that due to the different requirements of various types of data

Emerging Trend: Application-Driven Edge Infrastructure

As key edge use cases continue to mature, infrastructure providers are working with other technology companies to use the broad edge classifications described in this section as the foundation for fully integrated, purpose-built infrastructure systems that can then be easily configured to the specific requirements of an application. These “off-the-shelf” infrastructure solutions will be an important component in enabling enterprises and telecommunications providers to meet the demand for edge services.



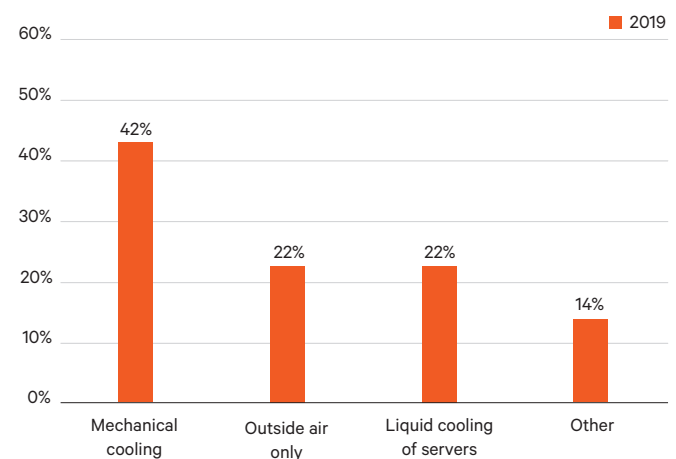
Mobile IP traffic was projected to increase seven-fold from 2016 to 2021, double the pace of growth in fixed IP traffic.

centers today, all of these technologies likely have a place in the data center ecosystem of 2025. This becomes clear when we analyze responses by data center types. Participants who define their data center type as colocation or enterprise/private cloud assigned the highest percentages to mechanical cooling (48 and 43% respectively). Liquid cooling was most popular with participants who define their data center type as hyperscale/public cloud (25%) and HPC (24%).

Participants who defined their facility type as hyperscale/public cloud also had the highest percentage for outside air only at 25%. “A one-size-fits-all approach to thermal management has never been effective,” said Vertiv expert, Steve Madara, VP of global cooling sales for Vertiv. “It’s always been necessary to tailor thermal management to the profile and environment of a particular facility. The difference today is the range of solutions and configurations that are available. Thermal engineers

Thermal Management Strategies

Figure 9: 2019 Data Center 2025 responses to the question, “In 2025, how will computing facilities remove the heat that servers produce?”



have a full suite of solutions at their disposal to tailor efficient, effective and intelligent thermal management systems to the specific requirements for density, efficiency, availability and management.”

Maintaining Availability

In terms of protecting the availability of data center services, AC UPS systems continued to be the strategy of choice for survey participants, growing from 30% in 2014 to 47% in 2019 (Figure 10).

As with thermal management systems, AC UPS systems continue to make advances with efficiencies in some operating modes approaching 99%. These systems have also added intelligence that increases their flexibility and maintainability. These capabilities make them less likely, in the minds of participants, to be displaced by competing technologies.

Somewhat surprisingly, considering the virtualization capabilities of cloud computing and the interconnection capabilities now offered by colocation providers, software failover saw a significant drop in confidence between the two surveys, as did momentary duty inverters. The reduced confidence in software failover could reflect challenges some early adopters have faced in implementing this strategy. It may also reflect the growing realization of the transmission bandwidth and cloud resources required to support software failover.

APAC (16%), EMEA and the U.S./Canada (both 15%) were the most optimistic regions in regard to software failover.

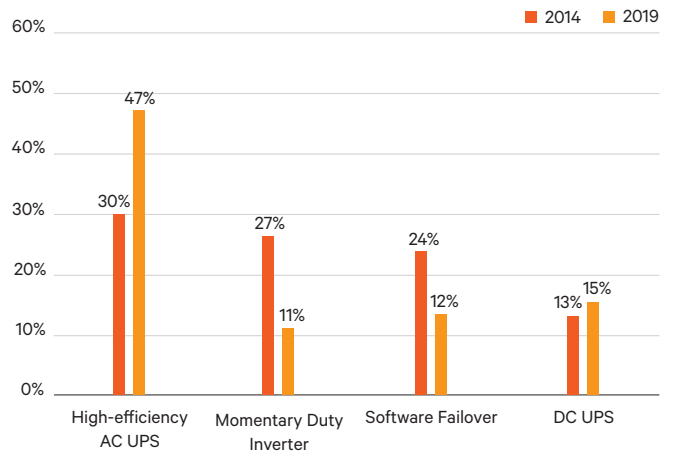
Infrastructure Management

While data centers become more diverse in size, functionality, density and architecture, one need that remains consistent across them all is the desire for increased visibility and automation. 2014 survey participants were optimistic they would not only get the visibility they needed (29%) but that their data centers would become self-healing (43%) and self-optimizing (25%).

In 2019, expectations for visibility have remained consistent as DCIM has evolved from an emerging to mainstream solution for data center management. The major shift in this year’s results compared to 2014

Backup Power Strategies

Figure 10: Comparison of 2014 and 2019 Data Center 2025 responses to the question, “What will be the primary means of delivering backup to data center equipment in 2025?”



is increased expectations around self-optimizing, which rose from 25% in 2014 to 39% in 2019. This increase came at the expense of reduced expectations around self-healing (Figure 11).

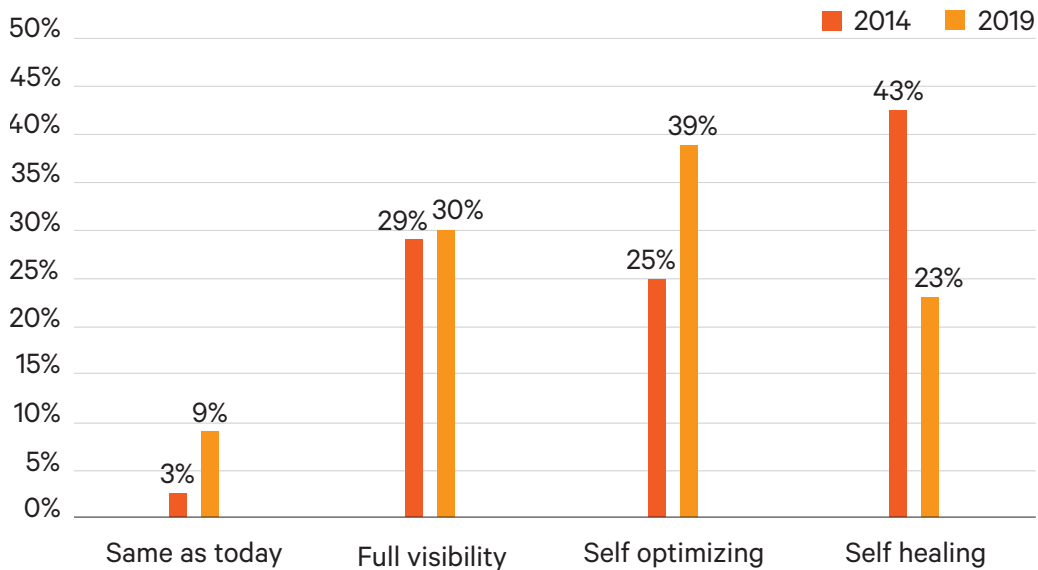
It’s likely that participants are seeing the ability of infrastructure systems today, guided by advances in intelligent controls, to self-optimize while the prospect of self-healing now seems like it may be beyond the 2025 horizon. The appeal of these strategies is enhanced by the reductions in operating resources many operators have experienced.

In general, the U.S./Canada was the most pessimistic about the future of data center management with one in five participants expecting it to be the same in 2025 as today. Stark differences by region also emerged in expectations around self-healing and self-optimizing data centers.

Fifty-three percent of participants from China expect self-healing data centers in 2025 while only 8% expect self-optimizing data centers. The opposite is the case in Latin America where 54% of participants anticipate self-optimizing data centers and 19% self-healing. In EMEA, the split is 43% choosing self-optimizing and 15% self-healing while the U.S./Canada is 35% self-optimizing and 14% self-healing.

Future of Data Center Management

Figure 11: Comparison of 2014 and 2019 Data Center 2025 responses to the question, “Which of the following best describes what you think data center management and control will be like in 2025?”



Emerging Trend: Lithium-Ion Batteries

While AC UPS systems will likely continue to be the primary method of power backup for the foreseeable future, the batteries these systems depend on to ride through short outages are undergoing an evolution. Increasingly users are turning away from traditional VRLA batteries to industrial lithium-ion batteries that provide longer lifecycles and reduced cooling costs. Based on these advantages, and increasingly competitive pricing, lithium-ion batteries could displace VRLA batteries in a majority of data centers by 2025.

According to Vertiv expert Patrick Quirk, VP and general manager of IT systems, “The ability of infrastructure systems to self-optimize is available today. Moving forward, we’ll see greater interconnectivity across critical infrastructure and IT systems and an increased use of machine learning, which will enable entire facilities to self-optimize based on workload and/or parametric driven metrics.”

In terms of protecting the availability of data center services, AC UPS systems continued to be the strategy of choice for survey participants, growing from 30% in 2014 to 47% in 2019.

A Revised View of 2025

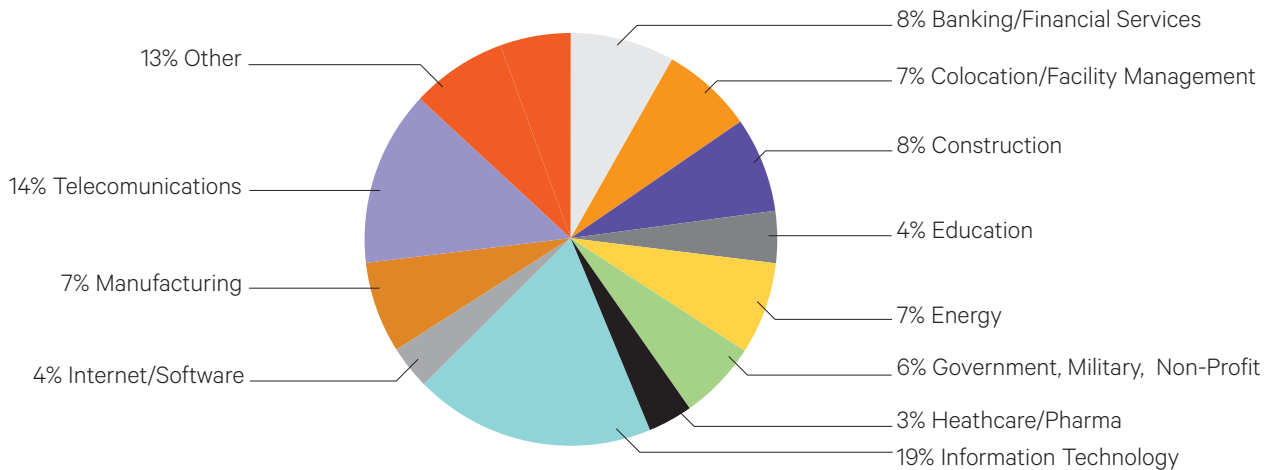
Five years into the Data Center 2025 initiative, the picture of what the data center of the future will look like is becoming clearer.

First, it will not, as has been made clear throughout this report, look the same across different types of facilities. Enterprise, HPC, edge, hyperscale and colocation facilities will have markedly different characteristics dictated by their role in a dynamic, interconnected network capable of handling the huge volumes of data being consumed and generated. The changing nature of each of these types of facilities may not be as dramatic as many in the original Data Center 2025 projected, but they are happening consistently and incrementally.

It's also clear we are at the front end of a significant shift not necessarily away from centralized computing but toward edge computing. We expect managing the growth in edge computing sites to be the single biggest challenge—and opportunity—data center professionals face between now and 2025.

Participant Profile

The 2019 Data Center 2025 survey included responses from more than 800 industry professionals with a variety of roles supporting the data center. Participants were fairly evenly distributed across a range of industries with the largest representation from Information Technology.



Emerging Trend: AI and Machine Learning



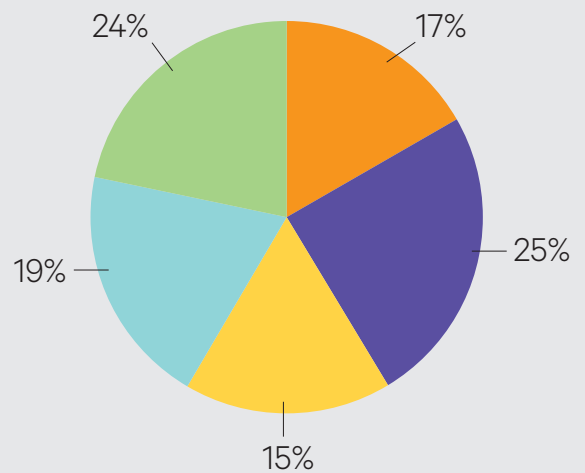
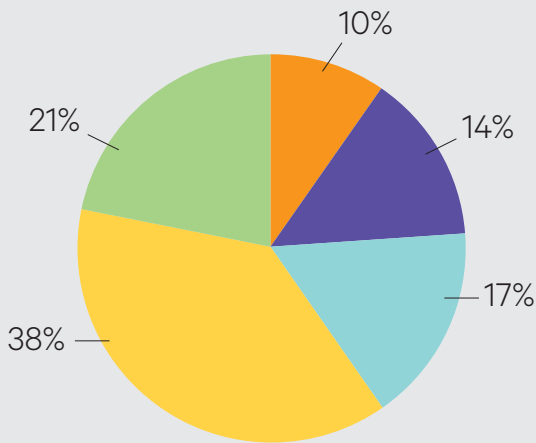
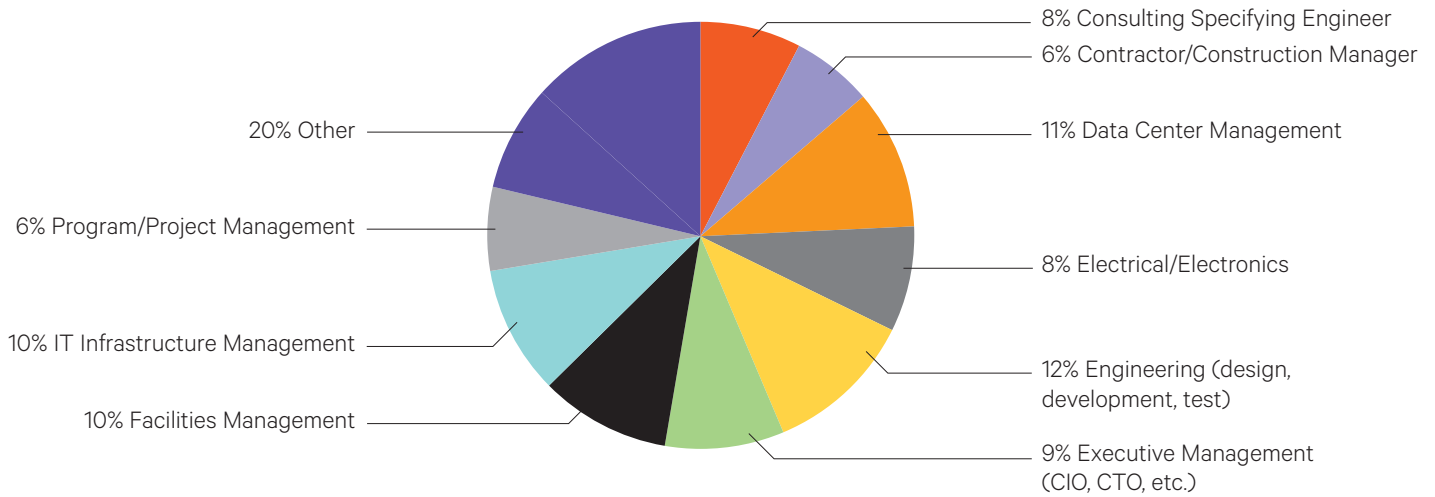
Data analytics is not only driving demand for increased computing within data center facilities, it's also emerging as a new tool for predicting failure and improving performance. Vertiv shared [research](#) that demonstrated the viability of applying machine learning to large volumes of historical battery data to accurately identify battery strings and units determined to be at risk, but otherwise not in an alarm state within the dimensions of manufacturer, model, age, voltage, temperature and ohmic readings. The results demonstrate the potential to use machine learning to increase battery reliability and extend life.

Participants also held a variety of positions with varying responsibilities for data center design and management, including data center, facilities and IT infrastructure management.

Participants were asked to define their dependence on data centers based on the impact of an outage to their business. Thirty-eight percent identified their data center as “critical to their

business” while an additional 21% said the business is “totally dependent” on the data center. Only 10% indicated their business could “operate for limited periods without computing.”

Geographically, participants represented all major world regions, with the highest numbers coming from Latin America and Asia Pacific.



- Our business can operate for limited periods without computing
- Our business relies on our data centers primarily for “back office” operations
- Downtime in our data centers hurts, but does not cripple our business
- Our data centers are critical to our ability to serve customers
- Our business is toally dependent on data center operations (cloud, colocation provider)

- US and Canada
- Latin and Central America
- China
- EMEA
- Asia Pacific

About Vertiv

Vertiv brings together hardware, software, analytics and ongoing services to ensure its customers' vital applications run continuously, perform optimally and grow with their business needs. Vertiv solves the most important challenges facing today's data centers, communication networks and commercial and industrial facilities with a portfolio of power, cooling and IT infrastructure solutions and services that extends from the cloud to the edge of the network. Headquartered in Columbus, Ohio, USA, Vertiv employs around 20,000 people and does business in more than 130 countries. For more information, and for the latest news and content from Vertiv, visit [Vertiv.com](https://www.vertiv.com).

