

UI Intelligence report 43

The people challenge: Global data center staffing forecast 2021-2025

Lead Author

Rhonda Ascierio, Vice President of Research, Uptime Institute

Contributing Authors

Chris Brown, Chief Technical Officer, Uptime Institute

Todd Traver, Vice President for IT Optimization & Strategy, Uptime Institute

Fred Dickerman, Senior Vice President of Management Services – Europe, Middle East and Africa, Uptime Institute

Rich Van Loo, Vice President of Facility Management Services, Uptime Institute

As the data center build-out continues across the globe, many more people will be needed to design, build and operate this critical infrastructure. The demand will exacerbate staffing shortages – but by how much? Uptime Institute has completed the first forecast of global data center workforce needs – by region, by data center type, and by education requirements. This report summarizes and analyzes its findings.



This Uptime Institute Intelligence report includes:

Executive Summary	3
Key findings	3
Acknowledgments	4
Introduction	4
About this research	5
Job domains	6
The rising tide of demand	7
Operations staff needed most	8
Regions to grow unevenly	8
Cloud versus enterprise staffing needs	10
Education requirements: Employers become more flexible	12
Requirements for technical staff	13
Uncertainty factors: What might change staffing demand?	14
Conclusions	18
Appendix: Methodology	19
Assumptions	22
Job domain definitions	25
Data center market size and forecast	25
Data center types and size ranges	26
Appendix: Additional reading for Uptime Institute members	27
About the authors	28
About Uptime Institute Intelligence	29
About Uptime Institute	29

EXECUTIVE SUMMARY

In every region of the world, data center capacity is being dramatically expanded, in a build-out of historic proportions. That requires people – and the availability (or lack) of specialist staff will be an increasing concern for all types of data centers, from mega-growth hyperscales to small, private enterprise facilities. Greater investment, more training and more creative approaches to employment will be needed.

Although data center recruitment needs are expected to rise steadily to 2025, the growth in demand does not need to represent a crisis. Individual employers can take steps to address the issue, and the sector can act together to raise the profile of opportunities and to improve recruitment and training.

KEY FINDINGS

- Global data center staff requirements are forecast to grow globally from about 2.0 million full-time employee equivalents in 2019 to nearly 2.3 million in 2025. This estimate covers more than 230 specialist job roles for different types and sizes of data centers, with varying criticality requirements, from design through operation.
- New staff will be needed across all geographic regions. Most demand will be in the Asia Pacific region, followed by North America and the Europe, Middle East and Africa regions.
- In the large and mature data center markets of the US and Western Europe, there is concern that many employees are due to retire around the same time, causing an additional surge in demand, particularly for senior roles.
- Globally, demand growth will come from cloud providers, internet giants and colocation data centers. Cloud/internet giants' data centers – those either owned or leased by major public cloud/internet companies – will require the most staff, by a significant margin.
- Enterprise data centers will continue to employ a large number of staff, but the number will decline, eventually being surpassed by cloud/internet giants after 2025.
- Most of the workforce will require a university/college or technical trade school degree or – critically – equivalent experience. More on-the-job training and sector-specific education will be key for the sector to meet future talent demand.
- Technical staff are notoriously difficult to recruit in data centers. Mechanical and electrical engineers in strategy and operations roles and all types of controls and monitoring employees are among the technical staff that will be increasingly needed through (at least) 2025.
- Innovations in artificial intelligence and automation are unlikely to flatten or reduce staff needs before 2025.

Acknowledgments

We would like to thank all contributors to this research over the past two years of its development. They include (in alphabetical order): Jeremy Brown, VP Critical Facilities, First Citizens Bank; Stephen Christie, Data Modelling & Design Ltd.; Bruno Fery, Head of Datacentre Services, EBRC; Bob Guski, Senior Manager, Data Center Operations, Target; Tony Jacob, VP Design and Construction, Interxion/Digital Realty; Sean McGinnis, Data Center and Colocations Platform Owner, Biogen; Tomas Rahkonen, Research Director, Distributed Data Centers, Uptime Institute; Mark Rouzan, Director of Operations, T5 Data Centers; Torrey Searles, Director of Talent Acquisition and Development, T5 Data Centers; James Wilman, CEO, Future-tech; plus the many others who are bound to remain anonymous due to confidentiality clauses. Your input has been invaluable.

Introduction

Digital growth continues to fuel strong new demand for data center capacity, and it shows no signs of slowing. The ongoing build-out of new data centers and networks is being driven largely by internet services, cloud, hosted and other as-a-service workloads, as more enterprises seek to outsource more of their IT. Across the board, the scale of capacity growth is stretching the critical infrastructure sector's talent supply.

Our research shows that the proportion of data center owners or operators globally that are having difficulty finding qualified candidates for open jobs rose to 50% in 2020. While there is hope that new technologies to manage and operate facilities will reduce staff burdens over time, Uptime Institute expects that their effect will be limited, at least until 2025.

There is now growing consensus across the sector that investment and concerted action will be needed to ensure a steady supply of talent. A

Data center staffing challenges: By the numbers

50%

The proportion of IT and data center managers who said they were having difficulty finding qualified candidates to fill open positions in 2020.

38% The proportion in 2018.

77%

The proportion of IT and data center managers who said they expect the increase in remote working prompted by the pandemic to still be in place in two to three years.

33% The proportion who think they'll have less operational staff coverage longer term.

34%

The proportion of IT and data center managers who said in 2020 they believe artificial intelligence will reduce their data center operation staffing levels in the next five years.

29% The proportion in 2019.

43% The proportion who said it will happen, but not in the next five years (it will take longer).

Sources: Uptime Institute Annual Global Survey of IT and Data Center Managers 2018 (n=202) and 2020 (n=218); Uptime Institute Survey – Long-term Impact of COVID-19 on Data Centers, July 2020 (n=277); Uptime Institute Annual Global Survey of IT and Data Center Managers 2020 (n=435)

UptimeInstitute® | INTELLIGENCE

lack of available staff is an increasing threat for all types of data centers, from mega-growth hyperscales to small, private enterprise facilities.

In some data center markets, such as the US and Western Europe, there is also concern of a “silver tsunami,” with a wave of retirements expected in the coming decade. In these markets, the sector is rallying to raise its profile among job seekers, educators and governments with new collective initiatives.

Globally, the biggest employers are investing in more training and education, not just by developing internal programs but also by working with universities/colleges and technical schools. The growing, long-term requirement for more trained people has caught the attention of private equity and other investors. More are backing data center facilities management suppliers, which offer services that can help overcome skills shortages.

Until now, the size and scope of the staffing challenge has not been known. This Uptime Institute research represents the first quantitative assessment of staff requirements of the global data center sector: we estimate the number of specialist data center infrastructure staff required, from designing and building facilities through ongoing operations, to 2025 – globally, by region, by data center type (cloud/ internet giants, colo and enterprise), and by minimum education requirements.

About this research

In a study of this kind, a large number of assumptions and methodological choices must be made. These are listed and discussed in **Appendix: Methodology**. But for immediate clarity, we make the following points:

- Estimates are based on what is required for a well-run, fully staffed data center. Many sites will operate less optimally, at least at times.
- The job roles discussed and estimated here are for full-time equivalent positions, based on a 40-hour workweek. In other words, a half-time role at two data centers is counted as one position. We include both fixed (on staff) and variable (contract) employees in our estimates.
- The data in this study is derived in part from a detailed proprietary bottom-up model of the number of data centers existing and planned from 2019 to 2025. Estimates of staff to design, build and operate these data centers, of varying types and sizes, have been calculated based on industry input and extensive expert advice.
- The jobs covered in this study span the life cycle of operational data centers, from design through build and operation, but excluding decommissioning.

By quantifying the challenge, our hope is that this research will help raise awareness of data center employment demand and job opportunities, including for:

- **Education institutions** - Being able to demonstrate job demand will help institutions devise and secure funding for data center-related curricula.
- **Job seekers** - A quantifiable measure of the strong demand in the coming years and the minimum education required will help attract more job seekers.
- **Governments and economic development planners** - Funding for new curricula and incentive packages for proposed new data centers often require an estimate of how many new jobs will be created, which – as our model shows – goes far beyond the staff working on-site in data centers.
- **Employers** - By understanding demand, employers may prepare to fund new training and skills-development programs.

Job domains

It takes many people to plan, design, build and operate a data center – sometimes hundreds. In total, we estimated the FTE (full-time employee) requirements of each data center against Uptime Institute's taxonomy of job roles. This has more than 230 data center-specific job roles, which we group into domains (key functional areas) and map onto five data center life cycle stages. Note that in smaller data centers, one person may carry out multiple job roles – for this reason, we have focused on roles (primary areas of responsibility) and not jobs.

Estimating employment needs can be complex. In large organizations, multiple people may have the same job role. In smaller organizations, a single person may have multiple job roles (will “wear many hats”) – we have assigned a portion of their work time to each applicable job role. But some job roles are required only periodically or at certain stages of a data center's life and are, therefore, calculated as a certain number of hours per year, over a limited duration. As such, our forecast is for FTE equivalents and not necessarily for FTEs (i.e., number of people). This was a necessary trade-off for better accuracy.

In order to estimate requirements, each job role is organized at the highest level into one of nine job domains that represent the various activities or responsibilities required during the life cycle of a facility. We define a job domain as a multidisciplinary group of staff that share a common function. (Note that while domains are largely composed of multidisciplinary teams, some staff are also interdisciplinary in how they operate.)

DATA CENTER JOB DOMAINS

1. Business Support
2. Strategy
3. Design
4. Construction
5. Operations Engineering
6. IT Hardware
7. Network and Connectivity
8. Controls and Monitoring
9. Operations

Table 1 shows the nine job domains mapped against the five broad life cycle stages of an operational facility. For descriptions of each job domain, see **Appendix: Methodology**.

Table 1. Data center staff: Functional domains					
Data centers employ multidisciplinary teams of staff that share a common function (i.e., job domain) spanning the different stages of a data center's life cycle.					
Job domain	Data center life cycle stage				
	Strategic requirements	Design	Construction (build)	Commissioning	Ongoing operations
Business Support	◆	◆		◆	◆
Strategy	◆	◆			
Design		◆	◆	◆	
Construction			◆	◆	
Operations Engineering	◆	◆		◆	◆
IT Hardware		◆			◆
Network and Connectivity	◆	◆		◆	◆
Controls and Monitoring	◆	◆		◆	◆
Operations					◆

More than 230 individual job roles

Source: Uptime Institute 2021 UptimeInstitute® | INTELLIGENCE

The rising tide of demand

Staff requirements for the global data center infrastructure sector is forecast to rise from about 2.0 million FTE equivalents in 2019 to nearly 2.3 million in 2025 (see Figure 1). This represents growth in more than 230 specialist job roles to design, build and operate different types and sizes of data centers.

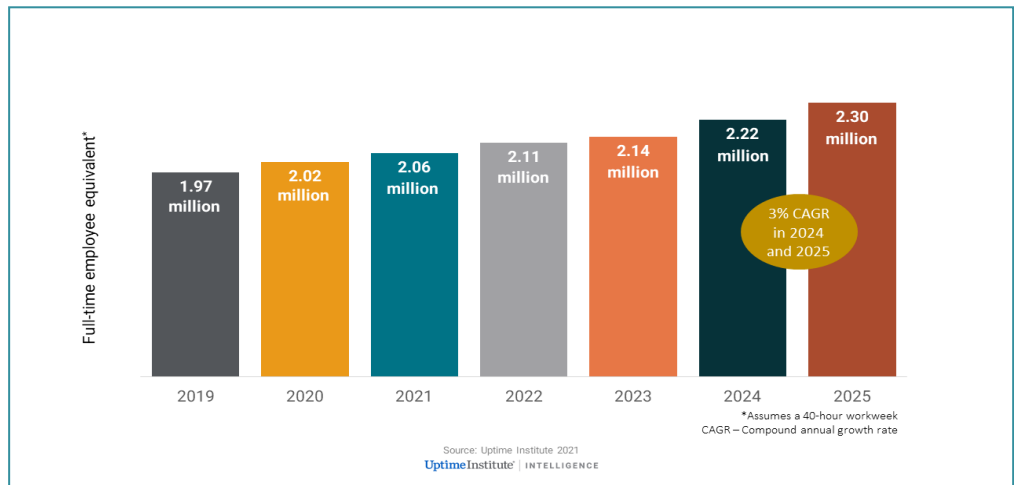


Figure 1. Global data center staff requirement projections

Although this represents only 2% compound annual growth rate (CAGR) over six years, the rising demand will still present some challenges. First, this figure represents new jobs, not net demand, which may be increased by other factors, such as the retirement of large number of existing workers. Second, the demand is unevenly spread, with the greatest needs clustered in certain regions.

Year-over-year CAGR is set to accelerate from about 2% through 2023 to 3% in 2024 and 2025.

Operations staff needed most

The bulk of staff will be needed for ongoing data center operations, with a smaller proportion required for the design and build of new or retrofitted facilities. This also presents challenges: Unlike many planning and design staff, some in the operations teams need to be at or near the data centers. These may be in areas where there are currently few skilled staff (many hyperscale data centers are sited in remote regions, where availability of power and cooling are the main priorities). Figure 2 shows the requirements for all operations staff (comprising several job domains) versus those for design and construction.




Figure 2. Global data center staff requirement projections: Operations versus Design plus Construction

Regions to grow unevenly

Staffing requirements across all major global regions will grow during the forecasted period, as shown in Figure 3. Most demand will be in the Asia-Pacific region, driven by expected cloud/internet giants and colo data center capacity growth in China, parts of Southeast Asia, Australia and elsewhere.




Figure 3. Data center staff requirement projections by region

Demand in North America and the Europe, Middle East and Africa regions is expected to be roughly similar, with notable growth forecast. In Latin America, growth is expected to be driven by several markets, including Brazil, Chile, Columbia and Mexico, among others.

Table 2 summarizes the net growth in staff required by 2025 in each region.

Region	Net new staff needed by 2025
Asia-Pacific	140,000
North America	84,200
Europe, Middle East and Africa	81,500
Latin America	21,000

Source: Uptime Institute Intelligence 2021

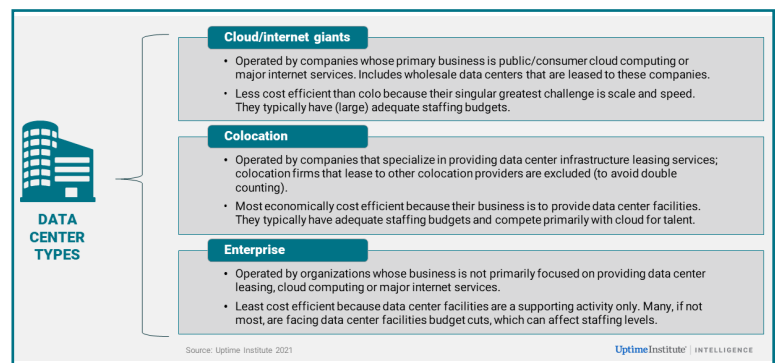
UptimeInstitute® | INTELLIGENCE

Cloud versus enterprise staffing needs

The explosion in public cloud services has been well documented. Cloud operators are creating and supporting new demand, especially in consumer-facing IT services, and attracting workloads (and staff) from existing enterprise data centers. The voracious growth of the cloud operators – as well as that of the internet giants – is, in turn, helping to drive leased data center sector growth. Measuring cloud data center demand, however, requires careful categorization.

We have assumed three broad types of data centers. Two are owner-operated types: enterprise and colocation. The third type, cloud/internet giants,

includes data centers that serve businesses that derive significant revenue from delivering internet and/or cloud



services, and that are owned by a cloud company/internet giant or a wholesale data center company and leased to a cloud/major internet company. (See **Appendix: Methodology** for definitional specifics.)

Cloud/internet giants can include major consumer internet companies (examples include eBay and Facebook), as well as enterprise public cloud providers (examples include Microsoft and Google). These businesses may also provide hosted application or infrastructure services to end users through, for example, software as a service (Salesforce.com, for example), bare metal hosting, managed hosting, web hosting, or infrastructure as a service (examples include Amazon Web Services and Rackspace).

These distinctions are important; we have defined data center types by the type of facility operator – this means cloud/internet giant data centers include leased facilities they operate (known as wholesale leased). Again, this is a trade-off for accuracy/insight: If we had classified wholesale leased data centers alongside colos (which are multi-tenant leased facilities), the size and trajectory for cloud/internet giants' staffing requirements would be materially lower, while the colo sector would be much larger.

As shown in Figure 4, forecasted demand for different types of data centers mirrors the long-term, fundamental and structural tilt toward ever-greater use of outsourced data center services, such as colo and cloud, and a relative decline in enterprise data center demand.




Figure 4. Global data center staff requirement projections by data center type

There are several reasons for the high staffing requirements in enterprise data centers. Enterprise facilities are currently the most numerous type of data center and are typically smaller than many colo and almost all cloud and internet giants’ facilities. This often means there are fewer opportunities for economies of scale – including for staff. For example, jobs roles such as IT hardware technician and electrician are required in all data centers, regardless of size (not including IT rooms and closets).

Similarly, cloud, internet giant and colo companies have and will invest much more in efficiency, using strategies such as value engineering (IT equipment is designed for reduced, easy and low-cost maintenance) and fail in place (allowing equipment to fail and be replaced or maintained later). Cloud, large internet and colo facilities often use more software-driven monitoring and automation technologies as well.

At the same time, enterprises are becoming more efficient. The trend toward consolidation of small regional data centers by enterprises into larger, centralized sites is well underway. This trend, alongside the move toward increased IT capacity in outsourced venues (colo and/or cloud), is driving the decline in enterprise staff needs over time.

Colo and cloud capacity, on the other hand, will experience continued rapid expansion (and so, too, will that of internet giants). The pace of growth can outweigh their economies of scale, which is pushing staff requirements for these types of data centers upward. We believe that in 2026 or so, total staff employed by cloud/internet giants’ data centers will surpass that of enterprise data centers.

Education requirements: Employers become more flexible

Education requirements – even for the same job role – can vary considerably across the sector, among different organizations, and between different regions. In some countries, industry-specific technical qualifications are required for certain job roles, although we are beginning to see requirements ease in some locations.

In reality, most jobs do not require a high level of formal education to carry out the role, even in positions where the employer may have initially required it. In other words, relevant experience, an internship/traineeship, or on-the-job training can often more than compensate for the lack of a formal qualification in most job roles.

We have assumed five broad minimum education requirements for data center roles. As Figure 5 shows, most of the workforce will require a university/college or technical trade school degree or, importantly, equivalent experience (as a substitute for a degree). Very few job roles – most of which are in technical engineering – require a university or college degree; equivalent experience will suffice for most of the tasks required in other roles. Only a limited number of roles require postgraduate certification (e.g., advanced engineering roles), but in those cases, the credentials are crucial.




Figure 5. Global data center staff requirement projections by minimum education or equivalent experience requirements

People from a wide range of educational backgrounds and experience often work in the same role. Also, skills transfer from other industries is becoming more common; equivalent experience could be from within or beyond the data center sector.

There are also different levels of job role maturity in different data centers. For example, an operations engineering program manager in a newly acquired data center may be tasked to create a new program, while an operations engineering program manager in another company may be tasked with maintaining an existing program; requirements for each may differ, even though it is the same job role. (There are also shared requirements in this case: for example, maintaining an existing program also involves new equipment, new processes, updated frequencies, and so on.) Job role maturity in smaller or newer organizations can also be a measure of how well-defined the roles are. In smaller companies, people often serve in many capacities and their job role may not be strictly delineated. As the company grows, the roles become more defined and more discipline around that emerges.

The industry is expanding and professionalizing, with a concentration of fewer, larger employers. This is helping to create more modern attitudes about education requirements. These larger employers are also investing more in training. Over time, we expect a gradual flattening of the highly varied levels of organizational and job role maturity across the sector.

Requirements for technical staff

We often hear concerns about rising demand for technical staff, particularly for the ongoing operation of data centers. Engineers can, notoriously, be among the most challenging technical job roles to fill.

In data centers, most engineering demand is for mechanical and electrical engineers in strategy and operations roles (a smaller proportion are needed for design and construction roles) – and it is forecast to grow steadily in the coming years, as shown in Figure 6. These roles all require an engineering degree from a university or college, at a minimum. Many data center engineers must obtain additional certifications, although it varies by region (for this reason, professional qualifications in addition to an engineering degree are not part of our estimates).




Figure 6. Global data center staff requirement projections: Mechanical and electrical operations engineers (university/college degree required)

Another concern for many managers is the availability of trained controls and monitoring staff, in line with the widespread intent of many operators to automate and operate more functions remotely. Figure 7 shows that demand for (various) job roles in data center controls and monitoring is expected to reach over 165,000 FTE equivalents in 2025.




Figure 7. Global data center staff requirement projections: Controls and monitoring

Uncertainty factors: What might change staffing demand?

All forecasts and models are prone to some inaccuracy. This is especially the case in sectors that are (relatively) new, growing rapidly, and are undergoing either structural or technological change. The mission-critical digital infrastructure industry is subject to all these. In this section, we briefly discuss five factors that may increase or decrease our forecast. While these factors were considered during the creation of the model, they were discounted because of the uncertainty around their impact.

1. Effect of automation and artificial intelligence

Our research shows that most data center operators and managers plan to increase their investment in remote and automation technologies, including those that involve artificial intelligence (AI). However, the impact on staffing needs is likely to be muted, at least in the short term. While more than a third of respondents to the Uptime Institute Global Survey of IT and Data Center Managers 2020 believe that AI will reduce their data center operations staffing levels in the next five years (as shown in Figure 8), a bigger proportion think it will have no impact during the period.




Figure 8. A third say artificial intelligence will reduce staffing levels within the next five years

2. Local edge staffing requirements

One of the most widely anticipated trends in IT and infrastructure is significant new demand for edge computing, fueled by technologies such as 5G, the internet of things and AI. However, net new demand for edge computing – processing, storing and integrating data close to where it is generated – has built slower than most investors (and suppliers) had anticipated. Our most recent annual survey data shows that only about one in five managers thought their organization’s edge computing demand would increase significantly in 2020 (Figure 9).




Figure 9. Most expect edge demand to increase

In our model we have not included staffing requirements for IT rooms and closets, nor for micro data centers at the edge. Because edge computing will rely heavily on prefabricated data center designs and will make considerable use of remote monitoring/operations, the staff requirements will likely be both lower than and different from those for centralized data centers. We expect that staffing requirements will

increase over time but the numbers are currently low and the demand timeline is difficult to predict.

3. Mission criticality requiring 24/7 operations staff

There are number of factors that are contributing to a move to greater levels of physical data center resiliency. Chief among these is the world’s growing dependency on digital services and, consequently, a growing aversion to digital risk; extreme weather and climate change risks; and most recently, the COVID-19 pandemic. As data centers move to a post-pandemic period, many will automate more (which could reduce net staff requirements), but some will respond by employing more to ensure 24/7 (24 hours a day, seven days a week) resilient operations (see Figure 10).




Figure 10. Some expect staffing costs to increase because of COVID-19

In our study, we have assumed that most, but not all, data centers that are 5,000 square feet or larger are mission critical, requiring 24/7 on-site staff (working in specific operations job roles; see **Appendix: Methodology**). The more important an operator views resiliency, the more likely they are to have more staff on site.

4. Data center design and construction timelines

The development of prefabricated data centers and related rapid-construction techniques (now several decades old) has led to much speculation that this will cause the loss of design and construction jobs. In reality, the impact has been imperceptible at an aggregate level. This is partly because of booming demand – helped by the very fact that data centers are becoming cheaper to build and can be built more quickly.

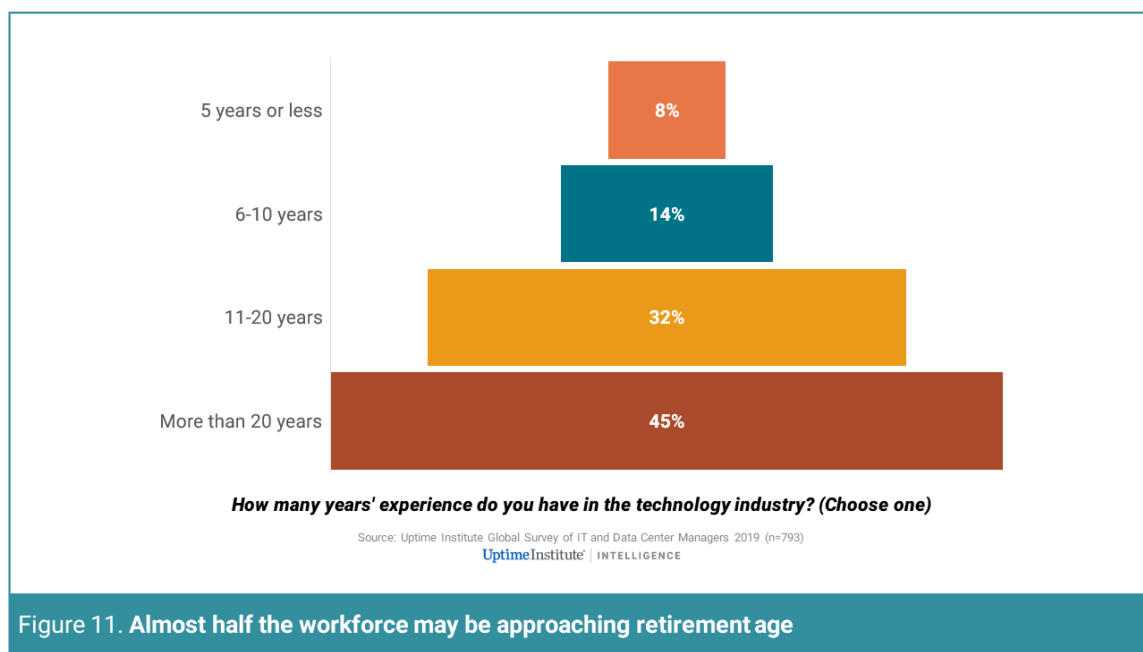
In this study, we have assumed different average time periods for the requirement of design and construction job roles, based on estimates

of project timelines for data centers of different sizes. This will vary considerably and, over time, those averages are likely to change (see our report [Best-in-class data center provisioning](#)).

5. An ageing workforce

In the large and mature data center markets of the US and Western Europe, there is concern that many employees are due to retire around the same time, causing a surge in demand. Many of those leaving are in senior roles. The lack of a succession plan for experienced data center staff is a significant problem for some organizations.

The strong preference for experienced job candidates in these mature markets means that even entry-level workers have often already completed stints in other sectors or professions. As a result, they are sometimes older than entry-level employees in other industries and may have shorter career spans. Uptime Institute does not have data to support this reasoning, but with 45% of respondents to Uptime's 2019 annual data center survey reporting 20 or more years of data center experience (see Figure 11), it can be extrapolated that the sector's workers are older than the US median (which was 42 in 2019, according to the [US Bureau of Labor Statistics](#)).



While many are predicting a silver tsunami – when a cohort of experienced professionals retire and leave behind numerous unfilled jobs and a vast experience gap (which may increase overall staffing requirements) – it seems likely that this will be spread over many years and will be focused mostly in North America and Western Europe.

Conclusions

The number, size and staffing needs of data centers around the world have grown rapidly – and oftentimes faster than recruitment practices and salaries have been able to adapt. With the strong build-out set to continue, there is widespread concern about skills and staffing shortages. Our research shows that while many managers are already struggling to fill open job positions, more than 325,000 net new FTE equivalent positions will be needed between 2019 and 2025 globally.

However, individual facilities experience the problem differently, depending on factors such as location, growth, local competition, staffing models and management practices (see our report [Staffing: Responding to the operations skills shortage](#)). For some, this will create significant difficulties and will require a concerted management response, but for most it will simply mean an increasingly competitive market for labor. What is clear is that the data center sector – and its workforce – is set to expand, which will exacerbate shortages where they already exist and could create new talent strains elsewhere.

As in every other aspect of data center operation, management must understand the local market, offer competitive compensation packages, remain nimble, avoid putting in place unnecessary impediments to hiring, and develop succession plans for each position. Our research shows that the most successful data center staffing strategies are those that lower barriers to employment (for example, by removing unnecessary education requirements) and that provide on-the-job and other training.

Appendix: Methodology

Data sources and estimates

This study is an estimate with values based on sampled real-world data from across the industry and from Uptime Institute's proprietary database. Values have been adjusted using statistical estimation approaches and assumptions, then modeled against our global data center market sizing and forecast. Our employment estimates have been adjusted to consider different staffing requirements at scale (including for organizations that have multiple data centers) and for data centers that have different levels of mission criticality.

Why bottom-up (not top-down)

This forecast is a bottom-up analysis, built on estimates of individual job roles across all functional departments (domains) in a data center.

Uptime Institute also undertook a top-down analysis, in which we searched the popular job site LinkedIn using different variants of the term "data center," including in different languages. Based on an in-depth analysis of those search results, which yielded some interesting insight, it was clear that many IT-related job roles outside of our scope are considered or used in association with this term. For example, roles in the IT hardware manufacturing sector are often known as being part of their supplier employer's data center division – and it was not feasible to reliably parse these from data center infrastructure job roles as we have defined them.

Sources of sampled data

Sampled data was derived from two primary sources:

- **Industry input**, including in-depth interviews with infrastructure and executive managers at companies across major geographies that combined own about 350 data centers. They represent the following vertical industries:
 - Data center real estate investment trusts
 - Data center design
 - Data center construction
 - Colocation
 - Cloud computing
 - Telecommunications
 - Financial services
 - Retail
 - Pharmaceutical

Interviewees represent a cross section of the sector, ranging from owners of some of the world's largest data center portfolios to operators with a single site in a mixed-use building. They represent multiple data center sizes and staffing models (a mix of in-house staff only, third-party staff only, and, more commonly, a combination of both) and different construction methods (including prefabricated modular, bricks-and-mortar, and retrofitting of existing structures), among other differences.

- **Uptime Institute's proprietary data and in-house expertise** - Our data includes detailed data center facility specifications, ranging from design to operational utilization and staffing levels, in nearly 100 countries over the past decade. Our in-house expert contributors included certified (professional) mechanical and electrical engineers and management and operations specialists, each averaging 20 years' practical data center experience.

Numbers for safe operation

These figures should be viewed as the number of employees required for the safe operation of reliable data centers, and not the actual number of employees that will be hired. In reality, many open job positions are unfilled, and many enterprise data centers are experiencing budget cuts.

Unknown factors

As with any estimate, there are several unknown factors that could affect accuracy, including macroeconomic developments. Historically, the data center sector has continued to expand, despite recessions – but this is something that is impossible to reliably predict. Other unknowns include the rate of IT capacity growth as a result of net new edge computing demand (and, therefore, demand for edge micro data centers and supporting staff). Similarly, the effects of AI and expanded automation approaches on employment requirements are not yet well understood; it is simply too early in their adoption for us to judiciously quantify. In order to provide a reasonable – if conservative – forecast, we have assumed a business-as-usual growth rate, in line with our data center capacity sizing and forecast model (for more on this topic, see **Uncertainty factors: What might change staffing demand?**).

Forecasted employment

Forecasted employment comprises:

- **Data center infrastructure staff**, which includes job roles spanning the beginning stages of new data center capacity – from location siting and negotiations to the design, construction, equipment installation and level 5 integrated systems testing – through the ongoing operations of servers, networking and connectivity, mechanical and electrical equipment, and the physical protection and cleaning of the site. End-of-life data center job roles, such as those involved in decommissioning,

are not included. Job roles further up the IT stack, for which forecasts already exist, are not included.

- **Sector specialist job roles**, which includes jobs that are specific to the design, build and operation of data center facilities, excluding nonspecialist job roles. A data center cleaner, for example, is included, but a cleaner at a data center owner's corporate office is not. A data center construction rigger is included, but the back-office roles of the construction firm are not. This means many job roles, such as accounting, human resources, marketing, sales, and so forth are not included.

This distinction between specialist and nonspecialist job roles was not always straightforward to define; notably, we have included colo pre-sales engineers (specific to data centers) but have not included colo salespeople – in part because the staffing of general business sales teams (with the business, in this case, being data center colocation) can vary considerably from one colo to another.

Fixed and variable employees

Our forecast includes fixed and variable employee types, including:

- **Fixed employees** – Those on the company payroll who perform a role full time or as a percentage of their time. Whether a role is fixed or variable can depend on different factors, such as the size of an employers' data center portfolio (see **Data sources and estimates**).
- **Fixed contractors** – Third-party employees who perform a given job role full time or as a percentage of their time. Again, whether a role is fixed or variable can depend on a range of factors.
- **Variable employees** – Contract or third-party employees who perform a job role either full time for a given duration or sporadically, calculated as hours per annum (and according to data center type, size and criticality, which we assume dictates requirements). Similarly, we have assumed that not all job roles are needed across all types and sizes of data centers (a data center strategy business analyst, for example, will likely not be required for an organization that has one small data center).

We do not specify whether employees required will be fixed, contractors or variable, although our research suggests that many data centers rely on at least some outsourced staff (see our report [Uptime Institute global data center survey 2020](#)).

Requirements for 24 hours a day, seven days a week employees are included, based on the number/capacity of facilities that have high levels of mission criticality. These round-the-clock job roles span operations, IT hardware, network and connectivity, and controls and monitoring job domains.

Assumptions

Our data center staffing forecast is based on the following assumptions, as agreed across Uptime Institute's technical team and with input from the data center owners/operators, design and/or construction firms that contributed to this study.

Forty-hour workweek

We have assumed that each FTE equivalent works eight hours a day, five days a week, equaling a 40-hour workweek.

Calculation of net utilized racks per data center

When calculating the average net utilized racks per data center, the year 2019 was chosen from our data center forecast. The totals for all four geographic regions for each data center type combination were averaged and then rounded.

Staffing requirements per square foot per rack

Staffing requirements will change per square foot per rack, but not to a significant degree. A reliable square foot per rack metric is notoriously difficult to establish because conditions vary considerably, including by region and by organization. We have assumed that most data centers are designed for 20 kilowatts per rack but, in reality, operate in a range of between 6-8 kilowatts per rack. Exceptions, such as high-performance computing environments (including liquid cooled) are unlikely to have a material effect on our model and are not included.

Local edge staffing requirements

We have not included staffing requirements for micro data centers. While the number of IT rooms and closets is estimated in our data center forecast, we have not included their staffing requirements in this model. This number of IT rooms and closets, as defined in our data center forecast, will decrease over time but will be largely offset by the increase in the deployment of micro data centers. This trend is hard to predict with any certainty. In our model we have assumed that the net effect will be largely neutral.

Effect of automation and artificial intelligence

The extrapolation of numbers from the base year (2019) onward do not consider the effect of any automation or AI advances that may significantly alter the staffing numbers of larger data centers over this timeline. As shown in Figure 8, more than a third of respondents to Uptime Institute's 2020 global data center survey believe that AI will reduce their data center operations staffing levels in the next five years.

Mission criticality requiring 24/7 operations staff

We have assumed that a vast majority of (but not all) data centers 5,000 square feet and above handle mission-critical workloads

(required for the ongoing operation of the organization) and require 24/7 staff either on site or on call. Round-the-clock job roles include physical security staff; critical operations management; critical facility engineers; heating, ventilation and air conditioning technicians; mechanics; controls and monitoring systems technicians; and data center infrastructure management specialists.

Data centers 5,000 square feet and below may or may not be mission critical but we have assumed they are primarily supported remotely and do not have 24/7 on-site staff requirements. While there will be circumstances when staff at these facilities are paid to visit the site outside of normal business hours, our model does not take this into account.

At-scale staffing requirements

The scale of an organization's data center portfolio determines the requirement for certain job roles, such as corporate-level roles (i.e., at-scale job roles). All of the roles in the strategy and business support domains are at-scale (corporate-level) roles. For our forecast, we have assumed:

- **Enterprise data centers** - Fewer than half of the total are multi-data center portfolios requiring certain at-scale job roles. We expect the proportion of multi-data center portfolios will increase over time; this incremental increase is not included in this model.
- **Colo data centers** - The vast majority of the total are multi-data center portfolios requiring certain at-scale job roles. We expect the proportion of multi-data center portfolios for colo businesses will increase over time; this incremental increase is not included in this model.
- **Cloud/internet giants** - All are multi-data center portfolios, requiring at-scale job roles.

Decommissioning end-of-life data centers

Staff required for decommissioning data centers are not included because we believe the practice of retrofitting existing buildings and data center sites will continue, if not accelerate, in the future – making it difficult to reliably estimate.

Differences in scale and data center types

The number of data centers under management by a single company is not known, yet this can affect staffing requirements for certain job roles. In addition to the previously mentioned assumptions, we estimated the financial viability of staffing requirements for job domains using Uptime Institute's data center operating costs model and our survey data on salaries as guidance for certain data center types and sizes.

These calculations were used by Uptime Institute's sector experts to reasonably determine how many staff each job domain could support

financially. In addition, the following high-level assumptions were applied:

- Colo data centers are the most cost efficient because their business is to provide data center facilities.
- Cloud and internet giants' data centers are less cost efficient than colo because their singular greatest challenge is scale and speed. They typically have (large) adequate staffing budgets.
- Enterprise data centers are the least cost efficient because data center facilities are a supporting activity only. Many, if not most, are facing data center facilities budget cuts, which can affect staffing levels.

Variable staffing requirements (job roles)

We have assumed FTE equivalent staffing requirements based on primary areas of responsibility (i.e., job roles). In smaller organizations and data centers, one person will likely fulfill several roles; we have assumed a proportion of a 40-hour workweek for each likely area of responsibility (role). In any size data center, some roles do not require a dedicated FTE and are carried out by contractors – we have assumed hours per year requirements for each “variable” role (according to data center type and size). Similarly, we have assumed that not all roles are needed across all types and sizes of data centers.

Data center design and construction timelines

We have assumed different time periods for the requirement of design and construction job roles, based on estimates of design and construction timelines for data centers of different sizes. In reality, this will vary considerably. Our estimates are based on reasonably assumed averages rather than best-in-class provisioning times.

Job titles versus job roles, and new job roles

This forecast is based on job roles (not job titles) because in a smaller data center a single FTE may have multiple job roles, whereas in a larger data center multiple FTEs may have the same (singular) job role. Also, job titles can vary considerably, including by geographic region and by organization (proprietary job titles). In determining how much time a single FTE spends carrying out multiple job roles, we have assumed a consistent mix of job roles in smaller data centers and a consistent amount of time spent on each. In reality, this will vary.

New job roles may emerge in the coming five years that have not been factored into our model. These are unlikely to have a material impact on the total forecast.

Job domain definitions

As discussed in the main report, a job domain consists of a multidisciplinary team of people, with various job roles, that together are responsible for specific functions. Uptime Institute proposes nine job domains, defined as follows:

- **Business Support** sustains and strengthens all other data center domains in interdisciplinary ways.
- **Strategy** guides the holistic planning, delivery and operational life of data centers.
- **Design** plans the detailed technical design, construction and retrofitting of data centers, often with a focus on availability and efficiency.
- **Construction** builds and manages the construction of new data centers and upgrades to existing facilities.
- **Operations Engineering** plans, optimizes and oversees data centers' critical systems, often with a focus on availability and efficiency.
- **IT Hardware** manages, installs, moves and maintains, and in some cases also designs, all IT equipment.
- **Network and Connectivity** connects all data center IT equipment, including between multiple data centers.
- **Controls and Monitoring** designs and maintains the systems that monitor and manage data center operations and equipment.
- **Operations** ensures the reliable operation of critical data center facility systems.

Note that while domains are largely composed of multidisciplinary teams, some staff are also interdisciplinary in how they operate.

Data center market size and forecast

The data inputs and assumptions were applied to Uptime Institute's data center market size and forecast model. This model is derived from:

- Longitudinal surveys of the owners/operators of enterprise, colo and cloud/internet giant data centers.
- A proprietary database of enterprise and colo data centers.
- Macroeconomic market data.
- Financial filings by major data center owners/operators (specifically, 10-Q and 10-K reports).

Our data center market size and forecast model is based on the following metrics:

- **Data center count (number of data centers):** Represents a view of the data center installed base, based on the number of total data center facilities.

- Data center space (total available technical square feet in the data center installed base).
- Data center power (total available technical kilowatts in the data center installed base).
- Data center racks (total number of racks that could be installed within the total square feet of data center space available).
- Data center utilized racks (total number of utilized racks in the data center installed base).

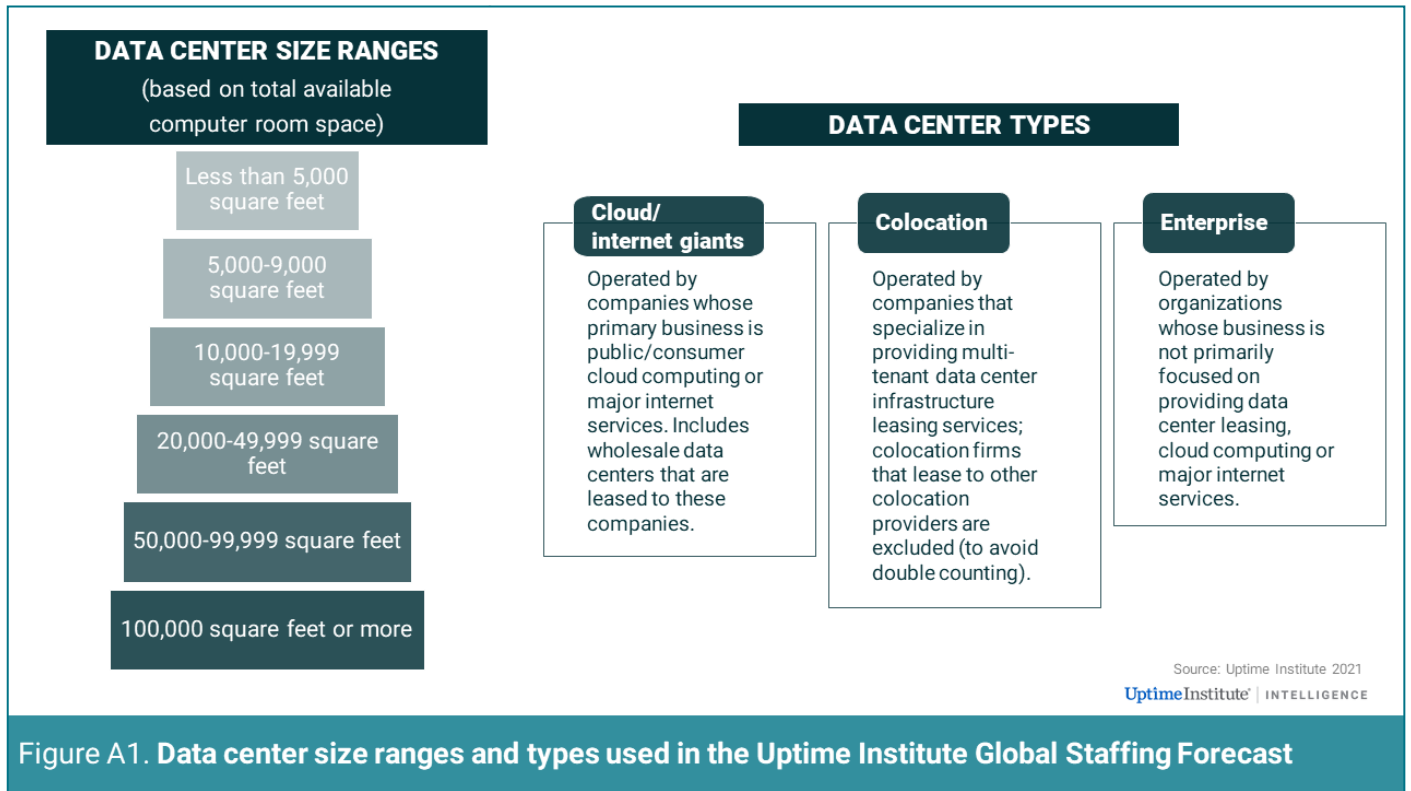
These metrics have been calculated for different data center types and size ranges, as described below.

Data center types and size ranges

Data center types are classified into three broad types of primary operator responsibility:

- **Enterprise** - Data centers owned and operated by organizations whose business is not primarily focused on providing data center leasing or cloud computing/major internet services. This includes traditional verticals such as healthcare, manufacturing, retail, public sector, financial services, transportation, construction, communications, media, energy and business services.
- **Colocation** - Multi-tenant data centers that are normally owned by firms that specialize in providing and operating facility infrastructure services, including data center leasing services to customers (individual enterprises). Some colo firms lease data center space to other colo providers; we have excluded these to avoid double counting.
- **Cloud/internet giants** - Data centers that are operated by businesses whose primary revenue stream, or a large portion of it, is derived from delivering internet and/or cloud services from their data centers. This can include major consumer internet companies (examples include eBay and Facebook), as well as enterprise public cloud providers (examples include Microsoft and Google). These businesses may also provide hosted application or infrastructure services to end users through, for example, software as a service (Salesforce.com, for example), bare metal hosting, managed hosting, web hosting, and infrastructure as a service (examples include Amazon Web Services and Rackspace). Data centers operated by cloud/internet giants but owned by a wholesale provider are included.

For each job role in each job domain, we calculated FTE equivalents globally for six different (and broad) data center size ranges (Figure A1).



ABOUT THE AUTHORS



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



Chris Brown is Uptime Institute's Chief Technical Officer. His career in critical facilities includes analyzing requirements, retrofitting existing systems, performing capacity planning and implementing new facilities design. He also has experience managing the operation and maintenance of critical facility infrastructure to include managing maintenance personnel and creating a preventive maintenance program. In his role at Uptime, Chris administers all of Uptime Institute's Tier Certification and training programs and helps ensure that these programs remain current and relevant to the market. Contact: cbrown@uptimeinstitute.com



Todd Traver is Uptime Institute's Vice President for IT Optimization and Strategy and has over 30 years' experience in all aspects of data center and IT planning, reliability design and operational efficiency. Todd was IBM Global Technology Services' Chief Engineer and Senior Technical Staff Member responsible for developing global data center portfolio technology strategy, tools, and operational and process optimization plans. Todd's passion is to expand Uptime Institute's offerings based on these experiences, providing a pathway for others to innovate new approaches to integrate facilities and IT departments, which reduces total cost of operation and enhances end-user application and data reliability and resiliency experience. Contact: ttraver@uptimeinstitute.com



Fred Dickerman is Uptime Institute's Senior Vice President of Management Services – Europe, Middle East and Africa, responsible for developing and delivering bespoke strategic consultancy services to Uptime Institute's global client base, with specific focus on mission-critical governance and standards, portfolio management and operations, and C-level advisory assignments. He is also tasked with creating new courses and service offerings to support data center owners and operators in meeting the challenges of high availability and cost-effective operation. Prior to joining Uptime, Fred worked with major data center developers in the US and Russia and accrued a project resume with over 20 successful data center builds. Contact: fdickerman@uptimeinstitute.com



Rich Van Loo is Uptime Institute's Vice President of Facility Management Services. In this capacity, Rich performs operational assessments and operational sustainability certifications. He is also an instructor for the Accredited Tier Specialist course and Accredited Operations Specialist course. Rich's work in critical facilities includes responsibilities ranging from project management of a major facility infrastructure service contract for a data center to space planning for the design/construct for several data center modifications and facilities IT support. As a contractor for the Department of Defense, Rich provided planning, design, construction, operation and maintenance services for mission-critical data center facilities throughout the world. His career includes 11 years as a facility engineer and 15 years as data center manager. Contact: rvanloo@uptimeinstitute.com

ABOUT UPTIME INSTITUTE INTELLIGENCE

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

ABOUT UPTIME INSTITUTE

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

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

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