Uptime Institute INTELLIGENCE

PLANNING & STRATEGY

UI Intelligence report 47

Uninterruptible power: Adoption trends to 2025

Authors
Rhonda Ascierto, Vice President of Research, Uptime Institute
Paul Johnson, Technical Consultant, Uptime Institute

Data center uninterruptible power supply (UPS) systems are evolving. New technologies are enabling various electrical approaches. But will UPS systems of the future meet the changing requirements of operators? This report discusses UPS adoption trends to 2025 for different types of data centers.

15-20 minutes to read

UII-47 v1.0PM published April 26, 2021 last updated April 26, 2021

This Uptime Institute Intelligence report includes:

ABOUT UPTIME INSTITUTE INTELLIGENCE

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

EXECUTIVE SUMMARY

To better understand the requirements of uninterruptible power supply (UPS) systems in the (near-term) future, Uptime Institute conducted in-depth interviews with 37 data center operators and their major engineering or operations partners, globally. Some clear trends emerge, notably that centralized UPS systems will likely continue to dominate in data centers with at least 1 megawatt of IT capacity, especially in those owned by enterprises and colocation providers.

KEY FINDINGS

- **Through 2025, few core data centers are likely to adopt novel electrical approaches, such as medium-voltage UPS, despite the advantages. Distributed uninterruptible power systems with batteries will also likely remain niche due to technical and commercial risk, either perceived or real.**
- **Centralized, 3-phase UPS systems will continue to dominate for the next few years, at least, even though problems with batteries, product reliability and safety, in addition to other factors, are likely to persist.**
- **Adoption of newer, high-efficiency UPS operational modes that go beyond established eco-mode functions will be largely driven by the ability of manufacturers to address concerns over risk.**
- **There is an appetite for data-driven UPS remote monitoring services, including condition-based maintenance, but IT security requirements threaten to slow — or stymie — broad adoption.**

Introduction

Power infrastructure requirements in data centers will be shaped by several factors in the coming years. For many, power architectures will be driven by pressure to lower capex without additional risk, while for others, regulatory concerns or changing customer requirements will play an increasingly prominent role. More operators will pursue higher levels of efficiency, adaptable resiliency, and integrated automation. The use of new battery technologies (and of renewable energy sources) will spread, particularly for data centers supporting very large critical loads — of which there will be many more.

This report focuses on just a few areas of potential change, including the extent to which distributed UPS systems with batteries and medium-voltage UPS systems play a role, and the ways in which centralized UPS systems will scale, operate efficiently, and be remotely monitored.

Findings are based on interviews with 21 data center operators (at cloud, colocation, telecommunications and other firms [enterprises]), as well as eight major operations and engineering firms, and eight Uptime Institute consultants (electrical engineers).

The focus of this investigation is on specific technologies and electrical topology approaches; there are many additional areas under investigation across Uptime Institute. Most lines of inquiry in this study are for larger "core" data centers, with at least 1 megawatt (MW) of provisioned IT capacity (the average cloud and colo site in our study is about 65 MW), as opposed to smaller "edge" data centers. (See the **Appendix** for details on study methodology.)

Distributed UPS with batteries

Some of the largest data center operators have adopted distributed uninterruptible power with batteries — for example, power protection with batteries at the rack. Examples of this approach are set out in alternative infrastructure designs, such as that proposed by the Open Compute Project, with batteries at the rack or row level. These types of "open architecture" approaches can enable operators to achieve efficiencies in several areas. Some adopters have reduced capital expenditure on the electrical topology alone by about 25%, compared with a traditional centralized approach.

Given the apparent cost advantages and the support of a distributed-with-batteries approach by some very large operators, its widespread use has long been thought to be an inevitability. However, uptake has, to date, lagged far behind industry expectations. Our research suggests adoption of the distributed approach may increase, but only slowly: just a fifth of the operators in our study expect to adopt or increasingly use the topology in their core data centers in 2025, and most operate cloud data centers in the Asia-Pacific region and North America. A small number of colos (in the Asia-Pacific region, Europe and North America) also expect to increase their usage. The most cited reasons for higher adoption were advances in newer battery technologies, such as lithium-ion (Li-ion); standardization; and, for colos, meeting the requirements of their (typically large) clients. Two respondents say batteries are likely to be built into servers in the future, instead of being part of a UPS system, which will drive adoption of distributed approaches.

Other reasons for an expected increase in distributed UPS (with batteries) include:

- An ability to better match the size of the UPS with variations in power densities (load fluctuations) in the rack.
- Increased availability and reduced mean time to repair faults. (While the UPS can statistically be less reliable when it is distributed, an ability to quickly repair faults drives up overall availability.)
- The ability to manage and optimize batteries on an individual basis.

Not all cloud facilities plan to increase usage — 43% of the cloud operators in our study do not expect their distributed UPS footprint will grow by 2025, nor do the majority of colo respondents (72%). Most in our study (21 participants) say they would not adopt the approach at all by 2025.

Some cloud respondents have a small portion (10% to 20%) of distributed power protection in their facilities today, while one cloud operator (in the Asia-Pacific region) reports 50%; none expect these proportions to change by 2025.

What, then, are the barriers to adopting or increasing distributed approaches? Many cite multiple deterrents, with the most common being client requirements for colos. Several colos say their clients do not want distributed power protection. Colos in our study often also say that because they have many small clients, a centralized system is preferred to allow power protection to be distributed across multiple client loads. Various types of operators have cost and safety concerns with newer battery technologies used in decentralized approaches.

Other concerns include:

- Placement of UPS systems in computer rooms/data halls, which could potentially cause heat-load issues in the data hall and, for colos, could take up valuable space (that would otherwise be leasable space).
- Placement of UPS systems in computer rooms/data halls could create security concerns (e.g., maintenance staff and vendors may be denied access to a data hall due to sensitive workloads and clients, possibly leading to rescheduled or deferred maintenance.
- More maintenance issues, greater complexity, less flexibility, and greater operations difficulties.

For the next few years, at least, it seems that concerns will outweigh any benefits of a distributed UPS approach with batteries for most in the study group. If our cohort is representative of most operators, this approach will remain a preferred option for relatively few — mostly large-scale cloud operators — through 2025.

Medium-voltage UPS

Today, most centralized 3-phase UPS systems operate at between 400 volts and 480 volts. UPS systems operating at a higher-than-typical voltage can offer efficiency advantages. These medium-voltage (MV) UPS systems usually operate between 6.6 kilovolts (kV) and 24 kV and are highly efficient. At this voltage, UPS systems can reduce facility build and operation costs because electrical currents, and therefore losses, are lower and cables can have a smaller cross section (diameter). MV UPS systems can reduce the need for companion components, such as switchgear, and, compared with low-voltage systems, can be placed farther away from computer rooms (e.g., in a substation or electrical room).

Yet the overwhelming majority in our study, 34 of 37 respondents, say they would not be likely to adopt MV UPS by 2025.

Of the three respondents likely to adopt (a mix of cloud, enterprise and colo operators), two say it would be for improved total cost of ownership (TCO). Specific reasons for planned adoption include:

- A reduced footprint (overall size).
- The ability to distance the UPS from the IT load.
- The ability to support high-density computing and fluctuating loads, such as artificial intelligence (AI) applications.

Why are so many others not planning to adopt? The reasons are shown below in order of commonality (see **Barriers to MV UPS adoption**). TCO or efficiency concerns are obstacles for half of the 34 respondents not planning to adopt MV UPS.

Barriers to MV UPS adoption

TCO and/or efficiency, including high initial cost and a perceived lack of financial return. Higher associated costs include requiring a MV generator (size plus the installation costs) or an additional step-up transformer.

Separation of UPS and load. A few respondents say having UPS systems (low voltage) and batteries as close to the load as possible means a lower chance of failure in the distribution path. MV UPS is viewed as being higher risk (compared with low-voltage UPS) if located farther from the server hall because of the increased possibility of distribution failures (due to faults in the system and also switching surges).

Client/risk. Several colos say that their clients are risk averse and that MV UPS is considered too high a risk for most. Some say their clients are used to seeing a traditional static UPS near the data hall providing security of supply. Having the power protection further upstream presents a perceived risk of disruption to significant quantities of load.

Safety risk, including the safety risks of handling high voltage levels. One respondent cites the risk of more stringent regulatory requirements in the future, which could limit their ability to use MV UPS.

Scalability. A few respondents say their loads are not large enough (nor will they be in the future) in any given facility to justify MV UPS. One respondent says MV UPS is not suited to colo because it is not as scalable. Low-voltage UPS has more granular steps, which means it can be scaled in smaller, more precise increments.

Too new. Respondents say that the technology is too new and, therefore, seen as unproven. A lack of product history has deterred some — specifically, an inability to compare the technology's efficiencies, footprints, reliability and availability.

Vendor selection. Some respondents are concerned about a limitation of MV UPS vendors, which could negatively impact costs and availability of products.

Skill set. The need for licensed MV electricians in many countries and the difficulty of finding a hybrid-authorized person or power conversion specialist are commonly cited. The (limited) availability of trained vendor or service provider staff is also a concern, as is (limited) local availability of and long lead times for spare parts.

Centralized UPS

Centralized 3-phase UPS has long been the system of choice for most larger (total UPS capacity of 1 MW or greater) enterprise and colo data centers. Centralized systems are typically viewed as having higher levels of reliability when compared with modular or distributed alternatives, thanks to the single static switch inside centralized models. In modular or distributed alternatives, several small static switches are used, which means faults can propagate across them — a concern for many operators of relatively large data centers. (Modular or distributed UPS systems are, however, often less expensive and easier to manage in facilities smaller than 1 MW.)

Centralized UPS systems are also fast and easy to deploy across large data centers and are often favored when IT space is expected to grow quickly over time. They are not, however, viewed by operators as inexpensive to purchase or operate. And, as discussed in **Pain points**, they are not without their issues.

Centralized UPS will likely continue to dominate in data centers with 1 MW or greater IT capacity, particularly those owned by enterprises and colos, for at least the next five years — and most likely for the next 10. This is partly because new developments in UPS technologies are typically slow to be released, and partly because they are slow to be accepted. Many operators and colo customers are averse to change, particularly when it involves well-established critical components.

Perhaps the biggest force of change to how UPS systems will be deployed in the future will be battery technologies — although, as discussed in **Distributed UPS with batteries**, there is still concern or skepticism over new approaches, such as the use of Li-ion batteries.

Pain points

Centralized 3-phase UPS systems are used in most data centers, but not without persistent issues. Many of these problems are expected to last until 2025 (at least). Some new issues are also likely.

Operators' primary pain points — issues that keep them awake at night — mostly relate to batteries, product reliability and safety. Battery pain points affect all data center types. In our study, product reliability and safety stand out as being particularly troubling to colo, cloud and telecom operators. Performance and output characteristics, as well as service and commissioning issues, are also common.

As Figure 1 shows, most interviewees cite multiple pain points today and most expect they will face multiple issues by 2025.

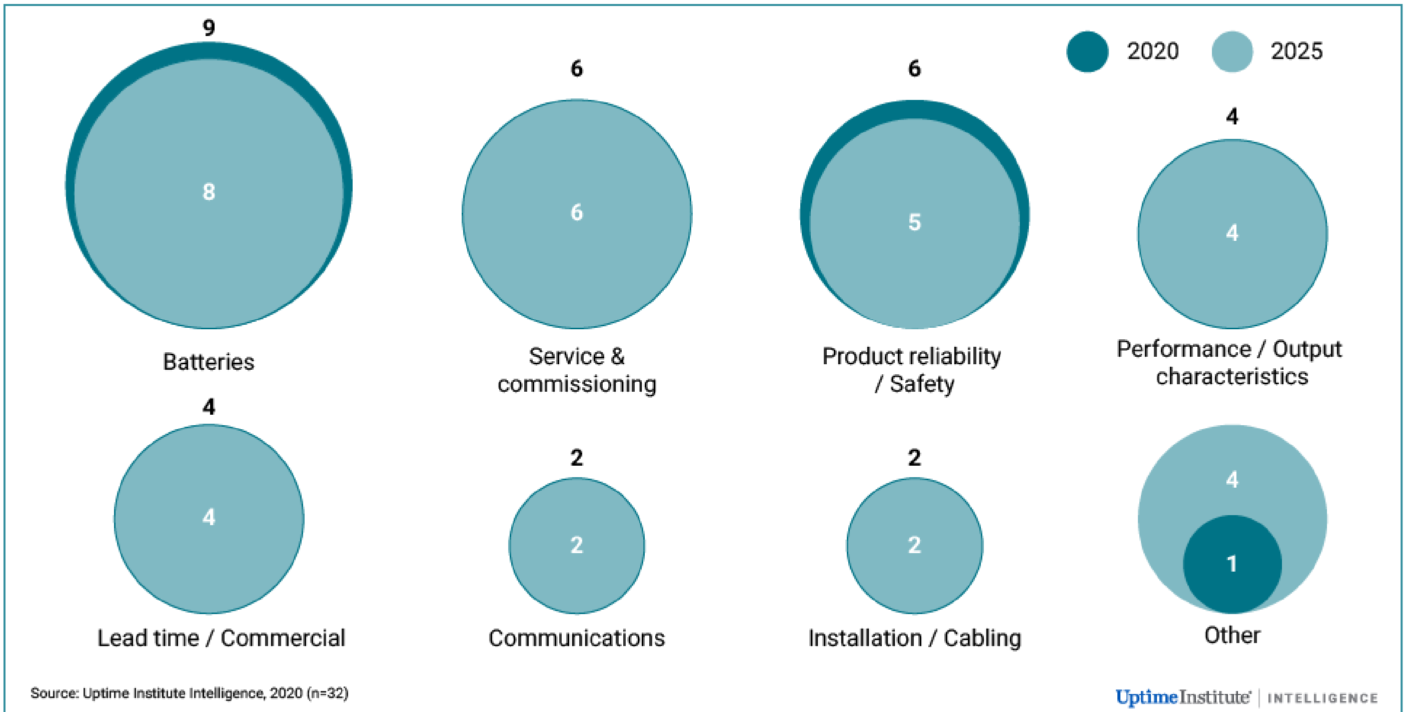

Figure 1. **Primary pain points of centralized 3-phase UPS systems in core data centers in 2020 and 2025**

The most common battery complaints today involve the high frequency of maintenance needed, and availability and performance issues (despite manufacturers' claims). A lack of commercially available integration of batteries is also a problem for some.

Looking ahead, many expect problems with Li-ion batteries, including installation, fire safety, monitoring and inspection issues. Others who had installed the technology are bracing for the possibility of high end-of-life costs.

Specific issues with product reliability and safety today mostly relate to a lack of service continuity and an increased number of outages — issues that most anticipate will persist through 2025.

Some users are experiencing high failure rates, including with UPS fans and, in older systems, with inverters. One respondent says its UPS fails when operating at full capacity, citing numerous internal component failures. Another hypothesizes that high overall demand for UPS systems in the last few years could have led to a lack of "burn-in" factory testing, so failures due to manufacturing, assembly or configuration issues may not be caught before products ship.

Looking ahead to 2025, some in our study are concerned about the reliability of UPS systems that will, by then, be close to the end of their expected life; they also worry about potential issues with ongoing product support from the systems' manufacturers. Others are concerned that, by 2025, centralized UPSs may not be available to adequately support expected (large) increases of AI and other highly variable workloads. This could lead to a situation where IT capacity cannot be used at peak times because the power available is insufficient — partly because of poor workload management. Greater efficiencies, smaller footprints and increased flexibility are other improvements over existing UPS systems that some respondents would like to see in the future.

Seven participants in our study, at various types of data centers and in different regions, say they have no pain points with their current UPS systems. In this group, five also do not anticipate issues arising in the coming five years. The two others expect that, by 2025, problems with batteries will arise, as will product reliability problems as systems age.

Outside of this group, two respondents who cited pain points today (performance issues due to paralleling too many units for one interviewee and UPS product reliability for the other) expect those issues will resolve by 2025.

Scaling preferences

How will operators at different types of core data centers modularly scale their centralized 3-phase UPS systems by 2025? With internal redundancy achieved by adding capacity within frames, or by adding new frames? A large majority would likely prefer to add new frames.

Cloud and colo interviewees mostly say they would prefer to scale by adding new frames, while telecom operators would mostly prefer to add capacity within frames. The picture among enterprises is mixed, as shown in Figure 2. (One enterprise and another cloud operator say they are undetermined; their decision will depend on UPS capacity or the design of the IT load distribution by 2025.)

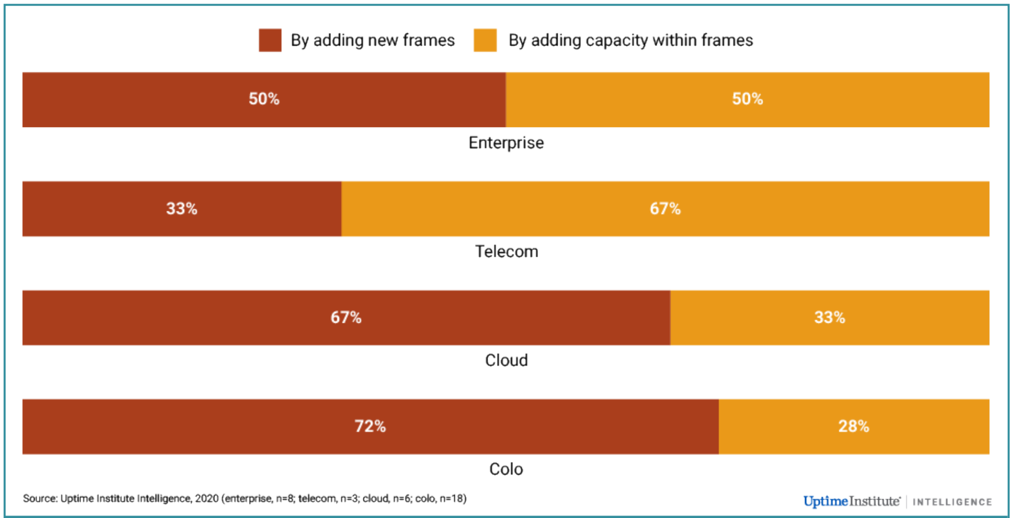

Figure 2. **Preferred way to scale centralized 3-phase UPS modularly by 2025**

Expanding by adding new frames is viewed by most as a way to lower risk. The approach could also increase safety at some sites because it avoids work in live enclosures.

Another driver for adding new frames is cost-effectiveness, in part because it can lower the need to over-provision a data center's UPS capacity at the outset. When frames are added, a UPS system has a dedicated battery arrangement per UPS (as opposed to having to expand the existing battery or oversizing the battery initially). Scaling with additional frames is also often viewed as fast to deploy. For respondents who would prefer to add modules within existing frames, most say it is because of lowered risk. Expanding within frames is viewed as increasing resiliency (therefore lowering risk) because redundancy can be maintained (leading to less downtime) and increased demand for IT loads can be better accommodated.

Speed of deployment is another common reason to expand within frames, specifically because there is no need to install cabling, cabinets or switchgear at a later date (after initial installation of the UPS system frames). One respondent is using smaller (less than 50 kilowatt, or kW) modules today to increase flexibility; at these sizes, modules are easy to store on-site. Some also think that expanding within frames has a favorable cost of ownership.

Newer, high-efficiency modes

Most centralized 3-phase UPS systems today achieve 94% to 96% efficiency in normal operational mode. Economy (eco) mode removes the double conversion process to enable higher levels of efficiency — up to 98%, typically.

In recent years, newer, high-efficiency operational modes that go beyond eco-mode have launched. In addition to normal operation (with the use of unconditioned power via the UPS bypass), newer UPS modes can actively manage power conversion by conditioning and supporting the output, but with reduced losses (all while maintaining the same level of power protection and the same level of battery backup).

However, for centralized 3-phase UPS systems that support only IT load (as opposed to those that support mechanical, or a mix of IT and mechanical), more than half (60%) of those in our study say they are unlikely to adopt newer, high-efficiency operational modes by 2025 (see Figure 3).

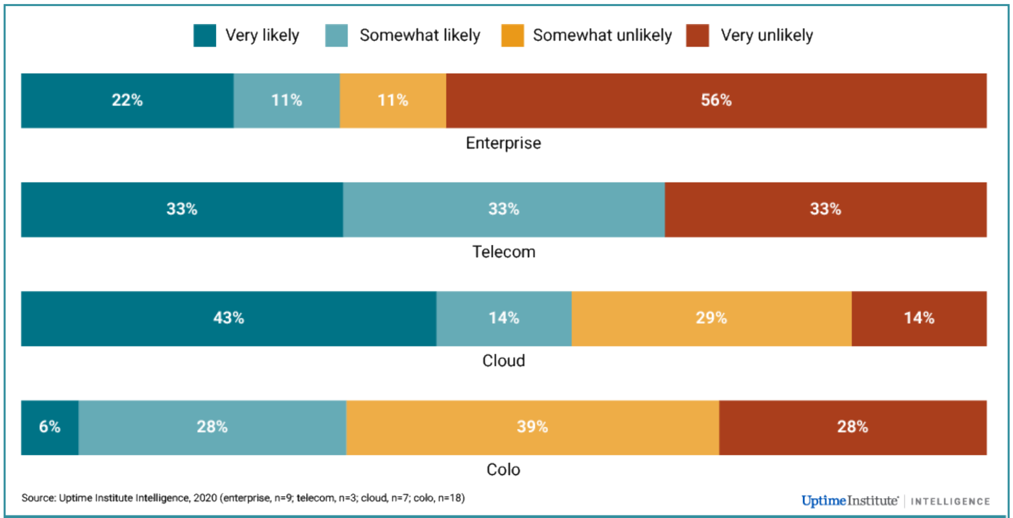

Figure 3. **Likelihood of adopting newer UPS high-efficiency modes by 2025**

Among our study subjects, increased risk is the overwhelming barrier to adoption, particularly for colo operators and their cloud clients. Newer, high-efficiency modes rely on a direct connection between input and output power terminals (via either a choke or switching device) to increase their efficiency — but this removes the galvanic isolation provided by double conversion, which means the load can be exposed to lightning or fast spikes and surges in voltage. Several colos say this lack of isolation is problematic for many of their customers. One European colo notes that even after factory demonstrations showing transfers between high-efficiency modes, clients still prefer full double conversion (normal operation).

A general reticence to being an early adopter of technology is also a factor for many respondents, including concerns over limited vendor choice.

One respondent cites a low level of confidence in the promise of reliable power protection without voltage and frequency independent (VFI) and voltage independent (VI) modes. Another notes that numerous servers are sensitive to supply voltages and require a conditioned voltage source from UPS equipment. Their concern is that any IT server disruptions, including the loss of a server, may happen unnoticed, potentially severely impacting the business.

Cloud and telecom operators are among the interviewees most likely to adopt UPS high-efficiency modes in the future. Cost efficiency is the primary driver, including improved power usage effectiveness —PUE — and lower overall energy costs.

Remote monitoring with condition-based maintenance

Some UPS manufacturers have launched data-driven remote monitoring services for UPS systems that include — or that will soon include — condition-based maintenance services.

By 2025, most in our study (65%) say their organization would be unlikely to use data-driven remote UPS monitoring services from the manufacturer (including condition-based maintenance services), either for their core data centers (5 MW or greater of IT capacity) or in distributed edge data centers (200 kW or greater of IT capacity), as shown in Figure 4.

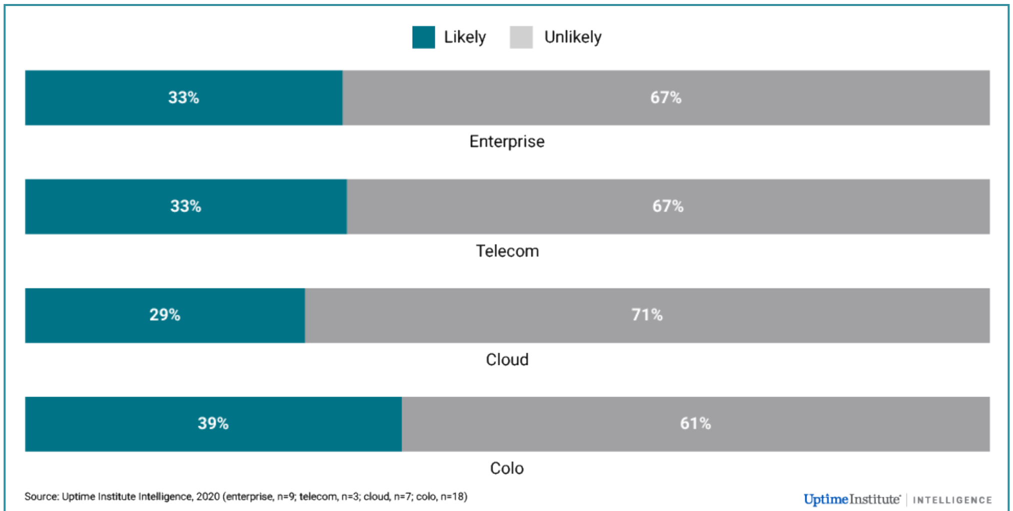

Figure 4. **Likelihood of adopting UPS remote monitoring with condition-based maintenance by 2025**

The most common barrier, by far, is security concerns, especially among cloud and enterprise respondents. They say exposing critical equipment directly to a vendor's online portal (and the vendor's cloud) is a technical and commercial risk. Most are concerned about technical risk, including inadequate firewall protections and potential cybersecurity attacks and ransom situations.

In our study, colo respondents in all geographies point to the perceived risk by clients, particularly those in highly regulated industries (e.g., finance and healthcare). The benefits of the technology would not outweigh the potential loss of confidence among clients.

Several interviewees say they would prefer to use their own monitoring system to retain control of the system (and data) and to avoid the risk of external exposure. Others do not see adequate value in condition-based monitoring and viewed on-site physical inspection as a superior approach ("there are things that can only be seen physically"). Among the 35% in our study likely to adopt remote monitoring services from the UPS manufacturer, condition-based maintenance is often viewed as more effective than scheduled approaches. Some say condition-based maintenance could potentially reduce the number of on-site interventions of a system, decreasing the possibility of human errors.

In the respondent group that are likely to use these services by 2025, there is interest from all types of data centers, particularly in the Asia-Pacific region and North America, as well as from colos globally. However, most would do so only if security concerns were addressed.

Some colos point to potential IT security weaknesses when transferring data from equipment on-site to the manufacturer's cloud. For others, firewall security issues would need to be addressed, and protection provided against potential unauthorized access into control systems. One enterprise operator in North America says they would adopt remote services, but the threat of hacking would likely mean the UPS manufacturer would have monitoring-only permissions (and no access to control functions).

Assuming that their security concerns are addressed, those respondents likely to adopt say remote services would offer operational value. Several point to the benefits of AI (and other big-data) approaches that can underpin remote UPS monitoring. They view real-time information and AI as enabling better-informed decisions than other approaches, leading to reduced risk. Several say AI is most effective in identifying potential issues before they happen, reducing costly incidents.

Cloud provider interviewees in the Asia-Pacific region say the lasting impacts of the COVID-19 pandemic is an impetus for them to adopt UPS remote monitoring services by 2025.

Conclusion

Operators are typically willing to adopt novel electrical approaches only if the rewards clearly outweigh the risks associated with doing so. Often the risks are technical or commercial. But, as our study shows, sometimes it is simply that an approach is too new, and therefore considered unproven, especially for enterprises and the colocation providers that serve them.

Our study suggests that for the foreseeable future — by 2025 — MV UPS systems are unlikely to be used by significantly more operators than today. Distributed uninterruptible power systems with batteries will also continue to be favored by only a relative few (mostly cloud operators).

If our cohort is representative of most operators, centralized, 3-phase UPS systems will remain the preferred option for most facilities — despite various and persistent issues ranging from problems with batteries and product reliability to safety.

If manufacturers can adequately address the risk concerns of many operators, greater adoption of newer, high-efficiency UPS operational modes that go beyond established eco-mode functions is likely. Similarly, some operators are interested in data-driven UPS remote monitoring services, including condition-based maintenance, but only if their IT security concerns are resolved. It may be that these newer approaches are, for many, an eventuality — but perhaps over a longer time period, beyond 2025, once they are no longer widely considered novel.

Appendix: Methodology

Uptime Institute subject matter experts (chartered or principal electrical engineers) conducted 30-minute telephone interviews in mid-2020 with data center owners and operators, as well as specialist third-party operations and/or engineering providers (serving multiple clients).

Interviewees were asked the same set of questions. They were provided with context and definitions to ensure a consistent understanding of terminology and uniform input for this study. Uptime Institute expert consultants, who provide professional services to data center owners and operators, also answered the interview questions, based on their real-world experiences with end users. The responses from third-party providers and Uptime consultants are based on their primary type of clients. Data center types are defined as:

- Cloud – Data centers that are privately owned by a cloud computing provider (also referred to as infrastructure-as-a-service).
- Colocation (colo) – Data centers owned by a multi-tenant data center provider.
- Enterprise – Data centers that are privately owned by a business in a vertical industry that is not cloud computing, colocation, or telecommunications.
- Telecommunications (telecom) – Data centers privately owned by a telecommunications provider.

The number of respondents per data center type is shown in Table A1.

Table A1. **Study participants by data center type**

Data center type	Number of respondents
Cloud	7
Colocation	18
Enterprise	9
Telecommunications	3
Total	37

Source: *Uptime Institute 2020*

Uptime Institute INTELLIGENCE

Respondents either provide input into decisions or had or shared responsibility for decisions regarding their organization's plans — either underway or under investigation — for their data centers' power protection requirements, specifically for 3-phase UPS systems. In addition to Uptime Institute expert consultants, respondents represent:

- Director, Senior Director, VP+, C-level in commercial data center facilities.
- Director, Senior Director, VP+, C-level in data center operations companies.
- Director, Senior Director, VP+, C-level in data center engineering companies.
- Director, Senior Director, VP+, C-level in mission-critical data center facilities.

Respondent demographics are shown in Figure A1. Colo and cloud respondents in this study have the largest number (count) of core data centers and approximate IT capacity (megawatts) today and expected in five years (consistent with overall market growth projections). Most data center capacity in this study is in North America and the Asia-Pacific region. However, the significant expansion plans of a single cloud operator in the Middle East-Africa (MEA) region means MEA is the region with the largest average growth of IT capacity per data center by 2025.

Enterprise	2020	2025
Number of core data centers	41	41
Total IT capacity per core data center (in megawatts)	1.7	2.3

Telecom	2020	2025
Number of core data centers	8	13
Total IT capacity per core data center (in megawatts)	5.0	6.9

Cloud	2020	2025
Number of core data centers	211	293
Total IT capacity per core data center (in megawatts)	66.7	74.9

Colo	2020	2025
Number of core data centers	211	314
Total IT capacity per core data center (in megawatts)	68.9	67.1

Source: Uptime Institute Intelligence, 2020 (enterprise, n=7; telecom, n=1; cloud, n=5; colo, n=11)

Uptime Institute | INTELLIGENCE

Figure A1. **Respondent demographics**

The average IT capacity at core data centers run by colocation providers in our study is expected to drop by 2025, but the number of sites will grow. Many expect to be operating more, smaller facilities, including edge sites under 150 kW.

While the study is too limited to determine sector-wide capacity growth by 2025, at either the core or the edge, if our sample is representative, expansion is likely across all types of colo, cloud and telecom data centers. Enterprises will build and operate new facilities but in increasingly declining numbers.

UptimeInstitute® | INTELLIGENCE

ABOUT THE AUTHORS

Rhonda Ascierto 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: rascierto@uptimeinstitute.com

Paul Johnson is a Technical Consultant in the Europe, Middle East and Africa region of Uptime Institute Professional Services, responsible for data center design and facility Tier certifications. Before joining Uptime Institute, Paul worked in various technical and operational management roles within power distribution and power protection manufacturers, covering field service management, project management, commissioning, application engineering, research and development and systems engineering. He holds a bachelor's degree in mathematics and materials science and is a Chartered Engineer and a member of the IET (Institute of Engineering and Technology). Contact: pjohnson@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