

UI Intelligence report 22

Smart energy in the data center

Data centers could be more energy efficient and make better use of their infrastructures, if existing and emerging smart energy technologies were in greater use. But adoption has so far not matched innovation.

Lead Analysts

Andy Lawrence (alawrence@uptimeinstitute.com)

Rhonda Ascierio (rascierio@uptimeinstitute.com)

This Uptime Institute Research report covers:

Uptime Institute Research & methodology	3
Summary: Inevitable, but distant.....	3
Key findings.....	3
Introduction.....	4
Smart energy: Role and benefits.....	5

The technologies	8
Energy/power management and DCIM.....	8
Intelligent IT power management.....	10
Microgrids.....	11
Software driven, software defined.....	11
Shared reserve/adaptable redundant power.....	13
Smart data centers and the grid.....	15

Drivers and barriers	17
Economics.....	17
Manageability.....	18
Availability.....	18
Data center modernization.....	18
Constraints.....	19
External factors.....	19

Appendix	20
-----------------------	-----------

Uptime Institute Research and methodology

Uptime Institute Research 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 Research and this report, see the appendix or visit <https://uptimeinstitute.com/ui-research>.

Summary: Inevitable, but distant?

Uptime Institute's view is that the data center industry will strongly embrace smart energy eventually. But it will take time: the relative immaturity of many of the component technologies and the overall complexity and uncertainties of smart energy technologies will likely slow adoption, as will cultural factors and a commitment to existing investments.

Over time, the acceptability of consistently operating data centers at far below their peak capacity will change, and the case for smart energy technologies as tools to enable higher utilization and re-use of capacity will be more compelling. Further, the ability to integrate and interoperate many of the technologies involved will increase the return on investment, enabling the gradual development of integrated, system-wide smart energy platforms (consisting of both software and hardware).

But operators will continue to be cautious. Adoption of these newer approaches will initially be seen at the extremes of the market—small, distributed and highly automated and remotely managed data centers at one end, and, at the other, large hyperscale facilities, or campus environments, where the rewards are greatest. Wider adoption will be patchy but eventually steady.

KEY FINDINGS

- Smart energy is not one system or single technology: It is a broad umbrella term for technologies and systems that can be used to intelligently manage power and energy in the data center. Some will be deployed tactically, and some as strategic platforms. In time, most will link to each other.
- Smart energy systems promise greater efficiency and use of capacity, reduced maintenance, more granular control of redundancy, and greater real-time management of resource use.

- Smart energy management for data centers is still in its infancy. Deployment will be slowed and complicated by technical, cultural, organization, and economic factors. Multiple uncertainties make it difficult to predict adoption rates or patterns.
- Technical complexity of some relevant technologies, uncertain product roadmaps, current low adoption, and uncertainty over power pricing are all factors deterring wider take up.
- Wider adoption of other key technologies is a prerequisite for smart energy. The use of Lithium-Ion (Li-Ion) batteries in data centers, one such technology, is rapidly becoming more common, driving innovation of software to manage these systems.
- Smart energy technologies could effectively end the clear physical differentiation between 2N configurations and less-expensive N+1 configurations. Redundancy could become a dynamic figure and might be expressed more as a floating percentage (i.e., $N = 156\%$).
- Smart energy technologies such as adaptive redundant power and software-defined power enable the introduction of N-1 power topologies. This means that it may still be possible to allocate and draw on remaining reserves to maintain critical applications after a major failure.
- The architecture of power distribution will likely evolve over the next decade, with greater use of transfer switches, intelligent PDUs and breakers, supercapacitors, in-row batteries, and other devices enabling the gradual introduction of Smart Energy architectures.
- Smart energy systems will play a key role in enabling a variety of power and IT technologies to be optimized in support of an industry-wide, long-term goal to reduce generator use.
- UPSs are expected to become smarter and support more functions. They are also expected to be linked to the grid and to cloud services for two-way resource and data sharing.

Introduction

Smart energy—one of many trending concepts in a field crowded with competing and complementary terms, technologies, and approaches—is a big, aspirational idea with many applications and potentially profound implications. Smart energy is best known in a consumer/utility context, where automated management systems and smart meters monitor, report, and optimize energy use and demand. The “smart grid” is a distributed, intelligent collaborative web of resources and devices, managed by utilities for the benefit of customers and suppliers.

In the data center, “smart energy” is not a widely known or applied concept. It may be conflated with other concepts, such as “intelligent power management” or “software-defined power,” that may or may not have the same meaning. However, all these ideas, taken at their broadest and most aspirational, are similar and share the same ultimate goals of enabling data centers to lower energy consumption, reduce over-provisioning and capital costs, and maintain redundancy

and efficiency at reduced costs. These goals will be achieved through intelligent and dynamic control of energy use, distribution, allocation, and storage, and when economically feasible, selling power back to the utility without increasing risk.

Uptime Institute views smart energy technologies as a part of IT-managed automation, which is deployed widely and ambitiously in most industries. Increasingly, but still quite rarely, automation involves elements of artificial intelligence (AI). But smart energy need not involve large data sets, exotic algorithms, or “self-learning”—which are still futuristic for most; rather, much of the optimization involves simpler management and control systems.

The mission-critical industry has experienced a long trend toward highly-efficient, responsive autonomic, self-managing data centers that integrate a range of intelligent control systems. Smart energy, therefore, involves integration, data management, and control software; it has a tendency to become complicated, especially given that live and potentially dangerous power is involved.

Smart energy may become part of the trend toward autonomic data centers, but ultimately it may only be a more prosaic development, perhaps involving the simple policy-driven openings of some breakers and transfer switches to make a little more power available elsewhere, carefully managed reduction of power to servers that are under-used, or re-direction of workloads to servers that are, from an energy point of view, able to handle the load more efficiently. In all cases, such simple decisions require data about the status of devices, about the power supplies, and possibly about the loads.

In this report, we discuss certain technologies and ideas that manage the supply, consumption, distribution, storage, and allocation of power, according to policies, rules, or objectives (automated or otherwise) in support of business needs and goals. In most cases, the major expected benefit is cost reduction, but the ability to maintain service and provide new services is also a possible outcome.

Smart Energy: Role and Benefits

The purpose of implementing smart energy is to make power available where it is needed in the most economical way so that availability matches demand. This can be done by using facility infrastructure tools (better use of energy resources, batteries, UPS, and more flexible power distribution), by using IT power management tools and workload management, and buying/selling energy more effectively, through the smart grid.

Smart energy technologies largely differ from, and go beyond, many energy management or energy monitoring systems. They not only measure, report, and analyze power use and energy resources, but

they switch power and move loads. They differ from traditional software-controlled switches, in that they use intelligence, multiple data sources, and automation. Smart energy software tools are in their infancy, certainly in terms of adoption. The potential benefits are greater, as they combine both intelligence and automation to help the data center meet its business goals. In practice, the data center owner or manager is mostly concerned with three imperatives when assessing and deploying technology, including smart energy technologies:

- Improving or maintaining availability (including safety)
- Reducing total cost of ownership
- Improving or maintaining agility and opportunity (meaning the ability to offer new services to customers or change infrastructure, or operating models.

Smart energy technologies promise to help organizations achieve greater goals in all three areas by reducing the overall capital expense of a data center relative to the load it can support or by increasing the load it can support using the same infrastructure.

Smart energy technologies largely differ from, and go beyond, many energy management or energy monitoring systems.

Smart energy technologies have yet to be proven on a large scale, and their success in data centers will likely depend not just on the promise of the technologies, but also on how well the products are designed and deployed.

Doing so will require changing how data centers are planned and built. Data centers today are almost always built to allow for expansion and to support the highest demand peaks with much over-provisioning. A more cost-effective way to plan and build a data center is to build to the expected average power density, and then, if there are peaks in demand or density, to find ways of accommodating these using tactical methods. But this approach has traditionally been viewed as limited and risky and is rarely adopted.

Current data center construction practices lock in a lot of capital in power infrastructure, which is almost always very under-utilized. This is partly intentional (in order to build in spare capacity and redundancy), but it is also imposed by the rigid design constraints of data center power infrastructure, which tends to strand large amounts of power that is not accessible for wider use. This thinking has been generally successful in periods of growing demand and high margins, but competition (including from third-party data center services such as colocation, managed hosting, and cloud) is now forcing data centers to reduce their costs and engineer with value in mind. Smart energy technologies may enable some of these facilities to run much

nearer their capacities from the outset.

Successful smart energy technology implementations promise to:

- Increase the amount of available power (for a given cost)
- Enable equipment to run in optimal states
- Provide the ability to better foresee issues, using better automation and management data
- Possibly reduce single points of failure

All of the above should improve availability.

The same design concepts that promise to improve availability and reduce capital lock-in also promise to increase agility. The intent of smart energy is to enable critical facility managers/owners to treat energy and power infrastructure as dynamic resources that can be directed to meet demand. That demand may be for more redundancy for some applications and loads, more density and capacity for others, or feed-in power for the grid. Better energy storage, better switches, and much better analytics make this more possible than in the past.

Pioneers of smart energy technologies foresee many new applications and services.

The Technologies

The management of power by software is not new: Programmable controllers, which no data center could do without, use switching logic to manage power distribution, and power management and power quality tools have been widely used for many decades. Other mature technologies such as server power capping have not yet been widely adopted.


At the same time, the increasing use of Li-Ion batteries makes energy storage a more dynamic resource. In addition, IT has become more homogenous, with virtual machines and software containers enabling workloads to be moved and consolidated more easily. Most important, management and control software has advanced significantly, with new tools and systems available and under development. These three developments increase the potential of management of energy by software.

Among the technologies that comprise the smart energy category and may be used as components are:

Energy/power management and DCIM

Energy management is a generic term usually used to describe the monitoring, management, reporting, and forecasting of energy use in facilities, including the data center. Given the importance of understanding energy use at the facility level, the hall level, and often at the row/rack level, most operators have systems in place to collect and display this data in real time. PUE figures are likely to depend on such systems.

Maturing technologies combine to enable "Smart" energy in data centers



Source: Uptime Institute, 2018

Key Benefits of Smart Energy

- Reduced capex costs if overall peak power is cut**
- Reduced capex if infrastructure is “virtually” matched to the redundancy requirements of load**
- Reduced opex if overall power use is cut**
- Reduced opex if power use/prices are optimized**
- Increased revenue if UPS/storage participates in utility ancillary services**

- Increased resiliency if energy sources are distributed, granular**
- Increased resiliency if loads are moved dynamically on demand, for maintenance or anticipate failures**
- Increased resiliency for some loads if infrastructure is “adaptable” to needs**

- Increased ability to offer reduced redundancy at reduced costs**
- Increased ability to free up power for higher density, higher redundancy if needed**

The systems used to collect and manage energy data vary from operator to operator: Some collect data from power management/monitoring tools, some may use building automation or monitoring systems, which can be extended to monitor power use, and others use DCIM systems, such as Schneider Electric’s StruxureWare for Data Centers or Nlyte Software’s DCIM product, which collect and analyze data from a wide range of sources (including meters, power management tools, and IT management systems).

The term data center power management is more commonly used to refer to tools that manage power quality—to track voltage drift, drop offs, harmonics, and surges, for example. These functions are likely to be carried out by an EPMS (electric power management system) product from suppliers such as ABB, Eaton, Schneider Electric, and Vertiv.

Power analytics is sometimes considered as a related product category because these products describe how power use behavior can be modeled against an ideal performance model to identify faults or configuration errors.

Most on-premises DCIM software incorporates sophisticated energy management tools for integration, normalization, capacity monitoring and monitoring, and forecasting. But few are ever used for control, for reasons of security, complexity, trust, and redundancy.

DCIM and similar systems, therefore, are likely to play important roles in overall smart energy approaches, but they are not the complete solution. DCIM (and similar software that may not be called DCIM) are becoming firmly established as integration backbones, connecting to cloud data, IT workload management systems, and server monitoring systems and cooling/building systems to provide critical and timely data and analytics.

Uptime Institute Data Center Maturity Model			
LEVEL	DESCRIPTION	OPERATING EFFICIENCY	SOFTWARE
Level 5: Self-optimizing, autonomic	AI-driven integrated management software adjusts data center behavior and makes best use of resources according to goals, rules and service requirements throughout its lifecycle.	HIGH	AI-driven, integrated DCIM with automation
Level 4: Optimizing	Physical and virtual IT and data center subsystems integrated; models used for prediction, service management and multiple views, optimizing in near real time. AI is applied to DCIM-based data lakes for advanced analytics.	MEDIUM	AI-driven, integrated DCIM
Level 3: Proactive	Physical data center equipment characteristics, location and operational status is tracked. Energy and environmental data is used to reduce risks, waste.	MEDIUM	Integrated DCIM
Level 2: Reactive	Software installed to monitor environmentals and equipment power use. Ability to adjust basic controls (i.e. cooling) to demand.	LOW	DCIM monitoring
Level 1: Basic	No integration of infrastructure data. Basic monitoring supplied with equipment. A reliance on BMS data. Simple alarming, error messaging.	LOW	Ad hoc

Source: Uptime Institute Research and 451 Research, 2018

Over time, data centers are gradually becoming more heavily instrumented and automated. The most automated data centers will be “self optimizing,” with smart energy systems eventually playing a key role.

Intelligent IT power management

Intelligent IT power management systems are intended to reduce power consumption by dynamically moving workloads, powering down equipment or putting it in lower power states, and by capping or reducing voltages and frequencies in such a way that services are not affected. Over the past two decades, vendors have developed and offered a range of technologies that enable one or more of these functions.

Reducing IT power demand not only saves energy, but also saves on cooling and power distribution losses. The most successful techniques for managing IT power have been virtualization and software containers, improved processors and memory chip design, and automatic scaling down of power consumption at low utilization rates at the chip level.

Adoption of intelligent IT power management systems has been very low, and innovators have not enjoyed commercial success. IT staff fears service risks from turning servers off or slowing them down with capping, and they rarely care about power consumption. Facility management often has no visibility into or control over IT systems, which might help them address these issues.

These technologies have not disappeared, but rather are being reintegrated, under the hood, into smart energy approaches. For example, suppliers and some operators are seeking to gain a better understanding of workload mission criticality and power requirements, and, in response to current needs or policies, either move workloads or move available power (reallocating or introducing redundant power) to match workloads. In modern environments, servers are (to a point) able to reduce their own power use automatically when workloads have been moved.

Microgrids

Microgrids may be defined as a localized group of electricity sources and loads that are synchronous with the traditional electrical utility grid (macrogrid) but can disconnect from the rest of the grid and function independently as conditions dictate.

At a simple level, it may be argued that all data centers that can operate without the grid are microgrids; however, few data centers can actively switch between power sources and manage loads as microgrids do. But there is some convergence with data center smart energy technologies and microgrids, especially in the largest sites, as interest in renewable onsite or near-site energy, transactive power, and demand response grows.

These technologies have not disappeared, but rather are being reintegrated, under the hood, into smart energy approaches.

Microgrids use control software and intelligence and can be programmed according to policies, with many of the controls embedded in hardware including transformers, inverters, switchgear, and batteries. They must manage switchovers between different sources (such as wind, solar, fuel cell, and the utility grid), transforming and rectifying current, transferring loads, and synchronizing phases, voltages, and frequencies.

A centralized control system (software) balances the power supply and demand and controls transfers and switching; but the need for resiliency and Concurrent Maintainability means that the most effective architectures will incorporate a distributed control system (DCS), thereby avoiding a single point of failure. This, in turn, is most effectively achieved by deploying distributed and resiliently networked control units in front of each of the main power sources and often the facility's PDUs.

It is not surprising that data centers have not developed or entered into microgrid arrangements, given the complexity and the cost of supporting and synchronizing multiple power sources and of applying policies and controls, especially with the current ethos of value engineering. But there are some applications where the huge and largely stable load of a data center can help to smooth out uneven demand and supply and justify investments in microgrid technology, especially in campuses and smart cities but also inside some large data centers.

Software driven, software defined

Software-defined power (SDP) is a relatively new term in the data center lexicon that arguably could be used interchangeably with smart energy. Its main proponents are leveraging the terminology, and indeed commercial success, of the software-defined data center (SDDC), an IT technology that treats processing, storage, and especially networking as virtualized and manageable resources.

Software-defined architectures operate at a layer of abstraction that makes it easier to match resources with changing needs using data, management tools, and automation. Some of the logic and complexity is stripped out of specific devices and handled by a management system, which lowers costs and improves performance. In a network and IT context, this is a significant and empowering change. The “control plane” communicates down, to devices through APIs, and up to applications that may gather data, control loads or routes, or apply business and technical policies.

In theory, significant cost savings can be achieved by treating power as a virtual resource that is managed through a similar control plane. Ultimately, this could mean that data center managers, using a suite of software tools (applications), have some centralized control over how power is used, distributed, capped, stored, or even sold to meet changing demands, service levels, or policies.

The small California software supplier Virtual Power Systems (VPS) is a strong proponent of SDP. But although its long-term ambitions lie in software, its Intelligent Control of Energy (ICE) platform is deployed as a mix of software and hardware—a tactical move that reflects the fact that most power devices in a data center are not designed to operate as part of an agile, software-driven environment.

VPS bases its system on the idea that there are two levers to pull in an SDP data center:

- The power sources themselves (supply side)
- The IT (the load or demand side).

The system includes distributed rack- and row-based Li-Ion batteries, which can be used to meet extra, peak power demands or in support of workloads that need power the most. In recognition that most data centers have a centralized UPS, the VPS system positions small, distributed ICE units in the racks, which can be used to shut UPS power to certain racks, allowing the batteries to take over. This preserves UPS power, allowing for it to be used elsewhere, or to provide a longer ride through time, or greater level of redundancy to other applications. VPS promotes the ability of its system to solve problems of stranded power (power can be “moved” from one rack to another) and for adding power redundancy to some racks at the expense of others that may need it less.

VPS provides a suite of tools that can be used to show the status and requirements of loads, the level of charge in batteries, and for the operator to set policies for use.

In order to operate most effectively, an SDP system, like any smart energy technology, needs to understand the loads and what protection they need. Other SDP-based applications may depend on knowing power prices, for example, at given times (for demand shifting), or the likely power draw of servers at different rates of utilization. The management system is the critical component. Such technologies will benefit from ever greater use of machine learning techniques—although this is not necessary for the simplest deployments.

At present, there is high interest but relatively low commercial adoption of SDP. This is still an early-adopter technology, with concerns over initial cost, uncertain RoI, and possibly the introduction of Li-Ion batteries in the white space.

SDP could effectively end the clear physical differentiation between 2N configurations and less-expensive N+1 configurations. Instead, the level of power redundancy would be set by policy, according to workload needs, and even adjusted in real time. Effectively, the N level becomes a dynamic figure that might be expressed more as a floating percentage (N = 156%).

Another benefit is that it becomes possible to go to N-1 in a failure.

For colocation data center suppliers, in particular, this floating N level may be attractive, enabling them to sell power more flexibly and profitably. But it presents sales challenges: most colocation customers are used to buying against simple and usually inflexible service level agreements and have not clearly defined what reduced redundancy might mean to them, especially at a workload level. This capability may yet prove to be a very important advance, but it may prove too complex, at least at first, for most colocation customers.

An informal Uptime Institute assessment suggests that an SDP system using Li-Ion batteries can be configured to conform to Uptime Institute Tier III (Concurrently Maintainable) and Tier IV (Fault Tolerant) requirements, although adding redundancy to a rack or row using diverted UPS or battery power does not necessarily make a system more available in the event of component failure—that will depend on the on-site configuration, including establishing and maintaining the necessary power paths required by the Tier objective to achieve Concurrent Maintainability or Fault Tolerance.

For some situations, the key to the success of SDP is prioritizing the right mix of workloads, which requires establishing a hierarchy of criticality—something that most data center operators do not necessarily have the visibility, tools, or processes to do. Data from several tools and systems must be integrated (typically from DCIM, ITSM, and virtual machine or container management, at a minimum), which often requires facilities and IT departments to work closely together. Interdisciplinary business issues can also arise.

Shared reserve/
adaptable
redundant power

Two companies, Vertiv, the data center equipment and services company, and i3, a data center engineering company, have both (separately) presented similar ideas around the idea of using transfer switches and an intelligent control system to pool and share the power of a combined number of large UPSs.

Vertiv described its Shared Reserve approach in 2017 and has been quietly working with some customers to enable a move from a 2N architecture for UPS to lower N+1 architectures, while still maintaining high levels of redundancy for workloads that are critical. This is achieved by sharing capacity between power sources under software control. In a more detailed paper (late 2018), i3 data center engineer Ed Ansett presented a similar approach, the adaptable redundant power architecture.

Adaptable redundant power takes the idea of a block redundant (N+1) architecture, makes use of static transfer switches to switch power from one UPS to another and adds extra transfer switches and a management system for allocating and connecting power to workloads. The effect is to create a “virtual” pool of power, made up of the combined parallel resources from multiple UPSs. The power can be intelligently and, in some situations, automatically allocated to particular racks or loads.

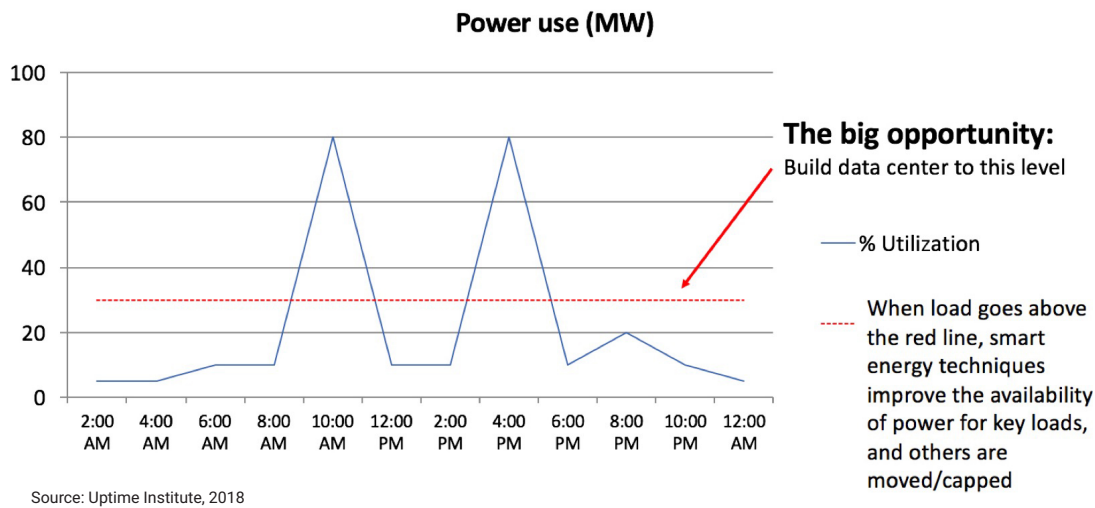
The data center industry is quickly moving toward a more transactive data center power management arrangement, using distributed energy resources.

This approach is somewhat similar to the SDP approach. Primarily, power infrastructure is no longer over-provisioned in some places and under-provisioned elsewhere: it reduces, if not eliminates, stranded power. In addition, it becomes possible to allocate dynamic and fractional levels of redundancy to particular loads, so one load may effectively be N+1, one 2N, and a third 1.5. This should, or could, mean that some loads get less protection, but are charged less.

Another benefit is that it becomes possible to go to N-1 in a failure. This means that once a load’s primary UPS has failed, or even several have run down, it is possible to favor a critical workload by allocating all or some of the remaining power to it. This prevents the hard shutdown that usually occurs, when stranded power remains stranded and cannot be directed to a workload.

Whereas VPS uses breakers controlled by its Intelligent Control of Energy units (and Li-Ion batteries to pick up the load), these shared reserve/adaptable power designs use intelligently controlled static transfer switches (STSS). This approach aligns more closely with current data center design in one way: batteries are needed in the white space. But it undoubtedly produces some challenges: STSS have caused many problems in the past, and although they are now much more reliable, they are seen as active devices that could fail. A well-thought out Concurrently Maintainable architecture, with dual power supplies for redundant pathways, will be important. This is certainly achievable, but as elsewhere in the field of resiliency, lowered costs and increased flexibility come at the price of increased complexity at the software/management level.

Smart energy can enable data centers to build to average, not peak, loads



Smart data centers and the grid

The data center industry is quickly moving toward a more transactive data center power management arrangement, using distributed energy resources. There are hundreds of new companies that specialize in demand response, in building and operating large banks of batteries, and in offering tools and services to manage and match supply and demand. Operators with significant investments in energy storage and generation (such as generators, Li-Ion batteries, and fuel cells) can participate in many ways—for example, trading power for micro seconds to help smooth power quality, or running generators at peak times to lessen the total load.

Data center operators have been able to participate in demand response or load reduction/shedding programs for several decades. Under these programs, usually organized through aggregators or utilities in power-constrained areas or where peaking power is expensive, end users are offered financial incentives to take their facilities off the grid at times of peak demand. In some countries, especially in the US, a large market has grown up for selling so-called negawatts, on-site power sold to the distribution utility for local use. For utilities, it is cheaper and often cleaner to pay generously to take demand off the grid by shedding loads, rather than operating peaking power plants, which are generally inefficient and expensive for utilities to run.

Very few data center operators participate in such programs. Until recently, the only practical way to help shed load from the grid in a financially meaningful way has been to run engine generators. Although some operators have been paid significant amounts for participating, for most an equally strong reason for participating is to exercise their generator failover processes. But many are concerned about increased risk, engine operating hours, and environmental issues. In recent years, environmental rules in areas of the US have forced some data center operators to drop out of demand response programs.

But there are signs that a significant revolution is on the way, with the possibility that data centers will participate more transactively and intelligently with the grid; and, furthermore, that cloud-based services will help to intelligently and optimally manage on site and behind-the-meter resources.

The rapid development and adoption of Li-Ion batteries offers new opportunities that go beyond generator-based demand response. When a battery energy storage system (BESS) or UPS is combined with smart energy management system, this may become a key component of a system that manages internal critical power and resources and supplies external (utility facing) services. Recent changes to the law in the US (FERC 841) enable owners of battery storage to participate in power markets with the same freedom and benefits as wholesale suppliers. Some see opportunities for the UPSs supporting multiple critical loads (for example, in adjacent data centers or across a campus) to be combined and managed as if they were a microgrid.

Some operators who have switched to Li-Ion UPS are working on projects to feed some power back to the grid. (Microsoft, for example, has a project in Seattle to explore doing this at scale). In Sweden, Eaton launched a UPS-as-a-reserve service with Nordic energy company Fortum. In the UK, Vertiv, in conjunction with the software company and energy aggregator Upside Energy, has launched a similar service to facilitate the sale of energy back to the grid. Suppliers are currently working with most data center customers under early stage non-disclosure agreements.

The use of such services in demand response and other programs is likely to be more complex and automated than generator-based demand response, which must be scheduled hours in advance. UPSs will be connected and managed by a service partner, and policies for the use of power are based on algorithms for calculating availability/capacity, as well as agreed policy. In effect, the power is discharged when needed for peak shaving for the grid or to stabilize the grid and purchased when it is cheap and stable, while software constantly monitors reserves and loads to ensure that availability is never compromised. The software can also forecast demand, calculate when it best to discharge economically and, ultimately, may even be able to link IT systems to initiate the movement of IT loads or traffic. If operators start to build complicated adaptive power architectures for their UPS configurations, the configurations will be challenging.

Drivers and Barriers

The vision and the promise of various smart energy technologies is considerable and far reaching. But proponents of the technology still face adoption challenges that are technical, cultural, and economic.

In our 2018 Uptime Institute/451 Research project, “Disruptive Technologies in the Data Center,” two separate panels of experts and end user/operators gave both SDP and microgrids low scores for their likely adoption and disruptive potential. We identified the following five areas that, taken together, will determine the rate of adoption of smart energy technologies.

Economics

The proponents of the technologies in this report argue that they significantly reduce either the capital budget of a data center, the operating costs, or both; further, the technologies can bring revenue. The reduction of over-provisioning or the use of over-provisioned capacity, in particular, is a benefit common to all smart energy technologies.

Uptime Institute has not carried out detailed TCO or ROI assessments of these technologies. But clearly, each situation is different; earlier studies with DCIM showed that, while some big returns are possible, some of the benefits may be soft or difficult to quantify. Moreover, the numbers in favor of one technology may depend on multiple use cases and benefits, which may need to be planned at the outset. Some benefits may depend on the deployment of another technology, such as Li-Ion batteries.

There are also different functions or classes of customer who may benefit more than others. Big operators, such as very large colos and internet giants, may see large savings on power infrastructure, achieved by reducing overall centralized power demand. These benefits may encourage adoption of dynamic shared reserve power, SDP, and reserve power as a service. Even so, to save on a scale that will justify investment, some large investments are needed that will have a long payback. As always, the best time to invest in a new data center technology is at the build stage.

But there may be pockets of interest elsewhere. Colocation customers, for example, may find they can quickly reduce their power use, and any penalty fees for exceeding their limits, by using power capping or SDP technology (using Li-Ion batteries). Such customers might prove an unlikely category of early adopters, if suppliers can make their case.

But this may depend very heavily on the agreements that they have with wholesale colocation operators.

Manageability

As data centers become bigger, more complex, and increasingly interconnected with each other, the ability of managers to see what is occurring, and to be able to be proactive, and reactive, can become more difficult. DCIM (or similar DCM-like software), arguably a complementary or enabling technology of smart energy, is now recognized as a necessary component for automated, efficient data centers for this reason.

Availability

Data center operators have continually proven that they consider availability above almost every concern, often in spite of very unfavorable economics. Smart energy technologies, therefore, will need to demonstrate that they are not adding significantly to risk and can, where needed, match the proven levels of availability achieved by simpler, hard-wired physical equipment. If there is any loss of planned availability, there needs to be a very clear economic justification. Suppliers of all types of smart energy technologies have gone to considerable lengths to ensure this is the case, with built-in and hard-wired policies and safety nets for those applications that cannot suffer any degradation of risk.

Even so, suppliers of technology that reduce the absolute levels of over-provisioning and redundancy face a marketing challenge. Even where the costs are significantly lower, many operators will still opt to deploy expensive and largely unused redundant capacity; further, their external or internal IT customers will continue to expect the highest availability on a sloping cost gradient and will often need to be persuaded to segment their applications by mission criticality needs in order to reduce costs.

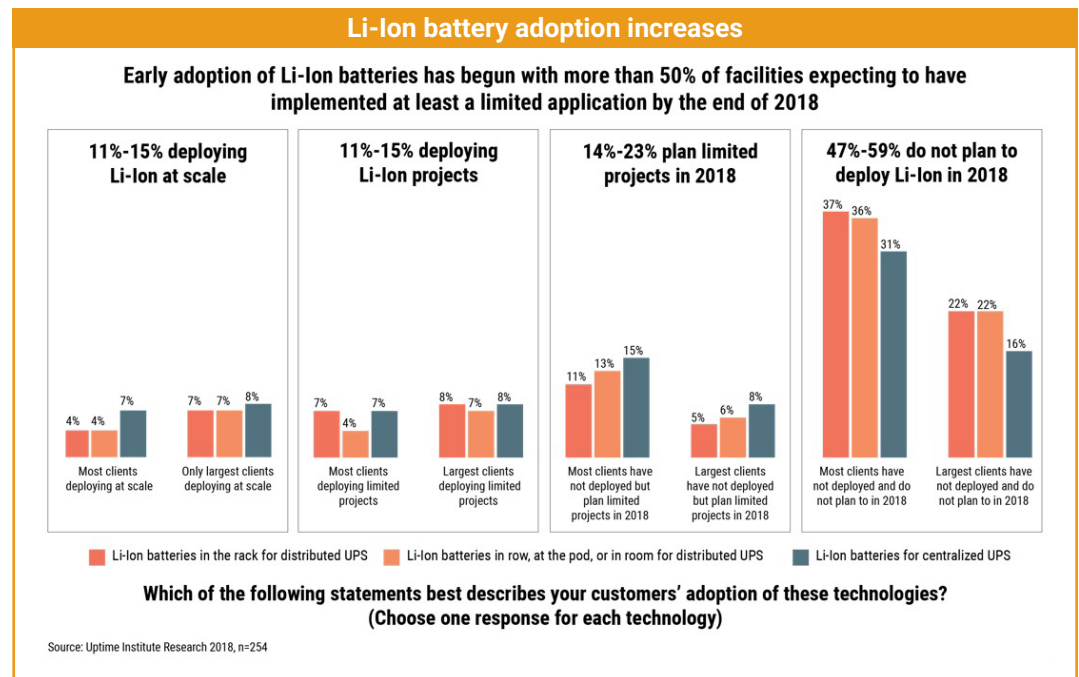
Data center modernization

Not every technology is installed with immediate cost justification in mind. One of the biggest drivers to the adoption of smart energy technologies is strategic—a desire by management to establish an engineering lead that will constitute a competitive advantage. This accounts for some of the investments made by at least two internet giants and one large colocation company.

Those who are already investing in on-site renewable power (as many now are, according to Uptime Institute's survey data) or distributed power architectures (such as the Open Compute Project) are the most likely to see value in intelligently managing power.

Constraints

Many data centers are either space or power constrained, with metropolitan colocation data center suppliers, in particular, finding their margins depressed by low power densities. The move to Li-Ion batteries is helping some of these increase their white space, while peak shaving and smart energy techniques can help to reduce power consumption (and costs). The move to Li-Ion storage, which is still more capital expensive than lead acid batteries, can be justified or supported by selling capacity, while the use of distributed battery can enable increases in total power capacity without expensive UPS and power system upgrades.



Uptime Institute Research shows that Li-Ion batteries, an enabler for smart energy, battery energy storage as a service, and software-defined power, are rapidly gaining traction.

External factors

External factors have so far not helped those offering advanced energy-efficiency technologies. Power prices have been largely stable, or have fallen, in many geographies, and (despite expectations) cloud technologies have created uncertainties that have held back investment in infrastructure, and environmental constraints and legislation have been minimal.

But other factors are likely to prove more favorable. Li-Ion and other battery technologies are becoming cheaper every year, lowering the cost of ownership, and there are renewed efforts to lower data center energy requirements by internet giants and various governments. Similarly, there is growing nervousness about the reliability of grid supplies and of grid capacity, which is encouraging operators, and investments, to promote more on-site energy security. The push to renewables also means that the macrogrid has a greater need for distributed, local storage. Smart energy technologies provide a platform to support improved both improved efficiency and less grid dependency.

Appendix

Uptime Institute Research and methodology

This report is written by Uptime Institute's research team, based on interviews with and visits to suppliers, data center operators, and other experts, as well as analysis of both public and private materials, including our own survey data. This report is funded entirely by Uptime Institute.

Uptime Institute Research is an independent unit of Uptime Institute dedicated to identifying, analyzing, and clearly explaining the trends, technologies, operational practices, and changing business models of the mission-critical infrastructure industry, so that executives, investors, and operators may make considered, well-informed decisions, adopt best practices, explore new opportunities, and reduce their exposure to risk.

Uptime Institute Research has a dedicated team of experienced researchers, supported by a global network of data center and IT consultants from Uptime Institute. Uptime Institute Research conducts surveys, holds regular reviews, interviews industry experts and advisors, and draws on insights from research colleagues at its sister company 451 Research.

Uptime Institute Research delivers information and insights in the form of reports, consulting, conferences, online portals, roundtable discussions, and webinars. Full, premium reports are primarily available through the Uptime Institute Inside Track portal, by direct subscription (contact Uptime Institute sales), or, in certain cases, on Uptime Institute's web site.



ABOUT THE LEAD ANALYSTS

Andy Lawrence is Uptime Institute's executive director of Research. Mr. Lawrence has built his career focusing on innovative new solutions, emerging technologies, and opportunities found at the intersection of IT and infrastructure.



Rhonda Ascierio is VP of Research at the Uptime Institute. She has spent nearly two decades at the crossroads of IT and business as an analyst, speaker, adviser, and editor covering the technology and competitive forces that shape the global IT industry.

Ms. Ascierio focuses on innovation and disruptive technologies in data centers and critical infrastructure, including those that enable the efficient use of all resources. Before joining the Uptime Institute, Rhonda was research director of Datacenters at 451 Research. She joined 451 Research from Ovum where she was a senior analyst in energy and sustainability IT. Previously, she had a successful career as an IT reporter and editor, mostly in Silicon Valley, at a number of well-known industry publications.

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

Uptime Institute – The Global Data Center Authority®, a division of The 451 Group, has office locations in the U.S., Mexico, Costa Rica, Brazil, U.K., Spain, U.A.E., Russia, Taiwan, Singapore, and Malaysia.

Visit www.uptimeinstitute.com for more information.

All general queries:

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