



WHITE PAPER

Data Centers and Advanced Microgrids

Meeting Resiliency, Efficiency, and Sustainability Goals
Through Smart and Cleaner Power Infrastructure

Published 4Q 2017

Commissioned by Schneider Electric

Peter Asmus

Principal Research Analyst

Section 1

INTRODUCTION: SETTING THE STAGE

1.1 The Impact of the New Digital Economy on Data Centers

The forecasted growth of Internet of Things devices and the rapid uptake of cloud-based applications to manage the emerging digital economy speak to a fundamental fact: the role data centers will play in daily life will continue to grow exponentially over time. Millennials are growing up in a world where data drives their lives. One million new devices are expected to come online per hour by 2020. By 2025, 60% of computing is expected to take place in the cloud. Entertainment, fundamental comfort, communications, and even transportation are all being interwoven, with systems dependent on massive amounts of high speed data transfers. Data centers need to maintain uptime under all circumstances.

The digitization of products and services is largely dependent upon them having reliable electricity. If done right, the shift to electricity can offer enormous advantages in terms of resource productivity and a more sustainable economy. If data centers do not take advantage of new energy technology innovations, massive inefficiency could doom the so-called fourth industrial revolution¹ promise of sustainability via digital innovations. The ramifications for developing countries and the rest of the world are significant in terms of climate change impacts if measured on an emissions per capita basis.

1.1.1 Data Centers Are Emerging as a Target for New Energy Innovations

One only needs to review the history of telecommunications to see what the future holds for utilities as well as data centers, a class of customers that has historically ranked among the most important customers of former electric utility monopolies. With large round-the-clock customer loads that sometimes require dual redundant feeds, what is not to like?

In today's world, however, notions that once held sway regarding energy are rapidly giving way to a new reality. The energy industry is being transformed; smaller and smarter is better. Business models that shift the focus from ownership of assets to real-time virtualization of energy as a service represent the highest value. Think of Uber, Airbnb analogies—a world of connectivity all residing on a mobile device. All these systems rely upon real-time delivery of data from an immense global network of different types of data centers.

¹ World Economic Forum, “Center for the Fourth Industrial Revolution,” World Economic Forum, www.weforum.org.

Looking to telecommunications, it appears that the changeover for new technology leaps forward every 5 years. That is how long it took for each step in the evolution from landline to cell phone to smartphone. Given recent advances in distributed energy—including advanced batteries, smart inverters, and automation controls—similar leapfrog developments can be expected in the electrical sector. Will data centers remain risk averse customers of the status quo—relying upon redundant utility feeds and onsite backup systems, including generators and uninterruptible power systems (UPS) for each critical load—or will they embrace more efficient and sustainable networks such as advanced microgrids?

No doubt, data centers are pioneers in corporate sustainability. In 2016, data centers signed contracts to purchase 1.2 GW of renewable energy. Consider the following examples of data centers that have turned to either renewables or clean burning natural gas for powering up their data centers:

- Salesforce: 100%
- Switch: 100%
- Apple: 87%
- Facebook: 74%
- YouTube: 71%

The Business for Social Responsibility's Future of Internet Power² program features 12 large data center companies representing a trillion dollars in market cap. Each of these Internet giants have committed to reducing their respective carbon footprints by adhering to green energy protocols embodied in a set of corporate colocation principles. On the surface, these statistics and efforts appear impressive. However, these achievements for the most part rely upon the purchase of renewable energy credits (RECs) to offset actual emissions from onsite distributed generation such as diesel generators and utility grid power still often led by fossil fuel combustion. One could argue such transactions represent financial innovation. However, would it not make more sense to rely upon onsite sustainable energy supplies and technologies, freeing up large wholesale renewable supplies to provide net-positive benefits not correlated with mitigation of polluting generation sources in the first place?

² <https://www.bsr.org/collaboration/groups/future-of-internet-power>.

“Smarts not parts,” is how Patrick Flynn, senior director of sustainability for Salesforce, sums up the necessary shift in thinking. “Software is the key. Don’t really need two of everything to insure resilience.” He is quick to add that the investments Salesforce has made to reach its goal of 100% renewable energy, net-zero greenhouse gas emissions target, and a carbon neutral cloud did stimulate the development of two new wind farms at a fixed price, rather than REC purchases from existing facilities. “When we look back in time 5 years from now, data centers will be viewed as pioneers with microgrids, but I believe they’ll also be forging tighter relationships with utilities, especially those in rural geographies such as Oregon and Texas, where data center loads and the jobs they create will be particularly appealing.”

While other facilities such as military bases, hospitals, and universities are exploring the concept of an advanced microgrid, data centers have been slow to warm up to the concept. Navigant Research expects this hesitation is about to change as the commodity costs of hardware assets drop while the capabilities of automation and software controls scale up. Two recent announcements from large data center operations underscore that early adopters may be seeing the light:

- Microsoft announced in September 2016 that it will run its data centers off of 50% renewable energy by 2018. More relevant to microgrids, the company also announced a pilot program with Primus Power, an advanced flow battery technology provider. In partnership with NRG Energy and the University of Texas at San Antonio, Microsoft aims to validate the ability of these long-duration batteries to act as a grid resource to improve the usability of onsite renewable energy resources.
- Equinix, the world’s largest provider of colocation services, announced in August 2017 that it will purchase 40 MW of Bloom Energy fuel cells, which could be considered single resource microgrids, to be installed by 2019. This shift to clean onsite distributed generation will reduce emissions associated with power otherwise purchased from utility grids by 20% to 45% while also saving 87 billion tons of water used for cooling centralized fossil power plants. Southern Company, a large utility previously focused on coal-fired centralized generation, is also a partner on the project.

1.1.2 Data Center Growth Trends

Tracking the growth of data centers is a bit of a zig-zag exercise, with some market segments shrinking while others pick up steam. If one focused on annual capacity additions,³ service provider data centers will add almost 7,000 MW of capacity in 2018, the

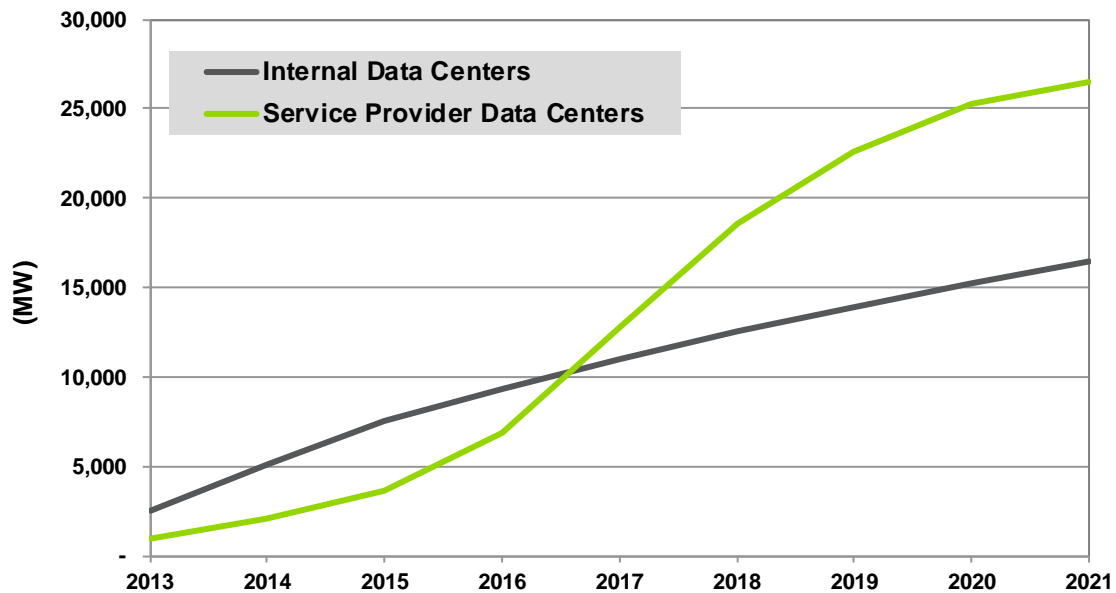
³ This analysis is based on data from the IDC report *Worldwide Datacenter Installation Census and Construction Forecast, 2017-2021*. The IDC data was converted from millions of square feet to capacity by using an average value of 150 watts/million square feet based on average energy consumption estimates provided by Schneider Electric.

fastest pace of expansion recorded. According to IDC, this pace of growth of the overall market will slow in the following year for both internal and service provider data centers.

Upon closer examination of cumulative market growth, there is a shift away from ownership (internal data centers) to reliance upon service provider data centers, which includes colocation facilities. This shift is driven by smaller data centers relying upon a public cloud, and wanting to shift responsibility for data centers to third-party experts. As one close observer put it, “Mom and Pop data centers are moving to the public cloud because data center infrastructure management is not their core competency. They don’t want to own their data center internally. On the other hand, large traditional companies, such as banks and insurance providers, had to build their own data centers, but now they too are looking at colocation and service provider options.” The sharp peak and then decline of service provider data centers identified in the IDC data on an annual basis reflects consolidation and a shift toward cloud-based applications, which translates into less physical infrastructure going into the ground.

The data center industry is at a crossroads. The types of data centers that will come online in the future may feature attributes much different than what was the status quo in the past. A clearer picture emerges if the data is sorted by cumulative capacity (see Chart 1.1). Upon examination of the entire data center fleet, the installed base of data centers is shifting away from internal data centers, with service provider data centers taking a leadership position from this point forward.

Chart 1.1 Cumulative Data Center Capacity, World Markets: 2013-2021



(Sources: Navigant Research, IDC)

Data centers are moving away from direct ownership of hard assets and seeking resiliency and efficiency via the cloud and the virtualization of energy services. This is setting the stage for creative solutions providers to offer up new technologies via new business models that can mitigate risks while increasing performance through options like advanced microgrids. This Navigant Research white paper explores the implications of this transformation on the data center industry.

Section 2

DATA CENTERS AND MICROGRIDS: FINDING COMMON GROUND

2.1 Mutual Goals of Data Centers and Microgrids

Navigant Research defines a microgrid in the following way:

A microgrid is a distribution network that incorporates a variety of possible distributed energy resources (DER) that can be optimized and aggregated into a single system that can balance loads and generation with or without energy storage and is capable of islanding whether connected or not connected to a traditional utility power grid.

Microgrids are often touted as networks of DER that can bolster resiliency. The 451 Group, a leading source of data for the data center industry, defines resiliency in the following way:

Resiliency describes the extent to which a system, digital infrastructure or application architecture is able to maintain its intended service levels, with minimal or no impact on users or business objectives, in spite of planned and unplanned disruptions. It also describes the ability of a system, infrastructure or application to recover full business operations after a disruption or disaster has occurred.

Data centers and microgrids: one in the same? Depends upon who you ask. IBM believes its global fleet of data centers qualify as microgrids. “We ensure continuous operation of the IT load, that is our design parameter for all of our data centers,” said Frank Dailey, global data center energy coordinator. The company is always exploring new technologies such as distributed generation (DG) technologies including fuel cells and solar PV at all of its sites, not just data centers. “We rank all potential investments in our data center portfolio based on innovation and return-on-investment,” added Dailey.

The average power outage cost for a data center in 2015 was \$740,357, a 38% increase in the cost of downtime compared to 2010. This figure comes from the Ponemon Institute’s annual report.⁴ In this case, 63 data center organizations in the United States were polled that have experienced an outage in the previous 12 months. The failure of uninterruptible power supply (UPS) systems is the number one cause of unplanned data center outages, accounting for one-quarter of all outages. Perhaps the most disturbing statistic found in the

⁴ Emerson Network Power and Ponemon Institute, “2016 Cost of Data Center Outages,” Ponemon Institute, <https://www.ponemon.org/blog/2016-cost-of-data-center-outages>

*Eaton Blackout Tracker*⁵ annual report is that the increase in maximum downtime costs rose 81% in 2015 compared to 2010, amounting to \$2.4 million.

A UPS has been viewed as an insurance policy toward grid outages. Yet it also does seem worthwhile to explore other avenues to the end goal of resiliency and uptime, including microgrids which can leverage onsite assets more efficiently, while also insuring uptime at reduced economic and environmental cost.

Not everyone agrees with this assessment. Chris Brown, chief technical officer for Uptime Institute, does not see a decline in reliance upon UPS systems and diesel generators in the near future. Uptime Institute has created de facto data center industry standards that range from Tier I to Tier IV, with the latter representing the highest possible resiliency. “Human beings have an almost emotional attachment to their diesel generators as they give data center owners and operators both comfort and a form of insurance. However, many data centers still need engine generators as they must be ready to operate through unforeseen conditions. As the data center industry embraces distributed computing more, there may be less of a technical need for engine generators for sites that are able to embrace distributed computing across regions, but engine generator usage will likely hang on as the emotional tie and the form of insurance will still be present.” He noted that the amount of space required for solar PV and the resulting power output may limit reliance upon such weather-dependent resources “to less critical loads.” He does, nevertheless, see a trend toward reliance upon larger batteries in part of the market that embraces distributed computing, “systems that can provide 1 hour of storage instead of just 5-10 minutes.” He also wonders if the whole industry is really moving to the cloud. “Government facilities may still have high security needs, so they may have reservations about the public cloud, and perhaps even a private cloud.”

Lance Black, founder and CEO of Vazata, a managed cloud provider for colocation data centers, is trying a new approach with a data center being developed in McKinney, Texas. The data center lacks both onsite backup diesel generators and a UPS. As an alternative, it is relying upon lithium ion (Li-ion) batteries charged up by the local utility. “We decided to take this route because it was the most efficient use of resources. It’s all about cost,” said Black. He added, “I had become frustrated by seeing batteries in the basement and an array of generators and the only time they were ever used was when we were doing maintenance mode. We thought there was a better way to reduce costs for our customers.” He investigated onsite solar PV, and may still go that route, but instead started with the batteries in this modular approach. The data center is starting at 0.5 MW, but the plan is to grow this modular system incrementally over time. “Two to 3 years from now, we hope to

⁵ “Blackout and Power Outage Tracker,” Eaton, <http://powerquality.eaton.com/blackouttracker/>.

be adding solar PV,” he acknowledged. “We are also selling back power to the local utility to reduce bills.”

2.1.1 Identification of Mission Critical Loads

Defining what is and what is not a mission critical load is fundamental to a good facility design. Upfront design and planning tools are therefore vital to developing an efficient and reliable power infrastructure.

One way to capture the performance of any power supply infrastructure, including data centers, is the term criticality. A data center’s criticality has arguably the strongest influence on lifetime total cost of ownership (TCO). For example, a fully redundant (2N) power architecture could more than double the 10-year TCO of a non-redundant (1N) power architecture. Although a significant cost penalty for 2N power is the doubling of electrical and mechanical equipment capital costs, the greater influence comes from energy costs associated with operating and maintaining this power equipment at 2N. Therefore, when choosing a data center criticality, a data center designer or owner needs to weigh both the costs and the criticality to establish a true cost/benefit analysis.

Many smaller microgrids may be better than a single large one. Likewise, redundancy is not the only way to protect mission critical loads. Diversity is another tool, as is efficiency.

2.1.2 Focus on Uptime

It is all about uptime when it comes to a data center. Given the high cost of an outage for an IT load, data centers have justified significant investment in multiple redundant systems including armies of backup assets included in the forms of UPS systems as backstops.

As one utility representative noted, “Every second is money to a data center.” This is one reason data centers do not often participate in demand response (DR) programs, he observed. Often, the amount of revenue generated from participation in DR event is minimal when compared to the large costs of a power outage perhaps created by participation in such a program. Furthermore, the acronym DR also means “disaster recovery” to most data centers operators, further tainting the term in their minds. Yet another challenge is that many backup diesel generators are only allowed to operate during an emergency due to their high emissions profiles. Utility DR programs are not considered an emergency.

2.1.3 Leveraging Existing Utility Grid Infrastructure

Data centers are often sited in locations where existing utility power lines offer pockets of superior reliability. In fact, data center owners often seek out sites that fall into the cracks between different utility service territories, allowing for interconnections by the data center to two separate utilities. The rationale for seeking out such sites is that if one utility grid goes down, chances are the other will still be up and running.

The incumbent electricity grid is a resource that should be leveraged as much as possible. While it may be economically prudent for a data center to shift to its onsite backup power systems to avoid utility demand charges or to anticipate possible threats to reliable power—such as the approach of a hurricane—utility and air quality rules and regulations often preclude such moves. The environmental impacts attached to most backup diesel generators severely limits the ability of data center owners/operators to take advantage of emerging revenue generation opportunities that are sure to grow as markets for utility ancillary services mature and proliferate.

2.2 How Advanced Microgrids Differ from Traditional Data Centers

At the surface level, many data centers may qualify as microgrids. Engineering purists may argue, however, that UPS systems installed at individual loads do not really meet the definition of a microgrid if they cannot operate parallel to the utility grid. Furthermore, most advanced microgrids deployed today usually include at least one form of renewable energy, and increasingly, some form of energy storage (often Li-ion batteries.) The adjective “advanced” also connotes leveraging the latest in software optimization, automation, and data analytics to transform microgrid assets into intelligent aggregations of diverse assets that can synergistically work together as a system of systems.

2.2.1 Resilience Focused on Intelligent Networking of Cleaner Onsite Assets

The need for resilience is shared by data centers and microgrids, but how to achieve this goal can vary immensely. Instead of load-specific backup power systems each featuring a nearly identical technology stack of diesel generators and lead-acid batteries, each sized conservatively to cover contingencies. In contrast, a microgrid seeks to more efficiently network a portfolio of DER assets that are typically dissimilar. In fact, the greater the diversity, the better, since diversity implies greater resilience. Diversity also tends to mean greater incorporation of variable renewable resources such as solar and wind. While the variability of these two renewable resources imply a threat to uptime, their value comes in the form of zero fuel costs (as well as need for storage of fuel or transportation of fuel). The key intermediary technology required to leverage the values of these renewables are advanced batteries.

2.2.2 Utility Relationship Evolves to Bidirectional Value Exchanges

The majority of microgrids connected to a utility grid run parallel to the grid well over 90% of the time. They shift to island mode only during times of emergency, or when it is economically advantageous to do so (i.e., during peak time of power demand). In the past, microgrids were often perceived to be a threat to utilities. Today, that perception by utilities is changing. Many utilities, such as Arizona Public Service, are deploying microgrids themselves for data centers. Among the reasons for this shift in thinking by utilities is that thanks to new sophisticated controls, microgrids can not only provide resiliency but grid services. Rather than being viewed purely as lost load or challenge to overall reliability,

microgrids can have it both ways. They can reduce costs and provide resiliency to a data center, while also supporting the local utility.

2.2.3 Smart Controls Optimize Diversity of DER

Most grid communications have been part of specialized systems, many of which are proprietary. However, vendors have realized the value in supporting standards efforts that now emphasize interoperability. The major challenge for microgrids with respect to interoperability involves the need to connect newer infrastructure (e.g., solar PV and advanced batteries) with older existing infrastructure (such as legacy diesel generators).

There is a growing consensus that the future of microgrid controls tilts toward distributed control approaches, given the complexity of variable DER and the need for DR and frequency regulation functionality. However, expected growth in advanced microgrids seeking to sell ancillary services and capacity into wholesale markets may also require a market interface function dependent upon longer-term supervisory controls. Navigant Research calls this emerging DER ecosystem of opportunity the Energy Cloud.⁶ This layer of modeling and controls would be able to anticipate changes in the weather, sending signals to energy storage devices and optimizing load over a longer-term operational planning horizon—as well as near instantaneous requests from grid operators. Confidence in this core function of smart microgrid controls is the key to data centers understanding the value of an advanced microgrid over the status quo.

2.3 Data Centers Drive Commercial and Industrial Microgrid Growth

One could argue that commercial and industrial (C&I) customers pioneered the concept of microgrids, though these self-sufficient networks featuring backup power assets are somewhat primitive by today's standards, often based on manual islanding and 100% fossil fuel supply portfolios. This stereotype needs a quick refresh. The C&I microgrid segment, which includes customers ranging from grocery chains to manufacturing facilities (and other forms of private enterprise) is evolving into one of the most innovative microgrid markets. It is quickly maturing as risk averse energy managers—including those that manage data centers—become more comfortable with the microgrid platform as it is validated by other customer segments. Among the unique challenges facing this microgrid market segment are the following:


- End users often enjoy the lowest electric rates of any customer class.

⁶ Navigant Research, *The Energy Cloud*, 2015, <http://www.navigantresearch.com/research/the-energy-cloud>.

- This market segment had to make a valid value proposition in the absence of federal or state funding support—other than incentives related to key enabling technologies (such as solar PV or advanced batteries).
- Internal competition for capital outlays is fierce within these companies and subject to intense scrutiny.

The C&I grid-tied microgrid market is poised for major growth. North America’s grid-tied C&I microgrid capacity is expected to grow from 215.5 MW to 870.8 MW annually between 2017 to 2026, a capacity compound annual growth rate of 16.8%. Corresponding implementation spending starts at \$857.0 million in 2017 and reaches \$3.3 billion by 2026.

Chart 2.1 *Grid-Tied C&I Microgrids by Region, World Markets: 2017-2026*



(Source: Navigant Research)

Navigant Research anticipates that data centers represent less than 15% of annual new grid-tied C&I capacity around the world today. In a decade, data centers could represent as much as 40% or more of the total grid-tied worldwide C&I microgrid market.

Section 3

DATA CENTER SEGMENTS AND SIZE DYNAMICS

3.1 Data Center Market Segment Impacts and New Notions on Distributed Resiliency

The data center industry is diverse and there are many ways to segment the market. The IDC forecasts looked at internal versus service provider segments. One can also look to segment the market by purpose, size, and other variables.

3.1.1 Centralized Data Center

Whether owned by an enterprise or part of the cloud, large, multi-megawatt centralized data centers are commonly viewed as highly mission critical, and as such, are designed with availability in mind. These are the data centers designed and often certified to the Uptime Institute's Tier 3 or Tier 4 standards. Colocation and cloud providers often tout these high availability design attributes as selling points to moving operation to their data centers. Should microgrids be categorized according to this tier structure? Perhaps that would help them gain traction within the world of data centers.

One variation on the centralized data center is the centralized cloud paradigm. Originally conceived for certain types of applications—such as email, payroll, and social media—where timing was not absolutely crucial, this portion of the data center market has evolved over time. As critical applications shift to the cloud (healthcare or self-driving vehicles are two examples), it has become apparent that latency, bandwidth limitations, security, and other regulatory requirements still have to be addressed.

3.1.2 Regional Data Centers

Regional data centers are closer to the endpoints where data is created and used and are smaller than the large centralized data centers. These data centers exist to bring latency- or bandwidth-sensitive applications closer to the point of use, and are strategically located to address high-volume needs. These data centers can be considered as the bridge between central data centers and on-premise, localized data centers. Similar to the large centralized data centers, regional data centers are typically designed with security and availability in mind. It is not uncommon to see Tier 3 designs in these kinds of facilities.

3.1.3 Localized or Edge Data Centers

A localized data center is one that is colocated with the users of the data center. They might range in size from 1 MW to 2 MW or as little as 10 kW to 20 kW. As enterprises outsource more and more of their business applications to the cloud or colocation providers, these data centers on the edge are trending toward the smaller end of that range, with sometimes only a couple of racks left in a small room or closet. In many of

these downsized data centers, the design practices often equate to a Tier 1 design, with little thought to redundancy or availability. Since on-premise applications include proprietary business critical applications and ensure network connectivity to the cloud, one could make the argument that these localized edge data centers, currently the weak link in the data center typology, are good candidates for advanced microgrids and/or cutting-edge IT innovation.

3.2 The Hybridization of Data Centers

The categories described in Section 3.1 represent where the data center market has come from. Although not an entirely new concept, the hybrid data center appears to represent the future. Instead of a focus purely on the cloud or purely on the size of the data center (or ownership), one can expect new mutations that encompass diverse forces at play, including the many hardware and software trends described previously in this report. As is the case with advanced microgrids, the most fundamental shift is toward software and cloud-based intelligence to serve as new forms of resilience, as described in the case study that follows.

3.2.1 Case Study: Next Generation Distributed Resiliency

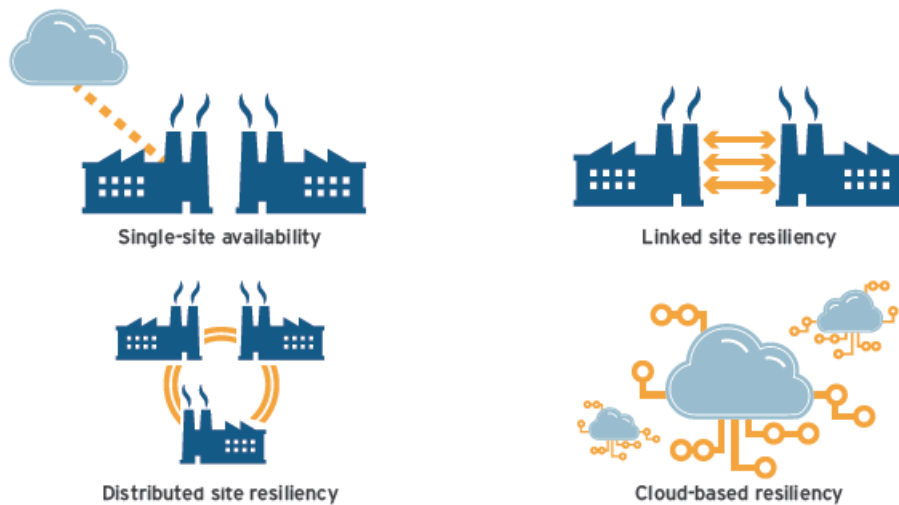
A recent report by the 451 Group, *Next Generation Resiliency*, offers a peek into the future, suggesting that the industry can no longer adhere to a business-as-usual approach to development and ongoing operations and maintenance (O&M) to maintain critical uptime. Among the new concepts to emerge within the data center space is distributed resiliency.

Under this new model of resilience, physical resilience insured by technologies such as UPS systems and, increasingly, microgrids, will be complemented by resiliency at the IT level via automated software controls. Much of this shift in thinking from hardware to software is a natural evolution as more and more applications move to the cloud. As these systems of distributed resilience evolve, incorporating increased levels of complexity and interdependency, failure (and recovery) will be less black and white/on and off. Instead, resiliency becomes autonomic, a self-managing automation process where loads and traffic can be shifted across geographies through software innovation that integrates into the physical assets that still serve as the foundation for smart networks such as microgrids. 451 Group sums up the shift as, “Single site, highly available data centers are still critical, but are becoming components in a distributed fabric, with software managing replication, availability, and integrity.” The group identified four types of distributed resiliency (see Figure 3.1):

- **Single site resiliency:** This has been the traditional approach to maintaining resiliency, focused on hardware redundancy in the past, but now shifting to the cloud, yet resiliency of service is still focused on a single site.
- **Distributed site resiliency:** Two or more independent sites using shared Internet/VPN networks to provide resilience through multiple asynchronously connected instances.

- **Linked site resiliency:** Multiple lower tier data centers are located close enough to each other that they are synchronized and are capable of processing data requests using a dedicated network.
- **Cloud-based resiliency:** Rather than multiple resilient buildings, one can also maintain critical uptime through multiple clouds, this scenario carries the concepts of distributed resiliency via cloud-based IT and software to the ultimate extreme.

Figure 3.1 *Types of Distributed Data Center Architectures*



(Source: 451 Group)

Section 4

WHY MICROGRID AS A SERVICE?

4.1 Key Challenges to Moving Forward with a Microgrid

If something is not broken, why try to fix it? This kind of thinking sums up the perspective of many data center owners and operators. If they feel comfortable with the technology or solution that has been in place for quite some time, the incentive to enact something new and different is small. Adopting technology that may be viewed by upper managers as a risk usually raises more questions than answers. Given the high stakes if something were to go wrong, one can understand that most data centers do not embrace novel technologies for energy supply. Like the military, there is comfort in sticking with the status quo, even if statistics point to these status quo systems as having high failure rates.

Data centers like to control their own destiny. That often means they want to own infrastructure. Yet just like solar leases and third-party power purchase agreements (PPAs) accelerated the solar PV industry at a critical point in time in its development path, similar models could also bring microgrids into the mainstream. Does such an approach hold promise in pushing broad adoption of microgrids as state-of-the-art data centers?

4.1.1 Parallel Paths for Data Center and Military Bases

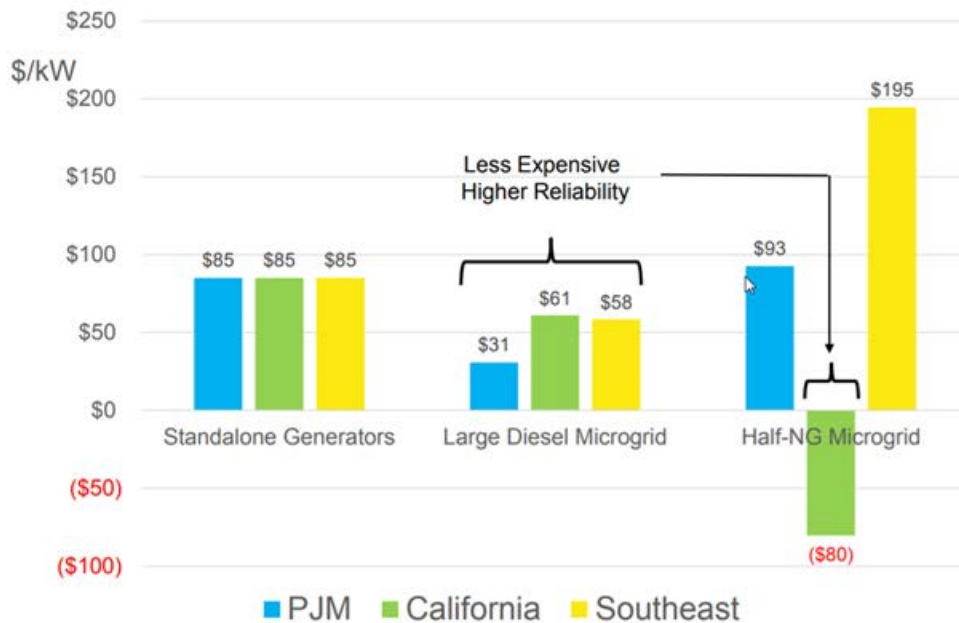
One analogy to the challenge facing data centers is military bases in the United States. A typical large-scale military base may feature from 100 to 350 backup diesel generators, each hard-wired to a single building, and in many instances, sized at more than 200% of each building's peak load as contingency for energy security. Just a simple networking of existing diesel generators into a microgrid can offer cost savings for military microgrids and data centers alike.

A study by Pew Charitable Trusts found, for example, that creating a microgrid instead of relying upon stand-alone backup diesel generators reduces the cost of resiliency significantly. Note that the savings vary by region, with greatest savings for those military microgrids deployed in the Pennsylvania-New Jersey-Maryland (PJM) independent system operator transmission control area.

Yet when displacing diesel backup generators with 50% diesel/natural gas fuel hybrid microgrid, California military bases boast the largest net savings. With a 50/50 portfolio of diesel/natural gas, microgrids in PJM and the southeast ironically show an increase in cost on a \$/kW basis if compared to current reliance upon diesel backup generators.

This is largely a result of low diesel fuel prices in those parts of the country, and arguably points to the need to diversify power generation sources with a microgrid beyond fossil fuels.

Figure 4.1 Annual Net Cost of Protection (\$/kW of Critical Load)



(Sources: Noblis-Pew Charitable Trusts)

4.1.2 Understanding the True TCO

What is the true TCO for data center infrastructure? And how does that compare to a microgrid? Both good questions. Neither has an easy answer. In short, it depends. Rather than focusing on the upfront capital costs, the majority of a data center’s TCO is in O&M expenses. This is also where most of the potential cost savings reside. Furthermore, energy costs represent the largest portion of the O&M costs.

In an effort to become more energy efficient, some data centers are reducing capacity safety margins and system redundancy. This in turn increases the importance of proactive maintenance. And advanced microgrid can help address these issues. High levels of facility automation and equipment performance data have created new opportunities for enhancing reliability and resiliency while reducing costs. Though this may sound too good to be true, if properly managed an advanced microgrid can reduce O&M costs while maintaining or even increasing resiliency.

Armed with data on the true TCO of data centers, owner/operators can at least explore the possibility of moving forward with a microgrid. Most microgrids are retrofits. They can be designed to squeeze the most value out of legacy assets while upgrading the power

network with the latest clean energy DG assets, smart inverters and meters, energy storage devices, and system controls. The data center industry is too important to stay stuck in a conservative and polluting approach to uptime. A creative blending of onsite hardware assets optimized by cloud-based IT can provide new strategic layers of resiliency. A mixture of new hardware assets, automation, and software can also generate new revenue streams without jeopardizing core mission-critical functions.

4.1.3 Lack of Experience with Variable Renewable Resource Integration

The large data centers committing to 100% renewable energy supply portfolios prefer the purchase of RECs and other vehicles of virtual greening up of their supply portfolios because it is so easy. They do not have to switch out any hardware or tinker with the nitty gritty details of electrical power infrastructure. Has the time come to challenge the traditional data center industry approach to resiliency and uptime?

Despite the delivery of a service vital to the digital economy, the data center industry is not taking advantage of new renewable energy resources they could touch, feel, and see literally in their own backyards. Furthermore, this lack of experience is shielding data centers from the latest integration technologies that could transform their operations into state-of-the-art smart buildings, which, in turn, become modular building blocks for microgrids. The longer the data center industry waits, the larger the pool of legacy assets becomes. Instead of trailing behind military bases and other sophisticated C&I customers, would it not be wise for the data center industry to embrace the clean energy revolution in a more direct and fundamental way?

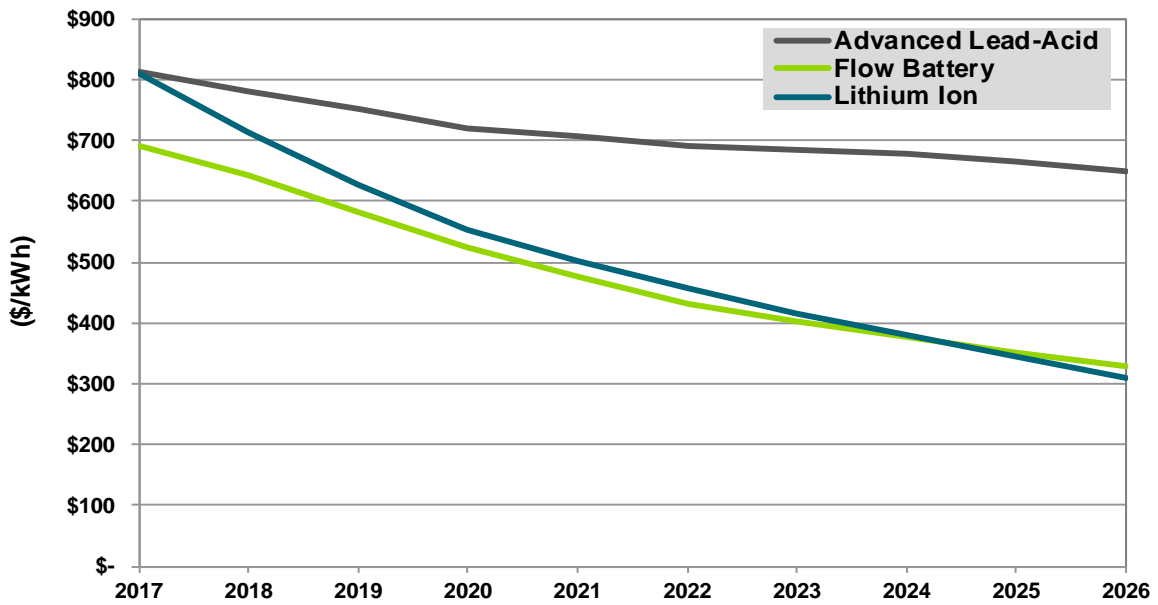
4.1.4 Knowledge of New Advanced Batteries is Limited

Similar to their experience with distributed renewables such as solar PV, the data center industry's knowledge about new batteries has been limited. Again, the industry has taken a conservative, copycat approach to UPS-based resiliency.

Today, the battery technology of choice for microgrids is Li-ion batteries—which have been experiencing slow but steady cost declines—because they can provide a range of energy storage services (both power and energy). As levels of renewable energy adoption increase over time, however, it is expected that longer duration battery technologies—such as flow batteries—will become increasingly attractive. Some microgrids today also feature a diversity of batteries.

Lead-acid batteries have long been the workhorse energy storage technology of choice for both data centers and microgrids. But this no longer is the case—at least with microgrids, and Chart 4.1 demonstrates why. Due to major investment in new advanced batteries, lead-acid batteries are forecast to be the most expensive batteries over the next 10 years, with costs for both Li-ion and flow batteries declining rapidly, further accelerating their adoption in microgrids and other grid-connected and off-grid applications.

Chart 4.1 C&I Energy Storage CAPEX Assumptions by Technology, Average Installed Costs, World Markets: 2017-2026



(Source: Navigant Research)

4.2 New Microgrid Business Models Emerging

Just as the data center industry is shifting from internal data center direct ownership business models to a service delivery approach, microgrids also appear to be moving in this direction. Previous research conducted by Navigant Research revealed that, among microgrids deployed in North America between 2015 and 2016, the PPA model was the most common business model deployed for microgrids based on both project number and total peak capacity.

4.2.1 Moving Closer to Plug-and-Play

One of the primary challenges facing the microgrid market today is the perception that each project is unique and therefore requires significant customized engineering. In fact, dozens of microgrids never seem to make it past the feasibility analysis phase, in part, due to this predicament. Vendors who bill by the hour for complex engineering tasks do not mind; however, when projects fail to move forward it does not build confidence in the microgrid platform.

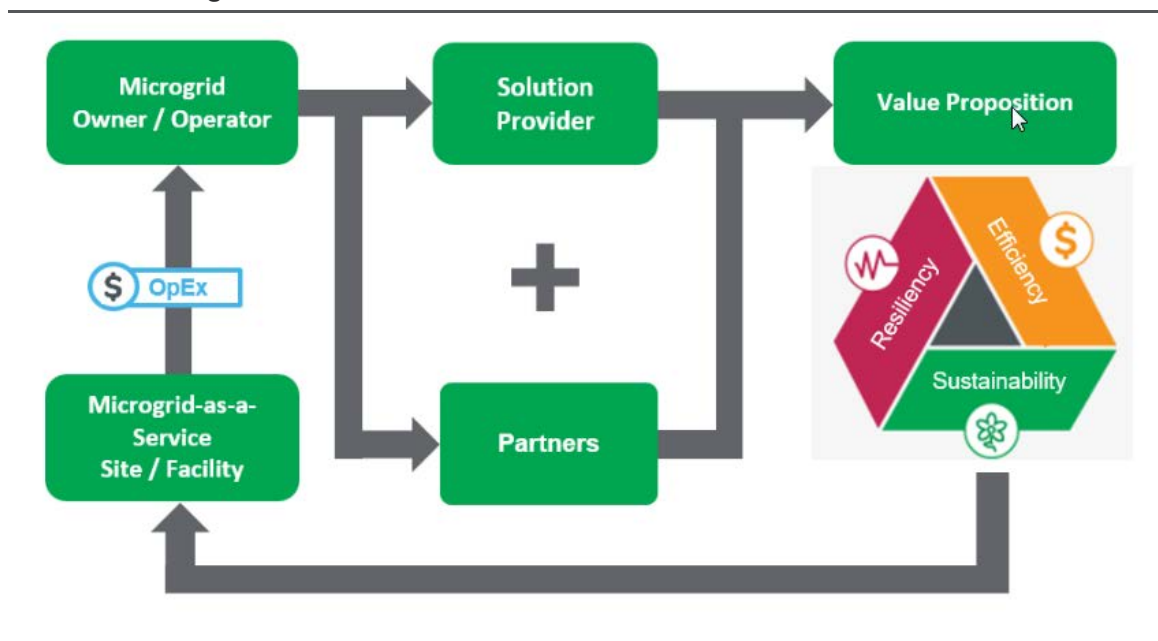
The idea of cookie-cutter microgrids seems impossible. A long line of vendors, however, among them Schneider Electric, are now offering products and services that are moving the market much closer to a plug-and-play paradigm. For the C&I market, this is key to market growth. This will especially be the case when it comes to data centers, which historically have relied upon tightly prescribed solutions to the challenges of uptime needs. Though few microgrids are identical, the key to standardization and easy replication of successful designs is software automation and the cloud. The very product enabled by data centers—cloud computing—is, ironically enough, the key to make data centers more resilient, efficient, and sustainable.

4.2.2 How Microgrid as a Service Works

The basic structure of the new microgrid as a service (MaaS) approach developed by Schneider Electric is depicted in Figure 4.2. The easiest way to understand how it works is to use a specific project example. In this case, Montgomery County, Maryland, is moving forward with two microgrids that help the county meet its policy goals to dramatically increase resiliency, efficiency, and sustainability without exposing the local government to large capital expenses. The reason it sought help from the private sector is that it wanted to address its aging infrastructure, but was budget constrained.

Figure 4.2 is a simple flow diagram illustrating the value streams that flow under the MaaS model. In the case of this Montgomery County project, Montgomery County is the MaaS site facility, Duke Energy Renewables is the owner/operator of the microgrid assets, and Schneider Electric is the solutions provider. The microgrid project is funded through the PPA partnership with Schneider Electric, which eliminates the need for upfront capital expenditure from Montgomery County. The PPA that is incorporated into this MaaS contract includes a unique capacity payment to cover the CAPEX expenses for the microgrid component, in addition to a volumetric energy charge that serves as the basis for a standard PPA.

Figure 4.2 *Microgrid as a Service*



(Source: Schneider Electric)

Montgomery County’s two microgrids, which are expected to be completed by 2018, will provide the following benefits to Montgomery County:

- Produce nearly all energy needed onsite (>80%)—3.6 MW hours of solar energy each year, enough to power more than 200 homes.
- Reduce greenhouse gas emissions by 3,629 metric tons each year, which translates to removing over 767 cars off the road or planting 94,000 trees.
- Allow the county to avoid \$4 million in outright CAPEX for medium and low voltage distribution line upgrades.
- Lock in known price of energy for 25 years.

4.2.2.1 *What if the Microgrid Vendor Defaults on Obligations?*

The primary fear among data center companies can be summed up succinctly: what if the solutions provider goes bankrupt or fails to deliver on the uptime requirements it expects and needs to remain in business? This is a legitimate concern, and it can be addressed in the structure of the contract, performance guarantees, and due diligence on both the solutions provider and technologies deployed within the microgrid. One could argue that there is also risk with an internal data center fully dependent upon UPS systems, that it might not fully understand and/or have the expertise to fix it in the event of a malfunction. Whether one wishes to own and operate or outsource to an experienced solutions provider, there will always be pros and cons. What must be determined is the appropriate risk tolerance given in-house capabilities and the availability of cost-effective, third-party solutions.

4.2.2.2 *Selection of an Experienced and Financially Sound Partner Is Key*

Given the risk profile of data center operations, it is paramount that if a data center operation wishes to outsource its energy asset management tasks to a third-party, it chooses wisely. The pool of microgrid innovators keeps growing. There is a mix of both large, established players as well as small, innovative startups. The beauty of the MaaS is that there is no CAPEX. Furthermore, the cost of the infrastructure upgrades is rolled into the O&M expense. As stated earlier, this is the primary point to attack to reduce TCO. Experience can cost money, as can a firm with a proven track record that makes them financially sound. The key to a winning project is finding the right combination of technologies, the right combination of vendor partners, and the right choice of the solutions provider and project integrator. Perhaps the operation of the data center could still be under the full control of the user, even if the infrastructure developed was performed by a third party.

Section 5

FIVE KEY TAKEAWAYS FOR DATA CENTERS

5.1 Why Microgrids are Better than Traditional Approaches for Meeting Uptime Goals

Advanced microgrids incorporate the latest hardware and software innovations to meet today's data center resiliency requirements at a lower TCO. They do so at a lower environmental cost. Here are five primary takeaways from this white paper.

5.1.1 How New Generation Assets Make Microgrids Attractive

Declines in the cost of solar PV, the low cost of natural gas, and even more importantly battery storage, translate in the ability to create a more sustainable network for data centers. Other options include biogas sources that could supply fuel cells, as well as other regionalized resources. Diesel is not dead. But it should serve as the last resort, not the primary pathway to resiliency.

5.1.2 How New Opportunities for Demand Management Translate into Revenue

Step one for a model data center resiliency program is energy efficiency. But DR—demand response, not disaster recovery—should also be on the table. If the assets deployed within the network (i.e., advanced microgrid) are cleaner, then they can offer more value since they can help the data center participate in utility grid service opportunities, which are often limited today due to air quality limitations placed on diesel generators.

5.1.3 How Digital Grid Trends Point to Microgrids as a Data Center Solution

Past fears about the variability of renewables can be addressed directly by recent advances in automation, data analytics, and artificial intelligence. These digital grid trends all make an advanced microgrid possible, blending a combination of onsite, cleaner hardware assets with remote cloud applications to create new forms of distributed resiliency.

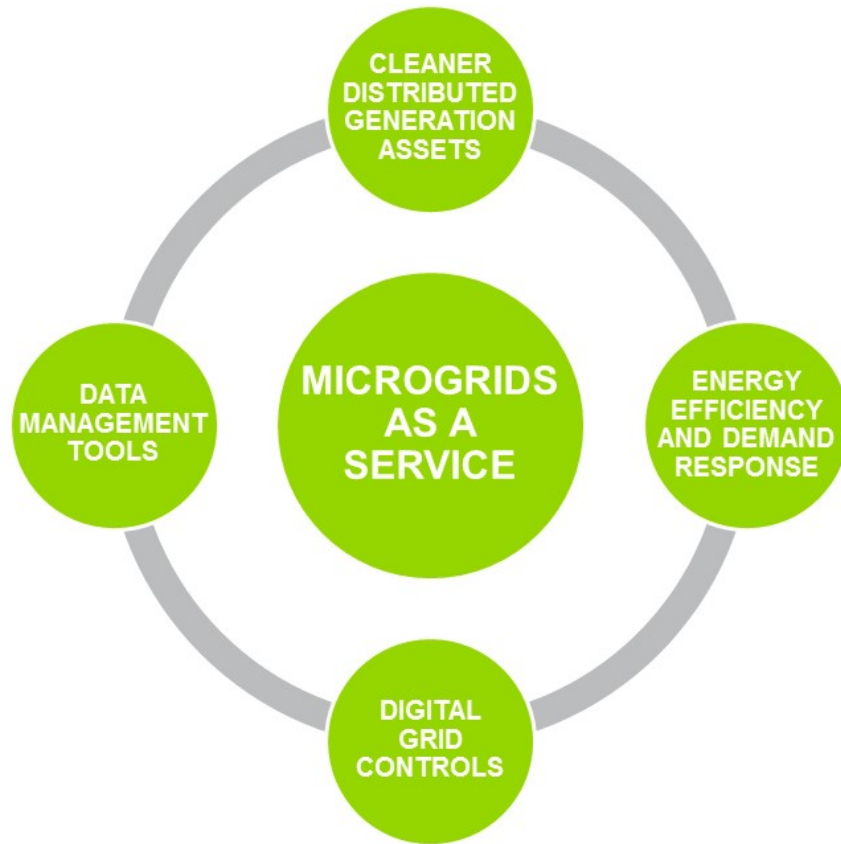
5.1.4 The Critical Role of Data Management Tools

New management software planning and implementation tools can vastly improve IT room allocation of power and cooling requirements, can provide rapid impact analysis when operations fail, and leverage historical data to improve future IT performance. Any DCIM program starts with good planning, including advanced microgrids.

5.1.5 MaaS Shifts Risk to the Experts

Translate CAPEX into OPEX and let an experienced solutions provider design and implement the best advanced microgrids for each clients' unique needs.

Figure 5.1 *Advanced Microgrids Advantages*



(Source: Navigant Research)

Section 6

ACRONYM AND ABBREVIATION LIST

1N.....	Non-Redundant
2N.....	Fully Redundant
C&I.....	Commercial and Industrial
CAPEX.....	Capital Expenditures
DCIM.....	Data Center Infrastructure Management
DER.....	Distributed Energy Resources
DG.....	Distributed Generation
DR.....	Demand Response
GW.....	Gigawatt
IT.....	Information Technology
kW.....	Kilowatt
Li-ion.....	Lithium Ion
MaaS.....	Microgrids as a Service
MW.....	Megawatt
O&M.....	Operations and Maintenance
OPEX.....	Abbreviation
PJM.....	Pennsylvania-New Jersey-Maryland
PPA.....	Power Purchase Agreement
PV.....	Photovoltaic
REC.....	Renewable Energy Credit
TCO.....	Total Cost of Ownership
UPS.....	Uninterrupted Power Supply

Section 7

TABLE OF CONTENTS

Section 1	1
Introduction: Setting the Stage	1
1.1 The Impact of the New Digital Economy on Data Centers	1
1.1.1 Data Centers Are Emerging as a Target for New Energy Innovations	1
1.1.2 Data Center Growth Trends	3
Section 2	6
Data Centers and Microgrids: Finding Common Ground	6
2.1 Mutual Goals of Data Centers and Microgrids	6
2.1.1 Identification of Mission Critical Loads	8
2.1.2 Focus on Uptime	8
2.1.3 Leveraging Existing Utility Grid Infrastructure	8
2.2 How Advanced Microgrids Differ from Traditional Data Centers	9
2.2.1 Resilience Focused on Intelligent Networking of Cleaner Onsite Assets	9
2.2.2 Utility Relationship Evolves to Bidirectional Value Exchanges	9
2.2.3 Smart Controls Optimize Diversity of DER	10
2.3 Data Centers Drive Commercial and Industrial Microgrid Growth	10
Section 3	12
Data Center Segments and Size Dynamics	12
3.1 Data Center Market Segment Impacts and New Notions on Distributed Resiliency	12
3.1.1 Centralized Data Center	12
3.1.2 Regional Data Centers	12
3.1.3 Localized or Edge Data Centers	12

3.2	The Hybridization of Data Centers	13
3.2.1	Case Study: Next Generation Distributed Resiliency	13
Section 4	15
Why Microgrid as A Service?	15
4.1	Key Challenges to Moving Forward with a Microgrid	15
4.1.1	Parallel Paths for Data Center and Military Bases	15
4.1.2	Understanding the True TCO	16
4.1.3	Lack of Experience with Variable Renewable Resource Integration	17
4.1.4	Knowledge of New Advanced Batteries is Limited	17
4.2	New Microgrid Business Models Emerging	18
4.2.1	Moving Closer to Plug-and-Play	19
4.2.2	How Microgrid as a Service Works	19
4.2.2.1	What if the Microgrid Vendor Defaults on Obligations?	21
4.2.2.2	Selection of an Experienced and Financially Sound Partner Is Key	21
Section 5	22
Five Key Takeaways for Data Centers	22
5.1	Why Microgrids are Better than Traditional Approaches for Meeting Uptime Goals	22
5.1.1	How New Generation Assets Make Microgrids Attractive	22
5.1.2	How New Opportunities for Demand Management Translate into Revenue	22
5.1.3	How Digital Grid Trends Point to Microgrids as a Data Center Solution	22
5.1.4	The Critical Role of Data Management Tools	22
5.1.5	MaaS Shifts Risk to the Experts	22
Section 6	24
Acronym and Abbreviation List	24

Section 7 25

Table of Contents 25

Section 8 28

Table of Charts and Figures 28

Section 9 29

Scope of Study 29

Sources and Methodology 29

Notes 30

Section 8

TABLE OF CHARTS AND FIGURES

Chart 1.1	Cumulative Data Center Capacity, World Markets: 2013-2021	4
Chart 2.1	Grid-Tied C&I Microgrids by Region, World Markets: 2017-2026	11
Chart 4.1	C&I Energy Storage CAPEX Assumptions by Technology, Average Installed Costs, World Markets: 2017-2026	18
Figure 3.1	Types of Distributed Data Center Architectures	14
Figure 4.1	Annual Net Cost of Protection (\$/kW of Critical Load)	16
Figure 4.2	Microgrid as a Service	20
Figure 5.1	Advanced Microgrids Advantages	23

Section 9

SCOPE OF STUDY

This white paper examines how trends in data center power infrastructure are setting the stage for an opportunity to develop advanced microgrids. Like every other microgrid market segment, there are unique opportunities and barriers. Since the data center industry has such a large impact on both economy and environment, innovation is key to reaching global sustainability goals. While technology advances with hardware such as solar PV and energy storage are a key part of this story, software and new business models that mirror the outsourcing of resiliency may be the most important factors in data center microgrid growth over the next 10 years. This report leverages Navigant Research reporting on both C&I microgrids as well as energy storage costs, as well as IDC data on future data center growth.

SOURCES AND METHODOLOGY

Navigant Research's industry analysts utilize a variety of research sources in preparing Research Reports. The key component of Navigant Research's analysis is primary research gained from phone and in-person interviews with industry leaders including executives, engineers, and marketing professionals. Analysts are diligent in ensuring that they speak with representatives from every part of the value chain, including but not limited to technology companies, utilities and other service providers, industry associations, government agencies, and the investment community.

Additional analysis includes secondary research conducted by Navigant Research's analysts and its staff of research assistants. Where applicable, all secondary research sources are appropriately cited within this report.

These primary and secondary research sources, combined with the analyst's industry expertise, are synthesized into the qualitative and quantitative analysis presented in Navigant Research's reports. Great care is taken in making sure that all analysis is well-supported by facts, but where the facts are unknown and assumptions must be made, analysts document their assumptions and are prepared to explain their methodology, both within the body of a report and in direct conversations with clients.

Navigant Research is a market research group whose goal is to present an objective, unbiased view of market opportunities within its coverage areas. Navigant Research is not beholden to any special interests and is thus able to offer clear, actionable advice to help clients succeed in the industry, unfettered by technology hype, political agendas, or emotional factors that are inherent in cleantech markets.

NOTES

CAGR refers to compound average annual growth rate, using the formula:

$$\text{CAGR} = (\text{End Year Value} \div \text{Start Year Value})^{(1/\text{steps})} - 1.$$

CAGRs presented in the tables are for the entire timeframe in the title. Where data for fewer years are given, the CAGR is for the range presented. Where relevant, CAGRs for shorter timeframes may be given as well.

Figures are based on the best estimates available at the time of calculation. Annual revenues, shipments, and sales are based on end-of-year figures unless otherwise noted. All values are expressed in year 2017 US dollars unless otherwise noted. Percentages may not add up to 100 due to rounding.

Published 4Q 2017

©2017 Navigant Consulting, Inc.
1375 Walnut Street, Suite 100
Boulder, CO 80302 USA
Tel: +1.303.997.7609
<http://www.navigantresearch.com>

Navigant Consulting, Inc. (Navigant) has provided the information in this publication for informational purposes only. The information has been obtained from sources believed to be reliable; however, Navigant does not make any express or implied warranty or representation concerning such information. Any market forecasts or predictions contained in the publication reflect Navigant's current expectations based on market data and trend analysis. Market predictions and expectations are inherently uncertain and actual results may differ materially from those contained in the publication. Navigant and its subsidiaries and affiliates hereby disclaim liability for any loss or damage caused by errors or omissions in this publication.

Any reference to a specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply an endorsement, recommendation, or favoring by Navigant.

This publication is intended for the sole and exclusive use of the original purchaser. No part of this publication may be reproduced, stored in a retrieval system, distributed or transmitted in any form or by any means, electronic or otherwise, including use in any public or private offering, without the prior written permission of Navigant Consulting, Inc., Chicago, Illinois, USA.

Government data and other data obtained from public sources found in this report are not protected by copyright or intellectual property claims.

Note: Editing of this report was closed on October 09, 2017.