Life Is On | Schneider Electric

Monetizing Energy Storage in the Data Center

White Paper 274

Version 1

by Christopher Thompson and Victor Avelar

Executive summary

Over the last several years, the electric utility markets have seen significant changes in the mix of generation supporting the grid. In particular, traditional generating forms such as coal, nuclear, and even natural gas are being replaced by intermittent generation sources such as wind and solar. Because of the increasing penetration of these intermittent sources, and their lack of controlled dispatchability, other fast-acting energy sources are becoming even more valuable in balancing supply and demand. This is one reason why utilities are increasingly providing financial incentives to their customers to more closely balance real-time supply and demand. Due to their generous capacity of underused UPS batteries, data centers are a prime candidate to take advantage of these incentives. In this paper, we describe approaches data center operators can use to monetize their UPS energy storage and explain the modes of UPS operation required for each method.

RATE THIS PAPER ★★★★★

Introduction

Most data centers around the world use UPS with battery backup to condition incoming power and ride through power outages. Due to the highly mission-critical nature of these facilities, almost all of the major subsystems are fully redundant. This includes subsystems such as UPS, batteries, generators, and cooling. In fact, even during extreme conditions such as natural disasters, most such facilities can run autonomously for days. These subsystems include redundant UPS batteries that represent a significant amount of on-site energy storage. In almost all locations these overprovisioned battery systems sit idle for most of the year. By one estimate, worldwide annual data center energy consumption will reach over 1,000 TWh [1] in 2020 (**Figure 1**), which we estimate to include about 22 GWh of battery energy storage [2]. What if this idle storage capacity could generate revenue for a data center operator by helping regulate the electric utility grid? In fact, over the last few years electric utilities have created a number of ancillary service programs that may incentivize data center operators to use this idle energy capacity to help the utility balance the supply and demand of electricity in real time [3].

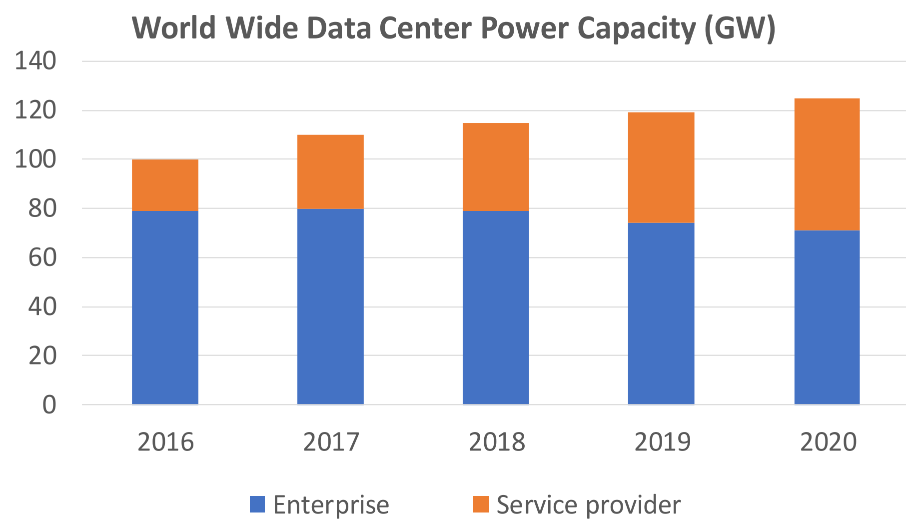

Figure 1
Estimated worldwide data center power consumption

While data centers have had battery-based energy storage for many decades, why would this make sense now? In this paper, we explain the trends behind this new ability for data centers to monetize underused capacity. We describe three basic ways data center operators can interact with electric utilities, and the UPS modes required for these interactions. Finally, we illustrate the pros and cons of each UPS mode.

Historically, electric utilities directly, and centrally, controlled diverse types of electricity generation to match with demand from their customer base. However, with the increase in renewable energy, utilities have limited control of these distributed energy sources due to their inherently intermittent nature. While these renewable generation sources are essentially free to operate once installed, they often don't

[1] IDC

[2] The report above estimates total data center load worldwide will be over 120 GW, at a PUE of 1.8. This represents over 67 GW of IT load (120 GW / 1.8). If we assume a UPS runtime of 20 minutes at 50% UPS load, this represents a total battery energy storage of over 22 GWh (67 GW x 0.33 hrs)

[3] Balancing the grid is a primary reason utilities are increasingly charging customers "time of use" pricing. Higher or lower rates during various parts of the day may help consumers save money while also improving grid stability by incentivizing users to change electric consumption patterns

provide other ancillary grid services[4] that are standard in previous generation forms and essential for grid stability. Examples of such ancillary services include frequency regulation, voltage control, VAR control, and inertia. Therefore, to keep total operating costs low, utilities look for other ways, including load control or battery storage, to provide such ancillary services.

An important point on this topic is that in addition to the financial benefits causing this market to grow, there are also regulatory drivers as well. In particular, the United States Federal Energy Regulatory Commission (FERC) issued Order 841[5] in 2017. This order directed regional transmission organizations (RTOs) and independent system operators (ISOs) to craft rules that would allow storage assets to participate in wholesale electricity markets. In the past, ancillary service participants faced rules that did allow for the capabilities of battery storage to be fully valued and only participants with large projects were able to participate. Today, curtailment service providers (CSP)[6], or scheduling coordinators, also offer a means for smaller facilities to participate through aggregation. Ancillary services markets will continue to grow as renewables continue to penetrate the market and as FERC Order 841 is implemented more broadly in the US. The European Union is also taking steps to increase energy storage incentives[7]. As the ancillary services market grows, so will the financial incentives to participate in these services. When you combine this with advances in Li-ion batteries, data center operators may realize a growing business opportunity.

The following points help to explain this:

- **Data centers represent a large relatively constant load year-round**
  Utility providers value large constant loads because it allows them to manipulate the balance of supply and demand in predictable ways and in larger steps sizes compared to residential loads.
- **Data centers are getting bigger and need more power**
  Similar to the point above, as individual data center capacities grow into the tens of MW, the financial incentives also grow.
- **Energy storage is already part of data center requirements**
  Data centers typically have a minimum 10-15 minutes of runtime at full load (1N UPS) and over 1 hour when partially loaded at 2N redundancy. This means most data centers can generate revenue from ancillary services without buying additional electrical equipment.
- **Data centers are migrating to high cycle-life batteries**
  Most data centers are migrating from Lead-acid to Lithium-ion (Li-ion) batteries and Li-ion battery technology is an ideal storage medium for these applications. Most Li-ion battery solutions on the market today are expected to last over 10 years and are also more reliable than lead-acid batteries. These batteries can also endure significantly more charge / discharge cycles than lead-acid batteries and they occupy less space. The short cycle life of lead-acid batteries and lack of financial utility incentives are perhaps the biggest reasons why data centers haven't historically taken advantage of their energy storage systems in this way.

Energy storage incentives

As regulations push for cleaner (zero emission) energy, more energy storage will be required to meet those goals. This has led to incentives for energy storage installations. One notable example is the state of California in the U.S. On September 10, 2018, the Governor signed into law, Senate Bill 100 which mandates 100% zero emissions electricity by 2045. To facilitate this mandate, Senate Bill 700 was signed into law which provides hundreds of millions of dollars for energy storage incentives. These types of incentives will help data centers worldwide justify the upfront cost of long-runtime UPS battery solutions in addition to the revenue for participating in utility programs.

[4] "Ancillary services are the services necessary to support the transmission of electric power from seller to purchaser given the obligations of control areas and transmitting utilities within those control areas to maintain reliable operations of the interconnected transmission system.", source: Wikipedia

[5] FERC 841 basically forces Regional Transmission Organizations (RTO) and Independent System Operators (ISO) to allow owners of small energy storage to participate in wholesale electricity markets.

[6] Curtailment service providers, act as agents to the utility, by aggregating customers and using their portfolio of customers to meet the amount of demand response they promised the utility provider.

[7] https://ease-storage.eu/first-steps-towards-es-strategy/ (last accessed June 1, 2020)

- **Cost of Li-ion energy storage is decreasing**
  Over the last several years, Li-ion batteries (as a general category) have become much more cost competitive with lead-acid batteries. This has led to significant adoption of Li-ion batteries in the overall data center market.
- **Diesel generator restrictions are becoming more complicated**
  In some locations, regulations are increasingly limiting the use of diesel generators. As an alternative, some data center operators are simply installing additional battery capacity to compensate.
- **Latent battery failures can be reduced**
  Using batteries more often increases the likelihood that you discover a latent battery issue before a power outage. In fact, one of the most common causes for facility outages is that batteries have failed without proper indication. So exercising them in a controlled way decreases the likelihood of failure during a mission critical event.

Despite these many reasons why data centers make potentially good grid participants, what's in it for data center operators, and are there any risks? The answer depends on how data center operators choose to participate, as we highlight in the next section.

General methods to monetize energy storage

Those who have investigated the idea of monetizing energy storage, typically associate this to participating in demand response programs offered by electric utilities. In these programs, utilities or distribution grid operators pay enrolled participants in exchange for altering their normal power consumption pattern. However, the term **demand response**, as defined by FERC [8] (and similarly in Europe [9]) is more general in that it also includes the idea of influencing (i.e. nudging) consumption patterns through changes in electricity prices. Basically, demand response is a way of allowing utility customers to participate in balancing the grid by controlling their use of utility power.

In this paper we focus on four key methods to monetize UPS energy storage.

- Time-of-use load management
- Demand charge management
- Third-party signaled or demand response event participation
- Third-party capacity aggregation

Time of use load management

In many regions the local utility may have electricity pricing that changes with the time of day, the day of the week, or even seasonally. To capitalize on such programs an owner of storage assets may choose to charge their batteries when electricity prices are low and discharge them when prices are high.

There are various reasons why electricity prices change over time. For example, as solar power generation has increased, there are growing periods of the day where there is actually overgeneration on the grid. During such conditions solar or wind

[8] "Changes in electric usage by demand-side resources from their normal consumption patterns in response to changes in the price of electricity over time, or to incentive payments designed to induce lower electricity use at times of high wholesale market prices or when system reliability is jeopardized." https://www.ferc.gov/industries/electric/indus-act/demand-response/dem-res-adv-metering.asp?csrt=3371792667015252339 (last accessed June 1, 2020)

[9] https://setis.ec.europa.eu/publications/setis-magazine/smart-grids/demand-response-empowering-european-consumer# (last accessed June 1, 2020)

plants may be called on to curtail energy production even though the available energy is free. On the contrary, of course, there are also periods of under generation. As discussed above, to help manage this daily cycle, many utilities now offer time-of-use pricing as shown graphically in **Figure 2**. The pricing structure incentivizes customers to use more electricity when solar production is at its peak by offering very low prices. Then during the early evening, when utilities have to procure more expensive traditional generation, the prices to customers are higher. With this method, data center operators can use their energy storage to load shift [10] as much or as little as they see fit, for optimizing the timing of their energy demands. As a side note, in addition to batteries, data center operators could also choose to over-produce ice or chilled water during periods of lower electricity prices.

To optimize an electric bill, an operator only needs to use energy storage in accordance with your electric rate structure. Note that bill optimization in this way is only sensible if time-of-use pricing is in effect. If your electric rate structure consisted solely of a fixed energy rate (i.e. $0.12/kWh), you couldn't lower your electric bill by using energy storage. In fact, you would slightly increase your bill due to the losses from charging and discharging the batteries. It is the variation in your electric rate, that provides the opportunity to lower your electric bill, also known as arbitrage.

Figure 2
Graphical representation of time of use electricity pricing

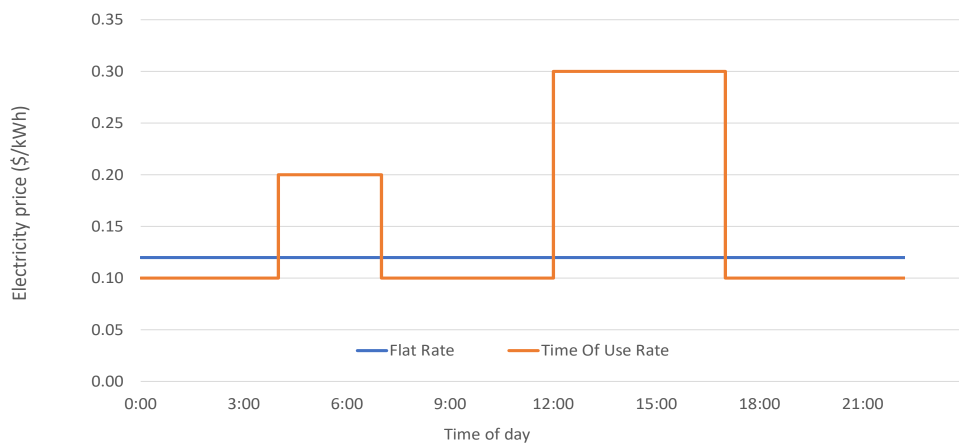

Demand charge management

Most commercial electric bills (see example in **Figure 3**), have two types of charges; the first is a charge for the maximum power drawn during the billing cycle, known as the **peak demand charge**. This is a power-related charge and is therefore assessed in $/kW. In addition, there is an **energy charge**, which is assessed in $/kWh of energy consumed over the billing cycle. Batteries can be discharged during periods of high demand to reduce the peak power required from the grid, hence lowering demand charges.

[10] "Load shifting essentially moves electricity consumption from one time period to another. For example, postponing an industrial process to another time" - GridBeyond (last accessed June 1, 2020)

nationalgrid

BILLING PERIOD	PLEASE PAY BY	AMOUNT DUE
Apr 7, 2019 to May 7, 2019	Jun 2, 2019	$ 31,881.68

Customer Charge				223.00
Dist Chg On Peak	0.01693999	x	144220 kWh	2,443.08
Dist Chg Off Peak	0.01093999	x	142151 kWh	1,555.12
Transition Charge	-0.00095	x	286371 kWh	-272.05
Transmission Charge	0.02063	x	286371 kWh	5,907.83
Distribution Demand Chg	5.76	x	834 kW/kVA	4,803.84
Energy Efficiency Chg	0.00685732	x	286371 kWh	1,963.74
Renewable Energy Chg	0.0005	x	286371 kWh	143.19
Distributed Solar Charge	0.00053	x	286371 kWh	151.78
Storm Performance Adjustment				-0.55
			Total Delivery Services	**$ 16,918.98**

Explanation of General Billing Terms

KWH: Kilowatt-hour, a basic unit of electricity used.
Off-Peak: Period of time when the need or demand for electricity on the Company's system is low, such as late evenings, weekends and holidays.
Peak: Period of time when the need or demand for electricity on the Company's system is high, normally during the day, Monday through Friday, excluding holidays.
Estimated Bill: A bill which is calculated based on your typical monthly usage rather than on an actual meter reading. It is usually rendered when we are unable to read your meter.
Meter Multiplier: A number by which the usage on certain meters must be multiplied by to obtain the total usage.
Demand Charge: The cost of providing electrical transmission and distribution equipment to accommodate your largest electrical load.

Supplier Service Charges are comprised of:

Generation Charge: The charge(s) to provide electricity and other services to the customer by a supplier.

Delivery Service Charges are comprised of:

Customer Charge: The cost of providing customer related service such as metering, meter reading and billing. These fixed costs are unaffected by the actual amount of electricity you use.
Distribution Charge: The cost of delivering electricity from the beginning of the Company's distribution system to your home or business.
Transition Charge: Company payments to its wholesale supplier for terminating its wholesale arrangements.
Transmission Charge: The cost of delivering electricity from the generation company to the beginning of the Company's distribution system.

Energy Efficiency Charge: The cost of energy efficiency program services offered by the Company.
Renewable Energy Charge: A charge to fund initiatives for communicating the benefits of renewable energy and fostering formation, growth, expansion and retention of renewable energy and related enterprises.
Distributed Solar Charge: Recovers the cost of the Massachusetts solar program, including payments to owners of solar systems.

Figure 3
Example of a commercial electric bill with time-of-use pricing and definitions

Demand charge reduction

There are various ways utilities can assess a data center's demand charge, but it's typically associated with the time period when demand for electricity is highest and measured over some defined interval (such as 15 minutes). These demand charges tend to range widely, for example, they can be from $4 to $51/kW in the United States [11]. A data center operator can reduce their demand charges by placing some or all of the IT load on UPS battery for a period of time as illustrated in **Figure 4**. Because demand charge reduction is implemented by reducing the peak demand, this approach is often referred to as "peak shaving". The profile of a data center's IT load is relatively flat, therefore peak shaving may only make sense during hot days when the cooling system peaks, which could set a higher demand charge for the entire month or even year. Note that total data center load increases slightly after peak shaving due to the battery charging losses.

In **Figure 4** the UPS is able to regulate the load as seen by the grid in a specifically controlled fashion. By doing so, peak power requirements are minimized, and demand charges are reduced. We describe below an example of how this process works.

[11] https://www.nrel.gov/solar/assets/pdfs/2017-us-demand-charges-webinar.pdf, page 11 (last accessed June 1, 2020)

Stage 1 Because the load is relatively low, the UPS is not contributing battery power.
Stage 2 As the load increases, the UPS contributes power from the battery to keep the maximum power demanded from the grid below a defined level. At this point the UPS is only contributing the exact amount of power needed (from battery) to keep the net load below the threshold.
Stage 3 At this stage the data center load is slightly below the threshold, but the UPS batteries require charging. Therefore, the charging power is carefully limited so that the combination of the data center load and the battery charger stay below the threshold. Note that no load contribution is provided by the battery at this point.
Stage 4 As the data center load drops further, the battery charging power can be increased to the maximum rating of the charger. If preferred, charging could also be scheduled at a later point in the day. This, for example, could allow for recharging when the electricity rates are the lowest.
Stage 5 At this point the load is below the threshold and the batteries are fully charged. The utility consumption now matches the UPS load as in Stage 1.

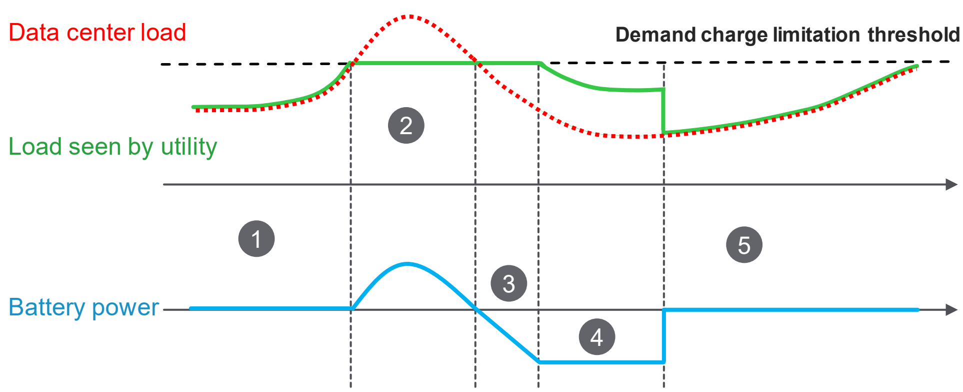

Figure 4
Reducing data center load during peak demand periods can reduce demand charges.

Combining peak shaving with time-of-use pricing
As described earlier, utilities may also have time-of-use pricing. Thus, while implementing peak shaving, an operator may selectively determine when to recharge the batteries for optimal total economic value. In fact, some utilities may have multiple programs that allow an operator to value stack the benefits of energy storage.

Third-party signaled or demand response event participation

From time-to-time grid operators may find there is a temporary imbalance between supply and demand on the grid. Examples of this are when the wind suddenly calms at a wind farm, the sun rapidly gets obstructed by clouds, or in an extreme case, a conventional power plant could go off-line.

In many cases this condition is predicted through proper forecasting but in some cases, it may be unanticipated and can be difficult to manage at the grid level. During such events, an enrolled participant will be asked to curtail their load to help reduce grid demand. Participants will then be compensated.

"Third-party signaled" events require that the energy customer enter into a contract directly with the utility or with a curtailment service provider. Basically, you commit to turning off some loads or, alternatively, placing some amount of your load on battery. This of course is done in exchange for a payment. There are various forms of such programs that could be as simple as responding to occasional events; however, there are also more involved programs with specific capacity commitments or even second-by-second frequency regulation.

Capacity programs and frequency regulation
When sudden changes occur within a utility grid there is a temporary mismatch between supply and demand (e.g. a power plant or large factory goes offline). During such periods the frequency increases or decreases from its baseline (e.g. 50.0Hz). Frequency increases when there is more energy generation (i.e. supply) than consumption (i.e. demand) and frequency decreases when there is less energy generation than consumption. These frequency disturbances are caused by rotating generators temporarily spinning slightly faster or slower in response to large load changes. When participating in such programs the capacity need could be as short as a few seconds or as long as a few minutes. Typically, high speed and short duration participation is called frequency regulation and multi-hour commitments are called capacity participation. Frequency regulation is probably the most demanding form of grid participation with energy storage and often requires high speed local controls and metering to meet the stringent requirements. At the same time, however, frequency regulation can also be the most lucrative.

Third-party capacity aggregation

One way to increase participation in grid services is for multiple smaller assets to get aggregated virtually so they can be bid into more programs. In such cases geographically distributed assets are controlled monolithically and bid as an aggregated asset. Sometimes this is referred to as a virtual power plant (VPP).

Three modes of UPS operation for monetizing energy storage

There are three modes of UPS operation for monetizing energy storage using the methods described above. These are:

- Grid isolated
- Grid parallel
- Bidirectional

The following sections describe each of these modes and their pros and cons.

Grid isolated mode

Most data center operators are familiar with what would be described as grid isolated mode. It’s the operational mode when the UPS has two primary states. One where the entire UPS load is fed from the grid and one where the entire UPS load is fed from the battery. The later could be due to a power outage, or in the case of energy storage monetization, commanding the UPS to put the load on battery. **Figure 5** shows the power path (red line) for both the normal operation, and grid isolated operation. In grid isolated mode, the entire UPS load is run off the UPS batteries.

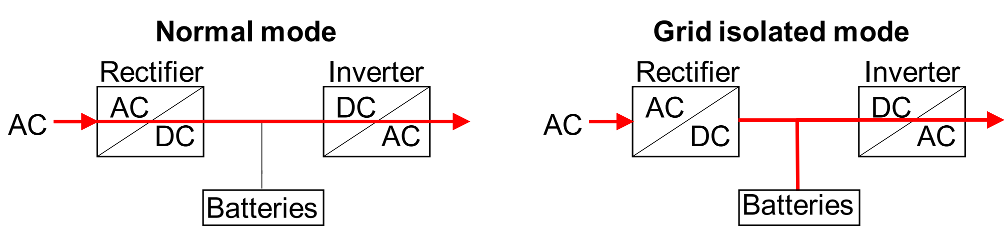

Figure 5
Difference between normal mode and grid isolated mode

For example, if you have 400 kW of IT load you can monetize 400 kW of IT load by placing it on battery for a given duration. Note that we chose a double conversion UPS topology to illustrate the modes of operation described in this paper. We call this mode “grid isolated” because the load is both electrically and physically isolated from the grid via contactors that open once the load is on battery. The purpose of these contactors is to prevent energy from potentially flowing back onto the grid.

Pros

- A standard mode of operation available in all UPS, therefore it is available at no additional cost.
- A proven mode of operation across all UPS topologies.

Cons

- There is no control of the exact portion of UPS load that is supported by the battery. It's either 100% fed from the grid or 100% from the battery.
- Because the exact amount of load on the batteries is uncontrollable, the amount of runtime is uncontrollable. This makes it harder to participate in many monetization programs.
- The capacity rating available for monetization is limited to the IT load on the UPS.

Grid parallel mode

The difference between this mode and the grid isolated mode is that, in this type of UPS, there is embedded control that allows for paralleling the battery and the rectifier on the DC bus in a controlled fashion. This allows for precisely controlled ratios of power to be provided from both the grid and the battery simultaneously. This is shown in **Figure 6**. In this mode, the UPS rectifier (fed from the grid) and battery work in parallel to feed the load. This is why this is referred to as grid parallel mode. This allows the operator to choose precisely how much load is curtailed. For example, if you have 400 kW of IT load and you're in agreement with the curtailment service provider to reduce your load by 300 kW for 2 hours, they can instruct the UPS to supply this 300 kW from battery while supporting the remaining 100 kW of load from the grid.

Figure 6
Difference between normal mode and grid parallel mode

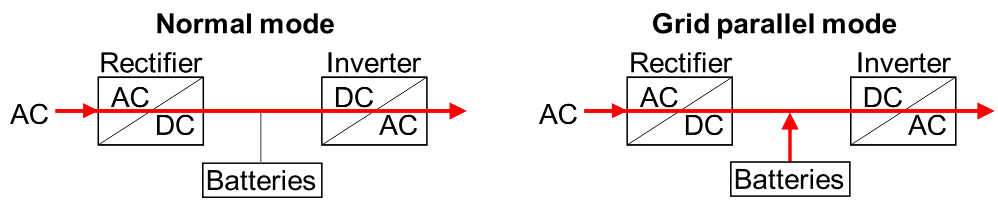

Pros

- This mode allows a specified portion of the UPS load to operate from the batteries.
- The battery life is increased because the load on the battery is typically lower in this mode. Consequently, the batteries experience lower operating stress.
- The battery runtime is increased due to discharging the batteries more slowly.
- This mode allows for a maximum number of combinations of load power and time on battery to maximize the utility incentive.
- The incremental cost of this mode is negligible because this feature is enabled through only software modifications in the UPS.

Cons

- This is a relatively new mode of operation that has not been UL verified to provide IEC 62040-3 Class 1 protection (by vendors other than Schneider Electric at this time).
- The capacity rating available for monetization is limited to the IT load on the UPS.

Bidirectional mode

In this mode, the UPS rectifier operates in both directions (bidirectionaly) as shown in **Figure 7**. In other words, it can convert AC to DC in normal operation or it can convert DC to AC, allowing it to take battery energy and inject it onto the grid. It's important to point out that this backfeeding mode of operation requires compliance with specific grid interconnection standards that are unique to each country. Examples include standards such as UL 1741 or IEC 62109. As a provider of power to the grid, rather than solely a load on the grid, this mode typically requires utility approval and some form of documented interconnection agreement.

Note: IEC 62109 and UL 1741 are specific grid-tied inverter safety and performance standards with several elements including inverter design safety requirements, power quality performance requirements, and grid interaction requirements. Each utility may have different revisions of these standards to which they may require compliance.

One potential benefit of bidirectional mode is that the entire UPS capacity is available for monetization. For example, if you have 1,000 kW of UPS capacity and 400 kW of IT load, you could place 400 kW of IT load on battery and inject 600 kW of battery energy into the grid for a given amount of runtime. In this scenario, you monetize 100% of the UPS capacity. Note that at this higher capacity, your runtime is significantly less than if your batteries supported only the 400 kW IT load.

Figure 7
Difference between normal mode and bi-directional mode

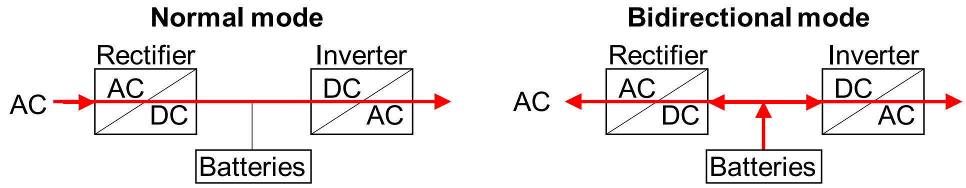

Pros

- Up to full UPS nameplate capacity available for participation.
- This mode allows for a maximum number of combinations of load power and time on battery to maximize the utility incentive.

Cons

- Requires UL1741, IEC 62109 or similar certification and no UPS vendors currently have this.
- Requires utility backfeed approval and some form of an interconnection agreement.
- Certification and utility approvals will slow down the installation process.
- Puts the batteries under the highest possible stress when backfeeding the grid and may later interfere with proper load protection.
- Induces earlier UPS wear-out due to asymmetrically loading the rectifier and inverter sections of the UPS.
- Increases UPS cost to enable bidirectionality in the rectifier and design for additional grid backfeed standards.

Decision criteria

All UPS, by definition, offer a "grid isolated mode". However, grid parallel and bidirectional modes are newer and not as common with UPS on the market today. The UPS mode that makes the most sense for you, is dependent on the monetization method chosen. For example, it wouldn't make sense to pay a premium for a UPS with bidirectional mode, if your monetization plan simply needed modest load reduction and didn't require backfeeding to the grid.

Decision criteria for choosing a UPS mode include:

- Specific requirements for participating in a demand response program
- Cost of implementing a UPS control system for demand response program participation
- Time required to plan, design, and install a UPS solution
- Cost and time required to secure an interconnection agreement
- Impact on the UPS battery warranty, lifetime, runtime, and reliability as a result of participating in demand response
- Incremental UPS component stress
- The cost premium for a UPS with grid parallel and/or bidirectional mode
- The battery runtime and UPS load required to participate in demand response
- Data center operations risk assessment for participating in demand response

All of the points above factor into the ultimate return on investment (ROI). In addition, utility programs and incentives vary widely by geography, so it's impossible to optimize the UPS mode for all conditions. However, based on certification complexity and the time requirement of utility interconnection agreements, as well as the highest stress on both the UPS and battery versus all other modes, **the bidirectional mode has a number of critical disadvantages and is not recommended**. Of the remaining two modes, the **grid-parallel operating mode is the best choice** due to the lowest amount of battery stress but the highest level of load flexibility.

Ultimately, a return on investment (ROI) analysis should inform a data center operator of the optimal UPS modes and future white papers will cover ROI calculations in greater detail.

Conclusion

As more and more renewable energy comes online, there will likely be a wider fluctuation of electricity prices; low when wind and solar are peaking, and high when they aren't, as well as increasingly available programs for loads to participate on the grid. This provides data center operators with both energy-saving and revenue earning opportunities. With incentives (i.e. tax breaks, rebates, etc.) to potentially help pay for batteries, and with the added revenue from participating in utility programs, data center operators may even want to add to their battery runtimes for purely economic reasons. In some cases, 1 hour of runtime may actually be cheaper than 10-15 minutes of runtime. Furthermore, with increasing growth in the electric vehicle market, Li-ion battery prices continue to drop. Ultimately this will drive more data centers to use Li-ion batteries instead of lead-acid batteries and make this option accessible to more operators.

There are three modes of UPS operation for monetizing energy storage using the methods described in this paper. Of these three modes, the **grid-parallel operating mode is the best choice** due to the lowest amount of battery stress but the highest level of load flexibility. Ultimately, a return on investment (ROI) analysis should inform a data center operator of the optimal UPS modes and future white papers will cover ROI calculations in greater detail.

About the authors

Chris Thompson is Vice President of the 3-phase UPS business at Schneider Electric and has over 20 years of diversified global experience with power conversion systems. Previous to his current role Chris was the Grid Power business unit manager at Eaton Corporation where he managed product lines for high power inverters that connect both solar and storage systems to the grid.
Chris also spent over 12 years at the APC and Xantrex Divisions of Schneider Electric in various roles in Japan, Singapore, Canada and the US.
Chris holds Bachelor's Degrees with highest honors in both Mechanical and Electrical engineering from Tufts University, an MBA in Marketing from the University of Chicago, and a Master's Degree in Electrical Engineering from the National University of Singapore where he served as a Fulbright Scholar.

Victor Avelar is the Director and Senior Research Analyst at Schneider Electric's Data Center Science Center. He is responsible for data center design and operations research, and consults with clients on risk assessment and design practices to optimize the availability and efficiency of their data center environments. Victor holds a bachelor's degree in mechanical engineering from Rensselaer Polytechnic Institute and an MBA from Babson College. He is a member of AFCOM.

RATE THIS PAPER ★★★★★

Resources

[The Different Types of UPS Systems](#)
White Paper 1

[Battery Technology for Data Centers and Network Rooms: Lead-Acid Battery Options](#)
White Paper 30

[FAQs for Using Lithium-ion Batteries with a UPS](#)
White Paper 231

[Browse all white papers](#)
whitepapers.apc.com

[Browse all TradeOff Tools™](#)
tools.apc.com

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
dcsc@schneider-electric.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm