



Find Your Product

## Data center design considerations

This article provides guidelines on distribution systems' levels of redundancy, the correct generator rating to use, and whether solar power can be used in a data center.

Debra Vieira, CH2M  
12/27/2017

Share



### Learning objectives:


- Know the difference between 2N, 3M2, and N+1 system topologies.
- Understand the characteristics of the system topologies.
- Learn the criteria for rating generators.
- Understand the different transformer types.

Over the past several years, mission critical clients seem to be asking the same series of questions regarding data center designs. These questions relate to the best distribution system and best level of redundancy, the correct generator rating to use, whether solar power can be used in a data center, and more. The answer to these questions is "It depends," which really doesn't help address the root of their questions. For every one of these topics, an entire white paper can be written to highlight the attributes and deficiencies, and in many cases, white papers are currently available. However, sometimes a simple and concise overview is what is required rather than an in-depth analysis. The following are the most common questions that this CH2M office has received along with a concise overview.

### What is the best system topology?



There isn't a single "best" system topology. There is only the best topology for an individual data center end user. The electrical distribution system for a data center can be configured in multiple topologies. While the options and suboptions can be myriad, the following topologies are commonly deployed (see Figure 1).



- **2N:** Simply designing twice as much equipment as needed for the base (i.e., N) load and using static transfer switches (STS), automatic transfer switches (ATS), and the information technology (IT) and HVAC equipment's dual cording to transfer the load between systems. The systems are aligned in an "A/B" configuration and the load is divided evenly over the two systems. In the event of failure or maintenance of one system, the overall topology goes to an N level of

redundancy.



Engineering Education

Power  
through-  
out the  
storm.



The natural  
choice.

NATURAL GAS  
GENERATORS

**GENERAC**  
INDUSTRIAL  
POWER



- **3M2:** This topology aligns the load over more than two independent systems. The distributed redundant topology is commonly deployed in a “three-to-make-two” (3M2) configuration, which allows more of the capacity of the equipment to be used while maintaining sufficient redundancy for the load in the event of a failure (see Figure 2). The systems are aligned in an “A/B/C” configuration, where if one system fails (e.g., A), the other two (B and C) will accept and support the critical load. The load is evenly divided with each system supporting 33.4% of the load or up to 66.7% of the equipment rating. In the event of a component failure or maintenance in one system, the overall topology goes to an N level of redundancy. In theory, additional systems could be supplied, such as 4M3 or 5M4, but deployment can significantly complicate the load management and increases the probability of operator error.
- **N+1 (SR):** The shared-redundant (SR) topology concept defines critical-load blocks. Each block is supported 100% by its associated electrical system. In the event of maintenance or a failure, the unsupported equipment would be transferred to a backup system that can support one or two blocks depending on the design. This backup system is shared across multiple blocks, with the number of blocks supported being left to the design team but typically in the range of 4:1 up to 6:1.
- **N+1 (CB):** The common-bus (CB) redundant system is like the shared redundant system in that the IT equipment’s A and B sources are connected to an N+1 uninterruptible power supply (UPS) source, but in the event of a failure or maintenance activities, the load is transferred to a raw power source via STS. The raw power source has the capability of being backed up by generators that are required to be run during maintenance activities to maintain the critical load.

The above topologies assume a low-voltage UPS installation. However, similar systems can be developed using a medium-voltage UPS. Beyond the redundancy configuration, these low-voltage UPS topologies also can be evaluated on ease-of-load management, backup power generation, their ability to deploy and commission initially and when expanding, first costs and total cost of ownership, physical footprint of the equipment comprising the topology, and time to construct the initial installation as well as expansion of the system.

A commonality between the different topologies presented is the need to transfer load between systems. No matter the system topology, the requirement to transfer load between electrical systems—either for planned maintenance activities, expansions, or failure modes—must be done. Load management refers to how the load is managed across multiple systems.


**2N topology.** The premise behind a 2N system is that there are two occurrences of each piece of critical electrical equipment to allow the failure or maintenance of any one piece without impacting the overall operation of the data center IT equipment. This configuration has a number of impacts:

- **Load management:** Among the topologies presented here, 2N has a relatively simple load-management scheme. The system will run independently of other distribution systems and can be sized to accommodate the total demand load of the IT block and associated HVAC equipment, minimizing the failure zone. The primary consideration for load management is to ensure the total load doesn’t overload a single substation/UPS system.
- **Backup power generation:** This topology uses a 2N backup generation with the simplest of schemes: having the generator paired to the distribution block. Each generator is sized for the entire block load and will carry 50% of the load under normal conditions. For large data centers, the option exists to parallel together multiple generator sets to create an “A” backup source and parallel together an equal number of generators to create a “B” backup source, distributing power via two different sets of paralleling switchgear. Typically, this is more expensive due to the addition of paralleling switchgear and controls. Selection of the voltage class usually depends on the size of load as well as physical space and cost to route cable from the generator to the switchgear. The ability to parallel generators tends to be limited by the paralleling switchgear bus ampacity ratings as well as short-circuit ratings. Beyond 6,000 amps at 480 V, consider using 15-kV-class generators.
- **Deployment:** Each 2N system can be designed to accommodate a discrete IT block. This allows multiple systems to be deployed independently, facilitating procurement, construction, commissioning, and operations with no impact to existing or future systems.
- **First cost/TCO:** The 2N system requires twice the quantity and capacity of electrical equipment than the load requires, causing the system to run at nominally 50% of nameplate capacity. Due to the nature of how electrical equipment operates, this tends to cause the equipment to run at a lower efficiency than can be realized in other topologies. An additional impact of the 2N system topology is that the first cost tends to be greater because of the quantity of equipment. Also, because there are additional systems in place, the ongoing operational and maintenance costs tend to be greater.
- **Spatial considerations:** Because it generally has the most equipment, the 2N configuration typically has the largest physical footprint. However, this system is the simplest to construct as a facility is expanded, thereby minimizing extra work and allowing the facility to grow with the IT demands.



- **Time to market:** As has been discussed, this system will have more equipment to support the topology, therefore there may be additional time to construct and commission the equipment. The systems are duplicates of each other, which allows for construction and commissioning efficiencies when multiple systems are installed, assuming the installation teams are maintained.

**Distributed redundant (3M2) topology.** The premise behind a 3M2 system is that there are three independent paths for power to flow, each path designed to run at approximately 66.7% of its rated capacity and at 100% during a failure or maintenance event. This configuration is realized by carefully assigning load such that the failover is properly distributed among the remaining systems.



This configuration has a number of impacts to the distribution:

- **Load management:** The load management for the 3M2 system should be carefully considered. The load will need to be balanced between the A, B, and C systems to ensure the critical load is properly supported without overloading any single system. Load management of a system like this can be aided by a power-monitoring system.
- **Backup power generation:** This topology follows the normal power flow and uses a 3M2 backup generation where the generator is paired to the distribution block. Each generator is sized for the entire block load and will carry 66.7% of their capacity under normal conditions. Parallel generator configurations are rarely used for 3M2 systems. Like 2N systems, the selection of the voltage class depends on the size of load as well as physical space and cost to route cable from the generator to switchgear (see Figure 3).
- **Deployment:** Each 3M2 system can be designed to accommodate a discrete IT block. Expansion within a deployed 3M2 system is exceptionally challenging and difficult, if not impossible to commission. Deployment of multiple 3M2 systems is the best option for addressing expansion and commissioning.
- **First cost/TCO:** The 3M2 system requires about 1.5 times the capacity of electrical equipment than the load requires and runs at 66.7% of its rated capacity. Because the equipment is running at a higher percentage, the 3M2 system tends to be more energy-efficient than the 2N, but less efficient than either of the shared redundant systems. An additional impact of the 3M2 system topology is that lower-capacity equipment can be used to support a similar size IT block, thereby causing the system to have a higher cost per kilowatt to install. However, if the greater capacity is realized by either sizing the IT blocks large enough to realize the benefits of this topology or by installing two IT blocks on each distribution system, then there will be a lower first cost. Essentially, the 2N system needs two substations and associated equipment for each IT block while the 3M2 system would need only three substation systems to support the IT block. First-cost savings is in addition to operational savings because there are fewer pieces of equipment to maintain. And the energy savings is because the equipment is running at a higher efficiency.
- **Spatial considerations:** Similar to the first-cost discussion above, the spatial layout can either be smaller or larger than a 2N system depending on how the topology is deployed and how many IT blocks each system supports.
- **Time to market:** The balance between the IT blocks supported by each system and the quantity of equipment will have an impact on the time to market, though the balance for this system is unlikely to be significant. The additional equipment should be balanced against smaller pieces of equipment, allowing faster installation time per unit. The systems are duplicates of each other, which allow for construction and commissioning efficiencies when multiple systems are installed, assuming the installation teams are maintained.

**N+1 shared redundant (N+1 SR).** The premise behind the N+1 SR system is that each IT block is supported by one primary path. In the event of maintenance or a failure, there is a redundant but shared module that provides backup support. The shared module in this topology has the same equipment capacities and configuration as the primary power system, minimizing the types of equipment to maintain.

For example, if six IT blocks are to be installed, then seven distribution systems (substations, generators, and UPS) will need to be installed for an N+1 system. This N+1 system can easily be reconfigured to an N+2 system with minimal impact (procuring eight systems in lieu of seven). This reconfiguration would allow the system to provide full reserve capacity even while a system is being maintained.



- **First cost/TCO:** For a large-scale deployment (i.e., exceeding two modules), the N+1 SR system has the lowest installed cost per kilowatt of the systems explored here that have full UPS protection for both the normal and redundant power distribution systems, due to the lower quantity of equipment. In addition, less equipment should also result in lower ongoing operation and maintenance costs.
- **Spatial considerations:** The N+1 SR layout will have the smallest spatial impact. Additional distribution is required between modules as well as a central location to house the redundant system.
- **Time to market:** The balance between the IT-blocks distribution system and the quantity of equipment will have an impact on the time to market. However, due to the fact that the N+1 SR has the smallest quantity of equipment, this configuration potentially has the shortest time to market of any system explored so far. This timing is further supported due to system duplicates, which should allow for construction and commissioning efficiencies on the subsequent installations, assuming the teams are maintained.

**N+1 common bus (N+1 CB).** The premise behind the N+1 CB system is there is one primary path that supports each IT block. This path also has an N+1 capacity UPS to facilitate maintenance and function in the event of a UPS failure. The system is backed up by a simple transfer switch system with a backup generator.

This configuration has a number of impacts on the distribution:

- **Load management:** Similar to the N+1 SR system, the load management for the N+1 CB is simple. As long as the local UPS/generator combination is not overloaded, the system will not be overloaded.
- **Backup power generation:** Like the previous topology, there is a generator paired to each distribution block including the redundant block.
- **Deployment/commissioning:** The deployment of the N+1 CB system is a modular deployment because each system functions independently. The only location where existing work has to be tested with the new equipment is on the common bus system.
- **First cost/TCO:** The N+1 CB system potentially has the lowest installed cost per kilowatt of any of the systems. This lower cost is due to a combination of lower quantities of UPS and generators coupled with simpler distribution. Additionally, less equipment means ongoing operation and maintenance costs should be lower as well.
- **Spatial considerations:** The N+1 CB layout will have a small spatial impact. Additional distribution is required between modules as well as a central location to locate the central bus system (transfer switches and generator).
- **Time to market:** Similar to the N+1 SR system, the N+1 CB has significantly fewer pieces of equipment than the 2N or 3M2 systems. This equipment count should support a faster time to market. However, it is difficult to determine which of the N+1 systems would have a quicker time to market.

The above topology descriptions only highlight a few systems. There are other topologies and multiple variations on these topologies. There isn't a ranking system for topologies; one isn't better than another. Each topology has pros and cons that must be weighed against the performance, budget, schedule, and the ultimate function of each data center.

### What generator rating should be used for a data center?

Generators need to be able to deliver backup power for an unknown number of hours when utility power is unavailable. To help select the appropriate generator, manufacturers have developed ratings for engine-generators to meet load and run time requirements under different conditions. The [International Standards Organization \(ISO\) Standard 8528-2005, Reciprocating Internal Combustion Engine Driven Alternating Current Generating Sets](#), tries to provide consistency across manufacturers. However, the ISO standard only defines the minimum requirements. If the generator is capable of a higher performance, then the manufacturer can determine the listed rating. To complicate generator ratings even more, some industries have their own ratings specific to that industry and application. These various ratings can make selecting the correct generator type complicated.



There are four ratings defined by ISO-8528:

1. **Continuous power:** designed for a constant load and unlimited operating hours; provides 100% of the nameplate rating for 100% of the operating hours.
2. **Prime power:** designed for a variable load and unlimited running hours; provides 100% of nameplate rating for a short period but with a load factor of 70%; 10% overload is allowed for a maximum of 1 hour in 12 hours and no more than 25 hours/year.
3. **Limited running:** designed for a constant load with a maximum run time of 500 hours annually; same nameplate rating as a prime-rated unit but allows for a load factor of up to 100%; there is no allowance for a 10% overload.
4. **Emergency standby power:** designed for a variable load with a maximum run time of 200 hours/year; rated to run at 70% of the nameplate.

The generator industry also has two additional ratings that are not defined by ISO-8528: mission critical standby and standby. Mission critical standby allows for an 85% load factor with only 5% of the run time at the nameplate rating. A standby-rated generator can provide the nameplate rating for the duration of an outage assuming a load factor of 70% and a maximum run time of 500 hours/year.

Data center designs assume a constant load and worst-case ambient temperatures. This does not reflect real-world operation and results in overbuilt and excess equipment. Furthermore, it is unrealistic to expect 100% load for 100% of the operating hours, as the generator typically requires maintenance and oil changes after every 500 hours of run time. Realistically during a long outage, the ambient temperature will fluctuate below the maximum design temperature. Similarly, the load in a data center is not constant. Based on [research performed by Caterpillar](#), real-world data center applications show an inherent variability in loads. This variability in both loads and ambient temperatures allows manufacturers to state that a standby-rated generator will provide nameplate power for the duration of the outage and it's appropriate for a data center application. However, if an end user truly desires an unlimited number of run hours, then a standby-rated generator is not the appropriate choice.



### What type of transformer is best?

The type of transformer to be used for a data center is constantly questioned and challenged by end users trying to understand if they should invest in a high-performance transformer or not. There are two categories of distribution transformers: dry-type and liquid-filled. Within each category, there are several different types. The dry-type category can be subdivided into five categories with the following features:

1. Open-wound transformers apply a layer of varnish on heated conductor coils and bake the coils until the varnish cures.
2. Vacuum-pressure impregnated (VPI) transformers are impregnated with a high-temperature polyester varnish, allowing for better penetration of the varnish into the coils and offering increased mechanical and short-circuit strength.
3. Vacuum-pressure encapsulated (VPE) transformer windings are encapsulated with silicon resin typically applied in accordance with a military spec and used in locations exposed to salt spray, such as shipboard applications with the U.S. Navy. VPE transformers are superior to VPI transformers, with better dielectric, mechanical, and short-circuit strength.
4. Encapsulated transformers have open wound windings that are insulated with epoxy, which makes them highly resistant to short-circuit forces, severe climate conditions, and cycling loads.
5. Cast-coil-type transformers have windings that are hermetically sealed in epoxy to provide both electrical and mechanical strength for higher levels of performance and environmental protection in high-moisture, dust-laden, and chemical environments.



For liquid-filled transformers, various types of fluids can be used to insulate and cool the transformers. These include less-flammable fluids, nonflammable fluids, mineral oil, and Askarel.

When put into the context of a mission critical environment, two transformers stand out: the cast-coil transformer due to its exceptional performance and the less-flammable liquid-immersed transformer due to its dependability and longevity in commercial and industrial environments. While both transformer types are appropriate for a data center, each comes with pros and cons that require evaluation for the specific environment.



Liquid-filled transformers are more efficient than cast coil. Because air is the basic cooling and insulating system for cast coil transformers, they will be larger than liquid-filled units of the same voltage and capacity. When operating at the same current, more material and more core and coil imply higher losses for cast coil. Liquid-filled transformers have the additional cooling and insulating properties associated with the oil-and-paper systems and tend to have lower losses than corresponding cast coil units.

Liquid-filled transformers have an average lifespan of 25 to 35 years. The average lifespan of a cast coil transformer is 15 to 25 years. Because liquid-filled transformers last longer than cast coil, they save on material, labor to replace, and operational impact due to replacement.





Recommended annual maintenance for a cast coil transformer consists of inspection, infrared examination of bolted connections, and vacuuming of grills and coils to maintain adequate cooling. Most times, cleaning of the grills and coils requires the transformer to be de-energized, which often leads to this maintenance procedure being skipped. The buildup of material on the transformer grills and coils can lead to decreased transformer efficiency due to decreased airflow.

Maintenance for a liquid-filled transformer consists of drawing and analyzing an oil sample. The oil analysis provides an accurate assessment of the transformer condition and allows for a scheduled repair or replacement rather than an unforeseen failure. This kind of assessment is not possible on a cast coil transformer. Additionally, omitting the oil sampling does not decrease the transformer efficiency.

Cast-coil-type transformers have a history of catastrophic failures within data centers due to switching induced transient voltages when switched by upstream vacuum breakers. There has been significant research by IEEE committees, which resulted in guidelines for mitigating techniques (i.e., resistive-capacitive [RC] snubbers) published in [IEEE C57.142-2010: IEEE Guide to Describe the Occurrence and Mitigation of Switching Transients Induced by Transformers, Switching Device, and System Interaction](#). Liquid-filled transformers seem less susceptible to this problem, as there is no published data on their failure. Regardless of the transformer type installed, best industry practice is to perform a switching transient study and install RC snubbers on the systems if warranted.

When a transformer fails, a decision must be made on whether to repair or replace it. Cast coil transformers typically are not repairable; they must be replaced. However, there are a few companies who are building recyclable cast coil transformers. On the other hand, in most cases, liquid-filled transformers can be repaired or rewound.

When a cast-coil transformer fails, the entire winding is rendered useless because it is encapsulated in epoxy resin. Because of the construction, the materials are difficult and expensive to recycle. Liquid-filled transformers are easily recycled after they've reached the end of their useful life. The steel, copper, and aluminum can be recycled.

Cast-coil transformers have a higher operating sound level than liquid-filled transformers. Typical cast coil transformers operate in the 64 to 70 dB range while liquid-filled transformers operate in the 58 to 63 dB range. A decibel is a logarithmic function and sound pressure doubles for every 3-dB increase.

Liquid-filled transformers have less material for the core and coil and use highly effective oil and paper cooling systems, which allow them to be small in physical footprint and weigh less than the corresponding cast coil unit. Because cast coil transformers are air-cooled, they are often larger than their liquid counterparts assuming the same voltage and capacity (kVA rating). Cast coil transformers have more core material, which implies higher costs and losses.

Dry-type transformers have the advantage of being easy to install with fire-resistant and environmental benefits. Liquid-filled transformers have the distinct disadvantage of requiring fluid containment. However, advances in insulating fluids, such as Envirotemp FR3 by Cargill, a natural ester derived from renewable vegetable oil, is reducing the advantages of dry-type transformers.

For indoor installations of transformers, cast coil must be located in a transformer room with minimum 1-hour fire-resistant construction in accordance with [NFPA 70-2017: National Electrical Code \(NEC\)](#) Article 450.21(B). However, if less-flammable liquid-insulated transformers are installed indoors, they are permitted in an area that is protected by an automatic fire-extinguishing system and has a liquid-confinement area in accordance with NEC Article 450.23.

Traditionally less-flammable liquid-filled transformers are installed outdoors. However, both types can be installed outdoors. This option of outdoor installation has the additional advantage of reducing data center cooling requirements. In this case, cast coil transformers need to have a weatherproof enclosure and cannot be located within 12 in. of combustible building materials per NEC Article 450.22. The liquid-filled transformer must be physically separated from doors, windows, and similar building openings in accordance with NEC Article 450.27.

The choice between a cast coil and a less-flammable liquid-filled transformer can be a challenging one to make. A liquid-filled transformer is a solid choice for a data center application because it is more efficient, physically smaller and lighter, quieter, recyclable, and has a longer lifespan. However, if the demand for high electrical and mechanical performance is of the utmost concern, then cast coil would be the appropriate choice.

#### **What IT distribution voltage should be used?**

By now it's well understood in the data center industry that 3-phase circuits can provide more power to the IT cabinet than a single-phase circuit. However, the choice of distribution voltage between 208 Y/120 V or 415 Y/240 V depends on the answers to several questions, such as:

- How much power needs to be delivered to each IT cabinet initially, and what does the power-growth curve look like for the future?
- What are the requirements of the IT equipment power supplies?
- Will legacy equipment be installed in the data center?
- Can the facilities team decide on the power supplies to be ordered when new IT equipment is purchased?

Let's start with the power of a 3-phase circuit. A 208 Y/120 V, 3-phase, 20-amp circuit can power up to a 5.7-kVA cabinet. Per NEC Article 210.20, branch-circuit breakers can be used up to 80% of their rating, assuming it's not a 100%-rated device. Therefore, a 208 V, 3-phase, 20-amp circuit can power a cabinet up to 5.7 kVA (20 amps x 0.8 x  $\sqrt{3}$  x 208 V). Now, if that same 20-amp circuit was operating at 415 Y/240 V, 3-phase, then that circuit could power a cabinet up to 11.5 kVA (20 amps x 0.8 x  $\sqrt{3}$  x 415 V). That's more than twice the power from the same circuit for no extra distribution cost.

If the specifications for the IT equipment can be tightly controlled, the decision to standardize on 415 Y/240 V distribution is a pretty simple one. However, if the IT environment cannot be tightly controlled, the decision is more challenging. Currently, most IT power supplies have a wide range of operating voltage, from 110 V to 240 V. This allows the equipment to be powered from numerous voltage options while only having to change the plug configuration to the power supply. However, legacy equipment or specialized IT equipment may have very precise voltage requirements, thereby not allowing for operation at the higher 240 V level. To address this problem, both 208 Y/120 V and 415 Y/120 V can be deployed within a data center, but this is rarely done as it creates confusion for deployment of IT equipment.

The follow-on question typically asked is if the entire data center can run at 415 V, rather than bringing in 480 V and having the energy loss associated with the transformation to 415 V. While technically feasible, the equipment costs are high because standard HVAC motors operate at 480 V. Use of 415 V for HVAC would require specially wound motors, thus increasing the cost of the HVAC equipment.

### **Must we install an emergency power-off system?**

Emergency power-off (EPO) buttons are the fear of every data center operator. With the push of a button, the entire data center power and cooling can be shut down. Because of the devastation that activation of an EPO can cause, EPOs typically are designed with a two- or three-step activation process, such as lifting a cover and pressing the button or having two EPO buttons that must be activated simultaneously. These multistep options assume that the authority having jurisdiction has provided approval for such a design. However, EPOs are not necessarily required. The need for an EPO is typically triggered by NEC Article 645.10, which allows alternative and significantly relaxed wiring methods in comparison with the requirements of Chapter 3 and Articles 708, 725, and 770. These relaxed wiring methods are allowed in exchange for adding an EPO system and ensuring separation of the IT equipment's HVAC occupancies from other occupancies. The principle benefit of using Article 645.10 is to allow more flexible wiring methods in the plenum spaces and raised floors. However, if the wiring is compliant with Chapter 3 and Articles 708, 725, and 770, the EPO is not required.

### **Can we use photovoltaic systems to power our data center?**

Corporations and data center investors are demanding sustainability be built into the data center. The positive impact on public relations by showcasing a sustainable data center can't be underestimated, especially considering how much of a power hog data centers can be. Additionally, many utility companies will offer incentives for the use of energy-efficient and sustainable technologies. An often-questioned item is whether photovoltaic (PV) systems can be used to meet some of the sustainability requirements in a data center environment. The answer is yes, but a good understanding of PV systems and the limitations and impacts on a data center are required prior to making the investment (see Figure 4).

The power production of PV equipment varies considerably depending on type and location of the system installed. There are three main types of solar panel technologies. Crystalline silicon (c-Si) is the most common PV array type, along with thin-film and concentrating PV. Thin-film is generally less efficient than c-Si, but also less expensive. Concentrating PV arrays use lenses and mirrors to reflect concentrated solar energy onto high-efficiency cells. Concentrating PV arrays require direct sunlight and tracking systems to be most effective and are typically used by utility companies.

Solar cells are not 100% efficient. In the infrared region of light, solar cells are too weak to create electricity; and in the ultraviolet region of light, solar cells create heat instead of electricity. The amount of power that can be generated with a PV array also varies due to the average sunshine (insolation, or the delivery of solar radiation to the earth's surface) along with the temperature and wind. Typically, PV arrays are rated at 77°F, allowing them to perform better in cold rather than in hot climates. As temperatures rise above 77°F, the array output decays (the amount of decay varies by type of system). Ultimately, what this means is that the power generation of an array can vary over the course of a day and year. Added to this are the inefficiencies of the inverter, and if used, storage batteries.

The physical space required to install the PV array can be significant. A simple rule is to assume 100,000 sq ft (about 2.5 acres) for a 1-MW PV-generating plant. However, this does not include the space required for access or other ground-mounted appurtenances. The total land required is better estimated at about 4 acres per MW. This estimate assumes a traditional c-Si PV array (without trackers). [Increase this area by 30%, for a total of about 6 acres, if thin-film technology \(without trackers\) is used due to the inefficiencies of the technology.](#)

A PV system may or may not provide power during a utility power failure, depending on the type of inverter installed. A standard grid-tied inverter will disconnect the PV system from the distribution system to prevent islanding. The inverter will reconnect when utility power is available. An interactive inverter will remain connected to the distribution system, but it is designed to only produce power when connected to an external power source of the correct frequency and voltage (i.e., it will come online under generator power). Typically, interactive inverters include batteries to carry the system through power outages, therefore the system should be designed such that there is enough PV-array capacity to supply the load and charge the batteries.

Most data centers do not have the necessary land to install a PV system, which substantially offsets the power demand. Then there is the question of what happens when the PV system is generating low or no power. Interactive inverters and deep-cycle storage batteries can be installed to cover these low-PV production periods, but this introduces new equipment, maintenance, and space requirements into the data center, thus creating more costs and more maintenance than may have been originally envisioned. Generally, data center sustainability is addressed more directly through efficient cooling and electrical distribution systems. Sustainability achieved through solar power, while nice to have, is generally not the focus of data center investments.

The trend is to provide a PV system that offsets some of the noncritical-administration power usage. These systems are typically small (less than 500 kW) and can be located on building rooftops, carports, and on the ground. They use a standard grid-tied inverter connected through the administration electrical distribution system, which ultimately ties into the site-distribution switchgear where the utility meter resides. A grid-tied inverter system will disconnect from the utility if there is a failure or when on generator power.

Because the grid-tied inverter connection is downstream from the utility revenue meter, a billing mechanism known as net metering generally is used. With net metering, owners are credited for any electricity they may add to the grid when the PV production is greater than the site usage. Although in most data centers, the critical load dwarfs the noncritical load; therefore it's rare that a PV system would generate power on the grid. There are differences between states and utility companies regarding the implementation, regulations, and incentives for net metering. Furthermore, there are some utility companies that perceive net metering as lost revenue and will not allow connection to their system.

A great resource for PV and renewable energy in general is the [National Renewable Energy Laboratory](#). The NREL website provides information on PV research, applications, publications as well as a free online tool, [PVWatts](#), which estimates the energy production and cost of energy of grid-tied PV systems throughout the world. The PVWatts tool easily develops estimates of potential PV-installation performance.

### **A medium-voltage alternative to low-voltage UPS**

Design topology evaluation also should consider the medium-voltage uninterruptible power supply (UPS). Like the topologies using the low-voltage UPS, the medium-voltage UPS can be deployed in 2N, N+1, and 3N/2 configurations. Regardless of the topology used, medium-voltage UPS systems offer advantages over low-voltage UPS systems. They generally are installed outdoors in containers, thereby minimizing the conditioned building footprint. While not required, medium-voltage UPS topologies are typically used for full-facility protection rather than using independent information technology and mechanical-cooling UPS systems, further reducing the building footprint. Medium-voltage UPS systems are large systems, starting at 2.5 MVA and scalable up to 20 MVA per UPS. Different manufacturers have different voltage offerings, but medium-voltage UPS systems can range from 5 kV up to 25 kV, with medium-voltage diesel rotary UPS systems going as high as 34.5 kV.

In early 2018, Michigan State University is expected to complete construction on a new 25,000-sq-ft data center with 10,600 sq ft of server space and initially hosting about 300 server racks. This \$46 million facility will use a medium-voltage UPS system, starting with 2.5 MW of critical power and the ability to increase in critical power as needed. The utility infrastructure is built to support an increase of load up to 10 MW. Figures 5, 6, and 7 highlight the outdoor power electronic switch, switchgear, and the medium-voltage UPS installed by the university.

---

**Debra Vieira** is a senior electrical engineer at [CH2M](#), Portland, Ore., with more than 20 years of experience for industrial, municipal, commercial, educational, and military clients globally.

---

