Chris Loeffler
Product Manager, Eaton

in

Ed Spears
Product Marketing
Manager, Eaton

in **y**

UPS basics

Everything you ever wanted to know about uninterruptible power supplies but were afraid to ask.

Executive summary

Budgeting for electricity, securing adequate supplies of it and finding ways to use less of it are all common topics of conversation among data center operators. Ensuring that the power their IT resources rely on is both dependable and clean, sadly, can sometimes be an afterthought.

In truth, however, power sags, surges and outages are not only unavoidable but more than capable of damaging valuable IT equipment and bringing productivity to a halt. That's why planning and deploying a robust power protection solution is absolutely vital.

An uninterruptible power system (UPS) is the central component of any well-designed power protection architecture. This white paper provides an introductory overview of what a UPS is and what kinds of UPS are available, as well as a comprehensive guide to selecting the right UPS and accessories for your needs.

Table of contents

F·A·T·N

Powering Business Worldwide

Why is power protection important?

No company can afford to leave its IT assets unprotected from power issues. Here are just a few of the reasons why:

- **Even short outages can be trouble.** Losing power for as little as a quarter second can trigger events that may keep IT equipment unavailable for anywhere from 15 minutes to many hours. And downtime is costly. Some experts believe the U.S. economy loses between $200 billion and $570 billion a year due to power outages and other disturbances.

- **Utility power isn't clean.** In practice, electrical power can vary widely enough to cause significant problems for IT equipment. According to current U.S. standards, for example, voltages can vary up to 8.3 percent from absolute specifications. That means that utility services promising 208-phase voltage may actually deliver 191 to 220 volts.

- **Utility power isn't 100 percent reliable.** In the U.S., in fact, it's only 99.9 percent reliable, which translates into a likely nine hours of utility outages every year.

- **The problems and risks are intensifying.** Today's storage systems, servers and network devices use components so miniaturized that they falter and fail under power conditions earlier-generation equipment easily withstood.

- **Generators and surge suppressors aren't enough.** Generators take time to startup, therefore there is a loss of AC power until the generator comes on-line, (typically 10 seconds), and they provide no protection from power spikes and other electrical disturbances. Generators however, are used to provide AC power during long utility power outages. Surge suppressors help with power spikes but not with issues like power loss, under-voltage and brownout conditions.

- **Availability is everything these days.** Once, IT played a supporting role in the enterprise, but now it's absolutely central to how most companies compete and win. When IT systems are down, core business processes quickly come to a standstill.

- **Availability is everything, but power costs must be managed.** The cost of power and cooling has spiraled out of control in recent years. Data center managers are typically held responsible for achieving high availability while simultaneously reducing power costs. Highly-efficient UPS systems can help with this goal, and products are available today that were not an option even a few years ago.

What is a UPS?

Put simply, a UPS is a device that:

1. Provides power conditioning and backup power when utility power fails, either long enough for critical equipment to shut down gracefully so that no data is lost, or long enough to keep required loads operational until a secondary AC source, like a generator, comes online.

2. Conditions incoming power so that all-too-common sags and surges don't damage sensitive electronic gear.

What are the main types of UPS?

UPSs come in three major varieties, which are also known as topologies:

Single-conversion systems

In normal operation, these feed incoming utility AC power to IT equipment. If the AC input supply falls out of predefined limits, the UPS utilizes its inverter to draw current from the battery, and

also disconnects the AC input supply to prevent backfeed from the inverter to the utility. The UPS stays on battery power until the AC input returns to normal tolerances or the battery runs out of power, whichever happens first. Two of the most popular single-conversion designs are standby and line-interactive.

- *Standby UPSs* allow IT equipment to run off utility power until the UPS detects a problem, at which point it switches to battery power. Some standby UPS designs incorporate transformers or other devices to provide limited power conditioning as well. Standby UPSs are typically the lowest cost of all UPS types, but they may sacrifice some performance specifications found in other UPS designs. One of the typical design traits is that the standby UPS may not create a sine wave output on battery, but may output a stepped-wave or modified sine wave, that looks more like a "square" wave. Small desktop PCs may be fine with this waveform, but larger servers, networking equipment and storage units, with more complex power supplies, may have issues and not operate properly, or at all, with this type of modified waveform.

- *Line-interactive UPSs* regulate input utility voltage up or down as necessary before allowing it to pass through to protected equipment. However, like standby UPSs, they use their battery to guard against frequency abnormalities.
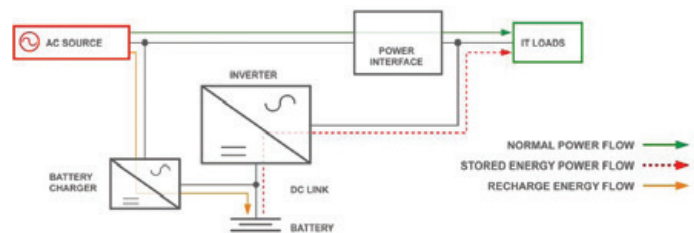

Figure 1. Internal design of a line-interactive UPS.

Double-conversion systems

As the name suggests, these devices convert power twice. First, an input rectifier converts AC power into DC and feeds it to an output inverter. The output inverter then processes the power back to AC before sending it on to IT equipment. This double-conversion process isolates critical loads from raw utility power, ensuring that IT equipment receives only clean, reliable electricity.

In normal operation, a double-conversion UPS continually processes power twice. If the AC input supply falls out of predefined limits, however, the input rectifier shuts off and the output inverter continues to operate, drawing power from the battery instead. The UPS continues to utilize battery power until the AC input returns to normal tolerances or the battery runs out of power, whichever occurs sooner. In case of a severe overload of the inverter, or a failure of the rectifier or inverter, the static switch bypass path is turned on quickly, to use unconditioned AC source power to support the output loads.
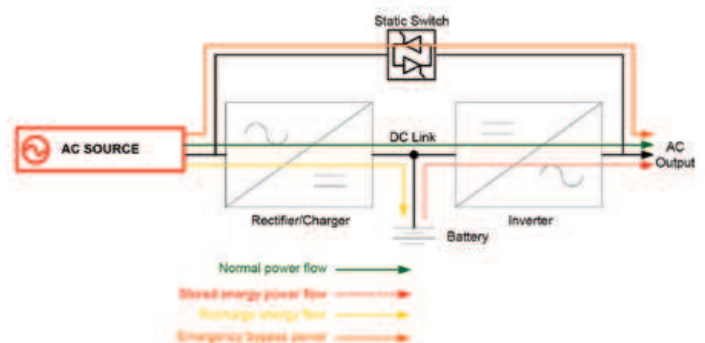

Figure 2. Internal design of a double-conversion UPS.

Multi-mode systems

Multi-mode UPSs combine features of both single- and double-conversion technologies to dynamically strike an ideal balance between efficiency and protection. Under normal conditions, they provide maximum efficiency by running in a highly-efficient energy saver mode. When problems occur, however, they automatically sacrifice some efficiency to deliver maximum levels of protection by automatically and immediately transitioning back to double-conversion mode. When the power quality returns to acceptable levels, the UPS transitions back into energy saver mode.

The end result is that data centers can save tens of thousands a year on energy without compromising data center performance or reliability.
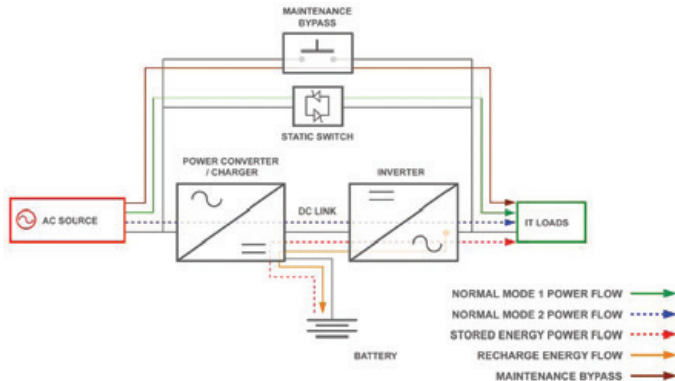

Figure 3. Internal design of a multi-mode UPS.

How do I choose the right UPS for the job?

To ensure that you always end up with the right UPS for your needs, be sure to consider these eight issues during your selection process:

Topology

Question number one is should you get a single-conversion, double-conversion or multi-mode UPS? The answer depends largely on how important energy efficiency is for your organization relative to protection.

Single-conversion UPSs are more efficient than double-conversion devices, but offer less protection. That makes them a good fit for loads with a higher tolerance for power anomalies, or where the incoming power is more stable. More specifically, standby UPSs (the most basic type of single conversion UPS) are generally the best option for smaller applications, like desktop and point-of-sale solutions, while line-interactive UPSs are typically preferable for smaller server, storage and network applications located in facilities with access to relatively trouble-free AC utility power.

Double-conversion UPSs, which provide the highest levels of isolation from the utility AC source, are less efficient but are usually the standard choice for protecting mission-critical systems. They are also the most preferred topology in electrically poor environments, such as industrial applications.

Multi-mode UPSs, although they may be more expensive than either single or double conversion systems, are the best choice for companies looking to achieve an optimal blend of both efficiency and protection

Single phase versus three phase

When the utility generates power, it is at the three-phase level. This type of power is available to almost all commercial and industrial customers, as they are typically large consumers of power. Three-phase power uses three separate "phase" wires, which allow higher power to be delivered most efficiently to a single point or load. Most homeowners only have single-phase power available to them, as homes are typically small power users. Single-phase power is delivered through either one or two phase wires, which are derived from the utility three-phase power system using transformers.

Single-phase UPSs tend to be a sensible and economical option for simpler, smaller applications with low kVA requirements, which are typically seen in the home, small businesses and in remote or satellite offices where computing power is less than 20,000 VA. Three-phase UPSs are generally the preferred choice for high kVA applications, which are typically more sophisticated and have high compute densities. Large multi-story buildings, data centers and industrial facilities protecting high-power processes are typical three-phase UPS customers, as they need to distribute large amounts of power over relatively long distances.
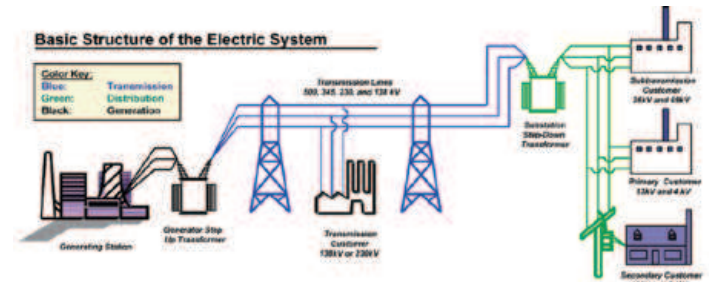

Figure 4. Three-phase power is generated and distributed to large commercial customers, with secondary customers, like homeowners, only receiving single phase power.

Power rating

A UPS's power rating is the amount of load, in volt-amperes (VA), that it's designed to support. UPSs are available with ratings as low as 300 VA and as high as 5,000,000 VA or more. Use this very basic procedure to determine the approximate UPS rating your organization requires:

1 Make a list of all the equipment your UPS will be protecting.

2 Determine how many volts and amps every device on the list draws.

3 For each device, multiply volts by amps to arrive at a VA figure.

4 Add all of the VA figures together.

5 Multiply that sum by 1.2, to build in room for growth.

The UPS you buy should have a rating equal to or greater than the final number you arrived at in step 5, unless you have more precise load data for the equipment you are protecting. Here are a few additional considerations to keep in mind:

- When deploying a centralized power protection architecture, you typically deploy larger kVA UPSs than you would deploy using a distributed power protection scheme.

- If your UPS will be supporting motors, variable-speed drives, medical imaging devices or laser printers, add more VA capacity to your requirements to account for the high power inrush that occurs when those devices startup. Your UPS vendor can assist in applying the proper UPS and rating for these types of applications.

- Companies that anticipate rapid near- or medium-term growth should use a multiple higher than 1.2 when building in room for growth in the procedure above. Organizations that expect to upgrade server hardware soon should also accommodate for more growth, as newer servers tend to have higher power requirements than older models.

Form factor

UPSs come in a range of form factors that fit into two master categories: rack-mounted and freestanding. The largest UPSs aren't available in rack-mounted form factors, so companies with substantial power requirements will use freestanding devices. For companies with more modest needs, deciding between rack-mounted and freestanding UPSs is largely a matter of data center design philosophy. Some organizations use rack-mounted UPSs in an effort to consolidate as much hardware as possible in their enclosures. Others prefer to maximize the amount of rack space available for servers by using freestanding UPSs. From a technical and financial standpoint, neither approach is inherently superior to the other.

Availability

Organizations can utilize a variety of deployment options, technologies and services to increase the reliability of their power protection solution. Here are some of the most effective ones:

Redundant deployment architectures: Deploying UPSs in redundant groups can increase availability by ensuring that critical loads remain protected even if one or more UPSs fail. There are three main kinds of redundant UPS architecture:

- **Zone:** In a zone architecture, one or more UPSs provide dedicated support for a specific set of data center resources. That way if a UPS fails during a power outage, the impact is limited to the zone that device supports.

- **Serial:** In a serial architecture, multiple UPSs are connected end to end such that if any one UPS in the string fails, the others can compensate automatically.

- **Parallel:** Parallel architectures use multiple independent parallel-connected UPSs to achieve increased redundancy. If any UPS fails completely, the other systems can keep protected information technology equipment (ITE) loads operational.
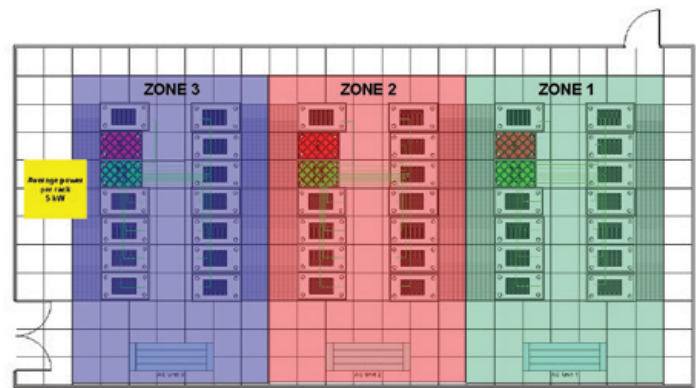

Figure 5. Zone protection uses separate protected "zones" and may use virtualization software to shift loads to other zones during failures or maintenance. Each zone has its own 60kW UPS system.
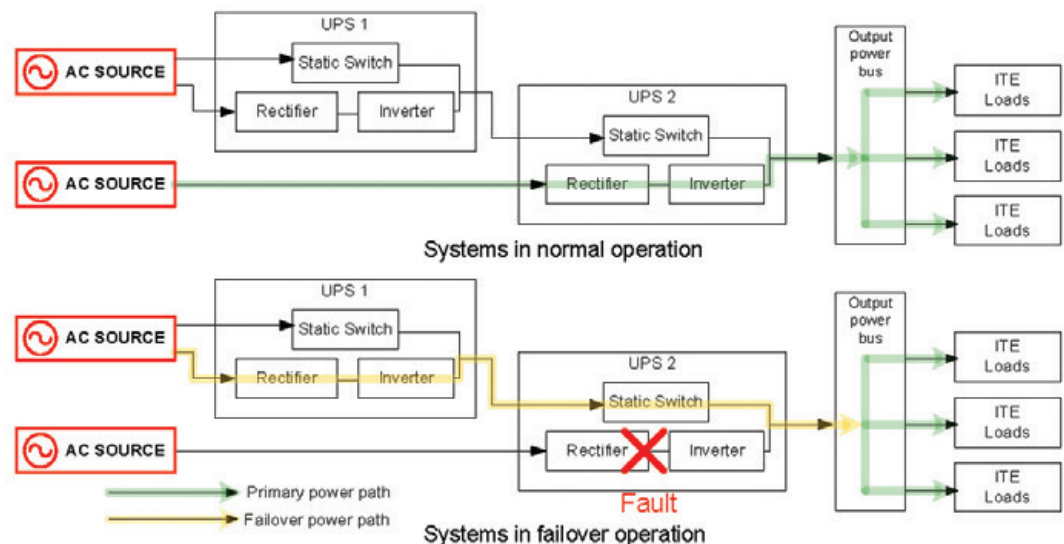

Figure 6. Serial (cascade) UPS architecture with changing power path if the UPS under load fails.
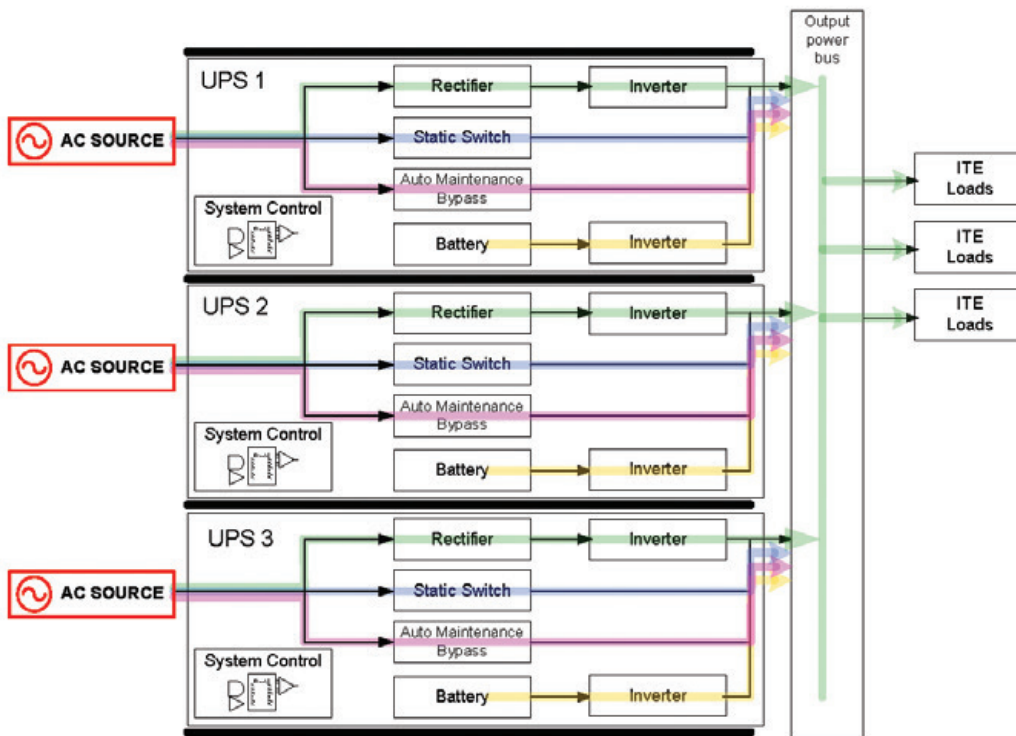
Figure 7. Parallel UPS systems all feed the output bus, so any single UPS module can be isolated for maintenance or in case of a failure. Parallel systems must be synchronized together to share the loads.

Hot-swappable components: Technicians can repair or administer a UPS that uses hot-swappable components without powering the UPS down, thereby exposing IT equipment to decreased risk of downtime.

Figure 8. Hot swappable battery modules allow service work to be performed while load equipment remains protected.

Extending battery runtime: A typical UPS battery provides five to 15 minutes of backup power. Organizations that need more than that can use supplemental external battery modules or cabinets to add as much as several hours of emergency runtime at full load.

Battery management: One of the most important parts of the UPS system is the energy storage system, which is typically a battery. Most power protection solutions get their emergency standby power from either sealed batteries (also known as Valve Regulated Lead Acid, or VRLA, batteries) or flooded batteries (also known as Vented Lead Acid, or VLA, batteries). Sealed batteries tend to be less expensive, but also wear out sooner. Flooded batteries generally require specialized installation and maintenance. Deciding which kind of battery is right for you basically comes down to whether you're willing to pay more upfront for batteries that you won't have to replace as often.

Many UPS systems continually "trickle-charge" the battery, which has a tendency to degrade the battery's internal chemical composition over time, reducing battery service life. While large banks of flooded electrolyte batteries for high power (greater than 500kVA) UPS systems need to have the battery trickle charged, most non-spillable VRLA batteries used in today's lower kVA UPSs can benefit from a charging technique where the charger turns off periodically and "rests" the battery like Eaton's Advanced Battery Management. This technique helps increase battery service life by as much as 50 percent with a three-stage charging process. Such technology can also provide notification prior to battery failure by running tests that detect significant change from the normal.. Batteries utilizing this technology last longer and are more reliable.

Scalability and modularity: Implementing a robust power protection solution takes time and money. To get as much value as possible from that investment, companies should estimate their needs for the next three to five years when comparing UPS options. If you're likely to have significantly larger power requirements in that timeframe, choose appropriately larger UPS hardware.

Data centers with fluctuating needs or future requirements that are difficult to forecast can employ two strategies to increase the scalability of their UPS deployment:

- **Deploy UPSs in parallel:** Parallel UPS architectures boost scalability as well as redundancy. As their power needs increase, organizations can simply add more modules to existing UPS systems, rather than replace current devices with new ones.

- **Use modular UPS products:** Some newer UPSs feature modular designs that allow you to add capacity incrementally as requirements increase. For example, some such systems provide up to 50 or 60 kW of capacity in 12 kW building blocks that fit in standard IT equipment racks. As requirements increase, another 12 kW unit can simply be plugged in. For mid-sized power requirements, building blocks of 50kW may be the best financial solution, allowing scaling from 50 or 100 kW up to 300 to 400kW very easily. Even the latest generation large UPS systems are available in modular 200 to 300 kW increments. That's a scalable and efficient approach to keeping up with escalating power needs that also lowers upfront capital spending and conserves equipment room floor space.

Keep in mind that, for either of the above options, the premises wiring should be sized for the maximum possible UPS size, so that future re-wiring will not be necessary as the UPS system grows.

Figure 9. The UPS management software in this example plugs into the administrator's console for
VMware vCenter, allowing the IT administrator to see and manage UPS-related events or alarms in the
same system as their IT management tasks.

Making the most out of your UPS selection

Offsite power monitoring:

The best way to address UPS problems is to prevent them. Remote UPS monitoring applications continually watch for warning signs of future trouble, such as deteriorating performance or battery weakening of failure, and send real-time notification when potential issues develop. This enables technicians to make repairs before serious breakdowns have a chance to occur. Data centers can perform remote monitoring themselves or contract with an external provider to do it for them. Another benefit to remote monitoring is reporting capabilities on the health of your UPS network.

Onsite power management:

If you do not monitor your power infrastructure (UPSs, rack power distribution units (PDUs) and bypass switches), you are at risk for blind spots on how the power at your location is behaving or the effect it has on your IT equipment. Even with a UPS, your IT system could still go down in case of an extended power failure or if the UPS gets overloaded for too long.

UPSs come with software that signals servers that AC power is lost and the UPS is on battery. If AC does not return and the battery energy is near depletion, the software closes all open applications to prevent any data loss. When AC power is restored, the IT systems can be programmed to automatically reboot, bringing the system back to its previous state. This solution may be sufficient on small PC servers protected by a single UPS, but are not sufficient in today's highly virtualized and consolidated IT environments.

Today's IT application complexity requires that UPSs be on the IT network to provide continual operational data. To accomplish this, the UPS needs an Ethernet port, or communications expansion card, which allows it to have an IP address on the network. These communication cards allow the UPS to transmit industry standard (SNMP, Modbus, BACnet IP, etc.) data packets, which can be monitored by a corresponding management platform. Some of these communication cards also allow web browsers to access the status of the UPS and includes the ability to send notifications via email or SMS (text) concerning power issues or UPS status and alarm conditions.

Leading-edge UPS manufacturers have developed software platforms that reduce management complexity by integrating the UPS software into virtualization management platforms like VMware's vCenter® or Citrix XenCenter®. In these environments, one single software installation can control and shutdown any cluster of servers by reading attributes from both power equipment and environmental data like temperature and humidity. At minimum, onsite power management needs to save data in case of an environmental issue or power outage, but better yet, advanced software allows MSPs to increase business continuity by keeping mission-critical applications running.

If power fails, onsite power management software saves data automatically by doing a live migration to a safe server in a virtualized environment, or initiates a shutdown sequence in a stand-alone server configuration. With advanced software, systems have more sophisticated features such as load shedding or virtually sending mission-critical applications to an offsite location.
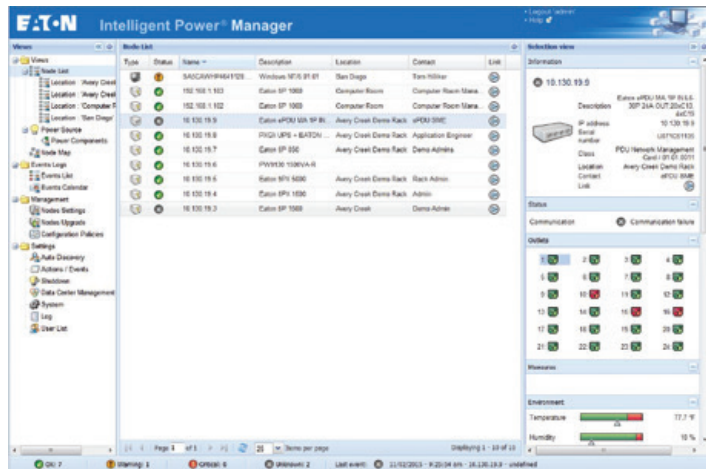
Load shedding reduces the workload on the IT devices so power consumption is minimized. Load shedding involves suspending or shutting down non-critical workloads (applications) and combining the higher critical workloads on fewer servers, then shutting down as many servers as possible. It can be used in both small and large data centers. In a small data center or network closet, it can be used to extend the battery runtime, especially when there is no generator to pick up the electrical load when power gets disrupted. In a large data center, it may involve a major disaster like a hurricane or ice storm where power may be out for several days or longer.

Software can also initiate a move of critical applications to a backup data center, also called a disaster recovery (DR) site. It can read environmental conditions (temperature over the limit) or UPS power alarms (utility power lost) from the UPS or rack PDU. If a small data center is using intelligent power equipment and receives an alarm that utility power is lost or the temperature is over the high limit, the advanced software initiates the process to move critical VMs/applications over to a DR site. It's important that the failover application first syncs the backup site to the primary site before starting the VMs at the DR site so there is no data loss. Since SMB data centers generally do not back up continuously, but daily or weekly, the synchronization is especially important to avoid data loss.
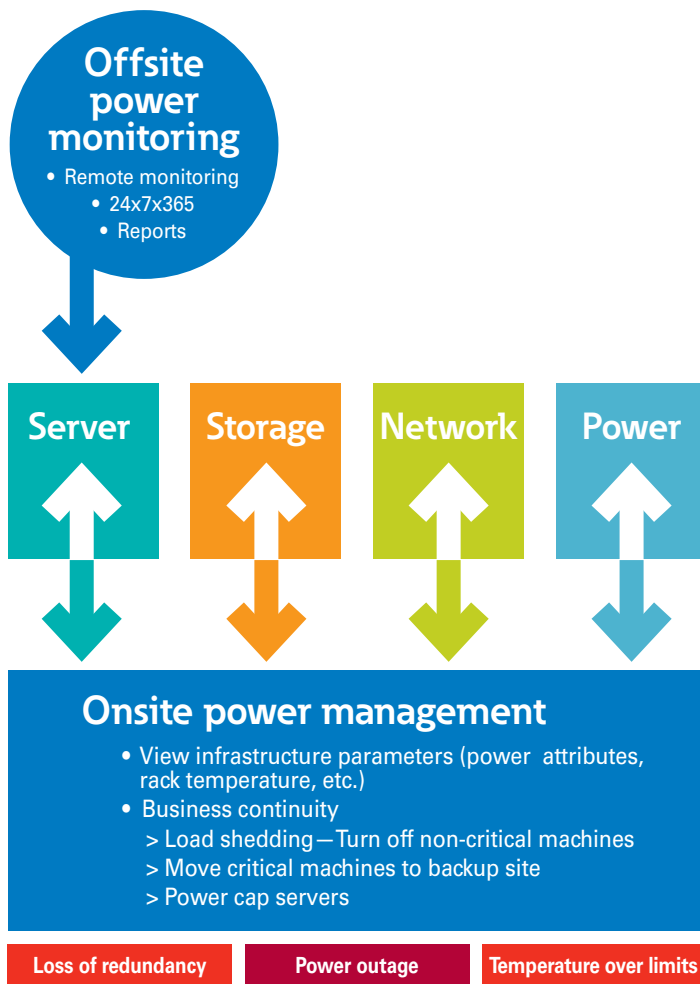

Figure 10. Onsite monitoring should be coupled with offsite management for a comprehensive power management plan.

To summarize, power management applications can alleviate companies from being burdened by the cost and damage of downtime with the following capabilities:

- Monitor and administer their UPSs from any location with Internet access
- Automatically notify key personnel of power alarms or environmental alerts
- Perform orderly, unattended shutdowns of connected equipment, or better, work with virtualization software to move virtual machines so as to maximize availability of key applications and hardware
- Selectively shut down non-critical systems to conserve runtime
- Analyze and graph trends, to predict and prevent problems before they happen
- Integrate with existing network and management systems via open standards and platforms

Power distribution units:

An essential component of any power quality infrastructure, PDUs distribute power to downstream ITE load equipment. Most companies use both floor-mount PDUs, which provide primary distribution to server racks, and rack-mount PDUs (also known as ePDUs), which distribute power to individual servers and other devices. PDUs can be equipped with optional devices like surge suppression and individual breaker (branch) monitoring systems to monitor energy use.

Services:

With proper servicing, a well-made UPS can operate safely and reliably for as long as 20 years. Without proper servicing, even the best UPS is significantly more likely to fail when you can least afford it. Companies in the market for UPS hardware, therefore, should also choose an appropriate UPS service plan from a service provider with the experience, know-how and resources to provide comprehensive, high-quality support. For more information on selecting the proper services for your system, please see the white papers "Choosing a UPS Service Plan" and "The not so shocking truth about UPS safety" available on www.eaton.com/pq/whitepapers.

Conclusion

Businesses today invest large sums of money in their IT infrastructure as well as the power required to keep it functioning. They count on this investment to keep them productive and competitive. Leaving that infrastructure defenseless against electrical dips, spikes and interruptions is a bad idea.

A well-built power protection solution, featuring high-quality, highly efficient UPS hardware, can help keep your business applications available, your power costs manageable and your data safe. By familiarizing themselves with the basics of what a UPS does and how to choose the right one for their needs, data center operators can ensure that mission-critical systems always have the clean, reliable electricity they need to drive long-term success.

About Eaton

Eaton's electrical business is a global leader with expertise in power distribution and circuit protection; backup power protection; control and automation; lighting and security; structural solutions and wiring devices; solutions for harsh and hazardous environments; and engineering services. Eaton is positioned through its global solutions to answer today's most critical electrical power management challenges.

Eaton is a power management company with 2014 sales of $22.6 billion. Eaton provides energy-efficient solutions that help our customers effectively manage electrical, hydraulic and mechanical power more efficiently, safely and sustainably. Eaton has approximately 102,000 employees and sells products to customers in more than 175 countries. For more information, visit Eaton.com.

About the authors

Chris Loeffler is a product manager in Eaton's Critical Power Solutions division. With more than 24 years of experience in the UPS industry, he has overseen product management for many Best Power, Powerware, and Eaton Power and Infrastructure solutions. He currently manages multiple power and power distribution solutions for IT, industrial, and healthcare applications. In addition, he has written a number of white papers on power systems, and application solutions for today's critical power environments. He may be reached at ChrisALoeffler@eaton.com

Ed Spears is a product marketing manager in Eaton's Critical Power Solutions division in Raleigh, North Carolina. A 35-year veteran of the power systems industry, Spears has experience in UPS systems testing, sales, applications engineering and training—as well as working in power quality engineering and marketing for telecommunications, data centers, cable television and broadband public networks. He may be reached at EdSpears@Eaton.com.

EAT•N

Powering Business Worldwide

For more information visit
Eaton.com/powerquality