White Paper

What your IT equipment needs from a UPS

The top five requirements that define "quality power" in the eyes of the power supplies in your IT systems

By Chris Loeffler
Data Center Applications Manager, Distributed Power Solutions, Eaton

Abstract

To support IT equipment as it was designed to operate, internal power supply units (PSUs) must have incoming power that meets five basic requirements, as defined by international standards set forth by the IT industry.

Designed to perform under local electric utility conditions around the world, modern PSUs are more robust than ever. They operate normally over a wide range of input voltages and frequencies. They have internal energy stores to ride through brief power interruptions. Most of them have input power factor correction circuitry and operate at a power factor close to unity. And they handle a broad range of transient power disturbances, as defined by industry standards.

In spite of their robust design, PSUs need protection from power quality problems generated by the electric utility or arising within the facility. To operate properly, IT equipment needs a consistent source of conditioned power that meets industry specifications. Providing that consistent, conditioned power is the job of an uninterruptible power system (UPS).

This white paper looks at five key attributes or requirements of power quality—as the PSU sees it—and the implications of each for UPS design and selection.

Contents

What your IT equipment needs from a UPS

The top five requirements that define "quality power" in the eyes of the power supplies in your IT systems

By Chris Loeffler
Data Center Applications Manager, Distributed Power Solutions, Eaton

Which uninterruptible power system (UPS) design is right for your application? The answer depends on a combination of factors that have been influenced by industry trends, technology advances and the degree of protection required. When selecting a UPS, it helps to look at the issue from the point of view of the "end-user"— that is, the power supply unit (PSU) inside the IT equipment. What does the power supply do, what does it need to do its job, and how does this influence UPS selection?

What does a power supply unit do?

The UPS supplies power via power distribution units (PDUs) to a PSU inside each server, storage device, network device or other IT system to be protected. So it is important to understand what these power supplies do and what they require from incoming power.

The PSU inside each piece of IT equipment converts incoming alternating current (AC) power to direct current (DC) power—typically 100–240V AC converted to 380V DC. The power supply then steps down the 380V DC to lower-voltage DC power (12 or 5V DC) to supply power to the computer's motherboard, where voltage regulators further step down the power to the voltages used by various chip sets in the IT equipment (3.3V, 1.8V, etc.). A few power supply designs include voltage regulation steps down to the chip level, eliminating the need for additional voltage regulators on the server's motherboard.

Sheathed in a sheet-metal enclosure, the PSU provides galvanic isolation between the input AC source and the DC power circuitry—a physical separation that improves the safety of the device and helps protect the IT equipment from damage or malfunction due to common-mode (ground-based) voltage or "noise."

When necessary, the PSU draws energy from its own internal energy storage device, typically a capacitor, to ride through short interruptions in input power—10–20 milliseconds (ms)—without locking up the IT equipment.

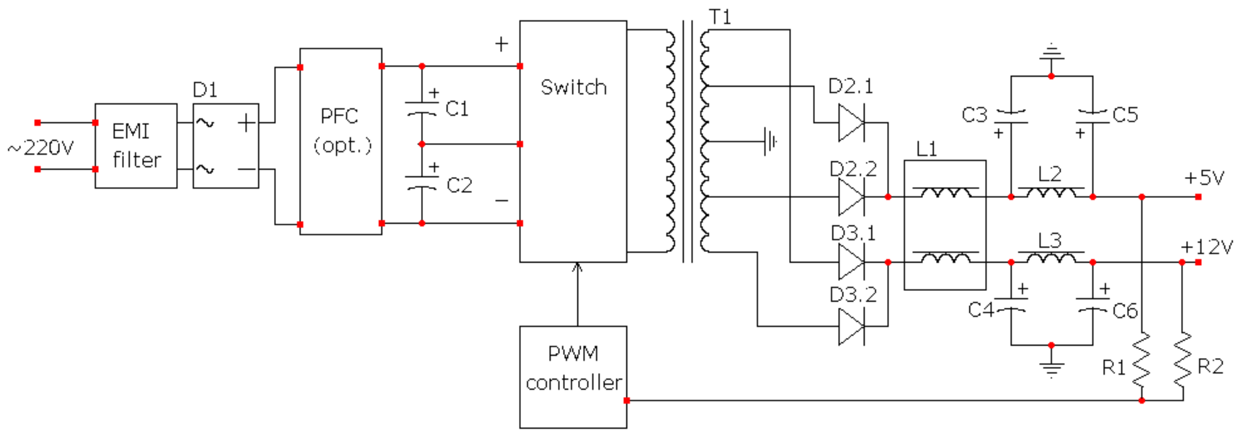

Figure 1. A power supply with energy storage capacitors (labeled C1 and C2 on this diagram) can ride through short interruptions in input power. (Source: X-bit laboratories)

A piece of IT equipment may have two, three or even six power supplies. High redundancy is common in the latest generation of blade servers and other high-density servers. With PSU redundancy, the device keeps running even if one or more power supplies fail. If those power supplies are fed by different power distribution paths, the device keeps running even if an input power source fails.

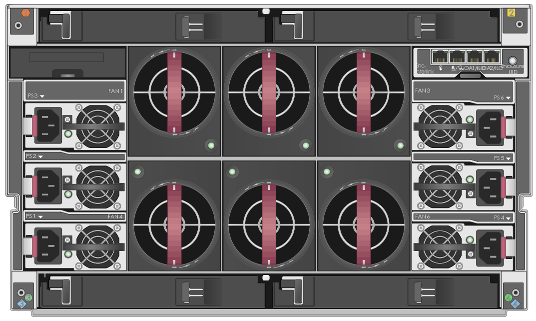

Figure 2. This HP BladeSystem C3000 server (rear view shown here) has six power supplies for redundancy.

The five essential power quality requirements for a PSU

To support IT equipment without malfunction or damage, the internal PSUs must have incoming power that meets some basic requirements, defined by international standards set forth by the IT industry:

- Input voltage within acceptable limits
- Input frequency within allowable ranges
- Sufficient input power to compensate for power factor
- Transfer to backup power faster than PSU “hold-up” time
- Protection from damaging power conditions

Let’s look at these elements that constitute “quality power” from the perspective of an IT equipment’s power supply—and the implications of each for UPS design.

Requirement One—
Input voltage within acceptable limits

To serve global markets, most equipment manufacturers use universal PSUs that can support the various input voltages and frequencies found in different countries. Just swap out the power cord and voila, you have a piece of equipment ready for sale and use in any country.

That means the PSU in your IT equipment is likely to support the low, 100V AC utility voltage used in Japan as well as the high, 240V AC utility voltage used on most other continents, including Europe, Asia and the Americas. In North America, the PSU may have to accommodate single-phase sources of 120V and 240V and three-phase sources with voltages of 120V, 208V and 240V. PSU manufacturers also ensure their designs can accommodate the voltage supplied by industry-standard power supply cords.

By standards set forth by the Server System Infrastructure (SSI) Forum (www.ssiforum.org), a PSU rated for 120–127V should operate normally at voltages ranging from 90V–140V. A PSU rated for 200V–240V should operate normally on input voltage from 180V–264V. Real design margins are even somewhat broader, again because of the need to handle input voltages from any country around the world. The power output from the PSU may even be automatically limited by input voltage, to protect the PSU and internal circuitry from damage if connected to the lower voltage range. But the bottom line is that today’s PSUs are more versatile, robust and tolerant than they were even five years ago.

Implications for UPS design

The UPS must be able to supply voltage within the specified range required by the PSU, for all voltage variations found in the AC power sources (utility mains or generator). For example, for higher watt rated power supplies requiring an input voltage of 200V–240V, the UPS must deliver power within the 180V–264V range. (*How* the UPS achieves that voltage regulation varies from one UPS topology to another. For more information, see Part 2 in this white paper series, “Which UPS Topology is Right for the Job?”)

Requirement Two—
Input frequency within allowable ranges

Once again, power supplies for IT equipment are typically designed for universal operations. That means a typical PSU can operate normally at frequencies from 47 Hz–63 Hz, (many supplies have a 45 Hz–65 Hz window) to accept utility power at both 50 Hz (the typical standard found in Europe, Australia and Asia) and 60 Hz (the standard in North America).

Implications for UPS design

The UPS must be able to regulate output frequency to meet the PSU’s specification range of 47 Hz–63 Hz for all frequency variations in the AC power source—whether that power is coming from utility mains or a generator.

Requirement Three—
Sufficient input power to compensate for power factor

Power factor (PF) is the ratio of real power to apparent power, where real power is the capacity of the circuit to perform work, and apparent power is the product of the current and voltage of the circuit, which also includes current affected by reactive components. Power factor is expressed as a number between zero and one that represents the ratio between the real power in watts and the apparent power in VA. A power factor of 1.0 indicates that the voltage and current peak together (the voltage and current sine waves are always the same polarity), which means that the VA and watt values are the same.

The power factor of a circuit is influenced by the type of equipment being powered:

- Circuits containing only resistive elements (such as filament lamps and cooking stoves) have a power factor of 1.0, or unity (Figure 3). This is the highest power factor that can be attained. Every incoming volt-ampere is available as useful wattage.

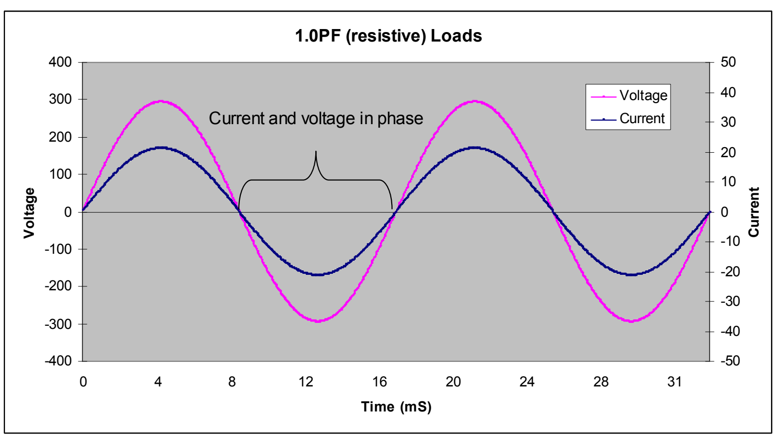

Figure 3. Power factor of 1.0, unity (60 Hz waveforms)

- Circuits containing inductive elements (motors, transformers) have what is known as a “lagging” power factor (Figure 4). Because the sine waves are out of sync, there is less real power.

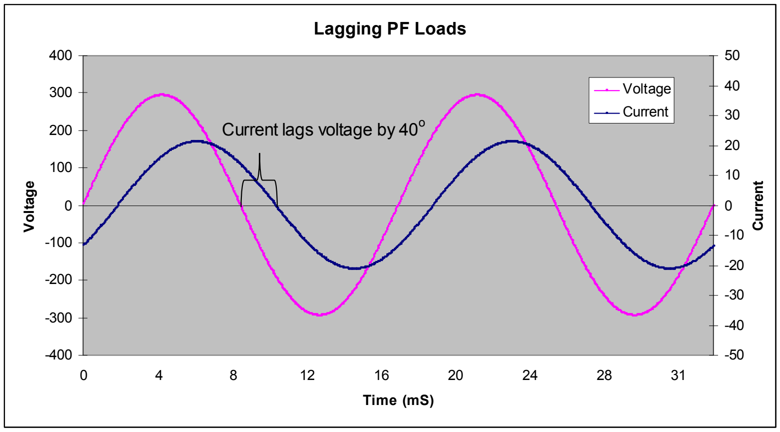

Figure 4. Lagging power factor

- Circuits containing many capacitive elements have a “leading” power factor (input current leads voltage).
- Circuits containing a mixture of reactive components (capacitors, inductors) typically have what is known as a distortion power factor. The current drawn on the input is a mixture of the fundamental frequency, as well as several harmonic frequencies (multiples of the base).

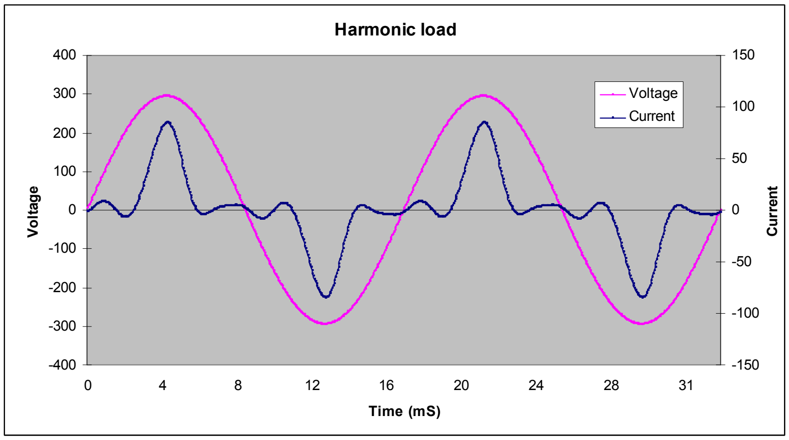

Figure 5. Harmonic load with high input THD (total harmonic distortion)

Power supplies used in IT equipment generally fit into this last category. However, the current drawn from the source is much less distorted than that shown in Figure 5. Harmonic distortion and power factor are directly related. The higher the power factor of a device, the lower the harmonic distortion (Figure 6).

PSUs used in IT equipment today have a power factor trending toward unity, because of the need to reduce harmonic current content in the AC source feeding the IT loads. As a result, in today’s power supply designs, a power factor of 0.9 would be considered acceptable, 0.95 would be typical, and a value of 0.99 would be excellent.

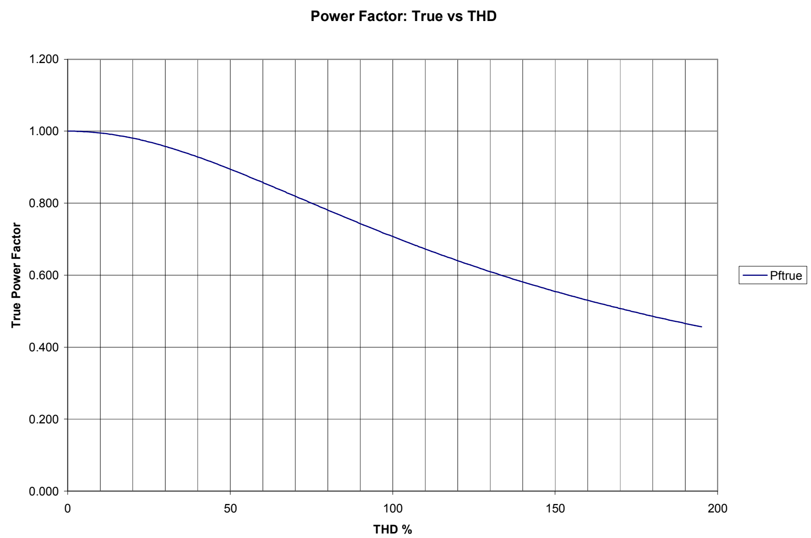

Figure 6. The higher the power factor of a device, the lower the harmonic distortion (distortion power factor loads only).

The performance implications of power factor and harmonic distortion

Because power factor in the typical data center is still less than unity, you need to supply slightly more apparent power to get the real power you need. To get a meaningful measure of the incoming VA needed to serve IT loads, take the watt rating of each piece of equipment and divide by the power factor found on the device's rating label. In a typical data center, you may find that you need three to six percent more apparent power to achieve the needed real power.

Poor power factor, due to high input total harmonic distortion (iTHD) (Figure 5), had been known to cause failed neutral conductors, overheated transformers, and in the worst cases, building fires. This was a concern with older switch-mode power supply designs used 10 years or more ago. These problems led to the creation of international design standards (EN61000-3-3, EN61000-3-2, IEC 1000-4-7), aimed at limiting allowable harmonic distortion on a power source.

Most modern power supplies above a 50-watt rating are designed to correct for poor power factor. They have an input power factor correction (PFC) circuit to raise power factor and lower input current distortion. In addition, most power supplies above 200 watts have an *active* PFC circuit that automatically adjusts the power factor based on the actual power required by the IT device. These capabilities enable such power supplies to drive power factor very close to unity.

Older designs with high iTHD have been outlawed in the European Union. Since manufacturers only want to build one design for use in any country, power factor correcting designs are widely used in North America as well.

With the proliferation of high-powered servers using redundant power supplies, each power supply might be running at less than 50 percent of capacity, so reactive components may have a more profound effect on overall power factor of the IT load. In some cases, this has even caused the current signature to go leading, or capacitive in nature, which can create issues with the UPS or generator supplying these PSUs if an entire data center of these loads exhibits this same trait.

Implications for UPS design

Some manufacturers' UPS designs can provide full-rated power into a leading power factor load. Other manufacturers recommend derating the power output of the UPS, or the UPS will experience excessive heating in the inverter and output filter causing unstable operation or shutdown of the system due to over-temperature.

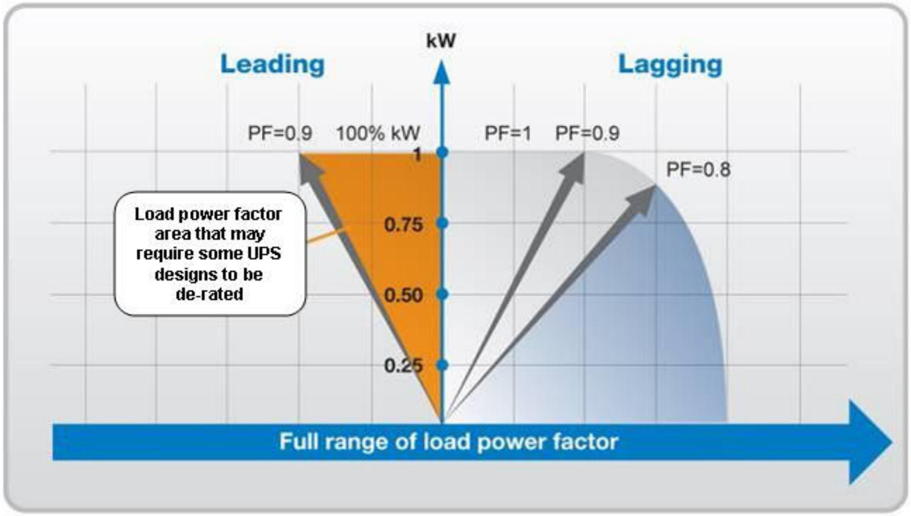

Figure 7. Power factor ranges that PSUs reflect back to the UPS

In sizing a UPS, the power rating (kW) is actually more important than the kVA rating (apparent power) due to the high power factor of the IT equipment. When assessing output power and battery backup time, make sure to use the real power (kW) rating of the UPS. If the kW is not apparent in the UPS specifications, it can be calculated by multiplying the UPS kVA rating times the output power factor rating of the UPS.

Requirement Four—
Transfer to backup power faster than PSU "hold-up" time

Power supply units inside IT equipment have an energy storage device, typically a capacitor, that stores enough energy to keep the device running during very brief power interruptions (milliseconds). This is known as "hold-up" time and depends on the internal capacitance of the power supply and the output power rating. At higher output power, the energy is drawn from internal capacitance faster than at lower output power.

According to IT equipment standards set forth by the SSI Forum, minimum hold-up time at fully rated output power is one cycle. At 50 Hz (most of Europe and Asia), this translates into 20 ms. At 60 Hz (the standard in North America), one cycle would be 16.7 ms. Since most IT equipment is designed for the global market, the minimum hold-up time is 20 ms and may be longer at lighter loads.

However, the trend is toward shorter hold-up times. Under pressures to reduce PSU size and cost, manufacturers are designing PSUs with smaller capacitors, which leads to shorter hold-up times. This effect is somewhat offset by the prevalence of redundant power supplies, since each power supply would be loaded to less than 50 percent of its capacity.

A related issue with respect to hold-up time is the peak inrush current (Table 1) required to charge up the capacitor that provides the ride-through capability. When first connected to an AC power source (or when powered up on an already connected source), the equipment temporarily draws a large inrush current that can last for 2–10 ms and be as much as 10–60 times the normal operating current.

Table 1: Example IT equipment nominal and peak inrush current

Equipment	Maximum Nominal Current	Peak Inrush Current
HP Proliant DL 360 G4 – IU server	2.4A	61A for 3 ms
HP Proliant e-class blade server	1.6A	100A for 2 ms
IBM BladeCenter, fully loaded	23.7A	200A for 4 ms
IBM x-series 260	4.9A	120A for 4 ms
Cisco 3825 Router	2A	50A for 10 ms

Similar to the start-up inrush current, there is also a surge current drawn to recharge the capacitors after short interruptions in power (Figure 8). If the power interruption was less than five ms, surge currents will typically last for half a cycle (10 ms) and will be less than 300 percent of nominal current. For interruptions of 10-15 ms, the surge current could be 700 to more than 1000 percent of nominal current, and can last for 1.0–1.5 cycles (20 ms–30 ms).

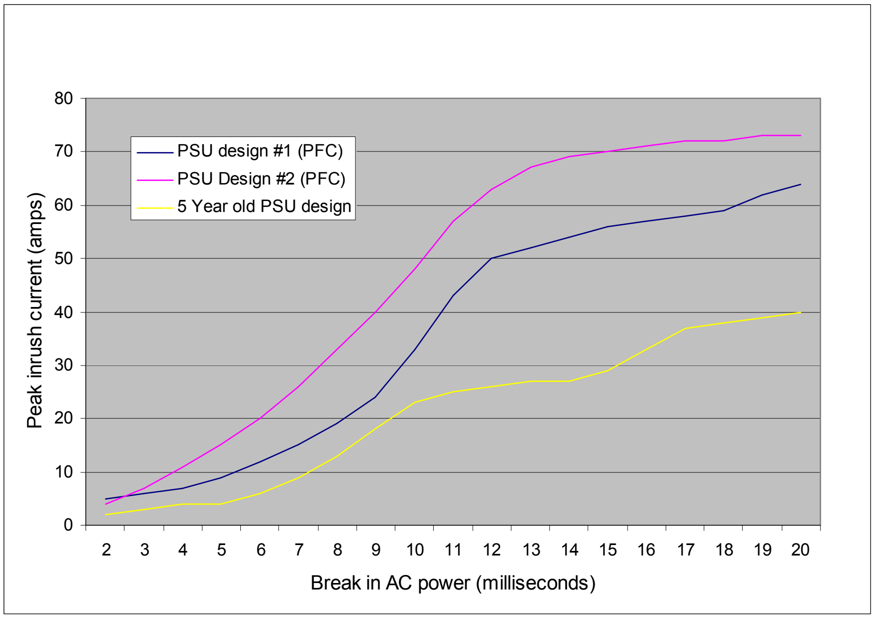

Figure 8. Inrush current of several different PSU designs

Implications for UPS design

UPS performance requirement: The UPS must ensure no interruption in its output that lasts longer than the hold-time of the IT equipment's PSU. This means that the UPS must have an acceptable transfer time for all transitions between different modes of operation—such as from normal operating mode to battery mode and back again, or between high-efficiency mode and double-conversion mode for new energy-saving UPSs.

Note that hold-up time will be different for single- or multi-corded servers, because the more PSUs on the IT equipment, the less power load on each PSU, and the longer the available hold-up time. Looking at it another way, single-corded IT equipment will need a UPS with faster transition times, to prevent unplanned shutdowns and reboots.

Transfer time should actually be much faster than the maximum allowable hold-up time, because the longer the PSU goes without power, the larger the surge current it will draw when it receives power again. In cases where the PSU is without power for more than 5–10 ms, the inrush current required by the PSU could easily exceed the maximum current output capacity of the UPS inverter, forcing the UPS to shut down to protect its own inverter components.

As a result, Eaton recommends using a UPS with a transfer time of less than 5 ms, if the UPS will be loaded to 70 percent or more of its capacity. Rapid transfer time helps prevent a UPS overload condition that could cause a system failure. To prevent problems with large inrush current during initial start-up or restart, most UPSs will feed the IT equipment directly from the utility source at that time.

Considerations for branch circuit breakers. The AC source supplying the PSU must also be able to support inrush current without tripping any breakers. This factor becomes especially critical if several pieces of IT equipment are served from the same circuit. All of this equipment draws inrush current simultaneously after a power outage. It is important to either size the breakers accordingly (using breaker types that have a trip curve compatible with the expected load), or sequence the start-up to reduce system-level inrush current.

Requirement Five—
Protection from damaging power conditions

PSUs are designed to handle voltage that sags 10 percent below nominal specification or surges 10 percent above, without loss of function or performance. If the nominal voltage range is 200V–240V, the PSU will operate normally when input voltage is as low as 180V or as high as 264V. The PSU is also required to handle surges of 30 percent from the midpoint of nominal (286V for a 220V PSU) for 0.5 cycles (8–10 ms).

For fast AC line transients, the power supply is designed to meet the EN61000-4-5 directive and any additional requirements in IEC1000-4-5:1995 and the Level 3 requirements for surge—and withstand capability without disruptions to normal operations.

These tolerances are well defined by the Information Technology Industry Council (ITIC) curve (Figure 9), published by ITIC Technical Committee 3 (TC3). The ITIC curve, which actually represents a stair step more than a curve, describes a voltage envelope that PSUs can typically tolerate without interrupting function.

The Y-axis shows the voltage range; the X-axis shows the duration of exposure. Voltage conditions within the upper and lower boundaries (the middle zone of this “curve”) are safe. Below this zone is a low-voltage area where the PSU would not be expected to operate normally, but it wouldn’t be harmed either. Above is the “prohibited region,” where voltage conditions could damage the equipment.

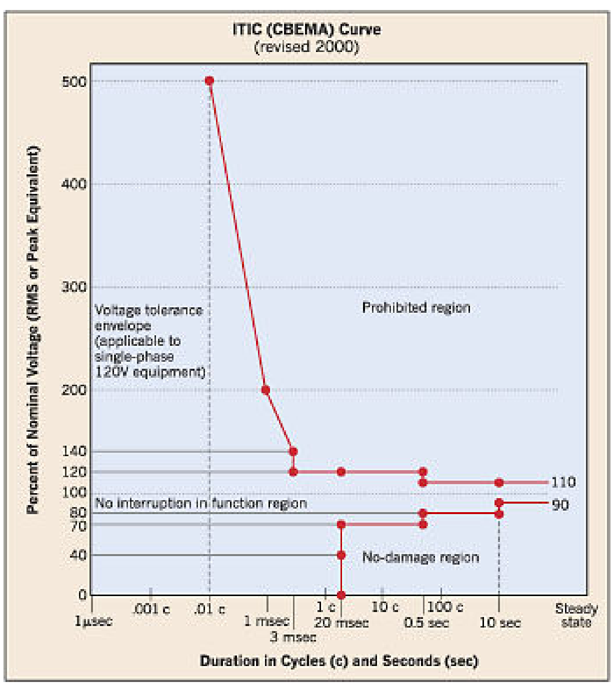

Figure 9. The ITIC (CBEMA) curve defines the range of power conditions that IT equipment can tolerate.

Implications for UPS design

UPS output voltage must be within the acceptable zone specified by the ITIC curve for all input AC line conditions. The UPS must especially be designed to ensure the voltage to the PSU is not in the prohibited range, since voltage in that range could damage the IT equipment. Also the UPS must be designed to handle high-speed impulses, such as lightning strikes or surge currents of longer duration, even though most PSU are designed to handle some level of surge current without damage.

Closing thoughts

Power supply units are more robust than ever. They operate normally over a wide range of input voltages and input frequencies. They have internal energy stores to ride through brief power interruptions. Most have power factor correction circuitry to ensure that power drawn from the AC source is sinusoidal (has a desired waveform) and a power factor close to unity. And they handle a broad range of transient power disturbances, as defined by the ITIC curve.

In spite of their robust design, power supplies need UPSs to protect them from utility conditions they were not designed to handle. The UPS provides a consistent source of conditioned power, backed up by battery or generator sources—and sufficient power to address power factor issues and inrush current.

Eaton manufactures, sells and supports UPSs that meet or exceed the requirements of IT power supply units for reliable, high quality power. Eaton offers modular, rack-mounted UPSs for distributed or zoned power protection, and larger UPSs for centralized protection—including the industry's most energy-efficient UPS, delivering up to 99 percent efficiency under normal operation.

For more than 40 years—from the first commercial UPS to today's high-efficiency, modular, scalable designs— Eaton UPSs have set the standard for power protection and backup power.

To find out more, visit www.eaton.com or call us at 800-356-5794.

About Eaton

Eaton Corporation is a diversified power management company with 2008 sales of $15 billion. Eaton is a global technology leader in electrical systems for power quality, distribution and control; hydraulics components, systems and services for industrial and mobile equipment; aerospace fuel, hydraulics and pneumatic systems for commercial and military use; and truck and automotive drivetrain and powertrain systems for performance, fuel economy and safety.

Eaton is your source for a comprehensive range of data center solutions, including power distribution, power protection, rack enclosures and accessories for network closets, computer rooms and data centers.

Eaton has approximately 75,000 employees and sells products to customers in more than 150 countries. For more information, visit www.eaton.com.

About the author

Chris Loeffler, *Data Center Applications Manager, Distributed Power Solutions, Eaton*

Chris Loeffler is the Global Applications Manager for Eaton Corporation, specializing in data center power solutions and services. With more than 19 years of experience in the UPS industry, he has overseen product management of more than 20 UPS products for data center and industrial applications.

Mr. Loeffler has held a variety of positions with Eaton, including roles in service engineering, application engineering, and more than 10 years within product management. Mr. Loeffler has authored a number of articles for trade publications and written several white papers on energy efficiency in the data center. He has also written articles on various UPS topologies for data center and industrial applications.

Tutorials on demand
Download Eaton white papers to learn more about technology topics or explain them to customers and contacts. Maintenance bypass, paralleling, UPS topologies, energy management and more are demystified in free white papers from our online library.
http://www.eaton.com/pq/whitepapers