

Which UPS is right for the job?

*Considerations in choosing standby, line-interactive, double-conversion designs—
and new high-efficiency, multi-mode capabilities—for your data center*

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

Abstract

Traditionally, data center managers and facilities managers could choose from three UPS topologies: standby, line-interactive and double-conversion—offering widely varying levels of efficiency, performance and protection.

The latest generation of double-conversion UPSs offers unique multi-mode capabilities. The UPS operates in a very high-efficiency mode unless power conditions warrant a switch to the higher protective level typical of double-conversion mode.

This white paper describes how various UPS topologies work and looks at the impact of operating mode on five key factors of UPS performance:

- Maintaining voltage within tolerances
- Transferring among modes without locking up IT equipment
- Transitioning gracefully to and from generator power
- Reliability and availability
- Energy efficiency

Contents

| | |
|---|-----------|
| A closer look at the different UPS topologies | 3 |
| How does UPS design and operating mode affect performance? | 5 |
| Maintaining voltage within tolerances | 6 |
| Transferring among modes without locking up IT equipment | 8 |
| Transitioning gracefully to and from generator power | 10 |
| How does UPS design affect reliability?..... | 11 |
| How does UPS topology affect energy efficiency? | 12 |
| Closing thoughts..... | 14 |
| About Eaton | 15 |
| About the author | 15 |

Which UPS is right for the job?

Considerations in choosing standby, line-interactive, double-conversion designs—and new high-efficiency, multi-mode capabilities—for your data center

By Chris Loeffler

Data Center Applications Manager, Distributed Power Solutions, Eaton

Which uninterruptible power system (UPS) design is right for your data center? The answer depends on a combination of factors that have been influenced by industry trends and technology advances—and complicated by marketing hype.

At its most basic level, a UPS performs two primary and complementary functions:

- *Conditioning incoming power* to smooth out the sags and spikes that are all too common on the public utility grid and other primary sources of power
- *Providing ride-through power* to cover for sags or short-term outages (say, five minutes to an hour), by dynamically selecting and drawing power from the utility grid, batteries, backup generators and other available sources

UPSs being marketed for data center applications provide these functions to meet the requirements of IT power supply units as set forth in industry standards and specifications. However, there are differences in the degree of protection they provide and the way they provide it. Three different UPS topologies are commonly sold to protect IT equipment:

- **Standby UPSs** allow the equipment to run off utility power until the UPS detects a problem, at which point the UPS switches to battery power to protect against sags, surges or outages.
- **Line-interactive UPSs** regulate voltage by boosting input utility voltage up or moderating (bucking) it down as necessary before allowing it to pass to the protected equipment—or resort to battery power.
- **Double-conversion UPSs** isolate equipment from raw utility power—converting power from AC to DC and back to AC again, to deliver the cleanest power and highest protection.

Among the latest generation of UPSs are models that offer multi-mode operation, often called “eco-mode” or “high-efficiency mode.” The UPS operates in high-efficiency mode unless power conditions warrant a switch to the higher protective level of double-conversion mode.

Other topologies, such as ferroresonant and rotary designs, are available for specific purposes, such as industrial or harsh environments. This paper addresses the typical static UPS topologies used in most IT environments.

Which design is right for your data center applications? The way a manufacturer designs a product to these topologies can have a significant effect on overall power performance, data center availability and total cost of ownership. This white paper provides an objective view of key issues that go into selecting the best UPS internal design for your requirements.

A closer look at the different UPS topologies

UPS topologies can be classified into two basic design categories: double-conversion and single-conversion.

- **Double-conversion UPS systems** process power twice. An input rectifier processes AC to DC, which then feeds into an output inverter that processes the DC back to AC power to send to the IT equipment. In normal operation, the UPS is always double-processing power through this combination of rectifier and inverter.
- If the AC input supply is out of predefined limits, the rectifier turns off, and the UPS draws current from the battery. Battery power passes through the output inverter and on to the IT equipment. The UPS stays on battery power until the AC input returns to normal tolerances (or until the battery runs out, whichever is sooner).

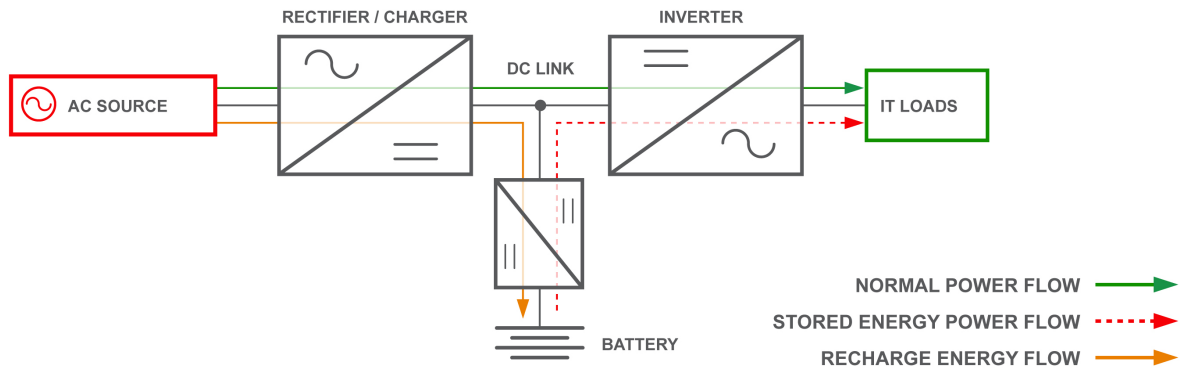


Figure 1. Internal design of a typical double-conversion UPS

- **Single-conversion UPS systems** use the incoming utility AC power to supply the required AC to the output loads in normal operation. Some systems use inductors or transformers to regulate output voltage. These designs also include some type of battery charging circuit to ensure the battery is fully charged.

If the AC input supply is out of predefined limits, the UPS turns on its inverter (in some designs, the inverter is always on, but without load). The inverter draws current from the battery, then disconnects the AC input supply to prevent backfeed from the inverter into the utility. The UPS stays on battery power until the AC input returns to normal tolerances (or until the battery runs out, whichever is sooner).

Two of the most popular single-conversion designs are the *line-interactive* and *standby* systems, which, in the past, were typically used in lower kVA applications. Line-interactive systems usually have a wider input voltage range than standby systems and can regulate voltage to within acceptable limits through the power interface without using the battery. Standby power systems simply pass incoming AC power to the connected equipment and switch to battery power when necessary. Some standby UPS designs incorporate transformers or other devices to provide limited power conditioning as well.

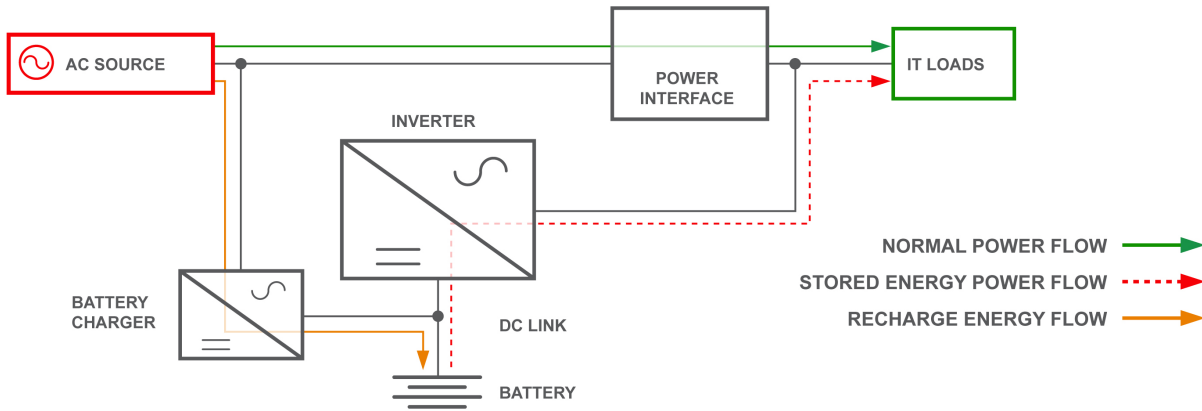


Figure 2. Internal design of a single-conversion design, line-interactive UPS

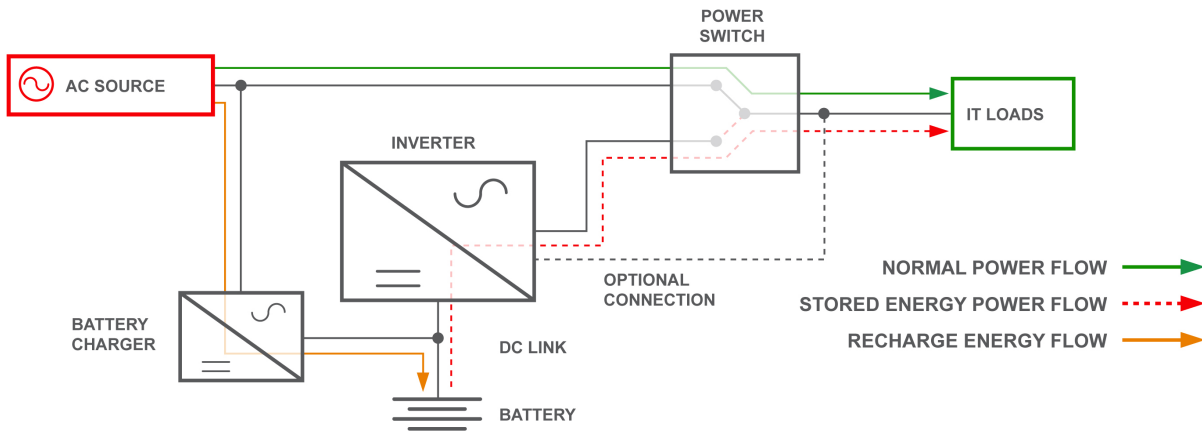


Figure 3. Internal design of a single-conversion design, standby UPS

New multi-mode UPSs combine the best of single- and double-conversion technologies to provide the benefits of each. Whenever the UPS's high-speed line-detection circuitry senses a change in condition, the system automatically changes modes accordingly:

- **Normal conditions:** When power conditions are within acceptable limits, the multi-mode UPS operates as a high-efficiency, energy-saving system—regulating voltage within safe tolerances and resolving common anomalies found in utility power.
- **Erratic power or fleeting disturbances:** If AC input power falls outside of preset tolerances for line-interactive mode, the system switches to double-conversion mode. The UPS processes incoming power through a rectifier and inverter, completely isolating IT equipment from the incoming AC source.
- **Power outage or sustained power anomalies:** If the AC input power is outside the tolerances of the double-conversion rectifier (or the power goes out altogether), the UPS uses the battery to supply energy to keep the output loads operating. When the generator comes online and supplies backup AC power, the UPS uses double-conversion mode until the generator has stabilized sufficiently to safely switch back into high-efficiency mode.

This new multi-mode technology provides exactly the level of power protection needed under the conditions of the moment. The unit does not operate in the low-efficiency/highest-protection mode unless necessary. With energy costs representing the single largest component of IT operations, the efficiency gains of this strategy are substantial. Even a small data center can save tens of thousands of dollars a year in utility bills. Larger data centers stand to save millions over time, without compromising data center performance or reliability.

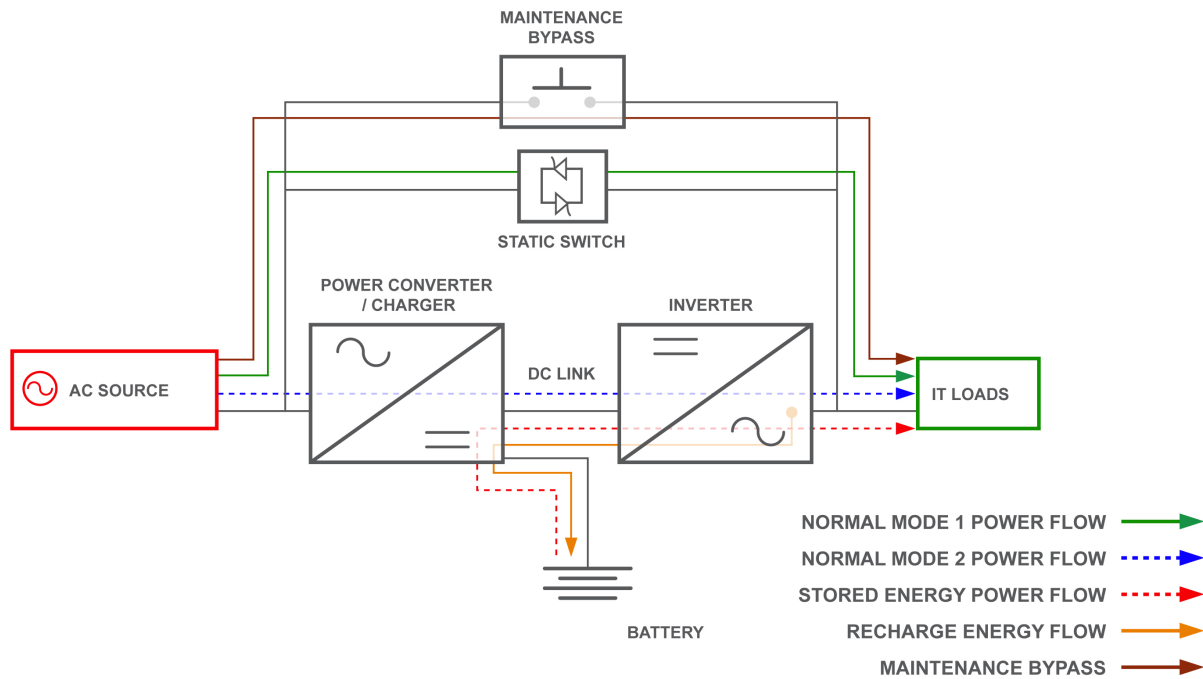


Figure 4. Internal design of an Eaton® multi-mode, high-efficiency UPS

How does UPS design and operating mode affect performance?

The key objective of the UPS is to ensure that the power supplied to the IT equipment remains within the specifications of the equipment’s power supply unit under all input AC conditions, including generator operation. Let’s take a look at how different designs fare on the key criteria:

- Maintaining voltage within tolerances
- Transferring among modes without locking up IT equipment
- Transitioning gracefully to and from generator power

The effect of UPS topology on performance:

Maintaining voltage within tolerances

UPS output voltage must be within the acceptable zone specified by the Information Technology Industry Council (ITIC) curve for all input AC line conditions.

Input voltage to the power supply unit (PSU) is shown on the vertical axis. The horizontal axis represents the duration of input voltage events (in AC cycles up to 10,000 cycles, about 28 minutes). The ITIC curve—more of a stair-step than a curve, actually—shows the acceptable input voltage envelope for a typical PSU design for IT equipment.

The UPS must ensure that voltage to the PSU is not in the prohibited range above the acceptable zone, since voltage in that range could damage the IT equipment. Voltages below the threshold could cause the PSU to shut down or exhibit erratic behavior.




Figure 5. (Right) The ITIC (CBEMA) curve shows power tolerances for IT equipment.

Almost all system designs provide some surge suppression to protect against high-frequency transients and intense spikes, such as from lightning or from damage to the facility's power plant.

- **Most small standby and line-interactive systems** use some form of transient clamping device, such as metal oxide varistors (MOVs) that shunt excess energy to ground or—if the energy level is too high—self destruct to absorb the overvoltage or transient hit. Since most UPSs of this type are small, designed to be located close to the equipment being protected, they may only have a minimal number of these clamping devices.
- **Double-conversion UPSs** operating in normal mode process power through an AC-to-DC-to-AC conversion process to prevent damaging input conditions from passing through the UPS to the connected load equipment. (However, if the UPS is on bypass mode, such as during system maintenance or a system fault, damaging input impulses would pass through the UPS bypass to the loads.)
- **Multi-mode, double-conversion UPSs** tend to be deployed closer to the utility input source, so they are generally designed with extra surge protection. These designs may include multiple parallel MOVs connected for three separate protection paths: line-to-line, line-to-ground and neutral-to-ground. The UPS may also have gas tube arrestors, surge coils or other filter circuits containing devices such as inductors and capacitors to eliminate damaging impulses before they can reach the critical load. In addition, these UPSs automatically switch from high-efficiency to double-conversion mode if incoming power conditions warrant, therefore isolating input transients from the loads. Most designs also ensure that even in bypass mode, the connected load equipment is protected from transient events. One way or the other, IT equipment is protected from intense surges and strikes.

Whichever UPS design is used, it is still recommended to have surge protection at the utility entrance, to protect the UPS input monitoring circuitry and to provide surge protection on the electrical circuit that feeds the UPS bypass.

The various UPS designs differ in how they handle less extreme voltage conditions, such as under- or over-voltage conditions:

- A **standby UPS** supplies the IT equipment with acceptable power to meet this specification as long as the input voltage is within predefined UPS tolerances. However, this band of normal operations is typically narrow (± 10 percent of the ITIC curve), so the UPS must resort to batteries frequently, which can reduce battery runtime and service life. Some standby systems allow a wider input voltage window, which helps conserve the battery but could cause lock-up or sporadic operational issues with the connected IT equipment.
- A **line-interactive UPS** supplies power within ITIC specifications as long as its input voltage is within preset UPS tolerances. However, the line-interactive system can provide some voltage regulation, using a tap changing transformer or a buck/boost circuit. That means it doesn't have to resort to batteries as often as a standby system, although it may use some battery power to support the transition between normal mode and voltage regulation mode. Battery usage is lower than a standby UPS but still higher than a double-conversion topology.
- A **double-conversion UPS** typically provides regulated output voltage, within 1–3 percent of nominal, under all input power conditions. When input voltage is within preset UPS tolerances, the output is regulated without going to battery. As such, the double-conversion UPS uses batteries less often, and for less time, than either standby or line-interactive designs. That can translate into longer battery runtime and service life. Many of today's double-conversion UPSs are intelligent enough that they allow an even wider input acceptance window if the UPS is less than 100 percent loaded.
- A **high-efficiency, double-conversion UPS with multi-mode operation** supplies power within ITIC specifications when input voltage is within preset UPS tolerances. When the input AC voltage is outside this range, the UPS automatically uses double-conversion mode to keep the output regulated to the ITIC specification. As a result, battery usage duration and frequency is similar to the double-conversion UPS and in some instances may even be lower.

Some larger system designs may allow adjustment of the output voltage window, so the system can also support non-IT power supplies that have a more limited input voltage range, while still gaining the benefits of higher operating efficiency.

As the figures show, all the UPS designs meet the input voltage requirements for IT equipment, as defined by the ITIC. The key difference is how the UPS achieves that result, which has a great influence on the frequency and duration of demands on the battery. Battery demands are highest for standby UPSs and lowest for double-conversion topologies.

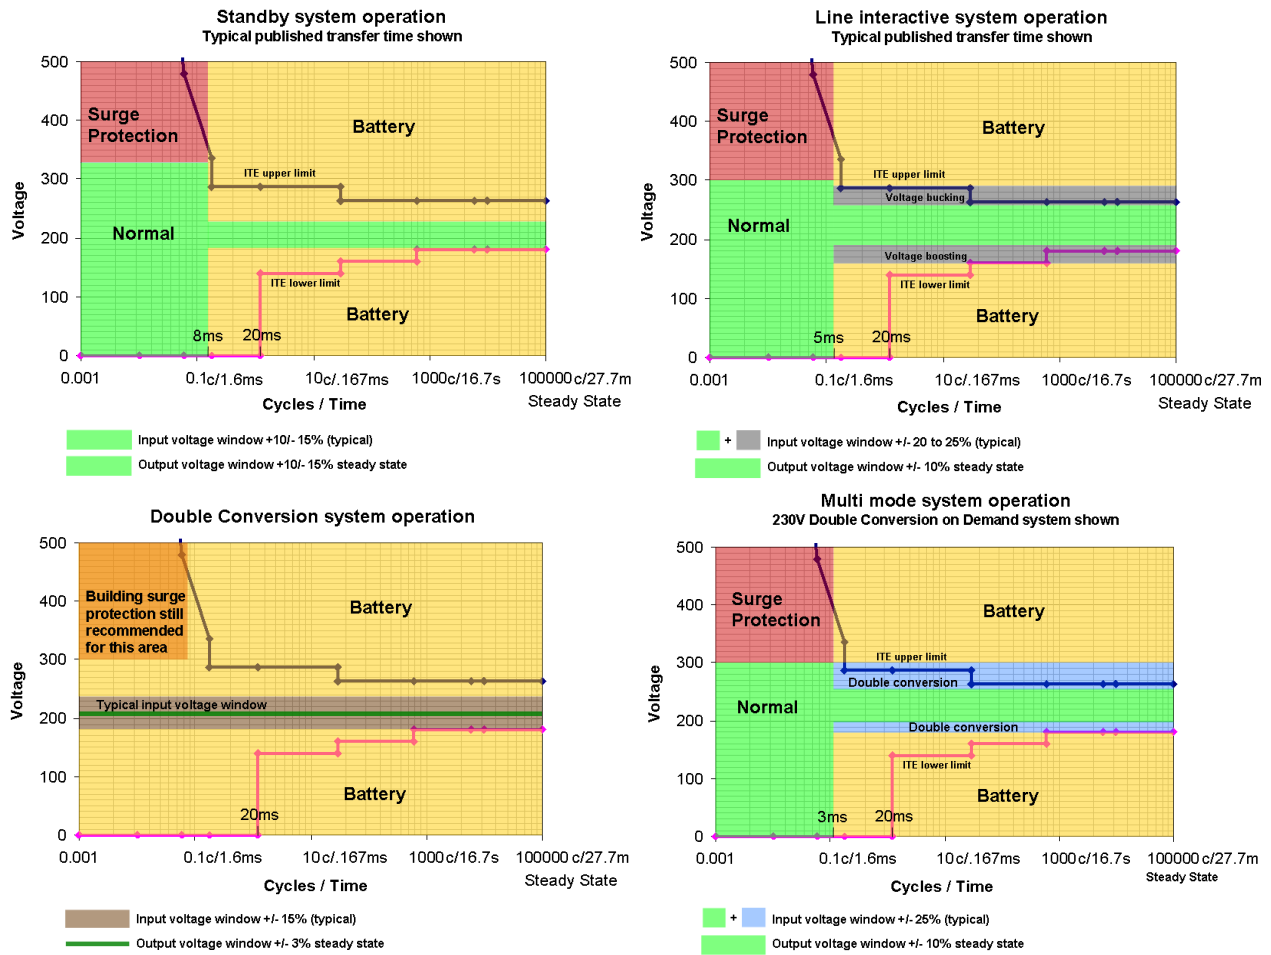


Figure 6. Performance of various UPS designs related to the ITIC curve

The effect of UPS topology on performance:

Transferring among modes without locking up IT equipment

By industry standards, power supply units inside IT equipment are designed to store enough energy to keep the device running for about 20 milliseconds (ms) of power interruption. This is known as “hold-up” time. That means the device can withstand brief interruptions in power while a UPS transitions between modes of operation, such as from normal operating mode to battery and back again.

However, transfer time should actually be much faster than 20 ms, because the longer the PSU goes without power, the larger the in-rush current it will draw when it receives power again. In-rush could exceed the current handling capacity of the UPS and cause it to shut down.

- **Standby UPSs** switch to battery mode in 5–12 ms (8 ms typical). Standby systems typically use a fast acting mechanical relay for the power switch, which extends the length of time before the transfer to battery can be made.

Most power supplies can tolerate this interruption. However, when transfer time is greater than 5 ms, the in-rush current may exceed the capability of the UPS inverter and cause a reset of the IT equipment, resulting in data corruption or shut-down. If the standby system allows output voltage to dip more than 10 percent below nominal current (say, below 108V on a 120V system), the PSU is likely to be in a state where it is pulling higher than normal current. For this reason, a prolonged loss of output increases the odds that the PSU will shut down.

Another concern about using standby systems for highly critical servers is the issue of output voltage waveform while on battery power. Many standby systems create a square wave or modified sine wave output, which today's power factor corrected power supplies may not be able to handle. If this is the case the power supply will almost always shut down once battery operation commences.

- **Line-interactive UPSs** switch to battery mode with a typical transfer time of 3–8 ms (5 ms typical), which is within acceptable limits for most power supplies. Some PSUs could exhibit in-rush currents exceeding 400 percent if the transfer time is longer than 5 ms; the UPS inverter could have problems supporting this high current requirement.
- **Double-conversion UPSs** begin drawing current from the battery with zero interruption (transfer time) in output power, therefore there is no risk of the transfer causing any in-rush.
- Advanced, **high-efficiency, double-conversion UPSs in multi-mode operation** typically switch to battery mode in 1–3 ms, well within the lowest portion of the in-rush curve of a typical PSU. The subsequent in-rush current is less than 200 percent of normal peak currents—easily handled by the battery and inverter for short periods.

Eaton multi-mode systems work very differently from a typical double-conversion UPS that advertises a “high-efficiency operation mode” or “eco-mode” option, in two important ways. Converted double-conversion UPSs typically:

- Operate in standby mode (rather than line-interactive) during the high-efficiency mode, so they offer less protection.
- Require 5–12 ms to transfer to double-conversion mode, due to transformers in the UPS design or delay in the detection circuitry's ability to identify a power problem. That transfer time could cause data corruption or shut down the IT equipment.

In standby mode, the UPS might not be able to synchronize the inverter immediately upon a loss of AC input, thereby delaying the transfer to battery power. Other issues arise if the inverter and rectifier are isolated from the incoming AC, which doesn't ensure proper transient protection of the critical load.

An effective multi-mode system must always track and synchronize the inverter to the AC input, so when AC is lost, the inverter can immediately take over the load with the slightest break in output power. In addition, the rectifier and inverter should always be online, ready to protect against transients and provide microsecond response in case of an AC outage.

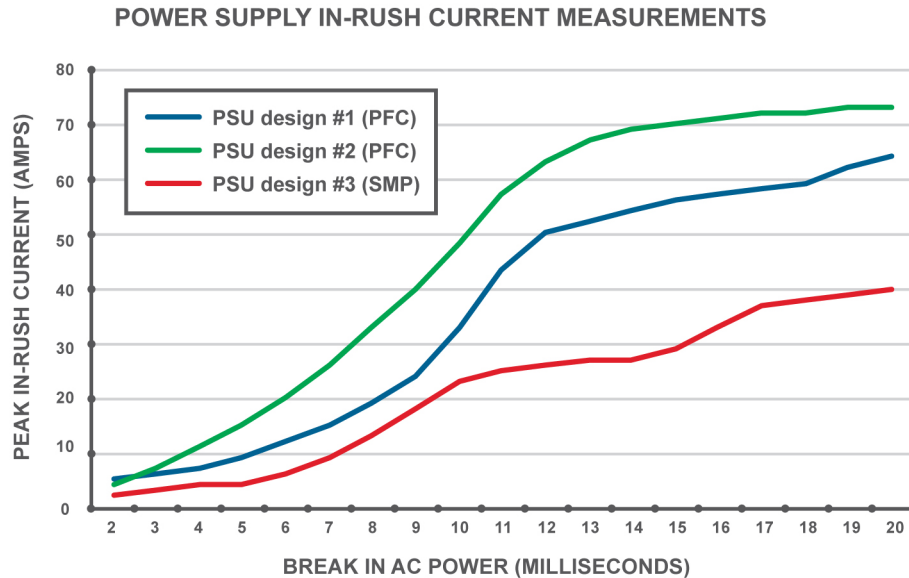


Figure 7. The longer the period of power loss, the higher the PSU in-rush current, but some PSU designs control this phenomenon better than others.

The effect of UPS topology on performance:

Transitioning gracefully to and from generator power

How does the UPS handle the transfer to backup generator during extended utility power outages? The transition may not be smooth, since the generator may have unstable voltage and frequency during start-up and warm-up periods.

The UPS must be able to handle generator output aberrations as the generator and its loads go through initial ramp-up to full operation. If this is not conditioned by the UPS, the unstable power can result in data corruption or shut-down of connected IT equipment. The UPS should make as few as possible transfers to and from battery operation, so there is less chance of a disruption in output power and less stress on the batteries.

Standby and line-interactive UPSs must first quantify the source and then synchronize the inverter to this source before switching the load over to the generator. These designs are also subject to switches back to battery operation if generator frequency or voltage deviates even slightly.

Double-conversion and high-efficiency, double-conversion UPSs with multi-mode operation ensure that even if the generator experiences unstable output voltage or frequency as the generator warms up (or due to other loads cycling on and off the generator), the UPS continues to operate on the rectifier, rather than switching to battery operation. Since the input rectifier is used to process AC to DC, these UPSs have the shortest period on battery power.

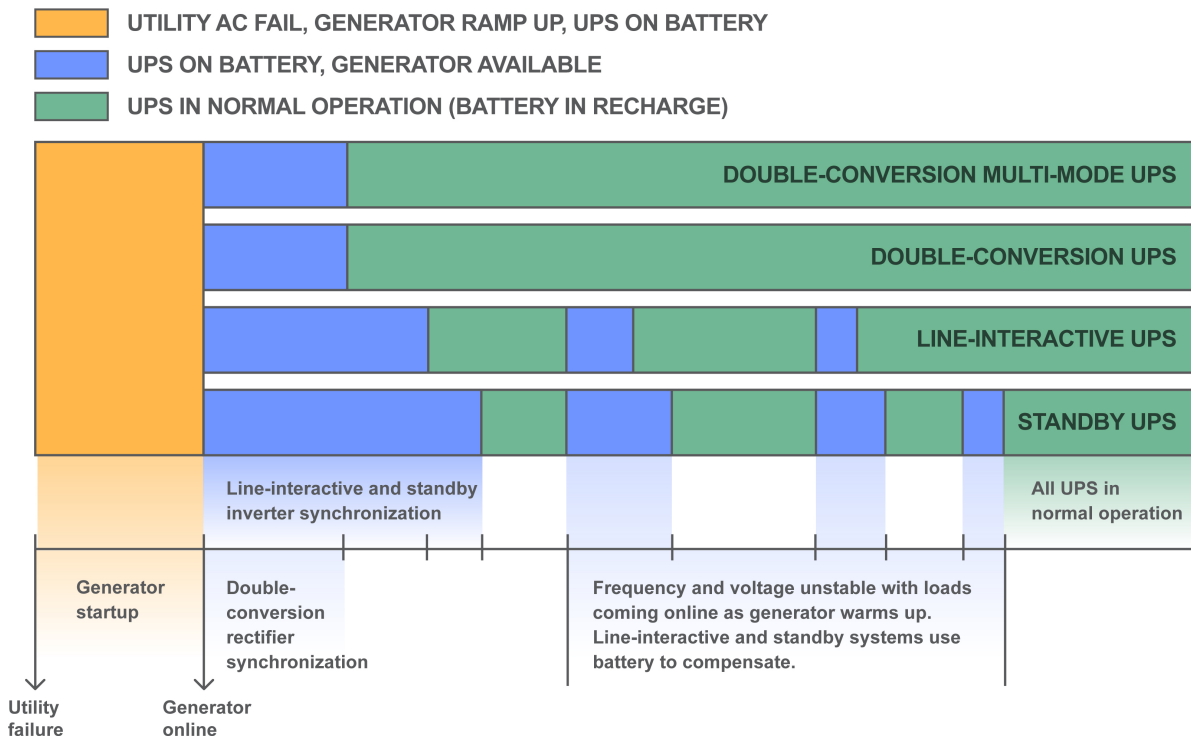


Figure 8. How different UPS designs handle generator start-up
Standby and line-interactive UPSs tend to revert frequently to battery during this process.

How does UPS design affect reliability?

The availability of a UPS configuration depends on several factors, such as the following:

Multiple power paths

A *standby UPS* typically has two power paths, but a single power switch serves them both. That means a failure in the power switch will cause the IT equipment to lose power. A *line-interactive UPS* has two power paths, but without that shared power interface. If the power interface fails, this UPS could still operate in battery mode—long enough to transfer to generator power or gracefully shut down connected equipment.

Double-conversion and *high-efficiency double-conversion UPSs with multi-mode operation* typically have two power paths—from utility/generator and battery sources—plus an electronic system bypass that is used to bypass failed components or to synchronize the system to a mechanical bypass system to perform planned maintenance. Leading multi-mode systems even provide an automated maintenance bypass system to ensure uninterrupted transfers during UPS servicing.

Redundancy through paralleling

You can increase reliability and availability by deploying multiple UPS systems to work together. In a parallel configuration, multiple UPSs feed a common output bus, which then supplies power to the IT equipment. If any UPS should fail, the others are sized to take over the load.

Because of the added costs of manufacturing systems that can be paralleled, this capability is reserved for higher-end UPSs where availability is critical—and that means *double-conversion* and *double-conversion UPS with multi-mode operation*.

Rapid mean time to repair (MTTR)

Mean time between failures (MTBF) is more a theoretical figure than a practical one, based on statistical extrapolation from component ratings and laboratory testing. It is really more important to know the unit's MTTR. When the UPS does require service, a product with a very low MTTR will be back in duty much sooner, which ultimately has a more profound impact on overall availability than MTBF.

MTTR is better in modular system designs and those with easily serviced components, such as hot-swappable battery and electronics modules. Modular systems are more costly to manufacture, so modularity is generally reserved for *line-interactive*, *double-conversion* and *double-conversion UPSs with multi-mode operation*.

Some standby UPSs have a very limited modularity—they can accept replacement batteries—but in general, standby systems are used in smaller, non-critical applications where the entire unit can easily be replaced without undue expense.

Battery health

The UPS design dictates the frequency of battery usage for any given power grid condition—which in turn affects battery runtime and service life. Battery drain is lowest in double-conversion and high-efficiency double-conversion with multi-mode designs. In addition, some manufacturers use multi-stage charging techniques that offer a battery rest period, significantly extending battery service life when compared to conventional trickle or float charging methods. This advanced battery technology is generally found in *line-interactive*, *double-conversion* and *double-conversion UPSs with multi-mode operation*.

How does UPS topology affect energy efficiency?

The more efficient the UPS, the less utility power you have to buy to run your data center. And since most of the lost power is dissipated as heat, the more efficient the UPS, the less you pay in air conditioning or other cooling systems to remove that heat. When overall data center infrastructure efficiency (DCiE) is high, cooling costs might only equal 50 percent of spending on energy to power the IT equipment. When energy efficiency is poor, it could cost almost as much to cool the data center as it does to run the equipment—as much as 80 to 100 percent of the cost of powering the IT equipment, according to industry studies.

So it's no surprise that data center managers are taking a close look at the efficiency of their power protection systems.

Fortunately, technology advances have dramatically improved UPS efficiency in the last three decades. In the 1980s, most UPSs were 75 to 80 percent efficient at best. For every dollar of utility power, you'd get only 75 to 80 cents worth of usable power. Energy losses are dissipated as heat, so that meant more expensive cooling.

By the 1990s, UPS efficiency had risen to 85 to 90 percent efficiency. The 2000s saw the efficiency increase to 94 percent. Then, as more pressure was felt from escalating energy costs, efficiency increased to 97 percent or better with UPSs optimized for today's IT equipment power supplies. The latest generation of UPSs is changing the game with energy-saving technologies that enable up to 99 percent efficiency without sacrificing reliability.

TECHNOLOGY ADVANCES CONTINUE TO BOOST THREE-PHASE UPS EFFICIENCY

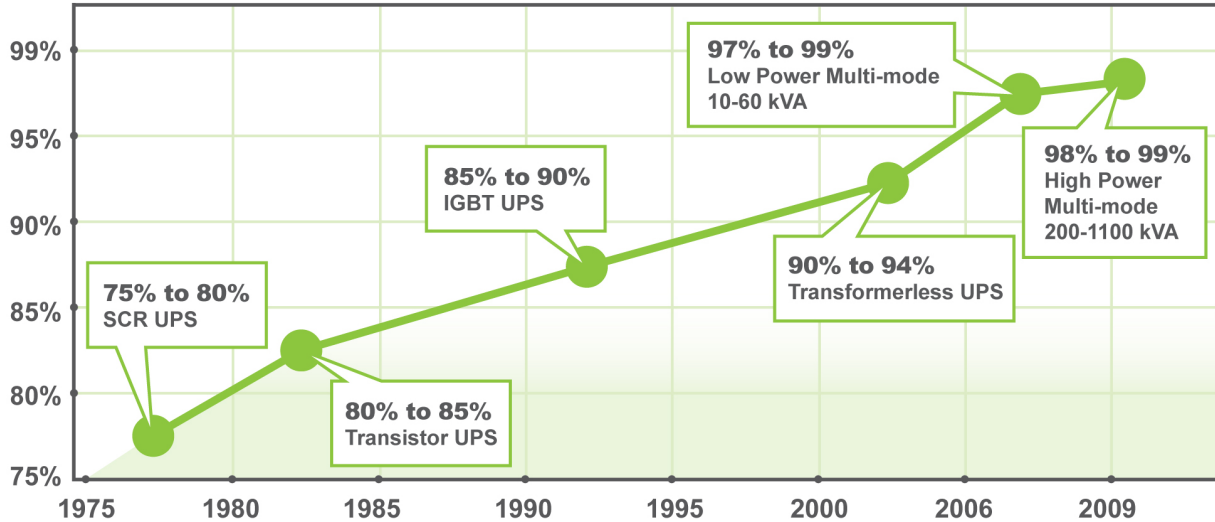


Figure 9. Technology advances have steadily improved UPS efficiency in the last three decades.

Efficiency is profoundly influenced by UPS design or operating mode. Single-conversion (*standby* and *line-interactive*) UPSs are more energy-efficient than *double-conversion* UPSs because there is no power conversion from AC to DC and then back to AC. New high-efficiency, *double-conversion UPSs with multi-mode operation* achieve very high efficiency because they only use the less efficient double-conversion mode when necessary, and operate as an energy-saving system the rest of the time.

Efficiency is also a factor of UPS size. Larger UPS modules typically have higher energy efficiency than smaller ones, because the support power required for control electronics and auxiliary components becomes a smaller portion of the total capacity of the UPS system. For example, a 500 kW UPS module of a given design would typically be more efficient than a 5 kW UPS module of the same design.

Relative energy efficiency of modern, transformer-less UPS designs (percent)

| UPS Module Size | Standby UPS | Line-interactive UPS | Double-conversion UPS | High-efficiency double-conversion multi-mode UPS |
|-----------------|-------------|----------------------|-----------------------|--|
| 5 kW | 95 | 96 | 91 | 96 |
| 100 kW | 98+ | 97+ | 94 | 98+ |
| 500 kW | 99+ | 98 | 95+ | 99+ |

Don't take vendors' efficiency ratings at face value. When evaluating a UPS, it's not enough to know the peak efficiency rating it can deliver at full load (the efficiency figure usually given). You are unlikely to be operating the UPS under full load. Since so many IT systems use dual power sources for redundancy, the typical data center loads UPSs at less than 50 percent capacity, as little as 20 to 40 percent in some cases. You would expect efficiency to be lower when the UPS is operated at partial loads, but to what degree?

Previous-generation UPSs (those bought before 1990) are markedly less efficient at low loads. Even most of today's UPSs are noticeably less efficient under the low loads typically expected of them. New multi-mode topologies with advanced power management features are changing that trend. You can now expect very high efficiency profiles of 95+ percent all the way down to 20 percent loading.

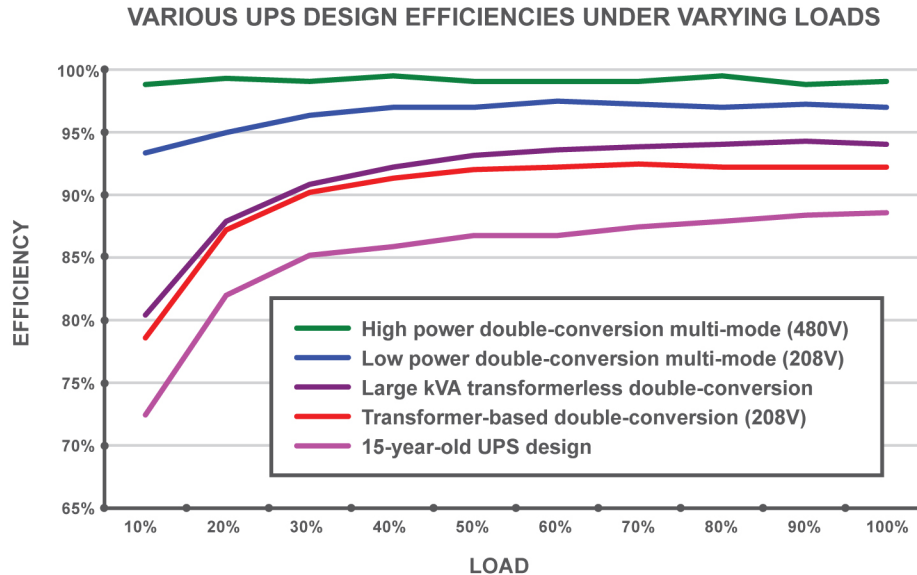


Figure 10. UPSs tend to be less efficient at light loads. New multi-mode UPSs can maintain very high efficiency profiles even at low loads.

Even small increases in UPS efficiency can quickly translate into thousands of dollars. In a one megawatt data center, a 10-year-old UPS is probably wasting about 150 kW of power and generating over 500,000 BTU of heat, which needs to be removed. Replacing that vintage equipment with new, high-efficiency UPSs can free up 120 kW or more of that power to support new IT equipment.

The savings can be dramatic. Imagine a typical data center with 1000 servers and a legacy UPS operating at 86 percent efficiency. Changing out that UPS for a new multi-mode system operating at 96 percent efficiency would save this data center more than \$70,000 in energy costs per year (at 10 cents per kwh). Higher UPS efficiency also provides more battery runtime for the same battery capacity and produces cooler operating conditions within the UPS environment, which in turn extends the service life of components and increases the overall reliability and performance of the data center.

Closing thoughts

In the past, the popular assumption was that for mission-critical data center applications, UPSs had to operate only in double-conversion mode. A system operating in double-conversion mode completely isolates IT equipment from the irregularities of input power, without undue stress on internal batteries. And it provides seamless operation from battery mode to generator operation and back again, so as not to even slightly interrupt power to the connected IT equipment.

However, data center managers now have viable and remarkably cost-effective new choices with high-efficiency, double-conversion, multi-mode UPSs that combine the best of single- and double-conversion topologies: exceptional efficiency plus the high protective level of double-conversion operation.

With best practices and the right choices in equipment, data center managers can reduce energy consumption by nearly 50 percent. That means that almost three-quarters of the power utility bill will fuel actual IT processing, compared to less than 50 percent of the power supplied to a normal data center today.

With a more efficient allocation of power, you not only reduce utility bills and total operating cost, but also achieve more with available backup power and cooling systems—delaying the point where those systems would have to be upgraded to support data center expansion.



Powering Business Worldwide

So, which UPS topology is right for your data center? Where there once was only one “right” answer, new technologies offer effective new choices specifically designed for high-efficiency, high-density data centers.

About Eaton

Eaton Corporation is a diversified power management company with 2008 sales of \$15 billion, with approximately 75,000 employees and customers in more than 150 countries.

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 manufactures, sells and supports UPSs that meet the demands of today’s data center equipment. All Eaton UPSs for data centers meet or exceed the specifications of IT power supply units for reliable, high quality power. For more than 40 years—from the first commercial UPS to the modular Eaton BladeUPS for high-density server environments and ultra-high-efficiency UPSs—Eaton has set the standard for power protection and backup power.

In addition to UPSs, 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.

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

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>

¹ *IEEE Recommended Practice for Emergency and Standby Power Systems for Industrial and Commercial Applications*, the “IEEE Orange Book,” (IEEE Std 446-1995), December 1995

² *IEEE Recommended Practice for Powering and Grounding Sensitive Electronic Equipment*, “IEEE Emerald Book,” (IEEE Std. 1100™ - 2005), (Revision of IEEE Std 1100-1999)

Eaton is a trademark of Eaton Corporation.