

Picking the Right UPS for Your Data Center
A White Paper on Availability



Introduction

Availability concerns are very familiar to data center managers. It comes down to simple economics: given the high cost of downtime, availability naturally became the most important metric on which data centers were evaluated in recent years. It wasn't uncommon for data center managers and corporations to operate under an "availability at any cost" approach. This mindset also dominated the selection criteria for uninterruptible power supply (UPS) systems.

New technologies and increasing demands are driving changes in the power systems that support data centers. While the importance of availability has in no way diminished, it has been joined by two new metrics: efficiency and flexibility (or scalability), as data center managers increasingly struggle to keep pace with growing capacity needs while working under the constraints of tightened budgets and energy efficiency initiatives.

This white paper is designed to provide data center managers with a clearer understanding of key factors and considerations involved in selecting the right UPS to meet current and future needs, and optimize the performance of their facilities.

UPS Topology

UPS topology, along with power system architecture, impacts data center availability and efficiency. Because the power system uses a relatively small amount of total data center power and is so critical to ensuring uptime, more robust topologies are generally preferred, especially because continued improvements in power technology allow robust power systems to deliver high efficiency.

UPS internal topology (configuration) determines the relationship between the UPS output power quality and incoming utility power and, ultimately, the effectiveness of the UPS at protecting against certain types of power disturbances. Topology is a primary determinant of system reliability and functionality, two key elements of availability. Historically, there have been three basic types of UPS topology used in the data center: passive standby, line interactive and on-line double-conversion.

Passive Standby

The simplest type of UPS is the offline or passive standby topology (Figure 1). Under normal operating conditions, AC power from the utility passes straight through the UPS to the critical load. A charger converts AC power to DC to charge the battery. The inverter is used to convert the DC power from the battery to create AC power to support the load when the utility fails. Normally the inverter is operating in a standby mode, with the battery charger maintaining the batteries. Should the utility power go out of limits, the inverter powers the load, drawing energy from the battery.

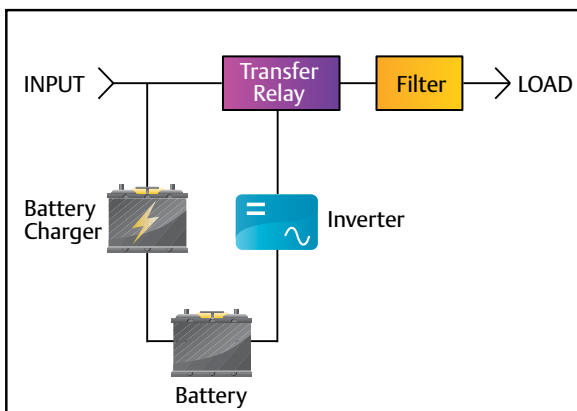


Figure 1. Passive standby UPS topology.

While this topology gets high marks for simplicity of design, the limitations of the design—such as the lack of power conditioning, brief power interruption during transfers, or generator compatibility (eg: gensets generally have poor step-load capabilities) make it unsuitable for three-phase applications. These types of UPS modules are typically only applied to single-phase (workstation-level) loads.

Line-Interactive

The next step upward is the line-interactive topology (Figure 2). It resembles passive standby, but inserts a transformer or inductor in series between the utility power source and the load. This inline inductor enables the UPS power converter to “interact” with incoming power and provide a measure of power conditioning to the load, which also helps with high and low input voltage conditions.

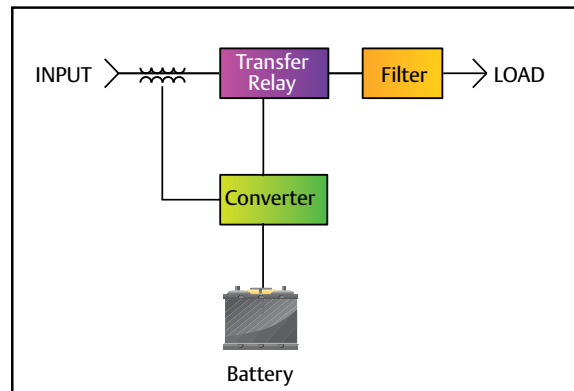


Figure 2. Line interactive UPS topology.

One major limitation of this topology is that it cannot completely isolate the critical load from the input line without operating on battery. Small perturbations in frequency and power quality can get passed directly to the critical load. Without electrical isolation, common-mode noise also passes right through to the load.

Like the passive standby topology, line-interactive can be lower-cost but effective, because it supports the entire critical load during power disturbances (for the duration of the battery). Line-interactive may be slightly less efficient, as result of the incremental losses associated with its power conditioning function. While similar to passive standby, line-interactive UPS systems are capable of soft-loading a standby genset, reducing the step-load effects.

 **Tweet This:**

White paper:
Choosing between
passive standby, line
interactive and
online double-
conversion US systems
<http://bit.ly/16BW3Ea>

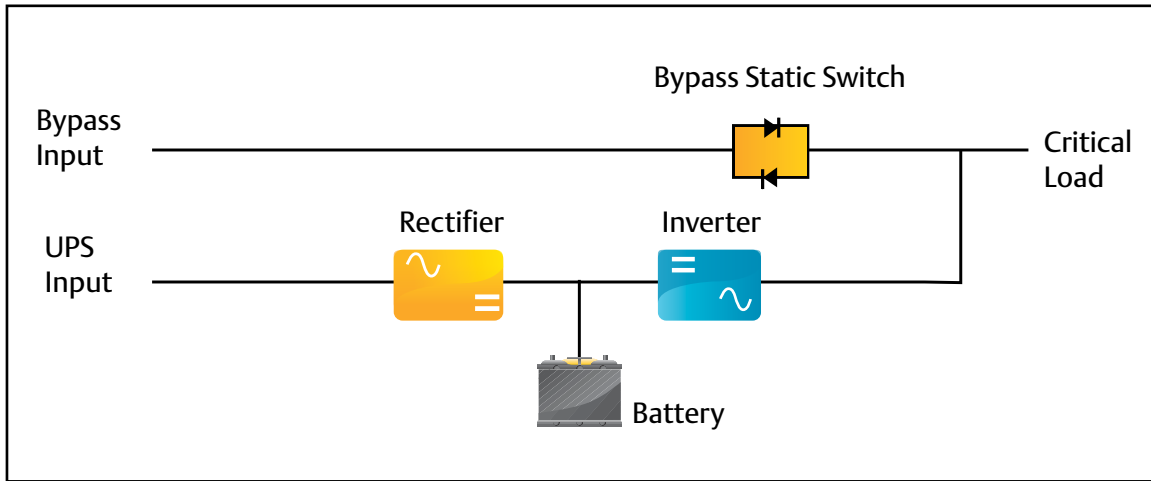


Figure 3. Online double-conversion UPS topology.

Online Double-Conversion

Online double-conversion topologies (Figure 3) have emerged as the preferred topology for the data center, displacing standby and line-interactive UPS systems in network access rooms as overall network availability requirements rise. Properly configured, it is the only topology that protects against the full range of power disturbances experienced by data center power systems.

With online double-conversion, incoming AC power is rectified to DC power to supply the internal DC bus of the UPS. The output inverter takes the DC power and produces regulated AC power to support the critical load. Batteries attached to the DC bus are float charged during normal operation. When the input power is out of limits, the batteries provide power to support the inverter and critical load.

The critical load is always being supplied by the output inverter, which is always being supplied from the internal DC bus. When input power fails, there is no transitional sag in the output voltage because the inverter is already operating on DC input.

Online double-conversion UPS topologies (Figure 4) deliver better protection than other types of UPS systems because they completely isolate sensitive electronics from the incoming power source, remove a wider range of disturbances and provide a seamless transition to backup power sources. With the combination of improved operating efficiencies and the use of new technologies, these robust power systems can achieve efficiencies greater than 96 percent without any impact on availability.

Power Problem Resolution by Topology			
Power Problems	Passive Standby	Line-Interactive	Online Double-Conversion
Line Noise			●
Frequency Variation			●
Switching Transients			●
Harmonic Distortion			●
Under Voltage		●	●
Over Voltage		●	●
Power Sag	●	●	●
Power Surge	●	●	●
Power Outage	●	●	●

Figure 4. Select your topology based on your local power profile and required availability.

Availability, Efficiency and Scalability

The mission-critical data center is one of the most dynamic and critical operations in any business. Complexity and criticality have only increased in recent years as data centers experienced steady growth in capacity and density, straining resources and increasing the consequences of poor performance.

As the data center continues to evolve and become more complex and heterogeneous, it is important that organizations now select those power systems that optimize the “trifecta” of availability, efficiency and scalability.

Availability

It is easy to understand how availability became the dominant metric on which data centers and power systems were evaluated. One significant outage can be so costly that it wipes out years of savings achieved through incremental efficiency improvements, and can severely impact an organization’s bottom line.

The actual costs associated with an unplanned outage can be alarming. A 2011 Ponemon Institute study, sponsored by Emerson Network Power, revealed that the mean cost for any type of data center outage is \$505,502, with the average cost of a partial data center shutdown being \$258,149. A full shutdown can, on average, cost more than \$680,000. In some applications, a power outage may result in the complete loss of the business enterprise.

According to survey respondents, more than 39 percent of data center outages reported were attributed directly to vulnerabilities in data center power. Among the general root causes of downtime related to power, UPS related failures (including batteries) proved to be the most costly at \$687,700.

One of the primary reasons power vulnerabilities are so costly for data centers is that a failure in the power infrastructure will likely result in a catastrophic, total unplanned outage. This means that in addition to any direct costs incurred to remedy the cause of the outage, indirect and opportunity costs also will be significant due to the fact that all stakeholders will be affected by the outage.

Mathematically, Availability “A,” is expressed:

$$A = (MTBF / (MTBF + MTTR)) \times 100\%, \text{ where}$$

MTBF = Mean Time Between UPS Output Failures

MTTR = Mean Time To Repair (the total time from the power failure to power restoration)

To improve availability, increase MTBF and/or reduce MTTR.

For more information on the causes and cost of downtime, see Emerson Network Power’s white papers:

- “Understanding the Cost of Data Center Downtime”
- “Addressing the Leading Root Causes of Downtime.”

 **Tweet This:**

Study: more than 39% of #datacenter outages attributed to vulnerabilities in power. Availability tips: [http:// bit.ly/16BW3Ea](http://bit.ly/16BW3Ea)

Does your UPS protect the load, or itself?

Many UPS systems boast reliability, but not all UPS models are designed to protect the connected equipment. Your UPS should be able to withstand multiple, simultaneous power and environmental events without forcing a switch to bypass.

What to look for:

- **Continuous-duty static switch:** Ensures maximum fault clearing capability in case of extreme overload or downstream faults.
- **No contactors in the bypass path:** Contactors are not designed for high withstand requirements.
- **Redundant fans:** Ability to tolerate more than a single fan failure before going to bypass.
- **High withstand rating (100kA or higher):** Allows the equipment to stay online during more severe short circuit or circuit overload conditions.
- **Start a like-sized Power Distribution Unit (PDU) without transfer to bypass:** Keeps the system on protected power when additional distribution is added.

Efficiency

Data center managers today are feeling increasing pressure to operate efficiently—in terms of both energy and management resources.

This is reflected in results from the Data Center Users Group (DCUG) study sponsored by Emerson Network Power, where for the last few years, availability and efficiency have been top concerns of respondents. For the fall 2012 DCUG survey, when asked to identify their top three facility/network concerns, 48 percent of respondents cited energy efficiency, while nearly 46 percent cited availability (Figure 5).

Power systems in use today in many data centers still remain below the efficiencies currently available. While server power supply efficiencies have improved significantly since the days of 79 percent power supply efficiencies, they continue to consume more energy than is necessary. Most estimates put the average power supply efficiency at 86.6 percent, well below the 93 percent that is available.

Emerson Network Power conducted a systematic analysis of data center energy use and the various approaches to reducing it to determine which were most effective. Using this analysis, Emerson

developed Energy Logic, a vendor-neutral roadmap of 10 strategies that can reduce a data center's energy consumption by up to 74 percent. One of the findings of the analysis was that increasing UPS efficiency from 86.6 percent to 93 percent reduces total data center power consumption by 7.1 percent.

For more information on Energy Logic, see Emerson Network Power's ebook: "Energy Logic 2.0: New Strategies for Cutting Data Center Energy Costs and Boosting Capacity"

As with other data center systems, server power supply efficiency varies depending on load (Figure 6). Figure 6 shows power supply efficiencies at different loads for two power supply models. At 20 percent load, model A has an efficiency of approximately 88 percent while model B has efficiency closer to 82 percent. Enabling them to provide higher efficiency, some new UPS modules now have flatter efficiency curves than traditional units even while running at partial loading.

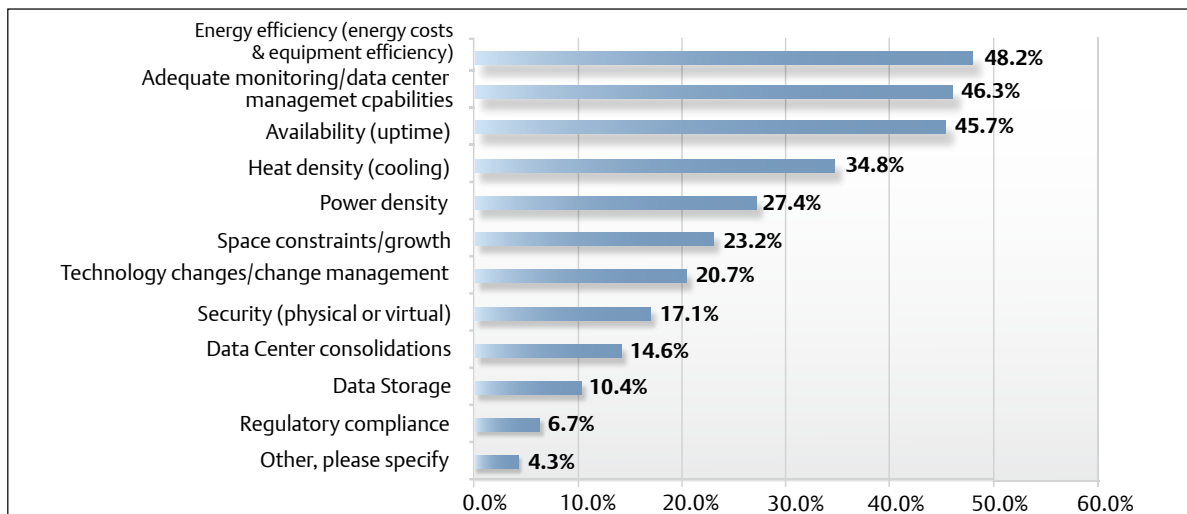


Figure 5. Respondents to 2012 DCUG survey identified top three facility / network concerns.

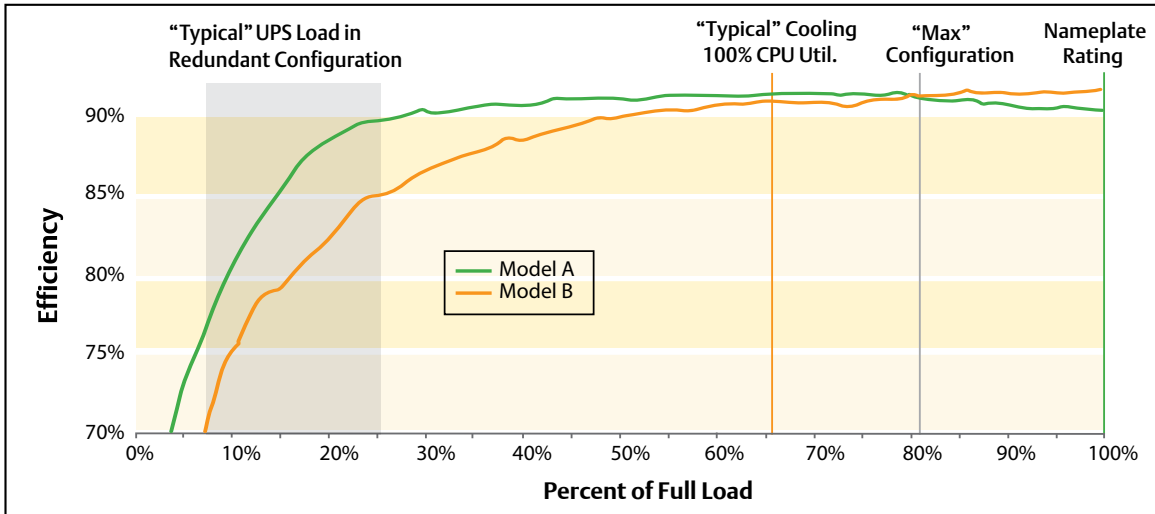


Figure 6. Power supply efficiency at various loads.

Scalability

For most in the industry, the days when data center managers had huge areas of unused floor space set aside for a decade’s worth of growth are long gone. This is evident in the Fall 2012 DCUG survey results, which show that maximum capacity issues continue to threaten to strain resources and negatively impact performance levels. Forty percent of respondents said they expect to run out of data center capacity by 2014, with another 29 percent expecting capacity constraints before 2017 (Figure 7).

Increasing flexibility in how IT systems and data centers are operated and managed is now the best way to achieve availability while maximizing opportunities to improve energy and operational efficiencies. It also allows data center managers to meet current requirements, while ensuring the ability to quickly and cost-effectively meet growing and fluctuating performance and capacity needs.

Today, this flexibility can be provided by a new generation of power system designs that are designed for simplified configuration changes and greater scalability, enabling systems to be right-sized during the design phase. Virtually every component of the data center physical infrastructure is now available in a modular design; this includes power systems.

These modular systems are typically integrated packages of equipment and controls that are built and pretested in a factory environment. Deployment times for some customers - depending on the size of the building block - can be days for smaller modules, or 4 to 6 weeks for even large UPS modules, assuming vendor managed inventory. The shorter deployment time, combined with the “build it as you need approach” offered by modularity, allows organizations to quickly meet fluctuating demands while lowering capital expenditure (CapEx) costs and maintaining higher operating efficiencies.




Figure 7. DCUG respondents project data center capacity constraints.

UPS Design

There are many options to consider in the area of power system design that can optimize efficiency, availability and scalability. There are a number of new technologies that expand upon the online double-conversion topology to lower CapEx and operational expenditure (OpEX) costs, while providing greater scalability and minimizing availability compromises.

Tweet This:

Transformer-based vs. Transformer-free UPS? Weighing the availability, efficiency & flexibility benefits <http://bit.ly/16BW3Ea>

Transformer-free UPS

There has been growing interest in using transformer-free UPS modules in higher power, three-phase mission critical power backup applications (e.g.; 200 kW to 5 MW).

In general, well-proven transformer-based UPS systems are highly robust and excel at providing the highest capacities and availability while simplifying external and internal voltage management and fault current control. The latest transformer-free designs offer better efficiency, smaller footprint and improved flexibility while providing high levels of availability. Driven by data center designer demand, most leading UPS suppliers offer both topologies.

Large transformer-free UPS systems are typically constructed of smaller, modular building blocks that deliver high power in a lighter weight, compact package. This modular design offers advantages when the timing of future load requirements is uncertain by allowing capacity to be more easily added as needed, either physically or via control settings.

For high-power enterprise data centers and other critical applications, a state-of-the-art transformer-based UPS still provides an edge in availability. However, if operational efficiency, expansion flexibility or limiting UPS footprint are of paramount importance, and other appropriate measures are instituted to provide an acceptable level of availability, transformer-free technology may be the optimal choice. In general, 300 kW is a threshold below which the space, weight and cost advantages of transformer-free UPS systems outweigh the robustness and higher capacity capabilities of transformer-based systems.

For more information on transformer-free UPS systems, see Emerson Network Power's white paper: "Comparing Transformer-free to Transformer-based UPS Designs" on EmersonNetworkPower.com.

Capacity-On-Demand

Some UPS technologies now enable modularity within the UPS module itself (vertical), providing capacity-on-demand that enables data center managers to meet current capacity requirements, while ensuring they can easily and cost-effectively adapt to future demands.

Certain UPS units in the market provide capacity on demand by using software to increase capacity without adding to the system footprint. UPS units with this sort of software functionality can be sized to current requirements and then easily scaled up to a larger capacity with a simple software key as power needs change. These types of modules are scalable in increments, such as scalable from 40 to 80 kVA, from 80 to 120 kVA, or from 160 to 200 kVA.

Some row-based, modular UPS units allow quick power capacity increases with the addition of internal power core hardware assemblies. These core assemblies allow the system to expand for capacity or redundancy in 15 kW increments up to 45 or 90kW within a single cabinet. It is important to look for power cores that incorporate distributed intelligence and scalable power in a common assembly and allow configuration of a completely redundant power and control system, sized to match the capacity of the protected equipment. When power requirements change, data center managers can easily add capacity without increasing the system footprint. This approach also allows for right-sizing of the UPS, resulting in improved energy efficiency and reduced power expenditures.

Does your UPS frustrate your data center and facilities employees?

A well-designed UPS should be user-friendly, not a source of frustration.

What to look for:

- **Ability to be fed from separate utility sources:** Higher availability can be achieved with separate source capability.
- **Compliance with emissions regulations, including FCC:** A UPS should meet the latest FCC standards for both radiated and conducted noise – interference with wired and wireless communications could be the result of a UPS that states it meets IEC requirements, but falls short of meeting FCC.
- **Large, informative user-friendly display:** Display that shows unit and system one-line in addition to measured parameters, alarm status and help screens reduce the risk of human error and simplifies methods of procedure.
- **Local 7 x 24 factory-trained support:** Scheduled maintenance, repair and advice from local domain experts.

Active Eco-Mode™

An economy mode of operation (Eco-Mode) changes the way a double-conversion UPS is employed to increase energy efficiency and reduce the operating expenses experienced in the data center.

Eco-Mode allows Double-Conversion UPS systems to achieve the highest possible efficiencies by transferring the UPS to static bypass during normal operation. When power problems are detected, the UPS instantaneously switches back to Double-Conversion Mode. Until recently, this transfer was not seamless, because when the critical load was being powered through the bypass, the rectifier and inverter were switched off. This resulted in a slight inverter start-up delay and a notch in the output waveform when the critical load returned to Double-Conversion Mode. While similar to the passive standby topology previously described, Active Eco-Mode provides a degree of bypass power conditioning as well.

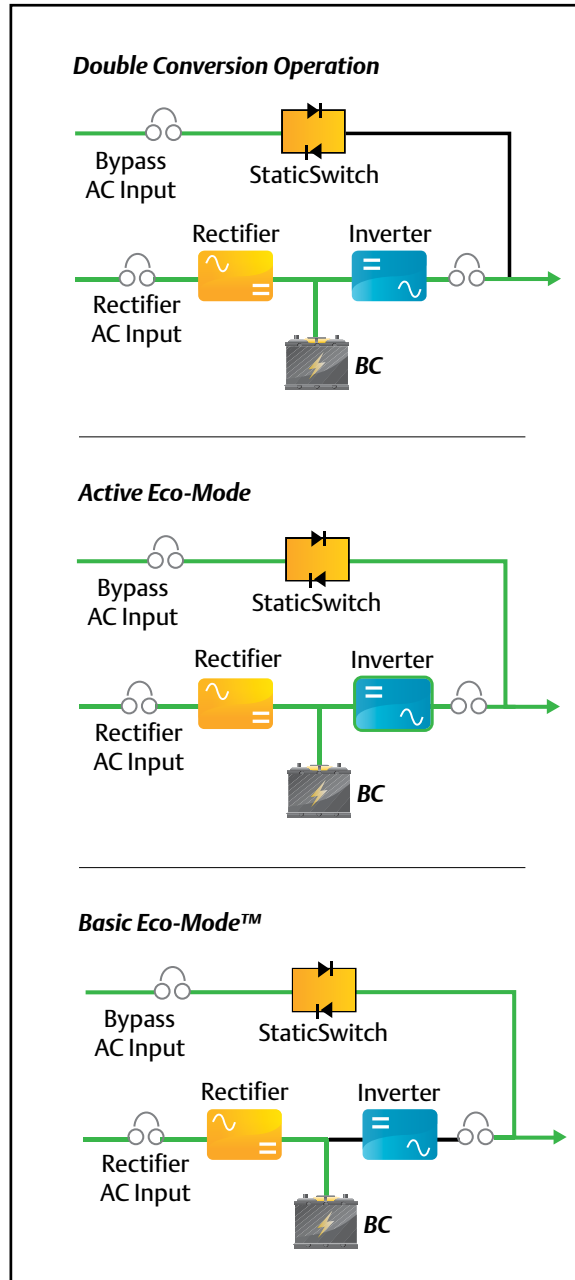


Figure 8. Double Conversion operation, Active Eco-Mode and Basic Eco-Mode configurations.

The Active Eco-Mode approach (Figure 8) eliminates prior eco-mode issues by keeping the inverter and rectifier in an active state. Because the inverter is kept in an active state, and is providing an output that matches the bypass, it is ready to accept the load instantly. As a result, the transfer to the inverter can be accomplished almost seamlessly.

Tweet This:

Active Eco-Mode can protect #datacenter availability while still achieving efficiency levels of 98%+ <http://bit.ly/16BW3Ea>

Tweet This:

White paper:
Tips and guidance
for picking the
right UPS for
your #datacenter
<http://bit.ly/16BW3Ea>

Keeping the inverter in a constant state of preparedness does require additional power; however, the power requirement is below 1 percent of the UPS rated power, creating potential savings of 4 to 4.5 percent compared with typical double-conversion operating modes. The result is a more continuous output wave form, a more seamless transfer and possible UPS efficiency levels above 98 percent. Also, keeping the rectifier in an active state ensures a constant float charge across the UPS battery, which is shown to have a significant positive impact on battery life.

Does your UPS sacrifice high availability for high efficiency?

UPS systems with Basic Eco-Mode™ offer high efficiency, but at the expense of availability. Models with Active Eco-Mode™ can deliver even higher efficiency without compromising availability. State-of-the-art efficiency and energy optimization modes deliver high efficiency without compromising availability.

What to look for:

- **High efficiency in typical load range:** Don't judge UPS efficiency at full load calculations – your system rarely runs at full load. Compare efficiencies in the 40-90 percent load ranges to get a true picture of operating costs.
- **Active Eco-Mode with no break operation:** Active inverter technology allows the inverter to instantaneously assume the load when switching from double conversion to Eco-Mode, and also provides a cleaner waveform.
- **Active Eco-Mode as a standard feature:** If a system employs a momentary-duty static switch, you will probably have to pay an upcharge to get a continuous-duty static switch and eco-mode.

Intelligent Paralleling

Another newer function enabled by UPS controls is intelligent paralleling (Figure 9), which improves the efficiency of redundant UPS systems by deactivating UPS modules that are not required to support the load and taking advantage of the inherent efficiency improvement available at higher loads.

For example, a redundant multi-module UPS system configured to support a 500 kVA load using three 250 kVA UPS modules can support loads below 250 kVA with only two modules, while maintaining redundancy and improving the efficiency of the remaining on-line UPS modules by enabling them to operate at a higher load level.

This feature is particularly useful for data centers that experience extended periods of low demand, such as a corporate data center operating at low capacity on weekends and holidays.




Figure 9. Intelligent paralleling of UPS modules.



Conclusion

As the complexity of data center environments continues to grow, data center managers are increasingly concerned about maintaining or improving availability in increasingly dense computing environments while increasing efficiency and reducing costs. Therefore, it is important they pick the right flexible UPS technology for their data center.

Historically, data center designers and managers have had to balance tradeoffs between availability and efficiency in the data center power system. Online double-conversion UPS modules provided the highest availability but could not deliver the same efficiency as a line-interactive topology.

Advances in double-conversion UPS technology have closed the gap in efficiency, and new features have been introduced that enable double-conversion UPS systems to achieve efficiencies on par with line-interactive systems while still providing the needed power quality. These technology advances and new features help organizations safeguard availability while increasing efficiency and advancing scalability in the data center.

Emerson Network Power-Global Headquarters

1050 Dearborn Drive
P.O. Box 29186
Columbus, Ohio 43229
Tel: +1 614 888 0246

liebert.com

Emerson Network Power.

EmersonNetworkPower.com

While every precaution has been taken to ensure accuracy and completeness in this literature, Liebert Corporation assumes no responsibility, and disclaims all liability for damages resulting from use of this information or for any errors or omissions.

© 2013 Liebert Corporation. All rights reserved throughout the world. Specifications subject to change without notice. All names referred to are trademarks or registered trademarks of their respective owners.

® Liebert and the Liebert logo are registered trademarks of the Liebert Corporation. Business-Critical Continuity, Emerson Network Power and the Emerson Network Power logo are trademarks and service marks of Emerson Electric Co. ©2013 Emerson Electric Co.

SL-24676 (R07-13) Printed in USA

EMERSON. CONSIDER IT SOLVED.™