

Evaluating the Opportunity for DC Power in the Data Center

by Mark Murrill and B.J. Sonnenberg, Emerson Network Power

Summary

With data center managers struggling to increase efficiency while maintaining or improving availability, every system in the data center is being evaluated in terms of its impact on these two critical requirements. The power system has proven to be one of the more difficult systems to optimize because efficiency and availability are often in conflict; the most efficient approach to critical power is rarely the most reliable.

One solution to power system optimization that deserves serious consideration is DC power. Since utility AC power must ultimately be converted to DC for use by IT system components and because stored energy systems (batteries, flywheel, etc.) provide DC power for backup, a DC power architecture requires fewer total conversions from grid to chip, creating the opportunity to reduce costs and increase efficiency.

A data center-optimized, row-based DC power protection system is now available to help data center operators take advantage of that opportunity. This system, combined with the availability of 48V DC-powered IT equipment from major manufacturers, makes DC power an ideal solution for small and midsize data centers seeking to optimize efficiency, reliability and scalability. Other applications include high-density equipment rows with a consistent footprint and pod-based data centers.

As the leading provider of AC and DC power systems, Emerson Network Power is uniquely positioned to help organizations evaluate the suitability of DC power and determine whether a row-based DC infrastructure is appropriate for a given application.

Introduction

The first decade of the twenty-first century was one of incredible growth and change for data centers. The demand for computing and storage capacity exploded, and many IT organizations struggled to deploy servers fast enough to meet the needs of their businesses. At the same time, the trend to consolidate data centers and centralize computing resources resulted in fewer opportunities for planned downtime while also increasing the cost of unplanned outages.

Data center operators were able to meet the demand for increased compute capacity by deploying more powerful servers—often in the same physical space as the servers being displaced—creating a dramatic rise in data center power consumption and density. Between 2004 and 2009, power and heat density became top

concerns among data center managers as they struggled to adapt to a 400 to 1,000 percent increase in rack density.

The dramatic increase in data center energy consumption created both financial and environmental challenges. Energy costs, which once had been relatively inconsequential to overall IT management, became more significant as the rise in consumption was exacerbated by a steady—and in some years significant— increase in the cost of electricity. In addition, increased awareness of the role that power generation plays in atmospheric carbon dioxide levels prompted the U.S. EPA to investigate large energy consumers such as data centers. In 2007 the EPA presented a report to the U.S. Congress that included recommendations for reducing data center energy consumption.

The industry responded with a new focus on energy efficiency and began implementing server virtualization, higher-efficiency server power supplies, and new approaches to cooling. Yet, while significant progress has been made in some areas, the critical power system has yet to be fully optimized. While individual components have been improved, the overall system complexity is high, which can create inefficiency and add operational risk. Faced with the choice of increasing system efficiency or adding risk, many continue to choose proven approaches that deliver high availability but do not deliver the highest efficiency.

However, a close examination of the available options reveals that, in many cases, efficiency can be improved without sacrificing overall availability.

Established Data Center Power Distribution Options

Traditional AC power distribution systems in North America bring 480V AC power into a UPS, where it is converted to DC to charge batteries, and then inverted back to AC. The power is then stepped down to 208V within the distribution system (PDU) for delivery to the IT equipment. The power supplies in the IT equipment convert the power back to DC and step it down to lower voltages that are consumed by processors, memory and storage [Figure 1].

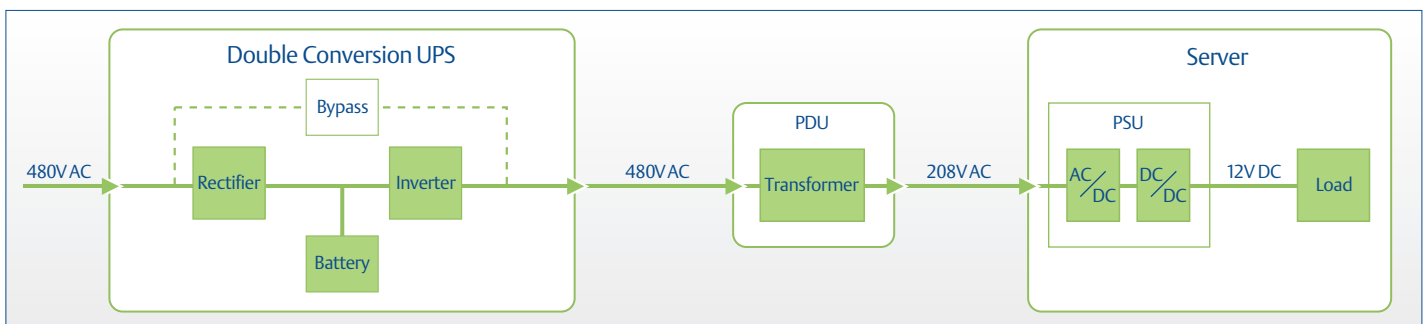


Figure 1. Typical 480V AC to 208V AC data center power system configuration.

In recent years, manufacturers have increased component efficiencies and developed new operating modes, such as the “eco-mode” available on some UPS systems, to improve efficiency.

Figure 2 shows AC UPS operation when in eco-mode. Efficiency is improved by bypassing the conversion in the inverter; however, this introduces the potential for reduced system reliability. Critical loads are no longer isolated from the mains or subject to the tight voltage regulation normally provided by the inverter, and complex synchronization circuits are needed to ensure reliable power transfer.

In the event of a power interruption, a switching function is required to bring the load back on to the inverter, and mains need to be quickly isolated to prevent power feedback into the distribution system. No matter how fast or reliably this switching occurs, there is still a chance it will fail when it is needed most.

There has been discussion in the industry about simplifying the AC distribution scheme by eliminating the step-down in the power distribution unit and delivering higher voltages to the IT equipment. One example, illustrated in Figure 3, takes 277V AC output from the UPS and delivers it directly to the server power supply. This approach may have potential in the future but is not viable today for a number of reasons.

First, by connecting loads phase-to-neutral, additional harmonic currents may be introduced, offsetting some of the anticipated efficiency gains. Second, because of the elevated voltages at the rack, arc flash concerns for IT personnel working in the rack will increase. Finally, there are no commercially available power supplies at this input level. All in all, this architecture needs further development before it can be considered a true alternative.

Consequently, 48V DC power may prove to be the best alternative to 480V/208V AC power distribution for organizations seeking to optimize reliability and efficiency today.

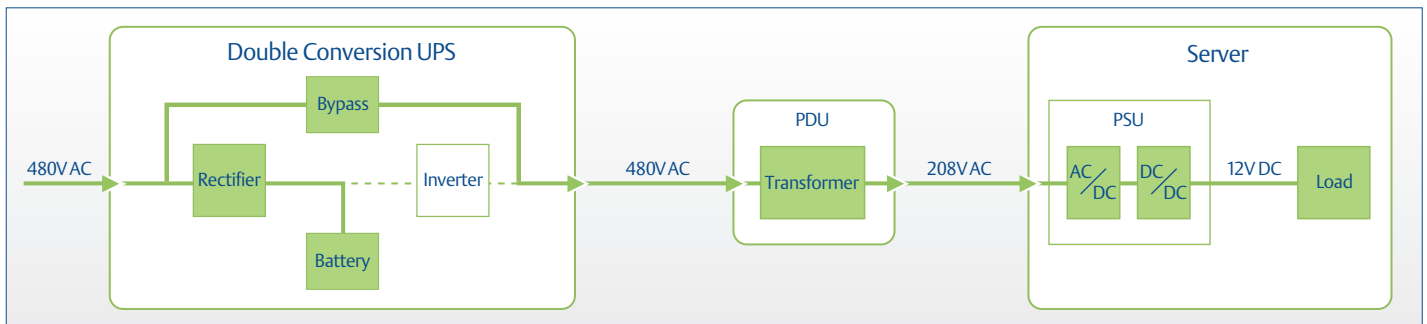


Figure 2. In eco-mode, incoming power bypasses the inverter to increase UPS system efficiency.

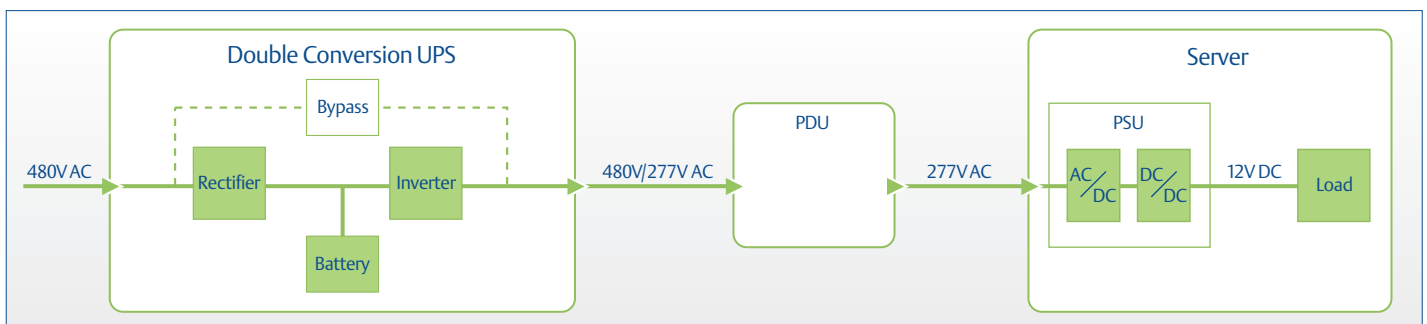


Figure 3. 480V AC to 277V AC data center power system configuration.

A Practical Approach to DC in the Data Center

Historical Background

48V DC power has a long history in telecommunication networks. It is inherently simple and reliable with few conversion stages to the point of use. In Alexander Graham Bell's day, 48V DC was chosen as the standard for two reasons:

1. DC power was felt to be more reliable than AC because it could be directly connected to backup batteries during grid outages;
2. 48V was considered the optimal tradeoff between transmission distance and human safety because it is considered safe to touch during maintenance or accidental exposure.

Today, telephone central offices (exchanges) are still powered by -48V DC. These facilities exhibit levels of availability as high as 9-nines, as confirmed in a study of 23,000 DC power systems over 10 years by NTT Facilities¹.

In most telecommunications installations, the 48V DC power system is deployed as three distinct elements as shown in **Figure 4**:

1. 480V AC to -48V DC modular power system
2. Battery banks for extended run time (at least 8 hours)
3. Load distribution cabinets (BDFB/PDU)

These elements are connected with large copper bus bars and wires routed around the facility to distribute power directly to

-48V power supplies at the point of use. The power supplies step the voltage down to 12V and lower DC voltages for internal uses.

It is easy to see the benefits of this approach when applied to the data center. Downstream of the AC/DC rectifiers, the power is completely isolated from the mains and is considered "safety extra low voltage" (SELV) per IEC/UL 60950 and can be maintained live by trained personnel. In addition, there is no need to derate capacity to account for phase balancing or harmonics, which are not a factor with DC power. Personnel safety is also improved when servicing equipment in the rack because arc flash is not a concern at -48V.

However, unlike telecommunications central offices, data centers are not designed with large copper bus bars to distribute DC power to racks. Fortunately, with new row-based DC topologies, these large copper runs are no longer necessary. In this row-based configuration, power is converted from AC to DC very close to the point of use, which decreases the conductor size and cost [see **Figure 5**].

Data centers are entering a new stage of maturity where the emphasis is on reliably delivering more computing capacity to the organization at lower cost and with reduced power consumption. Row-based DC power represents a practical and affordable solution for reducing data center complexity, increasing efficiency and enabling growth.

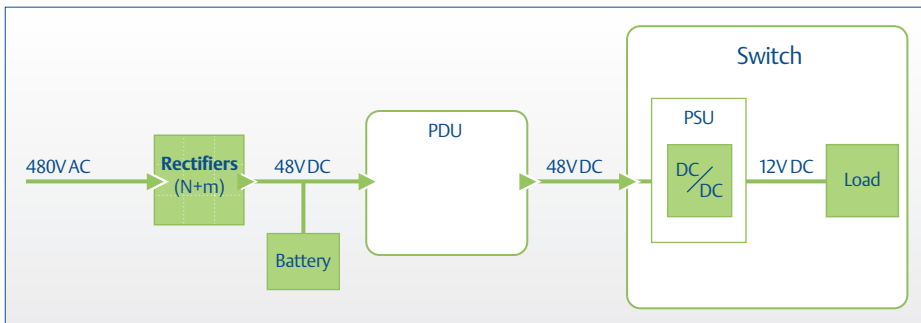


Figure 4. -48V DC power as typically implemented in telecommunications central offices.

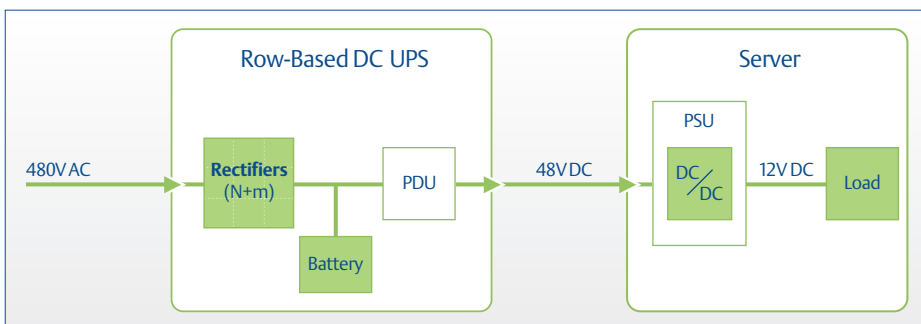


Figure 5. A row-based DC UPS eliminates the long copper runs and is optimized for the data center.

Factors Affecting Power Distribution Selection

When selecting an optimal power architecture for a particular facility, the number of factors data center designers and managers have to consider is substantial. Understanding the impact of each factor is essential to making an informed choice for your data center.

Efficiency

End-to-end power system efficiency can be difficult and time consuming to accurately calculate due to the complexity of variables involved, but it is still useful to consider and approximate efficiency based on specific site conditions and priorities.

In light of recent improvements in the efficiency of both AC and DC architectures, two vendor-neutral studies have been published by The Green Grid and Intel to provide an objective, yet simplified, comparison of the efficiency of various power system configurations.

The results from The Green Grid's study, "Quantitative Efficiency Analysis of Power Distribution Configurations for Data Centers"², for the power architectures discussed here, evaluated at 50 percent operating load, are summarized below:

- 480V AC to 208V AC (legacy): 63 percent efficient [Figure 1]
- 480V AC to 208V AC (current): 88 percent efficient [Figure 1]
- 480V AC to 277V AC: 89 percent efficient [Figure 3]
- 480V AC to 48V DC: 90 percent efficient [Figures 4 & 5]

Intel conducted a similar analysis, "Evaluation of 400V DC Distribution in Telco and Data Centers to Improve Energy Efficiency"³ and found that the highest overall system efficiencies are achieved with a 400V DC system; however, this technology is not yet commercially available. Overall, this study concluded that DC power (both 400V and 48V) was consistently more efficient than AC at all voltages.

Intel concluded that DC power (both 400V and 48V) was consistently more efficient than AC at all voltages.

While each study created scenarios that allowed for direct comparison of different configurations, many aspects of the power distribution system are complex and vary by application, making general comparisons difficult. Aspects which were not considered in either study include the degree of desired redundancy, the impact of harmonics and load variations, and the added load on cooling created by inefficient systems. In cases where these factors are present, the gaps between DC and AC would likely be wider, as the following analysis reveals.

■ Redundancy

The degree of redundancy in the power distribution system can impact both availability and efficiency. In the case of AC power distribution, these two factors may be mutually exclusive—greater redundancy increases availability but lowers overall efficiency.

The highest availability is achieved through 2N redundancy in which each UPS module is designed to run at no more than 50 percent load.

However, since the data center is rarely operating at peak capacities, the redundant systems are almost always operating at less than 50 percent capacity. With derating, maximum capacity may be closer to 40 percent, and, in normal operation, each unit may operate at 20 percent load or less. Most equipment efficiency curves drop dramatically below 30 percent load, resulting in actual operating efficiencies significantly less than ideal conditions.

In the case of DC power, redundancy can be seamlessly integrated into the DC UPS. In addition, Emerson Network

Power offers an efficiency-optimizing feature that maintains near-peak efficiencies at system loads as low as 5%, allowing the system to achieve both high availability and efficiency in real world conditions.

■ Harmonics

Harmonics are a distortion of the normal AC waveform generally transmitted by non-linear loads. In the data center, the switch-mode power supplies used in AC servers represent a non-linear load that can create harmonics. Harmonic currents accumulate in the neutral wire, causing distribution losses and increased heat generation. If the cumulative level of harmonics—known as total harmonic distortion—becomes too high, damage to sensitive electronics and reduced efficiency can result and may require equipment derating to overcome. These losses are difficult to predict in complex AC distribution systems, but could be quite significant. Harmonics are not present in DC systems because there are no waveforms to contend with.

■ Data Center Load Variations

Many data centers are operating at significantly less than 100 percent of design load much of the time, with real-time loads changing frequently. This variation further complicates efficiency modeling. As servers are turned on/off and utilized at different rates, the loading on each phase of a 3-phase power system changes, making alignment between the three phases difficult. Unbalanced loads lower total system efficiency and produce additional heat.

In the white paper “Phase Balancing: The Last Few Inches of a High-Efficiency Power System”⁴, Server Technology analyzes the impact of unbalanced loads on a 30A, 240/415V 3-phase circuit loaded in a Wye configuration. If the load is balanced, the current through each input phase is 8A, and losses can be calculated at 19 watts per 100 feet of cable. In the most unbalanced case, current through one phase is 24A and losses escalate to 115 watts per 100 feet of cable. This shows that power loss and heat generated in feeder cables can increase by as much as 600 percent because of unbalanced loads. Additionally, keeping track of which base loads are on which phase can be time consuming and tedious depending on the sophistication of the site and the tools available. Load balancing issues are also specific to three-phase AC power and are not a concern with DC.

■ **Power-Related Heat Loads**

After the IT equipment itself, the cooling system is the next largest user of power in the data center. Not only is cooling required for the heat generated from the IT equipment, but it also must offset the heat generated by power conversion losses. With fewer required conversions and greater overall efficiencies, a DC power system generates less heat than an AC system, reducing data center cooling energy consumption. As a rule of thumb, each watt of heat generation removed from the data center leads to an additional 1.4 to 2 watts saved in cooling.

Availability

Both AC and DC power systems can be designed to achieve high availability. The highest availability facilities will likely continue to use 2N redundancy because of its ability to support a dual-bus architecture that can eliminate downtime from system failures across the entire power distribution chain. However, that availability comes with a cost, both in terms of upfront equipment and ongoing operating efficiencies. Many organizations can achieve desired levels of availability from N + 1 redundancy in which redundancy is maintained at the module level.

At the system level, there are two ways that DC power systems offer high availability. First, the DC power conversion system has fewer components than a comparable AC system, which contributes to a higher mean time between failure (MTBF) rate and more uptime. In addition, the DC UPS uses an array of discrete rectifiers to deliver conditioned, isolated power to a distribution bus. These rectifiers provide built-in redundancy; the system can accommodate the failure of any individual rectifier without immediately affecting performance or capacity. Individual rectifiers can be safely hot-swapped out in the field without impacting system operation, thus minimizing system mean-time-to-repair (MTTR), a major contributor to unavailability.

Scalability

With equipment and rack density rising steadily, the power system can become a constraint to growth. As a result, modular approaches to data center design and expansion are gaining in popularity as demonstrated by the trend toward “containerized” data centers.

With a complete data center power and cooling infrastructure integrated into a fully enclosed container, these systems can be delivered to a site and quickly plugged into power, communications and chilled water systems to enable additional computing capacity. With its built-in redundancy and compact footprint, a DC UPS is an excellent solution for these applications.

This same philosophy is also manifesting itself in the move toward row-based equipment, which enables modular expansion of existing facilities or support for high density rows. Both AC and DC UPS systems provide scalability in a row-based format, but the DC UPS has the benefit of built-in redundancy and will occupy at least 50 percent less floor space than a comparable row-based AC system.

Cost

The main purpose of the critical power system is to eliminate power-related downtime; any cost comparison should consider the level of availability required and the cost of downtime. In general, row-based DC power applications will be less expensive to install, operate and maintain than a comparable AC system and will support redundancy levels of at least N + 1. Some savings may be offset by increased IT equipment costs in a DC system—list prices for servers with DC power supplies can be up to 10 percent higher than the more common AC servers. These costs may be negotiable and should come down with the higher volumes that accompany increased adoption. Overall, a DC-based end-to-end power architecture can be as much as 30 percent less expensive than an AC system, depending on many of the factors above.

NetSure™ ITM Row-Based DC UPS

The NetSure™ ITM is a data center-optimized row-based 48V DC UPS from Emerson Network Power. It is designed to support the move to higher efficiency, density and flexibility enabled by row-based power and cooling. Instead of deploying one large DC power plant, which would require DC power to be distributed across the facility, 480V AC power is delivered directly to the NetSure™ ITM, which then provides protected power to racks of 48V equipment in close proximity [Figure 6]. When additional racks of IT and communications equipment are required for data center growth, new DC UPS units can be added without disruption.

The NetSure™ ITM integrates power conversion, battery backup and branch distribution into a compact, highly reliable power protection system. (12) individual AC/DC power conversion units (PCUs) deliver 70kW per module of conditioned, isolated 48V power to an internal distribution bus that charges the batteries and distributes power directly to the IT or communications equipment. Backup time is 5-10 minutes at full load with the included 10 year-VRLA batteries; traditional DC systems may be more appropriate if more than 60 minutes of backup is required.

The system provides up to 22 branch circuits protected by 100A to 200A circuit breakers. Each branch is monitored in the controller for load and trip condition and is designed to support varying rack densities. Power distribution in the rack is provided by NetSure™ RDB Series zero-U distribution units that feature plug-and-play connectors for easy installation and tool-less maintenance [Figure 7].




Figure 6. The NetSure™ ITM is positioned within the row to provide power to multiple racks.

This DC UPS comes pre-assembled and factory tested to simplify installation. The only field connections required are the AC input to the module and the output to DC loads. In addition, individual components are user replaceable and can be swapped out while the system is operating, allowing live maintenance without bypass or generator switchover.

The NetSure™ ITM features an integrated control and monitoring system that provides data on input and output parameters (including branch circuit loads), as well as advanced Albér battery management. An Energy Optimization Mode featuring Intelligent Power Matching ensures near-peak efficiency for loads down to 5 percent



Figure 7. A zero-U power distribution unit provides power to 48V input loads.

by keeping a minimum number of PCUs near full load and putting the rest on standby with the ability to respond quickly to changing load conditions.

Because it eliminates the need for large copper runs and offers the reliability and scalability inherent in DC power, the NetSure™ ITM is ideal for a variety of data center applications, including:

- Sites with an existing constrained power infrastructure or problems with harmonics
- Data centers seeking to optimize efficiency and availability while establishing a path for growth by using DC power in concert with existing AC systems
- Modular or pre-configured data center layouts
- High-density scalable blade deployments

Conclusion

Direct current is already a fundamental part of your IT infrastructure: critical loads consume DC power and all backup sources generate it. Furthermore, the electric grid distributes AC power. So, the question becomes, where is the optimal point at which to convert AC to DC power while providing suitable protection from outages? If the conversion occurs too early, DC power must be transported long distances, which requires large conductors to reduce losses. If it occurs too late, additional conversions are introduced into the process which may compromise efficiency and reliability and increase costs.

In many applications, the ideal point for power conversion and energy storage is as close as possible to the load. A row-based, 48V DC UPS allows for optimal efficiency, reliability and flexibility in these data center environments.

The NetSure™ ITM from Emerson Network Power, combined with the availability of DC-powered IT equipment from major manufacturers, makes 48V DC power a practical, efficient and cost-effective solution for managing the conversion from AC grid power to DC power required by IT equipment.

The simplicity, efficiency and reliability inherent in the NetSure™ ITM system can reduce complexity and costs in a variety of applications and should be seriously considered for mid-size data centers and facilities with increasing capacity needs.

References

- [1] "Reliability Field Data of Power Systems." NTT Facilities. INTELEC 2007.
- [2] "Quantitative Efficiency Analysis of Power Distribution Configurations for Data Centers." The Green Grid. 2008.
- [3] Pratt, Annabelle et al. "Evaluation of 400 VDC Distribution in Telco and Data Centers to Improve Energy Efficiency." Intel. 2007.
- [4] "Phase Balancing: The Last Few Inches of a High-Efficiency Power System." Server Technology. 2010.

Emerson (NYSE:EMR), based in St. Louis, Missouri (USA), is a global leader in bringing technology and engineering together to provide innovative solutions for customers in industrial, commercial, and consumer markets through its network power, process management, industrial automation, climate technologies, and appliance and tools businesses. For more information, visit: Emerson.com.

Emerson Network Power, a business of Emerson (NYSE:EMR), is the global leader in enabling *Business-Critical Continuity*™ from grid to chip for telecommunication networks, data centers, health care and industrial facilities. Emerson Network Power provides innovative solutions and expertise in areas including AC and DC power, precision cooling, embedded computing and power, integrated racks and enclosures, power switching and controls, infrastructure management, and connectivity. All solutions are supported globally by local Emerson Network Power service technicians. For more information on Emerson Network Power's full suite of solutions specifically supporting the communications network infrastructure, visit: EmersonNetworkPower.com/EnergySystems.

Learn more about the NetSure™ ITM DC UPS from Emerson at: EmersonNetworkPower.com/NetSureITM.

This publication is issued to provide outline information only which (unless agreed by Emerson Network Power Energy Systems, North America, Inc. in writing) may not be used, applied or reproduced for any purpose or form part of any order or contract or be regarded as a representation relating to the products or services concerned. Emerson Network Power Energy Systems, North America, Inc. reserves the right to alter without notice the specification, design or conditions of supply of any product or service.

Emerson®, Emerson Network Power™, Business-Critical Continuity™ and NetSure™ are trademarks of Emerson Electric Co. and/or one of its subsidiaries.

Additionally, this brochure references the following company trademarks: NTT Facilities, The Green Grid, Intel and Server Technologies. Names of companies and products are trademarks or registered trademarks of the respective companies. Any questions regarding usage of trademark names should be directed to the original manufacturer.

Emerson Network Power

Energy Systems, North America

4350 Weaver Parkway, Warrenton, IL 60555

Toll Free: 800-800-1280 (USA and Canada)

Telephone: 440-246-6999 **Fax:** 440-246-4876

Web: EmersonNetworkPower.com/EnergySystems

Emerson Network Power.

The global leader in enabling *Business-Critical Continuity*™.

- | | | | |
|-------------------|--|------------------------------|-------------------------------|
| ■ AC Power | ■ Embedded Computing | ■ Outside Plant | ■ Racks & Integrated Cabinets |
| ■ Connectivity | ■ Embedded Power | ■ Power Switching & Controls | ■ Services |
| ■ DC Power | ■ Infrastructure Management & Monitoring | ■ Precision Cooling | ■ Surge Protection |