

Increasing Data Center Efficiency by Using Improved High Density Power Distribution

By Neil Rasmussen

White Paper #128



Executive Summary


A new approach to power distribution for high density server installations saves floor space, simplifies power cabling, saves capital cost, reduces weight, and increases electrical efficiency. This paper describes this distribution architecture and quantifies the benefits.

Note: The methods in this paper only apply in North America and are for problems that are unique to North America

Introduction

The standard power distribution system for large data centers in North America is a 277/480V three-phase power system supplying distributed Power Distribution Units (PDUs) which convert the voltage to the 208V and 120V single-phase branch circuits utilized by IT equipment. This arrangement is represented by the one-line diagram of **Figure 1**.

Figure 1 – One-line diagram showing standard power distribution system for North American data centers



High density server installations where racks can draw from 10 kW to 40 kW per rack place significant strains on this standard power distribution system. This increased rack power density is causing:

- Multiple branch circuits per rack, with corresponding breakers and cabling
- PDUs to become up to 30% of the floor space
- PDUs to become up to 30% of the floor weight loading
- More PDUs installed per IT rack, driving up capital costs

However, these problems are unique to North America. In Europe and the rest of the world the power distribution system is simpler and more efficient. This paper shows that it is practical to take the power distribution system used in the rest of the world and apply it to North American data centers to reduce cost, reduce weight, save floor space, and increase electrical efficiency.


In the past it was not practical to consider alternative higher voltage power distribution systems because a significant fraction of IT equipment operated from 120V. However in the modern high density data center virtually all rack-based equipment operates on 208V (or 240V outside North America). Even devices that can run on 120V are also capable of operating on higher voltage.

Virtually all IT equipment manufactured today is designed for worldwide compatibility. This means it can operate on the North American 120/208 volt system, the Japanese 100/200 volt system, and the 240 V system used in the rest of the world.

Alternative Power Distribution for North America

The alternate power distribution system proposed and described in this paper is an adaptation of the power distribution system used in the rest of the world. The majority of countries around the world use 230 VAC line to neutral. This paper proposes increasing the voltage to 240 VAC line to neutral since this is the highest voltage supported by most IT equipment. This arrangement is represented by the one-line diagram of **Figure 2**.

Figure 2 – One-line diagram showing alternative power distribution system for North American data centers



The differences between this diagram and the diagram of **Figure 1** are subtle, and include:


- Three-phase branch circuit power distribution that is 240 V (line to neutral) instead of 208 V (line to line)
- Elimination of the PDU transformers
- A transformer that is used to change 277/480 V to 240/415 V at the output of the UPS.

These items are explained in more detail in the following sections.

Three-phase branch power distribution – 240V (line to neutral) instead of 208V (line to line)

While the proposed power distribution method sounds like a minor change, it actually creates a very large impact on the power distribution system. First, the typical North American three-phase branch power distribution is at 120/208 Volts, while the proposed distribution uses 240/415 V. This is shown in **Figure 3**.

Figure 3 – Diagram showing the difference between three-phase 208 V distribution and 240 V distribution



Servers draw their power from the 208 V high voltage (line to line) connection in North America, while they draw their power from the 230 V low voltage (line to neutral) connection in almost every other country in the world. It is important to recognize that when comparing the “line to neutral” voltages of both distribution methods, the comparison should be between 120 V and 240 V and NOT between 208 V and 240 V. This “line to line” and “line to neutral” difference becomes readily apparent when the power capacity for a three-phase branch circuit is calculated. For example, assume 20 amp circuits are provided to the load in either case. The power capacity for the 120 V “line to neutral” distribution method is calculated as $(20 \text{ amps} \times 120 \text{ V} \times 3 = 7.2 \text{ kW})$, while the capacity for the 240 V “line to neutral” distribution method is calculated as $(20 \text{ amps} \times 240 \text{ V} \times 3 = 14.4 \text{ kW})$. Given the same circuit current rating, the 240 V distribution provides 100% more power than the 120 V distribution. This increases the power density capability per rack without adding extra circuit breakers as would be the case with the 120/208 V distribution. These additional breakers present additional points of failure resulting in decreased data center reliability.

Consider now the same comparison except this time assuming a fixed branch circuit power capacity. For example, assume 10 kW is available to the load in either case. The power capacity for the 120 V “line to neutral” distribution is calculated as $(27.7 \text{ amps} \times 120 \text{ V} \times 3 = 10 \text{ kW})$, while the capacity for the 240 V “line to neutral” distribution is calculated as $(13.9 \text{ amps} \times 240 \text{ V} \times 3 = 10 \text{ kW})$. Given the same power, the 240 V / 415V distribution method provides the same power capacity with nearly **half the current** required by the 120

V / 208V distribution method. This significant difference in current requires the branch circuit wires in the traditional North American system to be 300% larger in terms of their size and weight.

Elimination of the PDU transformer

Even more important than the cost, size, and weight savings in wiring is the savings in transformers. **The PDU transformers can be eliminated.** The PDU transformers are required in North American to change the voltage from the UPS output of 277/480 down to the 120/208 utilization voltage. In the rest of the world, the UPS output of 230V/400 V can be directly utilized by the IT loads without the need for an additional transformer. Similarly, the proposed North American distribution eliminates distribution transformers and supplies the IT loads with 240V.

In a North American data center with 30 kW racks, approximately 20% of space that would be available for racks is consumed by PDUs. In addition, the PDUs would make up more that 25% of the weight on the raised floor. Note that these impacts were much lower in data centers of the past which operated at approximately 2 kW per rack. **Therefore, the benefit of using this new power distribution approach increases dramatically as the rack power density increases.**

Transformer changes UPS output from 277/480 to 240/415

An additional transformer is used to change the UPS output voltage from 277/480V to 240/415V in the proposed new method. While it appears that this transformer replaces the PDU transformers, in fact it is 90% smaller and less expensive than the PDU transformers for the following reasons:

- It is wired as an autotransformer¹, and is therefore less than 20% of the size and weight of a standard transformer of the rated UPS capacity.
- It is sized to the UPS capacity, in contrast to the PDU transformers which as a group are typically sized at 1.5 to 3 times the UPS capacity.

In addition, this transformer can be located outside of the data center server space, such as in a back room, which further conserves valuable data center space.

The autotransformer described in this new power distribution approach for North America can be eliminated in some cases, resulting in further savings. For new installations the facility could be directly wired for 415V input from the utility company, and a 415V UPS can be used. If this is done, then the auto transformer is not needed. The disadvantage of this approach is that various related AC powered devices, such as pumps, generators, and switchgear are not readily available in North America at a 415V rating. Therefore at this time the preferred approach is to use the auto transformer as described in this paper.

¹ An autotransformer has only one winding (primary) as opposed to two in an isolation transformer (primary and secondary).

Performance Comparison

The alternate power distribution system proposed and described in this paper can save a significant amount of capital cost, space, and operating cost. A comparison of conventional power distribution design with the alternative design was performed based on an actual customer's data center. Their data center is 2,175 ft² and consists of 2 redundant modular UPSs at 600 kW each. A total of 104 racks range in density from 1.6 to 6.5 kW each. The results of this comparison are provided in **Table 1**.

Table 1 – Performance comparison for alternative power distribution in a 600 kW high density data center in North America

Characteristic	Standard 120/208 V System	Alternative 240/415 V System	Comment
UPS operating voltage	480 V	480 V	Same UPS system
UPS output transformer	None	Equivalent 164 kVA	A 82 kVA autotransformer is required for each side of the architecture.
PDU transformer	Equivalent 1.8 MW	None	8 x 225 kVA PDU transformers (4 on each side). Typically oversized by at least 1.6 and sometimes as high at 3.5 times the UPS capacity.
PDU / RDP / Transformer costs	\$145,091	\$71,773	The PDU in the standard system is replaced with autotransformers and rack based remote distribution panels (RDP) in the alternative system for a 51% cost savings.
PDU / RDP / Transformer weight	20,800lb	2,090 lb	Two auto transformers are used and a total of 16 RDPs are used for a weight savings of 90%.
PDU / RDP / Transformer area	135 sq ft	92 sq ft	Area savings of 32% which includes access areas.
Copper wire weight	5,658 lb	2,700 lb	Copper only. Estimated based on 104 racks with varying power densities of 1.6 to 6.5 kW each, fed by 8 x 225 kVA PDUs or 16 RDPs. Average whip length 42 ft each. Average of 29 whips per PDU and 15 per RDP. (52% weight savings)
Total distribution power loss	8,894 W	801 W	Includes power needed to cool the transformers and distribution copper losses. (Based on 500 watts of power required by the air conditioning system to remove 1 kW of heat.) (91% savings)
120 V support	Y	N	The alternative system requires a transformer for the few 120V devices that cannot operate on 240 V
Power distribution 10 year TCO	\$262,381	\$102,631	Based on 30% load (15% on each redundant UPS) (61% savings)

Note: Blue shading indicates best performance for the characteristic

Note the large savings in TCO and the large savings in space of the alternative approach of 61% and 32% respectively. The 10 year TCO savings by distributing at 240 volts is \$159,751 or 61% savings. Of the TCO savings approximately 47% is due to material cost savings and 53% is due to energy cost savings.

Special Cases

There are some special cases where the alternative power distribution system is perceived to have limitations. These cases and how the alternative distribution system addresses them are described here.

Occasional 120V loads

Data centers in North America encounter occasional 120 V loads, even in high density areas. These are dealt with as follows:

- Many 120 V loads are actually 120/240 V rated. These ratings are always on the devices. The vast majority of these devices have a detachable power cord. The solution is simply to change the power cord to a universal IEC C-13 or C-19 power cord. In some cases it may be necessary to change a switch setting on the device.
- In cases where it is not possible to avoid using 120 V for isolated device(s), a local rack-mount transformer can be used to step down 240 V to 120 V. These are available from APC and other suppliers.

Isolation and grounding

The conventional North American power distribution system provides isolation of the neutral wire at each power distribution unit. This characteristic has historically been considered desirable and is given various names including: galvanic isolation, separately derived source, and regenerated neutral. Many of the claims regarding the benefit of this arrangement are not based on science. In fact there is no data to suggest that data centers in North America are more reliable than data centers outside of North America, most of which operate according to principles outlined in this paper. This subject is discussed in more detail in APC White Paper #8, "[Inter-System Ground Noise: Causes and Effects](#)" and APC White Paper #87, "[Use of the Signal Reference Grid in Data Centers](#)".

Existing installations

The alternative system can be used alongside existing standard power distribution and is particularly well suited for data center expansions focused on high density server applications.

Connectors

The appropriate connectors for use with 240V circuits are IEC C13 and C19 types. This is the power cord connector provided by most OEM server manufacturers with rack-mount servers and storage devices. Therefore, most high density servers are already provided with the appropriate connector for use with a

240V system. Occasionally, a device may be provided with a power cord that uses a North American NEMA style twist lock connector. These are dealt with as follows:

- If the power cord is the detachable type with a IEC C13 or C19 connector on the chassis, then the cord should be replaced with a cord with IEC connectors at both ends. These are available from virtually all IT OEMs, and they are available from suppliers such as APC. Note that when substituting power cords it may be advisable to purchase a cord of shorter than standard length in order to simplify power cabling in the rack.
- In the case where the cord is permanently attached to the IT device, if the plug is 20A or less an adapter cable can be used.
- In the case where the connector is a three-phase connector, see the following section

Circuit breakers

The branch circuit breakers used in the system described in this paper operate at higher voltage than normal branch circuits in North American data centers. Most circuit breakers that are currently used in North American data centers are not rated for this voltage and cannot be used. Conversely, most European circuit breakers panels are not certified by UL for use in North America. Recently, manufacturers have introduced compact worldwide circuit breakers and circuit breaker panel boards that are rated and certified for high voltage use in North America. Companies like APC offer complete data-center PDUs, remote power panels, and rack power panels that are appropriately rated and certified for high voltage distribution in North America.

Three-phase IT loads

A very few IT loads are equipped with three-phase connectors. These include some Compaq blade servers, large EMC storage units, and mainframe-style servers. The key thing to understand here is that none of these devices actually require three-phase power. In fact all such devices actually use multiple power supplies drawing single phase power. Furthermore, all of these devices have versions that are sold outside North American and therefore are compatible with the alternative power distribution system. These devices are dealt with as follows:

- For large deployments of devices with three-phase connectors, specify to the supplier that the power input is to be wired for European voltage. This is often a simple internal jumper setting.
- Determine if the device is field re-wirable for European voltage operation. Many large devices are.
- If it is impractical to change the device or convert it, then consider using standard distribution to this device, or install a dedicated PDU to convert the voltage for this device.

The above information is provided to assist with implementation of the alternative power distribution system in special cases. Note that in the vast majority of cases, no modifications of any kind are needed to the IT

equipment, the power cords, or the rack power distribution devices. Devices which are known to work seamlessly in the alternative power system include: IBM blade servers, 1U servers, virtually all enterprise rack mount servers, and virtually all rack mount SAN and NAS storage devices.

Comparison With DC Distribution Alternatives

From time to time DC power distribution, either at low voltage or high voltage, has been proposed as an alternative to AC power distribution. The hypothetical benefits of a DC system are claimed to include substantial increases in electrical efficiency. The system proposed in this paper achieves many of the same benefits with much lower cost and complexity. Therefore, the high voltage AC distribution system is a preferred strategy to improve data center efficiency. This subject is discussed in more detail in APC White Paper #127, "[A Quantitative Comparison of High Efficiency AC vs. DC Power Distribution for Data Centers](#)".

Conclusion

There are significant advantages to using a higher voltage power distribution system in high density data centers in North America. Use of the proposed 240/415 V distribution system instead of the standard 120/208 system can save 61% in the lifetime cost of the distribution system, and save floor space and weight loading. The benefits are most apparent for high density installations. Furthermore, the alternative design can operate alongside conventional power distribution designs in existing data centers.

About the Author:

Neil Rasmussen is a founder and the Chief Technical Officer of American Power Conversion. At APC, Neil directs the world's largest R&D budget devoted to power, cooling, and rack infrastructure for critical networks, with principal product development centers in Massachusetts, Missouri, Denmark, Rhode Island, Taiwan, and Ireland. Neil is currently leading the effort at APC to develop modular scalable data center solutions.

Prior to founding APC in 1981, Neil received his Bachelors and Masters degrees from MIT in electrical engineering where he did his thesis on the analysis of a 200MW power supply for a Tokamak Fusion reactor. From 1979 to 1981 he worked at MIT Lincoln Laboratories on flywheel energy storage systems and solar electric power systems.