

A Scalable, Reconfigurable, and Efficient Data Center Power Distribution Architecture

White Paper 129

Revision 1

by Neil Rasmussen

> Executive summary

Significant improvements in efficiency, power density, power monitoring, and reconfigurability have been achieved in data center power distribution. The past techniques of massive transformer-based power distribution units feeding under-floor hardwired circuits via rigid or flexible conduit are shown to be obsolete. This paper explains some the newer approaches to power distribution including modular power distribution and overhead power busway, and shows their advantages when compared to the legacy approach.

Contents

Click on a section to jump to it

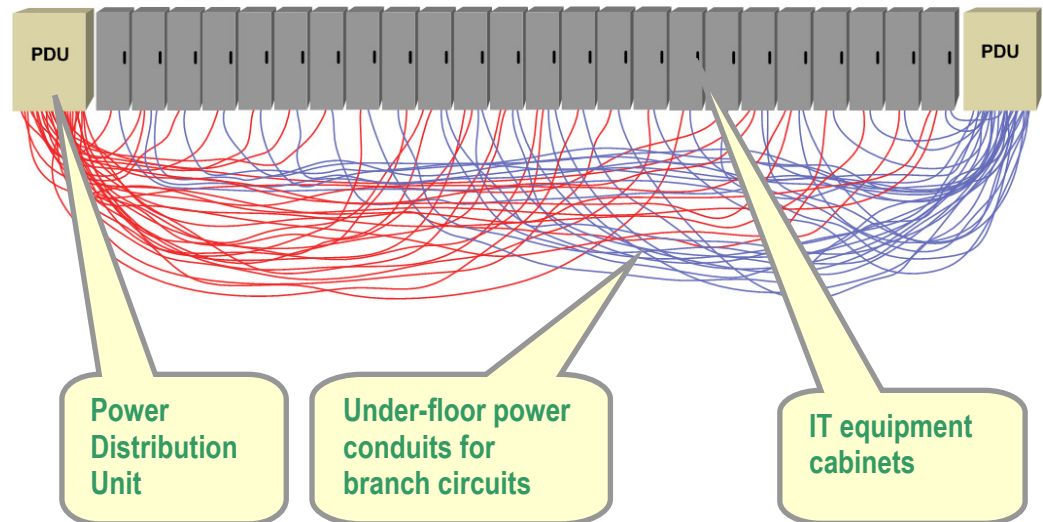
Introduction	2
Background	3
An optimized power distribution system	4
Alternative power distribution approaches	11
Conclusion	14
Resources	15

Introduction

Most existing data centers utilize the same power distribution architecture that was developed for data centers approximately 40 years ago. This system is illustrated in **Figure 1**.

Figure 1

Wiring of a traditional data center power center distribution system



In the traditional system, the main data center power is distributed to power distribution units (PDU - typically rated from 50 kW to 500 kW). The power distribution units may contain large power transformers to convert voltage or provide power conditioning. The power distribution units in turn distribute a quantity of branch circuits (typically rated from 1.5 kW to 15 kW) to the IT equipment. Each IT enclosure uses one or more branch circuits. The wiring to the IT enclosures is usually required to be in flexible or rigid conduit, typically located beneath the raised floor, as illustrated in **Figure 2**.

Figure 2

Under-floor power distribution in hard pipes



Since the introduction of this system, there have been dramatic changes in how power is utilized in data centers that have challenged this architecture, specifically driven by increasing power density, the increasing number of separate IT devices within the data center, and the need to add and remove IT devices on a continuous basis.

This paper explains why the evolution of the data center has made the historical power distribution architecture obsolete, and describes a more effective power distribution system that can be implemented today. The improved system allows IT racks and even whole PDUs to be installed or changed without any new wiring, distributes power overhead, supports racks up to 30 kW with a single flexible power feed, improves electrical efficiency, reduces copper consumption, is instrumented for power at the branch circuit, and has a standard capacity management system.

Background

When the traditional power distribution system was created, data centers consisted of a small number of large IT devices that were rarely changed except during the scheduled downtime of a major IT upgrade. The low power density of these facilities required a low volume of under-floor air, and typically less than one branch circuit for every three square meters of computer room. The modern data center has different characteristics which have challenged the traditional power distribution architecture:

- Instead of a few large IT devices, data centers may contain thousands of plug-in devices with separate power cords, requiring many more power receptacles
- IT devices within a rack enclosure are changed often within the lifetime of the data center, changing the power requirements or receptacle requirements at a rack location
- Due to changing power requirements, new power circuits must frequently be added to a live data center without disturbing nearby existing IT loads
- The per-rack power density has increased greatly, often requiring multiple branch circuits per cabinet
- The number of power feeds has clogged the under-floor air plenum with conduits, blocking the airflow and making changes very difficult
- The number of IT devices that may be connected to a branch circuit breaker is often much greater than one, making it difficult to size branch circuits or determine impending overload conditions
- Dual power path systems are commonly implemented, requiring assurance that no circuit is loaded above 50%, yet there is typically no mechanism to monitor or plan this

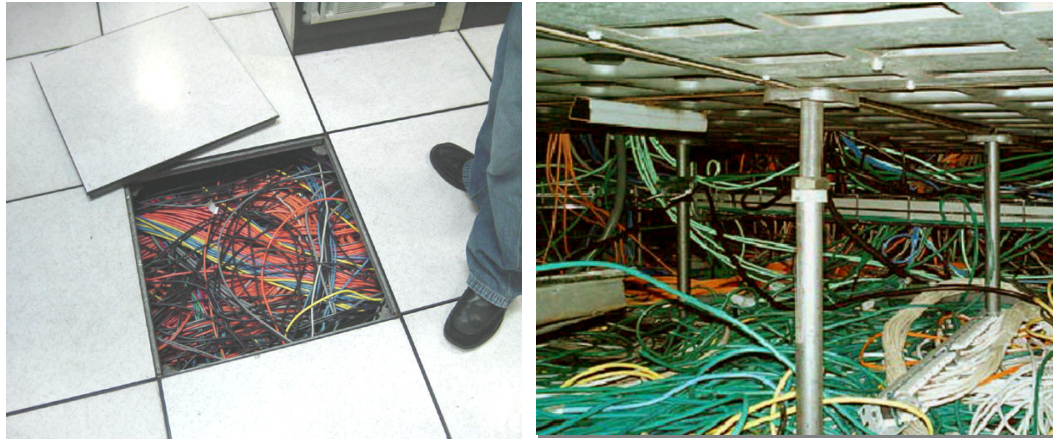
Although these issues are widely recognized, and there are a variety of products available to solve these problems, most data centers built today still use traditional methods and the result is that even newly built data centers suffer from the following unfavorable conditions:

- Data center operators are forced to make circuit changes on energized wiring (“hot work”)
- Data center operators cannot tell what branch circuits are near overload, or what circuits might overload in case of the loss of a power path
- Under-floor cooling plenums are blocked with cables, preventing the high volume air required by modern IT equipment
- Data center operators discover that the PDU units are taking up a significant fraction of the floor space and the floor load weight capacity
- Large transformer based PDU units cannot be fully utilized because they run out of branch circuits
- Large transformer based PDU units generate waste heat that must be cooled, decreasing data center efficiency

A few photographs of wiring conditions in actual operating data centers illustrate some of the above difficulties and are shown in **Figure 3**.

Figure 3

Wiring congestion in actual data centers



The power distribution architecture described in the next section addresses all of the problems described above.

An optimized power distribution system

An ideal power distribution system would have the following attributes:

- New circuits can be added or changed on a live system safely
- No under-floor cables needed
- All circuits monitored for power
- All breaker status remotely monitored
- IT zones and associated power distribution can be deployed over time
- All power levels supported using only a single cable to the IT enclosure
- Receptacle types can be changed at the IT enclosure by IT personnel
- Capacity and redundancy are managed on every circuit
- No excess copper is installed that is not needed
- High efficiency

Power distribution systems have been evolving in response to the needs of the modern data center, and various improvements have been introduced to the power distribution system over time, most notably:

- Branch circuit power metering
- Overhead cable tray with flexible power cords
- Overhead fixed busway with removable power taps
- High power, pluggable rack power distribution units
- Transformerless Power Distribution Units
- Power capacity management software

The power distribution system shown in **Figure 4** includes all of these elements in an architecture that is ideally suited for the modern high density data center.

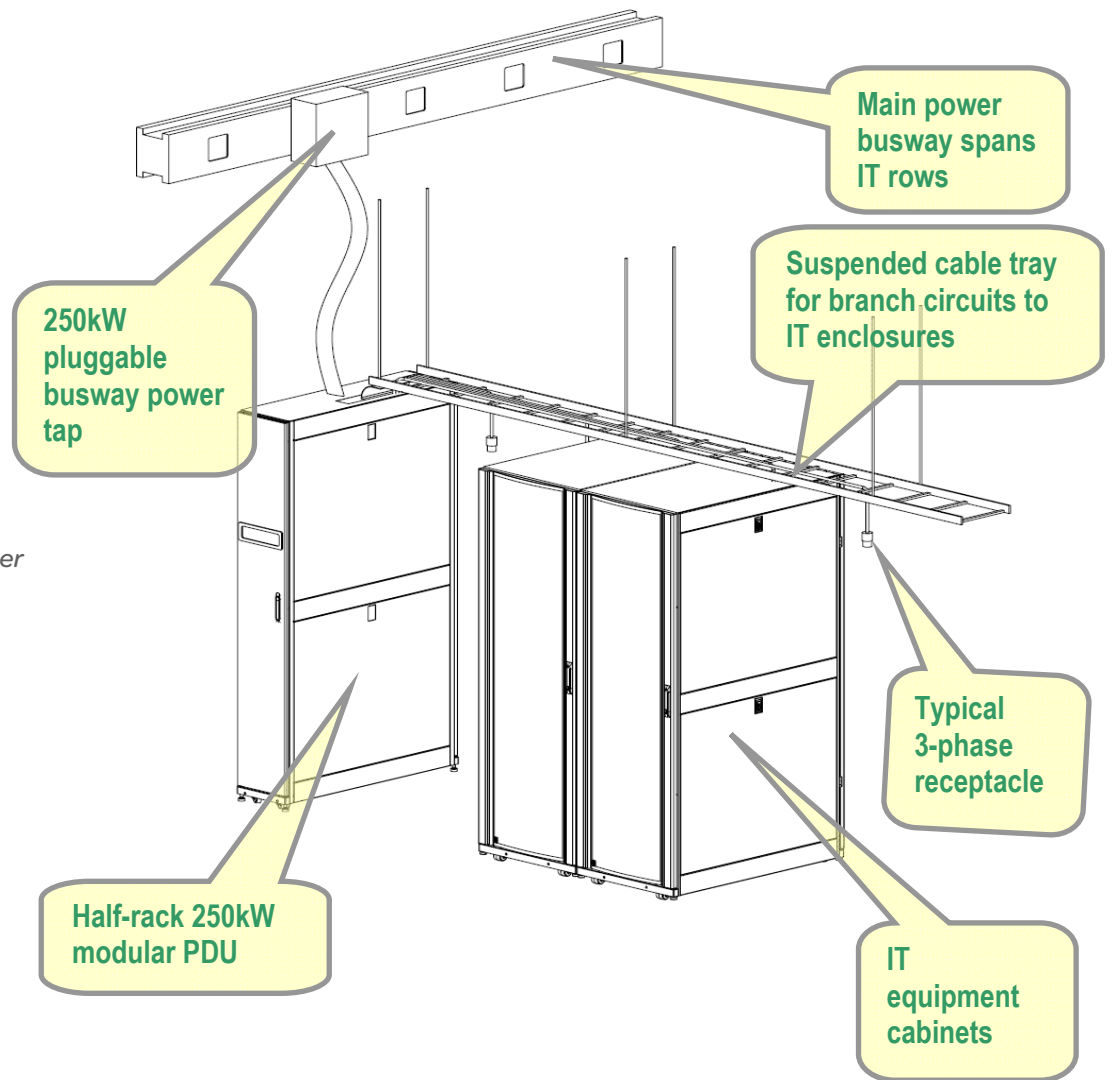


Figure 4

Illustration of modular power distribution system

Description of the system

The distribution system has two steps. For larger data centers the main critical bus power from the uninterruptible power supply (UPS) is distributed to IT rows using one or more overhead busways as shown at the top of **Figure 4**. The busways are installed up front and traverse the entire planned IT rack layout. When a group of racks is to be installed, a low-footprint modular PDU is installed at the same time and plugged into the overhead busway. The connection to the busway is also shown in **Figure 4**.

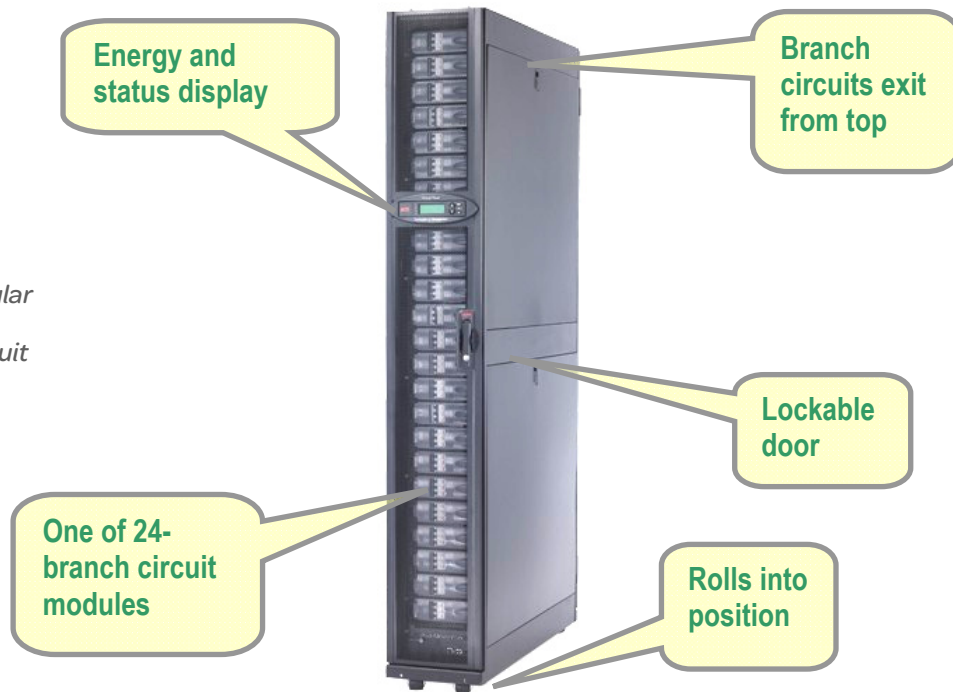


Figure 5

Example of a 250 kW modular PDU with a small footprint. Shown with 24-branch circuit modules installed

Instead of traditional circuit breaker panels with raw wire terminations, the modular PDU has a backplane into which pre-terminated shock-safe circuit breaker modules are installed. This arrangement allows the face of the PDU to be much narrower, and eliminates on-site termination of wires.

The modular PDU initially has no branch circuit modules installed. The power circuits from the modular PDU to the IT racks are flexible cable that are plugged into the front of the modular PDU on site to meet the requirements of each specific rack as needed. The branch circuit cables to the IT enclosures are pre-terminated with breaker modules that plug into the shock-safe backplane of the modular PDU. An example of a branch circuit module is shown in **Figure 6**.

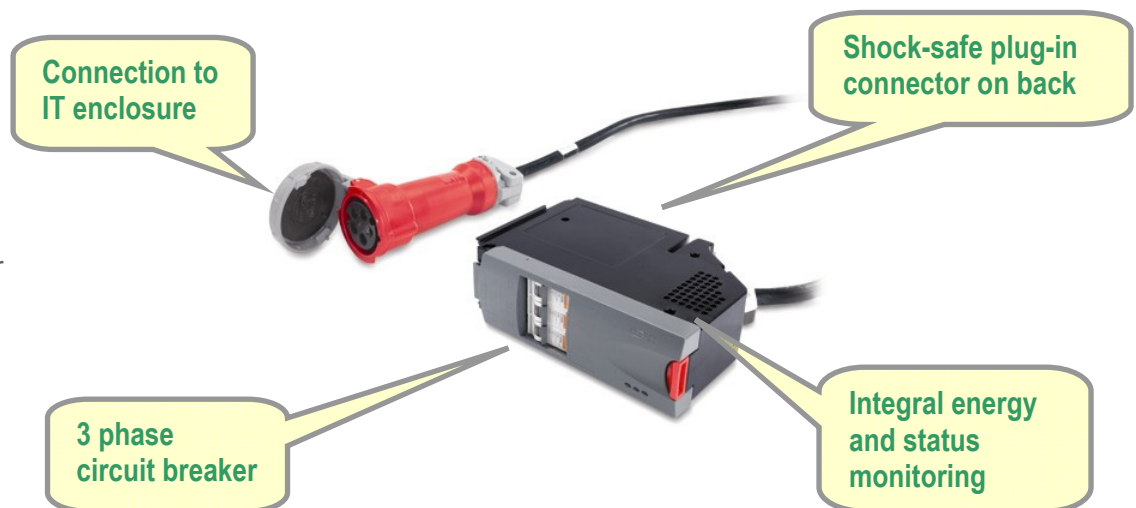


Figure 5

Branch circuit module that plugs into the modular power distribution unit

For equipment that requires a dedicated branch circuit, such as most blade servers, a single cable from the PDU carries one, two, or three branch circuits that plug directly into the blade server, with no additional rack PDU (i.e. power strip) required. When mixed equipment exists in the rack, rack PDUs are available that provide interchangeable receptacle and current ratings. A typical rack PDU is shown in **Figure 7**; the connector shown in this figure plugs into the mating connectors shown in the previous **Figure 6**.

Figure 7

Example of a 12 kW 3-phase rack outlet strip designed for vertical mounting in the rear of an IT enclosure



In this system, a PDU for a new row of IT enclosures, along with all of the associated branch circuit wiring and rack outlet strips, can be installed in an hour, without any wire cutting or terminations.

Small zones or very high density

In some cases there may be one or more zones within a data center where only a small number of branch circuits is needed. This can occur when there is a cluster of very high density racks, or when a small group of racks is isolated by room shape or other constraints. In these situations the full 24-branch circuit capability of the standard modular distribution system is not needed.

In these cases, the architecture includes a smaller version of the modular power distribution unit that can directly mount into an IT rack, consuming zero floor footprint and supporting up to 6 branch circuits. This PDU includes all of the status and energy monitoring capability of the larger floor standing unit in a 5U rack mount version.

Figure 8

A small modular PDU that mounts directly into an IT enclosure that can provide power to 6 IT enclosures (door closed and door open views)



Smaller data centers

The system in **Figure 4** is optimized for **larger** data centers with a large number of rows of IT racks, which may be deployed over time. The overhead busway is much easier to install than high-power wiring and allows PDUs to be added or changed. For smaller data centers of 300 kW or less, a simplified approach using the same components and same principles can be used.

For smaller data centers when the number of PDUs is restricted to one or two, the flexibility of deploying PDUs in phases is often unnecessary, and it may be cost effective to directly wire the modular PDU (**Figure 5**) to the critical bus with traditional conduit and wire. In this case the overhead busway is omitted. For very small data centers or data centers with irregular floor layouts the small modular PDU described in the previous section and illustrated in **Figure 8** above may be appropriate.

A further simplification is also possible in smaller data centers, where the modular PDU may be directly integrated with the UPS system into a compact arrangement that can be located in the IT room and integrated into an IT enclosure lineup. In this case the main bus wiring is eliminated and the need for a separate power room is eliminated. This convenient and popular approach for data centers of 200 kW or below is illustrated in **Figure 9**.

Figure 9

Modular power distribution unit integrated with a UPS for smaller data centers



Retrofit / upgrade applications

A large number of data center projects involve the upgrade of an existing data center, with common projects being the addition of capacity or installation of a high-density zone. The modular distribution system is particularly well suited to these types of retrofit projects, because installation is much less disruptive than installing a traditional PDU. There is a long list of challenges when installing a new traditional PDU in an existing data center, most of which are addressed by the design of the modular distribution system.

As a data center evolves, modular PDUs can operate alongside existing traditional PDUs. In these retrofit applications where traditional PDUs are often placed based on historical constraints, the overhead busway component of the architecture is not used and traditional pipe and wire is used to connect each PDU to the main bus.

One of the most important and overlooked benefits of the modular PDUs in the upgrade of a traditional data center is the fact that the installation adds no additional under-floor airflow blockages since the cables are run in overhead cable tray. This is often crucial in existing data centers where the floors are not deep and under-floor airflow is already a limiting factor in the cooling performance – and even in the overall electrical efficiency – of the data center.

Power and status monitoring

In a data center power distribution system there can be hundreds of circuit breakers that can be overloaded. This optimized distribution system uses higher capacity rack power feeds and uses 20-40% fewer breakers than a typical system, but there are still many circuits, which exist at 4 levels:

- UPS main bus
- PDU input
- Branch circuit

- Outlet

In the modular power distribution system, there is built-in current and energy monitoring on every circuit at all levels of the hierarchy (outlet-level monitoring is optional in some configurations). In addition, the branch circuit breakers in the PDU are monitored for status. All of the monitoring communicates via simple network management protocol (SNMP) open standards protocol. Capacity management software is used to monitor every circuit in the system and enforce safety margins, verify redundancy, and identify available capacity.

Voltage configuration

The architecture described in this paper is global and suitable for all data center operating voltages. However, in North America there are two operating voltage configuration options. The most electrically efficient IT operating voltage for North America is 415/240 V AC 3-phase power. This is the same distribution system used in Europe and most of the world, but is not commonly used in North America. The second option in North America is to use the traditional 208/120 V AC 3-phase power common in North American buildings. This option requires more copper and PDU units with large integral power transformers, and is less efficient and more costly.

For data centers in North America, the 415/240 V AC system has a number of important advantages, as explained in White Paper 128, *Increasing Data Center Efficiency by Using Improved High-Density Power Distribution*. It is the recommended voltage choice. However, the recommended power distribution system described in this paper can also be implemented in a traditional North American 208/120 V AC voltage configuration.



Link to resource
White Paper 128

Increasing Data Center Efficiency by Using Improved High-Density Power Distribution

Voltage configuration

The power distribution system described has been developed to address the many significant shortcomings of the traditional approach. The advantages of the new system are summarized in **Table 1**.

Table 1

Comparison of traditional power distribution with modular power distribution

Issue	Traditional distribution	Modular distribution
Adding branch circuits	New conduits must be installed, breakers sized and installed, and wires pulled and terminated. If system is live, electrician may need to work on live exposed electrical wiring. If circuit has power monitoring, new sensors and/or programming is typically required	Plug-in, pre-made branch circuits. Can be installed in an energized system without exposure to live electrical wiring. Power monitoring included in each branch circuit and automatically configures when plugged in
Removing branch circuits	If system is live, electrician may need to work on exposed electrical wiring. Conduit needs to be extracted from complex network of under-floor wiring. If circuit has power monitoring, programming may need to be changed	Branch circuit is unplugged from PDU, and may be re-used in another location
Ability to remove or change a rack cabinet	Branch circuit may need to be mechanically and/or electrically disconnected from the rack	Branch circuit unplugs at the rack, rack can be rolled out
Simplified planning	Typically the number and location of PDUs must be established early in the design, often before the final power density is known. In a raised floor installation, special under-floor mounting supports must be engineered	The number and location of PDUs do not need to be established up front. PDUs can be added later without special preparation
Reliability	Many wiring terminations are made in the field, resulting in loose connections and other defects. Mistakes during hot work can dislodge wires and trip breakers, disrupting other IT loads	Wiring terminations are pre-made in a controlled factory environment, improving reliability. Risk of interfering with other circuits during adds and changes is eliminated
Minimal floor footprint	Transformer-based PDUs consume approximately 2.5 m ² for every 100 kW of IT load, or approximately 7% of computer room space	Consumes approximately 0.7 m ² for every 100 kW of IT load, or approximately 2% of computer room space
Safety	Adding, removing, inspecting, and handheld current monitoring of branch circuits exposes operator to live electrical wiring	Shock-safe plug-in installation of branch circuits. No field wiring required
Airflow interference	Masses of cabling entering PDU from under-floor cause interference with under-floor airflow. Floor openings for wire create large bypass air paths that reduce air conditioner efficiency and capacity	No under-floor cabling. No floor air leakage caused by extra openings in the floor
Simplified engineering	Complex tradeoffs between row length, density, ampacity, and cost must be made in each installation in order to optimally locate the PDUs	Select from standard reference designs to meet requirements. Many decisions can be deferred to later phases
Up-front cost	PDUs are typically installed up front. Installation labor is significant. Built-out cost is incurred in the initial build	Most of the power distribution cost is incurred only when and if needed
Efficiency	Higher losses due to longer cable lengths. In North America 208/120 V AC nearly 10 times the amount of losses compared to 415/240 V AC system	Shorter cable lengths results in slightly lower losses

Alternative power distribution approaches

The power distribution architecture described in this paper is not the only approach to solving the problems of the traditional distribution system. There are two other variations that have been described in the literature and have been used in real data centers; they are busway to the rack, and DC power distribution. These two alternatives are briefly described and compared with the system described in this paper.

Busway to the rack

Until the architecture described in this paper was introduced in 2008, using overhead busway to the rack was the best alternative approach to delivering power to IT enclosures. In the busway to the rack system, the IT enclosures directly connect into the overhead busway via breaker boxes as shown in **Figure 10**.

Figure 10

Busway to the rack showing power taps attached to an overhead power bus



The busway is pre-installed over all IT equipment rows. This solves a number of the problems with traditional distribution, making changes easier and removing under-floor cabling. Busway was the first alternative to traditional distribution that achieved a flexible, reconfigurable distribution system. While busway to the rack remains a viable alternative to traditional approaches, there are a number of practical disadvantages of busway to the rack that are overcome by the new modular PDU architecture described in this paper. Busway to the rack is compared to modular distribution in **Table 2**.

Table 2

Comparison of busway to the rack with modular power distribution to IT enclosures (best performance highlighted in light green)

Issue	Busway to the rack	Modular distribution
Ability to handle mixed and changing power density	Busway must be sized in advance to the maximum density and capacity otherwise adding an additional busway in the future is disruptive and impractical	Power density is adjustable to suit the current configuration by adding or swapping branch circuits. Easier to install additional PDUs for extra capacity
Ability to handle specialized room layouts	Busway must be installed in advance over all expected enclosure locations	Flexible cable easily adapts to room obstructions, specialized IT cabinets, and constrained IT equipment floor plans
Safe & secure access to circuit breakers	Breakers are mounted on the overhead busway, requiring ladder access. This is not permitted due to local codes in many cases. Chain or other actuators may be required	All branch protectors behind a lockable door in one easily accessible location
Minimal floor footprint	Zero floor footprint	Consumes approximately 0.7 sq M for every 100 kW of IT load, or approximately 2% of computer room space
Standardized global solution	Regulations for busway vary geographically, requiring different physical configurations or current ratings or data communications in different locations	Standard architecture meets all international regulations and has a global standard for monitoring
Energy monitoring of each rack branch circuit	Busway systems typically only monitor the total power in the bus using optional equipment, and rely on Rack PDUs to report individual rack power	PDU supports auto-discovery for new branch circuits installed and has a single communication port for all branch circuits. Monitors power for each IT enclosure, even if the enclosure does not use a rack PDU, such as blade servers
Simplified engineering	Complex tradeoffs between row length, density, busway ampacity, and cost must be made in each installation, even in different rows within the data center, in order to optimize the result and ensure the busway is not overloaded	Select from standard reference designs to meet requirements. Many decisions can be deferred to later phases
Minimal copper use	Busway copper must be oversized to maximum power density	Branch circuit copper is only deployed when needed and at the capacity needed
Standard length final distribution cable	The busway is always the same distance from the IT enclosure so all cable drops are the same length, simplifying spares inventory	The distance from the PDU and the IT enclosure vary, requiring different length cables. Long cables can be cut and re-terminated by the user, but this is less convenient
Can be used in situations where overhead mounting is impractical	Busway cannot be installed under-floor in many locations due to regulations	Power distribution cable trays can be suspended from overhead, mounted to the tops of the IT racks, or installed under-floor
Minimal up-front cost	Most of the busway cost is incurred in the initial build	Most of the power distribution cost is incurred only when and if needed
Efficiency	Busway sized in advance for maximum capacity (i.e. the greatest amount of copper) resulting in slightly lower losses. Efficiency payback over 50 years given high cost of copper	Copper sized more closely to actual load resulting in slightly more losses

While busway to the rack is a dramatic improvement over the traditional approach, and the architecture of this paper uses busway to distribute bulk power in the data center, **Table 2** shows that the modular power distribution system has some advantages over busway for the final distribution to the IT enclosure. Busway for the final distribution has the advantage of zero footprint, but the modular distribution system is more scalable and adaptable to changing density, is standardized globally, and requires less up front planning and engineering.

In general, for distribution to the rack, busway is best suited for very large facilities with an open floor plan with a very well defined IT equipment layout. The modular distribution system has the greatest advantage when locations are not precisely defined in advance, the room is constrained in shape or has obstructions, or the power density is expected to vary significantly through the room. Either of these approaches is vastly superior to the traditional under-floor conduit system, and a summary of key factors to consider in the selection between the approaches is provided in **Table 3**.

Table 3

Factors to consider when choosing between busway and modular distribution to the IT enclosure

Factors suggesting busway to the rack	Factors suggesting modular distribution
No floor space, even 5% of the space or less, can be used by the power distribution system	<p>The IT enclosure layout is not well defined in advance</p> <p>The layout is not a simple rectangle with defined rows.</p> <p>The room will have zones of different density</p>
Busway to the rack may not be the preferred choice when:	Modular distribution may not be the best choice when:
<p>The locations of future IT enclosures are poorly defined</p> <p>The power densities of future zones are not well known in advance</p> <p>Overhead mounting is impractical due to ceiling construction or other constraints</p> <p>A global standard solution is required</p>	No IT floor space, even 5% of the space or less, can be devoted to the power distribution system

DC power distribution

DC power distribution has been proposed as an alternative to AC power distribution for data centers. There are actually four different DC approaches that have been described in the literature, using different DC voltage levels and wiring diagrams. The primary reason cited for moving to DC is improved electrical efficiency.

A number of studies have claimed substantial expected efficiency advantages for a DC power system, from 10% to 40%. However, these studies assume very low efficiency for AC power systems. Since these studies first came out, newer high efficiency AC power distribution architectures have become widely available with demonstrated efficiencies comparable with



Link to resource
White Paper 127

A Quantitative Comparison of High-Efficiency AC vs. DC Power Distribution for Data Centers

the hypothetical efficiency of DC systems. A quantitative analysis comparing AC vs. DC power efficiency can be found in the Green Grid White Paper 16, *Quantitative Efficiency Analysis of Power Distribution Configurations for Data Centers* and White Paper 127, *A Quantitative Comparison of High-Efficiency AC vs. DC Power Distribution for Data Centers*. These papers show that the best AC power distribution systems are nearly as efficient as DC, removing a key incentive to change the industry.

The most important problem with DC power distribution is the lack of availability of compatible IT devices. While some IT devices are available with a 48 V DC input option, this is the least efficient DC distribution voltage and consumes large amounts of copper wiring.

If DC were to ever become a standard distribution system for data centers, most of the same issues of circuit distribution and monitoring to the racks would remain. The fundamental approaches of busway and modular distribution could still be used, but new connector systems would need to be developed and the devices might need to become larger to comprehend the increased safety clearances required for high voltage DC.

The cost of transitioning the industry to DC would be very large and the lack of any significant cost or efficiency gains has removed the incentive to make this change. Therefore the industry is expected to remain AC based, although DC will be used as a distribution system within proprietary IT equipment including blade server chassis or racks, and within containerized servers.

Conclusion

This paper has described the significant limitations of the traditional power distribution architecture with under-floor conduit. Two alternative power distribution systems have been described, busway to the rack and modular distribution, both of which are major improvements over the traditional approach in terms of scalability, efficiency, reconfigurability, manageability, and power density. The modular power distribution system is shown to be particularly advantageous in practical data centers where layout is not precisely defined in advance, in retrofit applications, and in cases where the room layout includes unusual shaped floor plan or obstructions.



About the author

Neil Rasmussen is a Senior VP of Innovation for Schneider Electric. He establishes the technology direction for the world's largest R&D budget devoted to power, cooling, and rack infrastructure for critical networks.

Neil holds 19 patents related to high-efficiency and high-density data center power and cooling infrastructure, and has published over 50 white papers related to power and cooling systems, many published in more than 10 languages, most recently with a focus on the improvement of energy efficiency. He is an internationally recognized keynote speaker on the subject of high-efficiency data centers. Neil is currently working to advance the science of high-efficiency, high-density, scalable data center infrastructure solutions and is a principal architect of the APC InfraStruXure system.

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.



Resources

Click on icon to link to resource



Browse all
white papers
whitepapers.apc.com



**Increasing Data Center Efficiency by Using
Improved High-Density Power Distribution**
White Paper 128



**A Quantitative Comparison of High-Efficiency
AC vs. DC Power Distribution for Data Centers**
White Paper 127



Browse all
TradeOff Tools™
tools.apc.com



Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
DCSC@Schneider-Electric.com

If you are a customer and have questions specific to your data center project:

Contact your **Schneider Electric** representative