VOLTSERVER
DIGITAL ELECTRICITY™

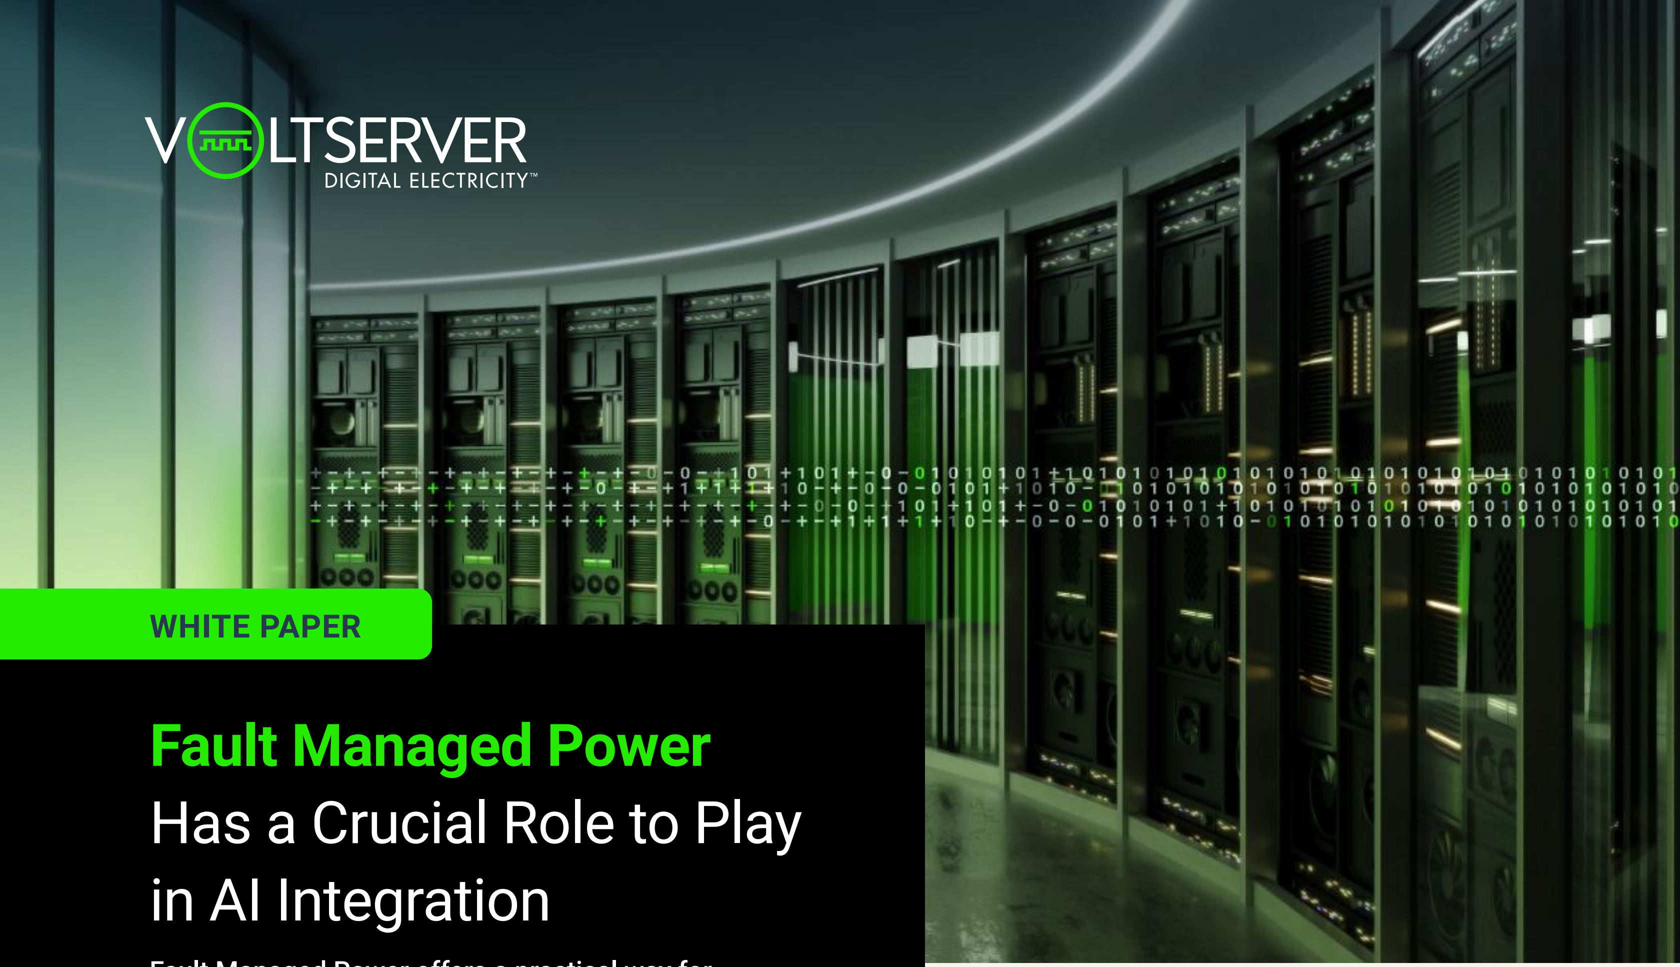

WHITE PAPER

Fault Managed Power Has a Crucial Role to Play in AI Integration

Fault Managed Power offers a practical way for operators to manage data center infrastructure costs while delivering high-performance AI capabilities.

By Steve Eaves and Ronna Davis

From now until 2030, the demand for AI-ready data center capacity is expected to climb by 33% each year.

To support this demand, data center power density is also rising: In fact, it has more than doubled in the last two years alone, jumping from 8 kW to 17 kW per rack. This trend is expected to continue in line with accelerations in AI.

Can the industry rely on legacy approaches to data center design and still keep up? Probably not.

From the inside out, it's clear that AI is transforming data center DNA. These deep, intrinsic changes call for more bandwidth closer to end devices to efficiently train and deploy complex machine learning models and support intensive AI workloads. Compared to traditional computing tasks, these require more power, cooling, and data.

The evolution of AI is driving major shifts in data center architectures, too. The equipment needed to process AI's vast amounts of data also changes physical space requirements. In the age of AI, data centers must be able to support:

- **Higher rack densities**, such as 100+ kW per rack
- **Advanced cooling systems**, such as liquid cooling, to remove heat directly from racks
- **Specific processor hardware**, such as GPUs (graphics processing units) and TPUs (tensor processing units)
- **High levels of scalability** to easily expand infrastructure as AI workloads grow
- **Innovative initiatives** like the Open Compute Project (OCP), which supports open standards for efficiency and flexibility

And as we discuss these inevitable changes, it's equally important not to forget the people working behind the scenes to make them happen. While operators juggle drastic pivots, they're also being called to reduce data center energy use and mitigate the environmental impact of data center facilities. Their ability to move quickly to prepare data centers for AI could impact the technology's deployment and adaptation rates.

This complex web of challenges not only forces the speed of progress but also brings new questions to the surface that operators have never had to answer before.

An Example: The Shift to 100 kW Racks

Faced with the pressure to respond to the relentless need for more computation power, intense storage requirements, and cutting-edge cooling demands, data center operators have no choice but to innovate. It's out with the old and in with the new to keep up with escalating demands in a cost-effective, sustainable manner.

Consider this scenario that represents the quandaries that operators will soon face.

After years of hosting standard 19-inch 15 kW racks to house server cabinets, a colocation data center is suddenly asked to host 100 kW racks that comply with Open Compute's Open Rack V3 (ORV3) to support modularity, higher power densities, and a 48V distribution system.

While capturing this new hosting opportunity is crucial for the colocation provider, it also poses a few problems:

- The existing overhead busway is already near capacity.
- The existing space offers no provisions for liquid cooling since the facility was designed and built a few years ago.
- Electrical infrastructure may need to be upgraded to support the higher power densities required by AI servers.
- Cages may need to be reconfigured to accommodate wider racks and proper airflow.
- Standard 19-inch racks aren't compatible with ORV3 specifications, which might require rack replacement.

How can operators safely deploy and manage these 100 kW racks without incurring major operational disruptions and costs?

With Data Center Innovation Comes Power-Related Questions

Not everything can transition at once. As AI takes hold, operators need to be able to oversee and maintain different types of installations, balancing:

- A mix of standard 19-inch and ORV3 rack types
- High-density, air-cooled racks up to 30 kW and very-high-density, liquid-cooled racks up to 150 kW
- Standard 42U racks and taller 48U and 52U racks
- Traditional centralized UPS (uninterruptible power supply) systems and ORV3 racks with internal batteries for backup

Scenes like these are already playing out in real life. In late 2024, NVIDIA announced that its Blackwell accelerated computing platform design will be used by OCP to drive the development of open, efficient, and scalable technologies in data centers that support higher compute density and networking bandwidth.

Its electro-mechanical design will drive the adoption of AI and includes rack architecture, compute and switch tray mechanicals, liquid-cooling and thermal environment specifications, and NVIDIA NVLink™ cable cartridge volumetrics. This will enable data center operators to build AI infrastructure faster and more cost-effectively while enabling the significantly faster performance and advanced workloads that AI demands.

An expansive group of partners—including Meta—are innovating and building on top of NVIDIA's Blackwell platform to develop equally flexible options that support high compute density and energy efficiency.

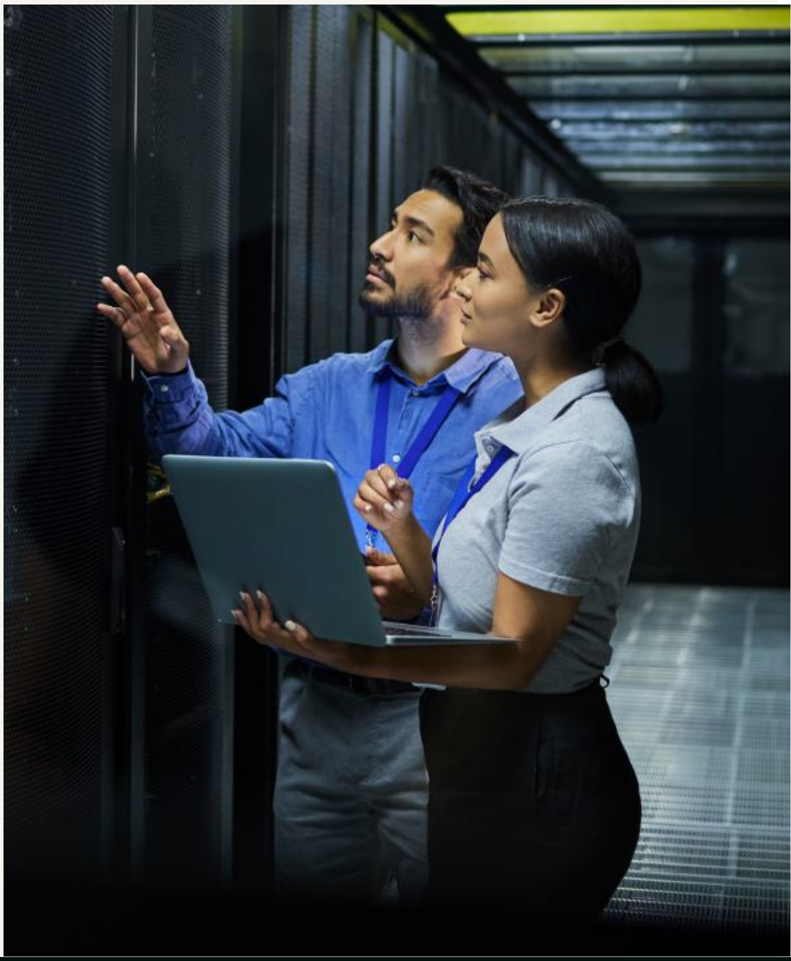

While this is a game-changing innovation, there are factors to consider before these types of platforms are deployed alongside more traditional systems:

- If available rack space is located far from the facility service entrance, how will data center operators move power and liquid cooling to the rack in an already-packed white space?
- How will safety and maintenance concerns about using liquid cooling in white space be addressed?
- How will data center operators manage the increased power requirements of the platform, which supports 100+ kW densities per rack?
- How will the overhead power distribution system be adapted to accommodate it?
- What are the structural implications of installing potentially wider and heavier racks that comply with ORV3 specifications?

Having answers to questions like these is critical before moving forward—and considering new types of power distribution systems that can support data center transformation in an AI-driven world is just as important.

The Role of Fault Managed Power in the Age of AI Data Centers

As a deeper understanding of power utilization becomes more important to manage power demands and reliability, data centers that support AI need smart power so loads can be managed intelligently.

Fault Managed Power (FMP) is a new class of electricity that addresses the challenges of designing data centers in an AI era. An intelligently managed power technology, it's a solution for operators who need to introduce AI workloads into their existing data center spaces but are constrained by pre-provisioned AC power distribution. FMP offers a faster, simpler way to welcome AI workloads without complete overhauls.

Adding FMP as a power distribution option can support fast scaling to help data centers keep up with AI and high-performance computing—and do it more economically.

FMP systems distribute power using built-in intelligence to manage and mitigate fault conditions. This makes power distribution faster, more flexible, more controllable, and safer than conventional electricity.

These systems are also much easier to deploy, which can significantly reduce project costs and allow fast device setup and powering. Unlike conventional electrical distribution, the inherent safety of Fault Managed Power systems means they can be installed using low-voltage installation practices.

Although VoltServer is the pioneer of this technology, FMP has grown exponentially in the last decade, with many other partners now making their mark. In 2024, these emerging leaders in power technology came together to form the Fault Managed Power (FMP) Alliance. Its members are working together to accelerate fault managed power adoption and help adjacent industries discover its benefits.

VOLTSERVER
DIGITAL ELECTRICITY™

Why FMP Is Uniquely Suited for AI

Compared to traditional AC power distribution, FMP is ideal for data center environments. Here are just a few reasons:

1. It eliminates the need for conventional overhead busways. Instead, FMP transmits power similarly to how data is transmitted over a network. It uses twisted-pair copper cabling that can be installed following the same installation practices as network cabling; it can be run through walls, under floors, etc.

2. It saves space and eliminates overhead busways because FMP and data can live in the same pathways or side by side with fiber or copper data lines.

3. Its intelligence and self-healing capabilities mitigate hazards. Consider leakage from liquid cooling systems, for instance. If a leak interferes with FMP, then the system immediately stops sending packets of energy. (As they're being sent from a transmitter to a receiver, these packets are always monitored for safety.) Power service continues once the hazard is remediated (this all happens within milliseconds).

4. It offers advanced, built-in capabilities for real-time remote power monitoring and control from a single interface. Issues can be detected and addressed quickly. This level of smart management means that data center systems always remain functional. When paired with a data center infrastructure management (DCIM) solution, FMP offers a comprehensive, actionable set of data about the status of power and data connections and environmental conditions.

5. It ensures reliability by alerting operators when incorrect power connections could compromise power distribution. Based on the alert, operators can investigate and take action.

6. It enables high levels of scalability and flexibility so operators can expand power-distribution infrastructure quickly as AI workloads grow.

Comparing Fault Managed Power and AC Power Distribution Systems

Let's compare an FMP design to a common AC power system in a large data center site.

In both scenarios, there are two areas to consider:

1. The **power pod** in an exterior container that houses incoming service, critical power backup, and the main switchboard

2. The **data room** or **white space**

Common AC Scenario

In this scenario, an AC UPS feeds the main output switchboard inside the power pod that connects to multiple distribution panels inside the data center.

The distribution panels feed power distribution units (PDUs) that contain transformers and circuit protection. The transformers provide electrical isolation to mitigate electrical harmonics and ground currents. The PDUs feed overhead busways that provide receptacles to accommodate server rack-mount PDUs.

AC power distribution systems like these require a large amount of equipment to provide redundancy. For example, since a fault in one busway could result in loss of power to multiple racks, two redundant pathways are usually provided. They are composed of electrical busways, distribution panels, and rack PDUs in what is known as a "2N architecture." All the equipment and connections call for capital expenses in addition to ongoing labor and maintenance.

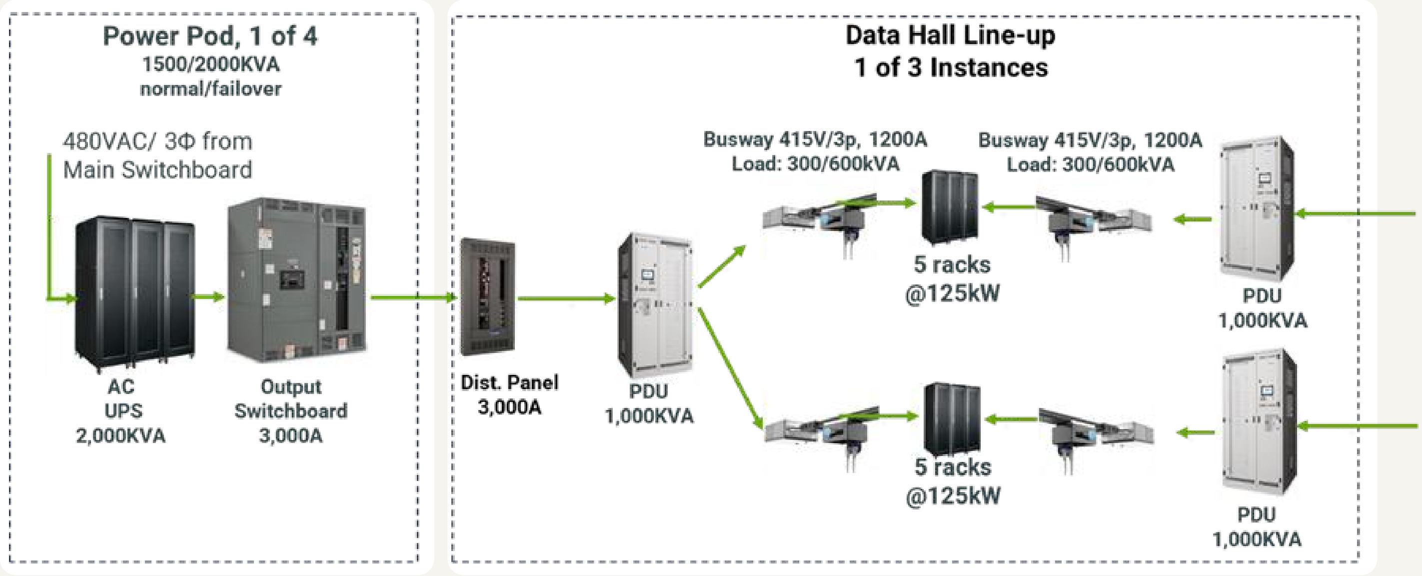

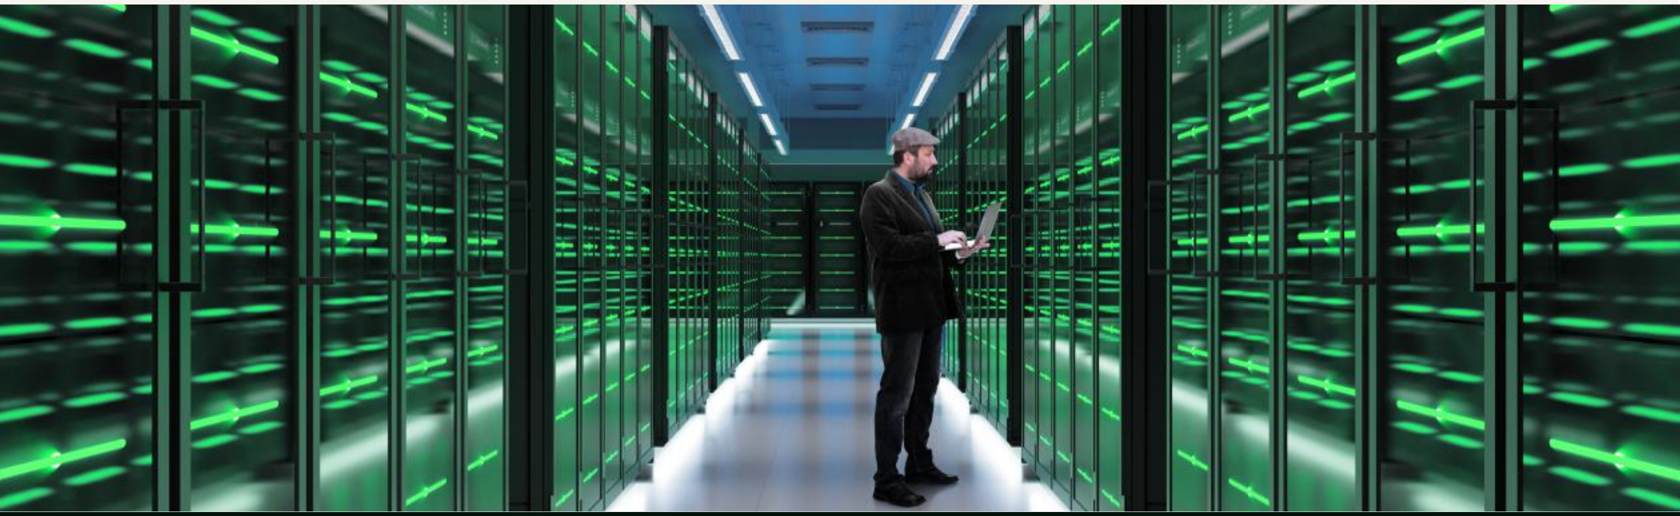

FMP Scenario

In this scenario, the power pod contains an FMP UPS, which is an integrated rack that contains UPS modules and FMP transmitters. One rack can be configured for 125 kW or more. The FMP UPS function eliminates the need for a separate UPS, PDU, distribution panel, and busway.

Safe, high-voltage FMP power is distributed by multiple sets of cables in an overhead cable tray; the cables connect to the FMP PDUs within the server rack, which, in turn, converts FMP to standard high-voltage DC for the server power supplies. Because FMP is a point-to-point architecture, a fault in a cable or FMP component disables only a small unit of power within the rack, requiring far less redundancy and equipment than a 2N AC architecture.

The use of FMP in the data center also eliminates electrical hazards and maintenance that come along with AC power, such as arc flash. In fact, standard IT practices similar to those used for Ethernet cables are used to install and maintain FMP cables. Moreover, FMP power contains embedded data to warn the operator of a misplaced or disconnected cable. It also provides inherent metering and on/off control as part of the standard offering.

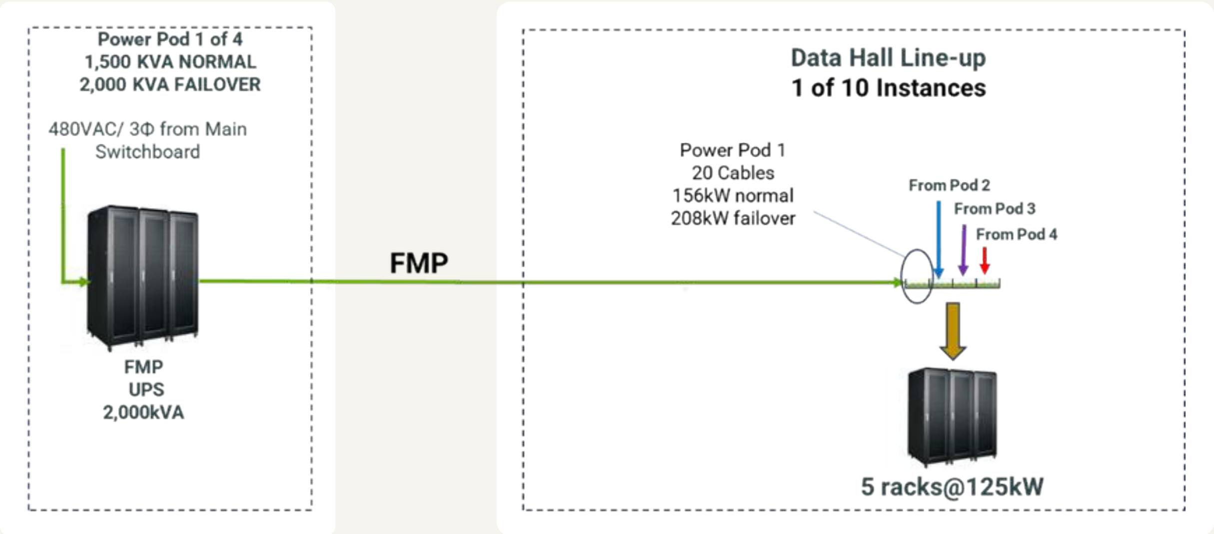

Below is a comparison of the cross-sectional area of a standard AC busway (415 VAC, 3 phase) configuration vs. the required cable and tray for FMP. This feeds five 125 kW server racks.

VoltServer Powers Data Centers

The AI revolution is far from over. Its relentless growth will continue to inspire data center operators to push the envelope.

FMP offers a practical way for operators to manage data center infrastructure costs while delivering high-performance AI capabilities. It saves space, reduces material and labor requirements, is faster to install, and is much safer than AC power can ever be.

Learn more about FMP in data centers.

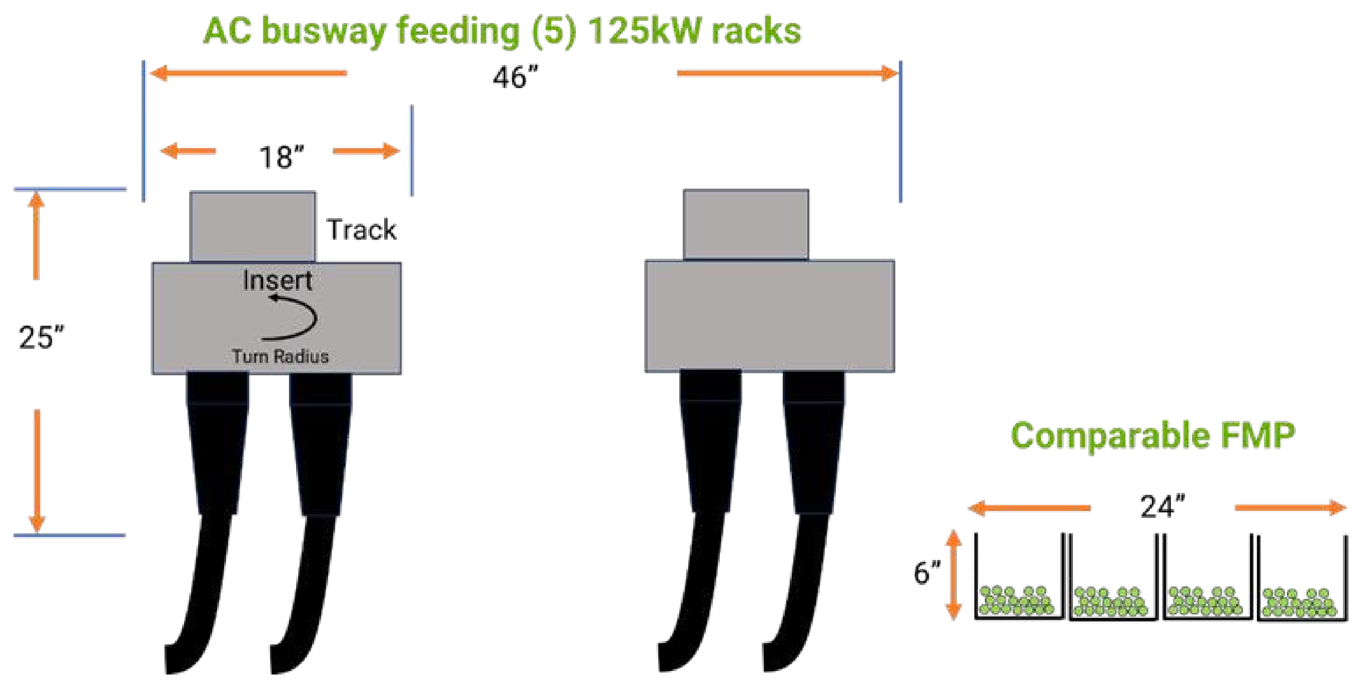

Steve Eaves is the founder and CEO of VoltServer. Ronna Davis is the CRO at VoltServer.