WHITE PAPER

Powering the AI Era:

Specifying and Deploying High-Density PDUs

A practical guide for infrastructure leaders deploying AI and high-performance computing environments

link-power.co.uk

01 Executive Summary

The rapid acceleration of AI and high-performance computing workloads is fundamentally changing how data centers are designed and operated. Rack power densities that once sat in the 5–10kW range are now routinely reaching 40–80kW, with deployments planning for 100kW and beyond.

In this environment, rack PDUs have moved from passive infrastructure to critical system components. Yet they are still routinely treated as a late-stage procurement item — specified after key design decisions have been made, sourced under time pressure, and deployed without adequate validation.

This approach introduces real risk: deployment delays driven by supply chain constraints, integration failures caused by physical incompatibilities, and operational inefficiencies from under-utilised intelligent features. In Europe specifically, the challenge is compounded by a limited pool of compliant, high-density PDU suppliers with proven field experience at scale.

> **The Core Message**
> *PDUs are the final critical connection between facility power and IT equipment. At high density, getting this wrong is expensive. This white paper provides the practical guidance needed to get it right.*

This document covers:

- What constitutes a modern high-density PDU and how the definition has evolved
- The key technical and operational challenges encountered during specification and deployment
- Practical, field-tested best practices to improve planning, execution, and long-term operation
- The emerging trends that will shape rack power distribution in the years ahead

02 The Shift to High-Density Compute

The data center industry is undergoing a structural transformation. AI, machine learning, and high-performance computing workloads demand significantly higher compute performance — and with it, a substantial increase in power consumption at the rack level.

From 10kW to 100kW: A New Normal

Enterprise data centers were historically designed around rack densities of 5–10kW. At those levels, power distribution was straightforward and PDUs were rarely a design constraint.

That has changed fundamentally. Today's AI infrastructure is redefining what is considered normal:

Deployment Type	Typical Rack Density
Traditional enterprise	5–10kW
Modern enterprise / edge	15–25kW
AI / HPC (current standard)	40–80kW
AI / HPC (near-term planning)	80–100kW
Ultra-high density (emerging)	100kW+

This increase in density is not linear in its impact. As power levels rise, complexity increases disproportionately across power distribution, thermal management, mechanical integration, and installation processes.

AC vs DC: Understanding the Divide

DC power distribution is not an emerging concept — for certain users, it is already mainstream. Hyperscalers who design and build their own data centres have adopted DC architectures at scale, constructing the facilities, processes, and operational expertise required to support them. Specialist colocation providers catering specifically for DC-powered deployments offer the same. For these organisations, DC is a deliberate, well-understood choice that delivers real advantages in efficiency and design flexibility.

For the majority of other organisations, the picture is different. DC distribution requires a fundamentally different approach to facility design, electrical infrastructure, safety processes, and operational management. The compatible IT ecosystem — while growing — remains more specific than the broad AC server market. Organisations without existing DC expertise face a significant barrier to entry: not because DC does not work, but because making it work requires a level of organisational investment and specialisation that most are not currently in a position to make.

AC power distribution therefore remains the practical default for the majority of deployments — driven by operational familiarity, established infrastructure, and the wide availability of AC-compatible server platforms from vendors such as Dell Technologies and Supermicro.

> **Key Point**
> *DC power distribution is viable and mainstream — for hyperscalers and specialist colocations who have built their infrastructure and processes around it. For most other organisations, AC remains the practical choice: not because DC is unproven, but because the barriers to entry are significant without the right expertise already in place.*

The Liquid Cooling Connection

Higher power densities are inseparable from the adoption of liquid cooling. Air cooling becomes insufficient at these densities, leading to the deployment of direct-to-chip cooling, manifolds, coolant distribution units, and associated infrastructure within the rack.

This has direct implications for PDU specification. Space for power distribution becomes more constrained, cable routing becomes more complex, and the case for remote monitoring and control becomes significantly stronger. When physical access to equipment is reduced, the ability to manage power remotely is no longer a luxury — it is a necessity.

Supply Chain Pressure

The rapid growth of AI infrastructure is placing real strain on the supply chain for high-density PDUs. A limited number of vendors are capable of producing these units, and fewer still have proven field deployments at scale. Large AI projects regularly require 500 to 5,000 PDU units, placing significant pressure on manufacturing capacity and delivery timelines.

This pressure is especially pronounced in Europe, where the pool of reputable, compliant suppliers is smaller than in North America or Asia. Supply chain strategy has become a critical component of PDU planning, not just a procurement detail.

03 Defining the Modern High-Density PDU

What was once a relatively simple device for distributing power within a rack has evolved into a critical infrastructure component. A high-density PDU must now deliver electrical capacity, physical flexibility, operational visibility, and control — often within a highly congested rack environment.

Electrical Specifications

The most common configurations for high-density deployments are:

Region	Standard Configuration
United States	60A, 415V three-phase
Europe	63A, 400V three-phase

These configurations support approximately 40–45kW per PDU. At higher rack densities, this typically results in two PDUs per rack (A and B feeds for redundancy), increasing the importance of phase balancing, load distribution, and coordination with upstream power architecture.

PDU Intelligence Levels

Modern PDUs span a significant range of functionality. The right level depends entirely on how the data and control will be used:

Level	Capability	Best Used When

MI — Metered Input	Overall PDU-level power monitoring	Basic load awareness only; limited value at scale
MPO — Metered per Outlet	Per-outlet power monitoring	Phase balancing and device-level visibility needed
MS — Metered + Switched	PDU monitoring + outlet switching	Overload prevention and controlled startup required
MSPO — Metered & Switched per Outlet	Full per-outlet monitoring and control	High-density, liquid-cooled or remote-managed environments

Field Observation

In liquid-cooled AI rack deployments, MSPO is rapidly becoming the default specification. When physical access is restricted, the ability to remotely monitor and control every outlet is not a premium — it is standard practice.

Outlet Strategy

Traditional PDUs use a mix of C13/C15 and C19/C21 outlets, requiring precise planning of device placement. In fast-moving, high-density environments, this introduces unnecessary rigidity and risk.

The C39 (4-in-1) outlet accepts C13, C15, C19, and C21 connectors in a single outlet. This eliminates the need to map outlet types to individual devices, simplifies deployment and reconfiguration, and reduces installation errors. The cost difference versus traditional outlets is minimal; the operational flexibility is significant.

Recommendation

Adopt C39 outlets as standard for high-density deployments. The marginal cost is negligible compared to the risk of specifying the wrong outlet type at scale.

Physical Characteristics

High-density PDUs place real physical demands on installation and rack design. A typical unit is:

- Up to 1.8 metres in length
- 20–25kg (45–55 lbs) in weight
- Equipped with heavy-gauge input cords extending 2–3 metres

These characteristics are frequently underestimated during planning. Combined with the congested rear-of-rack environment typical of liquid-cooled deployments, physical integration becomes a genuine design and installation challenge — not an afterthought.

04 Key Specification Challenges

Specifying the right high-density PDU involves navigating a range of interconnected technical, operational, and market challenges. Many of the issues encountered in real deployments are not caused by fundamental technology limitations, but by gaps in planning and evaluation.

Limited Supplier Choice with Proven Field Experience

High-density PDUs are not a commodity. At 60–63A, three-phase ratings with full compliance certification and real-world deployment experience, the supplier landscape is narrow — and in Europe, it is narrow to a degree that introduces genuine procurement risk.

Many vendors can quote a high-density specification. Far fewer have products with long-term field deployments in AI environments, third-party compliance certifications across multiple regions, and the manufacturing capacity to support large-volume projects at short notice.

Selecting a product that meets specifications on paper but lacks operational validation is a risk that becomes significantly more expensive at scale.

Outlet Type and Connector Compatibility

While overall outlet count is rarely the binding constraint, outlet type and flexibility are critical. Traditional C13/C19 configurations require precise pre-planning. In deployments where rack configurations can evolve — or where speed of installation is a priority — this introduces friction.

C39 multi-standard outlets remove this constraint, but the transition requires deliberate specification rather than defaulting to older approaches.

Cable Retention and Accidental Disconnection

In densely populated, liquid-cooled racks, the rear of the rack is a congested working environment. During installation and maintenance, power cables can be inadvertently dislodged — with potentially significant consequences at high density.

Locking outlet mechanisms address this risk, but not all approaches are equally suitable. Proprietary locking systems require specialized cables, introduce additional cost, and create supply dependencies. A server’s own power cord will not work with a proprietary lock — which means a $2 cable can halt an entire deployment if the proprietary alternative is out of stock.

Neutral locking outlets secure connections using standard in-box server cords, providing the same physical protection without the dependencies.

Phase Balancing at High Load

Phase balancing becomes significantly more challenging as individual device power increases. In AI racks, a single server or GPU module can represent a substantial portion of rack load. Imbalance can occur quickly, and the consequences are more severe at higher densities.

Intelligent PDUs with phase-level monitoring provide real-time visibility. Advanced MSPO units add the ability to control individual outlets, sequence startup to avoid inrush current events, and prevent unintended connections from disrupting load balance.

Monitoring Alignment: Capability vs Utilisation

A persistent gap in high-density deployments is the disconnect between specified capability and operational utilisation. Organisations invest in advanced MSPO functionality, then fail to integrate it into their monitoring workflows — paying for data they are not using.

> **Key Principle**
> *Data without action has no value. Intelligent PDU features are only worth their cost if they are integrated into operational processes. The question to answer before specifying is not 'what is technically possible?' but 'what will we actually use?'*

Environmental Monitoring in Liquid-Cooled Environments

Liquid cooling introduces a low-probability but high-impact risk: leaks. Modern PDUs can integrate leak detection, temperature, and humidity sensors — and, more importantly, can be configured to trigger automated responses such as controlled outlet shutdown when an anomaly is detected.

This represents a shift from reactive monitoring to integrated risk management within the power layer.

05 Deployment Challenges

Deployment is where specification decisions are tested against physical reality. Many of the most costly and time-consuming issues encountered in high-density PDU deployments are not technical failures — they are planning failures.

The Afterthought Problem

Across fast-moving AI data center projects, the pattern is consistent: compute platforms, cooling strategy, and high-level electrical design receive early and sustained attention. Rack PDUs are addressed last.

The consequence is that by the time PDU specification begins, design constraints are already set, timelines are compressed, and there is no capacity for evaluation or validation. PDUs are the final critical connection between infrastructure and IT equipment — they deserve better treatment than this.

Supply Chain at Scale

Large AI deployments require PDU volumes that test the limits of even established suppliers. Projects of 500 to 5,000 units are common, and the number of vendors genuinely capable of delivering compliant, high-quality units at this scale — particularly in Europe — is small.

When PDU procurement is initiated late, the risks compound: limited supplier options, extended lead times, and the potential need for expensive air freight. In a project where PDUs represent the final connection to IT equipment, a supply chain failure at this stage can bring everything to a halt.

Physical Handling and Installation

High-density PDUs are large, heavy, and physically demanding to install. A unit approaching 25kg and 1.8 metres in length, with a multi-metre power cord attached, is genuinely difficult for a single person to handle safely within a rack.

Best practice is two-person installation. In practice, this is frequently not planned for — resulting in improper installation, increased installation time, and real safety risk.

> **Safety Note**
> *Installing two or four heavy PDUs into an empty or partially populated rack, with heavy input cables hanging from the top, creates a rear-weight bias that can cause the rack to tip. Installation sequence and temporary stabilisation should be explicitly planned.*

Cable Entry and Physical Fitment

At higher power ratings, input connectors are large — up to 4.5 inches wide. A frequently overlooked but critical compatibility check is whether the connector can physically pass through the rack's roof or floor entry panel.

This is not a theoretical concern. Projects have encountered on-site rework and delay because the connector dimension was not verified against the rack panel design. It is a straightforward check that should happen at specification stage, not on-site.

Packaging and Waste

An underappreciated operational burden in large deployments is packaging waste. Each PDU arrives individually boxed and padded. At 500 to 5,000 units, this generates a significant volume of packaging material that requires removal and disposal on-site.

Vendors who actively support packaging removal and waste management as part of their service offering provide meaningful practical value. This is worth evaluating during vendor selection.

06 Best Practice Guidelines

The following guidance is drawn from real-world high-density PDU deployments. It is designed to be practical and direct, focused on the decisions that have the greatest impact on outcomes.

1. Engage PDU Vendors Early

Early vendor engagement does not need to be time-consuming. At a minimum, it should confirm whether a vendor can provide a detailed specification sheet, demonstrate valid compliance certifications, provide a credible quote at project scale, and give honest production and delivery timelines.

Published lead times are often boilerplate. Real lead times for high-specification, high-volume units can differ significantly. Understanding the actual production schedule — not the marketing estimate — is essential for project planning.

Where time permits, visiting a vendor's manufacturing facility is highly valuable. It provides direct insight into production capability, quality processes, and organisational maturity that no datasheet can convey.

2. Align Procurement with Deployment Schedules

PDUs should be ordered to arrive when needed — not speculatively early, and certainly not after the rack build has started. For large deployments, phased ordering aligned to deployment milestones is best practice.

A buffer of 5–10% is advisable. PDUs store easily, and a small buffer prevents deployment stalls caused by DOA units or minor logistics delays. When a DOA unit is identified, handle it directly with the vendor rather than returning it to the standard supply chain process — this avoids unnecessary delay.

3. Standardise with Clear Design Rules

Standardisation reduces cost, simplifies procurement, and improves interoperability — but it requires clear governance. A standardised 63A PDU rated for approximately 44kW will cause problems in a 50kW rack if the configuration rules are not enforced.

Define explicit rules linking rack IT configuration to PDU specification. With those rules in place, standardisation becomes a genuine advantage: volume-based cost reductions, faster procurement, and reduced deployment complexity.

4. Always Use C39 Outlets

The C39 (4-in-1) outlet accepts C13, C15, C19, and C21 connectors. There is almost no cost premium. The operational benefit — eliminating the need to map outlet types to device placement, and removing the risk of a configuration mismatch halting an installation — is significant.

In high-density, fast-moving environments, C39 should be the default. The cost of the alternative — discovering the wrong outlet type during deployment — is far higher.

5. Use Neutral Locking Outlets

Cable retention matters in dense racks. The approach to locking, however, matters equally. Proprietary locking systems require purchasing specific cables — if those cables are out of stock, deployment stops. The standard cord that ships with the server cannot be used.

Neutral locking outlets provide the same physical security using standard in-box cords. There is no legitimate operational reason to introduce proprietary cable dependencies in a high-density deployment.

6. Define a Monitoring Strategy Before Specifying Intelligence

The value of an intelligent PDU is entirely dependent on how its data is used. Before specifying MSPO capability, confirm that the organisation has the operational processes and tooling to act on the data it generates.

When used effectively, MSPO delivers real-time outlet-level power monitoring, phase balancing visibility, outlet sequencing to prevent inrush tripping, overload protection, and environmental alerting. This is genuinely valuable infrastructure.

When not used effectively, it is an unnecessary cost. The decision framework is simple: if the data will be used, specify the intelligence. If not, a simpler configuration is more appropriate.

> **The Autopilot Principle**
> *Data without action is like a plane without autopilot — all the sensors are working, but nothing is being done with the information. Intelligent PDUs only earn their cost when the data drives decisions, ideally automated ones.*

7. Plan for Phase Balancing and Dynamic Loads

AI workloads are not static power consumers. GPU and accelerator loads fluctuate, and phase imbalance can develop rapidly. Design for no more than approximately 80% phase utilisation as a ceiling, and monitor actively after commissioning.

Initial design based on nameplate ratings is a starting point, not a guarantee. Ongoing monitoring through intelligent PDUs enables continuous optimisation and early identification of capacity pressure.

8. Treat Power and Cooling as a Unified System

In liquid-cooled environments, the separation between power and cooling teams is no longer viable at the rack level. Supply and return manifolds, cooling hoses, A/B power feeds, and PDUs all compete for the same physical space. Planning decisions in one domain directly constrain the other.

Rack-level power and cooling should be designed together, by a team that owns both. Environmental monitoring integrated into PDUs — including leak detection and residual heat sensing — provides additional value in this integrated environment.

9. Plan Installation in Detail

High-density PDU installation is a physical task that requires planning. Key considerations:

- Two-person installation for units approaching 25kg
- Verify cable connector dimensions against rack roof/floor panel entry points before deployment
- Plan installation sequence to maintain rack stability during build
- Ensure sufficient staff and adequate time; these are not quick jobs
- Coordinate with the vendor on packaging removal — at scale, this is a real operational burden

10. Embed Vendors in the Project

The best PDU deployments happen when the vendor is treated as a project partner, not a purchase order. An embedded vendor understands the project timeline, anticipates demand fluctuations, resolves issues before they reach the installation team, and takes ownership of the small problems — packaging, DOAs, delivery coordination — that would otherwise consume project time.

The internal process from engineering approval to purchase order reaching a vendor can be surprisingly slow. An embedded vendor navigates this without escalation. A transactional vendor waits — and the project pays the price.

Three Rules for Successful High-Density PDU Deployment

1. Engage early and validate reality — understand real capabilities, real lead times, and real constraints before design decisions are locked. 2. Specify intelligence with intent — if you will use the data, invest in MSPO. If not, keep it simple. 3. Eliminate surprises — align design, supply chain, and deployment teams early. Without PDUs, IT equipment cannot be powered.

07 Future Trends

The trajectory of rack power distribution is clear: higher densities, greater integration between power and cooling, and increasing reliance on intelligent, action-driven infrastructure. Understanding these trends enables better decisions today.

Continued Growth in Rack Density

What is currently considered high density — 40–80kW per rack — is rapidly becoming standard. Racks in the 80–100kW range are expected to become increasingly common in the near term, with specialised deployments pushing further. This will place additional pressure on every aspect of rack-level power distribution: capacity, physical integration, cable management, and installation processes.

Convergence of Power and Cooling

As liquid cooling becomes widely adopted, the traditional separation between electrical and mechanical disciplines at the rack level will become untenable. Future deployments will increasingly treat rack power and cooling as a unified system. Vendors and infrastructure teams that are already working this way will have a meaningful advantage.

From Data Collection to Automated Action

The evolution of intelligent PDUs is moving from monitoring toward automation. MSPO configurations will become more common, but the real shift is in how the data is used — from manual review to automated response. Predefined actions triggered by sensor data (e.g., controlled outlet shutdown on leak detection, load shedding in response to phase imbalance) will become standard practice in well-run facilities.

Supply Chain as a Strategic Capability

As demand for high-density infrastructure grows, supply chain relationships will become a competitive differentiator. Organisations with embedded vendor partnerships, accurate forecasting, and procurement processes aligned to engineering timelines will deploy faster and at lower cost than those treating PDUs as a commodity purchase.

DC Architectures: Growing Adoption, High Barriers

DC power distribution is already mainstream within hyperscale environments and specialist DC-ready colocations. That adoption will continue to grow, and the compatible IT ecosystem will expand alongside it. For organisations already operating in DC environments, the trajectory is clear.

For everyone else, the barriers remain significant. Facility design, electrical infrastructure, operational processes, and team expertise all need to be built around DC from the ground up. This is achievable — but it is a major undertaking. Until the ecosystem matures further and the operational burden reduces, AC will remain the standard entry point for most organisations deploying high-density infrastructure.

08 Conclusion

AI infrastructure has fundamentally changed the requirements placed on every component of the data center — including the components that were once considered secondary. Rack PDUs are no longer passive or peripheral.

At 40–80kW and beyond, they are a critical interface between facility power and IT equipment, and their specification, availability, and deployment directly determine whether projects succeed or fail.

The good news is that most of the challenges described in this document are avoidable. They are not caused by technology limitations — they are caused by late engagement, insufficient planning, and underestimation of physical and operational complexity.

Organisations that treat PDUs as a critical infrastructure component from the earliest stages of project planning — engaging vendors early, aligning procurement with deployment, specifying intelligently, and coordinating power and cooling design — consistently achieve better outcomes.

Final Thought

The rack PDU is the last connection made before IT equipment is powered on. In a high-density AI deployment, it is also one of the most consequential. Treat it accordingly.

Link Power works with infrastructure teams across Europe to specify, source, and deploy high-density PDUs for AI and HPC environments. For further information, visit link-power.co.uk.