INTERVIEW

# From megawatts to gigawatts: Powering the AI-driven data center era

## Interview with Miguel A. Osuna-Perez, Global Segment Leader for Data Center applications, Transformers Business at Hitachi Energy

The global data center industry is entering a decisive phase in which the rapid expansion of digital infrastructure is colliding with the physical realities of power systems. Accelerated adoption of cloud computing, artificial intelligence, and high-density workloads is driving an unprecedented surge in electricity demand. What was once a megawatt-scale industry is increasingly becoming a gigawatt-scale challenge, fundamentally reshaping how digital infrastructure is planned, financed, and delivered.

This transformation is being led by hyperscalers—global cloud and AI platforms that are deploying infrastructure at extraordinary scale to support AI training, inference, storage, and mission-critical cloud services. Their requirements are redefining the industry's benchmarks: extremely large and concentrated power capacities, rapid deployment timelines, and electrical systems capable of supporting steep load ramps and evolving operational profiles driven by AI workloads.

In parallel, colocation providers are expanding aggressively to meet enterprise demand, host AI clusters, and serve as critical on-ramps to cloud platforms. While facing similar power and infrastructure constraints, colocation operators must also accommodate a diverse tenant base with varying requirements for redundancy, efficiency, and service-level performance. As a result, flexibility, scalability, and standardization are becoming as important as raw capacity.

Across both hyperscale and colocation segments, the industry is now confronting a dual challenge. First, developers must deploy electrical infrastructure capable of delivering very large amounts of power reliably and safely, often at hundreds of megawatts per site, and increasingly across multi-gigawatt campuses. This has accelerated the shift toward standardized, modularized electrical building blocks that reduce engineering cycles, improve repeatability, and enable faster deployment across multiple geographies.

Second, securing access to that power has become equally critical. Grid capacity constraints, long interconnection lead times, permitting complexity, and transmission availability are now decisive factors in site selection and project viability. In this environment, strong partnerships, supply-chain resilience, and a global manufacturing footprint are essential to ensure equipment availability, delivery certainty, and the ability to scale in parallel with customer demand.

While these dynamics are global, the United States is currently at the forefront of this transformation. The development of gigawatt-scale AI factories designed to power the next generation of large-language models and compute-intensive applications is reshaping long-standing grid assumptions and accelerating the need for resilient, reliable, and fast-to-deploy power infrastructure solutions.

Within this context, transformers play a central role. As the critical interface between utility grids and mission-critical computing environments, transformers must enable higher-power and higher-voltage grid connections while also supporting rising rack-level power densities inside the data hall. They are increasingly expected to deliver this performance within compressed timelines, without compromising safety, reliability, or long-term scalability.

This interview with **Miguel Osuna-Perez**, Hitachi Energy's **Global Segment Leader for Data Center Applications in the Transformer Business**, explores how these converging forces—AI-driven demand, power system constraints, and the need for speed and scale—are redefining technology choices, manufacturing strategies, and the future of power delivery for the digital economy.

**The global data center industry is entering a decisive phase in which the rapid expansion of digital infrastructure is colliding with the physical realities of power systems**

## As the critical interface between utility grids and mission-critical computing environments, transformers must enable higher-power and higher-voltage grid connections while also supporting rising rack-level power densities inside the data hall

### Welcome Miguel, to begin, could you please introduce yourself and share some insights about your professional background and experience in the data center and power-infrastructure industry?

I am the **Global Segment Leader for Data Center Applications** within Hitachi Energy's transformer business. In this role, I focus on aligning our transformer technologies, manufacturing capabilities, and global delivery frameworks with the rapidly evolving needs of cloud and AI infrastructure. A core principle of my work is maintaining very close proximity to customers, understanding their challenges firsthand and ensuring that the full breadth of our organization is deployed to support them effectively.

I was born and raised in Mexico, where my academic background is in **mechatronic engineering**, and where I also began my professional journey in **electrical-mechanical engineering and power systems**. Early in my career, I worked as a **design engineer for transformers and voltage regulators**, which gave me hands-on exposure to the core technologies that underpin grid infrastructure and industrial power applications. That experience grounded me in the fundamentals of reliability, system behavior, and long-term asset performance—principles that remain essential in everything I do today.

After moving to the United States, my career broadened beyond pure engineering into **product development, product and sales management, market development, and business leadership** roles. This transition allowed me to combine technical depth with a commercial and strategic perspective, particularly in the context of large-scale, mission-critical power infrastructure.

Over the years, my work has increasingly focused on the **intersection of rapid data center expansion and the power systems required to enable it**. I work closely with hyperscalers and colocation providers as they navigate rising load densities, compressed deployment timelines, and the emergence of **multi-gigawatt AI campuses**. By bringing together transformer technology, digitalization, and lifecycle services, my objective is to help ensure that developers and operators have **resilient, scalable power foundations** capable of supporting the next decade of global digital growth.

### Data center demand continues to expand rapidly. How do you view the role of power infrastructure, specifically transformers, in this evolution?

Data center demand is expanding at a pace the power system has not historically been designed to absorb. What we are seeing today—driven largely by AI workloads—is not incremental growth, but a structural shift in both scale and load behavior. Multiple industry analyses point to power availability, not land or capital, as the primary constraint on future data center expansion, with many campuses now planned at hundreds of megawatts and increasingly at gigawatt scale.

In this environment, power infrastructure moves from being a supporting function to a determining factor for whether projects can proceed at all. AI workloads introduce high power densities, rapid load ramps, and continuous operation profiles that place new demands on the grid and on-site electrical systems. Research from utilities, developers, and policy institutions consistently highlights that existing grid planning assumptions are being challenged by the size, speed, and concentration of data center loads.

Transformers sit at the center of this challenge. They are no longer passive components in the power chain; they are a critical interface between constrained grids and highly dynamic, mission-critical computing environments. As data centers scale, transformers increasingly influence how quickly power can be delivered, how reliably it can be operated, and how effectively sites can expand over time. Industry discussions now consistently link transformer design and availability to time-to-power, grid interconnection success, and long-term scalability.

We are also seeing a shift in how developers think about power architecture. Long interconnection queues and grid constraints are driving greater use of behind-the-meter generation, hybrid energy systems, and energy storage, all of which place additional technical and operational expectations on transformers—from handling more complex operating modes to supporting grid-interactive behavior. At the same time, higher rack densities and AI-driven load volatility are pushing the need for transformers with higher performance, improved efficiency, enhanced monitoring, and designs that support modular and phased deployment.

Ultimately, as data centers evolve into some of the largest and most critical electrical

**A core principle of my work is maintaining very close proximity to customers, understanding their challenges firsthand and ensuring that the full breadth of our organization is deployed to support them effectively**

loads on the grid, transformers become foundational to the digital economy itself. Their role is expanding from enabling connection to enabling confidence—confidence that power can be delivered reliably, scaled predictably, and adapted as AI workloads continue to evolve. That is why transformer technology, digitalization, and lifecycle support are increasingly central to how the industry thinks about powering the next generation of data centers.

**Ultimately, as data centers evolve into some of the largest and most critical electrical loads on the grid, transformers become foundational to the digital economy itself**

## Hitachi Energy is widely recognized for its global leadership in transformer technology. How does the company's experience translate into value for data center customers?

Hitachi Energy's experience translates into value for data center customers in a very practical way: it reduces execution risk at a time when power has become the critical path for growth. As AI-driven demand accelerates and data center campuses scale toward hundreds of megawatts and, increasingly, gigawatt-level architectures, customers are no longer just looking for equipment—they are looking for certainty around delivery, consistency, and long-term scalability.

Our leadership in transformer technology is built on decades of operating at the highest levels of voltage, power, and system criticality across grids and mission-critical infrastructure worldwide. For data center customers, that experience shows up first in **execution capability at scale**. Large programs often involve dozens—or hundreds—of transformers deployed across multiple sites and regions, under very compressed timelines. Our global design and manufacturing model, anchored by standardized design platforms such as **TrafoSTAR**, allows us to industrialize delivery while maintaining uniform quality and performance. This "design once, build everywhere" approach significantly reduces engineering cycles, improves predictability, and enables faster time-to-power across geographies.

Equally important is how this experience translates into **infrastructure that reflects real operating conditions**. AI workloads introduce high load factors, rapid ramp rates, and continuous utilization that place new stresses on power systems. Transformers are no longer selected purely on nameplate ratings; they must support efficiency at sustained high load, thermal robustness, power quality, and future expansion. Our large

INTERVIEW

## Our leadership in transformer technology is built on decades of operating at the highest levels of voltage, power, and system criticality across grids and mission-critical infrastructure worldwide

installed base and long operating history give us the insight to design solutions that perform reliably in these demanding environments, not just on paper.

Another key differentiator is **program-level execution**. Hyperscalers and colocation providers increasingly operate global deployment programs with standardized architectures and repeatable build models. Supporting that requires more than individual projects—it requires coordinated program management, factory alignment, and disciplined execution across the supply chain. Our standardized designs, global manufacturing footprint, and integrated program management capabilities allow customers to scale infrastructure with confidence while reducing project risk.

Finally, value extends well beyond delivery. As data centers grow larger and more complex, transformers become long-term enablers of availability and expansion. Our lifecycle approach—spanning application engineering, factory testing, digital monitoring, predictive diagnostics, and service—provides operators with real-time visibility and actionable insight into asset health. That transparency is becoming increasingly critical as uptime tolerances shrink and power systems operate closer to their limits.

Ultimately, the value we bring is the ability to **connect global expertise with local execution**. In a market where speed, reliability, and scalability are decisive, our standardized design frameworks, global manufacturing capability, and disciplined execution model give data center customers the confidence to deploy power infrastructure at scale—even as AI and digital demand continue to evolve.

### Transformers for data centers face unique challenges compared to other industrial applications. Could you share your perspective on what differentiates them?

Transformers serving data centers operate in a fundamentally different environment than those used in most traditional industrial applications. While industrial loads are often process-driven, intermittent, and linked to production cycles, data centers—particularly those supporting AI workloads—are **continuous,**

**mission-critical, and highly dynamic electrical environments**.

Transformers serving data centers operate in a fundamentally different environment than those used in most traditional industrial applications—both in scale and in how power is consumed. Modern data centers, particularly those supporting AI workloads, are defined by extremely high-power levels and load profiles shaped by GPU-dense architectures, continuous operation, and rapidly changing computational activity. As a result, transformers must meet a distinct and more demanding set of requirements.

One key differentiator is the combination of **high sustained loading and dynamic transitions**. Unlike many industrial facilities, where loads are process-driven and intermittent, data centers operate close to rated capacity for extended periods while also experiencing steep ramp rates driven by AI training and inference cycles. These rapid changes introduce elevated harmonic content and transient stresses that place significant demands on insulation systems, thermal design, and overall electromagnetic performance. Maintaining stable operation under these conditions requires precise engineering and designs optimized for continuous, high-duty operation rather than short-term overload capability.

Another critical distinction is the **exceptional reliability and power-quality requirement**. In data center environments, even minor voltage deviations, transient events, or unplanned interruptions can have immediate operational and financial consequences. As a result, transformers must be engineered not only for efficiency and capacity, but also for resilience—supporting stringent redundancy architectures, enhanced transient voltage protection, and harmonic-tolerant configurations that align with Tier III and Tier IV availability expectations.

**Scalability and predictability** further differentiate data center transformers from most industrial applications. Data centers are rarely designed as static facilities; they are built to scale in phases, often under tight timelines, while remaining fully operational. This places a premium on transformers that support modular expansion, standardized architectures, and predictable performance as capacity is added. Unlike industrial plants, where electrical systems are typically sized for a fixed production envelope, data center infrastructure must enable growth without disruption, reengineering, or extended outages. Predictable behavior under future load scenarios is therefore as important as initial performance.

**In data center environments, even minor voltage deviations, transient events, or unplanned interruptions can have immediate operational and financial consequences**

A final differentiator is the growing expectation for **digital-enabled assets**. Data centers increasingly require real-time visibility into the condition and performance of critical infrastructure. Transformers are no longer treated as passive components; they are expected to integrate into broader operational intelligence platforms. Through embedded monitoring and predictive diagnostics, transformers provide actionable insight that supports proactive maintenance, reduces unplanned downtime, and strengthens long-term availability.

Put simply, transformers for data centers are no longer passive elements in the electrical chain. They act as active enablers of **reliability, efficiency, scalability, and predictability**—supporting rapid deployment cycles, very high rack-level power densities, and the ability to evolve over time as data center architectures and AI workloads continue to advance.

## How does Hitachi Energy's portfolio and global scale directly benefit data center developers and operators?

Hitachi Energy's portfolio and global scale benefit data center developers and operators because they are grounded in deep infrastructure experience and deliberately shaped around the needs of this segment. Over time, we have built a **dedicated data center portfolio**, informed directly by how hyperscalers and colocation providers design, deploy, and operate power infrastructure at scale.

A key differentiator is our ability to **listen closely to customers and translate those insights into engineered solutions**, rather than adapting generic industrial designs. Data centers are evolving faster than almost any other infrastructure segment, and many requirements—higher short-circuit levels, tighter space constraints, new cooling concepts, and more dynamic load behavior—do not fit neatly into legacy design assumptions. Our R&D teams work closely with customers to develop and validate solutions for these emerging conditions, and our global organization allows us to **transfer that technology rapidly and consistently across regions**.

Our long history serving **TSOs, DSOs, utilities, and MV equipment OEMs** plays a critical role here. These are among the most demanding customers in the power industry, particularly when it comes to **short-circuit withstand capability, system robustness, and compliance with diverse grid codes**. That experience directly benefits data center customers, as campuses increasingly resemble utility-scale loads with very high fault levels and complex grid interfaces. Our industry-leading track record in short-circuit design provides confidence that transformers can operate safely and predictably as power levels continue to rise.

Another advantage comes from our broad experience across **multiple insulation systems and transformer technologies**. Data center environments vary widely outdoor substations, indoor installations, proximity to IT equipment, sustainability requirements, and regional regulations all influence design choices. Our portfolio spans liquid-filled, dry-type, and alternative insulation technologies, allowing customers to select solutions that best fit their application while maintaining a consistent design philosophy across their infrastructure.

Global scale then turns these capabilities into **repeatable execution**. Through

standardized design platforms and aligned manufacturing, we are able to take solutions developed for one customer or region and deploy them reliably across global build programs. This supports the universal, modular designs that data center developers increasingly seek—architectures that work in multiple scenarios, scale in phases, and deliver predictable performance as campuses expand.

Ultimately, the value of our portfolio and global scale lies in the combination of **engineering depth, customer-driven innovation, and industrialized execution**. It allows data center developers and operators to deploy power infrastructure that is not only robust and compliant today, but adaptable to future architectures, higher power densities, and evolving grid conditions.

## With the exponential growth of data centers, especially those supporting AI workloads, how is Hitachi Energy scaling to meet this demand?

The exponential growth of data centers—particularly those supporting AI workloads—requires a level of manufacturing scale, supply-chain resilience, and execution speed that goes well beyond traditional growth models. At Hitachi Energy, we are scaling deliberately and structurally to meet this demand, treating data centers as the critical infrastructure they have become.

A central element of our approach is a **strategic, multi-year global investment program** designed to expand capacity while strengthening business continuity. We are increasing transformer manufacturing capability across regions, reinforcing supply-chain security for critical materials, and ensuring that production is diversified rather than concentrated. This diversified footprint is essential in today's environment, as it provides customers with confidence that long-term power needs can be met reliably, even in the face of geopolitical, logistical, or market disruptions.

In the United States, for example, we have committed **nearly $1 billion to expand transformer manufacturing**, including the development of a new large power transformer facility in South Boston, Virginia. Once fully operational, it will be one of the largest of its kind in the country, significantly strengthening domestic capability to support AI-driven digital infrastructure. Similar investments across other regions reinforce our ability to serve global deployment programs with consistent quality and predictable timelines.

These investments are not only about adding volume; they are about **scaling the right capabilities**. AI-driven data centers increasingly require higher power ratings, higher short-circuit withstand levels, and more complex designs than traditional applications. We are expanding factory test capabilities, engineering resources, and industrial processes to support repeatable, utility-grade designs at scale. At the same time, our global design platforms allow solutions developed for one customer or market to be transferred efficiently across regions, avoiding fragmentation as demand accelerates.

Scaling also requires a different **execution model**. Many hyperscalers and colocation providers are pursuing large-scale, multi-site deployment programs rather than individual projects. Our investments support this shift by enabling standardized designs, aligned factories, and program-level coordination across engineering, manufacturing, supply chain, and service. This allows customers to execute parallel builds with greater predictability, even as volumes and complexity increase.

**Beyond manufacturing, we actively work with governments, utilities, and regional agencies to support the broader ecosystem required for digital infrastructure growth**

Beyond manufacturing, we actively work with **governments, utilities, and regional agencies** to support the broader ecosystem required for digital infrastructure growth. Accelerating grid access, strengthening local supply chains, and aligning on long-term infrastructure planning are all essential to sustaining the pace of expansion the industry is now experiencing.

Together, these actions—capacity expansion, diversified manufacturing, technology industrialization, and ecosystem engagement—enable Hitachi Energy to scale alongside the data center industry. Most importantly, they give developers and operators confidence that the high-capacity, utility-grade transformers required for the AI era will be available, repeatable, and supported over the long term.

## Looking ahead, what trends do you see shaping the data center ecosystem and the role of transformer technology within it?

Looking ahead, the data center ecosystem is entering a phase where **power architecture, not just compute architecture, becomes the defining design variable**. The rapid expansion of AI workloads is pushing power density, scale, and operational intensity beyond what traditional data center models were designed to support, and this is fundamentally reshaping both infrastructure design and the role of transformer technology.

One of the most visible shifts is the evolution of **grid-to-rack power architectures**. As rack-level power densities increase from tens to hundreds of kilowatts, operators are rethinking the entire power-delivery chain to reduce losses, simplify conversion stages, and maintain tighter voltage control under fast-changing computational loads. This is driving interest in **higher-voltage distribution concepts, including 800 VDC architectures**, which allow higher power density and improved efficiency closer to the IT load. In this context, **solid-state transformers (SSTs)** and hybrid architectures are increasingly


## Looking further ahead, nuclear energy—particularly small modular reactors (SMRs)—is becoming part of the strategic discussion for very large, AI-driven campuses

being explored as potential enablers of more compact, controllable, and digitally integrated interfaces between the grid and the data hall, while conventional transformers continue to evolve toward higher power ratings and higher primary voltages.

In addition, the industry is placing greater value on **reconnectable transformer platforms** that enable a single design to support multiple grid-connected voltage scenarios and provide additional high-voltage connectivity options when interconnection requirements evolve late in the project cycle. These platforms are also proving to be **reliable options when integrated into mobile substations**, supporting temporary or bridging power needs—an area seeing growing demand across the industry.

At the same time, data centers are increasingly being designed as **energy hubs rather than passive grid-connected facilities**. Grid constraints, interconnection delays, and sustainability targets are accelerating the integration of on-site generation, energy storage, and grid-interactive capabilities. In the near to medium term, **gas-fired generation** is playing a growing role as a source of firm, dispatchable power that can support rapid deployment and continuous operation. Looking further ahead, **nuclear energy—particularly small modular reactors (SMRs)**—is becoming part of the strategic discussion for very large, AI-driven campuses, given their ability to deliver carbon-free, 24/7 baseload power with a compact footprint. While regulatory and timeline considerations mean SMRs are not an immediate solution for most projects, they represent a credible long-term pathway for dedicated, utility-scale digital infrastructure.

These developments reinforce the need for transformer technologies that can operate reliably across **hybrid architectures**, support **bidirectional power flows**, and remain stable under a wider range of operating conditions than traditional applications. Transformers increasingly serve as the electrical backbone that connects generation, grid interaction, storage, and IT loads into a coherent system.

Another defining trend is **deep digital integration**. As data center campuses scale into the multi-gigawatt range, operators need real-time visibility across thousands of electrical assets to manage risk, optimize performance, and avoid unplanned outages. Transformers are becoming rich data nodes within this ecosystem, providing continuous insight into loading behavior, thermal performance, power quality, and asset health. This information is increasingly used not only for maintenance, but also for system-level optimization and long-term planning.

Finally, **modularization and standardization** will continue to shape the industry. Global, multi-site deployment programs require power infrastructure that can be replicated predictably while remaining adaptable to local grid conditions and regulatory environments. Transformer technology is evolving toward universal design platforms that support phased expansion, faster deployment, and predictable performance as data center architectures continue to change.

Taken together, these trends signal a future in which transformers are no longer viewed simply as grid-connection assets. They are becoming **dynamic enablers of efficiency, resilience, and adaptability**, supporting new power architectures, hybrid energy models, and the operational intelligence required to scale AI-driven digital infrastructure responsibly.

## What would you say to data center developers and operators who are evaluating long-term technology partners?

When evaluating long-term technology partners, my advice to data center developers and operators is to look beyond today's specifications and focus on **who can evolve with you over the full lifecycle of the infrastructure**.

The data center industry is entering a multi-year supercycle where power availability, execution speed, and resilience will define success. Long-term partners must demonstrate not only technical excellence, but also the **industrial scale, supply-chain resilience, and program discipline** required to deliver consistently as projects grow in size and complexity. Power systems are increasingly the critical path, and partners need to be able to commit capacity, manage risk, and support multi-site deployment models with predictable outcomes.

**The data center industry is entering a multi-year supercycle where power availability, execution speed, and resilience will define success**

Equally important is **system-level thinking**. Data centers are no longer isolated facilities; they are becoming integrated energy platforms that combine grid interconnection, on-site generation, storage, and advanced power architectures. Developers should seek partners with deep experience across utility-grade infrastructure, generation integration, and digital power systems—capable of supporting both current designs and emerging architectures as AI workloads continue to evolve.

Another key consideration is the ability to **co-innovate**. Technology cycles in data centers are accelerating, particularly around power density, efficiency, and digitalization. Long-term partners should have the R&D depth and engineering agility to listen closely, develop solutions collaboratively, and industrialize those innovations globally—without fragmenting designs or introducing execution risk.

Finally, lifecycle commitment matters. As facilities scale and operate closer to their limits, transparency, service capability, and digital insight become essential. Partners should be able to support assets over decades, providing visibility, predictive intelligence, and upgrade pathways that protect availability and future optionality.

In short, the right long-term partner is one that brings **confidence**—confidence that infrastructure can be delivered at scale, operated reliably, and adapted over time as power architectures, energy models, and AI requirements continue to change.

## As we close, what gives you the most confidence about the future of digital infrastructure as power systems and data centers continue to converge?

What gives me the most confidence is that the industry has clearly shifted from talking about ambition to **solving the hard, practical problems at scale**. Power and data centers are no longer being treated as separate domains; they are being planned, engineered, and operated as a single system—and that is a critical maturation point.

We are seeing progress and much stronger collaboration across the ecosystem: data center developers, utilities, equipment suppliers, and technology providers are working together earlier and more transparently to address grid constraints, energy availability, and delivery timelines. That level of coordination is essential, and it is already translating into more realistic designs, hybrid energy models, and infrastructure that can actually be built and operated reliably in today's environment.

I am also encouraged by the **pragmatism around power strategies**. Rather than relying on a single solution, the industry is embracing diversified approaches—combining grid connections, on-site generation, storage, and increasingly flexible operating models. This reflects a more mature understanding of how to balance speed, reliability, sustainability, and cost, especially as AI workloads continue to scale.

Finally, there is a growing recognition that execution capability matters as much as innovation. The technologies exist—or are rapidly emerging—but success will depend on standardization, disciplined program execution, on-time delivery and long-term partnerships that can deliver consistently under pressure. The fact that the industry is now focusing on these fundamentals gives me confidence that digital infrastructure can continue to scale responsibly, even as power systems and data centers become more tightly interconnected. ■