

Whitepaper

Building the Foundation for AI at Scale

Powering the Future of AI Data Center Development

HITACHI



KJ Joshi

Chief Business Officer and Head of Data Center Business, Hitachi, Ltd.

KJ Joshi is a data-driven, results-oriented business leader with 25 years of global experience spanning strategy consulting, cloud and AI data center infrastructure services. He has held roles at McKinsey and Bain, as well as executive leadership positions at Equinix and Dell. He currently leads strategy and growth initiatives across Hitachi's Data Center portfolio.



Waseem Ahmad

Vice President of Data Center Business, Hitachi, Ltd.

Waseem Ahmad is Vice President of Data Center Business at Hitachi Americas, focused on scaling next-generation data center and digital infrastructure ecosystems across AI, power, and digitalization. He leads strategic growth initiatives spanning power readiness, AI-scale infrastructure, data residency, and end-to-end digital solutions. With a foundation in electrical engineering and an MBA, Waseem brings experience across the U.S. and Japan, combining infrastructure expertise with commercial strategy and executive decision-making.

Contents

1 — Executive Summary	4
2 — AI Workloads Demand a Fundamentally Different Infrastructure Model	5
2.1 AI Is Driving a New Scale of Data Center Development	
3 — Power Has Become the Defining Constraint	6
3.1 AI Is Accelerating Demand for Electrical Infrastructure	
3.2 Grid Timelines Are Setting the Pace of Deployment	
4 — Traditional Development Models Are No Longer Sufficient	7
4.1 Power Must Move to the Front of the Development Lifecycle	
4.2 A Power-First Development Model	
5 — The Strategic Shift: Power-First Infrastructure	8
5.1 Diversified Energy Architectures Support AI at Scale	
5.2 Grid-to-Rack Integration and Digital Orchestration	
6 — Accelerating AI Infrastructure Delivery	9
6.1 Industrialized Construction Is Changing Infrastructure Delivery	
6.2 Standardization Enables Faster, More Scalable Growth	
6.3 Coordinated Execution Reduces Delivery Risk	
6.4 The One Hitachi Advantage	
7 — Preparing for the Next Era of AI Infrastructure	12
7.1 Building AI Infrastructure for Long-Term Growth	
7.2 An Integrated Strategy for the AI Era	
8 — Executive Action Summary	13
8.1 The Bottom Line	



1. Executive Summary

Power availability is emerging as the primary constraint to AI data center growth.

The rapid growth of artificial intelligence is driving unprecedented demand for data center capacity, but power availability is emerging as the primary constraint to that growth. While AI computing infrastructure can be designed and deployed at an accelerating pace, the electrical infrastructure required to support it cannot always keep up. As a result, power is becoming a defining factor in where AI data centers are built, how quickly they can be deployed, and how they are planned for long-term growth.



The growing disconnect between AI deployment timelines and grid readiness is reshaping development decisions.

Addressing this challenge requires a new model for AI infrastructure development. Energy availability, grid readiness, infrastructure architecture, and execution must be considered together from the earliest stages of planning. Diversified energy architectures, digital orchestration, standardized infrastructure, and industrialized delivery are enabling organizations to build more resilient, scalable, and adaptable AI environments.

Competitive advantage in the AI era will increasingly depend on the ability to align energy, infrastructure, digital capabilities, and execution into a coordinated strategy. This approach positions industry leaders to deploy AI capacity faster, reduce delivery risk, and scale with confidence.

2. AI Workloads Demand a Fundamentally Different Infrastructure Model

Artificial intelligence is reshaping data center development. Infrastructure models built for enterprise and cloud computing must now support significantly greater compute density, power capacity, cooling performance, and operational complexity. Scaling existing facilities alone is no longer enough. AI-ready environments require a different approach to how data centers are designed, powered, cooled, and operated.

Power distribution, cooling, networking, digital operations, and energy management are also becoming more closely connected. Supporting sustained, high-performance computing requires these systems to work as a coordinated infrastructure environment rather than as independent technologies.

SHIFT IN PHYSICAL COMPUTE DENSITY¹

LEGACY CLOUD DENSITIES	~10–15 kW per rack
CURRENT AI DEPLOYMENTS	80–150+ kW per rack next-gen designs exceeding 200 kW
DEPLOYMENT FOOTPRINTS	5–20 MW training pods up to 100–500+ MW campus-scale facilities
INFRASTRUCTURE INTEGRATION	Higher densities require advanced liquid cooling, specialized high-capacity power distribution, and AI-driven digital twin operations to manage complexity.

2.1 AI Is Driving a New Scale of Data Center Development

The physical requirements of AI workloads illustrate the magnitude of this change. Rack densities have increased from

approximately 10–15 kW in legacy cloud environments to 80–150+ kW in current AI deployments, with emerging designs expected to exceed 200 kW per rack. AI facility scale is increasing alongside compute density. Traditional 10–20 MW data centers are giving way to 100–500+ MW AI campuses¹ that require dedicated electrical infrastructure, modular power architectures, advanced cooling systems, and long-term utility coordination.

These demands are changing more than facility size. Electrical systems must accommodate substantially higher loads, advanced liquid cooling is becoming essential for high-density GPU environments, and networking technologies must move data efficiently across increasingly large AI clusters. Digital twins, AI-driven operations, and grid-interactive capabilities are also expanding the tools available to manage increasingly complex infrastructure.

The result is a new class of data center in which compute capacity and the infrastructure supporting it must be planned together. AI is not simply increasing demand for data centers. It is redefining what a data center must be.

That shift raises a critical question: Can the infrastructure that powers AI keep pace with the infrastructure being built to support it?

“AI is not simply increasing demand for data centers. It is redefining what a data center must be.”

U.S. POWER DEMAND FROM DATA CENTERS

CURRENT LOAD (2025)

~65-70 GW

REQUIRED BY (2030)

>165 GW

THERMAL CONSTRAINTS

AIR COOLING LIMIT


**~20-25
KW/RACK**

LIQUID REQUIREMENT

**>50
KW/RACK**

NEAR-TOTAL LIQUID

**~150
KW/RACK**



2025

AI demand is already reshaping data center infrastructure priorities.

2030

Expansion increasingly shifting toward power-advantaged regions.

GROWTH OUTLOOK

Power availability is emerging as the key determinant of deployment.

Figure 1. AI-driven growth is increasing power demand, changing infrastructure requirements, and shifting deployment toward regions with greater electrical capacity.

3. Power Has Become the Defining Constraint

Power availability is increasingly determining where, when, and how quickly AI capacity can be deployed. While data center design and construction continue to advance rapidly, the electrical infrastructure required to energize these facilities operates on a different timeline (Figure 1).

3.1 AI Is Accelerating Demand for Electrical Infrastructure

AI workloads are creating electrical demand at a scale that extends far beyond supplying additional electricity to individual facilities. Larger GPU clusters and campus-

scale deployments require expanded substations, transformers, transmission capacity, utility interconnections, and supporting infrastructure.

Higher electrical density also creates parallel thermal requirements. As rack densities move beyond the practical limits of traditional air cooling, liquid cooling architectures are becoming increasingly important for supporting high-density AI environments.

The challenge is therefore not simply generating additional power. It is ensuring sufficient electrical capacity can be delivered to the right location when new AI infrastructure is ready to operate.

3.2 Grid Timelines Are Setting the Pace of Deployment

AI data center projects are often expected to be ready for service within approximately 18–36 months, while utility upgrades, including substations, transmission infrastructure, and grid interconnections, may require five years or longer before sufficient electrical capacity is available.^{1 2 3 4}

That timeline gap changes the development equation. Completing construction no longer guarantees that a facility can begin operations. Deployment schedules may instead depend on utility readiness, interconnection timing, equipment availability, and regional grid capacity.

Power availability is consequently moving upstream in the decision-making process. Site selection, capital planning, expansion strategies, and project schedules increasingly depend on a clear understanding of when sufficient electrical capacity can be delivered.

For AI infrastructure, the critical development question is shifting from “Where should we build?” to “Where can sufficient power be secured on the timeline the business requires?”

4. Traditional Development Models Are No Longer Sufficient

Traditional data center development followed a relatively predictable sequence: select a site, design the facility, secure utility service, construct the data center, and begin operations. Power has always been a critical consideration, but it was generally viewed as an engineering requirement that followed site selection rather than a strategic factor shaping

“Power-first development is more than a change in project sequencing.”

where and how facilities were developed. That model worked when electrical infrastructure could expand alongside facility development and computing demand grew at a more predictable pace.

AI has disrupted that sequence.

KEY DRIVERS FOR FRONT-LOADING POWER INFRASTRUCTURE

01 GRID & INTERCONNECTION TIMELINES

Utility interconnections, transmission upgrades and substation expansion can influence project schedules.

02 EARLY RISK VISIBILITY

Evaluating capacity and infrastructure readiness earlier provides greater visibility into potential constraints before major commitments are made.

03 STRATEGIC SITE SELECTION

Available and future electrical capacity becomes a foundational consideration alongside traditional site-selection factors.

4.1 Power Must Move to the Front of the Development Lifecycle

This changes the role of power planning. It becomes a foundational input into development strategy rather than a downstream engineering activity (Figure 2).

4.2 A Power-First Development Model

Power-first development aligns site selection, infrastructure planning, facility

design, and long-term scalability around available and future electrical capacity. The objective is not simply to secure enough power for initial operations, but to understand how infrastructure decisions made today will affect expansion as AI workloads grow.

This model also creates greater visibility into project risk before construction begins. Utility timelines, regional capacity, infrastructure constraints, and expansion requirements can be evaluated alongside traditional development considerations rather than after sites and capital commitments have narrowed the available options.

Power-first development is more than a change in project sequencing. It moves energy considerations to the beginning of the development process, allowing power availability, infrastructure readiness, and long-term scalability to inform decisions before major investments are made.

5. The Strategic Shift: Power-First Infrastructure

A power-first approach changes more than when power is considered; it changes how AI infrastructure is designed. Supporting AI at scale requires energy architectures and digital capabilities that improve how power is delivered, managed, and optimized across the data center while increasing resilience, flexibility, and long-term scalability.

5.1 Diversified Energy Architectures Support AI at Scale

Supporting high-density AI environments increasingly requires a diversified approach to energy. Rather than relying on a single power source, operators are combining utility power with on-site generation, battery energy storage systems (BESS),

“Competitive advantage will increasingly depend on building energy architectures that are resilient, coordinated, and capable of adapting as infrastructure requirements change.”

and emerging technologies to create more resilient and adaptable energy architectures.


These approaches provide greater flexibility to balance reliability, operational resilience, sustainability, and future expansion. Coordinating multiple energy resources also allows operators to respond more effectively to changing workload demands while strengthening the long-term performance of AI infrastructure.

Battery energy storage systems play a strategic role within these architectures by supporting peak-demand management, improving operational flexibility, and strengthening resilience as infrastructure requirements continue to grow.

5.2 Grid-to-Rack Integration and Digital Orchestration

Power-intensive AI environments require greater visibility across the entire electrical

TRADITIONAL DEVELOPMENT MODEL



POWER-FIRST DEVELOPMENT MODEL




Figure 2. AI Is Changing the Sequence of Data Center Development.
Traditional Development Model vs. Power-First Development Model

chain, from the utility connection through substations, electrical distribution, cooling systems, and IT equipment. Understanding how decisions in one layer affect overall infrastructure performance is becoming essential to maintaining reliable operations at scale.

Digital orchestration provides that visibility. Simulation, digital twins, and AI-driven operational intelligence enable operators to monitor infrastructure holistically, optimize energy performance, anticipate operational constraints, and make more informed decisions throughout the infrastructure lifecycle. Connecting energy systems and IT operations into a coordinated operating environment improves resilience, operational efficiency, and scalability.

Competitive advantage will increasingly depend on building energy architectures


that are resilient, coordinated, and capable of adapting as infrastructure requirements change. Organizations that make this strategic shift will be better positioned to support AI growth with greater confidence and long-term flexibility.

6. Accelerating AI Infrastructure Delivery

Infrastructure strategy creates value only when it translates into deployed capacity. As AI development accelerates, execution has become a competitive differentiator. Delivering increasingly complex electrical and digital infrastructure within compressed timelines requires new approaches that improve speed, consistency, and scalability while reducing delivery risk.

01

Pre-Engineered Design



02

Factory Manufacturing



03

Transportable Units



04

Rapid On-Site Installation

Figure 3. Standardized Power Modules Accelerate AI Infrastructure Deployment

Addressing these challenges requires expertise that spans the entire infrastructure lifecycle, from power systems and digital technologies to industrialized delivery and long-term operations. The following approaches illustrate how integrated capabilities are helping accelerate AI infrastructure deployment while improving execution certainty.

6.1 Industrialized Construction Is Changing Infrastructure Delivery

Traditional site-built construction is giving way to more industrialized delivery models, where engineering, manufacturing, and system assembly occur in controlled factory environments before equipment arrives on site. This approach reduces

dependence on field labor, improves quality through standardized manufacturing, and compresses deployment schedules for increasingly complex AI infrastructure.

Factory-built, pretested power modules further accelerate deployment by shifting a significant portion of construction and integration work away from the project site. The result is greater consistency across facilities, reduced execution risk, and faster time to capacity (Figure 3).

6.2 Standardization Enables Faster, More Scalable Growth

Standardized infrastructure provides a repeatable foundation for AI expansion. Rather than engineering every deployment

from the ground up, repeatable infrastructure components can be adapted to individual site requirements while maintaining consistent performance and quality.

This approach supports phased expansion as demand grows, allowing operators to scale from megawatt- to gigawatt-scale campuses using proven infrastructure designs. Reducing engineering complexity while increasing deployment consistency enables faster expansion with greater confidence.

6.3 Coordinated Execution Reduces Delivery Risk

Delivering AI infrastructure at scale depends on close coordination across utilities, hyperscalers, engineering firms, OEMs, EPC contractors, and infrastructure providers. Each plays a critical role in determining project readiness, deployment schedules, and long-term operational success.

“The decisions made today will shape not only how quickly AI capacity can be deployed, but how resilient, adaptable, and scalable that infrastructure will be as AI continues to evolve.”

Aligning these stakeholders earlier in the development process improves visibility into project dependencies, reduces execution risk, and creates more predictable delivery timelines. It also enables energy systems, electrical

6.4 THE ONE HITACHI ADVANTAGE

Successfully delivering AI infrastructure requires expertise that extends across power systems, digital technologies, industrial infrastructure, and long-term operations.

Hitachi brings these capabilities together through expertise spanning energy, digital innovation, industrial systems, and infrastructure delivery. This breadth of expertise enables a more coordinated approach to developing, powering, and operating next-generation AI data centers. As infrastructure becomes more interconnected, the ability to

coordinate technologies, disciplines, and ecosystem partners becomes increasingly valuable. By aligning infrastructure strategy with execution, Hitachi simplifies complexity, improves delivery confidence, and supports the long-term scalability required for AI-ready infrastructure.

This integrated approach reflects Hitachi's Think Beyond the Data Center ambition, bringing together energy, infrastructure, digital capabilities, and execution to address the demands of AI-scale infrastructure.

infrastructure, digital technologies, and physical assets to be planned and delivered as a coordinated program rather than as independent workstreams.

7. Preparing for the Next Era of AI Infrastructure

Artificial intelligence is redefining the relationship between computing infrastructure and the systems that support it. Competitive advantage will increasingly depend not only on deploying more compute capacity, but on building the energy, infrastructure, and operational foundation required to sustain it.

As AI demand continues to accelerate, leaders will need to make infrastructure decisions in an environment where power availability, grid readiness, and execution capabilities increasingly shape where and how AI capacity can be deployed. Responding to this shift will require infrastructure strategies that anticipate future constraints rather than react to them.

7.1 Building AI Infrastructure for Long-Term Growth

Preparing for future AI demand begins with the infrastructure decisions made today. Planning exclusively for current requirements risks creating constraints as AI workloads, power demands, and operational expectations continue to increase.

Long-term success depends on infrastructure strategies that balance immediate deployment objectives with the flexibility to accommodate future growth. Resilient energy systems, adaptable infrastructure architectures, and scalable delivery models provide a stronger foundation for supporting future AI applications.

7.2 An Integrated Strategy for the AI Era

No single technology or infrastructure investment can address the demands of AI at scale.

Success will increasingly depend on aligning energy infrastructure, digital capabilities, industrial expertise, and disciplined execution into a coordinated strategy that spans the infrastructure lifecycle. Planning these capabilities together improves visibility into infrastructure risk, strengthens investment decisions, and creates greater flexibility as AI infrastructure requirements continue to evolve.

The next generation of AI will be shaped not only by advances in computing, but by the infrastructure strategies that make those advances possible. As power availability becomes an increasingly important determinant of where and how AI infrastructure can scale, leaders will need to view energy, infrastructure, and execution as strategic business priorities rather than operational considerations. The decisions made today will shape not only how quickly AI capacity can be deployed, but how resilient, adaptable, and scalable that infrastructure will be as AI continues to evolve.

8. Executive Action Summary

Key operational shifts required to build resilient, scalable, and adaptable AI data center environments.

INFRASTRUCTURE PILLAR	TRADITIONAL MODEL	POWER-FIRST MODEL
Primary Site Metric	Land acquisition and basic utility access	Grid capacity, transmission readiness, and interconnection timelines
Development Sequence	Site selection → Facility design → Power planning	Power & grid assessment → Site selection → Integrated design
Cooling Architecture	Standard air cooling (~10–20 kW/rack)	High-density liquid cooling & hybrid thermal readiness (80–150+ kW/rack)
Delivery Model	Sequential, site-built field construction	Pre-engineered, industrialized modular deployment
Operations & Control	Isolated facility management	Digital twin orchestration & grid-interactive energy management

8.1 THE BOTTOM LINE

Competitive advantage in the AI era will not be defined by compute volume alone but by the ability to align energy, digital capabilities, and industrialized execution into a unified strategy.

Sources:

1. Hitachi, Ltd. *Powering the Next Generation AI Data Center*. NVIDIA GTC Executive Dinner Presentation. March 18, 2026.
2. MarketsandMarkets. *Powering the Next Generation of AI Data Centers: Executive Insights from the Hitachi Roundtable*. March 18, 2026.
3. Idaho National Laboratory. *Data Center Playbook*. 2026. Available at: <https://inl.gov/content/uploads/2026/04/data-center-playbook.pdf>
4. Ernst & Young LLP. *Energy and Data Centers Report*. 2025. Available at: <https://www.ey.com/content/dam/ey-unified-site/ey-com/en-us/campaigns/data-centers/documents/ey-energy-data-center-report-vf.pdf>

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

hitachi.com

©Hitachi, Ltd 2026. All rights reserved.
Specifications subject to change without notice.