



## Data Centers:

**Breathtaking Drive: Exhausting Energy, Water and Brains**

August 2025



## **M Capital Group**

[www.mcapital-group.com](http://www.mcapital-group.com)

**“Our Word is Our Bond”**

<b>NEW YORK</b>	<b>LONDON</b>	<b>DUBAI</b>
1330 Av of the Americas Suite 23 New York, NY 10019 USA P: +1 212 634 6831	The Leadenhall Building Level 30 London, EC3V4AB UK P: +44 207 256 4246	Emirates Towers PO Box: 31303 Dubai UAE P: +971 4 319 746



## DISCLAIMER NOTICE

---

This memorandum has been prepared by M Capital Group (“MCG”). In no case is this document a solicitation, request and/or invitation for capital in any shape or form; or any securities offering in any legal jurisdiction.

This document is not intended to form the basis of any investment decision and should not be considered as a recommendation by MCG. This document does not constitute an offer to sell or an invitation to purchase or tender.

None of the information contained in this document has been verified by MCG. No party shall have any right of action against MCG or any other person in relation to the accuracy or completeness of the information contained in this document or any other written or oral information made.

In furnishing this document, MCG undertakes no obligation to provide the recipient with access to any additional information or to update this document or additional information or to correct any inaccuracies.

# TABLE OF CONTENTS

<b>CONTACTS</b>	<b>ii</b>
<b>DISCLAIMER NOTICE</b>	<b>iii</b>
<b>TABLE OF CONTENTS</b>	<b>iv</b>
<b>EXECUTIVE SUMMARY</b>	<b>1</b>
<b>INDUSTRY CONVULSION</b>	<b>5</b>
Unstoppable Dynamics	5
Key Trends and Indicators	14
Promising Growth Dynamics	18
<b>GLOBALIZATION</b>	<b>20</b>
Introduction	20
Global Drivers	20
US Still Leads	27
Europe Take Note	29
Asia Well Off	31
Middle East Taking Off	33
Africa Promising	36
Latin America Accelerates	37
<b>LOOKING FORWARD</b>	<b>41</b>

## EXECUTIVE SUMMARY

### INDUSTRY CONVULSION

M Capital Group (“MCG”) views this moment in the data center industry as a once-in-a-generation structural realignment. This report goes far beyond the headlines of “AI growth” or “digital acceleration”, it maps out how the physical architecture of intelligence is being redrawn at global scale. What excites MCG most is that this transformation is not simply about more servers or more cloud, it is about reconfiguring the way economies, governments, and capital interact in the age of artificial intelligence.

The global data center industry is undergoing a structural transformation, driven by unprecedented compute demand from generative AI, sovereign cloud mandates, and energy constraints. Between 2023 and 2030, global data center capacity is projected to grow from 60 GW to over 219 GW, with AI-specific workloads accounting for more than 70% of new demand. Electricity consumption, already at 176 TWh in the U.S. (2024), is forecasted to reach up to 580 TWh by 2028. These dynamics are forcing a redesign of infrastructure: legacy CPU-focused facilities are giving way to AI-grade campuses optimized for 50–600 kW per rack and liquid cooling technologies. Capital expenditure is surging in tandem; Meta alone plans to invest over US\$100 billion in AI-aligned infrastructure over the next decade.

## Key Trends and Indicators

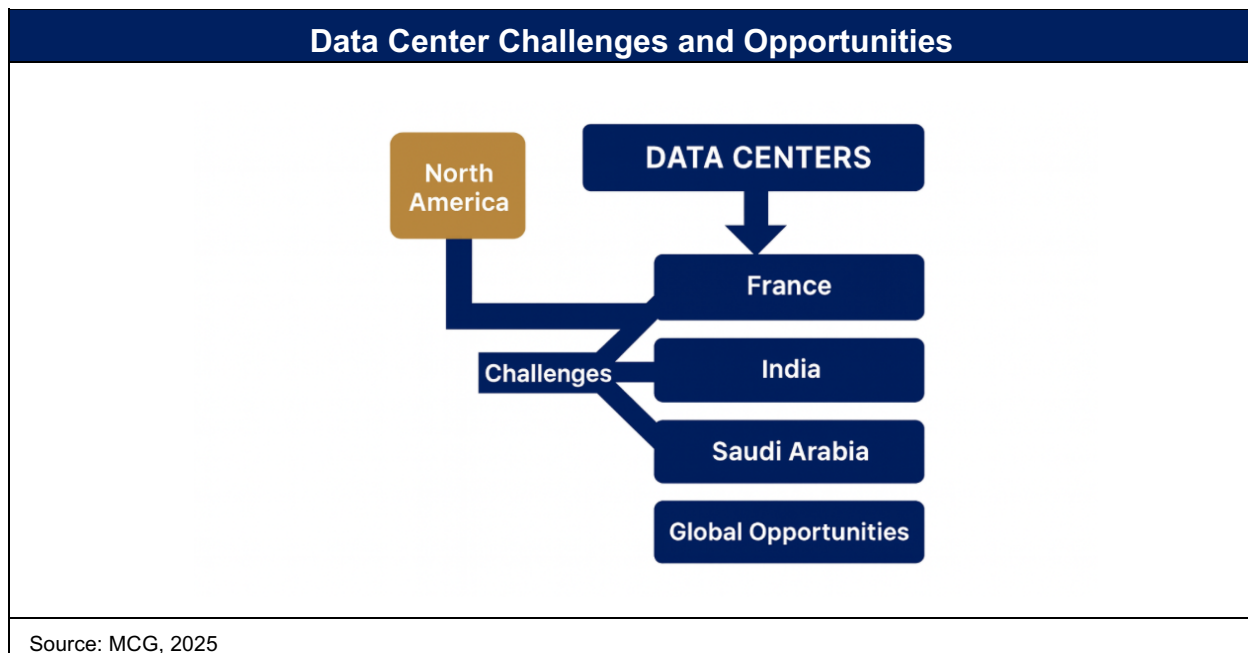
AI is the single most disruptive force shaping infrastructure. GPU power demand has grown 20x since 2017, and AI data centers now consume up to 3–5x more electricity per rack than traditional setups. Campus builds are escalating to multi-gigawatt scale, such as OpenAI/Oracle's Stargate (4.5 GW). Cooling, once a secondary concern, is now central to viability, driving a market projected to exceed US\$100 billion by 2035. Power scarcity has become a bottleneck in dense markets, forcing operators to secure multi-source PPAs and on-site generation years in advance. Regulation and ESG mandates are transforming both investor expectations and design logic, with Tier III+ compliance, sovereign alignment, and renewable integration now baseline requirements.

## Promising Growth Dynamics

The AI revolution is catalyzing the next multi-decade cycle of digital infrastructure growth. AI data center workloads are growing at a 30–40% CAGR through 2030, with energy needs rising in lockstep. Sovereign policy is unlocking new markets: the Middle East is seeing 27.9% annual growth, backed by national AI visions. Capital markets are responding accordingly, data center financing in the U.S. alone doubled from US\$30B (2024) to US\$60B (2025), while green bonds and ESG-linked investments surge. The market capitalization of green bonds has reached US\$2.9 trillion by the end of 2024, nearly a sixfold increase since 2018. Infrastructure is increasingly viewed as a real asset class, with long-duration returns and inflation hedges. AI-grade infrastructure, modular, high-density, and compliant, is now the gold standard.

## GLOBALIZATION

Data centers have become geopolitical assets, no longer mere IT utilities. National sovereignty and regulatory fragmentation are reshaping global deployment strategies. Over 60 countries now require domestic data localization, and AI is fast becoming a national policy pillar. In 2025, the global data center market is projected to reach US\$387 billion, with Asia and the Middle East growing fastest. India is targeting 100+ new centers by 2025; Saudi Arabia and Qatar are embedding AI into national infrastructure via sovereign cloud zones. Meanwhile, U.S. hyperscalers continue to lead in CAPEX, innovation, and power density, with North America still accounting for over 54% of global hyperscale capacity.



As data centers become strategic national assets, North America's innovation flows toward fast-growing markets like India, Saudi Arabia, Qatar, and France, driven by AI, localization mandates, and sovereign cloud initiatives amidst rising geopolitical and regulatory complexities.

## CONCLUSION

MCG anticipates a new AI-industrial era where computing, regulation, and sustainability converge to redefine global infrastructure. Tomorrow's competitive advantage will hinge not just on capacity, but on trust: compliance, power resilience, and sovereign fit. From the sovereign AI strategies of the Gulf to the regulatory leadership of Europe and the infrastructure boom in Asia, the winners will be those who localize, optimize, and operationalize at scale. The cloud era was software-native. The AI era will be infrastructure-native, and its blueprint is unfolding now.

The most striking insight from MCG is how quickly data centers have evolved from backend IT assets into frontline instruments of sovereignty, security, and strategic advantage. Around the world, MCG sees governments rewriting laws to keep critical data onshore. Hyperscalers are being forced into sovereign-compliant infrastructure partnerships. Infrastructure capital is moving not just toward the largest markets, but toward the most “trusted” and jurisdiction-ready opportunities. AI is the accelerant, but sovereignty, regulation, and sustainability are the scaffolding. This is no longer optional. It is the baseline.

M Capital Group's analysis shows that Asia markets are not only building sovereign cloud infrastructure, but they are also leapfrogging legacy paradigms by directly embedding national AI strategies into their data center blueprints. Meanwhile, in North America and Europe, trillion-parameter models are demanding facilities with 3–5x the power density of the last decade. Operators are now racing to secure grid access, long-term PPAs, and water-efficient cooling at a pace never seen in infrastructure. MCG believes this “AI-industrial era” is the next frontier for sovereign capital, infrastructure funds, and digital real estate developers alike.

For MCG, the strategic implications are enormous. This industry report outlines not just where the market is headed, but how to position ahead of it. MCG sees emerging platforms that combine high-density AI readiness, sovereign compliance, ESG alignment, and public-private financing as the new gold standard. Whether it's enabling secure compute for central banks, building national AI campuses, or investing in regional edge facilities governed by local ministries, MCG is moving into a world where infrastructure will be judged by trust, not just throughput. What makes this report essential reading is that it does not simply admire the growth. It dissects it. It links demand to regulation, regulation to design, and design to capital. For investors, policymakers, and developers who recognize that data is the new oil, but who understand that oil now requires an AI refinery, MCG believe this report is your blueprint.

The AI era will not be cloud-native. It will be infrastructure-native. And those who can localize, optimize, and operationalize at scale will define the next two decades of digital power.

## INDUSTRY CONVULSION

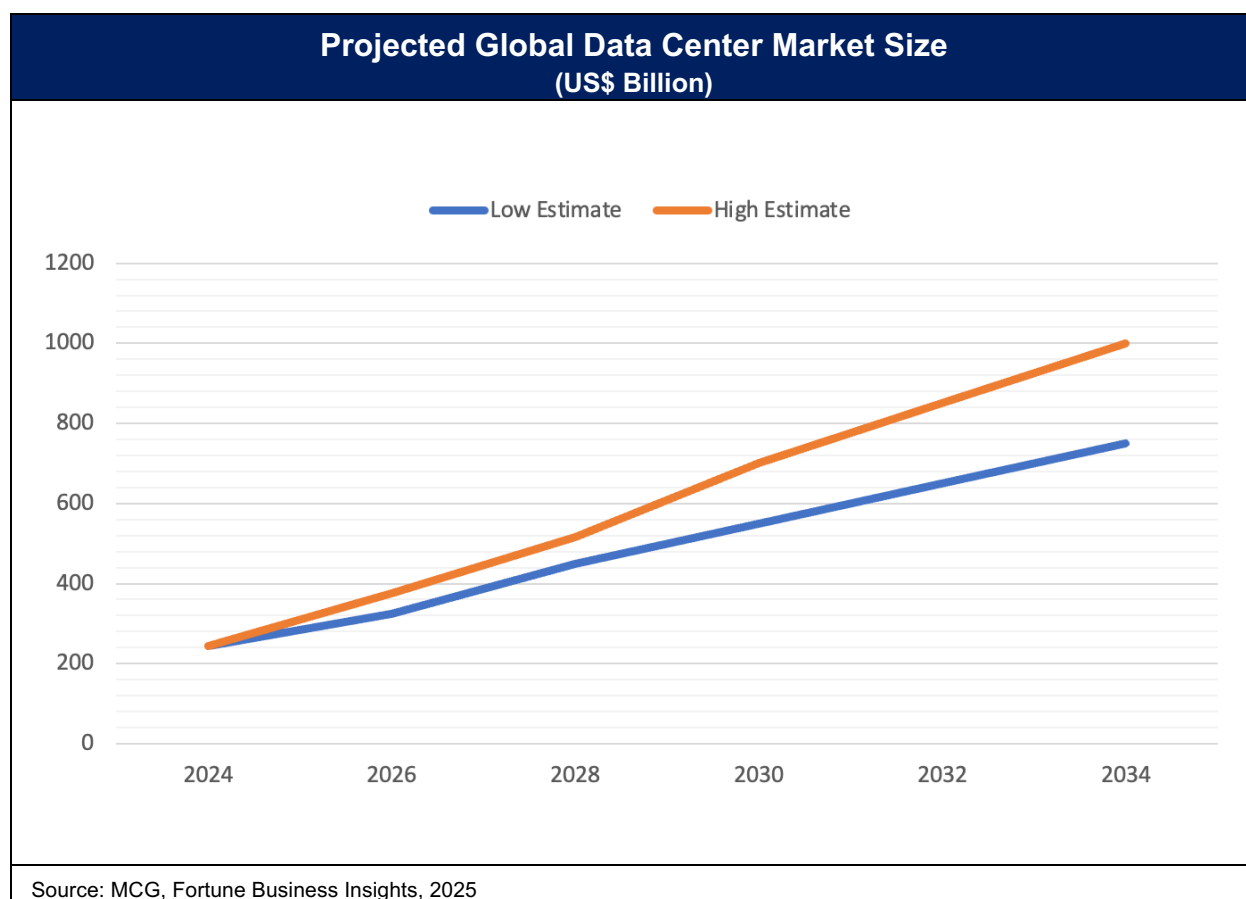
## UNSTOPPABLE DYNAMICS

### Background

#### A Foundational Pillar of the Digital Economy

The data center industry has emerged as one of the most vital infrastructure sectors globally, powering the backbone of cloud computing, artificial intelligence, enterprise IT, e-commerce, and digital media platforms. As digital transformation accelerates across every sector, the demand for secure, high-performance, and low-latency computing environments is growing at an unprecedented pace. This is driven not only by traditional enterprise workloads, but also by structural shifts in technology, including the proliferation of generative AI, the expansion of Internet of Things (“IoT”) networks, and the migration of software applications to cloud-native architectures.

The sector’s momentum is reinforced by sustained capital investment, geopolitical prioritization, and the rising role of data as a strategic asset. In 2024, the global data center market was valued at approximately US\$243 billion, and many forecasts suggest it could grow to between US\$585 billion and US\$1 trillion by year 2032 to year 2034, representing a compound annual growth rate of over 11%. This growth reflects not only rising compute demand but also a broader reclassification of data centers as critical infrastructure. Increasingly, they are viewed alongside utilities and transportation networks as essential to economic resilience, national security, and industrial competitiveness in the digital age.



### The Rise of Hyperscale and Cloud Platforms

Hyperscale providers such as Amazon Web Services, Microsoft Azure, and Google Cloud are central to this expansion. In the U.S., hyperscale and colocation data centers now account for nearly 80% of installed server capacity, a figure expected to exceed 90% by 2028. These large-scale deployments have also driven record levels of capital expenditure. For instance, Meta plans to invest hundreds of billions into compute and AI-supportive data center infrastructure.

- Equinix: A leading global interconnection and colocation provider that serves enterprise clients and cloud on-ramps, with a strong presence in edge and Tier II markets.
- Digital Realty: While sizable, it operates independently and caters to enterprise, cloud, and hybrid IT needs across global metro areas.
- CyrusOne: A U.S.-based colocation provider focused on delivering scalable, flexible infrastructure for high-density workloads and enterprise customers.

- Stack Infrastructure: Specializes in build-to-suit and hyperscale-ready data centers, particularly for emerging AI and machine learning needs.
- Aligned Data Centers: Known for its energy-efficient, sustainable infrastructure design and rapid deployment capabilities.
- EdgeConneX: Focuses on edge data centers located closer to end users, supporting low-latency applications and regional demand.
- TierPoint: Serves mid-market and regional enterprise clients with hybrid IT and disaster recovery solutions in underserved U.S. markets.

These independent and regional operators are becoming increasingly vital for meeting demand outside major cloud regions, supporting edge workloads, latency-sensitive applications, and enterprise hybrid deployments.

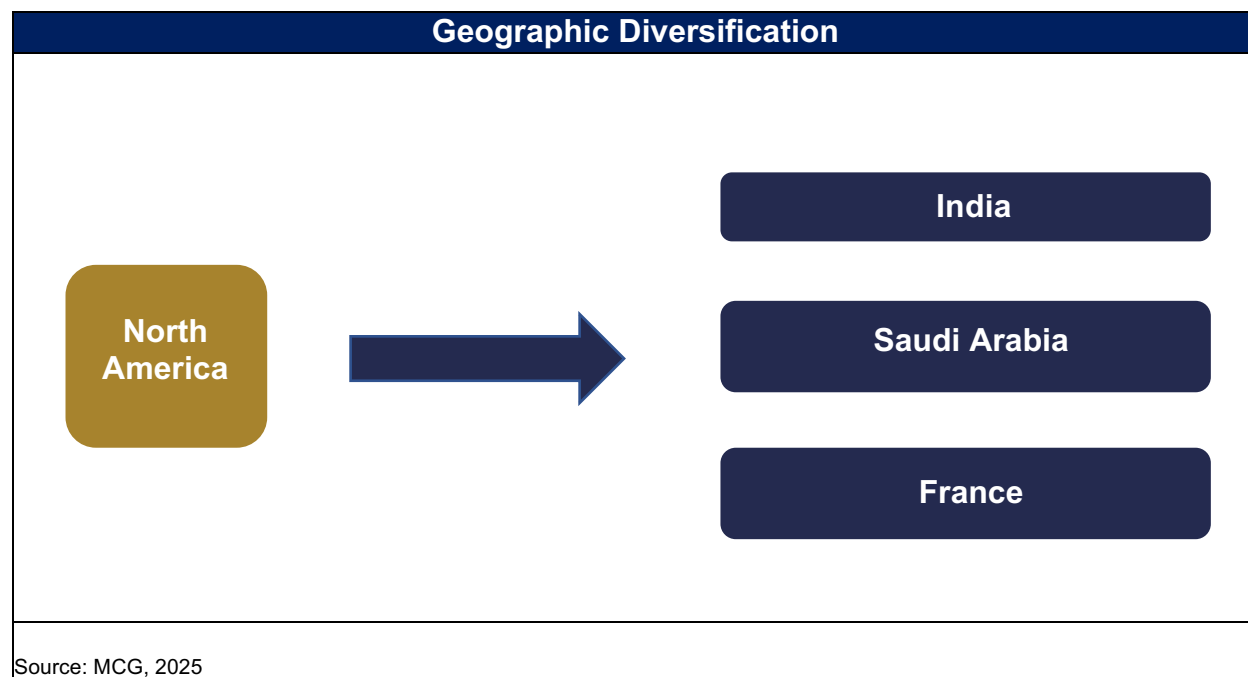
### **Energy Demand and Sustainability Challenges**

As capacity grows, so does energy consumption. In the U.S. alone, data centers consumed approximately 176 terawatt-hours (“TWh”) of electricity in 2023, roughly 4.4% of national usage. By 2028, this could rise to between 325 TWh and 580 TWh, or up to 12% of total demand, driven largely by accelerated AI deployment. Globally, data centers used around 460 TWh in 2022, and this figure is expected to double within a few years. A particularly sharp rise in energy demand has been observed among AI-enabled servers: in the U.S., GPU-accelerated compute workloads increased electricity use from under 2 TWh in 2017 to over 40 TWh in 2023. Recent estimates suggest AI workloads could consume up to 50% of all data center electricity by the end of 2025.

### **Efficiency Improvements and Energy Transition**

To address mounting environmental concerns, operators are improving energy efficiency and investing in renewables. The average industry Power Usage Effectiveness (“PUE”) ratio improved from ~1.6 in 2014 to between 1.3 and 1.4 in 2023. Leading operators are targeting further reductions to 1.15–1.25 by deploying technologies such as liquid cooling, modular architectures, and advanced airflow systems. At the same time, the proportion of renewable energy powering data centers has grown to 27% globally, with projections expecting more than 50% by 2030. Power purchase agreements (“PPAs”) for wind, solar, and hydroelectric sources are increasingly being used to secure clean energy supply, especially by hyperscalers.

### Geographic Diversification and Sovereign Investment



While North America maintains its installed capacity, emerging markets are becoming significant growth engines. In the Asia-Pacific region, the Middle East, and European countries like France, are pursuing ambitious digital infrastructure programs supported by sovereign capital, national policy, and rising cloud adoption. India alone is expected to see over 100 new data centers developed by 2025, backed by more than US\$2 billion in private and public investment, helped by strong demand for sovereign cloud and data localization policies.

### Future Growth Outlook

Industry analysts forecast sustained long-term growth, primarily driven by AI and high-density computing. McKinsey expects total global data center capacity to expand from 60 gigawatts ("GW") in 2023 to between 171 GW and 219 GW by 2030, with some aggressive scenarios reaching as high as 298 GW, implying a compound annual growth rate of up to 27%. Meeting this demand will require historic investments in power procurement, grid expansion, and land acquisition. In Europe alone, annual energy consumption from data centers may rise from 96 TWh to 236 TWh by 2035, prompting nearly €100 billion in annual energy and electrical infrastructure investment.

## **An Intersection of Digital and Industrial Transformation**

The data center sector now sits at the convergence of digital transformation, infrastructure finance, and the global energy transition. While it remains capital-intensive and increasingly complex, the sector also offers long-term contracted revenue streams, inflation protection, and exposure to powerful structural tailwinds. For operators, investors, and policymakers alike, data centers are no longer simply backend IT support, they are strategic assets underpinning the future digital economy.

## **AI as a Structural Infrastructure Driver**

### **A New Class of Infrastructure Demands**

Artificial Intelligence (“AI”) has rapidly evolved from a specialized computational workload to a foundational element of modern digital infrastructure. The explosion in generative AI, deep neural networks, and machine learning inference has pushed the limits of legacy data center architectures, which were originally designed for traditional cloud and enterprise IT workloads. Unlike these earlier workloads, AI training and inference demand sustained, high-throughput compute across tightly integrated clusters of GPUs or custom AI accelerators. These clusters not only require exponentially more power per rack but also necessitate advanced cooling systems, high-density interconnects, and purpose-driven internal designs.

As a result, a new class of “AI-grade” data centers has emerged, optimized for throughput per watt rather than conventional metrics like scale or redundancy. However, if the development of digital infrastructure fails to keep pace with the accelerating adoption of AI, a significant energy demand gap could arise. AI workloads could outstrip the capacity of existing power and cooling systems, leading to system bottlenecks, higher operational costs, and reduced computational efficiency. This mismatch could stall innovation, delay AI deployment at scale, and strain regional power grids, particularly in areas without sufficient renewable energy or next-generation grid readiness. Addressing this infrastructure-energy gap is therefore critical, not only for sustaining AI growth but also for ensuring energy-efficient and resilient digital ecosystems.

### **Shifts in Design and Deployment Logic**

AI-centric infrastructure is both compute-intensive and latency-sensitive. This is especially critical for real-time applications such as large language models or edge-based inference, where even microsecond delays impact performance. As a result, developers are moving beyond traditional siting logic, previously driven by access to cheap land or power, and are instead prioritizing low-latency chip-to-power pathways, node interconnect proximity, and layouts that reduce thermal bottlenecks. These shifts at the rack and floorplan level are influencing macro-level deployment strategy, driving investment in vertically integrated, high-performance data campuses.

### **Continuous Power Draw and Thermal Complexity**

Unlike cloud workloads with variable demand, AI training imposes a continuous, high baseline power load, especially during multi-day or multi-week training cycles. This operational intensity challenges conventional PUE metrics, which fail to capture the full scope of cooling and load consistency. Developers must now model facilities based on sustained power density, thermal output management, and chip-specific cooling behavior over time. Liquid cooling, immersion systems, and AI-native heat exchange design are becoming standard in high-density deployments, fundamentally changing data center economics and operational rhythms.

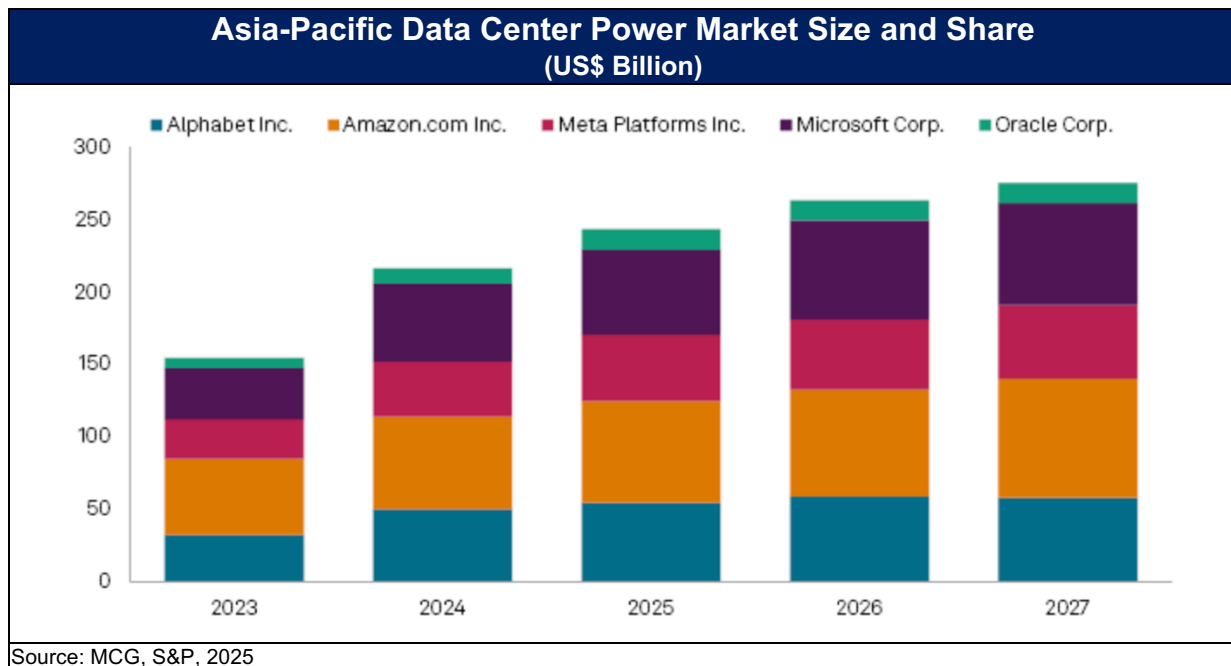
### **Capital Allocation Is Following Compute**

As AI reshapes infrastructure requirements, capital deployment strategies are evolving in parallel. Investors are no longer focused solely on real estate yield or tenant diversity, they are now evaluating AI-linked utility economics, energy procurement resilience, and long-term viability of cooling and backup systems under constant load. The result is a shift in how infrastructure assets are underwritten: from simple colocation facilities to advanced AI-grade platforms with long-term energy visibility, modular thermal capacity, and specialized tenancy profiles. This has redefined what constitutes “investment-grade” digital infrastructure in the post-AI era.

## **Hyperscale Consolidation and Strategic Self-Builds**

### **The Shift to Internalized, AI-Optimized Infrastructure**

Hyperscale operators, such as AWS, Microsoft, and Google, are accelerating a transition from leasing capacity to building and owning their own AI-optimized campuses. This self-build model is driving major changes in supply chains and financial dynamics. Hyperscalers are absorbing key inputs like power, and equipment, at scale, placing strain on electrical gear suppliers and driving up competition for grid access. Contractors must now deliver scalable, modular systems tailored to phased deployments rather than one-time builds. Local governments, in turn, are fast-tracking approvals to attract these capital-heavy projects. While cost-intensive upfront, these builds reduce long-term operating costs and offer direct control over mission-critical workloads. However, they also marginalize mid-size developers and colocation firms, pushing them toward edge deployments, sovereign partnerships, or custom solutions to stay competitive.



This chart depicts escalating annual capital expenditures among major hyperscalers, Alphabet, Amazon, Microsoft, Meta, and Oracle, specifically directed toward AI infrastructure between 2023 and 2026. From roughly US\$150 billion in 2023, total AI-linked capex swells to approximately US\$280 billion by 2027. These figures underscore the strategic pivot from leasing or colocation to constructing custom-built, AI-optimized campuses with advanced cooling, high-voltage power, and large-scale GPU clusters. Such a shift has multiple implications: first, the reliance on third-party data centers is reducing, with hyperscalers increasingly favoring in-house control over critical infrastructure. Second, site selection criteria, formerly dominated by land cost and network access, are now anchored in power availability, cooling integration, and inter-chip latency, all essential for hyperscaler-grade AI workloads.

For regional and mid-market operators, these developments create both challenges and opportunities. While markets may tighten, accessing hyperscaler supply chains means aligning with their standard, offering modular, liquid-cooled, and expandable rack infrastructure. In emerging regions, this also signifies the need for partnership models that enable compliance with local regulations while meeting global infrastructure expectations.

In summary, this capex trend redefines the data center ecosystem. It's not merely about scaling capacity, it's about creating facilities that meet the performance, scale, and resilience requirements of AI at hyperscaler levels. For infrastructure investors and developers, understanding and aligning with these evolving benchmarks is essential to navigate the competitive landscape.

## **Power and Cooling as Economic Bottlenecks**

### **Power Scarcity as a Growth Constraint**

In high-density AI environments, power is no longer a commodity, it's the core design constraint. Delays in securing sufficient grid capacity, especially in urban, tropical, or regulated markets, can stall deployments for years. In many cases, power availability, not capital, is the primary gating factor for new data center builds. To mitigate this, operators are securing dedicated substations, multi-source redundancies, and long-term PPAs years in advance.

### **Cooling Transitions from Efficiency to Survival**

Conventional air-based systems can't keep pace with racks exceeding 30 kW, forcing a shift to liquid-to-chip, immersion, or rear-door heat exchange systems. These cooling technologies demand higher capex, alter physical layout, and introduce operational complexity, but they're now essential to AI-grade workloads. Cooling is no longer a support function; it's integral to infrastructure viability.

### **Thermal Infrastructure as Investment Differentiator**

The global data center cooling market is projected to grow from US\$25.8 billion in 2024 to over US\$100 billion by 2035, highlighting its critical role. Investors increasingly prioritize sites with low PUEs, cooling redundancy, and renewable integration. Efficient thermal management has become a structural competitive advantage, driving both regulatory compliance and long-term asset value.

## **Jurisdictional Data Barriers and Localization Pressure**

### **Regulatory-Driven Infrastructure Localization**

Governments worldwide are enforcing data localization laws requiring that sensitive data, across sectors like healthcare, finance, and defense, remain stored domestically. This has forced global operators to pivot from centralized cloud delivery to in-country capacity and local partnerships, aligning infrastructure with national legal frameworks.

### **Design and Operational Complexity**

Complying with localization mandates often requires data center operators to decouple systems from their global networks. This involves deploying independent control planes, restricting outbound telemetry, and operating facilities in "air-gapped" or semi-isolated configurations. These changes reduce operational flexibility and complicate multi-region failover strategies. For example, a global cloud provider may need to replicate not just

storage, but also compute orchestration, identity management, and observability stacks in-region, doubling or tripling deployment costs. Microsoft has noted that sovereign cloud builds can cost 2–3x more than standard regional data center deployments due to these duplication requirements and compliance overhead.

The operational burden is also significant. Isolated environments make centralized patching, logging, and incident response more difficult, introducing greater risk of inconsistency or downtime. Redundancy must often be achieved locally, meaning mirrored infrastructure is required even in smaller, lower-demand regions, raising capital intensity and elongating deployment timelines. For example, Oracle's Sovereign Cloud for the EU operates entirely separate from its core cloud, with its own dedicated personnel and restricted administrative access, to comply with GDPR and national security policies. Similarly, Alibaba Cloud's isolated deployments in Southeast Asia are tailored to meet country-specific cybersecurity laws, requiring full-stack replication in jurisdictions like Indonesia.

These adaptations, while necessary for regulatory compliance, shift the data center model from global scale economies to localized, compliance-bound operations, driving up cost and slowing expansion timelines.

### **A New Strategic Advantage**

Despite the friction, legal clarity in mature jurisdictions has created attractive conditions for long-term deployment. Markets with stable regulatory frameworks are increasingly viewed as safe zones for digital investment. Localization, once seen as a hurdle, now serves as a strategic lever for investor confidence and regional differentiation.

## **Decentralized Workloads and Crypto-Induced Volatility**

### **Short-Term Demand, Long-Term Instability**

Decentralized workloads, primarily from crypto mining, prioritize cheap power and speed over stability. These tenants often operate on floating contracts with minimal redundancy, creating transformer and cooling stress while offering unreliable revenue. Regulatory shifts and token volatility make these deployments difficult to underwrite.

### **Repurposing and Technological Spillover**

Despite their instability, crypto facilities are being retrofitted for edge computing, AI inference, or content caching. ASICs and custom-cooling solutions developed for mining are being re-engineered for enterprise use. This creates pathways for distressed assets to evolve into productive infrastructure.

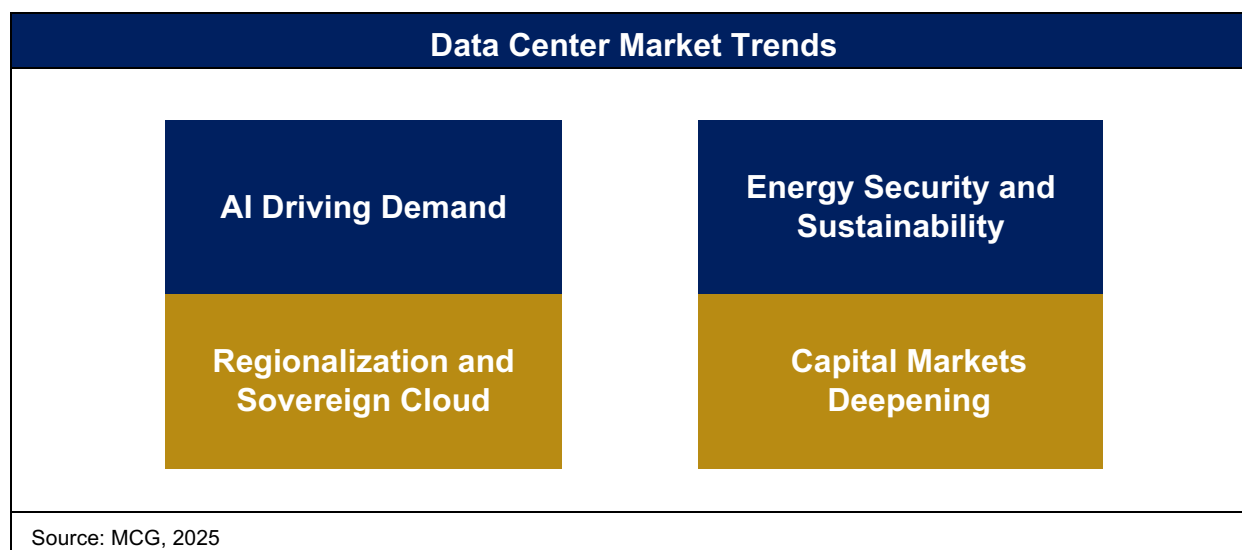
### **Caution Around Volatility**

Crypto's regulatory headwinds, environmental impact, and speculative economics make it a high-risk strategy for long-term operators. In ESG-conscious jurisdictions, new emissions standards are already marginalizing crypto workloads. Until governance and utility alignment improve, decentralized compute will remain an opportunistic, but volatile, subsector.

## **KEY TRENDS AND INDICATORS**

Since early 2025, the data center industry has continued to experience structural transformation, driven by the explosive rise of AI workloads, evolving regulatory frameworks, and a growing imperative for energy efficiency and sustainability. Here, M Capital Group explores the key trends that have underpinned the sector's evolution and will continue to shape its trajectory in the years ahead, particularly as technological complexity and power intensity increase across the industry.

These trends, such as AI-enabled automation, intelligent cooling systems, and distributed infrastructure, are allowing operators to manage higher-density workloads more efficiently, while also minimizing operating costs. With electricity prices surging across multiple regions and global data center power usage expected to nearly double by 2030, these innovations help preserve margins in a capital-intensive and highly competitive environment. Furthermore, institutional capital continues to deepen across the sector, reinforcing the data center's position as a core component of both digital transformation and real asset allocation strategies.

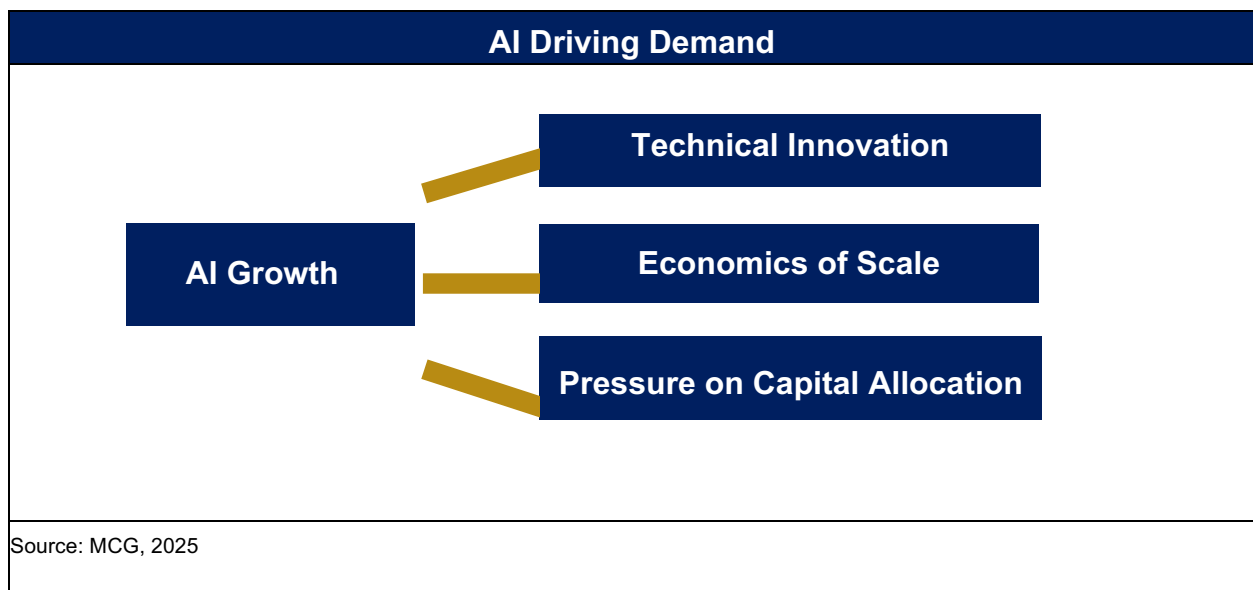


In 2025, structural shifts in the data center industry are being shaped by four critical forces: surging AI-driven demand, regionalization and sovereign cloud strategies, growing focus on energy security and sustainability, and deepening institutional capital markets. Together, these trends are redefining deployment models, enabling more efficient high-density operations, and reinforcing data centers as essential infrastructure in the global digital economy.

### AI and High-Density Compute Driving Infrastructure Demand

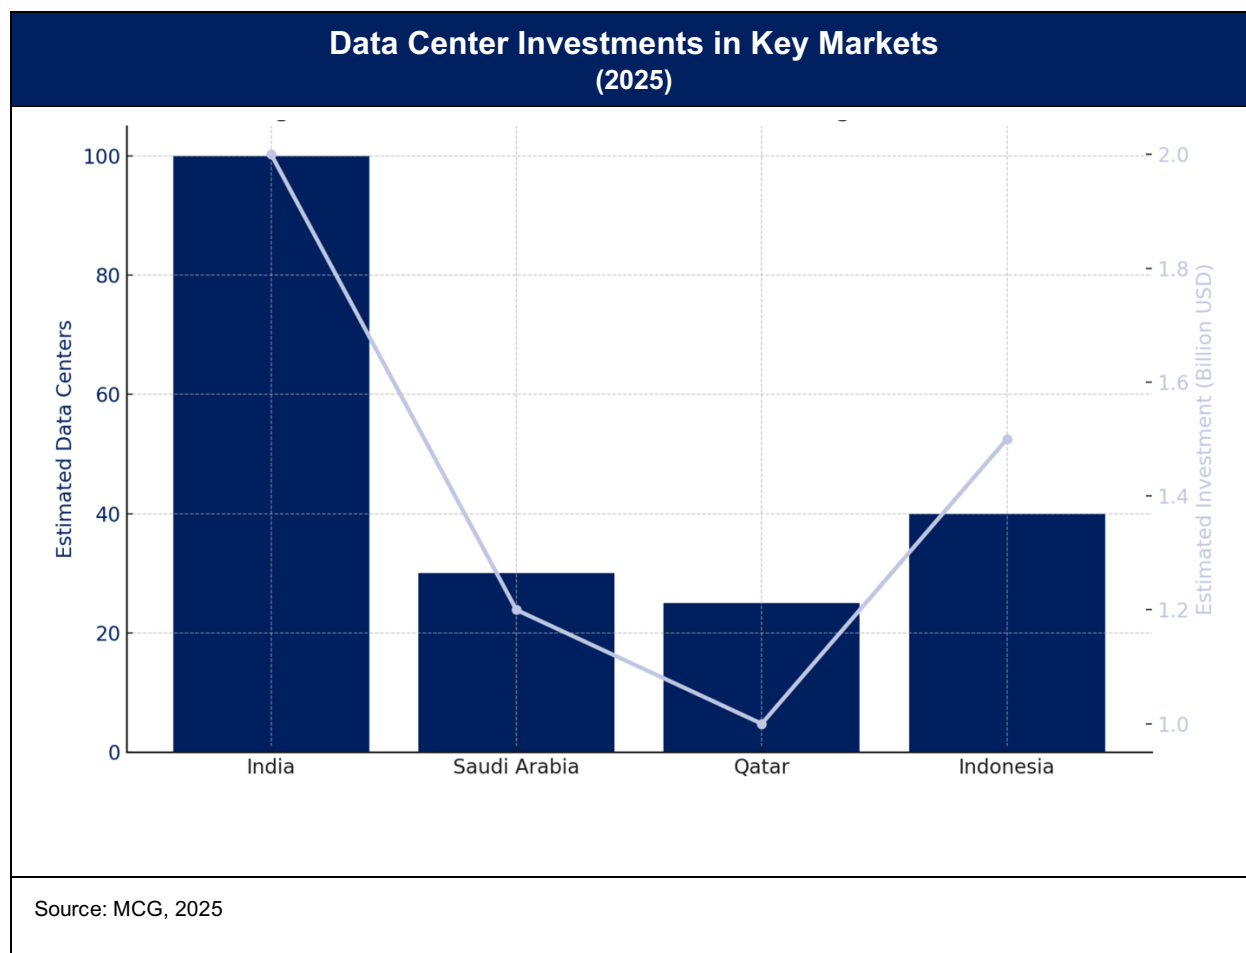
The explosive growth of generative AI and machine learning workloads continues to fundamentally reshape data center infrastructure, with significant implications for power demand, cooling, and facility design. In the United States, electricity consumption by GPU-accelerated servers surged from under 2 TWh in 2017 to over 40 TWh in 2023, accounting for a rapidly increasing share of total data center energy. Looking ahead, AI-related power demand is projected to drive global data center electricity consumption up 50 % by 2027 and up to 165 % by 2030 relative to 2023 levels, pushing annual usage toward 945 TWh or nearly 3 % of global electricity by the end of the decade.

This transformational shift is forcing data center operators to deploy next-generation infrastructure: power-dense racks in the range of 36 kW per rack in 2023, rising to 50 kW or more by 2027, and some designs pushing to 600 kW-per-rack or even 1 MW configurations to handle ultra-intensive AI clusters. Infrastructure projects are scaling accordingly. Leading hyperscalers and operators are building multi-gigawatt campuses to support “AI factory” workloads, as exemplified by OpenAI and Oracle’s Stargate initiative, which is targeting 4.5 GW of capacity initially and potentially 10 GW in total. Operators also face grid bottlenecks: U.S. data centers consumed 4.4 % of total U.S. electricity in 2023, tracked at 176 TWh, with projections suggesting consumption could rise to between 325–580 TWh (≈6.7–12 % of national electricity) by 2028.



These trends signal a structural shift in data center economics and design. Traditional layouts optimized for CPU-based workloads no longer suffice, operators must now invest in high-density power delivery, advanced cooling, and flexible campus design. The tradeoff: higher upfront capex but significantly greater compute throughput, offering economies of scale for hyperscalers and colocation providers. However, as power becomes a limiting constraint, securing reliable supply, whether through renewables, natural gas peaker plants, small-modular nuclear, or grid-scale upgrades, has become a strategic imperative. AI's insatiable power appetite thus not only drives technical innovation in thermal management and facility architecture but also exerts pressure on capital allocation strategies across the energy and infrastructure sectors.

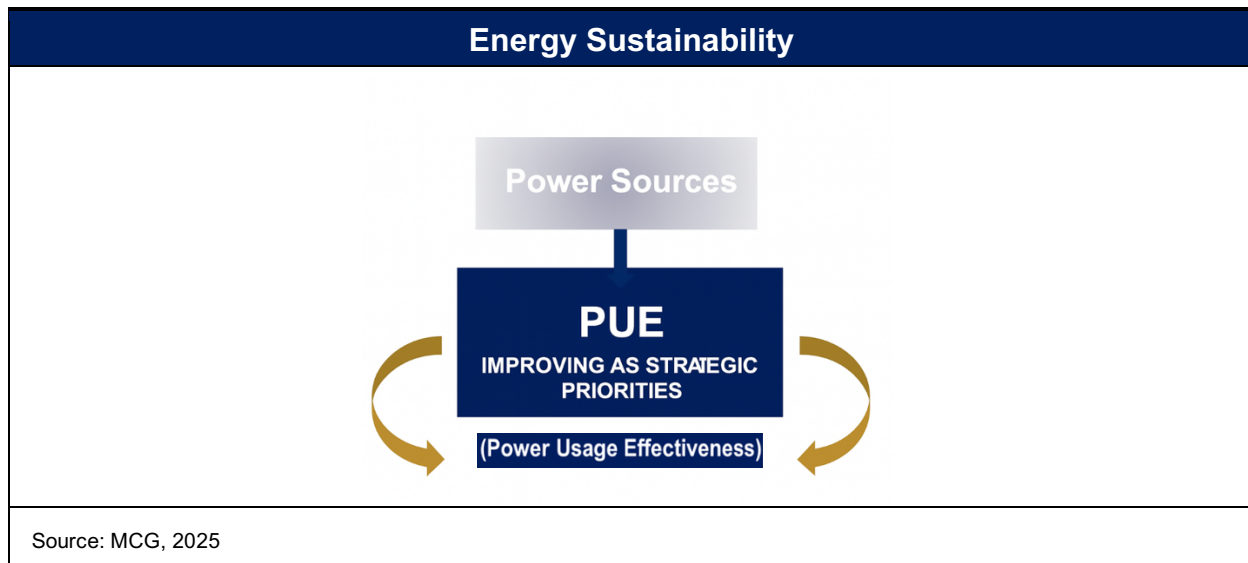
### Global Shift Toward Regionalization and Sovereign Cloud



Geopolitical uncertainty, data localization regulations, and latency requirements are encouraging regional data center deployments. Markets such as India, Saudi Arabia,

Qatar, and Indonesia are investing heavily in digital infrastructure to support sovereign cloud strategies and ensure compliance with local data laws. In India alone, over 100 new data centers are expected by 2025, representing US\$2+ billion in investment. Meanwhile, the Middle East is emerging as a global interconnection hub linking Europe, Africa, and Asia.

### Energy Security and Sustainability as Strategic Priorities



With data center energy consumption projected to exceed 1,000 TWh globally by 2026, the industry is under pressure to reduce its environmental footprint. Operators are investing in renewable PPAs, battery storage, nuclear-backed capacity, and on-site generation to secure cleaner and more stable power sources. Simultaneously, average PUE continues to decline, with leading operators targeting PUEs of 1.2 or lower. Regulatory frameworks, such as the EU's Energy Efficiency Directive and U.S. grid constraints, are accelerating this shift.

### Capital Markets Deepening and Institutionalization

The data center sector is attracting unprecedented levels of capital, including sovereign wealth funds, infrastructure funds, pension investors, and global REITs. Investors are drawn by the sector's long-term leases, inflation-linked revenue, and mission-critical profile. Private equity and cloud platforms are now forming large-scale joint ventures to co-develop multi-region campuses.

## PROMISING GROWTH DYNAMICS

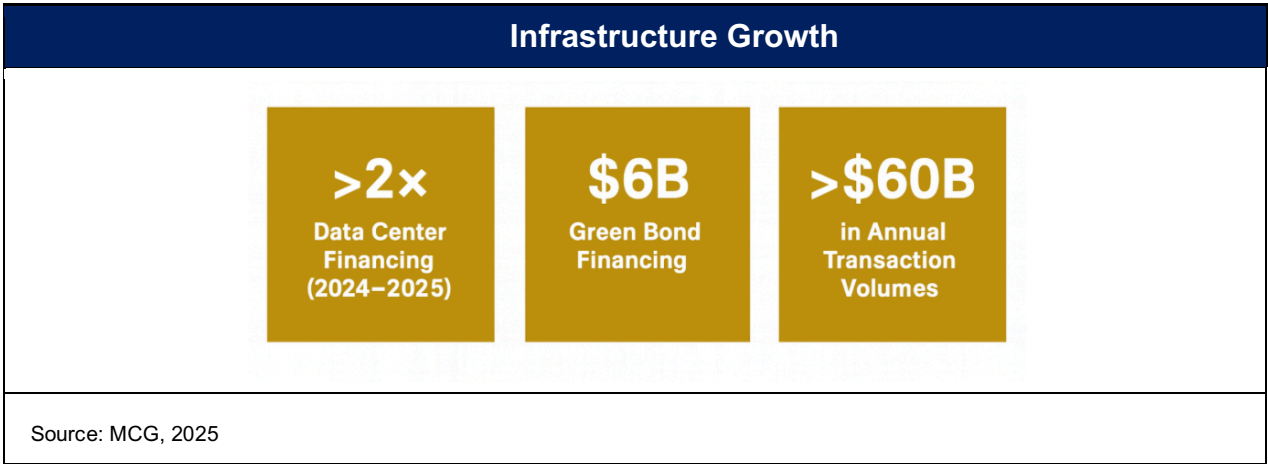
### AI-Led Demand Expansion

The most immediate and powerful growth driver in the data center industry remains the continued adoption of generative AI and large-scale machine learning. As enterprises and hyperscalers integrate AI models into their core operations, the demand for accelerated computing infrastructure, particularly GPU-dense clusters, is expected to rise exponentially. According to Goldman Sachs, AI-driven data center workloads are projected to grow at a 30–40% CAGR through 2030, with AI alone accounting for over 60% of new power demand in certain hyperscale environments. This sustained compute appetite translates directly into a long-term infrastructure expansion cycle, particularly for operators capable of delivering high-density, low-latency, and power-secure environments.

### Digital Sovereignty and Regionalization

Alongside rising compute demand, data localization and digital sovereignty are reshaping where and how infrastructure is deployed. Countries across Asia, the Middle East, and Europe are implementing strict data residency regulations to ensure national control over sensitive digital information. This trend is driving the need for localized, in-country cloud regions and colocation facilities. OpenAI, for example, recently launched sovereign cloud options in India and Southeast Asia in response to local requirements. Simultaneously, government-backed initiatives like the Transcendence AI program in Saudi Arabia and AWS’s planned US\$5.3 billion regional investment highlight how national policy is accelerating capacity buildouts. These regional hubs not only support compliance but also enable local AI training, latency reduction, and digital industrialization.

### Infrastructure as a Core Asset Class



As data centers become more essential to national economies and digital ecosystems, they are increasingly being treated as long-term, institutional-grade real assets. Infrastructure and private equity funds have committed billions to platform acquisitions, joint ventures, and hyperscale campus developments, driven by stable cash flows, high tenant retention, and long-duration contracts. In the U.S. alone, data center financing volume doubled from US\$30 billion in 2024 to a projected US\$60 billion in 2025, reflecting surging capital interest in digital infrastructure. STACK Infrastructure's recent US\$6 billion green bond raise further underscores investor appetite for ESG-aligned, scalable platforms. This wave of capital ensures that development pipelines can remain aggressive, even amid rising input costs and longer permitting timelines.

### **Energy Innovation and Sustainability Leadership**

A further pillar of sustained growth lies in the industry's proactive adaptation to the energy transition. While data centers face increasing scrutiny over power usage, leading operators are innovating rapidly across power procurement, cooling, and environmental management. The uptake of renewable energy PPAs, on-site solar, modular nuclear, and advanced thermal solutions is allowing developers to differentiate offerings and meet ESG requirements from both regulators and clients. The resulting alignment with national decarbonization goals strengthens the industry's social license to operate and expands access to ESG-aligned capital, effectively embedding data centers into the broader green transition narrative.

### **Technological Modernization and Automation**

Finally, the rise of software-defined infrastructure and AI-driven operational tooling is unlocking new efficiencies and business models. From autonomous data center management to dynamic energy load balancing, these innovations reduce operating costs while improving uptime and capacity utilization. Such advancements are especially valuable in edge deployments and hybrid cloud environments, where modularity, agility, and real-time responsiveness are essential. As data center infrastructure becomes increasingly intelligent and automated, providers will be better positioned to scale with demand while protecting margins in a high-cost, competitive landscape.

Taken together, these dynamics underpin MCG's conviction that the data center industry is entering a sustained multi-decade growth cycle, fueled by AI, grounded in real asset economics, and supported by both global digitization and energy transformation.

A photograph of server racks in a data center, overlaid with a digital network graphic consisting of glowing nodes and connecting lines. The word "GLOBALIZATION" is written in white capital letters across the top left of the image.

## GLOBALIZATION

## INTRODUCTION

The data center industry is undergoing a profound transformation globally, driven primarily by the exponential rise of artificial intelligence ("AI"), the sovereign push for national data infrastructure, and the expanding demand from crypto-related ecosystems. As AI and cloud computing become foundational to national competitiveness, digital infrastructure has shifted from being a commercial asset class to a strategic imperative. Below, M Capital Group outlines the globalization and evolution of this sector across key regions:

## GLOBAL DRIVERS

### Redefining the Digital Backbone

The global data center industry is undergoing an irreversible inflection point, driven by the explosive proliferation of artificial intelligence ("AI") workloads. Between 2020 and 2025, the volume of AI-generated data is expected to grow at a compound annual growth rate ("CAGR") exceeding 40%, outpacing traditional cloud services and accounting for an estimated 50% of all data center power consumption by year-end 2025. This is not merely a scale challenge—it's a fundamental reshaping of data infrastructure. Where legacy facilities were once passive, power-hungry warehouses for storage and basic compute, the new generation of AI-optimized data centers must support high-density GPU racks, liquid-cooled environments, and real-time latency thresholds under 10 milliseconds. In strategic terms, data centers have evolved into national infrastructure—on par with roads,

ports, and power grids—central to a country’s economic competitiveness, digital sovereignty, and military preparedness.

This transformation is both global and asymmetric. In mature markets like the United States and Northern Europe, hyperscalers such as Microsoft, Amazon, and Google have collectively announced over US\$75 billion in data center CAPEX for 2024–2026, with each campus exceeding 1 GW of power capacity and incorporating proprietary liquid cooling and silicon optimization for AI inference. These “second-wave” data centers are no longer just about uptime or scale—they are precision-engineered to run trillion-parameter models like GPT-5, Gemini, and Claude 3. Meanwhile, in the Middle East and Asia, the response is sovereign-led: the UAE’s G42 and Saudi’s NEOM have committed to building sovereign AI clouds, while India’s Digital Personal Data Protection Act is forcing foreign hyperscalers into joint ventures with local telcos. In Southeast Asia, countries like Indonesia and Vietnam are issuing exclusive AI hosting licenses to state-approved operators, creating a tiered ecosystem where regulatory compliance drives market access. Across all regions, a single truth persists: the future of AI is grounded in the physical limits and design choices of infrastructure.

**Data Center Market Size and Forecast 2025 to 2034**  
(US\$ Billion)



Source: MCG, Precedence Research, 2025

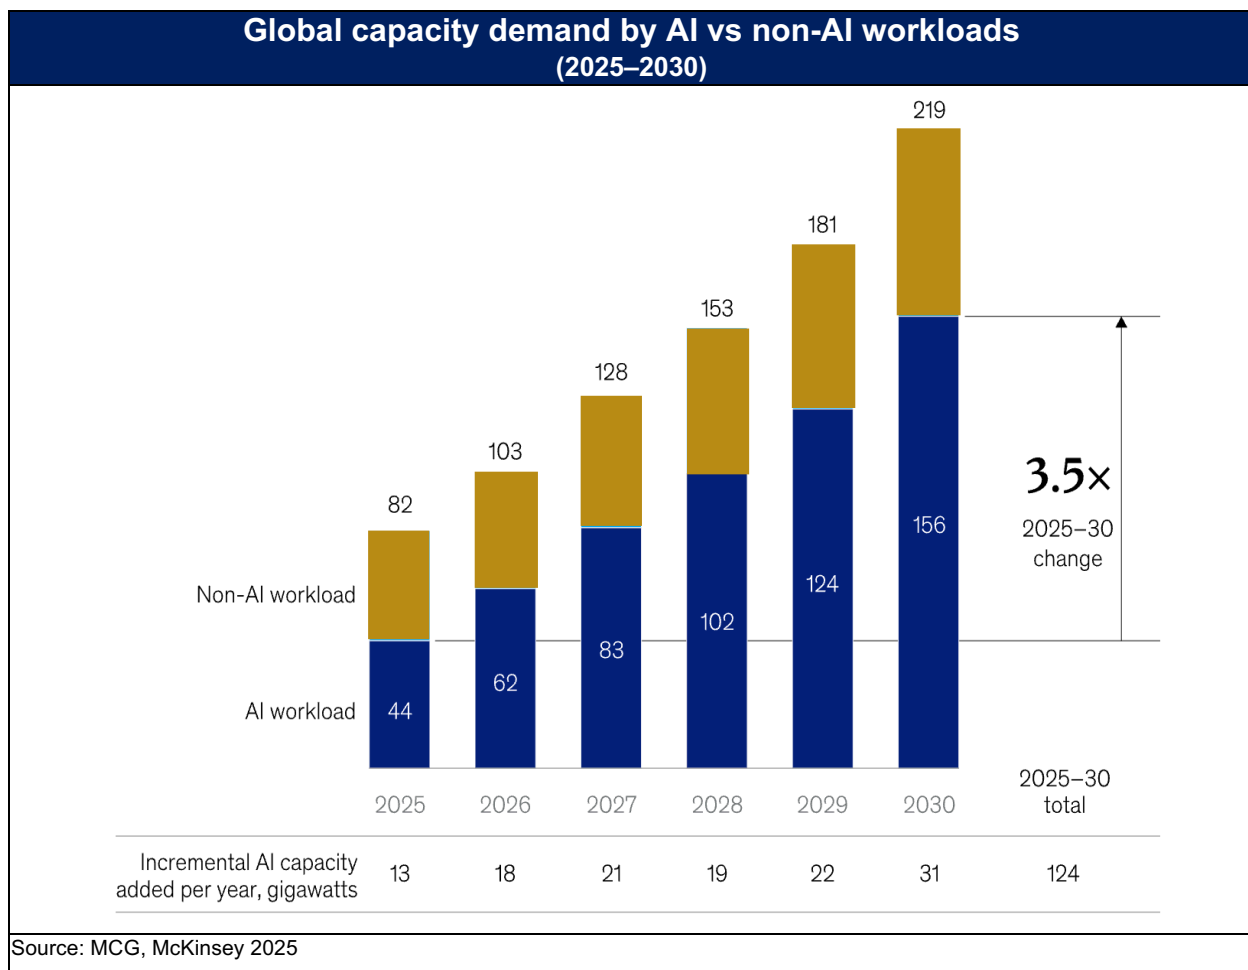
The global data center market is projected to reach US\$387 billion in 2025, with AI-native infrastructure representing the fastest-growing segment. According to industry projections, GPU server deployments will grow 5x from 2023 to 2026, with AI workloads demanding not only more energy, up to 20–50 kW per rack, but also regulatory navigation, particularly where cross-border data restrictions are enforced. Sovereign data mandates now apply in over 60 countries, with governments like Qatar, Germany, India, and Brazil requiring in-country storage for sectors including healthcare, finance, biometrics, and public administration. This shift has upended the economics of data center planning: global operators must adopt sovereign-compliant architectures that incorporate Tier III+ or IV specifications, air gap segmentation, audit trail monitoring, and government-cleared fiber interconnects. Investors must now recalibrate toward localized projects with dual regulatory and energy resilience planning, blending real estate, digital infrastructure, and geopolitical strategy into every underwriting decision. The stakes are no longer just yield; they are sovereignty, resilience, and AI readiness.

### **AI as Infrastructure’s Operating System**

Artificial Intelligence is no longer a vertical or a feature, it is the organizing principle of the digital economy. For infrastructure, this means AI is dictating hardware configuration, site selection, operational architecture, and risk underwriting. Unlike conventional SaaS or storage-driven demand, AI workloads require enormous concurrent processing capacity, high IOPS (input/output operations per second), and extremely low latency across distributed systems. As such, AI has become the single most important design constraint for data centers globally in 2025.

This shift is reflected in capital budgeting decisions. Operators must now pre-build facilities with high-performance compute (“HPC”) racks, pre-integrated GPU server bays, and flexible thermal control systems that can switch between air and liquid cooling. Notably, this also changes procurement lead times, with AI-capable hardware sometimes requiring 6–12 months of advanced ordering and vendor exclusivity agreements. Moreover, as AI evolves to more specialized use cases (e.g., inference at the edge, multimodal model deployment), infrastructure must adapt just as quickly, with modular architectures that allow seamless hardware and software upgrades.

The second-order effects of this AI transformation are just as important. Edge data centers are being redesigned to support real-time inferencing, pushing infrastructure closer to the user and placing pressure on network latency. Governments are building AI acceleration zones that combine compute credits with subsidized infrastructure and sovereign cloud compliance. Even enterprise software vendors are embedding LLMs into their platforms, pushing CIOs to adopt AI-readiness as a procurement criterion for third-party data center services. In every case, AI is redefining not only the ‘what’ of infrastructure, but also the ‘who’ and the ‘why’.



This graph highlights global data center capacity demand rising from ~82 GW in 2025 to 219 GW in 2030, with AI workloads increasing from 44 GW to 156 GW, while non-AI segments lag behind. It vividly illustrates that AI is the primary growth engine in data centers, requiring nearly three times more expansion capacity than all other workloads combined. The move toward GPU-driven clusters, deep learning frameworks, and generative AI architectures mandates dense compute infrastructure, no longer a niche but the mainstream of data center investment.

As a result, data center design must adapt: namesake metrics like PUE (Power Usage Effectiveness) are overshadowed by thermal resilience and throughput per watt, reflecting AI's demands. Developers and investors must now evaluate sites on AI-readiness, such as support for liquid cooling or GPU rack density, not just on connectivity or cost per kW. Failure to prioritize these criteria risks building obsolete assets in a rapidly evolving environment.

### **Crypto Mining**

While crypto-related demand is no longer in its 2021 peak, crypto mining remains a material driver of infrastructure allocation in certain geographies. This demand is unique in that it is ultra power-intensive, latency-agnostic, and capex-light. That is, crypto miners care most about cheap electricity and scalable rack capacity—not Tier IV reliability or sovereign compliance. In regions like Texas, Kazakhstan, or parts of the Middle East, crypto mining still accounts for 20–30% of deployed power capacity in newer facilities.

This presents a dilemma for infrastructure investors. On one hand, crypto tenants can drive rapid monetization of stranded assets, helping developers reach break-even faster on projects that might otherwise languish. On the other hand, crypto demand is highly volatile and often lacks regulatory clarity, making it difficult to underwrite over 5–10 year horizons. Operators must therefore employ tiered pricing, flexibility clauses, and exit ramps in their lease structures to manage tenant risk. Some have begun segregating crypto tenants into modular annexes, allowing for operational separation from more stable enterprise tenants.

Moreover, regulatory headwinds are mounting. Countries such as China and India have already implemented de facto bans, while others like the U.S. are pushing for energy disclosures and carbon taxation on mining operations. As a result, crypto-driven demand is migratory and opportunistic—useful for filling gaps, but not a reliable long-term anchor. For AI-focused operators, the challenge is to extract short-term value from crypto without compromising long-term design and reputation.

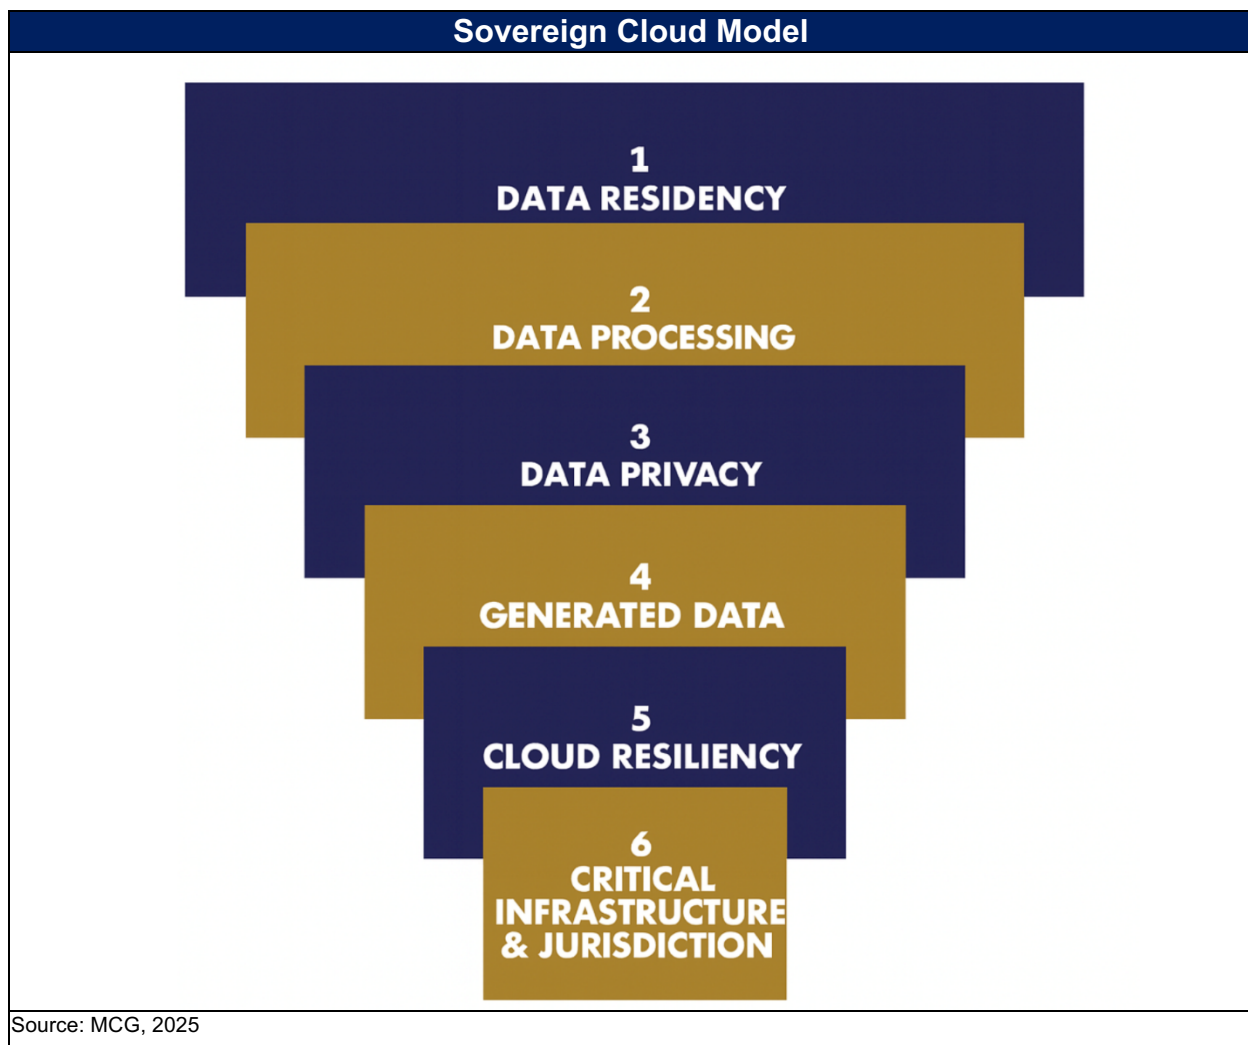
### **Data Sovereignty**

Perhaps the most enduring structural driver is the global pivot toward data sovereignty and localization. From the EU’s GDPR and Digital Markets Act to India’s PDP Bill and Qatar’s national data policies; governments are increasingly asserting that sensitive data must remain onshore. This includes not only healthcare and finance but also behavioral analytics, geospatial data, and citizen biometric records. These regulations are not just legal guidelines—they are shaping infrastructure blueprints.

In practical terms, this means that even hyperscalers must establish localized availability zones, sometimes in partnership with domestic telecoms or sovereign funds. For investors, this requires tailoring project finance structures to comply with local shareholding rules, approval cycles, and land restrictions. Sovereign compliance can also affect tenant quality, as certain multinationals may be unwilling or unable to deploy outside of approved sovereign cloud jurisdictions.

Localization also impacts design. Facilities must now integrate air gaps, hardware encryption, audit trail systems, and national fiber backbones, increasing both capex and time-to-market. In the most restrictive jurisdictions—such as Saudi Arabia and Qatar—

even public-private partnerships are subject to pre-clearance by central banks and information ministries. For developers and investors alike, this marks a paradigm shift: infrastructure projects must now be “sovereign-aware” by default, or risk being unbuildable altogether.



Though not always numerical, sovereignty requirements are enforced through compulsory localization, especially for financial, healthcare, and personal data. Visualization of national data localization policies across jurisdictions reinforces how infrastructure planning must be regionally compliant. A key example: Northern VA and EU quotas showcased in CBRE data center inventory graphs (Northern VA adding 523 MW in 2024, while European major markets grew by ~7% due to power constraints). These are driven not only by cost or demand—but by compliance and sovereignty constraints that mandate local compute.

## **Converging Data Center Drivers**

### **Physical Footprint vs. Power Density**

Traditionally, data centers expanded through sheer physical scale. Today, the metric has shifted: more power and performance must be extracted from tighter footprints. The need to deploy liquid-cooled racks, direct-to-chip cooling, and modular power systems is no longer an option but a baseline expectation. As a result, global operators face a unique convergence—AI demand is global, but the ability to meet that demand hinges on solving highly localized power, cooling, and land constraints.

### **Security and Compliance as Cost Centers**

These sovereign restrictions are not only technical; these restrictions significantly impact cost structure. Facilities must now build in hardware-level encryption, air gaps, audit trail infrastructure, and physical separation between public and restricted workloads. Sovereign design requirements are raising capital expenditure (“CapEx”) while elongating time-to-market, as each jurisdiction mandates bespoke permitting, inspections, and cybersecurity certifications. As AI models become integrated into sensitive domains—such as citizen services, financial systems, and healthcare—the regulatory perimeter for data center deployment has tightened. AI developers and infrastructure sponsors must now embed legal foresight into deployment plans: sovereign hosting may determine whether a contract is even legally viable. This has turned data sovereignty from a compliance item into a core infrastructure driver.

### **Sustainability as Regulatory and Tenant Requirement**

Beyond infrastructure, sustainability is fast becoming a license to operate. ESG-conscious investors and global tenants now demand traceable emissions reporting, renewable integration, and water reuse in operations. Data centers that fail to meet carbon neutrality targets risk exclusion from key AI workloads, especially in Europe and North America. Sustainability is no longer an advantage—it is a baseline threshold for global competitiveness in the AI age. Energy has emerged as the defining constraint in global data center expansion. AI-grade data centers now consume 3–5x more power per rack than legacy installations. Markets with favorable grid access, low-cost renewables, or sovereign-backed utilities are being prioritized. The ability to secure long-term power purchase agreements (“PPAs”) is as critical as land acquisition or financing.

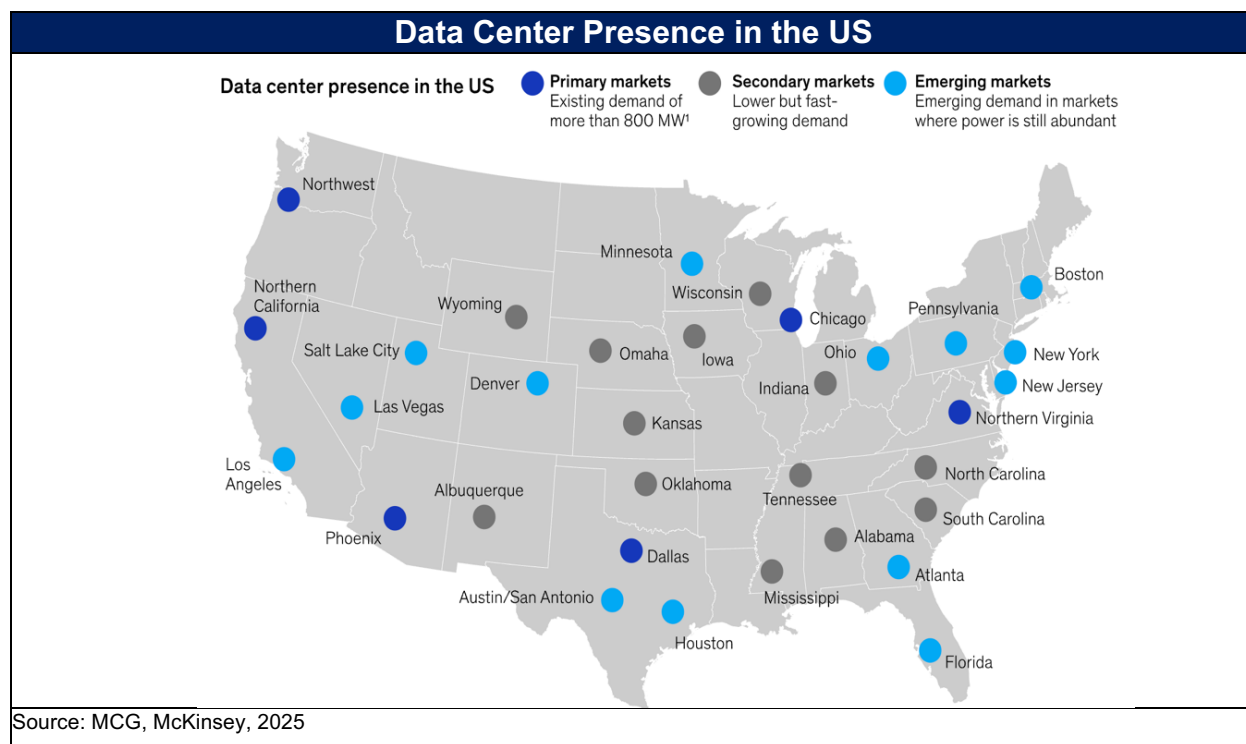
## US STILL LEADS

### Engineering Innovation Leadership

The United States continues to dominate the global AI data center landscape across nearly every metric: installed capacity, hyperscaler investment, infrastructure innovation, and geopolitical alignment. As of 2025, the U.S. accounts for over 54% of global hyperscale data center capacity, with more than 500 active hyperscale sites and a robust pipeline of expansions. This lead is supported by a highly developed capital ecosystem—comprising real estate investment trusts (“REITs”), private equity sponsors, utility-scale energy partners, and a wide range of engineering and design firms. Hyperscalers such as Amazon Web Services (“AWS”), Microsoft Azure, and Google Cloud are rapidly expanding proprietary AI campuses throughout the U.S., particularly in Virginia, Iowa, and Texas. According to a 2025 McKinsey report, AI workloads are projected to increase global data center capacity by over 33% annually through 2030, with U.S. hyperscalers expected to host over 60–65% of those workloads. Beyond public cloud, U.S.-based companies like Meta, Tesla, and OpenAI are also investing in vertically integrated infrastructure to support their private AI training operations, transforming the U.S. into the innovation sandbox of the global AI infrastructure race.

### Regional Market Dynamics

Equally significant is the U.S.'s leadership in engineering innovation around power, cooling, and operational efficiency, now the most expensive and technically complex components of running AI workloads. Average rack power density is projected to surpass 50 kW per rack within five years, far exceeding traditional designs. To address this, U.S. operators are pioneering liquid and immersion cooling, direct-to-chip heat dissipation, and predictive analytics-based load optimization. Microsoft and Google are piloting AI-assisted thermal management systems, lowering their PUE to below 1.1, far more efficient than the global average. Utilities in power-intensive states like Arizona and Texas are working directly with hyperscalers to design dedicated renewable microgrids, integrating AI workloads with grid-level demand response. Additionally, firms such as Google are deploying DeepMind-powered energy control systems across data center campuses, reducing cooling costs by up to 40%. These innovations not only mitigate rising OPEX but also fulfill ESG and carbon neutrality mandates. As the world shifts toward AI-native data center design, the U.S. stands far ahead in deploying smart infrastructure that scales efficiently, sustainably, and reliably.



## Geopolitical and Policy Alignment

What further distinguishes the U.S. is its alignment of AI infrastructure investment with national policy, security, and talent pipelines. The federal government has increasingly prioritized AI infrastructure as a strategic asset. Legislation such as the CHIPS and Science Act, alongside executive orders, has provided billions in funding for semiconductor manufacturing, data center expansion, and high-performance computing research. At the state level, momentum is also building. In the 2025 legislative session alone, all 50 states, Puerto Rico, the Virgin Islands, and Washington, D.C. introduced AI-related legislation. So far, 28 states and the Virgin Islands have adopted or enacted more than 75 new measures. This policy stability encourages private investment and accelerates hyperscaler timelines. On the talent side, the U.S. remains home to the world's top AI research institutions—MIT, Stanford, UC Berkeley—and AI labs such as OpenAI, Anthropic, and NVIDIA. This concentration of intellectual capital directly supports the co-design of chips, systems, and facility architecture, enabling real-time feedback loops between AI workloads and infrastructure evolution. Moreover, the U.S. is integrating AI data centers into its defense and scientific infrastructure, with agencies such as the DoD and NSA funding sovereign cloud and edge compute platforms. This convergence of commercial innovation, policy backing, national security, and world-class talent ensures that the United States will remain the undisputed global leader in AI data center infrastructure well into the next decade.

## EUROPE TAKE NOTE

### Regulatory Leadership & Sovereign Frameworks

Europe is emerging as one of the most policy-driven and sustainability-focused regions in the global data center industry. While the continent trails the U.S. and Asia in hyperscale capacity, it is outpacing other markets in establishing a regulatory and technological foundation that aligns with AI growth. The European Union's Gaia-X initiative exemplifies this approach, intended to build a federated, sovereign digital infrastructure across EU member states. This initiative emphasizes data sovereignty, local cloud provisioning, and strict GDPR compliance, significantly impacting where and how data centers are developed. In countries like Germany, France, and the Netherlands, hyperscale developments must now meet environmental benchmarks and operate within tightly regulated frameworks for cross-border data flow. These constraints, while slowing deployment velocity, are fostering a stable and compliant ecosystem that is becoming attractive for AI enterprises focused on trust, governance, and regulatory certainty.

### Strategic Infrastructure Expansion

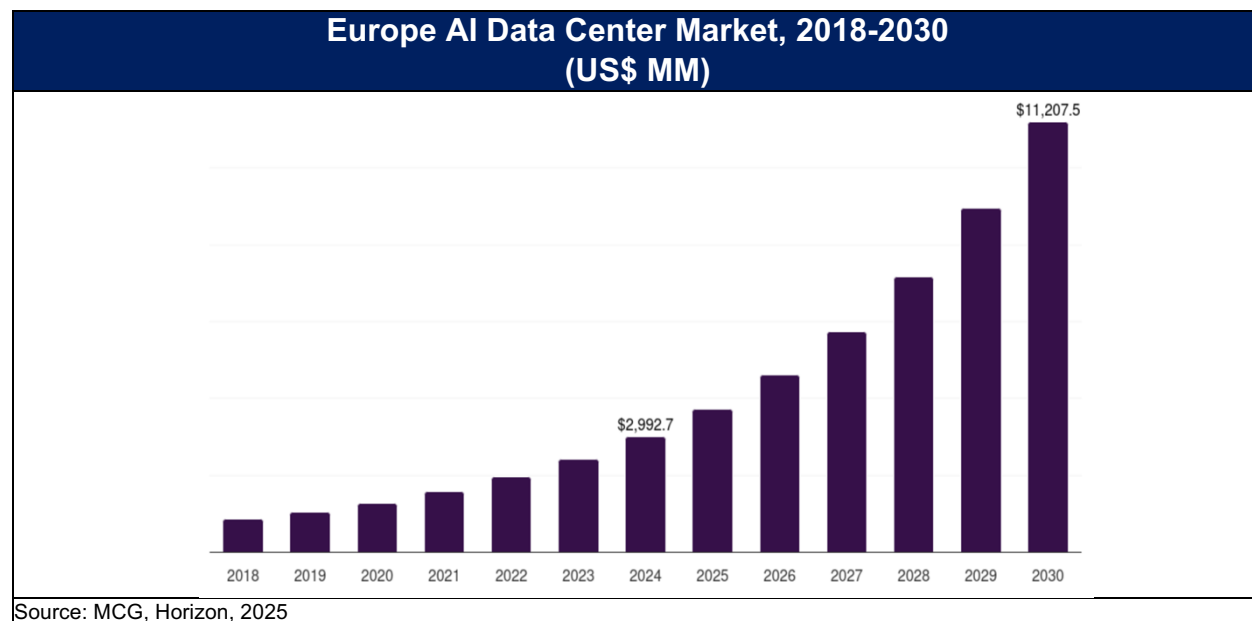
At the infrastructure level, Europe's AI-related data center expansion is highly fragmented but intensifying. Northern Europe, specifically Sweden, Norway, and Finland, is leading in green data center development due to abundant renewable energy, cool climate advantages, and government incentives. By contrast, regions like Southern and Eastern Europe are only beginning to scale up capacity but are drawing attention due to lower land and energy costs. Hyperscalers including Microsoft, Google, and AWS are investing in AI-optimized zones in Denmark, Spain, and Poland, while private equity-backed developers are launching edge projects in Tier II cities to meet latency and localization needs. The rise of AI-native workloads is pushing operators toward liquid-cooled, modular infrastructure, but local build timelines remain prolonged due to permitting delays and community opposition. The net result is a capacity expansion environment that is technically advanced, policy-constrained, and highly tailored to localized use cases.

### Sustainability as a Competitive Edge

Sustainability and ESG reporting pressures are also uniquely intense in Europe, directly shaping AI data center investment strategies. The EU's Corporate Sustainability Reporting Directive ("CSRD") and the EU Taxonomy for Sustainable Activities are now active, requiring large operators and institutional investors to disclose carbon, water, and biodiversity impacts of their infrastructure projects. This is transforming site selection, equipment design, and cooling strategies across the region. As AI compute density surges, Europe is prioritizing energy efficiency per watt of inference, with liquid immersion cooling, waste-heat reuse, and carbon-aware scheduling emerging as standard design criteria. Some operators are even required to submit urban integration and climate adaptation plans as a condition for zoning approval. Despite these challenges, European infrastructure funds and pension-backed platforms are committing billions to compliant

builds, banking on long-term stability, political alignment, and institutional demand for ESG-screened assets. In this respect, Europe is not only adapting to the AI revolution but shaping it around its values of digital trust, sustainability, and inclusion.

## Market Growth & Service-Layer Evolution



Supporting this, Grand View Research reports that the European AI data center market posted revenues of US\$3.0 billion in 2024, with a projected 24.7% CAGR between 2025 and 2030, reaching over US\$11.2 billion. The graph underscores an inflection point: hardware continues to dominate revenue shares, but burgeoning services revenue—particularly in areas like AI system integration, managed GPU infrastructure, and data localization solutions—is outpacing hardware growth. This trend hints at a maturing ecosystem where value-add layers beyond bare infrastructure are becoming essential. Advanced compute requirements, such as those for large language model (“LLM”) training and inference, are driving demand for bespoke service layers—pushing Europe toward a more holistic AI-centric data center model.

## ASIA WELL OFF

### Sovereign-Driven Expansion

Asia's digital infrastructure transformation has been one of the most dynamic globally, with countries such as China, India, Singapore, and Indonesia experiencing exponential demand for data center capacity. This expansion is heavily driven by the surge in AI adoption across financial services, e-commerce, and government systems. Sovereign clouds are becoming the default model in many markets, as regulatory regimes tighten control over data residency, security, and digital identity frameworks. In countries like India and Vietnam, governments now mandate that sensitive user and enterprise data must remain within national borders—effectively requiring new infrastructure to be built locally. This has given rise to hundreds of new AI-ready data center projects, many sponsored by domestic conglomerates in partnership with international capital.

### Regulatory Landscape

A defining characteristic of the region is the fragmented regulatory landscape, which presents both growth and complexity. While localization policies are universal, their interpretation and enforcement vary by country—adding legal risk to cross-border infrastructure strategies. For example, Indonesia has introduced a 2025 mandate requiring all financial and government platforms to host data domestically. In response, major domestic telcos and cloud firms have launched infrastructure investment arms to fast-track construction. In India, meanwhile, the Personal Data Protection Act has pushed foreign hyperscalers like Amazon and Microsoft to co-invest with local operators to ensure compliance. These dynamics have created fertile ground for private equity, infrastructure funds, and sovereign investors to enter through structured partnerships, often as co-owners with preferred local rights and exit mechanisms.

### Power & Cooling: Critical Constraints

However, the region also faces distinct technical and economic challenges, most notably in power availability, grid reliability, and cooling efficiency. According to Mordor Intelligence's 2025 report on Asia-Pacific Data Center Power, the region's market is expected to grow from US\$8.56 billion to US\$13.59 billion by 2030, indicating how rapidly operators are expanding grid access and backup capacity. Yet, the tropical climate in markets like Indonesia and Thailand severely strains cooling systems, especially as AI accelerates dense compute workloads. Liquid and immersion cooling, while increasingly discussed, remain cost-prohibitive at scale for many mid-market operators. As a result, capex inflation is a major concern, particularly in land-constrained hubs such as Hong

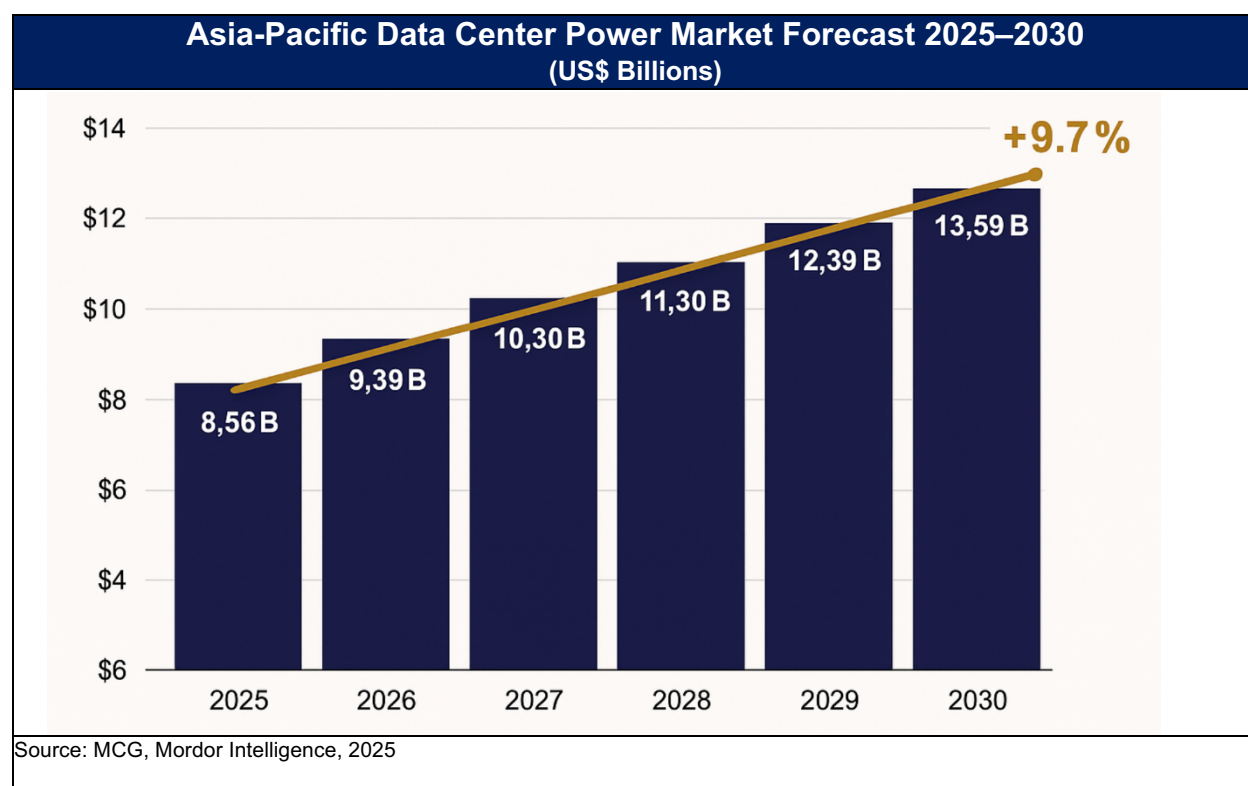
Kong, Tokyo, and Singapore, where land and permitting costs have risen over 40% since 2020.

## Emerging Demand Catalysts & Risks

### AI Training Firms: Next-Gen Tenant Class

Another major catalyst—and challenge—for the region is crypto-related data center demand. Cryptocurrency mining and blockchain validation services have created burst capacity needs in emerging hubs like Malaysia and the Philippines. While these tenants can rapidly fill rack space and generate strong early yields, these tenants often introduce higher volatility, regulatory scrutiny, and contract risk. AI training firms are also emerging as a new tenant class, often requiring complex configurations and GPU cluster access. As such, operators and sponsors must walk a fine line—maximizing power utilization without sacrificing resilience. In this landscape, investors like MCG have the opportunity to add value not just through capital, but by introducing operational best practices and tenant risk frameworks that have been tested in more mature markets.

### Risk Mitigation Strategies



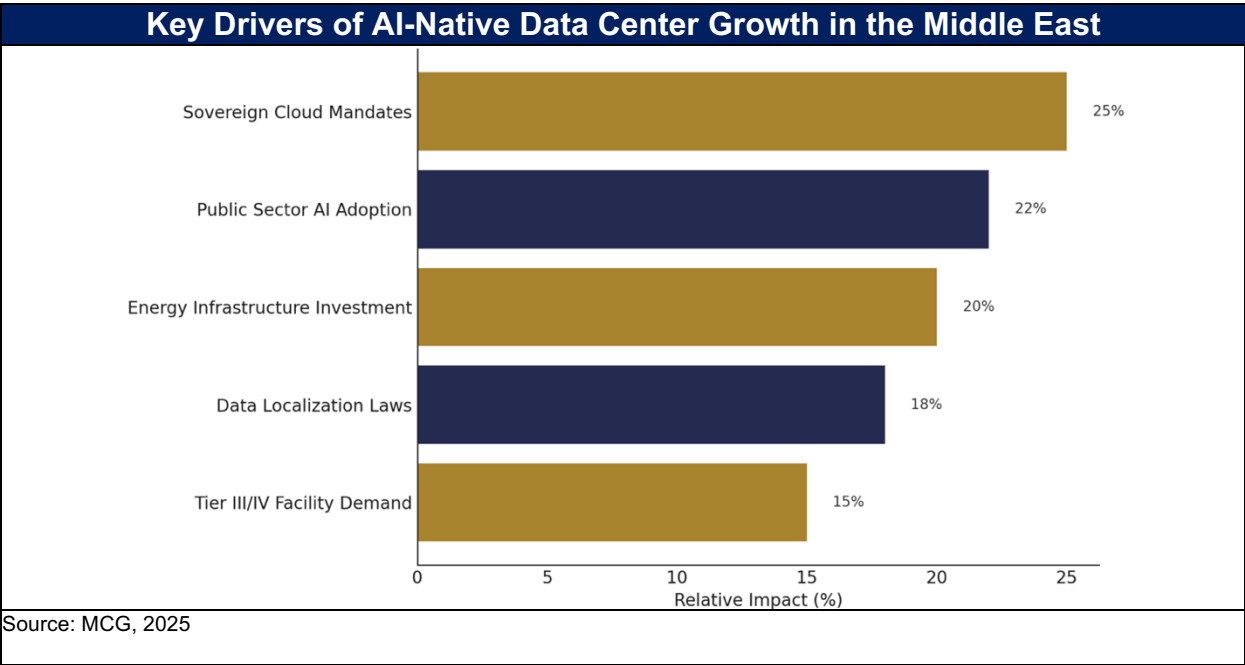
The Mordor Intelligence 2025 market report on Asia-Pacific's data center power sector provides critical insight into infrastructure scaling challenges. The region's data center power market is projected to expand from US\$8.56 billion in 2025 to US\$13.59 billion by 2030, reflecting a CAGR of 9.7%. The accompanying graph charts this steady upward climb, illustrating how power infrastructure has become a gating factor in AI data center expansion.

## MIDDLE EAST TAKING OFF

### Policy-Driven Sovereign Infrastructure

The Middle East has quickly emerged as one of the world’s most active regions in developing sovereign-backed digital infrastructure, particularly in AI-focused data centers. Countries like Qatar, Saudi Arabia, and the UAE have moved aggressively to capitalize on this trend by adopting data localization laws, investing public capital, and building out highly secure facilities. One of the defining regulatory pillars of this market is the prohibition of cross-border data storage, especially for government and financial services information. This policy, while protectionist in nature, has catalyzed demand for in-country Tier III and Tier IV data centers, many of which are being structured as joint ventures between public institutions and private equity-backed developers. This structural framework has created a favorable environment for strategic partnerships.

### Explosive Market Growth



A key driver of this regional momentum is the integration of AI into national economic agendas. Qatar’s National Vision 2030 and the UAE’s National AI Strategy have both identified artificial intelligence as a foundational component of public-sector transformation. Government agencies are increasingly embedding AI capabilities into

security, public health, finance, and energy systems—applications that require ultra-reliable, high-density computing environments. This trend is creating a steady pipeline of demand for sovereign-owned, AI-native data centers that support GPU-intensive workloads. Moreover, because these facilities must adhere to national data protection and cybersecurity protocols, public-sector contracts are increasingly favoring local data centers over international cloud providers. As a result, infrastructure operators are not only competing on price or uptime, but on compliance, trust, and strategic alignment.

### **Operational Headwinds**

#### **Power & Cooling: Dominant OpEx Burden**

However, this infrastructure evolution is not without challenges. One of the most pressing issues across the region is the high operational cost of power and cooling. According to Cushman & Wakefield's 2025 APAC Data Center Construction Cost Guide, power and cooling represent over 55% of total operating expenses in hot-climate markets like the Middle East. Traditional cooling technologies are being pushed to their limits as AI-related heat output escalates, particularly in GPU-driven data centers. Many operators are now turning to advanced thermal control systems—including liquid cooling and modular HVAC—to meet the unique demands of AI-driven workloads. These upgrades, while essential, drive-up capex significantly and have pushed local sponsors to seek foreign equity or technology partners, creating the environment for MCG to step in with its strategic capital and sector expertise.

#### **Tenant Volatility & Screening Imperatives**

Furthermore, the customer base itself is evolving in complexity and expectation. While traditional tenants have included telecoms and financial institutions, today's pipeline includes crypto miners, AI model developers, and sovereign digital infrastructure offices. Crypto, in particular, introduces both opportunity and volatility: miners are willing to pay high premiums for access to power and compute, but their financial sustainability can be unstable. This makes tenant screening and long-term lease structuring a critical risk-mitigation strategy. Simultaneously, operators must carefully navigate QCB's oversight and maintain flexibility for future regulatory changes, especially around data classification and foreign involvement. Altogether, the Middle East's AI data center evolution is a convergence of policy, technology, and capital strategy—offering long-term opportunity but requiring thoughtful execution.

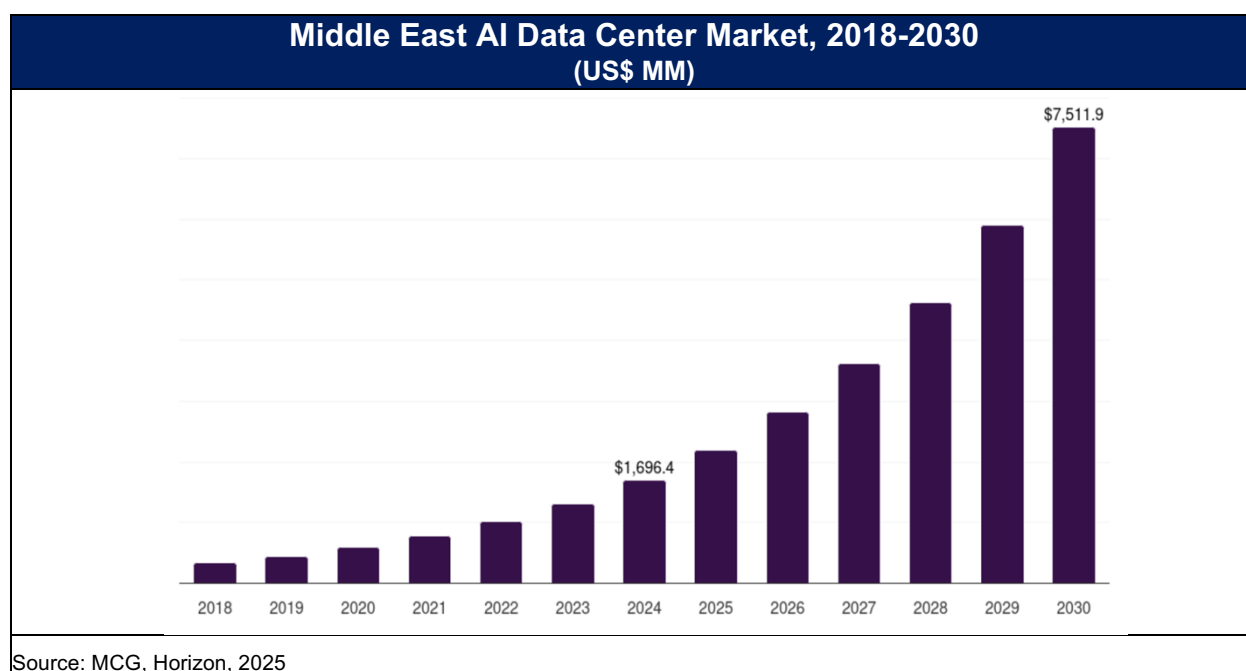
### **Evolving Tenant Landscape**

#### **Market Growth Trajectory**

According to Grand View Research’s 2025 outlook on the Middle East & Africa AI data center market, the regional market is expected to grow at a CAGR of 27.9% from 2025 through 2030. This trajectory reflects the rapid shift from traditional IT workloads to AI-optimized compute infrastructure. The Grand View graph highlights how countries like Qatar, UAE, and Saudi Arabia are scaling sovereign-grade infrastructure to support AI applications in banking, government, and logistics.

### **Government-Led Investment Catalysts**

Most notably, the data shows government-backed investments are driving this surge—enabled by favorable policies, capital inflows, and data sovereignty mandates. Qatar’s own regulatory environment, which prohibits outbound storage of sensitive data, is a primary driver of localized infrastructure like the QCB data center. The graph’s upward trendline underscores how AI—not just generic cloud—has redefined investment priorities in the region. Projects anchored by institutions such as Qatar Central Bank (“QCB”) and supported by public-private partnerships.



## **AFRICA PROMISING**

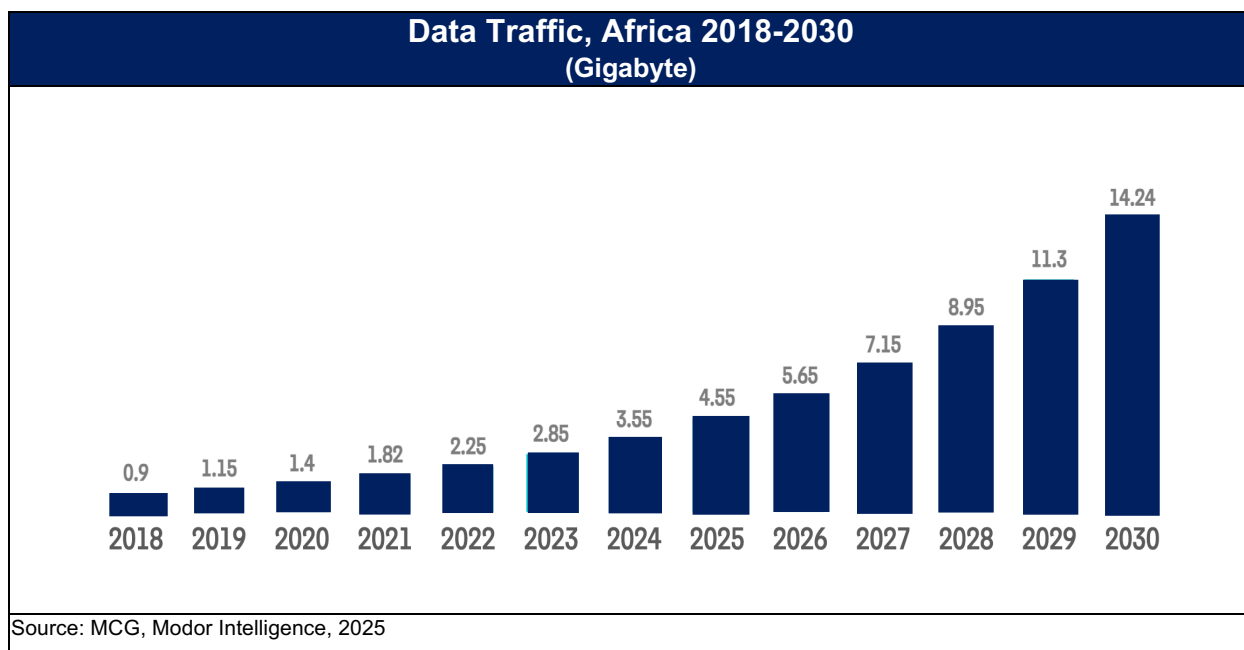
### **Policy-Driven Digital Transformation**

Africa's data center and AI infrastructure ecosystem are at an early but rapidly accelerating stage, driven by digital transformation policies, population-scale connectivity initiatives, and rising enterprise cloud adoption. With a population exceeding 1.4 billion, including one of the world's youngest demographics, Africa represents an immense latent demand for compute and digital services. Key governments—such as Nigeria, Kenya, South Africa, Egypt, and Morocco—have launched multi-billion-dollar national digital strategies emphasizing e-government, fintech inclusion, and public sector digitization. These initiatives are increasing demand for domestic, AI-capable data infrastructure. The continent had less than 250 MW of total IT load in 2020, but is expected to exceed 1,200 MW by 2030, according to Xalam Analytics. This growth is anchored by the rise of regional connectivity (e.g., Equiano and 2Africa subsea cables), local cloud service providers, and public-private partnerships that involve telecoms, governments, and international investors.

### **Infrastructure Scaling & AI Use Cases**

AI-specific infrastructure in Africa is still nascent, but critical use cases are emerging in agriculture, healthcare diagnostics, logistics, and public safety. Countries like Rwanda, Ghana, and Kenya are deploying machine learning for disease prediction, crop analytics, and fraud detection, often through partnerships with academic institutions and global nonprofits. These workloads require not only edge compute, but also regionalized, low-latency core data centers to aggregate, process, and serve model results. International cloud providers like Microsoft Azure and AWS have launched cloud regions in South Africa and are evaluating expansion in North and West Africa. However, concerns over data sovereignty, latency, and foreign control have prompted a new wave of local sovereign cloud initiatives. Nigeria's National Information Technology Development Agency ("NITDA"), for example, now mandates that all sensitive government data be hosted domestically. These shifts are driving demand for Tier III+ facilities with AI-ready configurations—supporting GPU clusters, sovereign storage, and compliance-layer tooling.

### Investment Landscape & Key Players



The investment landscape is evolving rapidly, with Africa attracting both global capital and local entrepreneurial energy. In 2023–2025 alone, more than US\$2.5 billion in new data center capital was committed across Africa, with projects led by Liquid Intelligent Technologies, Africa Data Centres, Raxio, Open Access Data Centres, and hyperscale-backed joint ventures. Notably, Equinix and Digital Realty have entered the African market through acquisitions or partnerships, seeing the continent as both a growth opportunity and a long-term hedge against saturation in Western markets. While funding AI-specific centers remains a challenge due to power constraints, supply chain delays, and land ownership issues, forward-looking investors are building modular campuses designed to scale into AI loads over time. Power reliability remains a critical issue: more than 50% of new sites are expected to require onsite renewable generation, backup storage, or energy-as-a-service models. In this landscape, early entrants with access to sovereign clients, sustainable design frameworks, and long-term capital will be best positioned to shape Africa’s AI-native digital infrastructure future.

## **LATIN AMERICA ACCELERATES**

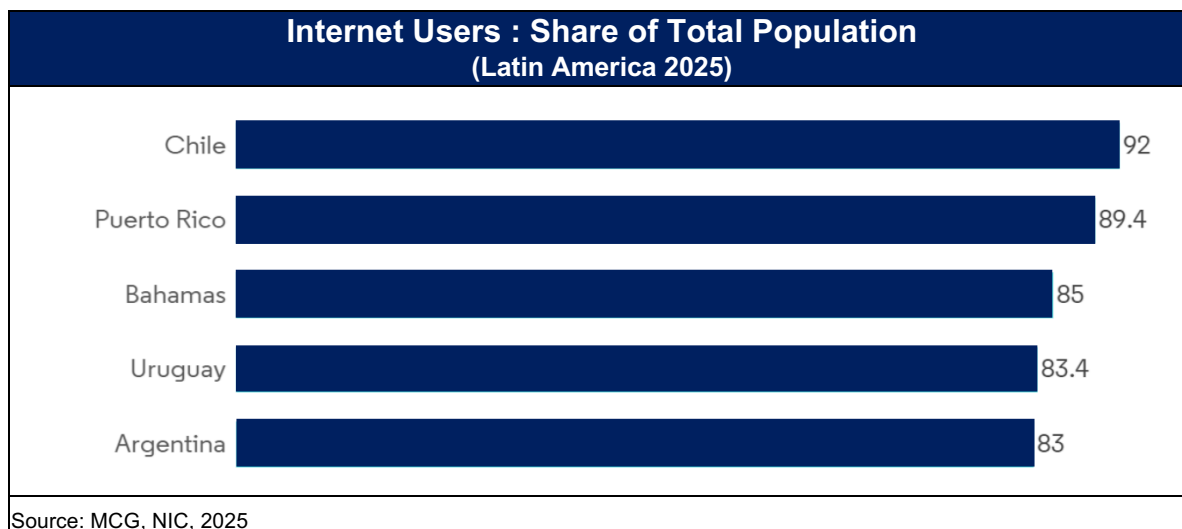
### **Market Drivers & National Priorities**

Latin America's data center industry is entering a high-growth phase, driven by surging internet penetration, e-government platforms, fintech innovation, and the early adoption of AI across both public and private sectors. Regional governments—including Brazil, Mexico, Chile, and Colombia—are investing heavily in digital transformation, which requires robust domestic infrastructure to support cloud computing, data analytics, and AI model training. Brazil leads the region, with over 60% of all hyperscale and colocation capacity, while Chile and Mexico are emerging as strategic connectivity hubs due to their stable macroeconomic policies and submarine cable access. The rapid expansion of digital banking, AI-powered logistics, and citizen-facing services (like tax and ID platforms) is creating a wave of demand for low-latency, sovereign-controlled, and AI-ready data centers. This is accelerating deployments in Tier II cities, while also prompting the entrance of global hyperscalers such as Google Cloud, AWS, and Oracle, who are co-locating AI infrastructure in the region to comply with latency and data residency requirements.

### **AI as a Development Catalyst**

AI is becoming a strategic imperative in Latin America, not only for commercial applications but also for national development goals. Countries like Brazil and Argentina are using AI for agricultural optimization, traffic management, energy forecasting, and fraud detection—creating highly localized data flows that require regional hosting and inference infrastructure. These use cases are producing data gravity—the tendency for compute to follow where data is generated—which in turn is shifting enterprise IT architecture toward edge data centers and modular regional nodes. The rise of Spanish- and Portuguese-language LLMs is further amplifying demand for local model training environments. Simultaneously, sovereign cloud mandates are evolving. Chile and Mexico have introduced guidelines for financial and public-sector data localization, while Brazil is proposing legislation to regulate AI system accountability and cross-border data transfers. Together, these legal frameworks are motivating developers and institutional investors to prioritize AI-native, jurisdiction-compliant infrastructure across urban and peri-urban markets.

## Investment & Infrastructure Expansion



Investment into Latin America's data center sector has reached new heights, with over US\$4.5 billion committed between 2022 and 2025 from a mix of global hyperscalers, infrastructure funds, and regional telcos. Equinix, Scala Data Centers, Ascenty (Digital Realty), and ODATA have announced major campus expansions, particularly in São Paulo, Santiago, and Querétaro, with plans to support AI compute clusters, renewable power integration, and liquid cooling technologies. The continent's energy landscape—rich in hydro, solar, and wind—also provides a natural foundation for sustainable AI infrastructure, especially in Chile and Colombia. However, challenges remain: permitting delays, land availability, and underdeveloped grid resilience in some areas can slow execution. Nonetheless, forward-thinking developers are pursuing modular design, pre-fabricated campuses, and smart grid partnerships to mitigate these risks. As Latin America moves into its next digital era, early-stage investors and infrastructure partners who embed AI-readiness, compliance, and power strategy into their deployment model will be best positioned to capture this underserved but fast-evolving market.

## LOOKING FORWARD

At M Capital Group (“MCG”), MCG view this moment in the data center industry as a once-in-a-generation structural realignment. This report goes far beyond the headlines of “AI growth” or “digital acceleration”, it maps out how the physical architecture of intelligence is being redrawn at global scale. What excites MCG most is that this transformation is not simply about more servers or more cloud, it is about reconfiguring the way economies, governments, and capital interact in the age of artificial intelligence.

The most striking insight from MCG’s research is how quickly data centers have evolved from backend IT assets into frontline instruments of sovereignty, security, and strategic advantage. Around the world, MCG see governments rewriting laws to keep critical data onshore. Hyperscalers are being forced into sovereign-compliant infrastructure partnerships. Infrastructure capital is moving not just toward the largest markets—but toward the most “trusted” and jurisdiction-ready opportunities. AI is the accelerant, but sovereignty, regulation, and sustainability are the scaffolding. This is no longer optional. It is the baseline.

MCG’s analysis shows that markets like Qatar, Saudi Arabia, and the UAE are not only building sovereign cloud infrastructure—these countries are leapfrogging legacy paradigms by directly embedding national AI strategies into their data center blueprints. Meanwhile, in North America and Europe, trillion-parameter models are demanding facilities with 3–5x the power density of the last decade. Operators are now racing to secure grid access, long-term PPAs, and water-efficient cooling at a pace never before

seen in infrastructure. MCG believe this “AI-industrial era” is the next frontier for sovereign capital, infrastructure funds, and digital real estate developers alike.

The strategic implications for ALL are enormous. M Capital Group sees emerging platforms that combine high-density AI readiness, sovereign compliance, ESG alignment, and public-private financing as the new gold standard. Whether it's enabling secure compute for central banks, building national AI campuses, or investing in regional edge facilities governed by local ministries, MCG are moving into a world where infrastructure will be judged by trust, not just throughput. What makes this report essential reading is that it does not simply admire the growth. It dissects it. It links demand to regulation, regulation to design, and design to capital. For investors, policymakers, and developers who recognize that data is the new oil, but who understand that oil now requires an AI refinery, MCG believe this report is the blueprint.

The AI era will not be cloud-native. It will be infrastructure-native. And those who can localize, optimize, and operationalize at scale will define the next two decades of digital power.

---

### Bibliography

<sup>1</sup>Sources: from various sources - M Capital Group Research, including data from Data Center Market Size, Share & Industry Analysis, July 2025; Data Center Market Size, Share, and Trends 2025 to 2034, May 2025; Data Center Services Market, June 2024; 215 Data Center Stats for 2025, July 2025; DOE Releases New Report Evaluating Increase in Electricity Demand from Data Centers, December 2024; AI is Eating Data Center Power Demand – and It's Only Getting Worse, May 2025; Data Centre Industry in India, July 2025; AI Power: Expanding Data Center Capacity to Meet Growing Demand, October 2024; Data Centers Aren't Just a U.S. Trend, July 2025; OpenAI Introduces Data Residency in Asia, May 2025; Unlocking the Data Centre Opportunity in the Middle East, April 2025; Data Center Financing Structures, June 2025; STACK Infrastructure's \$6B Green Investment Milestone Drives Sustainable Data Center Growth, May 2025; Precedence Research's "Data Center Market Size, Share, and Trends 2025 to 2034" (2025, May 23); McKinsey & Company's "Data Center Demands" (2025, May 20); Omedia's "Omedia Market Radar: Sovereign Cloud" (2025); McKinsey & Company's "AI power: Expanding data center capacity to meet growing demand" (2024, November 29); Horizon's "Europe AI Data Center Market Size & Outlook, 2024-2030" (2025); Modor Intelligence's "APAC Data Center Power Market Size & Share Analysis - Growth Trends & Forecasts (2025 - 2030)" (2025 June 16); Horizon's "Middle East & Africa AI Data Center Market Size & Outlook", 2024-2030" (2025); Modor Intelligence's "Africa Data Center Market SIZE & SHARE ANALYSIS - GROWTH TRENDS & FORECASTS UP TO 2030" (2025 January 9); Modor Intelligence's "Latin America Data Center Construction Market Size & Share Analysis - Growth Trends & Forecasts (2025 - 2030)" (2024 October 3)



## **M Capital Group**

**[www.mcapital-group.com](http://www.mcapital-group.com)**

<b>NEW YORK</b>	<b>LONDON</b>	<b>DUBAI</b>
1330 Av of the Americas Level 23 New York, NY 10152 United States Tel: +1 212 634 6831 Fax: +1 212 634 7474	The Leadenhall Building 122 Leadenhall Street London, EC3V 4AB United Kingdom Phone: +44 207 256 4246 Fax: +44 207 256 4050	Emirates Towers P.O. Box 31303 Dubai United Arab Emirates Phone: +971 4 319 7460 Fax: +971 4 330 3365

**Our Word is Our Bond**