ASIA PACIFIC FOUNDATION OF CANADA
FONDATION ASIE PACIFIQUE DU CANADA

ABAC
APEC Business Advisory Council

ABAC

Powering the Digital Economy:

THE DATA CENTER DILEMMA

TABLE OF CONTENTS

FOREWORD 3

ACKNOWLEDGMENTS 4

EXECUTIVE SUMMARY 5

SECTION 1: INTRODUCTION 8

SECTION 2: THE SUSTAINABILITY IMPERATIVE 10

SECTION 3: DIGITAL GROWTH'S ENERGY DILEMMA 13

- A. Energy Resources 13
- B. Grid Readiness 25

SECTION 4: UNLOCKING DATA CENTER EFFICIENCY 29

- A. Energy Efficiency Under Pressure 29
- B. Hardware and Software Efficiency 35
- C. Untapped Opportunity: Waste Heat Reuse 37

SECTION 5: ALIGNING POLICY AND REGIONAL COOPERATION 39

- A. Uneven Policy Instruments and Implementation 39
- B. Regional Cooperation and Public-Private Partnership 41

RECOMMENDATIONS 44

CONCLUSION 46

ANNEX 47

- A. AI and the Digital Economy 48
- B. Economy Chapters 50

FOREWORD

This year, the APEC Business Advisory Council (ABAC) is focused on "Bridge. Business. Beyond," emphasizing the need to bridge digital divides, foster business growth, and look beyond immediate challenges to achieve long-term prosperity in the APEC region.

In 2024, the digital economy represented 15% of global GDP and is its fastest-growing source. Artificial intelligence (AI) is accelerating this growth, transforming how we live and work, and offering both immense opportunities and complex challenges. ABAC's work is centered on creating the enabling conditions for businesses of all sizes to harness AI for inclusive economic growth.

One of the biggest barriers to AI adoption today is the lack of sufficient compute power — a constraint tied to the capacity of data centers and power grids. Energy demand within the digital economy is driven by how we use technology. The mass adoption of streaming platforms, online gaming, and cloud storage — combined with billions of devices in constant operation — has already fueled unprecedented demand for data centers. The rise of AI now compounds this trend: its compute-intensive workloads are accelerating both the scale and intensity of power consumption. Data centers are the infrastructure response and are one of the fastest-growing sources of electricity demand worldwide.

By 2030, power usage from AI-intensive data centers is expected to quadruple. Data centers could drive up to 50% of electricity demand growth in high-demand economies like the U.S. by 2030. This raises a critical question: How do we scale AI and the digital economy without compromising APEC's climate goals or deepening the digital divide?

The digital divide isn't just about access to technology anymore — it's about capital access, grid resilience, and human capacity. In economies without sufficient energy and digital infrastructure, adopting AI risks widening existing inequalities.

To help economies navigate these challenges, ABAC recently endorsed *The Burn-to-Earn Index*, a data insights tool developed by Tufts University and SGTech that benchmarks the economic and emissions impact of 126 digital economies. It provides an important tool for economies to start considering their competitive place in the digital economy and their sustainable path forward.

In line with APEC's 2025 theme, sustainably powering the digital economy has been at the forefront of discussions at ABAC and at other multilateral forums this year. In July 2025, ABAC business leaders endorsed a Declaration on Sustainable AI Investment and Infrastructure, calling on APEC to join this collective effort to advance a future-ready, responsible AI ecosystem across APEC.

Against this backdrop, this report takes stock of the data center dilemma, considers policies and regulations to support sustainable digital economy development, highlights best practices and innovative solutions being deployed in APEC economies, and proposes actionable strategies for both policymakers and private sector leaders.

We hope this report will act as a guide to powering and enabling the digital economy for all APEC members.

JAN DE SILVA

Chair, ABAC AI & Digital Innovation Working Group, ABAC Canada

RAYMOND KYUHO LEE

Chair, ABAC 2025, ABAC Korea

ACKNOWLEDGMENTS

This report was sponsored by the APEC Business Advisory Council (ABAC) and authored by the Asia Pacific Foundation of Canada. The project team is thankful for the close collaboration with ABAC Singapore and SGTech, and for the support and contributions of ABAC members and staffers from the nine focus economies. A special thank you to several contributors for their input, including:

Robert (RJ) Johnston, *School of Public Policy, University of Calgary*

Dennis Cha, *REV Renewables*

Innovation, Science & Economic Development

Ontario Power Generation

Natural Resources Canada

Wong Wai Meng

Dr. Yonggi Kim, *STEPI*

Uptime Institute

Finally, thank you to the Asia Pacific Foundation of Canada project team:

LEAD AUTHOR

Serena (Thi) To, *Consultant*

PROJECT LEAD

Amanda Doyle, *Senior Program Manager, APEC Networks*

EDITORS

Amanda Doyle

Erin Williams

Ted Fraser

Vina Nadjibulla

PROJECT SUPPORT

Ashley Brownlie

DESIGN

Chloe Fenemore

ABAC
APEC Business Advisory Council

ASIA PACIFIC FOUNDATION OF CANADA
FONDATION ASIE PACIFIQUE DU CANADA

ABOUT ABAC

Recognizing the integral role of business and the value of representative business advice on key issues, APEC leaders established the APEC Business Advisory Council (ABAC) in November 1995. ABAC is the sole private-sector arm of APEC with a mandate to advise leaders and other APEC officials on issues of interest to business.

This private-sector body presents recommendations to APEC leaders in an annual closed-door dialogue and advises APEC officials on business-sector priorities and concerns. ABAC meets four times a year; ABAC representatives also attend senior officials' meetings, the annual ministerial meeting and the sectoral ministerial meetings. ABAC members are appointed by their respective economic leaders and represent a range of business sectors, including micro, small and medium enterprises (MSMEs).

ABOUT THE ASIA PACIFIC FOUNDATION OF CANADA

The Asia Pacific Foundation of Canada (APF Canada) is an independent not-for-profit organization focused on Canada's relations with Asia. Our mission is to be Canada's catalyst for engagement with Asia and Asia's bridge to Canada.

APF Canada is dedicated to strengthening ties between Canada and Asia through its research, education, and convening activities, such as the Canada-in-Asia Conferences series, our Women's Business Missions to Asia, and the APEC-Canada Growing Business Partnership project. APF Canada serves as Canada's Secretariat for several APEC networks, including the APEC Business Advisory Council, Pacific Economic Cooperation Council, and serves as one of Canada's designated APEC Study Centers.

VISIT OUR WEBSITE ▸

EXECUTIVE SUMMARY

As digital transformation accelerates across the APEC region, data centers are becoming foundational to economic growth, innovation, and connectivity. These facilities now support a wide range of services: from AI development and cloud computing to government operations and critical digital infrastructure. Energy demand in the digital economy is fundamentally driven by patterns of use. The rapid uptake of streaming, online gaming, and cloud storage — layered on top of billions of devices running continuously — has already created extraordinary demand for data center capacity. As demand grows, so too does the need to understand how data center expansion intersects with energy systems, sustainability goals, and long-term infrastructure planning.

AI infrastructure is driving unprecedented electricity demand across APEC. A single large-scale AI cluster now requires as much power as an urban center such as Vancouver or Kuala Lumpur. Electricity demand from AI-specific workloads could quadruple by 2030, while two-thirds of the additional energy generation required is expected to come from fossil fuel sources. Today, data centers can consume more than 20% of grid capacity in some areas — and that's growing fast. Many economies lack an overall awareness of their energy demand gaps, and infrastructure is struggling to keep pace. Globally, one in five planned data center projects is at risk of delay because of power constraints. Grid connection wait times are getting longer, and building out grid infrastructure is slow — new transmission lines can take four to eight years in some APEC economies.

These developments point to the strategic importance of aligning digital infrastructure planning with energy, climate, and innovation agendas. Without clear, coordinated policy direction, this trajectory risks locking in high-carbon systems and excluding emerging economies from digital opportunity.

Within APEC, some markets already have 100,000 times more data servers than others. In many emerging economies, less than 0.1% of national electricity goes to digital infrastructure, while over one billion people still lack access to reliable electricity. The new digital divide is not about technology. It's about digital infrastructure — specifically capital access, grid resilience, and human capacity.

This report, *Powering the Digital Economy: The Data Center Dilemma*, was developed by the Asia Pacific Foundation of Canada (APF Canada) and commissioned by ABAC to provide policymakers and business leaders with a clear, evidence-based assessment of the challenges and opportunities associated with the growth of digital infrastructure in the region. The analysis focuses on nine APEC economies — Australia, Canada, Chile, China, Indonesia, Korea, Malaysia, Singapore, and the United States — chosen for their relevance to global data center markets and as a representation of APEC's diverse climates, geography, and energy mixes.

THE REPORT HIGHLIGHTS EIGHT KEY FINDINGS:

1. Clean energy remains a strategic goal for powering the digital economy across APEC. Yet many economies still depend on fossil fuels to power data centers due to reliability, cost, and political direction.

2. A reliable energy supply is a critical enabler of digital growth in APEC, but most power systems are not yet equipped to support the rapid expansion of digital infrastructure, with power planning and grid upgrades lagging behind demand.

3. Operational reliability is non-negotiable for data centers. While most hyperscale cloud firms have committed to 100% renewables, efforts to deploy 24/7 clean power are concentrated in regions where supportive policies, natural resources, robust infrastructure, and favorable market conditions make it feasible.

4. Data centers require around-the-clock stable energy sources to operate, causing nuclear energy, hydrogen, and geothermal to emerge as solutions offering a low-carbon, reliable baseload to balance intermittent renewables like wind and solar. However, policy uncertainty, site development challenges, and high capital costs hinder broader adoption across the region.

5. Energy efficiency gains are falling behind the pace of AI-driven data center growth. While advanced cooling technologies and more efficient use of IT resources (chips, servers, storage or software) offer room for improvement, high upfront costs and slow R&D adoption remain barriers. Waste heat remains a missed opportunity due to limited upstream integration, while fragmented operational metrics and inconsistent standards across APEC are hampering efforts to improve energy efficiency. Without coordinated public investment, stronger policy incentives, or procurement standards that value efficiency, many hardware innovations risk being underutilized.

6. While several governments are introducing targeted policies to promote cleaner, more efficient data centers, inconsistent regulations and uneven implementation across regions are creating a fragmented and unequal landscape for sustainable digital infrastructure.

7. Regionally, coordinated power-sharing and cross-border clean energy corridors are strengthening grid resilience, unlocking private investment, and supporting low-carbon digital growth, yet differences in regulatory frameworks, grid infrastructure, and market design continue to limit the effectiveness of public-private partnerships to scale cross-border digital infrastructure.

8. If the rapid expansion of data centers is left unmanaged, the unchecked growth could turn data centers from enablers of the digital economy into liabilities — threatening inclusive digital progress across APEC economies.

ABAC recommends APEC Ministers and Leaders work with the private sector to accelerate the development of sustainable AI infrastructure, including data centers, computing infrastructure, water supply, and power grids to support innovation and climate goals. To do so, APEC policymakers should work closely with ABAC, industry leaders, and relevant regional and global bodies to:

1. Measure, benchmark, and assess the economic value and environmental impact of the digital economy, using benchmark tools such as The Burn-to-Earn Index.
2. Collaborate with industry and energy providers to forecast data center demand against energy capacity and prioritize timely capacity development.
3. Monitor, support, and deploy sustainability and efficiency advancements in digital infrastructure, including advanced cooling, heat reuse, and integrated zone design, and pursue next-generation chips and semiconductors to drive operational efficiency. For policymakers, the opportunity lies in sending clear demand signals — through funding, regulation, or standards — to accelerate the adoption of technologies that can meaningfully reduce energy and water use across the data center sector.
4. Monitor and adopt global sustainability standards and set ambitious, measurable targets for green data centers in APEC.

The private sector plays a pivotal role in advancing energy-efficient, climate-aligned digital infrastructure across the APEC region. To support a resilient and sustainable digital economy, the private sector should lead by:

1. Catalyzing efficiency in AI systems by investing in low-power hardware and deploying software more efficiently.
2. Participating in utility-led grid flexibility programs by shifting workloads that are not time-sensitive (e.g., AI training, background processing) during peak hours to off-peak periods or renewable-rich windows, while maintaining 24/7 uptime for mission-critical services, supporting grid stability and decarbonization.
3. Embracing transparency by reporting energy consumption and emissions linked to AI training and digital operations, reinforcing accountability and enabling informed decisions across the value chain.
4. Adopting recognized data center-related certifications and standards and tracking performance-based metrics to improve operational efficiency and align with climate and sustainability targets.
5. Integrating heat recovery systems into facility design and partnering with local utilities or communities to convert server waste heat into usable energy, transforming a byproduct into a shared low-carbon resource

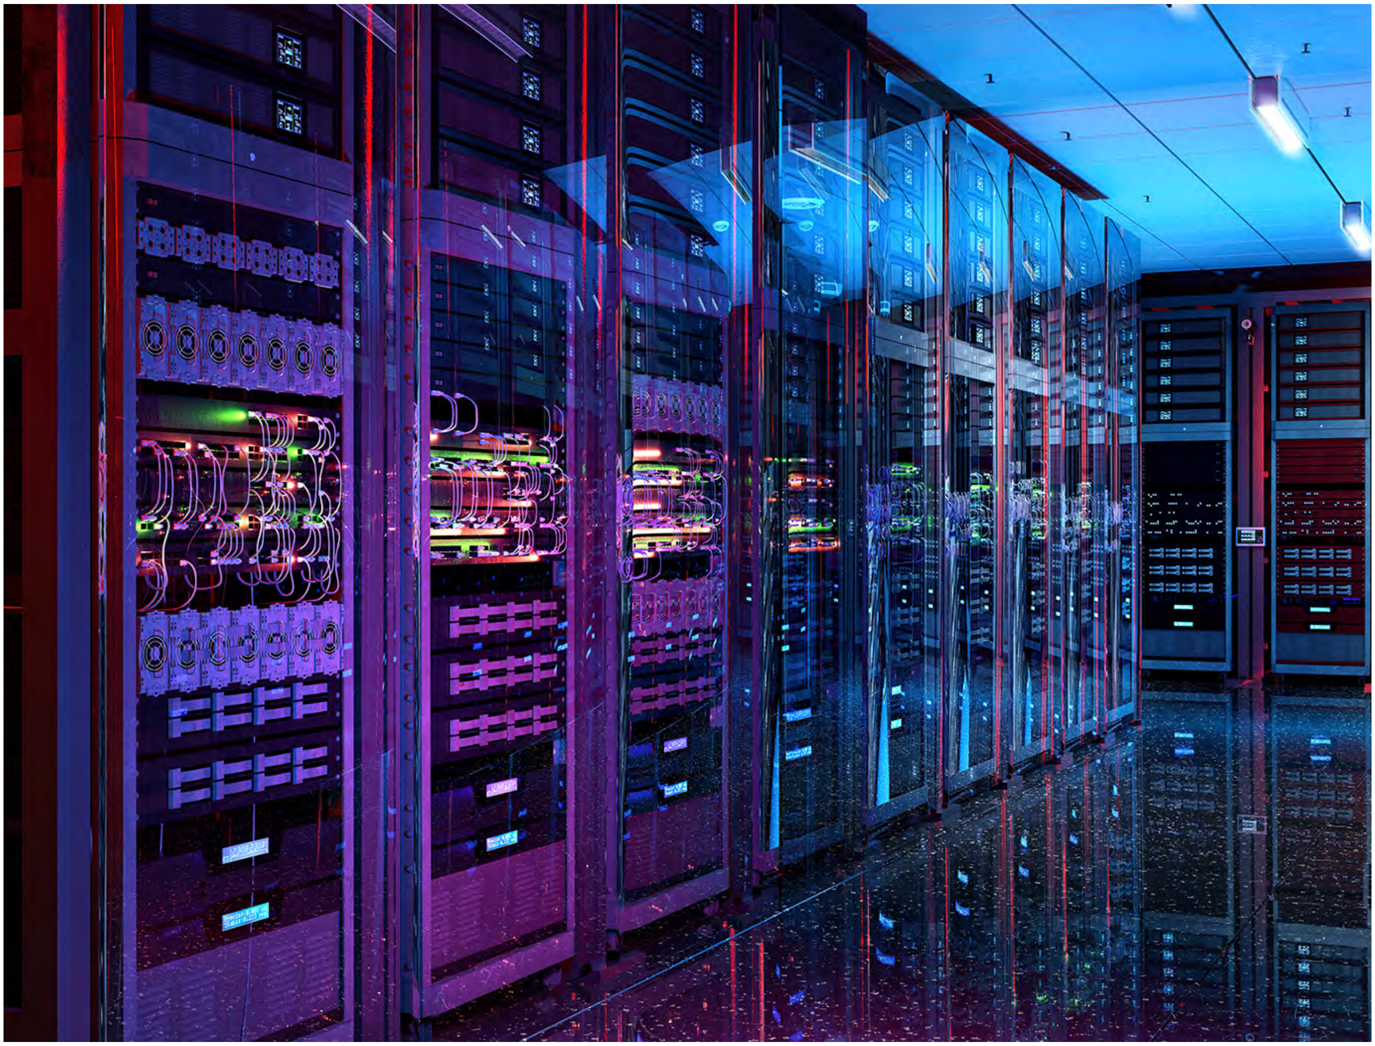

The future of AI in APEC is not just about speed or scale — it's about sustainability and shared prosperity. This report aims to support collaborative policy dialogue across APEC by offering practical insights into how digital infrastructure can be powered more sustainably. It emphasizes the importance of integrating data center planning with energy strategies, encouraging policy coherence, and advancing regional cooperation on metrics, standards, and best practices. By coordinating efforts across economies and sectors, APEC has the opportunity to shape a digital economy with infrastructure that is secure, sustainable, and aligned with long-term economic resilience.

SECTION 1:

INTRODUCTION

The digital economy today accounts for over 15% of global GDP (US$11–16 trillion) and is growing 2.5 times faster than the physical economy. This rapid growth is fueled by a wave of emerging technologies — from artificial intelligence (AI) and 5G to the Internet of Things (IoT) and cloud services — which drive soaring demand for data processing, storage, and real-time connectivity. To function at scale, these digital services rely on high-density, energy-intensive infrastructure — namely, data centers. Data centers are needed to deliver reliable computing power and 24/7 connectivity across borders.

As APEC economies deepen their digital transformation, the Republic of Korea's APEC 2025 key priority — "Building a Sustainable Tomorrow: Connect, Innovate, Prosper" — reinforces the imperative to integrate sustainability, innovation, and regional equity into the foundations of digital infrastructure. In this context, data centers are emerging as a critical backbone of the digital economy: one that must now confront rising energy and data demands, emissions constraints, and growing infrastructure pressures.

The digital economy presents immense opportunities for APEC economies, but that promise is unevenly distributed. While two-thirds of the global population is now online and almost 95% live within reach of mobile broadband, the next frontier of the digital divide lies not in connectivity, but in digital infrastructure readiness: server density, energy affordability, and grid readiness. Billions of dollars are being invested in data centers globally to meet the growing power generation demands of AI and emerging technologies, but the question remains: how can economies sustainably power them?

This report was developed by the Asia Pacific Foundation of Canada (APF Canada) on behalf of the APEC Business Advisory Council (ABAC) in a joint effort to explore how data center growth can be effectively aligned with energy system planning, sustainability objectives, and digital inclusion across APEC economies. It provides an evidence-based assessment of the infrastructure, regulatory, and environmental implications of scaling digital infrastructure, particularly in the context of AI-driven workloads. The report provides innovative solutions and best practices being piloted in APEC to make data centers more efficient and more sustainable.

Figure 1: Digital Economy Value per Internet User, 2023

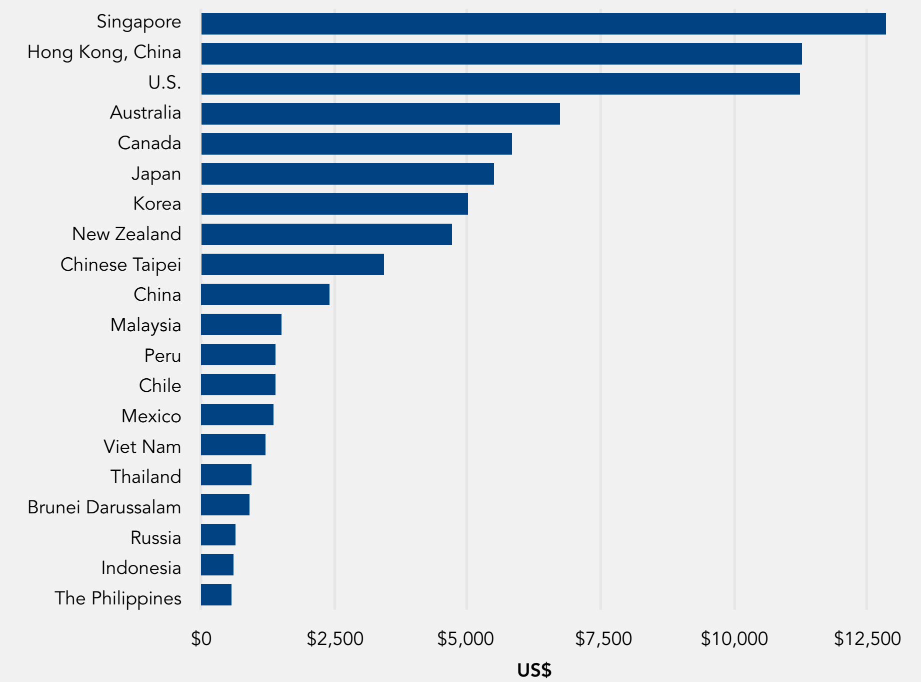

Source: The "Burn-To-Earn" Index Report · The data is sourced from The 'Burn-to-Earn' Index Report: The Global Digital Economy's Economic and Environmental Impact and Trade-offs. Data for Papua New Guinea is not available in the report.

Drawing on 150+ policy documents and expert interviews, the report focuses on nine APEC economies where data center growth is prevalent: Australia, Canada, Chile, China, Indonesia, Korea, Malaysia, Singapore, and the United States. The analysis is structured around key policy levers and system-level insights: digital economy growth, electricity supply and grid readiness, permitting and land use, efficiency and cooling, clean energy access, emissions and water use, and regulatory frameworks for data centers and sustainability. The report highlights both common challenges and regionally specific constraints, and provides targeted recommendations for APEC policymakers and the private sector to collectively advance coherent, future-ready digital infrastructure across the APEC region.

SECTION 2:

THE SUSTAINABILITY IMPERATIVE

Data centers are specialized facilities that host interconnected computing systems — servers, storage, networking, and cooling infrastructure — designed to process, store, and transmit vast volumes of digital information. They underpin every digital interaction, from mobile payments and video streaming to AI model training and medical diagnostics.

For years, data centers maintained a relatively stable energy footprint despite the expansion of cloud services, largely due to improvements in energy efficiency. However, in 2017, this balance began to shift. Between 2018 and 2022, electricity consumption by the 13 largest data centers operators — the majority of them located in APEC economies — more than doubled. Global data center capacity is set to increase nearly threefold in just five years, reaching 219 GW by 2030 at which point, global data centers are projected to consume as much electricity as Japan does today.

Over 95% of the world's internet traffic passes through data centers.

Figure 2: Global Data Center Demand Growth

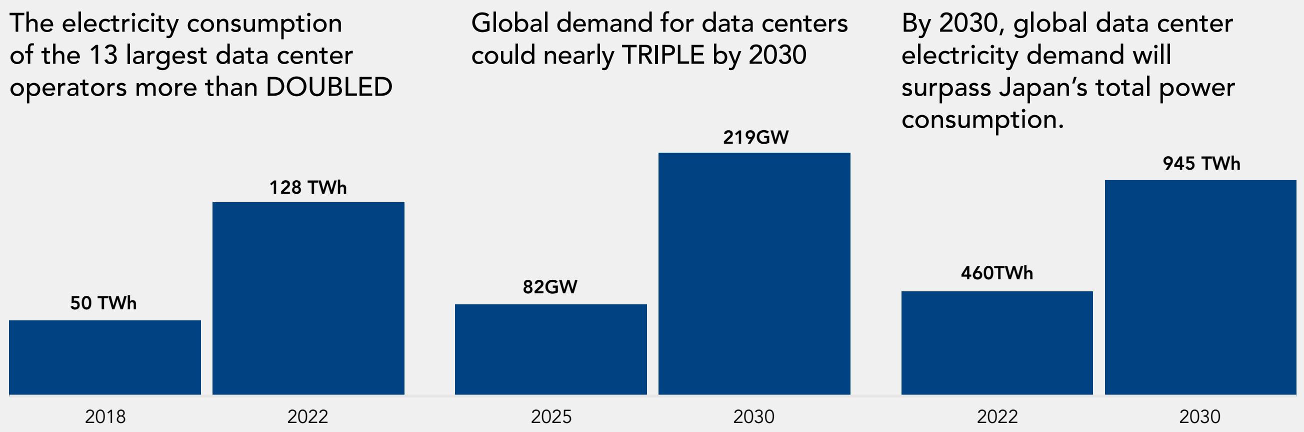

Source: McKinsey Quarterly "The cost of compute: A $7 trillion race to scale data centers", UNCTAD "2024 Digital Economy Report", IMF May 2025 Chart of the Week

A key driver of rising data center demand is the rapid growth of AI workloads. AI-optimized data centers are becoming markedly more energy-intensive per square foot, driven by the shift to ultra-dense Graphic Processing Unit (GPU) infrastructure, specialized semiconductor chips optimized for parallel computing in AI, but significantly more power- and cooling-intensive than conventional processors. By 2030, AI-related compute is expected to account for over 70% of total global data center capacity, with electricity demand from AI-intensive facilities projected to quadruple. Notably, inference workloads, the continuous "thinking" phase of AI, are anticipated to represent 80–90% of AI-related energy consumption, signaling a sustained and structural increase in baseline power demand.

ANNEX A - "AI AND THE DIGITAL ECONOMY" ▸

1 in 5 planned data center projects worldwide are at risk of delay due to power constraints.

These shifts are already placing significant strain on power systems. While data centers currently consume approximately 1.5% of global electricity, their localized impact is far more pronounced, accounting for 7% of national demand in Singapore, 4% in the United States, and a striking 26% in Virginia, the major global data hub in the U.S. In some high-demand jurisdictions such as Northern Virginia (the U.S.), data centers now use over 20% of grid capacity, and this share is rising. Yet supporting infrastructure is not keeping pace. Globally, one in five planned data center projects is at risk of delay due to power constraints. Grid connection wait times are increasing, while building new transmission infrastructure typically takes four to eight years, far slower than the pace of digital demand. Most economies have yet to fully assess or anticipate the energy demands of their evolving digital sectors.

The rapid expansion of digital infrastructure is also bringing a new wave of sustainability challenges to the forefront. Without policy intervention, AI-driven data center expansion could raise carbon emissions by 1.2% globally, and up to 5.5% in high-demand economies like the U.S by the end of the decade. High-density AI workloads are placing mounting pressure on freshwater systems used for cooling, critical materials required for specialized chips and embodied carbon resulting from rapid hardware turnover. These impacts are particularly acute in water-stressed regions and across global supply chains. Without integrated planning, the rapid expansion of digital infrastructure risks diverting scarce resources, straining ecosystems, and slowing digital transformation in already-constrained regions.

If left unmanaged, the unchecked growth could turn data centers from enablers of the digital economy into liabilities, threatening inclusive digital progress across APEC economies.

Figure 3: Growth Forecast for AI and Non-AI Workloads

■ AI workload ■ Non-AI workload

By 2030, AI workloads will grow 3.5x, and drive 70% of all data center power demand.

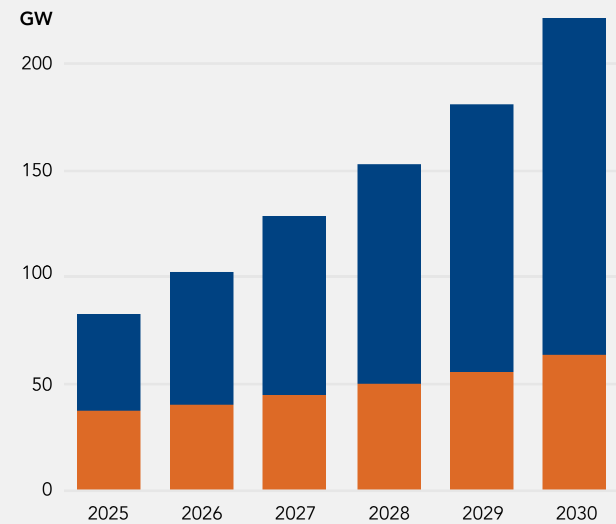

Source: McKinsey Quarterly "The cost of compute: A $7 trillion race to scale data centers"

Figure 4: APEC Data Centers by Region

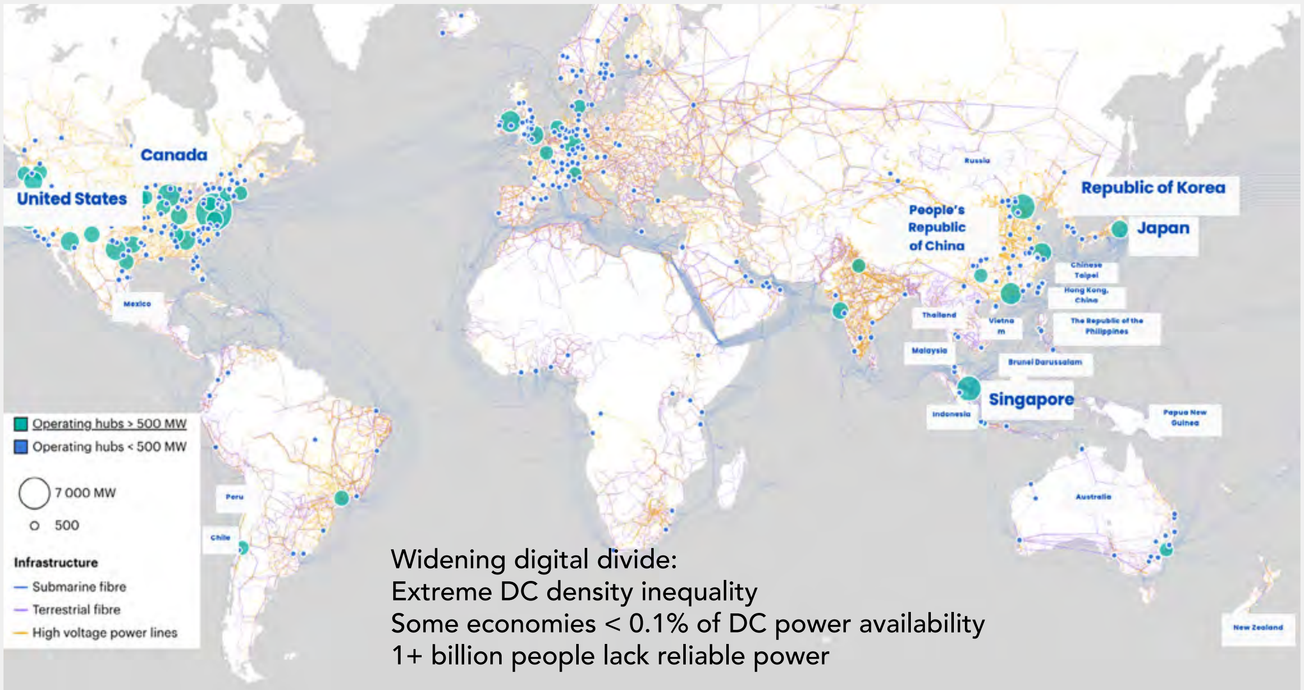

Source: International Energy Agency "Energy & AI Observatory", accessed September 2025, IDCA Global Digital Economy Report

The New Digital Divide in APEC: Digital Infrastructure

The digital economy presents immense opportunities for APEC economies, but that promise is unevenly distributed. Data centers are heavily concentrated in a few economies: the U.S. and China together account for approximately 70% of global installed data center capacity. Hyperscale hubs in these economies now regularly exceed 500 MW, and AI-driven facilities are reaching unprecedented levels, with recently announced projects targeting 1 GW+. By contrast, data center capacity is highly uneven across emerging economies: some still measure national capacity at under 50 MW, with individual facilities often just 5–20 MW at most.

In terms of electricity consumption, some emerging economies allocate less than 0.1% of national electricity to support digital infrastructure — underscoring the limited energy capacity to support data centers and cloud services. This divide is further reflected in compute density: several economies operate with fewer than 100 servers per million people — just 1/30th of the baseline (3,000+) needed to sustain a modern digital economy — Some advanced economies host up to 100,000 times more compute capacity. This stark imbalance highlights the urgency for policies that expand digital infrastructure more equitably across the APEC region, ensuring that the benefits of the digital economy are accessible to all.

ANNEX B: ECONOMY-SPECIFIC PROFILES ▸

SECTION 3:

DIGITAL GROWTH'S ENERGY DILEMMA

A. Energy Resources

UNEVEN RESOURCES, UNEQUAL OPPORTUNITIES

Figure 5: Electricity Generation and Energy Mix by Economy (GWh, 2022)

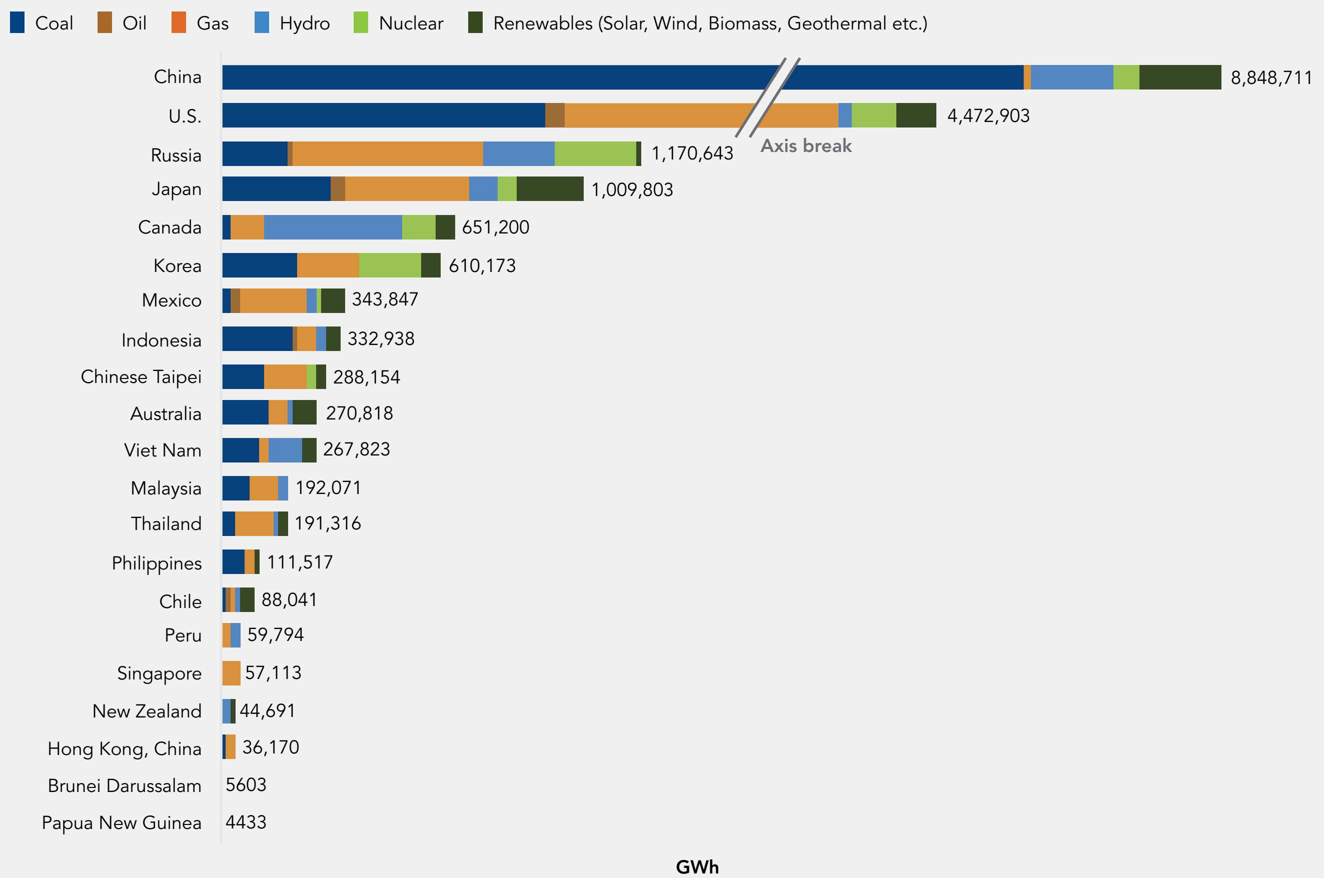

Source: APEC Expert Group Energy Data Analysis

The energy gap is deepening inequality and limiting participation in the digital economy. APEC economies start from vastly different positions when scaling data centers, shaped by disparities in energy self-sufficiency, energy mix, and land availability. Some APEC economies (e.g., the U.S., Canada, Australia) benefit from abundant domestic energy resources, giving them a stronger foundation to power energy-intensive data centers. Others (e.g., Japan, the Republic of Korea, Singapore, and several Pacific Island economies) rely more heavily on energy imports, making them more sensitive to global market dynamics and supply chain fluctuations. Land availability adds another layer of complexity. Economies with more geographic space are often better positioned to expand data center infrastructure, deploy large-scale renewables, and modernize power grids. In contrast, densely populated or geographically constrained economies (e.g., Hong Kong, Singapore, and the Republic of Korea) may face limitations that make it more challenging to integrate clean energy at scale and align digital growth with climate goals.

Figure 6: Energy resources and deployment timelines for data centers

Can match data center timeline   Outside of data center timeline

Bar definition: Estimated completion. Reference point: Assumes projects begin in 2025

Only a few options (e.g. natural gas, solar, onshore wind, and storage) align with data center timelines; most other energy sources lag by a decade or more.

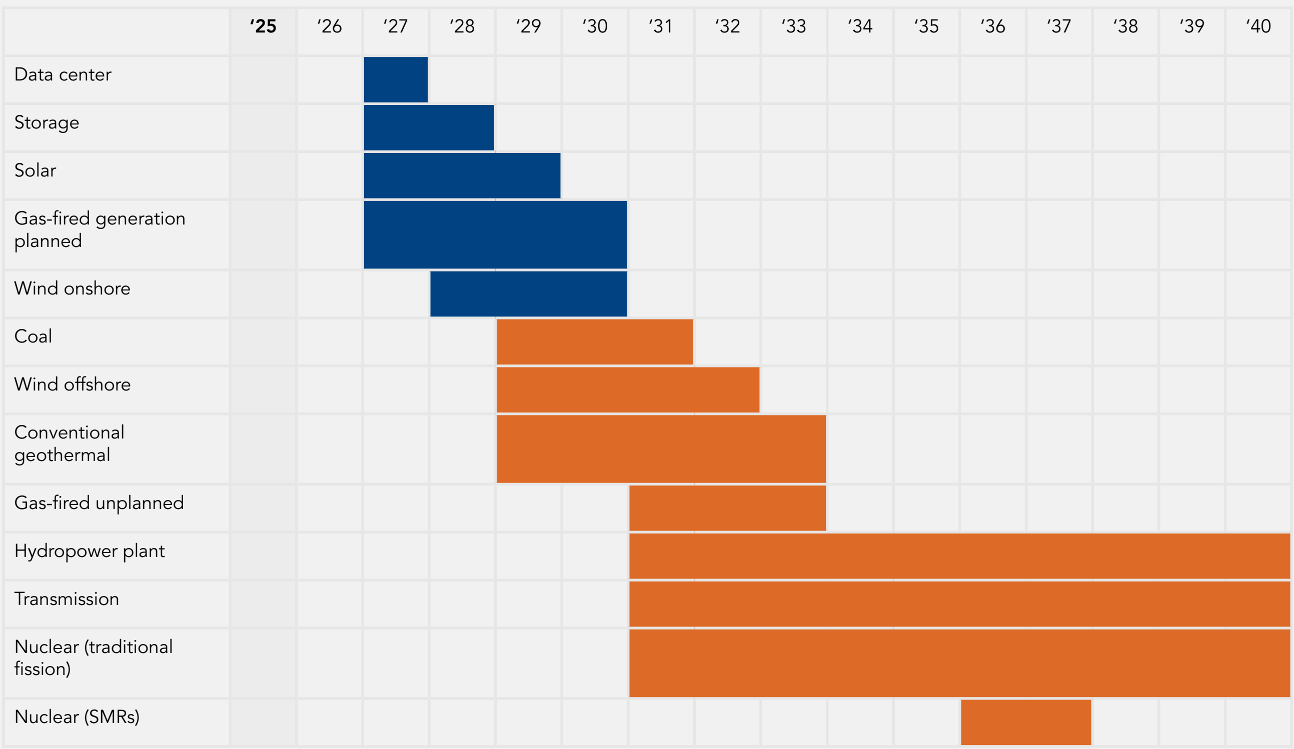

Source: Deloitte Research Center for Energy & Industrials · Assumes the project begins in 2025

Only a handful of energy sources — such as storage, solar, gas-fired generation, and some onshore wind — can be deployed quickly enough to align with typical data center build timelines. These are the few technologies that can realistically be brought online within two to three years. Everything else — clean or not — such as transmission upgrades, coal, hydropower, and nuclear, faces deployment timelines that stretch well beyond 2030. That's simply too slow to meet the breakneck pace of AI and cloud infrastructure growth.

As a result, both industry and governments are being forced to consider every option: from extending the life of gas plants to taking strategic bets on next-gen geothermal and small modular nuclear reactors. The clean energy transition is no longer just a carbon challenge — it's a timeline and scalability challenge.

> **"Data centers move fast—but energy infrastructure doesn't"**

FOSSIL FUELS WILL STILL DOMINATE

Clean energy remains a strategic goal for powering the digital economy across APEC. Yet many economies still depend on fossil fuels to power data centers due to reliability, cost, and political direction.

Clean energy is essential to align data center growth with climate goals. Yet as AI accelerates demand for compute, the focus is shifting from energy transition to sheer energy expansion, where reliability often takes precedence, even if it sidelines clean power in the short term.

> "Two-thirds of additional electricity generation through 2035 for data centers is expected to come from fossil fuels."

Figure 7: Global electricity generation for data centers

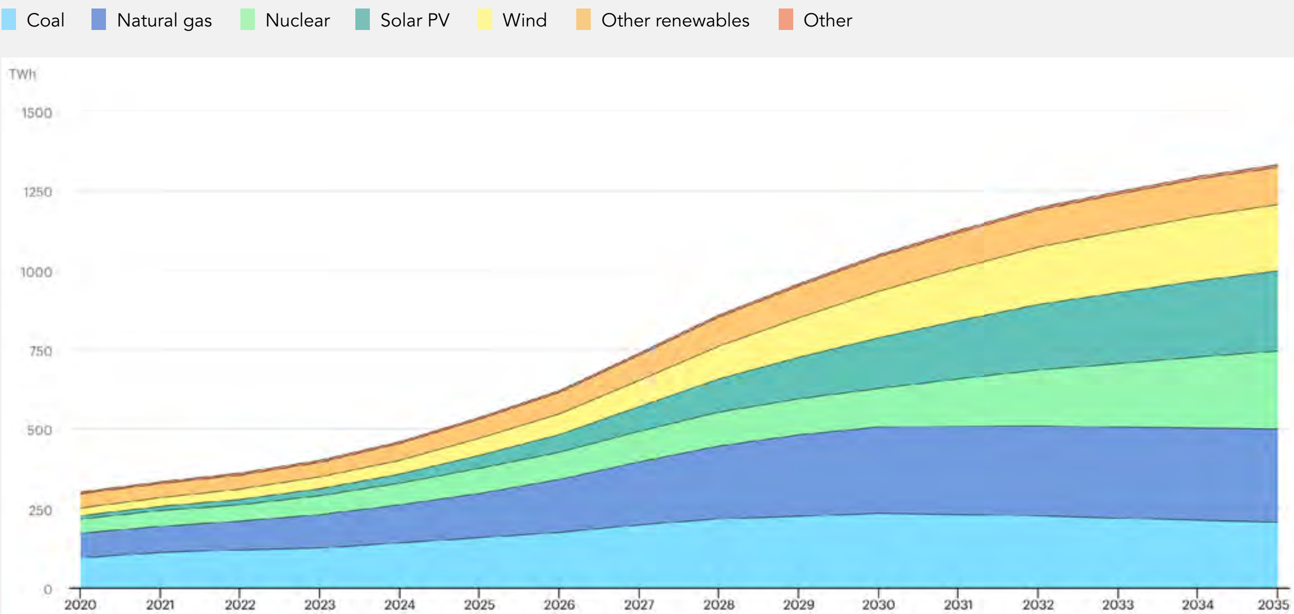

Source: International Energy Agency "Global data center electricity consumption, by equipment, Base Case, 2020-2030".

Despite optimistic projections from international agencies (e.g., IEA) that renewables could supply 50% of new electricity demand from data centers by 2035, fossil fuels, particularly coal and natural gas, still underpin most digital infrastructure across APEC economies today. This reliance is driven by legacy investments, affordability, and grid reliability concerns. In many jurisdictions, coal continues to supply over 60% of the national energy mix, meaning that data centers connected to these grids inherit a carbon-intensive power profile by default. While some economies, such as China, Australia, and Korea, have announced coal phaseout timelines and others are scaling renewable procurement (e.g., Chile, Malaysia), the pace of transition often lags behind the rapid growth of data center demand. Meanwhile, in the U.S., surging demand from AI hubs such as Virginia, Ohio, and Texas has been accompanied by notable increases in gas generation.

Figure 8: Share of coal and natural gas in total electricity generation by economy

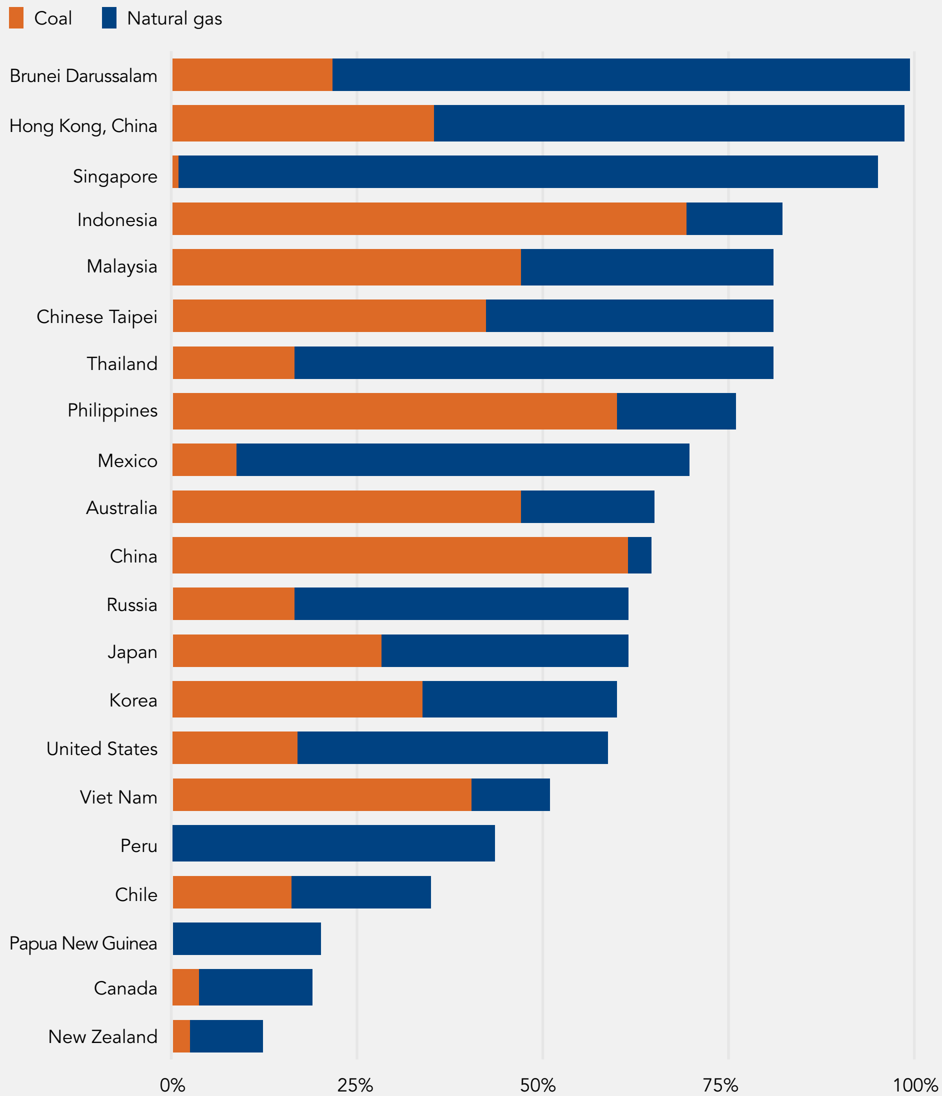

Source: Compiled by author based on International Energy Agency countries' electricity mix: Australia 2023, Canada 2023, Chile 2023, China 2023, Indonesia 2023, Malaysia 2022, Korea 2023, Singapore 2023, the United States 2023, Russia 2022, Mexico 2023, Chinese Taipei 2022, Vietnam 2022, Thailand 2023, Philippines 2022, Peru 2022, Hong Kong, China 2022, Brunei Darussalam 2022, Japan 2023, New Zealand 2023. Papua New Guinea's data is from the APEC Expert Group Energy Data Analysis

RENEWABLES ALONE CAN'T SCALE AI

APEC has set collective goals to reduce energy intensity by 45% by 2035 and to double the share of renewables in its energy mix by 2030 (from 2010 levels), as part of its broader commitment to sustainable and inclusive regional growth. The APEC region's shift toward clean-powered digital growth is underway, but implementation varies widely, shaped by policy, local resource availability, and market incentives.

Electricity demand varies by sector, and data centers, particularly those running AI workloads, are among the least flexible. Unlike EVs or some industrial loads that can shift operations to align with clean energy availability, AI systems require continuous, interruption-free power. Yet renewable sources such as solar and wind are inherently variable, averaging just six to nine hours of generation daily, making them unable to guarantee the uninterrupted power these AI facilities require. The result: renewables are essential but not sufficient to meet the 24/7 uptime demands of AI infrastructure. Despite renewables reaching over 46% of installed capacity in a record-breaking 2024, the world is still behind the COP28 goal of tripling renewables by 2030 (a shortfall of nearly 0.9 terawatts). The current pace simply isn't fast enough to keep up with the accelerating energy appetite of AI and digital infrastructure.

> **"Roughly 40% of new data center capacity additions are expected to come from renewables, yet their intermittency means that either gas or nuclear is still needed to secure 24/7 uptime for AI workloads."**

> **"The core bottleneck with renewables lies in integration."**

Structural barriers — such as transmission delays, limited land near demand centers, and supply chain disruptions — are slowing renewable deployment for data centers. Data centers can be built in one to two years; transmission lines often take nearly a decade. This mismatch constrains the integration of renewables, especially in regions where clean generation is remote from urban digital hubs. Limited land availability also constrains large-scale solar and wind deployment, particularly in urbanized or geographically constrained regions, further limiting the co-location of renewable energy and computing infrastructure.

Figure 9: Generation capacity for wind and solar varies widely throughout the year

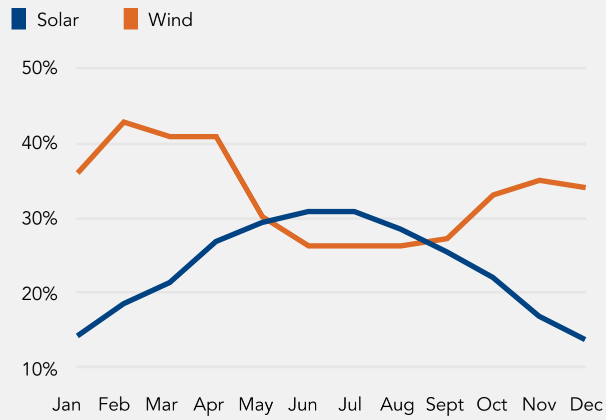

Source: Goldman Sachs "Is nuclear energy the answer to AI data centers' power consumption?" · US capacity factor averages for wind and solar in 2023

At the same time, supply chains for solar panels, grid equipment, and storage technologies remain vulnerable to global price shocks and trade disruptions, with material costs like copper and steel rising over 40% in the past five years. These barriers, when combined, make it far easier and faster to expand fossil fuel capacity, deepening the reliance on gas and coal to sustain digital growth in the near term.

A 24/7 renewable power solution (e.g. wind + solar + lithium battery storage to cover nighttime) can cost over [$200](#) per MWh. This is roughly double or triple the wholesale electricity price in many APEC markets. Using new technologies in energy storage could potentially cut that cost below $100/MWh, but even $100 is considerably higher than typical grid power rates.

"Round-the-clock renewable power and storage solutions are still financially out of reach in markets without subsidies or strong policy incentives."

Best Practice 1: Chile, Malaysia, Indonesia

CHILE: SCALING RENEWABLES TO ATTRACT DIGITAL INVESTMENT

With 60% of its electricity already sourced from renewables, Chile's energy transition — anchored by ambitious targets of 80% renewable generation and 2 GW of storage by 2030, and a full coal phase-out by 2040 — offers a best-practice model for building a low-carbon digital economy. The launch of its National Data Center Plan (2024) aims to triple the industry within five years, attract $2.5 billion in investment, decentralize facilities, and embed sustainability at the core, leveraging clean energy leadership to position Chile as a resilient, competitive, and climate-aligned hub for digital infrastructure.

MALAYSIA: INTEGRATE ON-SITE RENEWABLE GENERATION INTO DATA CENTER DESIGN

YTL's Johor Data Center Park — Malaysia's first and largest renewable energy–integrated data center campus — shows how embedding large-scale clean power into operations can redefine industry benchmarks. Spanning 275 acres with nearly 600 MW of capacity, the park dedicates roughly 90% of its footprint to solar infrastructure, delivering traceable, on-site generation for clients. This direct supply model offers greater transparency and carbon impact certainty than renewable energy credits or grid purchases, positioning on-site renewables as both a sustainability driver and a competitive differentiator in Asia and beyond.

INDONESIA: FLOATING SOLAR – SCALING CLEAN ENERGY WHILE PRESERVING LAND AND WATER

Indonesia is positioning floating solar as a best-practice model for balancing rapid renewable energy expansion with land-use efficiency, particularly in industrial hubs like Batam. With over 20 mapped reservoirs and retention ponds suitable for large-scale deployment, the approach maximizes underused water surfaces while reducing evaporation — a key resilience measure for drought-prone provinces. Flagship projects such as the 2.2-GW Duriangkang Floating Solar Plant, the world's largest by surface area, demonstrate scalability for both domestic supply and cross-border exports to Singapore via subsea cables.

Operational reliability is non-negotiable for data centers. While most hyperscale cloud firms have committed to 100% renewables, efforts to deploy 24/7 clean power are concentrated in regions where supportive policies, natural resources, robust infrastructure, and favorable market conditions make it feasible.

In practice, many data center firms still rely on on-site fossil fuel generation such as diesel or gas to guarantee the uninterrupted uptime AI requires. To guarantee net-zero accounting, some hyperscalers are piloting "24/7 Carbon-Free Energy" (CFE) initiatives: matching clean energy generation with data center demand in real time. These pilots depend on advanced tools: time-stamped Purchasing Power Agreements (PPA), energy storage, and local renewable integration.

Figure 10: Target Years for Net-Zero and Carbon Neutrality Among Data Center Operators

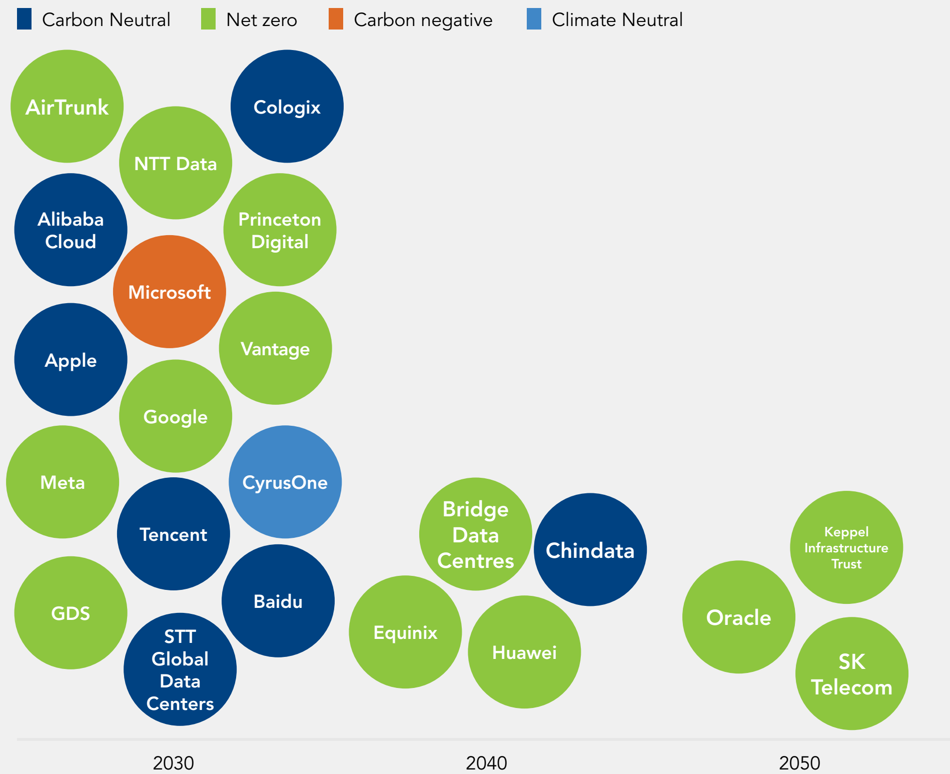

Source: Compiled by the author based on International Energy Agency's Energy and AI Report and companies' websites. The net zero and Carbon Neutrality targets are for Scope 1 and Scope 2 emissions.

But such solutions are concentrated in jurisdictions with enabling conditions: robust policy frameworks, flexible grid infrastructure, accessible renewable supply, and competitive markets. In Australia, Japan, Singapore, Chinese Taipei, and parts of Canada and the U.S., hyperscalers have launched renewable energy projects and time-matched procurement pilots. For example, Google's real-time solar PPA in Chinese Taipei was enabled by recent policy reforms; in Singapore, rooftop solar and Renewable Energy Certificates (RECs) help compensate for limited land and a fossil-heavy grid.

For many APEC economies, these enabling mechanisms either don't exist or remain heavily concentrated in capital cities or special economic zones. Regulatory frameworks often lack enabling mechanisms, storage standards, or systems for verifying renewable use. Electricity generation and supply remain largely controlled by state-owned, vertically integrated power providers, limiting opportunities for corporate clean energy procurement. Second- and third-tier cities, which increasingly host digital infrastructure, often face outdated grids and scarce renewable integration. This limits geographic flexibility and reinforces data center clustering in already congested energy hubs.

Table 1: Mechanism for renewable procurement

Mechanism	Definition	Real time clean energy
PPA (Power Purchase Agreement)	Long-term contract to buy electricity from a renewable generator	Not Real-time
Time-stamped PPA (24/7 PPA)	Advanced PPA with hourly tracking of clean energy generation and usage	Real-time
REC (Renewable Energy Certificate)	Certificate for 1 MWh of renewable energy produced	Not Real-time
Virtual PPA	A non-delivery financial contract: buyer pays a fixed rate, receives RECs, and settles price differences—no physical power exchanged.	Not Real-time
Carbon Offset	Credit for funding emissions reductions/removals elsewhere	Not Real-time
Energy Storage	Systems to store excess renewable energy for later use	Real-time
Local Renewable Integration	On-site solar, wind, or geothermal tied directly to data center	Real-time
Rooftop Solar	Solar PV panels installed on-site or on buildings	Real-time

Source: Compiled by author from publicly available definitions in International Renewable Energy Agency (IRENA) and Environmental Defense Fund (EDF) publications.

Best Practice 2: Singapore, Australia, Malaysia

SINGAPORE: CLEAN POWER WITHIN RESOURCE CONSTRAINTS"

Singapore offers a pragmatic model for clean digital operations in land- and resource-constrained settings. The country deploys rooftop solar to help compensate for limited land and a fossil-heavy grid. Renewable mechanisms such as RECs are also in place. For example, Google's data center maximizes its clean energy mix through a hybrid strategy, tapping local solar and supplementing with renewable energy certificates (RECs) from regional suppliers such as Malaysia and Indonesia. Amazon has two solar energy projects with Sembcorp and Sunseap that generate 79.6 MW, enough to power almost 200,000 Singaporean homes annually. Singapore is also advancing regional cross-border clean energy trade by forging long-term partnerships with neighboring economies to secure a diversified, low-carbon electricity supply.

AUSTRALIA: ON-SITE BATTERY INTEGRATION TO SUPPORT GRID RESILIENCE

Australia's Quinbrook "Supernode" project in Queensland exemplifies a forward-looking model for integrating digital infrastructure with clean energy systems. This development co-locates up to four hyperscale data centers with a 250 MW battery, situated adjacent to the South Pine substation. The battery is designed to absorb excess solar generation—particularly from Australia's high rooftop solar penetration—and discharge during peak periods, reducing strain on the grid. By directly buffering data center load, the project enhances energy reliability while demonstrating how on-site storage can play a strategic role in meeting Australia's digital and decarbonization goals.

MALAYSIA: CORPORATE RENEWABLE ENERGY PROCUREMENT TO DRIVE DECARBONIZATION

Malaysia is aligning its long-term climate goals—45% GHG emissions intensity reduction by 2030 (from 2005 levels) and 70% renewable capacity by 2050—with market-enabling mechanisms that expand corporate access to clean power. In 2022, the government launched the Corporate Green Power Programme (CGPP) and the Corporate Renewable Energy Supply Scheme (CRESS) in 2024 to facilitate large-scale solar adoption via corporate PPAs. These programs create a structured pathway for businesses to directly procure renewable electricity, accelerating private-sector investment in the country's energy transition.

NUCLEAR POWER IS GAINING RENEWED MOMENTUM ACROSS APEC

Data centers require around-the-clock stable energy sources to operate, causing nuclear energy, hydrogen, and geothermal to emerge as solutions offering a low-carbon, reliable baseload to balance intermittent renewables like wind and solar. However, policy uncertainty, siting challenges, and high capital costs hinder broader adoption across the region.

After years of uneven progress across regions, nuclear energy is regaining traction — particularly in the APEC region, where restarts, new builds of conventional reactors, and the emergence of SMRs are accelerating. Nuclear currently produces 10% of global generation and is the second-largest source of low-emissions electricity today after hydropower. Momentum has returned: At COP28, more than 20 economies pledged to triple global nuclear capacity by 2050 (from 2020 levels), while multilateral lenders such as the World Bank and Asian Development Bank are revising policies to enable nuclear financing for the first time in decades.

Across APEC, eight economies operate nuclear facilities. The level of development varies widely — from mature fleets in the U.S., Canada, Korea, and China to early-stage or exploratory programs in Indonesia and Malaysia. This diversity shapes the region's capacity to integrate nuclear power into digital infrastructure planning.

Figure 11: Share of nuclear energy in total electricity generation by economy

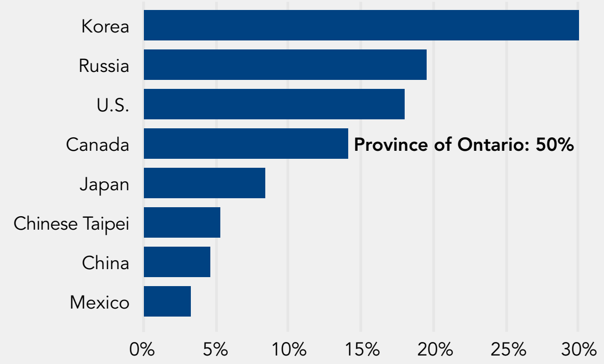

Source: Compiled by author based on International Energy Agency countries' electricity mix: Australia 2023, Chile 2023, China 2023, Indonesia 2023, Malaysia 2022, Korea 2023, Singapore 2023, the United States 2023. Canada's data is from Natural Resources Canada.

Nuclear's comeback reflects a strategic response to the surging energy demands of growing digital infrastructure.

Data centers — especially AI-driven ones — need clean, stable, round-the-clock power. Nuclear provides precisely that: capacity factors above 90%, predictable output, and the ability to stabilize grids with growing shares of intermittent renewables. A single large reactor can supply multiple hyperscale facilities, and pairing nuclear with renewables could deliver both reliability and carbon-free power.

The United States has emerged as the most active economy where hyperscalers are directly engaging with nuclear power to secure future supply. In 2024 alone, hyperscalers secured over 10 GW of nuclear PPAs, many tied to large reactor restarts and SMR developments. For example, Microsoft signed a 20-year PPA with Constellation Energy to restart the 835 MW Crane Clean Energy Center for its AI data center loads, and plans to repurpose the Three Mile Island nuclear site for AI infrastructure. AWS committed $20 billion to two data center hubs in Pennsylvania, including one adjacent to the Susquehanna nuclear plant, aiming to source up to 960 MW of nuclear power. The growing need for clean, firm power in the U.S. has also driven renewed interest in restarting reactors previously shut down for economic reasons. In addition to Microsoft's Three Mile Island deal, NextEra Energy is considering restarting the Duane Arnold nuclear power plant to help meet projected regional demand growth. A notable shift over the past two years is that U.S. data centers are no longer simply purchasing nuclear output — they are increasingly playing a direct role in enabling the development, restart, and commissioning of nuclear projects.

In China, nuclear expansion is accelerating as part of its broader industrial decarbonization strategy and growing digital infrastructure demand. Roughly half of all reactors under construction are located within China. Under the 14th Five-Year Plan (2021–2025), China aims to reach 70 GW of installed nuclear capacity by 2025, supported by parallel investments in SMR technology. In 2024 alone, China approved eleven new reactors, signaling the strongest political backing for nuclear. The flagship Xuwei nuclear complex in Jiangsu exemplifies this shift, combining two Hualong One reactors with a high-temperature gas-cooled reactor to supply both electricity and industrial steam. These developments align with China's push to stabilize the power supply for its "East Data, West Compute" strategy, where data centers in energy-abundant western regions increasingly require firm, zero-carbon baseload power. Although China has not publicly tied nuclear energy to its data center agenda, the pace of investment and technological innovation indicates that nuclear is re-emerging as a strategic backbone for powering AI and cloud infrastructure over the coming decade.

In other APEC economies, nuclear expansion strategies are tied to broader national electricity demand growth — driven by sectors such as EVs, data centers, and crypto — within government-led energy programs. For example, in Ontario (Canada), where nuclear supplies nearly 50% of power, the province is directing multiple large-scale builds. Projects under development include the Port Hope/Wesleyville site, which could host up to 10,000 MW of new capacity, enough to power the equivalent of 10 million homes, Bruce C to host up to 4,800 MW of power, and four SMRs at Darlington. Partnerships such as Microsoft's with Ontario Power Generation, which links clean energy credits to nuclear and hydro on an hourly basis, show that hyperscalers are engaging with local clean energy programs. Besides new builds, Ontario is also refurbishing the Darlington Nuclear Generating Station and Pickering Nuclear Generating Station, projects that will extend operations by up to 30 years.

Meanwhile, Japan has moved to extend and modernize its existing nuclear fleet. Amendments to the Electricity Business Act now permit reactors to operate beyond 60 years, enabling the restart of 3.3 GW of capacity in 2023–2024 and supporting discussions on future SMR deployment.

Nuclear Reality Check: Momentum meets hard timelines

Despite renewed momentum, nuclear power alone will not solve the near-term energy crunch facing AI-driven data center growth. Nuclear is increasingly viewed as a cornerstone of the next generation of clean, firm power—but it will complement, not replace other alternatives such as natural gas, renewables, and battery storage in meeting soaring electricity demand. Goldman Sachs estimates that meeting projected AI-related power needs by 2030 would require 85–90 GW of new nuclear capacity, yet less than 10% of that is expected to be operational by decade's end.

Even in the most efficient jurisdictions, nuclear power takes time. Construction timelines typically span around six years in China, seven to eight years in Korea, and more than ten years in the U.S. and Canada. However, data centers operate on much shorter timelines—typically requiring new power capacity within two to five years to keep pace with rapid demand growth. This mismatch highlights why nuclear power cannot currently address the short-term supply gap. The core bottlenecks are structural: complex licensing, capital-intensive financing, lengthy permitting, and public resistance over safety, waste disposal, and cost overruns.

When It Works: A few economies have shown that policy and regulation can help to contain costs and schedules. In Canada, its CANDU reactor refurbishments have largely stayed on time and within budget. In November 2024, Ontario Power Generation completed the refurbishment of Darlington Unit 1 five months ahead of schedule, adding enough clean power for 875,000 homes. China also maintains impressive efficiency, averaging six-year build cycles—a competitive edge rooted in standardized designs and state-led coordination. Korea has also managed comparatively modest delays under similar centralized oversight with its Saeul 1 and 2 reactors.

The regional picture is mixed. Across much of APEC—particularly in economies without established nuclear frameworks—deployment remains slow. Regulatory complexity, limited institutional capacity, and public skepticism continue to constrain progress, even as technology, financing tools, and project management practices improve.

"SMRs are viewed as a potential nuclear option for data centers, but none are likely to be operational before 2030."

Beyond large-scale reactors, hyperscalers are increasingly turning to small modular reactors (SMRs) as a potential long-term solution to their growing power needs. SMRs are designed as compact, factory-built units that can be co-located with data centers, delivering clean and continuous power. In theory, they could reduce transmission losses and align with the ultra-high uptime requirements of AI and cloud infrastructure.

Momentum is strongest in the U.S., where most announced SMR projects for the data center sector are concentrated. SMR developers are increasingly signing power supply agreements with utilities and technology companies, with over 32 GW of commitments announced since 2023. Notably, Google's partnership with Kairos Power aims to explore SMR deployment across multiple U.S. sites.

Figure 12: Global Announcements on Nuclear Energy for Data Centers

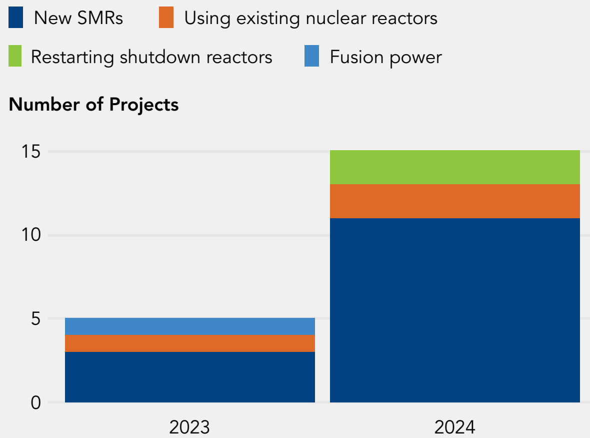

Source: International Energy Agency "The Path to a New Era for Nuclear Energy".

Investor interest is growing, supported by the expectation that SMRs could eventually play a role in powering the next generation of AI-driven infrastructure. However, the technology remains several years away from commercialization at scale. Cost projections suggest that SMRs may generate electricity at higher costs than both natural gas and conventional nuclear, and most estimates still reflect First-of-a-Kind (FOAK) economics rather than proven, repeatable deployment. Supply chain constraints add to these challenges. Most SMR designs rely on high-assay low-enriched uranium (HALEU), a specialized fuel with limited global production capacity.

Despite numerous exploratory agreements between hyperscalers and nuclear developers, as of August 2025, no SMR has entered commercial service for data centers, and their overall contribution is expected to remain limited before 2030.

Best Practice 3: Canada, the U.S.

CANADA – LEVERAGING NUCLEAR FOR STRATEGIC LOADS

Ontario's power mix — primarily nuclear and hydro — is over 90% carbon-free. Ontario's Energy for Generations 2025 Strategy positions data centers as strategic loads tied to clean baseload sources. Large-scale expansions at Bruce Nuclear and new siting work signal a long-term commitment to scalable, firm power. Ontario Power Generation (OPG) is also building the world's first grid-connected SMR at Darlington, with three more planned. Ontario now sees hyperscalers not as grid stressors but as demand anchors that support long-term planning. This shift is backed by nuclear-linked clean energy credits (CECs), which now include nuclear and support projects like Microsoft's 10-year CEC agreement.

THE U.S. – GOOGLE'S SMR DEPLOYMENT STRATEGY FOR 24/7 CARBON-FREE ENERGY

In October 2024, Google announced a strategic partnership with Kairos Power to deploy SMRs at multiple U.S. sites. The initiative is designed to provide 24/7 carbon-free electricity to support Google's data centers and corporate operations. The first unit is expected to be operational by 2030, with a total deployment goal of up to 500 MW by 2035. This project positions SMRs as a viable solution for meeting the continuous, high-capacity energy needs of hyperscale infrastructure while advancing national clean energy and innovation objectives.

HYDROGEN AND GEOTHERMAL ARE EMERGING AS NEXT-GEN POWER SOLUTIONS

Green hydrogen — produced using renewable electricity — can deliver clean, dispatchable, and scalable power through fuel cells. These systems are well-suited for remote sites such as Keppel's floating data centers in Singapore. Hydrogen fuels respond more quickly than diesel generators, helping data centers stay online during sudden power disruptions, and can back up critical workloads with a faster ramp-up than diesel. Hydrogen also complements intermittent renewables by filling in power gaps. Early projects like ECL's 1 MW hydrogen-powered data center in California, with plans to scale to 1 GW in Texas, show growing momentum. However, most projects still rely on grey hydrogen made from fossil fuels, limiting climate benefits. Wider adoption will depend on lowering costs, expanding green hydrogen supply, and building supporting infrastructure.

> **"Hydrogen is emerging as a promising zero-carbon energy source for AI data centers, especially in off-grid or grid-constrained areas."**

> **"New drilling and reservoir technologies are making geothermal energy a scalable, clean power option for data centers, especially in APEC regions with untapped geothermal potential"**

Figure 13: Data centers' place in the hydrogen economy

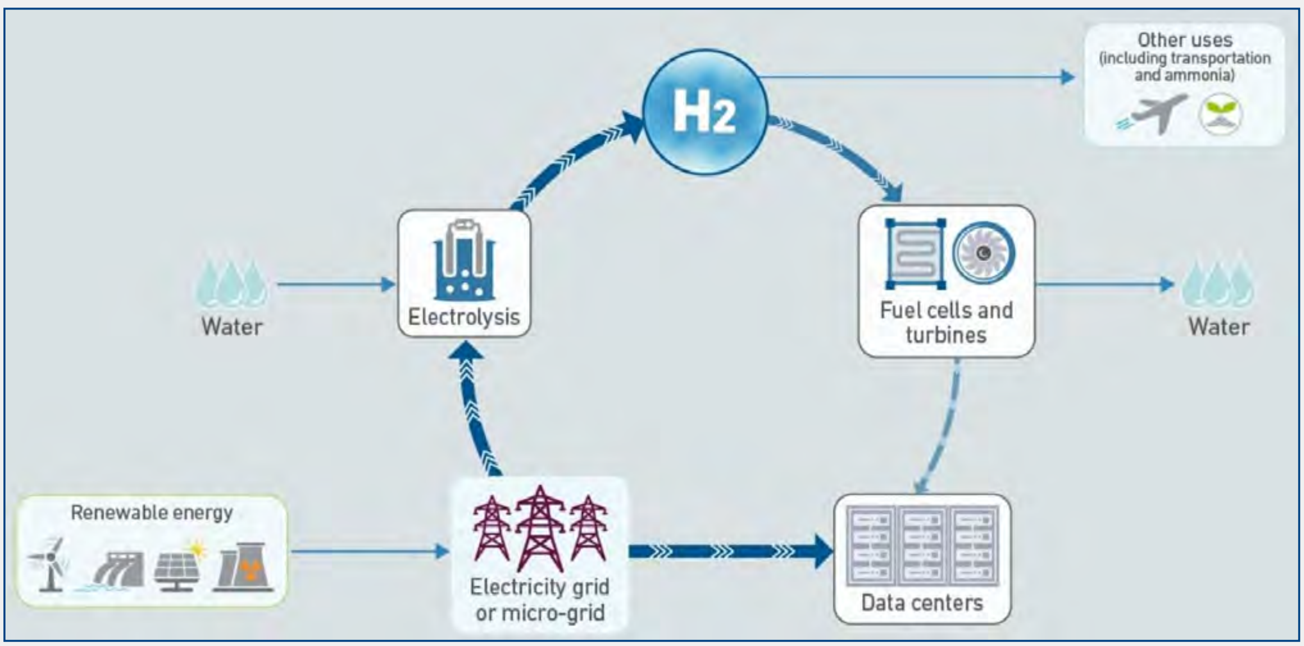

Source: Uptime Institute "Hydrogen in data centers: an introduction"

Next-generation geothermal energy is gaining traction as a reliable, carbon-free power solution for data centers, driven by advances in drilling and reservoir enhancement technologies. In 2023, Google and Fervo Energy launched a 3.5 MW pilot project in Nevada using Enhanced Geothermal Systems (EGS), successfully demonstrating that geothermal can provide consistent, 24/7 electricity. Larger commercial-scale projects are now underway, including Corsac's planned 115 MW facility and Meta's 150 MW Sage initiative. For APEC economies rich in geothermal resources — such as Indonesia, the Philippines, and New Zealand — these innovations open up new possibilities for powering digital infrastructure with steady, low-emission energy.

Best Practice 4: Korea and Singapore

KOREA: HYDROGEN-POWERED DATA CENTER DEVELOPMENT

Korea is advancing hydrogen as a next-generation energy solution for digital infrastructure. A flagship example is the $1.7 billion project in Dangjin-si, developed by Korea South-East Power Co. and Samsung C&T, which will power a hyperscale data center using hydrogen fuel cells supported by battery storage. This marks one of the first large-scale efforts in Asia to deploy hydrogen specifically for data center operations.

SINGAPORE: HYDROGEN-POWERED FLOATING DATA CENTER

Singapore is exploring innovative models for sustainable digital infrastructure through Keppel Data Centres' proposed Datapark+—a floating data center park powered by hydrogen. The project aims to source liquid hydrogen from Australia, with supply provided by Woodside Energy's H2Perth facility. AirTrunk is also actively exploring hydrogen power as a sustainable energy source to strengthen the environmental performance of its data centers.

B. Grid Readiness

THE GRID CONSTRAINT IS A SHARED CHALLENGE

A reliable energy supply is a critical enabler of digital growth in APEC. But most power systems are not yet equipped to support the rapid expansion of digital infrastructure, with power planning and grid upgrades lagging demand.

Data centers are specialized facilities housing interconnected servers, storage, networking, and cooling systems, designed to process, store, and transmit vast volumes of information. In essence, they underpin virtually every digital interaction. Today, data centers have evolved from back-end IT warehouses into mission-critical infrastructure. But their defining characteristic is clear: they are energy-hungry.

Global electricity demand from data centers is expected to more than triple, rising from under 300 TWh in 2020 to nearly 1,000 TWh by 2030, a load larger than Japan's entire electricity consumption today. As AI adoption accelerates across APEC economies, the ability to access reliable, cost-effective, and clean electricity is becoming a defining factor in site selection.

> **"A single 1 GW AI data center consumes 8.76 TWh/year = equivalent to 800,000 homes in Canada or 1.7 million homes in Malaysia. That's the same energy footprint as the entire city of Vancouver or central Kuala Lumpur."**

However, the grid infrastructure across APEC is already maxed out before the full wave of AI has arrived. The average age of regional power grids spans from 25-30 years in Latin America and Southeast Asia to 40 years in North America. In high-demand APEC economies, connection wait times for critical grid infrastructure have doubled in the past three years, with interconnection delays that can stretch up to eight years. As a result, grid stress has emerged as the industry's top concern: **79% of power and data center executives cited limited grid capacity as the main barrier to expansion through 2035.**

Table 2: Infrastructure Lead Time

Infrastructure Type	Lead Time (Established Economies)	Lead Time (Emerging Economies)
Electricity Transmission Lines	4–8 years	2–4 years
Electricity Distribution Infrastructure	2–4 years	1-3 years

Source: IEA's Energy and AI Report

Across APEC, grid limitations are quietly redrawing the data center map. In Korea, more than half of the 33 licensed data center projects in the Greater Seoul Area are facing grid connection delays — now averaging 12 months, compared to two to three months previously. Along with new capacity restrictions, grid constraints are pushing data center investment toward Ulsan and other cities with faster approvals and more diverse energy options. Singapore's 2019 moratorium has already reshaped the Southeast Asian regional landscape, catalyzing growth in Johor, Malaysia, and Batam, Indonesia, and fostering a cross-border digital infrastructure model. In Chile, persistent transmission bottlenecks — especially along the north-central corridor — limit the flow of renewable power to high-demand hubs such as Santiago, constraining reliable supply for large-scale digital infrastructure.

In the U.S., Northern Virginia's power demand is projected to double within the next decade, largely driven by data centers, and meeting even half of that growth will require massive new generation and transmission build-out. Dominion Energy, the utility responsible for supplying most of Virginia's electricity, has already slowed interconnections in Loudoun County, with some projects in eastern Loudoun effectively delayed until 2026 and beyond due to transmission limits, prompting hyperscalers to shift expansion toward other hubs such as Phoenix and Atlanta. Similarly, in Alberta, Canada, 29 pending data center connection requests total over 16 GW, already surpassing the province's current peak system demand of roughly 12 GW and underscoring mounting pressure on grid infrastructure. To manage this pressure, in June 2025, the Alberta Electric System Operator (AESO) set a cap of 1,200 MW for new large-load connections, such as data centers, through 2028 to protect grid reliability. Elsewhere, grid bottlenecks are combining with market shifts to reshape where capacity can realistically be added, marking a new phase in which energy access, not just demand, increasingly determines the map of the digital economy.

Best Practice 5: China

CHINA - RAPID, TARGETED GRID INFRASTRUCTURE BUILDOUT TO ENABLE DIGITAL INFRASTRUCTURE

China's ability to add 429 GW of new generation capacity in 2024 — paired with a rapid buildout of transmission and distribution infrastructure, and targeted deployment to renewable-rich regions prioritized for data center development such as Guizhou, Inner Mongolia, Gansu, and Ningxia — demonstrates how large-scale, coordinated grid expansion can remove capacity bottlenecks, meet surging AI-related electricity demand, and provide the flexibility to scale energy-intensive workloads without supply constraints.

Case Study 1: Grid Resilience Stress Test in Virginia's Data Center Alley

Northern Virginia's Data Center Alley is nearing 4 GW of demand. In July 2023, a grid fault in Northern Virginia's Data Center Alley triggered a mass disconnection of 60 data centers, dropping over 1,500 MW of load as facilities switched to diesel generators. While designed to protect sensitive equipment from voltage dips, this sudden offloading caused a major grid imbalance. PJM and Dominion Energy were forced to cut power plant output to avoid a cascading failure. This is a near-miss that exposed the growing vulnerability of grids to dense, fast-reacting digital loads. If the grid fault in Northern Virginia's Data Center Alley had spiraled into a cascading failure, it could have meant widespread blackouts affecting homes, hospitals, and businesses across the region.

This event highlighted a new class of grid disturbance. **According to an EPRI (2024) survey**, utilities report rising violations tied to data centers, including thermal overloads (22%), voltage violations (17%), harmonic distortions (9%), and ramp rate issues (26%). One critical concern is the load shedding-feedback loop: When grid disturbances occur, large data centers often switch instantly to backup power via uninterruptible power supply (UPS) systems. While this protects facility uptime, it causes an immediate drop in grid load, sometimes destabilizing local frequency or voltage. Reconnecting data centers after a disturbance presents its own operational risks. If ramp-up occurs too quickly, grid frequency and voltage can again be destabilized.

Case Study 2: Grid Strains in Korea

In October 2022, a fire at a data center in Pangyo (Gyeonggi Province, Greater Seoul Area), one of Korea's most prominent tech hubs, caused widespread outages to essential platforms, including messaging and fintech services, disrupting daily life, business operations, and even transportation systems for millions across the country. The incident highlighted the dangers of overconcentrating critical digital infrastructure in and around Seoul, where grid capacity is already under strain and disaster impacts are amplified. In response, the Korean government advanced policies to decentralize data center development, offering incentives for operators to locate facilities in regional areas. These measures aim to spread economic benefits, reduce systemic risk, improve disaster resilience, and ease energy and land-use pressures in the capital region.

WHY THE GRID CAN'T KEEP UP

Grid resilience is increasingly threatened by a convergence of supply chain bottlenecks, fragmented governance, and rising political and economic barriers

- **Supply constraints:** Grid capacity is not expanding fast enough or in the right locations to meet the pace and distribution of AI and digital infrastructure growth.
- **Governance gaps:** Interconnection processes, the procedure for linking a new electricity user or generator — such as a data center or solar farm — to the power grid, are often opaque and fragmented, with unclear timelines and approval pathways. Developers frequently submit speculative or duplicate applications to hedge against uncertainty, overwhelming queues and inflating demand signals. Many utilities, particularly in emerging APEC economies, lack the technical and staffing capacity to manage complex, high-density loads or plan for resilient, weather-hardened AI infrastructure. AI-driven workloads add further unpredictability, with volatile, traffic-dependent energy patterns that challenge traditional forecasting practices.
- **Economic and political barriers:** Grid congestion costs have surged across the region. Persistent supply chain bottlenecks and delays — such as a big jump in transformer shortages in 2024 — are making essential grid equipment far pricier, now costing about 50% more than in 2020. Meanwhile, new transmission development faces growing public opposition tied to environmental and social concerns.

> "Rising electricity costs are key risks. Hyperscalers saw their electricity bills nearly double between 2019 and 2023"

As grid constraints tighten, data center expansion is putting upward pressure on electricity prices across APEC economies. In several jurisdictions, utility deals with AI companies include discounted rates, shifting the financial burden to households and small businesses. For instance, in Virginia, residential bills could rise by up to $37.50 per month regardless of actual usage. If current trends persist, U.S. electricity prices alone could climb 8.6% by 2030, driven by surging AI demand and lagging renewable deployment. Without accelerated grid and supply-side investments, rising power costs could ripple across the region, undermining digital innovation and broader economic competitiveness.

Case Study 3: Malaysia

MALAYSIA'S ELECTRICITY PRICE REFORMS SIGNAL RISING COSTS AMID DATA CENTER GROWTH

Malaysia, particularly the Johor region, offers significantly lower electricity prices than Singapore, where rates are roughly three times higher. However, rapid data center growth is contributing to upward price pressure. Tenaga Nasional Berhad, Malaysia's sole electricity utility for Peninsular Malaysia, implemented a 14.2% base tariff increase, one of the steepest in recent years, citing the need to fund grid modernization and support the energy transition. Under a concurrent price restructuring, data centers were reclassified into a new "ultra-high voltage" category with higher average rates, estimated to raise industry electricity costs by 10–15%. While the reforms aim to better reflect usage patterns, sustained 24/7 data center demand could add to long-term cost pressures for consumers.

Case Study 4: Korea

TACKLING SPECULATION IN DATA CENTER GRID CONNECTIONS IN KOREA

As grid capacity tightens, especially in the Greater Seoul Area, electricity access has become a high-value asset, sparking an aggressive wave of speculative applications for power. Since 2020, many real estate developers and intermediaries have preemptively reserved large electricity blocks without confirmed data center construction plans, often to inflate land prices, secure grid rights for resale, or gain leverage in future negotiations. Between January 2020 and March 2023, nearly 68% of all data center power requests (678 out of 1,001) came from non-genuine end-users. In extreme cases, one individual filed for 28 different sites, and 33 approved projects failed to proceed to contract within a year. This artificial demand distorted national energy planning, stranded capacity in interconnection queues, and delayed legitimate operators, while increasing the risk of overbuilt transmission and substation assets that may never be used.

In July 2023, Korea Electric Power Corporation (KEPCO) introduced a landmark reform to curb speculation and improve grid allocation efficiency. The new rules prohibit non-end-user entities from reserving capacity without verifiable development intent, requiring applicants to provide proof of end-use, project timelines, and shovel-readiness. By aligning grid access with credible demand, KEPCO has prioritized ready-to-build projects, reduced congestion in the connection queue, and ensured that infrastructure investment is directed toward genuine needs. The policy offers a governance model for other fast-growing digital economies grappling with speculative behavior in grid interconnections.

Best Practice 6: Malaysia

MALAYSIA – ALIGNING GRID PLANNING WITH DIGITAL GROWTH

Malaysia is aligning its digital and energy strategies, anticipating that data centers could consume nearly 20% of future electricity demand. To prepare, the government is investing in grid capacity, expediting permits, and capitalizing on spillover demand from Singapore to attract hyperscale projects. Flagship initiatives such as the Green Lane Pathway and Kulai Fast Lane Program sharply cut approval times for construction and grid connections, strengthening Malaysia's position as a rising regional data center hub.

SECTION 4:

UNLOCKING DATA CENTER EFFICIENCY

A. Energy Efficiency Under Pressure

Energy efficiency in digital infrastructure has improved, but not fast enough to keep pace with the explosive growth of AI workloads and hyperscale deployments. In many data centers, servers and other IT hardware in data centers still operate at just 10–15% of their processing capacity. Even 100% renewable-powered data center facilities remain inefficient under such conditions, driving up unnecessary energy, space, and cooling demands. Yet little data is available on how much of the computing capacity is actually being used (IT utilization rate).

Figure 14: Power consumption proportionality in an average data center

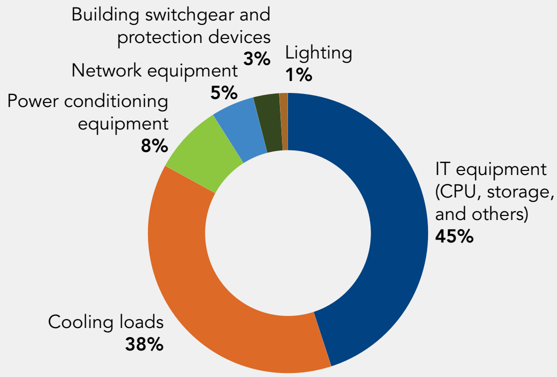

Source: APEC "Promoting Energy Efficient and Resilient Data Centres in the APEC Region" Report.

In 2023, APEC published Promoting Energy Efficient and Resilient Data Centres in the APEC Region report, aiming to analyze the standards, guidelines, energy data, new and renewable energies, and emerging innovative technologies to progress towards energy-efficient, resilient, and green data centers, serves as a reference for APEC member economies to boost energy efficiency and resilience in new data center installations.

Cooling accounts for roughly 40% of the electricity used by data centers. As GPUs get dramatically more power-hungry, traditional air cooling simply can't keep up. Liquid cooling is quickly becoming the new standard, especially in AI-heavy data centers. Most new high-density builds now use a hybrid setup, about 70% liquid cooling and 30% air, to manage rising thermal loads. However, retrofitting legacy data center sites to embrace liquid cooling remains a bottleneck. The main cooling methods — such as rear-door heat exchangers and direct-to-chip systems — are more energy efficient and are becoming industry standards, but upgrades in older data centers are often limited by a lack of space, power, or outdated layouts. These challenges make it harder to adopt newer systems quickly and continue to delay the broader transition toward low-carbon thermal management, especially across aging digital infrastructure in mature markets.

"Immersion cooling enables ultra-efficient cooling but faces adoption hurdles due to hardware weight and maintenance challenges"

Traditional and even advanced liquid cooling systems are approaching their performance limits, making way for immersion cooling to gain momentum as a next-generation solution. This technique involves fully submerging servers in a thermally conductive fluid, allowing heat to be removed far more efficiently than with conventional air-based systems. The dielectric liquids used in immersion cooling can have a thermal conductivity over 1,400 times greater than air, allowing heat to be transferred away from critical components far more efficiently. This makes it a compelling option for operators seeking to manage thermal loads while meeting sustainability and performance targets. However, fewer than one out of 10 data centers worldwide currently use the technology. Two key barriers stand in the way. First, retrofitting is often impractical: a single immersion tank filled with cooling fluid can weigh over 4 metric tons, surpassing the structural limits of standard commercial floors, especially in older or multi-tenant facilities. Second, operational complexity remains a concern. Immersion systems require specialized maintenance, and challenges such as fluid degradation, potential leaks, and difficult-to-access hardware can increase both risk and cost compared to conventional air cooling.

Best Practice 7: The U.S., Korea, Singapore

THE U.S.: GOOGLE'S OREGON DATA CENTER – INTEGRATED COOLING AND WATER REUSE

At its Oregon data center, Google employs liquid cooling to improve server efficiency and reduce energy consumption compared to traditional air-based systems. The facility also integrates an Aquifer Storage and Recovery (ASR) system, allowing it to store and reuse water sustainably. This approach not only supports the data center's operational resilience but also contributes to local water resource management, demonstrating how advanced cooling technologies can align digital infrastructure with environmental stewardship.

KOREA: SK TELECOM – IMMERSION COOLING PILOT IN INCHEON, KOREA

In a four-month pilot from June to October 2023, SK Telecom tested immersion cooling at its AI-focused data center in Incheon. The pilot achieved a 93% reduction in cooling power and more than 10% savings in server power, resulting in an overall energy reduction of 37%. Based on these results, full-scale deployment of immersion cooling began in 2024, positioning SK Telecom as a regional leader in high-efficiency AI infrastructure.

SINGAPORE: COOLING EFFICIENCY INNOVATIONS IN TROPICAL CLIMATES

Singapore's Tropical Data Centre Standard sets a benchmark for efficient operations in hot, humid environments. Developed over five years through government–industry–academic collaboration, it allows facilities to raise ambient temperatures from ~22 C to 26 C or higher, slashing cooling energy use by up to 50%. Integrated into the national Green Mark certification, the standard encourages zero freshwater cooling through innovations such as seawater-based systems in coastal or floating data centers. This climate-adaptive approach shifts efficiency limits from hardware tolerance to human comfort, positioning Singapore as a global model for tropical performance.

Building on this environment, ST Telemedia Global Data Centres (STT GDC) is scaling immersion cooling to support AI-ready infrastructure across Southeast Asia, beginning in Singapore and Thailand. After successful 2022 trials, STT GDC integrated both immersion and direct-to-chip liquid cooling into high-density GPU racks, removing heat directly at the source, cutting energy use, and reducing dependence on air-based systems to sustain intensive AI workloads.

EFFICIENCY GAINS ARE NOT WATER NEUTRAL.

> "By 2030, global data centers are projected to consume ~1,200 billion liters of water annually—enough to meet the water needs of nearly 15 million people, or the entire population of Ontario (Canada)"

Water is emerging as a critical and often overlooked resource constraint. Data center water use extends well beyond cooling: ~66% stems from electricity generation, ~25% from on-site cooling, and the rest from chip manufacturing. Data centers consume ~100 million liters per hour for cooling alone — up to 2 liters per kWh — placing significant strain on municipal water supplies and local ecosystems, particularly in drought-prone regions.

- **Water and Power Are Interlinked:** Roughly half of AI servers are powered by plants located in water-stressed areas. Many of these are thermal power plants — coal, gas, or nuclear — that require large volumes of water to operate. Even low-carbon sources such as hydropower and nuclear are water-intensive, meaning co-locating data centers near these facilities can amplify local water stress rather than alleviate it.

- **Not Just Water, Clean Water:** Data centers don't just need water; they need clean, treated water. Cooling systems rely on high-quality water to prevent corrosion and bacterial growth, raising operational risks and costs in arid climates. Meanwhile, semiconductor manufacturing, the backbone of AI, consumes vast quantities of ultra-pure water, accounting for ~90% of chip-related water use. With chip demand rising, water use per chip is climbing, and chip-related water demand for data centers is projected to increase 50% by 2030.

> "Efforts to cut electricity use will shift pressure onto local water systems."

There are efforts to cut electricity use, like raising operating temperatures, improving airflow, and investing in liquid cooling/immersion cooling. While these practices can reduce energy waste, they could also shift the burden to water systems. By 2030, half of global data center water consumption is projected to occur in APEC economies like Vietnam, Singapore, and Malaysia, where hotter climates drive higher cooling needs. Water usage effectiveness (WUE) can reach ~1.65 liters/kWh in these regions, more than three times the global average.

Water-saving technology exists, but not without trade-offs. For example, dry cooling reduces water demand but increases electricity use. Reclaimed or seawater systems offer alternatives, but are costly and still complex to deploy in many APEC economies. Free-air cooling works well in temperate zones (e.g., Canada), but is rarely viable in tropical APEC regions such as Southeast Asia.

Table 3: Global and Asia Pacific WUE

Region	Average WUE (L/kWh)
Global Average	0.5 L/kWh
Asia Pacific	1.65 L/kWh

Source: IEA's Energy and AI Report, 2025

As a result, sustainability pressures are rising, especially in Latin America, where hyperscale operators face growing scrutiny over water use. In Asia, regions such as Singapore, Malaysia, and northern China are seeing public concern over freshwater competition, worsened by tropical climates and reliance on water-intensive cooling.

Best Practice 8: Singapore, Malaysia, China

SINGAPORE: INTEGRATING WATER INNOVATION IN DATA CENTER STRATEGY

Singapore is piloting seawater cooling to reduce dependence on freshwater and land-based infrastructure, with Keppel's proposed Floating Data Centre Park at Loyang as a flagship project. By using seawater for heat rejection, the design conserves treated water, frees up scarce land, and can integrate off-site power generation for higher efficiency. The approach addresses the high cost and carbon footprint of desalination. Other current policy measures include using reclaimed NEWater and embedding WUE targets in standards like the Tropical DC Standard and BCA Platinum Green Mark.

Complementing these initiatives, AWS has invested in a 17.6 MW solar project with Sembcorp that feeds clean energy into the local grid and integrates a rainwater harvesting and treatment system capable of collecting up to 170,000 m^3 annually, enough to fill 68 Olympic-sized swimming pools. This reduces reliance on municipal freshwater supplies, showing how renewable energy infrastructure can be designed to deliver dual sustainability gains in both energy and water efficiency.

MALAYSIA: BALANCING DIGITAL GROWTH AND WATER SECURITY

Malaysia faces acute water stress, ranking highest among its regional peers — a concern amplified by the rapid expansion of the data center sector. Johor's current 16.9% water reserve margin masks growing tensions between industrial and residential demand, with 101 data center applications filed as of 2024 projecting over 808 million liters in daily water use. To pre-empt conflict between digital infrastructure growth and water security, the National Water Services Commission (SPAN) will, by mid-2025, roll out mandatory guidelines for Peninsular Malaysia requiring all new data centers to source 100% of their operational water needs from alternative supplies — including recycled, reclaimed, or harvested rainwater — within three years of commissioning. This regulatory shift aligns with broader efforts to integrate water sustainability targets into Malaysia's industrial development strategy, positioning the country to meet both digital economy and environmental objectives.

CHINA: UNDERSEA DATA CENTER HARNESSES SEAWATER COOLING AND OFFSHORE WIND

HiCloud is investing $223 million in an undersea data center off Shanghai that uses seawater cooling, pumping cold seawater through radiators behind server racks to absorb and carry away heat. The design is expected to consume at least 30% less electricity than comparable land-based facilities. The center will draw 97% of its power from a nearby offshore wind farm and, in its first phase, house 198 server racks—enough to support from 396 to 792 AI-capable servers.

POWER USAGE EFFECTIVENESS (PUE) IS NO LONGER SUFFICIENT.

By definition, PUE measures the ratio of total data center power to the energy used by IT equipment, effectively capturing energy lost to cooling and overhead systems. The global average PUE today stands at 1.55, meaning roughly 38% of electricity does not go toward computing work.

Among data center operators, PUE remains a widely used metric. Yet, PUE offers no insight into the carbon intensity of power sources, the volume of water consumed for cooling. A facility may boast a low PUE yet rely on fossil fuels or waste millions of liters of water. PUE also fails to reflect computing efficiency, such as energy used per unit of AI training or data storage. PUE mandates also vary widely across APEC economies, with most targeting values between 1.4 and 1.5 by 2025, while Singapore sets the most ambitious goal — 1.3 for new data centers.

> **"Besides PUE, water usage effectiveness (WUE), and carbon usage effectiveness (CUE) are useful, but they are often misapplied, overly simplified, or insufficient to fully capture efficiency performance."**

Water usage effectiveness (WUE): WUE measures how much water a data center uses to cool its IT equipment per kilowatt-hour of energy consumed. Despite the central role of water in energy-efficient design, water use reporting is rare and inconsistent. WUE metrics are less standardized than PUE and highly variable based on climate, cooling design, and water source, making it hard to compare across sites or countries. Just [39%](#) of data center operators reported water use in 2022. Big techs (e.g., Google, Amazon, Microsoft) are reporting water use and pledging to become water-positive. Data center firms in China such as GDS, Chindata, and Baidu Cloud are reporting WUE. Yet most small and medium-sized data centers in APEC lack water reporting and often exclude the indirect water footprint of electricity generation in their report.

Figure 15: Data Center Standards and Certifications

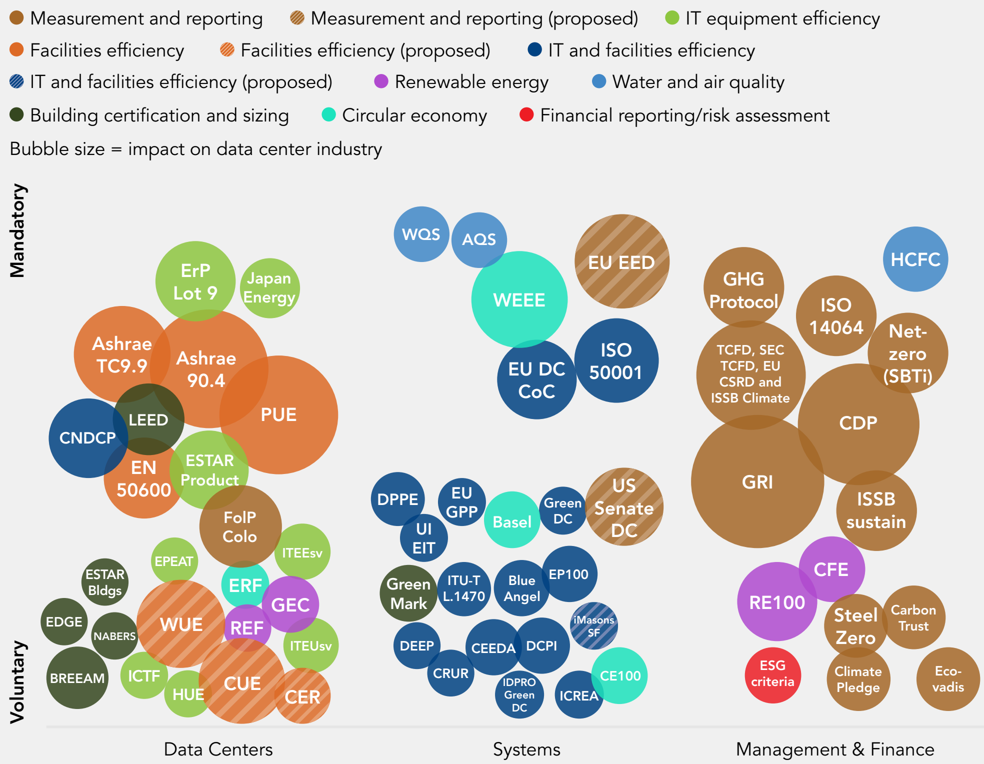

Source: Uptime Institute [Webinar - Digital Infrastructure Sustainability: Navigating Regulations and Standards.](#)

"Lacking a common sustainability framework, APEC data centers face a patchwork of over 150 overlapping standards. Without a common framework, it's difficult to track, compare, or improve sustainable performance across borders, holding back progress toward a more unified and efficient digital infrastructure in the region."

Carbon usage effectiveness (CUE): CUE measures the carbon emissions per unit of energy used in a data center. While useful, CUE is also difficult to apply in APEC economies with limited energy market flexibility. In countries like Vietnam, where electricity procurement is regulated, operators cannot choose renewable sources freely, meaning a poor CUE score may reflect policy constraints rather than operator performance.

Today, many energy and sustainability certifications apply to data centers with no shared taxonomy or baseline. Operators face a confusing mix of voluntary and mandatory schemes and often defer to cloud service clients to decide which benchmarks to follow.

Best Practice 8: Singapore, Australia, Canada, the U.S.

SINGAPORE: PUE AND WUE MANDATES

Singapore mandates a PUE threshold of 1.3 or lower for new data centers in designated zones, alongside mandatory WUE reporting to address water constraints in a city-state heavily reliant on energy-intensive desalination. The 1.3 PUE target — embedded in the BCA Platinum Green Mark certification — was developed through joint consultations between policymakers and industry, guided by a technical committee that included leading tech firms. The result is an ambitious yet technically grounded benchmark, exemplifying a collaborative policy model where government ambition and industry innovation move in lockstep.

AUSTRALIA: MANDATING HIGH-EFFICIENCY STANDARDS FOR GOVERNMENT DATA CENTERS

Australia is requiring all federal data centers — government-owned, operated, or contracted — to meet strict energy performance standards by July 2026. Under the ICT Sustainability Plan, facilities must achieve a 5-star NABERS Energy rating, use accredited GreenPower, comply with the NGER Act emissions thresholds, and target a PUE below 1.4.

** NABERS stands for National Australian Built Environment Rating System. It is a performance-based rating system that measures the environmental impact of buildings in Australia. A 5-star NABERS rating indicates excellent energy performance and is increasingly being used as a standard requirement for data centers serving Australian government agencies or aiming to meet sustainability targets.*

THE U.S AND CANADA: VOLUNTARY EFFICIENCY STANDARDS FOR MARKET SIGNALING

ENERGY STAR for data centers, adopted in both the U.S. and Canada, provides a voluntary yet influential benchmark for procuring energy-efficient equipment. Covering key categories such as servers, network hardware, storage, and uninterruptible power supplies (UPS), the program signals to the market which technologies meet best-in-class performance standards. While not mandatory, it harmonizes efficiency expectations across both countries, guiding procurement policies and incentivizing operators to adopt low-energy technologies as part of broader corporate ESG strategies.

B. Hardware and Software Efficiency

EFFICIENCY STARTS WITH HARDWARE

As high-performance AI chips significantly increase rack power density and drive up both energy and water demands for cooling, there's growing consensus that the next major leap in data center energy efficiency must come from rethinking the hardware itself. Yet despite this progress, R&D adoption remains slow as the market prioritizes speed over sustainability.

"By 2030, nearly $3.1 trillion, or 60% of the world's projected $5.2 trillion AI infrastructure investment, will be poured into next-generation computing hardware."

As AI workloads accelerate, custom chips and silicon-level innovations are emerging as a key frontier for improving data center sustainability. Unlike software optimization or facility-level retrofits, advances in hardware architecture offer structural gains, reducing energy use per computation at the source. Leading examples include Google's Tensor Processing Units (TPUs) and Apple's Neural Processing Units (NPUs), which deliver significantly higher performance per watt compared to traditional CPUs. Likewise, AMD's chiplet-based EPYC processors reduce energy waste by tailoring compute to workload types, while High Bandwidth Memory (HBM), as seen in NVIDIA's H100, slashes power lost in data movement by physically bringing memory closer to the processor.

Looking further ahead, photonic chips, which process data using light rather than electricity, may offer radical efficiency improvements, though they remain in early stages of development. More immediate solutions, such as low-power ARM-based processors and advanced cooling techniques, offer practical near-term efficiency gains.

Without coordinated public investment, stronger policy incentives, or procurement standards that value efficiency, many hardware innovations risk being underutilized. For policymakers, the opportunity lies in sending clear demand signals—through funding, regulation, or standards—to accelerate the adoption of hardware that can meaningfully reduce energy and water use across the data center sector.

Tech companies are already moving in this direction. Google's custom ASICs outperform conventional CPUs by up to 80 times in energy efficiency, while hyperscalers like Meta, Microsoft, and AWS are investing in proprietary chips tailored to the energy profiles of AI training and inference. In APEC economies such as Korea, national R&D programs now prioritize custom hardware, including NPUs and HBM integration, as a competitive pathway to high-efficiency digital infrastructure. However, the market continues to reward raw speed over energy savings, as seen in the premium placed on GPUs, where performance for AI load dominates investment decisions despite high power consumption. Similarly, cloud customers often select instance types based on raw throughput rather than total energy costs, reinforcing demand for ever-faster chips. This dynamic sidelines efficiency in procurement, yet it also creates space for policy and corporate initiatives, such as Korea's NPU programs, to reframe efficiency as part of competitiveness.

Best Practice 9: Korea

KOREA – ADVANCING EFFICIENCY THROUGH AI HARDWARE INNOVATION

Rather than regulating data center efficiency through PUE targets, Korea is investing in next-generation low-power AI hardware. National R&D programs prioritize Neural Processing Units (NPUs) and High Bandwidth Memory (HBM) integrated GPUs, technologies optimized for inference and edge workloads. This tech-led approach reduces energy consumption per computation at the chip level, positioning hardware innovation as a strategic pathway to national data center efficiency.

In August 2025, Korea launched a distinctive full-stack "sovereign AI" drive, mobilizing its biggest tech firms and startups to create a national foundational model built largely on homegrown technologies, from semiconductors and high-bandwidth memory to cloud platforms and AI software. Supported by the Ministry of Science and ICT, five consortia led by SK Telecom, LG, and Naver are advancing open-source models that capitalize on Korea's strengths in HBM production (SK Hynix, Samsung), AI chip design (Rebellions), and rapidly expanding data center infrastructure.

SOFTWARE EFFICIENCY IS UNDERUSED

"Improving software efficiency is emerging as a faster, lower-cost path to cutting data center energy use. Yet despite the clear benefits, software efficiency remains underused."

One of the biggest blind spots of digital decarbonization is data storage, which accounts for about 20% of a data center's energy use. A large share of that goes to "dark data," or information that's rarely or never accessed, such as outdated backups or unused logs. Dark data is estimated to make up over 50% of global storage and burns energy around the clock, quietly adding millions of tonnes of CO_2 to the atmosphere each year. Companies generate 1.3 trillion gigabytes of dark data daily, and storing it on non-renewable energy for a year produces CO_2 emissions equivalent to 3 million transatlantic flights between London and New York.

Regularly auditing storage, clearing out obsolete files, and building software that avoids creating unnecessary data in the first place can significantly shrink energy demand. Efficient storage tools, such as Solid-State Drives (SSDs) and cold storage, can also dramatically cut energy use when paired with smarter data practices.

However, unlike energy-efficient hardware, efficient software improvements attract less attention or drive commercial value. Few organizations have systems in place to measure or incentivize energy-efficient software. Managing "dark data" also requires time, coordination, and cross-functional effort, often without immediate returns, making it a lower priority in many data center operators.

Best Practice 10: China

CHINA – ADVANCING EFFICIENCY THROUGH AI HARDWARE INNOVATION

According to the State of Dark Data report, which draws on a global survey of 1,300+ business and IT decision makers conducted by TRUE Global Intelligence, China took the most aggressive stance toward the rising value of data, the reduction of dark data, and the role of AI in data management. A majority of Chinese respondents reported that less than half of their organization's data is "dark," signaling greater attention to managing unused data as a strategic priority. Combined with the highest global confidence in understanding AI and 84% of respondents identifying data skills as critical to future jobs (compared with 61% globally), China illustrates how prioritizing software efficiency — through data literacy, AI training, and dark data reduction — can strengthen competitiveness while curbing unnecessary energy use.

C. Untapped Opportunity: Waste Heat Reuse

Waste heat remains a missed opportunity due to limited upstream integration, while fragmented operational metrics and inconsistent standards across APEC are hampering efforts to improve energy efficiency.

Despite the immense heat AI workloads generate, most data centers in APEC economies still vent the heat into the air or ground, wasting a low-carbon energy source. Waste heat reuse remains largely untapped in APEC, even as urban centers face rising heating demand.

Figure 16: Equinix's Data Center Heat Reuse Case Study

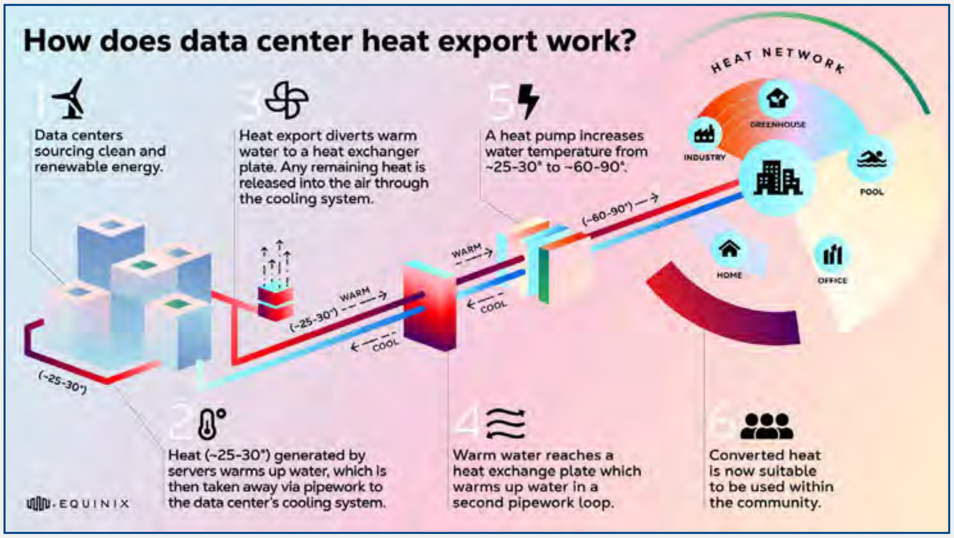

Source: Equinix "What Is Data Center Heat Export and How Does it Work?"

Data centers waste heat twice: first at the power plant, then again when AI workloads convert electricity into thermal exhaust. With the right infrastructure and planning, that byproduct heat could power district heating, desalination, or even hydrogen production. Heat reuse is especially viable in cold APEC cities and dense urban hubs where existing district energy networks — and nearby demand from buildings, greenhouses, or industry — make integration practical. **For economies such as Korea and Japan, and regions such as northern China and parts of Canada, data center heat recovery could fast-track the decarbonization of urban heating systems.**

- **Technologies are already commercially viable.** Heat recovery chillers provide both cooling and usable heat at the same time, making them highly efficient. They can be easily integrated into existing chilled water systems, allowing facilities to recover and repurpose waste heat rather than discard it. When paired with thermal storage, they enable load shifting and a more consistent heat supply. When connected to district energy networks, recovered heat can be distributed to nearby buildings. This integrated approach creates a reliable demand for recovered heat, maximizing overall energy efficiency and heat utilization across urban systems.
- **Cost isn't the problem either:** IEA estimates that heat reuse infrastructure costs just $200,000 – 275,000 per MW of heat — less than one-third the cost of building a new gas-fired heating plant (over $800,000 per MW).

Despite these strengths, adoption still lags. Key barriers include initial infrastructure costs, complex coordination across construction timelines, and regulatory and contractual challenges in deploying district systems. With integrated planning and strong public-private collaboration, waste heat reuse works in practice.

Europe is leading the policy shift. For example, Stockholm Data Parks has successfully connected over 20 data centers to the city's district heating network, supplying about 1.5% of Stockholm's total heating demand. Denmark is integrating data center waste heat into district heating networks. Facebook's Odense facility supplies heat to local homes, while atNorth partners with Wa3rm to repurpose excess heat for uses like greenhouse farming, showcasing circular energy use in digital infrastructure. Germany and the Netherlands now mandate heat reuse in new data centers. The EU Energy Efficiency Directive requires large facilities to recover waste heat or prove it's not feasible. No APEC economy has yet followed suit, despite rising demand for clean urban heat and growing electricity constraints.

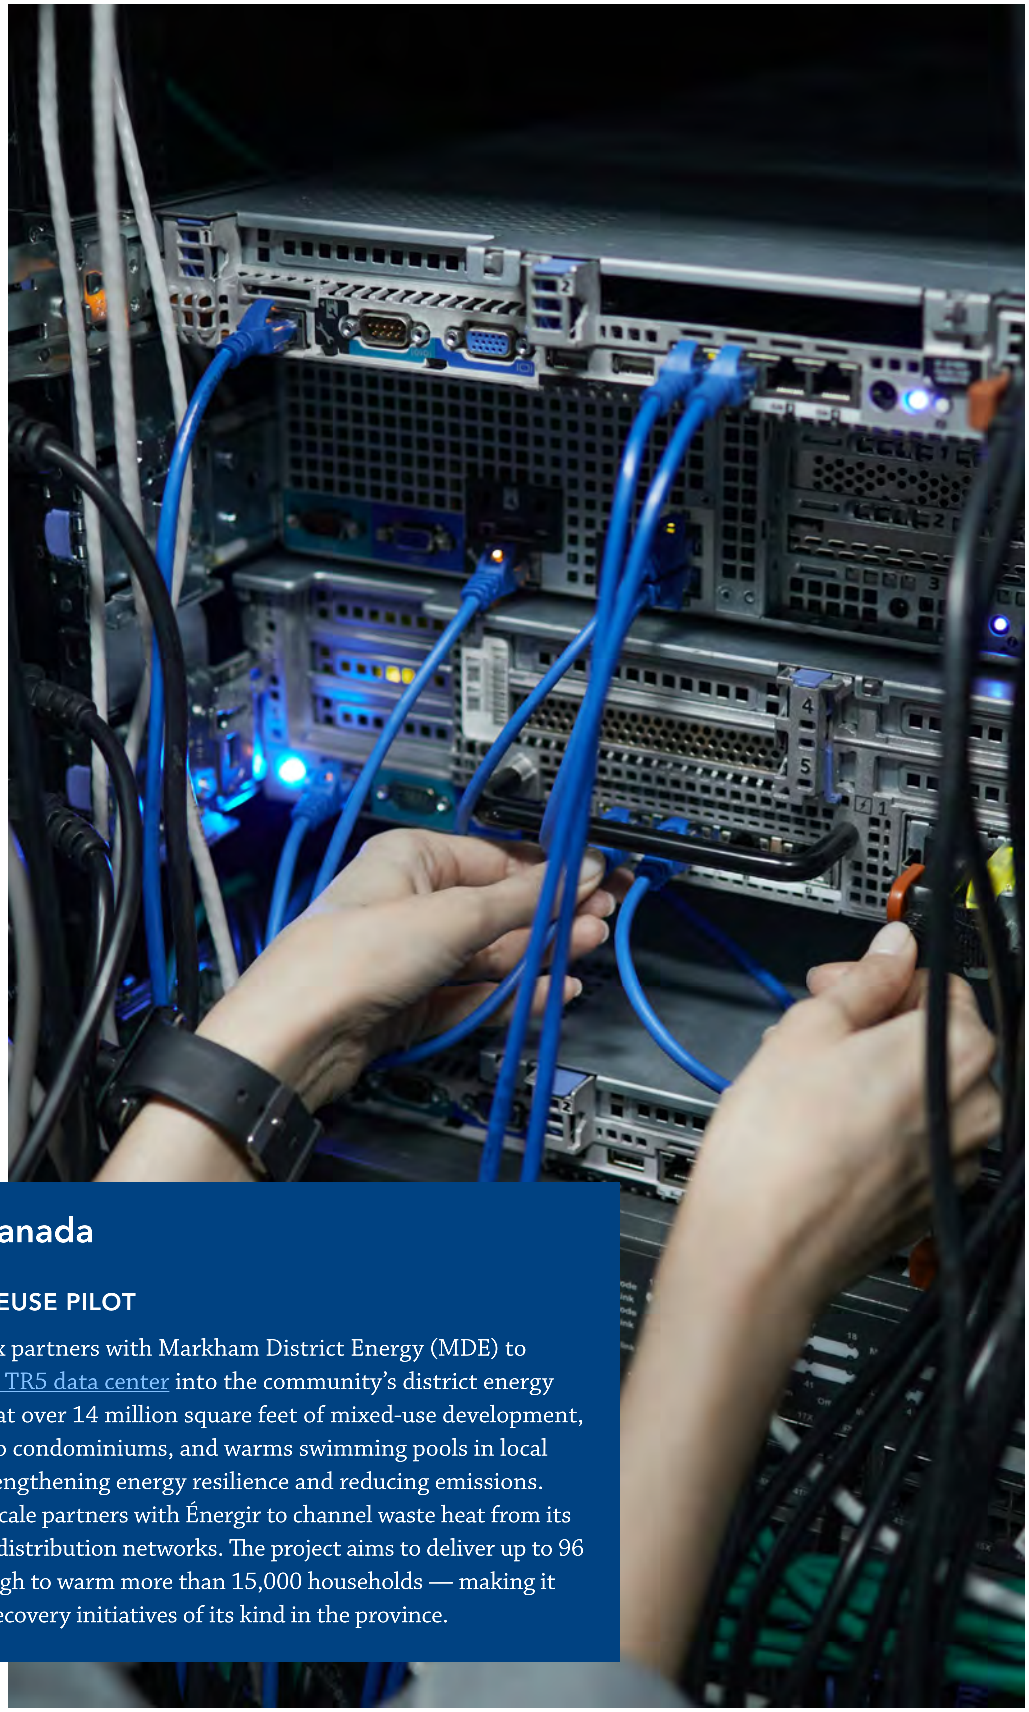

Best Practice 11: Canada

CANADA: WASTE HEAT REUSE PILOT

In Markham, Ontario, Equinix partners with Markham District Energy (MDE) to repurpose waste heat from its TR5 data center into the community's district energy network. The project helps heat over 14 million square feet of mixed-use development, supplies domestic hot water to condominiums, and warms swimming pools in local community centers, while strengthening energy resilience and reducing emissions. Similarly, in Lévis, Quebec, QScale partners with Énergir to channel waste heat from its Q01 campus into local energy distribution networks. The project aims to deliver up to 96 MW of recovered heat — enough to warm more than 15,000 households — making it one of the largest waste heat recovery initiatives of its kind in the province.

SECTION 4:

ALIGNING POLICY AND REGIONAL COOPERATION

A. Uneven Policy Instruments and Implementation

While several APEC economies are introducing targeted policies to promote cleaner, more efficient data centers, inconsistent regulations and uneven implementation across regions are creating a fragmented and unequal landscape for sustainable digital infrastructure.

Since early 2020, many economies have launched national strategies for AI and the digital economy, but timelines and scope vary widely. Canada was the earliest, introducing its first AI strategy in 2017 and introducing a new Sovereign AI Compute Strategy in 2024. On the data center front, several APEC economies are now advancing national strategies for sustainable digital infrastructure, with varying emphasis on green data center practices. Policy priorities differ — ranging from clean energy integration to hardware innovation incentives — and implementation stages remain uneven, reflecting divergent national contexts and readiness levels.

Policy tools to align digital infrastructure growth with clean energy availability vary widely across economies but generally fall into four categories: strategic siting, permitting incentives, innovation sandboxes, and sustainability reporting. Strategic siting — seen in China, Malaysia, and Indonesia — integrates data center development with renewable power generation and water resource planning. Permitting incentives in jurisdictions such as Johor (Malaysia), Alberta (Canada), Batam (Indonesia), and select U.S. states offer fast-track approvals or tax benefits for green-certified facilities, though environmental criteria remain inconsistent. Innovation sandboxes in Australia, Korea, and the U.S. provide a testing ground for next-generation solutions such as immersion cooling and advanced hardware efficiency. Meanwhile, mandatory sustainability reporting, though unevenly enforced, increasingly requires disclosure of energy use, emissions, and efficiency metrics, with some jurisdictions expanding to include Scope 1–3 GHG emissions, water consumption, and operational benchmarks such as PUE and WUE. While Singapore applies comprehensive mandates and Australia limits them to government-operated data centers, many economies still depend on voluntary reporting frameworks, underscoring the patchwork nature of current policy adoption.

Figure 17: Policy Landscapes across Focus Economies

● National ○ Subnational

Economy	National AI/ Data Center strategy	Government financial support			Reporting requirements		Performance mandates	
		R&D	Data centers	Chips	Emissions	Electricity consumption	PUE	WUE
Australia	●	●	●		●	●	●	
Canada	●	●	●	●	○	○		
Chile	●	●	●					
China	●	●	●	●		●	●	●
Indonesia	●	●	●	●				
Japan	●	●	●	●	●	●	●	
Korea	●	●	●	●				
Malaysia	●	●	●			●	●	
Singapore	●	●	●	●	●	●	●	●
U.S.	●	●	●	●	○	○	○	

Source: Compiled by author based on International Energy Agency Energy and AI Report and other government websites. · Note 1: Economies with national data center strategies: China's East-Data-West-Compute Strategy (2022), Chile's National Data Centers Plan (2025), Singapore's Green Data Centre Roadmap (2024), The U.S.'s AI Action Plan (2025). · Note 2: Economies with national AI strategies: Malaysia's National AI Roadmap 2021-2025 (2021), Canada's Sovereign AI Compute Strategy (2024) & Pan-Canadian AI Strategy (2017), Australia's National AI Capability Plan (upcoming, 2025), Korea's National AI Strategy Policy Directions (2024) & AI Framework Act (2025), Indonesia's National AI Strategy (2020), Japan National AI Strategy (2022) & AI Promotion Act (2025)

The current regulatory patchwork across APEC economies creates uncertainty for investors and hampers the ability to scale greener digital infrastructure across borders. Without cross-border standardization, hyperscalers and data center operators face fragmented compliance requirements, while inconsistent enforcement undermines accountability and weakens regional sustainability goals. Uneven mandates — such as selective certification rules or varying efficiency thresholds — further contribute to disparities in environmental performance, widening the gap between advanced and emerging digital markets.

B. Regional Cooperation and Public-Private Partnership

Coordinated power planning between governments and data center operators is becoming a cornerstone of sustainable digital infrastructure. Clean energy corridors, demand-side partnerships, long-term PPAs, and policy co-design are emerging tools to strengthen grid resilience, unlock private investment, and support low-carbon digital growth. Yet regulatory fragmentation, mismatched market rules, and uneven grid maturity continue to constrain the full potential of cross-border infrastructure and public–private alignment.

CROSS-BORDER CLEAN ENERGY CORRIDORS

Cross-border clean energy trade refers to the exchange of electricity, typically from renewable sources, between economies, enabling data centers to access cleaner power from neighboring grids when local power generation is limited or carbon-intensive.

Regionally coordinated clean energy corridors are advancing in Southeast Asia. For example, Indonesia's Batam Island is being developed as a data center zone in partnership with Singapore: Indonesian authorities, Singaporean investors, and utilities are negotiating cross-border electricity trade so that Batam data centers can export services while tapping Singapore's investments in renewables. In North America, U.S. and Canadian grid operators coordinate to transmit Quebec's surplus hydroelectric power to new data centers in the U.S. Northeast, benefiting from Canada–U.S. power trade agreements. Such regional collaboration ensures data centers have access to reliable, affordable, and clean energy beyond local constraints, leveraging wider geographic resources.

However, collaborative planning across the energy and digital sectors remains limited and uneven across APEC. The progress is also held back by region-specific issues such as subsea cable deployment (e.g., Singapore-Indonesia-Malaysia), uneven permitting timelines, grid interconnection bottlenecks, and fragmented governance across economies.

GRID FLEXIBILITY

Demand response programs allow data centers to reduce or shift electricity use during short, predictable periods of peak grid stress — such as evening peaks during heatwaves — in exchange for compensation from utilities. Data centers possess multiple flexibility levers, including onsite batteries

Best Practice 12: Singapore, Malaysia, Indonesia

SINGAPORE–MALAYSIA: JOHOR–SINGAPORE SPECIAL ECONOMIC ZONE

Singapore and Malaysia are advancing cross-border integration through the newly launched Johor–Singapore Special Economic Zone (JSSEZ) in 2025, designed to strengthen trade, investment, and infrastructure linkages. Envisioned as a 3,500 km^2 energy and innovation corridor, the JSSEZ connects Singapore's fast-growing digital economy with Johor's abundant renewable resources. The zone incorporates coordinated green power development, enhanced transport connectivity, and targeted investor incentives, creating a unified platform for manufacturing, digital infrastructure, and clean energy supply chains.

SINGAPORE – INDONESIA: INDONESIA–SINGAPORE GREEN ELECTRICITY EXPORT PARTNERSHIP

Leveraging its proximity to Singapore, deep-water marine infrastructure, and special economic zone advantages, Batam is central to the Indonesia–Singapore Green Electricity Export Partnership signed in September 2024. The agreement targets up to 20 GWp of installed solar capacity across Indonesia, with the bulk of generation transmitted to Singapore via high-voltage subsea cables. A flagship example is Marubeni's Galang Island project with Tuas Power, combining 2.55 GWp of solar capacity with a 7 GWp battery system to deliver stable, dispatchable power. Singapore's Energy Market Authority (EMA) has committed to importing 6 GW of low-carbon electricity by 2035, with up to 4 GW allocated to Indonesia-linked projects.

for frequency balancing, backup generation (such as natural gas or fuel cells), and adaptive cooling strategies like ice or thermal storage. Advanced workload management further boosts flexibility, with companies like Google shifting compute across time or geography to regions with cleaner power.

Some APEC economies are catching on—increasing grid flexibility. Korea, for example, is integrating hyperscale data centers into its demand response framework, rewarding operators for load shifting and treating them as grid-stabilizing assets rather than passive consumers. Australia has seen pilots where data centers feed stored energy from on-site batteries back to the grid at critical times, effectively acting as mini power plants. In the U.S, the Electric Power Research Institute (EPRI) launched DCFLEX, a pilot to integrate data centers into real-time grid operations in 2024. The Department of Energy also released a flexibility taxonomy and incentive design blueprint. Industry leaders like Google, Amazon, and Microsoft are offering interruptible compute at discounted rates, swapping uptime guarantees for grid-responsiveness.

"Data centers hold significant untapped potential for grid flexibility but most APEC flexibility programs remain tailored to legacy industries."

However, demand-side flexibility programs in many APEC economies remain rooted in legacy designs. Many grid systems were designed around conventional flexible loads like HVAC or industrial motors. Few programs offer mechanisms tailored to data centers' unique constraints, such as latency sensitivity, uptime requirements, or load predictability. This reflects a missed opportunity for true public-private partnerships to accelerate decarbonization and grid resilience

POWER PURCHASE AGREEMENTS

PPAs are long-term contracts where data centers buy clean energy directly from generators, enabling new renewable builds. As of 2024, PPAs have enabled over 120 GW of global renewable capacity, with an additional 60 GW in development, much of it driven by hyperscale data needs (IEA). Over 30% of all operational corporate PPAs now come from data center operators, with Amazon alone accounting for one-third of the global market.

PPAs are not adopted evenly across APEC. U.S. data centers lead the pack, accounting for 75% of all operational PPA capacity, including 20 GW of solar PV and 12 GW of onshore wind. In 2024, over 90% of new PPAs were signed in the U.S.

"Power purchase agreements (PPAs) are a powerful tool for scaling renewables in data center operations, but their climate impact hinges on policy design."

Meanwhile, many APEC economies lack standardized regulatory frameworks that support PPAs, especially in regulated or state-owned power markets. Some APEC countries started to reform regulations to allow direct contracts between companies and renewable producers. In Malaysia, for example, the Corporate Green Power Programme (CGPP) and Corporate Renewable Energy Supply Scheme (CRESS) have created clearer pathways for corporate renewable procurement. Indonesia's PPA framework remains in development; however, the country's 2060 net-zero commitment is driving progress, illustrated by Amazon AWS's 2022 offsite PPA with PLN to source power from four solar projects in Bali and Java. In Australia, physical PPAs face barriers due to grid access constraints and the structure of the National Electricity Market (NEM), making virtual PPAs the primary tool for data center renewable procurement. Chile's liberalized market has enabled transactions like Google's 100 MW solar-plus-wind PPA with Engie, supplying its data center entirely with renewable power.

Lacking clear PPA standards, many APEC economies miss a critical chance to accelerate renewable energy and foster real public-private partnerships. Even where PPA frameworks do exist, they often reward box-checking rather than cutting real, local emissions. Most traditional PPAs don't always line up with when power is used: data centers still burn fossil energy when renewables fall short. To cover the gap, many hyperscalers committed to RE100 by buying Renewable Energy Certificates (RECs) or Energy Attribute Certificates (EACs) from far-off projects. The lack of a time-matching PPA framework and local sourcing results in a disconnect between reported renewable use and actual grid decarbonization. Smarter policy design is needed to align incentives with actual climate impact.

POLICY CO-DESIGN

Governments are increasingly engaging data center operators, cloud providers, communities, and researchers in shaping digital infrastructure and energy policies. However, hyperscalers still dominate the conversation, while smaller firms and civil society often lack a seat at the table. Without clear legal mandates or long-term roadmaps, co-created policies risk becoming voluntary wish lists rather than enforceable strategies.

Best Practice 13: Singapore, Canada

SINGAPORE: INDUSTRY CONSULTATION ON PERFORMANCE-BASED FRAMEWORK

Singapore has established one of the APEC region's most rigorous performance-based frameworks for data centers, anchored in its Green Data Centre Roadmap and aligned with national decarbonization goals. Developed through extensive industry consultation, the framework demonstrates how high-ambition policy can be both enforceable and industry-backed.

CANADA: A COLLABORATIVE APPROACH TO CO-DESIGNING POLICY

Canada's 2024 Sovereign AI Compute Strategy, led by Innovation, Science and Economic Development Canada (ISED), highlights a collaborative, multi-stakeholder model for digital infrastructure planning. Developed through extensive consultations with universities, technology companies, startups, and public-sector experts, the strategy ensures federal policy is closely aligned with real-world AI compute needs and industry innovation. Importantly, Canada is also embedding Indigenous equity and co-ownership into energy and connectivity projects, advancing a more inclusive and sustainable model of national digital infrastructure.

RECOMMENDATIONS

ABAC recommends APEC Ministers and Leaders work with the private sector to accelerate the development of sustainable AI infrastructure, including data centers, computing infrastructure, water supply, and power grids to support innovation and climate goals. To do so, APEC policymakers should work closely with ABAC, industry leaders, and relevant regional and global bodies to:

1. **Measure, benchmark, and assess the economic value and environmental impact of the digital economy. Use benchmark tools such as The Burn-to-Earn Index.**

   APEC economies should establish a robust evidence base to inform sustainable digital infrastructure planning. This includes measuring both the economic contribution and environmental externalities — such as energy, water, and carbon footprints — of data centers, AI compute clusters, and cloud infrastructure. To enable fair comparison and informed policymaking, APEC should support the adoption of standardized benchmarking tools, such as The *Burn-to-Earn Index*, developed by Tufts University and SGTech. Public-private collaboration with industries, academic researchers, and local communities is essential for developing transparent, regionally relevant indicators.

2. **Collaborate with industry and energy providers to forecast data center demand against energy capacity.**

   To prevent future energy shortfalls, APEC governments should work closely with hyperscalers, colocation providers, and utilities to proactively align infrastructure buildout with grid capacity, clean power availability, and transmission planning. Forward-looking digital demand forecasting should be integrated into energy strategies, especially in jurisdictions facing grid congestion or fossil fuel dependency.

3. **Monitor, support and deploy sustainability and efficiency advancements in digital infrastructure, including advanced cooling, heat reuse, and integrated zone design, and pursue next-generation chips and semiconductors to drive operational efficiency.**

   For policymakers, the opportunity lies in sending clear demand signals — through funding, regulation, or standards — to accelerate the adoption of technologies that can meaningfully reduce energy and water use across the data center sector. APEC governments can accelerate progress by co-investing in integrated industrial zones and regional pilot projects that demonstrate scalable, climate-appropriate solutions.

4. **Monitor and adopt global sustainability standards. Set ambitious, measurable targets for green data centers in APEC.**

   To promote transparency and market confidence, APEC governments should support the adoption of regionally relevant, measurable targets for data center sustainability, including metrics for energy and water efficiency, carbon intensity, and clean energy use. To reduce regulatory fragmentation and strengthen cross-border collaboration, APEC should advance a shared sustainability framework that fosters transparency, accelerates market convergence, and unlocks access to green financing.

The private sector plays a pivotal role in advancing energy-efficient, climate-aligned digital infrastructure across the APEC region. To support a resilient and sustainable digital economy, the private sector should:

1. Catalyze efficiency in AI systems by investing in low-power hardware and deploying software more efficiently.
2. Participate in utility-led grid flexibility programs by shifting workloads that are not time-sensitive (e.g., AI training, background processing) during peak hours to off-peak periods or renewable-rich windows, while maintaining 24/7 uptime for mission-critical services, supporting grid stability and decarbonization.
3. Embrace transparency by reporting energy consumption and emissions linked to AI training and digital operations, reinforcing accountability and enabling informed decisions across the value chain.
4. Adopt recognized data center-related certifications and standards and track performance-based metrics to improve operational efficiency and align with climate and sustainability targets.
5. Integrate heat recovery systems into facility design and partner with local utilities or communities to convert server waste heat into usable energy, transforming a byproduct into a shared low-carbon resource.

CONCLUSION

Data centers are rapidly becoming the backbone of the digital economy across APEC. Yet without deliberate intervention, their accelerating energy footprint risks undermining regional sustainability goals, deepening energy security challenges, and straining already fragile infrastructure systems.

A coordinated, forward-looking response is not only feasible — it is imperative. Proven solutions such as clean energy integration, high-efficiency cooling, and advanced hardware efficiency are already within reach. The challenge now is scaling them at speed, guided by aligned regulations, long-term planning, and cross-sector collaboration. It's imperative for economies to forecast their digital economy growth trajectory against their energy mixes and grids.

Looking ahead, economies must proactively prepare for the next wave of compute-intensive technologies — including quantum computing, autonomous AI agents, immersive digital ecosystems, and blockchain-enabled platforms. These technologies are not distant possibilities — they are rapidly approaching deployment at scale, with the potential to transform productivity, health care, education, logistics, and public services. However, their emergence will bring exponential increases in processing demand, along with new pressures on electricity systems, cooling infrastructure, water availability, and land use. Without anticipatory planning, these pressures could outpace the capacity of current infrastructure and exacerbate regional inequalities.

The race to scale compute is also a race to ensure the power systems behind it are ready. APEC economies must seize this moment to align digital growth with robust, modernized energy infrastructure. By acting early and aligning investment with enabling policies, the region can ensure its energy systems evolve in step with digital growth, meeting rising demand while maintaining reliability and resilience. Acting now will secure the energy foundations needed to power AI, cloud, and next-gen services, positioning APEC not just to keep pace, but to lead the global digital economy.

ANNEX

A. AI AND THE DIGITAL ECONOMY

AI represents the most volatile and least predictable source of future electricity demand. Unlike sectors such as electric vehicles or industrial electrification — where deployment timelines, usage patterns, and load curves are relatively well-understood — AI compute requirements are accelerating rapidly, yet remain difficult to forecast with precision. Projections for AI's share in total data centers' energy demand in the future (2030) could range from 50% to over 70% according to different forecasting agencies. This uncertainty reflects a broad divergence in assumptions about future model complexity, deployment architectures, and the pace of software and hardware optimization. Despite these unknowns, investment signals point to high growth. AI is expected to drive over $5.2 trillion in global digital infrastructure investment by 2030, accounting for more than 70% of projected data center capital expenditure. Based on these trends, industry estimates suggest that AI-driven workloads could add more than 20 GW of electricity demand annually over the next three to five years.

AI data centers consume vastly more energy than legacy systems

Table 4: AI vs Traditional Workload

Metric	AI vs. Traditional
Energy consumption	10x higher
Energy per Rack	~5× higher
GPU vs. CPU Power Use	5–7× higher
GenAI Training Power Use	7–8× over standard compute

Source: Bain & Company, EPRI, MIT Technical Review, IEA

Modern AI models such as GPT-4 are trained on thousands of GPUs — most commonly Nvidia's H100s — which have become the backbone of AI compute since their 2022 launch. These GPUs consume five to seven times more power than traditional CPUs, resulting in AI data center racks drawing over five times the energy of legacy configurations. As chip designs advance from 7nm to 2nm nodes, they pack more computing power into less space. While this boosts chip performance, it also increases energy intensity and thermal loads. This growing density is straining current cooling systems and power delivery infrastructure, underscoring the need for forward-looking standards in facility design and energy planning.

Beyond power-hungry chips, AI data centers are systemically resource-intensive and environmentally demanding. While GPUs dominate headlines, the broader infrastructure is equally demanding. CPUs coordinate massive data flows, while industrial cooling systems — often relying on millions of gallons of water — run continuously to prevent overheating. System-wide inefficiencies further compound the issue: power losses occur across long cable runs, outdated rack designs, and repeated voltage conversions. Together, these elements create a 24/7 digital backbone that consumes energy and water at an industrial scale, frequently out of sync with local grid capacity or environmental constraints. As AI accelerates, addressing these systemic inefficiencies is critical to aligning digital growth with sustainability goals.

AI's energy footprint also varies significantly by application. Training large-scale generative AI models can consume up to 7–8x more electricity than conventional computing tasks. As applications expand beyond text to include image, audio, and video generation, energy demands grow exponentially. For instance, text generation consumes approximately 2-watt-hours per query, but video generation requires roughly 25 times more electricity per query. If generative AI were used for every Google search, the company's electricity consumption could exceed 29 TWh per year — comparable to the annual electricity usage of Ireland.

Unlike conventional cloud services, which tend to follow predictable daily usage patterns, AI systems—particularly large-scale models—operate under continuous strain. Training these models requires uninterrupted, high-performance computing, with downtime viewed as both technically inefficient and financially costly. If a company has invested in expensive, high-density AI infrastructure, the imperative is to maximize utilization — running these facilities 24/7 to recoup capital expenditure and ensure optimal return on investment. Even inference workloads, which are less intensive than training, remain largely inflexible. Real-time AI services — such as chatbots, search, and recommendation engines — must respond instantly, leaving little room for delays or throttling. This means inference can't simply be "turned off and back on" without degrading performance or user experience. These demands create a baseline of always-on compute that is difficult to shift or interrupt, posing new challenges for grid planning and load flexibility.

AI workloads are less flexible than traditional digital services.

AI infrastructure is evolving into a two-tiered architecture, with distinct siting strategies for training and inference workloads.

AI infrastructure is increasingly evolving into a two-tiered architecture, with distinct siting strategies for training and inference workloads. This emerging design mirrors patterns from the logistics sector—where centralized training centers function like regional warehouses and inference hubs serve as last-mile delivery points embedded in cities.

AI training facilities, which develop large-scale models such as GPT-4, require enormous and continuous power. To ensure stable, cost-effective electricity supply, these centers are increasingly located near major clean energy sources such as hydroelectric or nuclear plants. For example, Microsoft's $500 million investment in Québec leverages the province's abundant hydropower, while in the U.S., Microsoft signed a 20-year power purchase agreement to repower the decommissioned Three Mile Island Unit 1 in Pennsylvania, securing nuclear energy for its AI training campus. In contrast, AI inference centers—which handle real-time applications like search, chat, and content recommendations—are being built closer to population centers to minimize latency and improve user experience. Hyperscalers and data center operators are deploying edge infrastructure such as AWS Local Zones and metro data hubs like Equinix to embed compute capacity directly into urban networks.

This purposeful separation between training and inference is now a defining feature of global AI infrastructure planning. For APEC economies, understanding this shift is critical. It influences where energy investments are needed, how to coordinate grid upgrades, and how to design resilient digital infrastructure that balances performance, proximity, and sustainability.

B. ECONOMY CHAPTERS

An “economy chapter” was created for each of the report’s nine focus economies: **Australia**, **Canada**, **Chile**, **China**, **Indonesia**, **Korea**, **Malaysia**, **Singapore**, and the **United States**. These focus economies were selected for their relevance to both global data center markets and as a representation of APEC’s diverse climates, geography, and energy mixes.

THE FOLLOWING KEY INDICATORS ARE INCLUDED IN EACH ANALYSIS:

Digital Readiness Index (2025): The Global Digital Readiness Index evaluates over 140 economies on their progress toward building sustainable digital economies, ranking them across four phases: Pre-, Early-Stage, Substantially Developed, or Highly Developed. The Index measures performance on a 0–100 scale, where 100 is optimal, based on achievements relative to each nation’s resources and global standing, with current scores ranging from 8 to 85. It assesses key “vital ingredients” of digital economy maturity across four dimensions: Economy, Environment, Social, and Governance, providing a holistic view that sustainability is essential for long-term digital growth.

Burn-to-Earn Index (2025): The Burn-to-Earn Index ranks 126 economies by how efficiently they convert digital activity into economic value, calculated as grams of CO_2 equivalent emitted per dollar of digital revenue. This metric surfaces the trade-offs between economic output and environmental impact.

Electricity Mix (2023): As defined and tracked by the International Energy Agency, the electricity mix is the share and combination of different energy sources — such as coal, gas, nuclear, and renewables — used to produce electricity in a given economy.

AUSTRALIA

INTRODUCTION

Australia's data center sector is experiencing unprecedented growth, driven by rising demand for AI, cloud computing, and digital services. As of 2024, the economy hosts over 300 data centers. National capacity reached 2.18 GW in 2024 and is projected to nearly double to 4.07 GW by 2030. Data center investment reached $15 billion in 2024 and is projected to exceed USD$26 billion by 2028, reflecting strong institutional confidence in the sector's growth trajectory. This momentum is led by hyperscale developers — AWS, for instance, has committed $13.2 billion (AUD $20 billion) to expand data center infrastructure in Sydney and Melbourne, supported by the construction of three solar farms in Victoria and Queensland to supplement energy demand. Currently, 60% of national capacity is concentrated in Sydney, but geographic diversification is underway. Melbourne's market is expected to double within five years, while new investments are accelerating in Brisbane, Perth, and Canberra. Western Sydney alone has over 2 GW of capacity under construction, positioning it as a critical digital hub facing growing grid and land-use constraints.

However, the data center sector's surging energy demand presents growing challenges. In 2024, data centers consumed about 5% of national electricity. Projections suggest this could rise to between 8% and 15% by 2030, depending on AI uptake and facility concentration. The Australian government is responding with new environmental guardrails, including energy performance standards, integration protocols, and emissions reporting mechanisms aligned with Australia's 2022 Climate Change Acts and its Paris Agreement commitments.

Key Indicators

DIGITAL READINESS INDEX: SCORE 62/100

BURN-TO-EARN INDEX: 137.1 GCO2E PER $ - RANK 70/126

ELECTRICITY MIX: COAL 47%, NATURAL GAS 18%, SOLAR 15.3%, WIND 11.4%, HYDRO 6.1%

LOCAL DIGITAL INFRASTRUCTURE LANDSCAPE

Digital infrastructure's surging energy demand presents growing challenges for local power.

Australia's data center sector is grappling with mounting electricity access constraints as demand for digital infrastructure accelerates. Securing a physical grid connection typically requires navigating a 6–12 month negotiation process with licensed Network Service Providers (NSPs). However, some parts of the high-voltage network are already at capacity, limiting where large-scale data centers can be feasibly sited. These delays are further exacerbated by speculative behavior — developers reserving energy capacity without near-term build plans — creating artificial grid congestion and increasing pressure on regulators to streamline and reform the connection process.

Compounding this challenge is Australia's current electricity mix, which remains heavily reliant on fossil fuels: coal (47%) and natural gas (18%) together supply two-thirds of the grid, with solar (15%), wind (11.4%), and hydro (6.1%) making up the rest. The planned closure of major coal-fired power stations in the coming decade is expected to intensify grid pressure, just as data center loads surge in scale and frequency. Meanwhile, in Australia, nuclear energy remains a heated debate, with public opposition and a longstanding federal ban on domestic nuclear power generation. While some states and industry stakeholders are revisiting the potential of SMRs to meet growing electricity demand and decarbonization goals, federal policy still prohibits deployment, limiting near-term prospects.

To mitigate grid constraints, Australia is gradually advancing demand response capabilities, with pilot programs and regulatory reforms creating conditions for large energy users — such as data centers — to participate in grid balancing. The Australian Energy Market Operator (AEMO) has expanded its Wholesale Demand Response Mechanism since 2021. While formal demand response obligations for data centers remain limited, states like New South Wales and Victoria have introduced market-based schemes that reward load flexibility. However, uptake remains modest, highlighting a nascent scale of deployment and significant room for expansion (AEMC)

Australia has established renewable procurement mechanisms, but grid access constraints and the structure of the national wholesale electricity market limit direct clean energy sourcing for data centers.

In Australia, the primary mechanism for renewable energy procurement by data centers is through corporate virtual PPAs, as physical PPAs are largely constrained by grid access rules and the centralized structure of the National Electricity Market (NEM). NEM is Australia's main wholesale electricity market, covering the eastern and southern states.

Under the Renewable Energy Target (RET), a national policy aiming to reduce greenhouse gas emissions and increase the share of electricity generated from renewable sources, data centers are designated as 'liable entities'. It means that certain data centers must offset their electricity use by purchasing Renewable Energy Certificates and, during times of grid stress, secure backup capacity, under the Retailer Reliability Obligation (RRO). While Australia lacks a standardized 24/7 carbon-free energy (CFE) procurement framework, operators like AirTrunk are early movers — partnering with solar developers and tech firms to pioneer CFE alignment through hourly tracking and renewable PPAs.

To reduce reliance on grid electricity, operators are pursuing co-location strategies with renewables and storage. Projects like Quinbrook's "Supernode" in Queensland integrate hyperscale campuses with 250 MW batteries that help absorb and dispatch excess solar power, enhancing local grid stability. However, remote siting poses connectivity challenges, particularly for fiber infrastructure.

Best practices

EFFICIENCY STANDARDS AND ENVIRONMENTAL BENCHMARKS FOR DATA CENTERS:

Australia has made sustainability central to federal ICT procurement. From July 2025, all government data centers must achieve a 5-star National Australian Built Environment Rating System (NABERS) rating, setting a national benchmark for efficiency. The updated Data Centre Panel adds stricter rules on PUE, water, carbon reporting, and ISO compliance. The Net Zero in Government Operations Strategy also requires 100% renewable electricity for federal facilities by 2025, one of the most ambitious targets in the APEC region.

RENEWABLE PROCUREMENT AND INNOVATION:

Australia is advancing models for clean energy integration into its data center sector. Operators are leveraging corporate virtual power purchase agreements (vPPAs) and pioneering 24/7 carbon-free energy (CFE) tracking to align operations with decarbonization goals. Projects such as Quinbrook's Supernode — which integrates hyperscale campuses with large-scale batteries — demonstrate how renewable procurement is being paired with grid-supporting innovation.

Water is becoming a critical constraint for Australia's data center expansion

Water management is an often overlooked but increasingly complex constraint. With approximately 196 local water authorities across Australia, new data centers must navigate a fragmented permitting landscape. In jurisdictions like New South Wales, large-scale facilities deemed "State significant developments" are required to submit Integrated Water Management Plans as part of the environmental impact assessment — adding yet another layer of regulatory oversight to balance digital growth with sustainability.

POLICY LANDSCAPE

Australia has positioned itself to harness the benefits of AI and digital technologies, with a new National AI Capability Plan currently in development. Scheduled for delivery in late 2025, the plan builds on existing initiatives like the National Reconstruction Fund and AI Adopt Centers to position Australia as a competitive leader in AI. Its four objectives are to grow investment in AI, strengthen industry capabilities, boost workforce skills, and secure economic resilience through sovereign capacity and responsible adoption. Local governments are also stepping up — offering fast-track planning pathways (e.g., New South Wales) and public investments in land, infrastructure, and connectivity to attract new developments.

As AI-driven demand explodes, so does the footprint of energy-intensive data centers, triggering urgent action to keep environmental impacts in check. In response, the government is tightening the screws on energy performance and emissions compliance. From July 2025, all data centers handling federal workloads must meet a 5-star NABERS (National Australian Built Environment Rating System) rating — a shift that embeds energy efficiency into ICT procurement and makes NABERS the de facto national benchmark. The Data Centre Panel rules (updated in May 2023) go further, requiring data centers to hit PUE <1.4, report biannually on water and carbon metrics, and comply with ISO environmental standards. Facilities that fall short risk losing access to government contracts.

To address grid stability, the Australian Energy Market Commission (AEMC) is advancing reforms targeting large, AI-driven loads. Package 2 — expected in late 2025 — will require data centers to stay operational during grid disturbances and potentially provide ancillary support through technologies like grid-forming inverters or demand modulation. These changes are critical as data center clusters strain transmission infrastructure and limit capacity for new entrants.

At the system level, climate and energy frameworks are tightening: the Climate Change Act 2022 codifies Australia's 2030 and 2050 net-zero targets; the National Greenhouse and Energy Reporting (NGER) Scheme mandates emissions reporting; and the Safeguard Mechanism caps emissions for high-intensity sites.

Table 5: Digital Infrastructure Policies

Policy Name	Date	Details	References
National AI Capability Plan	Under development	The National AI Capability Plan, expected by the end of 2025, will harness artificial intelligence to grow Australia's economy, support local industry, and secure long-term prosperity. Its four objectives are to grow investment, strengthen industry capabilities, boost AI skills, and ensure economic resilience through sovereign capability and responsible adoption.	Department of Industry, Science and Resources
5-Star NABERS Mandate	July 2025	All data centers hosting federal workloads must achieve a 5-star NABERS rating. Previously limited to Digital Transformation Agency's Data Centre Panel providers, this mandate is now part of the broader Net Zero in Government Operations Strategy, setting a baseline for ICT procurement and energy performance across government infrastructure.	Link
Australian Energy Market Commission Access Reform – Package 1 & 2	May 2025	Following the finalization of generator access reforms under Package 1 in May 2025, the Australian Energy Market Commission is updating National Electricity Market rules to manage the impact of large AI-driven data centers. To manage AI-driven load growth, the AEMC's Package 2 will require data centers to remain grid-stable and potentially support grid resilience via technologies such as grid-forming inverters.	Australian Energy Market Commission
Data Centre Panel	May 2023	Updated panel requirements enforce sustainability standards such as a 5-star NABERS rating, PUE <1.4, WUE, CUE metrics, and adherence to ISO 14001 and ISO/IEC TS 22237-4:2018. Bi-annual reporting and penalties for non-compliance with the NGER Act further align data centers with national emissions and efficiency goals.	Digital Transformation Agency
Climate and Energy Frameworks	Ongoing	• Net Zero in Government Operations Strategy: Commits the Australian Government to source 100% renewable electricity for all operations by 2025. • Climate Change Act 2022: Codifies Australia's 2030 and 2050 net-zero targets; establishes the Climate Change Authority to monitor progress and advise on national decarbonization strategies. • National Greenhouse and Energy Reporting (NGER) Scheme: Mandates annual reporting of emissions, energy production, and consumption by large emitters meeting defined thresholds. • Safeguard Mechanism: Imposes emission limits on facilities emitting over 100,000 tonnes of CO_2-equivalent annually, complementing the NGER Scheme. • Climate-Related Financial Disclosure (CRFD) Regime: Requires large companies and investment funds to report on climate risks and emissions under AASB S2, with phased implementation between 2025 and 2028. • Sustainable Finance Taxonomy: Introduces voluntary criteria for classifying sustainable investments; Phase 1 completed with focus on climate mitigation and social safeguards.	Dentons

CANADA

INTRODUCTION

Canada was the first economy in the world to launch a fully funded national AI strategy in 2017, establishing itself as an early leader in AI development. That head start catalyzed a vibrant ecosystem anchored by world-class institutions such as the Montreal Institute for Learning Algorithms (Mila), the Vector Institute for Artificial Intelligence in Toronto, and the Alberta Machine Intelligence Institute (Amii) in Edmonton. The economy is home to roughly 10% of the world's top-tier AI researchers, ranking second globally behind the United States. This concentration of AI talent supports a dynamic and growing workforce capable of advancing AI and cloud computing at scale. Today, Canada is leveraging that head start with a compelling mix of climate, energy, and resource advantages that make it an increasingly attractive destination for sustainable digital infrastructure.

With more than 80% of its electricity generated from non-emitting sources — primarily hydro, nuclear, and wind — Canada offers one of the lowest-carbon energy mixes among APEC economies. Canada's cold climate further enhances energy efficiency by lowering cooling demands, translating into lower operational costs. Beyond energy, Canada plays a critical role in global clean tech supply chains, boasting rich reserves of essential minerals like uranium, nickel, cobalt, and lithium. It also maintains long-standing leadership in clean technology and energy innovation (e.g., nuclear, CCUS, smartgrids, energy storage and batteries, energy efficiency), reinforcing its strategic advantage in clean energy technologies that support non-emitting electricity generation. Canada further brings strong expertise in water and energy-efficiency technologies, positioning it to advance more efficient and sustainable AI data center development

As of 2024, Canada hosts roughly 336 data centers with a total capacity of approximately 750 MW, projected to grow to 1.16 GW by 2029. Ontario, Quebec, and Alberta are among Canada's leading data center hubs, each offering distinct advantages that are drawing significant investment in AI-ready infrastructure. Ontario leads in data center concentration, benefiting from a reliable nuclear power supply and its position at the heart of the Toronto tech corridor — one of North America's foremost hubs for AI and cloud innovation. Quebec offers abundant, low-cost hydroelectric power that appeals to energy-intensive operations, while Alberta leverages its deregulated electricity market and fast-tracked permitting under the 2024 AI Data Center Strategy to accelerate new developments. Together, these strengths reflect Canada's broader national advantage: a combination of clean energy resources, favorable policy signals, and a skilled workforce position the economy as a premier North American destination for sustainable digital infrastructure in the AI era.

Key Indicators

DIGITAL ECONOMY SCORE: SCORE 70/100

BURN-TO-EARN INDEX: 40.2 GCO2E PER $ - RANK 21/126

ELECTRICITY MIX: HYDRO 57.4%, NATURAL GAS 15.4%, NUCLEAR 13.5%, WIND 6.4%, COAL 3.5%

LOCAL DIGITAL INFRASTRUCTURE LANDSCAPE

While Canada's overall energy mix for data centers is largely clean, the economy's diverse provincial energy profiles shape distinct regional strategies for powering digital infrastructure.

Canada's energy mix is among the cleanest and most reliable in the world. As of 2024, the national grid is powered by approximately 57.4% hydroelectricity, 13.5% nuclear, 6.4% wind, 15.4% natural gas, and just 3.5% coal. While energy is abundant, Canada's energy mix varies significantly across provinces, shaping distinct pathways for data center development.

Ontario offers a low-carbon energy mix, with 50% of its electricity coming from nuclear and 25.1% from hydro. In recent years, data centers in Ontario — once viewed as potential stressors on the grid — have increasingly been recognized as consistent, round-the-clock loads that can help absorb surplus generation during off-peak periods and support long-term system planning. Ontario's electricity system operator, the Independent Electricity System Operator (IESO), recognizes data centers as a primary growth driver in the commercial electricity sector, factoring their rapid expansion into its long-term energy demand forecasts. The Energy For Generations Plan (2025) introduces the province's latest integrated energy planning framework, aligning generation, transmission, and distribution, while emphasizing demand response for industrial users, including data centers. Nuclear deployment is central to Ontario's digital infrastructure strategy.

Ontario's Energy for Generations Plan (2025)

UNLOCKING OPPORTUNITIES IN THE DIGITAL ECONOMY

Ontario is embracing the growth of the data center sector as a strategic opportunity to drive investment, innovation, and job creation, particularly in northern and rural communities where new investments could anchor new high-tech ecosystems. To ensure responsible growth that aligns with public priorities, the government has introduced legislation that, if passed, would allow Ontario to manage electricity connection requests and prioritize data centers that deliver real local, strategic, and economic benefits — not just power consumption. Ontario will prioritize data centers that serve the province's economic interests, including those that create high-quality jobs, support domestic data hosting, and strengthen Ontario's position in the digital economy. This approach ensures the sector enhances Ontario's competitiveness, energy security, and long-term prosperity.

Quebec has one of the cleanest electricity grids in Canada, with 94% of its power generated from hydro and 5% from wind. This low-emissions profile has attracted investment from major cloud providers, including Microsoft's $500-million investment in 2023 to establish new data center regions in the province. Hydro-Québec, Quebec's utility, anticipates a significant increase in electricity consumption from the data center sector by projecting an additional 4.1 TWh in demand between 2023 and 2032.

Alberta presents a different energy profile. The province's grid is currently powered primarily by natural gas (67.1%) and coal (10.7%). With limited hydro resources and no operational nuclear facilities, natural gas remains the main scalable option for meeting large, high-uptime power demands from data centers. Reflecting this context, Beacon AI Centers has announced plans to invest $7.3 billion (C$10 billion) to develop six hyperscale data centers near Edmonton and Calgary. These facilities are expected to require approximately 4.5 GW of electricity, sourced primarily from natural gas. Alberta's electricity system operator (AESO) includes data center demand in its "High Electrification" scenario, reflecting anticipated growth from AI and electrified industries.

Canada's accelerating data center growth is straining in-province grid infrastructure

While the economy is experiencing significant growth in data center development, generation and transmission expansion within provinces has not kept pace, as it takes years to plan and build new lines or generation assets. In Alberta, the challenge is particularly acute. As of June 2024, the Alberta Electric System Operator (AESO) — the province's independent system planner and operator — reported that 29 data center connection applications are currently progressing through its connection process, representing over 16 GW of demand. For context, Alberta's current peak system demand sits at approximately 12GW, underscoring the scale of anticipated load growth and the mounting pressure on grid infrastructure. The AESO, under Phase 1 of its interim approach to connecting large loads, has implemented a reliability cap limit of 1,200 MW for connecting projects up to 2027-28. Projects not selected under the cap to enter the connection queue will have to wait until the permanent framework is in place.

To address mounting grid constraints, provinces have increasingly encouraged data center developers to pursue on-site or co-located generation as a faster and more reliable alternative to centralized grid connections. This approach is gaining momentum, particularly in deregulated electricity markets like Alberta's, where in the past 5 years corporate power purchase agreements (PPAs) have enabled large buyers to contract directly with renewable energy producers. As of 2024, public sources indicate 30+ power purchase agreements have been made in Alberta with renewable energy providers since 2019, totaling nearly 3,000MW. More recently, with the advent of AI data center proposals with higher electricity demand profiles, project proponents have expressed interest in accessing a wider mix of power sources. Under Pillar 1 of Alberta's AI Data Centers Strategy, the province will prioritize both off-grid and grid-connected power solutions that must avoid compromising the affordability and reliability of electricity, utilize Alberta's natural gas reserves, and integrate renewable energy sources and carbon capture technology, where feasible.

Permitting remains a major bottleneck for clean energy and digital infrastructure development in Canada, with challenges varying by province and project type.

In Alberta, the provincial government has introduced a fast-tracked, concierge-style permitting process for data center development, aiming to accelerate approvals and attract investment. However, while siting digital infrastructure has become more efficient, permitting energy projects — particularly large-scale generation — takes several years. In response, power companies in Ontario, for example, are proactively front-loading the permitting process for large-scale nuclear projects to reduce future delays and ensure project readiness. This is crucial, as nuclear development is often delayed not by construction — typically completed in five to six years — but by the extensive pre-construction requirements. These include environmental assessments, Indigenous consultations, endangered species evaluations, and site compatibility studies, which can span a decade or more.

Best practices

CLEAN ENERGY ALIGNMENT:

Canada's abundant low-carbon power — hydro, nuclear, and renewables — offers a strategic advantage for powering sustainable data centers and aligning digital growth with climate goals.

PUBLIC-PRIVATE PARTNERSHIP IN POLICY DESIGN:

Through broad consultations across academia, industry, and SMEs, the federal government identified gaps in domestic digital infrastructure and stepped in to address bottlenecks through federal funding mechanisms.

INDIGENOUS PARTNERSHIP:

Canada is advancing inclusive digital infrastructure by embedding Indigenous equity and co-ownership into energy and connectivity projects.

GLOBAL TALENT ADVANTAGE:

Canada's AI and software ecosystem, anchored by institutions such as Mila, Amii, and Vector Institute, is increasingly attracting international attention. Foreign partners view Canada as a key research hub, not just for academic output but for high-quality talent.

As an example, Bruce Power's current expansion plans highlight the regulatory complexity associated with large-scale energy projects in Ontario, Canada. The company operates two units totaling eight reactors at its Tiverton site, which supply approximately 30% of Ontario's electricity. It is now exploring the development of a third unit within the existing site footprint. Although the site is already regulated and operational, current requirements still mandate a full set of impact assessments. This reflects a broader trend in which even projects that aim to densify established infrastructure must navigate lengthy and duplicative permitting processes.

Both federal and provincial governments have acknowledged the delays caused by regulatory complexity and are now pledging reforms to advance Indigenous partnership and streamline environmental and endangered species assessments to accelerate clean energy deployment. Ontario, for example, is taking steps to modernize

Indigenous engagement practices to incentivize Indigenous equity ownership in addition to adhering to consultation requirements. The province has adopted a points-based approval system that awards higher scores to energy projects with meaningful Indigenous ownership, aiming to align reconciliation with energy expansion.

On the federal level, due to Canada's jurisdictional responsibilities, the federal government does not directly engage in national-level planning for data centers. Instead, rising demand from data centers and AI is treated as part of the broader electricity demand expected across Canada. Through the national Clean Electricity Strategy, the federal government has committed to modernizing electricity systems by supporting strategic transmission investments and major clean power projects.

The recent passage of the One Canadian Economy Act (Bill C-5) in June 2025 represents a significant step forward, enabling the fast-tracking of nationally significant infrastructure projects and leading to the launch of the new Major Projects Office (MPO). The MPO is designed to serve as a single point of contact to accelerate nation-building projects. Its mandate is twofold: first, to streamline and expedite regulatory approval processes; and second, to help structure and coordinate financing where needed. Importantly, the MPO is tasked with advancing a Canadian sovereign cloud, placing digital infrastructure—compute capacity, AI, and quantum—on the same strategic footing as pipelines, nuclear plants, and ports. This signals a broader shift in Canada's nation-building agenda, where digital infrastructure is no longer viewed as ancillary but as foundational to Canada's competitiveness, security, and economic independence.

POLICY LANDSCAPE

Canada's digital infrastructure and energy governance are highly decentralized.

In Canada, energy and data center development fall under provincial and territorial jurisdiction. This decentralized governance model creates substantial variation in planning, permitting, and regulatory implementation across regions. In the absence of a coordinated federal–provincial framework, Canada faces growing challenges in aligning the rapid expansion of digital infrastructure with clean energy access, grid capacity, and long-term climate goals.

At the federal level, three institutions are playing increasingly active roles in shaping Canada's digital infrastructure landscape. Innovation, Science and Economic Development Canada (ISED) serves as the lead agency, guiding sovereign AI compute investments, digital infrastructure funding, and regulatory alignment across government. ISED is positioned to embed sustainability considerations into investments and to showcase a public program approach that prioritizes sustainable practices. Natural Resources Canada (NRCan) supports the digital infrastructure agenda through its ongoing program and policy work to ensure Canada's energy system is affordable, reliable, and low-carbon, as well as through the production of the Best Practice Guide for Canadian Data Centers (2024). The recently created Ministry of AI and Digital Innovation (2025) reflects the growing importance of AI and digital infrastructure in Canada's broader economic and innovation strategies.

Canada's first major step came through the 2017 Pan-Canadian AI Strategy, which positioned academic excellence as a foundation for global leadership in AI. The strategy established three national AI institutes and helped attract top-tier AI talent. Its second phase, launched in 2022, broadened the focus to commercialization, talent development, and ecosystem growth. Building on the foundation laid by Canada's Pan-Canadian AI Strategy, the 2024 Sovereign AI Compute Strategy represents a significant federal commitment to expanding domestic AI infrastructure. Backed by $1.45 billion (C$2 billion), the strategy aims to scale Canada's AI compute capacity while prioritizing environmentally responsible and energy-efficient infrastructure. Developed through broad national consultations, the strategy positions data centers, cloud services, and AI compute as critical national assets — integral to Canada's innovation capacity, digital sovereignty, and long-term competitiveness. Rather than imposing mandates, the federal government has opted to shape the growth of data center infrastructure through targeted incentives—most notably via the Strategic Innovation Fund and dedicated AI strategy funding. In practice, this incentive-based approach serves as a de facto national framework, channeling investment toward Canadian-controlled firms and projects that align with broader goals around sustainability, security, and technological leadership.

To support digital infrastructure growth, federal investment oversight has also evolved. Reforms to the Investment Canada Act (ICA) in 2024 introduced new provisions related to national security, foreign investment review, and strategic asset protection. Data centers and AI systems are now among the sectors requiring mandatory pre-closing notifications, with expanded review criteria that include cybersecurity, data governance, and protection of publicly funded IP.

In terms of AI demand forecasting, the federal government does not rely on a centralized authority to forecast electricity demand from digital infrastructure. Instead, the federal government leverages a distributed network of agencies to inform planning on energy demand, AI compute needs, and infrastructure readiness. Key players like CANARIE, the Digital Research Alliance of Canada (DRAC), and the national AI institutes (Amii, Mila, and Vector).

Canada's decentralized regulatory landscape results in varied approaches to sustainability standards, clean energy access, and demand-side flexibility for data centers

Currently, there is no mandatory federal standard for reporting sustainability metrics such as PUE, WUE or CUE. Similarly, policies related to heat reuse, water management, and energy efficiency are typically set at the provincial or municipal level. While many hyperscale operators voluntarily adopt international best practices — such as those from the Uptime Institute or ASHRAE — Canada lacks a unified national framework to ensure consistent application and benchmarking across regions.

Natural Resources Canada's Best Practice Guide for Canadian Data Centres (2024) establishes a baseline for energy-efficient operations, helping close transparency gaps and align Canada with international standards. It identifies opportunities in IT, cooling, power systems, and waste heat recovery, offering operators practical ways to reduce energy use and costs. For now, however, the Guide remains voluntary, with its impact dependent on industry uptake. Natural Resources Canada is working with provinces, regulators, and utilities to assess data tracking and explore how efficiency can be embedded into future planning. Potential next steps include making certification or labelling a prerequisite for project approvals, encouraging federally owned data centers to adopt these standards, and integrating them into future federal procurement.

Access to clean electricity also differs significantly between provinces. Alberta's deregulated electricity market enables direct procurement of renewable energy through corporate power purchase agreements (PPAs), giving it a relative advantage in attracting data center investment. In contrast, Ontario, Quebec, and British Columbia operate within more centralized electricity systems, where regulatory and market design constraints can limit corporate access to non-emitting power sources such as hydro and nuclear. Ontario's voluntary Clean Energy Credit (CEC) market is a promising tool for incentivizing clean energy procurement and enabling nuclear co-location for data centers, but its effectiveness is constrained by limited enforcement mechanisms and undefined eligibility standards. At the federal level, Canada does not yet have a dedicated policy for powering digital infrastructure with low-carbon energy. However, it has introduced a suite of policies and investment tools to accelerate clean power deployment more broadly. These include investment incentives such as the Clean Economy Investment Tax Credits; strategic financing through the Canada Infrastructure Bank (C$20B) and the Canada Growth Fund (C$15B); and targeted programming like Natural Resources Canada's Smart Renewables and Electrification Pathways Program (C$4.5B) to expand clean electricity infrastructure. While these are not specific to digital infrastructure, they establish a supportive policy and investment environment that could be leveraged as demand from data centers and AI accelerates.

Incentives for demand-side flexibility — such as load shifting, curtailment, or participation in flexibility markets — are available in some provinces but remain limited in scope. At the federal level, policy tools to support demand response for high-density, energy-intensive facilities — particularly AI clusters — are still under development.

On waste heat, CanmetENERGY—Natural Resources Canada's largest clean energy research center—is conducting an initial mapping exercise to assess recovery opportunities across multiple sectors, including data centers, to identify where integration may be most viable. Natural Resources

Canada is exploring a more detailed project to pinpoint specific sites where waste heat recovery from data centers may be viable, potentially through a comprehensive feasibility study. However, to date, no in-depth technical assessments have focused specifically on data centers, highlighting the need for targeted analysis to guide future policy and investment decisions.

Encouragingly, several provinces are advancing inclusive approaches to energy development through stronger Indigenous partnerships. Utilities and First Nations are increasingly embracing co-ownership and benefit-sharing models, enabling direct equity participation and long-term revenue generation for Indigenous communities. The various measures put forward have enabled significant Indigenous equity stakes in a number of recent transmission projects, even achieving a majority stake in some cases, reflecting a broader shift toward community-rooted, collaborative infrastructure development.

Table 6: Digital Infrastructure Policies

Policy Name	Date	Detail	References
Pan-Canadian AI Strategy	Phase 1 – 2017 Phase 2 – 2022	Launched with $90B (C$125M), this was the world's first national AI strategy, focused on building global research leadership through the creation of Mila (Montreal), Vector (Toronto), and Amii (Edmonton). A $320M (C$443M) Phase 2 expansion in 2022 shifted focus toward commercialization, workforce development, and ecosystem growth. The strategy now underpins Canada's shift toward sovereign AI compute infrastructure.	Innovation, Science and Economic Development Canada
Sovereign AI Compute Strategy	2025	A $1.45B (C$2B) five-year strategy aimed at scaling Canadian-owned AI compute infrastructure with strong sustainability and sovereignty requirements. It includes funding for compute access (targeting SMEs), support for energy-efficient, Canadian-owned data centers, and investment in public high-performance computing (HPC) for research and national security. All applicants must submit plans to minimize environmental impact.	Innovation, Science and Economic Development Canada
Strategic Innovation Fund (SIF)	2025	A flexible, performance-based fund supporting high-impact digital infrastructure projects, including AI data centers. Proposals are evaluated on sustainability (e.g., PUE), data sovereignty, cost efficiency, and long-term national value. Funding levels adjust based on metrics and public interest outcomes, and projects may include clauses on IP retention, data residency, or transparency.	Innovation, Science and Economic Development Canada
Investment Canada Act (ICA) Reform	2024	Recent reforms increase scrutiny on foreign investment in digital infrastructure sectors such as cloud, AI, and data centers. Measures include mandatory pre-closing notifications, expanded national security reviews, and interim orders. Emphasis is placed on Canadian ownership, cybersecurity, protection of IP, and alignment with national values.	Innovation, Science and Economic Development Canada
Powering Canada's Future: A Clean Electricity Strategy	2024	Frames a national vision to decarbonize the electricity grid while maintaining affordability and reliability. Clean electricity is positioned as a prerequisite for new high-demand projects, including data centers. Encourages load flexibility, energy efficiency, and partnerships with Indigenous and remote communities. "*A Best Practice Guide for Canadian Data Centres*" was released in 2025 to promote energy-efficient, sustainable design and operation.	Natural Resources Canada

CHILE

Key Indicators

DIGITAL READINESS INDEX: SCORE 61/100

BURN-TO-EARN INDEX: 118.7 GCO2E PER $ - RANK 62/126

ELECTRICITY MIX: HYDRO (26%), SOLAR PV 20%, NATURAL GAS 19%, COAL 16.1%, WIND 10.4%, BIOFUELS 6.3%

INTRODUCTION

Chile holds the top ranking in Latin America on the AI Index, underscoring its readiness for AI adoption and advanced digital infrastructure deployment. As of 2024, it hosts 33 operational data centers. While most facilities are retail centers offering short-term leases to multiple clients, the market is shifting rapidly toward AI-driven hyperscale facilities. With 34 greenfield projects or expansions already in development, Chile has surpassed its 2030 industry growth target years ahead of schedule. Several analysts project that the installed capacity of data centers in Chile will exhibit the highest growth rate throughout the Latin American region by 2030, consolidating the economy as the second-largest Latin American market, after Brazil. Leading global operators are expanding their presence, with Microsoft's launch of a data center region marking a major milestone. AWS has also committed $4 billion to regional development, reinforcing investor confidence in Chile's role as a long-term hub for cloud and AI infrastructure.

Chile offers a rare convergence of renewable energy abundance, critical mineral resource wealth, digital connectivity, and proactive policy, positioning itself as a climate-resilient, nearshore leader in digital infrastructure. Chile's digital economy is anchored in Santiago and Valparaíso, supported by a robust fiber-optic backbone and an advanced IT ecosystem. Chile is strengthening its role in global digital connectivity through a growing portfolio of subsea fiber-optic cables. Google operates the Curie Cable, a 10,500km subsea cable linking Chile to California with a branch to Panama. In August 2025, Google and Chile announced the Humboldt Cable, a 14,800km subsea cable connecting Valparaíso to Sydney via French Polynesia. Set for completion in 2027, Humboldt Cable will be the first direct subsea link between South America and Asia-Oceania, reinforcing Chile's status as a digital gateway between continents.

Chile's role extends beyond digital infrastructure into the clean technology supply chain. As one of the world's two largest lithium producers alongside Australia, it accounted for 72% of global production in 2023. Lithium is a critical input for batteries used in data centers and renewable energy storage systems — positioning Chile as both a leader in digital infrastructure and a cornerstone of the global energy transition.

LOCAL DIGITAL INFRASTRUCTURE LANDSCAPE

> **Chile is building a 60% renewable, storage-backed grid to power AI-ready data centers.**

Chile is rapidly transforming its energy system, reducing dependence on imported fossil fuels and leveraging abundant domestic renewable resources to power both economic growth and the next generation of digital infrastructure. Renewables now account for more than 60% of electricity generation—driven by hydro (26%), solar PV (20%), wind (10.4%), and biofuels (6.3%). Fossil fuels, primarily natural gas (19%) and coal (16.1%), still supply about 35%, but their role is steadily shrinking as clean energy capacity expands, and coal plants are retired. Between 2019 and 2024, Chile decommissioned 11 coal-fired power plants and aims to phase out or convert all remaining coal units by 2040 or earlier.

Chile's updated National Energy Policy (2015, revised 2022) outlines ambitious milestones: achieve 80% renewable electricity generation with 2 GW of storage by 2030. Thanks to these goals, Chile has encouraged massive investment in renewables. This renewable build-out is directly shaping the economy's appeal as a low-carbon data hub. In 2019, Google partnered with AES Chile to construct 23 wind turbines in the Biobío region as part of a hybrid wind-solar project powering its first Latin American data center. Now operational, the facility runs on roughly 80% carbon-free energy, demonstrating how clean energy investment can directly decarbonize digital operations.

Energy storage is also central to this strategy. The Oasis de Atacama battery project, launched in 2024, will deliver 1.2 GWh of storage capacity in its first phases and scale to 4.1 GWh by 2026 — making it the largest facility of its kind globally. Such large-scale storage enhances grid stability, supports high renewable penetration, and ensures reliable power for AI-driven hyperscale data centers and other mission-critical operations.

Looking ahead, Chile is advancing next-generation fuels to further secure and decarbonize the digital economy's energy supply. The Green Hydrogen Action Plan 2023–2030, released in April 2024, sets out 18 strategic pillars, 81 targeted measures, and 177 milestones to scale green hydrogen production and adoption. This initiative positions Chile as both a renewable powerhouse and a pioneer in integrating green hydrogen into backup and auxiliary power systems for data centers — linking sustainability with long-term energy resilience.

Best Practices

RENEWABLE RESOURCES:

Chile integrates data center policy with its broader clean energy agenda. The economy leverages abundant solar and wind resources, advances green hydrogen through the Green Hydrogen Action Plan 2023–2030, and ties data center development to clean power procurement and industrial decarbonization goals.

NATIONAL STRATEGY WITH DECENTRALIZED FOCUS & GUIDELINES:

Through its National Data Center Plan and linked policies (e.g., AI Policy, National Energy Policy), Chile promotes regional advantages with a supported framework, updated construction and environmental evaluation guides, professional training, and multi-stakeholder governance committees to ensure sustainable and interconnected growth.

Chile's renewable energy leadership is now challenged by grid constraints

Chile has long been viewed as a leader in renewable energy investment, consistently ranking among the top three most attractive emerging markets for clean power between 2015 and 2023. However, structural challenges in scaling renewable generation have become increasingly evident. In 2021, several renewable electricity producers experienced financial distress, prompting policy debate on electricity market reform. The rapid build-out of intermittent renewable sources has introduced operational risks such as intraday price volatility and curtailment, where renewable output must be reduced because the system cannot absorb it. Grid capacity in many regions remains insufficient or technically incompatible with large-scale renewable integration, while transmission and distribution expansion often lags behind generation growth. Severe weather events have further exposed these vulnerabilities. In 2024, a nationwide blackout plunged Santiago into darkness and disrupted the entire city, underscoring the fragility of electricity infrastructure amid rising demand from sectors such as AI-ready data centers.

In response, the government enacted Law No. 21.721, which strengthens the regulatory framework for electricity transmission, encourages competitive project development, and aims to improve system resilience. The law is designed to accelerate transmission buildout, enhance grid reliability, and better withstand extreme weather events — critical steps for ensuring a stable power supply to mission-critical digital infrastructure.

Chile's deepening water shortage has made data center cooling a contentious issue, fueling public pushback.

Chile's water crisis is structural and severe. Chile has been in a "megadrought" since 2010 — its longest in at least a millennium — with central Chile receiving 20–30% less rainfall over the past decade. Climate change is accelerating the retreat of Andean glaciers, Santiago's key freshwater source, while agriculture, mining, and urban growth intensify competition for dwindling supplies. The Dirección General de Aguas (DGA), Chile's General Water Directorate, has declared multiple river basins, including the Maipo that sustains Santiago, in a state of "extreme water shortage." In 2022, the capital even prepared rotating water rationing plans, a first in its history.

In this context, data center cooling demands have become politically charged. These facilities' high-volume water use competes directly with residential and agricultural needs in water-stressed regions. Community pushback is growing: in 2023, protests over water consumption forced Google to withdraw its permit for a second Santiago facility. The tech giant has since restarted its $200 million data center plans from scratch, redesigning the project to use air cooling instead of water. The dispute has become a flashpoint in Chile's broader struggle to reconcile rapid digital infrastructure growth with the urgent need to safeguard its scarce water resources.

POLICY LANDSCAPE

Chile is an early mover in national AI strategy. In 2020, the Ministry of Science, Technology, Knowledge, and Innovation led a nationwide consultation to capture public and industry perspectives, which shaped the Chilean AI Policy 2021–2030. The policy sets 70 priority measures and 185 initiatives to advance social, economic, and talent development goals over a decade. Building on this, 2024 marked the launch of the National Data Center Plan, a cross-government initiative to position Chile as a global hub. With 30+ projects in development, the plan targets tripling industry size within five years, securing $2.5 billion in new investment, promoting decentralization, and embedding sustainability in sector growth.

One focus is on improving the sustainability of data centers. Measures include introducing Clean Production Agreements to reduce resource use and environmental impact. The government will also regularly update construction and environmental evaluation guides to reflect changing conditions. In response to water constraints and grid pressures, a multi-stakeholder committee will link government, industry, and civil society to steer sustainable, efficient, and interconnected development.

Table 7: Digital Infrastructure Policies

Policy Name	Date	Detail	References
Chilean AI Policy 2021–2030	2021	Establishes an AI strategy across three pillars—enabling factors, AI development/adoption, and ethics/regulation/socioeconomic impacts	Link
National Data Center Plan	2024	Aims to triple the data center industry in five years, attract US$2.5B, decentralize facilities, and embed sustainability	UNCTAD Investment Policy Hub
Green Hydrogen Action Plan 2023–2030	2023	A roadmap with 18 action lines and 81 specific initiatives to develop green H_2 industry, attract investment, and streamline permitting	UNCTAD Investment Policy Hub
Law No. 21.721 – Electricity Transmission Reform	2024	Amends electrical services law to strengthen transmission infrastructure as a carbon-neutral enabler	DLA Piper, Legalink

CHINA

INTRODUCTION

China is advancing one of the world's most extensive digital infrastructure programs to underpin its AI ambitions, which it frames as a transformative frontier for both economic modernization and national competitiveness. National targets aim to grow the core AI industry to over [$140 billion] by 2030 and related sectors to $1.4 trillion in value. This trajectory is supported by a strategic combination of large-scale data center development, high-speed network deployment, and targeted incentives for AI adoption across key industries. Leading Chinese tech companies have committed massive capital to AI infrastructure. Alibaba Group has pledged more than $50 billion over the next three years for cloud computing and AI hardware, while ByteDance is earmarking roughly $20 billion for GPUs and data center expansion. Growth is further reinforced by ecosystem integration, linking technology companies, telecommunications providers, and leading research institutions to optimize computing resources and accelerate AI deployment. These capabilities are underpinned by an extensive 5G network, hyperscale cloud platforms, and high-capacity interconnections between regional hubs, enabling large-scale model training, real-time analytics, and rapid data exchange.

AI and digital infrastructure growth in China is reinforced by a significant domestic talent base, accounting for nearly half of the world's top AI researchers. Collaborations between universities and industry — such as DeepSeek's co-authored work with Tsinghua, Peking, and Nanjing universities — illustrate how academic expertise feeds directly into AI model development and compute-intensive applications. This integration of talent and infrastructure strengthens the utilization of China's expanding network of AI-optimized data centers, ensuring that high-capacity computing resources are matched with the skills to deploy them effectively.

Key Indicators

DIGITAL READINESS INDEX: SCORE 49/100

BURN-TO-EARN INDEX: 173.5 GCO2E PER $ - RANK 84/126

ELECTRICITY MIX: COAL 61%, HYDRO 13.5%, WIND 9.3%, SOLAR 6.1%, NUCLEAR 4.6%, BIOFUELS 2.1%, GEOTHERMAL 3.9%

LOCAL DIGITAL INFRASTRUCTURE LANDSCAPE

China's electricity mix remains dominated by coal (61.3%), with hydro (13.5%), wind (9.3%), solar (6.1%), and nuclear (4.6%) making up a smaller but rapidly expanding share. China is accelerating its transition from coal to renewable and nuclear power to meet the surging energy needs of its AI and data center sector. The launch of China's first Energy Law on November 8, 2024 — effective January 2025 — provides an institutional framework for the energy transition, marking a milestone in national energy legislation.

China is rapidly shifting from coal to renewable and nuclear power to secure low-carbon energy for its growing AI and data center sector.

China possesses the world's largest and most complete renewable energy industry chain, enabling rapid deployment of green power to support AI and data center growth. By 2030, the economy is projected to account for nearly 60% of all global renewable capacity additions, hosting almost half of the world's total renewable capacity by the decade's end. Hydropower remains the largest renewable source, but generation growth has been constrained by high temperatures and low precipitation, while wind and solar capacity have expanded at pace.

Nuclear power is also a strategic pillar for low-carbon, high-reliability electricity supply to energy-intensive data center facilities. China operates 56 reactors, primarily along the southeast coast, and has 30 reactors under construction. More than 50% of global new nuclear builds are now in China, and with rapid approvals, total nuclear capacity — operational, under construction, and approved — positions China as the global leader in nuclear power development. China's nuclear reactor construction time averages around 6 years—significantly faster than the 10+ years typical in most other economies. Looking ahead, small modular reactors (SMRs) are expected to begin contributing meaningfully to the data center power mix after 2030, further diversifying the low-carbon supply base.

China's national digital infrastructure planning, coupled with a market shift toward low-latency, inference-focused workloads, has resulted in a significant portion of capacity remaining underutilized.

Most data centers in China are concentrated in eastern provinces that rely on coal-fired generation, amplifying both carbon intensity and grid stress. To balance this, China runs one of the most orchestrated national coordination efforts in digital infrastructure within the APEC region through the "East Data, West Compute" policy. National and provincial governments are promoting new data center siting in renewable-rich western provinces through colocation mandates and targeted incentives. Between 2023 and 2024, more than 500 facilities were announced, backed by policy frameworks and economic support for developers.

However, a significant share of these facilities was developed on accelerated timelines, with varying levels of energy efficiency optimization and often in locations distant from major technology hubs. At the same time, the emergence of DeepSeek's open-source reasoning model marked a turning point in the economics of AI infrastructure, prompting a shift in demand from large-scale pretraining to inference-driven workloads. This transition favors low-latency, inference-optimized systems situated near major technology hubs, creating a structural mismatch with the remote, training-oriented clusters built during the earlier expansion phase. As a result, up to 80% of newly added computing resources are currently underutilized. While there is potential for consolidation and reallocation of resources, China's experience illustrates the challenge of aligning infrastructure planning with evolving technical requirements and AI application needs.

Best practices

INTEGRATED RENEWABLE ENERGY TARGETS FOR DATA CENTERS:

China integrates its renewable energy objectives with data center development by encouraging facilities in resource-rich regions to directly source green power, a strategy reinforced by the large-scale "East Data, West Compute" policy. The approach is designed to relieve power and land constraints in coastal cities, improve the utilization of renewable energy, and ease grid congestion.

PUSH FOR LOCAL TALENT AND R&D:

China demonstrates the value of aligning talent development with strategic R&D investment. By cultivating local AI engineers, researchers, and hardware designers while funding supercomputing centers and indigenous chip development, China strengthens self-reliance and reduces dependence on global supply chains.

China's unmatched grid expansion removed the capacity constraints limiting AI growth

China's AI and data center growth is underpinned by one of the fastest-expanding power systems globally. By 2030, data centers in China are projected to consume over 400 billion kWh of electricity — representing 3.7% of the nation's total power use. Backed by large-scale public investment, the national grid has the capacity to integrate new generation at a pace unmatched by most economies, removing a critical constraint on large computing hubs. In 2024, China added 429 GW of net new generation while directing a significant share of this capacity to renewable-rich regions prioritized for data center development. This scale and speed of grid expansion not only support rapid digital infrastructure deployment but also give domestic AI operators the flexibility to scale energy-intensive AI workloads without encountering the supply constraints seen in other markets.

However, sustaining this pace of grid-enabled digital infrastructure growth will depend on overcoming several system-level challenges. Key challenges include variability in regional electricity markets and a lack of standardization in PPAs. Large tech firms like Alibaba and Tencent are investing in both on-site renewables and provincial PPAs to hedge volatility. Key priorities in China include strengthening high-speed network links to connect geographically distributed computing hubs and encouraging broader multi-stakeholder coordination to align infrastructure, technology, and operational frameworks.

China is strengthening its digital backbone by advancing self-reliant technology ecosystems and accelerating the deployment of low-carbon data centers.

China is working to de-risk its digital backbone from geopolitical shocks by building self-reliant technology ecosystems — from semiconductor fabrication and AI hardware/software to physical infrastructure and algorithms — ensuring that no external disruption can cripple its core systems. For example, China is channeling significant public funding, including a recent $47 billion round from the National Integrated Circuit Industry Investment Fund, into R&D for advanced chipmaking technologies.

In parallel, China is pairing its rapid data center expansion with accelerated R&D and commercialization of sustainable technologies. The Dongjiang Lake Big Data Centre in Hunan Province, run by China Mobile, showcases China's push for green cooling innovation. Hosting 36,000 servers, it draws naturally cold water from Dongjiang Lake to regulate temperatures, reducing reliance on energy-hungry chillers. Powered mostly by hydro, wind, and solar, the system is designed to save 2 billion kWh of electricity a year while keeping its environmental impact low. Similarly, in Shanghai, Hicloud is developing a wind-powered underwater data center that aims to cut electricity use compared to conventional land-based facilities. The facility will use naturally cold seawater for heat exchange and pair the setup with offshore wind power to further lower its carbon footprint. Together, these initiatives demonstrate how China is converting R&D into scalable, low-carbon digital infrastructure, setting a new benchmark for sustainable technology deployment.

POLICY LANDSCAPE

China's strategy for AI infrastructure development is underpinned by a coordinated central–local governance model. At the national level, the central government formulates strategic visions and policy frameworks that define the trajectory of AI, while local governments translate these directives into tailored programs and investments aligned with regional priorities. This alignment has facilitated the redistribution of digital infrastructure across the economy and spurred the establishment of state-backed AI laboratories, pilot zones, and innovation clusters in major AI-driven hubs such as Shanghai, Shenzhen, and Beijing.

The central government's long-term commitment to AI development was first articulated in the New Generation AI Development Plan (2017), which outlined a three-phase roadmap aimed at positioning China as a global leader in AI by 2030. Building on this strategic foundation, subsequent policies have addressed more specific priorities. The 14th Five-Year Plan for National Economic and Social Development (2021) set quantitative targets to increase annual R&D expenditure by up to 7% and to expand the digital economy's share to 10% of national GDP. In 2024, the Cyberspace Administration of China (CAC) released the Guidelines for the Construction of a National Comprehensive Standardization System for the AI Industry, calling for the AI development of more than 50 national and industry standards to unify technical and operational practices across the sector.

A cornerstone of China's AI infrastructure strategy is the "East Data, West Compute" initiative, launched in 2021 through the Implementation Plan for the National Integrated Big Data Center Collaborative Innovation System. The policy aims to combine public and private cloud resources into a unified platform. The objective is to correct regional imbalances in computing capacity by enabling more efficient transfers of data, workloads, and algorithm-driven applications across the economy.

Under this framework, eastern provinces — home to advanced digital economies and high demand for compute — can shift workloads to western regions, which, while less digitally developed, possess abundant renewable energy potential. In practice, real-time, high-intensity computing will be concentrated in key eastern hubs such as Beijing–Tianjin–Hebei, the Yangtze River Delta, the Guangdong–Hong Kong–Macao Greater Bay Area, and Chengdu–Chongqing. Meanwhile, large-scale data centers in western provinces including Guizhou, Inner Mongolia, Gansu, and Ningxia will focus on non–time-sensitive processing, leveraging their lower energy costs and cleaner power supply.

Follow-up measures have expanded the scope of this strategy. The Guiding Opinions on Accelerating the Construction of a National Integrated Computing Power Network (2023) set a target for 60% of all new computing capacity to be located in the designated hub nodes by 2025. The Action Plan for the High-Quality Development of Computing Power Infrastructure outlined a goal of achieving 300 EFLOPS of total computing power by 2025, including 105 EFLOPS dedicated specifically to AI workloads.

Investment policy has evolved alongside infrastructure expansion. In January 2025, China introduced the $8.2-billion National AI Industry Investment Fund, complementing the $138-billion National Venture Capital Guidance Fund, which targets AI-adjacent sectors such as robotics and embodied intelligence. Local governments have also established their own AI-focused investment vehicles, with Shanghai, Hangzhou and Beijing among the most active in deploying capital to attract talent and scale industry capabilities. Shanghai, for instance, has issued China's first provincial-level regulation for industrial AI development, while also experimenting with subsidies and voucher systems to lower costs for algorithm research and model training. These local initiatives, combined with national-level strategy, are shaping an increasingly integrated, resource-optimized AI ecosystem across China.

China is pursuing a standards-driven, state-led strategy to green its rapidly expanding data center sector

China has adopted a state-led approach to improving the environmental performance of its expanding data center sector, combining mandatory targets with a range of technical standards and policy tools. Joint directives from the National Development and Reform Commission, the Ministry of Industry and Information Technology (MIIT), and other central bodies set an ambitious goal of reducing the average PUE of data centers to below 1.5 by 2025. Supporting this target are instruments such as the CIE Green Data Centre Evaluation Guidelines and the Green Data Centre Technology Catalogue. The Green Data Centre Technology Catalogue endorses 62 technologies across energy and resource efficiency, renewable and distributed energy use, e-waste and hazardous material reduction, and green operations and maintenance.

Policy efforts were reinforced in 2024 with the Guiding Opinions on Vigorously Implementing the Renewable Energy Substitution Initiative, which links clean energy deployment to new infrastructure planning. For data centers, the plan promotes siting in national hub nodes with cold water resources, upgrading older and smaller facilities with greener technologies, and steadily increasing the share of renewable energy in new projects. While the standards-led model offers detailed technical direction, China's regulatory framework for green data centers involves multiple agencies with intersecting mandates, which can make consistent enforcement and cross-jurisdictional coordination more complex.

Table 8: Digital Infrastructure Policies

Policy Name	Date	Detail	References
National AI Industry Investment Fund	2025	Allocates 8.2 billion to support AI start-ups and innovation projects across China, with a focus on strengthening domestic AI capabilities.	Link
Guiding Opinions on Vigorously Implementing the Renewable Energy Substitution Initiative ("New Renewable Energy Plan")	2024	Outlines measures to integrate renewable energy into new infrastructure, including data centers. Promotes facility siting in national hub nodes with cold water resources, green technology upgrades for older and smaller sites, and a gradual annual increase in the share of renewable energy used in newly built facilities.	Link
Guidelines for the Construction of a National Comprehensive Standardization System for the AI Industry	2024	Introduced by the Cyberspace Administration of China (CAC), the guidelines call for the creation of more than 50 national and industry standards to harmonize technical and operational practices in the AI sector.	Link
Guiding Opinions on Accelerating the Construction of a National Integrated Computing Power Network	2023	Sets a target for 60% of new computing capacity to be located within designated hub nodes by 2025, reinforcing the "East Data, West Compute" layout.	Link
Action Plan for the High-Quality Development of Computing Power Infrastructure	2023	Establishes a goal of reaching 300 EFLOPS of total computing capacity by 2025, including 105 EFLOPS allocated specifically for AI workloads.	Link
East Data, West Compute / National Integrated Big Data Center Collaborative Innovation System Implementation Plan	2021	Seeks to balance computing resources by directing real-time workloads to hubs in eastern China and non–time-sensitive workloads to large-scale data centers in western provinces, which benefit from lower energy costs and greater renewable energy availability.	Link
New Generation AI Development Plan	2017	Sets out a three-phase national roadmap to achieve global leadership in AI by 2030, with milestones in 2020, 2025, and 2030. Emphasizes sector-wide integration of AI technologies to support economic and industrial development.	Link

INDONESIA

INTRODUCTION

Indonesia's data center sector is expanding at pace, shaped by strong digital demand, strategic geography, and targeted policy incentives. The economy has more than 80 data center facilities. By 2030, investment in Indonesia's data center market is expected to hit $3.8 billion. Internet penetration of 77% provides a substantial domestic user base, while proximity to Singapore and other major Asian markets strengthens regional integration potential. On the demand side, data center growth is driven by Indonesia's booming digital economy — including native companies such as Gojek, Bukalapak, and Traveloka — alongside global cloud service providers and Chinese hyperscale platforms. This is reinforced by the adoption of AI-driven services, which are accelerating requirements for high-capacity, low-latency infrastructure. Sustained foreign investment — most notably Microsoft's $1.7-billion commitment for cloud, data center, and AI infrastructure — signals growing confidence in Indonesia's long-term digital trajectory.

Data centers in Indonesia remain concentrated in the Greater Jakarta Area, home to more than 35 facilities serving banking, financial services, and insurance industries. While Jakarta benefits from dense connectivity and a skilled workforce, escalating land costs and limited space are pushing hyperscale projects to suburban zones such as Bekasi, Karawang, and Bogor, where power availability and lower land prices allow for higher-capacity builds. The government is also facilitating a strategic shift toward secondary hubs, particularly Batam, which is just 20 km from Singapore. Batam is designated as a Special Economic Zone with a 10-year, 100% corporate tax reduction for qualifying investment. This has attracted approximately 300 MW of committed development from international and domestic operators, supported by submarine cable connectivity and renewable energy resources.

Key Indicators

DIGITAL READINESS INDEX: SCORE 42/100

BURN-TO-EARN INDEX: 274.3 GCO2E PER $ - RANK 109/26

ELECTRICITY MIX: COAL 69%, NATURAL GAS 12.9%, HYDRO 5.7%, BIOFUELS 6.1%

LOCAL DIGITAL INFRASTRUCTURE LANDSCAPE

> **Indonesia's fossil fuel–heavy grid is shifting toward renewables**

Indonesia is advancing its energy transition with a 2060 Net Zero Emissions target and the planned relocation of its capital to Nusantara, opening opportunities for modern, sustainable infrastructure. Yet its electricity mix remains dominated by fossil fuels, with coal at 69% and natural gas at roughly 13%. By 2025, the economy is projected to shift from a net LNG exporter to a net importer as domestic demand outpaces supply and reserves decline, potentially affecting energy costs and availability for power-intensive sectors such as data centers.

The renewable share is still small despite over 333 GW of potential across 1,500+ sites. Indonesia also holds the world's largest geothermal reserves — 27.79 GW — but only about 5% is currently utilized, representing a vast untapped resource for powering sustainable data center growth.

Perusahaan Listrik Negara (PLN), Indonesia's state-owned electricity utility and the sole national grid operator, has set out in its 2025–2034 Electricity Supply Business Plan (RUPTL) an ambitious target to add 39.5 GW of new generation capacity, with 76% expected from renewable sources, marking a significant pivot toward diversifying the economy's heavily fossil fuel–dependent power mix

Floating solar is emerging as a strategic renewable solution in Indonesia

Floating solar is rapidly gaining momentum in Indonesia as a dual solution to the economy's renewable energy expansion and land-use constraints, particularly in fast-growing data center hubs like Batam. Batam has over 20 mapped reservoir and retention pond sites capable of hosting large-scale floating solar power systems. Beyond generating clean power, these installations reduce reservoir evaporation—an important advantage for drought-prone provinces. The flagship 1,600-hectare Duriangkang Floating Solar Plant, developed by EDP Renewables, will have a capacity of 2.2 GWp, making it the largest floating solar project in the world by surface area. The project has the potential to supply both local demand and exports to Singapore via subsea cables. Complementing this, PLN and PT TBS Energi Utama are also advancing the 35 MWac Tembesi floating solar project, expected online by late 2025, which serves as a replicable model for smaller-scale, modular deployments on utility-managed reservoirs.

Indonesia's evolving regulatory framework is enhancing flexible power supply for data centers

Indonesia's regulatory framework is evolving to support a more reliable and flexible power supply for data centers. Data center operators in industrial estates can tap a dedicated power supply, as many estate developers operate or partner with Independent Power Producers (IPPs), giving them more flexibility in sourcing electricity. Facilities outside industrial estates rely primarily on PLN. Renewable Energy Certificates mechanisms are also currently available, through PLN's Green as a Service ("GEAS") Renewable Energy Certificate program. PPAs have also gained traction. In 2022, AWS signed Indonesia's first offsite PPA with PLN to source power from four solar projects in Bali and Java, signaling new opportunities for data centers to secure dedicated green energy. Though still under discussion, the draft Renewable Energy Bill could expand clean energy options for data centers by enabling power wheeling, allowing renewable producers to sell directly to consumers through PLN's transmission network.

Best practices

FLEXIBLE POWER SUPPLY AND GREEN ENERGY PROCUREMENT REFORMS:

Indonesia is reforming its energy policies to enable more flexible power supply and greener procurement for data centers, expanding access to RECs and PPAs while partnering with Independent Power Producers (IPPs) to enhance grid reliability and diversify sourcing beyond the state utility.

CROSS-BORDER RENEWABLE ENERGY CORRIDOR:

Indonesia, especially Batam region, is being positioned as the anchor of a clean energy export hub under the 2024 Indonesia–Singapore Green Electricity Export Partnership, with most of the solar-generated output transmitted to Singapore via high-voltage subsea cables.

FLOATING SOLAR DEVELOPMENT:

Indonesia is advancing large-scale floating solar projects to expand renewable capacity while addressing land constraints, with Batam at the forefront. The 2.2 GWp Duriangkang Floating Solar Plant, set to be the world's largest by surface area, will supply local demand and export surplus power to Singapore via subsea cables.

Batam is emerging as the linchpin of a cross-border renewable energy corridor

Batam is emerging as the anchor point of a new cross-border clean energy corridor linking Indonesia and Singapore. Its close maritime distance to Singapore, well-developed marine infrastructure, and Special Economic Zone incentives make it an ideal staging ground for large-scale renewable projects. Under the Indonesia–Singapore Green Electricity Export Partnership signed in September 2024, the two economies aim to install up to 20 GWp of solar capacity across Indonesia, transmitting most of the output to Singapore via high-voltage subsea cables. One of the most prominent initiatives is Marubeni's Galang Island project with Tuas Power, which will deliver 2.55 GWp of solar power supported by a massive 7 GWp battery system to ensure a steady supply. Construction began in 2024, with the first 600

MW of exports expected in 2027. With Singapore's Energy Market Authority targeting 6 GW of low-carbon electricity imports by 2035 — of which up to 4 GW will come from Indonesia — Batam is set to play a central role in generating, storing, and transmitting clean power, while creating thousands of jobs and strengthening Indonesia's domestic supply chains.

However, scaling cross-border energy in Batam faces practical challenges related to undersea cable capacity and logistics. The island's geography requires undersea infrastructure, which increases both the cost and complexity of energy delivery. Procurement and transportation of equipment are also more resource-intensive due to existing infrastructure constraints and Batam's distance from major supply hubs, adding to operational and capital expenditures.

Indonesia faces significant climate and environmental challenges alongside the rapid expansion of its digital infrastructure.

As the holder of 42% of global nickel reserves, the economy plays a pivotal role in the global supply chain for EV batteries and critical components of digital infrastructure. However, its mining-intensive economy contributes to water stress and pollution, deforestation, and biodiversity loss, compounding the environmental pressures. Indonesia is also one of the most disaster-prone economies, prompting the relocation of its capital from Jakarta to Nusantara in East Kalimantan — a move driven in part by Jakarta's severe land subsidence, earthquake exposure, and flood risks. For emerging data center hubs such as Batam, seismic risks underscore the importance of earthquake-resistant construction and resilient infrastructure planning. These environmental challenges are increasingly prompting policies that prioritize sustainable development and climate resilience.

POLICY LANDSCAPE

Indonesia has built a multi-layered policy framework to position itself as a leading digital economy. In Indonesia, unlike many economies, data centers are classified within the industrial sector rather than telecommunications or IT. This simplifies licensing and enabling operation in industrial estates. In 2018, the Making Indonesia 4.0 Strategy served as a cornerstone for the economy's digital transformation through measures such as 100% foreign ownership and long-term corporate tax holidays. The Digital Indonesia Vision 2045 (2019) sets a national ambition to transform Indonesia into a globally competitive digital economy, as part of the broader Golden Indonesia 2045 agenda. On the AI front, the National AI Strategy (2020) established AI adoption priorities in sectors such as healthcare, smart cities, education, and food security. Over the past decades, the government has also advanced foundational connectivity programs to support digital growth, notably the Palapa Ring Project, which delivered a 36,000km national fiber-optic backbone. In 2025, a new AI Strategy Roadmap is expected to be released, targeting foreign investment, showcasing AI's commercial potential, and guiding developers on infrastructure and AI cluster development, linking AI policy more directly to data center expansion.

Indonesia's digital economy vision is backed by a dual sovereign wealth fund architecture. The Indonesia Investment Authority (INA), launched in 2021, was Indonesia's first sovereign wealth fund, created to co-invest with global partners in strategic infrastructure, including renewable energy and hyperscale data centers. In 2025, Indonesia established its second fund — the Daya Anagata Nusantara Investment Management Agency (Danantara) — with an estimated $900 billion in state assets under management. Danantara is expected to channel long-term capital into strategic sectors, including AI infrastructure.

Demand growth is also policy-driven. The Electronic-Based Government System (SPBE) aims to fully digitize government services by 2025, increasing requirements for secure storage and processing capacity. The regulatory environment is anchored by Government Regulation No. 71/2019 (GR 71/2019), which mandates that public-sector electronic system operators store data domestically while allowing private-sector operators to host data overseas under defined conditions, in alignment with the Personal Data Protection (PDP) Law of 2022.

Sustainability is emerging as a complementary policy pillar. Green data center adoption remains voluntary, guided by standards such as the Indonesia Green Building Council's Greenship Data Centre rating, developed by the

Green Building Council Indonesia and IPUSTAH-ID. The rating offers a tailored rating system for energy-efficient, climate-resilient facilities. A Platinum-rated pilot at the National Data Center demonstrates the potential to cut emissions, reduce costs, and strengthen operational resilience. However, uptake is largely market-driven and voluntary, spurred by multinational clients.

Table 9: Digital Infrastructure Policies

Policy Name	Date	Detail	References
Daya Anagata Nusantara Investment Management Agency (Danantara)	2025	Indonesia's second sovereign wealth fund, with ~$900 billion in state assets; aims to channel long-term capital into digital sectors including AI infrastructure	Danantara
Personal Data Protection (PDP) Law	2022	PDP governs processing, consent, and cross-border transfers, including requiring public-sector data localization.	PDP
Greenship Data Center Certification	2022	National green rating system for energy-efficient, climate-resilient data centers, developed by Green Building Council Indonesia and IPUSTAH-ID; Platinum-rated pilot at National Data Center.	GBC Indonesia
Government Regulation No. 71/2019 (GR 71/2019)	2019	Requires public sector to store data domestically; allows private-sector to store data overseas under condition.	No 71/2019
Palapa Ring Project	Completed in 2019	36,000 km national fiber-optic backbone connecting all major islands, critical for broadband coverage and nationwide digital infrastructure.	Palapa Ring Project Completion
Digital Indonesia Vision 2045	2019	National strategy to position Indonesia as a globally competitive digital economy as part of the Golden Indonesia 2045 agenda.	Digital Indonesia Vision
Indonesia Investment Authority (INA)	2021	First sovereign wealth fund; co-invests with global partners in strategic infrastructure including renewable energy and hyperscale data centers.	INA
National AI Strategy	2020	Long-term AI roadmap to 2045; focuses on healthcare, smart cities, education/research, food security, and bureaucratic reform; supported by pillars on ethics, talent, infrastructure, and innovation.	National AI Strategy
Electronic-Based Government System (SPBE)	2018	Government program to fully digitize public services by 2025, increasing demand for secure data storage and processing.	SPBE
Making Indonesia 4.0 Strategy	2018	Industrial transformation strategy driving digital modernization; offers 100% foreign ownership, up to 20-year corporate tax holidays, and other incentives.	Making Indonesia 4.0 Strategy

KOREA

INTRODUCTION

Korea has been a digital infrastructure pioneer, launching the world's first commercial 5G network in 2019, which solidified its reputation as a leader in high-speed connectivity. By December 2024, the economy's data center power capacity had reached 1.9 GW, with a staggering 90% concentrated in the Greater Seoul Area (Seoul, Incheon, and Gyeonggi). Strong interest from institutional investors and global tech companies is expected to drive capacity growth to 4.8 GW by 2028, a 26.4% CAGR, with 4.1 GW likely concentrated in the Greater Seoul Area alone, where 40+ new facilities are planned. The demand surge is fueled by global hyperscalers such as AWS, Microsoft, and Google, all of which operate Korean cloud regions within the Greater Seoul Area. Local telcos (SKT, KT, LG U+) and Korean conglomerates (Samsung, LG, Kakao, Naver), as well as global data center operators (Equinix, Digital Edge), are also rapidly scaling up capacity.

However, Korea's digital infrastructure momentum is challenged by an underlying imbalance, as power grid bottlenecks and regional equity issues intensify from the overconcentration of data centers in the Greater Seoul Area. At the same time, Korea's digital economy, while significant, records high greenhouse gas intensity with a Burn-to-Earn Index of 99.4 gCO_2e per dollar (ranked 54th of 126), which implies structural challenges in the economy's growth model and energy mix.

Key Indicators

DIGITAL READINESS INDEX: SCORE 62/100

BURN-TO-EARN INDEX: 99.4 GCO2E PER $ - RANK 54/126

ELECTRICITY MIX: COAL 34%, NUCLEAR 30%, NATURAL GAS 26.3%, SOLAR 4.9%, HYDRO 1.2%

LOCAL DIGITAL INFRASTRUCTURE LANDSCAPE

The overconcentration of data centers in the Greater Seoul Area is creating growing pressure on the grid system and slowing permitting processes. Some projects are already facing risks of delays due to power supply constraints and pushback from local residents. While the pipeline of planned capacity exceeds 4 GW, there is a real possibility that not all of it will be delivered on schedule. The risks became clear in October 2022, when a fire at a Pangyo data center (Gyeonggi Province - Greater Seoul Area) disrupted critical services and briefly paralyzed daily life. In response, data center decentralization has become a national agenda since 2023, with government efforts now restricting power supply to the metropolitan area and encouraging regional development.

Grid competition is intensifying in the Greater Seoul Area.

As of early 2024, more than half of the 33 licensed data center projects in Greater Seoul were facing delays, driven by permitting issues, resident complaints, and the power demand collision with Korea's semiconductor industry. Korea Electric Power Corporation (KEPCO), the state power utility, is now taking 12 months (up from 2–3 months) to confirm power supply and has introduced new restrictions on additional capacity in Seoul. In 2024, the Korean government introduced the Special Act on Distributed Energy to incentivize data center development outside the capital region, pairing it with tax benefits to ease regional imbalances. As a result, new clusters are emerging beyond the Greater Seoul Area, including Ulsan, Busan, and Saemangeum/Gunsan. For example, in Ulsan, SK Telecom and AWS are building a ~100 MW facility (scaling to 1 GW by 2029), positioning Ulsan as a new high-performance computing hub.

However, the impact of such policies has been limited, as most large-scale projects continue to cluster in the Greater Seoul Area, where demand and infrastructure are concentrated. These challenges are compounded by delays in power grid investment, driven by public opposition to transmission lines and stricter system impact assessments, which make the government's data center decentralization policy difficult to realize without parallel expansion of transmission and distribution capacity.

At the same time, Korea's push to expand AI and digital infrastructure also overlaps with its aggressive semiconductor industry strategy, most notably the Yongin Semiconductor Cluster in Gyeonggi Province (Greater Seoul Area). The cluster is expected to consume up to 10 GW of power. This demand competes directly with data centers for scarce grid capacity, intensifying pressure on an already overburdened Greater Seoul Area.

As grid capacity tightens, especially in the Greater Seoul Area, securing electricity has become a strategic move for developers. Since 2020, however, this has given rise to a troubling trend: speculative power applications. Many developers have preemptively reserved electricity capacity not to build data centers, but to raise land value or secure grid access for future resale. Between January 2020 and March 2023, nearly 68% of all data center power applications were found to be from came from speculators rather than genuine developers. For example, one individual submitted applications for 28 different sites. It is reported that 33 approved projects never progressed to contract signing after one year. This artificial demand has distorted energy planning, blocked timely access for legitimate operators, and increased the risk of overbuilding transmission and substation infrastructure

Best practices

FOCUS ON HARDWARE-LEVEL ENERGY EFFICIENCY INNOVATIONS:

Korea's emphasis on chip-level energy efficiency — such as NPUs and custom AI processors — demonstrates an alternative path to sustainable digital infrastructure. Government-backed R&D and micro-cloud initiatives help bridge the gap between innovation and infrastructure deployment.

PUBLIC-PRIVATE GOVERNANCE PLATFORMS FOR DIGITAL INFRASTRUCTURE:

The creation of the National AI Committee (July 2024), chaired by the President and composed of stakeholders from across government and industry, exemplifies inclusive digital governance. This committee coordinates policies related to AI, infrastructure, and data center expansion, a model of integrated, high-level digital policymaking.

Grid pressures are compounded by a carbon-heavy power mix.

As of 2023, fossil fuels still dominate Korea's electricity mix: coal (34%), nuclear (30%), and natural gas (26.3%) together make up over 90% of generation, while solar sits below 5%. Korea's electricity mix generates a substantial carbon footprint for data centers. Policy priorities have also swung between renewables and renewed investments in nuclear and fossil fuels under different administrations in recent years, reshaping the mix and complicating efforts to establish a stable, long-term green strategy. With renewable energy's share still extremely low, the economy's heavy reliance on fossil fuels and LNG-based generation conflicts with carbon neutrality goals, underscoring the structural limitations facing Korea's data center sector.

Korea's 11th Basic Plan for Long Term Electricity Supply and Demand (2024–38) sets out a pathway to address these challenges. The policy sets a generation capacity target of approximately 157.8 GW by 2038 to meet surging demand from AI, semiconductors, data centers, and industrial electrification. Capacity additions include two large-scale nuclear reactors, plus one SMR (0.7 GW) by mid-2030s, boosting nuclear output to 35% of the energy mix by 2038. The Plan quadruples renewables capacity to ~121.9 GW by 2038 from ~30 GW in 2023, raising the share of renewables in the energy mix to ~29%. Combined nuclear and renewables aim for ~70% by 2038.

In the near term, natural gas, especially LNG, is filling the gap. LNG is gaining traction due to its dispatchability and lower emissions profile compared to coal. Hybrid solutions such as hydro-LNG co-firing are being piloted to enhance grid flexibility and reliability for high-load facilities. Yet, Korea still faces constraints in scaling LNG use for future data center clusters due to current infrastructure bottlenecks, as existing terminals and pipelines lack sufficient capacity without major upgrades.

Korea is exploring a strategy to move future data center developments closer to generation sites to reduce dependence on long-distance transmission. For example, Jeju Island is being considered as one emerging option for wind-powered data centers, utilizing local wind resources and setting a precedent for distributed green deployment. Still, renewable integration remains limited, and new transmission lines face political and environmental delays.

Nuclear will remain central in Korea's digital infrastructure strategy

Korea's nuclear is directly tied to growing electricity demand from data centers, electric vehicles, and semiconductors. Today, 26 reactors provide over 25% of the nation's electricity, with two other units under construction and additional capacity in the pipeline.

However, nuclear policy direction has shifted significantly under different administrations. From 2022 to 2025, under the previous administration, the government reaffirmed nuclear as a cornerstone of long-term stability and resilience. In July 2022, Korea's Ministry of Trade, Industry and Energy (MOTIE) released its New Energy Policy Direction, aiming to raise nuclear energy's share to at least 30% by 2030 as part of a broader effort to restructure the electricity market for long-term stability and resilience. The government's response included not only new full-scale reactors but also a targeted push for SMRs, which were gaining momentum as a decentralized solution. However, since 2025, under the new administration, policy emphasis has shifted toward accelerating renewable energy deployment, maintaining cautious oversight of nuclear development, and reassessing the role of SMRs within Korea's long-term energy strategy. Key debates persist around land-use efficiency compared to full-scale reactors, and commercial feasibility within a 7–10 year horizon.

Korea's energy efficiency strategy for data centers emphasizes advanced hardware innovation

Korea's approach to energy efficiency in digital infrastructure has leaned heavily toward hardware innovation rather than binding facility-level efficiency mandates. While frameworks like the Zero-Energy Building (ZEB) Grade 5 code exist, their application to data centers remains non-binding, and incentives for green infrastructure are still evolving, largely shaped by private certification schemes. Instead, national efforts have centered on semiconductor R&D — particularly in designing custom chips for AI that promise high performance with lower power consumption.

There is growing consensus among policymakers and industry that chip-level breakthroughs — such as Neural Processing Units (NPUs) and AI edge devices — hold the greatest promise for reducing energy intensity. Korean developers are increasingly active in translating these innovations into practice, with growing interest in micro-cloud and edge cloud architectures as distributed alternatives to centralized hyperscale models, particularly in densely populated or grid-constrained regions. The 2025 Domestic AI Semiconductor-Based Micro Data Center Dissemination (R&D) Project, backed by the Ministry of Science and ICT, is a flagship initiative designed to accelerate this trend. Yet, the project is still at an experimental stage. The practical deployment of micro data centers is uncertain, and questions remain about their economic viability and operational efficiency relative to large-scale facilities.

The Ministry of Science and ICT has also been actively promoting processing-in-memory (PIM) technology as part of its national AI semiconductor strategy. PIM can reduce energy consumption by allowing data to be processed directly within memory, minimizing the need for constant data transfer between memory and processors. However, the strategy faces structural challenges. Domestic AI chips lack the equivalent of NVIDIA's CUDA (Compute Unified Device Architecture), the dominant global ecosystem that enables developers to harness GPUs for AI and high-performance computing through a mature set of tools, libraries, and community support. Without a comparable platform, Korean chips face steep barriers to adoption. These challenges are compounded by limited access to investment capital compared to global competitors and a market that continues to prioritize speed and performance over energy efficiency. Recently, SK Telecom's AI chip unit SAPEON merged with Rebellion, reflecting consolidation pressures in the domestic AI semiconductor sector and signaling a search for a clearer strategic direction.

Grid flexibility is emerging as a key solution for Korea's data centers

Grid flexibility for data centers is encouraged by the Korean government. The government has included data centers in demand management programs to ensure grid stability during peak summer/winter loads. Korea's demand-side management strategy aims to reduce projected electricity generation by 11% by 2030, as outlined in the 10th Basic Plan for Electricity Supply and Demand (2022–2036). As part of this effort, the government is rolling out smart grid infrastructure to enable real-time energy management. In the data center sector, Samsung SDS Suwon Center joined Enel X Korea's Virtual Power Plant, adopting a customized demand response program to reduce peak load, stabilize operations, and support grid flexibility. This model demonstrates how large-scale data centers can balance reliability with sustainability by optimizing power use in coordination with the broader electricity system.

POLICY LANDSCAPE

Data center development in Korea is currently governed across three distinct ministries, each with its own mandate: the Ministry of Science and ICT (MSIT) oversees digital infrastructure and AI hardware; the Ministry of Trade, Industry and Energy (MOTIE) manages electricity supply and energy policy; and the Ministry of Environment (MOE) handles environmental assessments and sustainability. Many digital infrastructure areas, such as environmental review rules, building codes, and energy efficiency compliance, are managed by multiple ministries, sometimes resulting in overlapping authority and fragmented enforcement.

Korea's digital infrastructure policy landscape has become increasingly strategic, evolving from visionary ambition to regulatory precision. The foundation was laid in 2020 with the launch of the National Strategy for AI, which positioned artificial intelligence as a transformative force reshaping society, economy, and governance. The strategy set clear national targets: to rank third globally in digital competitiveness by 2030, to generate $350 billion (₩455 trillion) in economic surplus, and to elevate Korea into the top 10 economies in quality of life through digital innovation. While the strategy set ambitious targets, implementation had advanced at a more gradual pace, particularly in adoption rates and capital mobilization.

In 2024, Korea elevated its ambitions further. The National AI Strategy Policy Directions (September 2024), issued under the newly formed National AI Committee, reaffirmed Korea's intent to become a global top-three AI power by 2030. It aims to mobilize 49 billion USD (₩65 trillion) in private investment, expand national GPU infrastructure, and achieve 70% AI adoption in industry and 95% in the public sector. This momentum culminated in the Framework Act on Artificial Intelligence (December 2024)—also known as the AI Basic Act—which made Korea the first economy in Asia-Pacific and second globally (after the EU) to enact comprehensive AI legislation. The law establishes a dual mandate: to promote AI innovation through incentives and infrastructure, and to regulate high-risk applications through safeguards, trust frameworks, and ethical oversight. Still, with the Act not set to take full effect until 2026, its real impact on industry practice remains untested.

When a 2022 data center fire in Pangyo disrupted major services such as messaging apps and banking, the government responded with the Development Plan for Digital Platforms (December 2022), which expanded disaster resilience obligations for both data center operators (DCOs) and value-added service providers (VSPs). New mandates included decentralization of facilities, server redundancy, and mitigation of user harm. This was followed by a broader decentralization and grid management strategy. In March 2023, revisions to the Electric Utility Act gave KEPCO legal authority to reject electricity supply applications from large-scale users, including data centers, in congested areas. Simultaneously, the government launched a relocation policy encouraging data center development in underutilized regions like Jeju, Honam, and Gangwon, backed by tax incentives and flexible permitting. To support the shift toward localized energy systems, the Dispersed Energy Promotion Special Act was passed in June 2023. This law enables direct power trading, virtual power plants (VPPs), and mandates minimum distributed energy usage for large facilities. However, instances of 'grid speculation', where developers reserve capacity without moving forward with construction, have highlighted challenges in regulatory oversight. By July 2023, KEPCO introduced stricter supply application rules to curb speculative behavior by developers hoarding grid capacity without real construction intent.

At the time of writing, Korea embarked on a unique, full-stack approach to "sovereign AI," tasking its largest tech companies and startups to build a national foundational model using predominantly domestic technologies, from semiconductors and memory chips to cloud infrastructure and AI software. Backed by the Ministry of Science and ICT, five consortia led by firms such as SK Telecom, LG, and Naver are developing open-source models, leveraging Korea's global dominance in high-bandwidth memory (SK Hynix, Samsung), AI-focused chip innovation (Rebellions), and data center expansion.

Table 10: Digital Infrastructure Policies

Policy Name	Date	Detail	References
AI Basic Act	Dec 2024	Asia-Pacific's first comprehensive AI legislation, combining promotional tools with regulatory measures (e.g., risk management and ethics) to establish a trusted AI ecosystem.	Ministry of Science and ICT
National AI Strategy Policy Directions	Sep 2024	National strategy to become a global top-three AI power by 2030, mobilizing $49B in private investment, expanding GPU capacity 15x, and targeting 70% AI adoption in industry and 95% in government.	Ministry of Science and ICT
National Artificial Intelligence Committee (Presidential Decree)	Jul 2024	Establishes a presidential-level committee to coordinate national AI policy, R&D, regulation, and infrastructure, including data centers. Chaired by the President and comprised of 45 public and private leaders.	Ministry of Science and ICT
Artificial Intelligence Safety Institute (AISI)	Ongoing	National institute responsible for evaluating AI safety risks and researching mitigation technologies to inform future AI regulation.	AISI
Strategic Roadmap for AI and Advanced Biotechnology	Oct 2023	Aims to develop hyper-efficient AI and biotech systems, targeting over 50% energy savings, indirectly promoting greener, high-performance data centers.	Ministry of Science and ICT
KEPCO's Revised Electricity Supply Terms	Jul 2023	Introduced stricter vetting for grid applications to curb speculative land deals and phantom demand from data center developers. Includes cancellation rights for inactive projects.	Policy Study on the Location of Data Center Establishment in Korea
Dispersed Energy Promotion Special Act	Jun 2023 (Effective mid-2024)	Supports decentralized energy systems via virtual power plants (VPPs), direct power trading, and mandated distributed energy use. Large facilities such as data centers face new grid and environmental review thresholds.	
Electric Utility Act (Amended)	Mar 2023	Grants KEPCO authority to deny electricity access to high-load users like data centers if they threaten grid reliability or violate voltage standards.	
Measures to Ease Data Center Overconcentration	Mar 2023	The first national plan targeted the overconcentration of data centers in the Greater Seoul Area. The strategy restricts grid access for new high-load facilities (5MW+), mandates power grid impact assessments and incentivizes relocation through regulatory relief and electricity discounts in non-metropolitan regions (e.g., Jeju, Honam). The policy also promotes co-location with waste heat or surplus energy zones.	
Special Act on the Fostering of National Strategic Technology	Mar 2023	Provides legal foundation to prioritize and support strategic technologies (e.g., AI, semiconductors, quantum). While not data center-specific, it indirectly supports data centers through tax breaks, fast-track approvals, and R&D for AI infrastructure.	Link
Development Plan for Digital Platforms	Dec 2022	Post-Pangyo fire strategy to improve digital infrastructure resilience. Expands disaster obligations to data center operators and VAS providers, mandates redundancy, regular inspections, and stricter compliance under amended telecom laws.	MSIT, KCC, PIPC, KFTC
National Strategy for AI	Dec 2019	The strategy aims to position Korea as the third most digitally competitive economy by 2030, generate $350 billion in economic surplus through AI, and elevate the economy into the global top ten for quality of life.	OECD

MALAYSIA

Key Indicators

DIGITAL READINESS INDEX: SCORE 59/100

BURN-TO-EARN INDEX: 155.9 GCO2E PER $ - RANK 77/126

ELECTRICITY MIX: COAL 47%, NATURAL GAS 34%, HYDRO 16.6%, SOLAR 1.2%

Introduction

Malaysia is rapidly establishing itself as a leading digital infrastructure hub in the APEC region, attracting an estimated $43.6 billion in data center investments between 2021 and 2024. In Malaysia, the current data center capacity is 689 MW, with an additional 520 MW under construction. Electricity demand from the sector is projected to exceed 5 GW by 2035, representing over 11% of Malaysia's anticipated national capacity, a scale that highlights the sector's growing strategic importance.

Two regional clusters are driving this momentum. The Klang Valley — spanning Kuala Lumpur to Port Klang, including Cyberjaya — remains the administrative nucleus of Malaysia's data center ecosystem. In contrast, Johor, particularly the Nusajaya and Sedenak Tech Parks, has emerged as the dominant hyperscale corridor. As of 2024, Johor accounts for nearly 80% of Malaysia's live IT capacity, having grown from just 10 MW in 2021 to 1.3 GW, with projections to reach 2.7 GW by 2027. Johor offers significant advantages over other Southeast Asian regions, including lower latency, easier physical access, and more scalable power options, making it an ideal location for large-scale data center operations without the limitations of subsea cable infrastructure. Major global operators — including AWS, Microsoft, AirTrunk, Equinix, and YTL — are expanding campuses, many exceeding 100 MW, underscoring investor confidence in Malaysia's scalability.

Malaysia's data center growth is the result of a strategic convergence of demand-side trends, supply-side advantages, and regional developments.

Malaysia's data center demand is shaped by three dynamic and rapidly growing sectors. E-commerce and streaming companies, such as ByteDance and Shopee/Sea Group, were among the first movers in Johor, capitalizing on the region's cost-effective infrastructure and proximity to Singapore. These early entrants established a strong foothold, driving the need for scalable data infrastructure to support their high-volume operations. The next wave of demand comes from hyperscale cloud providers such as AWS, Google, Microsoft, and Oracle, whose predictable, long-term growth strategies require vast, redundant, and strategically located data centers. Finally, the AI-driven segment, the newest and fastest-growing, has introduced a more volatile demand pattern. This surge in AI-driven demand has been a significant factor in the recent explosion of data center projects in Johor, particularly over the past 12–24 months.

Malaysia's competitive edge lies in its strong supply-side fundamentals. Malaysia offers abundant and cost-effective land for construction, reliable grid infrastructure, and low-cost electricity, positioning it ahead of many regional peers. It also boasts robust connectivity, supported by nearly 30 submarine cable systems and 14 cable landing stations, and a government target of 66 Internet Exchange Point (IXP) providers by 2025, up from 12. Crucially, human capital complements physical infrastructure. Malaysia continues to invest in ICT education, producing a growing pool of skilled professionals in data management, cybersecurity, and network engineering — a vital component of long-term sector resilience.

A significant catalyst in Malaysia's rise has been the policy shift in neighboring Singapore. The Singapore 2019 moratorium on new data centers, followed by ongoing resource restrictions, redirected hyperscale interest to nearby alternatives. Malaysia quickly emerged as the logical overflow destination. Electricity in Malaysia is roughly three times more affordable than in Singapore, offering a competitive cost advantage. Johor, in particular, now serves as a functional extension of Singapore's digital footprint. The Johor-Singapore Special Economic Zone (JS-SEZ), launched in 2025, strategically focuses on data center development, leveraging Malaysia's proximity to Singapore. The zone aims to attract global data center operators, offering cutting-edge facilities in Sedenak Tech Park and Nusajaya Tech Park. The Invest Malaysia Facilitation Centre Johor (IMFC-J) streamlines the investment process, ensuring a business-friendly environment. As part of its vision, the JS-SEZ will deepen cross-border integration, enhancing Malaysia's role as a regional leader in the data economy and making it a top destination for hyperscale data centers in Southeast Asia.

LOCAL DIGITAL INFRASTRUCTURE LANDSCAPE

Malaysia's data center sector is growing rapidly, but solar power—its most scalable renewable option—faces major structural constraints

Malaysia's rapid rise as a regional digital infrastructure hub is unfolding within the constraints of a heavily fossil-fueled energy mix. As of 2022, coal and natural gas accounted for over 80% of national electricity generation — 46.8% from coal and 34.3% from gas. As a result, most data centers continue to rely on a carbon-intensive grid. However, Malaysia has made progress towards cleaner energy, with gas and renewables contributing more to the mix. In 2022, Malaysia's renewable energy share in electricity generation reached nearly 20%, marking a threefold increase from 2010 levels. Despite the establishment of a nuclear agency and periodic reviews of nuclear energy options, Malaysia has yet to develop any nuclear power plants, primarily due to policy decisions, public concerns, and the absence of the necessary infrastructure.

The Malaysian government has liberalized access to green energy. The government has introduced initiatives like the Corporate Green Power Programme (CGPP) to encourage clean energy procurement and open the national grid to direct corporate PPAs, enabling data center operators to source renewable electricity more flexibly. Private sector momentum is building: YTL Data Centers, for instance, is developing 600 MW of solar capacity to support its hyperscale park. Tech players like Citaglobal and Genetec are also advancing domestic battery storage solutions. Their joint project, MYBESS — Malaysia's first homegrown battery energy storage system — has already tendered 700 MWh and expects to scale to 1.5 GWh in 2024.

Best practices

STRATEGIC GOVERNMENT SUPPORT FOR DIGITAL INFRASTRUCTURE

Malaysia offers strong incentives to grow its digital infrastructure, including tax exemptions and investment allowances for data center providers. Fast-track programs streamline the approval process, reducing development timelines. At the same time, Malaysia is advancing its renewable energy targets, facilitating corporate PPAs to promote the use of green energy in data centers and other industries.

INTEGRATION OF RENEWABLE ENERGY SOLUTIONS

Malaysia is working to transition its energy mix towards more renewable sources with targets such as 40% renewable energy capacity by 2035. The government's initiatives like the Corporate Green Power Programme (CGPP) and the Corporate Renewable Energy Supply Scheme (CRESS) facilitate corporate PPAs enabling businesses to access green energy.

LEVERAGING COMPETITIVE REGIONAL ADVANTAGES AND CROSS-BORDER INTEGRATION

Malaysia's strategic location near Singapore, combined with lower latency, easier physical access, and scalable land and power options, makes it an ideal hub for data centers. Digital clusters like Johor's Nusajaya and Sedenak Tech Parks demonstrate how targeted infrastructure investments attract global operators. The Singapore-Johor Special Economic Zone (JS-SEZ) highlights Malaysia's focus on cross-border integration, aligning policies and infrastructure development with neighboring economies.

However, solar deployment for data centers in Malaysia is facing significant constraints. Meeting the energy demand from Johor's committed data center projects would require 59,000 acres, which is over 60% of the state's total landmass (ISEAS). While many new data centers are designed with solar-ready rooftops, their actual contribution to energy needs remains minimal. For example, AirTrunk's 150 MW JHB1 facility in Johor generates just 1 MW from rooftop solar, covering only 2% of its phase-one load. With Malaysia aiming for more renewable share, solar alone will not be sufficient to meet the energy needs of the growing data center sector due to land constraints.

Malaysia's accelerating data center growth is challenging grid planning and putting upward pressure on electricity prices.

Malaysia's electricity grid is nearing a critical inflection point as data center development accelerates at an unprecedented pace. By the end of 2024, 38 data center projects had secured Electricity Supply Agreements (ESAs) — formal contracts that allow operators to reserve power capacity ahead of operations — totaling a maximum demand of 5.9 GW. This represents more than 40% of the total contracted capacity of Tenaga Nasional Berhad (TNB), Malaysia's national electricity utility.

Although actual power usage from operational data centers remains relatively modest — estimated at 400–500 MW as of December 2024 — these figures reflect only the initial phases of deployment. As more facilities transition from planning to full operation, the cumulative energy demand could place significant strain on the grid.

Recognizing long-term infrastructure demands, in 2022, TNB announced plans to allocate approximately $4.5 billion annually in capital expenditure through 2050 to support energy transition and grid modernization. TNB also started a 14.2% increase in the base electricity price for Peninsular Malaysia by July 2025. While multiple factors contribute to rising costs, the growing capital and operational needs of maintaining a grid that supports 24/7 data center loads are increasingly relevant. Against this backdrop, studies have raised questions about whether Malaysia's data center expansion is proceeding in full alignment with national energy planning and grid capacity.

Malaysia is facing mounting water stress, with surging data center demand set to strain already limited water resources

Malaysia ranks highest among its regional peers in water stress, raising red flags as data center development accelerates. While Johor currently maintains a 16.9% water reserve margin (the difference between production capacity and usage), the mounting competition between industrial and residential users is hard to ignore. With 101 data center applications filed as of 2024, totaling over 808 million liters of daily water demand, the government is racing to head off a clash between digital growth and water security. In response, the National Water Services Commission (SPAN) will introduce mandatory guidelines by mid-2025 requiring new data centers in Peninsular Malaysia to rely entirely on alternative sources, such as recycled, reclaimed, or rainwater, within three years.

POLICY LANDSCAPE

Malaysia's policy landscape regarding digital infrastructure and sustainable data center development is evolving rapidly, fueled by strategic government initiatives aimed at positioning Malaysia as a leading digital economy. The Malaysia National AI Roadmap, launched in 2021, laid the groundwork for AI-driven infrastructure development, including data centers, by allowing 100% foreign ownership in the ICT and infrastructure sectors. This initiative is further supported by the MyDIGITAL Blueprint, which envisions Malaysia as a high-income, digitally driven nation, recognizing data centers as a critical component of its digital economy.

The government has introduced a series of incentives to attract both global players and local startups. The Digital Ecosystem Acceleration Scheme (DESAC), launched in 2022, provides substantial support for data center development through tax exemptions and investment allowances. With around $27 billion approved for 21 data center projects, primarily in Johor, the scheme plays a pivotal role in fostering a globally competitive digital ecosystem. Complementing this, initiatives like the Green Lane Pathway and Kulai Fast Lane Program streamline approval processes, significantly reducing the time required for construction permits and electricity supply connections, enhancing Malaysia's appeal as a key data center hub.

Additionally, the Malaysian Digital Economy Corporation (MDEC) has announced plans to enhance inter-data center connectivity in Johor through submarine optic fiber cables, strategically positioning Johor as a high-performance digital infrastructure hub. This development reflects Malaysia's dual focus on ensuring both power and bandwidth availability for its rapidly expanding data center sector.

Malaysia is committed to reducing its GHG emissions intensity by 45% by 2030, relative to its 2005 levels. Additionally, the economy has set a new target to achieve 70% renewable energy (RE) capacity by 2050 (APEC report). In line with its goal, the government has launched initiatives such as the Corporate Green Power Programme (CGPP) and the Corporate Renewable Energy Supply Scheme (CRESS) to facilitate the adoption of renewable energy, particularly solar, through corporate power purchase agreements (PPAs).

In terms of energy efficiency and sustainability, the Ministry of Housing and Local Government (MHLG) and the Malaysian Investment Development Authority (MIDA) have The Data Centre Planning Guidelines and Guidelines for Sustainable Development of Data Centres, emphasizing energy efficiency, carbon reduction, and responsible water usage in data center operations. The Malaysian Technical Standards Forum Bhd has issued a technical code establishing minimum requirements for green data center specifications, with a focus on energy efficiency and reducing carbon footprints. It also sets out best practices for data centers to follow in advancing a sustainable industry. Additionally, the National Investment Council has recommended a PUE target of 1.4 for hyperscale data centers to qualify for tax incentives. Johor's Investment, Trade, and Consumer Affairs Committee has stated that if data center operators do not demonstrate efforts to use green energy, the committee could recommend rejecting the building permission application for their data centers.

While this strong stance underscores Malaysia's commitment to aligning digital infrastructure development with sustainability goals, the guidelines remain recommendations for tax incentives and approval purposes, rather than enforceable standards. As a result, efficiency performance is often driven by internal targets and customer expectations, especially given that power costs are passed through to clients. Data center operators do not automatically pursue certifications such as Malaysia's Green Building Index (GBI), unless specifically requested by clients or contractually obligated (YTL Interview).

Table 11: Digital Infrastructure Policies

Policy Name	Date	Detail	References
Digital infrastructure investment related policies			
Malaysia National AI Roadmap 2021-2025	2021	Launched to attract AI investment, including data centers, with a focus on creating a conducive environment for digital infrastructure. The policy enables up to 100% foreign ownership in ICT and infrastructure sectors, but property acquisitions require state approval. It emphasizes strategic development of data centers to support AI ecosystem growth.	Ministry of Science, Technology and Innovation
MyDIGITAL Blueprint	2021	The MyDIGITAL initiative seeks to transform Malaysia into a high-income, digitally driven nation, positioning it as a regional digital hub. The framework focuses on creating a digital economy that includes the establishment of data centers as critical components.	Ministry of Economy
Malaysia Digital Status (MD)	2021	The MD status offers targeted incentives to data centers, including exemptions from foreign equity restrictions, tax exemptions, and customs duty exemptions for equipment.	Malaysia Digital Economy Corporation
Digital Ecosystem Acceleration Scheme (DESAC)	2022	Coordinated by the Ministry of International Trade and Industry (MIDA) and the Ministry of Communications and Multimedia (MDEC), the Digital Ecosystem Acceleration (DESAC) Scheme, launched in Budget 2022, aims to attract high-quality digital infrastructure projects, particularly data centers, to strengthen Malaysia's digital economy. The scheme targets Digital Technology Providers (DTPs) and Digital Infrastructure Providers (DIPs) like data centers, offering significant incentives such as tax exemptions (0%-10% income tax for up to 10 years for DTPs) and 100% investment tax allowances for DIPs. With $ 26.8 billion approved for 21 data center projects, primarily in Johor, the scheme fosters investment from MNCs, SMEs, and startups, creating a globally competitive digital ecosystem.	Malaysia Investment Development Authority
Green Lane Pathway Initiative	2023	The initiative fast-tracks electricity supply connections for data centers, reducing the typical timeline from 36-48 months to just 12 months. It includes a One-Stop Centre (OSC) for data center investors, providing streamlined support and coordination, ensuring faster implementation of data center projects in Malaysia.	TNB Green Lane
Kulai Fast Lane (KFL) Program	2023	The KFL program accelerates the approval process for business development projects in Johor, particularly for data centers. By reducing the time for construction permits and operating licenses, the program supports faster establishment of data centers in key locations like Johor.	KFL Program
Sustainable data center related policies			
Corporate Green Power Programme (CGPP)	2022	Facilitates virtual power purchase agreements (VPPAs) for companies to procure renewable energy, particularly solar, to meet ESG targets. Promotes the adoption of green energy via corporate PPAs.	Energy Commission of Malaysia
Corporate Renewable Energy Supply Scheme (CRESS)	2024	Launched by the Ministry of Energy Transition and Water Transformation to facilitate corporate PPAs, enabling data centers and other businesses to purchase renewable electricity via open access, promoting green energy adoption.	Ministry of Energy Transition and Water Transformation
Data Centre Planning Guideline	2024	Adopted by the Ministry of Housing and Local Government, it sets standards for sustainable siting, permitting, and operational practices for new data centers, emphasizing ESG principles, energy efficiency, carbon reduction, and responsible water usage.	Ministry of Housing and Local Government
Guidelines for Sustainable Development of Data Centres	2024	Released by Malaysian Investment Development Authority (MIDA) to promote energy efficiency, renewable energy adoption, and water conservation in data center investment, aiming to standardize development and attract sustainable investment.	Malaysian Investment Development Authority
Green Building Index (GBI) Certification	Ongoing	The GBI system rates buildings, including data centers, on sustainability criteria such as energy efficiency, water conservation, and materials.	GBI
PUE Recommendations	2024	The National Investment Council recommends that hyper-scale data centers (over 21 MW) achieve a PUE of 1.4 or lower to be eligible for tax incentives.	National Investment Council

SINGAPORE

INTRODUCTION

Singapore has long stood as a leading data center hub in the APEC region, hosting over 1.4 GW of operational capacity across more than 70 cloud, enterprise, and co-location facilities. These centers manage a diverse range of workloads, from e-commerce platforms and hyperscale cloud services to AI-intensive processing, placing Singapore among the densest data center markets in the world. The economy ranks third globally for connected data centers per capita (nine per 1 million people) and second for concentration, hosting an average of 151 companies per connected facility. This density reflects Singapore's role as a high-value interconnection hub, where rich fiber networks and multiple submarine cable systems enable seamless traffic exchange among tenants and drive regional digital leadership. It serves as a major nexus for submarine cables, with 26 systems landing at three sites. Under the Digital Connectivity Blueprint 2023, the government aims to double landing sites within a decade to enhance resilience.

Singapore's digital ecosystem is reinforced by a deep pool of skilled professionals and its ability to attract top-tier engineering talent from around the world. This ensures world-class operational standards, supported by a secure metropolitan environment that enables high-value digital operations. Businesses also benefit from a stable political climate, transparent regulations, low taxes, and efficient administration, making Singapore one of the most business-friendly jurisdictions in the region.

Sustainability is central to Singapore's long-term data center strategy. The economy is advancing low-carbon growth through renewable energy targets of 2 GWp of solar deployment by 2030, pioneering hydrogen research, and implementing energy-efficiency standards tailored to tropical-climate data centers. These initiatives position Singapore as a regional benchmark for sustainable, high-density, and innovation-driven digital infrastructure.

Key Indicators

DIGITAL READINESS INDEX: SCORE 56/100

BURN-TO-EARN INDEX: 48.6 GCO2E PER $ - RANK 28/126

ELECTRICITY MIX: NATURAL GAS 94%, SOLAR 2.4%, COAL 0.9%

LOCAL DIGITAL INFRASTRUCTURE LANDSCAPE

Digital infrastructure growth must align with sustainability in Singapore

Before 2019, Singapore experienced a surge of investment from global operators such as Digital Realty, Google, NTT, and Meta, driving rapid expansion in its data center capacity. The economy emerged as Southeast Asia's primary hub, hosting over 60% of the region's capacity. Between 2010 and 2015 alone, commercial data center space was projected to grow by 50%, but this swift expansion placed mounting pressure on the national grid. Large-scale facilities consume vast amounts of electricity — data centers in Singapore now account for nearly 7% of total national power demand. In a resource-constrained economy such as Singapore, where water is also heavily used for cooling, the environmental footprint of the sector became a growing policy concern.

In 2019, Singapore adopted a more controlled approach to data center growth, imposing a three-year moratorium on new builds to reassess policy and align the sector with climate and resource management goals. While the pause was lifted in July 2022, it was replaced by a more selective framework, requiring operators to meet high standards for both energy and water efficiency. Only "best-in-class" projects — meeting strict resource efficiency, PUE thresholds, and sustainability targets — are now approved.

Today, the market operates under tight conditions. Vacancy rates often fell below 2%, reflecting both the high demand for data center services and the limited pipeline of new capacity under the stricter regulatory regime. This scarcity is amplified by Singapore's position as one of the most power-constrained data center markets globally, underscoring the challenge of meeting growing digital demand within the economy's finite energy and water resources.

Singapore faces deep structural limits to scaling domestic renewables, making energy diversification, clean energy corridors, and next-generation fuels like hydrogen critical to powering its growing digital economy.

Singapore is a net energy importer, with its electricity mix heavily reliant on imported natural gas (93.7%) and solar accounting for just 2.4%. With constraints in land size, low wind speeds, and a lack of hydro and geothermal resources, Singapore faces structural constraints in expanding domestic renewable generation. Despite these challenges, Singapore is committed to a long-term transition toward a low-carbon energy mix over the next 30–40 years. Under its 2030 Nationally Determined Contribution (NDC), it aims to peak emissions before 2030 and reduce them to around 60 Mt CO_2e by that year, in line with its goal of achieving net-zero emissions by 2050.

To meet rising electricity demand — driven in part by the growth of AI, hyperscale cloud, and other digital infrastructure — while pursuing decarbonization, Singapore is adopting a diversified strategy.

Natural gas is expected to remain a cornerstone of Singapore's energy strategy in the medium term, supporting supply security and price stability. The current LNG terminal operated by Sembcorp and Singapore LNG Corporation can store enough fuel to power Singapore for several months, providing a strategic buffer against global supply disruptions. A second LNG terminal is planned by 2030. To

Best practices

SUSTAINABILITY AS A CORE MARKET CRITERION:

Singapore's sustainability-first framework positions environmental performance as a strategic enabler of digital growth. Through initiatives such as the Data Centre Call for Application, the Green Data Centre Roadmap, and the BCA Platinum Green Mark, new capacity is allocated only to projects with advanced cooling technologies, clean energy integration, and high efficiency standards — such as PUE ≤ 1.3. The Singapore-Asia Taxonomy further directs investment toward certified green assets, ensuring that each new development advances both economic competitiveness and long-term climate commitments.

LEVERAGING RESOURCE CONSTRAINTS TO ACCELERATE INNOVATION:

Singapore's water scarcity has driven the adoption of alternative cooling solutions, including reclaimed water systems, seawater-based cooling, floating data centers, and condensate reuse. By embedding water efficiency targets in certification schemes and aligning regulatory incentives with innovation, resource limitations become a catalyst for sustainable infrastructure advancement.

PUBLIC–PRIVATE PARTNERSHIP IN POLICY DESIGN:

Singapore's data center policies are co-developed with industry, ensuring ambitious targets are technically achievable and broadly supported. The 1.3 PUE benchmark emerged from joint consultation and is embedded in national certification standards, fostering a constructive regulatory environment where government ambition and industry innovation move in tandem.

ADVANCE REGIONAL CLEAN ENERGY CORRIDORS:

Singapore's strategy to import 6 GW of low-carbon electricity by 2035 underpins its ability to meet growing digital demand sustainably. By securing long-term clean energy agreements with Australia, Indonesia, Cambodia, and Viet Nam — and advancing cross-border projects like the Indonesia–Singapore Green Electricity Export Partnership, Johor–Singapore Special Economic Zone — the economy is creating a resilient, diversified energy supply.

reduce emissions, Singapore is integrating carbon capture and storage (CCS), targeting a 10% improvement in gas plant efficiency, and enabling plants to operate with at least 30% hydrogen by volume.

Solar deployment, while constrained by land availability, is advancing rapidly. Singapore reached its 2025 target of 1.5 GWp installed capacity by the end of 2024, one year ahead of schedule, and is on track to achieve at least 2 GWp by 2030. Beyond conventional rooftops, solar is being deployed through innovative applications such as floating systems on water bodies and canopies over canals, offering incremental renewable generation for urban data center operations. Nuclear power is not part of the near-to-medium-term strategy due to land constraints, urban density, and regional sensitivities, with deployment unlikely before 2050.

Hydrogen, particularly in liquid form, is emerging as a growing pillar of Singapore's future energy strategy. Seen as potentially more scalable and geopolitically secure than relying solely on subsea cable imports, liquid hydrogen offers synergies with Singapore's existing LNG expertise. The economy already has the capability to import, store, and re-gasify cryogenic fuels at scale, and similar infrastructure could be adapted for hydrogen using compressed vessels and offshore tanks. However, lacking renewable resources to produce green hydrogen locally, Singapore is a downstream player dependent on imports generated in renewable-rich economies. Given Singapore's emissions cap target, only green hydrogen, proven emissions-free at source, is likely to pass regulatory approval, making grey or blue hydrogen non-viable unless recognized as low-carbon. For data centers, hydrogen adoption depends on a fully developed value chain — producers, logistics providers, and utilities — plus regulatory steps such as Genco licensing (main power generation companies) and land-use approvals. While commercial deployment remains years away, industry players are future-proofing; for example, Keppel Data Centres has partnered with Australia's Woodside Energy to secure a liquid hydrogen supply by 2030 for its facilities.

Singapore is building a regional clean energy corridor to power its growing data center sector.

Singapore is advancing its vision for a regional clean energy corridor, leveraging its position as a stable, high-value market anchor to create market opportunities for renewable energy development across neighboring economies. By committing to long-term clean energy imports, it provides a commercial foundation for regional partners to scale solar and other low-carbon projects, fostering mutually beneficial trade and cooperation.

Singapore has pledged to import 6 GW of low-carbon electricity by 2035, representing about one-third of its future energy needs. To date, it has awarded 7.35 GW in conditional approvals and licences from a diverse range of sources, including Australia, Cambodia, Indonesia, and Viet Nam. Flagship developments include 1.75 GW from Australia's Northern Territory via a 4,300 km subsea cable by Sun Cable, 2.0 GW in conditional licenses from Indonesia, and further conditional approvals for 1.4 GW from Indonesia, 1 GW from Cambodia, and 1.2 GW from Viet Nam (APEC). Specifically, through the Indonesia–Singapore Green Electricity Export Partnership signed in September 2024, Singapore and Indonesia plan to deploy up to 20 GWp of solar capacity across Indonesia, with the bulk of the power delivered to Singapore via high-voltage subsea cables.

Complementing these efforts, Singapore and Malaysia have also signed an agreement to establish the Johor-Singapore Special Economic Zone (JSSEZ) to boost trade, investment, and connectivity. Positioned as a cross-border energy corridor, the JSSEZ links Singapore's high-demand digital economy with Johor's renewable energy potential. Spanning 3,500 km^2, it embeds green energy cooperation, streamlined transport links, and investor incentives to integrate manufacturing, digital infrastructure, and power supply chains, advancing both economies' energy security and decarbonization goals.

Singapore is positioning itself as a global testbed for resource-efficient data centers

While a regional clean energy corridor is technically feasible and strategically important to Singapore's energy strategy, progress hinges on close coordination between governments to navigate sovereignty concerns, energy security priorities, and commercial arrangements. Many projects require transmission cables to cross multiple jurisdictions — such as routes from Indonesia to Singapore via Malaysian waters or from Laos through Thailand and Malaysia — making trust-building diplomacy and sustained engagement essential to success.

Singapore is partnering with industry to drive next-generation efficiency in data centers — advancing hardware and software optimization, cooling innovation, and offshore deployment. Its efficiency roadmap sets a clear target: all data centers must achieve a PUE of 1.3 or better within the next decade, a benchmark embedded in the BCA Platinum Green Mark rating and reinforced through certification incentives.

Central to this push is the Tropical Data Centre Standard (SS 697:2023), launched in 2023 after more than five years of joint government–industry trials. The standard overturns the belief that high-performance facilities in hot, humid climates require energy-intensive cooling. Trials confirmed that modern IT equipment can operate reliably at 26–28°C, and in some cases up to 40°C with high humidity, without material degradation—provided airborne contaminants are controlled. Publicly available with empirical data, the standard offers a replicable benchmark for other hot-climate economies such as Vietnam and Thailand.

Singapore operators also use advanced cooling to handle the extreme thermal loads of high-density AI workloads, integrating solutions such as immersion cooling and direct-to-chip liquid cooling to remove heat at the source. For example, ST Telemedia Global Data Centres (STT GDC) is deploying immersion cooling at scale to support AI-ready infrastructure across Southeast Asia, beginning with Singapore and Thailand. Building on successful 2022 trials, the company now integrates immersion and direct-to-chip liquid cooling into high-density GPU racks. By submerging servers in dielectric fluid, STT GDC removes heat at the source, cuts energy use, and reduces reliance on air cooling—supporting sustained AI workloads.

Water scarcity is a critical constraint in Singapore

In resource-limited geographies like Singapore, where most freshwater is desalinated or highly treated, cooling water is a critical resource. The priority is to reduce dependency by using less of it or sourcing it through alternative means.

Singapore is piloting alternative cooling solutions to cut reliance on freshwater infrastructure. A flagship example is Keppel's proposed Floating Data Centre Park at the Loyang Offshore Supply Base, which would use seawater for heat rejection instead of freshwater or land-based cooling towers. This design frees up treated water for essential uses, releases scarce land for other developments, improves energy efficiency through integrated off-site power generation, and offers a future-proof model for sustainable data center growth in land-scarce cities. These moves are driven by the high cost and carbon footprint of desalination and water recycling, as well as the need to preserve limited freshwater reserves. However, planning for seawater-based cooling and offshore data centers requires new resilience strategies, factoring in risks such as tsunami exposure, sea-level rise, and the need for storm and flood protection.

Water efficiency is currently embedded in the government initiatives, such as Tropical DC Standard and BCA Platinum Green Mark certification, which require operators to track and minimize WUE. Under the Green DC Roadmap, the government has signaled its intention to work with industry toward net-zero water or zero-water cooling solutions over the long term. In the meantime, operators are improving WUE by using reclaimed NEWater instead of potable water (NEWater is Singapore's ultra-purified reclaimed water, produced from treated wastewater and used mainly for industrial and cooling purposes to conserve potable water). The long-term ambition is zero freshwater cooling, though this remains technically challenging and climate-dependent.

Public–private partnership in Singapore's data center sector aligns government ambition with industry innovation, built on joint consultation and a shared commitment to climate goals.

Singapore's push for greener data centers is being shaped through a public–private partnership model, with strong alignment between government and industry. Climate action is seen as a shared responsibility, and regulation is viewed as constructive rather than burdensome — a contrast to some markets where blanket PUE mandates face resistance due to legacy infrastructure. The 1.3 PUE benchmark was set through joint consultation between policymakers and operators, embedded in the BCA Platinum Green Mark certification, and developed by a technical committee that included tech firms. The result is an ambitious yet technically grounded target, reflecting a collaborative policy design where government ambition and industry innovation advance together.

POLICY LANDSCAPE

Singapore's data center policy framework has undergone a fundamental pivot—from managing growth through capacity restraints to steering it through a highly selective, sustainability-first model. Following the 2019–2022 moratorium on new builds, the government undertook a strategic reassessment of the sector, weighing environmental trade-offs against digital infrastructure needs. This policy recalibration resulted in the Data Centre Call for Application (DC-CFA), a competitive pilot process that allocates capacity only to projects that exceed stringent energy and sustainability thresholds.

The framework integrates technical and environmental requirements through multiple instruments. The updated BCA Green Mark for Data Centres embeds sustainability criteria into both new builds and retrofits, with clear differentiation between design PUE and operational PUE. Bonus credits are offered for advanced measures such as on-site renewable generation, heat recovery systems, and real-time monitoring and analytics. The Tropical Data Centre Standard codifies climate-specific cooling benchmarks to address the performance challenges in Singapore's hot and humid conditions. Meanwhile, the Monetary Authority of Singapore's Green Taxonomy channels capital toward certified green data center assets, lowering financing costs for compliant developments and making sustainability a financial as well as regulatory imperative. This policy–finance–technology nexus signals a decisive market shift: efficiency and sustainability are no longer competitive advantages — they are prerequisites for market entry.

In 2024, Singapore reinforced this trajectory through the Green Data Centre Roadmap, which earmarks at least 300 MW of new capacity to be developed under next-generation sustainability standards. These include hydrogen-ready power systems, advanced liquid cooling technologies, and AI-driven energy optimization. Crucially, the Roadmap is anchored to the Singapore Green Plan 2030 — the national blueprint for peaking emissions before 2030 and achieving net-zero by 2050— ensuring that digital infrastructure expansion aligns with the economy's long-term climate commitments.

Table 12: Digital Infrastructure Policies

Policy Name	Date	Detail	References
Green Data Centre Roadmap	2024	Launched by Infocomm Media Development Authority (IMDA) and Singapore Economic Development Board (EDB)in May 2024 to add at least 300 MW of new capacity over the next few years, prioritizing projects that deliver both sustainability and high economic value. Sets aggressive PUE and WUE efficiency targets and promotes advanced technologies such as liquid cooling, green fuels, on-site solar, and hydrogen-ready infrastructure.	Infocomm Media Development Authority
Tropical Data Centre Standard	2023	IMDA-developed operational standard allowing data centers to run efficiently at ambient temperatures of ≥ 26 °C, reducing reliance on energy-intensive cooling in tropical climates. Designed to optimize chiller loads without compromising reliability or uptime, enabling up to 30–40% cooling energy savings compared with conventional 22 °C settings.	Infocomm Media Development Authority
Singapore-Asia Taxonomy (SAT)	2023	Launched by the Monetary Authority of Singapore (MAS), it aims to define and classify environmentally sustainable and transition activities in an Asian context. It provides clear criteria for businesses and investors to identify green projects, channel capital toward credible climate solutions, and curb greenwashing.	Monetary Authority of Singapore
Data Centre – Call for Application (DC-CFA)	2022	Competitive allocation framework for new data center capacity following the moratorium. Awards based on sustainability, innovation, and economic contribution. The July 2023 awards went to AirTrunk–ByteDance, Equinix, GDS, and Microsoft. Evaluation criteria included achieving BCA–IMDA Green Mark Platinum, integrating international connectivity, anchoring AI/HPC workloads, and supporting the local talent pipeline.	Infocomm Media Development Authority
Singapore Green Plan 2030	2021	Whole-of-nation sustainability blueprint targeting peak emissions before 2030 and net zero by 2050. For the data center sector, emphasizes renewable energy sourcing, efficiency standards, and green certification.	Singapore Green Plan 2030
BCA–IMDA Green Mark for Data Centres	Ongoing	Voluntary green building certification (introduced 2012/2013; major update 2020). Platinum rating requires design PUE ≤ 1.5, robust cooling load management, water efficiency, and IT energy optimization. Distinguishes between design and operational PUE, with bonus points for on-site renewables, heat recovery, and advanced monitoring.	Infocomm Media Development Authority
Moratorium on New Data Centres	2019-2022	Three-year pause on new DC approvals due to concerns over grid strain, high energy/water intensity, and carbon impact. Lifted in July 2022 with a pilot allocation scheme initially capped at 60 MW/year (later raised to 80 MW/year) and stringent sustainability conditions, paving the way for the DC-CFA process.	Singapore Economic Development Board

UNITED STATES

Key Indicators

DIGITAL ECONOMY SCORE: SCORE 59/100

BURN-TO-EARN INDEX: 70.9 GCO2E PER $ - RANK 41/126

ELECTRICITY MIX: NATURAL GAS 42%, NUCLEAR 18%, COAL 16.7%, WIND 9.6%, HYDRO 6%, SOLAR 4.8%

INTRODUCTION

The United States sets the global benchmark for digital infrastructure development, hosting over half of the world's capacity and leading in hyperscale infrastructure. Cloud providers like Amazon, Google, and Microsoft drive 60–70% of the national data center power load. Enterprise-owned centers make up the remaining 20–30%.

Building on this dominance, the U.S. data center landscape is strategically concentrated in 15 states that together account for 80% of total demand, reflecting the intersection of power availability, fiber infrastructure, and market demand. Among those, Northern Virginia leads globally with over 4,000 MW of commissioned capacity across 300+ facilities. Dallas and Silicon Valley each surpass 1,600 MW, supported by mature cloud ecosystems and strong fiber connectivity. Phoenix, with 1,380 MW and more than 100 facilities, continues rapid expansion due to affordable energy and land availability. Chicago (700+ MW), New York Tri-State (600+ MW), and Atlanta (500+ MW) anchor critical enterprise and financial workloads with resilient infrastructure. On the coasts, Miami and Los Angeles serve as global connectivity hubs, enabling international traffic through interconnection points like NAP of the Americas and One Wilshire. Data center development is gradually extending beyond traditional hubs to regions such as Indiana, Iowa, and Wyoming, where available grid capacity and access to generation are favorable.

LOCAL DIGITAL INFRASTRUCTURE LANDSCAPE

U.S. data center electricity demand is projected to nearly triple by 2030.

In 2024, U.S. data centers consumed an estimated 200 TWh of electricity, accounting for about 4.4% of national demand. Driven by surging AI workloads and accelerated chip deployment, by 2028, the electricity share from data centers could nearly triple by 2030, to between 6.7% to 12% of total U.S. electricity consumption, depending on AI deployment rates and chip intensity. Data center power requirements in the U.S. are driven largely by AI. By 2028, AI-specific demand could soar to 165–326 TWh annually, surpassing all current data center usage and approaching the electricity needs of nearly a quarter of U.S. households.

Typical data center builds now exceed 50 MW, with more recent contracts hitting 1 GW and hyperscale developers planning 2–5 GW campuses — rivaling the capacity of the largest nuclear and gas plants. Meanwhile, vacancy rates are tightening across major hubs: Northern Virginia leads with 2,060 MW under construction and just 1% vacancy, while Silicon Valley (2.3%) and Atlanta (3.6%) highlight severe supply-demand imbalances despite rapid expansion in markets like Dallas, Phoenix, and New York.

In response to soaring AI-driven electricity demand, hyperscalers are investing directly in gas and nuclear infrastructure for reliability, while corporate PPAs continue to drive renewable energy expansion across the U.S. grid.

Natural gas continues to play a central role in powering U.S. data centers in 2024, driven by its dispatchability and the reliability demands of hyperscale operations. Looking ahead, U.S data center operators view natural gas as a short-term "bridge fuel" to cover demand over the next 5–10 years while other options, such as renewables and storage, scale. Regulatory restrictions mean few new gas turbines can be added before 2029–2030, though projects already approved are being built to serve near-term load. Utilities across 33 states have revised their Integrated Resource Plans (IRPs) — the formal, state-mandated blueprints for future energy supply — to reflect a sharp rise in AI and data center-driven demand, delays in transmission expansion, and growing concerns over grid stability. As a result, many IRPs now call for expanded natural gas-fired generation, alongside renewables and storage. Notably, tech companies like Meta and Amazon are directly co-investing in gas infrastructure, with some gas plants co-located at data center sites. Yet, a challenge in scaling gas lies in the rising cost of essential construction materials. Steel, aluminum, copper, timber, and cement prices have surged by 40% over the past five years, posing further constraints on scaling power and gas infrastructure. Besides natural gas, coal also remains part of the U.S. power mix in states like West Virginia, Wyoming, Indiana, and Kentucky, where legacy infrastructure and limited renewable integration keep coal generation above 40–80% of the electricity supply.

Renewables are expanding steadily, with corporate PPAs playing a central role in driving clean energy procurement. U.S. hyperscalers lead globally in PPA activity. Amazon alone has contracted nearly 35 GW of clean energy to date, cementing its position as the top corporate clean energy buyer worldwide. However, the pace of deployment is moderated by regional permitting, interconnection constraints, and supply chain disruption. In many state jurisdictions, transmission lines needed to deliver renewable power to data centers can take over a decade to develop, while land constraints often limit the ability to build sufficient capacity adjacent to demand centers. Another challenge is ongoing supply chain disruption for renewable deployment. While the U.S. has sufficient manufacturing capacity for solar modules, rising demand still relies on imported components such as wafers, cells, and inverters. As a result, renewable costs to power data centers in the U.S. are rising. Many of the most cost-effective sites have already been developed, meaning future renewable projects will likely face higher baseline costs. This makes past price levels increasingly difficult to replicate, though leading U.S. hyperscalers continue to demonstrate a willingness to pay modest premiums to advance their climate goals.

Best Practices

TAILORED INCENTIVES FOR DIGITAL INFRASTRUCTURE:

The U.S. has developed a suite of targeted incentives (e.g., fast-track permitting, R&D, semiconductor, nuclear) to support the build-out of energy-resilient, AI-specific infrastructure. These measures promote data center deployment, domestic chip innovation, next-generation nuclear, and high-performance grid integration.

PRIVATE SECTOR–LED CLEAN ENERGY PROCUREMENT:

U.S. hyperscalers lead globally in long-term clean energy procurement through Power Purchase Agreements (PPAs) and are pushing the frontier with 24/7 Carbon-Free Energy (CFE) matching, aiming to align electricity consumption with clean generation on an hourly basis.

GRID FLEXIBILITY THROUGH DEMAND RESPONSE:

Some U.S. States are advancing demand-side participation by integrating AI data centers into demand response programs. This enables delay-tolerant AI inference workloads to act as grid assets, providing flexibility during peak hours while improving energy cost management and system reliability.

Additionally, political and permitting considerations — such as restrictions on certain wind developments — remain a key factor in shaping the pace and scale of renewable deployment for data centers.

Nuclear power is regaining strategic prominence in the U.S. as a cornerstone of long-term grid reliability and decarbonization. As of 2024, the U.S. operates the world's largest nuclear fleet — 94 reactors providing nearly 97 GW of capacity — with nuclear supplying approximately 18% of the electricity consumed by data centers. To meet the dual challenge of surging AI-driven electricity demand and 2050 net-zero targets, the U.S. Department of Energy projects that nuclear capacity will need to triple. SMRs are gaining traction as a scalable, zero-carbon solution for powering future data centers, with over 20 GW of SMR capacity announced for post-2030 deployment by leading tech firms. Hyperscalers are increasingly exploring co-location with nuclear facilities to secure clean, reliable baseload power. Notably, AWS's $650 million acquisition of Talen Energy's 960 MW Cumulus data center campus — co-located with the Susquehanna nuclear plant — signals a decisive shift toward energy-resilient, climate-aligned hyperscale infrastructure. Federal policy is also catching up. The U.S. Nuclear Regulatory Commission (NRC), the federal agency overseeing nuclear safety and licensing, is streamlining approvals for SMRs to accelerate deployment and meet rising data center demand.

The rapid expansion of data centers is putting increasing pressure on local land use, infrastructure, and natural resources.

In high-density areas such as Santa Clara, California, data center clusters have raised zoning tensions as industrial development encroaches on residential zones. Water use is another critical flashpoint: U.S. data centers consumed over 500 million m^3 (132 billion U.S. gallons) of water in 2018, placing the sector among the top 10 industrial water users globally. In arid regions such as Arizona and Oregon, reliance on evaporative cooling has drawn scrutiny. Meanwhile, in fast-growing hubs such as Northern Virginia, interconnection delays and growing reliance on diesel backup generators during grid constraints have prompted public concern over emissions and noise. Economic concerns are also emerging, as communities question whether low-employment facilities on prime land align with broader development goals. Some jurisdictions, such as Loudoun County, Virginia, have responded by tightening zoning rules, requiring special exceptions and public hearings for new data centers.

In response, a mix of regulatory and market-based approaches is beginning to take shape. California is proposing legislation (SB 57, AB 222) to mandate zero-carbon electricity procurement, public energy reporting, and water-use disclosure for large data centers. Oregon has proposed requiring all data centers to shift to renewable power by 2040, while utilities would be restricted from passing infrastructure costs to ratepayers. Loudoun County is considering microgrids as a condition for permitting, addressing both resilience and local grid constraints.

On the industry side, major operators are increasingly participating in demand-response programs across some states — such as Virginia, Oregon, and California — curbing electricity use during peak demand to ease grid strain. Operators are testing several demand response strategies, from shifting workloads across regions to using advanced hardware controls to manage energy use. Some facilities even run on their own diesel generators during peak stress events, a practice known as "islanding," which eases pressure on the grid. These measures are driven by both cost and reputational benefits, as data centers aim to align with regulators and ratepayers. With Google announcing a dedicated data center demand response initiative in August 2025, the approach is quickly moving into the mainstream. Others, like Equinix, Switch, and Digital Realty, are investing in on-site microgrids that integrate solar, batteries, and low-carbon backup power. Importantly, demand response products vary significantly by state, which shapes how data centers can participate; some jurisdictions emphasize fast frequency response, while others incentivize capacity or peak shaving. In Virginia, data center firms and utilities are now collaborating through the Clean Energy Advisory Council to align infrastructure planning with sustainable energy objectives. These developments suggest a shift toward more proactive, grid-conscious, and environmentally responsible data center growth.

POLICY LANDSCAPE

Since January 2025, the U.S. federal policy landscape for powering AI and data center infrastructure has rapidly evolved under a new administration, reflecting the strategic urgency of aligning energy systems, permitting reform, and digital infrastructure development.

A pivotal shift in U.S. data center policy came with the AI Action Plan (July 2025), which proposed over 90 federal measures to accelerate the rollout of AI-related infrastructure—chiefly by dismantling permitting bottlenecks. The plan explicitly targets fast-tracking data center construction through streamlined environmental reviews under the National Environmental Policy Act (NEPA) and Clean Water Act (CWA), while easing emissions and remediation rules under the Clean Air Act (CAA) and Comprehensive Environmental Response, Compensation, and Liability Act (CERCLA). Reinforcing this deregulatory push, a July 2025 Executive Order authorized the

Department of Commerce to fund hyperscale data center development and directed agencies to revise legacy permitting frameworks to meet AI-era urgency. Months earlier, in April, the Department of Energy (DOE) issued a Request for Information (RFI) identifying 16 federal sites for expedited data center siting — signaling a proactive land-use strategy. These actions follow the president's January 2025 declaration of a national energy emergency, which, through the Unleashing American Energy Executive Order, prioritized fossil fuel development, permitting reform, and critical mineral supply chains — all under the banner of energy security. Together, these policies mark a significant pivot from the former administration's climate-forward approach to a more infrastructure-driven and energy security-focused agenda for accelerating digital infrastructure growth.

Before this transition, the former administration had established the White House Task Force on AI Data Center Infrastructure (September 2024), examining how advanced computing could accelerate the integration of clean energy into the grid. The Federal Energy Regulatory Commission (FERC) laid additional groundwork through Order No. 2023 (July 2023), which streamlined generator interconnection, and Order No. 1920 (May 2024), which mandated long-term transmission planning to support high-load growth scenarios such as AI data centers.

Policy pillars under the former administration also included semiconductor industrial strategy via the CHIPS and Science Act (2022), and robust renewable energy incentives through the Inflation Reduction Act (IRA), which introduced technology-neutral clean electricity credits, storage tax incentives, and funding for transmission. These tools helped drive record corporate clean power procurement and accelerated renewable buildout—especially in digital infrastructure zones. The DOE's Grid Resilience and Innovation Partnerships (GRIP) Program, launched in 2024, allocated over $4 billion to modernize transmission and add 20.5 GW of capacity, directly supporting clean energy interconnection.

Table 13: Digital Infrastructure Policies

Policy Name	Date	Detail	References
Policy since January 20, 2025			
Winning the Race: America's AI Action Plan	July 2025	AI Action Plan (July 2025) outlines 90+ federal actions to support AI infrastructure. Pillar II of the AI Action Plan specifically targets the overhaul of permitting and regulatory processes that delay critical AI infrastructure. It proposes streamlining environmental reviews under the National Environmental Policy Act (NEPA) and the Clean Water Act (CWA) and easing emissions and site cleanup rules under the Clean Air Act and the Comprehensive Environmental Response, Compensation, and Liability Act (CERCLA). It also directs federal agencies like the Department of Energy, Department of Defense, and Department of the Interior to identify federal lands suitable for large-scale AI data centers and associated energy infrastructure.	The White House
Accelerating Federal Permitting of Data Center Infrastructure	July 2025	The Executive Order authorizes the Department of Commerce to fund large-scale data centers and related infrastructure, while directing Environmental Protection Agency (EPA) and other agencies to fast-track permitting by revising key environmental regulations in line with the AI Action Plan.	The White House
AI Infrastructure on DOE Lands Request for Information	April 2025	DOE RFI (April 2025) explores AI infrastructure development on 16 federal sites with fast-track permitting potential.	The U.S. Department of Energy
Policy before 20 January 2025			
Recommendations on Powering Artificial Intelligence and Data Center Infrastructure	July 2024	Department of Energy's roadmap highlights soaring AI-driven data center power demands and outlines steps like an AI Data Center Testbed, utility collaboration, and rapid grid assessment.	The U.S. Department of Energy

Advancing United States Leadership in Artificial Intelligence Infrastructure	July 2025	The Executive Order directs federal agencies to prioritize permitting, siting, and funding for AI-related infrastructure—particularly data centers—and streamlines permitting for geothermal and nuclear energy to support their energy needs.	The White House
White House Task Force on AI Datacenter Infrastructure	September 2024	The White House Task Force on AI Data Center Infrastructure convened experts from utilities, grid operators, software firms, NGOs, and other stakeholders to examine how advanced AI computing can accelerate the integration of clean energy into the power grid.	The White House
Recommendations on Powering Artificial Intelligence and Data Center Infrastructure	July 2024	Department of Energy's roadmap highlights soaring AI-driven data center power demands and outlines steps like an AI Data Center Testbed, utility collaboration, and rapid grid assessment.	The U.S. Department of Energy
Data Center Optimization Initiative (DCOI)	Introduced 2016, updated 2019	DCOI mandates 24 federal agencies to consolidate inefficient data centers, improve energy performance, enhance cybersecurity, and transition to cloud and shared services. Federal agencies are required to measure PUE and achieve specific benchmarks, increase server utilization rates, and shut down under-utilized facilities.	U.S. Office of Management and Budget
FERC Order 1920: Building for the Future Through Electric Regional Transmission Planning and Cost Allocation	May 2024	Federal Energy Regulatory Commission (FERC), Order 1920 mandates regional transmission providers to conduct long-term (20-year) transmission planning, incorporating high-load growth scenarios — including AI data centers and electrification trends.	The U.S. Federal Energy Regulatory Commission
FERC Order 2023: Improvements to Generator Interconnection Procedures and Agreements	July 2023	Federal Energy Regulatory Commission (FERC), Order 2023 streamlines generator interconnection, accelerating capacity availability for data centers.	The U.S. Federal Energy Regulatory Commission
Grid Deployment Office (GDO)	2022	Founded in 2022, the Grid Deployment Office (GDO) supports the expansion and modernization of the U.S. electric grid by facilitating investments in generation facilities, enhancing high-capacity transmission systems, and advancing transmission and distribution technologies across the economy.	The U.S. Department of Energy
Other related targeted Policies			
Semiconductors (CHIPS and Science Act)	2022	The CHIPS and Science Act (2022 is a landmark U.S. law aimed at strengthening the U.S. semiconductor manufacturing, R&D, and national competitiveness. It authorizes around $280 billion in funding, with a major share directed toward boosting the semiconductor supply chain and related innovation.	McKinsey
Energy Policies		Under the Biden-Harris administration, the Investing in America & Permitting Action Plan (May 2022) focused on accelerating clean energy and infrastructure approvals with strong environmental safeguards. Additionally, the Department of Energy's 2024 Grid Resilience and Innovation Partnerships (GRIP) Program committed over USD 4 billion to transmission upgrades, enabling 20.5 GW of capacity and expediting clean energy interconnections	APEC Energy Overview 2025
Policies for Nuclear Energy Development		• Advanced Reactor Demonstration Program (ARDP) – Launched by the DOE in 2020 with $230 million to accelerate SMR deployment by companies like Kairos Power, X-Energy, and TerraPower. • Energy Act of 2020 – Authorized significant funding for next-generation nuclear technologies, supporting programs like ARDP. • Infrastructure Investment and Jobs Act (2021) – Provided over $5 billion in additional funding for advanced nuclear projects, including SMRs. • Inflation Reduction Act (2022) – Introduced tax credits supporting nuclear power: Sections 48E and 45Y (technology-neutral clean electricity credits) and Section 45U (zero-emission nuclear production credit). • DOE Loan Programs Office (2024) – Issued a $1.52 billion loan guarantee to restart the 800 MW Palisades Nuclear Plant in Michigan, with plans for two SMRs.	ING Research

ABAC

APEC Business Advisory Council